Chapter 7

Dynamics of linear branching processes

The previous chapter showed that some of the most interesting dynamics of a new mutation occur while it is still at a low frequency in the population ($f \ll 1$). In this limit, the single-locus model in Eq. (6.1) reduces to the linear SDE,

$$\frac{\partial f}{\partial t} = \underbrace{sf}_{\text{selection}} + \underbrace{\mu - \nu f}_{\text{mutation}} + \underbrace{\sqrt{\frac{f}{N}} \cdot \eta(t)}_{\text{genetic drift}} \tag{7.1}$$

also known as a ***linear branching process***.[1] The reasons for this linear behavior can be motivated by revisiting the microscopic Moran model from Chapter 5. When $f \ll 1$, most competitions involving the mutant occur against a wildtype individual, simply because the number of such pairs ($f \cdot 1$) is much larger than the number of mutant-mutant pairs ($f \cdot f \ll f \cdot 1$):

---

[1] Technically, it is a continuous-time, continuous-state branching process. Other versions exist that discretize the time and/or frequency dimensions. You will analyze one such example in Problem X of HW Y.

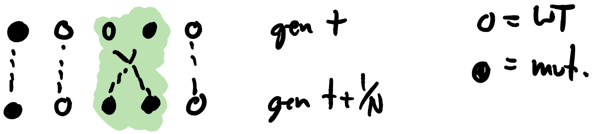

This suggests that selection and genetic drift can be approximated by assuming that the mutant is growing in an environment consisting *solely* of wildtype individuals. The descendants of any two mutant individuals must be independent in such a scenario, and the only way that this can occur is if the selection and drift terms in the SDE are linear functions of $f(t)$.

The independence assumption will clearly break down when the mutant reaches higher frequencies (e.g. 50%). For example, correlations between individuals are eventually critical for ensuring that the mutant frequency cannot exceed 100%. The linear model in Eq. (7.1), by contrast, allows the "frequency" to diverge to infinity. This unboundedness will not be an important feature for us here — we will always make sure to switch back to the full model in Eq. (6.1) long before the mutation reaches 50% frequency (see Section 7.3).

When the independence assumption is satisfied, the linear nature of Eq. (6.1) is simple enough that we will be able to gain a nearly complete picture of the ***temporal dynamics*** of mutations, in addition to the long-time limits (e.g. fixation probabilities and stationary distributions) that we explored Chapter 6. Understanding these dynamics will turn out to give us lots of useful intuition for thinking about evolutionary problems, and they will provide a natural starting point when we go on to consider more complicated scenarios later in the course. These temporal dynamics are also increasingly relevant for analyzing longitudinal data (e.g. ancient DNA, genomic surveillance of pathogens, laboratory evolution experiments, etc.), so a detailed understanding of this case will have useful practical benefits as well.

7.1 Dynamics of the mean and variance

For simplicity, we will first consider the case with no mutations ($\mu = \nu = 0$), where Eq. (7.1) reduces to

$$\frac{\partial f}{\partial t} = \underbrace{sf}_{\text{selection}} + \underbrace{\sqrt{\frac{f}{N}} \cdot \eta(t)}_{\text{genetic drift}} \,. \tag{7.2}$$

Since the selection term is now a linear function of $f(t)$, the moment equations no longer suffer from the "moment hell" that plagued our original model in Chapter 6. The mean frequency now satisfies the deterministic dynamics,

$$\frac{\partial \langle f(t) \rangle}{\partial t} = s \langle f(t) \rangle \,, \tag{7.3}$$

whose solution is a simple exponential growth function,

$$\langle f(t) \rangle = f_0 e^{st} \,. \tag{7.4}$$

Similar results can be obtained for higher moments as well. Repeating the steps in Chapter 6, one can show that the second moment now satisfies,

$$\frac{\langle f(t)^2 \rangle}{\partial t} = 2s \langle f(t)^2 \rangle + \frac{\langle f(t) \rangle}{N} \tag{7.5}$$

Since the mean is given by Eq. (7.4), we can integrate this linear ODE to obtain

$$\langle f(t)^2 \rangle = f_0^2 e^{2st} + \frac{f_0 e^{st} \left(e^{st} - 1 \right)}{Ns} \,. \tag{7.6}$$

The first term corresponds to the deterministic expectation, $\langle f(t)^2 \rangle \approx \langle f(t) \rangle^2$, while the second term is a new contribution due to genetic drift. It will be useful to express this result in terms of the ***coefficient of variation (CV)***,

$$c_V^2(t) \equiv \frac{\mathrm{Var}(f(t))}{\langle f(t) \rangle^2} = \frac{1 - e^{-st}}{N s f_0} \,, \tag{7.7}$$

The coefficient of variation is useful for visualizing the spread of a distribution in log space (i.e. how uncertain are we at an order-of-magnitude level). For example, for a “Case 1” distribution with $x = \langle x \rangle \pm \sigma$, we have

$$\log x = \log\left(\langle x \rangle \pm \sigma\right) \approx \log\langle x \rangle \pm c_V \tag{7.8}$$

when $c_V \ll 1$. When the coefficient of variation starts to exceed one, the average becomes a poor approximation for actual value of the mutation frequency.

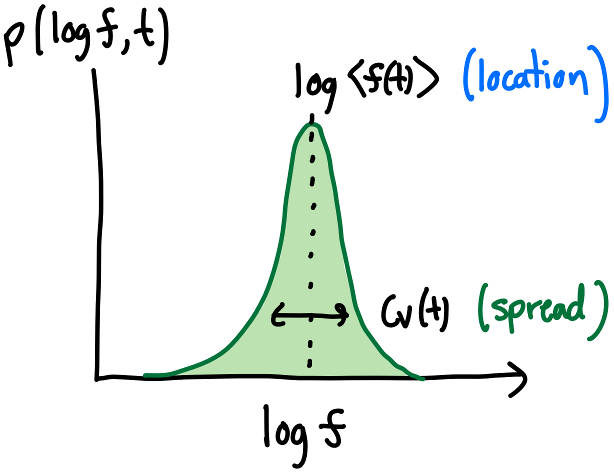

The coefficient of variation in Eq. (7.7) starts out with $c_V(0) \approx 0$, since we have assumed that the mutation begins at a fixed initial frequency. The behavior at later times strongly depends on the relative values of N, s, and f_0:

Case 1. For a positively selected mutation ($s > 0$), the coefficient of variation is bounded by its long-term value,

$$c_V^2(t) \leq \frac{1}{Nsf_0} \tag{7.9}$$

Thus, if the mutation starts out in the selection-dominated region of frequency space ($f_0 \gg 1/2Ns$), then coefficient of variation will remain small at all later times $[c_V(t) \ll 1]$. This implies that the frequency of the mutation will be well-approximated by its average value, $\langle f(t)\rangle = f_0 e^{st}$ (i.e. the distribution will be of the "Case 1" form from Chapter 2). These results are consistent with our fixation probability calculation from Chapter 6, which showed that beneficial mutations are guaranteed to fix when $f_0 \gg 1/2Ns$.

Case 2. In contrast, when a beneficial mutation starts out in the drift-dominated region of frequency space ($f_0 \ll 1/2Ns$), its coefficient of variation will initially be very small $[c_V(t) \ll 1]$, but it will eventually reach a point where $c_V(t) \gg 1$. Similar behavior will occur for neutral or deleterious mutations. The location of this transition can be defined by the critical time t^* where $c_V(t^*) \approx 1$.

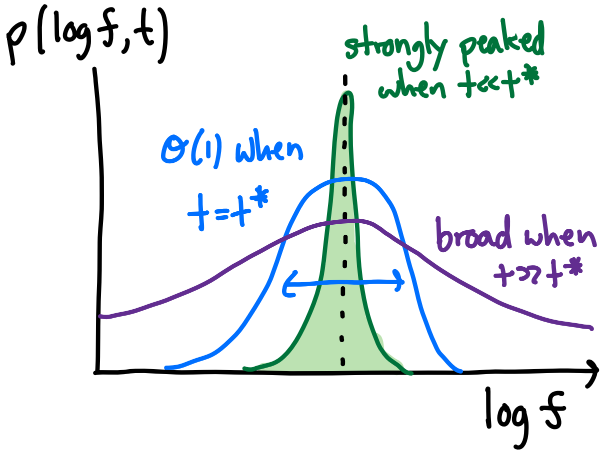

Solving for t^* yields

$$t^* \approx \begin{cases} Nf_0 & \text{if } f_0 \ll 1/N|s|, \\ \frac{1}{|s|}\log(N|s|f_0) & \text{if } s < 0 \text{ and } f_0 \gg 1/N|s|, \\ \infty & \text{if } s > 0 \text{ and } f_0 \gg 1/Ns. \end{cases} \tag{7.10}$$

When $t \ll t^*$, the coefficient of variation is very small $[c_V(t) \ll 1]$, and the frequency of the mutation can be well-approximated by its average value,

$\langle f(t)\rangle = f_0 e^{st}$. In contrast, when $t \gtrsim t^*$, the distribution of $f(t)$ will become extremely broad, and will approach a "Case 2" form whose properties we will derive below.

7.2 Solving for the full distribution

One of the most useful features of the branching process model in Eq. (7.2) is that it allows us to solve for the full distribution of $f(t)$. We could in principle do this by solving the Fokker-Planck equation,

$$\frac{\partial p(f,t)}{\partial t} = \underbrace{-\frac{\partial}{\partial f}\left[sfp(f,t)\right]}_{\text{selection}} + \underbrace{\frac{\partial^2}{\partial f^2}\left[\frac{f}{N}\cdot p(f,t)\right]}_{\text{genetic drift}}, \tag{7.11}$$

but the second derivative on the right-hand side makes this a difficult task (see the Appendix of Chapter 6). In this case, it will be much easier to work with the moment generating function of $f(t)$:

$$H(z,t) \equiv \langle e^{-zf(t)}\rangle \equiv \int e^{-zf} p(f,t)\, df\,, \tag{7.12}$$

which is governed by an analogous PDE,

$$\frac{\partial H}{\partial t} = \left[sz - \frac{z^2}{2N}\right]\frac{\partial H}{\partial z}\,. \tag{7.13}$$

subject to the initial condition $H(z,0) = e^{-zf_0}$. The main difference from our original model in Eq. (6.47) is that the branching process version contains only a single z derivative. PDEs of this form can be solved using a technique known as the ***method of characteristics***, which is a generalization of the trick that we used to solve for the fixation probability in Chapter 6. The details of this derivation are presented in the Appendix at the end of the chapter. For now, we will simply quote the final solution,

$$H(z,t) = \exp\left[\frac{-zf_0e^{st}}{1+\frac{z}{2Ns}\left(e^{st}-1\right)}\right] . \tag{7.14}$$

Formally, it is possible to invert this expression to obtain the corresponding probability distribution $p(f,t)$. However, the details are somewhat complicated, and the resulting expressions can be difficult to interpret in the general case.[2] Instead, we will see that one can actually learn a lot about $p(f,t)$ by examining the generating function $H(z,t)$ directly.

For example, using our results for the mean and variance above, we can rewrite $H(z,t)$ in the convenient form,

$$H(z,t) = \exp\left[\frac{-z\langle f(t)\rangle}{1+z\langle f(t)\rangle\cdot\frac{c_V^2(t)}{2}}\right] \tag{7.15}$$

By comparing this result to the generating function for a Gaussian random variable (Chapter 2),

$$\left\langle e^{-zx}\right\rangle = e^{-z\langle x\rangle+z^2\langle x\rangle^2\cdot\frac{c_V^2}{2}} \tag{7.16}$$

we can see that $f(t)$ is *not* normally distributed in general, but becomes approximately normally distributed in the limit that $c_V(t) \ll 1$. Our results in Eq. (7.7) show that this will be a good approximation at short times, but it will eventually break down for $t \gtrsim t^*$ in Eq. (7.10), when $c_V(t) \gtrsim 1$. What can we say about the distribution of $f(t)$ in these cases?

Extinction and survival probabilities

When the variation in $f(t)$ is as large as its mean [$c_V(t) \gtrsim 1$], we must consider the possibility that the mutant has gone extinct [$f(t) = 0$]. The probability of

---

[2] The details of this inversion are presented in an appendix at the end of this chapter.

this event can also be easily extracted from the generating function in Eq. (7.14). Recalling the definition of the generating function,

$$H(z,t) \equiv \int e^{-zf} p(f,t)\, df \tag{7.17}$$

we can see that the exponential factor acts like a crude version of a step function, approaching a uniform value for $f \ll 1/z$, and excluding contributions from frequencies with $f \gg 1/z$. In the extreme limit where $z \to \infty$, the only values of f that will contribute to the generating function integral are those with $f = 0$; all of the nonzero frequencies will have $e^{-zf} \to 0$. This implies that

$$\lim_{z\to\infty} H(z,t) = 1 \cdot p_{\text{ext}}(t) + 0 \cdot [1 - p_{\text{ext}}(t)] = p_{\text{ext}}(t) \tag{7.18}$$

where $p_{\text{ext}}(t)$ is the ***time-dependent extinction probability*** of the mutation (i.e., the probability that it has gone extinct *by* time t). Using our expression for $H(z,t)$ in Eq. (7.14), we find that

$$p_{\text{ext}}(t) = \exp\left[\frac{-2Nsf_0}{e^{st}-1}\right] = \exp\left[\frac{-2}{c_V^2(t)}\right] . \tag{7.19}$$

We can also define a corresponding ***survival probability***

$$p_s(t) = 1 - p_{\text{ext}}(t) = 1 - \exp\left[\frac{-2Nsf_0}{1-e^{-st}}\right] , \tag{7.20}$$

which denotes the probability that the mutant is still alive at time t.

These expressions show that the extinction and survival probabilities are intimately connected to the coefficient of variation in Eq. (7.7). At early times ($t \ll t^*$), the coefficient of variation is very small [$c_V(t) \ll 1$], so there is a negligible chance of extinction (i.e., the survival probability is close to 100%). However, once $c_V(t) \gtrsim 1$, there is a decent chance that the mutant has now gone extinct. This implies that the crossover time t^* in Eq. (7.10) can also be interpreted as a characteristic ***extinction time*** — i.e. the time at which the survival probability starts to drop below 100%. For a beneficial mutation that starts out in the selection-dominated regime ($f_0 \gg 1/2Ns$), the survival probability in Eq. (7.20) remains close to 100% at all later times. In all other cases, the survival probability will eventually become very small [$p_s(t) \ll 1$], and there is a large chance that the mutant has gone extinct.

Conditioning on non-extinction

Our results above suggest that when $t \gtrsim t^*$, the distribution of $f(t)$ will approach a "Case 2" form that contains a mixture of two different types of mutation trajectories: (i) ***extinct paths***, which have $f(t) = 0$, and (ii) ***non-extinct paths*** where $f(t) > 0$.

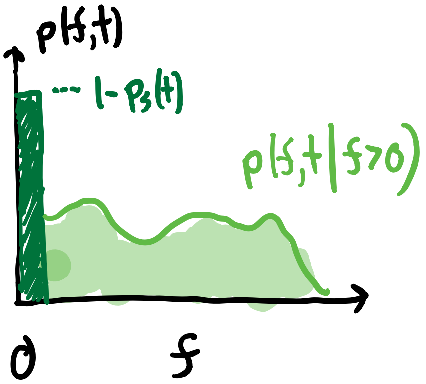

We can formalize this idea by writing $p(f, t)$ as a mixture of two components,

$$p(f,t) = \underbrace{[1 - p_s(t)]\,\delta(f)}_{\text{extinct paths}} + \underbrace{p_s(t) \cdot p(f,t|f>0)}_{\text{non-extinct paths}} \tag{7.21}$$

where $p(f, t|f > 0)$ denotes the ***conditional distribution*** of $f(t)$, given that it has survived for a time t. Since $p_s(t)$ is known, this conditional distribution contains all the non-trivial features of the full distribution $p(f, t)$. What can we learn about the frequencies of these surviving lineages?

One of the simplest things we can do is look at the mean of $p(f,t|f>0)$. Multiplying both sides of Eq. (7.21) by f and integrating, we find that

$$\langle f(t)\rangle = \underbrace{[1-p_s(t)]\cdot 0}_{\text{extinct paths}} + \underbrace{\langle f(t)|f>0\rangle \cdot p_s(t)}_{\text{non-extinct paths}}\,, \tag{7.22}$$

where $\langle f(t)|f>0\rangle$ denotes the ***conditional average*** of $f(t)$, given that it is alive at time t. We can rearrange this equation to obtain a formula for the conditional mean as a function of $\langle f(t)\rangle$ and $p_s(t)$:

$$\langle f(t)|f>0\rangle = \frac{\langle f(t)\rangle}{p_s(t)} \approx \begin{cases} f_0 e^{st} & \text{if } t \ll t^*, \\ \frac{e^{st}-1}{2Ns} & \text{if } t \gg t^*. \end{cases} \tag{7.23}$$

At short times ($t \ll t^*$), the survival probability remains close to 100%, and the conditional mean stays close to the unconditional average, $\langle f(t)\rangle = f_0 e^{st}$. In contrast, when the time starts to exceed the characteristic extinction time ($t \gg t^*$), the average frequency of a surviving lineage becomes much larger than $\langle f(t)\rangle$, since a large fraction of the lineages will have gone extinct.

In this latter case, the conditional mean follows a qualitatively different trajectory than the deterministic expectation $f_0 e^{st}$. In particular, the average frequency of a surviving lineage becomes independent of f_0. This initial frequency plays a crucial role in determining the survival probability of the mutation. But provided that it survives, its average frequency will lose all memory of where it started. We can gain some additional insight into this behavior by splitting Eq. (7.23) into three different cases depending on the relative values of s and t:

$$\langle f(t)|f>0\rangle \xrightarrow{t\gg t^*} \begin{cases} \frac{1}{2Ns}\cdot e^{st} & \text{if } s>0 \text{ and } t \gg 1/s, \\ \frac{t}{2N} & \text{if } t \ll 1/|s|, \\ \frac{1}{2N|s|} & \text{if } s<0 \text{ and } t \gg 1/|s|. \end{cases} \tag{7.24}$$

These three cases show that

1. Surviving beneficial mutations (eventually) grow exponentially with time, but with a different pre-factor compared to the deterministic expectation:

2. Beneficial and deleterious mutations look like neutral mutations on short timescales ($t^* \ll t \ll 1/|s|$). The frequencies of these mutations grow linearly with time, as opposed to the constant value we would expect in the absence of noise:

3. At longer times, deleterious mutations eventually saturate at a constant value, rather than declining exponentially with time.

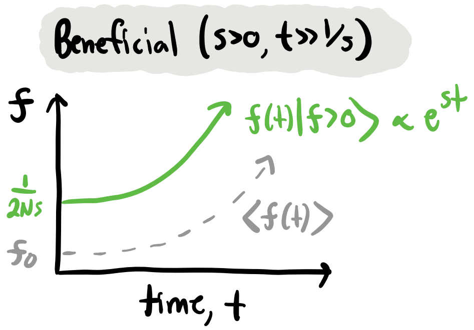

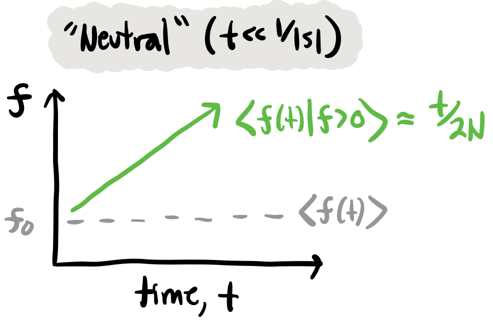

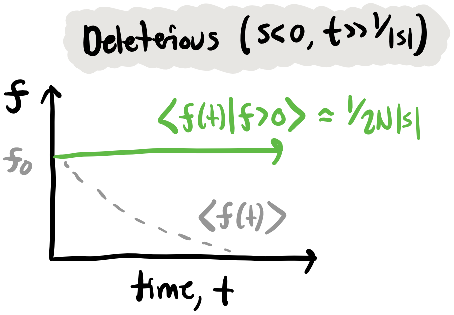

The corresponding survival probabilities in each of these three cases are given by

$$p_s(t) \xrightarrow{t \gg t^*} \begin{cases} 2Nsf_0 & \text{if } s > 0 \text{ and } t \gg 1/s, \\ 2Nf_0/t & \text{if } t \ll 1/|s|, \\ 2N|s|f_0e^{-st} & \text{if } s < 0 \text{ and } t \gg 1/|s|. \end{cases} \tag{7.25}$$

We can see that in each case, the survival probability and conditional mean are perfectly set up so that their product is equal to the unconditioned average $\langle f(t) \rangle$.

We can use a similar argument to calculate the full conditional distribution $p(f, t|f > 0)$. This is easiest to do by going through the generating function $H(z, t)$. Multiplying both sides of Eq. (7.21) by e^{-zf} and integrating, we find that

$$H(z,t) = \underbrace{[1 - p_s(t)] \cdot e^{-z \cdot 0}}_{\text{extinct paths}} + \underbrace{p_s(t) \cdot H(z,t|f>0)}_{\text{surviving paths}} , \tag{7.26}$$

where we have defined the ***conditional generating function***,

$$H(z,t|f>0) = \int e^{-zf} p(f,t|f>0)\, df \, . \tag{7.27}$$

We can rearrange this equation to obtain a formula for $H(z, t|f > 0)$:

$$H(z,t|f>0) = \frac{H(z,t) - [1-p_s(t)]}{p_s(t)} = \frac{e^{\frac{-z\langle f(t)\rangle}{1+z\langle f(t)\rangle c_V^2(t)}} - e^{-\frac{2}{c_V^2(t)}}}{1 - e^{-\frac{2}{c_V^2(t)}}} \, . \tag{7.28}$$

We can simplify this expression by noting that the argument of the first exponential is maximized when $z = \infty$, where $H(z, t) = e^{-2/c_V^2(t)}$. Thus, in the long time limit ($t \gg t^*$) where $c_V(t) \gg 1$, we can Taylor expand each of the exponential terms in Eq. (7.28) to obtain

$$H(z,t|f>0) \xrightarrow{t \gg t^*} \frac{\left(1 - \frac{z\langle f\rangle}{1+z\langle f\rangle\frac{c_V^2}{2}}\right) - \left(1 - \frac{2}{c_V^2}\right)}{1 - \left(1 - \frac{2}{c_V^2}\right)} = \frac{1}{1 + z \cdot \langle f(t)|f>0\rangle} \, , \tag{7.29}$$

where we have used the fact that $\langle f(t)|f > 0\rangle \approx 2\langle f(t)\rangle/c_V^2(t)$. By the ***"method of Wikipedia"***, we can see that this expression matches the generating function for an exponential distribution with a mean equal to the conditional average $\langle f(t)|f > 0\rangle$:

$$p(f,t|f>0) \approx \frac{e^{-f/\langle f(t)|f>0\rangle}}{\langle f(t)|f>0\rangle} \, , \tag{7.30}$$

The exponential distribution occupies an intermediate zone between the "Case 1" and "Case 2" distributions in Chapter 2. While it may look broadly distributed when we plot it on a linear scale, if we plot it in terms of $\log f$, we see that most of its probability is concentrated within one order of magnitude of the mean.

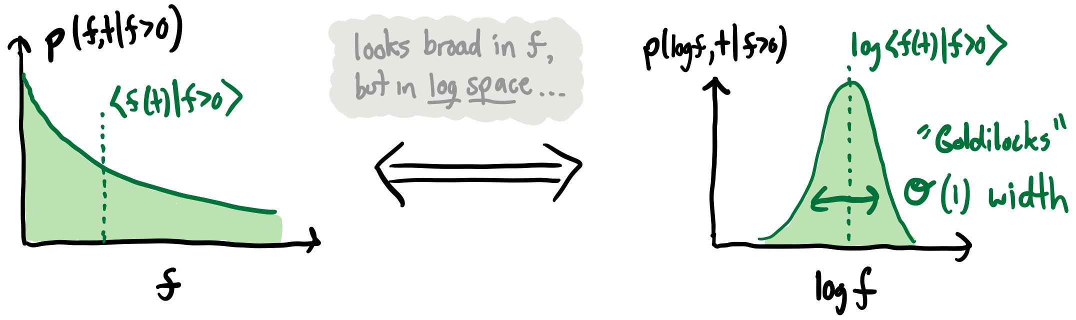

We can therefore think of the exponential distribution as a ***"Goldilocks case"*** where the mean is a *reasonable* summary of the typical behavior [i.e. one that is accurate up to an $\mathcal{O}(1)$ pre-factor].

Putting everything together, we can conclude that the temporal dynamics of $p(f, t)$ can be divided into two characteristic regimes:

- **Case 1:** At short times ($t \ll t^*$), the distribution of $f(t)$ will be strongly peaked around its deterministic expectation $\langle f(t)\rangle = f_0 e^{st}$, with Gaussian fluctuations of size $\pm\langle f(t)\rangle \cdot c_V(t)$.

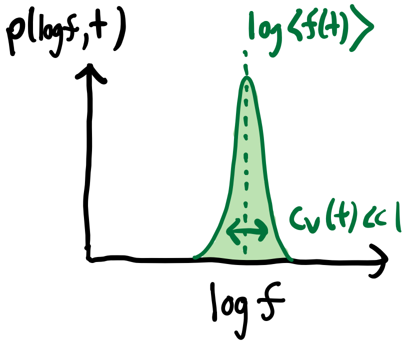

For a mutation that starts at a frequency $f_0 \ll 1/2N|s|$, this initial phase will last for $t^* \sim Nf_0$ generations, and the mean and variance are given by $f_0 \pm f_0 t/N$.

- **Case 2:** At longer times ($t \gg t^*$), the distribution of $f(t)$ will split into a bimodal shape, with a large fraction of mutations going extinct ($f = 0$). The surviving mutations will follow an exponential distribution, whose average size marches toward higher frequencies according to the conditional mean $\langle f(t) | f > 0 \rangle$ in Eq. (7.23).

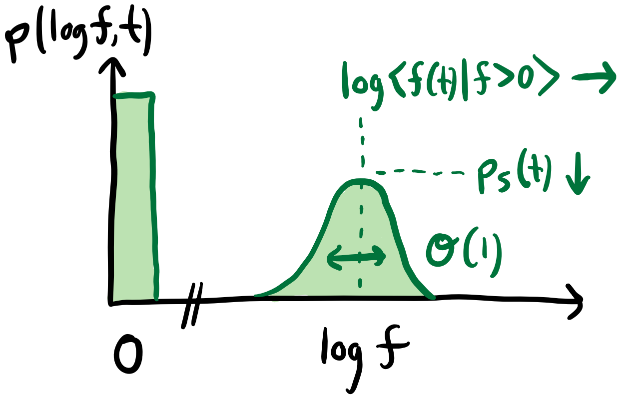

New mutations always start out in Case 2 (since $t^* \approx 1$), which implies that:

Transient distribution of a new mutation ($f \ll 1$)

For a new mutation with an initial frequency $f_0 = 1/N$, the frequencies at later times follow the mixture distribution,

$$p(f,t) = [1 - p_s(t)] \cdot \delta(f) + \frac{p_s(t)}{f_s(t)} e^{-t/f_s(t)}\,, \tag{7.31}$$

where $f_s(t) = (e^{st} - 1)/2Ns$ and $p_s(t) = 2s/(1 - e^{-st})$.

This leads to several interesting conclusions:

1. At early times ($t \ll 1/|s|$), the distribution of $f(t)$ is *indistinguishable* from that of a neutral mutation, even when $N|s| \gg 1$. This shows that our previous deduction from the fixation probability in Chapter 6 extends to the full dynamics of $f(t)$ as well.

2. In this initial phase ($t \ll 1/|s|$), the typical frequencies of the surviving mutations grow linearly in time,

$$\langle f(t) | f > 0 \rangle \approx \frac{t}{2N} \tag{7.32}$$

   Since the frequencies are exponentially distributed, this means that we would need to wait for $t \sim Nf$ generations to have an appreciable chance of observing a mutation with $f(t) \sim f$ when $t \ll 1/|s|$.

3. For a beneficial mutation, the linear growth in Eq. (7.32) is actually *faster* than the deterministic expectation, $f_0 e^{st}$. If we could measure its frequency trajectory in this early stage, its *apparent* fitness benefit [i.e. the slope of $\log f(t)$] would appear to be much larger than s:

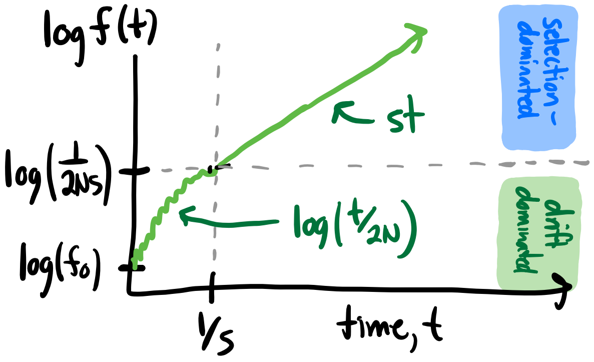

This linear growth is even more counterintuitive for neutral or deleterious mutations, since it implies that the average frequency of a *surviving* mutation tends to increase with time — even though the *overall* average is flat or declining. This paradoxical behavior can be reconciled by remembering that vast majority of the mutations have gone extinct by this point. By restricting our attention to the surviving mutations, we are biasing our observations toward the lucky trajectories that managed to avoid extinction by drifting to anomalously large frequencies.

4. Once $t \gg 1/|s|$, natural selection starts to exert its effect. Deleterious mutations are prevented from rising much higher than $1/2N|s|$, while beneficial mutations start to grow as $\frac{1}{2Ns}e^{st}$, with an overall prefactor set by the exponential distribution.

We can gain some additional insight by examining the whole ***ensemble of random paths***, $p[f(t_1), f(t_2), \ldots, f(t_n)]$, rather than just the frequency at the final timepoint. We can do this by recursively applying the results above. This yields an intuitive picture for the trajectory of a mutant lineage that starts at a frequency $f_0 \ll 1/2N|s|$:

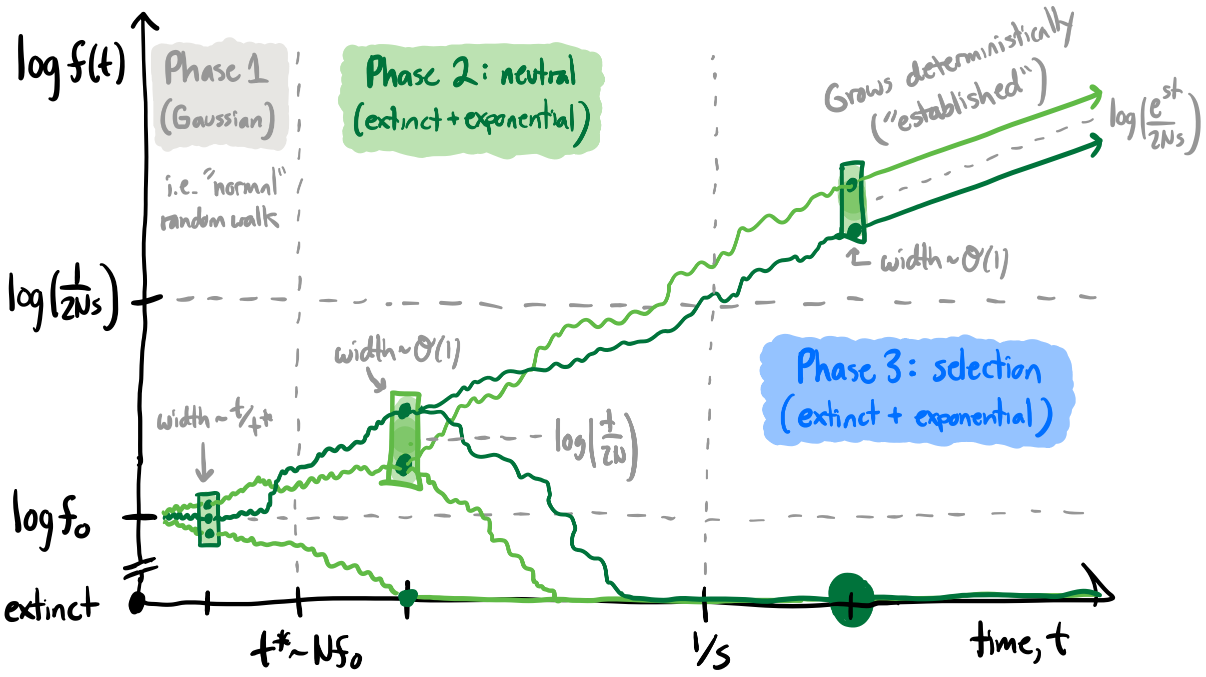

The mutation will initially behave like an ordinary random walk, with small fluctuations around f_0. After a time of order $t^* \sim Nf_0$, the mutation will forget its initial frequency and begin to grow linearly as $f(t) \approx t/2N$, while an increasing fraction will drift to extinction. There will be $\mathcal{O}(1)$ fluctuations around this mean, but these will be forgotten after another $Nf(t)$ generations. In each iteration, some of the current surviving paths will have a chance of drifting to extinction.

After a time of order $\sim 1/|s|$, a typical surviving mutation will have reached a frequency of order $\sim 1/2N|s|$. Deleterious mutations will get "stuck" at this point, while an increasing fraction will continue to drift to extinction. Beneficial mutations, on the other hand, will begin to grow deterministically at rate e^{st}. In this case, we say that the beneficial mutation has ***"established"***, since it will have a negligible probability of going extinct. The total probability that it reaches this point ($2Nsf_0$) is sometimes known as the ***establishment probability***.

Further fluctuations in the frequency will be small from this point forward, so most of the variation in $f(t)$ will come from the last round of fluctuations that occurred when the frequency was close to $\sim 1/Ns$. We can formalize this idea by introducing a new random variable $\nu(t)$ that factors out the expected time-dependence of $f(t)$:

$$f(t) \equiv \frac{\nu(t)}{2Ns} e^{st} \tag{7.33}$$

From the definition of the generating function, we have

$$H_\nu(z,t) = \langle e^{-z\cdot\nu(t)} \rangle = \langle e^{-z\cdot 2Nse^{-st} f(t)} \rangle = H_f(2Nse^{-st}\cdot z, t) \tag{7.34}$$

When $t \gg 1/s \gg t^*$, the conditional distribution of ν reduces to

$$H_\nu(z,t|f>0) \approx \frac{1}{1+2Nsz} \implies \nu \sim \text{Exponential}(1) \tag{7.35}$$

This suggests that the randomness in the frequency trajectory when $f(t) \gg 1/2Ns$ can be captured by a single $\mathcal{O}(1)$ pre-factor,

$$f(t) = \frac{\nu}{2Ns} e^{st} \, . \tag{7.36}$$

which is "frozen in" once $t \gg 1/s$.

7.3 Asymptotic matching at higher frequencies

An established beneficial mutation cannot grow exponentially forever. At some point, it will reach a large enough size that $f(t)$ will no longer be small compared to one, and our linear branching process approximation will break down. This will require us to switch back to the full single-locus diffusion model from Chapter 6. Fortunately, we have already shown that genetic drift will be negligible at these higher frequencies [since $f(t) \gg 1/2Ns$], so we can replace the full model with the deterministic version,

$$\frac{\partial f}{\partial t} = sf(1-f) \, . \tag{7.37}$$

We can implement the "handoff" between this model and the linear branching process in Eq. (7.2) using a technique known as ***asymptotic matching***. The basic idea applies anytime that we have two different approximations that overlap in a smaller region of parameter space. By matching the two approximations in the region where they are both valid, we can find a "global" approximation that is valid across the entire range.

In this case, the overlap region is the part of frequency space where $\frac{1}{2Ns} \ll f \ll 1$. The upper condition ($f \ll 1$) ensures that the linear branching process is a good approximation, while the lower condition ($f \gg 1/2Ns$) ensures that the deterministic approximation in Eq. (7.37) is also valid.

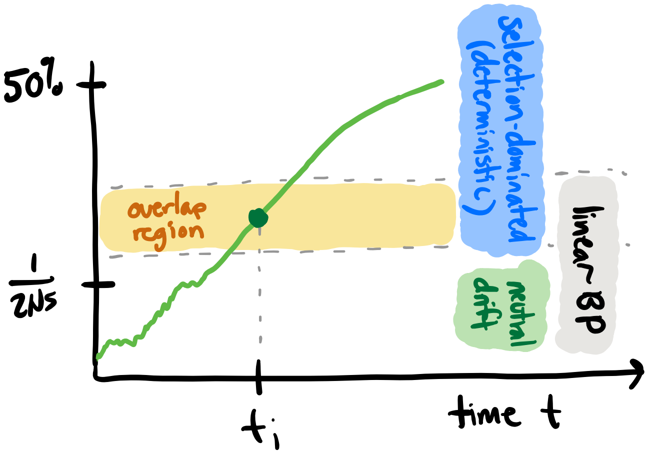

Having identified the relevant overlap region, we can implement our asymptotic matching procedure using the following steps:

- **Step 1.** To make things easier, we will first convert the overlap region from the y-axis (frequency) to the x-axis (time). We can do this by choosing an intermediate time t_i such that $\frac{1}{2Ns} \ll f(t_i) \ll 1$. When these conditions are satisfied, our results above imply that

$$f(t_i) = \begin{cases} \frac{\nu}{2Ns} e^{st_i} & \text{w/ prob } 2Nsf_0, \\ 0 & \text{else,} \end{cases} \tag{7.38}$$

where ν is an exponential random variable with mean one. The bottom case is easy to extrapolate to later times, since the mutation must remain extinct for all t. This means that we only need to consider the established case. In this case, Eq. (7.38) shows that our assumptions about $f(t_i)$ will be self-consistent if

$$\frac{1}{s} \ll t_i \ll \frac{1}{s} \log(Ns) \,. \tag{7.39}$$

When $Ns \gg 1$, there are (in principle[3]) many possible t_i values where this will be true. Let's choose one of them.

- **Step 2.** We can then use this intermediate timepoint as the starting point for the deterministic model in Eq. (7.37). This yields the frequency at any later timepoint:

  $$f(t) = \frac{f(t_i)e^{s(t-t_i)}}{1 - f(t_i) + f(t_i)e^{s(t-t_i)}} \approx \frac{f(t_i)e^{s(t-t_i)}}{1 + f(t_i)e^{s(t-t_i)}} \tag{7.40}$$

  where we have assumed that $f(t_i) \ll 1$.

- **Step 3.** We can then substitute our expression for $f(t_i)$ into Eq. (7.40) to obtain a "global" approximation,

  $$f(t) = \frac{\left(\frac{\nu}{2Ns}e^{st_i}\right) e^{s(t-t_i)}}{1 + \left(\frac{\nu}{2Ns}e^{st_i}\right) e^{s(t-t_i)}} = \frac{\frac{\nu}{2Ns}e^{st}}{1 + \frac{\nu}{2Ns}e^{st}} \,, \tag{7.41}$$

  that is valid for $t \gg 1/s$. Note that the intermediate time t_i has dropped out of our final answer. This is a good thing, since our choice of t_i was completely arbitrary, and any downstream predictions should not depend on the precise value that we chose.

---

[3] This is the "asymptotic" part of asymptotic matching – it requires us to assume that $Ns \to \infty$. In practice, since $\log(Ns)$ is a slowly growing function of Ns, the difference between $1/s$ and $\log(Ns)/s$ is never *that* big. Fortunately, like many of other asymptotic approximations we discussed in Chapter 2, we will see that the asymptotic matching approximation remains relatively accurate in practice, even for moderate values of $\log(Ns)$.

We can use our global approximation in Eq. (7.41) to calculate an interesting biological quantity: how long does it take for a brand new mutation to take over the population and fix? This is often known as the ***fixation time***. We can break this question down into two parts: (i) how long does it take a mutation to go from $f_0 = 1/N$ to $f = 50\%$? and (ii) how long does it take to go from 50% to 100%?

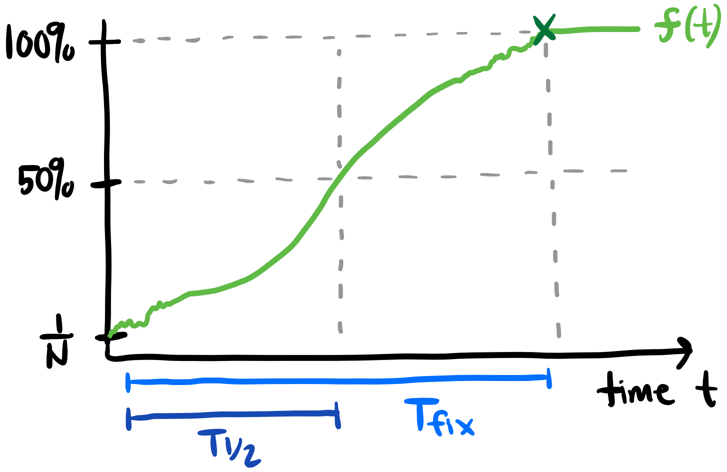

We can answer the first question using our expression in Eq. (7.41). Setting $f(t) = 1/2$ and solving for t yields:

$$T_{1/2} = \frac{1}{s} \log\left(\frac{2Ns}{\nu}\right) = \frac{1}{s}\log(2Ns) + \frac{1}{s}\log\left(\frac{1}{\nu}\right) \tag{7.42}$$

The second half of the trajectory is mirror symmetric with the first if we reverse the direction of time ($t \to -t$) and focus on the wildtype fraction ($f \to 1-f$). The total ***fixation time*** is given by

$$T_{\rm fix} = \frac{2}{s}\log(2Ns) + \frac{1}{s}\log\left(\frac{1}{\nu_1}\right) + \frac{1}{2}\log\left(\frac{1}{\nu_2}\right) \tag{7.43}$$

Since $Ns \gg 1$, the first term is much larger than the other two, which yields the leading order approximation,

Fixation time of a beneficial mutation $(Ns \gg 1)$

$$T_{\text{fix}} \approx \frac{2}{s} \log(2Ns)\,. \tag{7.44}$$

It is easy to imagine that evolution might be limited by the time it takes to find the right mutation. Equation (7.44) shows that there is also a fundamental "speed limit" on the time it takes for this newly produced mutation to take over the population and fix.

In large populations, the total fixation time can be much larger than the time it takes for a mutation to go from $10\% \to 90\%$, or even $1\% \to 99\%$. Our decomposition shows that most of this time elapses in the deterministic phase of the mutation's lifetime $[f(t) \gtrsim 1/2Ns]$, while it is still at a low frequency in the population $[f(t) \ll 1]$. You will have the chance to work through some concrete examples in Problem 5 of Homework 3. This speed limit on the fixation time of a new mutation will be very important when we consider longer genomes later in the course.

7.4 Heuristic picture

Because the linear branching process spans both the drift-dominated and selection-dominated parts of frequency space, it is able to *quantitatively* capture the complex transition between stochastic and deterministic growth. This transition is critical for understanding how new mutations take over a population. At the same time, we saw that obtaining these analytical results required some complex mathematical machinery (e.g., moment equations, generating functions, asymptotic matching) which can obscure some of the underlying physical intuition.

In this section, we will introduce an alternative ***heuristic approach*** for deriving many of the exact results we have obtained so far. These heuristic calculations may seem sloppy or arbitrary at first, but they are more precise than they look: with a bit of care, we will see that they can be done in a way that keeps track of the approximations in a controlled manner, while highlighting the key physical concepts. This can sometimes enable progress in more complicated settings where exact results are not available — you'll have a chance to work through a concrete example involving ***"recessive"*** mutations (relevant for genetic diseases in humans) in Problem 1 of Homework 3.

Warmup: Gaussian Random Walk

We will start by revisiting the simple Gaussian random walk from Chapter 6,

$$\frac{\partial x}{\partial t} = \mu + \sqrt{\sigma^2} \cdot \eta(t) \, , \tag{7.45}$$

which has the exact solution,

$$x(t) = x(0) + \mu t + \sqrt{\sigma^2 t} \cdot Z_t \, , \tag{7.46}$$

that is valid for any time interval t. In the spirit of our dominant balance approach from Chapter 2, we can then ask:

> *Do the stochastic or deterministic terms provide the dominant contribution to* $\Delta x = x(t) - x(0)$ *over a given time interval?*

This is secretly a question about timescales. Since the stochastic component is proportional to $\sqrt{t}$, it will always dominate over the deterministic contribution (μt) when $t \to 0$. For the same reason, the deterministic contribution will always dominate over the stochastic term when $t \to \infty$.

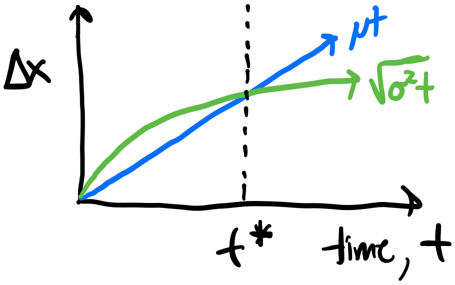

This implies that there will be a ***crossover time*** t^* where the two contributions will have equal magnitude:

$$\mu t^* = \sqrt{\sigma^2 t^*} \implies t^* = \frac{\sigma^2}{\mu} \tag{7.47}$$

When $t \gg t^*$, we can think about $x(t)$ as being mostly deterministic (i.e., $x(t) \approx \mu t$ plus a small correction). Conversely, when $t \ll t^*$, we can approximate $x(t)$ as being mostly stochastic ($x(t) \approx \sigma Z_t$ plus a small correction).

Back to evolution

We can now return to our evolutionary model,

$$\frac{\partial f}{\partial t} = \underbrace{sf(t)}_{\substack{\text{selection} \\ \textit{(deterministic)}}} + \underbrace{\sqrt{\frac{f(t)}{N}}\eta(t)}_{\substack{\text{genetic drift} \\ \textit{(stochastic)}}} \tag{7.48}$$

and ask the same question:

> *Do the deterministic (i.e. selection) or stochastic (genetic drift) terms provide the dominant contribution to $f(t)$ over a finite time interval t?*

This question is more difficult to answer than it was for the simple random walk example above, since the drift and selection terms now depend on $f(t)$ (which itself depends on the drift and selection terms at earlier timepoints, and so on). In the previous sections, we overcame this problem above by developing methods for "integrating" SDEs like Eq. (7.48), e.g. using the moment generating function. This was hard! And it only provided analytical results in simple cases like the linear branching process.

Heuristic calculations are a way to perform this same integration task approximately (i.e., ***"poor man's integration"***). The basic idea is to start from the infinitesimal version of Eq. (7.48) and ask how far we can get by extrapolating it to a *non-infinitesimal* time interval Δt:

$$\Delta f \approx s f(0) \cdot \Delta t + \sqrt{\frac{f(0)\Delta t}{N}} \cdot Z_0 \,, \tag{7.49}$$

This approach can't work forever, since $f(0)$ will eventually be a poor approximation for $f(t)$. However, if we restrict our attention to time intervals where Δf is of the same order-of-magnitude as $f(0)$, then the drift and selection terms in Eq. (7.49) will only be off by an $\mathcal{O}(1)$ factor. This implies that the value of Δf estimated from Eq. (7.49) will only be off by an $\mathcal{O}(1)$ factor as well. Thus, if we're willing to work to ***"logarithmic precision"*** — i.e. up to an $\mathcal{O}(1)$ prefactor in front of various quantities — then Eq. (7.49) will be a good approximation to our original model up to the point where

$$|\Delta f| = cf \,, \tag{7.50}$$

for some $\mathcal{O}(1)$ constant c. Since the exact value of c is not that important, we will often denote this using the "of-order" notation,

$$\text{“}|\Delta f| \sim f\text{”} \tag{7.51}$$

which is shorthand for Eq. (7.50).

Using this notation, we can define the ***"reset time"*** Δt^* to be the point where our short-time approximation in Eq. (7.49) first reaches $|\Delta f| \sim f(0)$. At this point, we can repeat the whole process starting from the new initial condition, $f(\Delta t) = f(0) + \Delta f$, and so on. This yields an iterative method for building up solutions for $f(t)$ when t is longer than a single Δt^*. It is conceptually very similar to ***Euler's method*** for numerically integrating differential equations — the main difference is that we are now trying to use it to obtain approximate analytical solutions instead.

This new framework lets us recast our original question in a more precise manner:

> *Do deterministic forces (selection) or stochastic forces (drift) provide the dominant contribution to the change in frequency,*
>
> $$\Delta f = \underbrace{sf \cdot \Delta t}_{\Delta f_{\text{sel}}} + \underbrace{\sqrt{\frac{f\Delta t}{N}} \cdot Z}_{\Delta f_{\text{drift}}}\,, \tag{7.52}$$
>
> *during a characteristic reset time Δt^*, where $\Delta f \sim f$?*

We can answer this question using the dominant balance approach from Chapter 2, by asking when each of the two possibilities are self-consistent:

Case 1. If deterministic forces are dominant ($\Delta f_{\text{sel}} \gg \Delta f_{\text{drift}}$), then the reset time is defined by

$$f \sim |\Delta f_{\text{sel}}| \sim |s| f \Delta t^*\,, \tag{7.53}$$

or

$$\Delta t^* \sim T_{\text{sel}} \equiv \frac{1}{|s|}\,, \tag{7.54}$$

where we have defined a characteristic ***selection timescale*** $T_{\rm sel} \equiv 1/s$. Recall that is shorthand for the statement that $\Delta t^* = c_1/|s|$ for some $\mathcal{O}(1)$ constant c_1. The precise value of c_1 will not be that important if we are working to logarithmic precision in s and t.

We can now ask when this assumption is self-consistent. Over the timescale $\Delta t^* \sim 1/|s|$, the corresponding contribution from genetic drift is given by

$$|\Delta f_{\rm drift}| \sim \sqrt{\frac{f\Delta t^*}{N}} \sim \sqrt{\frac{f}{N|s|}} \tag{7.55}$$

This will be small compared to $|\Delta f_{\rm sel}| \sim f$ if

$$f \gg \frac{1}{Ns} \tag{7.56}$$

This is equivalent to the selection-dominated regime we derived in the previous section [up to an $\mathcal{O}(1)$ constant].

We previously showed that the dynamics in this regime are well-approximated by the deterministic version of Eq. (7.48), which yields the exponential growth law $f(t) \approx f_0 e^{st}$. We can also obtain this result using the heuristic approach described here. After a single reset time ($\Delta t^* \sim 1/s$), the new initial frequency is given by

$$f(\Delta t^*) = f(0) + \Delta f \approx (1+c_1) f(0)\,, \tag{7.57}$$

where c_1 is an $\mathcal{O}(1)$ constant. After k such reset times, we obtain

$$f(t) \approx f(0)\left(1+c_1\right)^k = f(0)e^{c_2 st}\,, \tag{7.58}$$

where we have used the total time $t = k \cdot \Delta t^* {\sim} k/s$ to substitute for k. Thus, up to an $\mathcal{O}(1)$ constant in front of s, we can recover the exponential growth law above.

Case 2. If stochastic forces are dominant ($|\Delta f_{\text{drift}}| \gg |\Delta f_{\text{sel}}|$), then the reset time is instead defined by

$$f \sim |\Delta f_{\text{drift}}| \sim \sqrt{\frac{f \Delta t^*}{N}} \, . \tag{7.59}$$

which is equivalent to the characteristic ***drift timescale***,

$$\Delta t^* \sim T_{\text{drift}} \equiv N f \, . \tag{7.60}$$

During this time, the contribution from selection is given by

$$|\Delta f_{\text{sel}}| \sim s f \cdot \Delta t^* \sim N s f^2 \tag{7.61}$$

which will be small compared to $|\Delta f_{\text{drift}}| \sim f$ if

$$f \ll \frac{1}{Ns} \, . \tag{7.62}$$

This is equivalent to the drift-dominated regime we idenfified above.

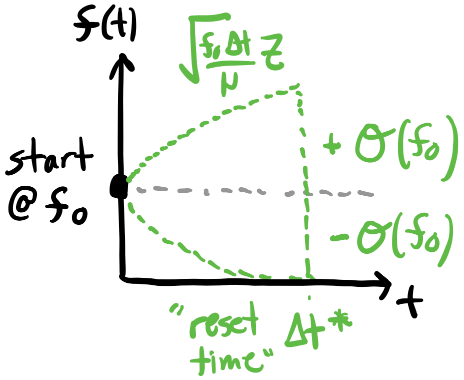

As above, we could use our exact results from Section 7.2 to predict the dynamics of $f(t)$ in this regime. We can also understand this behavior heuristically by modifying our iterative calculation in Eq. (7.58). After a single reset time

($\Delta t^* \sim Nf$), the mutant will have had an appreciable chance of going extinct (since $|\Delta f_{\text{drift}}| \sim f$). Let's call this probability $p_{\text{ext}} \approx e^{-c_1}$, though the precise value of c will not be that important. If the mutant didn't go extinct, it must have drifted to a slightly higher frequency

$$f(\Delta t^*) \sim \frac{f}{1 - e^{-c_1}} \tag{7.63}$$

so that

$$\langle f(\Delta t^*) \rangle = \underbrace{0 \cdot e^{-c_1}}_{\text{extinct paths}} + \underbrace{\left(\frac{f}{1 - e^{-c_1}} \right) \cdot \left(1 - e^{-c_1} \right)}_{\text{nonextinct paths}} = f \tag{7.64}$$

After k successive reset times, the probability that the mutation is still alive is

$$p_s(k) \approx \left(1 - e^{-c_1}\right)^k = e^{-c_2 \cdot k} \tag{7.65a}$$

and its typical size is

$$f(k) \approx f_0 \left(\frac{1}{1 - e^{-c_1}} \right)^k = f_0 e^{c_2 \cdot k} \tag{7.65b}$$

where $c_2 = \log\left(1 - e^{-c_1}\right)^{-1}$ is another $\mathcal{O}(1)$ constant.

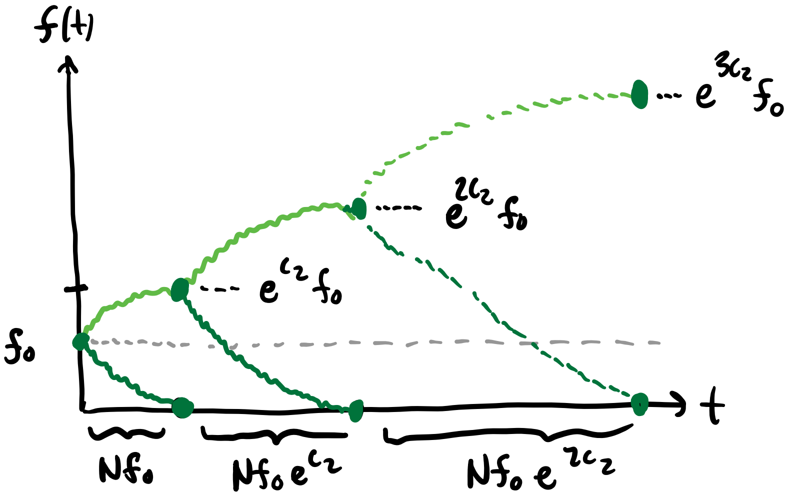

The exponential dependence on k is similar to what we observed in the selection-dominated case above. The main difference is that in the drift-dominated regime, the reset time gets progressively longer in each iteration, since $\Delta t^*(k) \sim Nf(k)$. The total time that elapses during k iterations is therefore given by

$$t(k) = \sum_{j=0}^{k-1} N f_0 e^{c_2 \cdot j} \sim c_3 N f(k) \tag{7.65c}$$

where $c_3 = (1 - e^{-c_2 \cdot k})/(1 - e^{-c_2})$ is another $\mathcal{O}(1)$ constant.

This relationship between $p(k)$, $f(k)$, and $t(k)$ can be summarized in two different ways. If we treat the total time $t = t(k)$ as our input parameter, we would say that:

Drift-dominated dynamics (time)

A mutation will survive for at least t *generations with probability* $\sim Nf_0/t$, *and will reach a typical size* $\sim t/N$.

Alternatively if we treat the final frequency $f = f(k)$ as the input parameter, we would say that

Drift-dominated dynamics (frequency)

With probability $\sim f_0/f$, *a mutation will drift to size* $\sim f$, *and will require* $\sim Nf$ *generations to do so.*

Patching everything together

These two regimes provide a simple picture for the dynamics of Eq. (7.48) when the frequencies are either much larger or much smaller than the ***drift threshold***, $f^* \equiv 1/Ns$. Our heuristic approach goes one step further, and assumes that the boundary between these two regimes is infinitely sharp, so that we can patch the solutions together at f^*. This is not exactly right — our analysis above showed that there is a finite crossover region $[\log f^* \pm \mathcal{O}(1)]$ where drift and selection both play a role. However, the errors involved in making this assumption will only introduce some additional $\mathcal{O}(1)$ prefactors, and will therefore be consistent with the approximations we have been making so far. By patching the drift-dominated and selection-dominated solutions together[4], we can obtain a simple picture for lifecycle of a new mutation:

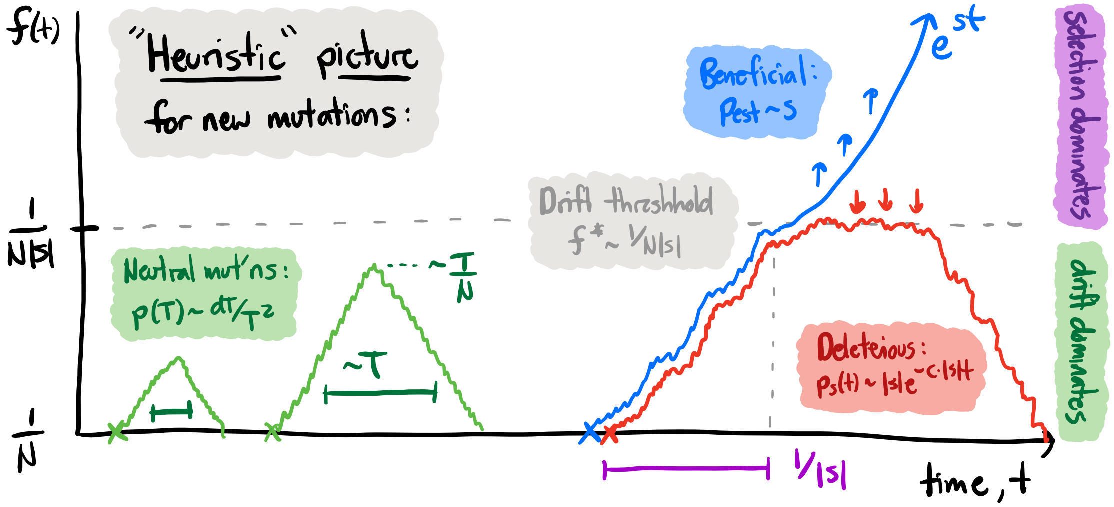

- A new beneficial mutation will start with an initial frequency $f_0{=}1/N$. With probability $\sim f_0/f^* {\approx} s$, its lineage will survive long enough to drift to size $f^* {\sim} 1/Ns$, and will take $\sim 1/s$ generations to do so. At this point, natural selection takes over, and the lineage will begin to grow exponentially at rate s.

- A deleterious mutation will behave similarly at first, and will have a probability $\sim |s|$ of drifting to size $f^* {\sim} 1/N|s|$. After this point, natural selection will prevent the mutation from rising much higher than $\sim 1/N|s|$,

---

[4] This is reminiscent of the asymptotic matching procedure from Section 7.3, but with a subtle and important twist: asymptotic matching is performed at a point where *both* approximations are valid, while our patching argument is gluing together solutions at a point where *neither* approximation is fully accurate (even asymptotically).

so it will continue to have an $\mathcal{O}(1)$ chance of drifting to extinction in each successive reset time ($\Delta t^* \sim 1/|s|$). The total probability of surviving for $t \gg 1/|s|$ generations therefore given by $p_s(t) \sim |s|e^{-c \cdot s \cdot t}$.

- The trajectories of neutral mutations (and beneficial or deleterious mutations below the drift threshold) will look like "triangles" with a height $\sim t/N$ and a width of t, and a probability of occurrence $\sim dt/t^2$.

7.5 Incorporating spontaneous mutations

So far, we have only considered cases where a mutation starts at a given initial frequency, and then we try to predict where it ends up at some later time t. In practice, we are often more interested in scenarios where there are no mutant individuals to begin with (e.g. a lab evolution experiment starting from a single ancestor strain), while mutation events randomly occur within the population over time:

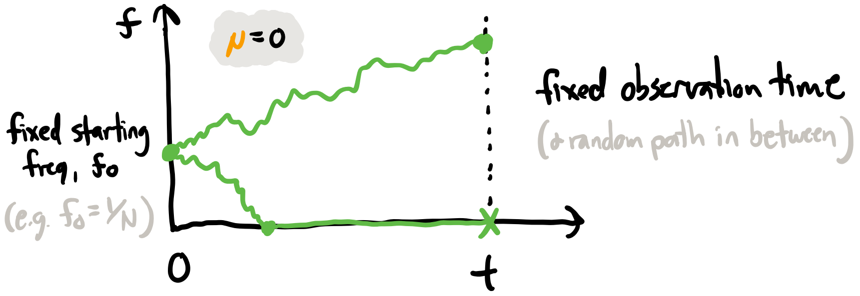

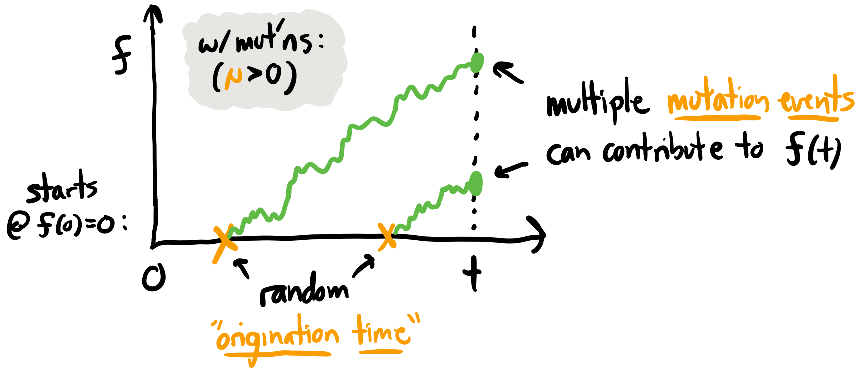

Spontaneous mutations introduce two new complications:

1. The ***origination time*** of the mutation (t_0) is now a random variable, rather than a fixed number as we considered above.

2. Multiple distinct mutation events can contribute to the mutant frequency at some later time t.

We can model these effects by adding a mutation term[5] to our linear branching process model,

$$\frac{\partial f}{\partial t} = \underbrace{\mu}_{\text{mut'n}} + \underbrace{sf}_{\text{selection}} + \underbrace{\sqrt{\frac{f}{N}} \cdot \eta(t)}_{\text{genetic drift}} \tag{7.66}$$

---

[5] In principle, we could add a back mutation term as well, but this is equivalent to shifting the selection coefficient $s \to s - \nu$.

which can introduce new mutant individuals even when $f(t) = 0$. Alternatively, we can also decompose the mutant frequency into a sum over individual mutation events,

$$f(t) = \sum_{t_0=1}^{t} \sum_{i=0}^{\theta(t_0)} f_i(t|t_0)\,, \tag{7.67a}$$

where $\theta(t_0)$ is a random variable representing the number of new mutations that occur in generation t_0,

$$\theta(t_0) \overset{i.i.d.}{\sim} \mathrm{Poisson}(N\mu)\,, \tag{7.67b}$$

and the $f_i(t|t_0)$ are independent mutation-free branching processes,

$$\frac{\partial f_i}{\partial t} = s f_i + \sqrt{\frac{f_i}{N}} \cdot \eta_i(t)\,, \tag{7.67c}$$

with the initial condition $f_i(t_0) = 1/N$. Both of these formulations will be useful for our analysis below.

Since Eqs. (7.66) and (7.67) are fundamentally linear models, we can continue to use the method of characteristics to solve for the corresponding generation function, $H(z,t)$. The details are listed in the Appendix at the end of this chapter. Inverting this solution, we find that the total frequency $f(t)$ follows a Gamma distribution,

Dynamic mutation-selection-drift balance ($f \ll 1$)

If the population is initially composed of wildtype individuals, then the frequency of the mutant at later times follows a Gamma distribution,

$$p(f,t) \propto f^{\alpha-1} e^{-f/f_s(t)}\,, \tag{7.68}$$

with shape parameter $\alpha = 2N\mu$ and scale parameter $f_s(t) = \frac{e^{st}-1}{2Ns}$.

This is a non-equilibrium version of the ***mutation-selection-drift balance*** from Chapter 6.

What does this distribution look like? From the properties of the Gamma distribution, we know that the mean and coefficient of variation are given by

$$\langle f(t) \rangle = \alpha f_s(t) = \frac{\mu}{s}\left(e^{st} - 1\right) \tag{7.69}$$

and

$$c_V^2(t) = \frac{1}{\alpha} = \frac{1}{2N\mu} \tag{7.70}$$

This suggests that we will again observe a transition depending on the magnitude of the scaled mutation rate $N\mu$:

Case 1 ("fuzzy noise"). When $N\mu \gg 1$, the distribution of $f(t)$ will again be strongly peaked around the deterministic expectation $\langle f(t) \rangle$:

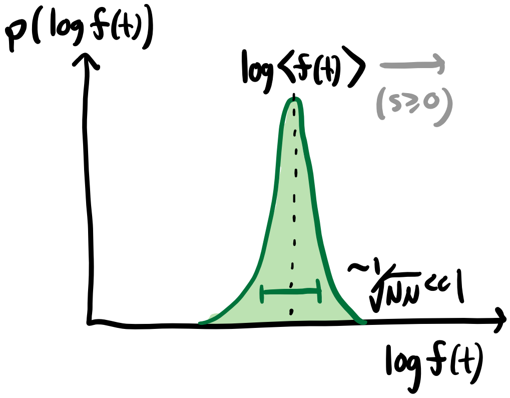

The emergence of this deterministic behavior can be rationalized by the decomposition in Eq. (7.67). We can rewrite this as

$$f(t) = \sum_{t_0=1}^{t} \theta(t_0) \cdot \left[\frac{1}{\theta(t_0)} \sum_{i=0}^{\theta(t_0)} f_i(t|t_0) \right] \tag{7.71}$$

When $\langle\theta(t_0)\rangle = N\mu \gg 1$, many new mutations will be introduced each generation. The law of large numbers then implies that

$$\frac{1}{\theta(t_0)}\sum_{i=0}^{\theta(t_0)} f_i(t|t_0) \approx \langle f_i(t|t_0)\rangle = \begin{cases} \frac{1}{N}\cdot e^{s(t-t_0)} & \text{if}\, t > t_0, \\ 0 & \text{else,} \end{cases} \tag{7.72}$$

so that

$$\langle f(t)\rangle \approx \sum_{t_0=1}^{t} N\mu \cdot \frac{1}{N} e^{s(t-t_0)} \approx \int_0^t \frac{\mu}{s} e^{s(t-t_0)}\, dt_0 = \frac{\mu}{s}\left(e^{st}-1\right) . \tag{7.73}$$

When $s \geq 0$, this average frequency will continue to increase with time. However, for deleterious mutations, it will eventually saturate at a constant value,

$$\langle f(t)\rangle = \frac{\mu}{|s|}\left(1-e^{-|s|t}\right) \longrightarrow \overline{f} \equiv \frac{\mu}{|s|}\,, \tag{7.74}$$

when $t \gg 1/|s|$, which matches the deterministic mutation-selection balance from Chapter 6. This shows that this equilibrium distribution is attained on a timescale $\sim 1/|s|$ when $|s| \gg \mu \gg 1/N$, and that the bulk of the contribution comes from mutations that occurred in the last $\sim 1/|s|$ generations.

We can also understand the fluctuations around this steady state by defining the shifted variable, $\Delta f(t) = f(t) - \overline{f}$, which satisfies,

$$\frac{\partial \Delta f(t)}{\partial t} = \mu - |s|\left(\overline{f} + \Delta f(t)\right) + \sqrt{\frac{\overline{f}+\Delta f(t)}{N}}\cdot \eta(t)\,. \tag{7.75}$$

When the fluctuations are small ($\Delta f(t) \ll \overline{f}$), we can Taylor expand this equation to obtain the simpler SDE,

$$\frac{\partial \Delta f}{\partial t} = -|s|\cdot \Delta f(t) + \sqrt{\frac{\mu}{N|s|}}\cdot \eta(t)\,, \tag{7.76}$$

which has the same functional form as the Brownian particle in a quadratic potential from Chapter 6:

$$x(t) \leftrightarrow \Delta f(t)\,, \quad r \leftrightarrow |s|\,, \quad 2D \leftrightarrow \frac{\mu}{N|s|} \tag{7.77}$$

Using our results from Section 6.1, we can conclude that the fluctuations in $\Delta f(t)$ will reach a typical magnitude $\sim \sqrt{D/r} = \overline{f} \cdot \sqrt{\frac{1}{2N\mu}}$ with a typical turnover time of $1/r = 1/|s|$.

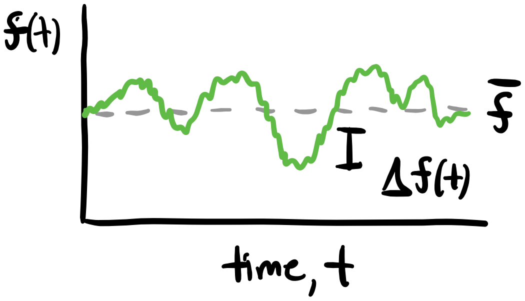

This constitutes a dynamical version of our "Case 1" behavior, where $f(t)$ remains close to its deterministic expectation, with a small amount of fuzziness due to noise.

Case 2 ("jagged noise"). In the opposite case, where $N\mu \ll 1$, the distribution of $f(t)$ transitions to an ***L-shaped form*** that is qualitatively different than the one above. This new shape is characterized by a $1/f$ decay at low frequencies and an exponential cutoff $f_s(t)$:

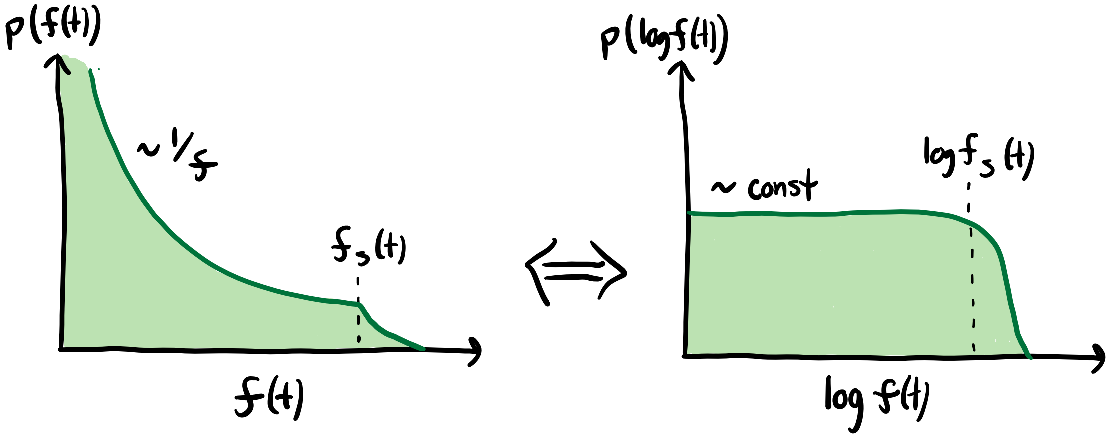

This resembles a truncated form of the U-shaped stationary distribution from Chapter 6. In this case, we can gain some additional intuition for the L-shaped "shoulder" by replotting the distribution in log space. The Jacobian factor transforms the $1/f$ decay in linear space to a nearly constant slope in log space. This implies that there is a roughly equal probability ($\sim N\mu$) of observing the mutant at a frequency with any order of magnitude up to $\log f_s(t)$.

We can gain more intuition for this behavior by considering different values of s. For example, for a neutral mutation ($s = 0$), the scale parameter reduces to $f_s(t) \approx t/2N$, and the frequency distribution becomes:

Dynamic mutation-drift balance ($N\mu \ll 1$)

$$p(f,t) \approx \frac{2N\mu}{f} e^{-2Nf/t}\,, \tag{7.78}$$

We can think about this distribution as modeling the contribution from at most *one* random mutation event:

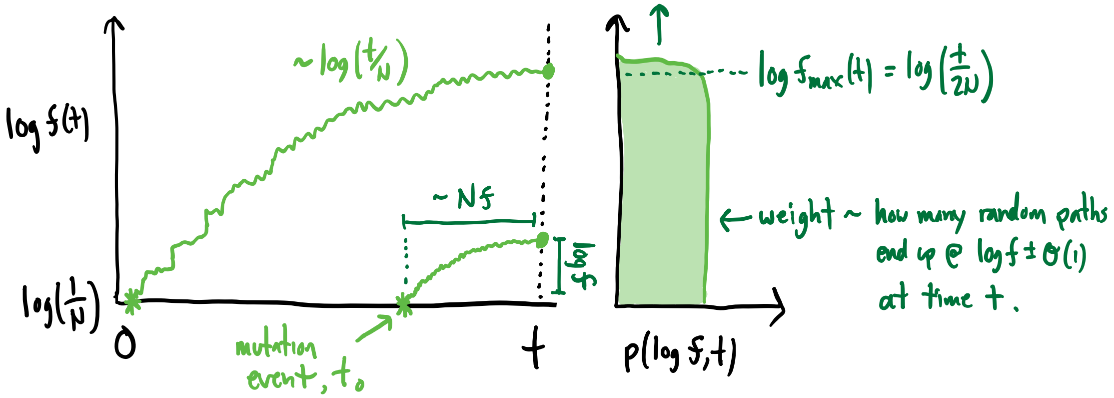

Note that the scale parameter $f_s(t) = t/2N$ corresponds to the average size of a neutral mutation that arose at $t = 0$, and survived until the present day. Thus, the exponential cutoff at $f_s(t)$ corresponds to the size that a mutation would have typically reached if it had occurred at the earliest possible time (i.e., the maximum typical size).

Similarly, the weight of the distribution at frequencies below $f_s(t)$ can be understood in terms of the number of random paths that end up in the frequency range $\log f \pm \mathcal{O}(1)$ by time t. Since the surviving paths increase linearly with time ($\sim t/N$), we can infer that the mutations that contribute to the frequency range $\log f \pm \mathcal{O}(1)$ must have originated $t_0 \sim Nf \pm \mathcal{O}(Nf)$ generations ago. These mutations are produced at rate $N\mu$ per generation, each one has a probability $\sim 1/Nf$ of drifting to its current frequency. Putting everything together, we can obtain a heuristic formula for the frequency distribution,

$$p(\log f, t) \cdot \underbrace{\Delta \log f}_{\mathcal{O}(1)} \sim \underbrace{N\mu}_{\substack{p(\text{mutation}) \\ \text{per gen.}}} \times \underbrace{\Delta t_0}_{\substack{\text{range of} \\ \text{origination} \\ \text{times}}} \times \underbrace{p\left(\tfrac{1}{N} \to f\right)}_{\substack{\text{probability of} \\ \text{drifting to } f}} \sim N\mu \times Nf \times \frac{1}{Nf} \sim N\mu$$

which yields

$$p(\log f, t) \sim N\mu \implies p(f,t) \sim \frac{N\mu}{f}, \quad (f \ll t/N) \tag{7.79}$$

as desired.

What about selected mutations? For sufficiently short times ($t \ll 1/|s|$), the scale parameter remains close to $f_s(t) \approx t/2N$, so the frequency distribution will be indistinguishable from the transient mutation-drift balance in Eq. (7.78). At longer times, the dynamics of $p(f,t)$ will strongly depend on the sign of s.

Deleterious mutations. When $s < 0$, the exponential cutoff eventually saturates at $f_s(t) \to f_{\max} \equiv 1/2N|s|$, so that the frequency distribution reduces to

$$p(f,t) \approx \frac{2N\mu}{f} e^{-2N|s|f}, \tag{7.80}$$

when $t \gg 1/|s|$. This expression matches the mutation-selection-drift balance we derived in Chapter 6 in the limit that $N\mu \ll 1$ and $N|s| \gg 1$. The intuition for this distribution is similar to the neutral case above, with at most one previous mutation event contributing the present day mutant frequency.

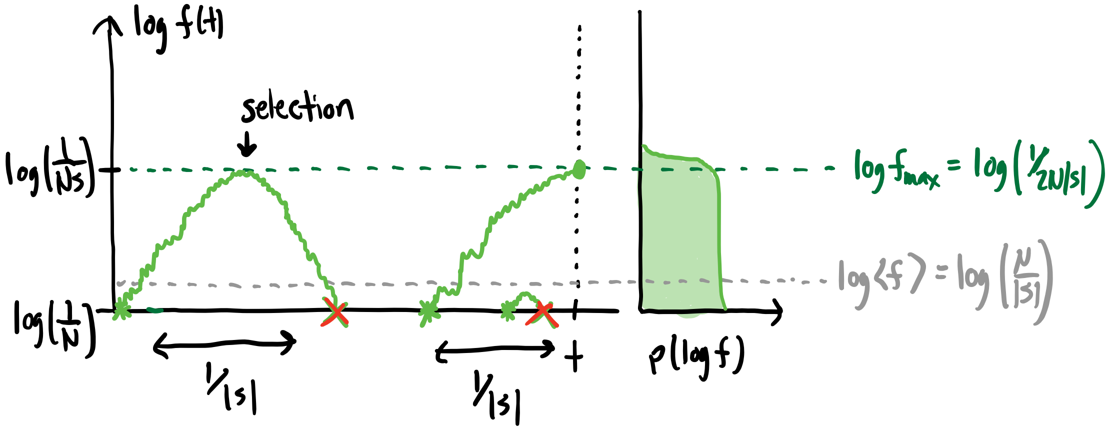

The main difference is that deleterious mutations can only reach a maximum typical size of $f_{\max} \sim 1/N|s|$ before they go extinct, so the frequency distribution will primarily reflect mutations that occurred within the last $\sim 1/|s|$ generations. This explains why the deleterious distribution resembles the neutral version evaluated at an effective equilibration time $t_{\text{eff}} = 1/|s|$.

Note that the *typical* frequencies in this distribution can be very different from the average value, $\langle f \rangle = \mu/|s|$. For example, an antibiotic resistance mutation with a 1% fitness cost in the absence of drugs, and a mutation rate of $\mu \sim 10^{-10}$, will be represented by an average of $N \cdot \mu/|s| = 0.01$ cells in a population of size $N = 10^6$. However, Eq. (7.80) shows that the mutant can reach a size as large as $\sim N \cdot 1/N|s| = 100$ cells by stochastically drifting to higher frequencies.

Beneficial mutations. When $s > 0$, the frequency distribution retains the same characteristic shape, but with a maximum size, $f_s(t) \approx e^{st}/2Ns$, that now increases exponentially with time:

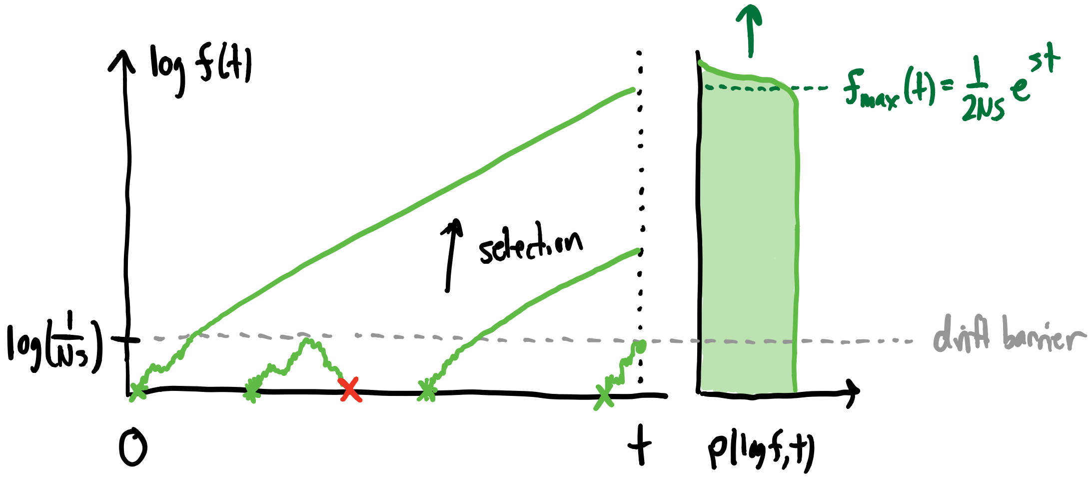

The distribution is still very broad, but in contrast to the last two cases, the contributing paths now grow largely deterministically once $f(t) \gg 1/Ns$. Motivated by the behavior we observed in the $\mu = 0$ case, we can once again try to capture this randomness in a single random number $\nu(t)$, by factoring out the expected time-dependence for a single established mutation:

$$f(t) \equiv \frac{\nu(t)}{2Ns} e^{st} \,, \tag{7.81}$$

Using the properties of the Gamma distribution, we can conclude that $\nu(t)$ will approach the time-independent distribution,

$$p(\nu, t) \propto \nu^{2N\mu - 1} e^{-\nu} \,, \tag{7.82}$$

when $t \gg 1/s$.

In this case, we will also find it useful to rewrite ν as an effective time variable τ_{est}:

$$f(t) = \frac{\nu}{2Ns} e^{st} \equiv \frac{1}{2Ns} e^{s(t - \tau_{\text{est}})} \implies \tau_{\text{est}} = \frac{1}{s} \log\left(\frac{1}{\nu}\right) \,, \tag{7.83}$$

which is sometimes known as the ***establishment time***. Intuitively, it corresponds to the time that the first successful mutation was able to escape genetic drift. Strictly speaking, the definition in Eq. (7.83) refers to the time that $f(t)$ *would have* reached frequency $1/2Ns$ if it was growing deterministically the entire time.

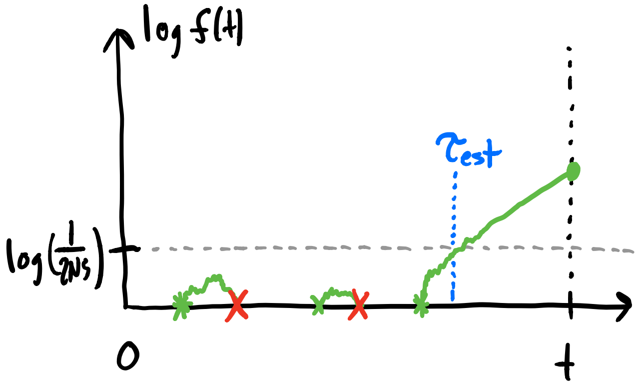

Using our result for ν in Eq. (7.82), we find that the distribution of τ_{est} is approximately exponentially distributed:

Establishment time for a new beneficial mutation ($N\mu \ll 1$)

$$\tau_{\mathrm{est}} \sim \mathrm{Exponential}\left(\frac{1}{2N\mu s}\right) \pm \mathcal{O}\left(\frac{1}{s}\right) \tag{7.84}$$

The first term in this expression reflects the randomness in the time it takes to produce the first successful mutation, while the second term represents the randomness in its trajectory while it drifts to a frequency $\gg 1/Ns$.

Both terms have a simple interpretation in terms of our heuristic picture above. The population produces new beneficial mutations at rate $N\mu$ per generation. Each of these mutations has a probability $\sim s$ of surviving genetic drift and establishing. This implies that successful beneficial mutations occur as a Poisson process with rate $\sim N\mu s$. A successful mutation will require $\sim 1/s$ generations to drift to frequency $\sim 1/Ns$, and will start to increase exponentially thereafter.

By substituting Eq. (7.83) into our asymptotic matching procedure in Section 7.3, we can obtain an analogous prediction for the total time it takes for the mutant to first reach fixation:

$$T_{\mathrm{fix}} = \mathrm{Exponential}\left(\frac{1}{2N\mu s}\right) + \frac{2}{s}\log(2Ns) \tag{7.85}$$

As above, the first term represents the time it takes for a successful mutation to first appear, while the second term represents the time required for it to sweep through the population once it arises.

When $N\mu \to 0$, the waiting time to produce a successful mutation will eventually grow much longer than the sweep time, so that $T_{\mathrm{fix}} \sim 1/N\mu s$. In this case, we would say that *"evolution is limited by the supply of new mutations"*, since increasing N or μ will decrease the waiting time by the same factor.

We can compare this behavior with the $N\mu \gg 1$ case, where an analogous calculation shows that

$$T_{\text{fix}} \approx \frac{2}{s} \log\left(\frac{s}{\mu}\right) \tag{7.86}$$

In this case, the waiting time for the mutation to reach $f(t) \approx 100\%$ is most strongly influenced by s, while changes in N or μ have a very weak effect. This behavior is driven by the fact that multiple beneficial mutations establish within the first $\sim 1/s$ generations after $t = 0$, so increasing or decreasing this number has only a small effect. We will revisit this phenomenon when we consider longer genomes later in the course.

Long-term approach to the stationary distribution

Finally, we can use our results for the transient distribution in Eq. (7.68) to answer a question we posed back in Chapter 6: how long does it take to reach the stationary distribution in Eq. (6.32)? Is it ever relevant in practice?

Our analysis above provides an answer for beneficial or deleterious mutations ($N|s| \gg 1$), so we will focus on the neutral case here. This is relevant, for example, if we are looking at synonymous mutations in the human genome, since deviations between synonymous and nonsynonymous frequencies can provide a powerful test for selection (see Chapter **??**).

If the population is initially composed of wildtype individuals, then the initial dynamics will follow the transient mutation-drift balance in Eq. (7.78), with the maximum possible frequency $f_s(t)$ — corresponding to a mutation arising at $t = 0$ — increasing toward higher frequencies at rate t/N. This gives rise to the left shoulder of the U-shaped distribution in Eq. (6.32):

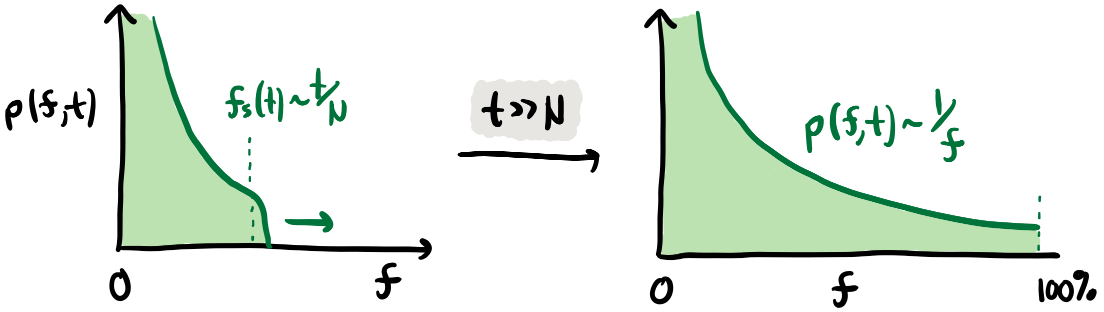

When $t \sim \mathcal{O}(N)$, the maximum frequency will approach 100%, and there is a chance that the mutant will have reached fixation. At this point, back mutations can create new wildtype lineages, which can fluctuate to higher frequencies through genetic drift. This ultimately creates the right-side of the U shape.

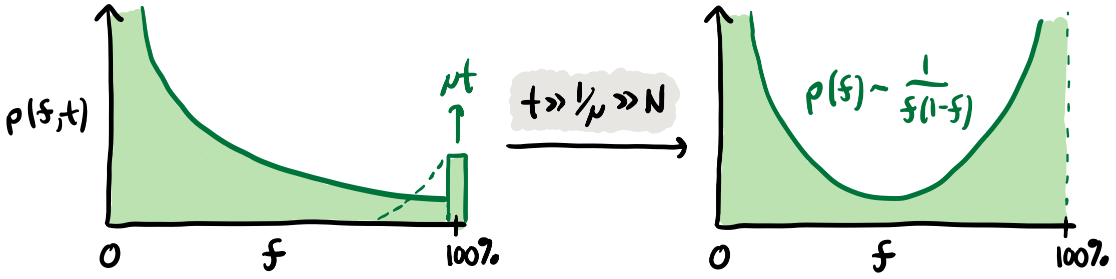

However, in contrast to forward mutations, which are introduced at rate $N\mu$, back mutations can only be produced in populations where the mutant allele has previously fixed. These accumulate at rate $N\mu \cdot \frac{1}{N} = \mu$, which implies that $\sim 1/\mu$ generations are required to reach the full U-shape.

This relaxation timescale can be extremely long. For example, in the human genome where $\mu \sim 10^{-8}$, the equilibration time for a synonymous site is 10^8 generations or ~2 billion years — far longer than humans have existed as a species. This means that the stationary distribution in Eq. (6.32) will not apply.

Instead, we will often be interested in scenarios where t is large compared to N but still small compared to $1/\mu$, so that only the left-half of the U is present:

Quasi-stationary distribution ($s = 0$, $N \ll t \ll 1/\mu$)

$$p(f,t) \approx \frac{2N\mu}{f} \tag{7.87}$$

While we have arrived at this formula by extrapolating our branching process results above (which technically required $f \ll 1$), we can also derive it more rigorously from the full diffusion model in Chapter 6.[6] We will also encounter this distribution from a genealogical perspective when we discuss ***neutral theory and the coalescent*** later in the course.

7.6 Appendix

7.6.1 Exact solution using the method of characteristics

In this section, we show how to solve the partial differential equation for the generating function of the linear branching process using the ***method of characteristics***.

No mutations ($\mu = \nu = 0$)

We will start by considering the case without mutations ($\mu = \nu = 0$), where the mutant starts at an initial frequency $f(0) = f_0$. The generating function satisfies the PDE in Eq. (7.13),

$$\frac{\partial H}{\partial t} = \left[sz - \frac{z^2}{2N}\right]\frac{\partial H}{\partial t}\,, \tag{7.88}$$

subject to the initial condition $H(z,0) = e^{-zf_0}$.

---

[6] See W J Ewens, “The Diffusion Equation and a Pseudo-Distribution in Genetics, ” *J Roy Stat Soc B* 1963 for more details.

The method of characteristics is a generalization of the trick that we used to solve for the fixation probability of the full single-locus model in Chapter 6. Recall that in that case, we found a special value of $z^* = 2Ns$ for which $\partial_t H(z^*, t) = 0$. This allowed us to relate the values of $H(z^*, t)$ at long times (where $f = 0, 1$) with the initial value $H(z^*, 0)e^{-z^* f}$.

We can generalize this idea by searching for a ***family of curves***, $z^*(t)$, along which

$$\frac{d}{dt}\left[H(z^*(t), t)\right] = 0\,. \tag{7.89}$$

When this condition is satisfied, we can again relate the values of $H(z, t)$ between the initial timepoint and any later time,

$$H(z^*(t), t) = H(z^*(0), 0) = e^{-z^*(0) f_0} \tag{7.90}$$

The line $z^*(t) = 2Ns$ is one such ***characteristic curve***, but there are infinitely many others. Using the chain rule on Eq. (7.89), we can write the total derivative as

$$\frac{dH(z^*(t), t)}{dt} = \frac{\partial H}{\partial t} + \frac{\partial H}{\partial z}\frac{dz^*}{dt} = \frac{\partial H}{\partial z}\left[sz^* - \frac{z^{*2}}{2N} + \frac{dz^*}{dt}\right]\,, \tag{7.91}$$

where we have used the equation of motion in Eq. (7.88) to replace $\partial H/\partial t$. This shows that if $z^*(t)$ satisfies the first order ODE,

$$\frac{dz^*}{dt} = -sz^* + \frac{z^{*2}}{2N} \tag{7.92}$$

then Eq. (7.90) will be satisfied. We can visualize this the following diagram:

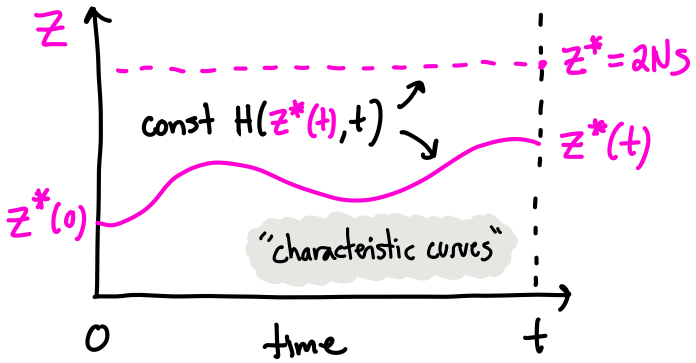

The curve $z^*(t) = 2Ns$ is one possible solution Eq. (7.92) corresponding to the initial condition $z^*(0) = 2Ns$. However, Eq. (7.90) shows that this only allows us to evaluate the generating function at a special value of $z = z^*$. To obtain the full generating function $H(z,t)$, we want to be able to choose the value of z that we will use to evaluate $H(z,t)$ in the present. In other words, we need to find the initial value $z^*(0)$ that produces a characteristic curve with $z^*(t) = z$.

This is easiest to accomplish by defining a corresponding curve in ***reverse time*** (i.e. working back from the final time t). In particular, if we define a function,

$$\phi(t') = z^*(t-t') \tag{7.93}$$

then $\phi(t')$ must satisfy the initial value problem

$$\frac{\partial \phi}{\partial t'} = s\phi - \frac{\phi^2}{2N}\,, \tag{7.94}$$

with $\phi(0) = z$, and the generating function is given by

$$H(z,t) = e^{-\phi(t) f_0}\,. \tag{7.95}$$

In this case, the solution to Eq. (7.94) is a simple logistic function,

$$\phi(t) = \frac{z e^{st}}{1 + \frac{z}{2Ns}\left(e^{st}-1\right)}\,, \tag{7.96}$$

so the generating function is given by

$$H(z,t) = \exp\left[\frac{-z f_0 e^{st}}{1 + \frac{z}{2Ns}\left(e^{st}-1\right)}\right] . \tag{7.97}$$

Incorporating spontaneous mutations

We can use a similar approach to account for spontaneous mutations. Repeating our derivation for the $\mu = 0$ case, we can show that the generating function for Eq. (7.66) must satisfy

$$\frac{\partial H}{\partial t} = \left[sz - \frac{z^2}{2N}\right]\frac{\partial H}{\partial t} - \mu z H \, , \tag{7.98}$$

which includes a new term, $\mu \cdot z \cdot H(z,t)$, when $\mu > 0$. In this case, we will often be interested in scenarios where the initial frequency is zero, so that $H(z,0) = e^{-z\cdot 0} = 1$.

PDEs of this form can be solved using a slight variation of method-of-characteristics approach above. In particular, if we define a new function,

$$\psi(t') = \log H(\phi(t'), t - t') \tag{7.99}$$

where $\phi(t')$ is the characteristic curve in Eq. (7.94), then $\psi(t')$ must satisfy the differential equation,

$$\frac{d\psi}{dt'} = \mu \cdot \phi(t') \, , \tag{7.100a}$$

with the initial and final conditions

$$\begin{aligned} \psi(0) &= \log H(\phi(0), t) = \log H(z,t) \\ \psi(t) &= \log H(\phi(t), 0) = 0 \end{aligned} \tag{7.100b}$$

Integrating from $t' = 0$ to t then yields

$$\psi(t) = \psi(0) + \int_0^t \mu\phi(t')dt' = \psi(0) + 2N\mu \log\left[1 + \frac{z}{2Ns}\left(e^{st} - 1\right)\right] \tag{7.101}$$

or

$$H(z,t) = \left(1 + z \cdot \frac{e^{st} - 1}{2Ns}\right)^{-2N\mu} . \tag{7.102}$$

Using the ***"method of Wikipedia"***, we can recognize this as the generating function of a Gamma distribution with shape parameter $\alpha = 2N\mu$ and scale parameter $f_s(t) = \frac{e^{st}-1}{2Ns}$:

$$p(f,t) = \frac{f_s(t)^\alpha}{\Gamma(\alpha)} f^{\alpha-1} e^{-f/f_s(t)} . \tag{7.103}$$