Risk management in biological evolution

Andreas Wagner*

Department of Biology, University of New Mexico, 167A Castetter Hall, Albuquerque, NM 817131-1091, USA

Received 22 January 2003; received in revised form 7 May 2003; accepted 22 May 2003

Abstract

I present a framework to study the evolution of traits that allow an organism to survive life-threatening but rare risks. Specifically, I am concerned with risks so rare that any one individual in a population may not experience the risk-causing event in its lifetime. A theory of rare risk management is virtually absent in evolutionary biology, although it is well developed in economics. This is surprising because of the great influence economics had on evolutionary biology, and because biology is full of examples for evolved risk management traits. They include the ability of bacteria to sporulate, of pathogens to survive antibiotic treatment, of temperate bacteriophages to enter a lytic life cycle, as well as traits that allow higher organisms to survive rare environmental disasters, such as sporadic wildfires and irregular flooding. I make predictions about the sustenance of risk management traits under two scenarios, one where the catastrophic events cause individual deaths, and another one where catastrophic events cause population extinction. A well-developed theory of risk management will not only predict the distribution of risk management traits, but may also serve other purposes, such as to reconstruct the spectrum of environments that an organism encountered in its evolutionary history from the record stored in its genome's memory.
© 2003 Elsevier Ltd. All rights reserved.

Keywords: Economic risk; Sporulation; Metabolic diversity; Fire adaptations; Flooding adaptations

1. Introduction

Evolutionary biology has been greatly influenced by economic thinking. The history of this marriage begins with Malthus' insights on population growth and their influence on Darwins work (Mayr, 1982, p. 82). More recent examples include the concept that life is constrained by resources that must be allocated optimally, the notion that organisms play evolutionary games (Maynard-Smith, 1982), and that organisms compete in biological markets (Noe and Hammerstein, 1995; Frank, 1998, p. 26). One key economic concept, however, has received relatively little attention in evolutionary biology. It is the concept of risk. Risk, uncertainty as to loss of 'assets' or life, is ubiquitous in both economic activity (Greene et al., 1992) and organismal life. In human history risk management originated early, being first documented in the Code of Hammurabi (about 1800 BC) (Bernstein, 1995), which contains the first cases of marine insurance against the

risk of maritime disasters. In modern times, risk management permeates all areas of economic life. It is employed in a multibillion dollar insurance industry to protect property and life, and in sophisticated mathematical tools to protect financial assets against volatility (Markowitz, 1991; Beckers, 1998). Its importance extends also to many other areas outside economics, most prominent among them technology risk management, where a vast body of both popular and technical literature exists (Chiles, 2001; Perrow, 1999; Pritchard, 2001; Vose, 2000; Zhou et al., 2002).

I will give the briefest overview over some economic risk management strategies for which parallels in non-human life are easy to find. Economic risk management techniques fall into several categories (Greene et al., 1992). The first of them is risk avoidance, exemplified by individuals who avoid taking airplanes or give up practicing medicine in the face of risking malpractice lawsuits. Loss control may present a second option when a risk is difficult to avoid. Loss control takes at least three forms, frequency reduction, severity reduction, and diversification. Frequency reduction aims at reducing the incidence of loss-causing events. It is

*Tel.: +1-505-277-2021; fax: +1-505-277-0304.
E-mail address: wagnera@unm.edu (A. Wagner).

exemplified by a driver who avoids busy roads to avoid accidents. Severity reduction, the second form of loss control, is illustrated by a driver who wears seat belts. Diversification, finally, is especially important in business and finance. For example, a firm may use several warehouses to store goods to protect against the risk of fire. A third main strategy of risk management, after risk avoidance and loss control, is risk retention. It involves the assumption of risk. A company or individual that practices risk retention does not try to avoid catastrophic events but generates back-up resources. And finally, there is risk transfer, where an institution different from the risk carrier assumes the risk. Most forms of insurance fall under this definition. Needless to say, the distinction between these risk management strategies is not clear-cut. For instance, self-insurance can be viewed as a form of either insurance or risk retention. And a medical practicioner undergoing continuing training engages both in frequency reduction and severity reduction with respect to malpractice lawsuits.

Countless examples for any of these risk management strategies can be found in non-human organisms. Diversification in an organism's diet assures survival if one of the food species suffers extinction. The ability to migrate can be viewed as a strategy to avoid the risks of climatic extremes. Food storage is nearly ubiquituous from microbes to mammals. It can be viewed as a risk retention strategy. However, the parallels between economics and evolution are limited. One example is the area of risk assumption. The analog of insurance companies, institutions whose sole purpose is to assume the risk of others, may be completely absent in non-humans. The same holds for the important economic distinction between pure and speculative risks. A pure risk is a risk due to an event that may cause a loss (and only a loss). In a speculative risk, a gain is equally possible. While it is possible to construe some biological risks as speculative risks (one organisms death in a catastrophic event may free resources available for better risk managers), the very notion of speculation carries a connotation of planning or foresight, whereas risk management in non-human organisms reflects past survival. I will thus limit myself to pure risks, risks of losses. Economics also distinguishes between static—unchanging—risks and dynamic risks. Examples include losses due to lightning or storms on one hand, and the risk of extinction in a habitat with changing species composition. While virtually all risks may be dynamic risks, modeling risk is often only possible if the incidence of losses does not change over time.

Any life-sustaining activity is fraught with risks. Foraging animals are subject to the risks of predation, plants are subject to the risks of herbivore damage, cooperation in social organisms is subject to exploitation, ectotherms are subject to the risk of extreme temperature variation, all of life is subject to the risk of parasitism, and so on. Countless organismal adaptations can be understood as evolved strategies to manage risks. Most of these adaptations already have adequate evolutionary explanations. Specifically, a considerable body of work on 'bet-hedging' regards adaptations to frequent risks that many organisms encounter (Gillespie 1973, 1977; Hopeer, 1999; Real, 1980; Stearns, 1992). However, current evolutionary thinking does not extend to an important class of risks, risks that have two distinguishing features. First, these risks have severe impact on the survival of an organism or a population: They threaten death or extinction. Two examples follow. Lysogenic bacteriophage can integrate their genomes into that of a host bacterium (Ackermann and DuBow, 1987). They remain dormant and replicate their genome along with the host for many generations. When the host is exposed to severe stress, the proviral genome enters the lytic life stage. Its genome becomes expressed, producing many phage particles in a short amount of time, which leave the sinking ship. The lysogeny-lysis switch can be viewed as a phage adaptation to the risk of host stress. Another case in point, from higher organisms, are adaptations to wildfires, such as are found in vascular plants in the form of fire-stimulated germination or regeneration through stump-resprouting. Wildfires can be rare events whose frequency may be as low as one per millennium (Clark et al., 1989; Cody and Mooney, 1978, Ogden et al., 1998). They illustrate the second feature of the risks I will consider here, namely their rarity. Unlike risks such as that of predation, wildfires and life-threatening host stress are not faced on a day-by-day basis. Loosely speaking, I consider as rare those catastrophic events that may affect individuals only once per generation or even less. It is worth noting that some risks that are normally too frequent to fall in this category may do so under special circumstances. A case in point is a population's risk to become exterminated by an infecting pathogen. Persistence of a pathogen in a host population depends on the host population size. If host population sizes are large, pathogens persist and any one host is likely to become infected during its lifetime. If host organisms are few, or if their density is low (that is, if they are widely dispersed), endemic infections cannot be sustained and a pathogen can thus become extinct. Renewed infections then depend on the reintroduction of the pathogen into the population, either through migration or through reinfection from a pathogen reservoir in a different species. Evidence from island populations, for example, suggests that human diseases such as measles cannot persist in populations of less than a quarter to a half million people. In populations smaller than that, they periodically go extinct and may become reintroduced, years later, through travelers (Anderson and May, 1991, p. 82). Similarly, many human diseases contributing

greatly to mortality in historic times, such as smallpox and cholera, could not have persisted in the small human populations of a pre-agricultural era (Anderson and May, 1991, p. 654).

Even rare risks might have enormous influence on organismal evolution. A potential case in point is the gene content of microbial genomes. Large-scale gene–deletion studies reveal that free-living microbes contain thousands of genes that are dispensable for life in a variety of laboratory environments (Winzeler et al., 1999). Such apparently dispensable genes may be critical in rare or unusual environments. Thus, a large fraction of a genome may embody risk management strategies. The meaning of 'rare' on an absolute time-scale does of course depend on the organism. It ranges from weeks for microbes to millennia for higher organisms.

Some risks are so rare, and their impact so severe that evolving adaptations against them is inconceivable (think of the proverbial meteor strike.) Others are so frequent or may have effects so slight that existing evolutionary theory has no trouble explaining them (Gillespie, 1973, 1977; Hopeer, 1999; Real, 1980; Stearns, 1992). The enormous gray area in between is the subject of this paper. An evolutionary framework to explain risk management strategies must answer one main question. Under what conditions can organisms evolve them? More specifically, how frequent must a risk be, and how severe its impact for risk management strategies to evolve? I will here make a first step towards answering this question for organisms with the simplest possible life histories and genetics.

2. Model and results

To illustrate the principles at work most clearly, I will restrict myself to the case of asexually reproducing organisms or, equivalently, genotypes that differ at one locus or at several tightly linked loci, such that recombination can be neglected. Fitness is malthusian, reflecting population growth rate in continuous time (Crow and Kimura, 1970, Chapter 1). I will consider two genotypes, G_0 and G_r, which may differ in one or more genes. G_r is the superior risk manager, that is, it is less susceptible to a particular rare catastrophic event. I will consider two evolutionary scenarios that lie at the extremes of a continuous spectrum. In the first, a rare catastrophic event affects individuals in a population independently. That is, a catastrophy is synonymous with individual death. Examples will typically occur in populations of macroscopic organisms at low population densities, where death may be caused by rare but localized resource shortages, or by the periodic introduction of an externally introduced pathogen. In the second scenario, which may be more important for large populations of small organisms at high densities,

the rare catastrophic event affects an entire population and lead to the extinction of all individuals of a genotype.

The models I present here are very simple. Specifically, they do not incorporate factors such as polygeny, sexual reproduction, dispersal, and complex life histories, all of which may affect the evolution of risk management traits. My motivation for studying such simple models is twofold. First, they illustrate the evolutionary principles at work most clearly. Second, their results can serve as points of reference and departure for further studies.

2.1. Scenario 1: Risks that cause individual deaths

The absolute frequencies of the genotypes G_0 and G_r are denoted by N_0 and N_r, respectively, and the population size as $N = N_0 + N_r$. The relative frequency of G_r is given by $p_r = N_r/N$. The population growth rate r_r of the genotype G_r is defined as

$$r_r = b_r - (d + d_r). \tag{1}$$

Here, b_r is the per capita birth rate, $d + d_r$ is the per-capita death rate, which has two components, the rate of deaths d from causes unrelated to the catastrophic event at issue, and the rate of deaths d_r due to such catastrophic events. One can partition the growth rate of G_0 analogously, i.e., $r_0 = b_0 - (d + d_0)$, where b_0 is the birth rate and d_0 is the death rate due to the catastrophic event. G_r is a better risk manager if $d_r < d_0$. It is perhaps most natural to express this relation as the ratio d_r/d_0, such that if $d_r/d_0 = 0.1$, G_r is ten times less likely to perish than G_0. Such superior risk management will usually carry costs in the form of a lower population growth rate b_r, that is, $b_r \leqslant b_0$. This cost could be caused by slower DNA replication, or various energy expenditures, such as that of expressing gene products associated with the risk management trait. Such costs are plausible and have been shown to exist in individual case studies (Cooper et al., 2001; Lenski et al., 1994, Schrag et al., 1997).

Models with similarly simple genetics have been explored in different contexts in population genetics, and pertinent mathematical results can thus be found in variety of texts (Crow and Kimura, 1970). The key difference of this model to other population genetic models is a conceptual one, namely that organisms here do not differ in fitness, on which selection acts continuously. Rather, they differ in their propensity to succumb to rare catastrophic events.

It is easy to see that in large populations, the frequency p_r of the superior risk manager evolves according to the differential equation

$$\dot{p}_r = (r_r - r_0)p_r(1 - p_r). \tag{2}$$

That is, the organism with reduced risk exposure will come to dominate the population, as long as $\Delta r := r_r - r_0 > 0$ or, equivalently,

$$b_0 - b_r < d_0 - d_r. \tag{3}$$

In other words, the cost of risk management in terms of reduced births, $b_r - b_0$, must be no greater than the reduction in deaths due to better risk management. How large must a population be for this relation to guarantee the success of G_r? If the cost of risk management is negligible, then $(N_0 d_0 - N_r d_r) \gg 1$ will be sufficient. That is, the difference in the expected number of individuals killed by a catastrophic event must be greater than one, otherwise effects of genetic drift (see also below) render this deterministic approach invalid. To give a numerical example, in a population with 10^5 organisms of either genotype, if $d_0 = 10^{-3}$ and $d_r = 10^{-4}$, the genotype with lower risk exposure will increase in frequency. The difference in expected deaths per unit time between the two genotypes is 90. If units of time are years, and if every individual leaves, on average, one offspring (i.e., population sizes are stable), then the expected waiting time for a catastrophic event to kill a member of a family lineage is greater than 1000 years in this example. This illustrates that selection can be effective against very rare risks if populations are sufficiently large. (It is a corollary of the well-known fact that selection can be effective in large populations even if fitness differences are very small.) If the cost $b_0 - b_r$ becomes non-negligible, then the necessary population sizes need to increase concomitantly. $N_r r_r - N_0 r_0 \gg 1$ is a safe margin for the deterministic model in this case.

The model above does not incorporate variations in growth rates over time. It lumps differences between genotypes G_0 and G_r into a different overall (average) growth rate. However, one might take the perspective that the very concept of risk management must take into account variation in performance (growth rate) over time. For example, genotype G_0 may have a faster maximal growth rate under favorable conditions ($r_{0,max} > r_{r,max}$), but it may do poorly in other environments where G_r does better. Which of these genotypes will fare better in the long run? This is known as the problem of 'bet-hedging', which is well-studied and thus needs only brief elaboration (Gillespie, 1973, 1977; Stearns, 1992). Because the abundances of genotypes in a population are calculated by multiplying their respective growth rates over many generations, the arithmetic mean growth rate—proportional to the sum of growth rates—can be a poor indicator of risk management. For any two genotypes with the same arithmetic mean growth rate, the genotype with the lower variance will come to dominate the population. This genotype is also typically the genotype with the lowest geometric mean growth rate. I show briefly in the

appendix that this relation between risk management and variance in growth rates does not extend to the geometric mean growth rates of genotypes. That is, if two genotypes have the same geometric mean growth rate, it is not guaranteed that the genotype with the lower variance in growth rate will win. In addition, the genotype with the higher geometric mean growth rate always has a greater than 50 percent chance of dominating the population, regardless of how wildly its growth rates gyrate. Thus, one must apply caution in choosing the appropriate measure of fitness variation when studying risk management.

Thus far, I have neglected that genotypes with lower risk exposure can arise and disappear from a population. For example, mutations may continually eliminate the genotype G_r. To sustain G_r, the growth rate difference $\Delta r = r_r - r_0$ then needs to be greater than in the absence of mutations. I denote the rate at which G_r arises by μ and that at which it disappears as v. μ and v will be minute if point mutations are responsible for creation and disappearance of a genotype. However, either one may be of moderate size, for example if a bacterial genotype arises through frequent horizontal gene transfer (as in the case of bacterial antibiotic resistance) or if it can disappear through abundant intragenomic recombination processes leading to gene deletions. If only μ and v influenced the population frequency of the genotype G_r, its equilibrium frequency would calculate as $\hat{p}_r = \mu/(\mu + v)$. It is best to consider only the case where $\mu < v$, because otherwise mutation alone would sustain G_r at a high frequency without any catastrophic events.

In large populations, the frequency p_r evolves under the influence of mutation and selection through rare catastrophic events according to

$$\dot{p}_r = \Delta r p_r (1 - p_r) + \mu(1 - p_r) - v p_r, \tag{4}$$

where $\Delta r = r_r - r_0$. This equation admits a unique (stable) equilibrium for $p_r \in (0, 1)$, which is

$$\hat{p}_r = (1/2\Delta r)$$
$$\times [\Delta r - (\mu + v) + \sqrt{(\Delta r - (\mu + v))^2 + 4\mu\Delta r}]. \tag{5}$$

It is easy to see, by setting $\mu = 0$, that the equilibrium frequency $\hat{p}_r$ observes

$$\hat{p}_r \geq 1 - \frac{v}{\Delta r}. \tag{6}$$

Equality holds if $\mu = 0$, which is the worst-case scenario for the genotype G_r. It means that the genotype G_r is not being continuously reintroduced through mutations or import, and mutation alone would only lead to its disappearance. In this case, the ratio of the rate of disappearance of the risk management trait and its relative advantage jointly determine its equilibrium frequency. Equality continues to hold approximately, even if $\mu > 0$, as long as v is small compared to Δr.
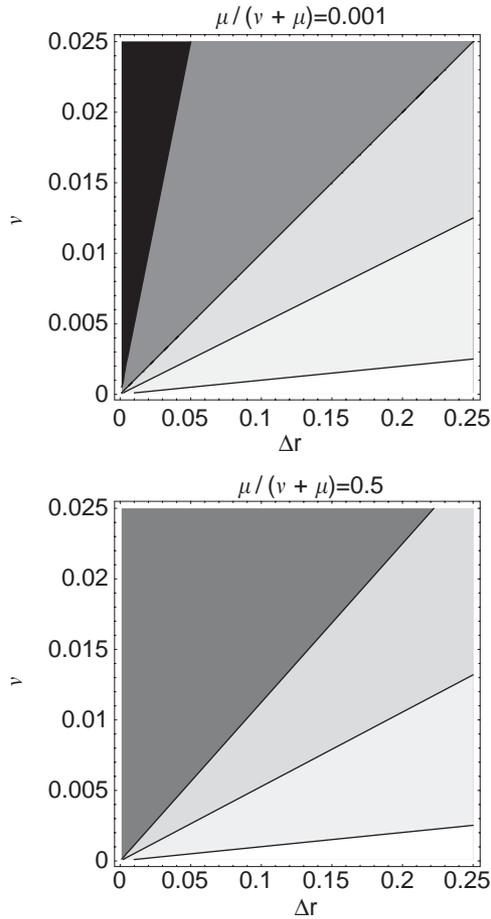
Fig. 1. Maintenance of risk management traits is not sensitive to the rate at which traits arise. Both panels show a contour plot indicating how the equilibrium frequency $\hat{p}_r$ (given by the solution of Eq. (5)) depends on the relative advantage Δr of the risk management trait and the rate v at which the trait spontaneously disappears from the population. Color codes are as follows. Black: $0 < \hat{p}_r < 0.5$; dark gray: $0.5 < \hat{p}_r < 0.9$; light gray: $0.9 < \hat{p}_r < 0.95$; white: $0.95 < \hat{p}_r < 0.99$. The rate μ at which the trait originates is chosen such that the mutation equilibrium $\mu/(v + \mu)$ equals $p_r = 0.001$ in (a) and $p_r = 0.5$ in (b).

Fig. 1a, based on Eq. (5) shows that whether the mutation-equilibrium frequency $v/(\mu + v)$ is equal to 0.001 or equal to 0.5, 500 times larger, it is still the ratio $v/\Delta r$ that dominates the equilibrium frequency in mutation-selection balance.

The above results hold in large populations. In small populations, genetic drift renders the evolution of G_r unpredictable. In this case, it is only possible to predict for a large (infinite) ensemble of populations, the probability $\phi(p)$ to find members of the ensemble where the frequency p_r of genotype G_r lies in the (infinitesimal) interval $(p, p + \mathrm{d}p)$. For the above parameters, this probability density is given by

$$\phi(p_r) = C p_r^{2N\mu-1}(1 - p_r)^{2Nv-1}\mathrm{e}^{2N(p_r-1)\Delta r}, \qquad (7)$$

where C is a normalization factor ensuring that $\int_0^1 \phi(p)\,\mathrm{d}p = 1$, and N is the (effective) population size (Crow and Kimura, 1970, 9.3). The question whether

risk management can evolve now has to be recast in probabilistic terms. For example, one could ask under what conditions G_r "wins", that is, when does the probability that more than a fraction ε of a population consists of G_r exceed some threshold δ. Formally,

$$\mathbf{P}(G_r \text{ wins}) = \int_\varepsilon^1 \phi(p_r)\,\mathrm{d}p_r > \delta. \qquad (8)$$

The behavior of Eq. (8) as a function of its parameters is well known to be complex, and no closed form solution exists (Crow and Kimura, 1970, Chapter 9). However, a few simple generalizations can be made. If $N\Delta r$, Nv (and thus also $N\mu$) are much greater than one, then the effects of genetic drift are negligible, and the above deterministic mutation-selection balance will be attained. If only $N\Delta r \gg 1$ then the frequency of G_r in equilibrium will be close to fixation ($1 - p_r \ll 1$), as long as $\Delta r > 0$. However, if $N\Delta r \ll 1$, then stable maintenance of G_r through natural selection is impossible, regardless of mutation rates. In other words, even though the probability that any one individual is affected by a catastrophic event may be exceedingly small, the risk-causing event must not be so rare that over a prolonged period of time no individuals at all in the population are affected.

2.2. Scenario 2. Risks that cause extinction of all genotypes in a population

In this scenario, it is impossible to predict the persistence of a risk reduction trait in one population. One can, however, predict whether the trait can dominate most individual populations in a large ensemble of populations. Consider again two genotypes G_0 and G_r. As opposed to scenario 1, where d indicated risks of individual deaths specific to a genotype, d will now indicate the risk of extinction of all individuals of a genotype in the population. The respective risks of extinction for genotypes G_0 and G_r can be modeled as a Poisson processes, where genotype G_0 is subject to extinction at an (infinitesimal) rate d_0 in the time interval $(t, t + \delta t)$. The expected waiting time for one such extinction event is $1/d_0$. Genotype G_r is subject to extinctions at the rate d_r, where $d_r < d_0$. (I will treat the cost of such risk management separately below.)

In this framework, the question whether risk management can be evolutionarily successful can be answered by considering the fraction of populations in an ensemble that contain G_r. If this fraction is greater than some arbitrary (but large) threshold ε, then risk management can be called evolutionarily successful. It is important to note that in this scenario, if one waits long enough ($t \to \infty$) all populations will go extinct. However, one needs only be concerned with the *fraction* of populations in which G_r persists.

I will first consider the simplest and idealized scenario where both genotypes initially occur in each population

of the ensemble, and where populations are so large that differences in reproductive rates cannot lead to the complete elimination of one genotype at any finite time t. I will treat two extreme cases that lie at the ends of continuous spectrum. In the first (easier) case, the probabilities of extinction in G_r and G_0 show a strong form of stochastic dependency, and in the second case they are stochastically independent.

Assume that some catastrophic events will cause all individuals of a population to go extinct, regardless of genotype. These events occur at the rate d_b. Other catastrophic events affect only G_0 but not G_r. They occur at a rate d_{00}. Genotype-specific extinction risks then are $d_r = d_b$ and $d_0 = d_b + d_{00}$, but the risk of extinction is highly correlated between the two genotypes, especially if d_{00} is small. Four different probabilities need to be distinguished. First among them is the probability that at time t none of the genotypes has suffered a catastrophic event. It is denoted by $\mathbf{P}_{0r}$, which is easily calculated from the waiting time for Poisson distributed (rare) events (Karlin and Taylor, 1975, Chapter 4). Second is the probability that G_0 but not G_r suffered extinction, denoted by $\mathbf{P}_r$. Third is the probability that G_r but not G_0 suffered extinction, denoted by $\mathbf{P}_0$. And fourth is the probability that both genotypes suffered extinction, denoted by $\mathbf{P}_{ex}$. It follows from the above assumption that $\mathbf{P}_0 = 0$, which makes it easy to calculate the probability for evolutionarily successful risk management, namely the fraction of (nonextinct) populations in which the genotype G_r occurs

$$\mathbf{P}(G_r \text{ wins}) = \frac{\mathbf{P}_r + \mathbf{P}_{0r}}{1 - \mathbf{P}_{ex}} = \frac{\mathbf{P}_r + \mathbf{P}_{0r}}{\mathbf{P}_r + \mathbf{P}_{0r} + \mathbf{P}_0} = 1, \qquad (9)$$

independent of t. The reason is that this scenario does not permit populations that consist only of genotype G_0. (Whenever G_r goes extinct, G_0 does as well.)

The second case of stochastically independent extinction risks d_0 and d_r is slightly more complicated. Again, four different cases need to be considered. First among them is the case where at time t none of the genotypes has suffered a catastrophic event. Under the Poisson assumption (Karlin and Taylor, 1975), its probability is $\mathbf{P}_{0r} = e^{-(d_0+d_r)t}$. Second, the probability that G_0 but not G_r suffered extinction is given by $\mathbf{P}_r = (1 - e^{-d_0 t})e^{-d_r t}$. Third, the probability that G_r but not G_0 suffered extinction is given by $\mathbf{P}_0 = (1 - e^{-d_r t})e^{-d_0 t}$. And finally, the probability that both genotypes suffered extinction is given by $\mathbf{P}_{ex} = (1 - e^{-d_r t})(1 - e^{-d_0 t})$. From these probabilities, the fraction of populations that contain G_r calculates as

$$\mathbf{P}(G_r \text{ wins}) = \frac{\mathbf{P}_r}{1 - \mathbf{P}_{ex}} = \frac{1}{1 + e^{(d_r-d_0)t} + e^{-d_0 t}}. \qquad (10)$$

It is easy to see that $\lim_{t \to \infty} \mathbf{P}(G_r \text{ wins}) = 1$ if and only if $d_r < d_0$. However, the formula also shows that at any finite t, time enters the relation only through its product with the extinction risk. That means that the process has a characteristic time scale determined by d_0 and d_r. Fig. 2 shows how the probability $\mathbf{P}(G_r \text{ wins}) = \delta$ depends on the ratio of extinction rates d_r/d_0 for different values of $d_0 t$. The bottom, middle, and upper curves correspond to $d_0 t = 5$, $d_0 t = 10$ and $d_0 t = 20$, respectively. The further time progresses (i.e., the larger $d_0 t$), the smaller a difference in extinction rates $d_0 - d_r$ is necessary to sustain G_r at a high frequency.

Everything said thus far assumes that both genotypes occur in each population of the ensemble. However, genotypes may both be introduced into a population and disappear from it for reasons unrelated to a catastrophic event. Disappearance may be due to three causes. First, a genotype may occur in a small number of individuals of a population, which makes it subject to elimination by mutation. Second, a genotype may be eliminated in a finite population due to genetic drift. And third, a genotype—especially one with superior risk management—may bear a cost in reproduction that will cause natural selection to eliminate it. Introduction of a genotype into a population is caused by processes similar to those in scenario 1 above, namely mutations and import, e.g., through horizontal gene transfer. I will denote the rates of introduction and disappearance of genotype G_r into the ensemble as μ_p and v_p, respectively. Notice that μ_p and v_p are different from the rates μ and v indicating the turnover rate of genotype G_r *within* a population. Specifically, v_p does not only include the effects of mutation alone, but also those of genetic drift and that of any fitness cost due to superior risk management. This also implies that neither μ_p nor v_p
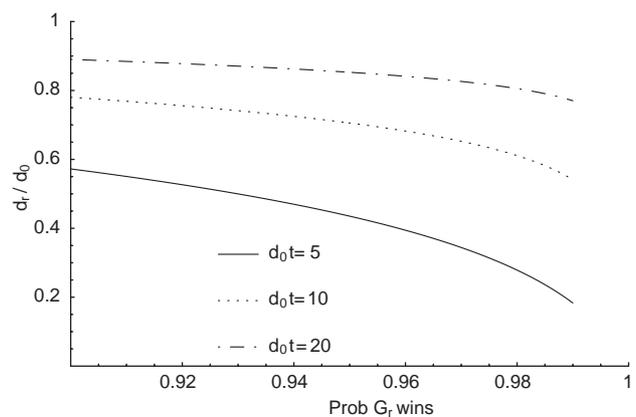

Fig. 2. In the long run, even a small reduction in extinction rates can effectively sustain superior risk management. The figure shows how much smaller the extinction rate d_r of genotype G_r has to be relative to d_0 of G_0 (expressed as the ratio d_0/d_r) such that a large fraction $\delta = \mathbf{P}(G_r \text{ wins})$ of a population ensemble contains only G_r. All three curves are numerical solutions of Eq. (10). The bottom, middle, and upper curves correspond to $d_0 t = 5$, $d_0 t = 10$ and $d_0 t = 20$, respectively. The further time progresses (i.e. the larger $d_0 t$), the smaller a difference in extinction rates $d_0 - d_r$ is necessary to sustain G_r at a high frequency.

are necessarily small. I also note that for the scenario of population extinction, the frequency of genotype G_r in a population is irrelevant, as long as the genotype occurs. (For example, if the extinction risk of a bacterial population is caused by exposure to an antibiotic, it does not matter how few resistant bacteria exist in a population.)

In the absence of catastrophic events, the fraction of populations that contain genotype G_r (at however small a frequency) will approach the equilibrium $\mu_p/(\mu_p + v_p)$. To calculate this equilibrium from within-population parameters such as μ and v, one can evaluate the fraction of populations of size N that contain G_r at a frequency greater than $1/N$ or $\int_{1/N}^{1} \phi(p)\,dp$, where $\phi(p_r)$ is distribution (7) of G_r in a population ensemble where $\Delta r = \Delta b$.

What happens to this equilibrium if catastrophic events occur? In the framework of Poisson processes, one can solve for $\mathbf{P}(G_r\text{wins})$ exactly, given μ_p and v_p, and the extinction rates d_0 and d_r. The solution, which I sketch in the appendix, is surprisingly complex. However, the key feature of the above, simpler case, is preserved, namely that

$$\mathbf{P}(G_r \text{ wins}) = \frac{\mathbf{P}_r + \mathbf{P}_{0r}}{1 - \mathbf{P}_{ex}} \to 1 \quad \text{for } t \to \infty \quad \text{if } d_r < d_0. \quad (11)$$

In other words, G_r wins regardless of the rates μ_p and v_p. The reason is that the different rates at which populations that contain only G_r or G_0 become extinct dominate the evolution of the population ensemble.

3. Discussion

I presented two simple mathematical models of risk management that stand for two extreme scenarios at the end of a continuous spectrum. In the first scenario, rare catastrophic events are responsible for deaths of individuals within a population. Here, natural selection can sustain genotypes that are responsible for superior risk management under two conditions. First, populations must be sufficiently large such that at least some members of a population are affected by a catastrophic event every generation. (However, the chances of any one individual being affected during its lifetime may be vanishingly small.) Second, the net benefit Δr of risk management must be greater than the rate v at which the genotype G_r conferring superior risk management disappears from the population. Specifically, G_r attains a population frequency greater than $1 - v/\Delta r$ (for $v < \Delta r$, and zero otherwise). The net benefit Δr of superior risk management is equal to the average difference in the death rate Δd due to the catastrophic events, discounted by the cost of risk management through a reduced birth rate Δb, i.e., $\Delta r = \Delta b - \Delta d$. These results are an outgrowth of standard population

genetic theory (Crow and Kimura, 1970), with one conceptual exception. In models of risk management, natural selection does not act continuously on genotypes whose fitnesses differ at all times. Instead, it acts only sporadically through rare catastrophic events.

In the second extreme scenario, rare catastrophic events cause extinctions of all individuals in a population, or of all individuals of one genotype. I showed that in the long-term limit, almost all populations in a population ensemble (a metapopulation) will contain genotypes with small extinction probability, regardless of the cost of such superior risk management, and regardless of the rate at which the genotype disappears from a population due to mutation or genetic drift. Note that the criteria for successful risk management are by necessity different in the two scenarios. In the first scenario, one requires most individuals in a population to be of genotype G_r. In the second scenario, one only requires that most populations contain at least one individual with genotype G_r, which will ensure the population's survival.

In practice, most catastrophic events may neither affect only one individual, nor all individuals in a population, but some fraction of individuals. In addition, although the models above can be applied to sexually reproducing organisms under certain conditions, they are designed with asexual populations in mind. In such populations, where the cohesion provided by recombination is lacking, it is often a matter of taste whether to call a group of individuals (such as a colony of bacteria) part of one population or one member population of a population ensemble. Taken together, these two observations suggest that the actual frequency of a risk management genotype G_r, whether determined from individuals in a population or from populations in a metapopulation will be an intermediate between the two extreme scenarios. Specifically, the lower bound $1 - v/\Delta r$ for G_r's equilibrium frequency still holds. The difference between the two scenarios lies in how this frequency is interpreted. In scenario one, it corresponds to the frequency of genotype G_r within a population. In scenario two, it corresponds to a frequency of (surviving) populations containing genotype G_r.

I focus here on largely qualitative predictions, because estimating population genetic variables important in quantitative predictions is notoriously difficult. However, I note that the above models contain all necessary ingredients to make quantitative predictions if the relevant factors, such as population sizes are known. Similarly, I am acutely aware that the biological assumptions I make are very simple. Especially in complex risk management traits of higher organisms (see also below) polygeny will be the norm. Such organisms also reproduce sexually, show varying degrees of ploidy, disperse to varying distances, and have complex life histories. Innumerable questions can be

asked about how each of these factors influence the evolvability of risk management. They present a wide open field for further investigation. In the following, I explore a number of potential empirical examples for risk management.

3.1. Extreme metabolic diversity as risk management

Many free-living microbes thrive on a bewildering array of substrates which serve as sources of carbon, nitrogen, phosphorus, and as electron acceptors (Neidhardt, 1996). Most of the time, the chemically complex environments of the wild will contain several substrates that can provide one and the same essential building block. Availability of only one substrate (a "minimal" medium) may be a potentially rare catastrophic event for a microbe that has lost the ability to thrive on this substrate. Put differently, the ability to grow on exotic substrates can be viewed as a risk management strategy for rare times where one such substrate becomes the only source of an essential cellular building block.

The results of several studies on microbial evolution in the laboratory (Cooper et al., 2001; Funchain et al., 2000; Lenski et al., 1991), although not designed for this purpose, help answer a specific question in this area. How fast do risk management capabilities get lost in environments where they are no longer needed? The question regards the reproductive cost of risk management traits, as well as the rates at which mutations cause their loss. Small or negligible costs of such traits seem to be the rule rather than the exception. In a few cases, individual studies estimated the cost of particular metabolic abilities. A case in point is the ability to use the five-carbon sugar arabinose as a carbon source, where Lenski and collaborators (Lenski et al., 1991) determined in competition experiments that its loss provided no detectable fitness advantage. The results of larger-scale studies involving many cell lineages and dozens of metabolic capabilities are also available. For example, Funchain and collaborators (Funchain et al., 2000) propagated 100 populations of the bacterium *Escherichia coli* on a rich medium for thousands of generations. They found that in populations with normal mutation rates, only some three percent of cell lineages had accumulated any detectable mutations impairing a metabolic function after 1000 generations of evolution (Funchain et al., 2000). The study monitored the activity of some 700 genes through multiple indicators of metabolic function, such as the ability to grow on a broad spectrum of carbon sources, and the capacity to synthesize amino acids. Populations were periodically subjected to severe bottlenecks, which favor the accumulation of mutations more than continuous growth. The low loss rate of metabolic capabilities is thus all the more remarkable.

However, there are exceptions to this rule, that is, there are metabolic abilities that carry substantial cost. A well-studied case is that of *E. coli*'s ability to metabolize D-Ribose, which carries a growth rate cost of 1–2 percent (Cooper et al., 2001). Following Cooper and collaborators (Cooper et al., 2001), it is possible to calculate the number of cell generations necessary for this trait to go from fixation to complete loss in a population growing in rich medium. I am here using the assumption of a population with $N = 3 \times 10^7$ individuals, comparable to that of the experiment (Cooper et al., 2001), and that of a mutation rate taking into account the total length (7000 bp) of the operon encoding this trait, $v = 3.5 \times 10^{-6}$ (Cooper et al., 2001). The time until the frequency of individuals capable to catabolize D-Ribose falls from one to $(1/N)$ calculates as 2792 generations. In other words, even traits with substantial reproductive costs can remain in a population for thousands of generations. If a limiting environment, in which the ability to catabolize D-Ribose becomes essential, occurs only every 2500 generations, this trait can be sustained indefinitely in a population.

One can also ask what percentage of individuals must be affected by a catastrophic event such that a risk management trait with an appreciable reproductive cost of 1 percent can be sustained at a high frequency such as $\hat{p}_r = 0.9$. The answer follows from solving $\hat{p}_r > 1 - v/(\Delta b - \Delta d)$ for Δd, i.e.,

$$\Delta d > \Delta b - \frac{v}{1 - \hat{p}_r} \tag{12}$$

(note that $\Delta d, \Delta b < 0$). Because v is much smaller than Δb, a Δd only slightly larger than Δb will suffice to sustain a risk management trait. For instance, if only as few as 1.1 percent of individuals in a population experience a life-threatening shortage of a particular chemical substrate every generation, then the respective metabolic risk management trait can be sustained at $\hat{p}_r > 0.9$.

Calculations based on such laboratory experiments have to be taken with a grain of salt, and not only because many metabolic capabilities carry much smaller cost than in the above example. Competition in laboratory studies is unusually intense over thousands of generations (Lenski et al., 1991), and it occurs in unchanging environments, unlike in the wild. Less intense competition in fluctuating environments may allow traits to be sustained at higher frequencies and for longer times as indicated by such calculations. Thus, an organism's genome contains a record of rare (and not so rare) environmental conditions that goes back thousands of generations. This "genome memory" may serve to reconstruct the breadth of environmental conditions that molded an organism on evolutionary time scales.

3.2. Antibiotic resistance and 'pure' risk management

From a bacterium's perspective, antibiotic resistance is a risk management trait. Some antibiotic resistance traits get easily lost from a population spontaneously, partly because they carry a substantial cost in the form of a reduced bacterial growth rate (Schrag et al., 1997; Seppala et al., 1997; Rahal et al., 1998). However, this does not hold for all antibiotic resistance traits. Schrag and collaborators studied the evolution of *E. coli* strains that harbor the streptomycin-resistance gene *rpsL*, which confers a serious growth rate deficiency on its host (Schrag et al., 1997). In some laboratory cultures, this resistance gene has not been lost after more than 10,000 generations in an antibiotic-free environment. The reason is that the bacteria harboring the gene have acquired a compensatory mutation which neutralizes some of the fitness disadvantages of the antibiotic resistance (Schrag et al., 1997). In addition, in these mutant bacterial strains, loss of *rpsL* now carries a serious disadvantage (Schrag et al., 1997). This is a case where a gene that has served one purpose originally may also have come to serve another (unknown) purpose in the absence of the antibiotic. In other, similar cases an antibiotic resistance trait confers a known benefit in the absence of the antibiotic (Baquero and Blazquez, 1997; Lenski et al., 1994; Blot et al., 1994).

Such secondarily acquired gene functions may effectively increase the memory of a genome, and retain genes much beyond their original purpose. They also raise an important questions. If most genes serve multiple purposes, then 'pure' risk management traits may not exist. Two lines of evidence speak against this possibility. First, organisms that experience very stable (risk-free) environments, such as endosymbionts or endoparasites, eventually lose a vast majority of their genes (Mira et al., 2001; Ochman and Moran, 2001). Second, under laboratory conditions thousands of genes—about one-third of the genome—in a microbe such as the yeast *Saccharomyces cerevisiae* appear completely dispensable, with no detectable fitness effect of their removal (Winzeler et al., 1999). These observation argue against the tight integration of most genes into all aspects of organismal life.

3.3. Risk management and dysfunctional DNA

Temperate bacteriophages can perpetuate their DNA indefinitely as part of the host genome. Far from damaging the host, this may even provide advantages to the host, such as immunity from superinfection and metabolic capabilities (Barksdale and Arden, 1974). The lysogenic state is extraordinarily stable, having been perpetuated for several decades in some bacterial strains kept in continuous laboratory culture (Barksdale and Arden, 1974). When the host is exposed to a variety of potentially life-threatening stressors, most notably among them ultraviolet radiation, prophages can enter a lytic state, where they produce many phage particles, which leads to cell death through lysis. (Some bacteriophages, where phage particles leave the cell through extrusion or budding, do not kill the host cell and may simply reduce the host's growth rate (Ackermann and DuBow, 1987, p. 66).) From a gene-centered perspective, that is, from the perspective of the prophage or proviral DNA, the ability to enter a lytic state upon host stress is a risk management strategy. At high phage concentrations, the lytic mode of reproduction can be extraordinarily efficient (and thus potentially self-defeating), killing most hosts in a population (Ackermann and DuBow, 1987, p. 52). This, together with the prevalence of lysogeny, suggests that lysogeny is an efficient reproduction mode under normal conditions.

Envision a population in which the vast majority of bacteria carry a prophage (and are thus immune to superinfection). Under normal conditions, phage production in this population is unfavorable, partly because a large fraction of the population is immune to infection, partly because producing phage particles adversely affects the healthy host. Spontaneous mutations at a rate v in the prophage DNA will lead some prophages to lose this risk management trait, i.e., they will become defective in their transition to the lytic stage. If individual cells in this population are frequently exposed to life-threatening stress, then prophages that have lost the ability to produce phage particles are at a comparative disadvantage. They are preferentially eliminated from the population through their host's death, because they cannot produce phage particles. However, if stressors are extremely rare, this disadvantage disappears. One would then expect that the frequency of such defective prophages increases in the population, eventually reaching a stable equilibrium greater than $1 - v/\Delta r$, as prescribed by Eq. (6). Here, Δr is a measure of the incidence of the stressor per unit time. The frequency of defective prophages will approach one when stressors occur more rarely than the rate v at which the ability to enter the lytic stage disappears. In other words, in populations extremely infrequently exposed to stress, individual genomes should be littered with defective prophages. Also, the kinds of mutations one observes in such defective prophages should not be scattered randomly in the prophage genome, but in early genes of the lytic cycle. The reason is that defective late genes of the lytic life cycle—when expressed—incur metabolic costs to the host, which may not get lysed as a result of the defect, but will grow more slowly.

How rare must life-threatening stresses be such that a high frequency of active prophages can be maintained in a bacterium like *E. coli*, with an estimated 100–300 (Savageau, 1983; Guttman and Dykhuizen, 1994)

generations per year in the wild? The mutation rate per base pair in *E. coli* has been estimated as 5×10^{-10} (Drake, 1991; Drake et al., 1998). The genome of a bacteriophage like that of phage λ comprises some 48,000 base pairs (Russell, 1998), of which about one-fifth encodes early genes of the lytic cycle, leading to a mutation rate of 5×10^{-6} in the relevant part of the genome. At least one-third of these mutations are silent, and many of the remaining amino acid substitutions will be neutral. Assuming a mutation rate of $v = 10^{-6}$ mutations towards defective prophages, $\Delta r = 10^{-5}$, or a rare event every 300–1000 years is sufficient to maintain active prophages at a frequency of 90 percent. In other words, the evolutionary record of life-threatening stresses stored in the genome of even a fast-dividing microbe may well span a millenium. The genome's memory is even longer if all individuals of the population are exposed to the stressor simultaneously, i.e., in scenario 2 from above. In this case, unless all functional prophages have disappeared from the population (which is extremely unlikely given the enormous stability and negligible cost of lysogeny) functional prophage frequency can remain high almost indefinitely.

I have focussed here on examples of risk management traits whose genetics is simple and well understood. However, many risk management traits are complex. Examples from prokaryotes include the formation of durable endospores in response to adverse environmental conditions (Sonenshein et al., 2002; Msadek, 1999; Stragier and Losick, 1996, Chapter 26). For instance, in the soil bacterium *Bacillus subtilis* more than 100 genes are necessary for successful sporulation (Stragier and Losick, 1996). Examples from higher organisms include plant adaptations to irregular flooding or to wildfires (Blom, 1999; Clark et al., 1989; Cody and Mooney, 1978; Ogden et al., 1998). For plants with short generation times, even fairly frequent environmental disasters may be sufficiently rare to meet the conditions on risk I posed here.

4. Conclusions

What purposes would be served by a well-developed evolutionary theory of risk management—of which I present only a first step? I see three main purposes. First, and most generally, such a theory would allow us to make specific predictions about the conditions, such as environmental variation, under which risk management traits like sporulation, the lytic switch, or fire adaptations can evolve. Second, such a theory may allow estimation of biological parameters important for successful risk management. A case in point is the cost of a risk management trait. Especially in microbes, where laboratory studies of evolution have been very successful, it is easy to subject populations to rare but life-threatening events. Other important parameters, such as population sizes and mutation rates, are either known or can be controlled in such populations. Through simple relations like those demonstrated here between genotype frequency, management cost, event frequency, and mutation rates, one can estimate the cost of the trait from other, measurable parameters. Third, to fully understand an organism's biology, it is necessary to understand the spectrum of environmental conditions under which the organism has evolved. Estimating the breadth of this spectrum is very difficult in practice. However, the genome contains a record even of rare environmental conditions, a record that, as I showed here, may extend over thousands of generations. A theory of risk management permits estimates of the time horizon of this genome memory. With an increasing number of fully sequenced genomes, ever improving genome annotations, and a framework to study risk management, we may thus be able to reconstruct the breadth of environmental fluctuations in an organism's past from its genome's record.

Appendix A

A.1. Reducing variance in fitness need not be an effective risk management strategy

I consider a discrete generation model, where a population consists of an equal number of individuals of two genotypes, G_0 and G_r. Their growth rates r vary over time, such that growth rates in consecutive generations are independent (or only weakly correlated) and drawn from otherwise arbitrary distributions with means μ_0 and μ_r, and variances σ_0^2 and σ_r^2 ($0 < r_i < \infty$, $0 < \sigma_i^2 < \infty$). For genotype G_0, the number of individuals at time t calculates as $N_0(t) = N_0(0) \prod_{i=1}^{t} r_0(t)$. (If $r_0(k) < 1$, then the population is shrinking in generation k.) Assuming, without loss of generality, that $N_0(0) = 1$, this is equivalent to

$$\log N_0(t) = \sum_{i=1}^{t} \log r_0(i). \tag{A.1}$$

With t large, $\log N_0(t)$ approaches a normal distribution with mean $\mu_{0t} := t \overline{\log r_0}$ and standard deviation $\sigma_{0t} := \sqrt{t}\sigma(\log r_0)$, where $\overline{\log r_0}$ and $\sigma(\log r_0)$ denote means and standard deviations of the distribution of $\log r_0$. I note

parenthetically that the mean logarithmic growth rate $\overline{\log r_0}$ is proportional to the geometric mean growth rate $(\prod_i r_0(i))^{1/t}$, because both logarithm and power functions are strictly monotonic functions. In sum, we have for t sufficiently large that

$$\log N_0(t) \sim N(\mu_{0t}, \sigma_{0t}). \tag{A.2}$$

With completely analogous notation, the logarithm of the population size $N_r(t)$ of genotype G_r is distributed as $N(\mu_{rt}, \sigma_{rt})$ for large t. The risk management strategy of a genotype G_r is effective if G_r comes to dominate a population. Cast in the terms of this model, the question is whether the probability that a large fraction of the population consists of G_r is greater than a prespecified threshold δ. Formally, this can be put as

$$\mathbf{P}\left(\frac{N_0(t)}{N_r(t)} < \varepsilon\right) > \delta. \tag{A.3}$$

For instance, $\varepsilon = 0.01$ and $\delta = 0.99$ are possible choices for these values. Because $(\log N_0(t) - \log N_r(t))$ is normally distributed as $N(\mu_{0t} - \mu_{rt}, \sqrt{\sigma_{0t}^2 + \sigma_{rt}^2})$, as long as growth rates $r_r(t)$ and $r_0(t)$ are no more than weakly correlated, it follows that

$$\mathbf{P}\left(\frac{N_0(t)}{N_r(t)} < \varepsilon\right) = \mathbf{P}(\log N_0(t) - \log N_r(t) < \log \varepsilon)$$
$$\propto \int_{-\infty}^{\log \varepsilon} \exp\frac{-(x - \mu_{0t} + \mu_{rt})^2}{\sigma_{0t}^2 + \sigma_{rt}^2}. \tag{A.4}$$

This expression shows that even if we know little about the actual distributions of growth rates of two genotypes r_i, questions about effective risk management can often be reduced to questions about random variates with normal distributions. I can now ask whether a genotype G_r with more stable growth rate ($\sigma_{rt}^2 < \sigma_{0t}^2$), perhaps at the price of a smaller mean growth rate ($\mu_{rt} \leqslant \mu_{0t}$), is an effective risk manager. That is, will G_r come to dominate a population? The following expression shows that this is not the case:

$$\mathbf{P}\left(\frac{N_0(t)}{N_r(t)} < \varepsilon\right)$$
$$= \frac{1}{\sqrt{2\pi(\sigma_{0t}^2 + \sigma_{rt}^2)}} \int_{-\infty}^{\log \varepsilon} \exp\frac{-(x - \mu_{0t} + \mu_{rt})^2}{\sigma_{0t}^2 + \sigma_{rt}^2}$$
$$\leqslant \frac{1}{\sqrt{2\pi(\sigma_{0t}^2 + \sigma_{rt}^2)}} \int_{-\infty}^{\log \varepsilon} \exp\frac{-x^2}{\sigma_{0t}^2 + \sigma_{rt}^2}$$
$$< 0.5. \tag{A.5}$$

The last inequality holds because $\log \varepsilon < 0$, and because the last distribution is symmetric around zero. Thus, even if G_r does not suffer a reduction in growth rate ($\mu_{rt} = \mu_{0t}$), the probability that it comes to dominate the population is less than 0.5. (It is even smaller if $\mu_{rt} < \mu_{0t}$.) In other words, reduction of variance in performance

need not be an evolutionarily effective risk management strategy.

This can also be seen in a different manner. Assume that the growth rate distributions of G_r and G_0 overlap little. That is, the mean growth rate of G_r is smaller than that of G_r ($\mu_{rt} < \mu_{0t}$), and the variances in growth rates are also small, e.g., $\mu_{rt} + 3\sigma_{rt} < \mu_{0t} - 3\sigma_{0t}$. In this case, the chances that G_r comes to dominate the population will be vanishingly small, because G_0 will grow faster almost every single generation. Now let σ_{0t}^2 grow, i.e. $\sigma_{0t}^2 \to \infty$. In this case, Eq. (A.5) implies that the probability that G_r wins approaches 0.5 from below, reaching this value only when the variance σ_{0t} reaches infinity. In sum, the genotype with the higher (logarithmic) mean growth rate always has at least a 50 percent chance of dominating the population, regardless of how wildly its growth rate gyrates.

A.2. Extinction risk in population ensembles

Denote as $\mathbf{P}_{0r}(t)$ the probability that both genotypes G_0 and G_r exist in a population at time t, as $\mathbf{P}_r(t)$ the probability that only G_r but not G_0 exists, as $\mathbf{P}_0(t)$ that only G_0 exist, and as $\mathbf{P}_{ex}$ the probability that both genotypes, that is, the entire population has become extinct. One can then write down a system of (stochastic) linear ordinary differential equations which describe how the vector $\vec{\mathbf{P}} = (\mathbf{P}_{ex}, \mathbf{P}_0, \mathbf{P}_r, \mathbf{P}_{0r})$ changes under the influence of rare catastrophic events that occur at genotype-specific rates d_0 and d_r, as well as under the rates μ_p and v_p at which G_r gets introduced and disappears, respectively, into populations of the ensemble:

$$\dot{\vec{\mathbf{P}}} = A\vec{\mathbf{P}}$$
$$= \begin{pmatrix} 0 & d_0 & d_r & 0 \\ 0 & -(d_0 + \mu_p) & 0 & v_p + d_r \\ 0 & 0 & -d_r & d_0 \\ 0 & \mu_p & 0 & -(d_r + d_0 + v_p) \end{pmatrix}$$
$$\times \begin{pmatrix} \mathbf{P}_{ex} \\ \mathbf{P}_0 \\ \mathbf{P}_r \\ \mathbf{P}_{0r} \end{pmatrix}. \tag{A.6}$$

It is easily verified that the temporal derivative of $(\mathbf{P}_{ex} + \mathbf{P}_0 + \mathbf{P}_r + \mathbf{P}_{0r})$ is equal to zero, such that the sum of the individual probabilities always remains one, as is required. The solution of Eq. (A.6) is

$$\vec{\mathbf{P}}(t) = \sum_{i=1}^{4} c_i v_i e^{\lambda_i t}, \tag{A.7}$$

where v_i is an eigenvector of the matrix A to the eigenvalue λ_i. The c_i's are constants depending on the initial condition $\vec{\mathbf{P}}(0)$, and can be obtained by solving

the linear system of equations $V\vec{c} = \vec{\mathbf{P}}(0)$, where V is a matrix whose column vectors consist of the eigenvectors v_i of A (Boyce and DiPrima, 1997). The eigenvalues of A are $\vec{\lambda} = (0, -d_0, -d_r, -(d_0 + d_r + \mu_p + v_p))$. I will denote as v_{ij} the j-th component of the i-th eigenvector of A. From Eq. (A.7), it follows that

$$\mathbf{P}(G_r \text{ wins}) = \frac{\mathbf{P}_r(t) + \mathbf{P}_{0r}(t)}{1 - \mathbf{P}_{ex}(t)}$$
$$= \frac{\sum_{i=1}^{4} c_i v_{i3} e^{\lambda_i t} + \sum_{i=1}^{4} c_i v_{i4} e^{\lambda_i t}}{1 - \sum_{i=1}^{4} c_i v_{i1} e^{\lambda_i t}}. \quad (A.8)$$

For a wide variety of initial conditions, including those where $\mathbf{P}_{0r}(0) = 1$, $\mathbf{P}_0 = 1$, or $\mathbf{P}_r = 1$, the constant $c_1 = 1$. This fact, in conjunction with $\lambda_1 = 0$ and $v_1 = (1, 0, 0, 0)$, yields the expression

$$\mathbf{P}(G_r \text{ wins}) = \frac{\mathbf{P}_r(t) + \mathbf{P}_{0r}(t)}{1 - \mathbf{P}_{ex}(t)}$$
$$= \frac{\sum_{i=2}^{4} c_i v_{i3} e^{\lambda_i t} + \sum_{i=2}^{4} c_i v_{i4} e^{\lambda_i t}}{-(\sum_{i=2}^{4} c_i v_{i1} e^{\lambda_i t})}. \quad (A.9)$$

(The equation describing $\mathbf{P}(G_r \text{ wins})$ in Eq. (10) of the main text derives as a special case if $\mu_p = 0$ and $v_p = 0$ in A.) Multiplying both numerator and denominator by $e^{\lambda_3 t}$, it is easy to see that

$$\lim_{t \to \infty} \mathbf{P}(G_r \text{ wins}) = \frac{v_{33} + v_{34}}{-v_{31}} = 1, \quad (A.10)$$

because $v_3 = (-1, 0, 1, 0)$. A more stringent definition of successful risk management, $\mathbf{P}(G_r \text{ wins}) = \mathbf{P}_r/(1 - \mathbf{P}_{ex})$, yields the same limit of one.

These results may seem surprising, especially in light of the fact that in the absence of extinction ($d_0 = d_r = 0$), $\mathbf{P}_0$ and $\mathbf{P}_{0r}$ attain a stable equilibrium prescribed by μ_p and v_p. The following qualitative perspective on the underlying dynamics shows that the reason for the eventual dominance of G_r is that the extinction dynamics overrides all other influences. Inspection of Eq. (A.6) shows that the temporal derivative of $\mathbf{P}_0 + \mathbf{P}_{0r}$ equals $\dot{\mathbf{P}}_0 + \dot{\mathbf{P}}_{0r} = -d_0(\mathbf{P}_0 + \mathbf{P}_{0r})$, and thus $[\mathbf{P}_0 + \mathbf{P}_{0r}](t) = \exp(-d_0 t)[\mathbf{P}_0 + \mathbf{P}_{0r}](0) =: \exp(-d_0 t)h$, independent of μ_p and v_p. Moreover, the third equation in Eq. (A.6) which describes the change in $\mathbf{P}_r$ shows that the solution to $\mathbf{P}_r(t)$ can be written as $\mathbf{P}_r(t) = \exp(-d_r t)\mathbf{P}_r(0) + f(t) > \exp(-d_r t)\mathbf{P}_r(0)$, where $f(t)$ is some positive function that depends on time t and $\mathbf{P}_{0r}(t)$. From these considerations it follows that

$$\frac{1 - \mathbf{P}_{ex}}{\mathbf{P}_0 + \mathbf{P}_{0r}} = \frac{\mathbf{P}_r + \mathbf{P}_{0r} + \mathbf{P}_0}{\mathbf{P}_0 + \mathbf{P}_{0r}}$$
$$= \frac{e^{-d_0 t}h + e^{-d_r t}\mathbf{P}_r(0) + f(t)}{e^{-d_0 t}h}$$
$$> \frac{e^{-d_0 t}h + e^{-d_r t}}{e^{-d_0 t}h} \to \infty \quad \text{for } t \to \infty. \quad (A.11)$$

It follows that the inverse of this ratio converges to zero, i.e.

$$\frac{\mathbf{P}_0 + \mathbf{P}_{0r}}{1 - \mathbf{P}_{ex}} \to 0 \quad \text{for } t \to \infty \quad (A.12)$$

and thus that

$$\frac{\mathbf{P}_r}{1 - \mathbf{P}_{ex}} = 1 - \frac{\mathbf{P}_0 + \mathbf{P}_{0r}}{1 - \mathbf{P}_{ex}} \to 1 \quad \text{for } t \to \infty. \quad (A.13)$$

Finally, it is worth noting that the model I discuss here allows for introduction of genotype G_r into populations consisting only of G_0, but not for introduction of genotype G_0 in populations consisting only of G_r. It is easy to incorporate this additional component into the matrix A, but the eigenvalues and eigenvectors become horrendously complex. However, the qualitative conclusions of Eq. (A.10) remain unchanged, because probability mass is only shifted between population containing G_r, such that $\dot{\mathbf{P}}_{0r} + \dot{\mathbf{P}}_r$ and thus Eq. (A.10) remains unchanged.

References

Ackermann, H.-W., DuBow, M., 1987. Viruses of Prokaryotes. General Properties of Bacteriophages, Vol. 1, CRC Press, Boca Raton, FL.

Anderson, R., May, R., 1991. Infectious Diseases of Humans: Dynamics and Control. Oxford University Press, Oxford, UK.

Baquero, F., Blazquez, J., 1997. Evolution of antibiotic resistance. Trends Ecol. Evol. 12, 482–487.

Barksdale, L., Arden, S., 1974. Persisting bacteriophage infections, lysogeny, and phage conversions. Ann. Rev. Microbiol. 28, 265–299.

Beckers, S., 1998. A survey of risk management theory and practice. In: Alexander C. (Ed.), Risk Management and Analysis, Vol. 1. Wiley, New York, pp. 39–60.

Bernstein, P.L., 1995. Against the Gods. The Remarkable Story of Risk. Wiley, New York.

Blom, C., 1999. Adaptations to flooding stress: from plant community to molecule. Plant Biol. 1, 261–273.

Blot, M., Hauer, B., Monnet, G., 1994. The tn5 bleomycin resistance gene confers improved survival and growth advantage on *Escherichia coli*. Mol. Gen. Genet. 242, 595–601.

Boyce, W., DiPrima, R., 1997. Elementary Differential Equations and Boundary Value Problems, 6th Edition. Wiley, New York, NY.

Chiles, J., 2001. Inviting Disaster: Lessons from the Edge of Technology. Harper Business, New York, NY.

Clark, J.S., Merkt, J., Muller, H., 1989. Post-glacial fire, vegetation, and human history on the northern alpine forelands, south-western Germany. J. Ecol. 77, 897.

Cody, M., Mooney, H., 1978. Convergence versus nonconvergence in mediterranean-climate ecosystems. Ann. Rev. Ecol. Systematics 9, 265–321.

Cooper, V.S., Schneider, D., Blot, M., Lenski, R.E., 2001. Mechanisms causing rapid and parallel losses of ribose catabolism in evolving populations of *Escherichia coli b*. J. Bacteriol. 183, 2834–2841.

Crow, J., Kimura, M., 1970. An Introduction to Population Genetics Theory. Harper and Row, New York.

Drake, J.W., 1991. Spontaneous mutation. Ann. Rev. Genet. 25, 125–146.

Drake, J.W., Charlesworth, B., Charlesworth, D., Crow, J.F., 1998. Rates of spontaneous mutation. Genetics 148, 1667–1686.

Frank, S., 1998. Foundations of Social Evolution. Princeton University Press, New York, NJ.

Funchain, P., Yeung, A., Stewart, J.L., Lin, R., Slupska, M.M., Miller, J.H., 2000. The consequences of growth of a mutator strain of *Escherichia coli* as measured by loss of function among multiple gene targets and loss of fitness. Genetics 154, 959–970.

Gillespie, J., 1973. Natural selection with varying selection coefficients—a haploid model. Genetical Research 21, 115–120.

Gillespie, J., 1977. Natural selection for variances in offspring numbers: a new evolutionary principle. Am. Nat. 111, 1010–1014.

Greene, M., Trieschmann, J., Gustavson, S., 1992. Risk and Insurance. South-Western Publishing Co., Cincinatti, OH.

Guttman, D., Dykhuizen, D., 1994. Clonal divergence in *Escherichia coli* as a result of recombination, not mutation. Science 266, 1380–1383.

Hopeer, K., 1999. Risk-spreading and bet-hedging in insect population biology. Ann. Rev. Entomol. 44, 535–560.

Karlin, S., Taylor, H., 1975. A First Course in Stochastic Processes, 2nd Edition. Academic Press, Boston, MA.

Lenski, R.E., Rose, M.R., Simpson, S.C., Tadler, S.C., 1991. Long-term experimental evolution in *Escherichia coli*. 1. Adaptation and divergence during 2;000 generations. Am. Nat. 138, 1315–1341.

Lenski, R.E., Simpson, S.C., Nguyen, T.T., 1994. Genetic-analysis of a plasmid-encoded, host genotype-specific enhancement of bacterial fitness. J. Bacteriol. 176, 3140–3147.

Markowitz, H., 1991. Portfolio Selection: Efficient Diversification of Investment, 2nd Edition. Blackwell, Cambridge, MA.

Maynard-Smith, J., 1982. Evolution and the Theory of Games. Cambridge University Press, Cambridge.

Mayr, E., 1982. The Growth of Biological Thought: Diversity, Evolution, and Inheritance. Belknap Press, Cambridge, MA.

Mira, A., Ochman, H., Moran, N.A., 2001. Deletional bias and the evolution of bacterial genomes. Trends Genet. 17, 589–596.

Msadek, T., 1999. When the going gets tough: survival strategies and environmental signaling networks in *Bacillus subtilis*. Trends Microbiol. 7, 201–207.

Neidhardt, F. (Ed.), 1996. *Escherichia coli and Salmonella*. ASM Press, Washington, D.C.

Noe, R., Hammerstein, P., 1995. Biological markets. Trends Ecol. Evol. 10, 336–339.

Ochman, H., Moran, N.A., 2001. Genes lost and genes found: evolution of bacterial pathogenesis and symbiosis. Science 292, 1096–1098.

Ogden, J., Basher, L., McGlone, M., 1998. Fire, forest regeneration and links with early human habitation: evidence from New Zealand. Ann. Bot. 81, 687–696.

Perrow, C., 1999. Normal Accidents. Princeton University Press, Princeton, NJ.

Pritchard, C., 2001. Risk Management: Concepts and Guidance. ESI International, London, UK.

Rahal, J.J., Urban, C., Horn, D., Freeman, K., SegalMaurer, S., Maurer, J., Mariano, N., Marks, S., Burns, J.M., Dominick, D., et al., 1998. Glass restriction of cephalosporin use to control total cephalosporin resistance in nosocomial klebsiella. Jama-J. Am. Med. Assoc. 280, 1233–1237.

Real, L., 1980. Fitness, uncertainty, and the role of diversification in evolution and behavior. Am. Nat. 115, 623–638.

Russell, P., 1998. Genetics, 5th Edition. Addison-Wesley Longman, Menlo Park, CA.

Savageau, M., 1983. *Escherichia coli* habitats, cell types, and molecular mechanisms of gene control. Am. Nat. 122, 732–744.

Schrag, S.J., Perrot, V., Levin, B.R., 1997. Adaptation to the fitness costs of antibiotic resistance in *Escherichia coli* Proc. Roy. Soc. London Ser. B-Biol. Sci. 264, 1287–1291.

Seppala, H., Klaukka, T., VuopioVarkila, J., Muotiala, A., Helenius, H., Lager, K., Huovinen, P., Kontiainen, S., Eskola, J., Korpela, J., et al., 1997. The effect of changes in the consumption of macrolide antibiotics on erythromycin resistance in group a streptococci in Finland. N. Eng. J. Med. 337, 441–446.

Sonenshein, A., Hoch, J., Losick, R. (Eds.), 2002. *Bacillus subtilis* and its Closest Relatives. From Genes to Cells. ASM Press, Washington, DC.

Stearns, S.C., 1992. The Evolution of Life Histories. Oxford University Press, New York.

Stragier, P., Losick, R., 1996. Molecular genetics of sporulation in *Bacillus subtilis*. Ann. Rev. Genet. 30, 297–341.

Vose, D., 2000. Risk Analysis: A Quantitative Guide. Wiley, New York.

Winzeler, E.A., Shoemaker, D.D., Astromoff, A., Liang, H., Anderson, K., Andre, B., Bangham, R., Benito, R., Boeke, J.D., Bussey, H., et al., 1999. Functional characterization of the s-cerevisiae genome by gene deletion and parallel analysis. Science 285, 901–906.

Zhou, T., Carlson, J., Doyle, J., 2002. Mutation, specialization, and hypersensitivity in highly optimized tolerance. Proc. Nat. Acad. Sci. USA 99, 2049–2054.