

# How Adaptive Learning Affects Evolution: Reviewing Theory on the Baldwin Effect

B. Sznajder · M. W. Sabelis · M. Egas

Received: 22 July 2011 / Accepted: 8 December 2011 / Published online: 25 December 2011  
© The Author(s) 2011. This article is published with open access at Springerlink.com

**Abstract** We review models of the Baldwin effect, i.e., the hypothesis that adaptive learning (i.e., learning to improve fitness) accelerates genetic evolution of the phenotype. Numerous theoretical studies scrutinized the hypothesis that a non-evolving ability of adaptive learning accelerates evolution of genetically determined behavior. However, their results are conflicting in that some studies predict an accelerating effect of learning on evolution, whereas others show a decelerating effect. We begin by describing the arguments underlying the hypothesis on the Baldwin effect and identify the core argument: adaptive learning influences the rate of evolution because it changes relative fitness of phenotypes. Then we analyze the theoretical studies of the Baldwin effect with respect to their model of adaptive learning and discuss how their contrasting results can be explained from differences in (1) the ways in which the effect of adaptive learning on the phenotype is modeled, (2) the assumptions underlying the function used to quantify fitness and (3) the time scale at which the evolutionary rate is measured. We finish by reviewing the specific assumptions used by the theoretical studies of the Baldwin effect and discuss the evolutionary implications for cases where these assumptions do not hold.

**Keywords** The Baldwin effect · Fitness landscape · Evolution of phenotype · Adaptive learning · Innate behavior · Phenotypic variation · Genetic variation

## Introduction

In theory the evolution of plastic phenotypic traits may lead to two extreme outcomes where the trait becomes either genetically fixed (and largely phenotypically invariable) or it is entirely shaped by environmental influences. In-between these extremes lies a spectrum of outcomes where traits contain a genetic component but they are also, to various degrees, modifiable in response to environmental influences. A recurring question in evolutionary biology is how phenotypic plasticity (i.e. the ability to modify phenotype in response to external or internal influences—see “**Box**”) may influence the outcome and the rate of evolution (Price et al. 2003; West-Eberhard 2005; Crispo 2008) by, for example, creating novel selectable forms that are entirely environmentally induced when there is not any genetic basis for such a variant (as in populations that colonize a novel environment). This question is particularly relevant in the face of the growing body of evidence that various forms of plasticity (such as the ability to learn—see “**Box**”) have a genetic basis (Mery and Kawecki 2002; Dukas 2004) and thus may evolve jointly with the genetically determined phenotype.

In this article we focus on the hypothesis that adaptive learning (i.e., learning that improves fitness; see “**Box**”) facilitates evolution of the genetic basis for phenotypic traits. This hypothesis has its origins in the arguments put forward by Mark Baldwin (1896, 1902), a contemporary of Charles Darwin. These arguments concern a population that finds itself in a new environment, and thus, presumably, does not contain a genetic basis for the complete phenotype that would be optimal in this new environment (i.e. the phenotype that achieves the highest possible fitness). Baldwin argues that adaptive plasticity allows sub-optimal individuals to acquire higher fitness. Hence, learning improves the survival

---

B. Sznajder · M. W. Sabelis (✉) · M. Egas  
Department of Population Biology, Institute for Biodiversity  
and Ecosystems Dynamics, University of Amsterdam,  
Amsterdam, The Netherlands  
e-mail: M.W.Sabelis@uva.nl

of the population of such individuals and thus it facilitates that the genetic evolution may proceed. Moreover, Baldwin observes that under these conditions there is direct selection for the ability to learn adaptively and, simultaneously, indirect selection for any heritable variation carried by the plastic individuals favored by direct selection. Clearly, this scenario applies to the evolution of a behavioral trait that is performed more than once in individual's lifetime such that there is an opportunity for the individual to modify this trait by learning. The central argument of Baldwin is that selection for the ability to acquire a fitter phenotype through learning may coincide with the genetic basis for the fitter phenotype (i.e., these indirectly selected genes provide a basis for a fitter phenotype). If this condition is fulfilled, then the selection for improved learning facilitates adaptive evolution of the genetic basis for the trait. Therefore, adaptive learning is predicted to accelerate evolution of this trait.

This hypothesis of Mark Baldwin, known in the literature as the Baldwin effect (Simpson 1953) has spurred numerous theoretical studies whose general approach is to measure the rate of evolution of a genetically determined trait, given different levels of a non-evolving ability to learn adaptively. Their results are ambiguous; some studies provide evidence for an accelerating effect of adaptive learning on evolution (Hinton and Nowlan 1987; Fontanari and Meir 1990; Mayley 1997; Ancel 2000—the norm of reactions models; Lande 2009), yet others show a decelerating effect of learning on genetic evolution (Papaj 1994; Anderson 1995; Ancel 2000—the quantitative genetic model; Dopazo et al. 2001; Borenstein et al. 2006).

In this article we analyze the theoretical studies of the Baldwin effect with the aim of explaining how learning yields these two contrasting effects. In order to do so we analyze how—in these studies—learning influences the relationship between different phenotypes and fitness and thereby influences the evolutionary response to selection. In fact, it is one of the underlying assumptions of the Baldwin effect that learning changes relative fitness differences among phenotypes such that it confers a larger fitness increase to those phenotypes (as well as underlying genotypes) that are already relatively closer to the fitness peak. In effect, the selection for the ability to acquire fitter phenotype through learning coincides with the genetic basis for the fitter phenotype. Moreover, we emphasize that the theoretical studies model two distinct evolutionary stages that are characterized by different evolutionary end-points in the Baldwin effect. The Baldwin effect concerns the evolution of a phenotypic trait towards a single and distant fitness peak; this process is initially realized through the selection of plastic phenotypes but it is finalized when these plastic phenotypes are substituted by a genetically determined and optimal phenotype (presumably because learning has a

fitness cost; Baldwin 1896; Simpson 1953). Hence, the theoretical studies of the Baldwin effect generally estimate the amount of time (in generations) needed for the completion of this entire process, but they also allow to separately analyze (1) the number of generations until the first genetically determined and optimal phenotype appears in the plastic population, and (2) the number of generations until this genetically determined optimal phenotype replaces the plastic phenotypes (which represents the general idea of staging the Baldwin effect, as first proposed by Simpson 1953). These two evolutionary stages may have different time scales and evolutionary dynamics. Therefore, it is reasonable to derive conclusions about the effect of learning on evolution separately for these two evolutionary stages.

We focus on theory pertaining to the first phase of the Baldwin effect, i.e. the evolution of genetically determined behavior towards a distant fitness peak, given that this behavior can be modified by learning. For the review of the second phase, i.e. genetic assimilation of learned behavior, we refer to Crispo (2007). Empirical evidence for the effect of adaptive learning on the rate of evolution is virtually absent, with the exception of the work by Mery and Kawecki (2004) on artificial selection on food preference in the presence or absence of learning opportunities. Our review of the theory aims to identify key assumptions about the nature of genetically determined behavior, the ability to change this behavior by learning, and the fitness cost of this ability. In doing so, we hope to stimulate empirical work on cognitive ecology and evolution. Empirical scrutiny of the often arbitrary assumptions underlying theory may prove relevant to the understanding of the evolutionary process.

In the next section, we begin by analyzing the results of the theoretical studies of the Baldwin effect, grouped with respect to the concepts of adaptive learning (or adaptive plasticity in general) and the assumed fitness function (see “Box”). We compare the evolutionary rates in the first phase of the Baldwin effect (as defined above) obtained in these studies, whenever these rates are available. By this review we aim to highlight that learning influences the rate of evolution by changing the fitness differences among phenotypes in a population. This in turn determines how the response to selection and thus the rate of evolution is affected by learning. Theoretical models differ in the way they allow learning to change relative fitnesses of different phenotypes in the population, and these differences explain their contrasting predictions.

### Models of the Baldwin Effect: The Concept of Adaptive Learning and the Choice of Fitness Function

The hypothesis on the Baldwin effect states that evolution of an innate trait (see “Box”) proceeds faster in populations

that harbor plastic individuals, than in populations that harbor none of such individuals. Therefore, the general approach in published studies of the Baldwin effect is to measure the rate of evolution of an innate trait given different levels of a non-evolving ability to learn. However, these studies vary with respect to assumed fitness functions; the fitness landscapes (see “Box”) they describe range from a single-peak ‘needle-in-haystack’ type to a single-peak landscape with a gradual slope, or a rugged landscape that contains many fitness peaks of varying heights. Moreover, although they all model adaptive learning (i.e., learning that leads to a change of phenotype in the direction of increased fitness) the reviewed studies of the Baldwin effect use different methods and assumptions to achieve such an effect of plasticity in their model systems. In this part we discuss these various modeling aspects and argue that the results of the studies are determined by the way they model adaptive learning; the models may be structured such that adaptive learning, in combination with assumed fitness function, tends to either reduce or increase fitness differences among modeled phenotypes. The degree of fitness differences due to phenotypic variation determines the degree of response to selection (Rice 2004), where reduced variation is associated with weaker response and increased variation with stronger response. This relationship provides a connection between the effect of adaptive learning on phenotype and the rate of evolution. To show how contrasting predictions may emerge, below we analyze the modeling approaches of different studies in terms of the effect of learning on the relative fitness of phenotypes.

In their seminal model reviving interest in the Baldwin effect, Hinton and Nowlan (1987) track the changes in the frequency of the alleles associated with fitness pay-off. The increase in the frequency of the allele associated with superior fitness is taken as a yardstick of adaptive evolution. In particular, genotypes are modeled as byte strings that consist of a number of loci. These loci can contain one of two types of alleles and the genotype that is completely homogeneous with respect to one particular type of the allele is taken to be the optimal one (i.e. confers the highest fitness). The assumed fitness landscape, therefore, is of the unimodal ‘needle-in-haystack’ type. Adaptive learning is introduced by another allele which is not fixed, but can be switched to the type that confers higher fitness based on a learning algorithm, and individuals are allowed to search for the correct setting of these alleles in a number of trials during their lifetime. The individuals that learn the optimal phenotype are preferentially selected for mating (where the probability of being selected for mating increases with decreasing the number of trials the individuals need to learn the optimal phenotype) and thus have more offspring. The model shows that such learning dramatically speeds up evolution in the population of individuals capable of

learning, a result corroborated by Fontanari and Meir (1990) who analyze evolution on the same fitness landscape, using the same learning protocol but assuming asexual reproduction. In fact, the population lacking an allele for learning (i.e. the ‘unspecified’ alleles that get fixed by learning) cannot find this evolutionary end-point. The explanation for these results is that at least some individuals harbor the set of fixed alleles that is not too different from the optimal one and hence they have a higher chance of finding the correct setting of all ‘unspecified’ alleles by learning within the time specified for learning. In other words, thanks to learning these genotypes (that are already closer to the fitness peak) gain a relatively higher fitness than do the plastic genotypes with fewer correct setting of alleles. These findings are also consistent with the argument of Baldwin that learning confers higher fitness gain to those genotypes that are already closer to the fitness peak and thus accelerates evolution of the genetic basis for the optimal phenotype. Nevertheless, the other observation of the model by Hinton and Nowlan is that ‘unspecified’ alleles are not entirely out-selected and remain in the population, indicating different evolutionary dynamics once the population evolves to the vicinity of the fitness peak. However, this result may also be attributed to the fact that learning in this model has no fitness cost.

The study by Mayley (1997) provides evidence that the cost of learning plays a critical role in the interplay between learning and genetic evolution. It also examines in more detail the relationship between the complexity of a fitness landscape and the effect of learning on evolution. In particular, the author compares the movement of a plastic population on unimodal and rugged (i.e. many fitness peaks of varying height) fitness landscapes. In his model, a genetically determined phenotype, represented by a point on the fitness landscape, is considered to evolve if it moves in the direction of the fitness peak. Mayley finds that there is no evolution on a unimodal fitness landscape if learning is cost-free because the optimal phenotype is acquired entirely by learning. Adaptive evolution on a unimodal fitness landscape is only possible, when there is a cost of learning. Yet, on a rugged fitness landscape the population evolves irrespective of the cost of learning. Mayley’s results demonstrate that learning is more likely to facilitate evolution on a rugged fitness landscape, i.e. where there is more than one fitness peak and/or, initially, learning allows the phenotypes to reach only the local fitness peaks but not the global fitness peak (i.e. modifies the phenotype such that it has the highest possible fitness). Moreover, in both the unimodal and rugged fitness landscapes the cost of learning is critical for the convergence of the population on the single optimal genotype, i.e. the genotype whose fitness cannot be improved by learning.

Borenstein et al. (2006) constructed a rugged fitness landscape characterized by a number of local fitness peaks of steadily increasing heights and one global fitness peak. In their model the population continues evolving towards the global optimum by crossing the intermediate fitness valleys and converging on local fitness peaks. The authors measure the rate of evolution as the time it takes the population to reach the global fitness peak and they approximate adaptive learning through the application of an algorithm which allows a learning genotype to repetitively explore the fitness landscape and to modify its phenotype according to the detected fitness gains. This learning process stops when continuation of sampling and learning cannot secure further fitness gains (i.e. the genotype has found the local fitness peak). As a consequence of this learning process all genotypes of the population acquire the same fitness, determined by the local fitness peak, because they all are equally capable of learning. This way of modeling the phenotypic effect of learning is more akin to the way learning is modeled in a series of models due to Hinton and Nowlan (1987) and Mayley (1997). One feature characteristic to this approach is that genotypes capable of learning can sample potentially large areas of the fitness landscape and modify their phenotypes accordingly. In the model of Borenstein et al. the learning process effectively smoothes the fitness landscape, i.e. it reduces fitness differences among genotypes. Model simulations carried out by Borenstein et al. confirm that such an effect of learning is associated with slower evolution on a unimodal fitness landscape. However, on a rugged fitness landscape the learning process results in faster evolution because the reduced fitness differences among genotypes help the population to cross fitness valleys, thereby allowing evolution towards the global fitness peak. At the same time, a population of individuals that cannot learn may never be able to cross the fitness valley and find the global optimum. These results prompt Borenstein et al. to conclude that the complexity of the fitness landscape, i.e., the presence of multiple fitness peaks and fitness valleys, determines whether the effect of learning on evolution is accelerating or decelerating.

A separate class of models using the quantitative genetics framework to measure the rate of phenotypic evolution assumes a unimodal fitness landscape (i.e. containing a single fitness peak), given by a Gaussian function (Anderson 1995; Ance 2000—quantitative genetics model). These studies introduce an adaptive effect of learning by an increase in the selection variance. Thus the learning process modeled is equivalent to (a small) adaptive shift of the genetically determined trait value of all sub-optimal individuals. This combination of the fitness function and the way of modeling learning results in decreased phenotypic variance and decreased fitness

differences among different phenotypes. Moreover, this evolutionary scenario approximates the second stage of the Baldwin effect: the stabilizing selection acting on the population in the vicinity of the fitness peak. Characteristically, these two studies show that learning extends the time required for convergence of the population on the optimal genotype as compared to the evolution in a population with individuals that cannot learn, thus supporting a decelerating effect of learning on evolution.

The same conclusion is drawn by Papaj (1994) in a model that measures the time required for the population to evolve a genetically determined, optimal phenotype (i.e. a genotype that has a highest possible fitness without any learning). This study also assumes a unimodal fitness landscape that is provided by a negative quadratic function (shape of inverted parabola) and, as a consequence of adaptive learning, different phenotypes eventually converge on the single fitness peak. Thus, in this study learning also effectively decreases the phenotypic variance and fitness differences among the phenotypes.

Another class of studies involves modeling adaptive plasticity as a norm of reactions. Ance (1999, 2000), in her norm of reaction model, explicitly addresses the rates of evolution in the two stages of the Baldwin effect, while varying the degree of plasticity reflected in the width of the norm of reaction. The mid-point of the norm of reactions represents the genetically determined trait value (i.e., the innate trait) while the phenotype with highest fitness within this range (based on the fitness function) represents the phenotype acquired through learning. Thus, all phenotypes are able to express the optimal phenotype if the norms of reaction of these phenotypes are wide enough to contain the fitness peak (as might be the case when the population is already in the vicinity of this fitness peak), even though there is variation in the innate value in such a population. On the other hand, setting the initial width of norms of reactions such that they do not contain the optimum, models a scenario where a population evolves towards a distant fitness peak. Ance (2000) examines how this plasticity affects the rate of evolution in two types of unimodal fitness landscapes: (1) a spiked landscape where a single genotype scores the highest fitness and all the other genotypes score the same flat fitness (also referred to as the 'needle-in-the-haystack' landscape, as in Hinton and Nowlan (1987), and (2) a Gaussian fitness function.

The novel aspect of Ance's model is that the width of the norm of reaction is allowed to evolve such that the upper and lower bounds of the norm of reactions may shift from one generation the next. For the two settings of the fitness function, Ance shows that costly adaptive plasticity generally accelerates the first stage of the Baldwin effect, i.e., it shortens the time required for the first optimal genotype to emerge in the population (Ance 2000). This

effect is associated with the initial selection for the wider norms of reactions (Ancel 1999). In contrast, plasticity decelerates the second stage of the Baldwin effect, i.e., it extends the time between the emergence of the optimal genotype and population convergence on this genotype because the wide norm of reaction effectively allow all individuals to learn the optimal phenotype (Ancel 2000).

These results of Ancel provide further evidence that adaptive learning accelerates evolution in the initial stages of the Baldwin effect, i.e. evolution towards a distant fitness peak. However, the decelerating effect of learning prevails in the second and final stage of the Baldwin effect. The results of Ancel obtained for the two stages of the Baldwin effect are corroborated by the study of Lande (2009) where plasticity is also modeled as a reaction norm evolving under the Gaussian fitness landscape. These two studies are a notable exception in the theory of the Baldwin effect by allowing phenotypic plasticity to evolve jointly with the innate trait (see also studies in the framework of artificial life/intelligence, e.g., Watson and Wiles 2002; Suzuki and Arita 2004).

Thus, the theoretical studies indicate that the effect of learning on evolution is not constant as the population evolves on a fitness landscape towards a distant fitness peak. Therefore it is reasonable to conduct a comparative analysis of the theoretical studies of the Baldwin effect on the studies that measure evolution within the same evolutionary stage (and at the same time scale). In fact, any long-term measure of evolutionary rate (such as, e.g., the time until a first genetically determined optimal phenotype appears in a population) is a net effect of the evolutionary responses occurring at each generation during evolution towards an evolutionary end-point. It is informative, therefore, to analyze how learning may influence this short-term rate of evolution occurring from one generation to the next. This is the approach used in the recent model by Paenke et al. (2007) who study the rate of evolution as determined by the degree of fitness differences due to phenotypic variation. In this way, they directly demonstrate the effect of learning on relative fitness and the rate of evolution. In the next section of this article we analyze the approach and results of this model.

### Adaptive Learning and the Response to Selection

Paenke et al. (2007) explore how a population's response to directional selection changes with improved adaptive learning (or some forms of developmental noise). To this end, the authors analyze how the relationship between phenotype and fitness changes as adaptive learning is improved. In particular, the authors compare the rate of evolution of the innate trait at two different and fixed levels

of plasticity and analytically demonstrate that improved adaptive plasticity strengthens the response to selection (and thus accelerates evolution) when it magnifies fitness differences among phenotypes: this is reflected in the steeper relationship between phenotype and fitness. Conversely, improved adaptive plasticity weakens the response to selection (and thus decelerates evolution) when it reduces fitness differences among phenotypes: this is reflected in the lower slope of the function relating phenotype and fitness.

By assuming a non-evolving learning ability, the authors entirely focus on the evolution of the innate trait (although in this model the evolution of adaptive learning may also be incorporated, thus introducing a second axis for the evolution of the phenotype). This allows them to derive a correspondence between their general result as presented above and specific properties of the fitness function (evaluated only in the direction of innate trait) reflected in the shape of the fitness function. In particular, the authors predict that learning magnifies fitness differences among phenotypes when the fitness landscape (evaluated in the direction of the innate trait) is convex. Conversely, adaptive learning reduced fitness differences among phenotypes when this fitness landscape is concave. The authors extend this analysis by assuming various specific functions for the innate phenotype and non-evolving adaptive plasticity, such as used in Ancel (2000) or Anderson (1995) to demonstrate that there exists a fitness landscape on which adaptive learning, as it is modeled, accelerates evolution.

The predictions of Paenke et al. (2007) are derived under the assumptions that the selection is directional (i.e. fitness consistently increases with the value of the innate trait) and non-evolving learning equally modifies different phenotypes (the authors point out that a form of learning that is dependent on the distance of the innate phenotype from the fitness peak may lead to novel predictions). Other assumptions of this framework include the assumption that there are no non-additive or dominance effects shaping the expression of the phenotype, or that there is no genetic covariance between the innate trait and adaptive plasticity.

The approach in the study of Paenke et al. (2007) provides an elegant demonstration of how adaptive learning influences the short-term rate of evolution, i.e., the response to selection measured from one generation to the next when learning is kept fixed. However, allowing the evolution of adaptive learning ability may change the long-term dynamics if the curvature of the fitness landscape is not overall uniform (which is assumed in the model of Paenke et al. 2007).

In summary, the results of Paenke et al. (2007) allow for the conclusion that the effect of adaptive learning on evolution depends on the shape of the fitness function as well as the model of adaptive learning. In particular,

adaptive change due to learning that is large relative to the distance of the innate trait from the fitness peak (such that optimal (or nearly optimal) phenotype can always be learned), is more likely to decelerate evolution of the innate trait irrespective of the curvature of the fitness landscape. This theoretical possibility may be unlikely, however, given that in a population adapted to an old environment (that is distant from the new fitness peak as argued in the Baldwin effect), low levels of plasticity are expected (Lande 2009), particularly if plasticity has a fitness cost. There may be selection to maintain high levels of plasticity in a population if there are frequent changes of environment (Stephens 1991). However, in such a theoretical situation the environment (and thus a fitness landscape) is dynamical, while the theoretical studies of the Baldwin effect generally assume a constant environment (and thus a constant fitness landscape).

## Discussion

The effect of adaptive learning on evolution of genetically determined traits is the subject of a long-standing debate and the theoretical treatments of this question provide contrasting results. Here, we discussed how these contrasting results can be partly explained from the different ways in which the theoretical studies measure the evolutionary rate. The traditional end-point of the Baldwin effect is the complete convergence of a population on an initially distant fitness peak associated with reduction in the level of adaptive learning. Adaptive learning is considered to accelerate evolution if it helps to reach this end-point faster. This measure however may fail to adequately describe the effect of learning on evolution if this effect is not constant but changes as the population evolves on a fitness landscape (particularly, a rugged fitness landscape). A measure of short-term evolutionary change as occurring from one generation to the next may be better suited to detect the variable effect of learning on evolution. The recent study by Paenke et al. (2007) provides such a framework where such a measure is employed to demonstrate how learning influences fitness differences among different innate phenotypes, thus either accelerating or decelerating the evolution of the innate phenotype. By relating the effect of learning on fitness differences among phenotypes to the shape of the fitness function (that determines these fitness differences) the authors demonstrate how theoretical predictions of the Baldwin effect depend on the choice of fitness functions. However, our analysis of this and other theoretical studies of the Baldwin effect indicates that the model of adaptive learning (i.e. how learning is modeled to change the innate phenotype) also matters to the theoretical predictions.

By definition adaptive learning modifies the phenotype so as to increase its fitness. However, adaptive learning may be characterized with respect to how much it modifies the innate phenotype given the distance of this innate phenotype from a fitness peak. In other words, the magnitude of the phenotypic modification due to learning can be modeled as either a small or a large step in phenotype space, depending on the size of the exploratory range attributed to the individuals. In particular, simulation models (Hinton and Nowlan 1987; Mayley 1997; Borenstein et al. 2006) employ learning which allows the genotype to sample large areas of a fitness landscape in search of a local fitness peak. In this process, phenotypes are allowed to experience many learning trials during their lifetime (as in Hinton and Nowlan 1987; Fontanari and Meir 1990) or adaptive search is repeated until phenotypic fitness can no longer be improved (Mayley 1997; Borenstein et al. 2006). Therefore, the optimal phenotype can be learned by all phenotypes. In contrast, in another class of models (Anderson 1995; Ancel 2000—quantitative genetic model) adaptive learning effectively involves a relatively small (with respect to the distance of the innate trait from the fitness peak) adaptive shift of the innate trait in the direction of increased fitness. Therefore the optimal phenotype (defined by the fitness peak) cannot be learned by all phenotypes, at least not in the first phase of the Baldwin effect (i.e. when the fitness peak is distant to the position of the population on the fitness landscape). This distinction between the two ways of approximating adaptive learning, based on the potential of learning to modify the phenotype, is relevant because each of these two modes of learning has distinct consequences for the relative fitness of individual phenotypes. We argue that adopting one or the other mode of learning may be particularly relevant in the case of evolution on a rugged fitness landscape. Adaptive learning that has a large potential to modify the phenotype is exemplified by unconstrained adaptive search of the fittest options on the fitness landscape. We argue that the effect of such learning on evolution is less likely to depend on the local curvature of the fitness slope because it allows genotypes to sample distant areas of the fitness landscape. On the other hand, adaptive learning modeled as a small shift of the phenotype is much less likely to allow the population to cross fitness valleys and find a global fitness peak.

One possible end-point of the Baldwin effect is the emergence of a completely genetically determined phenotype, implying the loss of plasticity once at the fitness peak. Such an extreme outcome however is quite unlikely in the real world, because the environment varies and thus the value for the optimal phenotype fluctuates in time. The ability to adjust behavior by learning may then confer sufficient fitness benefits. It remains to be explored how

adaptive learning influences evolution of the genetic basis for phenotypic traits on a fitness landscape that is dynamic due to environmental changes (see e.g. Anderson 1995) or a fitness landscape where the optimal phenotype depends on the frequency of other phenotypes in the population.

Theory shows that the cost of learning plays a crucial role in the evolutionary dynamics of traits modified by learning. Experimental evidence for costs of learning are only beginning to emerge (Mery and Kawecki 2003; or the cost of phenotypic plasticity see Auld et al. 2010), yet they are essential to motivate biologically realistic cost functions in the theoretical models of joint evolution of learning and innate behavior. Any cost of learning determines the evolution of learning and, therefore, it will play a particularly relevant role in any model of joint evolution of adaptive learning and innate behavior. Another common assumption awaiting empirical scrutiny is that all genotypes are equally capable of learning. This, however, need not be the case and theoretical predictions may change entirely if the level of learning is variable for different genotypes (for example, if there is a correlation between the genetically determined trait value and the level of learning as discussed in Mery and Kawecki 2004).

Current theory also assumes that learning is a fixed trait, and hence tends to concentrate on tracking evolution of the genetic basis alone. This assumption is challenged by the empirical evidence showing that adaptive learning can be successfully subjected to artificial selection (e.g. Mery and Kawecki 2002; Dukas 2004). It remains to be shown how the current theoretical predictions change if adaptive learning is allowed to evolve jointly with the genetically determined trait. Moreover, although not considered in the theory on the Baldwin effect, the mechanism of learning may not always be adaptive (as in the case of non-associative mechanisms of learning) and may give rise to entirely different evolutionary dynamics.

To date, empirical evidence for a role of learning in evolution is virtually absent (but see Mery and Kawecki 2004). An empirical approach requires a model system where (1) genetic variation for both a behavioral trait and the ability to learn are demonstrated, and (2) where the level of learning (Cahill et al. 2001) and the innate value of the behavioral trait (Samuels 2004) can both be quantified as separate traits. Evidence is growing that the above requirements are often satisfied in ecological systems involving, e.g., parasitoids and their hosts (Wang et al. 2003; Hoetjes et al. 2011; Takemoto et al. 2011) or predatory mites and their prey species (Egas and Sabelis 2001; Nomikou et al. 2003; Sznajder et al. 2011) or other species (Dukas and Bernays 2000; Behmer et al.; 2005). Behavioral responses in such ecological systems provide a model to study the role of learning in evolution as well as in ecological interactions.

Studies of brood-parasitic indigobirds (Payne et al. 2000) provide an ecological scenario where the first phase of the Baldwin effect may apply. In this species male chicks learn to perform the song of their hosts whereas female chicks learn to prefer the males exhibiting the song of their host. When a female indigobird lays her eggs in the nest of a novel host then its offspring will learn and exhibit host preference different to the preference learned and exhibited by the parent (Payne et al. 2000). Thus, new phenotypic variants emerge as a result of learning in the rearing environment and without a change in their genetic background. The prediction from theory on the Baldwin effect is that if genetic variation in the direction of the learned traits exists (or it subsequently occurs through mutation) then there is potential for learning to guide the evolution of the genetic basis for these traits (provided that they are indeed positively associated with fitness). Furthermore, given the reproductive isolation from the ancestor (due to different mating preferences mediated by song preferences) this novel variant may become a new species through the process of selection acting on new mutations (ten Cate 2000).

At the core of the Baldwin effect is the notion that learning changes the rate of evolution by influencing the way selection acts on the phenotypic variation and thereby the underlying genetic variation. While adaptive learning provides one source of phenotypic variation, applicable to behavioral traits in particular, there is a wealth of studies that report on other sources of new phenotypic variation that becomes available to natural selection in novel environments. This includes phenotypic plasticity in physiological processes involved in reproduction that provides the basis for colonization and evolution in novel environments; for example, timing and length of breeding season or hormonally regulated modifications thereof in dark-eyed juncos (Yeh and Price 2004) and house finches (Badyaev 2009). Another example includes phenotypic evolution driven by changes in diet and foraging habits in birds, as discussed by Price et al. (2003).

Another critical notion underlying the Baldwin effect is that new genetic variation emerges in novel environments, which then becomes exposed to natural selection. Traditionally genetic mutations are assumed to be the source of novel genetic variants. However, recent advances point to new processes. For example, environmental factors such as stress may induce the expression of hidden genetic variation (see e.g. Badyaev 2005, 2009, McGuigan et al. 2011). Other relevant examples come from epigenetic processes that change the expression of the genetic material thereby inducing the expression of hidden genetic variation (Bossdorf et al. 2008, Youngson and Whitelaw 2008) or that are responsible for some forms of phenotypic plasticity and parental effects (Bossdorf et al. 2008, Badyaev and

Uller 2009). This includes transgenerational transmission of phenotypes induced by e.g. predators (Agrawal et al. 1999) or stress (Badyaev 2005). A framework for the role of genetic and non-genetic inheritance in evolution is provided by Day and Bonduriansky (2011).

So far, we did not consider feedback from the phenotype into its environment. That such feedback exists is postulated within the perspective of niche construction (Day et al. 2003; Laland and Sterelny 2006). Learning may play a role in the choice of environment and in the way the phenotype modifies its local environment (Laland and Sterelny 2006). If so, there is a feedback from the phenotype that would introduce an ever-changing fitness function, unlike that assumed under the Baldwin effect. Thus, there is a need to extend current theory on the Baldwin effect to include niche construction and evolution in non-constant environments in general.

### Box: Explanation of Important Terms

#### Phenotypic Plasticity

The ability of one genotype to produce more than one phenotype in response to different biotic and/or abiotic environments (Scheiner 1993). Under a regime that imposes (long-lasting) selection on a trait phenotypic plasticity provides a means for a genotype to express the phenotypic variant of a trait that is favored by this selective regime. Thus plasticity is adaptive if in an altered environment it allows the genotype to express a phenotype that is fitter in this environment than the phenotype the genotype would have had if it were not adaptively plastic (Rice 2004; Garland and Kelly 2006). If adaptive, phenotypic plasticity may play an important role in the evolution of traits it modifies (Via et al. 1995; de Jong 2005; Crispo 2008).

#### Learning

In terms of its behavioral effects, learning is defined as the ability of an individual to modify its behavior due to experience that the individual remembers (Kawecki 2010). In terms of processes on the neural level, it is the acquisition of neural representation of new information (Dukas 2004). In terms of the effect of learning on Darwinian fitness, which is the focus of this review, learning is adaptive when it improves individual fitness. By way of illustration, Dukas and Bernays (2000) demonstrated the adaptive value of learning by showing in an experiment that grasshoppers that could employ associative learning for diet choice experienced higher growth rates than grasshoppers that were deprived of cues to learn associatively.

#### Innate Behavior

In this review this term refers to behavior that is genetically determined (i.e. there are genes that determine what this behavior is). The terms *genetically determined* and *innate* are interchangeably used in studies of the Baldwin effect to describe the phenotype that an individual expresses before it has an opportunity to learn, and to differentiate it from the phenotype that the individual expresses after it had learned. Thus an overall net effect of learning on phenotype is contrasted with the behavior as determined by genes alone. It is important to stress that in studies of the Baldwin effect the behavior under study is considered to have a genetically determined (innate) component as well as a learned component (see also Samuels 2004 for a discussion of the term *innate* in the context of other biological and cognitive questions). In summary, theoretical studies of the Baldwin effect define two components—the genetically determined one and the learned one, and track the evolution of the genetically determined component.

#### Fitness Function

The relationship between the value of a phenotypic trait of an individual and the fitness of that individual. The fitness function provides a measure of reproductive success of each specific phenotypic variant of the trait. Fitness can be measured as the number of offspring or the growth rate of a phenotypic variant relative to growth rate of other variants (Metz et al. 1992; Rice 2004). The values of the fitness function need not remain constant in time; the fitness of individuals with a given phenotype may depend on their frequency in the population (frequency-dependent selection). Fitness of particular phenotypes may also change as the environment changes, be it abiotic environment (temperature or humidity) or biotic environment (the phenotypes of resident population in the case of a rare mutant that invades this population).

#### Fitness Landscape

Fitness landscape (a metaphor proposed by Wright 1932) is a geometric representation of fitness function used to visualize evolution. In the simple case where phenotype consists of a single trait, the environment is constant in time, and the values of fitness are not frequency-dependent, then the fitness landscape can be evaluated along a single axis that relates all possible values of this trait to fitness. In a more realistic case all these simplifying conditions may not be met, in particular phenotype consists of many traits; each of these traits is related to fitness by its own fitness function. The relationship between such a multivariate phenotype and its fitness can be visualized as a



*multi-dimensional fitness landscape* (where the number of dimensions equals the number of traits under consideration). The surface of a fitness landscape is said to be *rugged* if it has many local *fitness peaks* of different heights; fitness peaks represent the optimal trait values (or the optimal combinations of trait values) i.e. the trait of higher fitness. If a single trait value has the highest fitness on a univariate fitness landscape (or a single combination of traits on a multivariate fitness landscape) then such a fitness landscape is termed *unimodal* or *single-peak landscape*. Finally, fitness function may also be expressed as a relationship between genotype and its fitness, where the genotype space is given by the set of all possible genotypes (for a review of the development of the notion of fitness landscape see Gavrillets 2004).

### Fitness Landscape in the Baldwin Effect

In the context of the Baldwin effect, evolution is thought to proceed on a two-dimensional fitness landscape that is determined by the relationship between fitness and two traits: the innate trait and the level of adaptive learning. A degree of adaptive learning changes the fitness function for the innate trait as compared to the situation where there is no learning at all, and therefore learning may influence the rate of evolution of the innate trait. However, to date a common approach in the studies of the Baldwin effect is to keep the level of learning fixed and to track the rate of evolution of the innate trait only (see this review). This approach of measuring the rate of evolution of the innate trait at the presence of a non-evolving learning is at odds with the growing evidence that the ability to learn itself has a genetic basis (McGuire and Hirsch 1977; Dukas 2004) and high and low levels of learning can be selected for (Mery and Kawecki 2002). The joint evolution of learning and the innate behaviour is also at the core of the original argument proposed by Baldwin.

**Open Access** This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

### References

- Agrawal, A. A., Laforsch, C., & Tollrian, R. (1999). Transgenerational induction of defences in animals and plants. *Nature*, *401*, 60–63.
- Ancel, L. (1999). A quantitative model of the Simpson-Baldwin effect. *Journal of Theoretical Biology*, *196*, 197–209.
- Ancel, L. (2000). Undermining the Baldwin expediting effect: Does phenotypic plasticity accelerate evolution? *Theoretical Population Biology*, *58*, 307–319.
- Anderson, R. W. (1995). Learning and evolution: A quantitative genetic approach. *Journal of Theoretical Biology*, *175*, 89–101.
- Auld, J. R., Agrawal, A. A., & Relyea, R. A. (2010). Re-evaluating the costs and limits of adaptive phenotypic plasticity. *Proceedings of the Royal Society B*, *277*, 503–511.
- Badyaev, A. V. (2005). Stress-induced variation in evolution: From behavioural plasticity to genetic assimilation. *Proceedings of the Royal Society B*, *272*, 877–886.
- Badyaev, A. V. (2009). Evolutionary significance of phenotypic accommodation in the novel environments: An empirical test of the Baldwin effect. *Philosophical Transactions of the Royal Society B*, *364*, 1125–1141.
- Badyaev, A. V., & Uller, T. (2009). Parental effects in ecology and evolution: mechanism, processes and implications. *Philosophical Transactions of the Royal Society B*, *364*, 1169–1177.
- Baldwin, J. M. (1896). A new factor in evolution. *American Naturalist*, *30*, 441–451.
- Baldwin, J. M. (1902). *Development and evolution*. New York: MacMillan & Co.
- Behmer, S. T., Belt, C. E., & Shapiro, M. S. (2005). Variable rewards and discrimination ability in an insect herbivore: What and how does a hungry locust learn? *The Journal of Experimental Biology*, *208*, 3463–3473.
- Borenstein, E., Meilijson, I., & Rupin, E. (2006). The effect of phenotypic plasticity on evolution in multi-peaked fitness landscapes. *Journal of Evolutionary Biology*, *19*, 1555–1570.
- Bosssdorf, O., Richards, C. L., & Piglucci, M. (2008). Epigenetics for ecologists. *Ecology Letters*, *11*, 106–115.
- Cahill, L., Mc Gaugh, J. L., & Weinberger, N. M. (2001). The neurobiology of learning and memory: Some reminders to remember. *Trends in Neurosciences*, *24*, 578–581.
- Crispo, E. (2007). The Baldwin effect and genetic assimilation: Revisiting two mechanisms of evolutionary change mediated by phenotypic plasticity. *Evolution*, *61*, 2469–2479.
- Crispo, E. (2008). Modifying effects of phenotypic plasticity on interactions among natural selection, adaptation and gene flow. *Journal of Evolutionary Biology*, *21*, 1460–1469.
- Day, T., & Bonduriansky, R. (2011). A unified approach to the evolutionary consequences of genetic and non-genetic inheritance. *The American Naturalist*, *178*, E18–E36.
- Day, R. L., Laland, K. N., & Odling-Smee, F. J. (2003). The niche-construction perspective. *Perspectives in Biology and Medicine*, *46*, 80–95.
- De Jong, G. (2005). Evolution of phenotypic plasticity: Patterns of plasticity and the emergence of ecotypes. *New Phytologist*, *166*, 101–118.
- Dopazo, H., Gordon, M. B., Perazzo, R., & Risau-Gusman, S. (2001). A model for the interaction of learning and evolution. *Bulletin of Mathematical Biology*, *63*, 117–134.
- Dukas, R. (2004). Evolutionary biology of animal cognition. *Annual Review of Ecology, Evolution, and Systematics*, *35*, 347–374.
- Dukas, R., & Bernays, E. A. (2000). Learning improves growth rates in grasshoppers. *Proceedings of the National Academy of Sciences*, *97*, 2637–2640.
- Egas, M., & Sabelis, M. W. (2001). Adaptive learning of host preference in a herbivorous arthropod. *Ecology Letters*, *4*, 190–195.
- Fontanari, J., & Meir, R. (1990). The effect of learning on the evolution of asexual populations. *Complex Systems*, *4*, 401–414.
- Garland, T., Jr., & Kelly, S. A. (2006). Phenotypic plasticity and experimental evolution. *The Journal of Experimental Biology*, *209*, 2344–2361.
- Gavrillets, S. (2004). *Fitness landscapes and the origin of species*. Princeton, USA: Princeton University Press.
- Hinton, G. E., & Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, *1*, 495–502.
- Hoetjes, K. M., Kruidhof, M. H., Huigens, M. E., Dicke, M., Vet, L. E. M., & Smid, H. M. (2011). Natural variation in learning rate

- and memory dynamics in parasitoid wasps: Opportunities for converging ecology and neuroscience. *Proceedings of the Royal Society B*, 278, 889–897.
- Kawecki, T. (2010). Evolutionary ecology of learning: Insights from fruit flies. *Population Ecology*, 52, 15–25.
- Laland, K. N., & Sterelny, K. (2006). Perspective: Seven reasons (not) to neglect niche construction. *Evolution*, 60, 1751–1762.
- Lande, R. (2009). Adaption to an extraordinary environment by evolution of phenotypic plasticity and genetic assimilation. *Journal of Evolutionary Biology*, 22, 1435–1446.
- Mayley, G. (1997). Guiding or hiding: Explorations into the effects of learning on the rate of evolution. In P. Husbands & I. Harvey (Eds.), *Proceedings of the fourth European conference on artificial life*. Cambridge, MA: MIT Press.
- McGuigan, K., Nishimura, N., Currey, M., Hurwit, D., & Cresko, W. A. (2011). Cryptic genetic variation and body size evolution in threespice stickleback. *Evolution*, 65, 1203–1211.
- McGuire, T. R., & Hirsch, J. (1977). Behavior-genetic analysis of *Phormia regina*: Conditioning, reliable individual differences, and selection. *Proceedings of the National Academy of Sciences*, 74, 5193–5197.
- Mery, F., & Kawecki, T. J. (2002). Experimental evolution of learning ability in fruit flies. *Proceedings of the National Academy of Sciences*, 99, 14274–14279.
- Mery, F., & Kawecki, T. J. (2003). A fitness cost of learning ability in *Drosophila melanogaster*. *Proceedings of the Royal Society B*, 270, 2465–2469.
- Mery, F., & Kawecki, T. J. (2004). The effect of learning on experimental evolution of resource preference in *Drosophila melanogaster*. *Evolution*, 59, 757–767.
- Metz, J. A. J., Nisbet R. M., & Geritz, S. A. H. (1992). How should we define “fitness” for general ecological scenarion? *Trends in Evolution and Ecology*, 7, 198–202.
- Nomikou, M., Janssen, A., & Sabelis, M. W. (2003). Herbivore host plant selection: Whitefly learns to avoid host plants that harbour predators of her offspring. *Oecologia*, 136, 484–488.
- Paenke, I., Sendhoff, B., & Kawecki, T. J. (2007). Influence of plasticity and learning on evolution under directional selection. *The American Naturalist*, 170, 47–58.
- Papaj, D. R. (1994). Optimizing learning and its effect on evolutionary change in behavior. In L. A. Real (Ed.), *Behavioral mechanisms in evolutionary ecology* (pp. 133–153). Chicago: University of Chicago Press.
- Payne, R. B., Payne, L., Woods, J. L., & Sorenson, M. D. (2000). Imprinting and the origin of parasite-host species associations in brood-parasitic indigobirds, *Vidua chalybeata*. *Animal Behaviour*, 59, 69–81.
- Price, T. D., Qvarnstrom, A., & Irwin, D. E. (2003). The role of phenotypic plasticity in driving genetic evolution. *Proceedings of the Royal Society B*, 270, 1433–1440.
- Rice, S. H. (2004). *Evolutionary theory. Mathematical and conceptual Foundations*. Sunderland, Massachusetts: Sinauer Associates, Inc. Publishers.
- Samuels, R. (2004). Innateness in cognitive sciences. *Trends in Cognitive Sciences*, 8, 136–141.
- Scheiner, S. M. (1993). Genetics and evolution of phenotypic plasticity. *Annual Reviews in Ecology and Systematics*, 24, 35–68.
- Simpson, G. G. (1953). The Baldwin effect. *Evolution*, 7, 110–117.
- Stephens, D. W. (1991). Change, regularity and value in the evolution of animal learning. *Behavioral Ecology*, 2, 77–89.
- Suzuki, R., & Arita, T. (2004). Interactions between learning and evolution: the outstanding strategy generated by the Baldwin effect. *Biosystems*, 77, 57–71.
- Sznajder, B., Sabelis, M. W., & Egas, M. (2011). The interplay between genetic and learned components of behavioral traits. *Journal of Plant Interactions*, 6, 77–80.
- Takemoto, H., Kainoh, Y., & Takabayashi, J. (2011). Learning of plant volatiles by aphid parasitoids: Timing to learn. *Journal of Plant Interactions*, 6, 137–140.
- Ten Cate, C. (2000). How learning mechanism might affect evolutionary process. *Trends in Ecology & Evolution*, 15, 179–180.
- Via, S., Gomulkiewicz, R., De Jong, G., Scheiner, S. M., Schlichting, C. D., & Van Tienderen, P. H. (1995). Adaptive phenotypic plasticity: Consensus and controversy. *Trends in Ecology & Evolution*, 10, 212–217.
- Wang, Q., Gu, H., & Dorn, S. (2003). Selection on olfactory response to semiochemicals from a plant-host complex in a parasitic wasp. *Heredity*, 91, 430–435.
- Watson, J., Wiles, J. (2002). The rise and fall of learning: A neural network model of the genetic assimilation of acquired traits. In: X. Yao (Ed.), *Proceedings of the 2002 congress on evolutionary computation* (pp. 600–605).
- West-Eberhard, M. J. (2005). Developmental plasticity and the origin of species differences. *Proceedings of the National Academy of Sciences USA*, 102, 6543–6549.
- Wright, S. (1932). The role of mutation, inbreeding, crossbreeding and selection in evolution. In D. F. Jones (Ed.), *Proceedings of the VI international congress on genetics* (Vol. I, pp. 356–366). Austin, Texas: University of Texas.
- Youngson, N. A., & Whitelaw, E. (2008). Transgenerational epigenetic effects. *Annual Review of Genomics and Human Genetics*, 9, 233–257.