

Phenotypic diversity as an adaptation to environmental uncertainty

Matina C. Donaldson-Matasci^{*†‡} Michael Lachmann[‡]
Carl T. Bergstrom[†]

May 18, 2008

Abstract

Question: What is the general quantitative relationship between adaptive phenotypic diversity, or bet-hedging, and the environmental uncertainty that selects for it?

Mathematical methods: Building on the fitness set approach introduced by Levins, we develop a graphical heuristic for determining the optimal amount of diversity in a fluctuating environment. We use as our optimality criterion the expected long-term growth rate of a lineage.

Key insights: Each of the phenotypes in a polyphenic population may be seen as investing a certain proportion of its reproductive effort in each of the possible environments. A bet-hedging lineage that produces the phenotypes in just the right proportions—so that the overall reproductive investment in each environment matches the environmental frequencies—grows faster on average than other lineages. How much faster it grows than the resident population, and thus the strength of selection towards the optimal bet-hedging strategy, depends on how far the residents are from the optimal investment profile.

Predictions: A rigorous empirical demonstration that bet-hedging is adaptive requires a comparison of the degree of phenotypic diversification in similar populations subject to varying levels of environmental uncertainty. We confirm that bet-hedging should be observed only within a certain range of environmental variation; when the environment is more predictable than this, a phenotypic generalist would

*Corresponding author: donaldso@eva.mpg.de

†University of Washington, Seattle, WA, USA

‡Max Planck Institute for Evolutionary Anthropology, Leipzig, Germany

do better. We furthermore give a simple method to calculate this range, based on the shape of the fitness trade-offs. Within this range, we predict a linear relationship between the frequency of phenotypes and the frequency of environments, independent of the shape of the trade-offs.

1 Introduction

Organisms inhabiting a changeable, unpredictable environment face a difficult adaptive challenge. In some cases they may evolve to specialize in one particular niche, making up for poor performance in some conditions by over-achieving in another. In other cases environmental uncertainty may select for a “jack-of-all-trades”, a generalist capable of overcoming any environmental hurdle but never reaching the peak productivity of a more specialized type. In this paper we focus on two different kinds of strategies that can function as adaptations to environmental uncertainty. A traditional generalist is an individual of a single, fixed type whose morphology and/or behavior is reasonably well equipped for any situation. A bet-hedging strategy, on the other hand, produces a phenotypically heterogeneous set of individuals, each of which may develop into one of several specialized types (Cooper and Kaplan, 1982; Seger and Brockmann, 1987). These two strategies may be viewed as two different ways of being a generalist: one strategy is employed by individuals, while the other is employed by genotypes. The fitness set framework introduced by Levins (1962) can be used to provide a basis for comparing the two kinds of generalists, and build intuition about the general circumstances giving rise to each.

When environmental uncertainty selects for a bet-hedging strategy, it creates non-genetic phenotypic diversity in the population. This may lead us to expect some relationship between the amount of environmental variation and the amount of phenotypic variation that is created as an adaptation to it. Such a connection has in fact been drawn for the simplest models, but a more general correspondence has been elusive (Bergstrom and Lachmann, 2004; Kussell and Leibler, 2005). We show here that the connection can be easily generalized when the phenotypes are described according to their relative fitness contributions in the different types of environments. These descriptions have a simple graphical interpretation in terms of Levins’s fitness sets, and may be interpreted as a measure of specialization. Using this measure, we develop an intuitive understanding of the quantitative relationship between environmental uncertainty and adaptive generalization.

2 Background: Levins’s fitness sets and evolution in an uncertain environment

Whether organisms adapt to environmental uncertainty by becoming specialists or generalists—or a mixture of both—depends on a number of interacting factors (Levins, 1962; Wilson and Yoshimura, 1994). First of all, the trade-off between adaptations to different environmental conditions may take different forms. Weak trade-offs, which allow intermediate types to perform fairly well in all conditions, tend to promote the evolution of generalists. Strong trade-offs, which make intermediate types perform poorly in all conditions, tend to promote the evolution of specialists. Secondly, the structure of environmental variation plays a key role. Environmental change that affects all individuals in the population at once on the time scale of a generation, like years of drought, favors the evolution of generalists. In contrast, environmental variation between individuals in a single generation, like random dispersal into small patches of different habitats, encourages the evolution of specialists. Although some authors use the terms “temporal” and “spatial” to distinguish these two kinds of environmental variation, we prefer the more general labels *population-level* and *individual-level* environmental variation.

Levins introduced the paired concepts of the fitness set and the adaptive function as a way to gain intuition about the way that fitness trade-offs and environmental variation interact to influence the evolution of specialists and generalists (Figure 1). Each achievable phenotype is characterized by its fitness profile over the different environments, and then plotted on a graph whose axes are the fitness in each environment. The set of all such points in fitness space is called the *fitness set*; the portion of its boundary with negative slope is a graphical representation of the trade-off between specializing in different environments. *Weak trade-offs* create convex fitness sets, while *strong trade-offs* create concave fitness set. The *adaptive function* describes how fitness in each environment contributes to the population’s overall reproductive rate. Individual-level environmental variation yields populations which reproduce according to the arithmetic mean fitness in every generation. Population-level environmental variation yields populations which vary in their productivity from generation to generation. Over the long term, such populations reproduce according to their geometric mean fitness (Dempster, 1955; Levins, 1962; Cohen, 1966). The fitness set and the adaptive function can be used together to identify a phenotype that is optimal—but only in the sense that it maximizes the population’s reproductive rate. (We will return to this point.)

One interesting result, which is easy to understand using Levins’s frame-

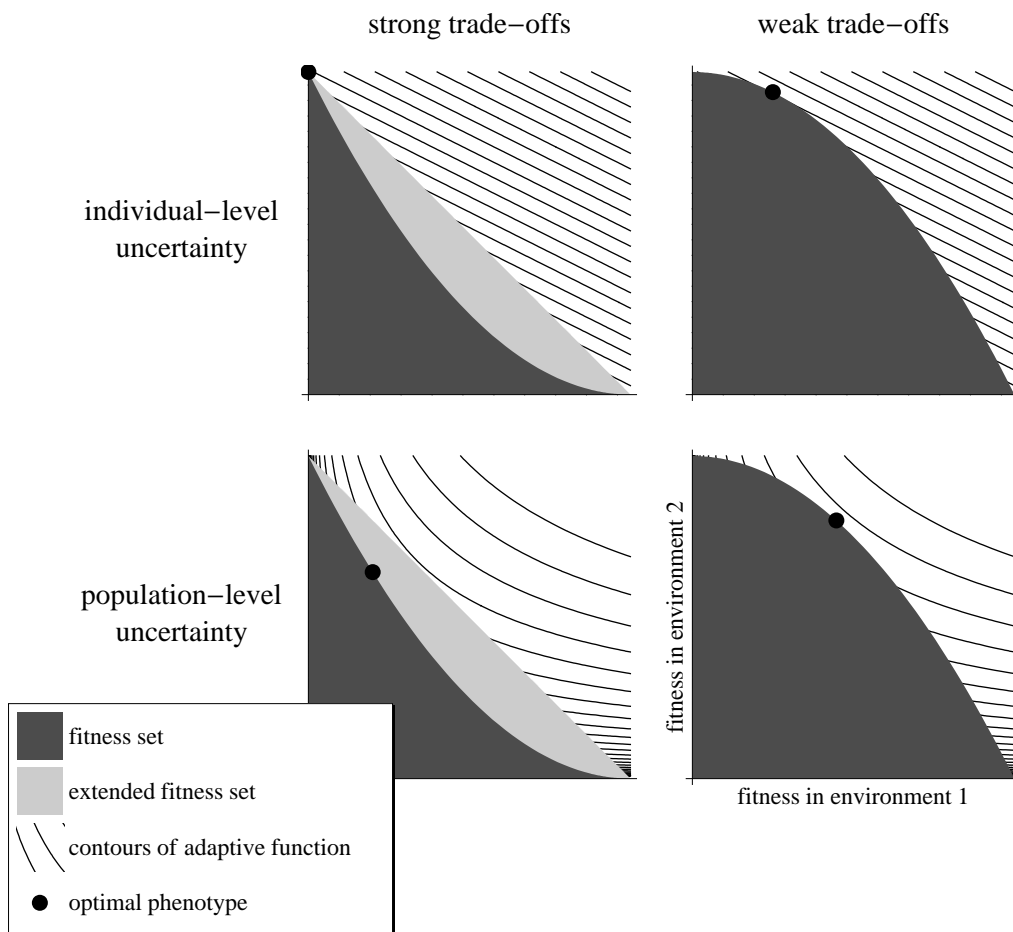


Figure 1: Levins's fitness set plots the fitness in each environment of all achievable phenotypes. Its boundary reflects the fitness trade-off between specializing in different environments. Trade-offs are described as weak if intermediate types can do fairly well in all environments, or strong if intermediate types do poorly in all environments. Contour lines of the adaptive function represent points in the fitness space where the population growth rates are equal. The achievable phenotype that maximizes the population growth rate is the point in the fitness set that lies on the highest contour. Specialists are favored when trade-offs are strong and environmental variation occurs at the individual level; generalists are favored when trade-offs are weak and variation occurs at the population level. The extended fitness set is created by considering in addition all mixtures of phenotypes. Allowing a mixed-phenotype population can increase the reproductive rate only when trade-offs are strong and environmental variation occurs at the population level (lower left.)

work, concerns populations containing a mixture of phenotypes. The population’s reproductive rate depends on the average fitness of the mixed population in each environment. The fitness profile of the mixed population is therefore a linear combination of the fitness profiles of the phenotypes that make up the population. The set of points in fitness space that can be achieved by a mixed population is known as the *extended fitness set*. It consists of all linear combinations of the original fitness set, which considers only single phenotypes (Figure 1). In mathematical terminology, the extended fitness set is the *convex hull* of the fitness set. Levins distinguishes two qualitatively different cases. Where trade-offs are weak, the extended fitness set does not create any new points (Figure 1, right column.) Where trade-offs are strong, however, the extended fitness set does expand the boundaries. Points along this new boundary represent fitness profiles that can only be achieved with a mixture of phenotypes (Figure 1, left column.) When environmental uncertainty occurs at the individual level, these points will never be optimal. The contours of the adaptive function and the new boundary of the extended fitness set are both linear (Figure 1, upper left.) However, when environmental uncertainty occurs at the population level, the contours of the adaptive function are no longer linear. In this case, a point along the linear edge of the extended fitness set may be better than any point in the fitness set (Figure 1, lower left.)

Levins initially suggested that population polymorphism could therefore be an adaptation to population-level environmental uncertainty, when being a generalist is not efficient—but some care is necessary in the interpretation. If, by population polymorphism, we mean genetic polymorphism, then we must consider not only the fitness of the population compared to other populations, but also the way that natural selection within the population may act to change the mixture. Levins addressed exactly this question using a simple mendelian trait in a later paper (Levins, 1964). He also suggested that long-term selection in a fluctuating environment might act on the genetic architecture to reduce the short-term effects of selection, thus maintaining some amount of genetic variation (Levins, 1965, 1968). Further work on the maintenance of genetic polymorphism as an adaptation to environmental uncertainty suggests that it is possible under certain conditions (Haldane and Jayakar, 1963; Gillespie, 1973; Sasaki and Ellner, 1995; Leimar, 2005). However, since natural selection does not generally maximize the geometric mean population fitness, Levins’s fitness set framework is poorly suited to this kind of analysis (Seeger and Brockmann, 1987; Godfrey-Smith, 1996).

On the other hand, phenotypic diversity in a population need not reflect genetic polymorphism. The notion of a stochastic developmental switch, which randomly produces one of several possible phenotypes, has long been

recognized as a potential mechanism for producing adaptive variation (e.g. Levins, 1968), and is central to the biological theory of bet-hedging. According to this theory, organisms may adapt to population-level uncertainty in their environment by randomly developing into one of several alternative phenotypes (Cohen, 1966; Cooper and Kaplan, 1982). Such a genotype may be thought of as a “developmental generalist”, because it produces a lineage which survives well in a variety of circumstances; depending on how strong trade-offs are, it may outcompete a phenotypic generalist (Wilson and Yoshimura, 1994). Because each point within the extended fitness set may be achieved by a single genotype that produces a mixture of phenotypes, and the genotype which produces the fastest-growing lineage is likely to outcompete all others, Levins’s diagrams are perfectly suited to distinguishing which of the two kinds of generalist is more efficient (Seger and Brockmann, 1987; Godfrey-Smith, 1996).

3 A model of the evolution of generalists

Consider a population of organisms with nonoverlapping generations. In each generation the environmental state k is drawn from some fixed probability distribution p_k , independently of the state in previous years. Each individual has a phenotype x that is fixed during development, and belongs to a continuous set of achievable phenotypes. Reproductive success depends both on the phenotype of the individual and the state of the environment. There may be individual-level variation in reproductive success within generations, but the average reproductive success of a phenotype, f_{xk} , must be consistent between generations with the same environmental state.

This is the classical model of evolution in fluctuating environments, where the phenotype with the highest geometric mean fitness is most likely to become fixed (e.g. Dempster, 1955; Levins, 1962; Cohen, 1966). One way to understand why natural selection tends to maximize the geometric mean in such cases is to look at long sequences of environments, and ask which genotype will take over the population in most of these sequences. Under the assumptions of our model, the genotype that will win in any particular sequence of environments is the one that had the largest growth rate. For a simple fluctuating environment, the strong law of large numbers implies that in almost all long sequences, each environmental condition is experienced approximately in proportion to its probability of occurring. In such sequences of environments, the expected long-term growth rate for a lineage expressing phenotype x is the growth rate averaged over the environmental

probabilities:

$$r(x) = \sum_k p_k \log f_{xk}, \quad (1)$$

which is the log of the geometric mean fitness. A natural extension to this approach also considers bet-hedging genotypes, which produce offspring with phenotypes given by some probability distribution g_x . Once a lineage is common enough, its average reproductive success is simply the weighted average reproductive success of the phenotypes it produces, so the long-term growth rate is given by:

$$r(g) = \sum_k p_k \log \sum_x g_x f_{xk} \quad (2)$$

(e.g. Seger and Brockmann, 1987; Yoshimura and Clark, 1991). The genotype that will be observed most of the time is the one that maximizes r —since its growth rate, over almost all sequences of environments, is larger than that of any other strategy.

In general, natural selection need not lead to optimization of the long-term growth rate. The model includes a number of important simplifying assumptions. In Section 4, we discuss the realism of these assumptions and the consequences of relaxing them.

3.1 The region of strong trade-offs

Given a model of the fitness trade-offs, we can use Levins’s fitness set diagram to graphically solve for the genotype with the highest long-term growth rate. It is easy to see that when trade-offs are weak, the best response to population-level variation will always be a single generalist phenotype. However, when trade-offs are strong, the situation is a little more complicated. A concave fitness set may show strong trade-offs only along part of its boundary. In this case, bet-hedging will be an optimal strategy only when the optimal single phenotype lies in a region of strong trade-offs. To illustrate, we begin with a simple example inspired by the ecology of amphibian metamorphosis.

Example. Frogs laying eggs in temporary ponds face a trade-off imposed by their time to metamorphosis. Late-metamorphosing tadpoles become larger frogs with higher fitness, but confront a higher risk that the pond will dry up before metamorphosis is complete. In our simple model, tadpoles are capable of metamorphosis beginning at some threshold number of days after hatching. After this, they grow linearly with the time to metamorphosis. However, if the pond dries up before metamorphosis, the tadpole

dies. Each frog lays its eggs in a different pond. In dry years, the time to drying of individual ponds is independent and normally distributed. In wet years, the time to drying of individual ponds is also independent and normally distributed, but with a different mean and standard deviation. From this we can calculate the expected fitness of a tadpole metamorphosing after a certain number of days, in a wet or dry year (Figure 2, see Appendix A for details of the model.)

The region of strong trade-offs is the part of the fitness set that makes it concave (Figure 3.) Since the adaptive function increases monotonically with fitness in each environment, the best pure strategy must lie on the boundary of the fitness set. Similarly, the best bet-hedging strategy must lie on the boundary of the extended fitness set. Whether a single generalist phenotype or a bet-hedging genotype is optimal depends on the overlap of these boundaries. We call the part where they overlap the region of weak trade-offs. The part where they do not overlap is the region of strong trade-offs; in this area, bet-hedging can improve upon any single phenotype. In particular, if the optimal strategy is a bet-hedging one, it will consist of a mixture of the phenotypes at the endpoints of the region of strong trade-offs.

We would like to determine, first of all, when bet-hedging is optimal, and secondly, if it is optimal, what is the best mixture of phenotypes. These questions can be answered for any specific model using the traditional fitness set approach, by identifying the point of intersection between the extended fitness set and the highest contour of the adaptive function (Levins, 1968; Yoshimura and Jansen, 1996). However, a much more general and intuitive approach is possible using the measure introduced in the next section.

3.2 A quantitative measure of specialization in different environments

When each phenotype can survive in only one environment, bet-hedging is the only way to survive: a lineage of any single phenotype is sure to go extinct. The optimal bet-hedging mixture of such phenotypes—regardless of fitness differences in different environments—is to match the probability of environments. For example, Cohen’s (1966) simplest model of seed germination, in which a germinating seed can only survive in good years, showed that the optimal germination fraction is equal to the probability of a good year. The generality of this result is well known in information theory (Kelly, 1956; Cover and Thomas, 1991), but has only more recently been applied in a biological context (Bergstrom and Lachmann, 2004; Kussell and Leibler, 2005). This case provides an important reference point, and a good basis for

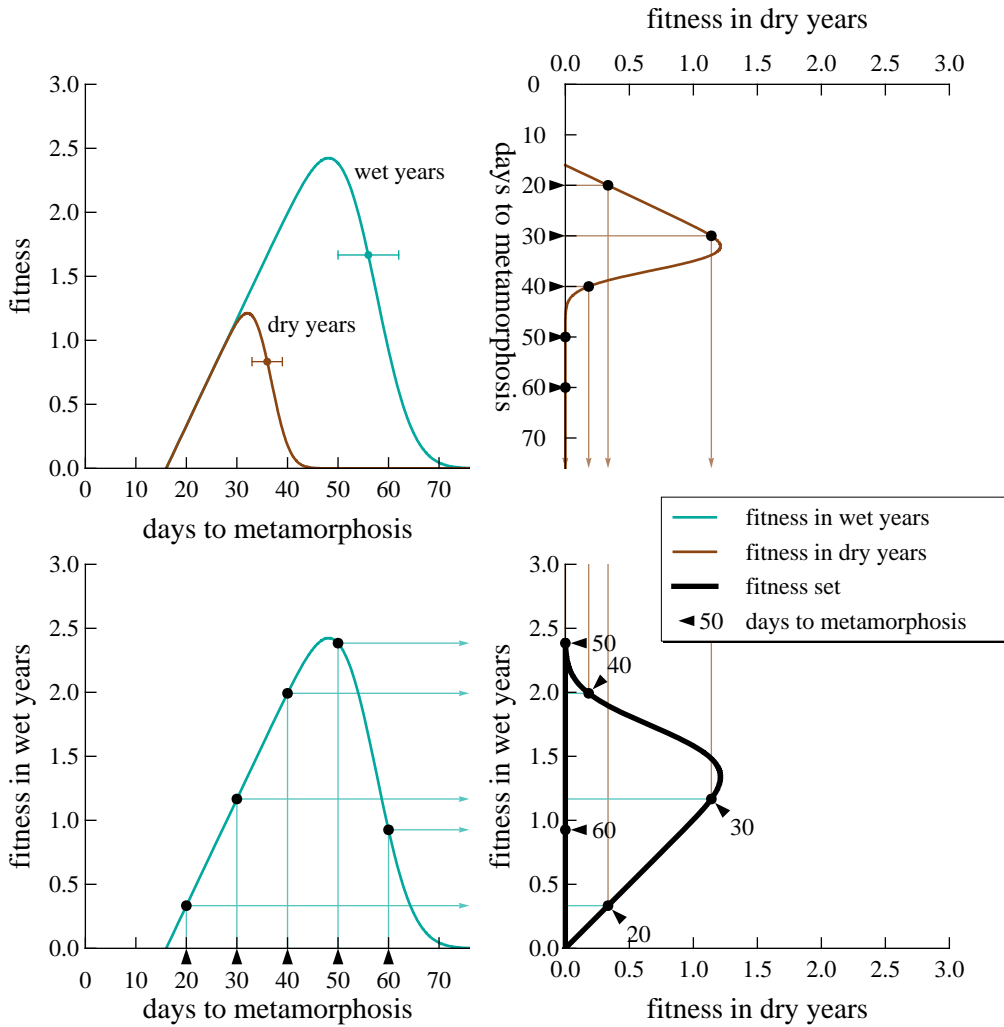


Figure 2: Levins's fitness set, lower right, is a parametric plot of the fitness of each phenotype in several environments. Shown on the upper left is a graph of average fitness in two kinds of years, for tadpoles that metamorphose at different ages. Error bars indicate the mean and standard deviation of time to drying. To illustrate how the parametric plot is derived, we show it alongside individual plots of fitness in each kind of year. Consider the phenotype of metamorphosis at 30 days. Its fitness in wet years, shown in the lower left plot, becomes the y coordinate in the fitness set. Its fitness in dry years, shown in the upper right plot, becomes the x coordinate in the fitness set. When this is done for every phenotype, a curve is traced out in fitness space. This is called the fitness set.

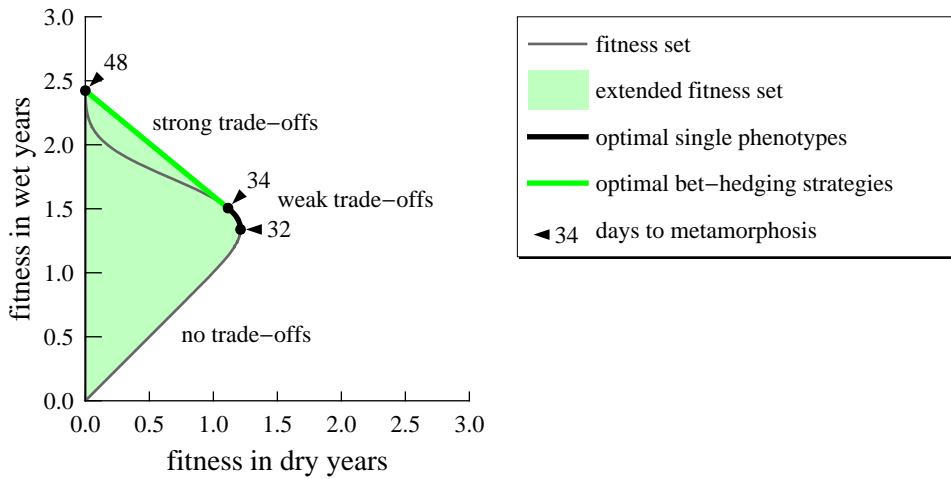


Figure 3: Levins’s extended fitness set consists of all points in the fitness set, plus all linear combinations of those points. The boundaries of this set reflect the trade-offs in the system. Until metamorphosis at 32 days, fitness increases in both kinds of years, so there is no trade-off at all. Between 32 and 48 days, fitness in dry years decreases while fitness in wet years increases. This region of trade-offs can be divided into two parts. Between 34 and 48 days, the boundary of the extended fitness set goes beyond the boundary of the fitness set. That means that if the best pure strategy lies somewhere in this range, a mixed strategy which combines tadpoles metamorphosing at 34 and 48 days would do even better. This is called the region of strong trade-offs. Between 32 and 34 days, the boundary of the fitness set and the extended fitness set overlap. If the best pure strategy lies somewhere in this range—the region of weak trade-offs—no mixed strategy can be an improvement.

our measure of specialization, precisely because it is so well understood. We will show that our measure can then easily be extended to treat the more general case, where phenotypes can survive in several environments.

We define a phenotype that survives in only one environment as *completely specialized* in that environment. Furthermore, a bet-hedging genotype that produces a mixture of completely specialized phenotypes is specialized in each environment according to that mixture. Consider a genotype g that allocates a fraction g_k of its offspring to a phenotype that is completely specialized in environment k . Then the fitness of that genotype in environment k is $f_{gk} = g_k d_k$, where d_k is the fitness of the completely specialized type in its environment. We say that the *degree of specialization* of the genotype in environment k is g_k , because that is the proportion of its reproductive effort that is invested in environment k . Note that, according to the result cited above, the optimal genotype should specialize in each environment according to its frequency.

The idea that a bet-hedging genotype divides its reproductive effort among environments can be extended to include phenotypes that survive in multiple environments. To illustrate how, we return to the example of amphibian metamorphosis. In the last section, we saw that any optimal bet-hedging strategy must be a mixture of just two types: a slow-metamorphosing tadpole (at 48 days) and a fast-metamorphosing tadpole (at 34 days.) Such genotypes lie along a straight line on the boundary of the extended fitness set (Figure 3.) Notice, however, that this line is a subsection of a longer line that extends all the way to both axes (Figure 4). This longer line represents all bet-hedging strategies that combine two completely specialized types: the slow-metamorphosing type, which has fitness 2.42 in wet years, and fitness 0 in dry years, and another type that has fitness 2.94 in dry years, and fitness 0 in wet years. Since any particular tadpole would do better in a wet year than a dry one, it is impossible to be completely specialized in dry years—the second phenotype cannot actually be achieved. However, we can still ask, if it were achievable, what would be the optimal mixture of these two completely specialized types? We know the answer must be to produce dry-year specialists as often as there are dry years, and wet-year specialists as often as there are wet years. If that point lies along the subsection of the line that can be produced by combining slow and fast metamorphosing tadpoles, then the optimal allocation of reproductive effort—matching the environmental frequencies—can still be achieved.

We therefore define the degree of specialization of a bet-hedging genotype in each environment in terms of the mixture of completely specialized types to which it corresponds. Consider a bet-hedging genotype g which combines several phenotypes that are not completely specialized. Suppose that each of




Figure 4: An optimal bet-hedging strategy matches its level of specialization in each environment to the frequency of that environment. In the left panel, we show how specialization levels for any bet-hedging mixture of two phenotypes can be calculated. The line between the two phenotypes plotted in fitness space is extrapolated to the axes. This defines a bet-hedging strategy that combines two perfectly specialized phenotypes. The best mixture of perfectly specialized phenotypes matches the frequency of different environments. However, in the example shown it is impossible to be perfectly specialized in dry years. Still, some of the same results can be achieved by using a mixture of the original phenotypes. In the right panel, we plot the best bet-hedging mixture of metamorphosis at 48 days and at 34 days, as a function of the percentage of dry years. The slope of this line reflects the sensitivity of the optimal strategy to changes in the environmental probabilities. The right-hand axis indicates what mixture of the two types attains the optimal level of specialization, while the left-hand axis indicates the average number of days to metamorphosis. We also show the best single-phenotype strategy, in days to metamorphosis. From 0–38% dry years, bet-hedging is better than the best single-phenotype strategy. Above this, bet-hedging cannot achieve the optimal level of specialization, so the optimal strategy is to produce a single phenotype.

those phenotypes has a fitness profile that can be represented as a different mixture of just one set of completely specialized phenotypes:

$$f_{xk} = s_{xk}d_k. \quad (3)$$

Any bet-hedging mixture of the original phenotypes is then also equivalent to a mixture of that same set of completely specialized phenotypes:

$$\sum_x g_x f_{xk} = s_{gk}d_k, \text{ where} \quad (4)$$

$$s_{gk} = \sum_x g_x s_{xk} \quad (5)$$

We define the degree of specialization of the bet-hedging genotype in terms of this mixture, s_{gk} . Haccou and Iwasa (1995) note that, if each fitness can be written as a product $s_{xk}d_k$ as in Equation 3, then the optimal distribution of types g_x^* is the one that achieves

$$\sum_k g_x^* s_{xk} = p_k \quad (6)$$

where p_k is the probability distribution over environments. Our definition of specialization in terms of an equivalent mixture of completely specialized phenotypes is therefore particularly convenient, because it means the right amount of specialization in any environment, s_{g^*k} , is simply the probability of that environment.

It remains to be shown when and how the appropriate decomposition of the fitness, as in Equation 3, can be accomplished. When there are just two or three different kinds of environments, this is easy to visualize using Levins's fitness sets, as illustrated in Figure 4. We use the region of strong trade-offs to identify the phenotypes that could be used in an optimal bet-hedging strategy. The number of phenotypes in that set is limited by the number of distinct environments, and will equal the number of environments as long as there are strong fitness trade-offs between all environments. In that case, the specialization levels of the individual phenotypes s_{xk} are uniquely defined, and can be found according to the method described in Appendix B.

Looking at genotypes in terms of their specialization in different environments provides the tools we need to answer the questions posed at the end of the last section: when is bet-hedging adaptive, and if it is adaptive, how much is optimal? As shown in Figure 4 (left panel), a tadpole that metamorphoses at 48 days is completely specialized in wet years, while a tadpole that metamorphoses at 32 days is only 38% specialized in dry years. Therefore,

if dry years occur more than 38% of the time, a bet-hedging strategy using only these two types could not be specialized enough in dry years. If dry years occur less often than this, then the optimal strategy is the combination of the two phenotypes that matches the degree of specialization in dry years to the probability of dry years (Figure 4, right panel). As dry years increase in probability from 0% to 38%, the optimal mixture of types changes from all metamorphosing at 48 days, to all metamorphosing at 34 days. That means that a fairly small increase in the percentage of dry years corresponds to a larger increase in the fast-metamorphosing type.

In general, bet-hedging with phenotypes that are not completely specialized in any environment can produce only a limited subset of possible specialization levels, defined by the specialization levels of the phenotypes. When environmental probabilities lie outside this range, a single, generalist phenotype will be optimal. Within this range, on the other hand, the optimal amount of phenotypic variation changes linearly, taking on all possible values, as environmental probabilities change (see Equation 6). The slope of this linear relationship, reflecting the sensitivity of the optimal bet-hedging mixture to the environmental probabilities, depends only on the size of the region of strong trade-offs. The general form of the relationship between the amount of adaptive diversification and the amount of environmental uncertainty that drives it is therefore quite simple and intuitive.

An additional advantage of describing genotypes in terms of their specialization in different environments is that it allows for a simple description of the strength of selection for a bet-hedging genotype, when it is optimal. We examine the difference between the optimal growth rate, which uses the bet-hedging genotype g^* , and the current long-term growth rate under the bet-hedging genotype g :

$$r(g^*) - r(g) = \sum_k p_k \log \frac{\sum_x g_x^* f_{xk}}{\sum_x g_x f_{xk}}. \quad (7)$$

Rewriting the fitnesses as in Equation 3, and substituting in Equations 4 and 6, we find

$$r(g^*) - r(g) = \sum_k p_k \log \frac{p_k}{s_{gk}} \quad (8)$$

which is the Kullback-Leibler divergence between the environmental probabilities p_k and the genotypic specialization level s_{gk} —the environmental probabilities for which the current genotype g would be optimally adapted (compare Haccou and Iwasa, 1995; Kussell and Leibler, 2005). In this sense, the amount of environmental uncertainty indicates how much of a generalist an optimally adapted genotype should be; the farther away a genotype is

from this ideal, the stronger the selection to improve. Furthermore, selection for moving from the best single-phenotype strategy to any bet-hedging strategy cannot exceed the Kullback-Leibler divergence between the specialization levels of the component phenotypes. If there are strong trade-offs only in a small portion of the fitness set, there will be only a narrow range of environmental uncertainty in which bet-hedging is optimal. Furthermore, selection for adopting a bet-hedging strategy will be strongly limited, because the best phenotypic generalist will do almost as well as the optimal bet-hedging genotype.

3.3 The effect of individual-level risk

Until now, we have addressed only the selective effects of population-level risk, because this is the only kind that can lead to adaptive phenotypic diversification. Now we ask, how does individual-level risk affect the strength of selection for bet-hedging that may be imposed by population-level risk? The method of dividing the variance in average genotypic fitness over generations into the variance in individual fitness and the correlation in fitness between individuals, introduced by Frank and Slatkin (1990), provides some insight into this question. They represent increasing individual-level uncertainty by decreasing correlations in fitness between individuals of the same genotype. This decreases the variance in average genotypic fitness, effectively decreasing the amount of population-level uncertainty. We may therefore expect that adding individual-level risk to our model will dampen the effects of population-level risk.

Our example of amphibian metamorphosis contains environmental risk at two different levels: at the population level, because dry years cause the ponds to dry up earlier on average, and at the individual level, because within each year the time to drying of individual ponds varies according to a normal distribution. In Levins's original formulation such combinations of risk at two different levels were represented by adjusting the adaptive function (Levins, 1962). However, we have instead included the individual-level risk in the fitness set: each phenotype is represented by a vector consisting of its average fitness in each environment. This is possible because, given the type of year, the ponds dry up independently of one another. Therefore, the variance over different years of the average genotypic fitness (a sample mean) will be negligible for large populations. In the example, individual-level uncertainty in the time to drying makes specialized types more generalist, by decreasing average fitness in the most favored circumstances and increasing it in the least favored ones. In the process, it reduces the size of the region of strong trade-offs, putting stricter and stricter limits on the amount of population-level

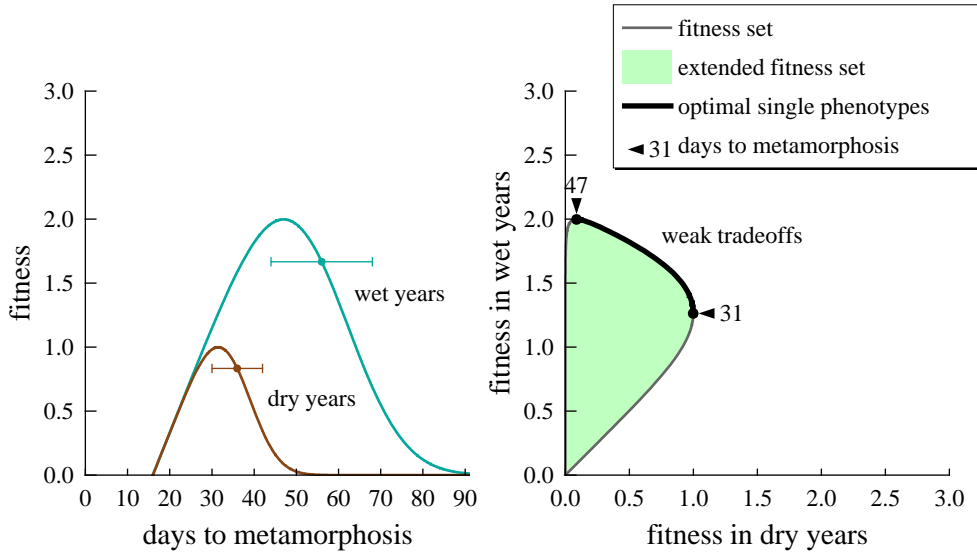


Figure 5: Individual-level uncertainty reduces selection for phenotypic diversification by changing the shape of the fitness set. The left panel shows fitness as a function of phenotype; the right panel shows the resulting fitness set and extended fitness set. Compared to the parameters used in Figures 2 and 3, the only difference here is a larger standard deviation in time to drying, in both kinds of years. This reduces the risk associated with being caught in the wrong kind of year, creating only a weak trade-off between specializing in wet and dry years over the entire range. In this situation, no matter the level of population-level risk, an individual-level generalist will always be favored.

uncertainty necessary to induce adaptive phenotypic variation. When we increase the variation within years enough that there is significant overlap in drying time between wet and dry years, the region of strong trade-offs disappears completely—making phenotypic variation non-adaptive (Figure 5.)

Our approach has the advantage of singling out population-level uncertainty as the driving force in the evolution of adaptive variation: it sets a target for the amount of specialization in different environments. Individual-level uncertainty may then alter the circumstances under which diversification is favored, by changing the shape of the fitness set. Individual-level environmental risk could affect the fitness profile of different phenotypes in many ways. Its impact depends on how it affects the average performance of different types. In our example, within-year variation in time to drying

of different ponds makes the average fitness of similar types more similar (the slopes in the left panel of Figure 5 are less steep than in the corresponding panel in Figure 2.) This effectively reduces the size of the region of strong trade-offs, making intermediate types better generalists, and making bet-hedging less important. On the other hand, if individual-level stochasticity affected only the variance in fitness of individual types, not the average fitness, it would not change the fitness set at all. Still, to the extent that individual-level uncertainty does make phenotypes more generalist, it will make bet-hedging both less likely to be adaptive, and less strongly selected for (see Section 3.2).

3.4 The effect of developmental plasticity costs

Our representation of the growth rate of a lineage stemming from a bet-hedging genotype (see Equation 2) implicitly assumes that there is no cost to plasticity. This assumption is important for our results, because it means that the fitness profile of a bet-hedging genotype is a simple linear combination of the fitness profiles of the phenotypes it produces. However, some kinds of cost can still be represented within our framework. If there is a fixed cost to being able to produce more than one phenotype, regardless of the proportions actually produced, the overall shape of the extended fitness set will not change. Consider a cost of plasticity that is proportional to an individual’s realized fitness. This simply scales the extended fitness set by some factor (Figure 6, left panel.) The boundary of the extended fitness set will still be linear wherever there are strong trade-offs, and will still correspond to the same mixed strategies. What changes is the comparison between the best pure strategies and the best mixed strategies. Mixed strategies are optimal over a smaller range of environmental uncertainties, and the range of optimal mixtures decreases as well. The linear relationship between the environmental frequencies and the optimal mixture of phenotypes still holds (Figure 6, right panel.)

Another possibility is that the cost of plasticity could depend on the amount of phenotypic diversity it generates. For example, a plastic developmental pathway might tend to build the less commonly produced phenotype with less accuracy, and thus lower average fitness. Then a bet-hedging strategy that almost always produces just one phenotype would incur very little fitness cost, while one that produces two phenotypes in equal proportions would incur a higher cost. In this case the boundary of the extended fitness set will no longer be linear. The graphical method of finding the point on the boundary which maximizes the adaptive function can still be used. However, the most important feature of our framework—the direct relationship




Figure 6: A fixed cost to plasticity reduces the range of environmental uncertainty in which bet-hedging is optimal. As in Figure 4, we use the extended fitness set to calculate the degree of specialization in dry years for a range of bet-hedging strategies. Because the cost is proportional to fitness, the shape of the extended fitness set does not change; it is merely scaled down. The phenotypes that are used to make all optimal bet-hedging mixtures are again metamorphosis at 34 and 48 days, and the specialization levels of each do not change. This means that the best bet-hedging strategy for any particular percentage of dry years, shown at right, also does not change. Likewise, the best single-phenotype strategy remains the same. What does change, however, is the comparison between single-phenotype strategies and bet-hedging strategies. The cost of bet-hedging means that the range of environmental uncertainty in which the best bet-hedging strategy outperforms the best single phenotype gets smaller: 0–25% dry years instead of 0–38%. This corresponds to mixtures where less than 66% of tadpoles metamorphose at 34 days. Mixtures outside this range are not worth the cost of plasticity.

between environmental probabilities and phenotypic diversity—is lost.

What does this mean for the applicability of our framework in natural systems? The answer depends on how common costs to plasticity are, and how those costs are related to the amount of phenotypic diversity that is created. To our knowledge, no one has yet empirically measured cost in a system where stochastic developmental plasticity has been demonstrated. Most empirical studies of the costs of plasticity have focused instead on adaptive plasticity in response to predictive cues (see DeWitt et al., 1998, for a review). Much of the cost of plasticity in these cases may be related to the ability to detect and respond appropriately to predictive cues. Such costs do not apply to stochastic developmental plasticity. However, it may still be more difficult for an organism to produce, along alternative developmental pathways, the same phenotypes that developmentally canalized organisms could produce. This difficulty may be due to limits on the range of plastic development or increased developmental instability. Empirical tests for these kinds of limitations to plasticity so far show no support for any association between increased plasticity and decreased phenotypic range or precision (e.g. DeWitt, 1998; Van Kleunen et al., 2000; Relyea, 2002). This suggests that the costs to stochastic developmental plasticity may be fairly limited, and not related to the amount of stochasticity. These are precisely the conditions in which our model is most useful.

4 Generality of the model

A more formal analysis of the effects of stochastic environments on natural selection would include explicit models of population dynamics. Our assumption that the genotype with the highest average log fitness is most likely to fix in the population depends on a number of key simplifications in the model. We consider each of these simplifications in turn.

First of all, we have made the assumption that the long-term growth rate of a lineage is a good predictor for the outcome of a competition between two genotypes. This is clearly true if both lineages are growing exponentially and completely independent of one another, but such a scenario is not biologically plausible. In most cases there will be some kind of density-dependent regulation of growth. We would like to know which strategy is likely to outcompete all others, in the long term. This depends on the way that the relative proportions of different strategies in the population change over time. Consider some form of density-dependent population regulation that is equally likely to affect all individuals, regardless of strategy or phenotype. This would maintain realistic population sizes, but would not affect

the relative proportions of different strategies. Therefore, the strategy with the highest long-term growth rate would still be most likely to outcompete the others (McNamara, 1995; Grafen, 1999).

On the other hand, some kinds of density-dependent population regulation can affect the proportions of different types of strategies. For example, individuals might specialize in the use of different limiting resources, whose availability varies from generation to generation. Then an individual's reproductive fitness will depend not only on its own phenotype, the environment, and the population size, but also on the frequency of other phenotypes in the population. This is a classic case of frequency-dependent selection, whose outcome cannot generally be predicted by principles of optimization. The concept of an evolutionarily stable strategy—a genotype that, once common, resists invasion from potential mutant types—can then be a useful tool for identifying adaptive strategies (Maynard Smith, 1982). In the context of fluctuating environments, the ability of a mutant genotype to successfully invade a resident is predicted by the long-term growth rate of the mutant in a population of the resident (Metz et al., 1992; Yoshimura and Jansen, 1996). Our framework can still be used to predict what would be the best bet-hedging strategy, for a particular distribution of phenotypes in the resident population. If this optimal strategy also produces the given distribution of phenotypes, then it is at least an equilibrium strategy, although it need not be stable.

The second important assumption is our use of the log fitness, averaged over different environmental states, as a proxy for the long-term growth rate of a lineage (see Equation 1). Whether this is a good approximation depends crucially on how the environment varies on several scales: within individuals, between individuals in the same generation, and between generations.

Our assumption that individuals live and reproduce in only one generation, and that each generation is characterized by a single environmental state, ignores the possibility of environmental variation within an individual's lifetime. In general, the effect of such within-individual variation is to reduce the impact of between-generation variation, because individual lifetime reproductive success becomes less variable (Frank and Slatkin, 1990; Sasaki and Ellner, 1995). In fact, organisms may in some cases evolve to survive and reproduce over multiple seasons as a response to an environment that fluctuates from generation to generation. Evolution of such a life history strategy may therefore sometimes be seen as a kind of “risk spreading” in its own right (Murphy, 1968; Goodman, 1984).

Variation between individuals in the same generation need not affect the long-term growth rate, as long as the average fitness of each phenotype within each environmental state does not vary much. This will be true as long as the

number of individuals of each phenotype is fairly large, and the reproductive success of different individuals is independent, conditional on the environmental state (Frank and Slatkin, 1990). We have included individual-level variation of exactly this sort in our example (see Section 3.3).

In this study, we have focused on the evolutionary impact of environmental variation between generations. We have made the important simplification that the environmental conditions in each generation do not depend on previous environmental history. In such cases, the asymptotic growth rate of a lineage simplifies to its average log fitness over generations. However, this simplification continues to hold even in more complicated scenarios: as long as the population is unstructured, and the environmental states constitute an ergodic process, the asymptotic growth rate is the log fitness of the lineage, averaged over the stationary distribution of environmental states (Tuljapurkar, 1990).

5 Discussion

Variable life history strategies have been studied as a potential adaptation to fluctuating environments in a wide variety of biological systems. The best-known example is delayed germination in desert annual plants. This was the inspiration for Cohen’s (1966) model, and has spawned a series of follow-ups (e.g. Philippi, 1993; Clauss and Venable, 2000; Evans et al., 2007; Venable, 2007). A similar phenomenon—an overwinter diapause—is observed in many insects and crustaceans (Saiah and Perrin, 1990; Bradford and Roff, 1997; Danforth, 1999; Hopper, 1999; Menu et al., 2000; Philippi and Seger, 2001). A highly variable time to metamorphosis, as described in our example, has been observed in some anurans breeding in temporary pools (Lane and Mahony, 2002; Morey and Reznick, 2004). In fish and amphibians, the trade-off between egg size and egg number may make variation in egg size an adaptation to environmental uncertainty (Capinera, 1979; Crump, 1981; Kaplan and Cooper, 1984; Koops et al., 2003). For organisms that switch between sexual and asexual modes of reproduction, like aphids and some plants, the timing of that switch may vary in response to uncertainty (Berg and Redbo-Torstensson, 1998; Halkett et al., 2004). Bacteria are rapidly becoming an important model system for the study of adaptive non-genetic phenotypic diversity, in part because the regulatory mechanisms underlying the heterogeneity are particularly amenable to study (reviewed in Avery, 2006; Smits et al., 2006). For example, *Escherichia coli* periodically exposed to antibiotics switch stochastically between a fast-growing antibiotic-sensitive type and a slow-growing, antibiotic-resistant type; the rate of switching varies between

strains and may be adapted for different frequencies of antibiotics exposure (Kussell and Leibler, 2005). Finally, in a few systems it is argued that a single, low-risk life history strategy is a better adaptation to fluctuating environments than phenotypic diversity would be (Boyce and Perrins, 1987; Einum and Fleming, 2004; Hassall et al., 2006; Simons and Johnston, 2003). This is sometimes called conservative bet-hedging, in contrast to diversified bet-hedging, which uses a variety of phenotypes (Seger and Brockmann, 1987; Philippi and Seger, 1989).

Empirical studies of life history evolution in response to environmental uncertainty use theory to make testable predictions in a number of ways. One approach is to show that the observed strategy maximizes the geometric mean fitness instead of the arithmetic mean fitness (Boyce and Perrins, 1987; Philippi and Seger, 2001). The difficulty with such quantitative predictions is that they are often quite sensitive to errors in observed parameters of the model, such as the frequency of different kinds of environments. Instead, most studies test qualitative predictions. For example, if it can be shown that a mixture of phenotypes performs better in the long term than any single phenotype, observed variation could be an adaptation to uncertainty (Saiah and Perrin, 1990; Menu et al., 2000; Evans et al., 2007). Stronger evidence is gained by comparing several populations or species with different amounts of uncertainty about the environment. If the amount of phenotypic diversity observed varies with the environmental risk as predicted by theory, that diversity is likely to be adaptive (Philippi, 1993; Clauss and Venable, 2000; Koops et al., 2003; Halkett et al., 2004; Venable, 2007).

The framework presented here can be used to make quantitative predictions about what phenotype or mixture of phenotypes would be best adapted to a particular set of conditions. In order to use it, a realistic, data-driven model of how different phenotypes fare in different environments is needed, and an accurate assessment of the frequency of different kinds of environments over the long term must be made. Then the procedure illustrated with the example of frog metamorphosis can be followed to predict the optimal response. While similar quantitative predictions could also be made without our framework—via numerical optimization or computer simulation—our graphical method of analysis provides a clearer picture of why bet-hedging may or may not be favored in any particular system. If bet-hedging is not adaptive, it could be because the region of strong trade-offs is so small that bet-hedging would almost never be adaptive. On the other hand, it might be because the environmental frequencies just happen to lie outside the appropriate range. Another reason for preferring our approach is that it gives a comprehensive picture of how sensitive the optimal strategy is to measurement errors in various parameters of the model. For example, small changes

in the fitness functions can change strong trade-offs into weak trade-offs (compare Figures 3 and 5), going from a situation in which bet-hedging is often favored to a situation where it is never favored.

Our framework generates several novel qualitative predictions that could be tested using a comparative approach. The first is that the range of environmental uncertainty in which bet-hedging is adaptive is limited by the size of the region of strong trade-offs (see Figure 4.) Other models have indicated that bet-hedging is only adaptive when the variance in which phenotype is optimal exceeds a certain threshold value (Slatkin and Lande, 1976; Bull, 1987), suggesting that a minimal amount of environmental variance is necessary. Haccou and Iwasa (1995) note that this minimal variance threshold depends on how generalist the phenotypes are, and calculate it for several specific functions describing the relationship between phenotype, environment and fitness. Our result is more general, because it is independent of the particular form of the fitness function. It is also more explicit, because it describes the environmental probability distributions themselves rather than simply the variance of those distributions. The second prediction is that the range of potentially optimal bet-hedging strategies is affected by the cost of plasticity. As long as there is no cost to plasticity, any given mixture of the right phenotypes can be optimal, for some particular amount of environmental uncertainty (see Figure 4, right panel.) On the other hand, if there is a cost to plasticity, some mixtures of the right phenotypes will never be optimal under any level of environmental uncertainty (see Figure 6, right panel.) Finally, we show that there is a linear relationship between the environmental frequencies and the optimal bet-hedging mixture of phenotypes. How sensitive the optimal mixture is to changes in environmental frequencies depends on how big the the range of potentially optimal bet-hedging strategies is, relative to the range of environmental frequencies in which bet-hedging is optimal. That means that if the region of strong trade-offs is small—say, making bet-hedging optimal only in the range from 50–52% dry years—even a tiny change in the environmental probabilities can have a big effect on the optimal mixture of phenotypes.

The idea that random phenotypic variation can lessen the negative impact of environmental stochasticity is often explained in terms of investment in the stock market. In a well-crafted stock portfolio, high-risk, high-return stocks are combined with low-risk, low-return stocks in a way that maximizes the expected rate of return. Similarly, bet-hedging genotypes often produce a combination of two life history strategies, one with high expected fitness, but high risk, and another with lower expected fitness and lower risk. Our framework builds upon this analogy. What we call a completely specialized phenotype is completely invested in one particular environment; in all other

environments that individual is counted as a total loss. Any phenotype that is not completely specialized, however, is effectively invested in several different environments at once. Its proportional investment in different environments is defined by comparing the relative fitness of different types. Bet-hedging is a way of fine-tuning the total investment in different environments, by producing a variety of offspring whose average specialization level matches the probabilities of the different kinds of environments.

Two important factors in the evolution of specialists and generalists are left for future exploration. One is the effect of competition for resources, which can promote the evolution of within-species polymorphism in resource use (see Skúlason and Smith, 1995, for a review). Examples of this are widespread in birds, amphibians, and fish and may arise through genetic divergence (e.g. Smith, 1993) or condition-dependent plasticity (e.g. Frankino and Pfennig, 2001). This mechanism for the generation of adaptive variation is driven by frequency-dependent selection, and can act independently of the kind of population-level environmental uncertainty we have considered here. As discussed in Section 4, our model cannot capture the full picture when there is frequency-dependent selection. However, recent work has extended the use of the fitness set into scenarios including frequency dependence (Rueffler et al., 2004; de Mazancourt and Dieckmann, 2004). Such an approach might be a useful extension for our model. Another exciting direction will be to consider predictive cues that can be used to direct phenotypic plasticity. For example, some tadpoles react to decreasing water levels in their pond by accelerating development (Denver et al., 1998). If the cues do not predict the environment perfectly, however, some environmental uncertainty may remain—making a combination of bet-hedging and plasticity a potentially useful strategy (DeWitt and Langerhans, 2004). Our framework can easily and naturally be extended to explore the relationship between the amount of information and population-level uncertainty in a cue, and the optimal balance between predictive and stochastic plasticity.

Acknowledgements

The authors would like to thank Franjo Weissing and Lauri Oksanen for helpful comments. This material is based upon work supported under an NSF Graduate Research Fellowship to M.C.D.

Appendices

A A model of frog metamorphosis

Here we specify a simple model of frog metamorphosis. A tadpole is capable of metamorphosis after a minimum number days, t . After this point, it grows linearly with slope k . This determines the size of the frog at metamorphosis, $f(x) = k(x - t)$. However, the longer it spends as a tadpole, the more likely it is that the pond will dry up before it manages to metamorphose. In wet years, the time of drying for individual ponds is independent and distributed normally, with mean μ_w and standard deviation σ_w . In dry years, the time of drying for individual ponds is independent and distributed normally, with mean μ_d and standard deviation σ_d . The probability that a pond does not dry up by time x in dry years is one minus the cumulative distribution function for the appropriate normal distribution:

$$s_d(x) = 1 - \frac{1}{2} \left(\operatorname{erf} \left(\frac{x - \mu_d}{\sqrt{2}\sigma_d} \right) + 1 \right), \quad (\text{A-1})$$

where $\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$. A similar expression holds for $s_w(x)$. The expected fitness for a tadpole metamorphosing after x days is $f_d(x) = k(x - t)s_d(x)$ in a dry year, or $f_w(x) = k(x - t)s_w(x)$ in a wet year. For Figures 2, 3, and 4, we use parameter values $t = 16$, $k = 1/12$, $\mu_d = 36$, $\sigma_d = 3$, $\mu_w = 56$, and $\sigma_w = 6$. For Figure 5 we increase the standard deviations to $\sigma_d = 6$ and $\sigma_w = 12$.

In Section 3.4, we add a cost of plasticity to the model. Compared to the same phenotype produced without plasticity, a plastically produced phenotype experiences a proportional decrease in fitness: $f'_d(x) = (1 - c_d)f_d(x)$ and $f'_w(x) = (1 - c_w)f_w(x)$. In Figure 6, we use $c_d = c_w = 0.05$. Because the cost is proportional to fitness, the extended fitness set created by all linear combinations of f'_d and f'_w , instead of f_d and f_w , shrinks but does not change shape. In addition, the extrapolation line used to calculate specialization levels is stretched without changing the proportions. This property of a fitness-proportional cost function makes it easier to compare the results with and without cost. However, any other kind of cost function could be used, as long as it is constant with respect to the frequency of phenotypes that are produced.

B Calculating specialization levels

Given a set of phenotypes (as defined by the region of strong trade-offs, see Section 3.1) we would like to write the fitnesses of any mixture in each

environment as a mixture of completely specialized phenotypes, having non-zero fitness in only one environment. This corresponds to writing the fitness f_{xk} of phenotype x in environment k as a linear combination of values d_k , each of which represents the fitness of a phenotype completely specialized in environment k . Each coefficient in these linear combinations, s_{xk} , is the specialization level of phenotype x in environment k . The problem can be formulated as a matrix equation:

$$\mathbf{F} = \mathbf{S}\mathbf{D}, \tag{A-2}$$

where \mathbf{S} is a stochastic matrix with rows summing to 1, and \mathbf{D} is a diagonal matrix.

We assume that \mathbf{F} is a square matrix, that is, there are as many phenotypes as environments. This will be true as long as there are strong fitness trade-offs between specializing in all possible environments. In this case (barring singularities) both \mathbf{F} and \mathbf{D} are invertible, so

$$\mathbf{D}^{-1} = \mathbf{F}^{-1}\mathbf{S}. \tag{A-3}$$

Multiplying on the right by a column vector of ones, we get the row sums

$$\text{row sum}(\mathbf{D}^{-1}) = \text{row sum}(\mathbf{F}^{-1}) \tag{A-4}$$

because a stochastic matrix has all row sums equal to 1. This defines the diagonal matrix \mathbf{D} completely, allowing the calculation of the specialization in each environment as

$$\mathbf{S} = \mathbf{F}\mathbf{D}^{-1}. \tag{A-5}$$

This method will yield positive specialization levels as long as the diagonal matrix \mathbf{D} has only positive entries. That is, returning to the graphical viewpoint illustrated for two dimensions in Figure 4, the hyperplane passing through all phenotypes plotted in fitness space must intersect each axis at a positive point. This will be true as long as there is a trade-off between fitness in all different environments, as we have assumed.

References

- AVERY, S. J. 2006. Microbial cell individuality and the underlying sources of heterogeneity. *Nature Reviews Microbiology* 4:577–587.
- BERG, H. AND REDBO-TORSTENSSON, P. 1998. Cleistogamy as a bet-hedging strategy in *Oxalis acetosella*, a perennial herb. *Journal of Ecology* 13:491–500.

- BERGSTROM, C. T. AND LACHMANN, M. 2004. Shannon information and biological fitness. *In* IEEE Information Theory Workshop 2004, pp. 50–54. IEEE. (See also arXiv.org:q-bio/0510007).
- BOYCE, M. S. AND PERRINS, C. M. 1987. Optimizing great tit clutch size in a fluctuating environment. *Ecology* 68:142–153.
- BRADFORD, M. J. AND ROFF, D. A. 1997. An empirical model of diapause strategies of the cricket *Allonemobius socius*. *Ecology* 78:442–451.
- BULL, J. J. 1987. Evolution of phenotypic variance. *Evolution* 41:303–315.
- CAPINERA, J. L. 1979. Qualitative variation in plants and insects: Effect of propagule size on ecological plasticity. *The American Naturalist* 114:350–361.
- CLAUSS, M. J. AND VENABLE, D. L. 2000. Seed germination in desert annuals: An empirical test of adaptive bet-hedging. *The American Naturalist* 155:168–186.
- COHEN, D. 1966. Optimizing reproduction in a randomly varying environment. *Journal of Theoretical Biology* 12:119–129.
- COOPER, W. S. AND KAPLAN, R. H. 1982. Adaptive “coin-flipping”: a decision-theoretic examination of natural selection for random individual variation. *Journal of Theoretical Biology* 94:135–151.
- COVER, T. M. AND THOMAS, J. A. 1991. Elements of Information Theory. Wiley Series in Telecommunications. John Wiley & Sons, New York.
- CRUMP, M. L. 1981. Variation in propagule size as a function of environmental uncertainty for tree frogs. *The American Naturalist* 117:724–737.
- DANFORTH, B. N. 1999. Emergence dynamics and bet hedging in a desert bee, *Perdita portalis*. *Proceedings of the Royal Society of London, Series B* 266:1985–1994.
- DE MAZANCOURT, C. AND DIECKMANN, U. 2004. Trade-off geometries and frequency-dependent selection. *The American Naturalist* 164:765–778.
- DEMPSTER, E. R. 1955. Maintenance of genetic heterogeneity. *Cold Spring Harbor Symposia on Quantitative Biology* 20:25–32.
- DENVER, R. J., MIRHADI, N., AND PHILLIPS, M. 1998. Adaptive plasticity in amphibian metamorphosis: Response of *Scaphiopus hammonii* tadpoles to habitat desiccation. *Ecology* 79:1859–1872.

- DEWITT, T. J. 1998. Costs and limits of phenotypic plasticity: Tests with predator-induced morphology and life history in a freshwater snail. *Journal of Evolutionary Biology* 11:465–480.
- DEWITT, T. J. AND LANGERHANS, R. B. 2004. Integrated solutions to environmental heterogeneity: Theory of multimoment reaction norms, pp. 98–111. In T. J. DeWitt and S. M. Scheiner (eds.), *Phenotypic Plasticity: Functional and Conceptual Approaches*. Oxford University Press, New York.
- DEWITT, T. J., SIH, A., AND WILSON, D. S. 1998. Costs and limits of phenotypic plasticity. *Trends in Ecology and Evolution* 13:77–81.
- EINUM, S. AND FLEMING, I. A. 2004. Environmental unpredictability and offspring size: conservative versus diversified bet-hedging. *Evolutionary Ecology Research* 6:443–455.
- EVANS, M. E. K., FERRIÈRE, R., KANE, M. J., AND VENABLE, D. L. 2007. Bet hedging via seed banking in desert evening primroses (*Oenothera*, *Onagraceae*): Demographic evidence from natural populations. *The American Naturalist* 169:184–194.
- FRANK, S. A. AND SLATKIN, M. 1990. Evolution in a variable environment. *The American Naturalist* 136:244–260.
- FRANKINO, W. A. AND PFENNIG, D. W. 2001. Condition-dependent expression of trophic polyphenism: effects of individual size and competitive ability. *Evolutionary Ecology Research* 3:939–951.
- GILLESPIE, J. 1973. Polymorphism in random environments. *Theoretical Population Biology* 4:193–195.
- GODFREY-SMITH, P. 1996. *Complexity and the Function of Mind in Nature*. Cambridge Studies in Philosophy and Biology. Cambridge University Press, Cambridge.
- GOODMAN, D. 1984. Risk spreading as an adaptive strategy in iteroparous life histories. *Theoretical Population Biology* 25:1–20.
- GRAFEN, A. 1999. Formal darwinism, the individual-as-maximizing-agent analogy and bet-hedging. *Proceedings of the Royal Society of London, Series B* 266:799–803.
- HACCOU, P. AND IWASA, Y. 1995. Optimal mixed strategies in stochastic environments. *Theoretical Population Biology* 47:212–243.

- HALDANE, J. B. S. AND JAYAKAR, S. D. 1963. Polymorphism due to selection of varying direction. *Journal of Genetics* 58:237–242.
- HALKETT, F., HARRINGTON, R., HULLÉ, M., KINDLMANN, P., MENU, F., RISPE, C., AND PLANTEGENEST, M. 2004. Dynamics of production of sexual forms in aphids: Theoretical and experimental evidence for adaptive “coin-flipping” plasticity. *The American Naturalist* 163:E112–E125.
- HASSALL, M., WALTERS, R. J., TELFER, M., AND HASSALL, M. R. J. 2006. Why does a grasshopper have fewer, larger offspring at its range limits? *Journal of Evolutionary Biology* 19:267–276.
- HOPPER, K. R. 1999. Risk-spreading and bet-hedging in insect population biology. *Annual Review of Entomology* 44:535–560.
- KAPLAN, R. H. AND COOPER, W. S. 1984. The evolution of developmental plasticity in reproductive characteristics: An application of the “adaptive coin-flipping” principle. *The American Naturalist* 123:393–410.
- KELLY, J. L. 1956. A new interpretation of information rate. *Bell System Technical Journal* 35:917–926.
- KOOPS, M. A., HUTCHINGS, J. A., AND ADAMS, B. K. 2003. Environmental predictability and the cost of imperfect information: influences on offspring size variability. *Evolutionary Ecology Research* 5:29–42.
- KUSSELL, E. AND LEIBLER, S. 2005. Phenotypic diversity, population growth and information in fluctuating environments. *Science* 23:2075–2078.
- LANE, S. J. AND MAHONY, M. J. 2002. Larval anurans with synchronous and asynchronous development periods: Contrasting responses to water reduction and predator presence. *Journal of Animal Ecology* 71:780–792.
- LEIMAR, O. 2005. The evolution of phenotypic polymorphism: Randomized strategies versus evolutionary branching. *The American Naturalist* 165:669–681.
- LEVINS, R. 1962. The theory of fitness in a heterogeneous environment. I. The fitness set and adaptive function. *The American Naturalist* 96:361–373.
- LEVINS, R. 1964. Theory of fitness in a heterogeneous environment. III. The response to selection. *Journal of Theoretical Biology* 7:224–240.

- LEVINS, R. 1965. The theory of fitness in a heterogeneous environment. V. Optimal genetic systems. *Genetics* 52:891–904.
- LEVINS, R. 1968. Evolution in changing environments. Princeton University Press, Princeton, New Jersey.
- MAYNARD SMITH, J. 1982. Evolution and the Theory of Games. Cambridge University Press, Cambridge.
- MCNAMARA, J. M. 1995. Implicit frequency dependence and kin selection in fluctuating environments. *Evolutionary Ecology* 9:185–203.
- MENU, F., ROEBUCK, J.-P., AND VIALA, M. 2000. Bet-hedging diapause strategies in stochastic environments. *The American Naturalist* 155:724–734.
- METZ, J. A. J., NISBET, R. M., AND GERITZ, S. A. H. 1992. How should we define ‘fitness’ for general ecological scenarios? *Trends in Ecology and Evolution* 7:198–202.
- MOREY, S. R. AND REZNICK, D. N. 2004. The relationship between habitat permanence and larval development in California spadefoot toads: Field and laboratory comparisons of developmental plasticity. *Oikos* 104:172–190.
- MURPHY, G. I. 1968. Pattern in life history and the environment. *The American Naturalist* 102:391–403.
- PHILIPPI, T. 1993. Bet-hedging germination of desert annuals: Variation among populations and maternal effects in *Lepidium lasiocarpum*. *The American Naturalist* 142:488–507.
- PHILIPPI, T. AND SEGER, J. 1989. Hedging one’s evolutionary bets, revisited. *Trends in Ecology and Evolution* 4:41–44.
- PHILIPPI, T. AND SEGER, J. 2001. Habitat ephemerality and hatching fractions of a diapausing anostracan (Crustacea: Branchiopoda). *Israel Journal of Zoology* 47:387–395.
- RELYEA, R. A. 2002. Costs of phenotypic plasticity. *The American Naturalist* 159:272–282.
- RUEFFLER, C., VAN DOOREN, T. J. M., AND METZ, J. A. J. 2004. Adaptive walks on changing landscapes: Levins’ approach extended. *Theoretical Population Biology* 65:165–178.

- SAIAH, H. AND PERRIN, N. 1990. Autumnal vs. spring hatching in the fairy shrimp *Siphonophanes grubii* (Dybowski) (Crustacea, Anostraca): diversified bet-hedging strategy? *Functional Ecology* 4:769–775.
- SASAKI, A. AND ELLNER, S. 1995. The evolutionarily stable phenotype distribution in a random environment. *Evolution* 49:337–350.
- SEGER, J. AND BROCKMANN, H. J. 1987. What is bet-hedging?, pp. 182–211. In P. Harvey and L. Partridge (eds.), *Oxford Surveys in Evolutionary Biology*, volume 4. Oxford University Press, Oxford.
- SIMONS, A. M. AND JOHNSTON, M. O. 2003. Suboptimal timing of reproduction in *Lobelia inflata* may be a conservative bet-hedging strategy. *Journal of Evolutionary Biology* 16:233–243.
- SKÚLASON, S. AND SMITH, T. B. 1995. Resource polymorphisms in vertebrates. *Trends in Ecology and Evolution* 10:366–370.
- SLATKIN, M. AND LANDE, R. 1976. Niche width in a fluctuating environment—density independent model. *The American Naturalist* 110:31–55.
- SMITH, T. B. 1993. Disruptive selection and the genetic basis of bill size polymorphism in the African finch *Pyrenestes*. *Nature* 363:618–620.
- SMITS, W. K., KUIPERS, O. P., AND VEENING, J.-W. 2006. Phenotypic variation in bacteria: the role of feedback regulation. *Nature Reviews Microbiology* 4:259–271.
- TULJAPURKAR, S. 1990. *Population Dynamics in Variable Environments*. Lecture Notes in Biomathematics. Springer-Verlag, Berlin.
- VAN KLEUNEN, M., FISCHER, M., AND SCHMID, B. 2000. Costs of plasticity in foraging characteristics of the clonal plant *Ranunculus reptans*. *Evolution* 54:1947–1955.
- VENABLE, D. L. 2007. Bet hedging in a guild of desert annuals. *Ecology* 88:1086–1090.
- WILSON, D. S. AND YOSHIMURA, J. 1994. On the coexistence of specialists and generalists. *The American Naturalist* 144:692–707.
- YOSHIMURA, J. AND CLARK, C. W. 1991. Individual adaptations in stochastic environments. *Evolutionary Ecology* 5:173–192.

YOSHIMURA, J. AND JANSEN, V. A. A. 1996. Evolution and population dynamics in stochastic environments. *Researches in Population Ecology* 38:165–182.