Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

The Ergodicity Gap: A Framework for Understanding When Population Evidence Describes Individual Lives

Shay Tsaban

shay.tsaban@gmail.com

Peres Academic Center

Research Article

Keywords: ergodicity, ensemble average, time average, heterogeneity, absorbing states, subjective well-being

Posted Date: August 31st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10788010/v1

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

Additional Declarations: No competing interests reported.

The Ergodicity Gap: A Framework for Understanding When Population Evidence Describes Individual Lives

Abstract

Ergodicity is the property that permits an average taken across a population to stand in for an average taken over time. Where it fails, the two averages measure different quantities, and evidence about one licenses no inference about the other. Economics has long asked whether an aggregate can be represented by a typical individual. This article asks the converse question, when an aggregate describes any individual's trajectory, and develops it for subjective well-being, where evidence gathered across people at a moment is used to answer questions about single lives unfolding over time. Two research traditions supply the ingredients. Molenaar established that psychological processes are non-ergodic because they are neither stationary nor homogeneous. Peters established that economic processes are non-ergodic because they compound. The ergodicity model of well-being (EMW) rests on four tenets: (a) well-being is a property of trajectories, so the quantity that matters to a person is a time average; (b) the resources that generate well-being compound, so the ensemble average tracks a shrinking minority of trajectories; (c) well-being trajectories contain absorbing states, and entering one tends to remove a person from observation; and (d) survivorship therefore fixes the sign of the gap, which is positive. Closed forms for three elementary processes show that compounding sets the magnitude of the gap, conditional on the curvature of the mapping from resources to reported well-being, while absorption with differential attrition sets its sign. Calibrated gaps are of the same order as the effects the field studies.

Keywords: ergodicity; ensemble average; time average; heterogeneity; absorbing states; subjective well-being

JEL classification: B41; C02; D31; I31

1. Introduction

Economics has spent three decades asking whether an aggregate can be stood in for by a representative individual. Kirman (1992) showed that a representative agent need not behave like any of the agents it is supposed to represent, and that aggregate relations can carry a sign that none of the individual relations composing them carries. Aoki and Yoshikawa (2012) showed that in a wide class of models the dispersion of an aggregate does not vanish as the number of agents grows, so that the mean characterizes no one in particular. The research programme that grew from these results treats heterogeneity and interaction as the source of macroeconomic regularities rather than as noise around a representative case (Delli Gatti et al. 2010; Gallegati and Kirman 2012). Each of these is a question about the step from individuals to an aggregate.

The human sciences make the reverse step constantly and defend it almost never. Their evidence is collected across many people at a single moment. Their questions concern what will happen to particular people as their lives continue. The move from the first to the second is not a matter of sampling, of power, or of statistical care. It is a question about the process under study, and the condition that settles it has both a name and a mathematics, neither of which the human sciences ordinarily invoke.

Well-being research is where the step bites hardest, and it is the case this article develops. The field has become one of the most policy-relevant branches of the behavioral sciences. National statistical agencies collect life-satisfaction data as a matter of routine, governments have been urged to adopt well-being as an explicit objective of policy, and employers report aggregate well-being figures beside conventional performance indicators (Adler and Seligman 2016; Diener et al. 2018; Frijters et al. 2020). The evidentiary base supporting this movement is

very largely a population base. It consists of surveys administered to many people, at one time or at a few times, from which averages, correlations and regression coefficients are computed.

That evidence is used, almost without exception, to answer a question of a different kind. A government asks whether a transfer will make its recipients happier. A physician asks whether a patient will adapt to a new disability. An employer asks whether a program will keep staff out of burnout. Each question concerns what will happen to particular people as their lives continue. The quantity of interest is a property of a trajectory. The quantity estimated is a property of a population. The step from one to the other is so routine that it is rarely stated, still less defended.

There is a name for the condition under which that step is legitimate. A process is ergodic when the average of many realizations at one time equals the average of one realization over a long time (Birkhoff 1931; von Neumann 1932). Ergodicity is not a technicality. It is the precise condition that licenses the inference from a population to a life, and it can fail.

Two research traditions have independently established that it does fail for the processes they study. Molenaar (2004, 2007) demonstrated that psychological processes are generally non-ergodic, because they are neither stationary within persons nor homogeneous across them, and drew the conclusion that inter-individual results cannot be transferred to the individual case. Peters (2019; Peters and Gell-Mann 2016) showed that economic processes are non-ergodic for a different reason, namely that wealth compounds, and that where dynamics are multiplicative the time-average growth rate of an individual differs systematically and computably from the growth rate of the population average.

These two arguments have not been joined, and well-being research has absorbed neither. The omission matters, because the two traditions supply different and complementary things. Molenaar's argument establishes that the inference fails. It does not say by how much, or in

which direction. Peters's argument supplies exactly that for a specified class of dynamics. A well-being science that took both seriously could say more than that population evidence is a poor guide to individual lives. It could say that the evidence errs in a particular direction, for identifiable reasons, and by an amount that depends on measurable features of the process.

This article develops that account. I propose the ergodicity model of well-being (EMW), an integrative framework built on the joint foundation. It treats the gap between population and individual well-being as a theoretical object with its own structure, rather than as noise or as a limitation to be acknowledged in a discussion section. The framework rests on four tenets. Well-being is a property of trajectories, so the quantity that matters to a person is a time average. The resources that generate well-being compound, so the ensemble average comes to track a shrinking minority of trajectories. Well-being trajectories contain absorbing states, and entering one tends to remove a person from observation. Survivorship therefore fixes the sign of the gap, which is positive: population well-being overstates what an individual can expect to live through.

What the article adds to the statement of the problem is a determinate form. Section 8 derives the gap in closed form for three elementary processes. The first shows that compounding sets the magnitude of the gap but not its sign, because whether an advantage in resources reaches reported well-being depends on the curvature of the mapping between them, and the logarithm is the knife-edge. The second and third show that absorption with exit from observation fixes the sign, whether absorption is permanent or recurrent, and that the gap vanishes exactly when everyone stays in view. Calibrating the closed forms with hazards and prevalences of the order reported for depression, long-term unemployment and disabling illness gives gaps of the same order as the effects the field treats as substantively important.

The argument proceeds as follows. Section 2 sets out the concept of ergodicity and its two averages, tracing it from statistical mechanics through the physics of ergodicity breaking and the economics of heterogeneity and multiplicative growth into its separate importations by psychology and economics. Section 3 establishes which of the two averages each common research design recovers. Section 4 defines hedonic ergodicity and the ergodicity gap and states the four tenets. Sections 5 to 7 develop the three mechanisms that generate the gap: multiplicative dynamics, absorbing states, and survivorship in the observed sample. Section 8 is the formal part. Section 9 distinguishes EMW from four adjacent critiques, Section 10 reinterprets four contested literatures, and Section 11 shows that the same structure appears in sciences well outside well-being research. Section 12 examines what institutions that govern by population well-being are led to optimize, Section 13 states the conditions under which the framework does not apply, and Section 14 specifies the evidence that would count against it.

2. Ergodicity and the Two Averages

2.1 Origins in statistical mechanics

The term entered physics through Boltzmann's attempt to justify statistical mechanics. A gas contains too many molecules to track individually, so its macroscopic properties are computed as averages over the ensemble of possible microscopic states. This procedure is only valid if a single system, followed for long enough, visits those states in the proportions the ensemble assigns them. Boltzmann assumed that it does. Birkhoff (1931) and von Neumann (1932) subsequently established the conditions under which the assumption holds, proving that for a measure-preserving transformation the time average of a function along almost every trajectory exists and, when the system is indecomposable, equals the ensemble average.

The result is worth stating plainly, because its structure carries over intact to the human sciences. There are two averages. One runs across many systems at a single moment, which will be called the ensemble average. The other runs along one system across many moments, which will be called the time average. What makes the two equal is ergodicity. It is not a property of the data or of the estimator. It is a property of the underlying process, and whether it holds is a substantive empirical question about that process rather than a matter of statistical technique.

Three conditions are required. The process must be stationary, so that its statistical properties do not change over time. It must be homogeneous across units, so that every individual is governed by the same process. And it must be indecomposable, so that no region of the state space traps a trajectory that enters it. Each of these can fail on its own, and each failure severs the equality in a different way.

2.2 Ergodicity breaking as a measurable property

Physics did not leave the matter at the level of a theorem. Bouchaud (1992) showed that systems whose trapping times have infinite mean break ergodicity weakly: a single trajectory never samples its state space in the proportions the ensemble assigns, however long it is followed, because the time spent in the deepest trap grows with the observation window. The literature on anomalous diffusion then turned the failure into a measured quantity. Metzler et al. (2014) review a family of processes in which the time-averaged observable of a single particle and the ensemble-averaged observable of many particles stay different at all times, and in which a dimensionless parameter comparing the two is estimated from data rather than assumed. Peters and Klein (2013) established the case that matters most here, geometric Brownian motion. In that process the ensemble average grows at one rate while almost every individual trajectory grows at

a lower rate, and the ensemble average is carried by a fraction of trajectories that shrinks toward zero as time passes.

Two lessons travel from this literature intact. Ergodicity breaking is ordinary rather than exotic. And the size of the break is something one computes or measures, not something one argues about.

2.3 Aggregation, heterogeneity and multiplicative growth in economics

Two routes that are rarely placed together led economists to the same structure. The first is the aggregation problem. Kirman (1992) asked whom or what the representative individual represents and answered that, in general, it represents no one. The aggregate of heterogeneous agents need not satisfy the restrictions each agent satisfies, and it can display preferences that no agent holds. Aoki and Yoshikawa (2012) gave the argument a statistical form. In models whose microstructure generates a sufficiently heavy-tailed distribution of agent states, the coefficient of variation of the aggregate does not fall toward zero as the number of agents grows, a property they called non-self-averaging. Where it holds, the mean is not a sufficient statistic for the system, and a policy designed for the mean is designed for a case that does not arise. The agent-based programme that followed treats heterogeneity and interaction as the generators of aggregate regularities rather than as departures from a representative case (Delli Gatti et al. 2010; Gallegati and Kirman 2012).

The second route is the study of multiplicative growth. Champernowne (1953) showed that incomes subject to proportional random shocks above a floor settle into a Pareto distribution, and the result generalizes to any quantity that grows by random proportions against a reflecting barrier (Gabaix 2009). Bouchaud and Mézard (2000) showed that multiplicative wealth dynamics with exchange condense wealth into a small number of agents when exchange is weak

relative to the randomness of returns, and Benhabib et al. (2015) obtained fat tails in wealth from heterogeneous returns to capital in an otherwise standard model.

Two recent contributions make the link to ergodicity explicit. Berman et al. (2021) tested the ergodic hypothesis in an established model of wealth in a growing and reallocating economy and found that, for recent decades in the United States, no stationary distribution exists, so the distribution is unstable on the timescales that matter. Bouchaud and Farmer (2023) showed that self-fulfilling beliefs can make wealth dynamics quasi non-ergodic, with inequality depending on history for periods long compared with a lifetime. In each case the aggregate is shaped by the trajectories that compound favorably. In each case the typical trajectory is somewhere else.

2.4 Two importations into the sciences of the person

Importing a formal construct from the physical sciences into psychology is an established and productive move, provided the import is disciplined. Hirsh et al. (2012) used the information-theoretic notion of entropy to reconceive uncertainty-related anxiety, and self-regulation has long been analyzed through the dynamics of self-organizing systems (Carver and Scheier 2002). More recent work continues in the same mode, reading learning through the formal apparatus of predictive processing (Trapp et al. 2026). What such transfers require is that the borrowed formalism do real work, generating claims that the receiving field could not otherwise state. Ergodicity meets that test unusually well, because it is not a loose analogy about systems and dynamics. It is a specific mathematical condition, and whether a given process satisfies it is decidable.

Molenaar (2004) brought this argument into psychology, and the conclusion he drew was uncompromising. Because human development is non-stationary and because people differ in the processes governing them, the classical inference from between-person variation to within-

person variation is invalid for most psychological constructs. Molenaar (2007) went further and argued that psychological methodology would have to be rebuilt around intra-individual variation, and Molenaar and Campbell (2009) set out the person-specific paradigm that would replace it.

The tradition that followed has largely accepted the diagnosis. Fisher et al. (2018) demonstrated the point empirically, showing that between-person and within-person estimates of the same association differ substantially in real data, often by a factor of two or more in the variance. Speelman and McGann (2013) made the complementary point about the mean itself, observing that psychology's central summary statistic is routinely treated as though it described a typical individual when it describes no one in particular. Work in the idiographic tradition has since developed methods for the single case that do not depend on the ergodic assumption (Beltz et al. 2016; Hurtado-Parrado and López-López 2015; von Eye et al. 2024). It has also debated how far nomothetic and idiographic knowledge can be reconciled at all (De Luca Picione 2015; von Fircks 2025; Zittoun 2017).

Peters (2019) brought the same argument into economics, and reached a conclusion with a different shape. Economic processes are non-ergodic principally because wealth compounds. A gamble repeated over time multiplies, and the growth rate an individual experiences is therefore governed by the expectation of the logarithm rather than by the logarithm of the expectation. Since the first is always at most the second, the individual's long-run growth rate falls below the growth rate of the population average, and it can be negative even where the population average grows without bound (Kelly 1956; Peters and Gell-Mann 2016). The same criterion resolves the St Petersburg paradox without appeal to utility (Peters 2011) and reproduces discounting from the dynamics alone (Adamou et al. 2021). Experiments show that human choice tracks the

distinction, with participants adjusting their risk-taking to whether the environment is additive or multiplicative (Meder et al. 2021; Skjold et al. 2026; Vanhoyweghen et al. 2022).

The framework remains contested among economists (Doctor et al. 2020), and the dispute concerns whether the time-average criterion supplants expected-utility theory or merely supplements it. That controversy is not at issue here. What matters for present purposes is the narrower and less disputed point, which is that the two averages come apart in a direction that can be derived from the dynamics.

Neither literature cites the other with any regularity, and the reason is instructive. Psychology encountered ergodicity as a threat to measurement and inference, and its response was methodological. Economics encountered ergodicity as a claim about what rational agents should maximize, and its response was decision-theoretic. Well-being research sits at the intersection of the two and has inherited the concerns of neither.

3. What Population Designs Actually Estimate

Before the framework can be stated, it is worth being specific about which of the two averages each common design recovers, because the answer is less obvious than it appears.

A single cross-section estimates the ensemble average directly and estimates nothing else. A repeated cross-section estimates a sequence of ensemble averages computed on different people, which describes the changing composition of a population as much as it describes change in anyone. Neither design contains information about a trajectory, and no amount of statistical care can extract information that was never collected.

Panel designs appear to solve the problem, and they solve part of it. A within-person fixed-effects estimator does recover a within-person quantity, which is a genuine advance over the cross-section. It recovers an average of within-person slopes across the people who remain in

the panel, however, and that average is itself an ensemble average taken over trajectories. The cross-lagged panel model, long the standard tool for questions of temporal precedence in well-being research, is more compromised still. Hamaker et al. (2015) showed that it conflates within-person and between-person sources of variance and can return conclusions about temporal ordering that reverse when the two are separated. The dominant workhorse of longitudinal well-being research therefore mixes precisely the two quantities that ergodicity distinguishes.

Intensive longitudinal designs come closest to the time average. Experience sampling and diary methods observe the same person repeatedly at short intervals, which is the only design that estimates a time average directly (Bolger et al. 2003). Their limitation is one of horizon rather than of principle. Such studies typically run for days or weeks, while the absorbing states that matter most in this argument operate over months and years. Work modeling the temporal dynamics of affect within persons has established that individuals differ systematically in the parameters governing their own dynamics, which is the homogeneity condition failing directly (Kuppens et al. 2010).

The practical position is therefore awkward. The designs that dominate well-being evidence estimate the ensemble average. The design that estimates the time average operates over horizons too short to observe the phenomena that make the two diverge. This is not an argument against any of these methods. It is an argument that the gap between them is a theoretical object requiring an account, which is what the remainder of this article supplies.

4. Hedonic Ergodicity and the Ergodicity Gap

Consider a population of individuals indexed by i, each with a well-being value $wi(t)$ at time t. Two averages can be formed. The ensemble average takes the mean across individuals at a fixed time. The time average takes the mean across time for a fixed individual, over an interval

of length T. Well-being research overwhelmingly estimates the first. The questions that motivate it overwhelmingly concern the second.

The individual side of the comparison can be read in two ways, and the distinction matters in Section 8. The expected time average is the mean, over the chances an individual faces, of that individual's time average. The typical time average is the time average of the trajectory at the middle of the distribution, which under multiplicative dynamics lies far from the mean trajectory. The first reading is the natural one where the issue is who stays in view. The second is the natural one where the issue is compounding, since compounding makes the mean trajectory unrepresentative of almost everyone.

Call the population **hedonically ergodic** when the individual time average, in either reading, equals the ensemble average, so that a cross-sectional finding transfers to a life without correction. Call the difference between them the **ergodicity gap**. The gap is defined as the observed ensemble average minus the individual time average over an interval of length T. A positive gap means that the population figure overstates what a person can expect to live through. Fig. 1 illustrates the two averages under dynamics that satisfy the ergodic conditions and under dynamics that violate them.

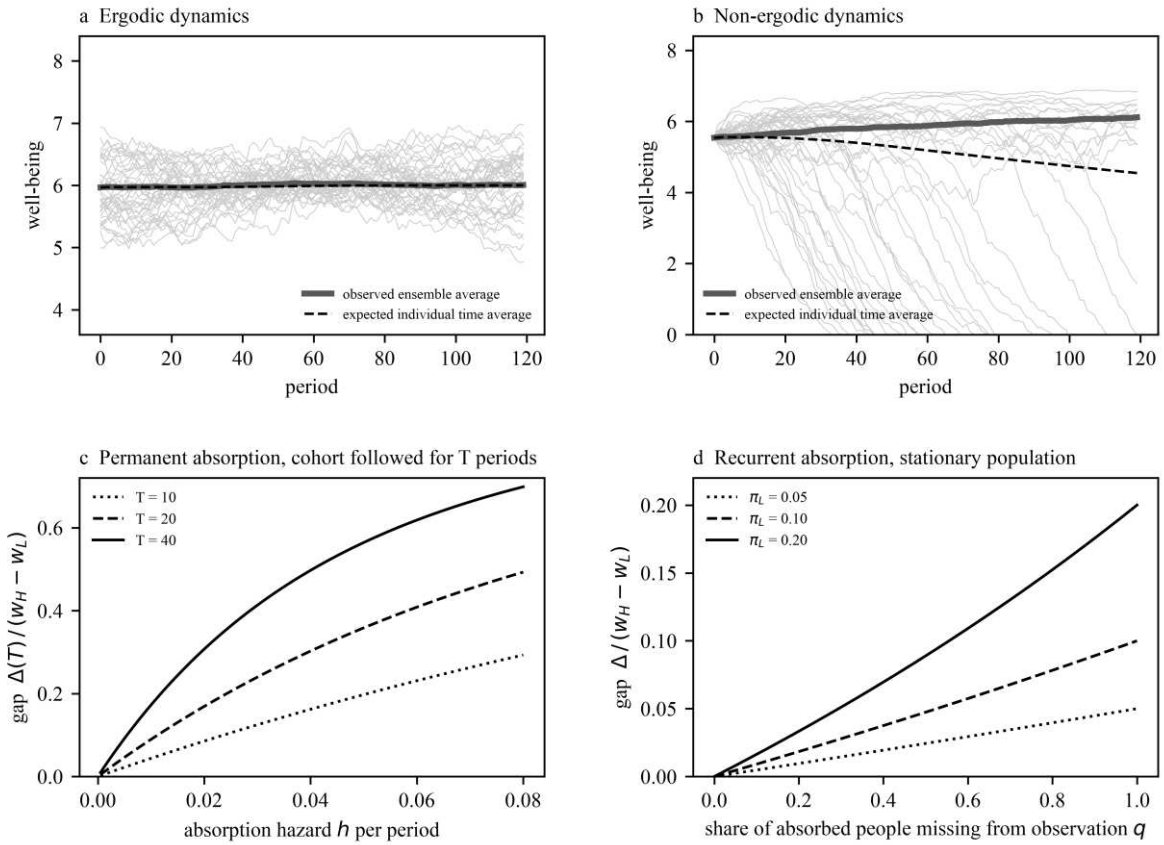


Fig. 1 Two averages under ergodic and non-ergodic dynamics, and the two survivorship gaps. Panel **a** shows an ensemble of simulated trajectories under additive, stationary dynamics with no absorbing region; the observed ensemble average and the expected individual time average coincide. Panel **b** shows trajectories generated by multiplicative resource dynamics in which individuals enter an absorbing region at a constant hazard and leave observation when they do; the observed ensemble average rises while the expected individual time average falls. Panel **c** plots the gap derived in Section 8.2 against the absorption hazard, for three observation windows. Panel **d** plots the gap derived in Section 8.3 against the share of absorbed people missing from observation, for three prevalences of the absorbed state. The simulations illustrate the framework's logic and are not offered as evidence for it

EMW consists of four tenets. The first is ontological. Well-being is a property of trajectories rather than of populations, because the object of concern is a life as it is lived, and a

life is a sequence rather than a point. This is not a novel claim in itself. It is implicit in the idiographic tradition, and it is explicit in accounts that treat well-being as the ongoing satisfaction of needs across the life course rather than as a level attained at a moment (Pincus 2024; Ryff 1989). The tenet's function within EMW is to fix which of the two averages is the theoretically privileged one. Once that is fixed, every ensemble estimate becomes a proxy whose adequacy must be argued rather than assumed.

The second tenet concerns the form of the dynamics. Well-being is generated by resources that compound, so the ensemble average of those resources comes to track a shrinking minority of trajectories. The third concerns the topology of the state space. Well-being trajectories contain absorbing and quasi-absorbing states, so the process is decomposable, and entering such a state tends to remove a person from observation. The fourth is the consequence of the other three. The ergodicity gap is systematically positive. Its sign is fixed by survivorship, and its magnitude increases with the hazard of absorption, with the share of the absorbed who leave observation, and, conditional on the mapping from resources to reported well-being, with the multiplicative component of the dynamics.

An important qualification belongs here, because it determines how much weight the framework can bear. Subjective well-being as ordinarily measured is bounded, typically on a scale from zero to ten, and a bounded quantity cannot follow a multiplicative random walk indefinitely. The multiplicative claim in the second tenet is therefore not a claim about reported well-being itself. It is a claim about the resources that generate it, which are not similarly bounded, and about the monotone mapping from those resources to reported well-being. This matters because it means the second and third tenets are logically independent. Multiplicativity

breaks ergodicity through the mechanism Peters identified. Absorbing states break it through indecomposability, and would do so even if every underlying process were perfectly additive.

The two routes have different consequences, which Fig. 2 summarizes and Section 8 derives. Compounding governs how far the ensemble average sits from the typical trajectory in the space of resources, and the curvature of the mapping governs how much of that distance reaches reported well-being. Absorption combined with exit from observation governs the sign. An objection to either route leaves the other standing, but only the second fixes the direction of the error.

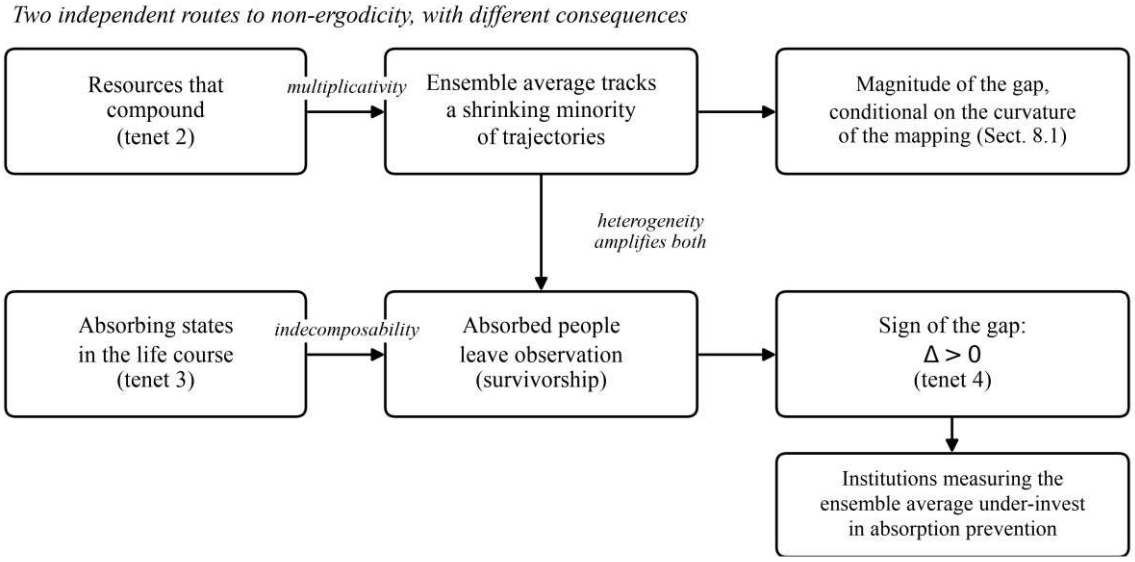


Fig. 2 Structure of the ergodicity model of well-being. Multiplicative resource dynamics and absorbing states are independent routes to non-ergodicity with different consequences. Compounding makes the ensemble average track a shrinking minority of trajectories and sets the magnitude of the gap, conditional on the curvature of the mapping from resources to reported well-being. Absorption combined with exit from observation fixes the sign of the gap, which in turn shapes what institutions measuring population well-being are led to optimize. Heterogeneity across individuals amplifies both routes

5. Multiplicative Dynamics in the Generation of Well-Being

Compounding resources are not a metaphor imported for the occasion. It is the central mechanism of two established theories.

Conservation of resources theory holds that resources beget resources, that those who possess them are best placed to acquire more, and that loss and gain therefore propagate as spirals rather than as independent increments (Hobfoll 1989; Hobfoll et al. 2018). Loss spirals are held to accelerate, because each depletion reduces the capacity to withstand the next demand. The broaden-and-build theory makes the complementary claim on the positive side, holding that positive emotions widen the repertoire of thought and action and thereby build durable resources that raise the probability of further positive experience (Fredrickson 2001). Both theories describe a process in which the effect of an event depends on the state the person is in when it arrives. That state-dependence is precisely what distinguishes multiplicative from additive dynamics.

The job demands-resources tradition supplies the same structure in an organizational setting. Demands deplete and resources buffer, and the buffering is conditional, so that identical demands produce different outcomes depending on the resources available to absorb them (Bakker and Demerouti 2017; Demerouti et al. 2001). Sonnentag et al. (2017) note that recovery is itself resource-dependent, so that the workers least able to recover are those who most need to, which is a compounding rather than an additive relation.

EMW draws the consequence that the psychological literature has not. Where the effect of an event scales with current state, the population mean of resources is a poor summary of what the process does to individuals within it. The mean is driven upward by the trajectories that compound favorably, and those trajectories become increasingly unrepresentative of the median

trajectory as time passes, which is the wealth-condensation result of Section 2.3 transposed to psychological resources. Whether the divergence reaches reported well-being depends on how reports respond to resources. If reports rise with resources less steeply than the logarithm, the ensemble average of reports also runs ahead of the typical report. If they rise exactly logarithmically, the two coincide. If they flatten faster than the logarithm, the order reverses. Section 8.1 derives the condition. What is unconditional is the claim about resources: the ensemble average of the things that generate well-being describes a life that almost no one is living.

The framework accordingly predicts that the correspondence between cross-sectional and within-person estimates of the same well-being relation will weaken as the multiplicative component of the dynamics grows, and that the weakening will be greater over longer observation windows. Fisher et al. (2018) found substantial group-to-individual discrepancies over relatively short windows, which EMW reads as a lower bound rather than as a typical value.

A second prediction follows for the timing of effects. Under additive dynamics, an intervention delivering a fixed benefit raises every recipient's trajectory by the same amount, and its population effect is a faithful summary of its individual effects. Under multiplicative dynamics it is not, because the same input is amplified for the resource-rich and damped for the resource-poor. EMW therefore predicts that the variance of individual responses to well-being interventions will exceed what additive models anticipate, and that the mean effect will be a progressively worse guide to the modal effect as the intervention's causal chain lengthens.

Influential accounts of intentional well-being change have been framed in additive terms. The architecture proposed by Lyubomirsky et al. (2005) partitions well-being variance into a genetic set point, circumstances, and intentional activity, and treats the last as the tractable

component available for deliberate improvement. Partitioning variance is a between-person operation, and a share of population variance does not translate into a share of any individual's attainable change. Under multiplicative dynamics the two quantities come apart, because the return on intentional activity depends on the resource state in which it is undertaken. EMW does not dispute that intentional activity matters. It disputes that its contribution can be characterized by a population variance share, and predicts that the same activity will yield systematically different returns across the resource distribution.

The point generalizes to how the field represents risk. Expected-utility reasoning, including its descriptive successors, evaluates a prospect by averaging over the states of the world that might obtain (Kahneman and Tversky 1979). That averaging is an ensemble operation performed in the imagination, and it is the appropriate operation when the prospect is one of many faced by many people. It is the wrong operation when a single person faces the prospect repeatedly and the outcomes compound, which is the situation that generates the divergence Peters identified. Well-being research imports the ensemble habit from decision theory without importing the conditions under which it is valid.

6. Absorbing States and the Asymmetry of Recovery

The second route to non-ergodicity requires no assumption about multiplicativity. It requires only that some states be difficult to leave.

A state is absorbing when the probability of leaving it is zero, and quasi-absorbing when that probability is low enough that most trajectories entering it do not exit within the observation window. Their presence makes a process decomposable, which is sufficient on its own to break the equality of the two averages. An ensemble average computed over a population containing both absorbed and unabsorbed members describes no individual trajectory. It describes a mixture.

Such states have been documented repeatedly in well-being research without being named as a class. The set-point tradition originally held that people return to baseline after major life events, an expectation that the earliest evidence appeared to support (Brickman et al. 1978). The longitudinal record has been considerably less accommodating. Lucas et al. (2003) found incomplete recovery following widowhood and divorce, and Lucas et al. (2004) found that unemployment shifts the set point itself, with life satisfaction failing to recover to its former level even after re-employment. Clark et al. (2001) documented the same scarring in an independent panel, and Knabe and Rätzel (2011) showed that the residue includes both the memory of past unemployment and the anticipation of future risk. Clark et al. (2008) traced the full time path of life satisfaction around major events and found that anticipation and lagged recovery differ markedly by event type. Oswald and Powdthavee (2008) found only partial adaptation to serious disability, with a substantial permanent decrement. Meta-analytic synthesis confirms that adaptation is event-specific and frequently incomplete, and that several classes of event leave lasting change (Luhmann et al. 2012).

Reviewing this record, Diener et al. (2006) and Lucas (2007) concluded that the strong form of set-point theory could not be sustained. Headey (2008) reached the same verdict from panel evidence, and Fujita and Diener (2005) found that although most people's satisfaction is stable, a substantial minority undergo durable change, which is the pattern a mixture of absorbed and unabsorbed trajectories produces.

Clinical dynamics supply the sharpest cases. Systems with alternative stable states exhibit generic warning signatures as they approach a transition, including critical slowing down, and these signatures have been identified across physical, ecological and social domains (Scheffer et al. 2009). Van de Leemput et al. (2014) detected exactly that signature in emotional time-series

data preceding both the onset of depression and recovery from it. Related work has modeled depression explicitly as a complex dynamic system, with a symptom network capable of settling into a self-sustaining configuration (Cramer et al. 2016; Wichers 2014). Methods for estimating such networks from psychological data have since made the approach tractable (Epskamp and Fried 2018). The same literature has argued for replacing static models of onset with dynamic ones (Nelson et al. 2017). Burnout has been characterized in similar terms, as a condition that becomes self-maintaining once established because it degrades the very resources required to escape it (Maslach and Leiter 2016).

At the limit, mortality is an absorbing state in the strict sense, and it is not independent of well-being. Diener and Chan (2011) and Steptoe et al. (2015) document a robust association between subjective well-being and subsequent longevity, while Case and Deaton (2015) and Blanchflower and Oswald (2020) trace rising mortality and extreme distress in populations whose average reported well-being moved comparatively little.

EMW treats these findings as instances of a single structural feature rather than as a list of exceptions to adaptation. The framework predicts that the ergodicity gap will be larger in domains where absorbing states are more prevalent. Discrepancies between population and individual estimates should therefore be greatest for exactly those outcomes that matter most, namely severe depression, chronic unemployment, disabling illness, and death. It also predicts that adaptation estimates derived from population panels will overstate the adaptation available to any individual, because the population figure averages over trajectories that recovered and trajectories that did not, and the latter are systematically less likely to remain in view.

That last clause is the mechanism to which I now turn, because it converts a claim about dynamics into a claim about the direction of the error.

7. Survivorship and the Sign of the Gap

A gap between two averages need not have a consistent sign. EMW claims that this one does, and the reason is that absorption and observation are not independent.

Entering an absorbing state tends to remove a person from the sample. The unemployed leave employer well-being surveys. The severely depressed are less likely to complete follow-up waves. Those who die leave every panel permanently. Panel attrition is well documented as non-random and as related to the very characteristics under study (Lugtig 2014; Watson and Wooden 2009). Econometrics has understood the general problem of estimating from a sample whose composition is determined by the outcome since Heckman (1979). Occupational epidemiology has confronted an especially clean version of it. The healthy worker survivor effect arises because workers who remain employed are systematically healthier than those who leave, so that a workforce sample grows less representative of the exposed population the longer it is followed (Buckley et al. 2015; Naimi et al. 2013).

The consequence for well-being evidence is direct. An observed ensemble average is computed over those still present. Those still present are disproportionately those who have not been absorbed. The observed population figure therefore exceeds the population figure that would obtain if the absorbed remained visible, and it exceeds by more the longer the observation continues.

EMW draws three predictions from this. First, longitudinal well-being estimates should be upward biased relative to the individual expectation, and the bias should grow with the length of the observation window rather than diminish. Second, and counter-intuitively, apparent adaptation in panel data should be partly an artifact of composition, since the trajectories that fail to adapt are the trajectories most likely to disappear. Third, studies that succeed in retaining

hard-to-follow participants should report both lower average well-being and more pessimistic adaptation estimates than studies that do not, and the difference should scale with retention. The third prediction is directly testable against existing panels and does not require new data collection.

This is where EMW departs most sharply from the standard treatment. In the standard treatment, attrition is a nuisance to be corrected by weighting, and the correction is judged successful when the weighted sample resembles the population on observed covariates. EMW holds that the problem is not primarily one of representativeness on covariates. It is that the quantity being estimated is the wrong quantity. Weighting a cross-section back to population proportions produces a better estimate of the ensemble average, which remains an average over a mixture of trajectories, and which no individual lives.

8. The Gap Under Simple Dynamics

The claims that the ergodicity gap has a sign, that it grows with absorption hazard and with the observation window, and that compounding sets its magnitude can be derived rather than asserted. Three elementary processes suffice, one for each mechanism. Their purpose is not realism. It is to show that the framework's central claims have determinate forms and are therefore open to being wrong.

8.1 Compounding and the curvature of the mapping

Let the resources that generate well-being follow geometric Brownian motion, $dx = x(\mu\, dt + \sigma\, dB)$, the canonical model of a quantity that grows by random proportions (Peters and Klein 2013). Two growth rates describe this process, and they differ. The ensemble average $E[x_t] = x_0 e^{\mu t}$ grows at rate μ. Almost every individual trajectory grows at rate $\mu - \sigma^2/2$, and the trajectory at the median of the distribution, $\tilde{x}_t = x_0 \exp[(\mu - \sigma^2/2)t]$, is the typical one. The

difference $\sigma^2/2$ is the ergodicity gap in the growth of resources, and it is not small. With an annual volatility of thirty percent it amounts to four and a half percentage points a year. The ensemble average is carried by a fraction of trajectories that shrinks toward zero as time passes, which is the formal content of the second tenet.

Reported well-being is not resources. Let it be an increasing function of them, $w = f(x)$, and index the curvature of the function by a single parameter through the family $f_\eta(x) = (x^\eta - 1)/\eta$ for $\eta \neq 0$, with $f_0(x) = \ln x$. The value $\eta = 1$ gives a linear mapping, $\eta = 0$ the logarithm, and negative values mappings more concave than the logarithm. The ensemble average of reported well-being at time t has the closed form

$$E[f_\eta(x_t)] = \frac{x_0^\eta \exp[\eta(\mu - \sigma^2/2)t + 1/2\,\eta^2\sigma^2 t] - 1}{\eta},$$

while the typical individual's reported well-being is $f_\eta(\tilde{x}_t)$. Their difference,

$$D_\eta(t) = \frac{x_0^\eta\, e^{\eta(\mu-\sigma^2/2)t}\left(e^{\eta^2\sigma^2 t/2} - 1\right)}{\eta},$$

has the sign of η and vanishes at $\eta = 0$. Averaging over an observation window preserves the sign, since the integrand has the same sign at every t.

Three consequences follow, and Fig. 3 shows them. Where reported well-being rises with resources less steeply than the logarithm, the ensemble average of reports runs ahead of the typical report, and the distance grows with $\sigma^2 t$. Where the mapping is exactly logarithmic, the ensemble average of reports equals the typical report at every t, even though the ensemble average of resources is racing away from almost everyone. Where the mapping flattens faster than the logarithm, the ensemble average of reports falls below the typical report, because the mapping compresses the lucky trajectories and stretches the unlucky ones.

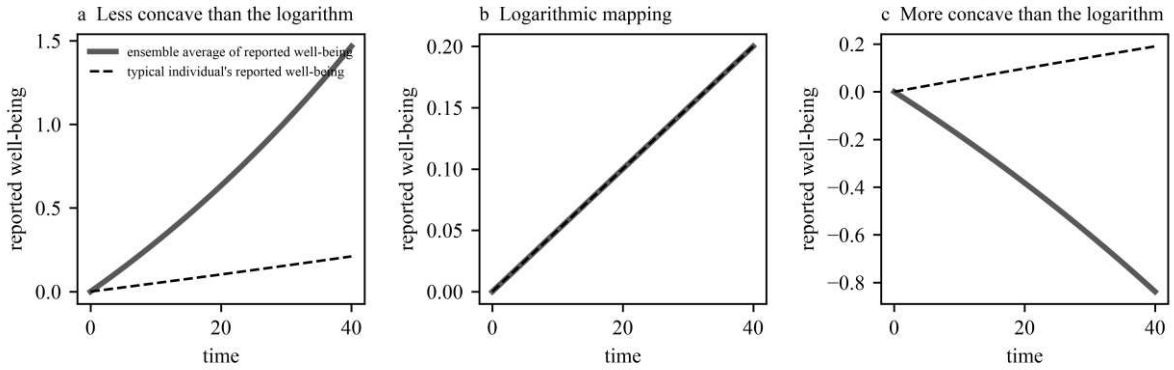


Fig. 3 The multiplicative route under three curvatures of the mapping from resources to reported well-being. Resources follow geometric Brownian motion with drift 0.05 and volatility 0.30 from an initial value of one. Each panel plots the ensemble average of reported well-being and the reported well-being of the typical individual, using the closed forms of Section 8.1. Panel **a** uses a mapping less concave than the logarithm, for which the ensemble average runs ahead of the typical individual. Panel **b** uses the logarithmic mapping, for which the two coincide at every time. Panel **c** uses a mapping more concave than the logarithm, for which the order reverses

The result has a pointed application. The best-replicated finding about income and life evaluation is that evaluation rises approximately linearly in the logarithm of income (Kahneman and Deaton 2010; Killingsworth 2021; Stevenson and Wolfers 2008). For the income channel, then, life evaluation sits at the knife-edge, and compounding alone produces no gap between the ensemble average of evaluations and the typical evaluation. The multiplicative route contributes to the gap in reported well-being only where the mapping is less concave than logarithmic, which is plausible for resources that enter well-being more directly than income does, and it reverses where bounded scales flatten at the top. That is why the framework does not rest its sign claim on compounding. Compounding determines how unrepresentative the ensemble average of resources is, and how far apart the two averages of reported well-being can be. The sign is fixed elsewhere.

Heterogeneity sharpens the conclusion rather than softening it. Suppose individuals differ in μ or σ, as Kuppens et al. (2010) found for the parameters of affective dynamics. The ensemble average of resources is then eventually dominated by the subpopulation with the highest drift, and the coefficient of variation of the aggregate does not shrink as the number of individuals grows. That is non-self-averaging in the sense of Aoki and Yoshikawa (2012), arrived at from the time axis.

8.2 Permanent absorption with exit from observation

Take a population in which each person occupies one of two states. In the well state the person reports well-being w_H, and in the absorbed state the person reports w_L, with $w_L < w_H$. Everyone begins in the well state. In each period a person in the well state enters the absorbed state with probability h, the absorption hazard, and absorption is treated as permanent. Suppose further that entering the absorbed state removes the person from observation, which is the limiting form of the survivorship mechanism described above.

Under these assumptions the observed ensemble average equals w_H in every period, because only unabsorbed people are ever observed. The expected time average for an individual followed over T periods is obtained by averaging the expected well-being at each period, which gives

$$a_T = w_L + (w_H - w_L)\,\frac{1-(1-h)^T}{T\,h}.$$

The ergodicity gap is the difference between the two:

$$\Delta(T) = (w_H - w_L)\left[1 - \frac{1-(1-h)^T}{T\,h}\right].$$

Three properties follow immediately, and each corresponds to a claim made informally above. As the absorption hazard approaches zero the bracketed term approaches zero and the gap

vanishes, which recovers the ergodic case in which population evidence transfers to individuals without correction. As the observation window lengthens with a positive hazard, the bracketed term approaches one and the gap approaches the full well-being decrement $w_H - w_L$. The gap is increasing in both h and T, so longer observation of a population containing absorbing states makes population evidence less representative of individual lives rather than more. Fig. 1c plots the gap against the absorption hazard for three window lengths.

The last property deserves emphasis, because it inverts an ordinary intuition. Longer panels are usually understood to give better access to individual experience than short ones. Under this model they give a progressively worse estimate of the individual time average, because the accumulating absorbed population is exactly the population that has left the sample.

8.3 Recurrent absorption with partial attrition

The cohort model makes three extreme assumptions. Absorption is permanent, exit from observation is certain, and everyone starts well. Relaxing all three at once shows what the survivorship mechanism does on its own.

Let a person in the well state enter the absorbed state with probability h in each period and leave it with probability r, so that absorption is recurrent rather than permanent. The two-state chain is irreducible when $r > 0$. By the ergodic theorem for Markov chains, every individual's long-run time average then equals the stationary mean, $a = \pi_H w_H + \pi_L w_L$, where $\pi_L = h/(h + r)$ is the long-run prevalence of the absorbed state and $\pi_H = 1 - \pi_L$. In a population that has reached its stationary distribution, the ensemble average over everyone equals a as well. The chain is ergodic, and if everyone were observed there would be no gap.

Now let a share q of the people in the absorbed state be missing from observation, because they have left the employer's survey, skipped the follow-up wave, or died, while

everyone in the well state is observed. The observed ensemble average is a weighted mean over the people still in view, and the ergodicity gap is

$$\Delta = \frac{(w_H - w_L)\, q\, \pi_H \pi_L}{1 - q\, \pi_L}.$$

The expression isolates the survivorship mechanism exactly. The gap is zero when $q = 0$ and positive whenever $q > 0$ and the absorbed state has positive prevalence. It increases in q, and it increases in the well-being decrement. Fig. 1d plots it against q for three prevalences.

It does not depend on the length of the observation window, because the population is stationary. What the window does is separate this permanent component from a transient one. A cohort that starts in the well state has an expected time average above a at first, since it has not yet experienced its share of absorption, and that excess decays at the rate $(1 - h - r)$ per period. Over a short window such a cohort looks better than the observed population, over a long window worse, and the difference converges to Δ. Longer observation does not remove the gap. It strips away the initial-conditions effect that conceals it.

8.4 Magnitudes

Tables 1 and 2 report the two closed forms as fractions of the well-being decrement $w_H - w_L$. For the cohort model the entries cover hazards of one, two and four percent per period and windows of ten, twenty and forty periods, which in annual data span the range of a working life. For the stationary model they cover prevalences of five, ten and twenty percent and attrition shares of one quarter, one half and three quarters. These values are of the order reported for major depression, long-term unemployment and disabling illness in working-age populations. The tables are calibrations rather than estimates, and the framework stands or falls on the derivations rather than on any entry.

Table 1 Ergodicity gap under permanent absorption with exit from observation, as a fraction of the well-being decrement, by absorption hazard per period and observation window

Absorption hazard	T = 10	T = 20	T = 40
h = 0.01	0.044	0.090	0.172
h = 0.02	0.085	0.169	0.307
h = 0.04	0.162	0.303	0.497

Note. Entries are $\Delta(T)/(w_H - w_L)$ from Section 8.2.

Table 2 Ergodicity gap under recurrent absorption with partial attrition, as a fraction of the well-being decrement, by prevalence of the absorbed state and by the share of absorbed people missing from observation

Prevalence of the absorbed state	q = 0.25	q = 0.50	q = 0.75
π_L = 0.05	0.012	0.024	0.037
π_L = 0.10	0.023	0.047	0.073
π_L = 0.20	0.042	0.089	0.141

Note. Entries are $\Delta/(w_H - w_L)$ from Section 8.3.

With a decrement of two scale points on an eleven-point scale, which is within the range of the event effects reviewed in Section 6, the cohort entries translate into gaps between about 0.1 and 1.0 scale points, and the stationary entries into gaps between about 0.02 and 0.3 points. The larger entries are of the same order as the cross-sectional associations the well-being literature treats as substantively important. The smaller ones are not negligible against the precision with which national well-being series are now reported. The point of the calibration is not that any entry is right. It is that a quantity this large has been living inside the population figures without a name.

The models are deliberately crude. Well-being is not two-valued, and resources do not compound at a constant volatility. Graded outcomes and time-varying volatility change the magnitudes while leaving the structure intact: the sign of the survivorship term depends only on $q > 0$, and the sign of the multiplicative term depends only on the curvature of the mapping.

What the three derivations establish together is that the sign of the ergodicity gap is not an artifact of the way the argument was phrased. It follows from who leaves the sample.

9. Position Within Adjacent Literatures

Four existing critiques resemble the one developed here, and EMW is easier to assess once it is distinguished from each.

The first is the idiographic critique itself. Molenaar (2004) and those following him argue that population results do not transfer to individuals, and EMW accepts that argument in full. The difference concerns what follows from it. The idiographic response has been largely methodological, calling for person-specific designs and analyses that estimate each individual's own dynamics (Beltz et al. 2016; Hurtado-Parrado and López-López 2015). That programme is valuable, and it is also expensive, slow, and difficult to scale to the populations that well-being policy addresses. EMW is not a call to replace population research. It holds that the discrepancy between the two averages is itself structured, that its structure can be theorized, and that a signed and bounded correction is a more practical objective than the abandonment of population evidence. Where the idiographic tradition says that the ensemble tells you nothing about the individual, EMW says that it tells you something predictably wrong, which is a more useful position to be in.

The second is the critique of the representative agent. Kirman (1992) and the heterogeneous-agents programme argue that an aggregate cannot be stood in for by a typical individual (Delli Gatti et al. 2010; Gallegati and Kirman 2012), and Aoki and Yoshikawa (2012) showed that the mean can fail to characterize a system at all. EMW is the converse of that critique and inherits its formal apparatus. The representative-agent question runs from individuals to the aggregate and asks whether the aggregate behaves like one of them. The

ergodicity question runs from the aggregate to an individual and asks whether any one of them lives like it. The two fail together whenever dynamics are multiplicative, because the ensemble average is then dominated by a vanishing fraction of trajectories. What EMW adds is the time axis, and the absorbing states that the agent-based literature has treated mainly as features of firms and markets rather than of lives.

The third is the literature on heterogeneous treatment effects. It is well understood that an average effect can conceal wide variation, and that subgroups may respond differently or in opposite directions. This is a claim about variance around a mean at one point in time. EMW makes a claim about the mean itself over time, and the two are independent. A well-being intervention could have a perfectly homogeneous effect across every individual and the ergodicity gap would remain, because the gap is produced by the dynamics of the process and the composition of the observed sample rather than by differences in responsiveness. Treatment-effect heterogeneity is a reason the mean may not describe a given person now. Non-ergodicity is a reason the mean does not describe anyone's life.

The last is measurement error. Discrepancies between within-person and between-person estimates are sometimes attributed to unreliability, on the reasoning that individual-level estimates are noisier and that noise attenuates associations. This explanation makes a directional prediction that can be checked. Measurement error attenuates, so longer averaging within persons should reduce it and bring the two estimates into closer agreement. EMW predicts the opposite, because both multiplicative divergence and cumulative absorption grow with the observation window. The two accounts are therefore separable with existing panel data, and the separation is the first of the studies proposed below.

10. Reinterpreting Four Literatures

A framework earns its place by making better sense of evidence already in hand. Four disputed literatures are reconsidered here.

Hedonic adaptation. The set-point literature has spent four decades adjusting a claim that people return to a baseline after major events. The adjustments have accumulated to the point where the original claim is barely recognizable, with adaptation now understood as partial, event-specific, and variable across individuals (Diener et al. 2006; Luhmann et al. 2012). EMW suggests that part of the difficulty is structural. Adaptation estimated from a population is the average behavior of a mixture containing both recovering and absorbed trajectories, and it is estimated on a sample from which the absorbed are progressively removed. Under those conditions, a finding of substantial average adaptation is consistent with a large minority who do not adapt at all, and the population estimate cannot distinguish the two situations. The observed drift of the literature toward weaker adaptation claims as panels lengthened is what EMW would predict, since longer panels accumulate more absorbed trajectories before losing them.

Income and well-being. The dispute over whether well-being rises with income above a threshold has been conducted almost entirely in cross-sections and repeated cross-sections (Easterlin 1995; Kahneman and Deaton 2010; Killingsworth 2021; Stevenson and Wolfers 2008). The most recent contribution resolved an apparent contradiction by identifying an unhappy minority for whom the relation flattens (Killingsworth et al. 2023). That resolution is instructive, because it locates the answer in the heterogeneity of trajectories rather than in the population mean. Bruno (2025) showed in an agent-based model that the sign of the aggregate relation between income and life satisfaction depends on whether individuals compare themselves with the mean or with the median of their reference group, and validated the pattern on two decades of

Eurobarometer data. That is the aggregation problem of Section 2.3 reaching well-being research by the cross-sectional route.

EMW supplies the longitudinal route. Income is the paradigm case of a multiplicative resource, its distribution has no stationary form on the timescales that matter (Berman et al. 2021; Bouchaud and Farmer 2023), and the relation between income and well-being over a life is therefore a relation between compounding quantities. A cross-sectional slope estimated across people at one moment need not equal the slope a person traverses over time. The logarithmic form of the evaluation slope, which Section 8.1 identified as the knife-edge, is exactly the case in which the cross-section is silent about the direction of the discrepancy. Clark (2018) notes that the field's central questions remain unsettled after four decades. EMW proposes one reason. The questions concern time averages and the evidence concerns ensemble averages.

Organizational well-being programs. Individual-level workplace well-being interventions perform poorly under evaluation. Fleming (2024) found no benefit across a large employer sample for most such programs. That result is consistent with earlier meta-analytic caution about intervention effects (Joyce et al. 2016; Richardson and Rothstein 2008), and with arguments that the field asks what works without asking for whom and under what circumstances (Nielsen and Miraglia 2017). EMW offers a structural account. These programs are designed to raise the ensemble mean and are evaluated on whether they do. They rarely target absorption hazard, which is the rate at which employees enter states from which they do not return. A program can raise the mean while leaving that hazard untouched, in which case the time average of the median employee is unchanged. If the program consumes time or attention that would otherwise have gone to recovery, the framework predicts that the time average will fall even as the measured mean rises. Research on human resource management and well-being has

documented that practices raising average reported outcomes may intensify demands at the same time (Guest 2017; Peccei and Van De Voorde 2019), which is the pattern EMW anticipates.

Inequality of well-being. Attention has turned to the dispersion of well-being rather than its level, on the argument that dispersion is a more complete measure of how a society is doing (Goff et al. 2018; Kalmijn and Veenhoven 2005). EMW supports the turn and identifies a limitation. Cross-sectional dispersion is dispersion across people at a moment. It is not dispersion across the lives those people are living, and the two coincide only under ergodicity. A society in which everyone spends a fifth of their life in the bottom quintile of well-being and a society in which a fixed fifth of people spend their whole lives there produce the same cross-sectional dispersion and are morally quite different. Distinguishing them requires the time average, and the framework specifies why the cross-section cannot.

11. The Same Structure Elsewhere in the Human Sciences

The argument has been developed for well-being because that is where the two averages are exchanged for one another most often and with least comment. The structure it describes is not peculiar to well-being, and saying where else it appears is one way of showing that the framework is not an artifact of a single literature's habits.

Occupational epidemiology met the survivorship mechanism first and gave it a name. The healthy worker survivor effect arises because workers who remain employed are systematically healthier than those who leave, so a workforce sample grows less representative of the exposed population the longer it is followed (Buckley et al. 2015; Naimi et al. 2013). Epidemiology treats this as a bias to be corrected by design, usually by modeling exposure and leaving jointly. EMW reads it as the general case rather than as a hazard of one setting. Wherever an outcome influences the probability of remaining under observation, the observed

ensemble average drifts away from the individual time average, and it drifts in a direction fixed by which trajectories disappear. Econometrics has possessed the general form of the problem since Heckman (1979), although it has been posed there as selection into a sample rather than as a divergence between two estimands.

Psychology, physics and economics supply the other components. Molenaar's (2004) demonstration that psychological processes are generally non-ergodic was not a claim about well-being, and Fisher et al. (2018) established the group-to-individual discrepancy across a range of psychological constructs. Work on critical transitions supplies the topology, with alternative stable states and the signatures that precede a shift documented in physical, ecological and social systems alike (Scheffer et al. 2009). The physics of ergodicity breaking supplies the measurement, in the demonstration that the distance between a time average and an ensemble average is a quantity to be estimated (Metzler et al. 2014). Ergodicity economics supplies the dynamics, in Peters's (2019) demonstration that compounding drives an individual's time-average growth rate below the growth rate of the population average. The heterogeneous-agents literature supplies the aggregation result that the mean of such a population describes none of its members (Aoki and Yoshikawa 2012; Kirman 1992). Each of these literatures is well established. What none of them has done is state the conditions jointly and ask what follows for a field that reasons from populations to lives as a matter of routine.

The consequence is visible wherever the two kinds of evidence are placed side by side. Case and Deaton (2015) documented rising mortality and distress in populations whose average reported well-being moved comparatively little, which is the divergence of this article read from the other end: an absorbing outcome rising while the ensemble average, computed over those still present, barely registers it. The claim being made is therefore not that well-being research is

careless. It is that a condition required by every science reasoning from populations to individual lives has gone largely unexamined, and that it has gone unexamined in the field whose subject matter is precisely how a life is going.

12. Governing by the Wrong Average

The final tenet has an institutional consequence, and it is the sharpest claim in this article.

Institutions optimize what they measure. This is among the best-established findings in the sociology of quantification, running from Campbell's (1979) observation that quantitative indicators become corrupted as they become consequential, through detailed accounts of how public measures reshape the conduct they purport to describe (Espeland and Sauder 2007). What EMW adds is a specification of the direction of the distortion in this particular case.

Institutions that measure well-being measure ensemble averages, because that is what surveys produce and what dashboards report. Consider what such an institution sees when it reduces absorption hazard. Preventing a fraction of employees from entering chronic burnout keeps those people in the workforce and in the sample. Their reported well-being is low, since they were at risk. The measured average therefore falls, or at best does not rise, while the time average of the affected individuals rises substantially. Now consider what the institution sees when it improves the experience of those already doing well. The measured average rises, and no absorption hazard has changed. An institution optimizing the ensemble average will prefer the second intervention to the first. EMW predicts that it will systematically under-invest in absorption prevention, and that the under-investment will be greater the more seriously the institution takes its own measurement.

The prediction is uncomfortable because it implies that improving well-being measurement, in the ordinary sense of making it more frequent and more accurate, does not fix

the problem and may worsen it. A more accurate ensemble average is still an ensemble average. The correction EMW implies is not better measurement of the same quantity but measurement of a different one: the rate at which people enter states they do not leave, and the duration of those states.

This connects the framework to a normative literature it does not otherwise engage. Accounts of human flourishing that specify sustained functioning across domains over time, rather than momentary affect, are asking for a time average even when they do not put it that way (VanderWeele 2017). EMW supplies the formal reason why such accounts cannot be operationalized through cross-sectional means, however carefully those means are constructed.

12.1 Implications for measurement and policy practice

Four practical consequences follow, and each is implementable with instruments that already exist. First, report the two averages separately. A national or organizational well-being series should state the ensemble average over those observed, the share of the target population not observed, and the reason they are absent, so that a reader can tell whether a rising series reflects better lives or a thinning sample. Second, collect the two quantities the framework identifies as missing: the rate at which people enter states they do not leave, and the time they spend there. Both are available from administrative records of long-term sickness absence, exit, hospitalization and death, and neither is recoverable from a satisfaction score.

Third, evaluate interventions on the hazard of entry and the duration of stay as well as on the mean. A program that raises the mean without moving the hazard has, on this framework, done nothing for the typical life. Fourth, where the share of the absorbed who leave observation can be estimated, use the closed form of Section 8.3 to convert an observed average into a bound

on the time average a typical member of the population can expect. The correction is one line of arithmetic, and it turns a figure that describes no one into a figure that describes someone.

13. Scope Conditions and Disconfirming Evidence

A framework that cannot fail explains nothing, so the conditions under which EMW does not apply should be stated as plainly as the conditions under which it does.

As the underlying dynamics approach additivity, stationarity and homogeneity, and as the hazard of absorption approaches zero, the ergodicity gap approaches zero. Where those conditions hold, cross-sectional and panel evidence licenses individual inference without correction, and the framework has nothing to contribute. Some well-being domains plausibly satisfy them well enough for practical purposes. Momentary affective responses to mild, frequent, reversible events are a reasonable candidate, since such events neither compound appreciably nor risk absorption. Diary and experience-sampling designs, which sample within persons at high frequency over short windows, are the methods best suited to those domains and are correspondingly least affected by the argument developed here (Bolger et al. 2003).

Two further scope conditions come directly from Section 8. The multiplicative route contributes to the gap in reported well-being only where the mapping from resources to reports is less concave than the logarithm. For life evaluation and income, where the evidence favors a logarithmic mapping, the route contributes nothing to the sign, and the framework's claims there rest on absorption and survivorship alone. And the survivorship route vanishes when everyone stays in view. A panel that retains the absorbed as faithfully as the well, and that observes the ensemble average converging on the long-run individual average, has not refuted the framework. It has found the case $q = 0$, which the framework predicts to be ergodic.

The framework would be weakened by several findings. If within-person and between-person estimates of well-being relations converged as observation windows lengthened, rather than diverging, the multiplicative claim would be in serious difficulty. If panels with unusually high retention of vulnerable participants reported average well-being and adaptation estimates no different from ordinary panels, the survivorship mechanism would lose its principal support. If interventions that demonstrably reduce absorption hazard produced no improvement in individual time averages, the fourth tenet would fail. Each of these comparisons can be made with data that already exist.

Two further limitations deserve statement. First, EMW is a framework about the relation between two estimands and not a theory of what makes life go well. It takes whatever substantive account of well-being a researcher brings and asks how that account should be aggregated over time and over people. It is compatible with hedonic, evaluative and eudaimonic conceptions alike, and it adjudicates between none of them. Second, the framework inherits an unresolved question about the correct time average. A simple mean over a life weights every period equally, which few would defend on reflection, and the alternatives carry substantive commitments about whether later periods, or periods of greater capacity, should count for more. Ergodicity economics has an answer for wealth, where the logarithm follows from the multiplicative dynamics themselves (Peters and Adamou 2022). No comparable derivation is available for well-being, and the analogy should not be pressed further than the argument supports.

14. An Empirical Program

Four studies follow directly from the framework, and three of them could be conducted with existing data.

The first compares within-person and between-person estimates of the same well-being relation across observation windows of increasing length in a long-running household panel. EMW predicts divergence that grows with window length. The prevailing alternative, that discrepancies reflect measurement error, predicts the opposite, since error attenuates with longer averaging.

The second exploits differential retention. Panels differ substantially in how successfully they follow participants who become unemployed, ill or institutionalized. Estimating average well-being and adaptation separately by retention rate, across panels or across waves of a single panel, tests the survivorship mechanism directly. EMW predicts that better retention yields lower averages and weaker adaptation, and the closed form of Section 8.3 states by how much for a given prevalence and attrition share.

The third concerns organizations. Where an employer holds both well-being survey data and administrative records of long-term sickness absence and exit, absorption hazard can be estimated alongside the reported mean. EMW predicts that the two move independently, and that programs raising the mean need not reduce the hazard. Since firms increasingly hold both kinds of record, the test requires linkage rather than new instrumentation.

The fourth is computational. Consider an agent-based implementation of the framework in which heterogeneous agents carry their own drift, volatility, absorption hazard and recovery rate, and in which observation depends on state. Such a model would allow the three mechanisms to be switched on separately, and their joint effect on the observed ensemble average to be traced over realistic horizons. The model of Bruno (2025) already shows that reference-group comparisons in such a population reverse the aggregate income-satisfaction relation. Adding the time axis and absorbing states to that architecture would show how much of

the ergodicity gap survives the heterogeneity that the agent-based tradition regards as the normal case, and it would generate the calibration that Tables 1 and 2 only approximate.

15. Summary

Ergodicity is the condition under which an average over a population equals an average over time. It is a strong condition, and the evidence from several independent research traditions is that human processes do not generally satisfy it. Psychology reached this conclusion by way of non-stationarity and heterogeneity. Economics reached it by way of compounding, and then found that the mean of a heterogeneous population can describe none of its members. Well-being research, which sits between them and has more practical need of the distinction than either, has proceeded largely as though the condition held.

This article has proposed the ergodicity model of well-being as a framework for taking the failure seriously. Its four tenets hold that a life is the unit that matters, that the resources generating well-being compound, that some states are effectively irreversible and tend to remove people from view, and that survivorship therefore makes population well-being a systematically optimistic guide to individual lives. Three closed forms give the claims a determinate shape. Compounding sets the magnitude of the gap, conditional on the curvature of the mapping from resources to reports, and the logarithm is the knife-edge. Absorption with exit from observation sets the sign, and the gap vanishes exactly when everyone stays in view. The framework reinterprets the long argument over hedonic adaptation, the unresolved dispute over income and well-being, the disappointing record of organizational well-being programs, and the recent turn toward well-being inequality. It also identifies a specific institutional pathology, in which measuring well-being as a population average creates a standing incentive to neglect the prevention of irreversible harm.

The practical conclusion is narrow and should not be overstated. Population well-being statistics remain the best instrument available for many purposes, and nothing here suggests abandoning them. What the framework does suggest is that their relation to individual lives is neither automatic nor benign, and that the direction of the discrepancy is predictable. Closing it would require two quantities that are not currently collected: the rate of entry into states people do not leave, and the time spent there. Placing well-being research within the mathematics of ergodicity makes that omission visible and states what it would take to repair.

Statements and Declarations

Competing interests: On behalf of all authors, the corresponding author states that there is no conflict of interest. The author has no relevant financial or non-financial interests to disclose.

Funding: No funding was received for conducting this study.

Ethics approval: This article is a conceptual and theoretical contribution. It does not report studies involving human participants or animals, and no empirical data were collected or analyzed. Ethics committee approval was therefore not required.

Data availability: No datasets were generated or analyzed for this article. The code that produces the figures, the tables and the numerical checks of every closed form is available from the author on request.

References

Adamou A, Berman Y, Mavroyiannis D, Peters O (2021) Microfoundations of discounting. Decision Analysis 18:257–272. https://doi.org/10.1287/deca.2021.0436

Adler A, Seligman MEP (2016) Using wellbeing for public policy: Theory, measurement, and recommendations. International Journal of Wellbeing 6:1–35. https://doi.org/10.5502/ijw.v6i1.429

Aoki M, Yoshikawa H (2012) Non-self-averaging in macroeconomic models: A criticism of modern micro-founded macroeconomics. Journal of Economic Interaction and Coordination 7:1–22. https://doi.org/10.1007/s11403-012-0088-3

Bakker AB, Demerouti E (2017) Job demands–resources theory: Taking stock and looking forward. Journal of Occupational Health Psychology 22:273–285. https://doi.org/10.1037/ocp0000056

Beltz AM, Wright AGC, Sprague BN, Molenaar PCM (2016) Bridging the nomothetic and idiographic approaches to the analysis of clinical data. Assessment 23:447–458. https://doi.org/10.1177/1073191116648209

Benhabib J, Bisin A, Zhu S (2015) The wealth distribution in Bewley economies with capital income risk. Journal of Economic Theory 159:489–515. https://doi.org/10.1016/j.jet.2015.07.013

Berman Y, Peters O, Adamou A (2021) Wealth inequality and the ergodic hypothesis: Evidence from the United States. Journal of Income Distribution. https://doi.org/10.25071/1874-6322.40455

Birkhoff GD (1931) Proof of the ergodic theorem. Proceedings of the National Academy of Sciences 17:656–660. https://doi.org/10.1073/pnas.17.2.656

Blanchflower DG, Oswald AJ (2020) Trends in extreme distress in the United States, 1993–2019. American Journal of Public Health 110:1538–1544. https://doi.org/10.2105/ajph.2020.305811

Bolger N, Davis A, Rafaeli E (2003) Diary methods: Capturing life as it is lived. Annual Review of Psychology 54:579–616. https://doi.org/10.1146/annurev.psych.54.101601.145030

Bouchaud JP (1992) Weak ergodicity breaking and aging in disordered systems. Journal de Physique I 2:1705–1713. https://doi.org/10.1051/jp1:1992238

Bouchaud JP, Farmer REA (2023) Self-fulfilling prophecies, quasi nonergodicity, and wealth inequality. Journal of Political Economy 131:947–993. https://doi.org/10.1086/722214

Bouchaud JP, Mézard M (2000) Wealth condensation in a simple model of economy. Physica A: Statistical Mechanics and its Applications 282:536–545. https://doi.org/10.1016/s0378-4371(00)00205-3

Brickman P, Coates D, Janoff-Bulman R (1978) Lottery winners and accident victims: Is happiness relative? Journal of Personality and Social Psychology 36:917–927. https://doi.org/10.1037/0022-3514.36.8.917

Bruno B (2025) An agent-based exploration of macroeconomic of happiness: New insights on inequality and the role of reference groups. Journal of Economic Interaction and Coordination 20:987–1020. https://doi.org/10.1007/s11403-025-00444-4

Buckley JP, Keil AP, McGrath LJ, Edwards JK (2015) Evolving methods for inference in the presence of healthy worker survivor bias. Epidemiology 26:204–212. https://doi.org/10.1097/ede.0000000000000217

Campbell DT (1979) Assessing the impact of planned social change. Evaluation and Program Planning 2:67–90. https://doi.org/10.1016/0149-7189(79)90048-x

Carver CS, Scheier MF (2002) Control processes and self-organization as complementary principles underlying behavior. Personality and Social Psychology Review 6:304–315. https://doi.org/10.1207/s15327957pspr0604_05

Case A, Deaton A (2015) Rising morbidity and mortality in midlife among white non-Hispanic Americans in the 21st century. Proceedings of the National Academy of Sciences 112:15078–15083. https://doi.org/10.1073/pnas.1518393112

Champernowne DG (1953) A model of income distribution. The Economic Journal 63:318–351. https://doi.org/10.2307/2227127

Clark AE (2018) Four decades of the economics of happiness: Where next? Review of Income and Wealth 64:245–269. https://doi.org/10.1111/roiw.12369

Clark AE, Georgellis Y, Sanfey P (2001) Scarring: The psychological impact of past unemployment. Economica 68:221–241. https://doi.org/10.1111/1468-0335.00243

Clark AE, Diener E, Georgellis Y, Lucas RE (2008) Lags and leads in life satisfaction: A test of the baseline hypothesis. The Economic Journal 118:F222–F243. https://doi.org/10.1111/j.1468-0297.2008.02150.x

Cramer AOJ, van Borkulo CD, Giltay EJ, van der Maas HLJ, Kendler KS, Scheffer M, Borsboom D (2016) Major depression as a complex dynamic system. PLOS ONE 11:e0167490. https://doi.org/10.1371/journal.pone.0167490

De Luca Picione R (2015) The idiographic approach in psychological research. The challenge of overcoming old distinctions without risking to homogenize. Integrative Psychological and Behavioral Science 49:360–370. https://doi.org/10.1007/s12124-015-9307-5

Delli Gatti D, Gaffeo E, Gallegati M (2010) Complex agent-based macroeconomics: A manifesto for a new paradigm. Journal of Economic Interaction and Coordination 5:111–135. https://doi.org/10.1007/s11403-010-0064-8

Demerouti E, Bakker AB, Nachreiner F, Schaufeli WB (2001) The job demands-resources model of burnout. Journal of Applied Psychology 86:499–512. https://doi.org/10.1037/0021-9010.86.3.499

Diener E, Chan MY (2011) Happy people live longer: Subjective well-being contributes to health and longevity. Applied Psychology: Health and Well-Being 3:1–43. https://doi.org/10.1111/j.1758-0854.2010.01045.x

Diener E, Lucas RE, Scollon CN (2006) Beyond the hedonic treadmill: Revising the adaptation theory of well-being. American Psychologist 61:305–314. https://doi.org/10.1037/0003-066x.61.4.305

Diener E, Oishi S, Tay L (2018) Advances in subjective well-being research. Nature Human Behaviour 2:253–260. https://doi.org/10.1038/s41562-018-0307-6

Doctor JN, Wakker PP, Wang TV (2020) Economists' views on the ergodicity problem. Nature Physics 16:1168. https://doi.org/10.1038/s41567-020-01106-x

Easterlin RA (1995) Will raising the incomes of all increase the happiness of all? Journal of Economic Behavior and Organization 27:35–47. https://doi.org/10.1016/0167-2681(95)00003-b

Epskamp S, Fried EI (2018) A tutorial on regularized partial correlation networks. Psychological Methods 23:617–634. https://doi.org/10.1037/met0000167

Espeland WN, Sauder M (2007) Rankings and reactivity: How public measures recreate social worlds. American Journal of Sociology 113:1–40. https://doi.org/10.1086/517897

Fisher AJ, Medaglia JD, Jeronimus BF (2018) Lack of group-to-individual generalizability is a threat to human subjects research. Proceedings of the National Academy of Sciences 115:E6106–E6115. https://doi.org/10.1073/pnas.1711978115

Fleming WJ (2024) Employee well-being outcomes from individual-level mental health interventions: Cross-sectional evidence from the United Kingdom. Industrial Relations Journal 55:162–182. https://doi.org/10.1111/irj.12418

Fredrickson BL (2001) The role of positive emotions in positive psychology: The broaden-and-build theory of positive emotions. American Psychologist 56:218–226. https://doi.org/10.1037//0003-066x.56.3.218

Frijters P, Clark AE, Krekel C, Layard R (2020) A happy choice: Wellbeing as the goal of government. Behavioural Public Policy 4:126–165. https://doi.org/10.1017/bpp.2019.39

Fujita F, Diener E (2005) Life satisfaction set point: Stability and change. Journal of Personality and Social Psychology 88:158–164. https://doi.org/10.1037/0022-3514.88.1.158

Gabaix X (2009) Power laws in economics and finance. Annual Review of Economics 1:255–294. https://doi.org/10.1146/annurev.economics.050708.142940

Gallegati M, Kirman A (2012) Reconstructing economics. Complexity Economics 1:5–31. https://doi.org/10.7564/12-coec2

Goff L, Helliwell JF, Mayraz G (2018) Inequality of subjective well-being as a comprehensive measure of inequality. Economic Inquiry 56:2177–2194. https://doi.org/10.1111/ecin.12582

Guest DE (2017) Human resource management and employee well-being: Towards a new analytic framework. Human Resource Management Journal 27:22–38. https://doi.org/10.1111/1748-8583.12139

Hamaker EL, Kuiper RM, Grasman RPPP (2015) A critique of the cross-lagged panel model. Psychological Methods 20:102–116. https://doi.org/10.1037/a0038889

Headey B (2008) The set-point theory of well-being: Negative results and consequent revisions. Social Indicators Research 85:389–403. https://doi.org/10.1007/s11205-007-9134-2

Heckman JJ (1979) Sample selection bias as a specification error. Econometrica 47:153–161. https://doi.org/10.2307/1912352

Hirsh JB, Mar RA, Peterson JB (2012) Psychological entropy: A framework for understanding uncertainty-related anxiety. Psychological Review 119:304–320. https://doi.org/10.1037/a0026767

Hobfoll SE (1989) Conservation of resources: A new attempt at conceptualizing stress. American Psychologist 44:513–524. https://doi.org/10.1037/0003-066x.44.3.513

Hobfoll SE, Halbesleben J, Neveu JP, Westman M (2018) Conservation of resources in the organizational context: The reality of resources and their consequences. Annual Review of Organizational Psychology and Organizational Behavior 5:103–128. https://doi.org/10.1146/annurev-orgpsych-032117-104640

Hurtado-Parrado C, López-López W (2015) Single-case research methods: History and suitability for a psychological science in need of alternatives. Integrative Psychological and Behavioral Science 49:323–349. https://doi.org/10.1007/s12124-014-9290-2

Joyce S, Modini M, Christensen H, Mykletun A, Bryant R, Mitchell PB, Harvey SB (2016) Workplace interventions for common mental disorders: A systematic meta-review. Psychological Medicine 46:683–697. https://doi.org/10.1017/s0033291715002408

Kahneman D, Deaton A (2010) High income improves evaluation of life but not emotional well-being. Proceedings of the National Academy of Sciences 107:16489–16493. https://doi.org/10.1073/pnas.1011492107

Kahneman D, Tversky A (1979) Prospect theory: An analysis of decision under risk. Econometrica 47:263–291. https://doi.org/10.2307/1914185

Kalmijn W, Veenhoven R (2005) Measuring inequality of happiness in nations: In search for proper statistics. Journal of Happiness Studies 6:357–396. https://doi.org/10.1007/s10902-005-8855-7

Kelly JL (1956) A new interpretation of information rate. Bell System Technical Journal 35:917–926. https://doi.org/10.1002/j.1538-7305.1956.tb03809.x

Killingsworth MA (2021) Experienced well-being rises with income, even above $75,000 per year. Proceedings of the National Academy of Sciences 118:e2016976118. https://doi.org/10.1073/pnas.2016976118

Killingsworth MA, Kahneman D, Mellers B (2023) Income and emotional well-being: A conflict resolved. Proceedings of the National Academy of Sciences 120:e2208661120. https://doi.org/10.1073/pnas.2208661120

Kirman AP (1992) Whom or what does the representative individual represent? Journal of Economic Perspectives 6:117–136. https://doi.org/10.1257/jep.6.2.117

Knabe A, Rätzel S (2011) Scarring or scaring? The psychological impact of past unemployment and future unemployment risk. Economica 78:283–293. https://doi.org/10.1111/j.1468-0335.2009.00816.x

Kuppens P, Oravecz Z, Tuerlinckx F (2010) Feelings change: Accounting for individual differences in the temporal dynamics of affect. Journal of Personality and Social Psychology 99:1042–1060. https://doi.org/10.1037/a0020962

Lucas RE (2007) Adaptation and the set-point model of subjective well-being. Current Directions in Psychological Science 16:75–79. https://doi.org/10.1111/j.1467-8721.2007.00479.x

Lucas RE, Clark AE, Georgellis Y, Diener E (2003) Reexamining adaptation and the set point model of happiness: Reactions to changes in marital status. Journal of Personality and Social Psychology 84:527–539. https://doi.org/10.1037/0022-3514.84.3.527

Lucas RE, Clark AE, Georgellis Y, Diener E (2004) Unemployment alters the set point for life satisfaction. Psychological Science 15:8–13. https://doi.org/10.1111/j.0963-7214.2004.01501002.x

Lugtig P (2014) Panel attrition. Sociological Methods and Research 43:699–723. https://doi.org/10.1177/0049124113520305

Luhmann M, Hofmann W, Eid M, Lucas RE (2012) Subjective well-being and adaptation to life events: A meta-analysis. Journal of Personality and Social Psychology 102:592–615. https://doi.org/10.1037/a0025948

Lyubomirsky S, Sheldon KM, Schkade D (2005) Pursuing happiness: The architecture of sustainable change. Review of General Psychology 9:111–131. https://doi.org/10.1037/1089-2680.9.2.111

Maslach C, Leiter MP (2016) Understanding the burnout experience: Recent research and its implications for psychiatry. World Psychiatry 15:103–111. https://doi.org/10.1002/wps.20311

Meder D, Rabe F, Morville T, Madsen KH, Koudahl MT, Dolan RJ, Siebner HR, Hulme OJ (2021) Ergodicity-breaking reveals time optimal decision making in humans. PLOS Computational Biology 17:e1009217. https://doi.org/10.1371/journal.pcbi.1009217

Metzler R, Jeon JH, Cherstvy AG, Barkai E (2014) Anomalous diffusion models and their properties: Non-stationarity, non-ergodicity, and ageing at the centenary of single particle tracking. Physical Chemistry Chemical Physics 16:24128–24164. https://doi.org/10.1039/c4cp03465a

Molenaar PCM (2004) A manifesto on psychology as idiographic science: Bringing the person back into scientific psychology, this time forever. Measurement: Interdisciplinary Research and Perspective 2:201–218. https://doi.org/10.1207/s15366359mea0204_1

Molenaar PCM (2007) Psychological methodology will change profoundly due to the necessity to focus on intra-individual variation. Integrative Psychological and Behavioral Science 41:35–40. https://doi.org/10.1007/s12124-007-9011-1

Molenaar PCM, Campbell CG (2009) The new person-specific paradigm in psychology. Current Directions in Psychological Science 18:112–117. https://doi.org/10.1111/j.1467-8721.2009.01619.x

Naimi AI, Richardson DB, Cole SR (2013) Causal inference in occupational epidemiology: Accounting for the healthy worker effect by using structural nested models. American Journal of Epidemiology 178:1681–1686. https://doi.org/10.1093/aje/kwt215

Nelson B, McGorry PD, Wichers M, Wigman JTW, Hartmann JA (2017) Moving from static to dynamic models of the onset of mental disorder. JAMA Psychiatry 74:528–534. https://doi.org/10.1001/jamapsychiatry.2017.0001

Nielsen K, Miraglia M (2017) What works for whom in which circumstances? On the need to move beyond the ‘what works?’ question in organizational intervention research. Human Relations 70:40–62. https://doi.org/10.1177/0018726716670226

Oswald AJ, Powdthavee N (2008) Does happiness adapt? A longitudinal study of disability with implications for economists and judges. Journal of Public Economics 92:1061–1077. https://doi.org/10.1016/j.jpubeco.2008.01.002

Peccei R, Van De Voorde K (2019) Human resource management–well-being–performance research revisited: Past, present, and future. Human Resource Management Journal 29:539–563. https://doi.org/10.1111/1748-8583.12254

Peters O (2011) The time resolution of the St Petersburg paradox. Philosophical Transactions of the Royal Society A 369:4913–4931. https://doi.org/10.1098/rsta.2011.0065

Peters O (2019) The ergodicity problem in economics. Nature Physics 15:1216–1221. https://doi.org/10.1038/s41567-019-0732-0

Peters O, Adamou A (2022) The ergodicity solution of the cooperation puzzle. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 380:20200425. https://doi.org/10.1098/rsta.2020.0425

Peters O, Gell-Mann M (2016) Evaluating gambles using dynamics. Chaos: An Interdisciplinary Journal of Nonlinear Science 26:023103. https://doi.org/10.1063/1.4940236

Peters O, Klein W (2013) Ergodicity breaking in geometric Brownian motion. Physical Review Letters 110:100603. https://doi.org/10.1103/physrevlett.110.100603

Pincus JD (2024) Well-being as need fulfillment: Implications for theory, methods, and practice. Integrative Psychological and Behavioral Science 58:1541–1579. https://doi.org/10.1007/s12124-023-09758-z

Richardson KM, Rothstein HR (2008) Effects of occupational stress management intervention programs: A meta-analysis. Journal of Occupational Health Psychology 13:69–93. https://doi.org/10.1037/1076-8998.13.1.69

Ryff CD (1989) Happiness is everything, or is it? Explorations on the meaning of psychological well-being. Journal of Personality and Social Psychology 57:1069–1081. https://doi.org/10.1037/0022-3514.57.6.1069

Scheffer M, Bascompte J, Brock WA, Brovkin V, Carpenter SR, Dakos V, Held H, van Nes EH, Rietkerk M, Sugihara G (2009) Early-warning signals for critical transitions. Nature 461:53–59. https://doi.org/10.1038/nature08227

Skjold B, Steinkamp SR, Connaughton C, Hulme OJ, Peters O (2026) Ergodicity transformations predict human decision-making under risk. PLOS Computational Biology 22:e1014409. https://doi.org/10.1371/journal.pcbi.1014409

Sonnentag S, Venz L, Casper A (2017) Advances in recovery research: What have we learned? What should be done next? Journal of Occupational Health Psychology 22:365–380. https://doi.org/10.1037/ocp0000079

Speelman CP, McGann M (2013) How mean is the mean? Frontiers in Psychology 4:451. https://doi.org/10.3389/fpsyg.2013.00451

Steptoe A, Deaton A, Stone AA (2015) Subjective wellbeing, health, and ageing. The Lancet 385:640–648. https://doi.org/10.1016/s0140-6736(13)61489-0

Stevenson B, Wolfers J (2008) Economic growth and subjective well-being: Reassessing the Easterlin paradox. Brookings Papers on Economic Activity 2008:1–87. https://doi.org/10.1353/eca.0.0001

Trapp S, Németh D, Janacsek K (2026) Learning under uncertainty: Predictive processing as an integrative framework for educational research. Integrative Psychological and Behavioral Science 60:55. https://doi.org/10.1007/s12124-026-10019-y

van de Leemput IA, Wichers M, Cramer AOJ, Borsboom D, Tuerlinckx F, Kuppens P, van Nes EH, Viechtbauer W, Giltay EJ, Aggen SH, Derom C, Jacobs N, Kendler KS, van der Maas HLJ, Neale MC, Peeters F, Thiery E, Zachar P, Scheffer M (2014) Critical slowing down as early warning for the onset and termination of depression. Proceedings of the National Academy of Sciences 111:87–92. https://doi.org/10.1073/pnas.1312114110

VanderWeele TJ (2017) On the promotion of human flourishing. Proceedings of the National Academy of Sciences 114:8148–8156. https://doi.org/10.1073/pnas.1702996114

Vanhoyweghen A, Verbeken B, Macharis C, Ginis V (2022) The influence of ergodicity on risk affinity of timed and non-timed respondents. Scientific Reports 12:3744. https://doi.org/10.1038/s41598-022-07613-6

von Eye A, Wiedermann W, Mun EY (2024) Log-linear and configural analysis of intra-individual time series under consideration of serial dependence. Integrative Psychological and Behavioral Science 58:759–770. https://doi.org/10.1007/s12124-023-09815-7

von Fircks E (2025) From nomothetic and idiographic to synthography. Integrative Psychological and Behavioral Science 59:89. https://doi.org/10.1007/s12124-025-09944-1

von Neumann J (1932) Proof of the quasi-ergodic hypothesis. Proceedings of the National Academy of Sciences 18:70–82. https://doi.org/10.1073/pnas.18.1.70

Watson N, Wooden M (2009) Identifying factors affecting longitudinal survey response. In: Lynn P (ed) Methodology of longitudinal surveys. Wiley, Chichester, pp 157–181. https://doi.org/10.1002/9780470743874.ch10

Wichers M (2014) The dynamic nature of depression: A new micro-level perspective of mental disorder that meets current challenges. Psychological Medicine 44:1349–1360. https://doi.org/10.1017/s0033291713001979

Zittoun T (2017) Modalities of generalization through single case studies. Integrative Psychological and Behavioral Science 51:171–194. https://doi.org/10.1007/s12124-016-9367-1