

The effect of phenotypic plasticity on evolution in multi-peaked fitness landscapes

E. BORENSTEIN,* I. MEILIJSON† & E. RUPPIN*,‡

*School of Computer Science, Tel-Aviv University, Tel-Aviv, Israel

†School of Mathematical Sciences, Tel-Aviv University, Tel-Aviv, Israel

‡School of Medicine, Tel-Aviv University, Tel-Aviv, Israel

Keywords:

drawdown;
 evolution;
 evolutionary convergence rate;
 multi-peaked fitness landscape;
 phenotypic plasticity;
 random walk;
 ruggedness.

Abstract

When facing the challenge of developing an individual that best fits its environment, nature demonstrates an interesting combination of two fundamentally different adaptive mechanisms: genetic evolution and phenotypic plasticity. Following numerous computational models, it has become the accepted wisdom that lifetime acclimation (e.g. via learning) smooths the fitness landscape and consequently accelerates evolution. However, analytical studies, focusing on the effect of phenotypic plasticity on evolution in simple unimodal landscapes, have often found that learning hinders the evolutionary process rather than accelerating it. Here, we provide a general framework for studying the effect of plasticity on evolution in multi-peaked landscapes and introduce a rigorous mathematical analysis of these dynamics. We show that the convergence rate of the evolutionary process in a given arbitrary one-dimensional fitness landscape is dominated by the largest descent (drawdown) in the landscape and provide numerical evidence to support an analogous dominance also in multidimensional landscapes. We consider several schemes of phenotypic plasticity and examine their effect on the landscape drawdown, identifying the conditions under which phenotypic plasticity is advantageous. The lack of such a drawdown in unimodal landscapes vs. its dominance in multi-peaked landscapes accounts for the seemingly contradictory findings of previous studies.

Introduction

Although traits acquired during the lifetime of an organism are not directly inherited in a pure Darwinian framework, they may change the individual's fitness and consequently dramatically alter the dynamics of the evolutionary process (e.g. via genetic assimilation of initially acquired traits through the Baldwin effect, Baldwin, 1896; Morgan, 1896). The interplay between evolution and phenotypic plasticity (e.g. lifetime learning, developmental plasticity, etc.) is thus far from being trivial and has been the subject of numerous biological (Waddington, 1942, 1953; Mery & Kawecki, 2002, 2004)

and computational (Hinton & Nowlan, 1987; Maynard-Smith, 1987; Belew, 1990; Gruau & Whitley, 1993; French & Messinger, 1994; Menczer & Belew, 1994; Mayley, 1996; Moriarty & Mikkulainen, 1996; Parisi & Nolfi, 1996; Floreano & Urzelai, 1998; Ancel, 1999; Nolfi & Floreano, 1999; Weber & Depew, 2003) studies. Clearly, the capacity of a phenotype to better adjust to its environment has an advantageous effect in a non-stationary environment, allowing individuals to acclimate to rapid changes that cannot be tracked by the slow evolutionary process (Littman & Ackley, 1991; Todd & Miller, 1991; Nolfi & Parisi, 1997). However, it has been argued that phenotypic plasticity may also be beneficial in static (or slowly changing) environments (Hinton & Nowlan, 1987; Nolfi & Floreano, 1999), facilitating the evolutionary search by smoothing the fitness landscape.

In recent years, a number of researchers have studied the complex interaction between phenotypic plasticity

Correspondence: Elhanan Borenstein, School of Computer Science, Tel-Aviv University, Tel Aviv 69978, Israel.
 Tel.: +972 3 640 5378; fax: +972 3 640 9357;
 e-mail: borens@post.tau.ac.il

(and specifically, learning) and evolution, employing a variety of methodologies. Comprehensive theories that can combine the two paradigms of evolution and phenotypic plasticity have been recently constructed (Schlichting & Pigliucci, 1998; West-Eberhard, 2003), demonstrating the importance of development and phenotypic response to environmental stimuli in evolutionary theory. Here, however, we focus on one specific question concerning adaptive evolution of phenotypes, examining the effect of plasticity on the convergence rate of the evolutionary process. In the seminal work of Hinton & Nowlan (1987), a simple computational model was introduced to demonstrate how learning can guide and accelerate evolution (see also Maynard-Smith, 1987). Despite its obvious limitations (Nolfi & Floreano, 1999), Hinton and Nowlan's model had successfully demonstrated a distilled model of this effect, bringing the interaction between learning and evolution back to the forefront of scientific research. A large body of work that followed Hinton and Nowlan's study (Belew, 1990; Littman & Ackley, 1991; Gruau & Whitley, 1993; French & Messinger, 1994; Menczer & Belew, 1994; Littman, 1996; Nolfi & Floreano, 1999; Dopazo *et al.*, 2001) further explored the beneficial effect of learning on evolution. Using various simulations of the evolutionary process, these studies demonstrated the benefit of combining learning and evolution in a wide range of stationary and non-stationary environments. Specifically, it has become the accepted wisdom that lifetime learning accelerates evolution in stationary environments by smoothing the fitness landscape and setting up favourable selection preferences for those individual whose genotypic configurations are in the vicinity of the optimal genotype.

However, even with this ever-growing body of evidence for the advantageous effect of phenotypic plasticity on evolution, rigorous theoretical analysis of this interaction is still scarce. Moreover, such analyses have often found that learning hinders evolution, leading to contradictory predictions (Mery & Kawecki, 2004): Fontanari & Meir (1990) performed a quantitative analysis of an asexual version of the Hinton and Nowlan's model, based on a classic population genetic approach. Corroborating the claims made by Hinton and Nowlan, they showed that learning contributes to the robustness of

the evolutionary process against high mutation rates. Studying a more general selection scenario and considering a one-dimensional Gaussian fitness function, Anderson (1995) found that while learning does have an obvious beneficial effect in changing environments, the advantage of learning in a fixed environment is transient. Representing lifetime acclimation as an increase in the variance of selection and using quantitative genetic models, he showed that learning actually slows the final convergence of the population to a maximal fitness solution. Ancel (2000) further demonstrated that when an extreme fitness scenario is not assumed, phenotypic plasticity does not universally accelerate evolution. Ancel's findings suggest that the Baldwin expediting effect (the term she used for this beneficial effect of learning) may thus not be sufficient to account for the evolutionary success of learning.

The findings of these analytical studies clearly disagree with the beneficial effect of phenotypic plasticity that has been demonstrated in the simulation studies cited above, leading to a long-standing debate. We believe that the source of the discrepancy lies in the structure of the fitness landscapes analysed. While most of the simulation studies explored relatively complex artificial environments, such that induce highly irregular fitness landscapes, the mathematical analyses have employed common population dynamics models, focusing on unimodal landscapes (Table 1). These relatively simple landscapes lack one of the key characteristics influencing the convergence rate of the evolutionary process – multiple local optima. The existence of multiple optima (and consequently, multiple domains of attraction) significantly slows down the evolutionary process and hence may make the effect of phenotypic plasticity (or any other mechanism that smooths the landscape) more important. Furthermore, complex genotype–phenotype mapping, developmental processes, epistasis, multiobjective optimization and frequency dependent selection may all render multiple optima genetic solutions (Wright, 1932), making such multipeaked landscapes a feasible model for biological landscapes and the subject of numerical (Kauffman & Levin, 1987) and experimental (Macken & Perelson, 1989; Korona *et al.*, 1994; Burch & Chao, 1999; Lenski *et al.*, 1999) studies. One noticeable

Table 1 Findings concerning the effect of phenotypic plasticity on evolution.

| Fitness landscape structure | Simulations results | Analytical results |
|------------------------------------|---|---|
| Extreme (one optimal phenotype) | accelerates evolution [1] | accelerates evolution [2] |
| Unimodal | slows evolution [3] | advantage is transient [4] slows evolution [5] |
| Multipeaked (complex environments) | accelerates evolution [6–10] improves evolving solution [8–10] | N/A |

[1], Hinton & Nowlan (1987); [2], Fontanari & Meir (1990); [3], Dopazo *et al.* (2001); [4], Anderson (1995); [5], Ancel (2000); [6], Littman & Ackley (1991); [7], Nolfi & Floreano (1999); [8], Gruau & Whitley (1993); [9], French & Messinger, 1994; [10], Littman (1996).

evidence for this characteristic of biological fitness landscapes is demonstrated by a recent study of laboratory evolution of *Escherichia coli* (Fong *et al.*, 2005). Using parallel, replicate adaptive evolution experiments and examining the evolution endpoints, it was shown that the fitness landscape includes distinct peaks of increased adaptive fitness. As shown by Table 1, a rigorous analysis of the effect of phenotypic plasticity on evolution in such landscapes is lacking.

To fill this gap, we focus on analysing the effects of phenotypic plasticity on evolution in arbitrary multi-peaked landscapes, characterizing the plasticity schemes and conditions under which phenotypic plasticity has an advantageous effect and identify its origins. The remainder of this paper is organized as follows. First, we use Random Walk (RW) theory to analyse a simple model of evolution and derive a rigorously quantitative measure of evolutionary rate on any given multi-peaked landscape. We then turn to examine the effect of plasticity on evolution: We first introduce the concept of innate vs. effective fitness landscapes, representing the effect of plasticity as a transformation, replacing the 'innate' fitness landscape that governs selection when no plasticity is present with an alternative 'effective', fitness landscape. Using this concept and the derived RW measure, the effect of phenotypic plasticity on evolution can be quantified by comparing the evolutionary convergence rate using the innate vs. the effective fitness functions for selection. We study various plasticity schemes by examining the effective fitness landscapes they induce and their effect on the evolutionary rate. In particular, we examine both deterministic and stochastic models of learning as well as a simple model of random phenotypic variation, and investigate the influence of varying plasticity rate. The paper concludes with a discussion of the implications of our findings and future work.

Mathematical analysis of evolutionary rate in multi-peaked landscapes

One-dimensional arbitrary multi-peaked landscapes

Analysing the dynamics of an evolutionary search on a given landscape has attracted considerable attention in recent years (Kallel *et al.*, 2001). Most efforts focused on studying the geometric properties of fitness landscapes, including multimodality (Goldberg, 1989), autocorrelation (Weinberger, 1990), and neutrality (Huynen *et al.*, 1996) and on strictly uphill adaptive walk dynamics (Kauffman & Levin, 1987), in an attempt to predict the difficulty of the search task (Stadler, 1995). Here, we provide a direct estimate of the time it will take a stochastic evolutionary process to reach the global optimum on an arbitrary landscape. To obtain a rigorous mathematical analysis of the evolutionary process dynamics we employ a canonical, one-dimensional model of asexual evolution in a fixed arbitrary environ-

ment. Each genotype is encoded by single integer value x , whose fitness value is given by $F(x)$. We assume that the genetic configuration in the first generation is 0 and let N denote the location of the global optimum. Evolution is represented as a simple random walk (RW) process wherein the probabilities p_i (taking a +1 step) and $q_i = 1 - p_i$ (taking a -1 step) for each location i are determined according to the differences between the fitness value of i and those of its neighbouring genetic configurations (see also Appendix S2). It should be noted that in contrast to genetic drift processes that are often modelled by simple *symmetric* random walks (i.e. $p_i = q_i = 1/2$ for every i), here, the +1 and -1 step probabilities depend on the fitness landscape structure and hence, p_i and q_i are not necessarily equal and may also vary for different i values. Consequently, we use a *nonsymmetric* random walk model (Spitzer, 2001), allowing us to represent also non-neutral selection schemes. The term *random* thus refers to the stochastic nature of the walk process, where in each point of time the step direction is selected at random with certain probabilities. Such nonsymmetric random walk models are commonly used in physics, engineering, economy and finance (Hughes, 1995).

Within this model, the expected first-passage time from 0 to N , E_0^N , can serve as a good measure for the progress rate of the evolutionary process (describing the time to first encounter of the global optimum) and can be explicitly calculated for any given one-dimensional landscape. Formally, consider a simple RW S_t (± 1 increments) in a changing environment on $\{0, 1, 2, \dots, N\}$. Let $p_i = P(S_{t+1} = i + 1 | S_t = i)$ and let $q_i = 1 - p_i = P(S_{t+1} = i - 1 | S_t = i)$. Let also ρ_i denote the odds-ratio $\frac{q_i}{p_i}$. Note that $\rho < 1$ indicates a positive selection pressure, $\rho > 1$ indicates a negative selection pressure, and $\rho = 1$ represents regions wherein $p = q$ and thus no selection pressures are exerted (neutral drift). Focusing on the time it takes evolution to reach the global optimum from an initial genetic configuration, it should be noted that the term *positive selection pressure* is used here in the sense that the +1 mutant, which is closer to the global optimum than the -1 mutant, is also fitter. Similarly, a *negative selection pressure* indicates regions wherein the -1 mutant is more fit than the +1 mutant. Let $p_0 = 1$ and assume that $0 < p_i < 1$ for all $0 < i < N$. As shown in Appendix 1, the expected first-passage time from 0 to N (i.e. the expected time to first hit N starting at 0) on a given landscape is

$$E_0^N = N + 2 \sum_{\substack{i \leq j \\ 0 \leq i, j < N}} \prod_{k=i}^j \rho_k.$$

For a multi-peaked landscape, define the *drawdown* R as the maximal element in this sum:

$$R = \max_{\substack{i \leq j \\ 0 \leq i, j < N}} \prod_{k=i}^j \rho_k.$$

This characterizing feature of the landscape has also been termed in the literature *gap* or *extent* (Noskowitz & Goldhirsch, 1990; Meilijson, 2003). It is shown that the expected first-passage time is dominated by R , whereby E_0^N is sharply bounded from above by $N^2(\frac{1+R}{2})$ (see Appendix 1 for a full analysis). Furthermore, as demonstrated in Appendix 1, this bound is reached when all ρ_i values contributing to the drawdown collapse to a single point in the landscape. Hence, crossing the fitness landscape from the initial genetic configuration toward the global optimum, R corresponds to the height difference between the record high fitness value to the consecutive record low fitness value (i.e. the largest descent along this fitness landscape), as illustrated in Fig. 1.

Multidimensional arbitrary multipeaked landscapes

While the one-dimensional case allows for a rigorous analysis of expected first-passage times, such an analysis is hard to obtain for a multidimensional fitness landscape. We show however that an analogous dominance of the drawdown on evolutionary rate can be inferred through numerical simulations. In this case, we consider all the possible pathways from the initial configuration to the global optimum configuration. Each of these pathways can be conceived as a simple one-dimensional landscape, with a specific drawdown value. We will term the pathway with the minimal drawdown value 'the *Principal-Pathway*' and the drawdown value it induces 'the *Principal-Pathway drawdown*'. We maintain that the Principal-Pathway drawdown dominates the random walk first-passage time on a given multidimensional landscape in a similar manner to that shown in the one-dimensional case (see Appendix S1 for details and discussion).

We validate the strong correlation between the Principal-Pathway drawdown and the first-passage time through numerical simulations, using common multi-

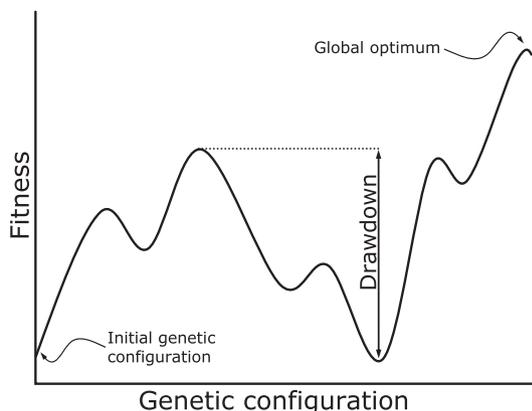


Fig. 1 A qualitative illustration of the fitness drawdown.

modal benchmark functions. These functions are utilized to generate numerous landscapes with one, two and three dimensions and varying drawdown values. The Principal-Pathway drawdown (or simple drawdown value in the one-dimensional case) and the expected first-passage time to the global-optimum, E_0^N , were evaluated for each landscape (see Appendix S1 for details). As demonstrated in Fig. 2, there is a strong correlation between the drawdown value in each fitness landscape and the first-passage time of the random walk process, corroborating the validity of the fitness landscape's drawdown as a measurable bound for the expected evolutionary progress rate.

The effect of phenotypic plasticity on evolution

Innate vs. effective fitness landscapes

Having established a measure of evolutionary rate, we turn to examine the dynamics of the evolutionary process in two modes: in the first, *nonplastic* mode, phenotypic plasticity is absent and the fitness value $F(\vec{x})$ assigned to each genotypic configuration \vec{x} is uniquely determined according to the innate survival and reproduction probability of the phenotype that it encodes, termed *innate fitness*. In the second, *plastic* mode, phenotypes can vary during their lifetime and as a result may effectively gain a different (and potentially higher) fitness value. Clearly, in this mode, selection operates according to the effective fitness value obtained by each individual. We denote this *effective fitness* value, $F_{efc}(\vec{x})$. It should be noted that both the innate and the effective fitness landscapes ultimately correspond to the common notion of fitness in natural selection. We use the terms *innate* and *effective* fitness to simply distinguish between the fitness values that govern selection without and with plasticity. Phenotypic plasticity hence manifests itself as a transformation of the fitness landscape, replacing the innate fitness that initially governed selection with an effective fitness landscape (Fig. 3a). The strength of this simple model lies in the fact that the complex dynamics of a hybrid process combining evolution and phenotypic plasticity can be studied by examining the simpler dynamics of a pure evolutionary process on the appropriate effective fitness landscape. The effect of phenotypic plasticity on the evolutionary convergence rate can be measured by comparing the time it takes the evolutionary process to obtain an optimal genotype using the innate vs. the effective fitness functions for selection. In particular, the random walk analysis and drawdown value presented in the section 'Mathematical analysis of evolutionary rate in multi-peaked landscapes' provide a measurable bound for the expected evolutionary progress rate for *any* given landscape, allowing for a direct quantitative comparison between the innate and the effective fitness landscapes.

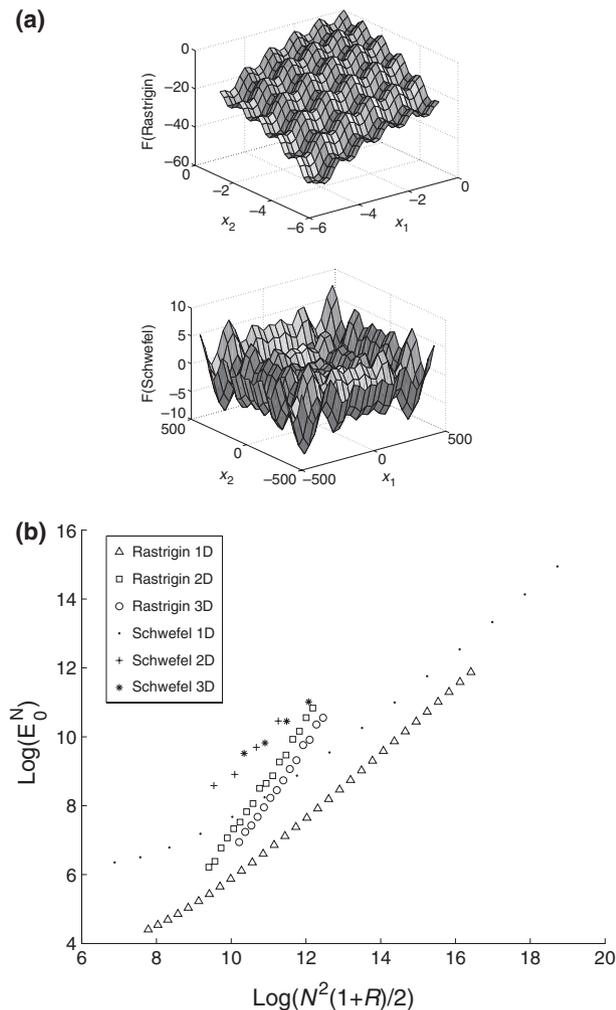


Fig. 2 (a) An example of two-dimensional landscapes generated by the Rastrigin function ($C_r = 3$) and the Schwefel function ($C_s = 0.01$) (see Appendix S1 for more details on these functions). (b) The correlation between the expected first-passage time and the landscape drawdown. Numerous landscapes of different dimensions and varying drawdown were tested. E_0^N was evaluated for each landscape through 100 random walk simulations. Landscapes for which not all simulations hit the global optimum within 250 000 steps were excluded from our analysis. For the one-dimensional landscapes, the expected first-passage time was analytically calculated as described before. The drawdown value is normalized according to the asymptotic bound found for the one-dimensional case, where N denotes the length of the pathway. As the drawdown values (and consequently, the first-passage times) exhibit a large variation, a logarithmic scale was applied. Linear regression analysis for the correlation between R and E_0^N in the Rastrigin landscapes yielded coefficient of determination values (per cent of variance explained) of 0.995, 0.9971 and 0.9961 for the one, two and three-dimensional landscapes respectively. In the Schwefel landscapes the resulting coefficient of determination values were 0.9809, 0.9718 and 0.9818 for the same dimensions.

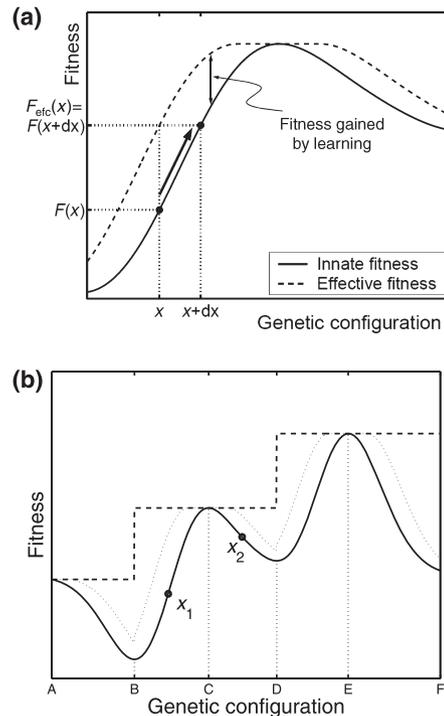


Fig. 3 The effect of deterministic learning on the fitness landscape. (a) An individual with genotype configuration x and innate fitness value $F(x)$ may acclimate by learning (illustrated here as a simple gradient ascent process) and gain a fitness value of $F(x + \Delta x)$. As the genotype of this individual remains unchanged, the effective fitness value $F_{\text{efc}}(x) = F(x + \Delta x)$ is applied to x . (b) The innate fitness function (solid line) and the effective functions obtained with partial learning, i.e. after a limited number of hill-climbing iterations (dotted line), and with ideal learning (dashed line). In the ideal learning scheme all configurations in the basin of attraction of a given local optimum (e.g. genotypes x_1 and x_2 in the interval $[B, D]$) acquire the same effective fitness value, that of the local optimum (C).

In this section we examine several models of phenotypic plasticity, for which the resulting effective fitness landscapes can be explicitly constructed and studied. We start with a simple model of deterministic learning, where phenotypic plasticity is manifested as a hill climbing process (either ideal or partial). Next, we examine a stochastic learning model in which the direction of each learning iteration is determined probabilistically. Finally, although traditionally not defined as a form of phenotypic plasticity, a model of random phenotypic variation is presented, where each genotype randomly realizes one of several alternative phenotypes. We compare the drawdown of the original, innate landscape with that of the resulting effective fitness landscapes induced by these plasticity schemes. Identifying the schemes and conditions under which the fitness landscape drawdown is reduced (and thus, the rate of evolution is accelerated) provides a characterization of the beneficial effect of phenotypic plasticity.

Ideal and partial deterministic learning

Individual learning, as a form of phenotypic plasticity, is modelled as an iterative process of phenotypic modifications aimed at increasing the individual's effective fitness, taking the form of a simple gradient-ascent process in the genotype/phenotype space. As in previous studies (Hinton & Nowlan, 1987; Anderson, 1995), we focus on the simple case of one-to-one mapping from genotype to phenotype and assume that learning and evolution both operate on the same fixed fitness landscape (see also the section 'Discussion'). We examine a simple model of learning where during each learning episode (iteration), an individual compares the innate fitness value of its current configuration with those of slightly modified configurations, and adopts a modified configuration if the latter's innate fitness value is higher (see Fig. 3a for an illustration). Such learning iterations may repeat, allowing the individual to adopt behaviours further away from its innate one, resulting in a modification of its effective fitness accordingly (although its genotype remains unchanged).

We first consider an *ideal deterministic learning* model, where each individual repeatedly employs such deterministic hill-climbing learning iterations until it reaches the nearest local-optimum and no further improvement of its effective fitness is possible. As demonstrated in Fig. 3b, in ideal learning all genetic configurations in the region forming the basin of attraction of a given local optimum will eventually acquire the same effective fitness value, equal to the innate fitness of the local optimum, totally suppressing selection pressures within each such region. In the one-dimensional case, ideal learning transforms each given consecutive pair of descending and ascending intervals in the innate fitness landscape into a single step function in the effective fitness landscape, whose height is equal to the difference between the extents of this descent and the consequent ascent (Fig. 3b). The drawdown characterizing the effective fitness landscape is hence smaller (or equal, in the worst case) than that induced by the original, innate fitness landscape (see also Fig. 4). As demonstrated above, a smaller drawdown value implies a shorter first-passage time, making the beneficial effect of this learning scheme evident.

The mathematical analysis presented above can account for the seemingly contradictory findings of previous studies concerning simple fitness landscapes (Anderson, 1995; Ancel, 2000; Dopazo *et al.*, 2001), in which learning was found to hinder the evolutionary process. While the evolutionary rate in multipeaked landscapes is dominated by the landscape drawdown, in the simple unimodal scenario, no drawdown exists ($R = 1$), and consequently the marked beneficial effect of learning demonstrated above is absent. In terms of our model, in such single peaked landscapes, the evolutionary process is scaled down to a simple random walk

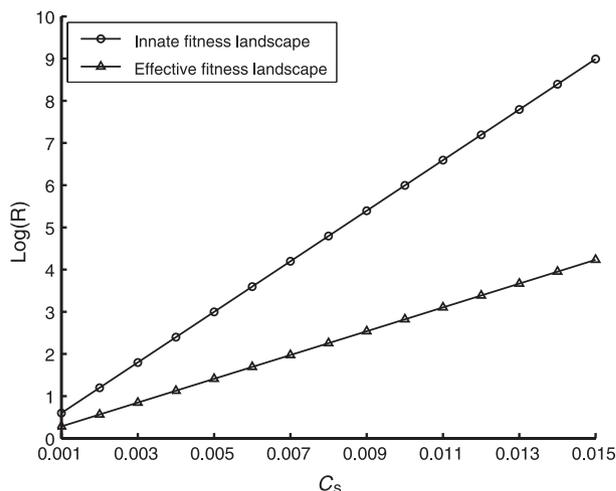


Fig. 4 The effect of ideal deterministic learning on the fitness landscape drawdown. The generalized Schwefel function described in Appendix S1 is used to generate two-dimensional fitness landscapes with varying ruggedness (tuned via the C_s parameter whose higher values denote increased levels of ruggedness). Ideal learning is then applied to produce the corresponding effective fitness landscape. The Principal-Pathway drawdown calculated for each landscape is illustrated. Evidently, the drawdown induced by the effective fitness is significantly smaller (note the logarithmic scale) than that induced by the innate fitness.

within a positive slope interval. As learning decreases the slope of the fitness function (causing $\frac{q}{p}$ to approach unity), our model clearly shows (see Appendix 1, Example 2, concerning a constant environment) that learning would slow down the convergence rate, as Anderson (1995) and Ancel (2000) have indeed found. In particular, learning schemes that cancel all selection pressures and produce a totally flat effective landscape (as is the case in ideal learning) result in a random drift process with quadratic first-passage times, markedly slower than the linear first-passage time in the innate single-peaked landscape (see Appendix 1, Example 1 and Example 2). It is only in a multipeaked landscape (Wright, 1932; Kauffman & Levin, 1987; Kauffman, 1993; Korona *et al.*, 1994; Burch & Chao, 1999; Lenski *et al.*, 1999; Fong *et al.*, 2005), where the overall evolutionary rate is dominated by the exponential passage time in the negative selection regions, that the beneficial effect of ideal learning is demonstrated.

When learning 'resources' are limited (e.g. learning is bounded by a certain cost) and individuals employ only a limited number of hill-climbing iterations, a *partial plastic mode* is obtained rather than ideal learning mode. In this mode, not all genetic configurations in the basin of attraction of each local optimum will inevitably gain the same effective fitness value. Individuals with innate genetic configurations farther from the local optimum configuration do improve through learning (and gain a higher effective fitness value), but may fall short of

reaching the local optimum's exact fitness level. The effective fitness landscape forms an intermediate state between the plastic and non-plastic modes, including both intervals with constant fitness and intervals with positive or negative slopes (Fig. 3b). Clearly, a partial learning scheme still reduces the extent of the innate landscape drawdown (and hence, will still accelerate evolution), though it does not cancel them altogether. It is thus expected that this learning mode will yield an intermediate convergence time, progressing slower than the ideal plastic mode, but still faster than the nonplastic one.

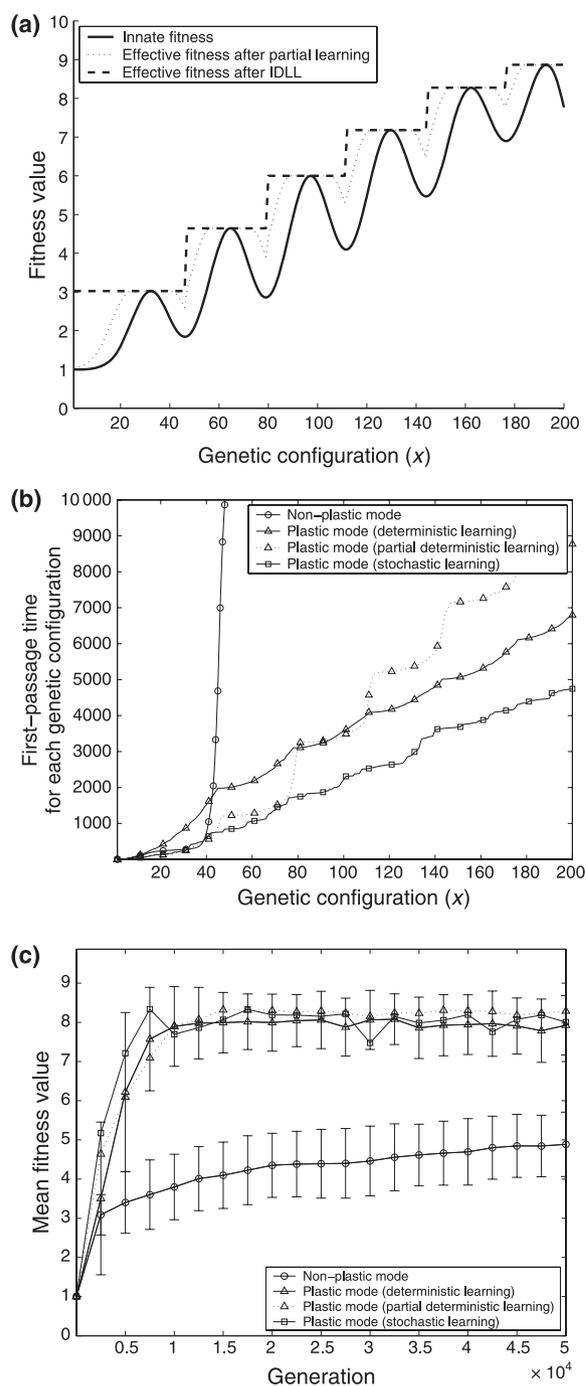
These effects of deterministic learning are validated numerically (see Appendix S2 for the simulation details). The mean first-passage time of each genetic configuration x (i.e. the expected time to first hit x starting at 0) is illustrated in Fig. 5b. The curves clearly agree with the results of our analysis. Figure 5c demonstrates the average **innate** fitness value of the evolving individual as a function of generation of the evolutionary process. Evidently, individuals evolving in the plastic mode converge much faster to the global optimum and gain higher fitness values. Although learning individuals using ideal learning do not converge to the exact global optimum, they successfully reach its basin of attraction and possess higher innate fitness values than nonplastic individuals who tend to get trapped in remote local optima. Examining the effect of partial learning, we find, as expected, that this mode yields an intermediate convergence time, progressing slower than the plastic mode, but still faster than the nonplastic one (Fig. 5b). Furthermore, as this form of learning does not entirely suppress the selection pressures in each optimum domain, it allows individuals that hit the global optimum

Fig. 5 The effect of ideal deterministic learning (IDLL), partial deterministic learning and stochastic learning on the evolutionary process in the one-dimensional case. (a) A one-dimensional innate fitness function was defined on the interval [1; 200] as a sum of several Gaussian functions, yielding a continuous, multi-peaked function $F(x)$ (solid line). Various plasticity schemes were then applied to produce the corresponding effective fitness functions (see Appendix S2 for more details). (b) The first-passage time of each genetic configuration x (i.e. the average time to first hit x). Each curve represents the average result of 100 runs for the deterministic learning stimulation and 10 runs for stochastic learning stimulation. In the plastic mode, using a deterministic learning scheme, all 100 simulation runs hit the global optimum within less than 16 200 generations. In the non-plastic mode, although the linear expected first hitting time in the positive slope intervals yields a fast progress, the exponential behaviour in the negative slope intervals dominates the dynamics of the random walk and hinders the evolutionary process. Out of 100 simulated evolutions, each running for a maximum of 200 000 generations, 16 never hit the third local optimum ($x = 97$) and 67 failed to hit the fourth ($x = 130$). (c) The mean *innate* fitness value as a function of generation. The standard deviation of the non-plastic mode and plastic mode with ideal deterministic learning is also illustrated.

basin of attraction to converge closer to the exact global optimum configuration, resulting in overall better average innate fitness values than those obtained with ideal learning (Fig. 5c).

Stochastic learning

Both the ideal and the partial learning schemes examined above embody two basic characteristics: locality and



accuracy; learning was assumed to exploit only local information about the fitness landscape on which it operates, and to do so with complete accuracy. It is these two features that guarantee the preservation of extrema domains in the effective landscape. However, the lack of complete environmental data, sensory input noise, imperfect information processing and nondeterministic decision making, all make a stochastic learning process more plausible as a model of learning in biology. Yet, as stochastic local search schemes are not bound to take the steepest ascent route, and can potentially discover remote local optima, the effective fitness function they yield may have a different regional structure than that of the original innate fitness. For example, in an extreme scenario, a stochastic learning algorithm may allow any innate genetic configuration to successfully reach the global optimum solution, totally suppressing genetic selection pressures. In this scenario the evolutionary process turns into a random drift, which, as was demonstrated in Appendix 1, yields a quadratic first-passage time.

To study the effects of stochastic learning, we use a simple variation of our model, where the hill-climbing learning algorithm is replaced with a simulated annealing (SA) optimization process (Kirkpatrick *et al.*, 1983) (see Appendix S2 for details). Clearly, stochastic learning does not guarantee a consistent fitness gain each time learning is applied. Consequently, the effective fitness value assigned to each genetic configuration varies from one learning process to the other and a deterministic effective fitness function cannot be explicitly constructed in advance. Yet, examining the characteristics of the *average* effective fitness function constructed by this stochastic scheme (see, for example, Fig. 7b), it is clear that the drawdown induced by the average effective landscape is smaller than that of the original innate landscape.

Numerical simulations of an evolving population applying stochastic learning (where the effective fitness is appropriately evaluated repeatedly for each individual in each generation) validate that this learning scheme indeed accelerates the evolutionary process. Evidently (Fig. 5b,c), also with a stochastic learning paradigm, learning individuals converge faster and gain significantly higher innate fitness values than those evolving in the nonplastic mode, obtaining values similar to those obtained with deterministic learning. Furthermore, it is shown that stochastic learning not only accelerates evolution in comparison to the nonplastic mode, but yields superior evolutionary convergence rates even in comparison to those obtained in the deterministic learning scheme examined above. The superiority of this scheme can be attributed to the resulting effective fitness landscape which is smoother (on average) than the one produced by a deterministic scheme. This can allow individuals near the boundary between basins of attraction to stochastically converge to either of the two adjacent local optima.

It should be noted that in the extreme case, stochastic learning can produce a totally flat effective fitness landscape, suppressing all selection pressures, even in a multi-peaked landscape. However, while the quadratic first-passage time induced by random drift on a flat landscape hinders evolution in comparison with the linear time on a single peaked landscape, it is still superior to the exponential first-passage time expected on a multi-peaked landscape.

Random phenotypic variation

The learning schemes discussed above represent a directed model of plasticity, aiming at increasing the individual effective fitness. However, phenotypic plasticity may also take the form of increased phenotypic variation (or developmental noise) in response to environmental fluctuations (Gavrilets & Hastings, 1994). Although biologists often refer to phenotypic plasticity as a *beneficial* response to the environment (rather than a random variation), here, as in Ancel & Fontana (2000), we wish to examine the effect of a random phenotypic flexibility scheme. This form of plasticity may be more common in molecular evolution, wherein a certain genotype may realize a range of phenotypic configurations according to the microenvironmental context.

Such random perturbations are clearly not necessarily in the direction of improved fitness. However, assuming some correlation between the phenotypic and genotypic spaces, on average, genotypes assigned with low innate fitness values will gain more by these perturbations than genotypes with high innate fitness values. In the extreme case, genotypes located in a local minima of the fitness landscape can only gain higher effective fitness by realizing phenotypes of neighbouring configurations, while genotypes located on local maxima will inevitably gain lower effective fitness. This dynamics, although stochastic, lead to a reduction in the fitness landscape drawdown, and hence, according to our analysis, accelerate evolution. Applying a simple model of genetic variation, where each genotype 'develops' into a phenotype associated with a randomly selected neighbouring genetic configuration within a predefined range, Δd (see Appendix S2 for more details), we validate the beneficial effect of this plasticity scheme (Fig. 6).

A simple example of such a plasticity scheme can be demonstrated in the RNA secondary structure. While the minimum free energy (MFE) secondary structure of an RNA sequence defines a simple mapping from genotypes to phenotypes, in practice, an RNA molecule may fold into a wide range of secondary structure configurations, providing that the energy barriers are sufficiently small. These phenomena can be conceived as the RNA equivalent of phenotypic plasticity or developmental noise (Ancel & Fontana, 2000). Furthermore, as demonstrated by Ancel & Fontana (2000), there is a significant correlation between the repertoire of thermodynamically

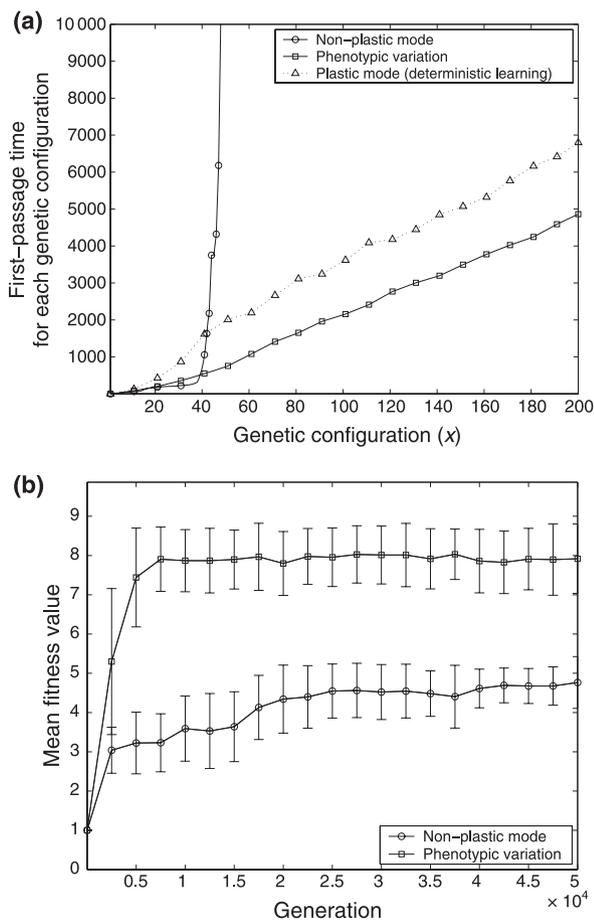


Fig. 6 The effect of random phenotypic variation with $\Delta d = 15$ on the evolutionary process. Each curve represents the average result of 100 simulation runs. (a) Mean first-passage time as a function of the genetic configuration x . The results obtained under the ideal deterministic learning scheme are depicted for comparison. (b) The mean and standard deviation of the innate fitness value as a function of generation.

accessible configurations and genetically accessible configurations (a phenomenon they term *plastogenetic congruence*). This correlation implies that RNA molecules that make such thermodynamic transitions can be described as effectively realizing MFE structures associated with neighbouring genotypic configurations. Interestingly, examining whether phenotypic plasticity expedites the evolutionary discovery of new structures in RNA, Ancel & Fontana (2000) find that no such expediting occurs due to intrinsic properties of the RNA genotype–phenotype map. Specifically, the high neutrality incorporated in the RNA genotype–phenotype map and its organization make the beneficial effect of plasticity restricted to relatively small regions in the genotypic space. Consequentially, the benefit gained by plasticity is negligible compared to the time it takes to discover these regions in the first place.

Varying learning rates

Evidently, different plasticity schemes yield different dynamics of the evolutionary process and result in different convergence rates. The number of phenotype acclimation iterations employed during life or the phenotypic variation range (which will both be referred to here as the *plasticity rate*) may also influence the convergence rate and the stability of the evolutionary process, as was demonstrated by the favourable effects of the partial learning scheme. To further explore and compare the effect of these learning schemes and in particular the effect of varying plasticity rates, an additional set of simulations was carried out. Figure 7a–c illustrate the effective fitness functions constructed by each scenario that was tested. To reduce the long computation time in stochastic learning simulations, the *mean* effective fitness was used as a constant effective fitness landscape, approximating a genuine stochastic learning paradigm.

Two measures were examined for each scenario: the *overall convergence time*, which was taken as the first-passage time of the global optimum ($x = 193$), and the *genetic stability* of the evolving individuals that was measured as the average genetic deviation from the global optimum configuration throughout 1000 generations following the first-passage time. As shown in Fig. 7d, the best convergence time for deterministic learning is obtained with 10 learning iterations. There is also a clear tradeoff between the convergence time and the genetic stability of the resulting evolutionary process. Figure 7e illustrates the results for stochastic learning schemes with varying number of SA iterations. Evidently, a low number of stochastic learning iterations results in faster convergence rates than those obtained with deterministic schemes and still yields relatively stable genetic solutions. Only when the learning process employs a considerably large number of stochastic iterations it diverges from the original structure of the innate fitness function (see Fig. 7b), reducing dramatically the evolutionary selection pressures and consequently reducing the genetic stability of the evolving individuals. A clear tradeoff between the convergence time and the genetic stability is also demonstrated in the phenotypic variation experiments (Fig. 7f). Applying a large variation (e.g. $\Delta d = 15$) results in a fast convergence rate, comparable with that obtained by deterministic learning, but dramatically reduces the genetic stability.

Numerical extensions of the random walk model

While the Random Walk model presented above allows for a rigorous analysis of first passage times, there are a few extensions that make this model more biologically plausible. In particular, we wish to examine whether the possibility to stay in the same genetic configuration over

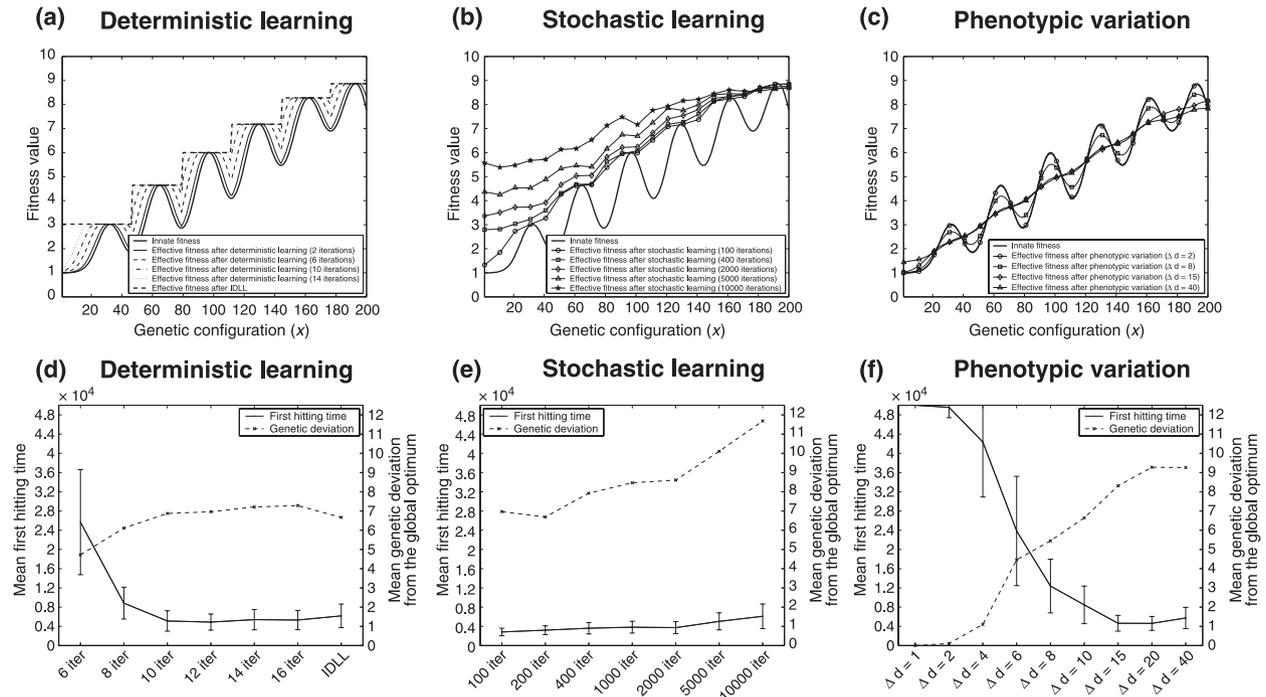


Fig. 7 The effect of various plasticity schemes and varying plasticity rate on evolution. (a–c) The average effective fitness functions resulting by a varying numbers of deterministic and stochastic learning iterations and by varying phenotypic variation range. Employing more than 14 deterministic learning iterations results with an effective fitness similar to the one obtained by ideal learning. The curves illustrated in (c) represent the average of 50 000 runs. (d–f) The average convergence rate, measured as the mean first-passage time of the global optimum (solid line) and genetic stability (dashed line) obtained for varying plasticity rates. Most simulation runs using less than six deterministic learning iterations did not converge to the global optimum.

several generations and the transformation from the fitness landscape values to fixation probabilities affect the resulting dynamics. The analytical treatment of the behaviour of the RW model with these extensions turns to be a difficult challenge, and hence these extensions are examined numerically in this section.

Random walk with static periods

As the RW model presented above assumes that in each point in time the walk process takes either a +1 or a –1 step, our measure of first passage time can be interpreted as measuring the number of mutation events required to reach the global optima. However, in practice, the population may stay in the same genetic configuration for many generations. Hence, if we wish to measure the convergence time in number of generations, we should also allow the walk process to stay in the same configuration with some probability.

As demonstrated in Fig. 8a, introducing this extension to the model (see Appendix S2 for the simulation details) does not significantly change the resulting dynamics, and in particular, preserves the superiority of the plastic mode vs. the non-plastic mode. Apparently, although the probability to stay in the same genetic configuration is

higher in neutral regions of the landscapes (which are abundant in the landscape induced by ideal plasticity) than in positive selection regions, the overall convergence time is still dominated by the slow exponential time it takes to cross negative selection regions (and, particularly, the fitness drawdown). The main effect of this extension is manifested by an overall, relatively constant delay in the time to reach each genetic configuration in comparison to the basic model.

Random walk with Kimura’s fixation probabilities

To further relax some of the differences between our model and traditional population biology models, we apply Kimura’s theory for population dynamics on neutral (or nearly neutral) landscapes (Kimura, 1983). Specifically, we use Kimura fixation probabilities rather than the Boltzmann scaling (see also Appendix S2) to determine the RW probabilities. Under this extension, our RW model is closely related to common evolutionary dynamics models on nearly neutral landscapes, taking also the population size into account. Moreover, assuming that the mutation rate is slow (and hence, each mutation becomes either fixed or extinct before the next mutation arises), the first passage time measure, E_0^N ,

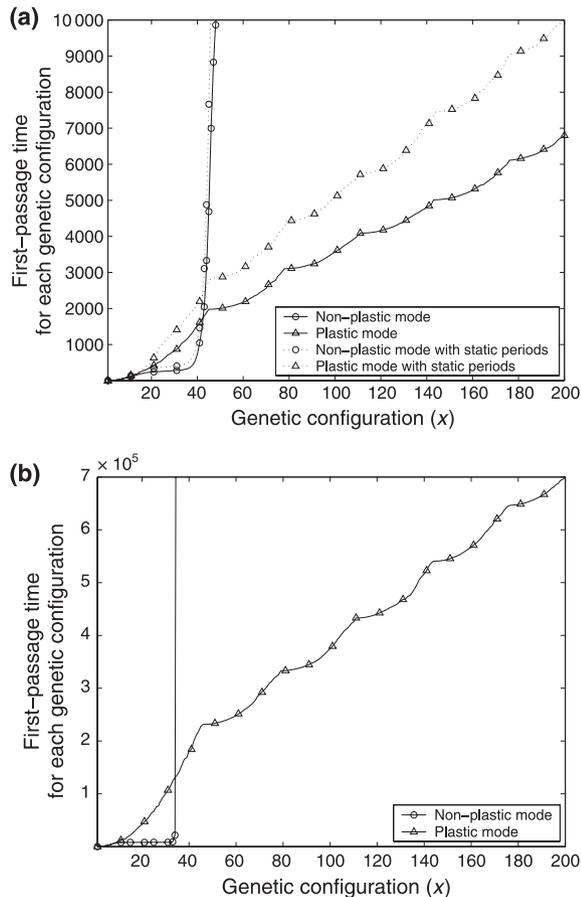


Fig. 8 Numerical extensions of the random-walk model. (a) First-passage times for each genetic configuration using a random walk model with static periods. (b) First-passage times for each genetic configuration using Kimura's fixation probabilities.

indicates the number of mutations that are required to *appear* until the population fixates on the landscape's global optimum and can be now easily translated to the expected number of generations simply by multiplying E_0^N by the mutation rate v . As in the previous section, this extension allows for the population to stay in the same genetic configuration over several generations if none of the neighbouring mutations is fixate.

As can be seen in Fig. 8b, using Kimura's probabilities dramatically slows down the entire process (note the first-passage time scale), mainly due to the significantly smaller probabilities of a neutral ($p = 1/2N_e$) or deleterious mutation to fixate. Clearly, in this model, fixation of a slightly beneficial mutation is markedly faster than that of a neutral mutation. Yet, as in our previous analysis, the probability of fixation of a slightly deleterious mutation becomes infinitesimally (exponentially) small as the population size increases and hence still dominates the overall convergence time, resulting in the superiority of the plastic mode over the nonplastic one.

Discussion

Our study focuses on the effects of phenotypic plasticity on the evolutionary convergence rate in stationary environments. We use random walk theory to derive a measure for the rate of evolution on arbitrary multi-peaked fitness landscapes, and demonstrate that the convergence rate is dominated by the landscape drawdown. Examining various phenotypic plasticity schemes we find that these schemes decrease the landscape drawdown and hence, accelerate evolution. These findings introduce a rigorous quantitative confirmation of the common hypothesis stating that phenotypic plasticity expedites evolution by smoothing the fitness landscape and identify the origins of this phenomenon. Our analysis provides a measure for both the convergence rate bottleneck induced by the landscape drawdown, and the benefit gained by smoothing the landscape and reducing the extent of the drawdown.

Our findings suggest two fundamental principles that affect the interaction between phenotypic plasticity and evolution. First, as the benefit of plasticity originates in its capacity to smooth the fitness landscape, this effect will be revealed only in multi-peaked landscapes wherein the evolutionary rate is dominated by the landscape drawdown. Conversely, if selection takes place only within the domain of a simple fitness function that does not include multiple local optima, plasticity hinders the evolutionary process as was also demonstrated in previous studies (Anderson, 1995; Ancel, 2000). Secondly, our analysis suggests that plasticity has a beneficial effect on evolution when genotypes with low innate fitness values (e.g. individuals at local fitness minima) gain more through phenotypic plasticity than genotypes with high innate fitness values (e.g. with local maxima configurations). These dynamics are governed mainly by the correlation between the genotypic and phenotypic spaces (determined by the genotype–phenotype mapping). As demonstrate in the section 'Random phenotypic variation', when such a correlation exists, phenotypic modifications due to plasticity are analogous to perturbations in the genetic space, and consequently plasticity yields, on average, higher gain for those individuals whose genotypes are located on local fitness minima.

Mayley (1997) discussed the effects of learning on evolution in a rugged landscape, arguing that such settings may give rise to two competing effects: a *guiding* effect, helping evolution to detect individuals located near superior local optima, and a *hiding* effect where learning suppresses the selection pressures within each local optima basin of attraction. Our analysis demonstrates that in this scenario, the guiding effect outweighs the hiding effect, resulting with an overall acceleration of the evolutionary process.

Clearly, the fixed landscape, the one-to-one genotype–phenotype mapping and the specific phenotypic plasticity

schemes examined in this study are a simplification of the dynamics that take place in natural systems. Living organisms incorporate a complex developmental process that may disturb the correlation between the phenotype and genotype spaces (Downing, 2004) and apply diverse and sophisticated plasticity methods. Associating plasticity costs may also affect the resulting fitness gain (Mayley, 1996). Moreover, the random walk model applied in this study restricts evolution to ± 1 increments, corresponding to small mutations. However, real mutations can come in a wide variety of increments, markedly influencing the evolutionary trajectory. Specifically, large mutations can help the evolutionary process to cross fitness valley barriers, reducing (though, most probably, not cancelling altogether) the deleterious effect of the fitness landscape drawdown. The effect of such mutations on the evolutionary dynamics in innate vs. effective fitness landscapes is of much interest and may vary with the exact structure of the landscape. Unfortunately, characterizing the structure of fitness landscapes of biological systems and the dynamics of biological plasticity is still an open question. We have thus focused on a simple mathematical model, with fully correlated landscapes, allowing us to explicitly construct and examine the resulting effective fitness function and to derive a rigorous, quantitative analysis of the expected convergence rate. However, the approach presented in the paper and the mathematical analysis of first-passage times can be utilized to examine the effects of additional plasticity paradigms and landscapes structures, as long as the effective fitness landscapes can be evaluated.

Future research may be able to predict scenarios in which phenotypic plasticity will be favoured by evolution. In particular, while the exact structure of the fitness landscape is usually hard to characterize, there are cases where multi-peaked landscapes are expected. For example, whenever the evolutionary process is required to optimize multiple objectives with a certain trade off function, multiple local optima usually exist (e.g. Oksanen & Lundberg, 1995). In such cases, phenotypic plasticity, being a mechanism that expedites the discovery rate of new optima, is valuable. We thus believe that increased phenotypic variation should be correlated with the existence and abundance of multiple optima.

An example of a biological experiment designed to directly examine the effect of learning on evolution has recently been presented (Mery & Kawecki, 2004). In this exciting study, populations of *Drosophila melanogaster* were exposed to various selection regimes concerning preference for oviposition substrate with and without the ability to use aversion learning. The results of this study showed that learning may evolve even in a stationary environment (in contrast to the common argument that learning should be favoured only in a changing environment). However, examining whether learning ability

affects the evolution of the innate component, they found that learning facilitated evolution in one direction of selection while hindering it in the other, leaving the controversy concerning the effect of learning on evolution unresolved.

Evidently, the dynamics governing the evolution of plastic individuals is still far from being completely characterized or understood. The variety of phenotypic plasticity mechanisms found in biological systems and their complexity render a challenging task. However, quantitative analysis of complex phenomena often requires simple models. While acknowledging its limitations, we believe that the framework presented in this paper can serve as a theoretical basis for future studies of issues concerning the interplay between phenotypic plasticity and evolution.

Acknowledgments

We are grateful to Marc Feldman, Lilach Hadany, Aaron Hirsh, Alon Keinan, Jeremy Kendal, Joanna Masel and Ben Sandbank for stimulating discussions and comments. We thank Ron Meir and Daniel Polani for reading earlier versions of this paper. E.B. is supported by the Yeshaya Horowitz Association through the Center of Complexity Science. E.R.'s research is supported by the Tauber Fund and the Israeli Science Foundation (ISF).

References

- Ancel, L. 1999. A quantitative model of the Simpson-Baldwin effect. *J. Theor. Biol.* **196**: 197–209.
- Ancel, L. 2000. Undermining the Baldwin expediting effect: does phenotypic plasticity accelerate evolution? *Theor. Popul. Biol.* **58**: 207–319.
- Ancel, L. & Fontana, W. 2000. Plasticity, evolvability, and modularity in RNA. *J. Exp. Zool.* **288**: 242–283.
- Anderson, R. 1995. Learning and evolution: a quantitative genetics approach. *J. Theor. Biol.* **175**: 89–101.
- Baldwin, J. 1896. A new factor in evolution. *Am. Nat.* **30**: 441–451.
- Belew, R. 1990. Evolution, learning, and culture: computational metaphors for adaptive algorithms. *Complex Systems* **4**: 11–49.
- Burch, C. L. & Chao, L. 1999. Evolution by small steps and rugged landscapes in the RNA virus $\phi 6$. *Genetics* **151**: 921–927.
- Dopazo, H., Gordon, M., Perazzo, R. & Risau-Gusman, S. 2001. A model for the interaction of learning and evolution. *Bull. Math. Biol.* **63**: 117–134.
- Downing, K. 2004. Development and the Baldwin effect. *Artif. Life* **10**: 39–63.
- Floreano, D. & Urzelai, J. 1998. Evolution and learning in autonomous robotic agents. In: *Bio-inspired Computing Systems* (D. Mange & M. Tomassini, eds), pp. 1–36. PPUR, Lausanne.
- Fong, S., Joyce, A. & Palsson, B. 2005. Parallel adaptive evolution cultures of *Escherichia coli* lead to convergent growth phenotypes with different gene expression states. *Genome Research* **15**: 1365–1372.

- Fontanari, J. & Meir, R. 1990. The effect of learning on the evolution of asexual populations. *Complex Systems* **4**: 401–414.
- French, R. & Messinger, A. 1994. Genes, phenes and the Baldwin effect: Learning and evolution in a simulated population. In: *Artificial Life IV* (R. Brooks & P. Maes, eds), pp. 277–282. MIT Press, Cambridge, MA.
- Gavrilets, S. & Hastings, A. 1994. A quantitative-genetic model for selection on developmental noise. *Evolution* **48**: 1478–1486.
- Goldberg, D. 1989. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Publishing Company, Reading, MA.
- Gruau, F. & Whitley, D. 1993. Adding learning to the cellular development of neural networks: evolution and the Baldwin effect. *Evol. Comput.* **1**: 213–233.
- Hinton, G. & Nowlan, S. 1987. How learning can guide evolution. *Complex Systems* **1**: 495–502.
- Hughes, B. 1995. *Random Walks and Random Environments*, volume 1. Clarendon Press, Oxford.
- Huynen, M., Stadler, P. & Fontana, W. 1996. Smoothness within ruggedness: The role of neutrality in adaptation. *Proc. Natl. Acad. Sci.* **93**: 397–401.
- Kallel, L., Naudts, B. & Reeves, C. 2001. Properties of fitness functions and search landscapes. In: *Theoretical Aspects of Evolutionary Computing* (L. Kallel, B. Naudts & A. Rogers, eds), pp. 175–206. Springer, Berlin.
- Kauffman, S. A. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, New York.
- Kauffman, S. & Levin, S. 1987. Towards a general theory of adaptive walks on rugged landscapes. *J. Theor. Biol.* **128**: 11–45.
- Kimura, M. 1983. *The Neutral Theory of Molecular Evolution*. Cambridge University Press, Cambridge, UK.
- Kirkpatrick, S., Gelatt, C. & Vecchi, M. 1983. Optimization by simulated annealing. *Science* **220**: 671–680.
- Korona, R., Nakatsu, C., Forney, L. & Lenski, R. 1994. Evidence for multiple adaptive peaks from populations of bacteria evolving in a structured habitat. *Proc. Natl. Acad. Sci.* **91**: 9037–9041.
- Lenski, R., Ofria, C., Collier, T. & Adami, C. 1999. Genome complexity, robustness and genetic interactions in digital organisms. *Nature* **400**: 661–664.
- Littman, M. 1996. Simulations combining evolution and learning. In: *Adaptive Individuals in Evolving Populations: Models and Algorithms* (R. K. Belew & M. Mitchell, eds), pp. 465–477. Addison Wesley, Reading, MA.
- Littman, M. & Ackley, D. 1991. Adaptation in constant utility non-stationary environments. In: *Proceedings of the Fourth International Conference on Genetic Algorithms* (R. Belew & L. Booker, eds), pp. 136–142, Morgan Kaufmann, San Mateo, CA.
- Macken, C. & Perelson, A. 1989. Protein evolution on rugged landscapes. *Proc. Natl. Acad. Sci.* **86**: 6191–6195.
- Mayley, G. 1996. Landscapes, learning costs and genetic assimilation. *Evol. Comput.* **4**: 213–234.
- Mayley, G. 1997. Guiding or hiding: explorations into the effects of learning on the rate of evolution. In: *Proceedings of the Fourth European Conference on Artificial Life* (P. Husbands & I. Harvey, eds), pp. 135–144. Bradford Books/MIT Press, Cambridge, MA.
- Maynard-Smith, J. 1987. Natural selection: when learning guides evolution. *Nature* **329**: 761–762.
- Meilijson, I. 2003. The time to a given drawdown in Brownian motion. *Séminaire de Probabilités XXXVII., Springer Lecture Notes in Mathematics* **1832**: 94–108.
- Menczer, F. & Belew, R. 1994. Evolving sensors in environments of controlled complexity. In: *Artificial Life IV* (R. Brooks & P. Maes, eds), pp. 210–221. MIT Press, Cambridge, MA.
- Mery, F. & Kawecki, T. 2002. Experimental evolution of learning ability in fruit flies. *Proc. Natl. Acad. Sci.* **99**: 14274–14279.
- Mery, F. & Kawecki, T. 2004. The effect of learning on experimental evolution of resource preference in *Drosophila melanogaster*. *Evolution* **58**: 757–767.
- Morgan, C. 1896. On modification and variation. *Science* **4**: 733–740.
- Moriarty, D. & Mikkulainen, R. 1996. Efficient reinforcement learning through symbiotic evolution. *Mach. Learn.* **22**: 11–32.
- Nolfi, S. & Floreano, D. 1999. Learning and evolution. *Auton. Robot.* **7**: 89–113.
- Nolfi, S. & Parisi, D. 1997. Learning to adapt to changing environment in evolving neural networks. *Adaptive Behavior* **1**: 99–105.
- Noskovicz, S. & Goldhirsch, I. 1990. First passage time distribution in random random walk. *Phys. Rev.* **A42**: 2047–2064.
- Oksanen, L. & Lundberg, P. 1995. Optimization of reproductive effort and foraging time in mammals: The influence of resource level and predation risk. *Evol. Ecol.* **9**: 45–56.
- Parisi, D. & Nolfi, S. 1996. The influence of learning on evolution. In: *Adaptive Individuals in Evolving Populations: Models and Algorithms* (R. K. Belew & M. Mitchell, eds), pp. 419–428. Addison Wesley, Reading, MA.
- Schlichting, C. & Pigliucci, M. 1998. *Phenotypic Evolution: A Reaction Norm Perspective*. Sinauer Associates Inc., Sunderland, MA.
- Spitzer, F. 2001. *Principles of Random Walk*, 2nd edn. Springer, New-York.
- Stadler, P. 1995. Towards a theory of landscapes. In: *Complex Systems and Binary Networks* (R. López-Peña, R. Capovilla, R. García-Pelayo, H. Waelbroeck & F. Zertuche, eds), pp. 77–163. Springer Verlag, Berlin, New York.
- Todd, P. & Miller, G. 1991. Exploring adaptive agency II: simulating the evolution of associative learning. In: *From Animals to Animals: Proceedings of the First International Conference on Simulation of Adaptive Behavior* (J. Meyer & S. Wilson, eds), pp. 306–315, MIT Press, Cambridge, MA.
- Waddington, C. 1942. Canalization of development and the inheritance of acquired characters. *Nature* **150**: 563–565.
- Waddington, C. 1953. Genetic assimilation for acquired character. *Evolution* **7**: 118–126.
- Weber, B. & Depew, D., eds 2003. *Learning and Evolution. The Baldwin Effect Reconsidered*. MIT Press, Cambridge, MA.
- Weinberger, E. 1990. Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biol. Cybern.* **63**: 325–336.
- West-Eberhard, M. 2003. *Developmental Plasticity and Evolution*. Oxford University Press, New York.
- Wright, S. 1932. The roles of mutation, inbreeding, crossbreeding, and selection in evolution. In: *Proceedings of the Sixth International Congress of Genetics*, volume 1, (D.F. Jones, ed.), pp. 356–366. Brooklyn Botanic Garden, Brooklyn, NY.

Appendix 1: mean first-passage times of a simple one-dimensional random walk in an arbitrary environment

First, we clarify the distinction between the well-established theory of random walk in random-environments (RWRE) and our analysis: the theory of RWRE deals with the scenario where the probabilities to take a +1 or -1 step are independently chosen from some distribution on (0,1), characterizing sublinear or linear speed

$$E_0^N = N + 2 \sum_{\substack{i \leq j \\ 0 < i, j < N}} \prod_{k=i}^j \rho_k.$$

Interestingly, this expression for the mean first-passage time from 0 to N can also be represented as the quadratic form

$$E_0^N = 1' \tilde{A} 1, \tag{2}$$

where

$$\tilde{A} = \begin{pmatrix} 1 & \rho_1 & \rho_1 \rho_2 & \rho_1 \rho_2 \rho_3 & \rho_1 \rho_2 \rho_3 \rho_4 & \cdots & \prod_{k=1}^{N-1} \rho_k \\ \rho_1 & 1 & \rho_2 & \rho_2 \rho_3 & \rho_2 \rho_3 \rho_4 & \cdots & \prod_{k=2}^{N-1} \rho_k \\ \rho_1 \rho_2 & \rho_2 & 1 & \rho_3 & \rho_3 \rho_4 & \cdots & \prod_{k=3}^{N-1} \rho_k \\ \rho_1 \rho_2 \rho_3 & \rho_2 \rho_3 & \rho_3 & 1 & \rho_4 & \cdots & \prod_{k=4}^{N-1} \rho_k \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \prod_{k=1}^{N-1} \rho_k & \prod_{k=2}^{N-1} \rho_k & \prod_{k=3}^{N-1} \rho_k & \cdots & \rho_{N-2} \rho_{N-1} & \rho_{N-1} & 1 \end{pmatrix}.$$

and large deviation questions. Our analysis, in contrast, relates to the simpler setup of a random walk on a finite set, but allows these probabilities to be arbitrary given constants.

1.1: general form and pertaining examples

Consider a simple random walk S_t in a changing environment on $\{0, 1, 2, \dots, N\}$. Let $p_i = P(S_{t+1} = i + 1 | S_t = i)$ and let $q_i = 1 - p_i = P(S_{t+1} = i - 1 | S_t = i)$. Let $p_0 = 1$ and assume that $0 < p_i < 1$ for all $0 < i < N$. Let E_i^j denote the mean first-passage time from i to j , i.e. the expected time to first hit j starting at i . We get $E_0^1 = 1$ and the recursion

$$E_i^{i+1} = 1 + p_i \cdot 0 + (1 - p_i)[E_{i-1}^i + E_i^{i+1}], \quad i = 1, 2, \dots, N-1. \tag{1}$$

Letting ρ_i denote the odds ratio $\frac{1-p_i}{p_i}$, eqn 1 may be written as

$$E_i^{i+1} = 1 + \rho_i [1 + E_{i-1}^i]$$

whose solution can be easily seen to be

$$E_i^{i+1} = 1 + 2 \sum_{j=1}^i \prod_{k=j}^i \rho_k.$$

As

$$E_0^N = \sum_{i=0}^{N-1} E_i^{i+1},$$

we finally obtain

In particular, consider the following examples, demonstrating the mean first-passage time in a constant environment.

Example 1: constant symmetric environment

In a constant symmetric environment, $p_i \equiv \frac{1}{2} (\rho_i \equiv 1)$ for all $0 < i < N$ and $\tilde{A} = 11'$, the matrix of ones ($N \times N$). The quadratic form (eqn 2) then yields $E_0^N = N^2$.

Example 2: constant nonsymmetric environment

In this scenario, $p_i \equiv p \neq \frac{1}{2} (\rho_i \equiv \rho \neq 1)$ for all $0 < i < N$, and

$$\tilde{A} = \begin{pmatrix} 1 & \rho & \rho^2 & \rho^3 & \cdots & \rho^{N-1} \\ \rho & 1 & \rho & \rho^2 & \cdots & \rho^{N-2} \\ \rho^2 & \rho & 1 & \rho & \cdots & \rho^{N-3} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho^{N-1} & \rho^{N-2} & \rho^{N-3} & \cdots & \rho & 1 \end{pmatrix}$$

The quadratic form (eqn 2) can be easily seen to yield

$$E_0^N = N \frac{1 + \rho}{1 - \rho} + 2\rho \frac{\rho^N - 1}{(\rho - 1)^2},$$

inducing a linear first-passage time (in N) if $\rho < 1$ and an exponential first-passage time if $\rho > 1$.

Now, consider the indices $x_1 < x_2 < \dots < x_K$, where $\rho_{x_i} \neq 1 (p_{x_i} \neq \frac{1}{2})$. Also define $x_0 = 0$ and $x_{K+1} = N$. Let $n_i = x_i - x_{i-1} (i = 1, 2, \dots, K + 1)$ denote the corresponding increments. As $\rho_{x_i+1}, \rho_{x_i+2}, \dots, \rho_{x_{i+1}-1}$ all equal 1 for each $0 \leq i \leq K$, \tilde{A} consists of rectangular blocks and $E_0^N = 1' \tilde{A} 1 = V' A V$, where $V = (n_1, n_2, \dots, n_{K+1})'$ and

$$A = \begin{pmatrix} 1 & \rho_{x_1} & \rho_{x_1}\rho_{x_2} & \rho_{x_1}\rho_{x_2}\rho_{x_3} & \rho_{x_1}\rho_{x_2}\rho_{x_3}\rho_{x_4} & \cdots & \prod_{k=1}^K \rho_{x_k} \\ \rho_{x_1} & 1 & \rho_{x_2} & \rho_{x_2}\rho_{x_3} & \rho_{x_2}\rho_{x_3}\rho_{x_4} & \cdots & \prod_{k=2}^K \rho_{x_k} \\ \rho_{x_1}\rho_{x_2} & \rho_{x_2} & 1 & \rho_{x_3} & \rho_{x_3}\rho_{x_4} & \cdots & \prod_{k=3}^K \rho_{x_k} \\ \rho_{x_1}\rho_{x_2}\rho_{x_3} & \rho_{x_2}\rho_{x_3} & \rho_{x_3} & 1 & \rho_{x_4} & \cdots & \prod_{k=4}^K \rho_{x_k} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \prod_{k=1}^K \rho_{x_k} & \prod_{k=2}^K \rho_{x_k} & \prod_{k=3}^K \rho_{x_k} & \cdots & \rho_{x_{K-1}}\rho_{x_K} & \rho_{x_K} & 1 \end{pmatrix}$$

Thus,

$$E_0^N = \sum_{i=1}^{K+1} n_i^2 + 2 \sum_{\substack{i \leq j \\ 0 < i, j < K+1}} n_i n_{j+1} \prod_{k=1}^j \rho_{x_k}. \tag{3}$$

In particular, consider the following example, compactly describing the first-passage time in a typical ‘flattened’ ideal learning landscape:

Example 3: symmetric random walk with K equally distant disturbances

In this scenario, we get

$$x_i = i \frac{N}{K+1} \quad \text{and} \quad n_i \equiv n = \frac{N}{K+1} \quad \text{for all } 0 < i \leq K+1.$$

Equation 3 then yields

$$E_0^N = \frac{N^2}{K+1} + 2 \frac{N^2}{(K+1)^2} \sum_{\substack{i \leq j \\ 0 < i, j < K+1}} \prod_{k=i}^j \rho_{x_k}.$$

1.2: an asymptotic bound

We now analyse an asymptotic bound for E_0^N . Define the ‘drawdown’ R of the random-walk process as the maximal element of A , i.e.

$$R = \max_{\substack{i \leq j \\ 0 < i, j < K+1}} \prod_{k=i}^j \rho_{x_k}.$$

We want to find bounds on E_0^N in terms of R . Clearly, $\max(N, 2R) \leq E_0^N \leq N^2 R$. In fact, as we shall now see, $E_0^N \leq N^2 (\frac{1+R}{2})$ and this bound is in some sense asymptotically sharp. We first consider the case where $R > 1$ and analyse the quadratic form $V'AV$ obtained above. Fix all K values $\rho_i \neq 1$ as given, but allow $Y = (y_1, y_2, \dots, y_{K+1}) = (1/N)(n_1, n_2, \dots, n_{K+1})$ to vary in its feasible set $FS_N = \{(y_1, y_2, \dots, y_{K+1}) | y_i \geq 0, \sum y_i = 1, N y_i \text{ positive integers}\}$, a subset of $FS_\infty = \{(y_1, y_2, \dots, y_{K+1}) | y_i \geq 0, \sum y_i = 1\}$. Then $\max_{Y \in FS_N} (Y'AY) \leq \max_{Y \in FS_\infty} (Y'AY)$, and the two are close to each other if $K \ll N$.

Claim: $\max_{Y \in FS_\infty} (Y'AY) = \frac{1+R}{2}$.

Proof:

- 1 For $\rho_i < 1$, fix all y_j except y_i and y_{i+1} and fix $y_i + y_{i+1} = T$. Then $Y'AY$ is a convex quadratic function of y_i so one of y_i and y_{i+1} should be zero. This means that all $\rho_i < 1$ should ‘collapse’ (i.e. be at distance zero) to a neighbouring $\rho_j > 1$. Furthermore, all $\rho_i < 1$ left of the leftmost $\rho_j > 1$ (be that at j_0) collapse to location zero and all $\rho_i < 1$ to the right of the rightmost $\rho_j > 1$ (be that at j_1) collapse to location N .
- 2 This leaves us with a reduced matrix \hat{A} of size $(\hat{K} + 1) \times (\hat{K} + 1)$, $\hat{K} \leq K$, with $\rho_i > 1$ for all $0 < i \leq \hat{K}$. We claim that $\max(\hat{Y}'\hat{A}\hat{Y})$ is achieved by collapsing all ρ 's from j_0 to j_1 into one point in the exact middle. This will give

$$\begin{pmatrix} \frac{1}{2} & \frac{1}{2} \end{pmatrix} \begin{pmatrix} 1 & R \\ R & 1 \end{pmatrix} \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix} = \frac{1+R}{2}.$$

Suppose that $\max(\hat{Y}'\hat{A}\hat{Y})$ is achieved at a vector $(y_1, y_2, \dots, y_{\hat{K}+1})$ with $y_i > 0$ for each $i = 1, 2, \dots, \hat{K} + 1$. Then this vector must achieve a zero of the derivative of the Lagrangian $\hat{Y}'\hat{A}\hat{Y} - \lambda(1'\hat{Y} - 1)$, yielding $\hat{Y} = C \cdot \hat{A}^{-1}1$. But if

$$\hat{A} = \begin{pmatrix} 1 & \rho_1 & \rho_1\rho_2 & \cdots \\ \rho_1 & 1 & \rho_2 & \cdots \\ \rho_1\rho_2 & \rho_2 & 1 & \cdots \\ \vdots & \cdots & \ddots & \vdots \\ \cdots & \cdots & \cdots & \rho_{\hat{K}} & 1 \end{pmatrix}$$

then \hat{A}^{-1} is a symmetric tridiagonal matrix with

$$\hat{A}^{-1}(1, 1) = \frac{-1}{\rho_1^2 - 1}, \hat{A}^{-1}(\hat{K} + 1, \hat{K} + 1) = \frac{-1}{\rho_{\hat{K}}^2 - 1},$$

$$\hat{A}^{-1}(i, i) = -(1 + \frac{1}{\rho_{i-1}^2 - 1} + \frac{-1}{\rho_i^2 - 1}) \quad \text{for all } 1 < i \leq \hat{K}$$

and

$$\hat{A}^{-1}(i, i + 1) = \hat{A}^{-1}(i + 1, i) = \frac{\rho_i}{\rho_i^2 - 1} \quad \text{for all } 1 < i \leq \hat{K}.$$

So

$$(\hat{A}^{-1}\mathbf{1})(1) = \hat{A}^{-1}(1, 1) + \hat{A}^{-1}(1, 2) = \frac{-1}{\rho_1^2 - 1} + \frac{\rho_1}{\rho_1^2 - 1} > 0,$$

$$\begin{aligned} (\hat{A}^{-1}\mathbf{1})(\hat{K} + 1) &= \hat{A}^{-1}(\hat{K} + 1, \hat{K}) + \hat{A}^{-1}(\hat{K} + 1, \hat{K} + 1) \\ &= \frac{\rho_{\hat{K}}}{\rho_{\hat{K}}^2 - 1} + \frac{-1}{\rho_{\hat{K}}^2 - 1} > 0, \end{aligned}$$

but

$$\begin{aligned} (\hat{A}^{-1}\mathbf{1})(i) &= \hat{A}^{-1}(i, i - 1) + \hat{A}^{-1}(i, i) + \hat{A}^{-1}(i, i + 1) \\ &= \frac{\rho_{i-1}}{\rho_{i-1}^2 - 1} - \left(1 + \frac{1}{\rho_{i-1}^2 - 1} + \frac{1}{\rho_i^2 - 1}\right) \\ &\quad + \frac{\rho_i}{\rho_i^2 - 1} < 0 \quad \text{for all } 1 < i < \hat{K} + 1. \end{aligned}$$

As we see, $\hat{K} = 1$ is the only feasible value and the problem reduces to

$$(y_1 \quad 1 - y_1) \begin{pmatrix} 1 & R \\ R & 1 \end{pmatrix} \begin{pmatrix} y_1 \\ 1 - y_1 \end{pmatrix} = 1 + 2(R - 1)y_1(1 - y_1),$$

which for $R > 1$ has a maximum at $y_1 = 1/2$, as claimed.

In the case of $R = 1$, let all $\rho_i < 1$ stick to 0 or N as before. We obtain a flat environment, for which $E_0^N = N^2$.

So, the final answer is

$$E_0^N \leq N^2 \left(\frac{1 + R}{2} \right)$$

and this is asymptotically sharp in the sense above, where $K \ll N$. Hence, R is the critical factor determining the passage time in a given landscape.

Supplementary Material

The following supplementary material is available for this article online:

Appendix S1: Mean first-passage times of a multi-dimensional random walk in an arbitrary environment

Appendix S2: Methods

This material is available as part of the online article from <http://www.blackwell-synergy.com>

Received 17 January 2006; revised 27 February 2006; accepted 28 February 2006