

Table of Contents

Note 1 — Quantitative descriptions of temporal scaling	2
1.1 Definitions and terminology	2
1.2 Temporal scaling	3
1.3 Accelerated Failure Time vs Proportional Hazards	4
1.4 Scaling and parametric models	6
1.4.1 On the parametrization of the Gompertz hazard	6
1.4.2 The parametric form of wild-type <i>C. elegans</i> hazard	6
Note 2 — Quantifying temporal scaling.....	8
2.1 Adapting the Kolmogorov-Smirnov two-sample test	8
2.2 Evaluating the performance of the generalized KS test	9
Note 3 — Heterogeneity and deviations from temporal scaling	13
3.1 Background: Weibull heterogeneity (frailty)	13
3.2 Quantifying heterogeneity	14
Note 4 — The effect of transient perturbations	17
4.1 Intuition	17
4.2 Definitions	18
4.3 No switching	19
4.4 Switching	20
4.5 Predictions	20
4.6 The effect of switching on distributions	20
Note 5 — Temporal scaling in phenomenological causal models.....	22
5.1 Competing risk models	23
5.2 Interdependency networks	25
5.2.1 Homogeneous interventions	26
5.2.2 Heterogeneous interventions	27
5.3 Conclusions	30
Note 6 — Vitality/Resilience models.....	30
6.1 Semi-infinite random walk	30
6.2 Drift-diffusion process	32
6.3 Statistical power and sensitivity of scaling in drift-diffusion processes	33
6.4 Drift-diffusion process with drift heterogeneity	35
6.5 Strehler-Mildvan with vitality drift-diffusion	35

Note 1 — Quantitative descriptions of temporal scaling

Summary

This Note introduces the demographic terms we use to describe and quantify the phenomenon of temporal scaling in our experimental data. We then relate our perspective to common alternative approaches, specifically the proportional hazards (PH) model and parameterized regression models.

Contents of Note 1

1.1	Definitions and terminology	2
1.2	Temporal scaling	3
1.3	Accelerated Failure Time vs Proportional Hazards	4
1.4	Scaling and parametric models	6
1.4.1	On the parametrization of the Gompertz hazard	6
1.4.2	The parametric form of wild-type <i>C. elegans</i> hazard	6

1.1 Definitions and terminology

The empirical survival function $\hat{S}[t]$ reports the fraction of individuals remaining alive at time t . It is an estimate of the continuous survival function $S(t)$ from which experimental data are sampled. Time (age) t is defined relative to a synchronization event that marks $t = 0$, which in our case is the transition from L4 to adulthood as determined in Methods. In probabilistic language, $S(t)$ is a cumulative distribution; it reports the probability that an individual is still alive at time t . If T is the random variable denoting the time of death, then $S(t) = \text{Prob}(T > t)$. The probability density function, or pdf, $l(t)$ is the probability density that an individual dies between t and $t + dt$, $l(t) = \text{Prob}(t < T < t + dt) = -dS(t)/dt$. This is also referred to as the lifespan distribution. Lifespan distribution and survival function are related through the hazard function $h(t)$, which is the instantaneous rate of death events, or in probabilistic terms, the conditional probability of dying between t and $t + dt$, given survival up to t : $h(t) = \text{Prob}(t < T < t + dt | T > t)$. The basic demographic equation in continuous time for non-reproducing populations synchronized at $t = 0$ is then given by

$$l(t) = -\frac{d}{dt}S(t) = h(t)S(t). \quad (1)$$

The hazard, $h(t) = -d \log S(t)/dt$, can be seen as playing the role of a time-dependent rate factor in the survival kinetics, equation (1). The survival kinetics should not be confused with the aging dynamics, which is some stochastic process that governs survival by determining $h(t)$. In our experiments we directly observe the survival kinetics, but not the aging process. In Fig. 1 of the main text, we estimate $h(t)$ from lifespan measurements and identify a striking regularity in the way $h(t)$ responds to interventions. This, then, allows us to infer properties of the aging dynamics that generates $h(t)$.

We occasionally refer to survival, hazard, and lifespan distribution as “endpoint statistics” because the stochastic aging process ends in an absorbing boundary (death). In this sense, lifespan is the first passage time of an individual process to this absorbing boundary [1].

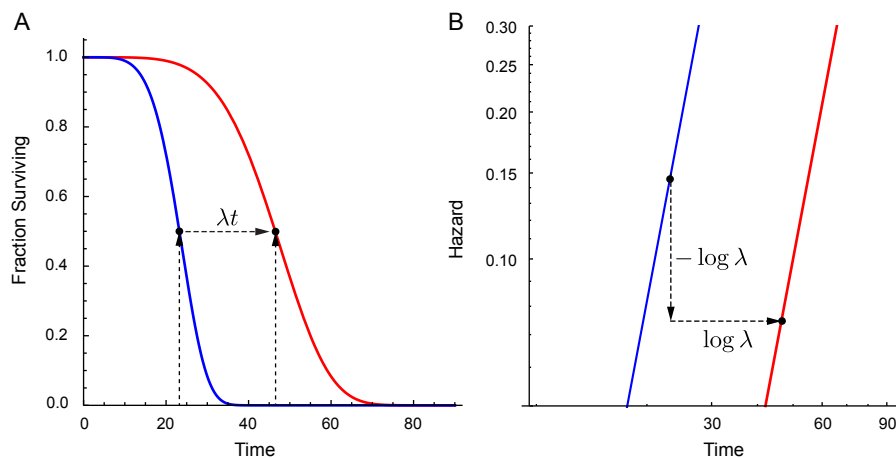
1.2 Temporal scaling

Temporal scaling refers to a situation in which the differences between lifespan distributions corresponding to populations in a reference state **0** and a perturbed state **1** consist in a rescaling of time, such that time t under condition **1** corresponds to time t/λ under condition **0** (Supp. Fig. 1-1), with λ a dimensionless factor:

$$S_1(t) = S_0(\lambda^{-1}t) \quad \text{or} \quad S_1(\lambda t) = S_0(t) \quad (2)$$

For example, if $\lambda = 2$, then 1 hour in system **1** corresponds to 1/2 hour in system **0**. We chose the definition of λ such that a $\lambda > 1$ means that time has slowed down (lifespan is extended by a factor of λ) in the new state and $\lambda < 1$ means that time has accelerated (lifespan is shortened by a factor of λ). Temporal scaling leaves the shape of $l(t)$, $S(t)$, and $h(t)$ invariant. One aspect of shape is the coefficient of variation, defined as the ratio of standard deviation to mean lifespan.

As an illustration, consider the Gompertz survival function for some reference state **0**, $S_0(t) = \exp(-a(\exp(b_0^{-1}t) - 1))$. If an intervention alters the time scale b_0 into b_1 , yielding $S_1(t) = \exp(-a(\exp(b_1^{-1}t) - 1))$, a dimensionless factor $\lambda = b_1/b_0$ rescales the time variable such that equation (2) holds: $\exp(-a(\exp(b_1^{-1}t) - 1)) = \exp(-a(\exp(b_0^{-1}\lambda^{-1}t) - 1))$.



Supplementary Figure 1-1: Scaling of survival and hazard functions (I). We speak of temporal scaling when two functions can be made identical by a dilation or contraction of the time axis. **A:** Each value of the blue survival function at time t , S_0 , corresponds to the value of the red function S_1 at time λt : $S_0(t) = S_1(\lambda t)$. The point singled out in this example is the median and $\lambda = 2$. **B:** In Fig. 1a of the main text, we show the scaling relation not for the survival function but the hazard function instead. Temporal scaling in the hazard means that $h_1(t) = \lambda^{-1}h_0(\lambda^{-1}t)$ (see equation 3 below) or, equivalently, $h_1(\lambda t) = \lambda^{-1}h_0(t)$. In a log-log plot, we first shift a point on $\log h_0$ (blue curve) corresponding to a particular survival quantile, say the median (black dot), vertically by $-\log \lambda$ (this is the $\lambda^{-1}h_0(t)$ part on a log-log scale) and then horizontally along the time axis by $\log \lambda$ to obtain the corresponding point of $\log h_1$ on the red curve (i.e. the $h_1(\lambda t)$ part on a log-log scale). This is the origin of the remark in the main text that any two hazard curves are shifted by an “equal and opposite amount in magnitude and time.” These statements are independent of any parametric form of the survival or hazard function. In this example we used for the purpose of illustration the Weibull distribution with survival function $S(t) = \exp(-(t/\beta)^\alpha)$ and hazard function $h(t) = \alpha/\beta (t/\beta)^{\alpha-1}$, which is linear on a log-log scale.

Using equation (2) in equation (1) yields:

$$-\frac{d}{dt}S_1(t) = h_1(t)S_1(t) \Rightarrow -\frac{d}{dt}S_0(\lambda^{-1}t) = h_1(t)S_0(\lambda^{-1}t)$$

$$\stackrel{\text{set } \tau = \lambda^{-1}t}{\implies} -\frac{d}{d\tau}S_0(\tau) = \lambda h_1(\lambda\tau)S_0(\tau) \Rightarrow h_0(\tau) = \lambda h_1(\lambda\tau). \quad (3)$$

Hence, a rescaling of survival, $S_1(t) = S_0(\lambda^{-1}t)$, is equivalent to a rescaling of the hazard as

$$h_1(t) = \lambda^{-1}h_0(\lambda^{-1}t) \quad (4)$$

and, by a similar argument, of the lifespan density as $l_1(t) = \lambda^{-1}l_0(\lambda^{-1}t)$. An intervention that yields $\lambda = 2$ slows down the original system twofold with respect to survival, which is the same as saying that the new system will exhibit half the hazard rate after twice the time.

1.3 Accelerated Failure Time vs Proportional Hazards

In survival analysis, equation (2) is known as an Accelerated Failure Time (AFT) model, in which the lifespan random variable T of population **0**, living at condition C_0 , is linearly transformed into the random variable λT of population **1**, living at condition C_1 . An AFT regression model with a single categorical covariate can be expressed as:

$$\log t_i = \beta \mathbf{x}_i + \varepsilon_i, \quad (5)$$

where t_i is the observed lifespan of individual i , X an explanatory categorical variable and \mathbf{x}_i a vector encoding this variable for individual i (e.g. whether it was living at C_0 or C_1); β is a vector of effects to be estimated, and ε_i is a residual representing the logarithm of the lifespan of individual i unexplained by the covariate. The relation to the scale factor λ is given by $\lambda = \exp(\beta \mathbf{x}_i)$.

A rescaling of the lifespan random variable T into a new random variable λT alters the expected value of T , $E(T)$, into $\lambda E(T)$ and the variance $\sigma^2(T)$ into $\lambda^2 \sigma^2(T)$. As a consequence, temporal rescaling leaves the coefficient of variation $\sigma(T)/E(T)$ unchanged.

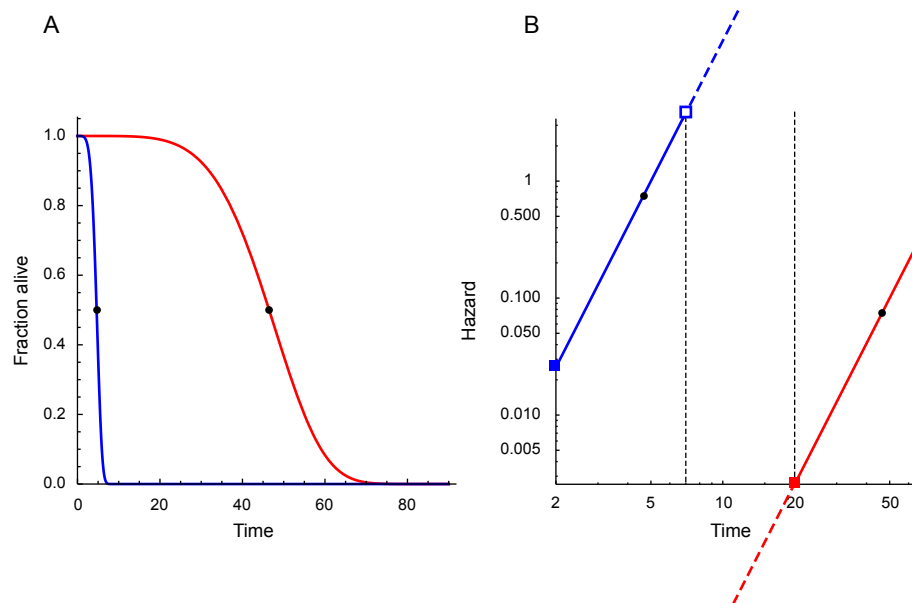
We apply the AFT model to link physical interventions in aging to their quantitative consequences on lifespan. When AFT models provide an excellent fit across sets of lifespan distributions, we conclude that the interventions producing such distributions result in temporal scaling. Temporal scaling suggests that interventions alter lifespan by altering the timescale of an aging process that determines the risk of death $h(t)$.

This approach contrasts with the widely used proportional hazards (PH) framework, which takes the perspective that interventions act directly on the risk of death, such that

$$\log h_i(t) = \beta \mathbf{x}_i + \log h_0(t), \quad (6)$$

in which $h_i(t)$ is expressed in terms of a baseline hazard $h_0(t)$. In the case of two populations ($i = 0, 1$ and assumed to consist of identical individuals), the PH approach thus quantifies the consequences of exposure to condition **1** as a proportionality factor $\exp \beta$ [2].

In the special case that mortality follows a Weibull distribution, AFT and PH are mathematically equivalent descriptions, because a power law is a homogenous function. Yet, even when data are not well fit by a Weibull, PH and AFT models can both be reasonable characterizations of the effect of interventions. Our preference for the AFT model over the PH model is founded on three



Supplementary Figure 1-2: Scaling of survival and hazard functions (II). **A:** The panel shows a synthetic situation (Weibull survival) analogous to Supp. Fig. 1-1a but with a tenfold slow-down ($\lambda = 10$) in time from blue to red. In the blue population everything has already happened before something happens in the red population. For example, in our temperature experiments of Fig. 1 in the main text, the scale factor between highest and lowest temperature is about 40 (!). As in Supp. Fig. 1-1, the black dots indicate the survival median and we used a Weibull distribution for illustration (thus the linear hazard in the log-log plot of panel B). **B:** The panel illustrates the corresponding situation for the hazards. The solid lines (blue and red) depict the hazard trajectory in the time interval starting when 1% (solid square) of the population have died and ending when 99% (open square) have died. In this synthetic case, the mathematical functions continue beyond these intervals (dotted lines) with extremely high or extremely low values, but seen from an empirical standpoint there are no meaningful data outside the solid regime. Vastly increasing the population size would only have a marginal effect on this picture. The red and blue empirical hazard functions have effectively no overlap and can therefore not be quantitatively compared as in equation (6), i.e. one hazard cannot be expressed in terms of the other as the baseline hazard. There is no empirically justified sense in which they could be “proportional”. This is not the case in situations characterized by relatively small effects, such as clinical drug applications, where the PH framework is routinely applied.

observations: (1) We find that our interventions have a permanent effect on lifespan when applied transiently early in life, before any deaths occur (Fig. 3b in the main text). This suggests that interventions act on some process determining T rather than on $h(t)$ directly. (2) In many of our experiments, populations differ so much in their aging timescale that there is no time interval during which both populations have survivors. While this presents no mathematical or computational problem, since $h_1(t) = \beta h_0(t)$ is well defined for all t , it does present an empirical problem, since $h_0(t)$ and $h_1(t)$ cannot be measured for any overlapping t (Supp. Fig. 1-2). (3) The parameter values obtained with the AFT model seem more physically realistic than those obtained with the PH model. For example, if one were to compare a population at 20 °C with one at 30 °C, the fold-increase in relative risk that we extrapolated for the PH model comes in at 4.5×10^5 . It is hard to imagine how a physical intervention could *directly* impact risk ratios

related to aging almost a 1,000,000-fold. It is more plausible that such an increase in risk is the *indirect* consequence of altering 18-fold the timescale of an underlying physical process, as assumed by the AFT model.

1.4 Scaling and parametric models

An often used method for characterizing the effect of interventions in aging is based on simple parametric distributions, such as the Gompertz, Weibull, log-logistic and inverse Gaussian distributions (see Note 1). For example, the Gompertz hazard has been used to classify interventions in *C. elegans* and other model organisms according to their effects on the parameters of the distribution [3–7]. Another function, quite common in the analysis of engineered systems, is the Weibull distribution, see Table 1 below.

1.4.1 On the parametrization of the Gompertz hazard

The wide-spread use of Gompertz warrants a comment on its parametrization. Scaling in the context of Gompertz fits is best understood using a different parametrization than the one typically employed. In the parametrization $h(t) = a \exp(t/b)$ (for example [3, 4]), the parameters a and b are not independent, but a depends inversely on b . If b is interpreted as a time scale parameter (or $1/b$ as a rate parameter), then a change from b to λb requires a change from a to $\lambda^{-1}a$, see equation 4. Another way of seeing this dependence is to calculate the survival function associated with the parametrization in question: $S(t) = \exp(-ab(\exp(t/b) - 1))$. The argument of the exponent should be a dimensionless quantity, which means that the time scale parameter b should only occur in conjunction with the time variable t . For this to be the case, the factor ab must be independent of b , which can be achieved by using a new parameter $a' = ab$, and thus $h(t) = (a'/b) \exp(t/b)$, where a' is now a constant that is independent of b , with associated survival function $S(t) = \exp(-a'(\exp(t/b) - 1))$. This second parameterization is common in some statistical packages and areas of study, but does not seem to be widely used in studies of aging. In the parametrization $h(t) = a \exp(t/b)$, any intervention that acts on b must necessarily affect a , confounding attempts to map effects to distinct parameters. Changes in the Gompertz a —vertical drops in a lin-log plot—should be interpreted only after the dependence on b is accounted for.

1.4.2 The parametric form of wild-type *C. elegans* hazard

Our quantification of deviations from temporal scaling (Note 2) does *not* proceed by fitting empiric data with simple parametric functions. This is a significant methodological advantage. Nevertheless, it is useful to identify a “best” parametrized form for the *C. elegans* hazard function when attempting to quantify hazard heterogeneity in the population (Note 3.2). To this end, we surveyed several parametric hazards derived from the Gompertz, Weibull, log-logistic, log-normal, and inverse Gaussian distributions, as well as Weibull and Gompertz hazards with a frailty term. These functions are listed in Table 1. Model parameters were estimated in R using the `flexsurv` package. Across several replicates containing between 650 and 1,200 wildtype individuals at 20 °C, we confirmed previous observations [8, 9] that neither the Gompertz nor the Weibull hazard without frailty provide good fits for *C. elegans* populations (Extended Data Item 1a). However, in all cases the Weibull distribution performs distinctively better than the Gompertz distribution (Extended Data Item 1b, c). The log-logistic and inverse Gaussian distributions both predict an intrinsic deceleration of the population hazard, and for that reason fit the data better than either the two-parameter Weibull or Gompertz distributions (Extended Data Item 1b–c). We find that incorporating a frailty correction reconciled the Gompertz and Weibull hazards with the experimental data; either frailty model provided excellent fits across all replicates. In all cases, the Weibull model required a smaller frailty term σ^2 (corresponding to the

assumption of a lower degree of hidden heterogeneity between individuals) to match the data than Gompertz (Extended Data Item 1b–d). This motivates our preference for the Weibull hazard with frailty for data analysis and simulation whenever a parametric form is needed.

$$\begin{aligned} \text{Gompertz: } h(t | a, b) &= \frac{a}{b} \exp\left(\frac{t}{b}\right) \\ \text{Gompertz with frailty: } h(t | a, b, \sigma) &= \frac{\frac{a}{b} \exp\left(\frac{t}{b}\right)}{1 + \sigma^2 a \exp\left(\frac{t}{b} - 1\right)} \\ \text{Weibull: } h(t | \alpha, \beta) &= \frac{\alpha}{\beta} \left(\frac{t}{\beta}\right)^{\alpha-1} \\ \text{Weibull with frailty: } h(t | \alpha, \beta, \sigma) &= \frac{\frac{\alpha}{\beta} \left(\frac{t}{\beta}\right)^{\alpha-1}}{1 + \sigma^2 \left(\frac{t}{\beta}\right)^\alpha} \\ \text{Log-normal: } h(t | \mu, \sigma) &= \frac{\phi\left(\frac{\log t - \mu}{\sigma}\right)}{\sigma t \left(1 - \Phi\left(\frac{\log t - \mu}{\sigma}\right)\right)} \\ \text{Log-logistic: } h(t | \alpha, \beta) &= \frac{\frac{\alpha}{\beta} \left(\frac{t}{\beta}\right)^{\alpha-1}}{1 + \left(\frac{t}{\beta}\right)^\alpha} \\ \text{inverse Gaussian: } h(t | \mu, \lambda) &= \frac{\sqrt{\frac{\lambda}{2\pi t^3}} \exp\left(-\frac{\lambda(t - \mu)^2}{2\mu^2 t}\right)}{1 - \Phi\left(\sqrt{\frac{\lambda}{t}} \left(\frac{t}{\mu} - 1\right)\right) - \exp\left(\frac{2\lambda}{\mu}\right) \phi\left(-\sqrt{\frac{\lambda}{t}} \left(\frac{t}{\mu} + 1\right)\right)} \end{aligned}$$

Table 1: Distributions used for fitting. ϕ and Φ denote the PDF and CDF, respectively, of the standard normal distribution; μ and σ denote the mean and standard deviation (in the case of the log-normal, the mean and standard deviation of the logarithm of x); λ , α , and a are shape parameters; β and b are scale parameters. In the case of frailty, individual hazards $h_i(t)$ are related to a baseline hazard by a random factor Z that follows a Gamma distribution with mean 1 and variance σ^2 .

Note 2 — Quantifying temporal scaling

Summary

We quantify deviations from temporal scaling by applying a modified Kolmogorov-Smirnov (KS) test to the residuals of Accelerated Failure Time (AFT) regression models (Fig. 1c-d). The KS test we deploy was generalized to work with censored data by Fleming et al. [10]. Like the standard KS procedure, it tests whether the AFT residuals come from the same underlying distribution, in which case we conclude that all statistical differences between two sets of lifespan data are indeed accounted for by a scale factor and the two distributions are therefore related by temporal scaling. The KS test also provides a distance metric between survival curves. We note that many familiar statistical methods for evaluating differences between lifespan distributions, for example the log-rank test, perform poorly when survival curves cross, a situation perhaps rare in raw empiric data but essentially guaranteed among AFT residuals which, by construction, will have nearly identical means and differ only at higher moments.

Contents of Note 2

2.1	Adapting the Kolmogorov-Smirnov two-sample test	8
2.2	Evaluating the performance of the generalized KS test	9

2.1 Adapting the Kolmogorov-Smirnov two-sample test

We use the test statistic proposed by Fleming et al. [10] to confirm or reject the null hypothesis that two survival datasets are sampled from the same underlying distributions (and thus related by temporal scaling). This test statistic is also useful as a distance metric for clustering survival curves (Extended Data Item 3).

We first fit two merged sets of lifespan data with the AFT model

$$\log y_i = \beta x_i + \epsilon_i \quad (7)$$

where the categorical covariate x_i codes for the experimental grouping of individuals, for example plate ID or scanner ID, as detailed in our Statistical Methods section. The resulting residuals ϵ_i are then grouped according to the experimental variable in which they differ, for example genotype or environmental condition. Given populations a and b , we refer to the CDFs of these residuals with $S_a^\epsilon(t)$ and $S_b^\epsilon(t)$, respectively, although these survival curves are never explicitly calculated.

In the absence of censoring, the canonical Kolmogorov-Smirnov test involves estimating

$$F_{a,b}(t) = S_a^\epsilon(t) - S_b^\epsilon(t) \quad (8)$$

and then computing the KS statistic

$$D_{a,b} = \sup |F_{a,b}(t)|, \quad (9)$$

where \sup is the supremum (maximum) function, to calculate a p -value pertaining to the null hypothesis. Given the nature of the KS test and our AFT pre-treatment, the null hypothesis *affirms* temporal scaling.

Fleming et al. propose an analogous statistics for censored data, with


$$Y_{a,b}(t) = \left(\frac{S_a^e(t) + S_b^e(t)}{2} \right) (\beta_a^e(t) - \beta_b^e(t)), \quad (10)$$

where $\beta^e[t]$ is the cumulative hazard at time t for each set of residuals, and the modified KS statistic

$$\bar{Y}_{a,b} = \sup |Y_{a,b}(t)|. \quad (11)$$

We will refer to $Y_{a,b}(t)$ as $Y(t)$ and $\bar{Y}_{a,b}$ as \bar{Y} when the two populations a and b are obvious from context. \bar{Y} quantifies the degree of heteroscedasticity in the AFT covariate X_i . A significant p -value corresponding to \bar{Y} suggests a significant deviation from perfect scaling between the two survival curves.

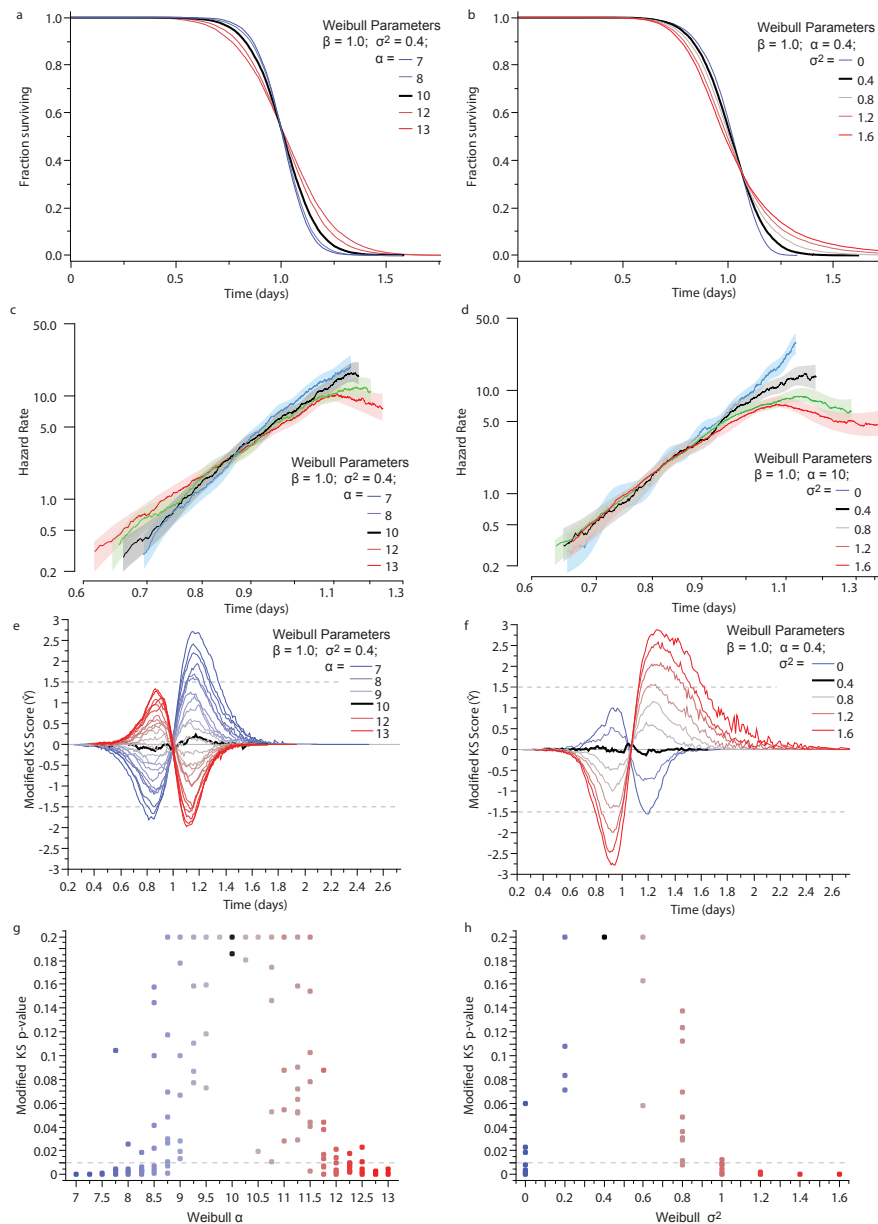
The p -value calculated from this modified KS test depends both on the significance level desired, the \bar{Y} obtained, and on the fraction of animals remaining alive at the end of an experiment (equation 2.4.iv in Flemming et al. [10]). Since this fraction is essentially zero in all experimental (and simulated) data sets presented here, the KS p -value simplifies to $\exp(-2 * \bar{Y}^2)$. Accordingly, any $\bar{Y} \geq 1.51$ or $\bar{Y} \leq -1.51$ will be significant at the $p \leq 0.01$ level. Because of how \bar{Y} is defined, any $Y(t)$ rising above above 1.51 or below -1.51 indicates a significant deviation from temporal scaling. We use this threshold throughout the Supplement when interpreting the values of $Y(t)$ and \bar{Y} .



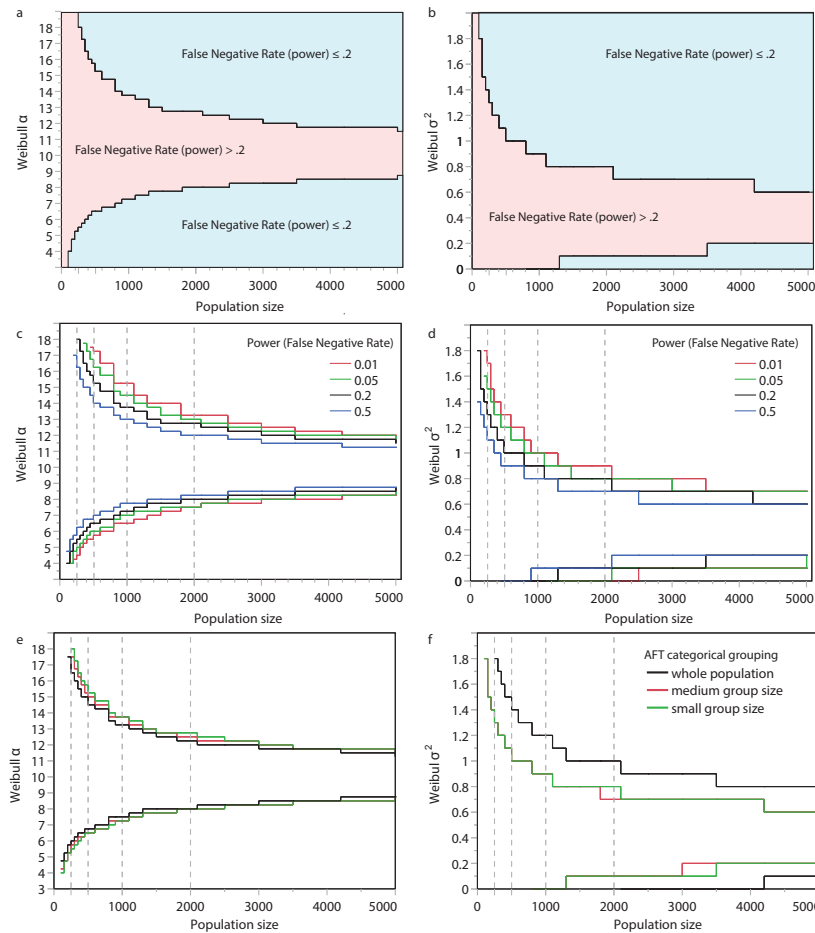
Supplementary Figure 2-1: Quantifying deviations from perfect scaling (schematic). **a:** To quantify deviations from perfect scaling, we first apply an Accelerated Failure Time model, to obtain the residuals of each curve (panel b). This removes any differences in time scale. If such differences are all there is, we have perfect scaling, which is to say that both residuals come from the same distribution. This is the null hypothesis that we test by quantifying the deviation from perfect scaling using a generalized Kolmogorov-Smirnov (KS) test [10]. **b:** The KS test evaluates the distance between two survival curves of AFT residuals (vertical gray bars) by assessing the probability of observing their maximal distance (arrow) by chance. This provides a way of quantifying the overall deviation from perfect scaling between the two curves.

2.2 Evaluating the performance of the generalized KS test

To evaluate the sensitivity of the AFT+KS-test combination, we compared sets of simulated lifespan distributions differing by progressively smaller parameter values (Fig. 2-2). In populations of 1,500 individuals, the test statistic reliably identified changes in the Weibull α of about 2.5 and increases in σ^2 of about 0.6. In this case and in subsequent power calculations, we censored 10% of death times to match the level typical in our experimental data.



Supplementary Figure 2-2: Evaluating the AFT+KS-test combination with simulated data. Sets of death times containing 1,500 samples were drawn from a Weibull distribution with frailty correction(1.4). Baseline parameters were chosen to approximate those observed in the AFT residuals of wild-type *C. elegans* lifespan distributions, $\alpha = 10$, $\theta = 0.4$, and $\beta = 1$ (α and β refer to parameters of the Weibull distribution, not the AFT regression). Any change in the shape parameter α or frailty parameter σ^2 will cause simulated populations to deviate from the baseline population, as is visible in the survival curves (panels **a–b**) and hazard plots (panels **c–d**) of the AFT residuals calculated for each population. **e–f:** Each simulated population was compared to the baseline population, plotting the generalized KS distance $Y(t)$, equation (10). **g–h:** Comparison of each population to the baseline across replicate trials with a calculation of the modified KS p -value to estimate the likelihood that observed differences occurred by chance. All values $p > 0.2$ were cropped at $p = 0.2$.



Supplementary Figure 2-3: Power of the AFT+KS-test combination with simulated data. **a:** The phase-space plot indicates the minimum population size required to detect a given effect size (a difference in the Weibull α to the reference population) at a statistical power of 0.8 (a false negative rate of 0.2) and significance level (false positive rate) of 0.01. The blue (red) area indicates effect and population size combinations that yield a statistical power greater (less) than 0.2. Populations of 1,000 individuals are required to detect differences in α of 2.5; increasing populations beyond 1,500 individuals appears to yield diminishing returns. **b:** Similar conclusions can be drawn from a phase-space plot for the frailty σ^2 . Increasing the population size beyond 1,500 individuals is met with diminishing returns in the capacity to discriminate differences in frailty. **c, d:** Same as in panels a and b, but for a statistical power of 0.5, 0.8, 0.95 and 0.99. **e, f:** The AFT model preceding the KS test can be run on different population size “units” afforded by the Lifespan Machine: plates (≈ 40 animals), whole scanners (≈ 320 animals), or the entire population. Here, each sample of the same unit size is a categorical variable in the AFT regression. After running the regression on all units, the residuals are pooled to form the distribution that is compared against the reference distribution with the generalized KS test. Depending on whether the differences are in α (panel e) or σ^2 (panel f), pooled populations will be subject to different distortions. It is therefore meaningful to check how such pooling impacts power. The results are shown in panels e and f for AFT pre-processing on subpopulations that were small (40), medium (320) or large (750). The size of subpopulation pooled does not have a big effect on power with regard to α (panel e), but it has some effect with regard to frailty (panel f).

To assess the statistical power of the AFT+KS-test combination, we generated sets of simulated death times as for Fig. 2-2, but varied the population size between 50 and 5,000 individuals; the Weibull α was varied between 4 and 18 and, separately, the frailty σ^2 between 0 and 2. Populations corresponding to each parameter set were compared to a baseline population of equal size with Weibull $\alpha = 10$ and $\sigma^2 = 0.4$, which closely matches *C. elegans* distributions. The p -value from each comparison was collected across 750 replicates. For each set of replicates and each population size, the fraction of significant p -values ($p \leq 0.01$) was then computed to estimate the power of the KS test at each population size (Fig. 2-3). Fig. 2-3 shows the α or σ^2 (the effect size) that can be detected at a given population size with a given false negative rate. We found that a minimum of 500 individuals per population were required to detect even substantial deviations from temporal scaling, but that increasing population sizes beyond 1,500 individuals provided diminishing returns in statistical power. As expected, substantial increases in population size were required to achieve greater statistical power: decreasing the false-negative rate of the test from 0.5 (half of all true deviations from survival curve shape are missed) to 0.05 (5% of all true deviations from survival curve shape are missed) required at least 2.5-fold larger populations at any effect size.

Note 3 — Heterogeneity and deviations from temporal scaling

Summary

We observe a change in hazard shape when animals are transferred from what we might call the “low temperature regime” (LTR, 20 °C–30 °C) to the “high temperature regime” (HTR, 30 °C–35 °C), across the 30 °C mark. This change appears to break scaling due to an increased deceleration of the hazard function in the HTR as compared to the LTR (Fig. 1d, e of the main text). A deceleration of the hazard can be accounted for in terms of “hidden” (i.e. not measured) heterogeneity also known as “frailty” in the demographic literature [11]. Such an approach distinguishes hazard functions that pertain to individuals from the hazard function that pertains to the population. From this perspective, an increased hazard deceleration could be due to an increased heterogeneity in the population. Furthermore, if heterogeneity were due to differences between individuals that nonetheless conform with a temporal rescaling of their aging process, the population hazard would still be affected in a way that deviates from temporal scaling. However, such a deviation would be incidental rather than fundamental. Heterogeneity of this kind might be an unavoidable side-effect of our experimental protocols, and a possible candidate are temperature inhomogeneities in our apparatus. We therefore sought to quantify the magnitude in heterogeneity that would be required to explain the change in hazard shape observed between 25 °C and 33 °C. We first quantify heterogeneity using a parametrized hazard function and then map it to a corresponding heterogeneity in temperature, since heterogeneity of this nature would be consistent with scaling at the individual level. The results of this analysis suggest that temperature inhomogeneities in our apparatus are unlikely to account for the change in hazard across the 30 °C transition. We cannot exclude other sources, some of which may be compatible with scaling of individual hazards, others may not. In sum, our experimental data do not allow us to adjudicate whether the change in hazard we observe is or is not a true scale breaker, but they suggest rejecting the hypothesis that potential temperature inhomogeneities in our apparatus are the cause of the change in hazard shape across the 30 °C transition.

Contents of Note 3

3.1	Background: Weibull heterogeneity (frailty)	13
3.2	Quantifying heterogeneity	14

3.1 Background: Weibull heterogeneity (frailty)

In the demographic literature, “frailty” refers to unobserved (in the sense of not measured) heterogeneity among individuals within a population [11]. The assumption is often made that individuals or subpopulations differ from one another with respect to hazard in a fashion that can be characterized by a PH model:

$$h_i(t) = Z_i h_0(t), \tag{12}$$

where $h_i(t)$ is the hazard of individual i , $h_0(t)$ is the common baseline hazard, and Z_i is the value of a random variable Z , often assumed to follow a Gamma distribution with mean 1 and variance σ^2 .

It has been shown (see any textbook, such as [1]) that under these assumptions the population

hazard $h(t)$ behaves like

$$h(t) = \frac{h_0(t)}{1 + \sigma^2 H_0(t)}, \quad (13)$$

where $H_0(t)$ is the cumulative baseline hazard, $H_0(t) = \int_0^t h_0(\tau) d\tau$.

For example, if the hazard follows a Weibull form, $\alpha/\beta(t/\beta)^{\alpha-1}$, the individual hazards in the frailty model just described become:

$$h_i(t | \alpha, \beta) = \frac{\alpha}{Z_i^{-\frac{1}{\alpha}} \beta} \left(\frac{t}{Z_i^{-\frac{1}{\alpha}} \beta} \right)^{\alpha-1} \quad (14)$$

In this case, the proportional frailty model, equation (12), is equivalent to a model in which individual differences (Z_i) multiply the timescale parameter β : in a Weibull frailty model, hidden heterogeneity within a population impacts lifespan through temporal scaling. The Weibull with frailty population-level hazard is listed in Table 1.

3.2 Quantifying heterogeneity

A deceleration in population hazard has been noted in model organisms and in human populations [6, 9, 12, 13]. Several distinct but mutually non-exclusive explanations for the phenomenon have been offered. One explanation invokes heterogeneity, which can give rise to a hazard deceleration because the more frail individuals die early, leaving a population with a declining hazard rate. This type of heterogeneity could be intrinsic to populations, or be produced by some extrinsic environmental heterogeneity, for example the slight temperature differences we can measure across different regions of our imaging apparatus [8]. Heterogeneity due to temperature differences at different locations in our apparatus would cause deviations from temporal scaling at the population level even though they only cause temporal rescaling at the individual level.

An alternative explanation is that Markov processes with an absorbing boundary (death) have, in general, absorption (hazard) rates that become constant in the long time limit, even if all individuals are identical initially. Such stochastic processes represent aging as a decline in a physiological state space. They reach a quasi-stationary distribution over physiological states [1, 14–16], which causes the hazard to level off. All examples in Supplementary Note 6, in which we illustrate scaling at the process level, exhibit this phenomenon. In particular, the inverse Gaussian distribution arises from a class of first-passage time models (drift-diffusion) [15] that fit our *C. elegans* data reasonably well. A Weibull model with frailty provides a better fit to our empiric data (Extended Data Item 1b–d), but a small frailty correction to the Gaussian model would entirely remove this advantage.

We next set out to determine whether the changes we observe in hazard curve shape could be reasonably explained by temperature heterogeneity. If the heterogeneity required to produce the observed degree of hazard deceleration were improbably large, we might reject that hypothesis.

In our data sets, the largest deviation from temporal scaling produced by temperature occurs at 30 °C. To estimate the magnitude of change in environmental heterogeneity required to explain this deviation from temporal scaling, we compared populations held at 25 °C and at 33 °C (Fig. 1d) by fitting both with a Weibull model with a frailty correction term (section 3.1). Instead of estimating model parameters independently for each populations and then comparing them, we

apply a single model to both populations simultaneously and quantify changes between them as relative factors. This enables more meaningful confidence intervals and therefore more informative estimates for the significance of differences. Our model represented the two curves as

$$h_{25^\circ}(t | \alpha, \beta, \sigma) = \frac{\frac{\alpha}{\beta} \left(\frac{t}{\beta}\right)^{\alpha-1}}{1 + \sigma^2 \left(\frac{t}{\beta}\right)^\alpha} \quad h_{33^\circ}(t | \alpha, \beta, \sigma, \Delta\alpha, \Delta\sigma^2) = \frac{\frac{\alpha\Delta\alpha}{\beta} \left(\frac{t}{\beta}\right)^{\alpha\Delta\alpha-1}}{1 + \sigma^2 \Delta\sigma^2 \left(\frac{t}{\beta}\right)^{\alpha\Delta\alpha}} \quad (15)$$

with $\Delta\alpha$ and $\Delta\sigma^2$ representing the effect of high temperature on the shape and scale of the curve.

We estimated model parameters using a maximum likelihood approach (R-package `flexsurv`). At 25 °C, $\alpha = 10.5$ (95% confidence interval $ci = [10.2 - 10.9]$). At 33 °C populations exhibited a 10% decrease in α , with $\Delta\alpha = 0.91$ (95% $ci = [0.86 - 0.96]$). At 33 °C exhibited a large increase in σ^2 , relative to 25 °C, with $\Delta\sigma^2 = 2.7$, (95% $ci = [2.2 - 3.38]$). A model assuming a fixed σ^2 across temperatures ($\Delta\sigma^2 = 1$) provided a substantially worse fit ($AIC = -7385$) to the data than a model assuming a fixed α across temperatures ($\Delta\alpha = 1$, $AIC = -7461$). Were the difference in survival curve shape caused by only a single parameter, it is more likely to result from an increase in σ^2 .

To map the magnitude of this increase in σ^2 to a biologically informative scale, consider that frailty (the distribution of the Z_i in equation 14) acts multiplicatively on the timescale β and can thus be interpreted as a distribution of individual scale factors λ_i in the population. A larger σ^2 suggests that individuals experience a greater diversity in their individual λ_i , as shown in Extended Data Item 8 a. At each set point, 25 °C and 33 °C, the data in Fig 4.b of the main text allow us to convert this heterogeneity into a corresponding distribution of temperatures accounting for the distribution of λ_i (and thus the estimated σ^2). This approach suggests that a relatively small heterogeneity in temperature is sufficient to explain σ^2 (Extended Data Item 8 b): at 25 °C only 25% of individuals would need to experience a temperature 0.18 °C less than average and only 25% would need to experience a temperature 0.10 °C greater than average. To explain the increased σ^2 at 33 °C, only 25% of individuals would need to experience a temperature 0.37 °C less than average, and only 25% would need to experience a temperature 0.31 °C greater than average.

Although these temperature differences seem small, we need to consider that the AFT algorithm used to estimate the σ^2 already accounts for any difference in temperature between scanners and between plates. Thus, these differences would have to correspond to thermal gradients *within* each plate. Seen from this angle, the putative temperature differences suddenly seem unlikely large. In addition, since individuals more or less randomize their location on plates over time, it seems improbable that a temperature gradient (of however unlikely magnitude) could produce a substantive heterogeneity-effect. We conclude, therefore, that static temperature heterogeneity does not account for the observed deviations from temporal scaling between populations held at 25 °C and 33 °C. Other sources of heterogeneity, for example a differential induction of the heat-shock response among individuals [17], must be invoked. Some might be compatible with temporal scaling at the individual level, others not.

Our frailty model assumes a baseline Weibull hazard. To determine whether our quantification of heterogeneity depended on the specific model used, we ran a simulation that artificially injected heterogeneity into experimental lifespan distributions. We randomly selected 56,000 death times from AFT-residuals of the data collected at 25 °C (Figure 1.d) and multiplied each sample by a random number drawn from a Gamma distribution corresponding to the increase in σ^2 observed

between populations at 25 °C and 33 °C (the $\Delta\sigma^2$ of Eq. 15). If the difference in shape observed between 25 °C and 33 °C shown in Fig. 1d of the main text resulted exclusively from an increase in frailty, then the distribution of the so-transformed residual death times should be equivalent to that measured at 33 °C. This is indeed what we see (Extended Data Item 8 c, d), suggesting that the deviations from perfect temporal scaling between 25 °C and 33 °C are consistent with an increase in heterogeneity within populations.

In final analysis, temperature differences within our apparatus seem not to account for the heterogeneity we observe across the 30 °C transition. It is therefore possible (though not certain) that whatever causes this heterogeneity is reflective of changes in mechanistic aspects of the stochastic aging process and thus breaks temporal scaling at the individual level.

Note 4 — The effect of transient perturbations

Summary

In Figure 1 of the main text, we identify interventions that, when applied throughout life, produce a temporal scaling of lifespan distributions. For such interventions, we derive expressions that predict the effects on the lifespan distribution when these interventions are applied transiently. We highlight the specific case where interventions are applied before any deaths occur in a population, in which case interventions are predicted to produce a temporal *shift* of the lifespan distribution of the affected population relative to a control. We then explain how these predictions are met by the experimental data presented in Fig. 3 of the main text.

Contents of Note 4


4.1	Intuition	17
4.2	Definitions	18
4.3	No switching	19
4.4	Switching	20
4.5	Predictions	20
4.6	The effect of switching on distributions	20

4.1 Intuition

We observe that diverse interventions produce a temporal scaling of lifespan distributions. Perhaps the simplest way for such scaling to come about would be if, throughout adulthood, all changes within a small time increment dt of those aspects of *C. elegans* physiology that determine the risk of death in control individuals occurred within a rescaled time increment λdt in exposed individuals. The intervention would then be producing a temporal scaling of a stochastic aging process that determines risk and generate, through this rescaled process, a correspondingly rescaled lifespan distribution. In other words, the cumulative physiological decline of individuals under intervention would accrue on a rescaled timescale (faster or slower, depending on the value of λ) than for individuals in the control population. This speed-up or slow-down in the physiological decline process determining risk would have to occur even if specific underlying molecular mechanisms changed as a result of the intervention.

If at some point the intervention were stopped and conditions reverted to those of control animals, then, from that moment onward, no additional differences would accumulate between populations. However, the existing differences would persist, making an exposed population “older” or “younger” (in terms of risk) than the control population, and, as a result, causing its lifespan distribution to be *shifted*, not scaled, toward earlier or later chronological times, respectively, compared to the control.

Within any population some individuals live longer than others. In our experiments, we leverage this variation within apparently homogeneous populations to test whether interventions are in fact rescaling a lifelong stochastic process that determines lifespan. We show that the variation in lifespan observed within homogenous populations is indeed shifted to a predictable extent by interventions that produce a temporal scaling of the underlying stochastic aging process. We explain the extent of this shift informally in the caption to Fig. 4-1; we derive it formally for expected lifespan in sections 4.2–4.5 and for the case of distributions in section 4.6).



Supplementary Figure 4-1: On shifting lifespan distributions. The distributions of a control population and a population that is transiently subject to an intervention will undergo a shift relative to each other, provided the intervention causes a rescaling of the lifespan distribution if applied throughout life. Consider a control population P_0 always held at condition C_0 and a test population P_1 held at C_1 for T time units, say from time $t = 0$ to $t = T$. As far as the aging process that governs risk of death is concerned, T time units pass for P_0 , while T/λ units pass for P_1 . (The scale factor λ has the same definition as in Note 1 and the main text. However, in our experiments we switch from a perturbed situation to a control and express time T in the exposed population relative to the control population; as a result, time T is divided by λ .) At time T , P_1 is switched to the same condition as P_0 . From that moment on, both age at the same speed, maintaining the temporal difference between control and exposed population: $T(1 - 1/\lambda)$. For illustration, consider two biased random walks, RW1 and RW2, in a physiologic state space that governs risk. Over time, RW1 and RW2 drift towards the absorbing boundary at 0. RW2 proceeds twice as fast as RW1 and both walks started in the same state. The plot shows the probability distributions over live states of these rescaled walks at two times. At time 10, the parameter values of RW1 are changed to those of RW2. Because of the scaling assumption, 5 time units ago the distribution over states for RW2 must have been exactly the same that RW1 has at the switching time 10 (blue curve). Thus, from time 10 onward, the switched RW1 will evolve exactly like RW2, but with a delay of 5 time units, i.e. a pure shift (e.g. orange curve).


4.2 Definitions

Throughout this Note, Y denotes the lifespan random variable. To predict the consequences of transient interventions on mean lifespan, we need the notion of a conditional mean lifespan. The condition will be used in section 4.4 to represent the switching time. The mean *residual* lifespan (mrl) at time t , $E(Y - t | Y > t)$, is the expected remaining lifespan of an individual that is alive at time t . Since the expectation of the (linearly) transformed random variable $Y - t$ is the expectation of Y minus t , we can state:

$$E(Y - t | Y > t) = E(Y | Y > t) - t. \quad (16)$$

where $E(Y | Y > t)$ is the conditional expectation of lifespan (computed at age zero) given survival at t : $E(Y | Y > t) = \int_t^\infty \tau l(\tau) d\tau / S(t)$, where $l(t)$ is the probability density function and $S(t)$ is the survival function. To declutter notation, let us write $\bar{l}(t)$ for $E(Y - t | Y > t)$. Fig. 4-2

illustrates these notions graphically using the drift-diffusion process (section 6.2) as the underlying aging model.



Supplementary Figure 4-2: Mean residual lifespan at time t , $\bar{l}(t)$, and conditional expected lifespan $E(Y | Y > t)$ according to the diffusion-drift aging process with lifespan distribution (31) and survival (32).

4.3 No switching

First we consider the scenario in which individuals W_1 and W_0 are held throughout life at environmental conditions C_1 and C_0 , respectively. (We shall use the indices 0 and 1 to refer to these conditions.) If we assume that scaling (equation 2) holds, then at time t , W_1 will have covered t/λ time units in its aging process compared to W_0 , who has covered t time units. With this assumption, the mrl functions $\bar{l}_1(t)$ and $\bar{l}_0(t)$ of W_1 and W_0 , respectively, can be related to each other:

$$\bar{l}_1(t) = \frac{\int_t^\infty \tau l_1(\tau) d\tau}{S_1(t)} - t = \frac{\int_t^\infty \tau/\lambda l_0(\tau/\lambda) d\tau}{S_0(t/\lambda)} - t = \frac{\int_{t/\lambda}^\infty x l_0(x) \lambda dx}{S_0(t/\lambda)} - t \quad (17)$$

$$= \lambda \left[\frac{\int_{t/\lambda}^\infty x l_0(x) dx}{S_0(t/\lambda)} - \frac{t}{\lambda} \right] \quad (18)$$

$$= \lambda \bar{l}_0(t/\lambda) \quad (19)$$

where, as before, $l_i(t)$, $i = 0, 1$ are the probability density functions and $S_i(t)$, $i = 0, 1$ the survival functions. We make use of temporal scaling in the second equality; in the third we make the variable transform $x = \tau/\lambda$; in the fourth we pull out λ . We therefore derive a scaling relation between mean residual lifespans: $\bar{l}_1(t) = \lambda \bar{l}_0(t/\lambda)$.

We can also calculate the difference in mean residual lifespan, $\Delta(t)$, between the populations

always living at C_0 and C_1 :

$$\begin{aligned}\Delta(t) &= \bar{l}_1(t) - \bar{l}_0(t) = \lambda \bar{l}_0(t/\lambda) - \bar{l}_0(t) = \lambda (E_0(Y | Y > t/\lambda) - t/\lambda) - E_0(Y | Y > t) + t \\ &= \lambda E_0(Y | Y > t/\lambda) - t - E_0(Y | Y > t) + t \\ &= \lambda E_0(Y | Y > t/\lambda) - E_0(Y | Y > t)\end{aligned}\tag{20}$$

$$\tag{21}$$

$\Delta(t)$ is pictured in Fig. 4-3 (left panel) for the case considered in Fig. 4-2. For early times t when the expected conditional lifespan does not effectively change, i.e. before any deaths occur in either population, we can write $E_0(Y | Y > t/\lambda) \approx E_0(Y | Y > t) \approx E_0(Y)$ and thus approximate the difference in mean remaining lifespan between the two populations by the simple relationship $\Delta(t) \approx (\lambda - 1)E_0(Y)$. For early times, $\Delta(t)$ is independent of t : $\Delta(t) = E_1(Y) - E_0(Y)$, and thus $E_1(Y) = \lambda E_0(Y)$, which is but the scaling relation for the mean life expectancy. This is illustrated on the left panel of Fig. 4-3 for the case $\lambda = 2$.

4.4 Switching

Now consider the scenario where W_1 is switched at time τ from condition C_1 to the condition C_0 at which W_0 is held. At τ , W_1 will have progressed in its lifespan by τ/λ relative to W_0 , but after the switch both are progressing along the same process with the same speed. In this scenario, the difference in mrl is:

$$\begin{aligned}\Delta(\tau) &= \bar{l}_0(\tau/\lambda) - \bar{l}_0(\tau) = E_0(Y | Y > \tau/\lambda) - \tau/\lambda - E_0(Y | Y > \tau) + \tau \\ &= E_0(Y | Y > \tau/\lambda) - E_0(Y | Y > \tau) + \tau(1 - 1/\lambda).\end{aligned}\tag{22}$$

$$\tag{23}$$

For early switch time τ , before any individuals have died in either population, $E_0(Y | Y > \tau/\lambda) \approx E_0(Y | Y > \tau)$. In this case,

$$\Delta(\tau) = \tau(1 - 1/\lambda).\tag{24}$$


For late times τ , $E_0(Y | Y > \tau/\lambda)$ grows like τ/λ and $E_0(Y | Y > \tau)$ like τ . Thus $\Delta(\tau)$ will converge to $\tau/\lambda - \tau + \tau - \tau/\lambda = 0$. Indeed, a situation in which the population on condition C_0 is switched after everyone in it has died will be indistinguishable from one in which no switching occurred at all. In that case, the mean remaining lifespans are not related by a shift, but by scaling.

4.5 Predictions

Equation (24) makes three predictions: (1) interventions early in life produce temporal shifts in lifespan distributions; (2) the magnitude of the shift increases proportionally with the time τ at which the switch occurs; and (3) the magnitude of the shift is inversely proportional to the scale factor λ that relates the environmental conditions before and after the switch. These predictions are illustrated in supplementary Fig. 4-3. All three predictions depend on the temporal scaling assumption made in the first step of equation (23). Since they appear to be in excellent agreement with the data presented in Fig. 3 of the main text and the Extended Data Item 5, we conclude that temporal scaling is an appropriate characterization of the effect that temperature has on *C. elegans* aging.

4.6 The effect of switching on distributions

The general analysis of switching experiments in the previous section did refer to mean residual lifespan. The scale and shift factors we calculated for the mean residual lifespan also hold for the residual lifespan *distribution* by virtue of our assumption of temporal scaling. However, it is



Supplementary Figure 4-3: Mean residual lifespan differences as a function of switch time τ (left and middle) and λ (right). Left panel: $\bar{l}_1(\tau) - \bar{l}_0(\tau)$ vs τ , without switching. This is the difference in mean residual lifespan at τ of W_1 always at condition C_1 and W_0 always at condition C_0 . Middle: $\bar{l}_0(\tau/\lambda) - \bar{l}_0(\tau)$ vs τ , with switching from condition C_1 with $\lambda = 2$ to condition C_0 at time τ . Right: Difference in mean residual lifespan when switching occurs at the times indicated ($\tau = 2, 3, 4, 10$) vs $1/\lambda$, where λ is the scale factor of condition C_1 relative to C_0 . The slope is $-\tau$, and the intercept is τ , when τ is small enough that $E(Y|Y > \tau) = E(Y|Y > \tau/\lambda) \approx E(Y)$. If λ is very large, such that $E(Y|Y > \tau) \approx \tau$ but $E(Y|Y > \tau/\lambda) \approx E(Y)$, the intercept becomes $E(Y)$ (and the slope is still $-\tau$). This is the case when λ is such that one can transfer W_1 before any member of its population, P_1 , has died, but after all members of the population P_0 (represented by W_0) have died. The underlying aging model is the drift-diffusion process with lifespan distribution (31) and survival (32); section 6.2. See Fig. 4-2 for the survival function.

useful to numerically illustrate the situation for the full lifespan distribution.

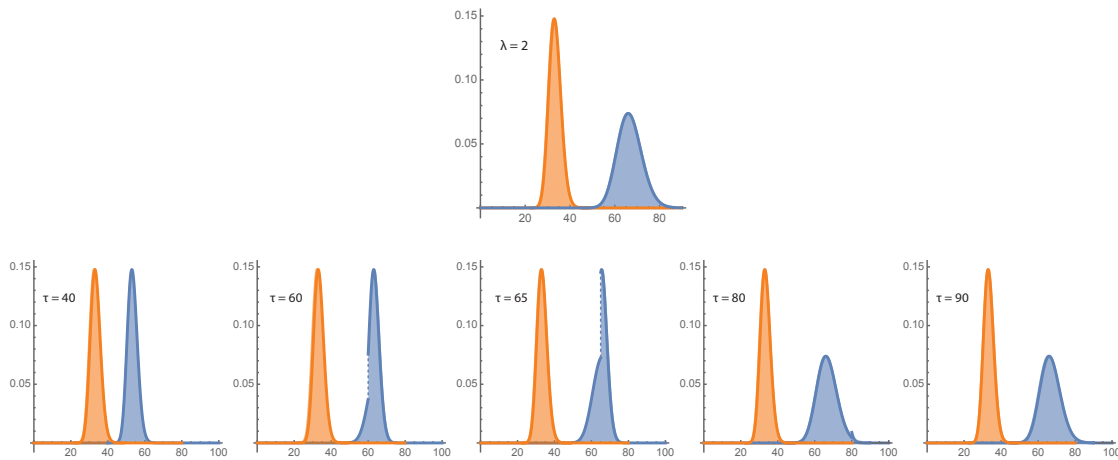
For the purpose of illustration we use a simple drift-diffusion aging model. Consider two populations P_1 and P_0 that age according to equation 30 with parameters $C_1 = \{\mu_1, \sigma_1^2\}$ and $C_0 = \{\mu_0, \sigma_0^2\}$, respectively. We switch P_1 from C_1 to C_0 at time τ . Let the lifespan distribution of the shifted population be $p_{1 \rightarrow 0}(X(t) = 0 | X(0) = c)$:

$$p_{1 \rightarrow 0}(X(t) = 0 | X(0) = c) = \begin{cases} \int_0^\infty p_0(X(t) = 0 | X(\tau) = v) p_1(X(\tau) = v | X(0) = c) dv & \text{if } t \geq \tau \\ p_1(X(t) = 0 | X(0) = c) & \text{if } t < \tau \end{cases} \quad (25)$$

$$= \begin{cases} \int_0^\infty p_0(X(t - \tau) = 0 | X(0) = v) p_1(X(\tau) = v | X(0) = c) dv & \text{if } t \geq \tau \\ p_1(X(t) = 0 | X(0) = c) & \text{if } t < \tau \end{cases} \quad (26)$$

In (26) we made use of the fact that the process is time-homogeneous.

$p_0(X(t - \tau) = 0 | X(0) = v)$ and $p_1(X(t) = 0 | X(0) = c)$ are given by (31) and $p_1(X(\tau) = v | X(0) = c)$ by (39). The integral can be expressed in “closed” form using error functions, but the lengthy expression is not particularly informative to write down explicitly. It is more useful to only present the result graphically, Fig. 4-4.



Supplementary Figure 4-4: Two populations P_1 (blue) and P_0 (orange) that are held at constant conditions C_1 and C_0 , respectively, yield lifespan distributions that differ by a scaling factor (here $\lambda = 2$, blue relative to orange). In the second row, P_0 (a high-T population, say) stays at high T all the time, while P_1 is switched from low T to high T in different simulated experiments at times $\tau = 40, 60, 65, 80, 90$. The blue distribution is the resulting lifespan distribution for all of P_1 . If the change in T occurs before any member of P_1 has died (up to about $\tau = 50$ in this example), one obtains a shift in the distribution relative to P_0 , as seen in the case of $\tau = 40$. When the switch occurs at $\tau = 90$ (rightmost), there is nothing to switch, because all members of P_1 have already died, yielding P_1 's original lifespan distribution—a scaling relative to P_0 . In between, P_1 's distribution will be piecewise—one piece resulting from any deaths that have occurred at low T before P_1 was switched and the other piece resulting from the deaths that occurred at high T after the switch. The piece of P_1 's lifespan distribution at high T is a shifted version of the corresponding piece of P_0 , while the piece coming from low T is a scaled version. The shift factors are the same as calculated for the mean residual lifespan case, $\tau(1 - 1/\lambda)$, regardless of the specific aging model.

Note 5 — Temporal scaling in phenomenological causal models

Summary

A mechanistic grounding of temporal scaling is hard to achieve at present, because little is known about the causes of death in *C. elegans* and the processes that lead up to them (such as bacterial pathogenesis [18] or germ-line proliferation [19], both of which seem to kill late in life). In this section we explore the conditions for temporal scaling in the context of an aging process that renders animals increasingly susceptible to the proximate causes that kill them. We analyze two types of phenomenological models. One type explores the failure of systems consisting of independent parts (“competing risks” model); the other explores the failure of networks consisting of interdependent parts (dependency network model). We numerically explore the constraints that each type of model places on temporal scaling, the main empirical phenomenon we observe in *C. elegans*. The dependency model in particular illustrates the concept of an emergent state variable whose dynamics determines lifespan. It also suggests how interventions that affect only a part of a system nonetheless impact the dynamics of the overall state variable in a fashion that produces temporal scaling. The two models we discuss here can be seen as quantitative versions of the schematics shown in Extended Data Item 9.

Contents of Note 5

5.1	Competing risk models	23
-----	-----------------------------	----

5.2	Interdependency networks	25
5.2.1	Homogeneous interventions	26
5.2.2	Heterogeneous interventions	27
5.3	Conclusions	30

5.1 Competing risk models

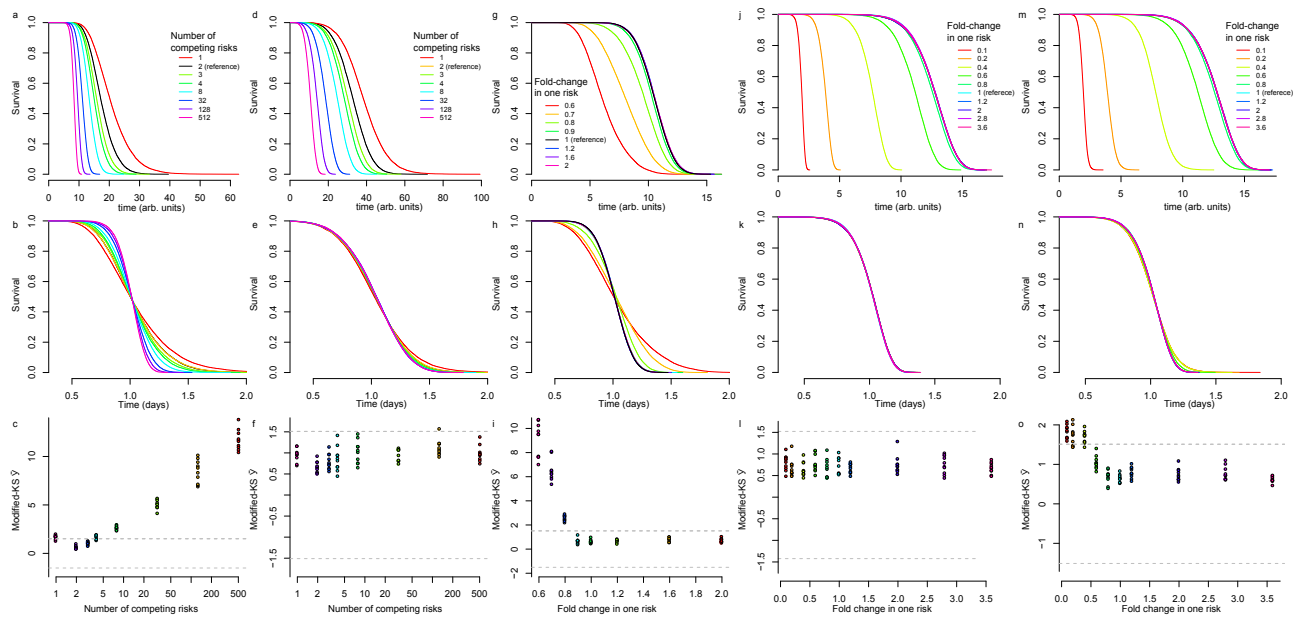
We first consider the possibility that animals might die from one of several possible causes that occur independently of one another. In the statistics literature, this assumption forms the basis of competing risks models [20]. In reliability theory, this assumption corresponds to an architecture in which systems rely on multiple critical components that operate in series, i.e. a “weakest-link” model. Both views posit that individuals/systems remain alive as long as no component has failed, and, conversely, that the first component to fail causes the entire system to fail. If each component i has survival function $S_i(t)$, $i = 1, \dots, n$, where n is the number of components, and hazard function $h_i(t)$, $i = 1, \dots, n$, the survival and risk functions of the whole system, $S(t)$ and $h(t)$, respectively, will be:

$$S(t) = \prod_{i=1}^n S_i(t) \quad \text{and} \quad h(t) = \sum_{i=1}^n h_i(t) \quad (27)$$

Obviously, systems made of strong components (i.e. components with greater mean lifespans) will fail later than systems made of weak components. It is also evident that systems with fewer components will last longer, all else being the same, than systems with a greater number of components, as each additional component introduces an additional possibility of failure. A component that is substantially weaker than all the others ($S_1(t) \ll S_i(t)$ for all t and all $i > 1$) will dominate the failure statistics of the system: $S(t) \approx S_1(t)$.

We explored the impact of parameters of the competing risks model with numerical simulations, summarized in Fig. 5-1. First, we considered the simple case where all components have the same distribution of failure times $S_i(t)$, specifically a log-normal distribution (Fig. 5-1a). In this case, the system’s behavior conforms with Extreme Value Theory (EVT) for minima bounded below by zero (as is the case with lifespans). A central result of EVT is that the survival distribution $S(t)$ of a system consisting of a large number of components with independently and identically distributed (iid) failure time distributions will take the form of a Weibull distribution [21]. The Weibull distribution will arise regardless of the specific form of the probability density of failure describing the components.

In simulations, we confirmed that when the total number of components is small, increasing the number of components will shorten lifespan in a manner that deviates significantly from temporal scaling (Fig. 5-1b, c). Such a change is therefore inconsistent with our experimental results. When the total number of components is large, changes in the number of components have little effect on lifespan. This again disagrees with our experimental findings, where we show that substantial lifespan increases occur by temporal scaling. Additionally, any intervention that breaks the iid assumption will be inconsistent with our experimental results: Even with large numbers of components (i.e. independent death causes), for example $n = 50$, we observe that altering the shape of the risk of one single cause shortens lifespan in a manner that deviates from temporal scaling (Fig. 5-1g–i). This will be true in particular for any system whose competing death causes are not iid to begin with. Were any intervention to emphasize one cause over any other, it would distort the overall lifespan distribution accordingly. In summary, we conclude that



Supplementary Figure 5-1: Independent competing risks. **a:** We presume a series of independent death causes (vital components) with log-normal distribution. Reducing the number of causes from two (black) to one (red) extends lifespan, while increasing the number of causes decreases lifespan (green-purple). **b:** As seen from the AFT residuals, altering the number of causes, yields a change in survival shape. **c:** The generalized KS-test confirms the significance of these changes; values of $|\bar{Y}|$ greater than 1.51 (gray line) correspond to a significant difference in shape ($p < 0.01$). **d:** The same model as in panel **a** but all components failing according to a Weibull distribution with a frailty correction chosen to match experimental data. As before, changing the number of components will increase or decrease the mean lifespan (panel **e**), but with little effect on the survival curve shape (panel **f**). **g:** Assuming again a log-normal failure time distribution for each component and fixing the number of components at 50, we alter the μ parameter of a single component, changing its mean lifespan while keeping its coefficient of variation fixed. Changing μ 0.9-fold or less shortens the lifespan of affected populations (green-red), but increasing it 1.2-fold or more has little effect, as the remainder of shorter-lived components dominate. **h–i:** Decreasing the lognormal μ of a single component by 0.8-fold or less produces statistically significant deviations from perfect temporal scaling. **j–l:** If component failure follows a Weibull distribution with no frailty correction, decreasing the time scale β of a single component changes lifespan without producing any deviation from temporal scaling. **m–o:** When a frailty term is incorporated to allow the reference population to better match *C. elegans* data, significant deviations from temporal scaling can be seen at large effect sizes if the frailty term σ is not changed in proportion to the scale factor β . See text for details.

to reconcile our experimental data with a competing risks model requires that interventions alter the risk profiles associated with all independent causes of death in exactly the same manner.

There is one exception to the iid assumption, however. If the failure risks of all components happened to be Weibull risks, then an intervention that affected only the time scale parameter of a subset of such components would still exhibit temporal scaling at the systems level, (Fig. 5-1j–l), which would remain consistent with our experimental data showing scaling. Incorporating a frailty term so that the reference population more closely matches our empirical data requires that the frailty term σ^2 also be scaled in proportion with the Weibull β , or else significant deviations

from temporal scaling are produced (Fig. 5-1m–o; Note 6.4). Altering the number of components would produce effects consistent with our experimental data (Fig. 5-1d–f). All this reflects the fact that the sum of Weibull risks that only differ in their time scale parameters is still a Weibull risk. With regard to interpreting our experimental results, a component Weibull scenario would not be very helpful, as it still requires an explanation of why each system component exhibits an extreme value distribution (Weibull). This distribution arises from a large number of *independent* components that exhibit *identically* distributed failure times. It is not obvious why all biological risk factors should meet such criteria. Moreover, it is unclear why an intervention on any so-structured risk factor would affect only its time scale parameter. Finally, while lifespan-shortening interventions might be consistent with our results, any substantial increase in average lifespan would require that the risk profiles of many components be rescaled by the same intervention.

We believe the competing risks case to be overly simplistic. Indeed, for interventions to produce temporal scaling within a competing risks framework they would have to affect the likelihood of all death causes to the same extent. This would be the case, however, if there was a physiological state variable of the kind we postulate, whose stochastic decline determined lifespan.

5.2 Interdependency networks

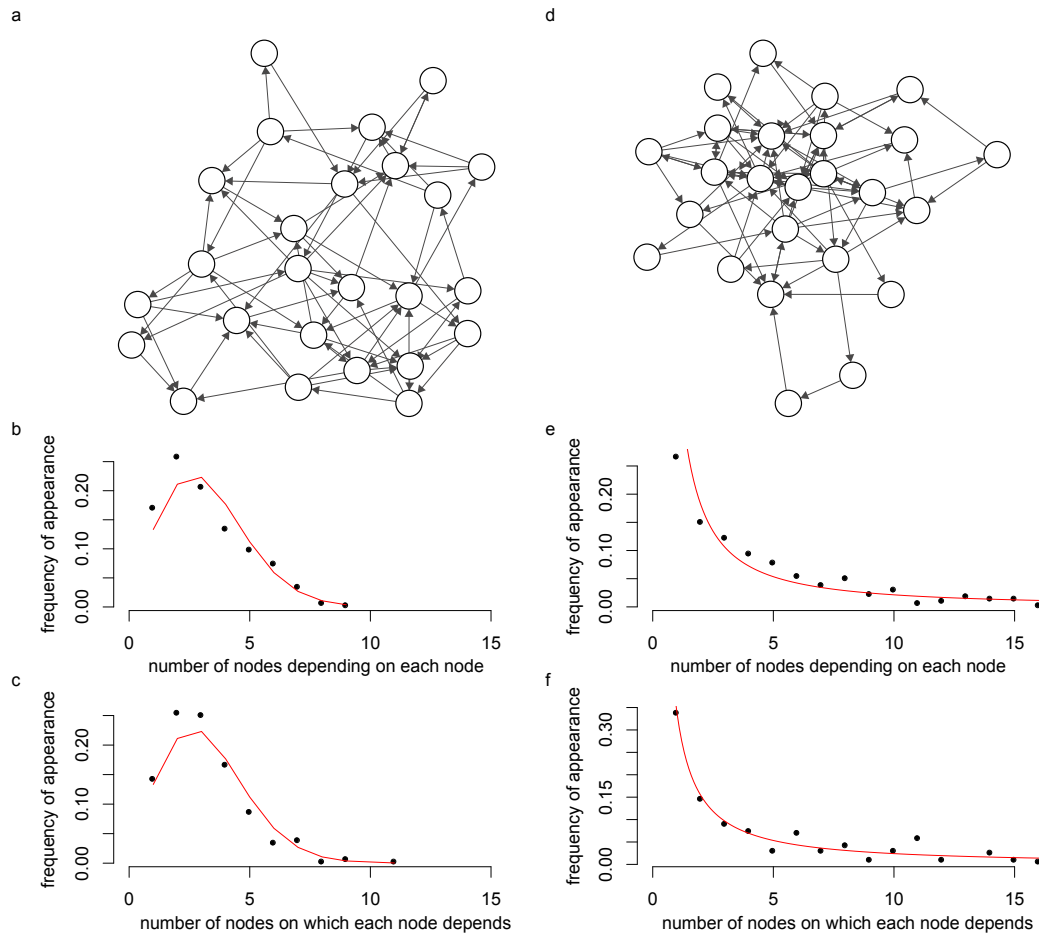
It seems unlikely that *C. elegans* die from causes that are independent. Interdependence is an essential aspect of life, as organ systems, tissues, cellular components and biochemical pathways all depend on each other. Although we do not understand these dependencies at a detail that allows us to suggest mechanistic models, progress might be made at a more abstract level by considering a recent aging model based on a dependency network [22].

In the Vural-Morrison-Mahadevan (VMM) model [22], nodes (representing functional units like cells, metabolic networks, organ systems) provide each other with physiological functions. A node has two states: functional or failed. A node fails with an intrinsic failure probability and, additionally, when a majority of nodes it depends upon have failed. This interdependency is represented as a set of directed edges that weave the nodes into a network. An arrow $i \rightarrow j$ means that node i contributes to maintaining the functionality that node j provides. Nodes may also have an intrinsic repair probability. The network fails as a whole when less than a specified fraction of nodes remains functional. The time to failure is the lifespan of an individual (network).

The VMM model requires that networks are initially connected, meaning that each node must depend on at least one other node, and each node must have at least one dependent. A node is subject to spontaneous (intrinsic) failure with a constant probability. Failed nodes can be repaired, again with a constant probability. A node experiences an extrinsic failure when more than half the nodes it depends upon have failed. This latter mechanism can cause networks to collapse by cascading failures. The entire network represents an individual. It fails when 99% of its nodes have failed. The model behavior is not sensitive to this threshold, as most networks collapse completely once more than 50% of nodes have failed.

AFT regression models and KS tests both assume that death times are not discretized. In order to use these statistical tools to analyze the results of simulations, we implemented the VMM model as a continuous-time Monte Carlo simulation [23]. We confirmed (data not shown) that our continuous-time version matched the behavior of the discrete version in [22].

Vural et al. demonstrate that specific parameter sets can cause interdependency models to yield hazard functions that correspond closely to empiric estimates of *C. elegans* lifespan distributions. Here, we do not address whether the VMM model fits our empiric data, because in Note 1.4 we show that a broad class of functions can provide reasonable fits to experimental data. Instead, we



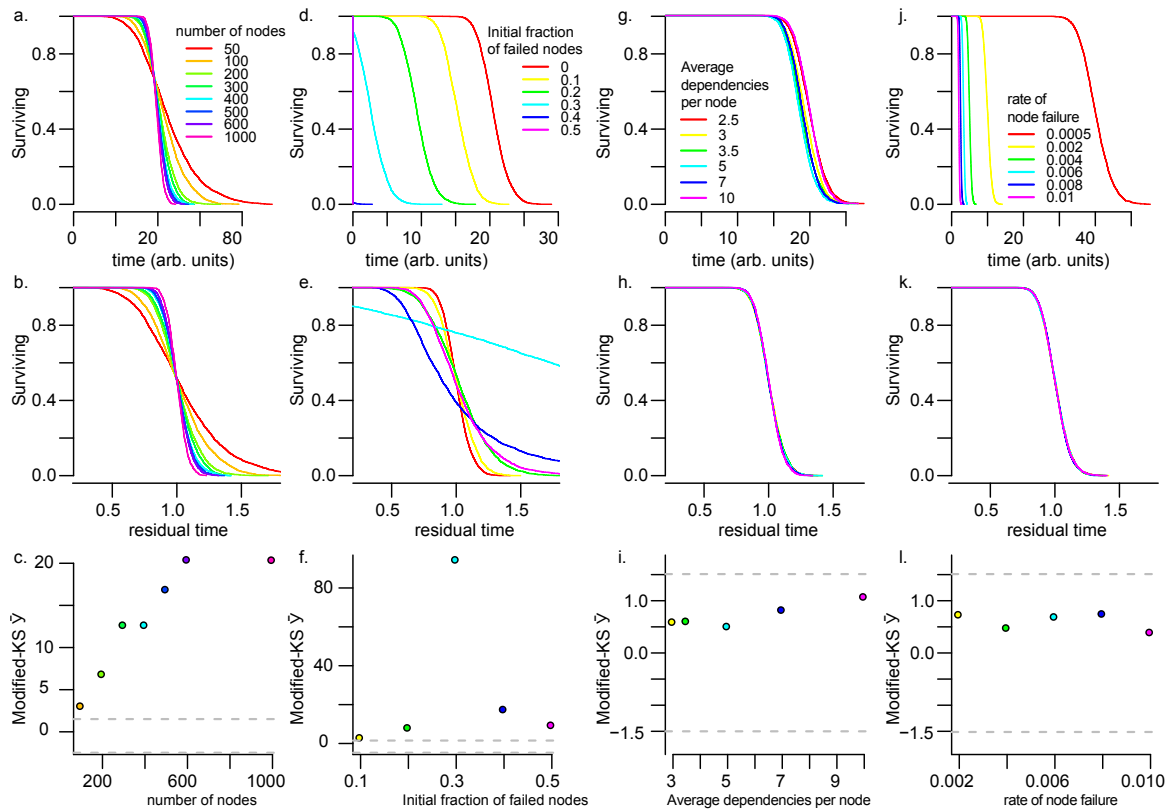
Supplementary Figure 5-2: Interdependency networks. **a, d:** Vural *et al.* [22] describe interdependency networks with random and scale-free connectivity. See text for details. **b–c:** Random dependency networks exhibit a binomial distribution of incoming edges (number of neighbors on which any given node depends) and outgoing edges (number of neighbors that depend on any given node). **e–f:** Scale-free networks exhibit a power-law distribution of incoming and outgoing edges.

examine VMM models to determine which changes to network architecture and failure dynamics, if any, are capable of producing a temporal scaling of lifespan distributions.

5.2.1 Homogeneous interventions

We find that most changes to VMM network architecture and failure dynamics, including changes to the number of nodes, the initial number of failed nodes, or the number of dependencies per node, either do not produce a substantial effect on lifespan or alter lifespan distributions in a manner not consistent with temporal scaling (Fig. 5-3). In contrast, changes to the rate of intrinsic failure or repair of all nodes by the same factor produced a temporal rescaling of mortality statistics (Fig. 5-3). Regardless of how interdependent the nodes might be, a universal change of intrinsic failure rates simply speeds up or slows down the overall system dynamics. This temporal rescaling of system dynamics then produces a temporal rescaling of mortality statistics.

These results were consistent across both random and scale-free network architectures, and across multiple random permutations of network dependencies.

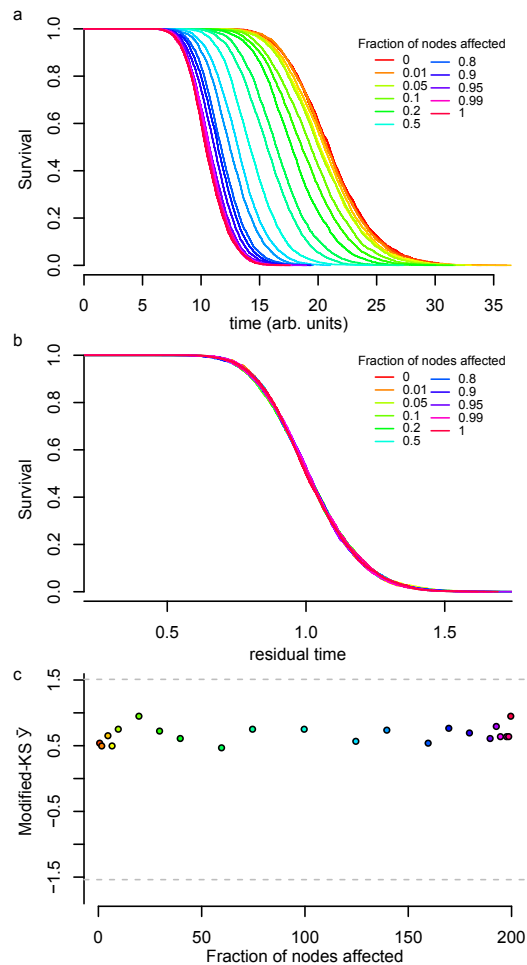


Supplementary Figure 5-3: Temporal scaling in interdependency networks. We recapitulate the perturbations by Vural *et al.* [22] for the purpose of studying their effect on temporal scaling. Populations were simulated by generating 1,500 random networks of the same size with the same underlying distribution of node degrees and with all nodes sharing the same rate parameters. Each network underwent a continuous-time Monte Carlo analog of the dynamics defined in Vural *et al.*. Nodes fail (or are repaired) with a constant probability and subject to the majority rule described in the text. The lifespan of the network is the time it takes for all nodes to fail. **a:** Changing the number of nodes that constitute individuals did not alter the average lifespan of populations. **b:** However, it does alter the characteristic shape of population survival curves as can be seen in the AFT-residuals. **c:** All changes to the network size yielded a KS test statistic greater than 1.51 (grey lines), corresponding to a statistically significant ($p < 0.01$) deviation from perfect temporal scaling. **d:** Increasing the fraction of (randomly chosen) nodes that have failed at the outset did shorten the average lifespan of the population. **e–f:** Yet, it produced significant deviations from perfect scaling. **g:** Increasing the interdependence within networks only marginally increased lifespan. **h–i:** This intervention showed no significant deviation from temporal scaling. **j:** Changing the net failure rate of each node altered lifespan. **k–l:** As expected, such a homogeneous intervention produced perfect temporal scaling.

5.2.2 Heterogeneous interventions

We asked whether modifying the failure rate of only a subset of nodes would break temporal scaling. Such a change might mimic interventions in *C. elegans* that act only on a subset of lifespan determinants. We found that a two-fold increase in the failure rate of any fraction of

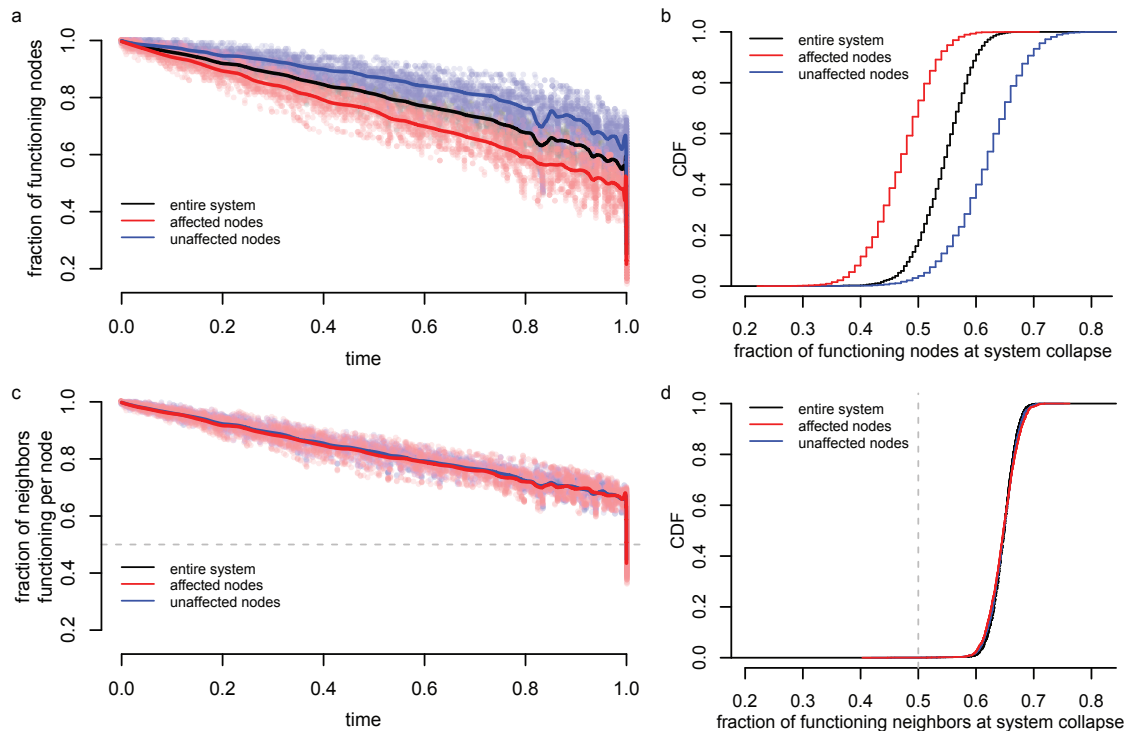
nodes altered lifespan distributions in a manner consistent with temporal scaling, with the magnitude of the scale factor depending on the fraction of nodes affected (Fig. 5-4). These results were consistent across both random and scale-free network architectures, and across multiple random permutations of network dependencies.



Supplementary Figure 5-4: Scaling of heterogeneous interventions in dependency networks. We investigated the effect of interventions that alter only a subset of all nodes. In populations of 1,500 random networks, we increased the failure rate of a subset of nodes, thus “weakening” this subset. **a:** Not surprisingly, increasing the fraction of weakened nodes decreased network lifespan. **b, c:** Very surprisingly, the AFT residuals and the generalized KS test statistic indicated no significant deviations from temporal scaling ($p < 0.01$, gray lines)

We sought to explain this surprising result using the model introduced in the main text: that diverse interventions produce temporal scaling because they influence a single system-wide state variable that determines the risk of death. In VMM networks, networks collapse when the fraction of nodes still operational (henceforth the operational node fraction) drops below a minimum threshold. At the start of each simulation, this operational node fraction begins high and then exhibits an almost linear decline for a long period, decreasing steadily until a catastrophic failure [22] occurs. The operational node fraction, therefore, makes an obvious candidate for a state variable with which to describe the risk of death (network collapse). A reliability model [24] posits that aging is a decrease in the number of redundant system

components over time; perhaps older *C. elegans* have fewer functioning parts.



Supplementary Figure 5-5: State variables in dependency networks. **a:** The panel shows the dynamics of the fraction of functioning nodes (operational node fraction) in the entire system (black), the weakened subnetwork (red) and the non-weakened subnetwork (blue). All times were plotted as a fraction of lifespan, allowing the results of 500 simulations to be overlaid (colored circles) and summarized by a LOESS regression model (lines). **b:** As in panel a, but the quantity shown is the fraction of functional links per node (operational link fraction). All three subnetworks are indistinguishable with regard to this quantity. **c,d:** To better characterize the diversity within a population of the quantities tracked in panels a and b, we collected their value immediately before system collapse, leading to the death of the individual network, and plotted their cumulative distribution function.

Despite its intuitive appeal, we find that the operational node fraction is not the best state variable with which to describe the dynamics of VMM network aging. In Fig. 5-5, we show individual trajectories of the state variable that correspond to the lifespan distributions in Fig. 5-4, where half of the nodes operate with a two-fold increased failure rate. We find that the operational node fraction calculated specifically for the subsets of “weak” and “strong” nodes diverge over time. The operational node fraction of the entire network, as a state variable describing network dynamics, does not capture this divergence. Furthermore, though all networks exhibited a low operational node fraction immediately prior to network collapse, this fraction was highly variable, limiting its predictive value with regard to the timing of network collapse.

A different systemic property, the average fraction of dependencies remaining operational across all nodes (henceforth the operational link fraction), appears to be better suited. In VMM networks, a node fails when the fraction of nodes it depends on drops below 50% of its initial value. The operational link fraction might therefore be a useful state variable with which to characterize the rate of dependency-induced failure events in a network. We find that the

operational link fractions calculated for the subsets of “weak” and “strong” nodes do not diverge over time (Fig. 5-5d). While “weak” and “strong” nodes differ in their intrinsic failure probabilities, they are equally likely to have “weak” or strong” neighbors and therefore both groups must exhibit equivalent operational link fractions at all times. This makes for a systemic property that is consistent across all subnetworks. It also exhibits much less variability immediately prior to network collapse than the operational node fraction. This suggests that the operational link fraction is more predictive of lifespan than the operational node fraction (Fig. 5-5d).

5.3 Conclusions

In this study of dependency networks, we accomplish three tasks. First, we provide an example of how a single state variable, such as the operational link fraction, can describe the state of a complex system and predict its risk of collapse. Second, we illustrate that such a state variable can have a simple dynamics governed by a single time scale despite emerging from a complex network of heterogeneous components. Finally, we highlight one way in which interventions acting only on a subset of a system might produce a temporal scaling in the mortality statistics.

In our simulations of interdependency networks, temporal scaling phenomena arise from the basic assumption that nodes remain operational as long as *any* 50% of their dependencies remain operational. Because the dependencies of a node are fungible and contribute equivalently to a node’s operation, the consequences of local perturbations can be distributed across the whole network. We do not know how such a homogeneity might arise in *C. elegans*.

Note 6 — Vitality/Resilience models based on diffusion

Summary

In this section, we explore how our experimental findings relate to a specific type of stochastic process often invoked to describe aging phenomenologically: random walks and diffusion in a physiological state space [1, 15]. In the context of these models, we explore the constraints that temporal scaling places on the effects of interventions on model parameters. A concrete process model also enables us to explore the sensitivity of temporal scaling to violations of these constraints and to gain intuition on how large populations would have to be to detect such violations with the statistical tools that we used to analyze our data.

Contents of Note 6

6.1	Semi-infinite random walk	30
6.2	Drift-diffusion process	32
6.3	Statistical power and sensitivity of scaling in drift-diffusion processes	33
6.4	Drift-diffusion process with drift heterogeneity	35
6.5	Strehler-Mildvan with vitality drift-diffusion	35

6.1 Semi-infinite random walk

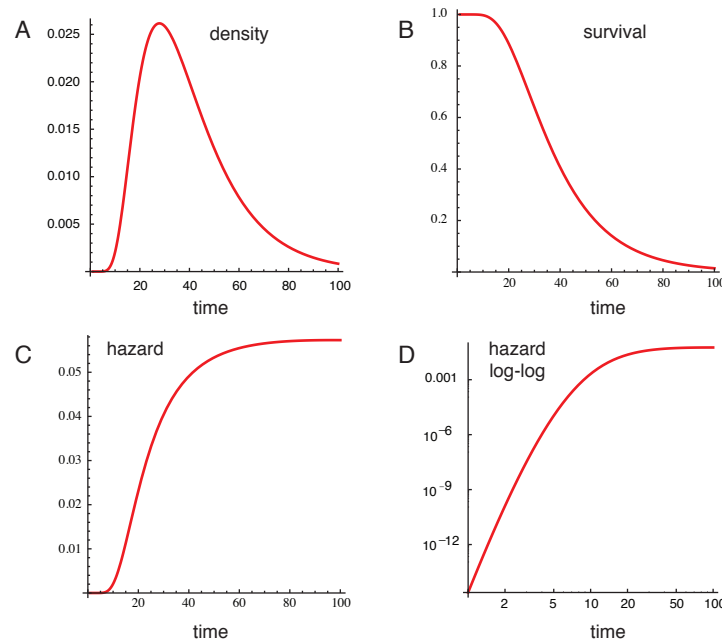
The aging process $X(t)$ considered here is continuous in time t and has infinitely many discrete states $i = 0, 1, 2, \dots$ with an absorbing boundary at 0. At time $t = 0$ the random walker starts in state i : $X(0) = i$. The probabilities of moving up or down the state space are β_1 and β_{-1} , respectively, independent of i . If $\beta_{-1} > \beta_1$, absorption occurs with certainty.

The lifespan distribution (probability density of absorption times), $p(X(t) = 0 | X(0) = i)$ is given by [1, 25] (Fig. 6-1):

$$p(X(t) = 0 | X(0) = i) = i \left(\frac{\beta_{-1}}{\beta_1} \right)^{\frac{i}{2}} \frac{1}{t} \exp(-(\beta_{-1} + \beta_1)t) I_i(2t\sqrt{\beta_{-1}\beta_1}), \quad t > 0 \quad (28)$$

with $I_i(x)$ the modified Bessel function (of the first kind) of order i , for which an integral representation is $I_i(x) = (1/\pi) \int_0^\pi \exp(x \cos \theta) \cos(i\theta) d\theta$.

To rescale time in the survival curve requires multiplying *each* rate parameter β_{-1} and β_1 with the same factor λ^{-1} . Consider the survival function $S(t)$:



Supplementary Figure 6-1: **A:** Lifespan density: $l(t) \equiv p(X(t) = 0 | X(0) = 20)$ with $\beta_{-1} = 1.5$ and $\beta_1 = 1$, as per equation (28). **B:** Survival: $S(t) \equiv 1 - \int_0^t p(X(\tau) = 0 | X(0) = 20) d\tau$. **C:** Hazard: $l(t)/S(t)$. The hazard decelerates (in fact, asymptotes to a constant) because the process attains a quasi-stationary distribution over the states $j > 0$, $q_j = \lim_{t \rightarrow \infty} p(X(t) = j | X(t) > 0, X(0) = i)$, with $q_1 = (1 - \sqrt{\beta_1/\beta_{-1}})^2$ and $q_j = q_1 j (\beta_1/\beta_{-1})^{(j-1)/2}$, $\forall j > 1$. As $t \rightarrow \infty$, the hazard asymptotes to $h(\infty) = \beta_{-1}q_1 = (\sqrt{\beta_{-1}} - \sqrt{\beta_1})^2$, which, for the values used in this example, is .0505... **D:** Log-log plot of **C**. If we were to multiply β_{-1} and β_1 by 2, the survival curve **B** would look exactly the same up to a rescaling of the time axis by 1/2.

$$S_1(t) = 1 - \int_0^t p(X(\tau) = 0 | X(0) = i) d\tau \quad (29a)$$

$$= 1 - i \left(\frac{\lambda^{-1}\beta_{-1}}{\lambda^{-1}\beta_1} \right)^{\frac{i}{2}} \int_0^t \frac{1}{\tau} \exp(-(\beta_{-1} + \beta_1)\lambda^{-1}\tau) I_i(2\sqrt{\beta_{-1}\beta_1}\lambda^{-1}\tau) d\tau \quad (29b)$$

integration by substitution: set $\tau' = \lambda^{-1}\tau$, $d\tau' = \lambda^{-1}d\tau \dots$

$$= 1 - i \left(\frac{\beta_{-1}}{\beta_1} \right)^{\frac{i}{2}} \int_0^{\lambda^{-1}t} \frac{1}{\lambda\tau'} \exp(-(\beta_{-1} + \beta_1)\tau') I_i(2\sqrt{\beta_{-1}\beta_1}\tau') \lambda d\tau' \quad (29c)$$

$$= 1 - i \left(\frac{\beta_{-1}}{\beta_1} \right)^{\frac{i}{2}} \int_0^{\lambda^{-1}t} \frac{1}{\tau} \exp(-(\beta_{-1} + \beta_1)\tau) I_i(2\sqrt{\beta_{-1}\beta_1}\tau) d\tau \quad (29d)$$

$$= S_0(\lambda^{-1}t) \quad (29e)$$

Likewise for the probability density (28):

$$p_1(X(t) = 0 | X(0) = i) = \lambda^{-1}p_0(X(\lambda^{-1}t) = 0 | X(0) = i).$$

6.2 Drift-diffusion process

The random walk with continuous state space is a drift-diffusion process with initial value c , drift coefficient $-\mu$ ($\mu > 0$), and diffusion coefficient σ^2 :

$$X(t) = c - \mu t + \sigma W(t), \quad (30)$$

where $W(t)$ is a Wiener process, defined by $W(0) = 0$ and increments that are independent and normally distributed with $E\{W(s+t) - W(t)\} = 0$ and $\text{Var}\{W(s+t) - W(t)\} = s$. In terms of increments: $P(W(s+t) \in [y, y+dy] | W(t) = x) = N((y-x)/\sqrt{t}, 0, 1)/\sqrt{t}$, where $N(x, 0, 1)$ is the density of the standard normal distribution. Eqn. (30) is sometimes written as a stochastic differential equation, $dX(t) = -\mu dt + \sigma dW(t)$, with initial condition $X(0) = c$.

The lifespan distribution $p(X(t) = 0 | X(0) = c)$ is an inverse Gaussian [1]:

$$p(X(t) = 0 | X(0) = c) = \frac{c}{\sigma\sqrt{2\pi}} \frac{1}{t^{3/2}} \exp\left[-\frac{(c-\mu t)^2}{2\sigma^2 t}\right]. \quad (31)$$

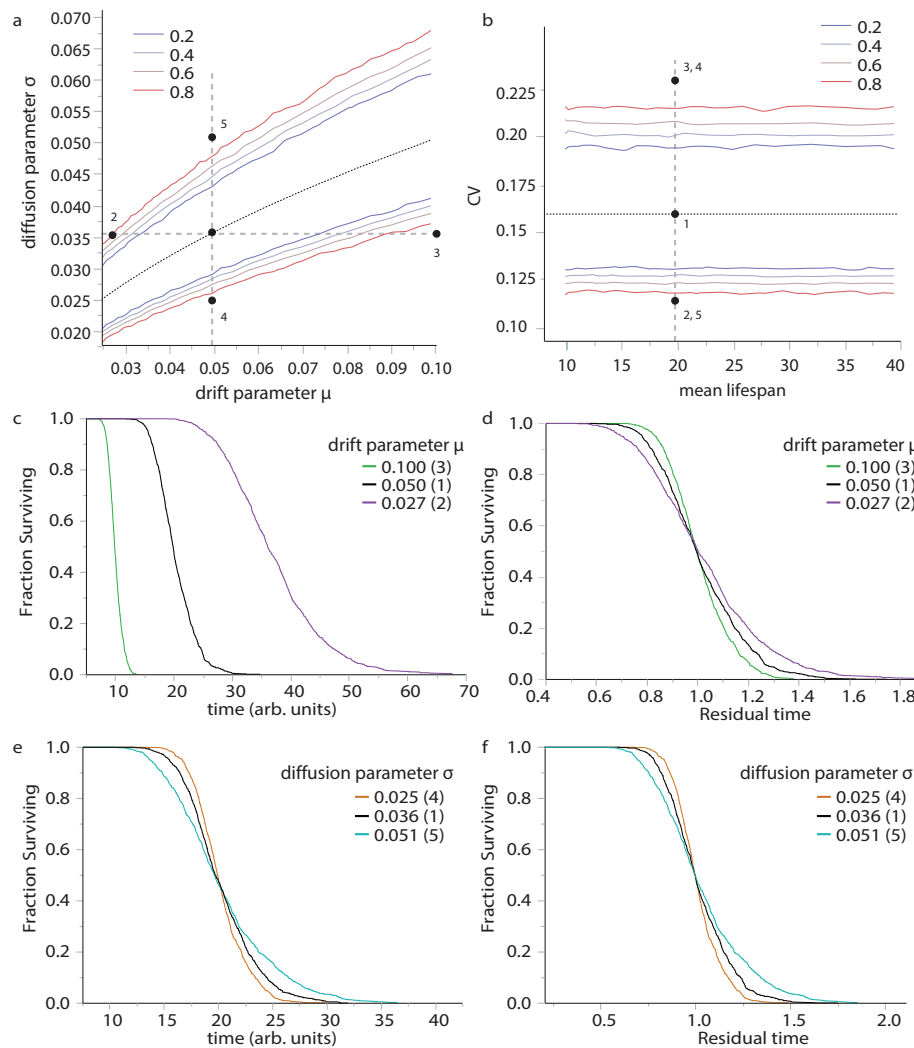
The survival function is $S(t) = 1 - \int_0^t p(X(t) = 0 | X(0) = c)$:

$$S(t) = \Phi\left(\frac{c-\mu t}{\sigma\sqrt{t}}\right) - \exp\left(\frac{2c\mu}{\sigma^2}\right) \Phi\left(\frac{-c-\mu t}{\sigma\sqrt{t}}\right) \quad (32)$$

with $\Phi(x)$ the cumulative standard normal distribution, $\Phi(x) = (1/\sqrt{2\pi}) \int_{-\infty}^x \exp(-t^2/2) dt$.

Rescaling time means multiplying μ and σ^2 with the same factor λ^{-1} :

$$S_1(t) = \Phi\left(\frac{c-\mu\lambda^{-1}t}{\sigma\sqrt{\lambda^{-1}t}}\right) - \exp\left(\frac{2c\mu}{\sigma^2}\right) \Phi\left(\frac{-c-\mu\lambda^{-1}t}{\sigma\sqrt{\lambda^{-1}t}}\right) = S_0(\lambda^{-1}t). \quad (33)$$



Supplementary Figure 6-2: Detecting deviations from perfect scaling. **a:** The plot shows the parameter combinations (μ, σ) for which we can detect significant deviation from scaling relative to a reference set ($\mu_0 = 0.049, \sigma_0 = 0.036$) with a given statistical power (color) for population size 1,000. See text for details. Contour colors indicate increasing power. The area between the blue (inner) contours, corresponds to parameter combinations with a less than 20% probability, at $p < 0.01$, of detecting a deviation from scaling using our approach. Parameter combinations outside the red (outer) contours have a greater than 80% probability, at $p < 0.01$, of being detected. The dotted curve plots the line along which populations exhibit a coefficient of variation equal to that of the reference population (point 1) **b:** We plot the same results in terms of the coefficient of variation of lifespan vs the mean lifespan $1/\mu$, instead of directly the process parameters. (Scaling implies a constant coefficient of variation.) **c:** The panel shows survival curves corresponding to the marked positions 1 and 2 in panels a–b. **d:** The panel shows the corresponding AFT residuals on which we ran the Kolmogorov-Smirnov test. **e–f:** As in panels c and d, but for the marked positions 3 and 4 in panels a–b.

6.3 Statistical power and sensitivity of scaling in drift-diffusion processes

The drift-diffusion process we describe in section 6.2 exhibits a first passage time distribution that is the so-called inverse Gaussian distribution, equation (31). Without loss of generality, we

can fix $c = 1$. The mean lifespan then is given by μ^{-1} , its variance is σ^2/μ^3 , and its coefficient of variation is $\sigma/\sqrt{\mu}$. One can see that, to produce temporal scaling, a change from μ to μ/λ must be accompanied by a corresponding change from σ to $\sigma/\sqrt{\lambda}$.

The extent to which our combined experimental and statistical approach can constrain drift-diffusion models depends on its capacity to detect deviations from scaling, which in turn depends on the size of the changes in each parameter and the size of the populations used to detect them. In other words, the issue is not to detect just changes in, say, μ (a wilcoxon or log-rank will work well even with relatively small populations of a hundred individuals), but to detect *disproportionate* changes in μ relative to a given σ . This requires much larger population sizes. Figure 6-2 shows the results for a population size of 1,000, which corresponds to our typical experiments.

In Figure 6-2, the parameters μ and σ always refer to the parameters of the drift and diffusion component of the *process*. (The mean and the standard deviation of the inverse Gaussian then are $1/\mu$ and $1/\sigma$.) We generated a reference sample (μ_0, σ_0) of 1,000 death times sampled from an inverse Gaussian probability density with (process) parameters $\mu = 0.049$ and $\sigma = 0.036$. These parameters were obtained by fitting an inverse Gaussian to our experimental data to obtain realistic parameters. (This parametric form is not the best fit for our experimental data—we prefer a Weibull distribution; but the inverse Gaussian provides a reasonable fit, as seen in Extended Data Item 1a–d.) We then generated sets (μ, σ) of 1,000 death times each, sampled from inverse Gaussians whose μ and σ parameters were systematically varied over a range ($\mu \in [0.025 - 0.1]$ and $\sigma \in [0.017 - 0.071]$). We compared each parameter setting (μ, σ) to (μ_0, σ_0) using the AFT+KS combination (section 2.1) to calculate a p -value with which each pair (μ_0, σ_0) and (μ, σ) could be said to exhibit a significant deviation from perfect scaling at a significance level 0.01. By generating 750 replicates of each comparison, we determined the fraction of significant p -values, which provides us with an estimate of the statistical power at population size 1,000, and thus the effect sizes that can be distinguished.

Because we apply the same statistical test to evaluate temporal scaling in both simulated models and experimental data, figure 6-2 can be used to assess our ability to draw strong inferences from our data. The only assumption we make here is that *C. elegans* lifespan distributions can be reasonably approximated by inverse Gaussian distributions.

Given a population of 1,000 individuals, we ask how much the drift parameter μ of a drift-diffusion process must change, at constant σ , for our statistical approach to reliably detect deviations from perfect scaling at a significance 0.01 and power 0.8. The answer from Fig. 6-2 is that the drift parameter must be altered to shorten or extend lifespan by about two-fold or more. Tests for smaller effects will yield a higher rate of type II errors, and thus be less reliable. Increasing population sizes to 5,000 individuals (requiring heroic efforts in the lab) would allow one to identify deviations from perfect scaling resulting from a constant σ when lifespan is altered only 0.5-fold.

This simulation suggests that, were aging determined by a Wiener process with drift, a fixed σ can be reasonably excluded for interventions that have a large effect on lifespan, for example exposure to tBuOOH at 1.5 mM, exposure to live bacteria (main Fig. 1 g-h), exposure to temperatures above 27 °C (main Fig. 1.b,c), and the *daf-2(e1368)* allele. For interventions with smaller effects on lifespan, data obtained with the population sizes we consider serve to limit the magnitude with which interventions affect σ . For example, in the case of *hif-1(ia4)*, our data would constrain the ratio of σ and μ in mutants to within 0.6 and 1.4 of wildtype. If this does not seem impressive, remember that our data also exclude other types of change to the aging process that might occur outside the Wiener process model (for example a disproportionate change in one risk determinant

as in Fig. 5-1d-f that produces a more prominent change in the survival distribution, e.g. a deviation from the hazard function associated with an inverse Gaussian lifespan distribution).

This all serves to illustrate a more general point: when scaling interventions produce large effects on lifespan (e.g. a 50% increase or decrease in a population comprising more than 1,000 individuals), they can be used positively to suggest that the lifespan extension itself emerges from a simultaneous effect on multiple aspects of the aging process, σ and μ in the present model or all risk factors of equation (27). Interventions with smaller effects on lifespan are best interpreted as providing limits on the magnitude of any disproportionate effect across the aging process, such as causing a large change in σ without a corresponding change in μ , or adding or removing a risk factor in equation (27).

6.4 Drift-diffusion process with drift heterogeneity

Here we consider the diffusion process of the previous section, but introduce heterogeneity in the drift coefficient μ . Let $p_\mu(V(t) = v)$ be the probability of being in state $v > 0$ given a drift parameter $-\mu$ chosen from a normal distribution with mean $\bar{\mu} > 0$ and deviation σ_μ :

$$p_\mu(X(t) = v) = \int_{-\infty}^{\infty} p(X(t) = v | \mu) p(\mu) d\mu. \quad (34)$$

A tedious calculation gives

$$p_\mu(X(t) = v) = \frac{1}{\sqrt{2\pi t}(\sigma^2 + \sigma_\mu^2 t)} \left[1 - \exp\left(-\frac{2cv}{\sigma^2 t}\right) \right] \exp\left(-\frac{(v - c + \bar{\mu}t)^2}{2t(\sigma^2 + \sigma_\mu^2 t)}\right) \quad (35)$$

Integrating over $v > 0$ yields the survival curve of the drift-diffusion process with drift heterogeneity:

$$\begin{aligned} S(t) &= \int_0^\infty p_\mu(X(t) = v) dv = \\ &= \Phi\left(\frac{c - \bar{\mu}t}{\sqrt{t(\sigma^2 + t\sigma_\mu^2)}}\right) - \Phi\left(-\frac{c\sigma^2 + t\bar{\mu}\sigma^2 + 2ct\sigma_\mu^2}{\sigma^2\sqrt{t(\sigma^2 + t\sigma_\mu^2)}}\right) \\ &\quad \times \exp\left(-\frac{\sigma^2\sigma^2(c - t\bar{\mu})^2 + (c\sigma^2 + t\sigma^2\bar{\mu} + 2ct\sigma_\mu^2)^2}{t\sigma^2\sigma^2(\sigma^2 + t\sigma_\mu^2)}\right) \end{aligned} \quad (36)$$

Again, temporal scaling obtains when the parameters σ^2 , $\bar{\mu}$, and σ_μ are scaled with λ^{-1} . Note that this agrees with our competing risk simulations in Fig. 5-1, where we present the phenomenon quantitatively.


6.5 Strehler-Mildvan with vitality drift-diffusion

In 1960, Strehler and Mildvan proposed a now-prominent model of aging [26, 27] that posits a separation of the aging process from the proximate causes of death. The former generates the (putatively) linear loss of vitality (represented by the physiological state variable $X(t)$ in the above processes), while the latter are events sampled from an age-independent distribution of (internal or external) challenges that cause death when their magnitude exceeds the vitality of the process at a given time.

Specifically, the waiting time τ between challenges has exponential density, $p(\tau) = \gamma \exp(-\gamma\tau)$. Thus, $P(\tau \geq t) = \exp(-\gamma t)$, $t \geq 0$. The magnitude m of a challenge is also exponential with parameter ν : $p(m) = \nu \exp(-\nu m)$ and $P(m \geq M) = \exp(-\nu M)$. If the vitality process $V(t)$ is deterministic, these assumptions imply for the hazard $h(t)$ [27]:

$$h(t) = \gamma P(m \geq V(t)). \quad (37)$$

That is, $h(t)$ is the average rate at which shocks occur (γ) times the probability that a shock kills.



Supplementary Figure 6-3: Hazard function of Strehler-Mildvan model with vitality drift-diffusion. **A:** The plots show equation (41) as a function of time and ν (blue: $\nu = 1$; red: $\nu = 0.1$; green: $\nu = 0.001$). Parameters: $c = 20$, $\mu = 1.5$, $\sigma^2 = 1$, $\gamma = 1$. **B:** The lin-log plots illustrate the decidedly Gompertzian flavor of the model due to the imposed shock statistics.

We generalize equation (37) for the case in which vitality $X(t)$ is undergoing a drift-diffusion process as in section 6.2:

$$h(t) = \gamma \int_0^\infty P(m \geq v | X(t) = v) p(X(t) = v) dv. \quad (38)$$

The probability density that a drift-diffusion process that started at location c is at location $v > 0$ at time t is given by [1]:

$$p(X(t) = v) = \frac{1}{\sigma\sqrt{2\pi t}} \left\{ \exp\left(-\frac{(v-c+\mu t)^2}{2\sigma^2 t}\right) - \exp\left(\frac{2c\mu}{\sigma^2}\right) \exp\left(-\frac{(v+c+\mu t)^2}{2\sigma^2 t}\right) \right\}. \quad (39)$$

By integrating eqn (39) over v we obtain the probability density that the process has not yet been absorbed in state 0, which is the inverse Gaussian survival curve (32); $p(X(t) = v)$ loses probability into the absorbing state and we therefore need to have a separate term accounting for the probability accumulating in the absorbing 0 state. As usual, the integral can be expressed in

terms of cumulative standard normals (error functions):

$$h(t) = \gamma \int_0^{\infty} P(m \geq v | X(t) = v) p(X(t) = v) dv \quad (40)$$

$$= \gamma \left[\exp\left(\frac{1}{2}\nu(-2c + 2\mu t + \nu\sigma^2 t)\right) \times \left\{ \Phi\left(\frac{c - (\mu + \nu\sigma^2)t}{\sigma\sqrt{t}}\right) - \exp\left(\frac{2c(\mu + \nu\sigma^2)}{\sigma^2}\right) \Phi\left(-\frac{c + (\mu + \nu\sigma^2)t}{\sigma\sqrt{t}}\right) \right\} + 1 - \Phi\left(\frac{c - \mu t}{\sigma\sqrt{t}}\right) + \exp\left(\frac{2c\mu}{\sigma^2}\right) \Phi\left(-\frac{c + \mu t}{\sigma\sqrt{t}}\right) \right]. \quad (41)$$

Therefore, for temporal scaling to occur in Strehler-Mildvan models, all parameters— μ , σ^2 , and also the hit rate γ , but not ν , the parameter of the exponential density of hit magnitudes—must be multiplied with λ^{-1} .

References

- [1] Aalen, O., Borgan, O. & Gjessing, H. *Survival and Event History Analysis : A Process Point of View* (New York, NY : Springer New York, 2008).
- [2] Cox, D. Regression Models and Life-Tables. *Journal of the Royal Statistical Society, Series B* **34**, 187–220 (1972).
- [3] Atlan, H., Mique, I. J., Helmle, L. & Dolkas, C. Thermodynamics of aging in *Drosophila melanogaster*. *Mechanisms of Ageing and Development* **5**, 371–387 (1976).
- [4] Mair, W., Goymer, P., Pletcher, S. & Partridge, L. Demography of dietary restriction and death in *Drosophila*. *Science* **301**, 1731–1733 (2003).
- [5] Wu, D., Rea, S., Cypser, J. & Johnson, T. Mortality shifts in *Caenorhabditis elegans*: remembrance of conditions past. *Aging Cell* **8**, 666–675 (2009).
- [6] Johnson, T. E., Wu, D., Tedesco, P., Dames, S. & Vaupel, J. W. Age-specific demographic profiles of longevity mutants in *Caenorhabditis elegans* show segmental effects. *J Gerontol A (Biol Sci Med Sci)* **56**, B331–339 (2001).
- [7] Samuelson, A., Carr, C. & Ruvkun, G. Gene activities that mediate increased life span of *C. elegans* insulin-like signaling mutants. *Genes Dev.* **21**, 2976–94 (2007).
- [8] Stroustrup, N. *et al.* The *Caenorhabditis elegans* Lifespan Machine. *Nature Methods* **10**, 665–70 (2013).
- [9] Vanfleteren, J., De Vreese, A. & Braeckman, B. Two-parameter logistic and Weibull equations provide better fits to survival data from isogenic populations of *Caenorhabditis elegans* in axenic culture than does the Gompertz model. *J Gerontol A Biol Sci Med Sci* **53**, B393–403; discussion B404–398 (1998).
- [10] Fleming, T. R., O’Fallon, J. R., O’Brien, P. C. & Harrington, D. P. Modified kolmogorov-smirnov test procedures with application to arbitrarily right-censored data. *Biometrics* **36**, 607–625 (1980).

- [11] Vaupel, J., Manton, K. & Stallard, E. The impact of heterogeneity in individual frailty on the dynamics of mortality. *Demography* **16**, 439–454 (1979).
- [12] Vaupel, J. W. *et al.* Biodemographic trajectories of longevity. *Science* **280**, 855–860 (1998).
- [13] Baeriswyl, S. e. a. Modulation of aging profiles in isogenic populations of *Caenorhabditis elegans* by bacteria causing different extrinsic mortality rates. *Biogerontology* **1**, 53–65 (2010).
- [14] Aalen, O. & Gjessing, H. K. Understanding the Shape of the Hazard Rate: A Process Point of View. *Statistical Science* **16**, 1–22 (2001).
- [15] Weitz, J. & Fraser, H. Explaining mortality rate plateaus. *Proc Natl Acad Sci USA* **98**, 15383–15386 (2001).
- [16] Steinsaltz, D. & Evans, S. Quasistationary Distributions For One-Dimensional Diffusions With Killing. *American Mathematical Society* **359**, 1285–1324 (2007).
- [17] Wu, D., Rea, S., Yashin, A. & Johnson, T. Visualizing hidden heterogeneity in isogenic populations of *C. elegans*. *Exp Gerontol.* **3**, 261–70 (2006).
- [18] Garigan, D. *et al.* Genetic analysis of tissue aging in *Caenorhabditis elegans*: a role for heat-shock factor and bacterial proliferation. *Genetics* **161**, 1101–1112 (2002).
- [19] Riesen, M. *et al.* MDL-1, a growth- and tumor-suppressor, slows aging and prevents germline hyperplasia and hypertrophy in *C. elegans*. *Aging* **2**, 98–117 (2014).
- [20] Kalbfleisch, J. D. & Prentice, R. L. *The Statistical Analysis of Failure Time Data. Second Edition.* Wiley Series in Probability and Statistics (Wiley, 2002).
- [21] Rinne, H. *The Weibull Distribution: A Handbook* (Chapman and Hall/CRC, 2008), 1 edn.
- [22] Vural, D. C., Morrison, G. & Mahadevan, L. Aging in complex interdependency networks. *Phys. Rev. E* **89**, 022811 (2014).
- [23] Gillespie, D. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry* **81**, 2340–2361 (1977).
- [24] Gavrilov, L. & Gavrilova, N. The reliability theory of aging and longevity. *Journal of Theoretical Biology* **213**, 527–545 (2001).
- [25] Heathcote, C. R. & Moyal, J. E. The Random Walk [In Continuous Time] And Its Application To The Theory Of Queues. *Biometrika* **46**, 400–411 (1959).
- [26] Strehler, B. & Mildvan, A. General Theory of Mortality and Aging. *Science* **132**, 14–21 (1960).
- [27] Wagner, P. Vitality heterogeneity in the Strehler-Mildvan theory of mortality. *MPIDR Working Papers* **WP-2011-012** (2011).