

Is a Scalar Health State Sufficient for Health Dynamics? Evidence from Health Deficit Accumulation *

Astrid Krenz [†]

Holger Strulik [‡]

September 28, 2026

Abstract. Can a scalar measure of current health provide a sufficiently rich state for modeling health dynamics? Using longitudinal data from the Health and Retirement Study, we represent health by a count of 31 deficits, yielding a scalar state with 32 ordered values. Building on a stochastic deficit-accumulation model, we study transitions among deficit states and death and test whether information outside the current state improves prediction. For morbidity among survivors, current deficit count is an approximately, but not strictly, sufficient predictive state. Chronological age and observed demographic and socioeconomic characteristics add little predictive information once current health is known, but recent health history retains a quantitatively meaningful contribution. Mortality reveals a different boundary: recent deficit-count history adds little, while chronological age, selected individual characteristics, and deficit composition remain informative.

Keywords: health dynamics, health deficits, frailty index, mortality

JEL Codes: I10, I14, C33

* We thank Volker Grossmann for helpful comments.

[†] London School of Economics and Political Science, Data Science Institute, Houghton Street, London, WC2A 2AE, UK, email: a.krenz@lse.ac.uk. Declarations of interest: none.

[‡] University of Göttingen, Department of Economics, Platz der Göttinger Sieben 3, 37073 Göttingen, Germany; email: holger.strulik@wiwi.uni-goettingen.de. Declarations of interest: none.

1. INTRODUCTION

Dynamic models in health economics require a state variable: a description of current health that contains enough information to predict the distribution of future health outcomes. The underlying information is inherently multidimensional. Individuals differ in diseases, symptoms, functional limitations, disabilities, prior shocks, accumulated exposures, and latent robustness. Yet tractable economic models generally require a much simpler representation. The central question is therefore not only whether current health predicts future health, but whether multidimensional health information can be compressed into a scalar state without sacrificing most of the information needed to predict its subsequent evolution.

The deficit-accumulation approach provides an unusually direct setting in which to study this question. A wide range of diseases, symptoms, functional limitations, and disabilities are coded as health deficits and summarized by their count, n . Dividing that count by the number of measured items produces the frailty index. Developed by Mitnitski, Rockwood, and various coauthors across a series of papers building on Mitnitski et al. (2001), the frailty index captures the share of a broad set of aging-related health conditions that an individual exhibits, spanning deficits from minor to life-threatening. To qualify for inclusion in the index, a health deficit generally must become more prevalent with age, correlate with overall health, avoid saturating too soon, and represent a range of organ systems (see Searle et al., 2008, for a methodological review).

Several studies have shown that the particular deficits chosen for the index have relatively little bearing on its predictive performance, provided the list is sufficiently broad (Mitnitski et al., 2005; Goggins et al., 2005; Theou et al., 2013). This robustness has been related to the interconnected and redundant nature of health deficits, with individual deficits reflecting broader clusters of health problems (Mitnitski et al., 2002; Rutenberg, 2018). Related work emphasizes the network structure of disease and biological damage (Barabasi et al., 2011), the accumulation and propagation of molecular damage (Belikov, 2019), and interactions in aging across organ systems (Tian et al., 2023). The frailty index has consequently been proposed as a summary measure of overall health and biological aging (Rockwood and Mitnitski, 2007; Howlett et al., 2021).

For dynamic economic modeling, the important feature of this approach is the distinction between the dimensionality of the underlying health information and the dimensionality of the state variable. In our application, 31 binary health deficits form a multidimensional description of health, but they are mapped into a single integer, $n_t = \sum_{j=1}^{31} d_{jt}$. The individual health state is therefore scalar, while

its state space remains comparatively rich: n_t can take 32 ordered values from zero to 31. In the stochastic model, these health states are supplemented by death as an absorbing state. This differs from representing health by only a few categories such as good, fair, or poor. Both approaches use a one-dimensional health state, but the deficit count retains a much finer ordering of observed health. The question we ask is whether this compression is sufficient: once the current deficit count is known, how much predictive information about subsequent health remains elsewhere?

Mitnitski et al. (2006) took an important step toward answering this question by modeling transitions between deficit-count states as a Markov chain. Conditional on the current number of deficits, the subsequent deficit count followed a simple state-dependent stochastic law, while mortality rose with deficit burden. In the Canadian Study of Health and Aging (CSHA), the model closely fits the transition matrix, and the estimated relationships were similar across two successive five-year intervals. The original paper also explicitly asked how chronological age influences transitions relative to the number of deficits and found deficit count to be more important than age in predicting mortality. The model therefore provides a natural starting point for asking whether current deficit count can serve more generally as the state variable for subsequent health dynamics.

This question matters for dynamic models of health, insurance, labor supply, retirement, medical spending, and other economic decisions in which health is one component of a larger state space. A scalar health state can substantially reduce the dimensionality of such problems. But scalar need not mean coarse. The deficit approach potentially combines a one-dimensional state variable at the individual level with a comparatively fine-grained representation of heterogeneity in observed health. Its usefulness depends on how much relevant information is lost in the compression. If age, individual characteristics, deficit composition, or health history continue to predict transitions conditional on n_t , then current deficit count is not a strictly sufficient state and particular applications may require an augmented representation.

We investigate these questions using longitudinal data from the Health and Retirement Study (HRS). We first reproduce the stochastic architecture of Mitnitski et al. (2006) using a count of 31 health deficits and death as an absorbing state. The repeated approximately biennial HRS observations provide multiple transition periods and substantially greater support in the upper part of the deficit distribution than the original CSHA application. This allows us both to examine the stability of the basic transition law and to stress-test restrictions that are difficult to evaluate over the narrower support of the original data. We introduce minimal refinements that respect the finite state space,

distinguish mortality not captured by measured deficit count from deficit-dependent mortality, and accommodate modest persistence in the aggregate deficit state.

Our main contribution is to use this framework to test whether current deficit count is a sufficient predictive state. The comparison is deliberately asymmetric. A useful state variable may already embody information generated by chronological age, previous health, socioeconomic circumstances, and persistent individual heterogeneity. We therefore do not ask how predictive information shared by n_t and another variable should be divided between them. Instead, we ask whether each candidate information set contributes additional predictive information once current n_t is observed. This distinction separates strict sufficiency from approximate predictive sufficiency and avoids giving the resulting conditional associations a causal interpretation.

We consider three increasingly demanding sources of information outside the current state. We begin with chronological age, both because it is the benchmark closest to Mitnitski et al. (2006) and because conditional age comparisons illustrate the role of dynamic selection. An older individual observed with few deficits is a selected survivor and need not be comparable to a younger individual with the same n_t . We next examine sex, education, and race/ethnicity as predictors of heterogeneous transition processes rather than causal treatments. Race and ethnicity in particular are interpreted as markers of social stratification and accumulated exposures, not biological types. Finally, we examine recent health history. If n_t is a sufficient first-order Markov state, the immediately preceding deficit state should contain no additional information about n_{t+1} once current health is observed.

Four main findings emerge. First, the stochastic deficit-count architecture travels well to the HRS. A small number of parameters closely reproduce the distribution of subsequent health states, and the basic survivor transition relationship is stable across repeated HRS wave pairs. Extending the model over the richer HRS state space also reveals restrictions requiring modest refinement, particularly finite-support coherence, mortality at low measured deficit counts, and excess persistence at the current count.

Second, current deficit count strongly compresses the predictive information contained in standard observable characteristics for subsequent morbidity. On a common three-wave survivor sample, exact current deficit count and wave fixed effects explain 73.79 percent of individual variation in the next deficit count. Conditional on that baseline, chronological age has a partial R^2 of 0.0130, education 0.0036, race/ethnicity 0.0002, and sex 0.0001. Formal conditional-invariance tests frequently reject strict sufficiency, but the incremental predictive contribution of these characteristics is small. Chronological

age in particular contributes little additional information about subsequent morbidity once current deficit count is known.

Third, current health does not fully compress recent health history. Adding the immediately preceding deficit count raises R^2 from 0.7379 to 0.7568, corresponding to a partial R^2 of 0.0721. Recent history is therefore by far the most informative observed addition to the current state. Strict first-order sufficiency of n_t for survivor morbidity is rejected, even though current deficit count remains highly predictive. This distinction between strong compression and literal Markov sufficiency is central to our interpretation.

Fourth, mortality reveals a different boundary of the scalar representation. Recent deficit-count history adds little to mortality prediction conditional on current health, but chronological age, sex, race/ethnicity, and information about which particular deficits are present can retain predictive content. The scalar count is therefore more nearly sufficient for the evolution of morbidity among survivors than for the complete morbidity-and-mortality process. Predictive sufficiency is outcome-specific rather than an intrinsic property of the health measure alone.

Our study contributes to several strands of the economics literature on health dynamics. Dalgaard and Strulik (2014) develop a deterministic health-deficit model in which current deficits contribute to the production of future deficits. Early applications treat death as occurring when sufficiently many deficits have accumulated, while later applications allow mortality to be stochastic and deficit-dependent. Calibrated using the frailty index, this framework has been applied to a range of questions in health behavior and health policy (e.g. Dalgaard and Strulik, 2017; Dalgaard et al., 2021; Schüenemann et al., 2017, 2020; Grossmann et al., 2024). Grossmann and Strulik (2019) provide a relatively rare economic application of a Markov state-transition formulation. Our analysis differs by making the empirical sufficiency of the observed deficit state itself the object of investigation.

A broader literature models stochastic health dynamics using either latent health capital, building on Grossman (1972), or a small number of discrete observed health states. Examples of the former include Hugonnier et al. (2013), Margaritis and Wallenius (2024), and Ozkan (2025), while examples using relatively coarse observed health categories include French and Jones (2011), Braun et al. (2017), De Nardi et al. (2025), and Mahler and Yum (2024). In these approaches, future health generally depends on additional individual characteristics or latent components rather than on a single observed contemporaneous health measure alone. The deficit-count approach provides a different possibility: retain a scalar observed health state, but increase substantially the resolution with which that state represents current health.

Two recent papers are particularly closely related. Hosseini et al. (2022) construct a frailty index from accumulated health deficits and estimate a stochastic process for its evolution over the life cycle. Their approach allows individual heterogeneity together with persistent and transitory health components and is designed to reproduce the life-cycle variance and autocovariance structure of health. The relevant predictive state is consequently richer than the contemporaneous observed frailty index alone (see Hosseini et al., 2026, for an application). Our approach asks a complementary question: how far can the observed deficit count itself serve as the predictive state before latent or historical information must be added?

Borella et al. (2025), using the HRS, identify persistent health types from individual health deficit trajectories and show that these types predict subsequent health and mortality beyond standard observable characteristics. Their analysis also demonstrates the predictive value of health history, but uses rich health deficit information differently. Their subsequent transition model represents contemporaneous health using a small number of self-reported-health categories, while information from earlier deficit trajectories enters through estimated health types. We instead place the high-resolution deficit count directly in the contemporaneous state and ask whether age, observed characteristics, and recent health history continue to predict transitions once that richer current state is known. The approaches are therefore complementary ways of asking how rich health information should enter the state representation. Section 5 returns to this comparison and to the implications of our results for alternative representations of health dynamics.

The remainder of the paper proceeds as follows. Section 2 introduces the HRS deficit-count state, documents the empirical transition patterns, reproduces the stochastic transition model of Mitnitski et al. (2006), and examines its stability across HRS waves. Section 3 stress-tests the model over the extended HRS state space and develops the finite-support, mortality, and persistence refinements. Section 4 asks whether current deficit count is a sufficient predictive state by examining the incremental information contained in chronological age, demographic and socioeconomic characteristics, deficit composition, and recent health history. Section 5 discusses the interpretation of approximate predictive sufficiency and relates the results to alternative approaches to modeling health dynamics. Section 6 concludes.

2. THE BASIC STOCHASTIC HEALTH-STATE MODEL

2.1. Data and the deficit-count state. We study whether a rich description of current health can be represented by a simple state variable without losing most of the information needed to predict

subsequent health dynamics. The deficit-accumulation approach provides a particularly useful setting for this question. We observe 31 binary indicators of disease, functional limitation, and disability. These indicators form a multidimensional description of health, but the health state variable used in the stochastic model is simply their integer sum, $n_t = \sum_{j=1}^{31} d_{jt}$. Thus, health is represented at the individual level by a scalar state variable with 32 ordered values, $n_t \in \{0, \dots, 31\}$. The complete stochastic process additionally includes death, D , as an absorbing state, so that its state space consists of 32 living health states and one absorbing death state, $\mathcal{S} = \{0, \dots, 31, D\}$. Throughout the analysis we refer to n_t as the deficit count or current health state and distinguish it from the absorbing death state.

We constructed the 31-item deficit count to correspond as closely as possible to the formulation used by Mitnitski et al. (2006) and the general methodology outlined in Searle et al. (2008). Each deficit was coded as present or absent. Observations with missing information on any of the 31 components were excluded from the complete-case deficit count. Appendix Table A.1 reports the individual deficits and their summary statistics.

The primary analysis uses HRS waves 10 through 15, for which measurement of the 31 deficits is sufficiently stable. The resulting deficit distribution is substantially broader than in the original Canadian Study of Health and Aging (CSHA) application. The 75th, 90th, 95th, and 99th percentiles of the baseline deficit count are 12, 16, 19, and 24, respectively, and observed counts extend to 30. This broader support later allows us to examine restrictions of the original stochastic specification that were largely irrelevant in the CSHA, where fewer than 4% of participants had more than 11 deficits.

For each eligible baseline wave, we identify status at the subsequent HRS wave. Death before the next assessment is treated as an absorbing destination. Among respondents surviving to the subsequent wave, a deficit transition is observed when all 31 follow-up deficit indicators are available. Follow-up nonresponse and incomplete deficit measurement are distinguished from death rather than incorporated into the death state. We define a complete survivor transition as an adjacent-wave observation in which the respondent is alive at the subsequent wave and has a complete deficit count at both assessments. These observations identify the distribution of subsequent deficit states conditional on survival. Respondents known to survive to the next wave but lacking a complete follow-up deficit count contribute to the mortality analysis but not to the survivor-transition distribution.

2.2. Descriptive transition patterns. Before imposing a stochastic law, Table 1 describes how deficit counts change between adjacent HRS waves among respondents who survive and have complete measurements at both dates.

TABLE 1. Direction of change in deficit accumulation across adjacent HRS waves

Transition	<i>N</i>	Baseline deficit count, n_t			Follow-up Mean n_{t+1}	Direction of change		
		Mean	SD	Median		Improve	Same	Deteriorate
10–11	9,602	7.16	5.41	6	7.61	0.324	0.248	0.428
11–12	9,079	7.31	5.48	6	7.96	0.291	0.254	0.455
12–13	8,106	7.60	5.48	7	8.13	0.309	0.237	0.453
13–14	9,451	7.33	5.65	6	7.76	0.318	0.247	0.435
14–15	8,060	7.38	5.54	6	7.65	0.343	0.250	0.407

Notes: The table includes respondents who survived to the subsequent HRS wave and had complete 31-deficit measurements at both baseline and follow-up. Improvement is defined as $n_{t+1} < n_t$, no change as $n_{t+1} = n_t$, and deterioration as $n_{t+1} > n_t$. The three transition probabilities therefore sum to one within each wave pair. The analysis is descriptive and conditions on survival and observation of the follow-up deficit count.

Three features are immediate. First, the deficit count is not an irreversible accumulation process. Improvement is common: between 29 and 34 percent of observed survivor transitions end with fewer measured deficits than at baseline. Second, deterioration is nevertheless more frequent than improvement in every wave pair, accounting for roughly 41–46 percent of survivor transitions. Third, approximately one quarter of respondents remain at exactly the same deficit count. Mean deficit burden therefore rises between adjacent waves even though individual trajectories move in both directions.

These patterns make clear that the relevant object is not simply expected change in health. For every current state n , we require a distribution over possible future health states, together with the probability of death. The appeal of the deficit-count representation is that this potentially rich transition distribution is indexed by a single observed state variable.

2.3. The stochastic transition model. Mitnitski et al. (2006) proposed a parsimonious stochastic law for transitions between deficit-count states. Let P_{nk} denote the probability that a respondent with n deficits at baseline has k deficits at the subsequent assessment, and let P_{nD} denote the probability of death before that assessment. Conditional on survival, the original model assumes a Poisson distribution with state-dependent mean

$$\lambda_n = a_1 + b_1 n, \quad (1)$$

while mortality is described by

$$P_{nD} = \exp(a_2 + b_2 n). \quad (2)$$

The unconditional probability of transition to a living state is therefore

$$P_{nk} = \frac{(a_1 + b_1 n)^k}{k!} \exp[-(a_1 + b_1 n)] [1 - \exp(a_2 + b_2 n)]. \quad (3)$$

The model is economical in its state representation and in its parameterization. The individual health state is the scalar count n , and four parameters describe the complete transition row: the distribution over subsequent living deficit states and the absorbing probability of death. The model is stochastic rather than deterministic; individuals with the same current count may move to different future states even if n is the relevant state variable.

We first estimate the model jointly over baseline states $n = 0, \dots, 11$. This range is prespecified as the CSHA-comparable benchmark because the original study reported that fewer than 4% of participants had more than 11 deficits. Parameters are estimated from empirical state-transition probabilities using nonlinear least squares, paralleling the nonlinear fitting approach of Mitnitski et al. (2006).

2.4. Replication in the HRS. The HRS reproduces the central empirical feature of the original model: transition probabilities are strongly structured by the current number of accumulated deficits. In the CSHA-comparable range $n = 0, \dots, 11$, joint estimation gives

$$\hat{a}_1 = 0.468, \quad \hat{b}_1 = 0.999, \quad \hat{a}_2 = -5.064, \quad \hat{b}_2 = 0.207, \quad (4)$$


with $R^2 = 0.947$.

Figure 1 illustrates the survivor component. Each panel fixes the baseline deficit count and plots the empirical distribution of follow-up counts together with the fitted stochastic law. As baseline deficits rise, the distribution shifts progressively to the right, while retaining positive probability of improvement, no change, and deterioration.

Figure 2 shows the corresponding absorbing-state probability. Mortality is low at small n and rises steeply with the baseline deficit count. Together, the distribution over living destination states and the probability of death constitute the complete transition row.

Our estimated parameter values differ from those reported for the CSHA, as expected given differences in population, measurement, and the approximately biennial rather than five-year transition interval. More important for replication is that the same four-parameter stochastic architecture accounts for a large proportion of the variation in the HRS transition matrix. In the original combined CSHA analysis, the corresponding estimates were $a_1 = 1.771$, $b_1 = 0.823$, $a_2 = -2.153$, and $b_2 = 0.160$, with $R^2 = 0.979$.

FIGURE 1. Observed and fitted transition distributions by baseline deficit count




Dots: data; lines: fitted model prediction. Each panel fixes the baseline state n and shows the distribution of follow-up deficit counts among the living destinations.

The interpretation of this goodness-of-fit measure is important for what follows. The R^2 here describes fit to empirical state-to-state transition probabilities, not predictability of individual realized health outcomes. A stochastic transition law can reproduce the conditional distribution of future states closely even though considerable uncertainty remains about the particular state realized by an individual. Section 4 later asks a different question: whether information beyond current n_t improves prediction of individual transitions.

The replication therefore establishes the first part of the paper's argument. A multidimensional set of observed health conditions can be compressed into a scalar deficit count with a relatively fine state space, and a parsimonious stochastic law indexed by that count closely describes the resulting transition matrix.

2.5. Stability across HRS waves. A useful state representation should not depend on a single idiosyncratic transition interval. We therefore estimate the benchmark model separately for each

FIGURE 2. Observed and fitted mortality by baseline deficit count



Observed and fitted mortality by baseline deficit count in the benchmark state range. Death is treated as an absorbing destination state in the stochastic transition model.

adjacent HRS wave pair. Table 2 shows that the survivor deficit process is strikingly stable across the five transitions.

Estimates of b_1 range from 0.977 to 1.021, while estimates of a_1 range from 0.448 to 0.494. Interval-specific R^2 values lie between 0.932 and 0.947. Mortality parameters vary more across intervals than the survivor-transition parameters, but display no monotonic temporal pattern.

This stability strengthens the state-variable interpretation of the deficit count. The association between current deficit burden and the distribution of subsequent health states is not a feature of one particular wave pair. Across repeated HRS transitions, essentially the same stochastic architecture organizes the survivor transition distribution. The next question is therefore not whether the basic model works, but how far it can be pushed before its simplifying restrictions become empirically consequential.

3. ROBUSTNESS AND REFINEMENTS

3.1. Stress-testing the extended HRS state space. The broader support of the HRS deficit distribution provides a more demanding test of the original model than was possible in the CSHA. The original formulation contains two restrictions that are relatively innocuous when observed deficit counts remain far below their upper bound. First, the Poisson survivor distribution has infinite support even though no respondent can possess more than 31 of the 31 measured deficits. Mitnitski et al. (2006)

TABLE 2. Stability of the stochastic deficit and mortality processes across HRS waves

<i>Panel A: Deficit process</i>					
Transition	$\hat{a}_1$	$\hat{b}_1$	Survivor N	Joint R^2	Joint RMSE
Pooled	0.4678 (0.0205)	0.9995 (0.0101)	34,430	0.9473	0.01635
10–11	0.4608 (0.0206)	0.9906 (0.0101)	7,582	0.9468	0.01652
11–12	0.4480 (0.0229)	1.0207 (0.0117)	7,070	0.9320	0.01871
12–13	0.4872 (0.0228)	1.0137 (0.0112)	6,217	0.9377	0.01772
13–14	0.4943 (0.0211)	0.9967 (0.0102)	7,328	0.9461	0.01645
14–15	0.4478 (0.0214)	0.9768 (0.0106)	6,233	0.9417	0.01750
<i>Panel B: Mortality process</i>					
Transition	$\hat{a}_2$	$\hat{b}_2$	Mortality N		
Pooled	-5.0637 (0.6258)	0.2069 (0.0668)	46,962		
10–11	-4.9965 (0.7089)	0.1669 (0.0787)	10,166		
11–12	-5.1819 (0.7004)	0.2346 (0.0733)	9,322		
12–13	-4.6503 (0.4995)	0.1840 (0.0545)	8,625		
13–14	-5.6842 (0.8553)	0.2661 (0.0877)	10,177		
14–15	-5.0272 (0.6654)	0.2004 (0.0715)	8,672		

Notes: Each row reports estimates from the joint stochastic transition model over baseline deficit states $n = 0, \dots, 11$. The pooled specification combines HRS transitions 10–11 through 14–15. Panel A reports parameters of the conditional survivor deficit process, $E[n_{t+1} \mid n_t = n, \text{survive}] = a_1 + b_1 n$. Panel B reports parameters of the mortality process, $P_{nd} = \exp(a_2 + b_2 n)$. Standard errors are reported in parentheses. R^2 and RMSE refer to the complete joint model and are therefore reported only in Panel A. Survivor N is the number of complete survivor transitions; mortality N is the number of transitions with known mortality status. Each row estimates the four-parameter stochastic transition model separately for the indicated HRS wave pair. Joint R^2 and RMSE measure goodness of fit across the full set of modeled transition probabilities, including transitions among living deficit states and the absorbing death state, with the survivor-transition and mortality components estimated simultaneously.

explicitly recognized this limitation. Second, the exponential mortality probability is not globally bounded above by one.

We therefore re-estimate the original joint model over progressively broader baseline ranges, $n \leq 11$, $n \leq 18$, $n \leq 20$, and $n \leq 22$. The primary extension range is $n = 0, \dots, 20$, which retains the large majority of observed HRS transitions while avoiding reliance on the sparsest extreme states. The $n \leq 18$ and $n \leq 22$ specifications provide sensitivity checks.

Expanding the estimation range does not produce a collapse in conventional goodness of fit. The original specification retains R^2 values of approximately 0.94 as progressively higher baseline states are incorporated. The basic stochastic architecture therefore continues to describe observed HRS transitions well beyond $n = 11$.

What the broader range reveals instead is a problem of probabilistic coherence. Because the Poisson survivor kernel has infinite support, it increasingly assigns probability to follow-up counts above the feasible maximum of 31. At $n = 20$, this impossible probability mass remains below 1%, but it rises rapidly as the model is extrapolated toward the upper end of the state space. In addition, the original mortality probability can eventually exceed one. The total fitted probability assigned to feasible living states plus death therefore need not equal one at high deficit counts; Appendix Figure A.1 illustrates this behavior.

Table 3 separates goodness of fit from probabilistic coherence. The original model continues to fit observed transition probabilities well, but its probability laws are not globally valid over the finite deficit state space. A minimal repair truncates and renormalizes the Poisson survivor distribution over $k = 0, \dots, 31$ and writes mortality through a bounded hazard transformation. These changes leave empirical fit almost unchanged while ensuring that every transition row is a valid probability distribution.

The distinction is substantive for interpretation. The HRS does not show that the original stochastic architecture fails outside the CSHA range. Rather, richer upper-tail support exposes restrictions that were largely latent in the original application. This motivates refinements designed to preserve the model’s parsimony while making the transition law coherent over the full finite state space.

3.2. Refining the mortality process. We first isolate the mortality component. To guarantee a valid probability, mortality is written through an integrated hazard,

$$P_D(n) = 1 - \exp[-H_D(n)]. \quad (5)$$

Our preferred specification is

$$H_D(n) = h_0 + cn^\beta, \quad (6)$$

where $h_0 > 0$ captures mortality not explained by the measured deficit count and cn^β is the deficit-dependent component. This formulation permits positive mortality when $n = 0$, allows mortality to rise nonlinearly with deficit burden, and guarantees $0 < P_D(n) < 1$. The power specification is

TABLE 3. Joint estimation of the original and repaired stochastic transition models

Upper bound	Model	Parameters	R^2	RMSE	Weighted RMSE	Row sum at bound	Row sum at 31	$P_D(31)$
11	original	$a_1 = 0.468, b_1 = 0.999, a_2 = -5.064, b_2 = 0.207$	0.9473	0.01635	0.01773	1.0000	2.3834	3.8548
11	repaired	$a_1 = 0.468, b_1 = 0.999, \alpha = -5.075, \beta = 0.211$	0.9473	0.01635	0.01773	1.0000	1.0000	0.9862
18	original	$a_1 = 0.480, b_1 = 0.980, a_2 = -4.418, b_2 = 0.134$	0.9463	0.01494	0.01675	0.9983	0.8949	0.7629
18	repaired	$a_1 = 0.480, b_1 = 0.980, \alpha = -4.444, \beta = 0.139$	0.9463	0.01493	0.01674	1.0000	1.0000	0.5826
20	original	$a_1 = 0.483, b_1 = 0.976, a_2 = -4.362, b_2 = 0.128$	0.9442	0.01498	0.01675	0.9932	0.8614	0.6811
20	repaired	$a_1 = 0.484, b_1 = 0.977, \alpha = -4.397, \beta = 0.134$	0.9443	0.01497	0.01674	1.0000	1.0000	0.5466
22	original	$a_1 = 0.490, b_1 = 0.971, a_2 = -4.489, b_2 = 0.137$	0.9437	0.01492	0.01673	0.9811	0.9134	0.7955
22	repaired	$a_1 = 0.490, b_1 = 0.971, \alpha = -4.561, \beta = 0.146$	0.9438	0.01490	0.01672	1.0000	1.0000	0.6196

Notes: Each model is estimated jointly by unweighted nonlinear least squares on the same empirical transition cells: 32 feasible living destinations and one death destination for every included baseline state. The original model uses an ordinary Poisson survivor kernel and $P_D(n) = \exp(a_2 + b_2 n)$. The repaired model renormalizes the Poisson survivor kernel over $k = 0, \dots, 31$ and uses the bounded hazard probability $P_D(n) = 1 - \exp[-\exp(\alpha + \beta n)]$. The primary extension range is $n \leq 20$; $n \leq 11$ bridges to the frozen core replication and $n \leq 18, 22$ are range sensitivities. Weighted RMSE uses the number of mortality-known transitions in each origin state as a diagnostic only. Row sums refer to the probability assigned to the 32 feasible living states plus death; the repaired model sums to one by construction.

consistent with previous evidence of a strong power-law relationship between the frailty index and mortality rates (e.g. Mitnitski et al., 2002b; Dalgaard et al., 2022; Krenz and Strulik, 2026).

The term h_0 should not be interpreted as identifying a particular biological or cause-of-death mechanism. It is residual with respect to the measured deficit count and may capture acute events, infections, accidents, unmeasured disease, or other sources of mortality risk. Likewise, the deficit-dependent component is an empirical relationship between measured deficit burden and mortality rather than a cause-specific hazard.

We compare this specification with the original exponential mortality probability, a bounded exponential-hazard model, and a shifted power-law hazard. Table 4 shows that all four alternatives fit the complete transition matrix similarly. The advantage of the background-plus-power specification is therefore local rather than global. Under the original exponential-probability model, predicted mortality at $n = 0$ is 0.0128, compared with an observed probability of approximately 0.0055. The

background-plus-power model predicts 0.0080. Across low baseline deficit states, the background-plus-power specification reduces mortality prediction error, while global R^2 and RMSE change very little.

TABLE 4. Joint stochastic models with alternative mortality laws

Model	a_1	b_1	Mort. intercept	Mort. slope	$P_D(0)$	R^2	RMSE
Original exponential probability	0.4835	0.9765	-4.3618	0.1283	0.0128	0.9442	0.01498
Bounded exponential hazard	0.4836	0.9766	-4.3971	0.1343	0.0122	0.9443	0.01497
Fixed shifted power-law hazard	0.4900	0.9760	-7.4158	1.8462	0.0006	0.9444	0.01496
Background + deficit power law	0.4863	0.9764	-7.6864	1.9549	0.0080	0.9445	0.01496

Notes: All models are jointly estimated by unweighted nonlinear least squares on the same transition cells for baseline $n = 0, \dots, 20$. The primary new mortality model is $H(n) = \exp(\gamma) + \exp(\alpha)n^\beta$, with $P_D(n) = 1 - \exp[-H(n)]$. Its first hazard component represents mortality not captured by measured deficit count and is present even at $n = 0$; its second is a deficit-dependent power law. Living transitions use the finite-support Poisson kernel in the three bounded-hazard extensions.

The fitted deficit-dependent hazard rises nonlinearly with deficit count, with an estimated exponent of approximately $\beta = 1.95$. The estimated mortality probability not captured by measured deficits at $n = 0$ is approximately 0.0080 over one HRS transition interval. The deficit-dependent component is initially small, becomes approximately equal to the residual component near four deficits, and dominates at higher deficit counts. Appendix Figure A.2 shows that the competing mortality laws differ most clearly at the bottom of the deficit distribution.

The refinement should therefore not be interpreted as a search for a materially higher global R^2 . Its advantages are that it represents mortality at low deficit counts more accurately, distinguishes mortality not captured by measured deficits from deficit-dependent mortality, and remains a valid probability specification throughout the state space.

3.3. Persistence in survivor deficit dynamics. The mortality refinement leaves the conditional distribution of follow-up deficit counts among survivors largely unchanged. The principal remaining discrepancy is instead in the probability of remaining at exactly the same aggregate deficit count. For several low baseline states, the observed diagonal probability is systematically greater than predicted by the finite-support Poisson kernel.

We therefore introduce a single count-state persistence parameter. Let $q_P(k \mid n, S)$ denote the Poisson survivor distribution truncated and renormalized over the feasible states $k = 0, \dots, 31$. For $n > 0$, the refined survivor distribution is

$$q(k \mid n, S) = \pi I(k = n) + (1 - \pi)q_P(k \mid n, S), \quad (7)$$

where $I(k = n)$ equals one if the follow-up deficit count equals the baseline count. At $n = 0$, the persistence component is constrained to zero. The corresponding unconditional transition probability to a living state is $P_{nk} = [1 - P_D(n)]q(k | n, S)$, with death remaining the absorbing state.

The parameter π represents persistence of the aggregate deficit count, not persistence of the same individual health deficits. An individual can remain at the same count while some deficits resolve and others arise. Adding this single parameter produces a larger improvement than the changes in mortality specification. The estimated persistence probability is $\hat{\pi} = 0.0606$. Global R^2 rises from 0.9445 under the preferred mortality model with a finite-support Poisson survivor kernel to 0.9535 after adding count-state persistence. RMSE declines from 0.01496 to 0.01409.

Diagnostics focused specifically on the survivor distribution show the same pattern. Averaged over baseline states $n = 1, \dots, 5$, conditional-survivor total-variation distance declines from 0.1565 to 0.1260. Mean diagonal error – observed minus fitted probability of an unchanged deficit count – falls from 0.1048 to 0.0589, while off-diagonal L_1 error declines from 0.2083 to 0.1930. The improvement therefore does not arise solely from mechanically shifting probability to the diagonal.

Appendix Figure A.3 illustrates the result directly. For baseline counts between one and five deficits, the simple Poisson kernel substantially underpredicts the probability of observing the same count at follow-up. The persistence mixture moves fitted probabilities toward the observed values while also reducing overall survivor-distribution error. At zero deficits, where the persistence component is constrained to zero, the two specifications remain similar.

3.4. The preferred extended stochastic model. Taken together, the evidence supports a parsimonious extension rather than a replacement of the original stochastic framework. The preferred specification retains the scalar deficit count as the observed health state and death as an absorbing destination, but incorporates three refinements: finite support of the survivor state space, mortality not captured by measured deficit burden alongside deficit-dependent mortality, and modest persistence of the current aggregate deficit count.

The mortality process is

$$H_D(n) = h_0 + cn^\beta, \quad P_D(n) = 1 - \exp[-H_D(n)], \quad (8)$$

and, for $n > 0$, the survivor process is

$$q(k | n, S) = \pi I(k = n) + (1 - \pi) \frac{\Pr\{\text{Poisson}(\lambda_n) = k\}}{\Pr\{\text{Poisson}(\lambda_n) \leq 31\}}, \quad (9)$$

where $\lambda_n = a_1 + b_1 n$. The complete transition law is therefore

$$P_{nk} = [1 - P_D(n)]q(k | n, S), \quad (10)$$

together with the absorbing death probability $P_D(n)$.

The refinements improve the probabilistic coherence and local empirical performance of the model without altering its central attraction. A multidimensional vector of 31 measured health indicators is represented by a single integer state, yet that scalar retains a high-resolution state space and organizes a rich distribution of subsequent outcomes through a small number of parameters. The transition model therefore provides a useful benchmark for the more demanding question addressed next: whether current deficit count is merely highly predictive, or whether it is sufficiently informative that age, individual characteristics, and recent health history contribute little once the current state is known.

4. IS CURRENT HEALTH A SUFFICIENT PREDICTIVE STATE?

The preceding sections establish that a scalar deficit count with a high-resolution state space organizes the distribution of subsequent health outcomes remarkably well. Parsimony, however, is not the same as sufficiency. A stronger question is whether information outside the current deficit state improves prediction of the next health state.

We define sufficiency in predictive rather than causal terms. For an observed information set Z_t , current deficit count is sufficient with respect to Z_t if

$$P(n_{t+1} | n_t, Z_t, w, S_{t+1} = 1) = P(n_{t+1} | n_t, w, S_{t+1} = 1),$$

where w denotes survey wave. The comparison is deliberately asymmetric. Current deficit count may already embody information associated with chronological age, persistent individual characteristics, accumulated exposures, or previous health. We do not attempt to separate these sources of shared predictive information. Instead, we ask whether Z_t contributes additional predictive information once current health is known.

For survivor morbidity, our common empirical baseline therefore contains exact current deficit-count indicators and survey-wave fixed effects. We then add candidate information sets separately and measure the resulting increase in R^2 and partial R^2 . This approach distinguishes strict predictive sufficiency from approximate predictive sufficiency. In a large sample, even small departures from conditional independence may be statistically detectable. The economically more relevant question is how much predictive information remains outside the current state.

4.1. Does Chronological Age Add Predictive Information? Chronological age provides the natural first test and is closest to the original motivation of Mitnitski et al. (2006), who asked how age influences transition probabilities relative to accumulated deficits. Age and deficit burden are strongly related: mean and median age generally rise with n , as shown in Appendix Figure A.4. The relevant state-sufficiency question, however, is whether age continues to predict subsequent health after the current deficit count has been observed.

For survivor morbidity, remarkably little age information remains. In the common three-wave sample reported below in Table 5, a model containing exact current deficit-count indicators and survey-wave fixed effects explains 73.79 percent of individual-level variation in the next deficit count. Adding a quadratic in chronological age raises R^2 only from 0.7379 to 0.7413. The corresponding partial R^2 is 0.0130: age explains 1.30 percent of the residual variation left after current deficit count and wave are known.

Formal tests nevertheless reject literal invariance of morbidity transitions across age groups conditional on n (Appendix Table A.2). Thus, chronological age is not irrelevant. Rather, it contributes little incremental predictive information about subsequent deficit accumulation once current deficit count is known. This distinction between statistical rejection and predictive magnitude is central to our interpretation of approximate state sufficiency.

Mortality provides a clear qualification. Appendix Figure A.5 shows that, conditional on the same measured deficit count, mortality rises substantially with chronological age. The detailed Brier-score and deviance comparisons reported in Appendix Table A.3 lead to the same conclusion: current deficit burden compresses age-related information much more effectively for subsequent morbidity than for mortality. The Brier score is the mean squared error of predicted death probabilities. We report the proportional reduction in the Brier score when the indicated characteristic – here chronological age – is added to the basic model. Deviance gain is the analogous proportional reduction in binomial deviance.

These conditional age comparisons should not be interpreted causally. An 80-year-old observed at $n = 0$ has survived to an advanced age while avoiding all 31 measured deficits; a 50-year-old with the same count has not undergone the same mortality and morbidity selection. Equal current deficit counts can therefore conceal differences in latent robustness, accumulated exposures, cohort, and prior health. The finding is about predictive information: chronological age adds little to short-run morbidity prediction conditional on current n , but substantially more to mortality prediction. The oldest-old are also more sparsely represented at some deficit states; Appendix Tables A.4 and A.5 document this support and show that the main morbidity conclusion is not driven by respondents aged 85 and older.

TABLE 5. Incremental predictive information beyond current health

Predictive information	Model R^2	ΔR^2	Partial R^2	N
Current deficit count (baseline)	0.7379	–	–	29,859
+ Age	0.7413	0.0034	0.0130	29,859
+ Education	0.7389	0.0009	0.0036	29,859
+ Race/ethnicity	0.7380	0.0001	0.0002	29,859
+ Sex	0.7380	0.0000	0.0001	29,859
+ Recent health history	0.7568	0.0189	0.0721	29,859

Notes: All rows use the same three-wave survivor sample and the same outcome, next deficit count. The baseline model contains exact current deficit-count indicators and survey-wave fixed effects. Each subsequent row adds the indicated information to this same baseline. ΔR^2 is the increase in ordinary R^2 relative to the baseline. Partial R^2 is the proportional reduction in residual sum of squares from adding the indicated information to the baseline. Current and lagged deficit counts enter as exact state indicators; age enters quadratically; education enters as years of schooling; sex and race/ethnicity enter categorically. Recent health history is the deficit count at the immediately preceding HRS wave. The comparison is deliberately asymmetric: it asks whether each information set adds predictive content once current health is known, allowing current deficit count to retain any information it already embodies about age, persistent heterogeneity, or prior health. The R^2 values in this table measure individual-level variation in realized next-period deficit counts and are not directly comparable to the goodness-of-fit R^2 for state-to-state transition probabilities reported in Section 2.

4.2. Do Demographic and Socioeconomic Characteristics Add Predictive Information?

We next ask whether sex, education, and race/ethnicity contain predictive information not already represented by current health. Diagnostic specifications condition on exact current deficit count, chronological age, and survey wave, with standard errors clustered by respondent. Education is measured in years in the primary specification, with educational categories used as a nonlinear robustness check. Race and ethnicity are represented by mutually exclusive categories for non-Hispanic White, non-Hispanic Black, Hispanic, and non-Hispanic Other respondents. These characteristics are treated as predictors of heterogeneous health dynamics rather than as treatments. In particular, race and ethnicity index social stratification and accumulated exposures rather than biological types.

Formal tests reject strict sufficiency with respect to most of these characteristics (Appendix Table A.6). Race/ethnicity predicts mortality and each survivor-transition outcome conditional on current n , age, and wave. Sex predicts mortality and most survivor outcomes, while education predicts the survivor outcomes but not mortality in its separate specification. With large samples, however, these tests say little about the magnitude of the information omitted from current health.

The common-sample comparison in Table 5 shows that this omitted information is extremely small for subsequent morbidity. Starting from the same baseline R^2 of 0.7379, education raises R^2 by 0.0009, race/ethnicity by 0.0001, and sex by less than 0.0001. Their partial R^2 values are 0.0036, 0.0002, and 0.0001, respectively. Current deficit count therefore compresses almost all of the morbidity-relevant predictive information contained in these observed individual characteristics.

Mortality is again less compressible. Education contributes essentially no additional mortality prediction once current deficit count, age, and wave are known, while race/ethnicity contributes a small amount. Sex is more important: adding sex reduces the Brier score by about 1.1 percent relative to a model containing current deficit count, age, and wave. The reciprocal predictive comparisons and likelihood-based diagnostics are reported in Appendix Table A.7.

The detailed mortality decompositions provide further evidence that a scalar deficit count does not exhaust mortality-relevant information. Using the refined hazard specification, Appendix Table A.8 allows sex, education, and race/ethnicity to shift mortality not captured by measured deficits, deficit-dependent mortality, or both. The strongest heterogeneity appears for sex, where most of the improvement comes through the deficit-dependent component, although allowing both components to vary fits best. Race/ethnicity also exhibits heterogeneity in both components. Education is much weaker and criterion-sensitive. Corresponding fitted mortality probabilities are reported in Appendix Table A.9.

One reason demographic characteristics retain predictive information for mortality is that the scalar count discards information about deficit composition. Two individuals with $n = 8$ need not possess the same eight health deficits, and particular deficits may carry different mortality risks. Appendix Table A.10 shows that replacing the scalar count with indicators for the 31 individual deficits reduces the incremental Brier-score contribution of sex by 38.5 percent and that of race/ethnicity by 52.5 percent. The individual deficit indicators themselves reduce the mortality Brier score by approximately 3 percent relative to the scalar count, conditional on age and wave.

This diagnostic illustrates both the strength and the cost of the deficit-count representation. The underlying health information is multidimensional, but the individual state variable is a single integer. That scalar has a high-resolution state space – 32 possible values with 31 measured deficits – while deliberately discarding information about which deficits constitute a given count. For subsequent morbidity, little predictive information in standard demographic and socioeconomic characteristics survives this compression. For mortality, deficit composition and selected individual characteristics remain more consequential.

4.3. Does Recent Health History Add Predictive Information? Observed characteristics provide useful tests of the state representation, but recent health history provides the more direct test of its Markov interpretation. If current deficit count is a sufficient first-order state, the path by which an individual arrived at n_t should not improve prediction of the next state. For survivor morbidity,

this restriction implies

$$P(n_{t+1} \mid n_t, n_{t-1}, w, S_{t+1} = 1) = P(n_{t+1} \mid n_t, w, S_{t+1} = 1).$$

This restriction does not require health transitions to be deterministic. It requires the conditional distribution of the next state to be independent of the preceding state once the current state is observed.

We test this restriction using the deficit count at the immediately preceding HRS wave, n_{t-1} . The exercise requires three consecutive observations and restricts both n_{t-1} and n_t to the supported range from zero to 20. The baseline model contains exact current deficit-count indicators and survey-wave fixed effects; the augmented model adds exact indicators for n_{t-1} . Standard errors for the corresponding diagnostic tests are clustered by respondent. This is deliberately a local test of first-order state sufficiency: one lag can establish that current health does not fully summarize recent history, but cannot establish whether it summarizes the individual’s complete health trajectory.

Recent history contains clear additional information about survivor morbidity. Conditional on current deficit count and wave, lagged deficit count jointly predicts the subsequent deficit count among survivors ($p < 0.0001$). It also predicts whether the respondent improves or deteriorates ($p < 0.0001$ in both cases), although it does not significantly predict remaining at exactly the same deficit count ($p = 0.165$). These results are reported in Appendix Table A.11 and are essentially unchanged when a quadratic in chronological age is added to the conditioning set. Individuals who arrive at the same current deficit count through different recent health trajectories therefore face systematically different subsequent morbidity transitions.

The magnitude of this departure is more revealing than the statistical rejection. Table 5 puts recent health history and the observable characteristics on exactly the same footing. On the common three-wave survivor sample, the baseline model containing current deficit count and wave has $R^2 = 0.7379$. Adding n_{t-1} raises it to 0.7568, an increase of 0.0189. The corresponding partial R^2 is 0.0721: recent health history explains 7.21 percent of the residual variation left by current deficit count and wave.

This is substantially larger than the incremental contribution of any other observed information set. Age has a partial R^2 of 0.0130, education 0.0036, race/ethnicity 0.0002, and sex 0.0001. Recent health history is therefore by far the most informative observed addition to the current deficit state.

The comparison is intentionally asymmetric. The state-sufficiency question is not how predictive information shared by current deficit count and another variable should be attributed between them. A useful state variable is valuable precisely because it summarizes information inherited from the past. Current deficit count may therefore embody information about previous health, cumulative influences

associated with chronological age, and persistent individual characteristics. The relevant question is whether those variables contain additional predictive information once the current state is observed.

This comparison also clarifies the Markov interpretation. The fact that n_{t-1} predicts n_t , or that age predicts deficit accumulation, is not evidence against a Markov representation. A state variable can summarize the past precisely because past health and other persistent influences are reflected in the current state. The stronger Markov restriction is that these variables should cease to improve prediction once n_t is observed. Our results show that this compression is substantial but incomplete.

For mortality, recent deficit-count history appears much less important. Conditional on current deficit count and wave, lagged deficit count is not jointly significant in predicting death ($p = 0.305$); adding a quadratic in age produces the same conclusion ($p = 0.206$). This does not imply that current deficit count is sufficient for mortality. As shown above, chronological age, selected demographic characteristics, and deficit composition retain predictive information beyond the current count. Rather, the information omitted from n_t that matters for mortality appears not to be captured primarily by one additional lag of the aggregate deficit count.

5. DISCUSSION

The central question of this paper is whether a multidimensional description of health can be compressed into a scalar state without sacrificing most of the information needed to predict its subsequent evolution. The deficit-accumulation approach provides an unusually direct way to investigate this question. It begins with 31 indicators of disease, functional limitation, and disability, but represents an individual’s current health by a single integer, n_t , with 32 ordered values. Our results suggest that this combination of scalar representation and high-resolution measurement is empirically useful, but not literally sufficient.

For morbidity among survivors, current deficit count is an approximately, but not strictly, sufficient predictive state. Once n_t is known, chronological age adds relatively little to prediction of the next deficit count, while education, race/ethnicity, and sex add still less. Recent health history provides the important exception: the immediately preceding deficit count retains a quantitatively meaningful association with subsequent morbidity conditional on current health. Mortality marks a different boundary. Recent deficit-count history adds little once n_t is observed, but chronological age, selected individual characteristics, and the composition of accumulated deficits retain predictive information beyond the scalar count. Sufficiency therefore depends on the outcome being modeled and the information that the state is expected to summarize.

This interpretation rests on an asymmetric notion of predictive sufficiency. We do not attempt to attribute information shared by current deficit count and another variable to one source or the other. A useful state variable is valuable precisely because it can embody information generated by earlier circumstances. Chronological age, socioeconomic conditions, previous health shocks, and persistent individual characteristics may all have contributed to current deficit burden. Their influence does not have to disappear simply because they are absent from a transition equation: to the extent that their consequences are reflected in n_t , those consequences propagate through the current state. The relevant empirical question is whether these variables contain additional predictive information once n_t is observed.

Chronological age illustrates both the usefulness and the limits of this interpretation. Deficit accumulation is strongly age-related, but age contributes little incremental predictive information for subsequent morbidity conditional on current deficit count. This does not imply that chronological aging is unimportant for health dynamics. Reaching age 80 with few measured deficits reflects a different history of morbidity and mortality selection from reaching age 50 with the same count. Equal n_t can therefore conceal differences in latent robustness, previous exposures, cohort, and unmeasured health. The narrower empirical result is nevertheless useful for economic modeling: much of the age-related information relevant for near-term morbidity is already represented by current deficit burden.

Recent health history provides a more demanding test. If n_t were a strictly sufficient first-order Markov state, previous deficit states would cease to improve prediction once the current state was observed. They do not. Among survivors, n_{t-1} predicts the subsequent deficit count, improvement, and deterioration conditional on n_t , and its incremental predictive contribution is substantially larger than that of age, sex, education, or race/ethnicity. Two individuals occupying the same current state but arriving there along different recent trajectories therefore need not face the same subsequent transition distribution.

The history result does not contradict the strong fit of the stochastic transition model. A Markov process is stochastic rather than deterministic, so even under exact state sufficiency individuals occupying the same state can realize different subsequent outcomes. Moreover, the two sets of R^2 statistics in the paper measure different objects. The goodness-of-fit measure for the stochastic model describes how closely the parametric transition law reproduces cell-level state-to-state probabilities. The state-sufficiency analysis asks how much individual variation in realized next-period deficit counts is predicted by current health and how much additional information is contained in other observables. A

parsimonious transition law can therefore reproduce the aggregate transition matrix closely even when individual transitions remain uncertain and recent history contains additional predictive information.

The distinction between a scalar state and a coarse state is central to the economic interpretation. Low dimensionality does not imply crude measurement. A model with three categories of self-reported health and one with a 31-deficit count both use a single contemporaneous health variable, but the latter provides 32 ordered health states. The deficit approach can therefore preserve considerably more observed health heterogeneity without adding dimensions to the individual health state. Its compression occurs elsewhere: the count discards information about deficit composition, recent trajectory, and latent heterogeneity.

This perspective helps locate our approach relative to recent models of health dynamics. Hosseini et al. (2022) also exploit the informational richness of the frailty index, but model its dynamics using latent individual heterogeneity together with persistent and transitory health components. Their relevant predictive state is therefore richer than the contemporaneous observed frailty index. Our approach asks a complementary question: how much state augmentation is empirically necessary if the observed deficit measure itself is used as the contemporaneous state? The finding that n_{t-1} retains predictive information conditional on n_t is consistent with their richer representation: the scalar observed state captures substantial persistence but does not reveal all persistent heterogeneity or health dynamics.

Borella et al. (2025) use rich health information differently. They extract persistent health types from health deficit trajectories and show that these types predict subsequent health and mortality, while their transition analysis uses a small number of self-reported-health categories as the contemporaneous health state. We instead place the high-resolution deficit count directly in the contemporaneous state and ask whether further information remains once that richer state is observed. Their longer trajectories are designed to recover persistent heterogeneity, whereas our one-lag exercise provides a deliberately local test of history dependence. Finding that recent history remains informative even conditional on a high-resolution current state reinforces the broader conclusion that contemporaneous health need not fully reveal persistent differences in health dynamics. It does not establish that adding one lag is the preferred augmentation: latent persistent components, trajectory-based types, deficit duration, or other summaries of health history may provide more efficient state representations.

Mortality reveals another cost of scalar compression. Individuals with the same deficit count need not have the same deficits. The mapping from a set of deficits to the deficit count deliberately discards information about deficit composition, and our mortality diagnostics show that some of this information matters. Replacing the scalar count with indicators for individual deficits improves mortality prediction

and reduces the residual predictive contribution of sex and race/ethnicity. Equal counts can also conceal differences in deficit severity and duration, comorbidity, treatment, environmental exposure, and unmeasured health. This helps explain why the scalar state is less sufficient for mortality than for subsequent deficit accumulation. It also illustrates why there is no outcome-independent ranking of candidate state representations: recent trajectory is particularly informative for morbidity, whereas age and deficit composition matter more for mortality.

The deficit-count approach consequently occupies a useful middle ground in the representation of health dynamics. Conventional categorical measures obtain a scalar state by dividing health into only a few categories. Rich latent-state models preserve additional persistence and heterogeneity but require a larger state representation. The deficit approach instead compresses multidimensional observed health information into a scalar state with a high-resolution state space. This is potentially valuable in quantitative economic models in which health is only one component of the state vector, alongside variables such as assets, employment, insurance, or family circumstances.

The practical implication is not that deficit count should always be the only health state, but that state-space augmentation should be disciplined by the outcome and horizon of the application. For near-term morbidity, current deficit count provides a strong scalar benchmark and standard demographic and socioeconomic characteristics add little predictive content. If persistence beyond the current state matters, recent history, trajectory-based heterogeneity, or latent persistent components may warrant inclusion. If survival is central, chronological age and richer information about the underlying deficits may be necessary. Approximate predictive sufficiency of current health is therefore best viewed as a modeling benchmark: begin with the scalar high-resolution state and enlarge the state space when the omitted information is predictively important for the outcome under study.

6. CONCLUSION

This paper revisits the stochastic deficit-accumulation model of Mitnitski et al. (2006) using the HRS and uses it to ask a broader question about the representation of health in dynamic economic models. Thirty-one indicators of disease, functional limitation, and disability can be compressed into a single deficit count with 32 ordered health states. Despite this substantial dimensionality reduction, a parsimonious stochastic law closely reproduces the distribution of subsequent health states over a much broader range of deficit counts than in the original application. Finite-support, persistence, and mortality refinements improve the coherence of this representation while preserving its scalar health state.

Current deficit count emerges as an approximately, but not strictly, sufficient predictive state. For morbidity among survivors, it compresses most of the predictive information contained in chronological age and observed demographic and socioeconomic characteristics, but does not fully compress recent health history. For mortality, recent deficit-count history adds little, while chronological age, selected individual characteristics, and deficit composition retain predictive information beyond the scalar count. The deficit approach therefore provides a high-resolution scalar benchmark for modeling health dynamics, while the systematic departures from sufficiency identify when the state should be augmented for the outcome and horizon under study.

REFERENCES

- Barabasi, A. L., Gulbahce, N., and Loscalzo, J. (2011). Network medicine: a network-based approach to human disease. *Nature Reviews Genetics* 12(1), 56-68.
- Belikov, A. V. (2019). Age-related diseases as vicious cycles. *Ageing Research Reviews* 49, 11-26.
- Borella, M., Bullano, F., De Nardi, M., Krueger, B., and Manresa, E. (2025). Health inequality and health types. *The Econometrics Journal* 28(3), 341-384.
- Braun, R. A., Kopecky, K. A., and Koreshkova, T. (2017). Old, sick, alone, and poor: A welfare analysis of old-age social insurance programmes. *Review of Economic Studies* 84(2), 580-612.
- Dalgaard, C. J., and Strulik, H. (2014). Optimal aging and death: understanding the Preston curve. *Journal of the European Economic Association* 12(3): 672-701.
- Dalgaard, C. J., and Strulik, H. (2017). The genesis of the golden age: Accounting for the rise in health and leisure. *Review of Economic Dynamics* 24, 132-151.
- Dalgaard, C. J., Hansen, C. W., and Strulik, H. (2021). Fetal origins – A life cycle model of health and aging from conception to death. *Health Economics* 30(6), 1276-1290.
- Dalgaard, C. J., Hansen, C. W., and Strulik, H. (2022). Physiological aging around the world. *PloS one*, 17(6), e0268276.
- De Nardi, M., Pashchenko, S., and Porapakarm, P. (2025). The lifetime costs of bad health. *Review of Economic Studies* 92(3), 1987-2026.
- French, E., and Jones, J. B. (2011). The effects of health insurance and self-insurance on retirement behavior. *Econometrica* 79(3), 693-732.
- Goggins, W. B., Woo, J., Sham, A., and Ho, S. C. (2005). Frailty index as a measure of biological age in a Chinese population. *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences* 60(8), 1046-1051.
- Grossman, M. (1972). On the concept of health capital and the demand for health. *Journal of Political Economy* 80, 223-255.
- Grossmann, V., and Strulik, H. (2019). Optimal social insurance and health inequality. *German Economic Review* 20(4), e913-e948.
- Grossmann, V., Schünemann, J., and Strulik, H. (2024). Fair pension policies with occupation-specific ageing. *Economic Journal* 134(663), 2835-2875.
- Hosseini, R., Kopecky, K. A., and Zhao, K. (2022). The evolution of health over the life cycle. *Review of Economic Dynamics* 45, 237-263.
- Hosseini, R., Kopecky, K., and Zhao, K. (2026). How important is health inequality for lifetime earnings inequality? *Review of Economic Studies* 93(1), 556-599.
- Howlett, S. E., Rutenberg, A. D., and Rockwood, K. (2021). The degree of frailty as a translational measure of health in aging. *Nature Aging* 1(8), 651-665.

- Hugonnier, J., Pelgrin, F., and St-Amour, P. (2013). Health and (other) asset holdings. *Review of Economic Studies* 80(2), 663-710.
- Krenz, A., and Strulik, H. (2026). MICHELE: The Computation of Healthy Life Expectancy from Mortality Schedules. Cesifo Discussion Paper.
- Mahler, L. and Yum, M. (2024). Lifestyle Behaviors and Wealth-Health Gaps in Germany, *Econometrica* 92(5), 1697-1733.
- Margaris, P. and Wallenius, J. (2024). Can Wealth Buy Health? A Model of Pecuniary and Non-Pecuniary Investments in Health, *Journal of the European Economic Association* 22(3), 1097–1138.
- Mitnitski, A.B., Mogilner, A.J., and Rockwood, K. (2001). Accumulation of deficits as a proxy measure of aging. *Scientific World* 1, 323-336.
- Mitnitski, A. B., Graham, J. E., Mogilner, A. J., and Rockwood, K. (2002a). Frailty, fitness and late-life mortality in relation to chronological and biological age. *BMC Geriatrics* 2(1), 1-8.
- Mitnitski, A. B., Mogilner, A. J., MacKnight, C., and Rockwood, K. (2002b). The accumulation of deficits with age and possible invariants of aging. *The Scientific World Journal* 2, 1816-1822.
- Mitnitski, A., Song, X., Skoog, I., Broe, G. A., Cox, J. L., Grunfeld, E., and Rockwood, K. (2005). Relative fitness and frailty of elderly men and women in developed countries and their relationship with mortality. *Journal of the American Geriatrics Society* 53(12), 2184-2189.
- Mitnitski, A., Bao, L., and Rockwood, K. (2006). Going from bad to worse: a stochastic model of transitions in deficit accumulation, in relation to mortality. *Mechanisms of Ageing and Development* 127(5), 490-493.
- Ozkan, S. (2025). Income differences and health disparities: Roles of preventive vs. curative medicine, *Journal of Monetary Economics* 150, 103698.
- Rockwood, K., and Mitnitski, A. (2007). Frailty in relation to the accumulation of deficits. *Journals of Gerontology Series A: Biological Sciences and Medical Sciences* 62(7), 722-727.
- Rutenberg, A. D., Mitnitski, A. B., Farrell, S. G., and Rockwood, K. (2018). Unifying aging and frailty through complex dynamical networks. *Experimental Gerontology* 107, 126-129.
- Schünemann, J., Strulik, H., and Trimborn, T. (2017). The gender gap in mortality: How much is explained by behavior? *Journal of Health Economics* 54, 79-90.
- Schünemann, J., Strulik, H., and Trimborn, T. (2020). The marriage gap: Optimal aging and death in partnerships. *Review of Economic Dynamics* 36, 158-176.
- Searle, S. D., Mitnitski, A., Gahbauer, E. A., Gill, T. M., and Rockwood, K. (2008). A standard procedure for creating a frailty index. *BMC Geriatrics*, 8(1), 24.
- Theou, O., Brothers, T. D., Mitnitski, A., and Rockwood, K. (2013). Operationalization of frailty using eight commonly used scales and comparison of their ability to predict all-cause mortality. *Journal of the American Geriatrics Society* 61(9), 1537-1551.

Tian, Y. E., Cropley, V., Maier, A. B., Lautenschlager, N. T., Breakspear, M., and Zalesky, A. (2023). Heterogeneous aging across multiple organ systems and prediction of chronic disease and mortality. *Nature Medicine* 29(5), 1221-1231.

APPENDIX

TABLE A.1. Health deficits used to construct the deficit count

No.	Health deficit	Mean	SD	N
1	Reports high BP this wv	0.615	0.487	73,139
2	Reports diabetes this wv	0.257	0.437	73,139
3	Reports cancer this wv	0.148	0.355	73,139
4	Reports lung disease this wv	0.105	0.307	73,139
5	Reports heart prob this wv	0.239	0.426	73,139
6	Reports stroke this wv	0.074	0.262	73,139
7	Reports arthritis this wv	0.567	0.496	73,139
8	Health problems limit work	0.372	0.483	73,139
9	Any Diff-Walk sev blocks	0.320	0.467	73,139
10	Any Diff-Jog one mile	0.687	0.464	73,139
11	Any Diff-Walk across room	0.072	0.259	73,139
12	Any Diff-Sit for 2 hours	0.213	0.409	73,139
13	Any Diff-Get up fr chair	0.387	0.487	73,139
14	Any Diff-Clmb sev flt str	0.471	0.499	73,139
15	Any Diff-Clmb 1 flt stair	0.198	0.399	73,139
16	Any Diff-Stoop/Kneel/Crch	0.468	0.499	73,139
17	Any Diff-Lift/carry 10lbs	0.259	0.438	73,139
18	Any Diff-Pick up a dime	0.073	0.261	73,139
19	Any Diff-Rch/xtnd arms up	0.182	0.386	73,139
20	Any Diff-Push/pull lg obj	0.283	0.450	73,139
21	Any Diff-Dressing	0.109	0.312	73,139
22	Any Diff-Bathing, shower	0.072	0.259	73,139
23	Any Diff-Eating	0.035	0.185	73,139
24	Any Diff-Get in/out bed	0.074	0.261	73,139
25	Any Diff-Use a map	0.146	0.353	73,139
26	Any Diff-Managing money	0.072	0.259	73,139
27	Any Diff-Use telephone	0.045	0.206	73,139
28	Any Diff-Take medications	0.040	0.195	73,139
29	Reports psych prob this wv	0.191	0.393	73,139
30	Any hospitalization	0.248	0.432	73,139
31	Regular use of prescription drugs	0.814	0.389	73,139

Notes: The table reports the 31 binary health deficits used to construct the deficit count. Each deficit is coded 0/1, with one indicating the presence of the deficit. Means therefore equal the proportion of person-wave observations in which the deficit is present. Statistics are calculated among complete 31-item measurements in HRS waves 10–15. Because respondents may contribute observations in more than one wave, N refers to person-wave observations rather than unique individuals.

TABLE A.2. Does chronological age predict transitions conditional on current deficit state?

Outcome	Age 50+ p -value	Age 50–84 p -value
Death	0.00000	0.00000
Next deficit count among survivors	0.00000	0.00000
Improvement among survivors	0.00000	0.00000
Unchanged count among survivors	0.00145	0.00141
Deterioration among survivors	0.00000	0.00000

Notes: All models condition on exact baseline deficit-count fixed effects and survey-wave fixed effects for $n \leq 20$. Age 50+ uses groups 50–64, 65–74, 75–84, and 85+; the age 50–84 specification omits the oldest-old and tests the remaining three groups. Standard errors are clustered by respondent. Improvement, unchanged count, and deterioration are estimated as separate linear-probability diagnostics.

TABLE A.3. Relative predictive importance of chronological age and current deficit state

Outcome and metric	Sample	Age contribution	n contribution	Age/ n
Next deficit count: partial R^2	Age 50+	0.0131	0.7186	0.018
	Age 50–84	0.0070	0.7247	0.010
Death: Brier-score gain	Age 50+	0.0433	0.0319	1.360
	Age 50–84	0.0176	0.0263	0.669
Death: deviance reduction	Age 50+	0.0655	0.0628	–
	Age 50–84	0.0359	0.0650	–

Notes: For the survivor next-count outcome, the age contribution is the partial R^2 from adding age-group indicators to a model containing exact n and wave fixed effects; the n contribution reverses this comparison by adding exact n indicators to age group and wave. For mortality, the primary magnitude measure is the proportional reduction in Brier score from the analogous nested logit models. The deviance-reduction rows provide a likelihood-based sensitivity metric. The Age/ n column is the ratio of the two incremental measures within an outcome and sample.

TABLE A.4. Support for age-specific transition analysis

Age group	Mortality-known	Complete survivors	Min. mortality N_n	Min. survivor N_n	Contiguous cap
50–64	31,335	24,004	301	206	20
65–74	14,789	10,694	141	77	20
75–84	11,662	7,270	85	56	20
85+	3,014	1,345	19	11	.

Notes: Counts refer to baseline states $n = 0, \dots, 20$. The contiguous cap is the largest value c such that every state from $n = 0$ through $n = c$ contains at least 20 mortality-known transitions and at least 20 complete survivor transitions within that age group.

TABLE A.5. State support among respondents aged 85 and older

n	Mortality-known	Complete survivors	Death probability
0	19	11	0.0526
1	20	13	0.0500
2	35	15	0.1143
3	94	52	0.0319
4	96	54	0.0417
5	147	71	0.1156
6	159	81	0.1195
7	205	101	0.1268
8	201	98	0.1692
9	196	95	0.1480
10	205	96	0.1610
11	238	125	0.1849
12	227	90	0.2203
13	203	84	0.2266
14	197	88	0.2183
15	152	58	0.2697
16	161	67	0.2981
17	141	43	0.3404
18	115	42	0.3652
19	117	36	0.3504
20	86	25	0.4419

Notes: This table is diagnostic. The oldest-old are retained in the descriptive analysis, but sparse state cells and survival selection motivate reporting age tests both with and without respondents aged 85 and older.

TABLE A.6. Do demographic and socioeconomic characteristics predict transitions conditional on current health?

Characteristic	Outcome	p -value	N
Education	Death	0.9719	60,493
Education	Next deficit count	<0.0001	43,106
Education	improve	0.0013	43,106
Education	same	<0.0001	43,106
Education	worsen	<0.0001	43,106
Race/ethnicity	Death	<0.0001	60,679
Race/ethnicity	Next deficit count	<0.0001	43,233
Race/ethnicity	improve	0.0116	43,233
Race/ethnicity	same	<0.0001	43,233
Race/ethnicity	worsen	<0.0001	43,233
Sex	Death	<0.0001	60,800
Sex	Next deficit count	<0.0001	43,313
Sex	improve	0.0002	43,313
Sex	same	0.0572	43,313
Sex	worsen	<0.0001	43,313

Notes: Tests condition on baseline deficit-count fixed effects, a quadratic in chronological age, and survey-wave fixed effects for $n \leq 20$. Survivor outcomes are defined among respondents with an observed next-wave deficit count. Standard errors are clustered by respondent.

TABLE A.7. Relative predictive importance of demographic and socioeconomic characteristics and current deficit state

Characteristic	Outcome	Metric	Group	n	Group/ n	N
Education	Mortality	Brier gain	-0.0000	0.0280	-0.000	60,493
Education	Mortality	Deviance gain	0.0000	0.0565	0.000	60,493
Education	Survivor count	Partial R2	0.0043	0.7053	0.006	43,106
Race/ethnicity	Mortality	Brier gain	0.0013	0.0292	0.046	60,679
Race/ethnicity	Mortality	Deviance gain	0.0024	0.0585	0.041	60,679
Race/ethnicity	Survivor count	Partial R2	0.0019	0.7117	0.003	43,233
Sex	Mortality	Brier gain	0.0108	0.0343	0.316	60,800
Sex	Mortality	Deviance gain	0.0137	0.0660	0.208	60,800
Sex	Survivor count	Partial R2	0.0004	0.7101	0.001	43,313

Notes: For the survivor-count outcome, entries are incremental partial R^2 measures. For mortality, entries are proportional Brier-score and deviance improvements from nested models. Each group contribution is evaluated conditional on current deficit state, quadratic age controls, and wave; the n contribution reverses the comparison on the same estimation sample.

TABLE A.8. Group heterogeneity in components of the mortality hazard

Characteristic	Model	Log likelihood	Δ AIC	Δ BIC	N
Education	M0	-8788.37	23.84	0.00	60,493
Education	MA	-8787.70	28.52	31.71	60,493
Education	MG	-8776.97	7.04	10.24	60,493
Education	MGA	-8770.44	0.00	30.22	60,493
Race/ethnicity	M0	-8800.06	73.94	19.86	60,679
Race/ethnicity	MA	-8777.82	35.46	8.42	60,679
Race/ethnicity	MG	-8782.21	44.24	17.20	60,679
Race/ethnicity	MGA	-8757.09	0.00	0.00	60,679
Sex	M0	-8814.49	243.12	225.09	60,800
Sex	MA	-8697.00	10.14	1.12	60,800
Sex	MG	-8741.98	100.09	91.07	60,800
Sex	MGA	-8690.93	0.00	0.00	60,800

Notes: M0 imposes a common mortality curve; MG allows group shifts in the background component; MA allows group shifts in the deficit-dependent component; and MGA allows both. All specifications include common quadratic age and survey-wave adjustments. Δ AIC and Δ BIC are measured relative to the best-fitting model for the characteristic under the corresponding criterion.

TABLE A.9. Adjusted fitted mortality probabilities by group and deficit state

Characteristic	Group	Model	$n = 0$	$n = 5$	$n = 10$	$n = 15$	$n = 20$
Education	College or more		0.0052	0.0195	0.0545	0.0954	0.1580
Education	High school		0.0052	0.0195	0.0545	0.0954	0.1580
Education	Less than high school		0.0052	0.0195	0.0545	0.0954	0.1580
Education	Some college		0.0052	0.0195	0.0545	0.0954	0.1580
Race/ethnicity	Hispanic		0.0047	0.0138	0.0350	0.0610	0.1040
Race/ethnicity	NH Black		0.0125	0.0287	0.0573	0.0883	0.1394
Race/ethnicity	NH Other		0.0046	0.0175	0.0517	0.0944	0.1601
Race/ethnicity	NH White		0.0040	0.0178	0.0561	0.1040	0.1762
Sex	Female		0.0035	0.0132	0.0394	0.0732	0.1276
Sex	Male		0.0066	0.0247	0.0724	0.1306	0.2159

Notes: Entries are fitted next-wave death probabilities from the age- and wave-adjusted mortality specifications, evaluated at selected baseline deficit counts. Education uses the parsimonious common mortality curve; sex and race/ethnicity use the preferred heterogeneous specifications.

TABLE A.10. Does deficit composition account for residual mortality information?

Characteristic	Group gain conditional on n	Group gain conditional on 31 deficits	Remaining ratio	Reduction in group gain	Composition gain
Sex	0.0108	0.0067	0.615	0.385	0.0299
Race/ethnicity	0.0013	0.0006	0.475	0.525	0.0299


Notes: Group gain is the proportional reduction in the Brier score when the indicated characteristic is added to a mortality model containing quadratic age and wave controls and either exact current deficit count n or the 31 individual deficit indicators. The remaining ratio is the group gain conditional on the 31 deficits divided by the group gain conditional on n ; the reduction in group gain is one minus this ratio. Composition gain is the proportional Brier-score reduction from replacing n by the 31 individual deficits in the model without the indicated characteristic. All comparisons within a row use a common estimation sample.

TABLE A.11. Does recent health history predict the next transition conditional on current health?

Outcome	Test statistic	p -value	N
Death	22.68	0.30465	40,187
Next deficit count	72.43	0.00000	30,029
Improvement	82.51	0.00000	30,029
Unchanged count	1.30	0.16479	30,029
Deterioration	73.78	0.00000	30,029


Notes: Each row tests the joint significance of exact lagged-deficit-count indicators in a model containing exact current deficit-count indicators and survey-wave fixed effects. The lagged count is measured at the immediately preceding HRS wave, so the analysis requires three consecutive observations of health. Mortality uses a logit model and reports a cluster-robust Wald χ^2 ; survivor outcomes use linear models and report cluster-robust F -tests. Standard errors are clustered by respondent. Current and preceding deficit counts are restricted to 0–20. The specification is deliberately asymmetric: the question is whether recent history adds predictive information after current health is known. Quadratic age is added separately in the common-sample comparison and in a robustness specification.

FIGURE A.1. Structural stress test of the original stochastic model




Panel A shows probability assigned by the ordinary Poisson survivor kernel to impossible destination states above 31 deficits. Panel B shows the total probability assigned to feasible living destinations plus death.

FIGURE A.2. Observed and fitted mortality at low baseline deficit counts




Observed and fitted mortality at low baseline deficit counts under alternative mortality laws. The background-plus-power specification permits positive mortality when no deficits are measured while retaining a nonlinear deficit-dependent mortality component.

FIGURE A.3. Survivor transition refinement




Panel A shows conditional-survivor total-variation distance for the finite-support Poisson and persistence-mixture specifications. Panel B compares observed and fitted probabilities of an unchanged deficit count.

FIGURE A.4. Age composition across baseline deficit states



The figure plots mean and median chronological age within each baseline deficit count.

FIGURE A.5. Mortality conditional on baseline deficit count, by age



Mortality conditional on baseline deficit count, by age. Points are empirical next-wave death probabilities within baseline deficit states. The figure is descriptive: age differences conditional on n need not represent causal effects of chronological age.