Reaction norms with bifurcations shaped by evolution

Tom J. M. van Dooren

Department of Biology, University of Antwerp, Universiteitsplein 1, 2610 Antwerp, Belgium (vdooren@uia.ua.ac.be)

Two versions of a model for the evolution of seasonal polyphenism investigate the evolution of reaction norm bifurcation and branching. The first version is without a specific submodel for morphological development and the second has an explicit developmental map. Version 1 is evolutionarily relatively unconstrained: (i) reaction norms are specified by matrices containing the probabilities of occurrence of environment–phenotype combinations, (ii) all conceivable reaction norm matrices are reachable through a sequence of mutations, and (iii) small as well as large mutational effects occur. This version is used to find the evolutionarily stable strategy favoured by the population ecology that is characterized by stabilizing viability selection with a cyclically fluctuating selection optimum. When the strength of selection is large and when the lag between initiation of development and selection on mature phenotype is not a multiple of half the period of the environmental cycle, a branching reaction norm evolves. In the second model version, branching reaction norms occur for certain parameter combinations of the developmental submodel, but the evolution of this pattern is often constrained. The evolutionary trajectory becomes trapped in a local selective optimum for the parameters of the developmental system. Substantial developmental noise evolves, but mutations that produce a selectively advantageous branching pattern do not occur from there.

Keywords: evolution of reaction norms; fluctuating environments; developmental dynamics; evolutionary constraint

1. INTRODUCTION

A reaction norm function describes which phenotypes a single genotype will produce over a range of environmental conditions (Woltereck 1909). The term developmental reaction norm (Schlichting & Pigliucci 1998) has been coined to stress the fact that each morphological phenotype is a state of an ontogenetic or developmental dynamical system. Reaction norms for morphological traits that have completed their development show the attracting equilibrium states of the developmental dynamics over a range of environments. Discontinuities can occur (Roff 1996) and they correspond to environmentally controlled bifurcations in the developmental dynamical system. Such bifurcations correspond to qualitative changes in the dynamical behaviour that occur when a control variable is slightly altered. Thus, reaction norms for these traits are bifurcation diagrams that show the effect of an environmental control variable on the equilibrium states of the developmental dynamics. With so-called sharp or catastrophic bifurcations (Thom 1972; Kuznetsov 1998), the equilibria of the developmental dynamics change location abruptly, producing a discontinuous jump in the reaction norm. Bifurcations in reaction norms can also be non-sharp. This is the case, for example, when the bifurcation is of pitchfork type. By tracing such a reaction norm on a graph, one can then get from one branch onto another without making jumps. If developmental noise or measurement error blurs the deterministic part of the developmental dynamics, the non-sharpness makes it difficult to decide where one branch actually stops and the other one starts. With sharp reaction norm bifurcations, one can still assign phenotypes to a low number of discrete alternative morphs or reaction norm branches. Quantitative genetics links these

few alternative phenotypes to an ‘underlying’ liability scale that is continuous again (Gianola 1982) by assuming thresholds that control development into one of the alternatives. In this paper, the evolution of the bifurcation pattern is studied with phenotypes that are allowed to vary in a continuous manner. Contrary to current quantitative genetic methods, the method applied is not *a priori* restricted to sharp or catastrophic reaction norm bifurcations, and allows for non-sharp bifurcations as well.

The model presented is constructed around the phenomenon of seasonal polyphenism, i.e. substantial phenotypic variation with dependence on the state of a seasonally fluctuating environment. A famous example of seasonal polyphenism is the dry and wet seasonal forms in *Bicyclus* butterflies from East Africa (Brakefield & Reitsma 1991). When individuals use observations of an environmental variable with noisy cyclic variation as a cue for predicting the environmental state in the future, the reliability of that cue can vary (figure 1). Individuals observing the instantaneous environmental state at the top or the bottom of the cycle can use this cue in order to predict the trend of the immediate future rather accurately. If the environmental variable has an extremely high value, then it will most probably decrease in the near future. Individuals observing intermediate values of the environmental variable cannot reliably predict whether the cycle is in the increasing phase or decreasing phase. This type of variation in the reliability of the cue that individuals use to steer development and which provides information on the future state of the selective environment can make differences evolve in the number of morphological phenotypes developed per environment (Haccou & Iwasa 1995; Sasaki & Ellner 1995). Such differences imply reaction norm bifurcations with a branching aspect.

Figure 1. Cyclic environment and life histories. The environmental variable fluctuates according to a sinusoid function with uniform noise superimposed. The lag between the moment when individuals observe the state of the environment and the moment of optimizing selection determines whether the future is predictable for all observed environmental states (life history I) or whether predictability of the future will vary between different observation values (life history II).

The model presented in the following sections has clearly delineated modules for population ecology and the development of individual phenotypes. It is primarily strategic in nature, but the modular structure should make it easy to falsify the model in the context of any specific real-world system. It also facilitates subsequent updating of the modules with population ecological or developmental and genetic data. Two different versions of the same model are studied. The first version is without an explicit model for morphological development and is used to find the global evolutionarily stable strategy (ESS) of the population ecology. The second version has an explicit developmental map. It is used to investigate the effects and occurrence of evolutionary constraints.

2. ECOLOGICAL MODEL

In the population ecological part of the model, stabilizing viability selection on phenotype occurs. The viability of an individual is decided from a bell-shaped (Gaussian) curve. The width of the selection function is fixed per simulation of an evolutionary process, but the phenotype that is assigned maximum viability changes over time. The fluctuations in this optimal phenotype are of a cyclic character, with noise superimposed (as in figure 1). This fluctuating selection optimum is the environmental variable e_t in this model. At a certain point in time, all individuals from a cohort make an observation of the environmental variable and use this observation in order to predict what the selection optimum will be at a fixed distance in the future where they will experience viability selection. After selection, individuals reproduce and die. Clonal reproduction is assumed and generations are discrete. The generation time T is the sum of the lag l between observation and selection on the same individuals plus the time between selection and the moment the cohort of offspring individuals will observe the state of the world. This last interval is fixed at 41 time-units ($T = l + 41$) such that the cycle of the environmental variable will be sampled rather evenly in simulations.

The external environment has time-dependent states e_t ($t \in \mathbf{N}$) that are a combination of a deterministic sinusoid cycle and a random variable u_t :

$$e_t = 1.5 + \sin(2\pi t/360) + u_t. \quad (1)$$

The length of the deterministic part of the environmental cycle is 360 time-units. In order to produce a realistic environmental cycle, the superimposed noise is drawn from a bounded interval. Otherwise extreme outliers will occur in long time-series that are impossible for any physical environmental variable. Instead of using a more complex truncated distribution with several parameters for the environmental noise u_t , I opted for a simple uniform distribution on a bounded interval. The random effect u_t is drawn from a uniform distribution on the interval $[-0.4, 0.4]$, such that the range of e_t becomes $[0.1, 2.9]$. Many distributions become more rectangular and uniform-like when the tails of the distribution are truncated.

From the results of Haccou & Iwasa (1995) and Sasaki & Ellner (1995), a global ESS is expected for the population ecology of this model. When the environmental variable is drawn from a bounded interval and for viability selection according to a Gaussian curve, the ESS reaction norm will be composed of one or several discrete strategies per environment, i.e. one or several reaction norm branches. Whether branching reaction norms will be produced or not does not seem to depend critically on the distribution of the environmental noise u_t . When the intensity of stabilizing selection can be made arbitrarily small or large, such as with a Gaussian selection curve, the results of Haccou & Iwasa (1995) suggest that, for many distributions of environmental noise, a reaction norm with a single branch or several ones can be produced if the selection intensity parameter is sufficiently decreased or increased, respectively.

3. FIRST MODEL VERSION: DETERMINATION OF GLOBAL EVOLUTIONARILY STABLE STRATEGIES

This first implementation investigates the dependence of reaction norm evolution on the population ecology without restricting the developmental system that produces phenotypes to a specific functional form. There is no explicit genotype–phenotype or developmental map in the first model version and reaction norms are modelled as a grid or a matrix of probabilities of observation–prediction combinations of selection optima. A specific reaction norm function might be insufficiently flexible and unable to take the shape favoured by the population ecology.

(a) Strategy matrices

Each strategy is specified by a matrix $P = [p_{\text{pred,obs}}]_{25 \times 25}$ of probability weights on different observation–prediction combinations of values of environmental states (see figure 2). For instance, $p_{1.5,0.375}$ is the probability that a selection optimum with a value of 1.5 is predicted when environmental value $e_{\text{obs}} = 0.375$ is observed. For each value of e_{obs} that corresponds to a column in the matrix, probabilities $p_{e_{\text{pred}},e_{\text{obs}}}$ add up to 1. Since strategies are specified on a

Figure 2. Mutational mappings picture gallery. In the first model version, strategies are specified by matrices with probability weights of observation–prediction combinations of environmental states. Such matrices can be drawn as a grid of combinations with the intensity of filling directly proportional to the probability weight in that matrix cell. Mutations are modelled as random ‘deformations’ of probability weights in the matrix. Each mutation mapping is given a name that describes the way a probability weight is moved across cells. (a) A progenitor reaction norm and (b) a number of one-step mutants are drawn. The progenitor was obtained from a flat reaction norm by applying two successive and well-aimed bubble mappings.

matrix or a grid, values of the environmental variable e_t have to be assigned to a category in a list of possible observed values e_{obs} . The length of this list is arbitrary and is determined by the resolution in the approximation to the reaction norm function that one wants to obtain. In this case, the set contains 25 equally spaced values, ranging from $e_{\text{obs}} = 0$ to $e_{\text{obs}} = 3$. When an environmental state occurs in between two observation values specified in the strategy matrix, the prediction from one of the two nearest specified observation values is used, with relative probabilities inversely proportional to the distances between actual observation and specified observation. Thus, the environmental state e_t from the original distribution (equation (1))

changes through observation into state e_{obs} , which equals one of the categories in the strategy matrix. For example, the maximum value $e_t = 2.9$ that occurs in the environmental cycle modelled (equation (1)) is assigned to $e_{\text{obs}} = 2.88$ with probability 0.83 and to $e_{\text{obs}} = 3$ with probability 0.17. For phenotypic values above $e_{\text{pred}} = 2.9$ (and below a value of zero), the phenotypes from the range between 0.1 and 2.9 will always be closer to the selection optimum and better predictions in terms of viability. Therefore, the range of values modelled in the strategy matrix will contain all phenotypes in the evolutionarily attracting or ESS reaction norm up to the approximation error introduced by the discretization of the reaction norm range.

Mutational effects are modelled as simple random deformations of the pattern of probabilities on the grid (figure 2 and Appendix A). Mutational mappings on matrix strategies all work in a similar manner. They move the probability mass between cells in the list by means of scrolling cells with non-zero content back or forth between different prediction values corresponding to a single observation value (this means scrolling cells some rows up or down in the same matrix column), by collecting the probability mass of different cells within the same column in a single matrix cell or by dividing the probability mass from a single cell over different ones. No probability mass is ever moved across the minimum or maximum e_{pred} cell in a column. When there are non-zero cells in the first or the last row, the scrolling is stopped (see, for example, mutation tilt in figure 2).

The maximum number of cells scrolled, focal cell and range of environmental states [$e_{\text{obs}1}$, $e_{\text{obs}2}$] affected are chosen at random for each mutation. Sometimes the randomly picked range has such a small value that probability masses only change in a single column of the matrix, i.e. for a single observation value of the environment.

Therefore, with the set of mutational maps discussed in Appendix A all distributions of probability masses over the matrix can be reached through a series of mutations. In most cases the range of observation values affected is much larger than one matrix column. This is done to speed up evolution towards a reaction norm that at least approximates a shape that would be favoured in the evolutionary process. Moreover, most mutational deformations also have locally smooth effects such that mutations can modify a pattern without losing most of the information in it.

Considering the set of mutational mappings, constraints with respect to branching of reaction norms are unlikely to be present. Both small adjustments and large mutational jumps occur such that the evolutionary process will not get trapped easily in local selective optima.

(b) *Viability selection*

In this first model version, for an environmental state e_t observed as $e_{\text{obs},t}$ at time t and selection occurring at time $t+l$, the selection function S (equation (2)) is the sum of probability weights $p_{\text{pred},\text{obs},t}$ put on each prediction value multiplied by a Gaussian function used as a distance measure between prediction and the actual state of the environment:

$$S(P, e_{\text{obs},t}, e_t + l) = \sum_{e_{\text{pred}}} p_{\text{pred},\text{obs},t} \exp(- (e_{\text{pred}} - e_{t+l})^2 / 2\sigma^2). \quad (2)$$

The inverse of the variance parameter of the Gaussian, $\rho = 1/\sigma^2$, corresponds to the strength of selection towards the time-dependent selection optimum.

(c) *Evolutionary dynamics of strategy matrices*

In this kind of ecological scenario with variation in the reliability of an environmental cue, a branching reaction norm will appear when the ESS or the evolutionarily attracting reaction norm produces a single phenotypic strategy for some environmental states and a mixture of several strategies in other environments. In terms of a

strategy matrix, this implies that columns will obtain different numbers of cells with non-zero probability mass.

In order to assess the invasion fitness (Metz *et al.* 1992) of a mutant type in a resident population, the following quantity is compared between mutant P_m and resident P_r strategy matrices:

$$F(P) = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_0^{t-1} \delta_{0,t \bmod T} \ln S(P, e_{\text{obs},t}, e_{t+l}). \quad (3)$$

This is the long-term average of the selective values S of a strategy originating at time $t=0$. It is assumed that the mechanism of density dependence does not affect selection between different strategies and that fecundity parameters are fixed. The average in equation (3) is taken over different generations with an observed environment $e_{\text{obs},t}$ and a selection optimum e_{t+l} for each generation. These have a fixed lag of l time-units between them. For this reason, Kronecker delta $\delta_{0,t \bmod T}$ is inserted into the equation. This function picks out the selection functions at time-steps that are equally spaced by the generation time T .

The ESS of this evolutionary random walk maximizes F (Metz & Mylius 2001) and, at the same time, is evolutionarily attracting (Rand *et al.* 1994). It is a global ESS with a set of alternative strategies per observation environment whenever the reliability of that observation value for prediction is low (Haccou & Iwasa 1995; Sasaki & Ellner 1995). Discontinuous reaction norms with a single phenotypic strategy per environment and an abrupt change, as in classical threshold traits (Gianola 1982; Roff 1996), are not expected to be the ESS in this ecological scenario. The cycle in the environmental variable is smooth, such that an evolved reaction norm should either track the environmental variable smoothly or hedge its bets over different phenotypes. At extreme values of the environmental variable, the reliability of the observed cue will often be larger than for intermediate values of e_{obs} , as explained in §1. If branching reaction norms occur, then they will hedge their bets (Seger & Brockmann 1987) for intermediate values of the environmental variable while still producing a single phenotype at the extremes.

An approximation to a maximum of $F(P)$ was found by simulating evolutionary random walks. A ‘flat’ reaction norm with prediction $p_{1.5,\text{obs}}=1$ for all possible observation values e_{obs} was used as the initial state of the evolutionary system. All mutation mappings were applied to the strategy in each evolutionary step, giving 13 new mutants. Invasion fitnesses were estimated over 150 generations and the best of all new mutants and the former ‘resident’ became the new state of the evolutionary random walk. This approach was chosen in order to make the simulations faster in approaching the (global) ESS. Thirty such steps were made in each evolutionary random walk. This is sufficient for observing whether evolution favours a branching reaction norm or not.

Simulations were run for different combinations of selection strength ($\rho=0.5-10$) and time-lag ($l=20-400$) between observation and selection (figure 3a). The likelihood that an evolved strategy puts probability weights on different prediction values when observing an intermediate state of the environment increases with selection

Figure 3. The occurrence of bet-hedging in reaction norms. (a) Ten evolutionary random walks were simulated with strategies specified by matrices for each combination of selection variance and lag between observation and prediction. The proportion of strategies with bet hedging at intermediate states of the environment out of the ten simulated ones is indicated. (b) Reaction norms for numbered pairs of selection strength and time-lag (selection strength 2 or 8) $\times$ (lag 60 or 160 time-units). They show the grid-based reaction norm with the highest invasion fitness out of the ten simulated ones used to draw each point in (a).

strength. Figure 3a shows that they only evolve when the selection intensity is sufficiently strong, such that a small deviation from the selection optimum already results in a low viability. Bet-hedging does not evolve when the lag between observation and selection is a multiple of half the period of the environmental cycle. The reliability of the environmental cue then stays the same for whatever state of the environment is observed. Individuals from a clonal line in successive generations do not perceive the sinusoid environmental cycle and only experience the environmental noise u_t . Since the generation time T is

kept constant, the pattern of occurrence of bet-hedging repeats itself with a phase equal to half the cycle of the environmental variable. All evolved reaction norms have a single phenotype at the extremes of the environmental distribution and, therefore, all strategies with bet-hedging at intermediate e_{obs} correspond to bifurcating and branching reaction norms. Part 1 of figure 3b shows that such a reaction norm puts probability weights on two discrete alternatives and suggests that—when accounting for the restricted resolution imposed by the strategy specification—the reaction norm bifurcation is non-sharp. If the selection strength is made much larger than the range pictured in figure 3a (when $\rho > 50$ approximately), bet-hedging occurs for all values of lag between observation and selection. These reaction norms often produce two distinct reaction norm branches over the whole range of environmental values e_{obs} . As such, they are not bifurcating within the specified range of e_{obs} . With extremely large selection intensities, the penalty for a deviation from the selection optimum becomes so large that, evolutionarily, it starts paying off to always hedge your bets.

4. VERSION TWO: EVOLUTIONARY CONSTRAINTS ON REACTION NORM BRANCHING

Once the possibility of branching is demonstrated, it remains fruitful to model the evolving system with a developmental submodel. This will enable us to assess the nature of evolutionary constraints that might arise when a specific model for development is assumed and when evolution operates more locally by modifying (in small steps) the parameters of the dynamical system.

(a) Developmental dynamics

The developmental dynamical system assumed should be able to switch between producing a single phenotype per environment and at least two distinct ones such that the occurrence of reaction norm branching is not excluded *a priori*. I modified a model previously used in understanding the mechanism behind spatial patterning in embryos (Lewis *et al.* 1977). It models the effect of a signalling substance on the synthesis of a gene product and can be interpreted as a model for the establishment of the equilibrium level of a morphogenetic substance in a butterfly wing focus. Diffusion of this substance throughout surrounding tissue can then control the size of a wing spot (Lewis *et al.* 1977; Brakefield *et al.* 1996) and different wing spot sizes are typically associated with seasonal polyphenism in butterflies.

Assuming that the equilibrium level of morphogen controls a morphological trait, this stable equilibrium level can be seen as a phenotypic trait in itself. The morphology on which selection acts is assumed to correlate perfectly with it such that, for simplicity, one can assume that selection occurs on the equilibrium level of morphogen. The developmental model for individual phenotype then becomes a nonlinear ordinary differential equation (ODE) for the concentration g of a morphogenetic gene product, with equation

$$dg/dt = as + bs^2 + c - mg + fg^2/(h + g^2). \quad (4)$$

Figure 4. Developmental noise or reaction norm branching. These reaction norms evolved in version 2, with an explicit developmental map (ecological parameters $\rho=8$ and $l=60$, corresponding to part 1 in figure 3*b* and point 1 in figure 3*a*). In order to make comparison of figures 3 and 4 easier, reaction norms are represented as the probabilities of occurrence of e_{obs} , e_{pred} combinations. The development of 40 individuals is simulated for each value of e_{obs} in order to generate these probabilities. Some individuals ($< 10\%$ for each e_{obs}) have a phenotype larger than $e_{\text{pred}}=3$ and are not represented. For the strategy in (a), reaction norm branching is constrained. The strategy in (b) is unconstrained with respect to branching, but still does not reach the shape that is found in the grid-based approach by far (see figure 3*b*). (a) Developmental noise: initial reaction norm (1.5, 0.5, 0, 5, 5, 5 and 0.1) and parameters evolved (1.5, 0.5, 0.2, 3.1, 4.6, 2.7 and 0.4). (b) Bifurcation and developmental noise: initial reaction norm (0, 3.5, -0.3 , 10, 11, 1 and 0.1) and parameters evolved (0, 5.4, -0.3 , 10, 11, 0.35 and 0.61).

In this equation, s is the level of a globally present internal signal. It could be a hormone titre for example. The other parameters in equation (4) are assumed to be properties of the developmental system under genetic control. The last term with parameters f and h represents autocatalysis of gene product synthesis. Parameter m is the sum of the decay rate of gene product and the diffusion rate into surrounding tissue. In comparison with the model of Lewis *et al.* (1977), the dependence of gene product on signal level s is elaborated more as a quadratic function with parameters a and b . This functional form allows a flexible dependence of gene product synthesis on signal level. When parameters a and b evolve to zero values, synthesis becomes independent of signal. With a zero value for b , the dependence is linear but synthesis can increase ($a > 0$) or decrease ($a < 0$) with increasing signal level. When b is non-zero, the amount of synthesis can be maximal or minimal for intermediate signal levels. Note that this ODE can have multiple steady states with positive values even with independence of the signal s , i.e. when parameters a and b are zero and production rate c is non-zero.

It is also assumed that behaviour and physiology affect the translation of the environmental state e_t external to the organism into an internal signal variable s . This effect is modelled as a stochastic effect ε per individual such that $s = \varepsilon e_t$. When many independent small stochastic effects affect signal translation in a multiplicative manner, this results in a log-normally distributed error. The observation translation effect ε is therefore a random number drawn from a log-normal distribution with parameters $\mu_{\ln(\varepsilon)} = 0$ and $\sigma_{\ln(\varepsilon)}^2 = \nu$ for each individual. Parameter ν represents the amount of stochasticity that signal translation introduces and is assumed to be under genetic control. A model (equation 4) without this stochasticity would still be able to produce a reaction norm with a

sharp bifurcation, but a multimodal distribution of phenotypes per value of e_t would not be possible. The translation noise is required to set development on trajectories with very different outcomes for a single state of the external environment e_t .

Different strategy types are thus different parameter combinations (a , b , c , m , f , h and n): these are the parameters of the ODE (equation (4)) and the variance ν introduced by signal translation. A steady state g^* of this ODE, which is determined from a specific strategy specification, an individual observation of the environment $s = \varepsilon e_t$ and initial condition $g(0) = 0$, is used as a prediction of the environmental state, i.e. $e_{\text{pred}} = g^*$. Negative states of gene product phenotype g are set at zero ($g^* \in \mathbf{R}^+$) and the maximum phenotype is set at 10. Phenotype values above 10 have very low and similar viabilities for any value of selection optimum occurring such that the effect of this upper limit on the evolutionary trajectory is negligible.

Summarizing, the qualitative types of reaction norms that can evolve are (i) continuous with low developmental variation, (ii) unimodal distributions of phenotypes per environment with high developmental variation but continuous appearance, (iii) unimodal with a sharp bifurcation and low developmental variation, (iv) unimodal with a sharp bifurcation and high developmental variation, and (v) a sharp bifurcation with branching and with high developmental variation.

Although a bifurcation analysis shows that reaction norm branching is possible and that discontinuous reaction norms can also occur, the reaction norm shapes that this developmental system (equation (4)) allows are restricted. For instance, a branching pattern will involve a sharp bifurcation (see Lewis *et al.* 1977) and branching reaction norms cannot occur for very low values of signal translation variance ν .

(b) *Evolutionary dynamics*

The development of 40 individuals is simulated for each generation of the population ecology in order to generate an empirical probability distribution of prediction values of the environmental variable. All individuals have the same strategy parameters and experience the same environment e_t . A new value of the translation error e_i is drawn for each individual and the equilibrium g_i^* ($i=1, \dots, 40$) of each developmental dynamics is determined by means of numerical approximation. In this model version, the selection function per generation is the average of the selective values $\exp(-(g_i^* - e_{t+1})^2/2\sigma^2)$ of these 40 individuals that observed the environmental state e_t at time t . Invasion fitnesses are again estimated over 150 generations.

Evolutionary random walks are of 2000 steps. One mutant is generated in each step by mutating a single parameter of the haplotype. Non-zero mutational covariances are not modelled such that this type of constraint cannot be present. Mutational effects are either from a normal distribution or a log-normal one. The parameters of the developmental dynamics that are allowed to vary over positive and negative values, parameters a , b and c , are mutated by adding a normally distributed effect (mean value 0 and variance 0.25). Other strategy parameters are restricted to positive values, namely the signal translation variance ν , the decay rate m and the parameters f and h of the autocatalysis term in equation (4). These parameters are mutated by multiplying them with a log-normally distributed factor (the logarithm of this mutational random variable is $N(0, 0.25)$). Evolutionary random walks were studied for different ecological parameters and different strategies were used as initial states of the evolutionary random walks.

In this second version of the evolutionary model, the same reaction norms are obtained as with the grid-based approach when a continuous reaction norm is favoured by the population ecology (as in parts 2–4 of figure 3*b*). Bifurcating reaction norms almost never evolve when they are expected from the results in figure 3*a*, although many regions in the trait space of the parameters of the developmental model do produce them. Instead of the separate branches of a bifurcation, a large amount of random variation in phenotypes evolves over a range of environments (figure 4*a*). This variation is clearly adaptive for two reasons: (i) the reaction norm has been selected in an evolutionary process, and (ii) the developmental noise disappears again when the observation–selection lag is set at values that do not produce branches in the grid-based approach. Because we studied the grid-based version first, we know that the selection of this large amount of random variation must be the consequence of implicit constraints from the developmental model and not the best possible solution in terms of the population ecology. I calculated invasion fitnesses for a range of mutations in each parameter of the strategy in figure 4*a* (within a ± 0.5 change in each parameter). The reaction norm appears to be at a local selective optimum for all these parameters. It is stuck at a local compromise ESS imposed by the developmental system and will never reach the global ESS favoured by the population ecology.

It is possible to make a bifurcation evolve. When the initial reaction norm has a zero initial value for parameter c

and this parameter is not allowed to evolve, branching sometimes appears together with substantial developmental noise over a range of environments (figure 4*b*). This constraint in mutational variation makes the developmental dynamics very dependent on the value of the environmental signal. Bet-hedging does not occur at completely intermediate values of the environmental variable for the reaction norm in figure 4*b* and a lot of developmental noise is observed for large values of e_t . Although the evolution of reaction norm branching has not been constrained (paradoxically, by introducing an additional mutational constraint), the shape of the reaction norm does not match the ecologically most favoured shape by far (compare part in figure 3*b* and figure 4*b*). A perfect match is impossible with the model used.

5. DISCUSSION

The model versions in this paper integrate concepts from game theory (Maynard Smith 1982) and adaptive dynamics (Metz *et al.* 1996), with either a purely algorithmic approach (version 1) or with a nonlinear model for the developmental dynamics (version 2). Using results in both versions, it is established that bifurcating reaction norms or developmental noise can evolve in variable environments as a kind of environment-specific stochastic polyphenism (Walker 1986). When the reliability of the cue that is used to steer development varies sufficiently over the range of possible environments and provided that selection is sufficiently strong but not too strong, bifurcating reaction norms can evolve.

However, evolutionary constraints induced by a specific developmental system can halt evolution at a reaction norm with substantial developmental noise and no branching. In the second model version, the constraint with respect to branching originates from the existence of local selective optima that occur for all seven parameters of the developmental dynamical system. In comparison with the first model version, evolution is not as free in the mutational steps that can occur. The parameters' values are modified in small steps, while the functional specification of the ODE (equation (4)) remains fixed. Mutations always affect phenotypes developed over a range of environments. Constraints are considered (i) with respect to the evolution of reaction norm branching, and (ii) with respect to perfect matching of the ecologically favoured reaction norm. Without comparison of this version with the matrix-based approach, it would be difficult to spot constraints. The only available information would then come from the evolution of parameters in the specific developmental system and with the effects of the developmental map and population ecology tangled in expressions for invasion fitness. The specified developmental dynamics is first of all used to investigate constraints with respect to branching (type (i)). The model simply cannot match the favoured branching reaction norms completely. For instance, branching requires developmental noise, which also reduces the match of figure 4*b* with part (1) of figure 3*b* at large values of the environmental variable. In order to allow the evolution of a reduced developmental noise at these large values, the model structure should be different. However, the constraint with respect to branching (constraint type (ii)) can be controlled within

the model structure assumed. Branching did not occur when a signal was not required to initiate development (i.e. when $c > 0$). The constitutive synthesis rate c and the other parameters then evolve in such a way that an ecologically more advantageous branching reaction norm is not within mutational reach. The constraint comes into being by a mixture of selective and genetic reasons precluding the production of a mutant that approximates the global ESS better (Maynard Smith *et al.* 1985). If non-zero mutational covariances were modelled starting from the local ESS shown in figure 4a, a successful mutant that decreases constitutive synthesis while increasing the slope of the dependence on signal level might cause a branching reaction norm. Non-zero mutational covariances might even make it that the evolutionary trajectory would not be attracted to the compromise ESSs of figure 4 at all.

When the evolution of branching was constrained, a large amount of stochastic developmental noise evolved. It can be considered a compromise or an alternative when branching is globally favoured but locally not possible anymore.

In the context of stochastic strategies, it remains a major issue to determine how individuals hedge their bets. How do they perform adaptive coin flipping (Kaplan & Cooper 1984)? I conjecture that sloppy observation or low robustness of biochemical pathways (Barkal & Leibler 1997) will be the rule when bet-hedging is necessary. Amplifying a small amount of random variation that is already present is probably evolutionarily easier to obtain than constructing a stand-alone biological random number generator. The origin of developmental noise in the system was modelled with such a mechanism in mind. Note that the developmental noise is not in any way related to the presence of stressful factors during development (Møller & Swaddle 1997). It also needs to be noted that the strategies of the developmental submodel are never unstable (Simons & Johnston 1997), but stable equilibrium attractors of the developmental dynamics. Different initial conditions with a small amount of stochasticity simply produce different stable phenotypes.

The polynomials previously used to characterize reaction norms (Kirkpatrick & Heckman 1989; De Jong 1990; Gavrillets & Scheiner 1993; Sasaki & de Jong 1999) do not allow for reaction norm bifurcations directly. The quantitative genetic method that assumes a threshold trait with an underlying continuous liability scale applies to phenotypes with a discrete distribution and with a low number of discrete morphs. These correspond to reaction norms with sharp bifurcations. Quantitative genetics cannot deal easily with non-sharp bifurcations in reaction norms. The usual assumption of normally distributed phenotypes or liabilities per environment simply cannot always hold good and reducing the data to a description in terms of few phenotypic categories implies a loss of phenotypic information with potential ecological relevance. It will be necessary to extend existing methods with, for instance, the Bayesian approach used for the analysis of mixtures of distributions (Richardson & Green 1998). Moreover, in order to understand the effects of environmental variation and constraints better, it might be more fruitful to estimate parameters of nonlinear developmental dynamical systems and their

genetic and environmental variation instead of parameters of linear functions of mature phenotypes.

6. CONCLUSION

Population ecological scenarios exist that favour branching reaction norms with non-sharp bifurcations. Within the limits of a specific developmental system, the evolution of branches can be constrained and, if possible, substantial developmental noise is likely to evolve instead.

Ulf Dieckmann, Tom van de Castele and an anonymous referee gave helpful comments on a previous version of this manuscript. This research was supported by the Fund for Scientific Research (FWO-Flanders) and by grant GOA-BOF-UA-1999-2003.

APPENDIX A

Mutational mappings are represented in figure 2 and are listed here for reference.

The mutational mappings shift(up) and shift(down) scroll probabilities towards higher (lower) values of e_{pred} . The cells in a matrix column are scrolled one row up or down at a time and in columns over a randomly drawn range. Scrolling continues until the first (last) row has non-zero content or until the effect size of the mapping in terms of the maximum number of cells scrolled is reached.

The mappings tilt and parabolic produce effects that approximate the effect of changing the slope of a reaction norm or changing the curvature. The mappings tilt(up) and tilt(down) pick a focal cell and scroll cells with smaller and larger e_{obs} in different directions. For example, tilt(up) scrolls cells with a larger e_{obs} towards larger values of e_{pred} . The effect increases with distance from the focal e_{obs} . Parabolic(up) and parabolic(down) work in a similar manner, but all cells are scrolled in one direction.

The following mutations change the number of cells with probability content per environment. Bubble divides the probability mass over two cells and scrolls them away from a focal cell. The scrolled number of cells decreases to zero at a random distance from the focal e_{obs} , which defines the range of bubble. Lump collects the probability mass from two cells into a single cell over a random range of e_{obs} categories. Split(up) and split(down) pick a focal cell with non-zero content. Rows on either side of the focal e_{pred} are scrolled towards the top and bottom rows. A range of affected observation values is randomly chosen. A random portion of the cell content is scrolled upwards in the focal row, whereas the rest is scrolled to the bottom row. Split(up) moves most of the content into cells at larger e_{pred} values.

The mappings stretch and shrink change the maximum distance (in terms of prediction values) between cells with non-zero content. Stretch divides the range of e_{pred} into two intervals. Cells are scrolled towards the edges. Shrink decreases the range of e_{pred} over which a non-zero probability mass occurs by scrolling cells in a column towards a focal matrix row.

Allocate redistributes the probability mass between cells with non-zero content.

The randomness in a mutational effect is introduced as follows. A focal cell is chosen according to a uniform

distribution over the strategy matrix (either a row or a column or both). The maximum and minimum effect sizes are fixed integers for each type of mutation. When effects are arbitrarily large, very large effect sizes often scroll all the probability mass into the cells that will produce extreme phenotypes. The resulting strategies are not expected to be selectively favoured and a fixed maximum effect size avoids them. Mutations that often leave the matrix unchanged simply waste simulation time. A minimum effect is specified such that almost all mutations have a visible effect. The distribution of the integer difference between applied effect and minimum effect is binomial with a mean value of 0.5 (maximum effect minus minimum effect).

REFERENCES

- Barkal, N. & Leibler, S. 1997 Robustness in simple biochemical networks. *Nature* **387**, 913–917.
- Brakefield, P. M. & Reitsma, N. 1991 Phenotypic plasticity, seasonal climate and the population biology of *Bicyclus* butterflies in Malawi. *Ecol. Entomol.* **10**, 291–303.
- Brakefield, P. M., Gates, J., Keys, D., Kesbeke, F., Wijngaarden, P., Monteiro, A., French, V. & Carroll, S. B. 1996 Development, plasticity and evolution of butterfly eyespot patterns. *Nature* **384**, 236–242.
- de Jong, G. 1990 Quantitative genetics of reaction norms. *J. Evol. Biol.* **3**, 447–468.
- Gavrilets, S. & Scheiner, S. M. 1993 The genetics of phenotypic plasticity. V. Evolution of reaction norm shape. *J. Evol. Biol.* **6**, 49–68.
- Gianola, D. 1982 Theory and analysis of threshold characters. *J. Anim. Sci.* **54**, 1079–1096.
- Haccou, P. & Iwasa, Y. 1995 Optimal mixed strategies in stochastic environments. *Theor. Pop. Biol.* **47**, 212–243.
- Kaplan, R. H. & Cooper, W. S. 1984 The evolution of developmental plasticity in reproductive characteristics: an application of the 'adaptive coin-flipping' principle. *Am. Nat.* **123**, 393–410.
- Kirkpatrick, M. & Heckman, N. 1989 A quantitative genetic model for growth, shape, reaction norms, and other infinite-dimensional characters. *J. Math. Biol.* **27**, 429–450.
- Kuznetsov, Y. A. 1998 *Elements of applied bifurcation theory*, 2nd edn. New York: Springer.
- Lewis, J., Slack, M. W. & Wolpert, L. 1977 Thresholds in development. *J. Theor. Biol.* **65**, 579–590.
- Maynard Smith, J. 1982 *Evolution and the theory of games*. Cambridge University Press.
- Maynard Smith, J., Burian, R., Kauffman, S., Alberch, P., Campbell, J., Goodwin, B., Lande, R., Raup, D. & Wolpert, L. 1985 Developmental constraints and evolution. *Quart. Rev. Biol.* **60**, 265–287.
- Metz, J. A. J. & Mylius, S. 2001 When does evolution optimize? In *Adaptive dynamics in context* (ed. J. A. J. Metz & U. Dieckmann). Cambridge University Press. (In the press.)
- Metz, J. A. J., Nisbet, R. M. & Geritz, S. A. H. 1992 How should we define fitness for general ecological scenarios? *Trends Ecol. Evol.* **7**, 198–202.
- Metz, J. A. J., Geritz, S. A. H., Meszéna, G., Jacobs, F. J. A. & Van Heerwaarden, J. S. 1996 Adaptive dynamics: a geometrical study of the consequences of nearly faithful reproduction. In *Stochastic and spatial structures of dynamical systems* (ed. S. J. van Strien & S. M. Verduyn Lunel), pp. 183–231. Amsterdam: KNAW Verhandeligen, North Holland.
- Møller, A. P. & Swaddle, J. P. 1997 *Asymmetry, developmental stability, and evolution*. Oxford University Press.
- Rand, D. A., Wilson, H. B. & McGlade, J. M. 1994 Dynamics and evolution: evolutionary stable attractors, invasion exponents and phenotype dynamics. *Phil. Trans. R. Soc. Lond.* **B343**, 261–283.
- Richardson, S. & Green, P. 1998 On Bayesian analysis of mixtures with an unknown number of components. *J. R. Stat. Soc. B* **59**, 731–792.
- Roff, D. A. 1996 The evolution of threshold traits. *Quart. Rev. Biol.* **71**, 3–35.
- Sasaki, A. & de Jong, G. 1999 Density dependence and unpredictable selection in a heterogeneous environment: compromise and polymorphism in the ESS reaction norm. *Evolution* **53**, 1329–1342.
- Sasaki, A. & Ellner, S. 1995 The evolutionarily stable phenotype distribution in a random environment. *Evolution* **49**, 337–350.
- Schlichting, C. D. & Pigliucci, M. 1998 *Phenotypic evolution: a reaction norm perspective*. Sunderland, MA: Sinauer Associates.
- Seger, J. & Brockmann, H. J. 1987 What is bet-hedging? *Oxford Surv. Evol. Biol.* **4**, 182–211.
- Simons, A. M. & Johnston, M. O. 1997 Developmental stability as a bet-hedging strategy. *Oikos* **80**, 401–406.
- Thom, R. 1972 *Stabilité structurelle et morphogénèse*. New York: Benjamin.
- Walker, T. J. 1986 Stochastic polyphenism: coping with uncertainty. *Florida Entomol.* **69**, 46–62.
- Woltereck, R. 1909 Weiterer experimentelle Untersuchungen über Artveränderung, Speziell über das Wesen Quantitativer Artunterschiede bei Daphniden. *Versuch. Deutsch. Tsch. Zool. Gesellschaft* **19**, 110–172.

As this paper exceeds the maximum length normally permitted, the author has agreed to contribute to production costs.