Received: 15 April 2022 | Accepted: 21 December 2022

DOI: 10.1111/2041-210X.14065

Methods in Ecology and Evolution

RESEARCH ARTICLE

An exact version of Life Table Response Experiment analysis, and the R package exactLTRE

Christina M. Hernández[1] | Stephen P. Ellner[1] | Peter B. Adler[2] | Giles Hooker[3,4] | Robin E. Snyder[5]

[1]Department of Ecology and Evolutionary Biology, Cornell University, Ithaca, New York, USA

[2]Department of Wildland Resources and The Ecology Center, Utah State University, Logan, Utah, USA

[3]Department of Statistics and Data Science, Cornell University, Ithaca, New York, USA

[4]Department of Statistics, University of California Berkeley, Berkeley, California, USA

[5]Department of Biology, Case Western Reserve University, Cleveland, Ohio, United States

Correspondence
Christina M. Hernández
Email: cmh352@cornell.edu

Funding information
National Science Foundation, Grant/Award Number: DEB-1933497, DEB-1933561 and DEB-1933612

Handling Editor: Torbjørn Ergon

Abstract

1. Matrix population models are frequently built and used by ecologists to analyse demography and elucidate the processes driving population growth or decline. Life Table Response Experiments (LTREs) are comparative analyses that decompose the realized difference or variance in population growth rate (λ) into contributions from the differences or variances in the vital rates (i.e. the matrix elements). Since their introduction, LTREs have been based on approximations and have not included biologically relevant interaction terms.
2. We used the functional analysis of variance framework to derive an exact LTRE method, which calculates the exact response of λ to the difference or variance in a given vital rate, for all interactions among vital rates—including higher-order interactions neglected by the classical methods. We used the publicly available COMADRE and COMPADRE databases to perform a meta-analysis comparing the results of exact and classical LTRE methods. We analysed 186 and 1487 LTREs for animal and plant matrix population models, respectively.
3. We found that the classical methods often had small errors, but that very high errors were possible. Overall error was related to the difference or variance in the matrices being analysed, consistent with the Taylor series basis of the classical method. Neglected interaction terms accounted for most of the errors in fixed design LTRE, highlighting the importance of two-way interaction terms. For random design LTRE, errors in the contribution terms present in both classical and exact methods were comparable to errors due to neglected interaction terms. In most examples we analysed, evaluating exact contributions up to three-way interaction terms was sufficient for interpreting 90% or more of the difference or variance in λ.
4. Relative error, previously used to evaluate the accuracy of classical LTREs, is not a reliable metric of how closely the classical and exact methods agree. Error compensation between estimated contribution terms and neglected contribution terms can lead to low relative error despite faulty biological interpretation. Trade-offs or negative covariances among matrix elements can lead to high relative error despite accurate biological interpretation. Exact LTRE provides reliable

and accurate biological interpretation, and the R package exactLTRE makes the exact method accessible to ecologists.

KEYWORDS
demography, life history, Life Table Response Experiment, matrix population model, population growth rate

1 | INTRODUCTION

Matrix population models are widely used for demographic studies of both plant and animal populations, including for conservation and management. These models relate population structure (e.g. age, size or developmental stage) to vital rates and population growth, yet are simple to build from field-collected census data. Once built, the models can be used to calculate a variety of population- and individual-level metrics. The asymptotic population growth rate (λ), given by the largest eigenvalue of the population projection matrix, is of particular interest; others include the stable population structure, the expected number of offspring over an individual's lifespan (R_0), and the expected time between generations.

In addition to studying the dynamics of a single population, matrix population models are valuable tools for comparative demography. They enable us to investigate the processes driving variation in populations across locations, treatments, or points in time. An important analysis for comparative demography is Life Table Response Experiment, or "LTRE," which relates observed variation in vital rates to observed variation in a population-level metric. A direct comparison of the vital rates between populations can be misleading because a large difference in a given vital rate will not necessarily produce a large difference at the population level (e.g. λ). On the other hand, LTREs account for both the *sensitivity* of λ to the vital rates and the *observed* changes in those same rates, including the *covariation* among different vital rates (Caswell, 1989). For example, the λ of a matrix population model might be highly sensitive to seed survival in the seedbank, but this may explain little of the observed variation in λ if seed survival showed very limited variation in the populations being analysed. Because of their grounding in observations that account for covariation among vital rates, the results of LTRE analysis can be used to understand the processes driving population dynamics (e.g. Caswell, 1996; Fréville & Silvertown, 2005). This can be particularly important for understanding human impacts on endangered, declining, or harvested species (e.g. Bruna & Oli, 2005; Oro & Doak, 2020).

Fundamentally, an LTRE analysis decomposes the difference or variance in an observed population outcome into the contributions from a set of parameters or vital rates that vary across two or more matrices. The population growth rate (λ) is most commonly used and will be our focus, but any population outcome (e.g. R_0, expected lifespan, generation time) can be analysed with these methods. We define the *contribution* of a given vital rate or matrix element to be the observed effect on λ when all else is held constant. For example, the contribution of the observed difference in juvenile survival is the difference in λ that results from allowing only juvenile survival to differ between two matrices of interest, with all other matrix elements held at appropriate baseline (or standard-of-reference) values. The vital rates are rarely independent from one another (e.g. trade-offs between survival and reproductive output), and λ tends to be a strongly non-linear function of the matrix elements, so we would expect interactions between/among vital rates to contribute to observed changes in λ. These interaction terms may have important biological meaning. For instance, higher adult reproductive output will have a stronger effect on λ when juvenile survival is high. There also may be important interactions among adult survival and fertility, juvenile survival and first-year reproductive output, and growth of multiple pre-reproductive size classes (because of the effect on time to first reproduction).

LTREs were invented and popularized for ecology by Caswell (1989, 1996). The analyses as formulated therein depend on Taylor series expansions to approximate λ as a function of the matrix elements (Caswell, 1989).[1] As such, these approximations do not consider the effects on λ of many of the possible interactions among matrix elements, but instead provide estimated contributions from a limited range of terms (main effects and, sometimes, two-way interaction terms). Thus, classical LTRE methods are both incomplete (they omit potentially important higher-order interactions) and approximate (the terms that they do include are based on approximations that may not be accurate across the range of matrices being analysed).

Functional analysis of variance (fANOVA; Ellner et al., 2019; Hooker, 2007) provides a theoretical framework for designing an exact version of LTRE in which the contributions of each matrix element, and their interactions, are obtained directly from the difference or variance in λ caused by each combination of matrix elements, accounting for lower-order terms. Specifically, the response of interest is decomposed as a sum,

$$\begin{aligned}\Delta\lambda \text{ or } \mathrm{Var}(\lambda) \;&= \sum_{i,j} \text{Main effect of matrix entry } a_{i,j} \\ &+ \sum_{i,j,l,m} \text{Interaction between } a_{ij} \text{ and } a_{l,m} \\ &+ \sum_{i,j,k,m,n,o} \text{Three-way interaction among } a_{ij}, a_{l,m}, a_{n,o} \cdots \\ &+ \sum k\text{-way interactions} \\ &+ \text{Total contribution of all interactions of order } (k+1) \text{ or higher.}\end{aligned} \tag{1}$$

Below, we suggest $k = 2$ or 3, because 4-way and higher-order interactions are hard to interpret. The methods we present enables

the calculation of all higher-order interactions for any set of matrices, though custom code may be required for large matrices. Our R package is limited by the maximum allowable vector size in R (currently 2×10^9), so it can calculate all possible interaction terms for matrices of up to 30 matrix elements that vary, and can calculate terms up to $k = 3$ for up to 2289 matrix elements that vary.[2] Relative to classical LTRE, Equation (1) includes interaction terms (up to some order) that classical LTRE omits (which may be its main benefit for some users), and it is exact in that each term is computed exactly (up to the precision of computer arithmetic) and their sum is exactly equal to $\Delta\lambda$ or $\mathrm{Var}(\lambda)$.

In what follows, we focus on the two most flexible and frequently used "classical" LTRE analyses: one-way fixed design and random design. There are additional forms of LTRE that are appropriate for different experimental settings or data sets. For example, regression LTRE analyzes a continuous curve giving the response of λ to a continuous treatment variable x; it uses the chain rule to exactly decompose $d\lambda / dx$ in terms of the sensitivity of each vital rate to the treatment, and the sensitivity of λ to each vital rate (Caswell, 2001, pp. 273–274). Factorial LTRE accounts for interactions among multiple treatments, using a first-order Taylor approximation with respect to matrix entries (Caswell, 1989, 2001, p. 263). The fANOVA-based methods that we present here can be applied to both regression and factorial LTRE designs, but we do not focus on those applications in this paper. Appendix E provides a "recipe" for how our methods and R package could be used to calculate an fANOVA-based exact factorial design LTRE.

The next three sections of this paper present methods and theory for exact LTRE. We first review the classical methods (Section 2). We then introduce the formulas and algorithms for performing exact LTRE (Section 3) and an R package that contains functions for both classical and exact LTRE methods (Section 3.2 and Table S1). After that, we present a meta-analysis of LTREs from a wide array of matrix population models for both plants and animals. We first characterized the scale and distribution of errors in classical LTRE methods, and then focused on three primary research questions: (1) How do errors arise in classical LTRE and how do they affect interpretation?, (2) Is the relative error a useful measure for evaluating the accuracy of classical LTRE?, and (3) How important are higher-order terms in LTRE analyses? Section 4 describes the design of this meta-analysis and how we defined and identified errors. We then present the results and several instructive examples of LTREs with errors (Section 5) and discuss the implications of our meta-analysis for the use of exact LTRE methods (Section 6).

As we show, the errors from the approximate classical methods can be large and can change our interpretation of population dynamics. Second-order terms matter a lot, and higher-order terms can also be quite important. Although the classical methods work well when there are small differences between the matrices being examined, the relative error is an unreliable measure of accuracy. Meanwhile, the exact method will always yield the true contributions, and our R package makes evaluation of exact LTRE easy.

2 | BACKGROUND: CLASSICAL LTRE CALCULATIONS

Many of the analytical techniques for matrix population models rely on the existence of a unique real-valued eigenvalue that is larger than all others, and which has corresponding right and left eigenvectors that have real non-negative values. The Perron-Frobenius Theorem guarantees the existence of such an eigenvalue for non-negative matrices that are irreducible and primitive (Caswell, 2001).

The calculations presented here were introduced in Caswell (1989) and Brault and Caswell (1993). We closely follow the notation in the Caswell (2001) monograph (p. 261 and 269). Terminology for the matrices used in LTRE analyses is summarized in Table 1.

2.1 | One-way fixed design

A fixed design LTRE is used when the particular treatment levels or conditions that a population faced are themselves of interest. A one-way fixed design LTRE decomposes the *difference* in λ between a pair of matrices, one identified as the *reference* matrix and the other as corresponding to *treatment* level "m". The difference in λ can be attributed to the differences in the elements of the matrices according to

$$\Delta\lambda = \lambda^{(m)} - \lambda^{(r)} \approx \sum_{i,j} \left(a_{ij}^{(m)} - a_{ij}^{(r)}\right) \left.\frac{\partial\lambda}{\partial a_{ij}}\right|_{\mathbf{A}^\dagger}, \tag{2}$$

where a_{ij} is the (i,j) entry of a population projection matrix, and $\frac{\partial\lambda}{\partial a_{ij}}$ is given by the (i,j) entry of the sensitivity matrix evaluated at the *pivot matrix* $\mathbf{A}^\dagger$. The term involving a_{ij} is called the *contribution* of the (i,j) matrix entry to $\Delta\lambda$.

In principle, the pivot matrix $\mathbf{A}^\dagger$ can be any matrix between the reference and treatment matrices, but it is generally recommended to choose $\mathbf{A}^\dagger$ equal to the mean matrix,

$$\mathbf{A}^\dagger = \overline{\mathbf{A}} = \frac{1}{2}\left(\mathbf{A}^{(m)} + \mathbf{A}^{(r)}\right), \tag{3}$$

to maximize the expected accuracy of the linear approximation for λ as a function of matrix entries (see Appendix A.1 for the technical details). The sensitivity matrix, the derivatives of λ with respect to each entry of $\mathbf{A}^\dagger$, is calculated from the right eigenvector ($\mathbf{w}$) and left eigenvector ($\mathbf{v}$) corresponding to λ,

$$\mathbf{S} = \frac{\mathbf{v}\mathbf{w}^\top}{\mathbf{v}\cdot\mathbf{w}}. \tag{4}$$

2.2 | Random design

A random design LTRE decomposes the *variance* in λ into contributions from the *variance and covariance* in the matrix elements, according to

$$\mathrm{Var}(\lambda) \approx \sum_{i,j}\sum_{k,l} C(ij,kl)s_{ij}s_{kl}, \tag{5}$$

TABLE 1 Definitions of matrices used in Life Table Response Experiment (LTRE) calculations

Matrix name	Analyses used in	Meaning
Treatment matrix $\mathbf{A}^{(m)}$ and Reference matrix $\mathbf{A}^{(r)}$	Classical fixed design	The two matrices being compared in a classical one-way fixed design
Pivot matrix $\mathbf{A}^{\dagger}$	Classical fixed design, classical random design	The matrix at which sensitivities are evaluated. We introduce this term here, consistent with terminology for the point about which a Taylor series is expanded
Mean matrix $\overline{\mathbf{A}}$	Classical fixed design, classical random design, exact fixed design, exact random design	Calculated by taking the mean of each matrix element across all matrices in a set of interest
Observed matrices	Classical random design, exact fixed design, exact random design	The set of matrices for which the difference or variance is being decomposed
Baseline matrix	Exact fixed design, exact random design	The matrix from which we evaluate how changes in matrix elements and their interactions affect lambda, in the functional analysis of variance-based exact LTRE method

where $C(ij, kl)$ is the covariance of a_{ij} and a_{kl}, and the sensitivities s_{ij} and s_{kl} are the (i, j) and (k, l) entries of the sensitivity matrix $\mathbf{S}$ evaluated at the pivot matrix. For random design LTRE, the pivot matrix is always chosen to be the mean matrix across the observed matrices for which variance is being decomposed.

3 | EXACT LTRE, FROM fANOVA PRINCIPLES

Like the classical methods, a fixed design exact LTRE decomposes the *difference* in λ, while a random design exact LTRE decomposes the *variance* in λ.

As stated previously, LTRE decomposes the difference or variance in a population-level outcome (such as λ) into contributions from different vital rates and their interactions. From here on, "vital rates" will refer to elements of the projection matrix, and the functions in our R package also operate at the level of projection matrix elements. However, we note that the methods we present can be applied to underlying parameters that define matrix elements. For example, if a matrix element $a_{4,3}$ is computed as the product of survival probability s_3 and growth probability $g_{4,3}$, while $a_{5,3}$ is analogously computed as $s_3 g_{5,3}$, our methods can be applied to s_3, $g_{4,3}$, $g_{5,3}$ and other parameters defining other matrix entries, rather than to the matrix entries themselves. The methods can also be applied to models other than matrix population models.

We carry out exact LTRE analyses using fANOVA (Ellner et al., 2019; Hooker, 2007; see Appendix A.4 for a detailed general introduction to fANOVA) to decompose the observed difference or variance in λ into contributions from each individual matrix element, and from their interactions. We start with a set of observed matrices and select a reasonable baseline matrix. In random design LTRE, the baseline is always the mean matrix, that is, the matrix composed of mean values for each matrix element across all observed matrices used in the calculation. In fixed design LTRE, we recommend using a "control," "undisturbed," or some other edge case population as the baseline; if such a matrix is unavailable, then we recommend using the mean matrix ($\overline{\mathbf{A}}$ in Equation 3: see Section 3.1 for more discussion of choosing a baseline in fixed design exact LTRE). We then change matrix elements from the baseline one by one, and in all possible combinations, and calculate the difference (fixed design) or variance (random design) in λ under each hypothetical situation. fANOVA then gives us a general recipe for converting the set of responses to changes in matrix elements into a set of main effects of each element, and a set of higher-order interactions, whose sum equals the effect of simultaneously changing all matrix elements from the baseline to the observed values.

For example, suppose that we have several projection matrices for a population with two stages, using a pre-breeding census. Juveniles (stage 1) are approximately 1 year old, and adults (stage 2) include all individuals that are approximately 2 years old or older. New juveniles are produced by reproductive activity of adults. The projection matrix takes the form

$$\mathbf{A} = \begin{bmatrix} 0 & f_a \\ s_J & s_a \end{bmatrix},$$

where f_a is the per-capita reproduction by adults, s_J is the probability of a juvenile surviving to the adult stage, and s_a is the probability of adults surviving for another year.

Everything that follows in this section will be written for fixed design LTRE, a decomposition of the difference in λ between two matrices. For a random design LTRE, $\Delta\lambda$ should be replaced with $\mathrm{var}(\lambda)$. The example below is a *directional* analysis, meaning that one of the two observed projection matrices is used as the baseline for perturbations; an alternative *symmetric* analysis is explained in Section 3.1 and Appendix A.5.

Specifically, suppose that we have projection matrices for two laboratory populations, one a control population (represented by $\mathbf{A}^{(c)}$) and the other exposed to a pollutant ($\mathbf{A}^{(p)}$) that has a negative effect on all non-zero matrix elements,

$$\mathbf{A}^{(c)} = \begin{bmatrix} 0 & 3 \\ 0.6 & 0.9 \end{bmatrix} \quad \mathbf{A}^{(p)} = \begin{bmatrix} 0 & 1 \\ 0.35 & 0.5 \end{bmatrix}.$$

The control population is growing ($\lambda = 1.87$)[3] and the pollutant-exposed population is declining ($\lambda = 0.89$). The difference in λ due to the pollutant is – 0.98. How much of this difference comes from the three individual effects of decreasing one matrix element, and how much from interactions among those decreases?

To calculate these contributions using our fANOVA approach, we construct hypothetical (or "counterfactual") matrices where only some of the matrix elements differ from the baseline. Since we are interested in understanding how the pollutant caused changes from the control conditions, we use the control matrix as the baseline matrix. Thus, we calculate the main effect of the change in juvenile survival by

$$c^{s_J} = \Delta\lambda^{s_J} = \lambda\left(\begin{bmatrix} 0 & 3 \\ 0.35 & 0.9 \end{bmatrix}\right) - \lambda\left(\begin{bmatrix} 0 & 3 \\ 0.6 & 0.9 \end{bmatrix}\right) = -0.296. \quad (6)$$

The first matrix above (with $s_J = 0.35$) does not correspond to any real population. But it tells us exactly what *would have happened* to λ if only juvenile survival differed between treatment and control.

Likewise, the main effect of changes in adult fertility is what *would have happened* to λ if only adult fertility differed in the treatment population,

$$c^{f_a} = \Delta\lambda^{f_a} = \lambda\left(\begin{bmatrix} 0 & 1 \\ 0.6 & 0.9 \end{bmatrix}\right) - \lambda\left(\begin{bmatrix} 0 & 3 \\ 0.6 & 0.9 \end{bmatrix}\right) = -0.519. \quad (7)$$

To evaluate the interaction between adult fertility and juvenile survival, we need to know the effect of changing both,

$$\Delta\lambda^{s_J,f_a} = \lambda\left(\begin{bmatrix} 0 & 1 \\ 0.35 & 0.9 \end{bmatrix}\right) - \lambda\left(\begin{bmatrix} 0 & 3 \\ 0.6 & 0.9 \end{bmatrix}\right) = -0.672. \quad (8)$$

In this case, the interaction term c^{s_J,f_a} does not equal $\Delta\lambda^{s_J,f_a}$. Rather, the interaction of juvenile survival and adult fertility is defined as the difference between their combined effect, and the sum of their individual main effects,

$$c^{s_J,f_a} = \Delta\lambda^{s_J,f_a} - \left(c^{s_J} + c^{f_a}\right) = 0.143. \quad (9)$$

The interaction is positive, which has an intuitive biological interpretation. The intuition is that a decrease in adult fertility f_a has a bigger effect when juvenile survival s_J is high, than it does when juvenile survival is low. The lower s_J in the treatment population thus reduces the negative impact of the lower f_a, hence the interaction term is positive (i.e. the interaction turns a negative number into a smaller negative number). We could also interpret the interaction by saying that a change in juvenile survival has a bigger effect on λ when adult fertility is high than it does when adult fertility is low—both ways of looking at it are equally valid.

This biological interpretation illustrates an important general point: interaction between juvenile survival and adult fertility is a genuine component of the change in λ between treatment and control. A decomposition of the change that only included main effects—such as the classical LTRE—is biologically incomplete because meaningful components of the change are omitted.

The main effect of s_a and the other pairwise interactions are calculated following the same pattern. Finally, the contribution from the three-way interaction of juvenile survival, adult survival, and adult fecundity, is

$$c^{s_J,f_a,s_a} = \Delta\lambda^{s_J,f_a,s_a} - \left(c^{s_J} + c^{f_a} + c^{s_a} + c^{s_J,f_a} + c^{s_J,s_a} + c^{f_a,s_a}\right) = -0.005.$$

That is, the three-way interaction is the effect of changing all three matrix entries from control to treatment values, above and beyond the sum of the three main effects and the three pairwise interactions.

This pattern extends to higher orders of interaction. In Appendix A.5, we present a fully worked symmetric fixed design exact LTRE with three varying parameters, with all terms and matrices in fANOVA notation.

To handle situations where any number of matrix elements can vary, a general operator matrix can be defined that calculates all main effects and interaction terms from the vector of $\Delta\lambda$ values for the different counterfactual matrices. In Appendix B, we present these calculations using standard notation from fANOVA, and provide a proof of the form of our general operator matrix. In Appendix C, we present a method, developed for the analysis of epistasis (Poelwijk et al., 2016), that is computationally efficient when all possible interaction terms are to be calculated, but requires more memory and is therefore limited to a smaller number of varying matrix entries. Our R package uses either the epistasis form of the operator matrix or our general form, depending on the requested interaction order and the number of matrix elements that vary.

It is often useful to calculate a subset of contributions up to a chosen interaction order, for two reasons. First the number of interaction terms increases very rapidly with the number of varying matrix elements. Second, contributions from the interactions of three or more matrix elements become difficult to interpret biologically. However, when interactions above a chosen order are not calculated, it is important to ask whether the omitted terms are too important to neglect. Because exact LTRE defines a set of contributions that sum to the exact difference or variance in λ, the discrepancy between the sum of computed contributions and the observed difference or variance, equals the net effect of all uncomputed interactions. Users should check that this discrepancy is small (i.e. less than 5%–10% of the total difference). A small discrepancy is not a guarantee that all higher-order interactions are small, because terms of opposite sign can cancel each other out. But a large discrepancy implies that some higher-order interactions are too large to neglect.

3.1 | Directional versus symmetric fixed design LTRE

The classical method for one-way fixed design LTRE, as presented here and Caswell (2001), is a symmetric analysis in the following

sense (see Appendix A.1 for more details). Because it is recommended (for good reasons) that eigenvalue sensitivities are calculated at a pivot matrix equal to the average of the treatment and reference matrices ($\overline{\mathbf{A}}$), interchanging the treatment matrix $\mathbf{A}^{(m)}$ with the reference matrix $\mathbf{A}^{(r)}$ changes the sign of all contributions but does not affect their magnitude. The relative importance of contributions from different matrix entries remains exactly the same. Neither matrix has a privileged role in determining the relative importance of different contributions. The analysis thus provides an answer to the question, *why are the treatment and reference matrices different from each other*? This is appropriate when none of the treatment matrices has a special status as a "control" (or undisturbed, or some other edge case such as lowest elevation, etc.) while the others represent different levels of some treatment. For example, if a fixed design LTRE is used to compare two nearby lakes without any experimental manipulations (Figure 1), neither one should play a privileged role and a symmetric analysis is appropriate. Either lake can be chosen as reference with the other lake as treatment, and the choice is immaterial so long as the $\overline{\mathbf{A}}$ is used as the pivot matrix.

But experimental designs are often not symmetric in this sense, instead comparing a "control" with one or more treatment populations (e.g. applications of various pesticides). In such cases there is a meaningful difference between the populations, with the control population privileged as the standard of reference to which the treatment population is compared. The research question then becomes *why is the treatment population different from the control*? More precisely, how do each of the changes in matrix entries from control value $a_{ij}^{(r)}$ to treatment value $a_{ij}^{(m)}$ contribute to $\lambda^{(m)} - \lambda^{(r)}$?

Unlike the classical method, the fANOVA method for fixed design LTRE has the option of being formulated as either symmetric (why are these populations different from each other?) or asymmetric and directional (why is the treatment different from the control)? This is determined by the choice of baseline matrix for the analysis. If the mean matrix $\overline{\mathbf{A}}$ is used as the baseline for fANOVA, the analysis is symmetric. Each term in the decomposition measures the effect on the $\Delta\lambda$ between the two observed matrices when some entries in each matrix take their true (observed) values, rather than the average value. As in the classical LTRE, this implies that interchanging the two observed population matrices only affects the sign of contributions, not their absolute magnitude. The two observed matrices play equivalent roles with neither being privileged. For a fully-worked example of a symmetric exact LTRE, see Appendix A.5.

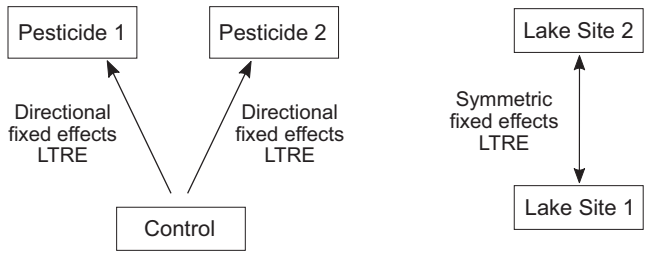

FIGURE 1 Example scenarios calling for directional versus symmetric fixed effects Life Table Response Experiment (LTRE) analyses.

However, if one of the observed matrices is used as the baseline in the fANOVA decomposition, the results are directional. Each term in the decomposition measures the $\Delta\lambda$ relative to the baseline matrix that results from taking some elements of the baseline matrix, and replacing them with the corresponding elements of the observed treatment matrix (see the worked directional example in Section 3 above). Each term in the fANOVA decomposition can therefore be interpreted as reflecting one (or some) of the ways that the treatment affected the reference population. The fANOVA exact LTRE thus has the added flexibility of allowing the user to choose between a symmetric analysis comparing two populations with neither playing a privileged role and a directional analysis with one population identified as the baseline standard of reference.

In our meta-analysis comparing the exact LTRE methods with the classical methods (Section 4), we always used the symmetric fixed design for the exact LTREs. However, we also determined whether the LTREs we identified were more appropriate for symmetric or directional analysis. We classified studies as directional if they compared a control with a treatment, an unharvested or unperturbed population with a human-impacted population, or if the comparison was between the center and an edge of the population range. In examples from animal species, 56 of the 97 (58%) fixed design LTREs were appropriate for directional analysis. For the data on plant species, we found 230 out of 645 (36%) fixed design LTREs were appropriate for directional analysis.

3.2 | The R package exactLTRE

The R package exactLTRE contains the functions required to evaluate LTREs on an arbitrary set of matrices (Table S1). It includes functions for both the approximate and exact methods as presented here. The functions include options to select random or fixed design, and to select the maximum interaction order to be calculated for exact LTRE. For fixed design exact LTRE, the user can specify whether the analysis will be directional or symmetric.

The exactLTRE R package also includes functions to calculate the fundamental matrix, net reproductive rate (R_0), lifespan, generation time, and the variance–covariance matrix for a set of population projection matrices (Table S1). The package documentation includes examples of usage, and our code for analysing the large number of LTREs from the databases of animal and plant matrix population models is available in Supplemental Materials (Hernández et al. 2022).

4 | METHODS: META-ANALYSIS COMPARING EXACT AND CLASSICAL LTRE

Exact LTRE is preferable to the classical methods in principle, but ecologists may wonder if switching is necessary in practice.

Additionally, they might wonder if past classical analyses, by themselves or other researchers, need to be re-visited. Because the classical approximate methods are familiar and have seen wide use in ecology, we were interested in characterizing the scale of errors in the classical methods and the prevalence of qualitative (interpretation) errors. It would also be valuable if we could identify situations where the classical methods are reliably accurate, to help in evaluating previous LTRE analyses.

The exact and classical LTREs each generate a vector of contributions. The exact LTRE extends to a user-selected order in Equation (1) (here we use $k = 3$) and uses exact evaluation of each term ($\mathbf{c}$) by the methods explained in Section 3. The general formulas for the terms are derived and explained in Appendices A.4 and A.5. The vector $\mathbf{c}$ includes the final term in Equation (1), the sum of all interactions of order $k + 1$ or higher that are not computed individually, so the sum of all terms in $\mathbf{c}$ is exactly equal to the response, $\Delta\lambda$ or $\mathrm{Var}(\lambda)$. The classical LTRE extends to first (fixed design) or second (random design) order, and uses approximate evaluations of terms ($\hat{\mathbf{c}}$) as explained above in Section 2.

We define the total or overall error (E) of the classical LTRE as the distance between the vectors of contributions estimated by the classical and exact methods

$$E = \| \mathbf{c} - \hat{\mathbf{c}} \|_1, \tag{10}$$

where $\hat{\mathbf{c}}$ has been extended to the same length as $\mathbf{c}$ by adding a zero value for each of the higher-order terms that are missing from $\hat{\mathbf{c}}$. The 1-norm of a vector is the sum of the absolute values of the vector entries.

The overall error E comes from the combined effect of *approximation error* (the discrepancy between approximate and exact values of terms that are present in both methods) and *truncation error* (error due to the higher-order terms not included in the classical method). The approximation errors represent the mismatch between a Taylor expansion approach to LTRE that relies on calculating sensitivity $\frac{\partial\lambda}{\partial a_{ij}}$, and the direct calculation of how each observed change in a_{ij} caused changes to λ. We compute the approximation and truncation errors by taking the 1-norm of the corresponding portions of $(\mathbf{c} - \hat{\mathbf{c}})$.

We also wanted to investigate whether these errors show any patterns with other characteristics of the matrices or species life history. As covariates, we considered the distance between or among matrices, difference or standard deviation in λ, the value of λ at the mean matrix, the matrix dimension, the net reproductive output (R_0), lifespan, and generation time. The distance between or among matrices is calculated as

$$D_{\text{between}} = \sum_{ij} | a_{ij}^{(1)} - a_{ij}^{(2)} | \quad \text{or} \quad D_{\text{among}} = \sum_{ij} \mathrm{var}(a_{ij}). \tag{11}$$

Lifespan and R_0 are calculated using the fundamental matrix, which gives the expected number of timesteps that an individual spends in each stage/age/size class. We used the Bienvenu and Legendre (2015) definition of generation time as the average time between two birth events in an individual's ancestral genealogy, calculated as

$$T = \frac{\lambda \mathbf{v}^{\mathsf{T}} \mathbf{w}}{\mathbf{v}^{\mathsf{T}} \mathbf{F} \mathbf{w}}, \tag{12}$$

where $\mathbf{v}$ and $\mathbf{w}$ are the left and right eigenvectors, respectively, of the projection matrix $\mathbf{A}$, and $\mathbf{F}$ is a matrix containing all the fertility transitions and zeros elsewhere.

In the past, researchers have often evaluated how well the classical method performed by calculating the relative error of the sum of contributions compared to the observed difference or variance in λ

$$E_{rel} = \left| \left(\frac{\Delta\lambda - \sum_i (\hat{c}_i)}{\Delta\lambda} \right) \right|, \tag{13}$$

where $\hat{c}_i$ is one of the contribution terms in the classical (approximate) LTRE. In the case of random design LTRE, $\Delta\lambda$ would be replaced by $\mathrm{var}(\lambda)$. This relative error has the advantage that it can be calculated from only the results of the classical method. We used the Pearson correlation coefficient to evaluate how well the relative error (Equation 13) predicts the overall error of the contribution vector (Equation 10).

4.1 | Selection of matrices for the meta-analysis

We used the COMADRE and COMPADRE databases (Salguero-Gómez et al., 2015, 2016) to compile a list of all LTREs (fixed or random design) that could reasonably be conducted. These databases archive published matrix population models in a form that is easily accessed from R, with a variety of accompanying metadata. COMADRE contains matrix population models on animals, and COMPADRE contains those on plants. We will present the results from LTREs on animal and plant species separately, because there are a few notable differences between these groups. There are far more published matrix population models on plants than on animals. Plant life cycles are more likely to be size-classified and to include shrinkage, while animals are typically age- or stage-classified without a possibility of regression. Plant models are also more likely to include multiple types of offspring, such as reproduction into the seed bank and reproduction directly to seedlings.

Matrices were chosen from COMADRE version 4.21.1.0 (release date 25 January 2021) and COMPADRE version 6.21.1.0 (release date 25 January 2021). During our screening and analysis the databases had ongoing edits for quality control, so the analyses presented here were performed using COMADRE version 4.21.8.0 (release date 20 August 2021) and COMPADRE version 6.21.8.0 (release date 20 August 2021).[4]

We first performed automated screening to identify potential LTRE analyses. We excluded matrices with missing values, and required matrices to be ergodic, primitive, and irreducible. We only considered studies with at least 2 matrices from the same species. We removed studies that did not measure fertility transitions, or

that did not observe fertility in any matrix (where each matrix corresponds to one time period, site or treatment) for the study populations. Furthermore, we removed studies that observed clonal reproduction of plants, unless the publication noted that clonal reproduction was rare. Because such studies often do not have a consistent definition of an individual (i.e. genets vs. ramets), estimates of population growth rate from the resulting matrix model are less reliable.

These automated screening steps (given in the file comadre-compadre_inventory.R) generated spreadsheets of matrices that could potentially be included in LTRE analyses, which were then manually screened. We eliminated matrices that represented post-hoc manipulations of collected data rather than unique measurements (e.g. a hypothetical scenario where fertility is reduced by 10% and adult survival increased by 10%), matrices where not all life stages were measured (for example, where seed viability was measured in 1 year and applied to multiple years), as well as matrices with errors in the definition of the life cycle (Kendall et al., 2019).[5] We identified all possible LTRE decompositions for the set of matrices from each study and species. This could include multiple fixed and random design comparisons, across time, space, or experimental treatments. In the plant matrices, we encountered issues with calculating lifespan and generation time due to apparent immortality, arising from very high survival (including retrogression) in the oldest/largest class. Among the 2181 plant population matrices that we selected for LTRE analyses, there were 478 (22%) where the final column sum of the survival matrix **U** was above 0.99. We decided to adjust these survival matrices by scaling down the final column so that the survival probability matched a more reasonable estimate from the same study and species. More details on these corrections are in Appendix D.

We set the maximum interaction order for these analyses to three for the exact LTRE, meaning that we calculate contributions from all main effects, two-way interactions, and three-way interactions, and the sum of all higher-order (4+ order) interactions is calculated as one additional term.

5 | RESULTS: COMPARISON OF APPROXIMATE AND exact LTRE

We calculated 186 LTREs for animal species, including 97 fixed design and 89 random design LTREs. These came from 76 species/study combinations (note that a single species could be in multiple studies, and a single study could include multiple species). The fixed design LTREs included 52 species, and the random design LTREs included 67. For fixed design, the majority (72%) of the approximate analyses had an overall error of less than 0.05 (Figure 2a).[6] The highest error for fixed design LTREs on animals was 4.63, in a decomposition of the effect of pollutants on a laboratory population of *Caenorhabditis elegans*. Likewise for the random design LTREs for animal species, the overall error of the classical method was below 0.01 for most of the analyses (79%; Figure 3a).[7] The three largest errors for the random design LTREs on animals were 0.80 and 0.33 for two temporal decompositions on woolly sculpin (*Clinocottus analis*) and 0.23 for a comparison among treatment levels for the same model of *C. elegans* that had high errors in fixed design LTRE.

For plant species, we calculated 1487 LTREs, 643 fixed design and 844 random design. In total there were 209 species/studies combinations. Of these species, 136 were represented in the fixed design LTREs and 205 were represented in the random design LTREs. This sample size was much larger than for animals, and we observed more large errors (Figures 2 and 3), and a larger range in the covariates (Figures S1 and S2). The maximum overall error observed in fixed design LTREs on plants was 2.27 in a herbivory exclusion experiment for Tideström's lupine (*Lupinus tidestromii*). In fixed design LTREs on plants, only 41% of the decompositions we examined had an error of 0.05 or less. Likewise for classical LTREs on plants, 41% of the LTREs examined had an overall error under 0.01. The plant random design LTRE with the highest error was 5.9 in a temporal comparison for the upright prairie coneflower (*Ratibida columnifera*).

Truncation errors seemed to be driving overall error in fixed design LTRE, while approximation and truncation errors played similar roles in driving overall error in random design LTRE. This suggests that second-order terms, which are present in random design LTRE but not in fixed design, are very important. In fixed design LTREs, we found that the distribution of overall errors more closely follows the distribution of truncation rather than approximation errors (Figure 2). The correlation between overall and truncation error (0.937 and 0.966 for animals and plants, respectively; Figure S3) is stronger than that between overall and approximation error (0.861 and 0.269 for animals and plants, respectively). On the other hand, in random design LTREs, the distributions of overall, approximation, and truncation errors look more similar to one another, without an indication that one kind of error is playing a stronger role than the other. Likewise, the correlation between approximation and truncation errors in plant LTREs is stronger and much closer to the 1:1 line in random design ($R^2 = 0.387$; Figure 3h) than in fixed design ($R^2 = 0.124$; Figure 2h).

We found that our reduced decomposition (up to 3-way interaction terms) performed very well, with the exception of random design LTREs on plant populations. We calculated the proportion of the true difference or variance in λ for all the LTREs that had at least 4 matrix elements that varied. We specified that higher-order terms were too large to neglect if the 4+ term accounted for 10% or more of the total difference or variance in λ. We found that higher-order terms could not be neglected in 3% of fixed design LTREs on animals, 6% of fixed design LTREs on plants, and 4% of random design LTREs on animals. However, the 4+ term contributed more than 10% of the variance in λ for 32% of the random design LTREs in plants. In fact, the 4+ term contributed more than 50% of the variance in 13% of the random design LTREs in plants.

We tested a number of covariates that could be predictive of large errors, and the only relationships that seem notable are those that arise from the Taylor expansion. For fixed design LTREs in both plants and animals (Figures S1 and S4), the overall error was

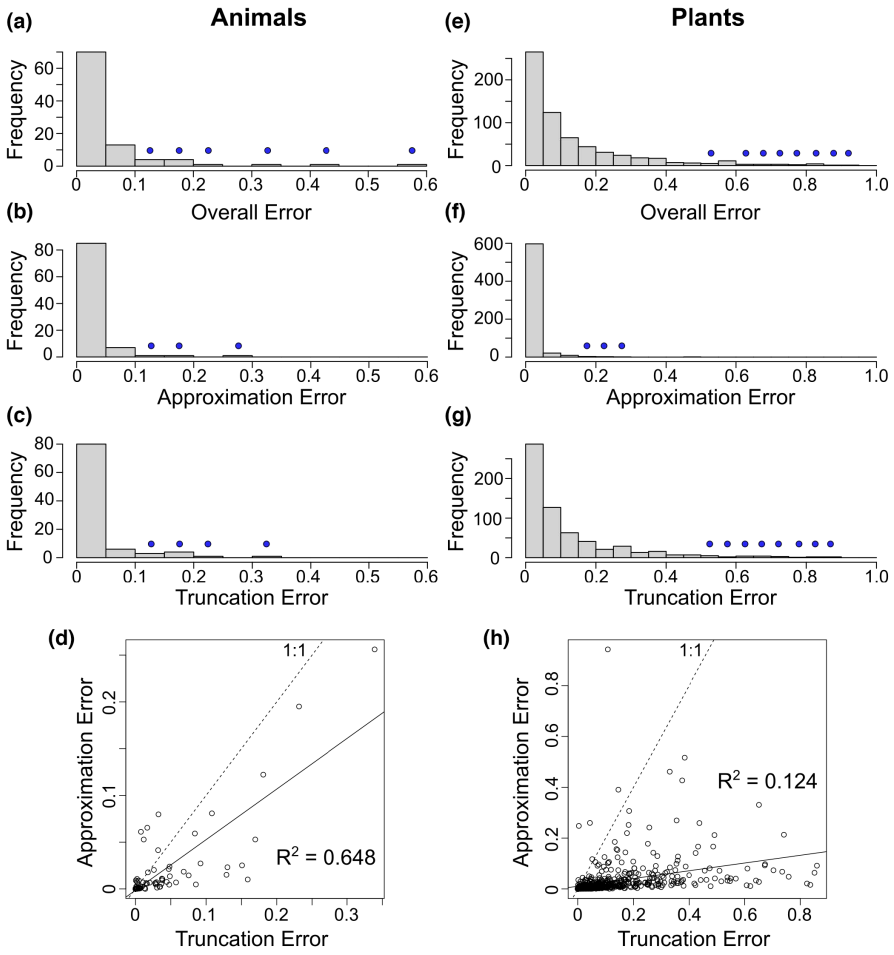

FIGURE 2 Error of the classical method for fixed design Life Table Response Experiment (LTRE) for animals (left column; panels a, b, c, and d) and plants (right column; panels e, f, g, and h). The overall error (a, e) is defined as the distance (using the 1-norm) between the contribution vectors for the exact and classical LTRE methods (Equation 10). The approximation error (b, f) is the 1-norm of the difference between terms that are present in both methods. The truncation error (c, g) is the 1-norm of terms that are neglected by the classical method. Histogram bins have width of 0.05. To aid visualization, circles are plotted above bins with 1–5 observations. In panels d and h, we show the relationship between truncation and approximation errors and a linear best-fit line (solid line). The dashed line is the 1:1 line. For animals, 2 (out of 97) LTREs with $E > 1$ were excluded; for plants, 9 (out of 643) LTREs with $E > 1$ were excluded.

positively correlated with both the distance between matrices ($\tau = 0.526,0.380$ for animals, plants) and the absolute value of the difference in λ ($\tau = 0.537,319$ for animals, plants). Similarly, in random design LTREs (Figures S2 and S5), the overall error was positively correlated with both the sum of variances of matrix elements ($\tau = 0.514,0.361$ for animals, plants) and the standard deviation of λ ($\tau = 0.651,0.524$ for animals, plants). These are all, essentially, measures of how far apart the matrices being tested are—when the matrices being tested are far apart, the Taylor expansion will lead to greater errors. For random design LTREs in plants, there was also a negative relationship between generation time and overall error (Figure S2G, $\tau = -0.400$). In very long-lived and late-maturing (so presumably slow-growing) species, the time scale of variation in matrix elements may be much longer than the observation period, such that species with long generation times are more likely to have matrices that are very close together. In accordance with this, we observed a negative relationship between generation time and the sum of variances in matrix elements (Figure S2H, $\tau = -0.252$).

Importantly, the relative error cannot be considered a reliable measure of how well the classical method has performed. There is a weak correlation between the relative error (calculated only using results from the classical method) and our overall error (E) that measures the distance between the classical and exact methods. The Pearson correlation coefficient ranges from 0.02 to 0.48 (Figure 4). In the random design LTREs on animal projection matrices, there does seem to be a weak relationship. However, with the much larger sample size in plant projection matrices, that relationship does not hold.

The first reason that relative error can be a poor proxy for accuracy of the classical method is compensation between approximation and truncation errors, that is, some effects of interactions are instead attributed to main effects. This is generally what will happen in a classical fixed design LTRE when the pivot matrix is selected such that the main effect terms sum exactly to the difference in λ. In a fixed design LTRE for the sand olive (*Dodonaea angustifolia*; Bekele, 2000), the relative error was 2.2% (Figure 5). The largest contributions are four first-order terms, two of which are underestimated by the classical method. In the exact LTRE, the larger first-order terms are offset by negative interaction terms. The exact and classical (approximate) LTREs give the same total effect $\Delta\lambda$, but do so for different reasons. We saw a similar pattern in another fixed design LTRE for sand olive (Figure S7) and one for *Eisenia fetida* (Figure S8).

The other reason that relative error can be a bad proxy is when large positive and negative contributions from different matrix elements or interactions nearly cancel out, giving a small difference or variance in λ.This can result in a very large relative error in the overall response, despite a small absolute error and uniformly small errors in each of the contributions. This is seen in a fixed design LTRE for ground squirrels (*Spermophilus armatus*, Oli et al., 2001), where the

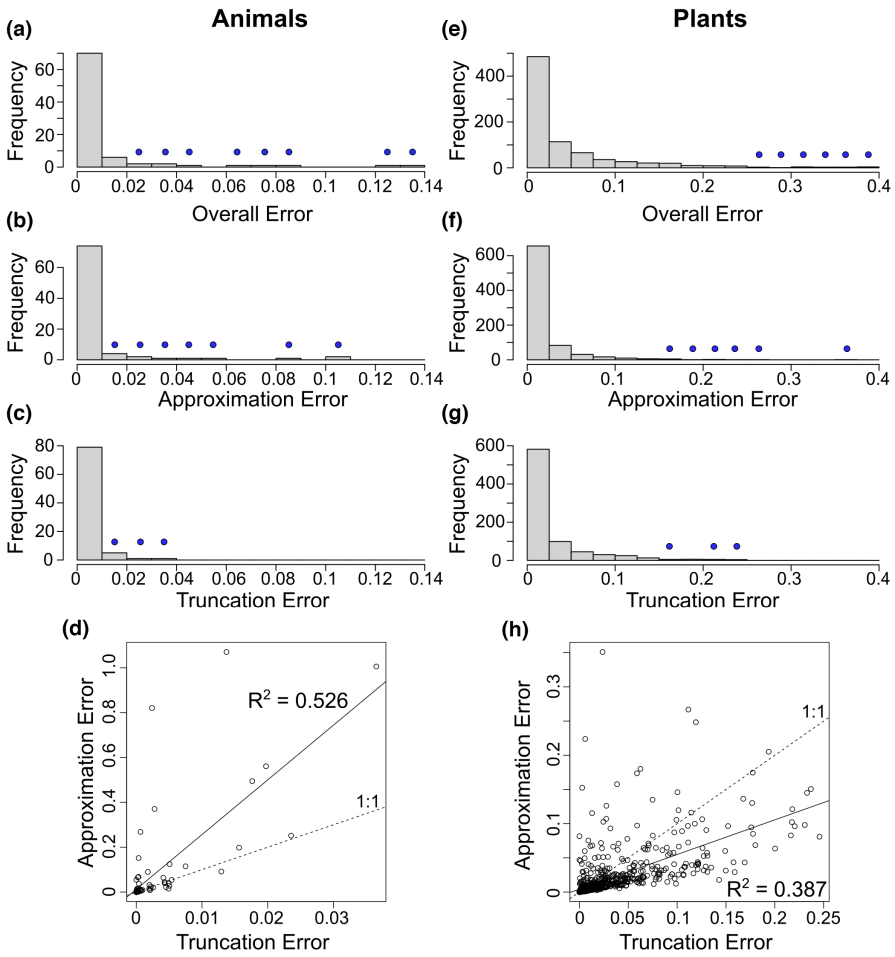

FIGURE 3 Error of the classical method for random design Life Table Response Experiment (LTRE), as in Figure 2. Histogram bins have a width of 0.01 for the left column (animals; panels a, b, c, and d) and a width of 0.025 for the right column (plants; panels e, f, g, and h). To aid visualization, circles are plotted above bins with 1–5 observations. For animals, 3 (out of 89) LTREs with $E > 0.2$ were excluded; for plants, 29 (out of 843) LTREs with $E > 0.4$ were excluded.

main effects almost entirely cancelled one another out, as did the contributions of two-way interactions (Figure S9). The sum of main effect (order 1) contributions was two orders of magnitude smaller than the main effects themselves. As a result, despite close agreement in interpretation between the classical and exact methods, the relative error of the classical method is large (in fact, the sum of approximate contributions has the opposite sign as $\Delta\lambda$). Similarly, in a random design LTRE for garlic mustard (*Alliaria petiolata*), the observed variance in λ is very tiny, and the relative error is nonsensically high: 674,042% (Figure S10). Most of the main-effect contributions are positive, while most of the two-way interaction contributions are negative. The classical method overestimates the largest first-order term and underestimates the two largest second-order terms, and the neglected third-order terms are as large as some of the first- and second-order terms. But again, relative error in the overall response to the variation in all matrix elements is not informative about how well the classical method estimates the individual contributions.

6 | DISCUSSION

We have introduced an exact version of LTREs, a widely used demographic analysis tool. The classical methods in use for several decades rely on Taylor series approximations that require lower computational effort, at the cost of approximating lower-order terms and neglecting higher-order interactions, which can be biologically meaningful.

We found that errors arising in the classical method were usually small (Figures 2 and 3). But the error distributions had long tails—the error is sometimes large. The overall error was most strongly related to the distance between/among the matrices being analysed: the Taylor series is less accurate when the matrices being compared are very different (Figures S1, S2, S4, and S5). In fixed design LTREs, the overall error was driven strongly by truncation errors, because the classical one-way fixed design LTRE evaluates only main effects (Figure 2; Figure S3). In random design LTREs, truncation and approximation errors contributed more evenly to overall error (Figure 3; Figure S6). This difference between fixed and random design LTREs underscores the importance of the two-way interaction terms.

While the two- and three-way interaction terms are likely to make important contributions to changes in λ, the 4+ order term tended to be very small. In most cases that we investigated, calculating the exact contributions of main effects, two-way interactions, and three-way interactions was sufficient for decomposing 90% or more of the total effects on λ. This does not guarantee that the higher-order terms are small, but it suggests that they are either small or cancel one another out. However, we note that many of the random design LTREs on plants (32% of the studies with 4 matrix

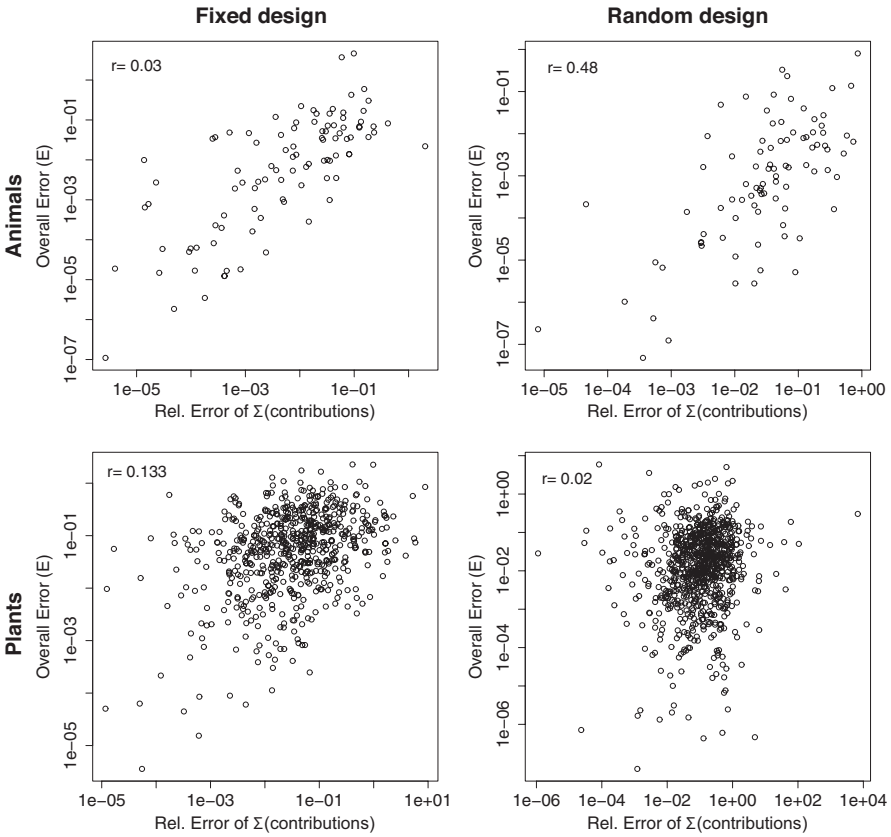

FIGURE 4 Relationship between the relative error of the classical method, and the overall error between the classical and exact methods. The relative error compares the sum of contributions from the classical (approximate) Life Table Response Experiment to the observed difference or variance in λ (Equation 13). The overall error is the distance between the vectors of contributions from the classical and exact methods (Equation 10). All axes are log-scaled.

elements or more) had a large discrepancy between the sum of contributions up to three-way interactions and the observed variance in λ (i.e. the 4+ term was >10% of var(λ)).

In past LTRE applications using the classical methods, relative error was sometimes (but not universally) reported as a measure of how well the approximation performed. Our results suggest that relative error is not a reliable measure. When there is compensation between approximation and truncation errors, important (but neglected) interaction effects may be subsumed into the approximate contributions, causing the sum of approximate contributions to be very close to the observed difference or variance—thus the classical method will have a small relative error but incorrect interpretation. Alternatively, small relative term-by-term errors can cause a very large relative error when terms are large but trade off against one another, such that the true difference or variance of λ is very small—in this case, the classical method will have a large relative error but generally correct interpretation.

Ultimately, interpretation is more important than quantitative errors, because the goal of LTRE is to understand how observed variation in vital rates drives population dynamics via effects on λ (and, for that matter, effects on the structure of the stable population). The results of LTRE analyses can also inform conservation and management (Bruna & Oli, 2005; Oli et al., 2001; Oro & Doak, 2020), by providing a more mechanistic understanding of the effects of interventions or variation on population dynamics. Therefore, the numerical contribution of changes in each vital rate to λ may be less important than understanding the relative roles of all of the vital rates or stage classes. Exact LTRE removes any need to decide whether the classical LTRE results are reliable.

Another important way that classical LTRE can lead to faulty interpretation is when a directional experiment is analysed using a symmetric LTRE. In our meta-analysis we compared the results from the approximate fixed design method with a symmetric exact fixed design, because the standard application of the classical method is symmetric (Equation 2). However, we found that 58% and 36% of the published fixed design LTREs that we identified for animal and plant species, respectively, were directional in design. Using the exact directional rather than symmetric analysis will change the resulting contributions and is likely to change the interpretation of the difference in conditions. For example, in the ground squirrel LTRE discussed earlier, the two matrices being compared were before and after an experimental manipulation of population density, so a directional fixed design LTRE would be more appropriate. In Figure S11, we show that the contributions of the terms differ substantially between directional and symmetric LTRE. When the analysis matched the experimental design, we see important positive contributions from second-year fertility and survival, and important negative contributions from first-year survival and the two-way interaction of adult survival and first-year fertility.

We have applied LTREs at the level of matrix elements in our meta-analysis and R package, in line with how the classical methods are typically introduced (Caswell, 1989, 2001). There are two primary reasons for this. First, the matrices archived in the COM(P)ADRE databases record simply the numerical values of matrix elements,

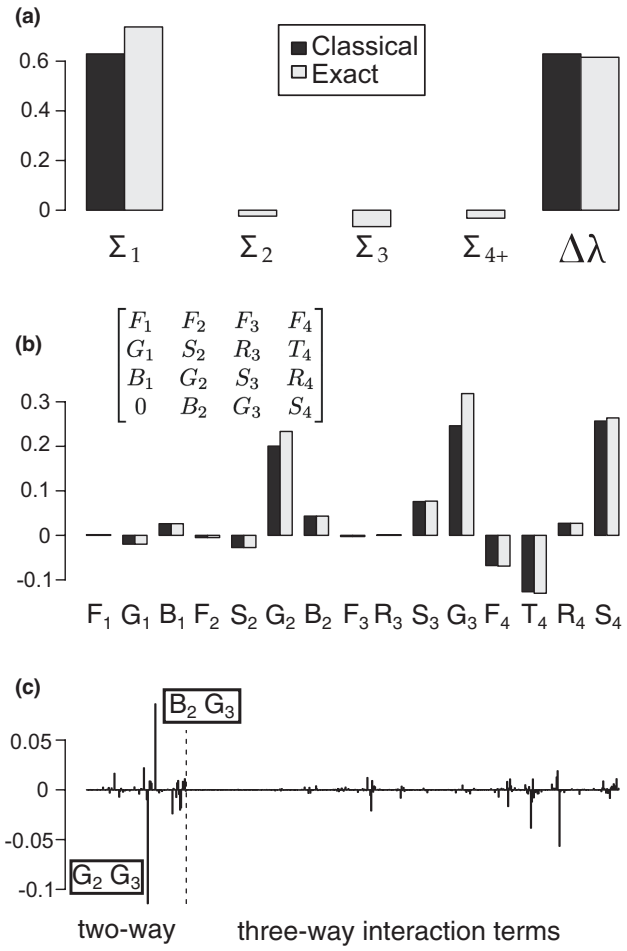

FIGURE 5 Comparison of the classical and exact methods for fixed design Life Table Response Experiment (LTRE) in *Dodonaea angustifolia*, comparing a protected open habitat patch and a disturbed slope habitat patch. (a) Comparison of the sum of terms by order; Σ_1 is the sum of all main effects (first-order terms); Σ_2 is the sum of all pairwise interactions (second order terms), and so on; $\Delta\lambda$ for the classical method is given by the sum of all estimated contributions. (b) Term-by-term comparison of the first-order terms for the classical and exact methods. (c) Two-way and three-way interaction terms which are present only in the exact LTRE. The inset matrix shows the structure of the projection matrix, with parameters corresponding to the labels in panels b and c. Note the difference in scale between the three panels. The overall error of the classical method was $E = 1.09$ and the relative error was 2.2%.

so significant effort would be required to look through the source materials and formulate matrix elements as functions of lower-level parameters. Secondly, analysis of lower-level parameters would either require unique functions for each matrix model, or more sophisticated functions utilizing symbolic programming of matrices based on lower-level parameters. For our meta-analysis and for users who are interested in performing LTRE at the level of matrix elements, our code is sufficient. However, for matrices that can be formulated in terms of underlying parameters, the decomposition can (and probably should) be carried out in terms of lower-level vital rates. This is particularly important when considering fertility parameters, which are not independent from adult survival in a post-breeding census design. Our methods can certainly be used to perform a decomposition of contributions from lower-level vital rates and our R package may provide some useful utilities for the calculation, but it would require bespoke analysis code.

In conclusion, we suggest that exact LTRE methods be preferred for future calculations of one-way fixed design and random design LTRE, because the exact methods allow direct calculation of the contributions of parameters and their interactions. Although the classical methods may often lead to accurate biological interpretation, the relative error of the classical LTRE method is an unreliable predictor of accuracy and therefore is there is no way to "check" the accuracy of the classical method. With the introduction of the exactLTRE package in R, together with existing tools for working with matrix population models (e.g. popdemo and Rage packages), the use of exact LTRE methods is straightforward and accessible. Given the utility of LTRE for management and conservation, we should aim for interpretations of observed population dynamics that are as accurate as possible. Future directions of this work could include extensions to analyse effects on other population metrics (e.g. generation time, R_0) in addition to λ, finer partitioning of interaction terms in random design LTRE into contributions from joint variation and from covariance (Ellner et al., 2019), and extension of the exactLTRE package to analyse integral projection models (IPMs).

AUTHOR CONTRIBUTIONS

All authors contributed to conceiving the project, discussed all aspects of the research, and contributed to writing and revising the paper. Christina M. Hernández wrote the code and R package, performed the meta-analysis and wrote the first draft.

ACKNOWLEDGEMENTS

This research was supported by US NSF grants DEB-1933497 (CMH, SPE and GH), DEB-1933612 (RES), and DEB-1933561 (PBA). We are very grateful to Rob Salguero-Gomez and the COMPADRE team for building and maintaining databases of matrix population models. We thank Hal Caswell and Yngvild Vindenes for their thoughtful peer reviews, as well as Megan Greischar, Tim Lambert, Anna Poulton, Jordi Ripoll, Gregor-Fausto Siegmund and Jürg Spaak for comments on earlier versions of the manuscript.

CONFLICT OF INTEREST STATEMENT

The authors declare none.

PEER REVIEW

The peer review history for this article is available at https://www.webofscience.com/api/gateway/wos/peer-review/10.1111/2041-210X.14065.

DATA AVAILABILITY STATEMENT

Code and data to repeat the analyses presented in this paper are available at https://doi.org/10.5281/zenodo.6461580 (Hernández et al., 2022).

ORCID

Christina M. Hernández https://orcid.org/0000-0002-7188-8217
Stephen P. Ellner https://orcid.org/0000-0002-8351-9734
Peter B. Adler https://orcid.org/0000-0002-4216-4009
Giles Hooker https://orcid.org/0000-0003-2648-1167
Robin E. Snyder https://orcid.org/0000-0002-6111-0284

ENDNOTES

[1] We present much more detailed background information, including the Taylor series expressions, in Appendix A.

[2] We note that computation times can become quite long for matrices with more than ~100 elements that vary and $k = 3$. Users should consider what terms are of interest to them and what computational time they are comfortable with. We found that an exact LTRE with $k = 3$ took only a few minutes on a good 2022 desktop computer for the maximum size included in our meta-analysis: 59 matrix elements that vary.

[3] This λ value is high but not unrealistic for a laboratory population that is well-fed and protected from predation and disease.

[4] The database corrections produced a few changes in matrix identification numbers. The identification numbers provided in the Supplemental Data Files comadre_ltres_torun.csv and compadre_ltres_torun.csv match the latter versions of the databases (i.e. the versions of the databases that we used for calculations).

[5] In the course of screening, we manually corrected matrices from three animal species where the publication represented the life cycle correctly but there had been a mistake in digitization: *Puma concolor*, *Ursus americanus* subsp. *floridanus* and *Esox lucius*. See exactLTRE_LoadDatabases.R for details of these manual corrections.

[6] The error ε is in the same units as λ or $\mathrm{var}(\lambda)$, and is the sum of the absolute value of differences between the results of the exact and classical methods. To help with interpreting these error values, we have stated the value of ε for all of the figures for example LTRE comparisons (Figure 5; Figures S7–S10).

[7] Note that, in general, values of $\mathrm{var}(\lambda)$ and ε in random design LTREs tended to be much smaller than values of $\Delta\lambda$ and ε in fixed design LTREs.

REFERENCES

Bekele, T. (2000). *Plant population dynamics of Dodonaea angustifolia and Olea europaea ssp. cuspidata in dry Afromontane forests of Ethiopia* (PhD thesis). Acta Universitatts Upsaliensis.

Bienvenu, F., & Legendre, S. (2015). A new approach to the generation time in matrix population models. *American Naturalist*, *185*, 834–843.

Brault, S., & Caswell, H. (1993). Pod-specific demography of killer whales (*Orcinus orca*). *Ecology*, *74*, 1444–1454.

Bruna, E. M., & Oli, M. K. (2005). Demographic effects of habitat fragmentation on a tropical herb: Life-table response experiments. *Ecology*, *86*, 1816–1824.

Caswell, H. (1989). Analysis of life table response experiments: Decomposition of effects on population growth rate. *Ecological Modelling*, *46*, 221–237.

Caswell, H. (1996). Analysis of life table response experiments II. Alternative parameterizations for size- and stage-structured models. *Ecological Modelling*, *88*, 73–82.

Caswell, H. (2001). *Matrix population models* (2nd ed.). Sinauer Associates.

Ellner, S. P., Snyder, R. E., Adler, P. B., & Hooker, G. (2019). An expanded modern coexistence theory for empirical applications. *Ecology Letters*, *22*, 3–18.

Fréville, H., & Silvertown, J. (2005). Analysis of interspecific competition in perennial plants using life table response experiments. *Plant Ecology*, *176*, 69–78.

Hernández, C. M., Ellner, S. P., Adler, P. B., Hooker, G., & Snyder, R. E. (2022). *Supplemental data and code for "an exact version of life table response experiment analysis, and the R package exactLTRE"*. Zenodo.

Hooker, G. (2007). Generalized functional ANOVA diagnostics for high-dimensional functions of dependent variables. *Journal of Computational and Graphical Statistics*, *16*, 709–732.

Kendall, B. E., Fujiwara, M., Diaz-Lopez, J., Schneider, S., Voigt, J., & Wiesner, S. (2019). Persistent problems in the construction of matrix population models. *Ecological Modelling*, *406*, 33–43.

Oli, M. K., Slade, N. A., & Dobson, F. S. (2001). Effect of density reduction on Uinta ground squirrels: Analysis of life table response experiments. *Ecology*, *82*, 1921–1929.

Oro, D., & Doak, D. F. (2020). Breeding transients in capture–recapture modeling and their consequences for local population dynamics. *Scientific Reports*, *10*, 15815.

Poelwijk, F. J., Krishna, V., & Ranganathan, R. (2016). The context-dependence of mutations: A linkage of formalisms. *PLoS Computational Biology*, *12*, 1–19.

Salguero-Gómez, R., Jones, O. R., Archer, C. R., Bein, C., de Buhr, H., Farack, C., Gottschalk, F., Hartmann, A., Henning, A., Hoppe, G., Römer, G., Ruoff, T., Sommer, V., Wille, J., Voigt, J., Zeh, S., Vieregg, D., Buckley, Y. M., Che-Castaldo, J., … Vaupel, J. W. (2016). COMADRE: A global data base of animal demography. *Journal of Animal Ecology*, *85*, 371–384.

Salguero-Gómez, R., Jones, O. R., Archer, C. R., Buckley, Y. M., Che-Castaldo, J., Caswell, H., Hodgson, D., Scheuerlein, A., Conde, D. A., Brinks, E., Buhr, H., Farack, C., Gottschalk, F., Hartmann, A., Henning, A., Hoppe, G., Römer, G., Runge, J., Ruoff, T., … Vaupel, J. W. (2015). The COMPADRE plant matrix database: An open online repository for plant demography. *Journal of Ecology*, *103*, 202–218.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

Appendix S1. This file contains additional text and figures, arranged as: Appendices (A) Background on LTRE and fANOVA; (B) Proof of the parity formula for fANOVA effects; (C) A computationally efficient method for calculating all possible interactions in an exact LTRE; (D) Corrections to the survival matrices U; (E) Exact two-way factorial design; and (F) Supplemental Figures and Tables.

How to cite this article: Hernández, C. M., Ellner, S. P., Adler, P. B., Hooker, G., & Snyder, R. E. (2023). An exact version of Life Table Response Experiment analysis, and the R package exactLTRE. *Methods in Ecology and Evolution*, *14*, 939–951. https://doi.org/10.1111/2041-210X.14065