Composite virulence: useful metric or conceptual trap?

Luís M. Silva [1,2*]

Tiago G. Zeferino [1,3]

[1] Institute of Biology, University of Neuchâtel, Switzerland.

[2] Department of Zoology, The University of British Columbia, Vancouver, Canada.

[3] Max Planck Institute for Chemical Ecology, Jena, Germany.

* Corresponding author: luis.silva@ubc.ca (LMS)

ORCID

LMS: 0000-0002-8203-4006

TGZ: 0009-0004-0767-5200

Abstract

Virulence, *the harm an infection causes to its host,* is a cornerstone concept in ecology and evolution, yet it remains difficult to quantify because infection impact is multidimensional, dynamic, and context-dependent. Infections can reduce host performance through multiple, partially redundant routes (including mortality, fecundity loss, behavioural impairment, and physiological disruption), whose relative importance depends on the host's life history and ecological conditions. Composite virulence measures have emerged to address this complexity by combining information across traits, but they can embed strong assumptions: component choice may reflect availability rather than biological relevance, correlated traits can be double-counted, and combining parasite drivers or host-state mediators with life-history outcomes can obscure interpretation and introduce circularity. Here we synthesize the conceptual foundations of composite virulence and outline a biologically grounded workflow for building interpretable composites. We treat composite virulence as a measurement strategy anchored to an explicit biological question, with trait selection conducted to minimize redundancy while preserving timing and context-dependency. Used carefully, composite approaches can clarify which routes to harm dominate across hosts, parasites, and environments and thereby strengthen evolutionary inference; used casually, they risk becoming a conceptual trap that yields clean rankings while concealing mechanisms.

1. The composite virulence dilemma

1.1. Virulence as a core concept

Parasites, ranging from viruses to helminths, are infectious agents that live in or on a host using its resources to survive and reproduce [1]. Infection can impose costs on the host, and a key way to describe these *costs* is through virulence, defined here as *the harm inflicted by infection on the host* [2,3] **(Box 1)**. Virulence is central to evolutionary ecology, public health, and livestock management [4–7], as it is used to characterize host-parasite interactions [8–12], track evolutionary change [13–16], assess risks from circulating and emerging infections [17–20], prevent future disease outbreaks [21,22], and evaluate interventions that mitigate disease severity and spread [23–27].

Yet, despite its importance and intuitive appeal, virulence remains difficult to quantify because it is neither a single trait nor a fixed property of a parasite [2,9]. Rather, virulence is an emergent outcome produced by interactions among host and parasite genetics and their shared environment (a classic genotype × genotype × environment interaction) [2,28]. Infection can reduce host fitness through multiple routes, including increased mortality, reduced fecundity (or even sterility), delayed development, reduced condition or altered motility and behaviour [1]. These costs to the host can unfold on different timescales and in different ecological contexts [2,28]. As a result, authors have used several measures of virulence to address the impact of infection on the host. The most appropriate virulence metric will always depend on the biological question being asked and on which components of harm (e.g., fitness measures such as survival or reproduction; conditions, such as resources available or development) are most limiting in the focal system.

This multidimensionality creates a measurement problem. Single proxies are attractive because they are often comparable and straightforward, but they can misrepresent overall costs when harm is expressed primarily through other host fitness and health measures. Mortality, for example, is unambiguous [29–31] and directly captures survival, but it can be a late-stage outcome and insensitive to sublethal effects such as sterilization, delayed or impaired reproduction and physiology [14,32–34]. Physiological markers can be misleading when used as proxies for virulence because changes in host state are not inherently costly [35]. Infection can elevate markers such as oxidative stress or inflammation without detectable effects on survival or reproduction, meaning that “physiological virulence” may overestimate harm if fitness consequences are not demonstrated. Similar "disease severity" markers can reflect pathology, adaptive host reallocation, or context-dependent trade-offs. Hence, infection-associated harm spans across *drivers* (parasite growth and damage mechanisms), host-state *mediators* (pathology, inflammation, dysbiosis, sickness behaviours), and fitness *outcomes* (survival and reproduction). These layers may change together, but they do not always align. As a result, they capture different parts of the infection process and should not be treated as interchangeable measures of virulence.

Virulence is a conceptual construct, so it is typically assessed through proxies [2]. The most common

measures are mortality and fecundity, but the best proxy depends on the question and the system, and authors often use more specific terms to clarify exactly what is being measured. **Intrinsic virulence** refers to the harm caused by a single parasite genotype when measured in a standardized host background and specific controlled environment (e.g., a host survival curve, fecundity loss, or pathology score). By design, it minimizes host heterogeneity and environmental variation and is therefore widely used for comparative purposes within and between laboratories [11,27,36]. **Realized virulence**, in contrast, refers to the total cost a parasite imposes on its host in a particular ecological and evolutionary context. It explicitly accommodates genotype-by-genotype-by-environment interactions, variation in exposure and host condition, and other sources of context dependence, and is therefore useful for linking virulence to epidemiological dynamics and selection in the wild [37,38]. **Proximate virulence** is often used in microbial, clinical and applied settings to describe mechanisms and intermediate processes that contribute to harm (e.g., toxins, inflammatory markers, tissue damage). These measures can provide rapid mechanistic insight, but they are not equivalent to downstream outcomes unless their consequences are explicitly demonstrated in the focal context [39,40].

Taken together, these terms clarify which facet of infection impact is being emphasized: controlled comparisons of parasite genotypes (intrinsic), context-dependent costs in real ecological settings (realized), or intermediate mechanisms and host-state shifts (proximate). In practice, each is often quantified using a single proxy (or a small number of proxies), which can be useful but incomplete: infection-induced harm is typically expressed across multiple traits, unfolds over time, and may first manifest as changes in host state before translating into survival or reproduction. This sets up a central measurement challenge: how should we represent multidimensional harm in a way that remains interpretable for the question at hand?

1.2. Multiple routes to harm

A second challenge is that the same endpoint infection measure can be produced by different biological processes. For instance, similar reductions in survival or reproduction may arise from parasite growth and host (resource) exploitation, but also from per-parasite pathogenicity (PPP): mechanisms that generate harm at a given parasite burden through cell and tissue injury, inflammatory collateral damage, and toxin- or effector-mediated dysfunction [9,41–44]. Mechanistically, infection can affect hosts through at least three partially separable processes: (i) host exploitation, whereby parasites divert resources and damage tissues through replication and nutrient sequestration [9,41,45]; (ii) direct injury, in which invasion and growth cause cell lysis, barrier disruption, vascular obstruction, or organ dysfunction even at similar parasite loads [43,44,46–48]; and (iii) host-mediated damage, whereby immune activation generates immunopathology such as oxidative stress or cytokine-driven tissue damage [49,50]. These processes can covary, but they do not always, and they can differ in timing, producing diverse disease trajectories and trade-offs.

These distinctions matter for interpretation and for evolution. In malaria infections, for example, *Plasmodium spp.* increase harm not only through replication but also through erythrocyte destruction, cytoadherence-mediated microvascular obstruction, and heme-driven oxidative damage, which amplifies inflammation and tissue dysfunction [51,52]. Consequently, two infections can show similar parasite densities yet differ markedly in anaemia, inflammatory pathology, and risk of severe disease because harm reflects different combinations of exploitation, direct injury, and host-mediated damage. However, the cause of the outcome does matter: mortality alone collapses these processes into a single endpoint measurement, obscuring both the costs of infection and the mechanistic routes through which parasites inflict harm, and potentially evolve it [14,53].

1.3. Why composites emerged: a brief history of trait integration

Historically, virulence measurement expanded whenever a single proxy proved too narrow to capture the impact of infection. In clinical and veterinary settings, "severity" scores have long combined multiple signs and symptoms into standardized summaries, reflecting the practical reality that disease burden manifests across several dimensions [14,15,45,54–56]. Plant pathology offers particularly influential precedents because it formalized multidimensional disease impact into tractable summaries [57–59]. Disease impact has long been quantified by combining disease incidence (how many hosts or tissues are infected) and disease severity (how much damage occurs per infected unit), explicitly recognizing that frequency and intensity capture distinct dimensions of harm [58,60]. Likewise, longitudinal disease pressure is frequently summarized using the area under the disease progress curve (AUDPC/AUC), which integrates repeated measurements of disease severity over time into a single estimate of cumulative burden [61]. These approaches underscore two general principles: infection impact is inherently multifaceted, and collapsing multiple measurements into a summary is a modelling choice with interpretive consequences.

Evolutionary ecology adopted multivariate measurement for related reasons, but with emphasis on life-history outcomes and the pathways linking them to parasite traits. In plant–virus systems, for instance, tolerance is quantified by tracking infection-induced changes across multiple traits (*e.g.*, growth and reproduction), revealing that genotypes can differ in how costs are distributed across life-history components rather than in total damage alone [62,63]. Analogous logic motivates the routine use of trait panels in animal and microbial systems, where infection effects on survival and reproduction are often analyzed alongside host-state mediators (*e.g.*, pathology, oxidative stress, behaviour) and parasite traits (*e.g.*, burden, growth, persistence) [14,41,53,64,65]. The rationale is straightforward: mortality is a downstream outcome of infection, whereas selection acts on parasite traits through their effects on host survival, reproduction, and transmission opportunity. In this context, multivariate measurement is not a statistical flourish; it is often the most direct way to represent the biology of infection trajectories, trade-offs, and context dependence.

When multiple traits are compressed into a single index, a series of assumptions are inevitably made: which traits to include, how to weight them, whether to treat them as independent or interchangeable, and which biological level the index is intended to represent (*e.g.,* health, fitness, or mechanism). Once collapsed into one number, however, these choices become largely invisible to the reader. The resulting index can appear objective and comparable across studies, while in fact embedding unspoken assumptions about what virulence represents, which traits matter most, and how different processes relate to one another. In this sense, a composite index does not merely summarise data; it also implicitly packages a theoretical view of virulence.

As composite indices embed assumptions about trait hierarchy, weighting, and timing, two studies analyzing the same infection system could reach different conclusions about which parasite genotype is "more virulent" simply because they constructed different composites. In this sense, composite virulence does not merely summarise biology; it can silently reshape evolutionary inference.

1.4. The dilemma: composites clarify, but can also conceal

Crucially, these assumptions extend to time. The same trait can have very different consequences depending on when it is expressed or measured during infection, ranging from no detectable fitness cost to substantial harm. Early physiological or behavioural changes may be transient or compensated, whereas similar changes later in infection can translate directly into survival or reproductive costs. As a result, combining traits without specifying how timing is handled implicitly assumes that effects are constant through infection. Composite measures, therefore, either require integrating trait values across the infection trajectory or making an explicit choice about which phase of infection a given trait is intended to represent.

As multivariate trait panels became more common, a further step followed naturally: multiple measurements were increasingly compressed into a single “composite virulence” score, often via additive indices. Such summaries can be useful. They acknowledge multidimensional effects of infection, reduce the burden of multiple comparisons, and facilitate ranking or comparing treatments such as parasite genotypes, host genotypes, or environments. However, compression also increases the risk that a composite becomes a convenient label rather than an interpretable representation.

The core concern is not whether summarising is “allowed”, but whether the resulting core corresponds to a clearly defined facet of infection impact. Composite scores often embed strong assumptions that are rarely stated. First, they assume that the included components describe the same construct, even when they mix life-history outcomes with host-state mediators or mechanistic measures. Second, they assume that correlated traits contribute independent information, even when traits share variance because they reflect the same underlying process or because one constrains another. They also assume that traits are comparable in scale, such that differences in numerical range do not determine their influence on the composite. In practice, traits with large ranges can dominate an index, while

biologically important traits with narrower ranges are effectively down-weighted or obscured. Third, they assume that timing can be ignored or safely collapsed, even though infection effects frequently unfold dynamically (*e.g.*, early sterilization *versus* late mortality) and the order of effects can matter as much as their magnitude. Fourth, they treat context dependence as nuisance variation, even though environmental conditions and host state can reorder which traits are most limiting and can alter which routes to harm dominate.

This broader concern is not limited to composite scores. Recent work by Surasinghe & Ogbunugafor (2026) using time-series causal discovery has shown that the relationship between virulence and transmission can reverse direction or weaken entirely depending on ecological scale and context, undermining static or univariate interpretations of trait associations [66]. Such findings reinforce the idea that trait relationships are dynamic and context-dependent rather than fixed properties of pathogens. Composite measures that ignore this structure risk summarizing patterns that are themselves emergent outcomes of shifting causal architectures.

These assumptions have practical consequences. Different scaling and weighting choices can yield different composite virulence rankings, especially when traits are correlated, when effects occur on different timescales, or when environments reorder trait importance. In the worst case, a composite can look clean while obscuring the biology: it can double-count the same information, merge mechanistic drivers with downstream outcomes, or conceal shifts in infection trajectory that would change evolutionary interpretation.

This paper addresses that dilemma. We argue that composite virulence should be treated as a model rather than a statistical black box: a deliberate measurement strategy anchored to a clearly defined representation of infection-induced harm, with explicit assumptions about: (i) which facet(s) of infection impact are being combined, (ii) redundancy and dependence among components, (iii) infection dynamics and timing, and (iv) ecological and evolutionary context. Treating composites in this way does not preclude single-number summaries; instead, it makes them interpretable and testable. In the sections that follow, we outline common pitfalls and provide a biologically grounded workflow for constructing composite measures that retain meaning for evolutionary inference.

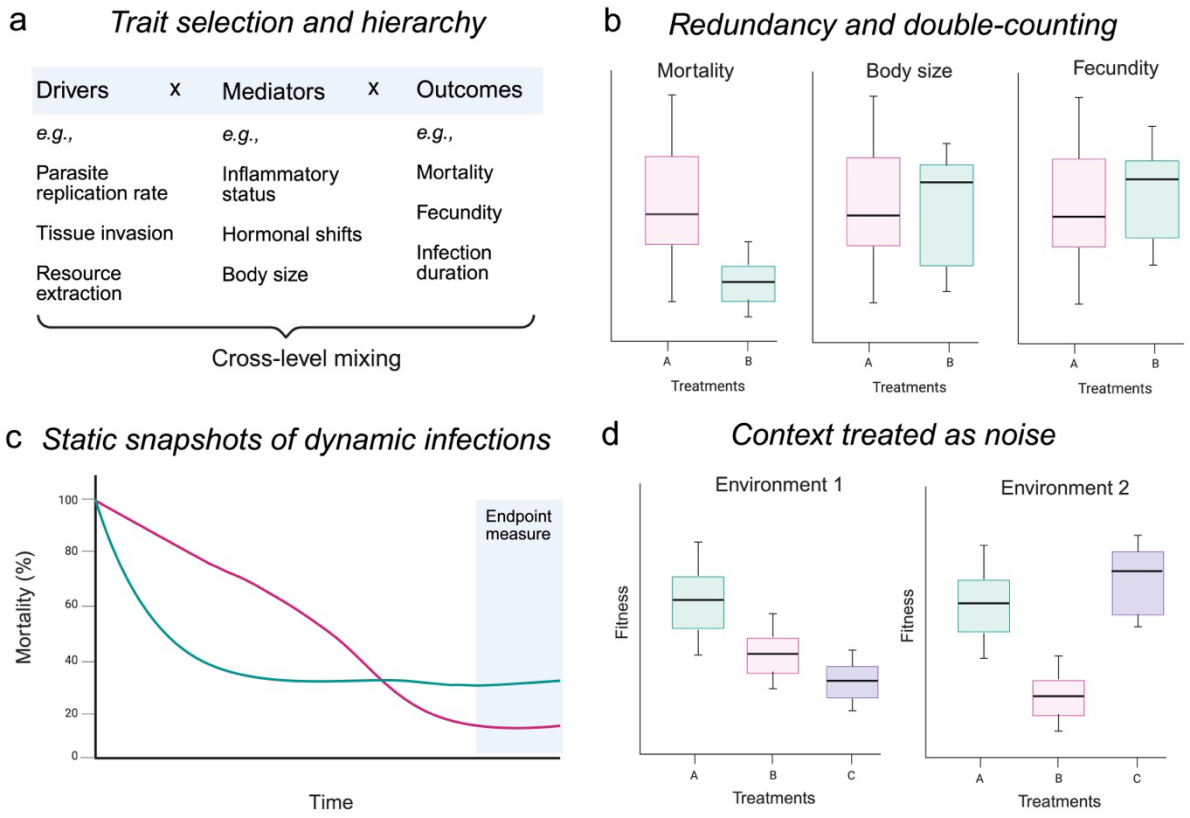

Figure 1. Common pitfalls of composite virulence metrics. (a) Trait selection and hierarchy. Composite indices often combine variables that operate at different biological levels: parasite drivers (*e.g.*, replication rate, tissue invasion, resource extraction), host-state mediators (*e.g.*, inflammatory status, hormonal shifts, body size), and fitness outcomes (*e.g.*, mortality, fecundity, infection duration). Mixing these levels obscures causal structure and conflates mechanism with consequence. **(b)** Redundancy and double-counting. Correlated traits are frequently aggregated without accounting for shared variance. When mortality, body size, and fecundity covary, combining them can overweight a single biological process and artificially inflate differences among treatments. **(c)** Static snapshots of dynamic infections. Virulence unfolds over time, yet composite measures often rely on endpoint values. Distinct infection trajectories can yield similar final values, masking differences in timing, peak severity, or recovery dynamics that are relevant for selection. **(d)** Context treated as noise. Environmental structure reshapes trait expression and treatment rankings. When ecological context is ignored or averaged across, biologically meaningful genotype–environment interactions are misinterpreted as residual variation rather than determinants of virulence expression. Together, these pitfalls illustrate how composite virulence indices can generate clean numerical rankings while concealing causal structure, temporal dynamics, and ecological contingency.

2. Pitfalls and a biologically grounded workflow for composite virulence

Composite metrics are motivated by a biological reality: infection affects multiple traits in a dynamic, context-dependent manner. The risk is that composites are often built bottom-up, pooling traits that happen to be available rather than traits selected to represent a defined facet of infection impact in a particular host–parasite system. When the relationships among traits are not considered, such

aggregation can mask effects or generate artefactual null results, for example, when components respond in opposite directions. This can make results look clean while concealing strong assumptions about what the composite represents, which components carry unique information or obscure signals when aggregated, and how infection trajectories and contexts shape downstream costs.

In this paper, we treat **composite virulence** as a measurement strategy rather than a new concept: *a pre-specified, multi-trait representation that combines at least two non-redundant measures of infection impact within a stated construct* (*e.g.,* life-history costs, disease burden, or mechanistic severity). The workflow below makes the assumptions behind any composite explicit and therefore testable **(Box 2)**.

2.1. Trait selection and causal hierarchy

Composite metrics often become ambiguous because they combine fundamentally different variables drawn from different levels of the infection process **(Fig. 1a)**. This ambiguity arises both when variables are combined across layers and within a single layer. Mixing variables from different levels can obscure interpretation because they play distinct causal roles, but pooling multiple variables from the same layer can also be uninformative when components are redundant or capture opposing effects. This issue is particularly well recognized for fitness-related outcomes, where aggregation can mask trade-offs among survival, reproduction, or related measures. Parasite burden, replication rate, or persistence typically act as drivers. Pathology, immune activation, physiological disruption, and dysbiosis often function as host-state mediators. Survival or reproduction are fitness or life-history outcomes. When a single score collapses these layers, interpretation becomes unclear: a high composite virulence could reflect higher parasite growth, stronger immunopathology, higher PPP, or simply differences in the timing of trait measurement. It can also become circular because a mediator included within the composite cannot be evaluated as an explanatory process, since, by definition, it has already been counted as part of harm.

A more robust approach begins by pre-defining what the composite is meant to represent. In evolutionary ecology, the default targets are often fitness and life-history costs because selection acts through survival and reproduction. Lifetime reproductive success (LRS) would be the ideal measure, but it is rarely tractable; survival and partial measures of reproduction are therefore common and defensible outcomes because they contribute directly to fitness. In other settings, the target of the composite may instead be overall disease burden or severity, or a population-level impact; such composites are legitimate, but they require different component choices and support different claims.

Once the composite is specified, candidate traits should be organized into a simple pathway that distinguishes which traits define the composite from those that explain its variation. For a host-fitness composite, this typically means building the composite from survival and reproduction alone, while treating burden- and mechanistic-related traits as parallel descriptors or predictors analyzed alongside the composite to interpret why outcomes differ. For a disease-burden composite, multiple host-state mediators may be combined within the composite itself, whereas drivers and life-history outcomes are

kept external and analyzed in parallel. The key principle is that a composite should be anchored to a single conceptual level, with traits from other levels used for explanation rather than combined into the same index.

2.2. Redundancy and double-counting

Even when the biological target is clear, composites can exaggerate differences by treating correlated measures as independent evidence of harm. In multi-trait panels, redundancy is common: multiple symptoms and physiological readouts can track shared sickness (*e.g.*, inflammatory markers, body condition, tissue damage), and life-history traits can covary through shared constraints. Additive indices built from many overlapping readouts effectively double-count the same biology **(Fig. 1b)**. Data-driven reductions can also mislead when the dominant axis reflects whichever trait varies most, rather than whichever trait best represents the chosen biological target.

A pragmatic way forward is to integrate sparsely and transparently. One useful principle is to group measurements into a small number of biologically meaningful categories of infection impact (*e.g.,* survival costs, reproductive costs, developmental delays, and performance impairment). Each category is then represented by a single defensible summary rather than by many overlapping readouts. This preserves multidimensionality while limiting double-counting, and it makes clear which facet of infection impact each component is intended to capture.

Where life-history outcomes are available, redundancy can be evaluated directly by asking whether a proposed category explains unique variation in survival or reproduction after accounting for other categories and, where appropriate, parasite burden. In practice, this provides a concrete criterion for deciding whether an additional trait adds a genuinely new dimension of infection impact or merely re-measures information already captured.

Classical plant pathology provides an intuitive parallel: incidence and severity are routinely distinguished because high incidence with mild lesions can yield losses comparable to low incidence with severe lesions [57–59]. When these components are combined, the resulting index is only interpretable if the weights reflect how each component contributes to the loss in the relevant context. Similar issues arise in animal systems. For example, parasites that castrate hosts can impose profound reproductive costs with little effect on survival [32,33]; composites that heavily weight mortality will therefore underestimate the impact of infection in such systems. A single-number index can still be useful, but it should be presented as a transparent summary of the category and accompanied by robustness checks that show the conclusions are not artifacts of scaling or weighting choices, especially when traits are correlated.

A concrete example of why redundancy matters arises in cases where different components of fitness outcome yield different virulence rankings. In a natural host–parasite system, Silva and Koella (2025)

found that parasites selected for late transmission increased mortality-inferred virulence in some host stages, yet conclusions changed when virulence was inferred from fecundity [14]. This is precisely where dense composites can mislead: if several correlated performance readouts (*e.g.*, activity, growth, condition) are all included alongside reproduction, one route to harm can be effectively counted multiple times, inflating its apparent contribution. Category-based summaries reduce this risk by collapsing correlated measures within the same level of the infection process and then testing whether each category provides distinct information about life-history outcomes once other categories and parasite burden are accounted for.

2.3. Static snapshots of dynamic infections

A second common pitfall is treating virulence as static **(Fig. 1c)**. Many composite indices are derived from single time points or endpoints, yet infections are trajectories: parasite growth, host responses, and damage profiles shift through time, often nonlinearly [14,41,67]. Hosts and parasites can therefore differ not only in the magnitude of infection [68], but also in its schedule [69]. This matters because selection operates through life-history windows [14,53,70]. When reproduction mostly happens early in infection, costs expressed later can become effectively invisible to selection, whereas delayed reproduction or transmission can make those same late costs decisive. Two infections can show similar endpoint pathology or similar mortality risk yet impose different cumulative costs because effects are concentrated early *versus* late, or because hosts partially recover.

Dynamics also change the mechanistic interpretation. The same host-state mediator can be protective early and damaging later, and traits that appear minor at one time point can be decisive when evaluated over the window that matters most. Conversely, end-stage pathology can be dramatic yet evolutionarily marginal if it occurs after most reproduction has already occurred. Collapsing time can therefore erase precisely the variation that distinguishes parasite strategies and host responses.

To keep composites interpretable, time should be treated as part of the phenotype rather than a nuisance. When feasible, longitudinal outcomes provide the cleanest summaries: survival curves, age-specific fecundity, or repeated performance measures can be summarised as cumulative reproduction, hazards within a defined window, or AUC-style measures. When full tracking is infeasible, a window-based approach often captures most of the signal: define biologically justified phases of infection (e.g., establishment, peak, late-stage/recovery) and report phase-specific summaries rather than timeless values. A minimal but powerful addition is to report timing features (onset, peak, recovery, time-to-threshold) alongside magnitude, because timing often differentiates strategies and strongly shapes evolutionary consequences.

The core principle is alignment: the time window used to quantify a composite should match the window that matters for the host's life history and the parasite's transmission schedule. Without that alignment, "overall virulence" can become a property of the sampling design rather than the infection

process.

2.4. Context treated as noise

Composite virulence is often presented as an intrinsic property of a host–parasite pair (*i.e.*, a genotype × genotype interaction), measured under a single set of laboratory conditions and implicitly assumed to generalize. Yet virulence is routinely context-dependent (*i.e.*, genotype x genotype x environment) **(Fig. 1d)** [2,71–73]. Host genotype, age, nutrition, temperature, microbiome state, density, co-infection, and exposure route can each change not only the magnitude of infection impact, but also which dimension dominates, or even whether a given trait has any detectable effect at all. In practice, context can re-order rankings because it shifts which component is limiting: the same infection may primarily reduce reproduction under nutritional stress, primarily reduce survival at high temperature, or primarily alter performance when transmission depends on contact and movement [74,75]. Importantly, this context-dependence generates structured, often time-dependent variation in infection impact rather than random noise, with different traits becoming limiting at different stages or under different conditions. In coevolving host-parasite systems, this variation is expected, as virulence often reflects the phenotypic outcome of reciprocal adaptation under specific ecological constraints. Treating this dependence as nuisance variation risks building composites that are precise but non-transferable.

A path analysis (or structured equation modelling) helps because it makes context-dependent interpretations rather than merely descriptive [76–78]. Contexts often act on specific points of the infection process: temperature may amplify parasite replication (drivers), nutrition may alter tolerance and repair capacity (host-state mediators), microbiome state may shift immune activation and dysbiosis (host-state mediators), and exposure route may change which behavioural traits influence transmission opportunity (outcomes relevant to epidemiological impact) [28,79]. Under this view, context dependence is not random variation around a single virulence value; it is a structured variation in how infection translates into impact across time and ecological constraints, and therefore a key to interpreting when and how costs emerge.

Where feasible, the most informative representation is a reaction norm rather than a single number, because it makes context dependence explicit [80]. Rather than focusing on a single estimate, composites (or their component categories) can be evaluated across ecologically meaningful contexts, allowing comparisons of slopes and means. Even modest factorial designs can be sufficient to reveal whether differences are robust (parallel norms) or context specific (crossing norms), and whether variation arises primarily from main effects (host or parasite) *versus* interactions (host × parasite, environment × parasite, or host × parasite × environment). When full crossing is impractical, measuring key contextual variables and modelling them explicitly may be preferable to pooling across them.

The guiding principle is that context should be treated as part of the virulence phenotype. Composites that ignore context risk averaging away structured, time-dependent variation in infection impact that

drives transmission and selection, whereas context-aware composites can preserve this variation and reveal when and why particular routes to life-history costs dominate.

3. Conclusions

Composite measures of virulence arise from a genuine biological challenge: infection impact is multidimensional, dynamic, and context-dependent. No single proxy can capture all relevant facets of harm, and multivariate measurements are often necessary. The problem is therefore not whether composites should be used, but how they are constructed and interpreted.

We argue that the value of a composite depends on whether it is anchored to a clearly specified biological target and whether its components are chosen to represent that target rather than pooled by convenience. When this distinction is ignored, composites can obscure biology by mixing drivers, host-state mediators, and life-history outcomes, double-counting correlated traits, collapsing dynamic trajectories, or averaging across contexts that reorder which routes to harm dominate. In such cases, a single number may appear clean while concealing the mechanisms and trade-offs that matter for evolutionary and ecological inference.

Treating composite virulence as a measurement strategy rather than a new definition resolves much of this tension. Composites can legitimately summarise different facets of infection impact (life-history costs, disease burden, mechanistic severity, or epidemiological impact), but they cannot do so simultaneously without loss of interpretability. Making the biological target explicit, organizing traits by their role in the infection process, and preserving information on timing and context allow composites to remain informative rather than reductive.

This perspective also clarifies the relationship between evolutionary and applied uses of virulence metrics. In clinical and applied settings, composites that summarise proximate severity can be powerful tools for diagnosis, triage, and intervention. In evolutionary ecology, composites that integrate effects on survival and reproduction can illuminate how parasites and hosts are shaped by selection. These uses are complementary, not competing, provided the target of inference is stated, and the limits of each summary are acknowledged.

More broadly, the framework outlined here reframes composite virulence as an opportunity rather than a liability. When built transparently, composites can expose redundancy, reveal trade-offs among fitness components, and highlight when and why infection impact changes over time or across contexts. When built opaquely, they risk flattening complex biology into rankings that are difficult to interpret or generalize.

As infectious disease research increasingly integrates multi-trait, longitudinal, and context-rich data,

the pressure to summarise will only grow. The challenge is therefore not to resist composites, but to ensure that they remain biologically grounded representations of infection impact rather than analytical conveniences. Doing so will strengthen inference across systems and help align virulence measurement with the ecological and evolutionary questions it is meant to address. Without explicit biological anchoring, virulence risks drifting from an evolutionary parameter grounded in life-history costs to a numerical label whose meaning depends on undocumented modelling choices.

Acknowledgments

We thank Marine Mamin and Tommy J. Travers-Cook for feedback on earlier drafts of the manuscript.

Funding

LMS and TGZ were supported by the SNSF grant 310030_192786. Additionally, LMS was supported by an NSERC Canada Excellence Research Chair grant awarded to Kayla C. King. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Author contributions

LMS and TGZ conceptualized, formalized, wrote and reviewed the manuscript.

Data availability

No data was produced for this manuscript.

Competing interests

The authors have declared that no competing interests exist.

References

1. Schmid-Hempel, P. *Evolutionary parasitology: the integrated study of infections, immunology, ecology, and genetics*. (Oxford University Press, 2021).
2. Silva, L. M. & King, K. C. Ecology and evolution of virulence. *Curr. Biol.* (2025). doi:http://doi.org/10.1016/j.cub.2025.08.029
3. Read, A. F. The evolution of virulence. *Trends Microbiol.* **2**, 73–76 (1994).

4. Thomas, S. R. & Elkinton, J. S. Pathogenicity and virulence. *J. Invertebr. Pathol.* **85**, 146–151 (2004).

5. Day, T. On the evolution of virulence and the relationship between various measures of mortality. *Proc. R. Soc. London. Ser. B Biol. Sci.* **269**, 1317–1323 (2002).

6. Bull, J. J. & Lauring, A. S. Theory and empiricism in virulence evolution. *PLoS Pathog.* **10**, e1004387 (2014).

7. Casadevall, A. & Pirofski, L. Host-pathogen interactions: the attributes of virulence. *J. Infect. Dis.* **184**, 337–344 (2001).

8. Read, A. F. & Schrag, S. J. The evolution of virulence: experimental evidence. (1991).

9. Råberg, L. & Stjernman, M. The evolutionary ecology of infectious disease virulence. *Ecoimmunology* **548**, 78 (2012).

10. Smith, C. A. & Ashby, B. Tolerance-conferring defensive symbionts and the evolution of parasite virulence. *Evol. Lett.* **7**, 262–272 (2023).

11. Wood, J. & Ashby, B. Hyperparasitism and the evolution of parasite virulence. *Evolution (N. Y).* **77**, 2631–2641 (2023).

12. Casadevall, A. & Pirofski, L. Accidental virulence, cryptic pathogenesis, martians, lost hosts, and the pathogenicity of environmental microbes. *Eukaryot. Cell* **6**, 2169–2174 (2007).

13. Day, T. Parasite transmission modes and the evolution of virulence. *Evolution (N. Y).* **55**, 2389–2400 (2001).

14. Silva, L. M. & Koella, J. C. Complex interactions in the life cycle of a simple parasite shape the evolution of virulence. *PLoS Pathog.* (2025). doi:https://doi.org/10.1371/journal.ppat.1013294

15. Silva, L. M. & Koella, J. C. Virulence evolution: thinking outside of the host. *Evol. Appl.* **18**, e70136 (2025).

16. Walther, B. A. & Ewald, P. W. Pathogen survival in the external environment and the evolution of virulence. *Biol. Rev.* **79**, 849–869 (2004).

17. Sorci, G. & Faivre, B. Age-dependent virulence of human pathogens. *PLoS Pathog.* **18**, e1010866 (2022).

18. Liang, Y. Pathogenicity and virulence of influenza. *Virulence* **14**, 2223057 (2023).

19. Alizon, S. & Sofonea, M. T. SARS-CoV-2 virulence evolution: avirulence theory, immunity and trade-offs. *J. Evol. Biol.* **34**, 1867–1877 (2021).

20. Hawley, D. M. *et al.* Parallel patterns of increased virulence in a recently emerged wildlife pathogen. *PLoS Biol.* **11**, e1001570 (2013).

21. Possas, C. *et al.* Highly pathogenic avian influenza: pandemic preparedness for a scenario of high lethality with no vaccines. *Front. Public Heal.* **13**, 1613869 (2025).

22. Chen, Y. *et al.* Pathogen virulence genes: Advances, challenges and future directions in infectious disease research. *Int. J. Mol. Med.* **56**, 173 (2025).

23. Hoang, K. L., Read, T. D. & King, K. C. Incomplete immunity in a natural animal-microbiota interaction selects for higher pathogen virulence. *Curr. Biol.* (2024).

24. Read, A. F. *et al.* Imperfect vaccination can enhance the transmission of highly virulent pathogens. *PLoS Biol.* **13**, e1002198 (2015).

25. Read, A. F., Gandon, S., Nee, S. & Mackinnon, M. J. The evolution of pathogen virulence in response to animal and public health interventions. *Infect. Dis. host–pathogen Evol.* 265–292 (2004).

26. O'Neill, X. *et al.* Superspreading and the evolution of virulence. *PLOS Comput. Biol.* **21**, e1013517 (2025).

27. Gupta, S. Darwin review: the evolution of virulence in human pathogens. *Proc. R. Soc. B* **291**, 20232043 (2024).

28. Wolinska, J. & King, K. C. Environment can alter selection in host–parasite interactions. *Trends Parasitol.* **25**, 236–244 (2009).

29. Duneau, D. *et al.* Stochastic variation in the initial phase of bacterial infection predicts the probability of survival in D. melanogaster. *Elife* **6**, (2017).

30. Longdon, B. *et al.* The causes and consequences of changes in virulence following pathogen host shifts. *PLoS Pathog.* **11**, e1004728 (2015).

31. Iqbal, Z. *et al.* Comparative virulence studies and transcriptome analysis of Staphylococcus aureus strains isolated from animals. *Sci. Rep.* **6**, 35442 (2016).

32. Jensen, K. H., Little, T., Skorping, A. & Ebert, D. Empirical support for optimal virulence in a castrating parasite. *PLoS Biol.* **4**, e197 (2006).

33. Lafferty, K. D. & Kuris, A. M. Parasitic castration: the evolution and ecology of body snatchers. *Trends Parasitol.* **25**, 564–572 (2009).

34. Silva, L. M., Belli, A. & Koella, J. C. Testing the redox theory of aging under parasitism.

35. Zeferino, T. G., Rojas Mora, A., Vallat, A. & Koella, J. C. Dietary shifts in infected mosquitoes suggest a form of self-medication despite benefits in uninfected individuals. *Proc. B* **292**, 20251659 (2025).

36. Casadevall, A., Fang, F. C. & Pirofski, L. Microbial virulence as an emergent property: consequences and opportunities. *PLoS Pathog.* **7**, e1002136 (2011).

37. Anderson, R. M. & May, R. M. Coevolution of hosts and parasites. *Parasitology* **85**, 411–426 (1982).

38. Cressler, C. E., McLeod, D. V, Rozins, C., Van Den Hoogen, J. & Day, T. The adaptive evolution of virulence: a review of theoretical predictions and empirical tests. *Parasitology* **143**, 915–930 (2016).

39. Casadevall, A. & Pirofski, L. Virulence factors and their mechanisms of action: the view from a damage–response framework. *J. Water Health* **7**, S2–S18 (2009).

40. Choisy, M. & de Roode, J. C. Mixed infections and the evolution of virulence: effects of resource competition, parasite plasticity, and impaired host immunity. *Am. Nat.* **175**, E105–E118 (2010).

41. Acuña-Hidalgo, B., Silva, L. M., Franz, M., Regoes, R. R. & Armitage, S. A. O. Decomposing virulence to understand bacterial clearance in persistent infections. *Nat. Commun.* **13**, 1–14 (2022).

42. Wollein Waldetoft, K., Råberg, L. & Lood, R. Proliferation and benevolence—A framework for dissecting the mechanisms of microbial virulence and health promotion. *Evol. Appl.* **13**, 879–888 (2020).

43. Opota, O. *et al.* Monalysin, a novel ß-pore-forming toxin from the Drosophila pathogen Pseudomonas entomophila, contributes to host intestinal damage and lethality. *PLoS Pathog* **7**,

e1002259 (2011).

44. Rubinić, M., Alagesan, K. & Iatsenko, I. A Pore-Forming Toxin Monalysin Contributes to Infection-Induced Suppression of Defecation in Female Drosophila. *bioRxiv* 2012–2025 (2025).

45. Silva, L. M., Vallat, A. & Koella, J. C. Mechanisms of host exploitation by a microsporidian parasite. *Proc. R. Soc. B* (2025). doi:https://doi.org/10.1098/rspb.2025.1123

46. Lovegrove, F. E. *et al.* Parasite burden and CD36-mediated sequestration are determinants of acute lung injury in an experimental malaria model. *PLoS Pathog.* **4**, e1000068 (2008).

47. Vazquez, B. P. *et al.* Inflammatory responses and intestinal injury development during acute Trypanosoma cruzi infection are associated with the parasite load. *Parasit. Vectors* **8**, 206 (2015).

48. Silva, L. M. The transcriptome of the mosquito host Anopheles gambiae upon infection by different selected lines of the microsporidian parasite Vavraia culicis. *bioRxiv* 2009–2024 (2024). doi:https://doi.org/10.1101/2024.09.22.613703

49. Graham, A. L., Allen, J. E. & Read, A. F. Evolutionary causes and consequences of immunopathology. *Annu. Rev. Ecol. Evol. Syst.* **36**, 373–397 (2005).

50. Lazzaro, B. P. & Tate, A. T. Balancing sensitivity, risk, and immunopathology in immune regulation. *Curr. Opin. Insect Sci.* 100874 (2022).

51. de Villiers, K. A. & Egan, T. J. Heme detoxification in the malaria parasite: A target for antimalarial drug development. *Acc. Chem. Res.* **54**, 2649–2659 (2021).

52. Seixas, E. *et al.* Heme oxygenase-1 affords protection against noncerebral forms of severe malaria. *Proc. Natl. Acad. Sci.* **106**, 15837–15842 (2009).

53. De Roode, J. C., Yates, A. J. & Altizer, S. Virulence-transmission trade-offs and population divergence in virulence in a naturally occurring butterfly parasite. *Proc. Natl. Acad. Sci.* **105**, 7489–7494 (2008).

54. Buczinski, S., Boccardo, A. & Pravettoni, D. Clinical scores in veterinary medicine: what are the pitfalls of score construction, reliability, and validation? A general methodological approach applied in cattle. *Animals* **11**, 3244 (2021).

55. Hayes, G., Mathews, K., Kruth, S., Doig, G. & Dewey, C. Illness severity scores in veterinary medicine: what can we learn? *J. Vet. Intern. Med.* **24**, 457–466 (2010).

56. Talbot, S. R. *et al.* Robustness of a multivariate composite score when evaluating distress of animal models for gastrointestinal diseases. *Sci. Rep.* **13**, 2605 (2023).

57. Montes, N., Vijayan, V. & Pagán, I. Host population structure for tolerance determines the evolution of plant-virus interactions. *New Phytol.* (2021).

58. Baucom, R. S. & de Roode, J. C. Ecological immunology and tolerance in plants and animals. *Funct. Ecol.* **25**, 18–28 (2011).

59. Strauss, S. Y. & Agrawal, A. A. The ecology and evolution of plant tolerance to herbivory. *Trends Ecol. Evol.* **14**, 179–185 (1999).

60. Madden, L. V, Hughes, G. & Bosch, F. van den. *The study of plant disease epidemics.* (2007).

61. Shaner, G. & Finney, R. E. The effect of nitrogen fertilization on the expression of slow-mildewing resistance in Knox wheat. *Phytopathology* **67**, 1051–1056 (1977).

62. Pagán, I. & García-Arenal, F. Tolerance to Plant Pathogens: Theory and Experimental Evidence. *Int. J. Mol. Sci.* **19**, 810 (2018).

63. Jeger, M. J. Tolerance of plant virus disease: Its genetic, physiological, and epidemiological significance. *Food Energy Secur.* **12**, e440 (2023).

64. Day, T. & Burns, J. G. A consideration of patterns of virulence arising from host-parasite coevolution. *Evolution (N. Y).* **57**, 671–676 (2003).

65. Bonsall, M. B. Parasite replication and the evolutionary epidemiology of parasite virulence. *PLoS One* **5**, e12440 (2010).

66. Surasinghe, S. & Ogbunugafor, C. B. Learning virulence-transmission relationships using causal inference. *arXiv Prepr. arXiv2602.05196* (2026).

67. Franz, M., Armitage, S. A. O., Rolff, J. & Regoes, R. R. Virulence decomposition for bifurcating infections. *Proc. R. Soc. B* **290**, 20230396 (2023).

68. Simms, E. L. & Triplett, J. Costs and benefits of plant responses to disease: resistance and tolerance. *Evolution (N. Y).* **48**, 1973–1985 (1994).

69. Macedo da Silva, L. M. Pathogen infection dynamics and the evolution of host resistance and tolerance. (2021).

70. Stearns, S. C. Trade-offs in life-history evolution. *Funct. Ecol.* **3**, 259–268 (1989).

71. Li, J. *et al.* Excess mortality of infected ectotherms induced by warming depends on pathogen kingdom and evolutionary history. *PLoS Biol.* **22**, e3002900 (2024).

72. Vale, P. F., Wilson, A. J., Best, A., Boots, M. & Little, T. J. Epidemiological, evolutionary, and coevolutionary implications of context-dependent parasitism. *Am. Nat.* **177**, 510–521 (2011).

73. Howick, V. M. & Lazzaro, B. P. Genotype and diet shape resistance and tolerance across distinct phases of bacterial infection. *BMC Evol. Biol.* **14**, 56 (2014).

74. Silva, L. M., King, K. C. & Koella, J. C. Dissecting transmission to understand parasite evolution. *PLoS Pathog.* **21**, e1012964 (2025).

75. Leggett, H. C., Cornwallis, C. K., Buckling, A. & West, S. A. Growth rate, transmission mode and virulence in human pathogens. *Philos. Trans. R. Soc. B Biol. Sci.* **372**, 20160094 (2017).

76. Stage, F. K., Carter, H. C. & Nora, A. Path analysis: An introduction and analysis of a decade of research. *J. Educ. Res.* **98**, 5–13 (2004).

77. Streiner, D. L. Finding our way: an introduction to path analysis. *Can. J. Psychiatry* **50**, 115–122 (2005).

78. Cohen, C., Einav, M. & Hawlena, H. Path analyses of cross-sectional and longitudinal data suggest that variability in natural communities of blood-associated parasites is derived from host characteristics and not interspecific interactions. *Parasit. Vectors* **8**, 429 (2015).

79. King, K. C., Hall, M. D. & Wolinska, J. Infectious disease ecology and evolution in a changing world. *Philosophical Transactions of the Royal Society B* **378**, 20220002 (2023).

80. Bedhomme, S., Agnew, P., Sidobre, C. & Michalakis, Y. Virulence reaction norms across a food gradient. *Proc. R. Soc. London. Ser. B Biol. Sci.* **271**, 739–744 (2004).

Box 1. Glossary

Virulence: infection-induced harm to the host. In evolutionary contexts, this harm is often inferred from effects on LRS-related fitness traits, such as survival and reproduction.

Intrinsic virulence: harm caused by a parasite genotype under controlled conditions, typically for comparative purposes.

Realized virulence: the total impact of the infection expressed in a specific ecological and evolutionary context, by considering host, parasite and environmental variation.

Proximate virulence: Mechanistic and host-state (or condition) traits impacted by infection, which contribute to harm but are not equivalent to downstream fitness outcomes unless explicitly demonstrated.

Composite virulence: A measurement strategy that integrates multiple traits into a single representation of infection cost for a specific biological target/question.

Biological target: The specific question or facet of infection that a composite measure represents.

Drivers of harm: Parasite traits or processes that generate stress or damage (*e.g.,* replication rate).

Host-state mediators: Host responses or conditions that transmit or modify damage (*e.g.,* pathology, immune activation, dysbiosis).

Life-history outcomes: Downstream effects of infection on host survival, reproduction, or development that determine evolutionary consequences. Evolutionary measure of virulence, by quantifying fitness costs.

Host exploitation: Parasite growth-dependent cost of infection to the host, often measured as a reduction in resources and energy reserves in proportion to parasite growth.

Per-parasite pathogenicity (PPP): Harm caused per unit of parasite, reflecting mechanisms that are independent of parasite density, such as inflammation and toxin production.

Box 2. Diagnostic criteria for interpretable composite virulence

Composite virulence is best treated as a measurement strategy rather than a definition. Its interpretability depends on making a small set of biological decisions explicit.

1) Specify the biological target/question. State which group of traits your composite represents, and which question(s) it should answer.

2) Respect trait hierarchy. Traits occupy different positions in the infection process, including parasite drivers (*e.g.,* replication or burden), host-state mediators (*e.g.*, pathology or performance impairment), and life-history outcomes (*e.g.,* survival or reproduction). A composite should be built from traits that define a single biological target at one of these levels. Combining traits from different levels into a single score blurs interpretation and can introduce circularity (*e.g.,* by combining causes and consequences). Traits from other levels should instead be analyzed alongside the composite to explain why it varies.

3) Minimize redundancy. Correlated traits often measure the same biology and can be double-counted. Prefer sparse integration: group measurements into a few categories (*e.g.*, survival, reproduction, performance) and represent each with a defensible summary; assess whether each categories contribute information beyond others.

4) Treat time and context as part of the phenotype. Infection impact is dynamic and context-dependent. Where feasible, use longitudinal outcomes or biologically justified windows and report timing features (onset/peak/recovery) alongside magnitude. Evaluate composites (or categories) across key contexts and compare reaction norms rather than relying on single-condition rankings.

5) Check robustness. Report components/categories alongside any single-number index and show that conclusions are not artifacts of scaling, weighting, or a single correlated trait.