The adaptive value of developmental plasticity in response to individually variable cues

Matina C. Donaldson-Matasci [*†]   Carl T. Bergstrom [‡]

Michael Lachmann [§]

February 8, 2011

Abstract

Developmental plasticity is one important way that organisms manage the risk created by a constantly changing environment. Information collected during development, if it helps predict environmental changes, can be used to guide development towards a suitable phenotype. This kind of predictive plasticity can be quite successful. On the other hand, if the environment is completely unpredictable, a last-ditch strategy is bet-hedging—randomizing the phenotype so that at least some offspring will get it right by chance. In situations where only partially reliable information is available, a well-adapted developmental strategy may strike a balance between predictive plasticity and bet-hedging, depending on how well developmental conditions predict environmental change. In this article, we develop a theoretical model which shows that it is not just the predictive value of developmental cues that matters, but also the correlation between those cues within and between generations. For example, an early spring rain may prompt many seeds to germinate early, with potentially disastrous results for the parent plant's lineage if that year's growing conditions turn out to be poor. However, if there are subtle differences in microclimate that affect the timing of germination, even if some seeds germinate early, others are likely to germinate late. In such circumstances a strongly differentiated response to individually variable developmental conditions can achieve two goals at once: it helps match the phenotype to the predicted environment, and at the same time creates phenotypic diversity—hedging bets in case that prediction is wrong.

1   Introduction

Many species living in highly variable environments also show remarkable variability in phenotype. This phenotypic diversity often stems not from genetic variation but is rather a result of adaptive plasticity during development. For example, acorn barnacles

---

[*]Corresponding author: `matina@email.arizona.edu`

[†]University of Arizona, Tucson, AZ, USA

[‡]University of Washington, Seattle, WA, USA

[§]Max Planck Institute for Evolutionary Anthropology, Leipzig, Germany

develop either a bent or conical shell shape, depending on the presence of predatory snails (Lively, 1986*b*). In desert annual plants, a fraction of the seeds produced by any individual will go into diapause, delaying germination until another year—just in case the growing conditions in the current year are poor (Cohen, 1966). For the barnacles, the decision of which phenotype to produce is controlled by predictive developmental cues, while for the desert annuals it seems to be at least partially random. The latter is a classic example of what is known as bet-hedging: developmental stochasticity as an adaptation to an unpredictable environment (Seger and Brockmann, 1987; Cooper and Kaplan, 1982). What influences the adaptive evolution of these two different modes of developmental control? One important factor is how well the selective environment can be predicted using cues available during development. Another is the way that variation in the environment is distributed among individuals, within and between generations (Frank and Slatkin, 1990). Drought, for example, affects all the desert annuals that have germinated that year, while predatory snails affect only barnacles that settle in certain areas. A number of theoretical models have been developed to explore the combined impact of these two factors (Cohen, 1967; Levins, 1968; Moran, 1992; Haccou and Iwasa, 1995), inspiring numerous empirical tests (e.g. Philippi, 1993; Clauss and Venable, 2000; Lane and Mahony, 2002; Halkett et al., 2004; Kussell et al., 2005; Venable, 2007).

In all this research, however, a potentially important feature of developmental cues has been overlooked. The way that developmental cues vary among individuals, within and between generations, could be just as important as variation in the selective environment. Consider a cue that seeds might use to predict drought, like the amount and timing of the rain that falls during the germination season. First of all, seeds in different locations may be subject to subtly different microclimates, so the rainfall detected by one seed may not precisely reflect the pattern overall. Secondly, weather patterns are variable enough that even a perfect record of the global patterns of rainfall during the germination season would not be enough to predict a drought with certainty. These two sources of error have very different implications for the way that cues are distributed among individuals, yet traditionally this distinction has been ignored.

In this article, we show that the adaptive evolution of developmental plasticity depends on both kinds of variation in developmental cues, but in qualitatively different ways. Errors in prediction that are the result of unusual patterns in the global environment—like early spring rains followed by an extended drought—can select for developmental strategies that produce high levels of within-generation phenotypic diversity. However, the strategies used to generate that diversity may be quite different, depending on how much error there is in individual observations of developmental cues. This has several surprising implications. Intuitively, one might predict that the less reliable a cue is, the weaker the developmental response to it should be. This intuition holds as long as the error in the cue affects all individuals in a generation: a strong developmental response to the cue would mean that any mistake in prediction could lead to disaster for the whole population. However, when each individual makes independent errors in observation of the cue, less reliability can actually select for a stronger developmental response to the cue. Furthermore, even though each individual may do a very poor job of predicting the environment, the developmental strategy may perform quite well on average. Surprisingly, this means that there will often be no selection for individuals to improve their accuracy in observing environmental predictors, even though by doing so they could improve their

ability to predict future environmental conditions.

2 An example: two kinds of error in a desert annual's prediction of drought

Cohen's classic model of bet-hedging was based on the biology of desert annual plants (Cohen, 1966). Because they live in such a harsh environment, and have only one chance to reproduce, these plants are very strongly affected by the possibility of drought. Consider the fate of a plant whose life history strategy is to produce seeds that always germinate the following spring. If the next year happens to be a drought year, the plant will produce no grand-offspring at all—its lineage will come to an end. Even if the next year is not a drought year, sooner or later one will come along. Since all the descendents of the original plant have the same life history strategy, all the seeds in its lineage will germinate despite the drought and fail to reproduce, again ending the lineage. Many desert annual plants get around this problem by bet-hedging: each seed "chooses" randomly whether to germinate or not. In good years, those seeds that do germinate are likely to reproduce successfully. In drought years, even though some seeds will germinate and die, others will stay dormant, retaining the chance to carry on the lineage in some future year. Cohen showed that, as long there is no chance of survival in bad years, the optimal fraction of seeds germinating each year is equal to the fraction of good years. Thus, if severe drought occurs about one year in five, we would expect approximately 80% of seeds to germinate each season.

Bet-hedging is a reasonable strategy as long as there is no information available about the quality of the coming growing season. However, when environmental cues are available to the seed that can help predict future environmental conditions, it may be possible to use that information to do even better (Cohen, 1967). For example, early and plentiful spring rains may indicate that a good growing season is more likely, while late or sparse spring rains may indicate the opposite (Philippi, 1993). Cohen showed that the optimal fraction of seeds germinating in response to a cue is equal to the conditional probability of a good year, given that cue. Say that severe drought occurs about half the time overall, but in years where there are early spring rains there is just a one in five chance of drought, while with late spring rains there is a four in five chance of drought. In this situation we would expect about 80% of seeds to germinate in response to early spring rains, but only about 20% of seeds to germinate in response to late spring rains. Such a strategy combines elements of predictive plasticity, because seeds change their probability of germination in response to cues, and bet-hedging, because not all seeds will germinate even under the best conditions.

The above calculations assume that all individuals have exactly the same perception of the cue—if the spring rains do come early, then all seeds will discern that and respond accordingly, and vice versa. However, different seeds may detect slightly different amounts of rain due to differences in their microhabitat. Say that, in years where the first spring rains sufficient for germination come early, 20% of seeds miss them and germinate late, while in years where the first spring rains sufficient for germination come late, 20% of seeds germinate early anyway due to particularly moist microhabitats. This is effectively the same strategy as described above, but now the within-generation variation is caused

not by developmental stochasticity but by perceptual error in detecting the developmental cue.

This very simple, but rather contrived example illustrates three important points. First of all, variability between the cues that individuals detect can create some level of within-generation phenotypic diversity, spreading the risk among individuals and potentially obviating the need for bet-hedging. Second, even though individual errors in perception make the cue less reliable, the optimal strategy is actually to respond more strongly to the cue, which seems rather counterintuitive. Finally, the phenotypic distributions created by the optimal strategies with and without individual errors in perception may in some cases be the same—which suggests that even though they make cues less reliable, these errors may not actually be costly. In the following section we develop a model of bet-hedging and plasticity in an uncertain environment that is designed to highlight the role of individual errors in perception of cues, and will allow us to explore these three points more thoroughly.

3   A model of developmental plasticity in response to unreliable cues

We model a species with discrete, non-overlapping generations, inhabiting a fluctuating environment. We consider population-level risk, in which the environmental state varies from one generation to the next, affecting every individual in the population. In each generation, the state of the environment e is drawn from a distribution $\Pr(e)$ over a discrete set of possible environmental states.

An individual's phenotype is determined by its developmental trajectory, which is shaped by its genotype, the conditions it experiences during development, and developmental noise. We model this by assuming that the individual's genotype encodes a developmental strategy g. Each individual i receives a developmental cue c_i, and develops a phenotype x_i, which remains fixed over its lifetime. For simplicity, we focus on discrete phenotypes. The developmental strategy g specifies the probability of developing each phenotype, conditional on receiving each developmental cue: $g(x|c)$. Each strategy shows a particular balance of stochasticity and predictive plasticity. One way to quantify the amount of developmental stochasticity, drawn from information theory (Shannon, 1948), is the amount of uncertainty about the phenotype once the cue has been observed: the conditional entropy, $H(X|C)$. Similarly, a measure of predictive plasticity is how much the knowing the cue reduces the amount of uncertainty about the phenotype: the mutual information between the cue and the phenotype, $I(X;C)$.

Plasticity is only useful if individual cues can be used to help predict the environmental state. If the cues do not perfectly reflect the coming environmental conditions, then they contain some error. This error may be either shared among all individuals in the population, or it may be different for different individuals. For example, if early spring rains are usually associated with a good growing season, then the exception—early spring rains followed by a poor growing season—is a misleading cue that could potentially cause many seeds to germinate and then die without being able to reproduce. On the other hand, if different seeds detect slightly different amounts of rain due to differences in their microhabitat, each seed's error in observation affects only its own choice about whether

to germinate.

To represent these two types of error, we posit a shared population-level cue q which varies from one generation to the next, and is correlated with the environmental conditions e. Individuals do not directly observe the population-level cue, but each make their own independent observation of it, which we call the individual-level cue c_i. Mathematically, the process can be described as follows. In each generation, an environmental state e is chosen from the distribution $\Pr(e)$. Second, a population-level cue q is chosen from the conditional distribution $\Pr(q|e)$, which describes the population-level error in the cue. Finally, for each individual i in the population, an individual-level cue c_i is chosen independently from the conditional distribution $\Pr(c|q)$, which describes the individual-level error in the cue. In the example of desert annual plants, the environmental state e would reflect the amount of rain that falls during the growing season, the population-level cue q would reflect when the first rain sufficient for germination falls, and the individual cue c_i would reflect when each individual seed detects sufficient rain for germination. Individual error in observation always reduces the predictive value of a cue: an individual's information about the environment, $I(E; C)$, is always less than the information she would have if she could directly observe the population-level cue, $I(E; Q)$ (by the data processing inequality; e.g. Cover and Thomas, 1991).

The fitness of each individual is determined by its phenotype x_i and the shared environmental state e according to the fitness function $f(x, e)$. In a fluctuating environment, the phenotype with the highest mean fitness is not necessarily most likely to fix in the population. Because the environment varies from generation to generation, in some generations one phenotype will be favored, while in other generations another phenotype will be favored. Over the long term, the genotype that is most likely to fix is the one whose long-term growth rate is greatest. As long as there is no interaction between phenotypes, a good estimator of the long-term growth rate of a genotype is the mean log fitness (which is the log of the geometric mean fitness) (Dempster, 1955; Cohen, 1966; Cooper and Kaplan, 1982). This depends not just on the mean fitness of individuals of that genotype, but also on the variation in mean fitness from one generation to the next (Frank and Slatkin, 1990; Yoshimura and Clark, 1991). The general form for the long-term growth rate of the lineage stemming from genotype g is the mean, over generations, of the log mean fitness within generations. Since generations differ from one another in their environmental state e and the population-level cue q, the long-term growth rate can be written as:

$$
\begin{aligned}
r(g) &= \sum_{e,q} \Pr(e, q) \log \bar{f}(g|e, q) \\
&= \sum_{e} \Pr(e) \sum_{q} \Pr(q|e) \log \frac{1}{n} \sum_{i=1}^{n} f(x_i, e)
\end{aligned}
\tag{1}
$$

Since, in our model, an individual's phenotype x_i is determined by its developmental strategy g and the cue it receives, c_i, a good approximation when the number of individuals in the lineage is large is:

$$
r(g) \approx \sum_{e} \Pr(e) \sum_{q} \Pr(q|e) \log \sum_{c} \Pr(c|q) \sum_{x} g(x|c) f(x, e).
\tag{2}
$$

The mean fitness within a generation depends both on the environmental state, which directly affects the fitness of all individuals, and the population-level cue, which influences the developmental cues that individuals receive, and, in turn, their phenotypes and their fitness. Note that the individual-level error, represented by the term $\Pr(c|q)$, and the population-level error, represented by the term $\Pr(q|e)$, will have very different effects on the long-term growth rate.

This model explicitly assumes that all individual variation in developmental conditions stems from independent errors in observing the same population-level cue. More generally, we might want to consider cases in which individual developmental conditions are correlated with the environmental state and with each other, but there is no obvious population-level cue that each individual is observing. In Appendix A we argue that under wide variety of correlational structures, we can nonetheless mathematically identify a global feature of a particular generation which can be treated like a population-level cue. This feature is the frequency distribution of individual cues in the population, i.e. the proportion of individuals that observe each possible value of the developmental cue.

4   Model Results

Using the model described above, we look at the effect of individual-level error in a cue on the optimal developmental strategy and the fitness it can achieve. Under what conditions is a combination of plasticity and bet-hedging better than plasticity alone—that is, when does the optimal strategy involve some randomization? How does the optimal strategy change as the amount of individual-level error changes? Under what conditions might there be selection to improve individual perception, reducing the amount of individual-level error?

To solve for the optimal developmental strategy, we must find the conditional probability distribution $g(x|c)$ that maximizes the long-term growth rate given in Equation 2. We base our method on the principle of proportional betting: when several outcomes are possible, and only betting on the correct outcome yields any payoff, the optimal scheme for long-term investment is to place money on each outcome according to its probability of occurring—regardless of the associated payoff. For example, in Cohen's classic model of diapause in desert annual plants, the lineage that maximizes its long-term growth rate is the one whose conditional probability of germination in response to a cue equals the conditional probability of a good year (Cohen, 1967). This simple, intuitive result depends on two special assumptions: first, that there is no chance of reproduction in a bad year, and second, that all individuals receive the same cue. In a previous paper we have shown how to extend this principle to cases where phenotypes can survive in multiple environments (Donaldson-Matasci et al., 2008). Because our focus in this article is relaxation of the second assumption, here we will assume that each phenotype survives in just one environment in order to keep things simple. However, in Appendix B we show that our results still hold when the first assumption is relaxed as well.

With individual-level error, because different individuals may receive different cues even within the same generation, the principle of proportional betting does not directly apply. However, we can still solve for the optimal response to the population-level cue, as if individuals could observe it perfectly, and then ask whether that response could still

be achieved by a developmental strategy that could only respond to cues with individual-level error. According to the principle of proportional betting, the optimal strategy in response to the population-level cue should match the probability of each phenotype to the conditional probability that its environment occurs:

$$\hat{g}(x|q) = \Pr(e_x|q). \tag{3}$$

The effective response to the population-level cue is determined both by the developmental response to individual cues, and the probability of observing those cues given the population-level cue:

$$\bar{g}(x|q) = \sum_c g(x|c) \Pr(c|q). \tag{4}$$

In the following sections we will use row stochastic matrices to simplify the notation. The developmental strategy is represented by a matrix $\mathbf{G}$ with entries $g(x|c)$, the individual error is represented by a matrix $\mathbf{C}$ with entries $\Pr(c|q)$, and the population-level error is represented by a matrix $\mathbf{P}$ with entries $\Pr(e|q)$. Putting together Equations 3 and 4 and using matrix notation, we find the following criterion for the optimal developmental strategy in response to individual cues, $\hat{g}(x|c)$:

$$\mathbf{C}\hat{\mathbf{G}} = \mathbf{P}. \tag{5}$$

This simple formulation will allow us to make several interesting observations about the effect of individual-level error on the optimal developmental strategy and the resulting long-term growth rate.

4.1 Individual error may make bet-hedging unnecessary

Depending on how variable environmental conditions are, individual level error in the perception of developmental cues may make bet-hedging superfluous. Figure 1 shows an example for the case with two phenotypes and two possible values for the cue. For low levels of individual-level error, bet-hedging in response to both cues is almost always advantageous, while for higher levels of individual-level error, bet-hedging in response to both cues is only helpful when there are high levels of population-level error as well.

In general, as the level of individual-level error increases, the amount of population-level error necessary to drive the evolution of bet-hedging increases as well. Equation 5 describes the conditions under which full bet-hedging can be favored; when it cannot be satisfied with a stochastic matrix $\hat{\mathbf{G}}$, the best achievable strategy will be a boundary solution involving only partial bet-hedging in response to some cues, or no bet-hedging at all. This occurs when the point in conditional probability space defined by $\mathbf{P}$ falls outside of the parallelopiped created by $\mathbf{C}$ (e.g. the grey parallelogram in Figure 1.) The volume of that parallelopiped, $|\det(\mathbf{C})|$, is 1 when individuals can observe the population-level cue directly, is 0 when the observation is completely independent of the population-level cue, and decreases as the amount of individual-level error increases.

4.2 A stronger response to cues with individual-level error

Any type of error in the developmental cue an individual perceives reduces the amount of information in that cue, and makes it harder to predict environmental conditions. One
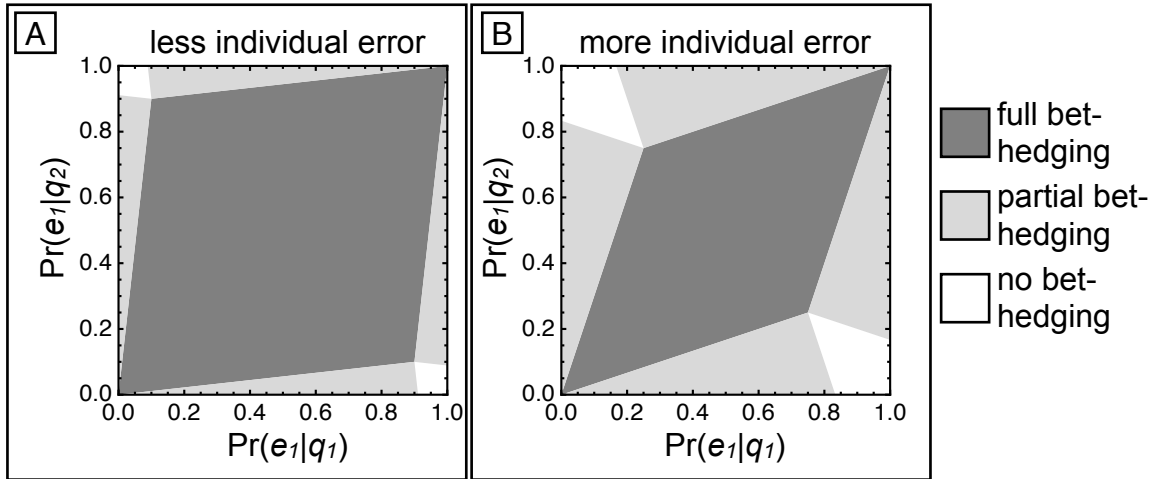
Figure 1: Individual error may make bet-hedging unnecessary. Given a particular level of environmental uncertainty, represented by the conditional probabilities on each axis, the optimal developmental strategy may involve full, partial or no bet-hedging. In the central dark grey regions, full bet-hedging—randomizing in response to both values of the cue—would be optimal; in the lighter grey regions, partial bet-hedging—randomizing in response to just one cue—would be optimal; in the white corner regions, bet-hedging provides no advantage at all. With just a 10% chance that an individual will mistake the value of the cue (panel A: $\Pr(c_1|q_1) = \Pr(c_2|q_2) = 0.90$), there is a broader range of conditions where full bet-hedging is optimal than there is in with a 25% chance of an individual mistake (panel B: $\Pr(c_1|q_1) = \Pr(c_2|q_2) = 0.75$.)

might expect individuals to pay more attention to more reliable cues, and less attention to less reliable cues. This is in fact the case—as long as any changes in the cue's reliability are created by changes in the amount of population-level error in the cue (see Figure 2A.) However, when the cue becomes less informative because individual-level error in the cue is increased, the best strategy actually strengthens its response to the cue by developing in a more directed, less randomized way (see Figure 2B.) To see this more generally, we know from Equation 5 that $|\det(\mathbf{C})| \cdot |\det(\hat{\mathbf{G}})| = |\det(\mathbf{P})|$. As argued earlier, an increase in individual error will decrease the absolute value of the determinant of $\mathbf{C}$; to compensate, the absolute value of the determinant of $\hat{\mathbf{G}}$ must increase. This corresponds to stronger predictive plasticity, with more differentiation in the developmental responses to different cues.

Why the difference between the two types of error? Bet-hedging is an adaptation to environmental uncertainty shared by an entire population. As long as all individuals in the population receive the same misleading cue, the uncertainty created by any error in the cue is population-level uncertainty. It is this kind of uncertainty that determines how a developmental strategy should invest its effort into different phenotypes within and between generations. On the other hand, if some individuals perceive a common developmental cue incorrectly, this creates uncertainty about the environment at the individual level—for which bet-hedging is no use. This explains why individual error does not increase bet-hedging, but why should it actually decrease the optimal amount
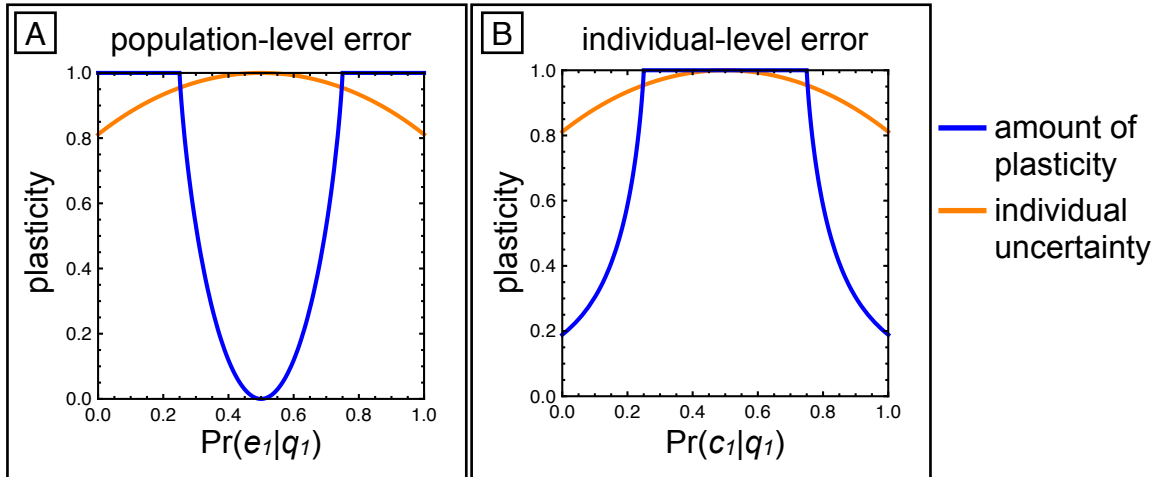
Figure 2: Individual error increases the optimal strength of response to a cue. Population-level error (A) and individual-level error (B) both increase the amount of uncertainty an individual has about the state of the environment, making it harder to predict which phenotype to develop. When individual error is fixed, an increase in population-level error means that a more stochastic, less plastic developmental strategy is better (panel A: $\Pr(c_1|q_1) = \Pr(c_2|q_2) = 0.75$.) On the other hand, increasing individual-level error while population-level error remains fixed means that a less stochastic, more plastic developmental strategy is better (panel B: $\Pr(e_1|q_1) = \Pr(e_2|q_2) = 0.75$.) Individual uncertainty is measured as the conditional entropy of the environment given a cue, $H(E|C)$, while the amount of plasticity is measured as the mutual information between the cue and the phenotype, $I(X;C)$.

of bet-hedging? Developing different phenotypes in response to a cue with individual error can actually achieve the same purpose as bet-hedging in response to a cue with population error; in both cases, different individuals within the same generation develop differently.

4.3   Individual error is not always detrimental

How does individual error in a cue affect fitness, measured in terms of the long-term growth rate? One way to measure the value of a cue is to compare the optimal growth rate with the cue to the optimal growth rate with no cue at all. This indicates the cue's potential to increase the growth rate of a lineage, if the developmental strategy manages to take full advantage of the information it contains. According to this measure, the value of a cue is equal to the amount of information it contains about the environment, $I(E;Q)$, as long as that cue is observed by all members of the population and the optimal response can be achieved (Donaldson-Matasci et al., 2010). How does the value of a cue with individual error relate to the amount of information it carries about the environment? As an example, if the population-level error in the cue is 25%, as long as the individual-level error is below 25%, there is some developmental strategy that can do just as well as a lineage that could observe the population cue perfectly (see Figure 3B.) In general,
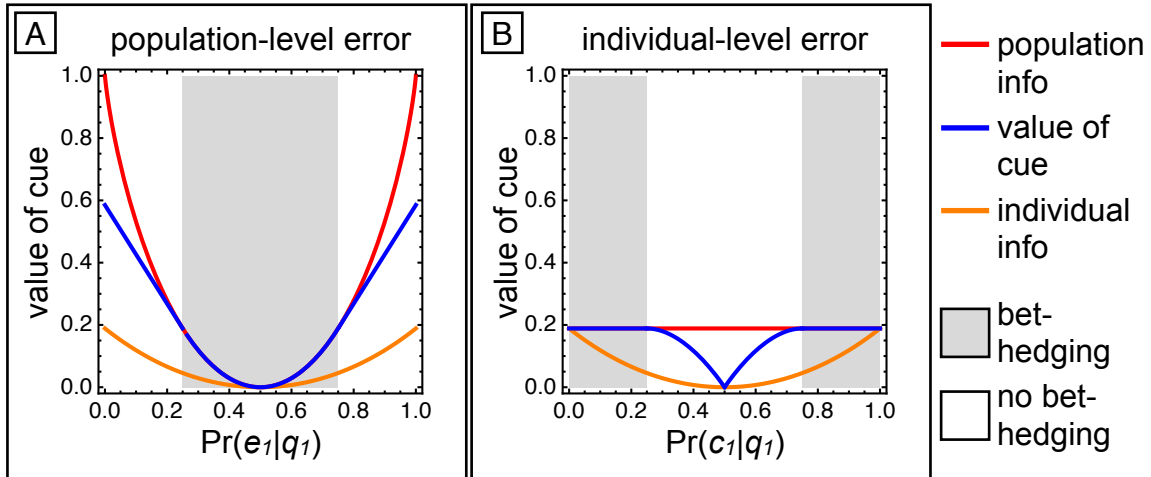
Figure 3: Individual error does not always decrease the value of a cue. When individual-level error is fixed (panel A: $\Pr(c_1|q_1) = \Pr(c_2|q_2) = 0.75$), any increase in population-level error decreases the value of the cue, measured as the difference in optimal growth rate with and without the cue. In contrast, when population-level error is fixed (panel B: $\Pr(e_1|q_1) = \Pr(e_2|q_2) = 0.75$), as long as individual-level error is lower, it does not decrease the value of the cue. Note that wherever full bet-hedging is optimal, the value of the cue is exactly equal to the amount of information in the population-level cue about the environment, $I(E;Q)$. In contrast, the value of the cue is always greater than the amount of information individuals actually gain about the environment from their developmental cues, $I(E;C)$.

adding individual-level error in the observation of a cue has no effect on the cue's value at all, as long as the error is small enough. This is because whenever the conditions for Equation 5 are met, the optimal response to the population-level cue can actually be achieved, even though individuals do not observe the cue directly. Thus, as long as full bet-hedging is still optimal, individual-level error does not decrease the value of a cue. This is surprising, because it means that the growth rate is not limited by the amount of information that individuals actually have, $I(E;C)$; it is limited by the amount of information that they would have if they could make perfect observations without any error, $I(E;Q)$.

So far we have considered the evolution of the developmental strategy in response to a cue, assuming that all sources of error in that cue are immutable. However, if individual error in the cue is due to something like mistakes in perception, or an insufficiently accurate sensor, there may be an opportunity for its accuracy to improve. One might imagine that this would always be advantageous, because any amount of error in a cue would always make it harder to predict the best course of action. However, as long as there is enough error in the cue at the population level, a certain amount of individual-level error is tolerable and perhaps even useful. In Figure 4A we show the fitness landscape that results when the probability of individual errors in perception and the developmental strategy are allowed to evolve simultaneously. When individual-level error is large enough that bet-hedging is no longer favored, there will be selection to improve individ-
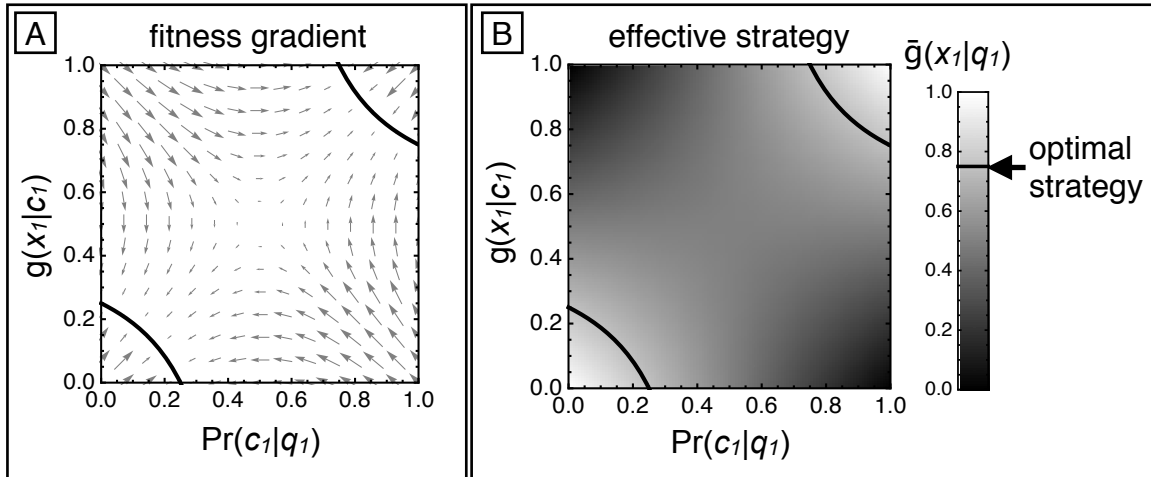
Figure 4: A certain level of individual error may actually be advantageous. Here we allow change in both the chance of individual error and the developmental strategy. To make visualization easier, we stipulate that both the error and the developmental strategy change symmetrically, that is, $\Pr(c_2|q_2) = \Pr(c_1|q_1)$ and $g(x_2|c_2) = g(x_1|c_1)$. The fitness gradient, as measured by the derivatives of the long-term growth rate $r(g)$, shows two lines of equilibria combining various levels of individual error and developmental stochasticity (panel A.) Together, each combination of individual error and developmental strategy achieve an effective strategy $\bar{g}(x|q)$, which is the average response to the population-level cue q (panel B.) Along the lines of equilibria, all combinations achieve exactly the same effective strategy, $\bar{g}(x|q) = \Pr(e|q) = 0.75$.

ual perception, because it sets a constraint on the growth rate that can be achieved. On the other hand, when individual-level error is small enough that bet-hedging is favored, evolution may act to modify either the amount of error, the developmental strategy, or both. A stable state is reached when the amount of stochasticity in perception and development together create just the right levels of phenotypic diversity in response to the population-level cue (see Figure 4B.)

5 Discussion

The importance of the distinction between environmental variation within and between generations is well recognized where fitness functions are concerned (e.g. Seger and Brockmann, 1987; Frank and Slatkin, 1990; Moran, 1992; Robson, 1996). It stands to reason that, for developmentally plastic organisms, distinguishing within- and between-generation variation in developmental cues might be just as important. Previous models have generally considered only developmental cues that vary at the same level as the environment, either both at the individual level (e.g. Lively, 1986a; Moran, 1992), or both at the population level (e.g. Cohen, 1967; Moran, 1992; Haccou and Iwasa, 1995). The only exception, to our knowledge, is an interesting mathematical paper which shows, as discussed in Section 4.3, that the value of a cue with both individual-level and population-level error is limited by the amount of information in the population-level cue (Rivoire

and Leibler, 2010). In all of those models, the usefulness of the cue is considered to lie in its ability to allow individuals to develop a phenotype that is well-matched to the environment. However, there is another potential role that developmental cues can play. One early article on the evolution of development in response to population-level environmental uncertainty called the production of random phenotypic diversity "adaptive coin-flipping" (Cooper and Kaplan, 1982). The authors suggested that in some cases, organisms could evolve to be particularly sensitive to features of their developmental environment that are completely independent of the population-level environmental variation (also see Simons and Johnston, 1997, 2006). Those features carry no predictive information, but are nonetheless useful because they can act as the coin flip on which the organism's developmental decision is based. These two viewpoints suggest that responding to developmental cues can be adaptive either because they provide useful predictive information, or because they provide a source of randomness. We show that in fact a single developmental cue can potentially play both roles.

An optimal developmental strategy, when it can be achieved, balances predictive plasticity and bet-hedging in proportion to the amount of information available in population-wide environmental cues and the uncertainty that remains once those cues have been observed. Predictive plasticity creates variation in phenotype across generations, helping match up individuals' phenotypes to environmental fluctuations, while bet-hedging creates variation in phenotype within generations, just in case predictions are incorrect. Individual-level variation in developmental cues can affect this scenario in two ways. On one hand, individual mistakes in the observation of developmental cues reduce the amount of information each individual collects about the environment, potentially interfering with an individual's ability to successfully predict environmental conditions. On the other hand, individual-level error in a developmental cue can contribute directly to the production of phenotypic diversity, essentially acting as a coin flip for any developmental strategy that responds plastically to that cue. As long as the individual-level error is not too high, it may be used to create adaptive phenotypic diversity for the purpose of bet-hedging. In this case, the error acts solely as a coin-flip to generate within-generation diversity. Even though individuals may do a poor job of prediction, the right developmental strategy on average can do just as good a job of responding plastically to the population-wide environmental cue as another one would do without individual-level error. However, once the amount of individual-level error gets too large, no developmental strategy can be responsive enough to population-level variation in the cue; too much phenotypic variation will be created within generations, and not enough variation between generations. In this situation, we argue that there could be selection to reduce individual variation in cues, for example by improving individual perception.

A large number of empirical tests have now suggested that bet-hedging could play an adaptive role in a variety of life history traits, such as delayed germination of seeds (Philippi, 1993; Evans et al., 2007; Venable, 2007; Petru and Tielboerger, 2008; Simons, 2009), timing of metamorphosis in anurans (Lane and Mahony, 2002; Morey and Reznick, 2004; Richter-Boix et al., 2006), arthropod diapause (Saiah and Perrin, 1990; Danforth, 1999; Philippi et al., 2001; Menu and Desouhant, 2002), egg size and number (Koops et al., 2003), and bacterial persistence (Balaban et al., 2004; Kussell et al., 2005; Acar et al., 2008). The most quantitatively rigorous of these show a correlation between delayed germination and year-to-year variation in fitness (Venable, 2007), or use simulations

to show that the observed amount of variation in the timing of germination could promote long-term survival (Evans et al., 2007; Simons, 2009). The difficulty of such quantitative evaluations is that the fitness functions and environmental probabilities are very difficult to measure, since they may vary on quite long time scales. Empirical studies that look at both plasticity and bet-hedging as potential adaptations to environmental uncertainty are still quite rare, but a few examples suggest strategies that combine the two (e.g. Danforth, 1999; Richter-Boix et al., 2006; Sadeh et al., 2009; Khatchikian et al., 2010). Most of these studies focus on identifying features of the global environment that could be acting as predictive cues, like the effects of temperature and precipitation on germination fraction (Freas and Kemp, 1983; Philippi, 1993; Khatchikian et al., 2010). Some have also used experimental manipulations to pinpoint which features might directly cause the developmental response (Clauss and Venable, 2000; Adondakis and Venable, 2004; Morey and Reznick, 2004; Simons and Johnston, 2006; Sadeh et al., 2009). Our theoretical results have two important implications for these studies. We show that even when the developmental mechanisms are unknown, environmental variation can select for certain levels of within- and between-generation phenotypic variation. This means that an adaptive fit can be assessed just by looking at the effective response to population-level cues like the weather, even if there is variation in the way individuals experience those cues. However, we argue that in order to understand the adaptive role of the developmental mechanisms at play, it is necessary to take into account natural variation in developmental cues, both within and between generations. A strong developmental response to a cue with high variation within generations may just be another kind of adaptive coin-flipping.

Appendices

A   A population-level cue describes the distribution of individual cues

In our model, we assume that the developmental cues that individuals receive are produced by a two-step process. The first step is the generation of a population-level cue, which is common to all individuals within a generation. In the second step, each individual's cue is generated independently of all other cues, conditional only on the population-level cue. In many cases, however, it may be possible to identify only the individual-level cues, which are generated according to some unknown process. In this appendix, we argue that our two-step model is nonetheless often a good description, and show that the population-level cue can be defined as the distribution of individual cues.

Say that in each generation, we can observe the environmental state e, as well as the developmental cue c_i for each individual i. In each generation, we will keep track of the fate of every individual in a single lineage by giving each individual an index i from 1 to n. We write a sequence representing one developmental cue for each individual in the lineage as $c^n = c_1, ..., c_n$. Similarly, we use a sequence $x^n = x_1, ..., x_n$ to represent the phenotypes x_i adopted by all individuals in the lineage. The overall strategy for the lineage, $g(x^n|c^n)$, represents the probability that the sequence of phenotypes x^n is produced, given the sequence of developmental cues c^n. The total reproductive output of the lineage in one generation, $f(x^n, e)$, depends on the phenotype of each individual and

on the environmental state.

Very generally, we can describe the relationship between the environmental state and the individual cues using the joint distribution $\Pr(e, c^n)$. These two things together can be used to define the type of a generation, in terms of its effect on the per-capita fitness of a genotype. The expected long-term growth rate for the strategy g is therefore:

$$
\begin{aligned}
r(g) &= \sum_{e,c^n} \Pr(e, c^n) \log \bar{f}(g|e, c^n) \\
&= \sum_{e,c^n} \Pr(e, c^n) \log \frac{1}{n} \sum_{x^n} g(x^n|c^n) f(x^n, e).
\end{aligned}
\tag{A-1}
$$

To find a simpler form, we would like to distinguish between generations where the per-capita fitness is different, but lump together those that are the same. We will argue that, under quite reasonable assumptions, it is not important to keep track of exactly which individual received which developmental cue.

The first assumption is that each individual's fitness depends only on its own phenotype, not on the phenotypes of other individuals. The second assumption is that each individual's phenotype depends only on its own developmental cue, not on the cues received by others. Given these two assumptions, the per-capita fitness of a strategy within a particular generation can be rewritten:

$$
\bar{f}(g|e, c^n) = \frac{1}{n} \sum_{i=1}^{n} \sum_{x} g(x_i = x|c_i) f(x, e).
\tag{A-2}
$$

Notice that this quantity does not depend on which individual received which cue, but rather what fraction of all individuals received each different type of cue. It is therefore sufficient to distinguish between generations that have different distributions of developmental cues. We use an index θ to describe a distribution of individual-level cues, yielding a simpler form for the per-capita fitness of a strategy within a generation:

$$
\bar{f}(g|e, \theta) = \sum_{c} \Pr(c|\theta) \sum_{x} g(x|c) f(x, e).
\tag{A-3}
$$

A comparison to Equation 2 in the main text shows that θ plays the same role as the population-level cue q.

A famous theorem due to de Finetti proves a related result for exchangeable variables (see e.g. Feller, 1966). A finite sequence of random variables $C_1, ..., C_n$ is exchangeable if every permutation of these variables has the same probability distribution, and an infinite sequence is exchangeable if every finite subsequence is exchangeable. De Finetti's theorem states that any infinite, exchangeable sequence of binary random variables can be understood as a draw from some probability distribution of a single random parameter Θ, followed by a sequence of independent Bernoulli trials weighted according to the parameter θ. When generalized to discrete-valued (rather than only binary) random variables, the parameter θ describes the probability distribution of individual cues (Hewitt and Savage, 1955). We argue that even if the developmental cues received by different individuals in the same generation are not exchangeable variables, they nonetheless act as if exchangeable, as far as per-capita within-generation fitness is concerned. The only

caveat comes when we consider competition between lineages. If the developmental cues are more strongly correlated within a lineage than they are between two competing lineages, this simple breakdown into population-level and individual-level cues may not be warranted.

B   Results also hold when phenotypes can survive in multiple environments

In the main text, we have assumed that each phenotype can survive in just one environment; however, most real phenotypes are not that specialized. Generalizing the fitness function so that each phenotype can survive in multiple environments means that the principle of proportional betting cannot be used directly to identify the optimal phenotypic proportions. The key to the method developed in Donaldson-Matasci et al. (2008) is the identification of a set of hypothetical phenotypes, each of which survive in just one environment, but which, when combined in the right proportions, yield the same average fitness in each environment as the real phenotypes. This relationship can be nicely captured in matrix notation as follows:

$$\mathbf{F} = \mathbf{SD}. \tag{B-1}$$

Each entry of the fitness matrix $\mathbf{F}$ is the fitness $f(x, e)$; the diagonal matrix $\mathbf{D}$ represents the fitness of each hypothetical phenotype, $f(y_e, e)$; and the specialization matrix $\mathbf{S}$ describes the combination $s(y|x)$ of hypothetical phenotypes y that corresponds to each real phenotype x. The specialization function $s(y|x)$ can be interpreted as a measure of how invested each real phenotype is in each environment where it can survive. The advantage of rewriting the fitness in terms of these hypothetical phenotypes is that the principle of proportional betting applies to them: we want to create a combination of phenotypes x which is effectively like matching the proportion of y to the probability of the corresponding environment: $\sum_x s(y|x)\hat{g}(x) = \Pr(e_y)$.

The same idea works when we condition on cues that are common to all individuals in the population. If we write the developmental strategy $g(x|q)$ as a matrix $\mathbf{G}$, then the within-generation expected fitness $\bar{f}(g|q, e) = \sum_x g(x|q)f(x, e)$ can now be written conveniently in matrix notation as follows:

$$\mathbf{GF} = \mathbf{GSD}. \tag{B-2}$$

The principle of proportional betting says that the optimal developmental strategy in response to cues with no individual-level error, $\hat{g}(x|q)$, should meet the following criterion:

$$\mathbf{\hat{G}S} = \mathbf{P}. \tag{B-3}$$

Finally, we would like to add individual-level error in the cue. In this situation, the within-generation expected fitness is $\bar{f}(g|q, e) = \sum_c \Pr(c|q) \sum_x g(x|c)f(x, e)$, now written in matrix notation as:

$$\mathbf{CGF} = \mathbf{CGSD}. \tag{B-4}$$

We know that the optimal strategy in response to the population-level cue is given by Equation B-3; can the same effective strategy be achieved in response to an individual-level cue? This will be possible when there exists a stochastic matrix $\mathbf{\hat{G}}$ such that

$$\mathbf{C\hat{G}S} = \mathbf{P} \tag{B-5}$$

(compare to Equation 5 in the main text.)

In Section 4.1, we describe the conditions under which an optimal strategy can actually be achieved. These conditions are more restrictive when phenotypes can survive in multiple environments. The matrix $\mathbf{S}$ defines a hyperrectangular region in the conditional probability space $\Pr(e|q)$ where bet-hedging is advantageous as long as there is no individual-level error. When individual-level error is present, the matrix $\mathbf{C}$ acts as a linear transformation on that region, changing its volume by a factor equal to the absolute value of its determinant. As argued in the main text, as the individual-level error gets larger, the determinant of $\mathbf{C}$ moves towards zero, shrinking the potential range that $\mathbf{P}$ can fall into. Thus the first point generalizes: the amount of population-level error necessary to drive bet-hedging increases as individual-level error increases, even with an arbitrary fitness function.

In Section 4.2, we argue that increasing individual-level error will strengthen the optimal developmental response to the cue. Trading Equation B-5 for Equation 5, we see that the optimal strategy must satisfy $|\det(\mathbf{C})| \cdot |\det(\hat{\mathbf{G}})| = |\det(\mathbf{PS^{-1}})|$. Thus, as long as the fitness function (which determines $\mathbf{S}$) and the population-level error (which determines $\mathbf{P}$) remain the same, any decrease in the determinant of $\mathbf{C}$ will induce a corresponding increase in the determinant of $\hat{\mathbf{G}}$. Thus the second point generalizes for arbitrary fitness functions as well: the optimal amount of developmental plasticity increases as the amount of individual-level error in the cue increases.

In Section 4.3, we argue that whenever full bet-hedging is optimal, the value of a cue is equal to the information its population-level component carries about the environmental state. This argument draws on the work in (Donaldson-Matasci et al., 2010), where we show that the optimal strategy is the one that effectively invests in each environment according to its conditional probability. The argument here is exactly the same, but now the effective strategy is influenced not only by the fitness function, which determines the investment each phenotype puts into each environment, $\mathbf{S}$, but also the individual error $\mathbf{C}$, which scrambles the response to the population-level cue. Whenever full bet-hedging is favored, that means that the conditions of Equation B-5 are met, and the optimal strategy can actually be achieved. Under these conditions, there may be no selection to reduce individual variation in the developmental cues received. Thus the third point also generalizes to arbitrary fitness functions.

References

Acar, M., J. T. Mettetal, and A. van Oudenaarden. 2008. Stochastic switching as a survival strategy in fluctuating environments. Nature Genetics 40:471–475.

Adondakis, S., and D. Venable. 2004. Dormancy and germination in a guild of Sonoran Desert annuals. Ecology 85:2582–2590.

Balaban, N., J. Merrin, R. Chait, L. Kowalik, and S. Leibler. 2004. Bacterial persistence as a phenotypic switch. Science 305:1622–1625.

Clauss, M. J., and D. L. Venable. 2000. Seed germination in desert annuals: An empirical test of adaptive bet-hedging. American Naturalist 155:168–186.

Cohen, D. 1966. Optimizing reproduction in a randomly varying environment. Journal of Theoretical Biology 12:119–129.

———. 1967. Optimizing reproduction in a randomly varying environment when a correlation may exist between the conditions at the time a choice has to be made and the subsequent outcome. Journal of Theoretical Biology 16:1–14.

Cooper, W. S., and R. H. Kaplan. 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 J. A. Thomas. 1991. Elements of Information Theory. Wiley Series in Telecommunications. John Wiley & Sons, New York.

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.

Dempster, E. R. 1955. Maintenance of genetic heterogeneity. Cold Spring Harbor Symposia on Quantitative Biology 20:25–32.

Donaldson-Matasci, M. C., C. T. Bergstrom, and M. Lachmann. 2010. The fitness value of information. Oikos 119:219–230.

Donaldson-Matasci, M. C., M. Lachmann, and C. T. Bergstrom. 2008. Phenotypic diversity as an adaptation to environmental uncertainty. Evolutionary Ecology Research 10:493–515.

Evans, M. E. K., R. Ferrière, M. J. Kane, and D. L. Venable. 2007. Bet hedging via seed banking in desert evening primroses (Oenothera, Onagraceae): Demographic evidence from natural populations. American Naturalist 169:184–194.

Feller, W. 1966. An Introduction to Probability Theory and Its Applications, vol. II. Wiley, New York.

Frank, S. A., and M. Slatkin. 1990. Evolution in a variable environment. American Naturalist 136:244–260.

Freas, K. E., and P. R. Kemp. 1983. Some relationships between environmental reliability and seed dormancy in desert annual plants. Journal of Ecology 71:pp. 211–217.

Haccou, P., and Y. Iwasa. 1995. Optimal mixed strategies in stochastic environments. Theoretical Population Biology 47:212–243.

Halkett, F., R. Harrington, M. Hullé, P. Kindlmann, F. Menu, C. Rispe, and M. Plantegenest. 2004. Dynamics of production of sexual forms in aphids: Theoretical and experimental evidence for adaptive "coin-flipping" plasticity. American Naturalist 163:E112–E125.

Hewitt, E., and L. J. Savage. 1955. Symmetric measures on Cartesian products. Transactions of the American Mathematical Society 80:470–501.

Khatchikian, C. E., J. J. Dennehy, C. J. Vitek, and T. P. Livdahl. 2010. Environmental effects on bet hedging in aedes mosquito egg hatch. Evolutionary Ecology 24:1159–1169.

Koops, M. A., J. A. Hutchings, and B. K. Adams. 2003. Environmental predictability and the cost of imperfect information: influences on offspring size variability. Evolutionary Ecology Research 5:29–42.

Kussell, E., R. Kishony, N. Balaban, and S. Leibler. 2005. Bacterial persistence: A model of survival in changing environments. Genetics 169:1807–1814.

Lane, S. J., and M. J. Mahony. 2002. Larval anurans with synchronous and asynchronous development periods: Contrasting responses to water reduction and predator presence. Journal of Animal Ecology 71:780–792.

Levins, R. 1968. Evolution in changing environments. Princeton University Press, Princeton, New Jersey.

Lively, C. M. 1986a. Canalization versus developmental conversion in a spatially-variable environment. American Naturalist 128:561–572.

———. 1986b. Predator-induced shell dimorphism in the acorn barnacle *Chthamalus anisopoma*. Evolution 40:232–242.

Menu, F., and E. Desouhant. 2002. Bet-hedging for variability in life cycle duration: bigger and later-emerging chestnut weevils have increased probability of a prolonged diapause. Oecologia 132:167–174.

Moran, N. A. 1992. The evolutionary maintenance of alternative phenotypes. American Naturalist 139:971–989.

Morey, S. R., and D. N. Reznick. 2004. The relationship between habitat permanence and larval development in California spadefoot toads: Field and laboratory comparisons of developmental plasticity. Oikos 104:172–190.

Petru, M., and K. Tielboerger. 2008. Germination behaviour of annual plants under changing climatic conditions: separating local and regional environmental effects. Oecologia 155:717–728.

Philippi, T. 1993. Bet-hedging germination of desert annuals: Variation among populations and maternal effects in *Lepidium lasiocarpum*. American Naturalist 142:488–507.

Philippi, T., M. Simovich, E. Bauder, and J. Moorad. 2001. Habitat ephemerality and hatching fractions of a diapausing anostracan (Crustacea : Branchiopoda). Israel Journal of Zoology 47:387–395.

Richter-Boix, A., G. A. Llorente, and A. Montori. 2006. A comparative analysis of the adaptive developmental plasticity hypothesis in six mediterranean anuran species along a pond permanency gradient. Evolutionary Ecology Research 8:1139–1154.

Rivoire, O., and S. Leibler. 2010. The value of information for populations in varying environments. arXiv/1010.5092.

Robson, A. J. 1996. A biological basis for expected utility. Journal of Economic Theory 68:397–424.

Sadeh, A., H. Guterman, M. Gersani, and O. Ovadia. 2009. Plastic bet-hedging in an amphicarpic annual: an integrated strategy under variable conditions. Evolutionary Ecology 23:373–388.

Saiah, H., and N. Perrin. 1990. Autumnal vs. spring hatching in the fairy shrimp *Siphonophanes grubii* (Dybowski) (Crustacea, Anostraca): diversified bet-hedging strategy? Functional Ecology 4:769–775.

Seger, J., and H. J. Brockmann. 1987. What is bet-hedging? Pages 182–211 *in* P. Harvey and L. Partridge, eds. Oxford Surveys in Evolutionary Biology, vol. 4. Oxford University Press, Oxford.

Shannon, C. E. 1948. A mathematical theory of communication. The Bell System Technical Journal 27:379—423, 623—656.

Simons, A. M. 2009. Fluctuating natural selection accounts for the evolution of diversification bet hedging. Proceedings of the Royal Society B: Biological Sciences 276:1987–1992.

Simons, A. M., and M. O. Johnston. 1997. Developmental instability as a bet-hedging strategy. Oikos 80:401–406.

———. 2006. Environmental and genetic sources of diversification in the timing of seed germination: Implications for the evolution of bet hedging. Evolution 60:2280–2292.

Venable, D. L. 2007. Bet hedging in a guild of desert annuals. Ecology 88:1086–1090.

Yoshimura, J., and C. W. Clark. 1991. Individual adaptations in stochastic environments. Evolutionary Ecology 5:173–192.