Surviving the apocalypse: resisting extinction via the Baldwin Effect

Ryan Lynch*
Massachusetts Institute of Technology, 185 Albany St, 02139 Cambridge USA

The Baldwin Effect is a means of explaining genetic assimilation in evolutionary theory: a process by which phenotype plasticity can eventually become hard-coded as genetic instinct. Hinton and Nowlan propose an evolutionary model in which the genetic assimilation predicted by the Baldwin effect is observed for static fitness landscapes. In this paper, we explore the behavior of the Hinton-Nowlan model in a dynamic fitness landscape, where the optimal-fitness phenotype evolves in time. Specifically, we examine the ability for Hinton-Nowlan populations to successfully adapt to this changing fitness landscape and prevent extinction events. We find that plasticity helps to accelerate evolutionary selection, and we observe genetic assimilation even in a changing fitness landscape. We see that as a result of this genetic assimilation, plasticity alone is not enough to prevent extinction events, but that genetic mutations are also needed to maintain some non-zero level of genetic variance. Finally, we argue that in order for a population to resist extinction, we should have the characteristic timescale of evolutionary selection $\lesssim$ the characteristic timescale of genetic mutation $\lesssim$ the characteristic timescale of the changing fitness landscape.

I. INTRODUCTION

A fascinating concept that has persisted throughout the history of evolutionary theory is Lamarckian inheritance. Lamarckism is a form of epigenetic inheritance where traits learned by parents can be passed onto their children without altering the parents' genotypes. Although August Weismann was able to disprove many elements of pure Lamarckism theory with his 1893 experiment [1], James Baldwin (1896) was able to come up with a Darwinist explanation of why some Lamarckian-esque phenomena are observed in nature [2]. The so-called "Baldwin Effect" argues that a genetically-inheritable phenotype may have some plasticity associated with it, and that this ability for an individual to explore several phenotypes during its lifetime not only effects its own fitness but also effects the genetic makeup of future generations.

Although somewhat controversial throughout the 20^{th} century, the Baldwin Effect is now generally recognized as part of the modern synthesis of evolutionary genetics. This acceptance is in part due to a 1987 paper by Hinton and Nowlan [3] in which they demonstrate via computer simulation that having phenotype plasticity can rapidly accelerate the process of fixing the desired phenotype to be instinctual for future generations. The intuition behind this phenomena can be illustrated by considering a complex task, such humans' ability to hunt. In order for hunting to be instinctual, the many components of hunting need to be instinctual as well, such as the ability to track prey, the ability to build weapons, etc.. Hinton and Nowlan modeled the genotype encoding of these traits as a string of binary alleles, 0's and 1's. Requiring evolution to simultaneously hard-code all of these necessary traits in the proper combination may require the timescales of random global searches over alleles ($2^{\text{string length}}$), especially if partially-correct phenotypes are destructive (such as the ability for humans to create weapons without the ability to track prey). In other words, evolution may need to find a phenotype of 10 1's, but in a landscape where the fitness is only non-negligible when all 10 1's are found simultaneously.

Instead, the Baldwin effect allows individuals to have the ability to explore several different phenotypes in their lifetimes, i.e., it allows for several local searches over traits in between global searches [4]. In the Hinton and Nowlan model, plastic alleles are modeled as ?'s, which are randomly replaced with either a 0 or 1 during each phenotype trial. The total number of attempts is determined by the population's learning rate. In this example scenario, individuals that learn the target phenotype the quickest have the highest fitnesses, since they can quickly learn to hunt and provide for both themselves and for others. Hinton and Nowlan modeled the fitness of an individual, F, explicitly as

$$F = L - \frac{L-1}{\lambda}\min(k, \lambda) \tag{1}$$

where L is the length of the target string, λ is the maximum number of phenotype trials each individual has, and k is the number of trials it takes to find the target phenotype. Because higher fitnesses are achieved when the target phenotype is found faster, evolutionary selection favors the replacement of any plastic alleles with the target allele. In other words, Hinton and Nowlan showed that mating and mutation can eventually replace ?'s with the correct 0's or 1's, so that the target phenotype is eventually hard-coded in the genotype. This fixation process is known as genetic assimilation (see [5] for a in-depth study of the conditions under which genetic assimilation is possible).

Although genetic assimilation is a remarkable phenomena, it is not necessarily beneficial to a population's survival. Since plasticity is a valuable trait, its removal from the population's collective genotype decreases genetic variation and the ability to react to evolving fitness

* ryan.lynch@ligo.org

landscapes. Thus, even though a population with minimal plasticity might possess a larger absolute fitness than a more plastic population, the latter will be better prepared to stave off extinction in a changing environment. In the remainder of this paper, we explore this last point quantitatively by simulating how the Hinton-Nowlan evolutionary model behaves under a changing target phenotype. We refer the reader to [6] for a recent review of other evolutionary models that incorporate the Baldwin Effect, such as those making the unimodal fitness landscape of Hinton-Nowlan more rugged.

II. DESCRIPTION OF THE EVOLUTIONARY MODEL

We implement a version of the Hinton-Nowlan evolutionary model that evolves from one generation of M individuals with length-L genotypes (containing a combination of 0's, 1's, and ?'s as alleles) to another generation of the same size using the following steps:

1. Two parents are selected from the size-M population in proportion to their relative fitnesses, namely with probability $\frac{F_i}{\sum_{i=1}^{M} F_i}$. The individuals are assumed to be hermaphrodites, so that any individual can reproduce with any other individual (including itself).

2. We simulate recombination by randomly assigning each allele of the child's length-L genotype to be a random selection of the parents' alleles at the same position.

3. We then simulate genetic mutation of the child's genotype by randomly changing each allele in the genotype with probability μ. There are no preferred mutations among the 0's, 1's, and ?'s: each allele remains the same with probability $1-\mu$, and changes to each of the two different values with probability $\frac{1}{2}\mu$

4. The above three steps are repeated until M children have been produced.

5. We now simulate a change in the fitness landscape that the children will experience. This is done by randomly mutating each allele in the target phenotype in a similar manner as the genetic mutation described above, only with mutation probability γ instead of μ. We stress that the target phenotype consists of only 0's and 1's.

6. Finally, we calculate the fitness of each individual in the child generation with respect to the new target phenotype using Eq. 1. This involves doing at most λ joint random replacements of all ?'s with 0's and 1's for each individual in search of the target phenotype.

For this paper, we run the above algorithm for 500 generations with a fixed population size of $M = 1000$ individuals and genotype lengths of $L = 10$. In the first generation, each string is initialized by randomly choosing between 0's, 1's, and ?'s for each allele with equal probability. With this initialization of the genotypes, the initialization of the target phenotype does not matter. In order to summarize the behavior of each trial, we compute two figures of merit. The first is the number of generations it takes the randomly initialized population to first achieve a mean fitness greater than 2, which we define as the "convergence time". The second is the number of generations measured from the convergence time that it takes for the population to have a mean fitness of exactly 1. The mean fitness of 1 can only be achieved if no member of the population is able to find the target phenotype during its learning, and thus we call this statistic the "extinction time". We perform the above simulation over a grid of values for the genetic mutation probability μ (0., 0.001), the fitness-landscape mutation probability γ (0., 0.001, 0.01), and the learning rate λ (2, 10, 100). We run 50 trials of each configuration in order to accumulate statistics, and we stop any given trial as soon as an extinction event (mean fitness of exactly 1) occurs. We analyze the results of these simulations below.

III. SIMULATION RESULTS

We first analyze the results of simulations with no genetic mutation, i.e., $\mu = 0$. The timescale figure of merits are given in Table I and the time-series evolution of example runs are displayed in Fig. 1. Let us first consider the classical Hinton-Nowlan model where the target phenotype does not change over time, i.e., $\gamma = 0.$. Randomly initialized populations with higher learning rates λ converge to non-negligible fitnesses quicker than those with low learning rates, showing that plasticity can in fact accelerate the convergence of evolution. We also observe that genetic assimilation does indeed occur. This is made evident by examining how both the standard deviation of the population's fitness and the mean number of ? alleles in the population decay over time in Fig. 1. Note that once the population's mean fitness converges to the optimal value of 10, both of these quantities decay to zero, meaning every member of the population eventually has the exact target phenotype as its genotype (see Eq. 1). As we expect, this decay occurs rapidly when the learning rate λ is small since each plastic allele ? is unlikely to find its optimal value during the λ local-search trials. On the contrary, the decay occurs much more slowly when λ is large since an individual can still achieve a near-optimal fitness even if it contains many ?'s as a result of its large number of local-search trials.

Once we add a time-evolving fitness landscape to the model by allowing γ to be non-zero, we begin to observe extinction events. However, for all values of γ, the extinc-

TABLE I. Convergence and extinction timescales for $\mu = 0$, averaged over 50 trials. The convergence (i.e, selection) time is defined as the number of generations is takes for a randomly initialized population to reach a mean fitness of 2. The extinction time is the number of generations it takes for a converged population to reach a mean fitness of exactly 1 (the minimal allowed fitness value). We define "percent extinct" to be the number of trials that experience extinction events within the first 500 generations.

γ	λ [Trials per Generation]	Mean Convergence Time [Generations]	Percent Extinct [%]	Mean Extinction Time [Generations]
0.0	2	32	0	N/A
0.0	10	6	0	N/A
0.0	100	3	0	N/A
0.001	2	31	100	102
0.001	10	6	98	113
0.001	100	3	96	204
0.01	2	44	100	15
0.01	10	6	100	34
0.01	100	3	100	73

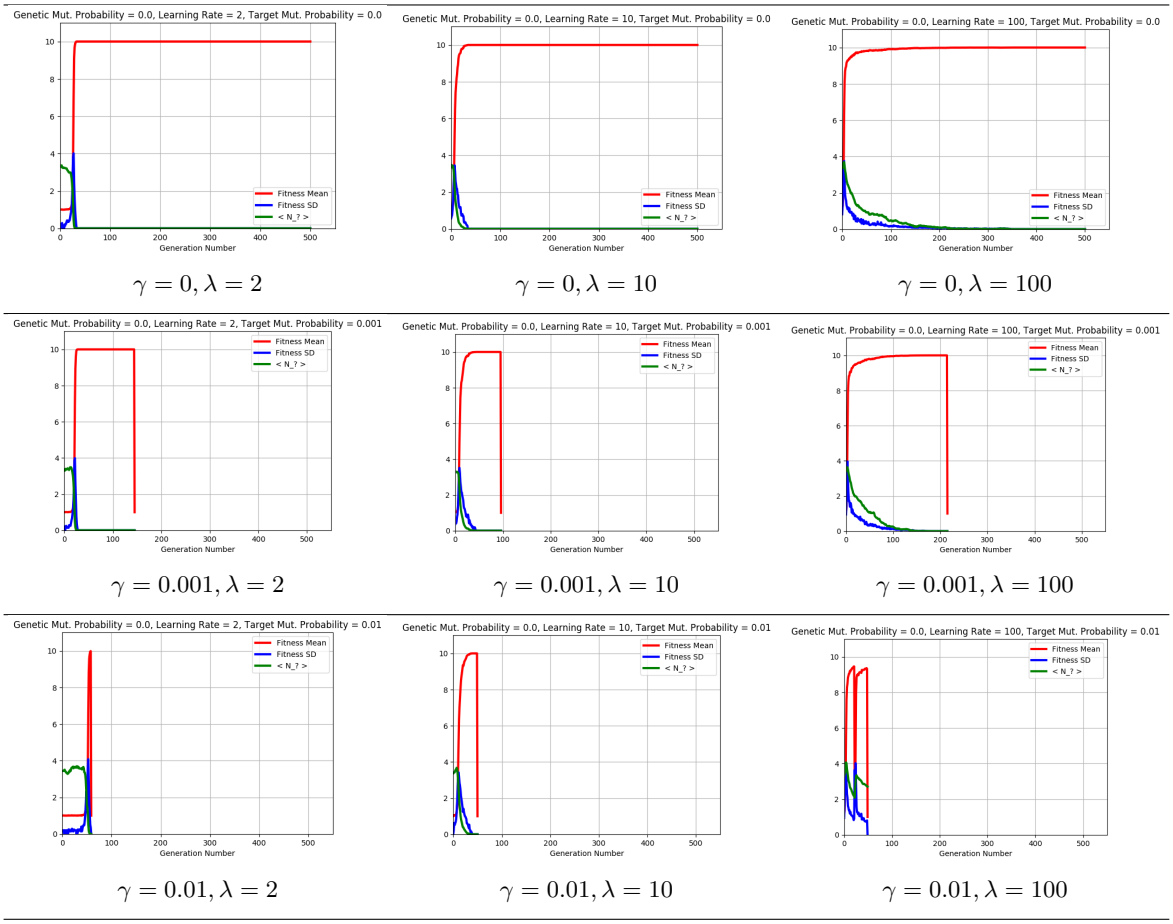

$\gamma = 0, \lambda = 2$ $\gamma = 0, \lambda = 10$ $\gamma = 0, \lambda = 100$

$\gamma = 0.001, \lambda = 2$ $\gamma = 0.001, \lambda = 10$ $\gamma = 0.001, \lambda = 100$

$\gamma = 0.01, \lambda = 2$ $\gamma = 0.01, \lambda = 10$ $\gamma = 0.01, \lambda = 100$

FIG. 1. Time series evolution of a population (as a function of generation) for $\mu = 0$, for single example trials. We plot the mean and standard deviation (SD) of the fitness, calculated over the 1000 individuals in the population. We also plot the mean number of ?-alleles ($\langle N_? \rangle$) in each $L = 10$ genotype, again calculated over the 1000 individuals in the population. We stop plotting the time series after an extinction event occurs or if 500 generations is reached without an extinction.

tions occur less quickly as the learning rate λ increases. As evidence, there is close to a factor-of-2 increase in the mean extinction time when going from $\lambda = 10$ to $\lambda = 100$ for each non-zero value of γ. For low-λ populations, genetic assimilation occurs very quickly, meaning there is little chance that the population has any plasticity or genetic variation to adapt with when the target phenotype mutates. On the other hand, high-λ populations are more likely to have enough plasticity to adapt to a given change in the target phenotype since their genetic assimilation occurs more slowly. One example of this adaptation is shown in the $\gamma = 0.01, \lambda = 100$ panel

of Fig. 1. The target phenotype changes around generation 25, leading to a dip in the population's mean fitness and rises in the the population's fitness variance and mean plasticity. These behaviors are representative of the population encouraging plasticity, which is needed to adapt to the new target phenotype. We find no evidence of a high fitness-landscape mutation rate helping to prevent extinction, as the extinction time decreases as γ increases.

We now analyze the results of simulations with non-zero genetic mutation, where $\mu = 0.001$. The timescale figure of merits are given in Table II and the time-series evolution of example runs are displayed in Fig. 2. We still observe: 1.) plasticity accelerating the evolutionary selection, as higher values of λ converge faster, and 2.) genetic assimilation, as all simulations asymptote to having near optimal-fitness and low numbers of plastic ? alleles. The difference with these $\mu = 0.001$ runs is that there is a non-zero level of genetic variance once the assimilation has occurred, meaning some minimal level of plasticity is always maintained within the population. While the $\gamma = 0$ panels of Fig. 2 show that the mean number of ?'s is close to 0 over the population of 1000 individuals, the non-zero fitness variance ensures that there are at least a few individuals with plasticity.

We also observe that despite this post-assimilation genetic variance, extinctions still do occur. However, comparing the fraction of simulations that experience extinction and mean extinction times in Table II to those in Table I, we see that extinctions occur significantly less frequently when $\mu = 0.001$ than when $\mu = 0$. This reduction in extinctions is due to the increased genetic variance discussed above, which gives the populations a better ability to adapt to the changing fitness environment. This increased ability can be seen by noting that there are many more recovery events (the dips and spikes described above) seen in Fig. 2 than in Fig. 1. The best evidence that these recovery events are due to the non-zero μ is that recovery events are not just seen for high-λ configurations, but also for low-λ configurations where genetic assimilation occurs quickly. Thus populations have the ability to adapt to a mutated target phenotype regardless of if the population has assimilated to the pre-mutation phenotype or not.

At large values of μ (as compared to the fitness-landscape mutation rate γ), it is possible that enough plasticity is maintained to eliminate extinction events almost entirely. Nevertheless, if μ gets too large, the variance will eventually become large enough that genetic-drift extinction events will occur even for comparatively static fitness landscapes. As is known from drift-diffusion approximations of evolutionary models (e.g., see [7] for an application to the Hinton-Nowlan model), when the genetic mutation rate μ grows large enough, the distribution of each character in the genotype string will be distributed between 0, 1, and ? according to their relative mutation rates (here, uniformly across all 3 alleles). In the Hinton-Nowlan model, only the exact target phenotype is rewarded with non-minimal fitness. Without a selection force, the probability that an individual in the population achieves the optimal phenotype by chance decreases exponentially as $e^{-\alpha L}$, where α is a learning-rate dependent parameter that would have a value of $\ln 2$ if no plasticity is allowed (i.e., if we only had 1's and 2's with equal probabilities). As a result, extinctions become exponentially likely in the high-μ regime. In summary, our limited set of simulations suggest that the frequency of extinction events will be minimized if the three relevant timescales in the problems have the following ordering:

$$\tau_{\text{selection}} \lesssim \tau_\mu \lesssim \tau_\gamma \tag{2}$$

where $\tau_{\text{selection}}$ is the evolutionary selection (i.e., convergence) timescale, τ_μ is the genetic-mutation timescale ($\sim \frac{1}{ML\mu}$), and τ_γ is the fitness-landscape mutation timescale ($\sim \frac{1}{L\gamma}$).

IV. CONCLUSIONS

The Hinton-Nowlan evolutionary model has quite stringent requirements for fitness-based selection, since non-minimal fitness values are only achievable if a individual can obtain the exact target phenotype within its lifetime. Thus, even slight changes to the target phenotype can possibly cause the extinction of a finite-sized population. Simulating the Hinton-Nowlan model, we are able to observe the two key features of the Baldwin Effect: 1.) that plasticity can help accelerate evolutionary selection, and 2.) that genetic assimilation occurs, which eventually replaces plasticity with instinctual knowledge of the target phenotype. Both of these phenomena occur even in an evolving fitness landscape where the target phenotype evolves in time. We find that because of genetic assimilation, having a high plasticity alone does not guarantee that a population can survive changes in the fitness landscape, exactly because this plasticity will eventually disappear as a result of selection forces. Thus, some level of genetic mutation is needed in order to maintain non-zero plasticity in post-assimilation generations. With genetic mutation, populations become much more effective at adapting to changes in the fitness landscape and thus at preventing extinction. We finally argue that the frequency of extinction events should be minimal if the rate of evolutionary selection $\gtrsim$ the rate of genetic mutations $\gtrsim$ the rate of fitness-landscape mutations.

[1] A. Weismann, *The Germ-plasm: a theory of heredity. Translated by W. Newton Parker and Harriet Rnnfeldt.* (New YorkScribner) p. 514,

TABLE II. Convergence and extinction timescales for $\mu = 0.001$, averaged over 50 trials. The convergence (i.e, selection) time is defined as the number of generations is takes for a randomly initialized population to reach a mean fitness of 2. The extinction time is the number of generations it takes for a converged population to reach a mean fitness of exactly 1 (the minimal allowed fitness value). We define "percent extinct" to be the number of trials that experience extinction events within the first 500 generations.

γ	λ [Trials per Generation]	Mean Convergence Time [Generations]	Percent Extinct [%]	Mean Extinction Time [Generations]
0.0	2	30	0	N/A
0.0	10	6	0	N/A
0.0	100	3	0	N/A
0.001	2	34	88	133
0.001	10	6	40	305
0.001	100	3	10	317
0.01	2	50	100	24
0.01	10	7	100	59
0.01	100	3	98	158

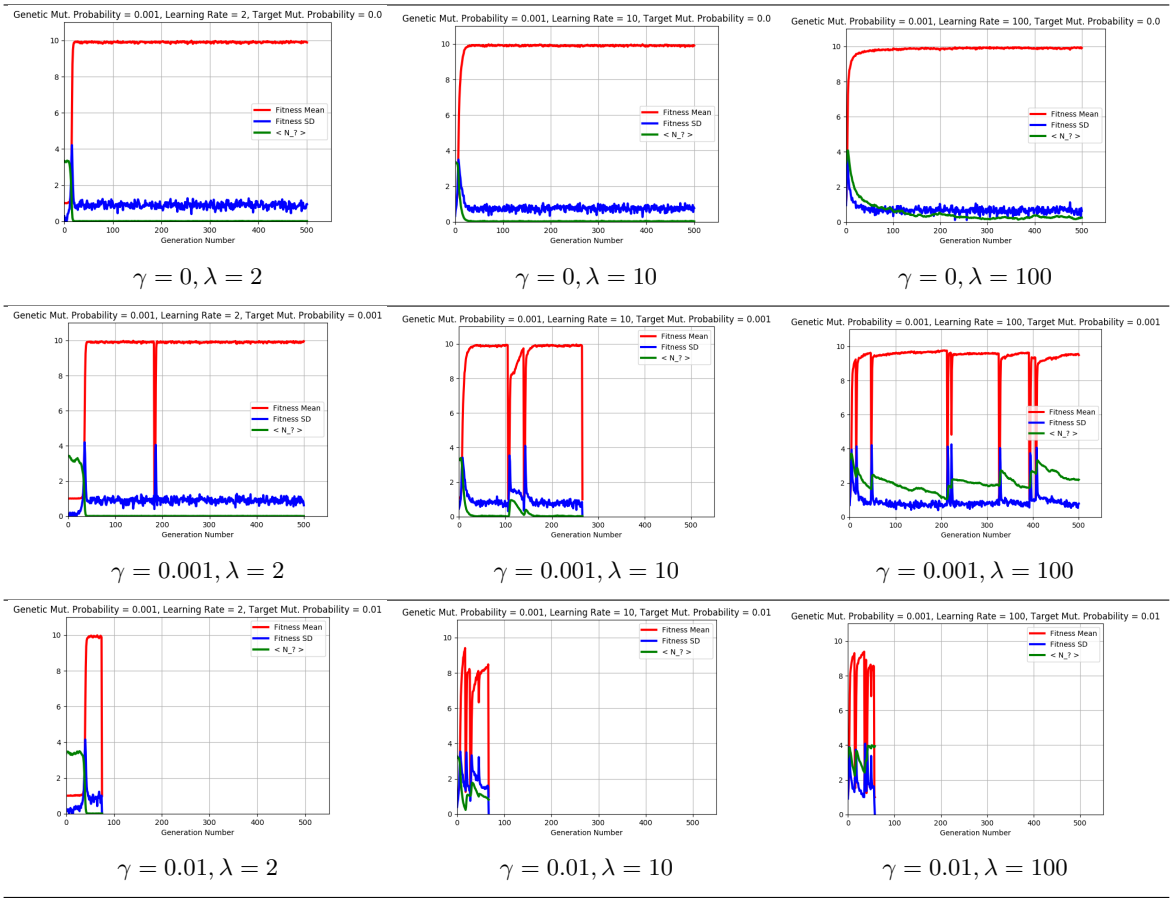

$\gamma = 0, \lambda = 2$ $\gamma = 0, \lambda = 10$ $\gamma = 0, \lambda = 100$

$\gamma = 0.001, \lambda = 2$ $\gamma = 0.001, \lambda = 10$ $\gamma = 0.001, \lambda = 100$

$\gamma = 0.01, \lambda = 2$ $\gamma = 0.01, \lambda = 10$ $\gamma = 0.01, \lambda = 100$

FIG. 2. Time series evolution of a population (as a function of generation) for $\mu = 0.001$, for single example trials. We plot the mean and standard deviation (SD) of the fitness, calculated over the 1000 individuals in the population. We also plot the mean number of ?-alleles ($\langle N_? \rangle$) in each $L = 10$ genotype, again calculated over the 1000 individuals in the population. We stop plotting the time series after an extinction event occurs or if 500 generations is reached without an extinction.

https://www.biodiversitylibrary.org/bibliography/25196.
[2] J. M. Baldwin, The American Naturalist **30**, 441 (1896).
[3] G. E. Hinton and S. J. Nowlan, Complex Systems **1**, 495 (1987).
[4] J. Shapiro, "Genetic algorithms in machine learning," in *Machine Learning and Its Applications: Advanced Lectures*, edited by G. Paliouras, V. Karkaletsis, and C. D. Spyropoulos (Springer Berlin Heidelberg, Berlin, Heidelberg, 2001) pp. 146–168.
[5] G. Mayley, Evolutionary Computation **4**, 213 (1996), https://doi.org/10.1162/evco.1996.4.3.213.
[6] B. Sznajder, M. W. Sabelis, and M. Egas, Evolutionary Biology **39**, 301 (2012).

[7] E. U. K, S. F. (ed, I. Harvey, and I. Harvey, in *Proceedings of the Fifth International Conference on Genetic Algorithms* (Morgan Kaufmann, 1993) pp. 15–22.