

Longitudinal analysis of body weight reveals homeostatic and adaptive traits linked to lifespan in diversity outbred mice

Anil Raj

anil@calicolabs.com

Calico Life Sciences LLC

G.V. Prateek

Calico Life Sciences LLC

Zhenghao Chen

Calico Life Sciences LLC

Kevin Wright

Calico Life Sciences LLC

Andrea Francesco

Calico Life Sciences LLC

Vladimir Jojic

Calico Life Sciences LLC

Gary Churchill

Jackson Laboratories <https://orcid.org/0000-0001-9190-9284>

Article

Keywords:

Posted Date: July 10th, 2024

DOI: <https://doi.org/10.21203/rs.3.rs-4596377/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: Yes there is potential Competing Interest. GVP, ZC, KW, ADF, VJ, and AR were employees of Calico Life Sciences, LLC during the time of their contribution to the study. The authors declare no other competing interests.

Longitudinal analysis of body weight reveals homeostatic and adaptive traits linked to lifespan in diversity outbred mice

G.V. Prateek¹, Zhenghao Chen¹, Kevin Wright¹, Andrea Di Francesco¹, Vladimir Jojic¹,
Gary A. Churchill² and Anil Raj¹

¹Calico Life Sciences LLC, South San Francisco, United States

²The Jackson Labs, Bar Harbor, United States

Thursday 13th June, 2024

Abstract

Dense temporal measurements of physiological health using simple and consistent assays are essential to characterize biological processes associated with aging and evaluate the effectiveness of interventions on these processes. We measured body weight in 960 genetically diverse mice, every 7-10 days over the full course of their lifespan. We used a state space model to characterize the trajectories of body weight throughout life and derived 10 novel traits capturing the dynamics of body weight that are both associated with lifespan and heritable. Genetic mapping of these traits identified 12 genomic loci, none of which were previously mapped to body weight. We observed that the ability to stabilize body weight, despite fluctuations in energy intake and expenditure, is positively associated with lifespan and mapped to a genomic locus linked to energy homeostasis. Our results highlight the importance of dense longitudinal measurements of physiological traits for monitoring health and aging.

Introduction

Aging is characterized by the progressive loss of physiological integrity (López-Otín et al. (2013); Freund (2019); Chen et al. (2022)). Homeostasis, the ability to maintain physiological integrity in response to intrinsic and extrinsic changes, is considered to be an important determinant of aging (Cannon (1929); Moldakozhayev and Gladyshev (2023)). Longitudinal phenotyping throughout the organism's lifespan enables us to measure homeostasis by quantifying temporal

relationships in one or more physiological traits, and associate these measures to healthspan and lifespan. In general, non-invasive procedures are required for longitudinal studies, as they allow for more frequent measurements with minimal stress to the organism. A variety of technologies can be used to collect longitudinal data with high temporal density, including body weight and composition, blood pressure, heart rate, sleep patterns, and frailty (Palliyaguru et al. (2021); Kuo et al. (2022)). Longitudinal data collected at high temporal density requires specialized computational tools to capture the time dynamics of physiological changes.

Body weight can be longitudinally measured at high frequency and changes in body weight have been linked to age-related changes in metabolism, disease outcomes, response to stress, and lifespan (Goodrick et al. (1990); Lissner et al. (1991); Bou Sleiman et al. (2022)). Body weight is influenced by a complex interplay of genetic and environmental factors (Wright et al. (2022)), and one such environmental factor is diet (Yang et al. (2014)). Over-consumption of calories and unhealthy diets can increase the risk of obesity and related diseases (Dixon (2010)). Conversely, a nutritious diet can help maintain healthy body weight and reduce the risk of age-related diseases (Willett et al. (2019)). Studies in model organisms have shown that certain dietary interventions, such as caloric restriction (CR) and intermittent fasting (IF), can extend lifespan and have beneficial health effects such as reduced inflammation, improved metabolic function, and increased cellular repair and regeneration (Fontana et al. (2010); Di Francesco et al. (2018); Longo and Anderson (2022)). However, it is unclear how diet, age and genetics influence the dynamics of body weight trajectories measured throughout life and whether summaries describing these dynamics are associated with lifespan.

We measured the body weight of 960 genetically diverse mice from weaning until death and characterized the dynamics of this classic complex trait throughout life. We developed a computational model for these body weight trajectories and derive novel phenotypes from the dynamics of body weight throughout life. We demonstrated the utility of this model to capture important functional changes that occur with aging such as the ability to maintain stable body weight in response to stress. Additionally, we examined how diet modulates these derived phenotypes as mice age, and explored whether these phenotypes are useful predictors of lifespan. Finally, we identified genetic determinants of these body weight-derived traits and quantified their heritability under different diets. We expect that body weight measured at high temporal density can reveal patterns of change that are associated with longevity and healthy aging.

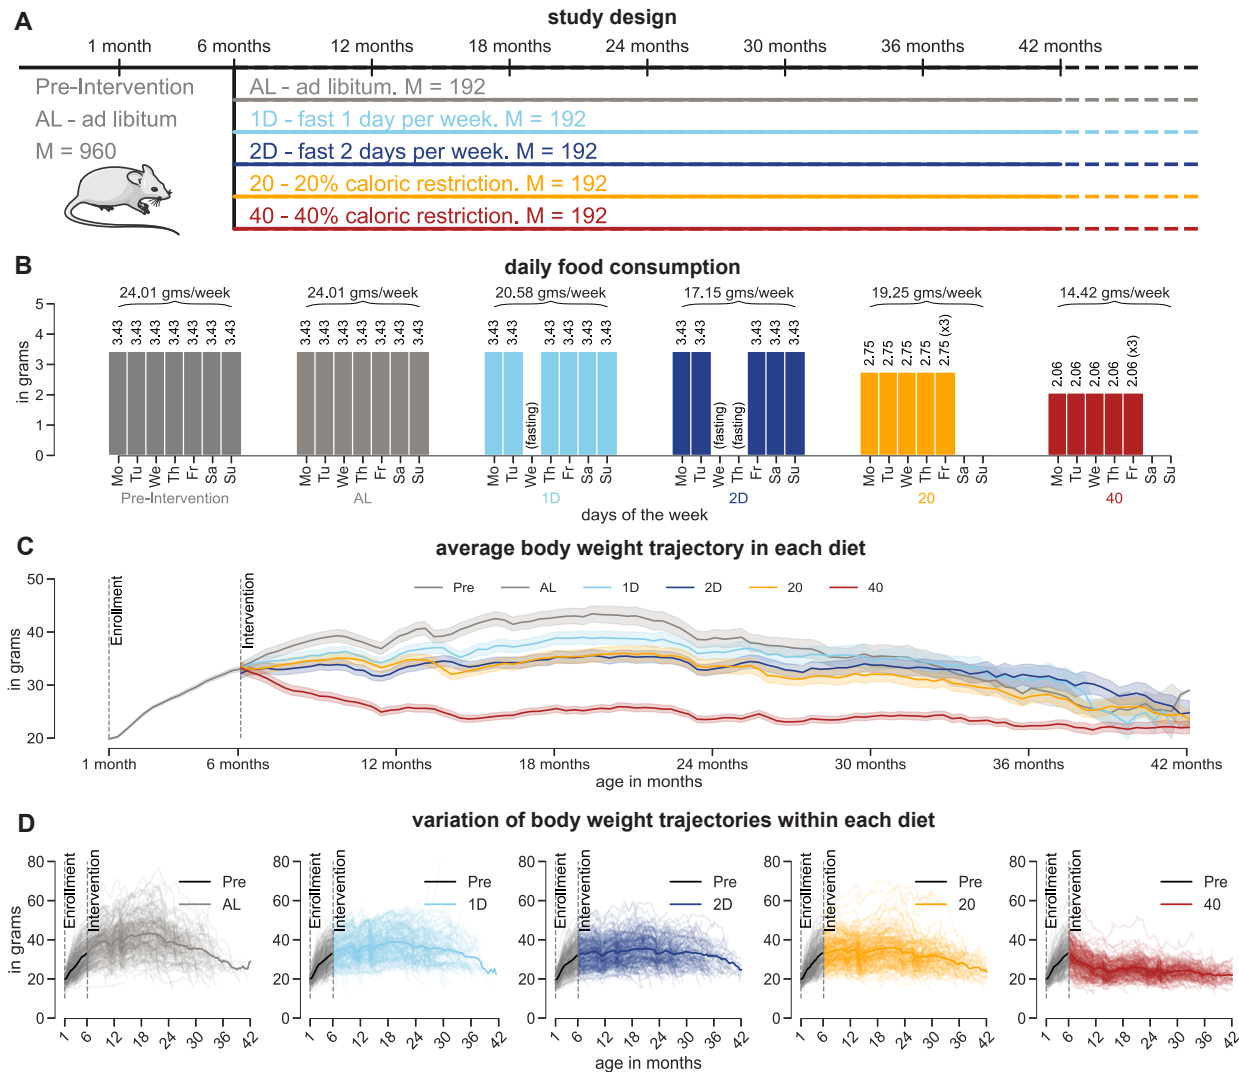


Figure 1: Overview of the study design. (A) At six months age, mice were randomly divided into equal proportions of five dietary groups: ad libitum (AL), one-day fasting (1D), two-day fasting (2D), 20% calorie restricted (20), and 40% calorie restricted (40). Prior to the dietary intervention, mice were on ad libitum diet. (B) Average daily food consumption (in grams) per day of the week in the pre- and post-intervention phases across diet groups. Mice belonging to the 20% and 40% groups were given three times (3x) their daily average intake every Friday. (C) Average body weight (solid lines) and standard errors (shaded region) of mice under different dietary interventions. (D) Body weight trajectories of mice in each dietary group in the pre- and post-intervention phases. The colored solid line represents the average body weight trajectory of mice in each diet.

Results

Outline of study design

Body weights of mice can be easily and accurately measured. They can be altered by diets, diseases, aging, stressors, and environmental conditions. We enrolled 960 diversity outbred (DO) female mice derived from eight inbred founder

strains (AJ, B6, 129, NOD, NZO, CAST, PWK, and WSB) (Svenson et al. (2012)), and used the longitudinally measured body weight, sampled at every 7-10 days over the full course of their lifespan (Francesco et al. (2023)). From enrollment (3 weeks) until 6 months of age, all mice were fed *ad libitum* (AL) on a diet of standard rodent chow (5KOG LabDiet). At 6 months of age, randomly assigned to groups of mice were switched to dietary treatments (192/group): *ad libitum* feeding (AL), 20% calorie restriction (20), 40% calorie restriction (40), fasting 1 day per week (1D), and fasting 2 consecutive days per week (2D) (Figure 1A). Individual body weight measures were taken approximately once a week (Methods).

For mice on the 1D and 2D IF diets, fasting was imposed weekly from Wednesday noon to Thursday noon and from Wednesday noon to Friday noon, respectively. These mice had unlimited food access (similar to AL mice) on their non-fasting days (Figure 1B). For mice on the 20% and 40% CR diets, a predetermined amount of food of 2.75 grams/mouse/day and 2.06 grams/mouse/day, respectively, were given every afternoon from Monday to Thursday, and were triple fed (x3) on Friday afternoon to last until Monday afternoon (Figure 1B). The predetermined amount of food was determined from a previous internal study at the Jackson Laboratory where the average amount of food consumed by female DO mice at 6 months of age was estimated to be 3.43 grams/day.

Modeling the dynamics of body weight

Following the onset of interventions, diet was the largest source of variation in average body weight across the groups (Figure 1C). However, within each diet group, we observed substantial variation across mice and ages (Figure 1D). In order to model the temporal variation in body weight, we developed an autoregressive hidden Markov model (ARHMM) that captures the dynamics of body weight using discrete latent physiological states (Figure 2A).

The ARHMM combines an autoregressive model and a hidden Markov model to model a body weight trace (time-series) as a sequence of latent states. The latent states evolve according to a Markov process and capture the switching dynamics of body weight states over time. The autoregressive model captures the relationship between the current and past observations conditional on the current latent state (Methods). We trained the ARHMM with 70% of body weight traces from each dietary group. To determine the number of latent states of the ARHMM, we computed the deviance information criteria (Spiegelhalter et al. (2002)) on the held-out body weight traces for different numbers of latent states and starting from different random initializations (Supplementary Figure 1A). We obtained the smallest deviance information criteria for a model of order $K = 3$. For the given model order, the random initialization which resulted in the highest evidence-based lower bound (ELBO) on the training set was selected as the best model (Supplementary Figures 1B-1D). Based on the sign and magnitude of the intercept parameters for the three latent states (ϕ_0^1 , ϕ_0^2 , and ϕ_0^3), we named the states as *decline*, *steady*, and *growth* (Supplementary Table 1).

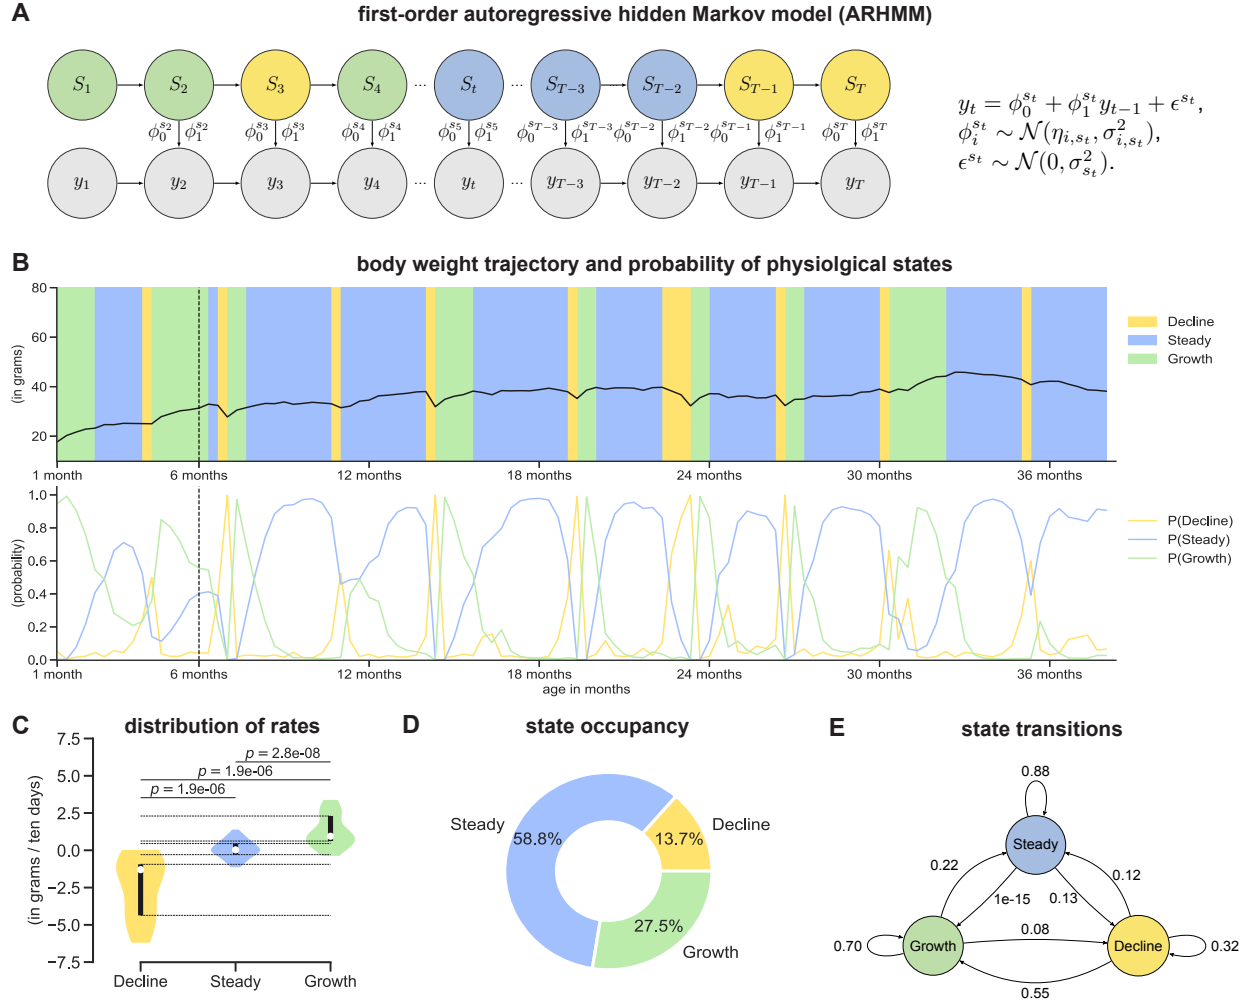


Figure 2: ARHMM and body weight-derived phenotypes. (A) A graphical illustration of an autoregressive hidden Markov model (ARHMM). (B) Sample body weight trajectory of a mouse on a one-day fasting diet. The start time of the intervention is indicated with a dashed vertical line. The underlying physiological states (growth, steady, and decline) are represented using background colors (top panel). The posterior probability trace of a physiological state for the corresponding body weight trajectory and its inferred dynamics (bottom panel). (C) Distribution of the rates conditioned on the physiological states. (D) The percentage of time spent in each physiological state. (E) Empirical state transition probabilities computed using the inferred states for the sample mouse.

For each body weight trace, at each time point, we inferred the posterior probability that the mouse belonged to one of these three physiological states, using the best ARHMM model, and assigned the state with the highest posterior probability to that time point (Figure 2B). We found little overlap in the inter-quartile range of estimated rates of change of body weight within each state, indicating that these states can be well-resolved throughout the lifespan of a mouse (Figure 2C). Our inferences under the ARHMM model allowed us to derive several state-related traits, such as *state occupancy* and *state transitions*. For every mouse, the *state occupancy* was defined as the percentage of time the mouse spent in each state (Figure 2D). Similarly, the *state transition* was defined as the frequency with which the mouse

switched from one state to another (**Figure 2E**). We accounted for uncertainties in state assignments when deriving these traits by using a weighted average approach, where the posterior probabilities of the state assignments were used as weights. In the next sections, we explore the effects of age, diet, and genetics on these traits and how informative they are in predicting mouse lifespan.

Calorie restriction improves body weight homeostasis throughout life

Homeostasis is the maintenance of physiological processes in response to intrinsic and extrinsic changes experienced by an organism (Billman (2020)). Body weight homeostasis is defined as a state of unchanging body weight in response to environmental perturbations (Ravussin et al. (2014)). Modeling the dynamics of body weight allows us to generalize the definition of body weight homeostasis to include *maintenance of constant body weight dynamics*. We hypothesize that homeostasis of body weight dynamics can be influenced by a combination of various factors including diet, genetics, energy imbalance, and environmental stress.

To characterize homeostasis of body weight dynamics across diets as mice age, we divided the time-axis into six-month non-overlapping intervals starting from birth to 42 months. For each age bin, we defined *steady state homeostasis* as the fraction of time spent in the steady state (**Figure 3A**, left panel). We observed that, on average, mice spent over 60% of time post-intervention in steady state homeostasis. However, we observed large variation in steady state homeostasis across diets; on average, over a 6-month period, 40% CR mice could maintain body weight for nearly 4.5 months while AL mice could do so for just 3.5 months. When we changed the time-axis from six-month interval bins to ten percent interval bins of proportion of life lived, we observed that 40% CR mice spent nearly 75% of their time in steady state homeostasis into the last decile of life (**Figure 3B**, left panel). We could redefine steady state homeostasis in terms of the empirical probability of not transitioning out of steady state. The higher this probability, the more stable the steady state is in the mouse. We observed similar trends in the effects of age and diet on steady state homeostasis when it was redefined in terms of state stability (**Supplementary Figures 2A-2B**); left panels).

Similar to steady state homeostasis, we defined *growth state homeostasis* as the fraction of time spent in the growth state (**Figure 3A**, center panel) and *decline state homeostasis* as the fraction of time spent in the decline state (**Figure 3A**, right panel). As expected, we observed that AL mice, on average, gained body weight for at least 4 months out of their first 6 months of life; this growth state homeostasis dropped to approximately 20% at 18 months of age. Unsurprisingly, CR and IF significantly reduced the ability of a mouse to sustain growth, with 40% CR mice spending less than 3 weeks gaining body weight over the year following dietary intervention. On the other hand, we observed that decline state homeostasis in AL mice substantially increased with age after 18 months, with mice spending nearly 30% of their time on average losing body weight until the end of their life. In contrast, calorie restriction severely reduced decline state

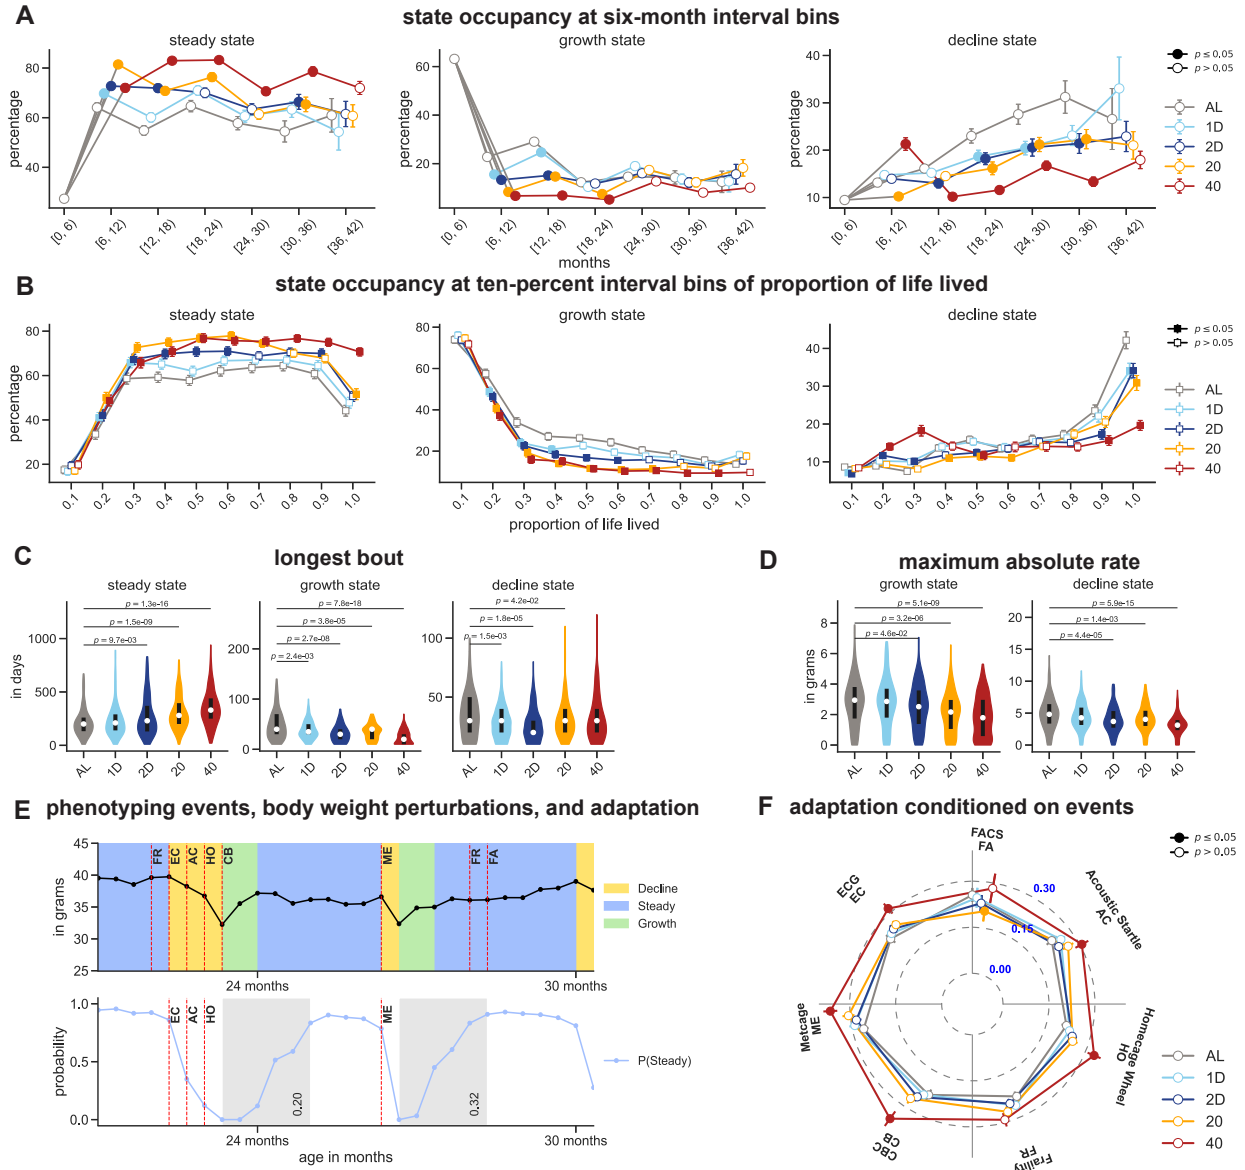


Figure 3: Influence of diet and age on body weight homeostasis. (A) The mean and standard errors of the state occupancy in steady, growth, and decline states, at the six-month interval bins, are represented with a circle and vertical line, respectively. The pre-intervention phase, i.e., [0, 6) months, is indicated in gray color. (B) The mean and standard errors of the state occupancy in steady, growth, and decline states, at ten percent interval bins of proportion of life lived, are represented with a square and vertical line, respectively. (C) The longest continuous bout (in days) in the growth, steady, and decline states in the post-intervention phases. (D) The maximum absolute rate (in grams) of the growth and decline states in the post-intervention phases. (E) A zoomed-in plot of the body weight trace of a sample mouse on one-day fasting diet. The approximate time of a phenotyping event is indicated using red-dashed line (top panel). The posterior probability of steady state for the corresponding body weight trace. The time period of adaptation to stress are indicated using gray boxes. (F) Radar chart of the average value and the standard-error of adaptation to stress conditioned on the phenotypic assay. In (A), (B), and (F), solid squares or circles indicate p -values < 0.05 , where the p -values were obtained by performing a Mann-Whitney test between a diet group and the AL group conditioned on the same interval bin or phenotyping assay. In (C) and (D), p -values were obtained by performing a Mann-Whitney test between two diets, with the AL diet as the reference group.

homeostasis, with 40% CR mice spending only 10-15% of their time losing body weight after 18 months. Similar trends were also observed in the effects of age and diet on growth and decline state homeostasis when the time-axis was

changed to proportion of life lived and homeostasis was defined in terms of state stability rather than state occupancy (Supplementary Figures 2C-2F)).

We quantified age-independent summaries of body weight trajectories post-intervention such as longest continuous bout and maximum rate of change of body weight conditioned on the latent state, and explored the effects of diet on these measures. Calorie restriction significantly increased the longest continuous bout in steady state while IF had no significant effect on this trait (Figure 3C, left panel). The average longest stretch of maintaining body weight lasted longer than 30% of the post-intervention lifespan for the average CR mouse compared to approximately 25% for the average AL mouse. (Supplementary Figure 3A, left panel). On the other hand, both CR and IF significantly shortened the longest continuous bout of gaining or losing body weight (Figure 3C, Supplementary Figure 3A, center and right panels). While both CR and IF affected the lifespan-normalized time of onset of the longest stretch of body weight gain, only CR affected the lifespan-normalized time of onset of the longest stretch of body weight loss (Supplementary Figures 3B-3C, center and right panels).

Steady state homeostasis is disrupted when mice experience a sudden change in body weight. To quantify this disruption, we calculated the maximum rate of change of body weight in the growth and decline states. AL mice achieved significantly faster post-intervention gain and loss of body weight compared to mice on other diets (Figure 3D). However, the effects of diet on these phenotypes were no longer significant when the rates were normalized by body weight, measured at the time when the maximum rates were recorded (Supplementary Figure 3D). Nonetheless, the average mouse could achieve an 8% per week increase in body weight and a 15% per week drop in body weight, suggesting that stressful life events can cause strong perturbations to body weight. However, diet did not have a substantial effect on the time at which the maximum rate of change in body weight were recorded (Supplementary Figures 3E-3F).

Calorie restriction increases adaptation to stress

Mice in this study experienced stress due to handling during an annual one-month long phenotyping event (Francesco et al. (2023)) (Supplementary Figure 4A). Handling involved interactions with lab technicians and instruments, changes between group and single housing, periodic availability of running wheels, and other stressors which resulted in the disruption of steady state homeostasis (Figure 3E, top panel). We defined *adaptation to stress* as the rate at which a mouse returns to steady state dynamics following a deviation from steady state caused by one or more phenotyping events. We identified time windows immediately following a phenotyping event (or group of events) where a mouse is returning to steady state (shaded portions on bottom panel of Figure 3E), and measured the rate of adaptation to these stressor events. Specifically, we modeled the steady state posterior probability following a

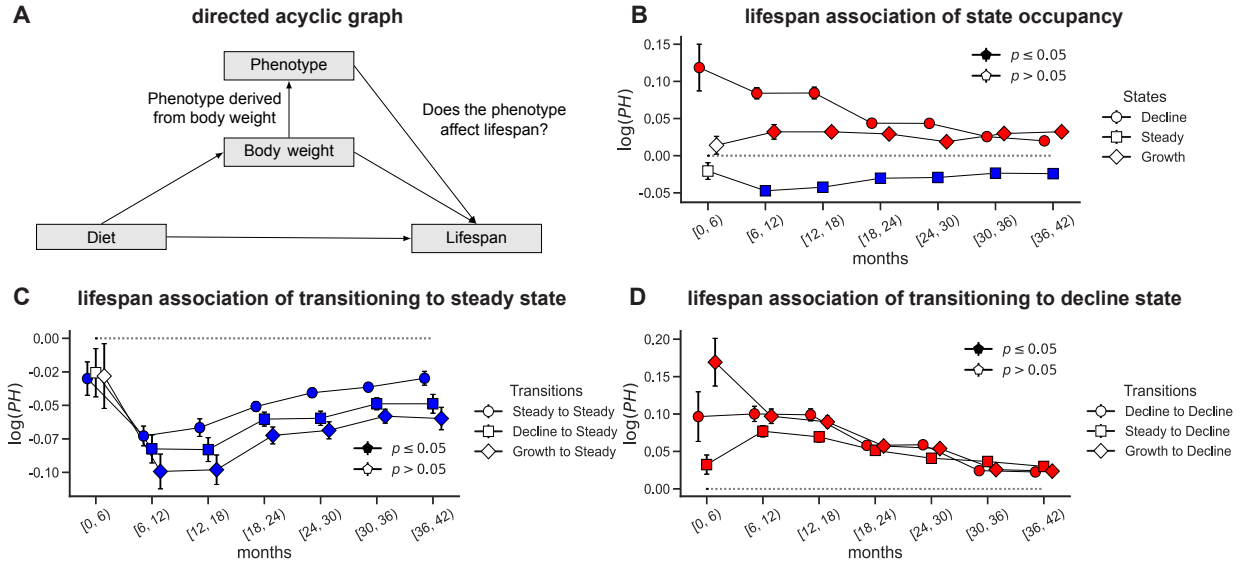


Figure 4: Influence of body weight homeostasis on lifespan. (A) A directed acyclic graph capturing the dependencies between covariates and lifespan. We used a time-varying Cox proportional hazard model to determine the association between body weight-derived traits and lifespan, while also accounting for confounding factors such as body weight and diet. (B) Effect size and standard error of state occupancy in steady, growth, and decline states, and its association with lifespan at every six-month interval bins. (C) Effect size and standard error of transitioning to steady state and its association with lifespan at every six-month interval bins. (D) Effect size and standard error of transitioning to decline state and its association with lifespan at every six-month interval bins. In (B), (C), and (D), lifespan associations which were statistically significant (p -value < 0.05) were indicated with solid colors. If the effect size of the trait at a given time interval lies above the dashed gray line (red filled markers), then an increase in the value of the trait decreases lifespan. Conversely, if it lies below the dashed gray line (blue filled markers), then an increase in the value of the trait increases lifespan.

phenotyping event using a cumulative exponential distribution function and selected the parameter that generated the best fit to this function as the rate of return to steady state (**Methods**). The higher the adaptability to stress, the faster a mouse returns to steady state.

We measured adaptation to stress following seven of the ten phenotyping assays which substantially disrupted mice from steady state (**Supplementary Figure 4B**). Across all assays, on average, 40% CR mice showed the fastest adaptation to stress (**Figure 3F**). To quantify age-related effects on adaptation to stress, we divided the age-axis into two non-overlapping bins: (6-18) months and (18-30) months; the age-bins were picked to ensure that the number of assay events in each age-bin were equal. Across most assays, we noticed that AL mice recovered to steady state more rapidly in late life than in mid life (**Supplementary Figures 4C-4D**), while mice on other diets did not show a significant difference in stress adaptation between mid-life and late-life. This suggests that a short duration (3-4 months) of dietary restriction induced and maintained an equivalent amount of adaptability to stress as that induced by the stressors associated with a battery of phenotyping assays spread over a year.

Steady state homeostasis is associated with reduced mortality throughout life

Body weight in mice is known to be predictive of lifespan in an age-dependent manner (Goodrick et al. (1990)); however, it is unknown whether the dynamics of body weight at different ages are also predictive of lifespan. We estimated the age-dependent effect of our derived homeostatic traits on mortality hazard (Figure 4A), after controlling for the effects of other lifespan determinants such as diet and body weight. Specifically, we applied a time-varying Cox proportional hazard model (Zhang et al. (2018)) to describe the survival time as a function of time-independent covariates such as diet and generation, and time-dependent covariates such as body weight (densely measured) and age-binned measures of homeostasis (sparsely quantified) (Methods).

We observed that steady state homeostasis was positively associated with lifespan, the strength of association was highest in mid-life (6-18 months) and gradually decreased with age (Figure 4B). Conversely, decline state homeostasis and growth state homeostasis were both negatively associated with lifespan, with a much stronger effect for decline state homeostasis. Although mice on average spent only 2 weeks in the decline state in early life (0-6 months) (Figure 3C), the association with lifespan was strongest for this age bin, highlighting that loss of body weight during early life in response to stressful stimuli has severe detrimental effects on lifespan.

Because increased steady state homeostasis was associated with increased lifespan, we sought to quantify the effect of transitioning to steady state (from other states) on lifespan. We observed that the rate of transitioning to steady state was positively associated with lifespan as well, although the effect sizes were dependent on the physiological state from which the mouse transitioned (Figure 4C). Similar to our earlier observation, high rates of transitioning to decline and growth states were associated with a decreased lifespan (Figure 4D, Supplementary Figure 5A). Additionally, transitions from growth to decline during early life showed the strongest negative effect on lifespan (Figure 4D), highlighting that perturbations to steady growth in early life that result in a loss of body weight are severely detrimental to mortality. In contrast to our homeostatic traits, adaptation to stress due to the phenotyping events showed no significant association with lifespan (Supplementary Figure 5B).

Genetic determinants of body weight dynamics are distinct from those of body weight

We carried genetic mapping of our derived body weight traits on the DO mice using the GxEMM model (Wright et al. (2022)). We tested for the additive effects of genotyped variants on our derived homeostatic traits (Supplementary Table 2). We computed these traits using the entire post-intervention phase and focused on 18 traits that were nominally associated with lifespan (p -value < 0.05) (Supplementary Figures 5C-5D). Grouping the traits by their direction of effect on lifespan, we identified 10 traits with at least one significant variant following group-specific false discovery

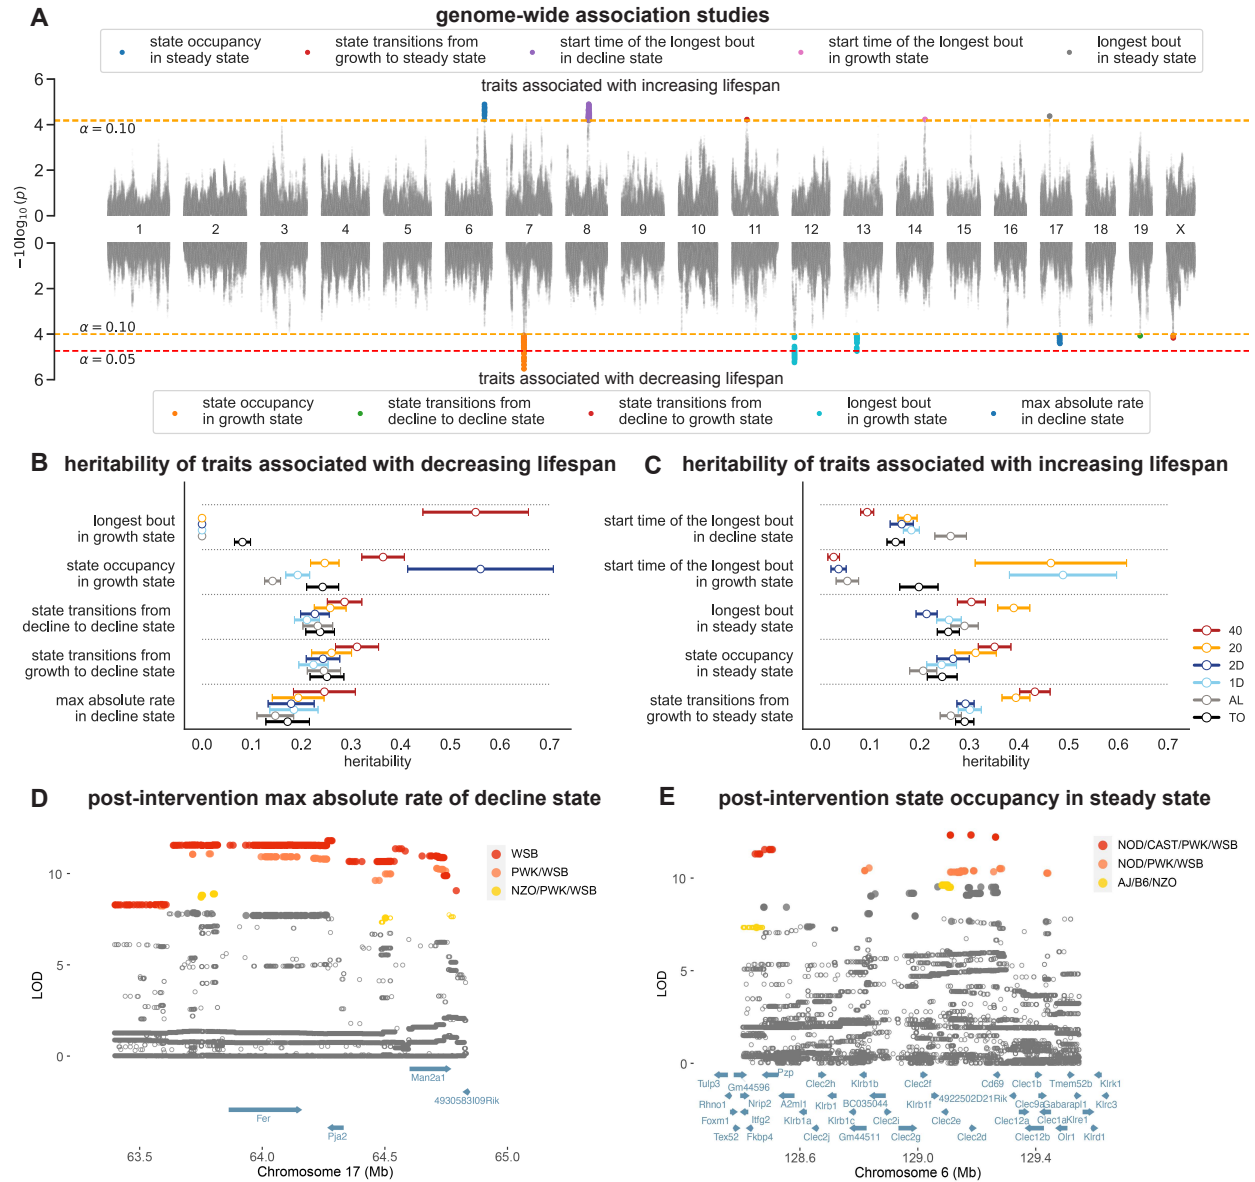


Figure 5: Influence of genetics and environment. (A) Combined Manhattan plot of traits in the post-intervention phase which were either associated with increasing or decreasing lifespan. The red and orange dashed lines represent the group-specific threshold at a false discovery rate of α . (B) and (C) Forest plot of the diet-specific heritability and corresponding standard error of different body-weight derived traits, divided into two groups: associated with decreasing lifespan (left) and increasing lifespan (right). Total heritability of the trait is indicated in black. (D) and (E) Fine-mapping loci associated with maximum absolute rate of decline state and steady state homeostasis traits on chromosome 17 and 6, respectively. Solid circles indicate significant variants. Colors denote variants with shared founder allele patterns (FAP). Ranks 1, 2, and 3 by logarithm of the odds (LOD) score are colored red, orange, and yellow, respectively.

rate (FDR) correction, and a total of 12 loci significantly associated with these traits (Figure 5A; see (Methods) for details). None of these loci overlapped with genetic loci previously mapped to body weight of mice from the same experiment (Wright et al. (2022)). We tested for additive genetic effects on 7 lifespan-associated traits computed using the pre-intervention phase but identified no significant loci. Finally, we also tested for genotype \times diet interaction effects

on each of these 25 lifespan-associated traits but identified no significant interaction effects.

We evaluated whether the genetic contribution of the ten lifespan-associated traits differed between various dietary contexts. We quantified the total and diet-dependent heritability for each trait using the GxEMM model, controlling for diet and generation effects. The lifespan-associated traits were less heritable than average body weight post-intervention. Additionally, we observed heritabilities of these traits differed between dietary contexts, with traits typically being more heritable under the 40% CR diet. The higher heritability of these traits in the 40% CR diet can be largely attributed to increased genetic variation relative to other diet groups, suggesting that the dynamics of body weight are tightly regulated under severe calorie restriction. We computed the partial correlation, controlling for diet and cohort effects, between pairs of traits and their genetic correlation using a matrix-variate linear mixed model (Furlotte and Eskin (2015)). We observed the genetic correlation between these traits to be structured, suggesting that the lifespan-associated heritable traits are regulated by shared cellular and molecular processes ([Supplementary Figure 6A](#)).

To interpret these genetic associations, we fine-mapped each of the 12 significant trait-locus pairs using allele-dosage at both genotyped and imputed variants within a 1-megabase (Mb) window centered at each locus. Our fine-mapping analysis confirmed 7 candidate trait-loci pairs, of which 4 were associated with decreasing lifespan ([Table 1](#)). We highlight the fine-mapping results of two traits: (a) maximum absolute rate in decline state, which was associated with decreasing lifespan, and (b) steady state homeostasis, which was associated with increasing lifespan ([Figures 5D-5E](#)). At the chromosome 17 locus associated with *maximum absolute rate of decline*, we grouped fine-mapped variants based on their founder-allele-pattern (FAP) and ranked groups based on the largest LOD score among its constituent variants, to identify variants and haplotypes most likely responsible for the association at this locus ([Methods](#)). The top FAP group contained variants with minor alleles specific to the WSB strain ([Figure 5B](#), [Supplementary Figure 6B](#)). All significant variants were located in a genomic interval containing three genes: *Fer*, *Pja2*, and *Man2a1* ([Figure 5D](#)). Intersecting the WSB-specific significant variants with regulatory elements assayed across a number of mouse tissues (Cusanovich et al. (2018); Gorkin et al. (2020)), we observed that the strongest variants (p -value $< 10^{-5}$) were located within regulatory elements around *Fer* and *Man2a1* genes, active across adipose, intestine, and muscle tissues ([Supplementary Figure 6C](#)). These observations suggest that this locus likely influenced the rate of loss of body weight by regulating gene expression in tissues directly related to metabolism. Similarly, at the chromosome 6 locus associated with *steady state homeostasis*, the top FAP group contained variants with minor alleles shared by the NOD, CAST, PWK, and WSB strains ([Figure 5E](#), [Supplementary Figure 6D](#)). All significant variants were located in a genomic interval containing the following genes: *Pzp*, *A2ml1*, and *Nkrp1-Clr* cluster ([Figure 5E](#)). The strongest candidate variants were located within regulatory elements at the *Pzp* and *Clec2d* genes, active across most tissues ([Supplementary Figure 6E](#)).

Table 1: Fine mapping results of body weight-derived phenotypes associated with either increasing or decreasing lifespan in the post-intervention phase.

Phenotype (Lifespan Association)	Chr	FAP	FAP Rank	P-value	FAP Position		Candidate Genes
					Start	End	
Decreasing lifespan							
maximum abs rate in the decline state (↓)	17	WSB	1	1.04E-06	63.863049	64.75511	Fer, Pja2, Man2a1
state occupancy in the growth state (↓)	7	AJ/129	3	3.22E-06	74.438821	74.438821	Slco3a1
state transitions from decline to growth (↓)	X	PWK	1	8.47E-07	47.809368	47.966428	Smarca1, Ocr1
longest bout in the growth state (↓)	13	AJ/129 PWK/WSB	1	2.14E-06	44.729474	45.002147	Jarid2, Dtnbp1
Increasing lifespan							
state occupancy in the steady state (↑)	6	NOD/CAST PWK/WSB	1	7.74E-07	128.483567	129.186535	Pzp, A2ml1, Nkrp1-Clr Cluster
state transitions from growth to steady (↑)	11	CAST/WSB	1	6.55E-06	49.609263	49.714004	Flt4, Scgb3a1, Cnot6
longest bout in the steady state (↑)	17	AJ/129 NOD/NZO	2	8.03E-06	31.944769	32.389439	Hsf2bp, Rrp1b, Notch3 Ephx3, Brd4, Akap8

Discussion

In this study, we derived novel phenotypes describing the dynamics of body weight measurements taken at high temporal resolution throughout life, evaluated their utility for predicting lifespan, and quantified their genetic determinants. To quantify temporal dynamics in body weight, we developed an autoregressive hidden Markov state space model to extract states representing growth, decline, and maintenance of steady body weight (Figure 2). Previous studies investigated age-related trends in body weight from longitudinal measurements using simple parametric models. Due to the non-monotonic nature of body weight dynamics, studies have fit an asymmetrical *inverted-U* pattern either by piece-wise linear regression or polynomial regression and dividing the age-axis into early (growth), mid (steady), and late (decline) life based on the estimated slopes (Wagener et al. (2013)). Such approaches are well suited to capturing long-term trends in mean body weight, but do not capture short-term, transient changes in body weight within each phase such as loss or gain and subsequent stabilization of body weight in response to stressful life events. In contrast, our model of body weight trajectories as a sequence of states allowed for the detection of multiple, transient switches between these states that occur throughout life and captured the continuous nature of body weight fluctuations (i.e., did not require fixed breakpoints of the age axis). Using these inferred states, we derived novel traits that quantify homeostasis of body weight dynamics and adaptation to stress (Figure 3). We further demonstrate that these traits are

modulated by dietary restriction and predictive of lifespan.

Prior research has demonstrated that body weight during early life predicts lifespan (Miller et al. (2002); Roy et al. (2021)). Similarly, it is well-recognized that diet has an effect on body weight (Fontana et al. (2010); Di Francesco et al. (2018); Longo and Panda (2016); de Cabo and Mattson (2019)). However, no previous studies have quantified the effect of diet on body weight trajectories throughout life and how features derived from such trajectories are associated with lifespan. In our work, we showed that both CR and IF increased the proportion of time spent in a state without dramatic weight loss or gain (steady state). This was not solely due to an inability to gain body weight due to caloric constraints or a physiological lower limit on body weight preventing weight loss. We observed that mice under 1D IF were consuming the same number of calories as ad libitum mice and that mice on CR and IF did spend brief periods of time in both growth and decline states but returned more quickly to steady state compared to ad lib mice. We also demonstrated that time spent in steady state was strongly associated with lifespan at multiple ages, even after conditioning on the effects of diet and body weight (Figure 4). Deviation from steady state, captured by the time spent in either growth or decline states, were influenced by diet and phenotyping events, and also associated with lifespan.

One interpretation of deviation from steady state is a perturbation to body weight due to extrinsic stress such as a phenotyping event. This is supported by observation that deviations often coincided with weeks where phenotyping was performed. We hypothesized that the rate of return to steady state is a measure of adaptation to stress. We note that this measure of adaptation to stress is different from *resilience*, a phenotype traditionally used in the field of aging. In biology, *resilience* is often defined as the ability to recover following an acute stress, whereas our measure of adaptation to stress does not require recovery of the original body weight, but rather a return to homeostasis at a new and potentially different body weight (Kirkland et al. (2016)). Our definition of adaptation to stress is more aligned with the definition of resilience in the field of systems dynamics, where the resilience of a system is the rate at which the system converges to an equilibrium state (steady state) after a disturbance (Hirsch and Smith (2006)). Thus, when a mouse returns to steady state following a phenotyping event, it is not required that the mouse regains the body weight lost during the phenotyping event.

While most studies used inbred strains to evaluate the metabolic effects of CR and IF, in this study, we used mice from a genetically diverse outbred stock which enables us to estimate genetic effects on body weight dynamics. Out of 25 lifespan-associated traits, genetic mapping identified 10 traits mapped to 12 genomic loci, including time spent in or away from steady state (Figure 5). We estimated diet-specific heritability for each of the 10 lifespan-associated traits and found six had significant differences in heritability between diets. We highlight two traits of interest and their respective associated genetic loci. One trait, *maximum rate of decline in body weight*, was negatively associated with lifespan and fine-mapped to the gene *Man2a1*, with the strongest candidate variants located within regulatory elements

that were active in adipose, intestine, and muscle tissues. Mannosidase alpha class 2A member 1 (*Man2a1*) encodes a glycosyl hydrolase that has previously been associated with the pathogenesis of inflammatory bowel disease (Suzuki et al. (2018)), and its inhibition ameliorated rapid body weight decline and symptoms of ulcerative colitis in mouse models. Another trait, *proportion of time spent in steady state*, was positively associated with lifespan and fine-mapped to the gene *Pzp*, where the fine-mapped variants colocalized within a regulatory element active across several tissues relevant to metabolism including adipose and liver tissues. Pregnancy zone protein (*Pzp*), a member of the alpha-2 globulin family of proteins, was recently identified as a key hepatokine regulating factor for fasting-refeeding triggered energy homeostasis through inter-organ cross talk between liver and brown adipose tissue (Lin et al. (2021)).

Our study provides valuable insights into the changes of body weight dynamics. We examine their relationship with age and diet as well as lifespan and genetics. There are some limitations to consider. It is known that body composition and adiposity are also strongly influenced by diet and energy expenditure, and are predictive of lifespan. Future studies would ideally also include dense longitudinal measurements of body composition, food intake, and energy expenditure. Given the importance of body weight changes in response to extrinsic events, a study design that samples body weight more densely around such events (e.g., daily or more often) would allow for a more accurate characterization of adaptation to stress. A measure of how well body weight is regulated in response to stressful events could help identify factors that promote resilience. Additionally, combining sparse temporal sampling of molecular and cellular data (epigenetics, proteomics, metabolomics, and lipidomics) with the dynamics of densely sampled body weight could provide new insight into the mechanisms underlying homeostasis and aging.

Methods

Mouse housing, feeding, and body weight measurements

The DO mice were generated by breeding eight founder inbred strains (AJ, B6, 129, NOD, NZO, CAST, PWK, and WSB) to produce an outbred heterozygous population with a random assortment of genetic variation. In this study, 960 DO female mice, sampled at generations 22-24 and 26-28, were enrolled into the study after wean age of 3 weeks old, and maintained at the Jackson Laboratory in Maine. No mice in the study were siblings and maximum genetic diversity was achieved. There were two cohorts per generation for a total of 12 cohorts and 80 animals per cohort. Enrollment occurred in successive quarterly waves starting in March 2016 and continuing through November 2017. Mice were housed in pressurized, individually ventilated cages at a density of eight animals per cage with random cage assignments. They were subject to 12 hours of continuous light and dark cycles beginning at 06:00 AM and 06:00 PM,

respectively. Animals exit the study upon death. All animal procedures were approved by the Animal Care and Use Committee at The Jackson Laboratory.

For mice under calorie restriction (CR), food was weighed out for an entire cage of eight. Observation of the animals indicated that the distribution of food consumed was roughly equal among all mice in a cage across diet groups. As the number of mice in each CR cage decreased over time, the amount of food given to each cage was adjusted to reflect the number of mice in that cage. Nonetheless, individual feeding behavior was not controlled among the CR mice and it is possible that some of the variability observed in the CR treatment groups was due to varying degrees of individual caloric intake.

The body weight measurements for this analysis were collated in October 2021, at which point 960 mice (97.35%) had measurements at 6 months, 884 (92.08%) at 12 months, 803 (83.64%) at 18 months, 639 (66.56%) at 24 months, 416 (43.33%) at 30 months, 187 (19.47%) at 36 months, and 43 (4.47%) at 42 months. We included all body weight measures for each mouse up to 1632 days of age (>42 months). We excluded four mice from the downstream analysis (1 from AL group, 2 from 1D group, and 1 from 2D group) because these mice died early in the study and had fewer than seven body weight measurements in total, resulting in 956 mice in our analyses. To obtain a uniform sampling interval for the body weight dynamics, we rounded the time (in days) of measurement of the body weight to the nearest ten. When two body weight measurements were available at the nearest ten, we took the average of these measurements and assigned it to the nearest ten. This strategy of creating a uniform sampling of ten day sampling interval resulted in 4.83% missing values at specific time points. Because the number of missing values is a tiny fraction of the entire database, we used a one-dimensional linear interpolation filter to fill the missing values.

Autoregressive hidden Markov model

Autoregressive models are based on the idea that the current value of the time-series, y_t measured at time t , can be explained as a function of ℓ past values, where ℓ is the lag-order and determines the number of steps into the past needed to forecast the current values. The simplest form of an autoregressive model can be represented as $AR(\ell)$ which takes the form $y_t = \phi_0 + \phi_1 y_{t-1} + \dots + \phi_\ell y_{t-\ell} + \epsilon_t$, where $(\phi_0, \dots, \phi_\ell)$ are the autoregressive coefficients and ϵ_t is the additive white Gaussian noise distributed as $\mathcal{N}(0, \sigma^2)$ (see [Figure 2A](#)). A hidden Markov model is a state-space model which is characterized by the hidden states and observations generated by the hidden states. The hidden states are assumed to follow a first-order Markov chain and can only be detected through the observed sequence as they emit observations on varying probabilities. A Gaussian hidden Markov model where $s_t \in \{1, \dots, K\}$ represents latent state at time t is specified with an initial probability distribution $\pi_0 \in \mathbb{R}^{K \times 1}$, a transition matrix $A \in \mathbb{R}^{K \times K}$ with each $a_{i,j}$ representing the probability of moving from state $s_{t-1} = i$ to state $s_t = j$ such that the row sum equals one, and

a sequence of observation likelihoods, also called as emission probabilities, each expressing the probability of an observation y_t being generated from state $s_t = i$ drawn from a Gaussian distribution $\mathcal{N}(\eta_{s_t}, \sigma_{s_t}^2)$ (see [Figure 2A](#)).

The ARHMM combines an autoregressive model and a hidden Markov model. In an ARHMM, the observations are generated by a few autoregressive time-series models of fixed lag-order, where the switching between these models is controlled by the hidden states which follow a first-order Markov chain. We denote the ARHMM of l -th lag-order as $\text{ARHMM}(\ell)$ and is defined as $y_t = \phi_0^{s_t} + \phi_1^{s_t} y_{t-1} + \dots + \phi_\ell^{s_t} y_{t-\ell} + \epsilon_t^{s_t}$, where s_t represents one of the K possible latent states, $(\phi_0^{s_t}, \dots, \phi_\ell^{s_t})$ are the autoregressive coefficients corresponding to state s_t , and $\epsilon_t^{s_t}$ is the state-dependent additive white Gaussian noise distributed as $\mathcal{N}(0, \sigma_{s_t}^2)$. To capture the variation within each latent state, we assumed that the autoregressive coefficients are drawn from a Gaussian distribution of unknown mean and variance given as $\phi_{s_t}^\ell \sim \mathcal{N}(\eta_{\ell, s_t}, \sigma_{\ell, s_t}^2)$, where ℓ is the lag-order of the autoregressive model in state s_t (see [Figure 2A](#)).

The physiological states are represented by the discrete latent states that capture the underlying body weight dynamics at an organismal level. Furthermore, we only considered autoregressive models of lag-order one and set the autoregressive coefficient of the first lag-order to be one, i.e., $\phi_{s_t}^1 = 1, \forall 1 \leq s_t \leq K$. Therefore, $y_t = \phi_0^{s_t} + y_{t-1} + \epsilon_t^{s_t}$, which implies that the body weight at time t depends on body weight at time $t - 1$, a state-dependent random variable $\phi_0^{s_t}$ which is drawn from $\mathcal{N}(\eta_{0, s_t}, \sigma_{0, s_t}^2)$, and a state-dependent error which is drawn from $\mathcal{N}(0, \sigma_{s_t}^2)$. These simplifications made the latent states extracted using the ARHMM interpretable because the physiological states are determined based on the sign and magnitude of the mean of the $\phi_0^{s_t}$ autoregressive coefficient. We implemented the ARHMM using a variational inference approach. The details of the derivation are provided in the Appendix.

Adaptation to stress

We used the posterior probability of steady state as a proxy to measure the rate of recovery because mice spend a significant proportion of their lifespan in the steady state. To compute the rate of recovery, first we identified regions of homeostasis. A region of homeostasis starts when the posterior probability of steady state goes below 0.05, and ends when the probability of steady state stops monotonically increasing or is monotonically increasing and goes beyond 0.95. Thereafter, for each region, we fit a cumulative distribution function of an exponential distribution, which is given as $(1 - e^{-\lambda x})$, where λ is the rate parameter of an exponential distribution. Finally, for any given region of homeostasis longer than thirty days duration, we had at least 3 measurements to estimate λ and the value of λ that generated the best fit was considered as the rate of recovery.

Time-varying Cox proportional hazard model

In a traditional Cox model, the hazard ratio is dependent on the covariates and independent of time. In contrast, in an extended Cox proportional model or a time-varying Cox proportional model, the hazard ratio is dependent on time, i.e, the covariates vary with time. The time-varying covariates are classified into two types: internal (dependent on individuals in the study) and external (independent of the individuals in the study). In our proposed model for lifespan analysis, we used body weight as an internal time-varying covariate and the body weight-derived trait measured in a given time-interval bin as an external time-varying covariate. We considered the following time-varying Cox proportional model:

$$h(t, \mathbf{Z}) = \underbrace{h_0(t)}_{\text{baseline hazard}} \exp \left\{ \underbrace{\sum_{j=1}^N \beta_j^{(T)} Z^{(T)}(B_j)}_{\text{trait : age bin}} + \underbrace{\sum_{j=1}^N \beta_j^{(B)} \sum_{t \in B_j} Z^{(B)}(t)}_{\text{body weight : age bin}} + \underbrace{\sum_{j=1}^N \sum_{d \in D} \beta_j^{(D)} Z^{(D)}(B_j)}_{\text{diet : age bin}} \right. \\ \left. + \underbrace{\sum_{j=1}^N \sum_{d \in D} \beta_j^{(B,D)} \sum_{t \in B_j} Z^{(B,D)}(t)}_{\text{body weight : diet : age bin}} + \underbrace{\sum_{j \in G} \beta_j^{(G)} Z^{(G)}}_{\text{generation}} \right\},$$

where $h_0(t)$ is the baseline hazard, N is the number of non-overlapping and continuous age bins represented as $\{B_1, \dots, B_N\}$, and $\beta_j^{(T)}$, $\beta_j^{(B)}$, $\beta_j^{(D)}$, $\beta_j^{(B,D)}$, and $\beta_j^{(G)}$ are the effect sizes of (trait : age bin) interaction covariates, (body weight : age bin) interaction covariates, (diet : age bin) interaction covariates, (body weight : diet : age bin) interaction covariates, and generation covariates, respectively. We used the lifelines package implemented in Python to estimate the effect sizes of the body weight-derived trait at every six month interval.

Genotype measurements

We collected tail clippings and extracted DNA from 954 animals. Samples were genotyped using the 143,259-probe GigaMUGA array from the Illumina Infinium II platform by NeoGen Corp. We evaluated genotype quality using the R package: qtl2. We processed all raw genotype data with a corrected physical map of the GigaMUGA array probes. After filtering genetic markers for uniquely mapped probes, genotype quality, and a 20% genotype missingness threshold, our dataset contained 110,807 markers. Next, examined the genotype quality of individual animals. We found seven pairs of animals with identical genotypes, which suggested that one of each pair was mislabeled. We identified and

removed a single mislabeled animal per pair by referencing the genetic data against coat color. Next, we removed a single sample with missingness in excess of 20%. The final quality assurance analysis found that all samples exhibited high consistency between tightly linked markers: log odds ratio error scores were less than 2.0 for all samples. The final set of genetic data consisted of 946 mice. For each mouse, starting with its genotypes at the 110,807 markers and the genotypes of the 8 founder strains at the same markers, we inferred the founders-of-origin for each of the alleles at each marker using the R package: qtl2. This allowed us to test directly for association between founder-of-origin and phenotype (rather than allele dosage and phenotype, as is commonly done in QTL mapping) at all genotyped markers. Using the founder-of-origin of consecutive typed markers and the genotypes of untyped variants in the founder strains, we then imputed the genotypes of all untyped variants (34.5 million) in all 946 mice. Targeted association testing at imputed variants allowed us to fine-map QTLs to a resolution of 1–10 genes.

Genetic mapping with genotyped markers

For each phenotype, we tested for additive effects between the founder-of-origin at each genotyped variant and the phenotype using the GxEMM model (Wright et al. (2022)), controlling for diet and cohort fixed effects and allowing for genotype-x-diet random effects. Expecting that the set of phenotypes in our genetic analyses were not all independent, we grouped the phenotypes into four groups to determine a group-specific genomewide-level significance threshold. First, we applied a SNP-level false discovery rate of α to all the phenotypes within each group. Thereafter, we generated a group-specific p -value array consisting of the maximum group-specific p -value if the SNP-level null hypothesis is rejected and a randomly drawn group-specific p -value if the SNP-level null hypothesis is not rejected. Finally, we applied a false discovery rate of α to the p -value array of each group separately to obtain a group-specific genome-level significance threshold. This procedure resulted in a conservative and moderate threshold of 1.94×10^{-5} (at $\alpha = 0.05$) and 8.99×10^{-5} (at $\alpha = 0.10$), respectively, for those phenotypes associated with decreasing lifespan in the post-intervention phase. Similarly, we obtained a moderate threshold of 6.75×10^{-5} (at $\alpha = 0.10$) for those phenotypes associated with increasing lifespan in the post-intervention phase. Pre-intervention phenotypes, either associated with increasing or decreasing lifespan, did not result in any significant variants at FDR of $\alpha = 0.05$ and $\alpha = 0.1$.

Genetic fine-mapping with founder-allele-patterns

To more precisely fine-map the genomic interval of each QTL, we imputed all SNPs and insertion-deletion variants from the fully sequenced DO founders for a 5 Mb interval centered at the lead genotype-marker. For each imputed variant, we identified the founder-of-origin for the major and minor allele. To illustrate the process, consider the bi-allelic A/G

variant. If allele A was specific to founders AJ, NZO, and PWK, and allele G was specific to the other five founders, then we assigned A to be the minor allele and defined a founder-of-origin allele pattern (FAP) of AJ/NZO/PWK for this variant. We identified variants and founder haplotypes most likely responsible for the association at each locus by grouping variants based on their FAP and ranked groups based on the largest logarithm of odds (LOD) score among its constituent variants. We hypothesized that the functional variant(s) responsible for trait-specific variation were among those in the lead FAP group because they exhibit the strongest statistical association and it is unlikely any additional variants are segregating in this genomic interval beyond those identified in the full genome sequences of the eight founder strains. By focusing on the FAP groups with the largest LOD scores, we significantly reduced the number of putative causal variants, while representing the age- and diet-dependent effects of these loci in terms of the effects of its top FAP groups. We further narrowed the number of candidates by intersecting the variants in top FAP groups with functional annotations (e.g., gene annotations, regulatory elements, tissue-specific regulatory activity, etc.). This procedure identified candidate regions containing 1 – 3 genes.

Code Availability

The code for extracting novel body weight-derived traits using an autoregressive hidden Markov model, as well as the subsequent lifespan and genetic analyses of these traits, is publicly available through the following GitHub repository: https://github.com/calico/do_bwd. The repository also contains trained models, processed datasets, and scripts that can be utilized to train models from scratch and generate processed datasets.

Acknowledgements

We thank the JAX Nathan Shock Center Animal and Phenotyping core team for their assistance with animal husbandry, data collection, and curation. We would also like to thank Eugene Melamud for helpful feedback and discussion. This work was funded by Calico Life Sciences LLC.

References

- Billman, G. E. (2020). Homeostasis: The underappreciated and far too often ignored central organizing principle of physiology. *Frontiers in Physiology*, 11.
- Bou Sleiman, M., Roy, S., Gao, A. W., Sadler, M. C., Von Alvensleben, G. V. G., Li, H., Sen, S., Harrison,

- D. E., Nelson, J. F., Strong, R., Miller, R. A., Kutalik, Z., Williams, R. W., and Auwerx, J. (2022). Sex- and age-dependent genetics of longevity in a heterogeneous mouse population. *Science*, 377(6614):eabo3191.
- Cannon, W. B. (1929). Organization for physiological homeostasis. *Physiological Reviews*, 9(3):399–431.
- Chen, Z., Raj, A., Prateek, G., Di Francesco, A., Liu, J., Keyes, B. E., Kolumam, G., Jojic, V., and Freund, A. (2022). Automated, high-dimensional evaluation of physiological aging and resilience in outbred mice. *eLife*, 11:e72664.
- Cusanovich, D. A., Hill, A. J., Aghamirzaie, D., Daza, R. M., Pliner, H. A., Berletch, J. B., Filippova, G. N., Huang, X., Christiansen, L., DeWitt, W. S., Lee, C., Regalado, S. G., Read, D. F., Steemers, F. J., Distech, C. M., Trapnell, C., and Shendure, J. (2018). A single-cell atlas of in-vivo mammalian chromatin accessibility. *Cell*, 174(5):1309–1324.e18.
- de Cabo, R. and Mattson, M. P. (2019). Effects of intermittent fasting on health, aging, and disease. *New England Journal of Medicine*, 381(26):2541–2551. PMID: 31881139.
- Di Francesco, A., Di Germanio, C., Bernier, M., and De Cabo, R. (2018). A time to fast. *Science*, 362(6416):770–775.
- Dixon, J. B. (2010). The effect of obesity on health outcomes. *Molecular and Cellular Endocrinology*, 316(2):104–108.
- Fontana, L., Partridge, L., and Longo, V. D. (2010). Extending healthy life span — from yeast to humans. *Science*, 328(5976):321–326.
- Francesco, A. D., Deighan, A. G., Litichevskiy, L., Chen, Z., Luciano, A., Robinson, L., Garland, G., Donato, H., Schott, W., Wright, K. M., Raj, A., Prateek, G., Mullis, M., Hill, W., Zeidel, M., Peters, L., Harding, F., Botstein, D., Korstanje, R., Thaiss, C. A., Freund, A., and Churchill, G. A. (2023). Regulators of health and lifespan extension in genetically diverse mice on dietary restriction. *bioRxiv*.
- Freund, A. (2019). Untangling aging using dynamic, organism-level phenotypic networks. *Cell Systems*, 8(3):172–181.
- Furlotte, N. A. and Eskin, E. (2015). Efficient Multiple-Trait Association and Estimation of Genetic Correlation Using the Matrix-Variate Linear Mixed Model. *Genetics*, 200(1):59–68.
- Goodrick, C., Ingram, D., Reynolds, M., Freeman, J., and Cider, N. (1990). Effects of intermittent feeding upon body weight and lifespan in inbred mice: interaction of genotype and age. *Mechanisms of Ageing and Development*, 55(1):69–87.
- Gorkin, D. U., Barozzi, I., Zhao, Y., Zhang, Y., Huang, H., Lee, A. Y., Li, B., Chiou, J., Wildberg, A., Ding, B., Zhang, B., Wang, M., Strattan, J. S., Davidson, J. M., Qiu, Y., Afzal, V., Akiyama, J. A., Plajzer-Frick, I., Novak, C. S., Kato, M., Garvin, T. H., Pham, Q. T., Harrington, A. N., Mannion, B. J., Lee, E. A., Fukuda-Yuzawa, Y., He, Y., Preissl, S., Chee, S., Han, J. Y., Williams, B. A., Trout, D., Amrhein, H., Yang, H., Cherry, J. M., Wang,

- W., Gaulton, K., Ecker, J. R., Shen, Y., Dickel, D. E., Visel, A., Pennacchio, L. A., and Ren, B. (2020). An atlas of dynamic chromatin landscapes in mouse fetal development. *Nature*, 583(7818):744–751.
- Hirsch, M. and Smith, H. (2006). Chapter 4 monotone dynamical systems. volume 2 of *Handbook of Differential Equations: Ordinary Differential Equations*, pages 239–357. North-Holland.
- Kirkland, J. L., Stout, M. B., and Sierra, F. (2016). Resilience in Aging Mice. *The Journals of Gerontology: Series A*, 71(11):1407–1414.
- Kuo, P.-L., Schrack, J. A., Levine, M. E., Shardell, M. D., Simonsick, E. M., Chia, C. W., Moore, A. Z., Tanaka, T., An, Y., Karikkineth, A., AlGhatrif, M., Elango, P., Zukley, L. M., Egan, J. M., De Cabo, R., Resnick, S. M., and Ferrucci, L. (2022). Longitudinal phenotypic aging metrics in the Baltimore longitudinal study of aging. *Nature Aging*, 2(7):635–643.
- Lin, J., Jiang, X., Dong, M., Liu, X., Shen, Q., Huang, Y., Zhang, H., Ye, R., Zhou, H., Yan, C., Yuan, S., Wu, X., Chen, L., Wang, Y., He, M., Tao, Y., Zhang, Z., and Jin, W. (2021). Hepatokine pregnancy zone protein governs the diet-induced thermogenesis through activating brown adipose tissue. *Advanced Science*, 8(21):2101991.
- Lissner, L., Odell, P. M., D’Agostino, R. B., Stokes, J., Kreger, B. E., Belanger, A. J., and Brownell, K. D. (1991). Variability of body weight and health outcomes in the Framingham population. *New England Journal of Medicine*, 324(26):1839–1844.
- Longo, V. D. and Anderson, R. M. (2022). Nutrition, longevity and disease: From molecular mechanisms to interventions. *Cell*, 185(9):1455–1470.
- Longo, V. D. and Panda, S. (2016). Fasting, circadian rhythms, and time-restricted feeding in healthy lifespan. *Cell Metabolism*, 23(6):1048–1059.
- López-Otín, C., Blasco, M. A., Partridge, L., Serrano, M., and Kroemer, G. (2013). The Hallmarks of Aging. *Cell*, 153(6):1194–1217.
- Miller, R. A., Harper, J. M., Galecki, A., and Burke, D. T. (2002). Big mice die young: early life body weight predicts longevity in genetically heterogeneous mice. *Aging Cell*, 1(1):22–29.
- Moldakozhayev, A. and Gladyshev, V. N. (2023). Metabolism, homeostasis, and aging. *Trends in Endocrinology & Metabolism*, 34(3):158–169.
- Palliyaguru, D. L., Vieira Ligo Teixeira, C., Duregon, E., di Germanio, C., Alfaro, I., Mitchell, S. J., Navas-Enamorado, I., Shiroma, E. J., Studenski, S., Bernier, M., Camandola, S., Price, N. L., Ferrucci, L., and de Cabo,

R. (2021). Study of longitudinal aging in mice: presentation of experimental techniques. *The Journals of Gerontology: Series A*, 76(4):552–560.

Ravussin, Y., Leibel, R., and Ferrante, A. (2014). A missing link in body weight homeostasis: The catabolic signal of the overfed state. *Cell Metabolism*, 20(4):565–572.

Roy, S., Sleiman, M. B., Jha, P., Ingels, J. F., Chapman, C. J., McCarty, M. S., Ziebarth, J. D., Hook, M., Sun, A., Zhao, W., Huang, J., Neuner, S. M., Wilmott, L. A., Shapaker, T. M., Centeno, A. G., Ashbrook, D. G., Mulligan, M. K., Kaczorowski, C. C., Makowski, L., Cui, Y., Read, R. W., Miller, R. A., Mozhui, K., Williams, E. G., Sen, S., Lu, L., Auwerx, J., and Williams, R. W. (2021). Gene-by-environment modulation of lifespan and weight gain in the murine bxd family. *Nature Metabolism*, 3(9):1217–1227.

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., and Van Der Linde, A. (2002). Bayesian Measures of Model Complexity and Fit. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 64(4):583–639.

Suzuki, K., Yamada, T., Yamazaki, K., Hirota, M., Ishihara, N., Sakamoto, M., Takahashi, D., Iijima, H., and Hase, K. (2018). Intestinal epithelial cell-specific deletion of alpha-mannosidase II ameliorates experimental colitis. *Cell Structure and Function*, 43(1):25–39.

Svenson, K. L., Gatti, D. M., Valdar, W., Welsh, C. E., Cheng, R., Chesler, E. J., Palmer, A. A., McMillan, L., and Churchill, G. A. (2012). High-resolution genetic mapping using the mouse diversity outbred population. *Genetics*, 190(2):437–447.

Wagener, A., Müller, U., and Brockmann, G. A. (2013). The age of attaining highest body weight correlates with lifespan in a genetically obese mouse model. *Nutrition & Diabetes*, 3(3):e62–e62.

Willett, W., Rockström, J., Loken, B., Springmann, M., Lang, T., Vermeulen, S., Garnett, T., Tilman, D., DeClerck, F., Wood, A., Jonell, M., Clark, M., Gordon, L. J., Fanzo, J., Hawkes, C., Zurayk, R., Rivera, J. A., De Vries, W., Majele Sibanda, L., Afshin, A., Chaudhary, A., Herrero, M., Agustina, R., Branca, F., Lartey, A., Fan, S., Crona, B., Fox, E., Bignet, V., Troell, M., Lindahl, T., Singh, S., Cornell, S. E., Srinath Reddy, K., Narain, S., Nishtar, S., and Murray, C. J. L. (2019). Food in the Anthropocene: the EAT–Lancet Commission on healthy diets from sustainable food systems. *The Lancet*, 393(10170):447–492.

Wright, K. M., Deighan, A. G., Di Francesco, A., Freund, A., Jojic, V., Churchill, G. A., and Raj, A. (2022). Age and diet shape the genetic architecture of body weight in diversity outbred mice. *eLife*, 11:e64329.

Yang, Y., Smith, D. L., Keating, K. D., Allison, D. B., and Nagy, T. R. (2014). Variations in body weight, food intake and body composition after long-term high-fat diet feeding in C57BL/6J mice: Variations in Diet-Induced Obese

- 502 C57BL/6J Mice. *Obesity*, 22(10):2147–2155.
- 503 Zhang, Z., Reinikainen, J., Adeleke, K. A., Pieterse, M. E., and Groothuis-Oudshoorn, C. G. M. (2018). Time-varying
- 504 covariates and coefficients in cox regression models. *Annals of Translational Medicine*, 6(7).

Supplementary Tables

Table 1: Estimated parameters of the variational inference-based autoregressive hidden Markov model ($y_t = \phi_0^{s_t} + \phi_1^{s_t} y_{t-1} + \epsilon^{s_t}$, $s_t \in \{0, 1, 2\}$).

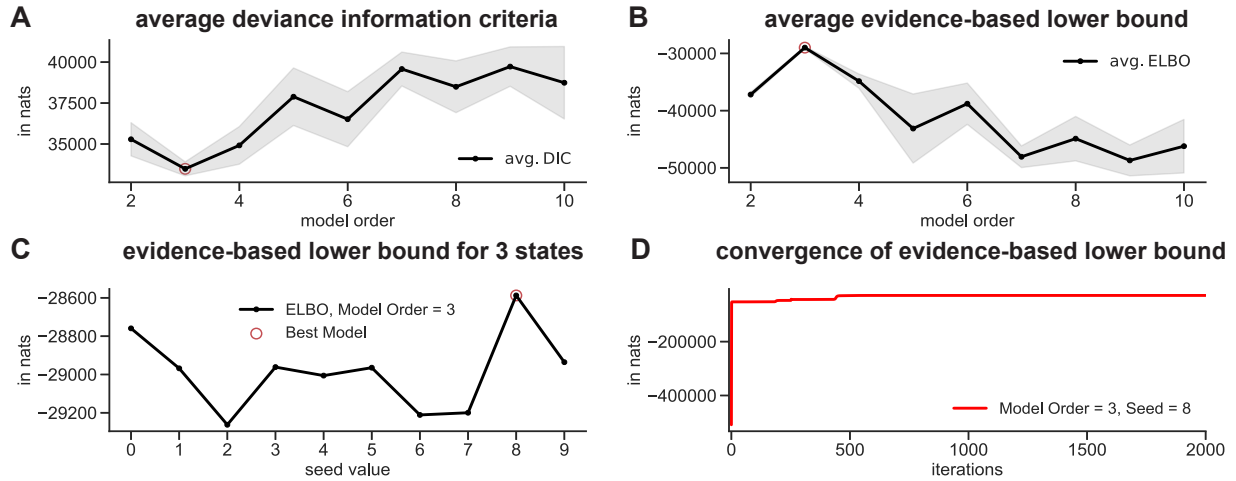
State	Initial Probability	Transition Probability from State X			AR(0)	Error
		P(DS X)	P(SS X)	P(GS X)		
Decline (X = DS)	$\pi_0 = 0.42$	0.42	0.21	0.37	$\phi_0^0 \sim \mathcal{N}(-1.6, 9.1 \times 10^{-4})$	$\epsilon^0 \sim \mathcal{N}(0, 2.00)$
Steady (X = SS)	$\pi_1 = 0.07$	0.07	0.93	≈ 0.00	$\phi_0^1 \sim \mathcal{N}(+0.1, 4.4 \times 10^{-6})$	$\epsilon^1 \sim \mathcal{N}(0, 0.64)$
Growth (X = GS)	$\pi_2 = 0.14$	0.14	0.12	0.74	$\phi_0^2 \sim \mathcal{N}(+1.4, 3.4 \times 10^{-5})$	$\epsilon^2 \sim \mathcal{N}(0, 1.83)$

Table 2: Body weight-derived traits and definitions

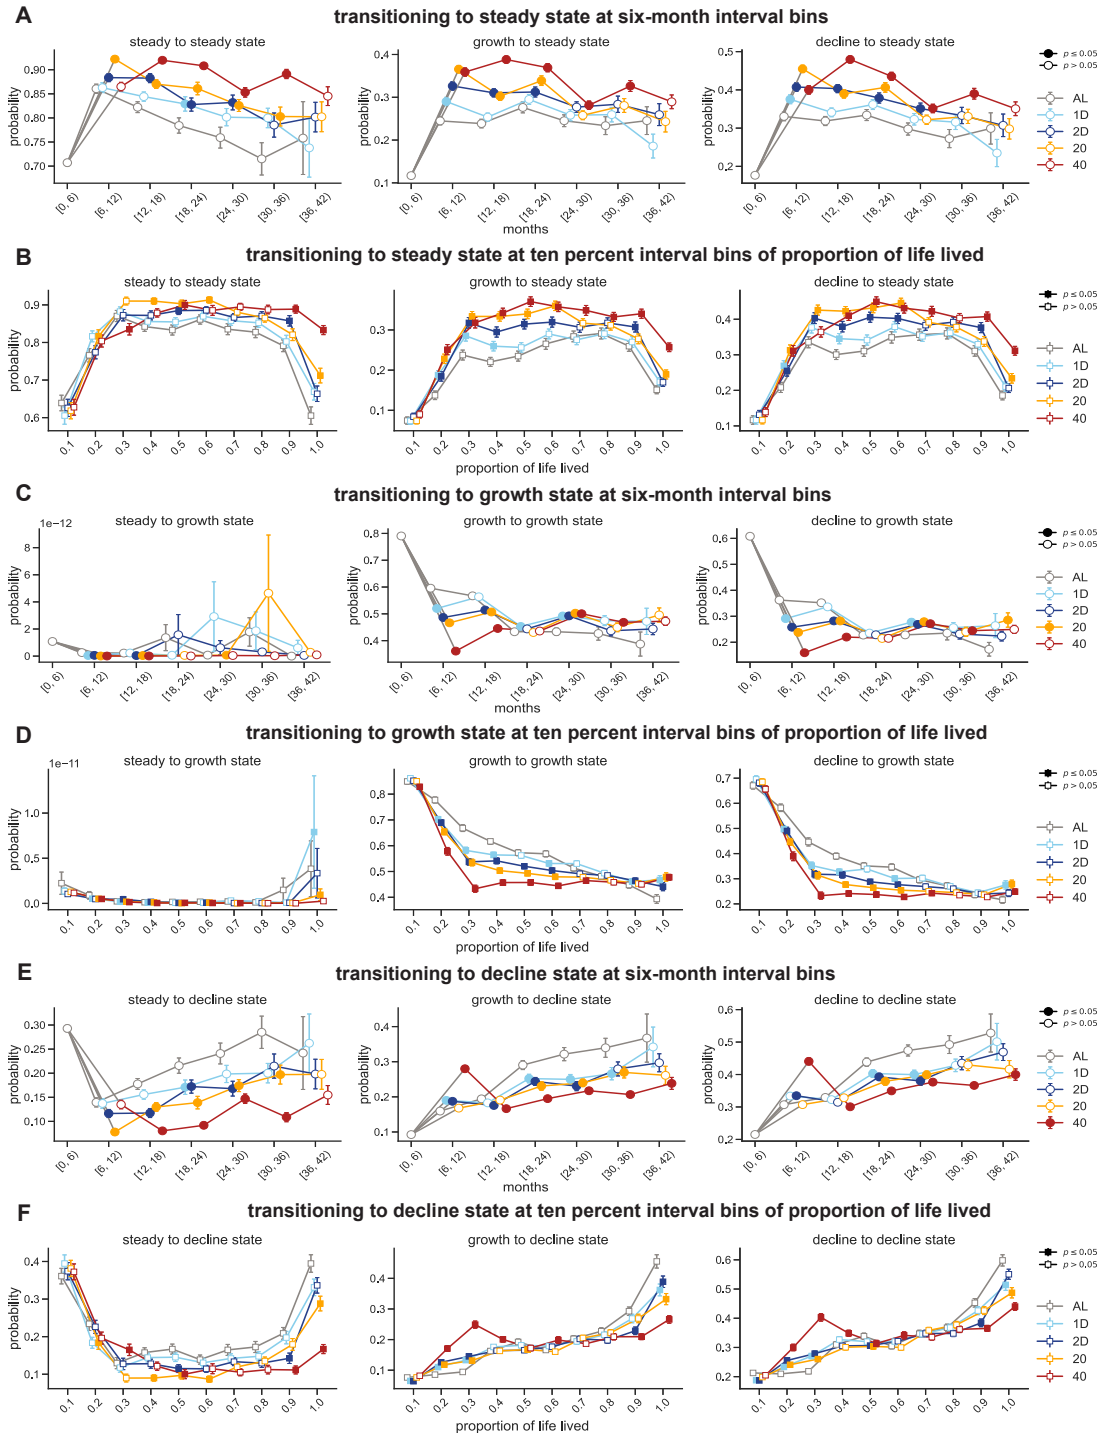
Phenotype	Definition
State occupancy	It is a measure of the duration of a physiological state in a given interval of time and can be used to assess the stability of a physiological system.
State transitions	It is a measure of changes in the state of a system over time. They can be described as the probability of switching from one state to itself or another state in a given interval of time.
Maximum absolute rate	It is the maximum rate of change experienced by a physiological state in a given interval of time. These frequently capture the maximum magnitude of the perturbation experienced by a system.
Longest bout	It is the longest continuous duration of a physiological state within in a given interval of time and can be used to assess the robustness to perturbations of a physiological system.
Start time	It is the time instance at which an event, such as maximum absolute rate or longest bout, was recorded. It can be useful in understanding abrupt changes in a system.

506

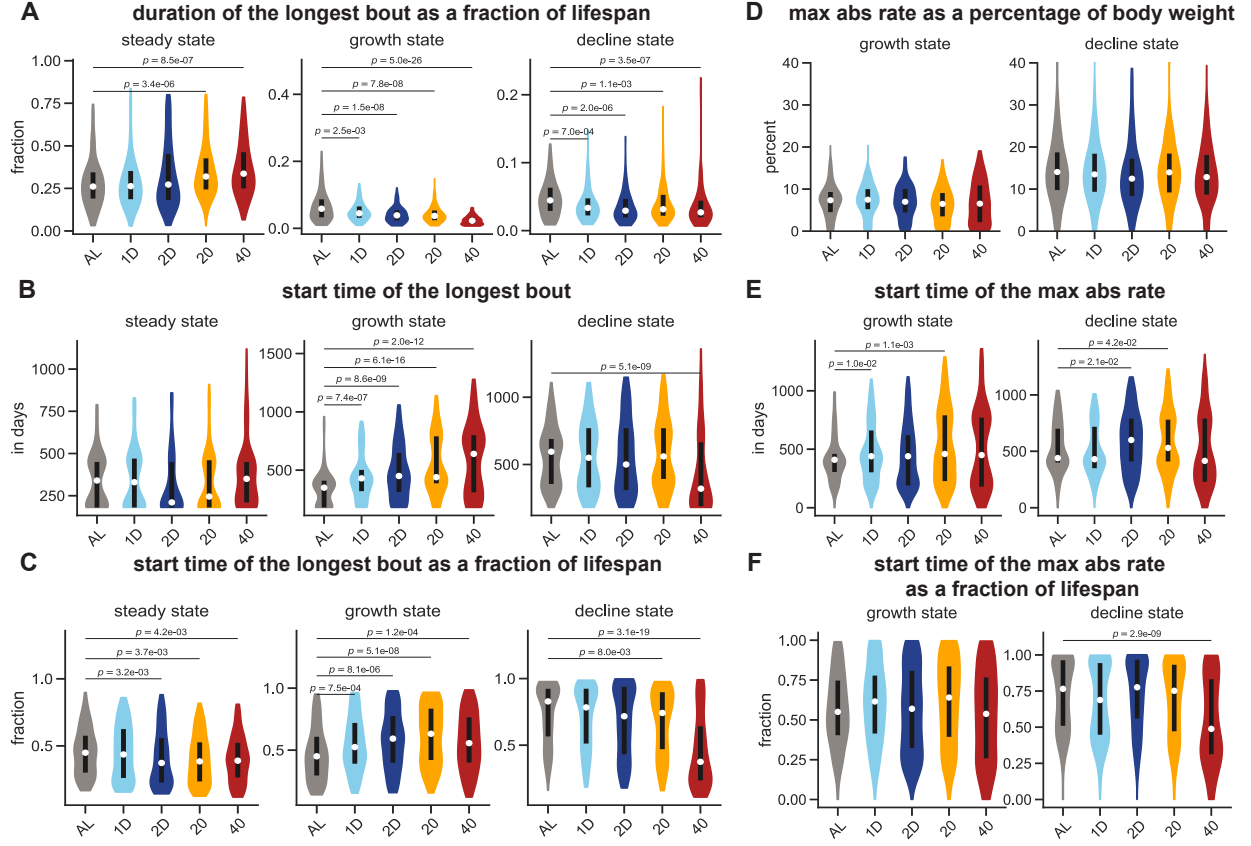
Supplementary Figures



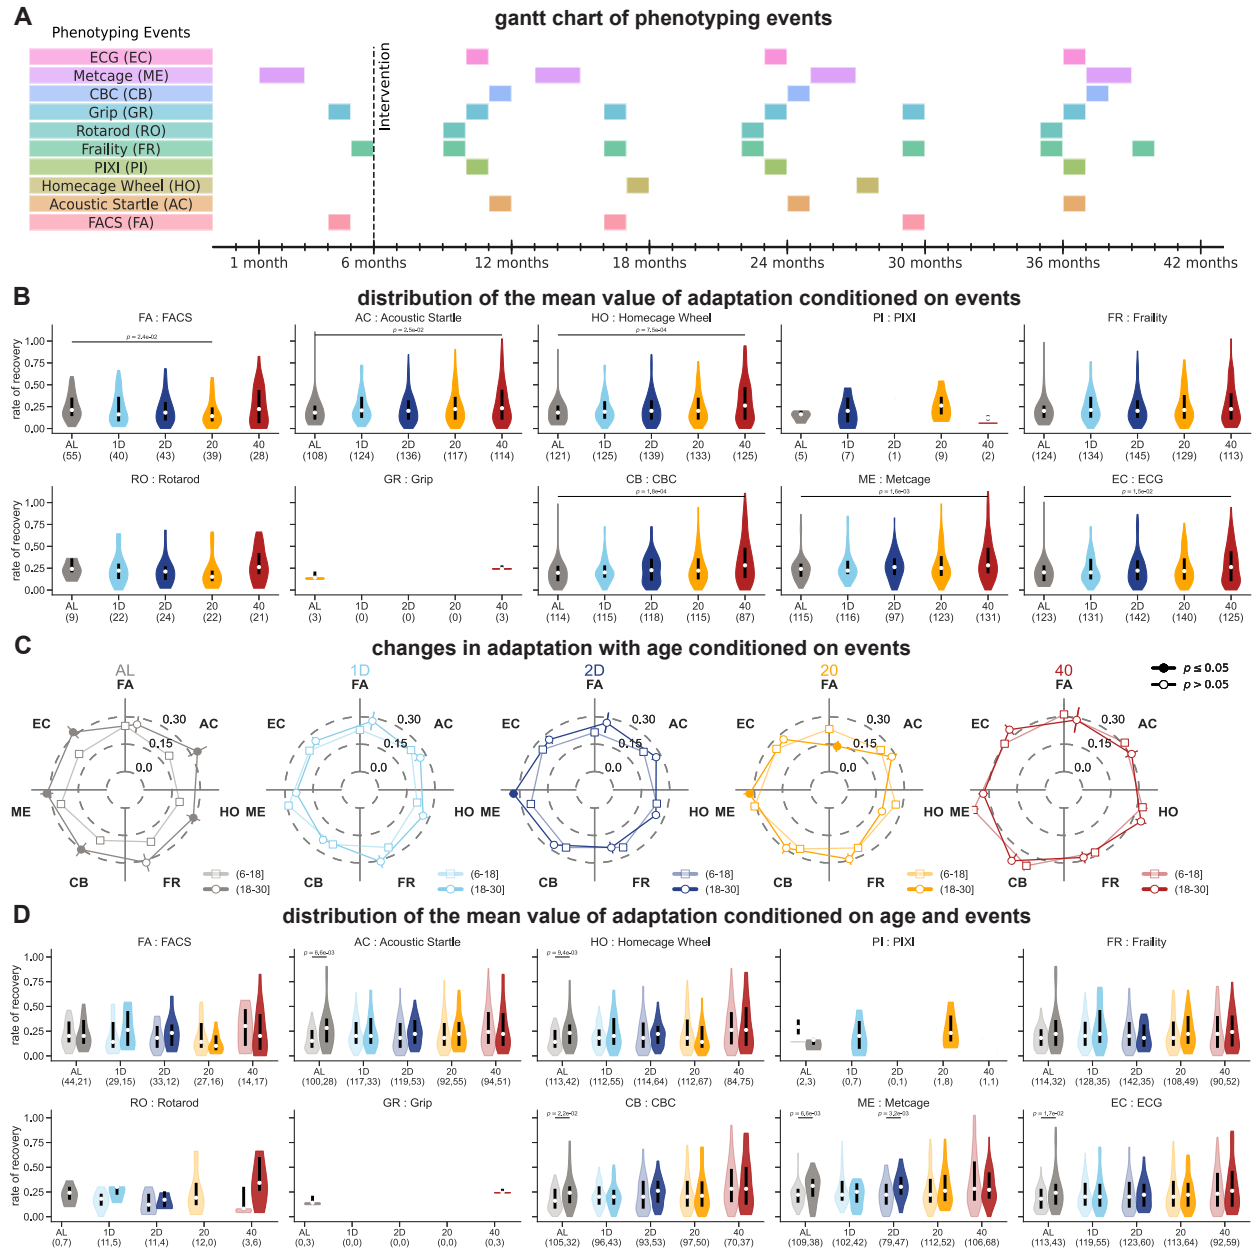
Supplemental Figure 1: **Variation inference-based autoregressive hidden Markov model (VI-ARHMM)**. (A) The mean value of the deviance information criteria (DIC) and standard error computed on the validation set for different model orders $K \in \{2, \dots, 10\}$ at ten different random state initializations. Model order $K = 3$ resulted in the smallest DIC. (B) The mean value of the evidence-based lower bound (ELBO) and standard error computed on the training set for different model orders at ten different random state initializations. Once again, model order $K = 3$ resulted in the highest ELBO. (C) The ELBO of the training data for model order $K = 3$ at different random state seed values. The highest ELBO was obtained at seed value of 8 (best model). (D) The VIARHMM approach showed convergence of the ELBO for the best model.



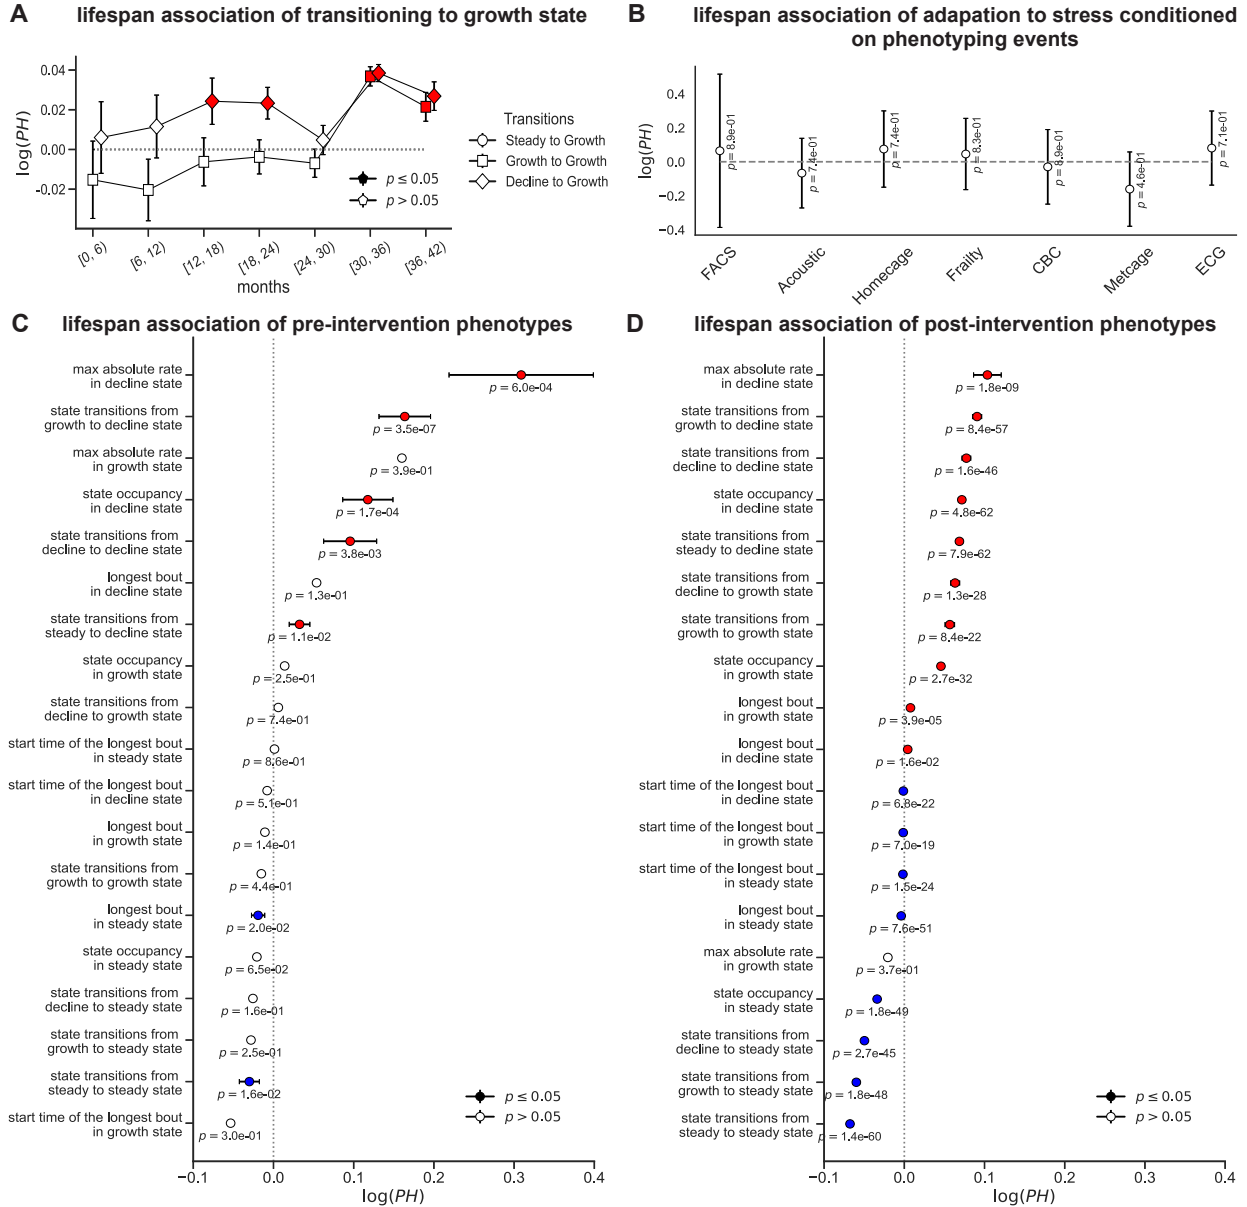
Supplemental Figure 2: **Influence of diet and age on the state stability.** (A) and (B) Transitioning into steady state at six month and ten percent of proportion of life lived interval bins, respectively. (C) and (D) Transitioning into growth state at six month and ten percent of proportion of life lived interval bins, respectively. (E) and (F) Transitioning into decline state at six month and ten percent of proportion of life lived interval bins, respectively. In (A)-(F), solid squares or circles indicate p -values < 0.05 , where the p -values were obtained by performing a Mann-Whitney test between a diet group and the AL group conditioned at the same interval bin.



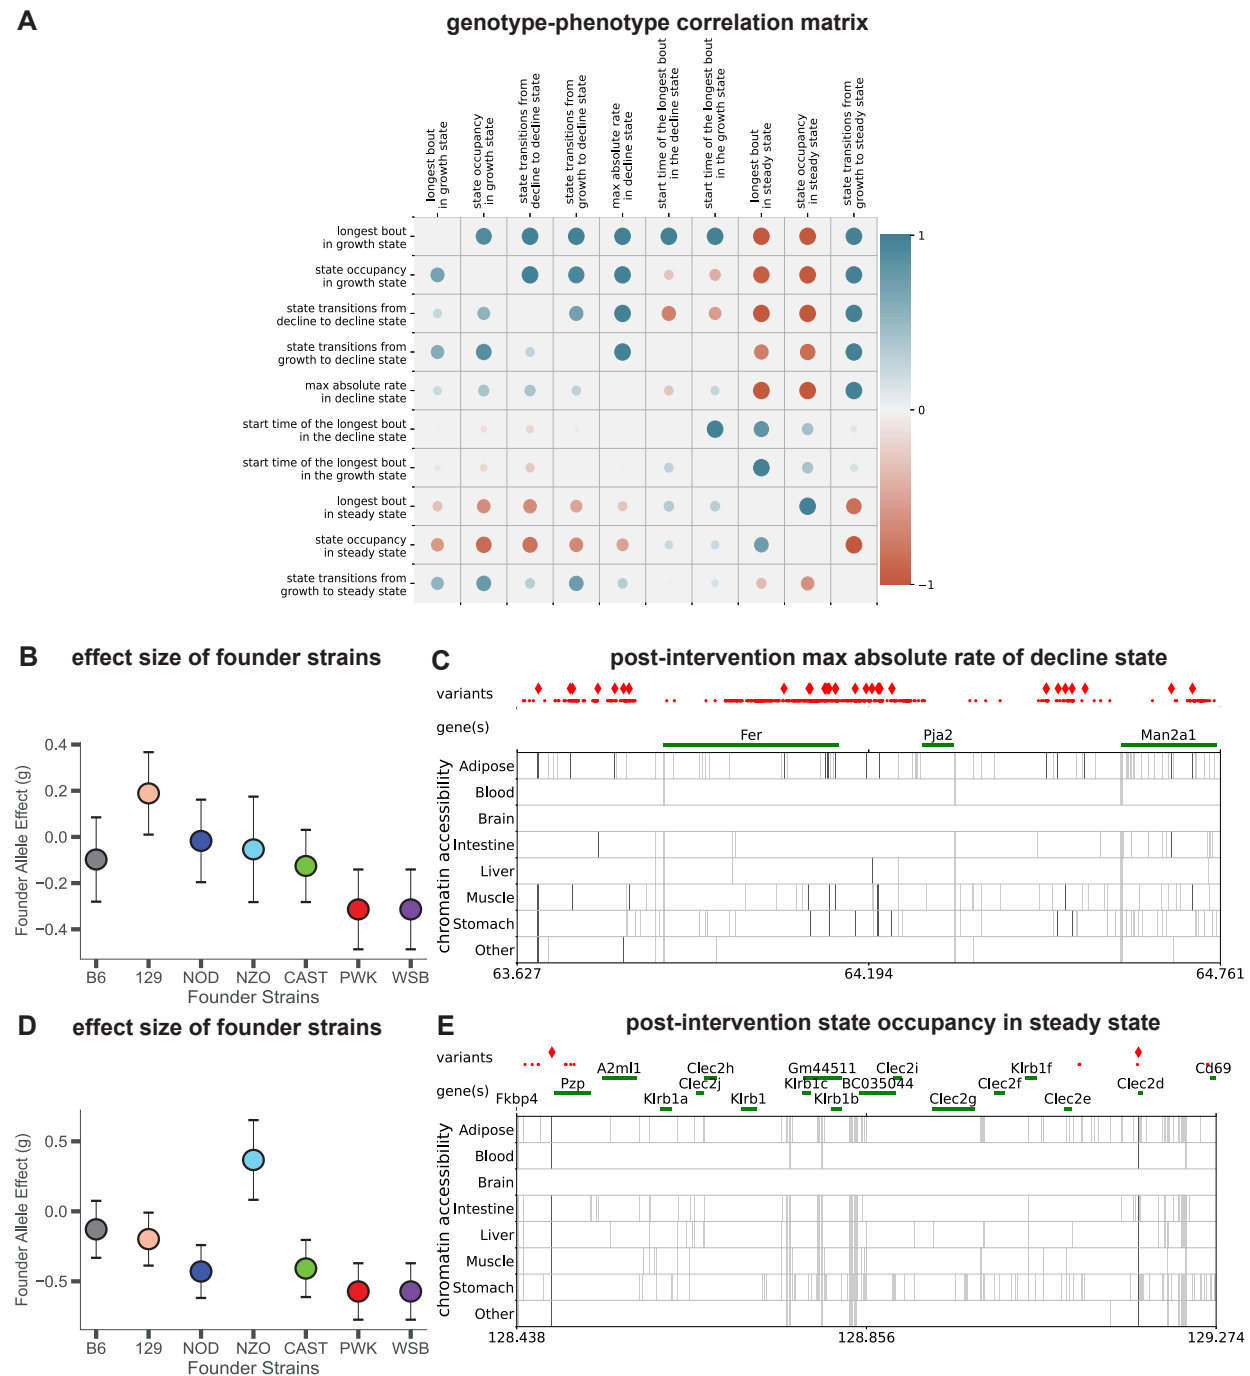
Supplemental Figure 3: **Robustness and perturbations due to body weight homeostasis.** (A) Fraction of life spent in the longest continuous bouts of growth, steady, and decline states. (B) The start time of the longest bout in growth, steady, and decline states. (C) Normalized time at which the start time of the longest bout were recorded. (D) Percentage of body weight gained and lost at the time of the maximum absolute rates of growth and decline states, respectively. (E) The time (in days) at which the maximum absolute rates of growth and decline states were recorded. (F) Normalized time at which the maximum absolute rates of growth and decline states were recorded.



Supplemental Figure 4: Adaptation to stress is influenced by diet and age. (A) A gantt chart of the different phenotypic assays a mouse undergoes in its lifespan. The approximate duration it takes to assay all mice for a particular assay is indicated using a color bar corresponding to the phenotypic assay. (B) Violin plots of the average values of the rate of recovery, conditioned on the phenotypic assay. (C) Radar charts of the average value and the standard-errors of the rate of recovery conditioned on phenotypic assay for each diet at two non-overlapping age-bins. (D) Violin plots of the average values of the rate of recovery, conditioned on the phenotypic assay and time interval. Intervals [6–18] and [18–30] months are represented in lighter and darker shades, respectively. In (B) and (D), the number of mice for which a perturbation event was registered following a phenotypic assay in the post-intervention phase is mentioned in the parenthesis below the diet. Only a small number of perturbation events were registered after PIXI, rotarod, and grip strength phenotypic assays.



Supplemental Figure 5: **Body weight-derived traits and its association with lifespan.** (A) Effect size and standard error of transitioning to growth state and its association with lifespan at every six-month interval bins. (B) Effect size and standard error of the mean value of adaptation to stress conditioned on the phenotyping event and its association with lifespan. (C) and (D) Effect sizes and standard errors of the pre- and post-intervention body weight-derived traits that are associated with either increasing or decreasing lifespan. The traits are arranged in descending order of the effect sizes. In (A)-(D), lifespan associations which were statistically significant (p -value < 0.05) were indicated with solid colors. If the effect size of the trait lies above or to the right of the dashed gray line (red filled markers), then an increase in the value of the trait decreases lifespan. Conversely, if it lies below or to the left the dashed gray line (blue filled markers), then an increase in the value of the trait increases lifespan.



Supplemental Figure 6: **Phenotypic-genotypic correlations and evidence supporting QTL.** (A) Pairwise genetic (upper triangle) and phenotypic (lower triangle) correlations of the body weight-derived traits in the post-intervention phase. (B) and (D) Effect size of founder allele strains observed at the variant with the highest LOD score for max absolute rate of decline and steady state occupancy, respectively, in the post-intervention phase. (C) and (E) Significant variants, colored by their founder allele pattern (FAP) group and the tissue-specific activity of regulatory elements near these variants (shown in gray). Significant variants that lie within regulatory elements are highlighted as diamonds, and regulatory elements that contain a significant variant are highlighted in black.

Appendix

Variational Inference-based Autoregressive Hidden Markov Model

960 diversity outbred female mice were subject to 5 dietary interventions. The body weight of each mouse was measured
 510 each week (approximately) starting from one month of age. Using these high resolution temporal measurements,
 511 we wish to learn interesting physiological and developmental stages in the life of a mouse. To identify the different
 512 physiological and development stages in the life of a mouse, we develop a variational inference-based auto-regressive
 513 model.

- If M denotes the number of mice enrolled in the study, we represent the body weight trace of the m -th mouse as
 514 $\mathbf{y}_{1:T_m} = \{y_1, \dots, y_{T_m}\}$, where T_m is the total number of body weight measurements for the m -th mouse and the
 515 body weight measurements are sorted based on the age of the mouse at the time of the measurement.
- Let $\mathbf{s}_{1:T_m} = \{s_1, \dots, s_{T_m}\}$ denote a sequence of hidden physiological states corresponding to the body weight
 517 measurements $\mathbf{y}_{1:T_m}$, generated by a first-order Markov process, where $s_i \in \{1, \dots, K\}$ and K is the total number
 518 of physiological states.
- Let $\boldsymbol{\pi}$ denote the initial probability vector of the first-order Markov process where,

$$\pi_j = p(s_L = j), \quad \forall 1 \leq j \leq K.$$

- Let \mathbf{A} denote the state transition matrix where transition between states is governed by Markov chain whose
 realizations take on values $\{1, \dots, K\}$ and the elements of the state transition matrix are given as:

$$a_{m,n} = p(s_t = n \mid s_{t-1} = m) \text{ such that } \sum_{n=1}^K a_{m,n} = 1, \quad \forall 1 \leq m \leq K,$$

520 where $a_{m,n}$ is the probability of transitioning from $s_{t-1} = m$ to $s_t = n$.

- The dynamics of the body weight measurements at a given physiological state s_t can be defined using an
 auto-regressive model as:

$$y_t^{s_t} = \phi_0^{s_t} + y_{t-1}^{s_t} \phi_1^{s_t} + \dots + y_{t-L}^{s_t} \phi_L^{s_t} + \epsilon^{s_t}, \quad \forall L < t \leq T_m,$$

where L is the lag order, $\phi_l^{s_t}$ is the auto-regressive coefficient for the l -th lag order, and $\epsilon^{s_t} \sim \mathcal{N}(0, \sigma_{s_t}^2)$ is zero mean additive white Gaussian noise.

- We assume that the mouse-specific parameters are unobserved and propose a hierarchical model on the dynamics parameters. If $\phi_l = [\phi_{l,1}, \dots, \phi_{l,K}]^\top$ denotes a vector of autoregressive coefficients for the l -th lag order, then $\phi_l \sim \mathcal{N}(\eta_l, \Sigma_l)$, where $\eta_l = [\eta_{l,1}, \dots, \eta_{l,K}]^\top$ and Σ_l is a diagonal covariance matrix with diagonal elements $\{\sigma_{l,1}^2, \dots, \sigma_{l,K}^2\}$. For notation convenience, we concatenate all auto-regressive coefficients into a long vector $\phi = [\phi_0^\top, \dots, \phi_L^\top]^\top$.

The joint likelihood can be written as:

$$p(\mathbf{y}_{1:T_m}, \mathbf{s}_{1:T_m}, \phi) = p(s_L) \prod_{t=L+1}^{T_m} p(s_t | s_{t-1}) P(y_t | y_{t-1:t-L}, s_t, \phi) \prod_{l=0}^L p(\phi_l | \eta_l, \Sigma_l). \quad (1)$$

The latent variables $\mathbf{s}_{1:T_m}$ and ϕ are not observed. Therefore, we integrate the latent variables and rewrite the complete log-likelihood as:

$$\begin{aligned} \log p(\mathbf{y}_{1:T_m}) &= \log \int_{\phi} \sum_{\mathbf{s}_{1:T_m}} p(\mathbf{y}_{1:T_m}, \mathbf{s}_{1:T_m}, \phi) d\phi, \\ &= \log \int_{\phi} \sum_{\mathbf{s}_{1:T_m}} \frac{p(\mathbf{y}_{1:T_m}, \mathbf{s}_{1:T_m}, \phi)}{q(\mathbf{s}_{1:T_m}, \phi)} q(\mathbf{s}_{1:T_m}, \phi) d\phi, \\ &= \log \mathbb{E}_q \left[\frac{p(\mathbf{y}_{1:T_m}, \mathbf{s}_{1:T_m}, \phi)}{q(\mathbf{s}_{1:T_m}, \phi)} \right], \\ &\geq \mathbb{E}_q \left[\log \frac{p(\mathbf{y}_{1:T_m}, \mathbf{s}_{1:T_m}, \phi)}{q(\mathbf{s}_{1:T_m}, \phi)} \right], \quad (\text{Jensen's inequality}) \\ &= \mathbb{E}_q [\log p(\mathbf{y}_{1:T_m}, \mathbf{s}_{1:T_m}, \phi)] - \mathbb{E}_q [\log q(\mathbf{s}_{1:T_m}, \phi)]. \quad (\text{Evidence-based lower bound}) \end{aligned} \quad (2)$$

The log-likelihood is maximized when $q(\mathbf{s}_{1:T_m}, \phi) = p(\mathbf{s}_{1:T_m}, \phi)$. However, the joint probability of $p(\mathbf{s}_{1:T_m}, \phi)$ is intractable. The goal of variational inference algorithm is to find the distribution $p(\mathbf{s}_{1:T_m}, \phi)$ from the chosen variational family that maximizes the lower bound to the log marginal likelihood. We want to optimize the evidence-based lower bound (ELBO) over a chosen space of variational distributions, to find the variational distribution closest to the true posterior distribution $p(\mathbf{s}_{1:T_m}, \phi)$. Therefore, we replace $p(\mathbf{s}_{1:T_m}, \phi)$ with $q(\mathbf{s}_{1:T_m}, \phi)$ such that $q(\mathbf{s}_{1:T_m}, \phi)$ is tractable and $p(\mathbf{s}_{1:T_m}, \phi) \approx q(\mathbf{s}_{1:T_m}, \phi)$. Further, to perform this optimization, we assume that (a) the states and the autoregressive coefficients are independent, (b) autoregressive coefficients are independent, and (c) the transition between states is

governed by a first-order Markov process. Under these assumptions,

$$\begin{aligned} q(\mathbf{s}_{1:T_m}, \boldsymbol{\phi}) &= q(\mathbf{s}_{1:T_m})q(\boldsymbol{\phi}), \\ &= q(\mathbf{s}_{1:T_m}) \prod_{l=0}^L q(\boldsymbol{\phi}_l \mid \tilde{\boldsymbol{\eta}}_l, \tilde{\boldsymbol{\Sigma}}_l), \end{aligned}$$

where $q(\boldsymbol{\phi}_l \mid \tilde{\boldsymbol{\eta}}_l, \tilde{\boldsymbol{\Sigma}}_l) \sim \mathcal{N}(\tilde{\boldsymbol{\eta}}_l, \tilde{\boldsymbol{\Sigma}}_l)$, $\tilde{\boldsymbol{\eta}}_l = [\tilde{\eta}_{l,1}, \dots, \tilde{\eta}_{l,K}]^\top$ is the mean, $\tilde{\boldsymbol{\Sigma}}_l$ is a diagonal covariance matrix with diagonal elements $\{\tilde{\sigma}_{l,1}^2, \dots, \tilde{\sigma}_{l,K}^2\}$,

$$q(\mathbf{s}_{1:T_m}) \propto \prod_{j=1}^K q(s_L = j)^{\delta(s_L=j)} \prod_{t=L+1}^T \prod_{m=1}^K \prod_{n=1}^K q(s_t = n \mid s_{t-1} = m)^{\delta(s_{t-1}=m, s_t=n)},$$

and $\delta(\cdot)$ acts as a masking function which is set to one if the condition is true else zero. We expand the terms in the ELBO expression in (2) and rewrite it as:

$$\begin{aligned} \text{ELBO} &= \sum_{j=1}^K q(s_L = j) \log p(s_L) + \sum_{t=L+1}^{T_m} \sum_{m=1}^K \sum_{n=1}^K q(s_t = n, s_{t-1} = m) \log p(s_t \mid s_{t-1}) \\ &\quad + \sum_{t=L+1}^{T_m} \sum_{j=1}^K q(s_t = j) \langle \log p(y_t \mid y_{t-1:t-L}, s_t, \boldsymbol{\phi}) \rangle_{q(\boldsymbol{\phi})} + \sum_{l=0}^L q(\boldsymbol{\phi}_l \mid \tilde{\boldsymbol{\eta}}_l, \tilde{\boldsymbol{\Sigma}}_l) \log p(\boldsymbol{\phi}_l \mid \boldsymbol{\eta}_l, \boldsymbol{\Sigma}_l) \\ &\quad - \sum_{j=1}^K q(s_L = j) \log q(s_L) - \sum_{t=L+1}^{T_m} \sum_{m=1}^K \sum_{n=1}^K q(s_t = n, s_{t-1} = m) \log q(s_t \mid s_{t-1}) \\ &\quad - \sum_{l=0}^L q(\boldsymbol{\phi}_l \mid \tilde{\boldsymbol{\eta}}_l, \tilde{\boldsymbol{\Sigma}}_l) \log q(\boldsymbol{\phi}_l \mid \tilde{\boldsymbol{\eta}}_l, \tilde{\boldsymbol{\Sigma}}_l). \end{aligned} \quad (3)$$

The above expression can be further simplified as:

$$\begin{aligned} \text{ELBO} &= \langle \log p(s_L) - \log q(s_L) \rangle_{q(s_L)} + \sum_{t=L+1}^{T_m} \langle \log p(s_t \mid s_{t-1}) - \log q(s_t \mid s_{t-1}) \rangle_{q(s_t, s_{t-1})} \\ &\quad + \sum_{t=L+1}^{T_m} \langle \langle \log p(y_t \mid y_{t-1:t-L}, s_t, \boldsymbol{\phi}) \rangle_{q(\boldsymbol{\phi})} \rangle_{q(s_t)} + \sum_{l=0}^L \langle \log p(\boldsymbol{\phi}_l \mid \boldsymbol{\eta}_l, \boldsymbol{\Sigma}_l) - \log q(\boldsymbol{\phi}_l \mid \tilde{\boldsymbol{\eta}}_l, \tilde{\boldsymbol{\Sigma}}_l) \rangle_{q(\boldsymbol{\phi}_l)}, \end{aligned} \quad (4)$$

where $\langle \cdot \rangle_{q(\cdot)}$ is the expectation with respect to $q(\cdot)$. Furthermore, based on the normality assumptions made on $q(\boldsymbol{\phi}_l \mid \tilde{\boldsymbol{\eta}}_l, \tilde{\boldsymbol{\Sigma}}_l)$, we get

$$\langle \phi_{l,j} \rangle_{q(\phi_{l,j})} = \tilde{\eta}_{l,j}, \quad (5)$$

$$\langle \phi_{l,j}^2 \rangle_{q(\phi_{l,j})} = \tilde{\eta}_{l,j}^2 + \tilde{\sigma}_{l,j}^2. \quad (6)$$

Substituting (5) and (6) in (4) gives the complete expression of ELBO which is given as:

$$\begin{aligned}
\text{ELBO} = & \langle \log p(s_L) - \log q(s_L) \rangle_{q(s_L)} + \sum_{t=L+1}^{T_m} \langle \log p(s_t | s_{t-1}) - \log q(s_t | s_{t-1}) \rangle_{q(s_t, s_{t-1})} \\
& - \sum_{t=L+1}^{T_m} \sum_{j=1}^K q(s_t = j) \left[\frac{1}{2} \log \sigma_j^2 + \frac{1}{2\sigma_j^2} \left(y_t^2 - 2y_t \tilde{\eta}_{0,j} - 2y_t \sum_{l=1}^L \tilde{\eta}_{l,j} y_{t-l} + \right. \right. \\
& \left. \left. \tilde{\eta}_{0,j}^2 + \tilde{\sigma}_{0,j}^2 + \sum_{l=1}^L (\tilde{\eta}_{l,j}^2 + \tilde{\sigma}_{l,j}^2) y_{t-l}^2 + 2 \sum_{l=1}^L \tilde{\eta}_{l,0} \tilde{\eta}_{l,j} y_{t-l} + 2 \sum_{p=1}^K \sum_{\substack{q=1 \\ p \neq q}}^K \tilde{\eta}_{p,j} \tilde{\eta}_{q,j} y_{t-p} y_{t-q} \right) \right] \\
& + \sum_{l=0}^L \sum_{j=1}^K \frac{1}{2} \left[1 + \log \tilde{\sigma}_{l,j}^2 - \log \sigma_{l,j}^2 - \frac{\tilde{\sigma}_{l,j}^2 + (\tilde{\eta}_{l,j} - \eta_{l,j})^2}{\sigma_{l,j}^2} \right]
\end{aligned} \tag{7}$$

528 Maximizing the variational parameters of the ELBO by keeping the model parameters fixed gives us the variational
529 E-step and M-step, and maximizing the model parameters of the ELBO by keeping the variational parameters fixed
530 gives us the M-step. These steps are repeated until the convergence criteria is met.

531 Variational Bayes E-Step

In the variational Bayes E-step, the mouse-specific and model parameters are fixed and the ELBO is maximized over the mouse-specific variational marginal and joint probabilities. We used the forward-backward algorithm to compute the marginal and joint variational probabilities $q(s_t = j)$ and $q(s_t = n, s_{t-1} = m)$, respectively. The compact expression of ELBO in (2) can be rewritten in a functional form as:

$$\begin{aligned}
\text{ELBO} &= \int d\boldsymbol{\phi} \sum_{\mathbf{s}_{1:T_m}} q(\mathbf{s}_{1:T_m}, \boldsymbol{\phi}) \log \frac{p(\mathbf{y}_{1:T_m}, \mathbf{s}_{1:T_m} | \boldsymbol{\phi}) p(\boldsymbol{\phi})}{q(\mathbf{s}_{1:T_m}, \boldsymbol{\phi})} \\
&= \int d\boldsymbol{\phi} q(\boldsymbol{\phi}) \left[\log \frac{p(\boldsymbol{\phi})}{q(\boldsymbol{\phi})} + \sum_{\mathbf{s}_{1:T_m}} q(\mathbf{s}_{1:T_m}) \log \frac{p(\mathbf{y}_{1:T_m}, \mathbf{s}_{1:T_m} | \boldsymbol{\phi})}{q(\mathbf{s}_{1:T_m})} \right] \\
&= \mathcal{F}(q(\boldsymbol{\phi}) q(\mathbf{s}_{1:T_m})).
\end{aligned} \tag{8}$$

Taking the functional derivative of (8) with respect to $q(\mathbf{s}_{1:T_m})$ gives:

$$\log q(\mathbf{s}_{1:T_m}) = \langle \log p(\mathbf{y}_{1:T_m}, \mathbf{s}_{1:T_m} | \boldsymbol{\phi}) \rangle_{q(\boldsymbol{\phi})} - \log \tilde{\mathcal{Z}}(\mathbf{y}_{1:T_m}) \tag{9}$$

where $\tilde{\mathcal{Z}}(\mathbf{y}_{1:T_m})$ is a normalization constant. The normalization constant enables the feasibility of computing $q(s_t = j)$ and $q(s_t = n, s_{t-1} = m)$ using forward-backward algorithm. Unlike the hidden Markov model which requires the joint log-likelihood, the re-parameterized implementation of the forward-backward algorithm requires the expectation of the

partial log-likelihood. Moreover, based on the assumptions made on the variational parameters, we can easily compute the closed-form expression of the expectation of the partial log-likelihood with respect to $q(\phi)$ as follows:

$$\begin{aligned} \log q(s_{1:T}) &= \left\langle \log \left[p(s_L) \prod_{t=L+1}^{T_m} p(s_t | s_{t-1}) p(y_t | y_{t-1:t-L}, s_t, \phi) \right] \right\rangle_{q(\phi)} - \log \tilde{\mathcal{Z}}(y_{1:T}) \\ &= \sum_{j=1}^K q(s_L = j) \log \pi_j + \sum_{t=L+1}^{T_m} \sum_{m=1}^K \sum_{n=1}^K q(s_{t-1} = m, s_t = n) \log a_{m,n} \\ &\quad + \sum_{t=L+1}^{T_m} \sum_{j=1}^K q(s_t = j) \langle \log p(y_t | y_{t-1:t-L}, s_t, j, \phi) \rangle_{q(\phi)} - \log \tilde{\mathcal{Z}}(y_{1:T}), \end{aligned}$$

where

$$\begin{aligned} \langle \log p(y_t | y_{t-1:t-L}, s_t, j, \phi) \rangle_{q(\phi)} &= - \left[\frac{1}{2} \log \sigma_j^2 + \frac{1}{2\sigma_j^2} \left(y_t^2 - 2y_t \tilde{\eta}_{0,j} - 2y_t \sum_{l=1}^L \tilde{\eta}_{l,j} y_{t-l} + \right. \right. \\ &\quad \left. \left. \tilde{\eta}_{0,j}^2 + \tilde{\sigma}_{0,j}^2 + \sum_{l=1}^L (\tilde{\eta}_{l,j}^2 + \tilde{\sigma}_{l,j}^2) y_{t-l}^2 + 2 \sum_{l=1}^L \tilde{\eta}_{l,0} \tilde{\eta}_{l,j} y_{t-l} + 2 \sum_{p=1}^K \sum_{\substack{q=1 \\ p \neq q}}^K \tilde{\eta}_{p,j} \tilde{\eta}_{q,j} y_{t-p} y_{t-q} \right) \right]. \end{aligned} \quad (10)$$

The output of the forward-backward algorithm generates forward and backward probabilities, which are used to compute the variational marginal and joint probabilities for a specific mouse. The main advantage of the reparameterization step is that we can avoid computing the partial derivatives of the marginal and joint probabilities, which are often computationally intensive to solve.

Variational Bayes M-Step

In the variational Bayes M-step, the mouse-specific variational probabilities and model parameters are fixed and the ELBO is maximized over the variational model parameters.

- Variational initial probability $q(s_L)$: To get $q(s_L)$, we solve the following optimization problem:

$$\frac{\partial}{\partial q(s_L)} \left[\sum_{j=1}^K q(s_L) \log p(s_L) - \sum_{j=1}^K q(s_L) \log q(s_L) + \lambda \left(\sum_{j=1}^K q(s_L) - 1 \right) \right] = 0.$$

On solving, we get:

$$q(s_L = j) = \frac{p(s_L = j)}{\sum_{j=1}^K p(s_L = j)} \quad \forall j \in \{1, \dots, K\}. \quad (11)$$

- Variational transition probability $q(s_t | s_{t-1})$: To get $q(s_t | s_{t-1})$, we solve the following optimization problem:

$$\frac{\partial}{\partial \tilde{a}_{m,n}} \left[\sum_{t=L+1}^{T_m} \sum_{n=1}^K \sum_{m=1}^K q(s_t = n, s_{t-1} = m) \log a_{m,n} - \sum_{t=L+1}^{T_m} \sum_{n=1}^K \sum_{m=1}^K q(s_t = n, s_{t-1} = m) \log \tilde{a}_{m,n} + \lambda \left(\sum_{n=1}^K \tilde{a}_{m,n} - 1 \right) \right] = 0.$$

On solving, we get:

$$q(s_t = n | s_{t-1} = m) = \tilde{a}_{m,n} = \frac{\sum_{t=L+1}^{T_m} q(s_t = n, s_{t-1} = m)}{\sum_{t=L+1}^{T_m} \sum_{n=1}^K q(s_t = n, s_{t-1} = m)}, \quad \forall m, n \in \{1, \dots, K\}. \quad (12)$$

- Variational variance $\tilde{\sigma}_{l,j}^2$:

$$\frac{1}{\tilde{\sigma}_{0,j}^2} = \frac{1}{\sigma_{0,j}^2} + \frac{1}{\sigma_j^2} \sum_{t=L+1}^{T_m} q(s_t = j), \quad \forall j \in \{1, \dots, K\}, \quad (13)$$

$$\frac{1}{\tilde{\sigma}_{l,j}^2} = \frac{1}{\sigma_{l,j}^2} + \frac{1}{\sigma_j^2} \sum_{t=L+1}^{T_m} q(s_t = j) y_{t-l}^2, \quad \forall j \in \{1, \dots, K\}, l \in \{1, \dots, L\}. \quad (14)$$

- Variational mean $\tilde{\eta}_{l,j}$:

$$\tilde{\eta}_{0,j} = \frac{\frac{\eta_{0,j}}{\sigma_{0,j}^2} + \frac{1}{\sigma_s^2} \sum_{t=L+1}^{T_m} q(s_t = j) (y_t - \sum_{l=1}^L \tilde{\eta}_{l,j} y_{t-l})}{\frac{1}{\sigma_{0,j}^2} + \frac{1}{\sigma_j^2} \sum_{t=L+1}^{T_m} q(s_t = j)}, \quad \forall j \in \{1, \dots, K\}, \quad (15)$$

$$\tilde{\eta}_{l,j} = \frac{\frac{\eta_{l,j}}{\sigma_{l,j}^2} + \frac{1}{\sigma_s^2} \sum_{t=L+1}^{T_m} q(s_t = j) y_{t-l} (y_t - \tilde{\eta}_{0,j} - \frac{1}{2} \sum_{p=1}^L \tilde{\eta}_{p,j} y_{t-p})}{\frac{1}{\sigma_{0,j}^2} + \frac{1}{\sigma_j^2} \sum_{t=L+1}^{T_m} q(s_t = j) y_{t-l}^2}, \quad \forall j \in \{1, \dots, K\}, \forall l \in \{1, \dots, L\}. \quad (16)$$

539 M-Step

540 In the M-step, the mouse-specific variational parameters are fixed and the ELBO is maximized over model parameters.

541 We used subscript i to indicate the variational parameters from the i -th mouse.

- Initial probability $p(s_L)$:

$$p(s_L = j) = \frac{\sum_{i=1}^M q_i(s_L = j)}{\sum_{i=1}^M \sum_{j=1}^K q_i(s_L = j)}, \quad \forall j \in \{1, \dots, K\}. \quad (17)$$

- Transition probability $p(s_t | s_{t-1})$:

$$p(s_t = n | s_{t-1} = m) = a_{m,n} = \frac{\sum_{i=1}^M \sum_{t=L+1}^{T_m} [\tilde{a}_{m,n}]_i q_i(s_{t-1} = m)}{\sum_{i=1}^M \sum_{t=L+1}^T \sum_{n=1}^K [\tilde{a}_{m,n}]_i q_i(s_{t-1} = m)} \quad \forall m, n \in \{1, \dots, K\}. \quad (18)$$

- Error variance σ_j^2

$$\sigma_j^2 = \frac{\sum_{i=1}^M \sum_{t=L+1}^{T_m} q_i(s_t = j) [\psi_{j,t}]_i}{\sum_{i=1}^M \sum_{t=L+1}^T q_i(s_t = j)}, \quad \forall j \in \{1, \dots, K\}, \quad (19)$$

where

$$[\psi_{j,t}]_i = \left([y_t^2]_i - 2[y_t]_i [\tilde{\eta}_{0,j}]_i - 2[y_t]_i \sum_{l=1}^L [\tilde{\eta}_{l,j}]_i [y_{t-l}]_i [\tilde{\eta}_{0,j}^2]_i + [\tilde{\sigma}_{0,j}^2]_i + \sum_{l=1}^L ([\tilde{\eta}_{l,j}^2]_i + [\tilde{\sigma}_{l,j}^2]_i) [y_{t-l}^2]_i + 2 \sum_{l=1}^L [\tilde{\eta}_{l,0}]_i [\tilde{\eta}_{l,j}]_i [y_{t-l}]_i + 2 \sum_{p=1}^K \sum_{\substack{q=1 \\ p \neq q}}^K [\tilde{\eta}_{p,j}]_i [\tilde{\eta}_{q,j}]_i [y_{t-p}]_i [y_{t-q}]_i \right).$$

- Auto-regressive coefficient parameters $(\eta_{l,j}, \sigma_{l,j}^2)$:

$$\eta_{l,j} = \frac{1}{M} \sum_{i=1}^M [\tilde{\eta}_{l,j}]_i, \quad \forall j \in \{1, \dots, K\}, \forall l \in \{1, \dots, L\}, \quad (20)$$

$$\sigma_{l,j}^2 = \frac{1}{M} \sum_{i=1}^M [\tilde{\sigma}_{l,j}^2]_i + ([\tilde{\eta}_{l,j}]_i - \eta_{l,j})^2, \quad \forall j \in \{1, \dots, K\}, \forall l \in \{1, \dots, L\}. \quad (21)$$

542 Model Selection

We used deviance information criteria (DIC) to select the model order for the variational-inference based auto-regressive hidden Markov model (Spiegelhalter et al. (2002)). The DIC combines model complexity and fit, where the model complexity is obtained by taking difference between the posterior mean of the deviance and the deviance calculated at the posterior mean $\tilde{\phi}$, and fit is obtained from the log-likelihood computed at the posterior mean $\tilde{\phi}$. The DIC is expressed as:

$$\text{DIC} = 2p_D - 2 \log p(\mathbf{y} | \phi),$$

where

$$p_D = -2 \int q(\boldsymbol{\phi}) \log \frac{q(\boldsymbol{\phi})}{p(\boldsymbol{\phi})} d\boldsymbol{\phi} + 2 \log \frac{q(\tilde{\boldsymbol{\phi}})}{p(\tilde{\boldsymbol{\phi}})}.$$

Based on the assumptions made about the model and variational parameters,

$$\log p(\boldsymbol{\phi} \mid \boldsymbol{\eta}, \boldsymbol{\sigma}^2) = -\frac{1}{2} \sum_{l=0}^L \sum_{j=1}^K \left[\log(2\pi\sigma_{l,j}^2) + \frac{1}{\sigma_{l,j}^2} (\phi_{l,j} - \eta_{l,j})^2 \right] \quad (22)$$

$$\log q(\boldsymbol{\phi} \mid \tilde{\boldsymbol{\eta}}, \tilde{\boldsymbol{\sigma}}^2) = -\frac{1}{2} \sum_{l=0}^L \sum_{j=1}^K \left[\log(2\pi\tilde{\sigma}_{l,j}^2) + \frac{1}{\tilde{\sigma}_{l,j}^2} (\phi_{l,j} - \tilde{\eta}_{l,j})^2 \right]. \quad (23)$$

Using (22) and (23), p_D can be simplified as:

$$p_D = \sum_{l=0}^L \sum_{j=1}^K \left\{ \left[\frac{1}{\sigma_{l,j}^2} (\tilde{\sigma}_{l,j}^2 + (\eta_{l,j} - \tilde{\eta}_{l,j})^2) \right] - \left[\frac{1}{\tilde{\sigma}_{l,j}^2} (\eta_{l,j} - \tilde{\eta}_{l,j})^2 \right] \right\}.$$

We computed the deviance information criteria at ten different seed values for each model order. The model order with the smallest deviance information criteria is considered as the best model order. To select the best seed value at a given model order, we computed the ELBO and choose the seed value with the highest ELBO.