

EFECT – A Method and Metric to Assess the Reproducibility of Stochastic Simulation Studies.

T.J. Sego<sup>1\*</sup>, Matthias König<sup>2</sup>, Luis L. Fonseca<sup>1</sup>, Baylor Fain<sup>1</sup>, Adam C. Knapp<sup>1</sup>, Krishna Tiwari<sup>3</sup>, Henning Hermjakob<sup>3</sup>, Herbert M. Sauro<sup>4</sup>, James A. Glazier<sup>5</sup>, Reinhard C. Laubenbacher<sup>1</sup>, Rahuman S. Malik-Sheriff<sup>3,6†</sup>

<sup>1</sup>Department of Medicine, University of Florida, Gainesville, FL, USA

<sup>2</sup>Institute for Theoretical Biology, Humboldt University Berlin, Berlin, Germany

<sup>3</sup>European Molecular Biology Laboratory, European Bioinformatics Institute (EMBL-EBI), Hinxton, Cambridge, UK

<sup>4</sup>Department of Bioengineering, University of Washington, Seattle, WA, USA

<sup>5</sup>Department of Intelligent Systems Engineering and Biocomplexity Institute, Indiana University, Bloomington, IN, USA

<sup>6</sup>Department of Surgery & Cancer, Faculty of Medicine, Imperial College London, London, UK

\* timothy.sego@medicine.ufl.edu

† sheriff@ebi.ac.uk

## Abstract

Reproducibility is a foundational standard for validating scientific claims in computational research. Stochastic computational models are employed across diverse fields such as systems biology, financial modelling and environmental sciences. Existing infrastructure and software tools support various aspects of reproducible model development, application, and dissemination, but do not adequately address independently reproducing simulation results that form the basis of scientific conclusions. To bridge this gap, we introduce the Empirical Characteristic Function Equality Convergence Test (EFECT), a data-driven, statistically robust method to quantify the reproducibility of stochastic simulation results. EFECT employs empirical characteristic functions to compare reported results with those independently generated by assessing distributional inequality, termed EFECT error, a metric to quantify the likelihood of equality. Additionally, we establish the EFECT convergence point, a quantitative metric for determining the required number of simulation runs to achieve an EFECT error value of *a priori* statistical significance, setting a reproducibility benchmark. EFECT supports all real-valued and bounded results irrespective of the model or method that produced them, and accommodates stochasticity from intrinsic model variability and random sampling of model inputs. We tested EFECT with stochastic differential equations, agent-based models, and Boolean networks, demonstrating its broad applicability and effectiveness. EFECT standardizes stochastic simulation reproducibility, establishing a workflow that guarantees reliable results, supporting a wide range of stakeholders, and thereby enhancing validation of stochastic simulation studies, across a model's lifecycle. To promote future standardization efforts, we are developing open source Stochastic Simulation Reproducibility software library (libSSR) in diverse programming languages for easy integration of EFECT.

## Introduction

Various disciplines have acknowledged a crisis of reproducibility in science (Baker, 2016), to which computational research is not immune (Stodden et al., 2018). Reproducibility has been argued to be the minimum standard for assessing scientific claims in computational science (Peng, 2011). Reproducibility in stochastic modeling is essential for ensuring that environmental predictions and assessments are reliable and can be consistently validated across different studies.

In the life sciences, efforts towards computational model reproducibility have been enumerated in several different fields, including computational neuroscience (McDougal et al., 2016), quantitative systems pharmacology (Kirouac et al., 2019), and computational drug discovery (Schaduangrat et al., 2020). Basic stages of reproducible model results have been proposed for modeling communities like systems biology modeling (Porubsky & Sauro, 2023), for which there are supporting tools for model construction, annotation, execution, and data exchange with varying degrees of support for reproducible models (Ghosh et al., 2011; Shin et al., 2023; Sordo Vieira & Laubenbacher, 2022). Resources like BioModels, the world's largest repository of curated biological models (Malik-Sheriff et al., 2020), and BioSimulators, a free online registry of biological model simulators (Shaikh et al., 2022), leverage foundational infrastructure like the Systems Biology Ontology (SBO, (Courtot et al., 2011)), Systems Biology Markup Language (SBML, (Keating et al., 2020)), and Simulation Experiment Description Markup Language (SED-ML, (Waltemath et al., 2011)) to promote reusability and reproducibility of biological models. Such ecosystems enable practices that assess and improve the credibility of individual tools, and thus of the models they support, like verifying that ecosystem tools produce the same outputs for the same inputs (Bergmann & Sauro, 2008). The same can be done with reported models. A 2021 BioModels study reported that results from 49% of 455 assessed BioModels entries could not be directly reproduced (Tiwari et al., 2021).

Recommendations for best practices in reproducible computational research have stated that, at minimum, a researcher should be able to reproduce their own results (Sandve et al., 2013). Proposed best practices in reproducible systems biology models have recommended that all structured unprocessed simulation results be stored, all data underlying reported graphs and tables be shared, and model verification and validation be automated and documented (Porubsky et al., 2020), even to calling for journal editors to enforce publication of raw simulation data (Mendes, 2018). Other proposed guidelines for reproducible systems biology models have recommended that researchers should be able to regenerate statistically identical simulation results, and that workflows should be expanded to verify the statistical repeatability of simulation results (Medley et al., 2016). The majority of reproducibility guidelines and recommendations are based on deterministic models and analyses. However, advanced model-based research across diverse fields such as systems biology, financial modeling, and environmental sciences employs the use of stochastic models to capture complex, random processes that deterministic models cannot adequately represent (Casdagli, 1992; Friedrich et al., 2011; Goldenfeld & Kadanoff, 1999).

In emerging technologies like digital twins that will likely use stochastic modeling, there are recent calls for verification and validation of simulation results (Fuller et al., 2020). Likewise, recent emphasis has been placed on the need for strong focus on results distributions during model validation (Read et al., 2020). Reproducibility in stochastic modeling is essential for ensuring that environmental predictions and assessments are reliable and can be consistently

validated across different studies. For example, reproducibility is particularly important for climate models where variability and uncertainty must be systematically accounted for to enhance credibility (Stainforth et al., 2005). By reproducible, we mean whether one obtains the same results as those produced by an experiment when conducting an independent study with procedures as closely matched as possible (*i.e.*, reproducibility of results, (Goodman et al., 2016; Mendes, 2018)). By stochastic results, we mean results that are not deterministic whether due to intrinsic stochasticity or stochastic sampling of model inputs (*e.g.*, randomly sampling of parameters, initial conditions). The inherent variability in stochastic models, where each simulation run can yield different results, necessitates the development of approaches and metrics to quantify reproducibility. Establishing such standards is crucial for ensuring that findings from stochastic simulations are reliable and can be consistently validated across different studies.

Reproducing individual trajectories of stochastic results is theoretically possible with pseudorandom number generators, though not without potential pitfalls due to differences in software implementations (Medley et al., 2016). More importantly, such an approach has at least two major limitations: 1) reproducing individual trajectories does not account for under- or selective model sampling, inhibiting improvement to the strength of conclusions drawn from a model by means of reproducing results motivating them; and 2) storing and sharing data for all individual trajectories does not scale to community-level practice like those associated with BioModels. For comparing stochastic results by their distribution, tools exist to compare summary statistics of results distributions but require decision-making or assumptions on model results (*e.g.*, normality) that diminish relevance outside of engineering contexts. For example, SBML Test Suite (Evans et al., 2008) appropriately uses summary statistics for comparing stochastic results produced by SBML model simulators when suitable test data is chosen *a priori*. However, summary statistics like mean and standard deviation are not unique and so are unreliable for general identification of differences in stochastic model results (*e.g.*, when initial conditions have different distributions with the same summary statistics, Source Code S3.21). Previous work developed a test and heuristics to quantify aleatoric uncertainty for determining a minimum sample size that sufficiently represents a stochastic model (Reiczigel et al., 2005; Vargha & Delaney, 2000) but did not provide a means to support standardized reproduction of results.

In this work we developed a method and metrics spanning the lifetime of a model by which stochastic simulation results can be quantitatively reproduced. Our method tests for equality of stochastic results by evaluating whether results converge to the same unique, unknown distribution using empirical characteristic functions. Hence, we call our method the Empirical Characteristic Function Equality Convergence Test (EFECT). We showed the broad relevance of our method by demonstrating its application with multiple stochastic modeling applications, modeling methodologies, and BioModels entries. Our method quantifies differences in stochastic simulation results independent of modeling method and with universal meaning such that outcomes from reproducing results are generally interpretable. We designed a workflow built on our method to accommodate data exchange within modeling communities for stochastic model reproducibility through outlets like online repositories and journal publications.

## Results

EFFECT-based reproduction of stochastic simulation results consists of two components: first, generation of reliable and reproducible stochastic simulation results by the modeler; second, testing whether the stochastic simulation results have been sufficiently reproduced using a well-defined quantitative metric by a model curator. We refer to the individual who produces the results as the Modeler and to the one who independently analyzes the model and reproduces the results as the Curator. We require that EFFECT assesses similarity of results using a quantitative metric, called the *EFFECT error*, with meaning that is both independent of a model's mathematical construct and scale of results, and robust to support a broad range of models and sources of results' stochasticity. In particular, we require that EFFECT does not rely on any knowledge about the model from which results were produced. Likewise, we require that the EFFECT error supports results with variance produced by simulation stochasticity and/or stochastic selection of model parameters and initial conditions. EFFECT assesses the similarity of stochastic simulation results by testing for equality in distribution of results through comparison of empirical characteristic functions (ECFs) generated from results distributions. Throughout, we refer to a set results collected from one or more executions of a stochastic simulation as a *simulation sample*, and likewise to the number of executions as the *sample size*. The EFFECT error quantifies similarity of simulation samples over a model- and scale-independent range, where an EFFECT error of zero demonstrates equality of simulation samples, and of two demonstrates maximum difference. Using EFFECT, a Modeler tests whether the distribution of their simulation sample is reproducible by randomly dividing their simulation sample into two subsamples of equal size and measuring the EFFECT error of the subsamples. The Modeler improves the reproducibility of their simulation sample by repeating this process while increasing sample size until the distribution of the sampled EFFECT error satisfies a reproducibility criterion, called the *EFFECT convergence point*. We refer to the minimum simulation sample size to satisfy the EFFECT convergent point as the *EFFECT sample size*. Given the size of the subsamples, the ECFs generated from one of them, and some additional information (described below) accompanying a reported model, a Curator tests whether they can reproduce a simulation sample by comparing to ECFs from their own implementation of the reported model. A Curator tests a reported simulation sample to their own for equality in distribution according to the null hypothesis that samples of the EFFECT error when comparing simulation samples and when quantifying reproducibility are taken from the same unknown distribution. A Curator assesses whether they reproduce a simulation sample by rejecting the null hypothesis at an *a priori* significance level according to Chebyshev's inequality for unknown population mean and variance. For details, see *Materials and methods*.

### *Similarity Quantification of Stochastic Simulation Results*

We first tested the capability of EFFECT to quantify the similarity of simulation samples using a variety of stochastic differential equation (SDE) models with different sources of stochasticity (Figure 1). Test models consisted of either deterministic simulation while sampling model parameters and/or initial conditions using a normal distribution for one (Source Code S3.1, S3.6, S3.11, S3.12; (A. M. Smith & Smith, 2016)), two (Source Code S3.2), and ten (Source Code S3.3) parameters and/or initial conditions, or using a uniform distribution sampling for one (Source Code S3.5, Source Code S3.11) or ten (Source Code S3.4) parameters and/or initial conditions; or of stochastic simulation using the Gillespie Stochastic Simulation Algorithm (Source Code S3.10, S3.13; (Gillespie, 1976; Wearing et al., 2005)), including three BioModels entries (Source Code S3.7, S3.8, S3.9; (Kareva et al., 2010; Liu et al., 2012; Lo et al., 2005)). For all test models subject to a source of stochasticity, we increased sample size and quantified

reproducibility of the simulation sample by evaluating the similarity of equally sized subsamples as previously described for our Modeler’s test for reproducibility. Similarity of simulation samples was quantified using the EFECT error defined in [ 3 ]. We found that the EFECT error decreased with increasing sample size for all modeler’s test cases according to a power law. A representative result of this general observation is shown in Figure 1D.

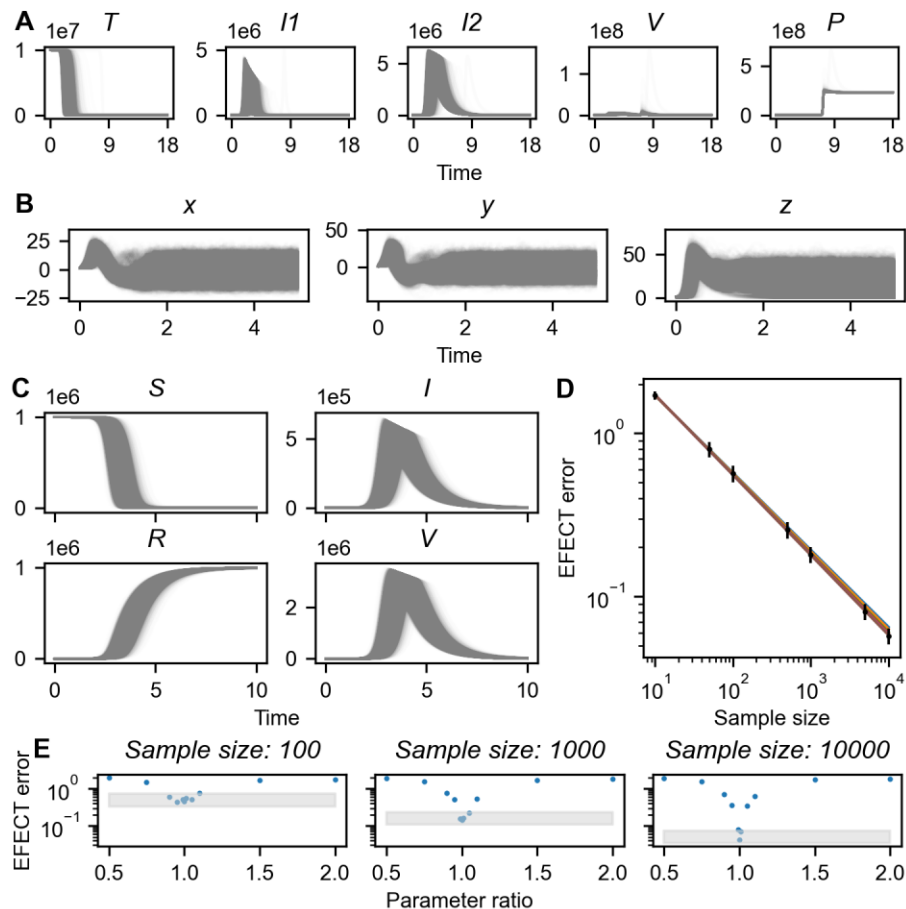


Figure 1. Testing similarity of stochastic biological simulation results from ordinary differential equation models. A: Trajectories while randomly sampling a parameter of a model of viral-bacterial coinfection (Source Code S3.1). Sample size of 10k trajectories shown. B: Trajectories while simulating a Lorenz system with Gillespie stochastic simulation algorithm (Source Code S3.10). Sample size of 10k trajectories shown. C: Trajectories while randomly sampling a parameter of a model of viral infection. Sample size of 10k trajectories shown. D: EFECT error quantifying reproducibility of a sample from the model shown in panel C as a function of sample size. An EFECT error of zero demonstrates equality, and of two demonstrates complete dissimilarity. Lines show fits of a typical log-linear relationship, where similarity increases with increasing sample size. Dots show mean evaluations of the EFECT error. Bars show one standard deviation of EFECT error evaluations upward and downward. E: EFECT error per sample size from the model shown in panel C while comparing samples from models of the same structure but a difference in parameter (as measured by the parameter ratio). Shaded regions show the mean  $\pm$  three standard deviations of the EFECT error when testing for reproducibility.

We then tested the capability of EFECT to detect structural and parametric differences in SDE models by comparing simulation samples generated from them. We tested detection of structural differences by comparing simulation samples of various sample sizes generated from different models. When comparing simulation samples from models with structural differences, we found that the EFECT error does not decrease with increasing sample size but instead tends to measure near maximum difference (Source Code S3.14). We tested detection of parametric

differences by comparing simulation samples of various sizes generated from the same model but with a difference in the distribution from which a parameter value was drawn. Specifically, we assessed the likelihood of a false-positive according to [ 6 ] when testing for equality in distribution for a given sample size and difference in parameter distribution. Generally, EFECT can detect a difference if it is unlikely to yield a false-positive. In this test, we perform the following procedure: Step 1: generate a simulation sample; Step 2: quantify reproducibility of the simulation sample; Step 3: extract a subsample of half the size (as would be done when reporting a model); Step 4: generate a second simulation sample of the same half size and with a modified parameter distribution; and Step 5: test how well the simulation sample from Step 4 reproduces the simulation sample from Step 3.

For the model shown in Figure 1C, we scaled the mean and standard deviation of the normal distribution for a single parameter by factors ranging from 0.5 to 2. Using a significance level of 0.05, we found that EFECT could detect differences in distributions for a sample size of 10,000 (reproducibility test EFECT error of  $0.0537 \pm 6.47 \times 10^{-3}$ ) when scaling downward or upward by as little as 5% to 1%, respectively (p-value of  $6.00 \times 10^{-4}$  and 0.0258 for 5% downward and 1% upward scaling of parameter distribution mean and standard deviation, respectively). The same test and model with a sample size of 1,000 (reproducibility test EFECT error of  $0.168 \pm 0.0192$ ) could detect 5% scaling differences (p-value of 0.0220 and  $7.98 \times 10^{-3}$  during 5% downward and upward scaling, respectively). With a sample size of 100 (reproducibility test EFECT error of  $0.534 \pm 0.0634$ ), the same test and model could detect 25% downward or 50% upward scaling differences (p-value of 0.0196 in both cases). Further tests with ODE models performed the same analysis with different ODE models (detected 5% or more scaling differences for sample size 10,000, Source Code S3.18), with constant parameters with a scaling difference and stochastic simulation (detected 5% or more scaling differences for sample size 10,000, Source Code S3.19), or while individually scaling distribution mean (detected 5% or more scaling differences for sample sizes 5,000 and 10,000, Source Code S3.17) or variance (detected scaling differences of 25% or more downward, and 50% or more upward, for sample sizes 5,000 and 10,000, Source Code S3.20). In all tests, upward and downward scaling differences of between 5% and 50% were detectable with a significance level of 0.05 when the reproducibility test EFECT error mean was below 0.1. Hence, EFECT can detect even a single parametric difference between models by comparing samples from them, and increasingly so with decreasing EFECT error.

To demonstrate whether EFECT supports reproducibility of stochastic results from methods other than SDEs, we tested whether we could perform the same analysis with three implementations of an agent-based model (ABM, Figure 2). We considered the Grass-Sheep-Wolf (GSW) ABM implemented in NetLogo ((Wilensky & Reisman, 2006), Figure 2A), and implementations derived from it in MATLAB and Python (Source Code S4). We generated samples with sizes up to 20,000 (Figure 2B) and again found a decrease in EFECT error with increasing sample size according to a power law when testing for reproducibility (Figure 2C). Comparing results from different implementations, EFECT showed that no two implementations produce results that are equal in distribution (Figure 2D). Thus, we found that results from the three ABM implementations were reproducible, and that none of them produced the same results, according to EFECT. We also tested a stochastic Boolean network model of cell fate (Calzone et al., 2010) implemented in MaBoSS (Stoll et al., 2017) and found that EFECT supports evaluation of the reproducibility of stochastic Boolean network model results (Source Code S2.1, S3.27).

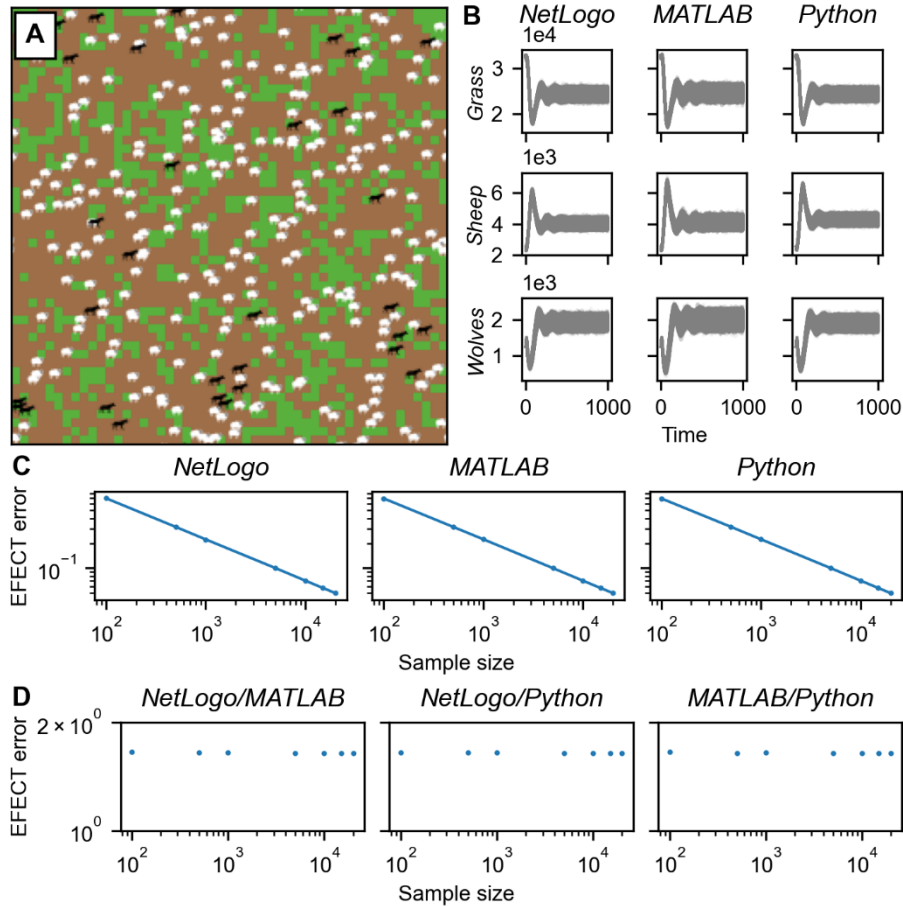


Figure 2. Testing reproducibility of results from three implementations of a Grass-Sheep-Wolf agent-based model. A: Screenshot of the original agent-based model implemented in NetLogo. B: Trajectories of model Grass (top), Sheep (middle), and Wolves (bottom) from 20,000 executions of the NetLogo (left, Source Code S4.2), MATLAB (center, Source Code S4.1), and Python (right, Source Code S4.3) implementations. C: Testing for reproducibility of results distributions from each implementation (Source Code S3.22). Results from all implementations converged in distribution. D: Testing for equality in distribution of pairs of samples from all implementations. No two implementations produced samples that were equal in distribution, as demonstrated by non-converging EFECT errors.

Finally, further evaluation showed that testing for reproducibility with EFECT is generally insensitive to the number of model variables (Source Code S3.31, Figure S1), significant differences in scale between model variables (Source Code S3.28, Figure S2), or order of magnitude of standard deviation (Source Code S3.32). This demonstrates that EFECT is a robust method to test model reproducibility. We also found that the runtime of our implementation of EFECT increases linearly with the number of model variables (runtime of  $545 \pm 336$  s,  $912 \pm 322$  s,  $1,183 \pm 379$  s,  $1,295 \pm 640$  s, and  $2,489 \pm 1,019$  s when testing reproducibility for 2, 3, 4, 5, and 10 model variables on a 2021 Apple M1 Max, Source Code S2.2).

### Workflows for Reproducibility of Stochastic Models

After testing that EFECT supports quantifying whether simulation results are reproducible and reproduced, we developed a Stochastic Simulation Reproducibility (SSR) workflow that defines the steps required by all stakeholders (*i.e.*, the Modeler and Curator) over the lifetime of a model

to generate stochastic results that are independently reproduced (Figure 3). Generally, the workflow involves the Modeler producing a published model and supporting data, which the Curator uses to test and report whether they are able to reproduce results according to an *a priori* significance level. Generally, neither stakeholder is required to use the same simulator that implements the model (though it may be beneficial to use different simulators).

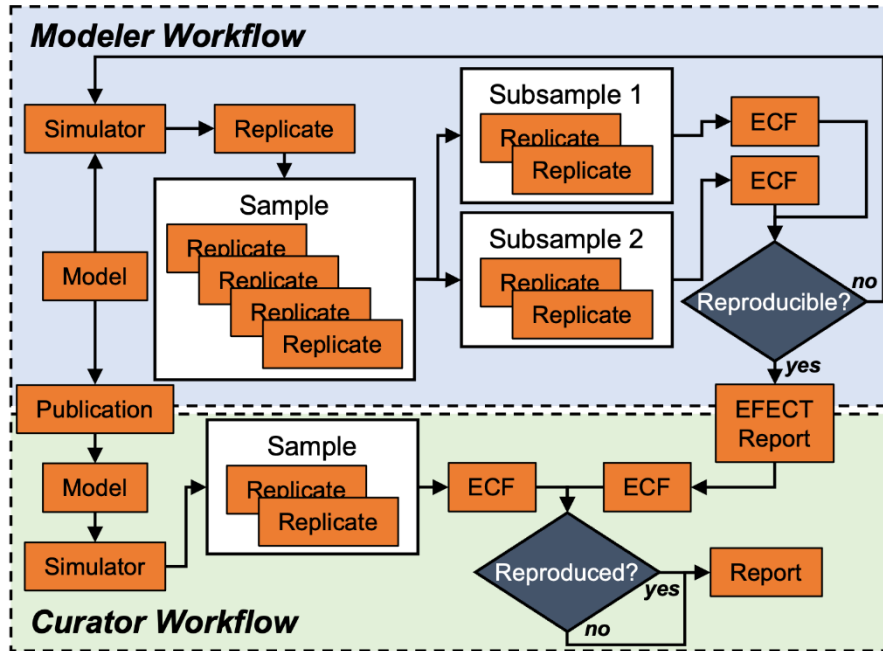


Figure 3. Process diagram of reproduced stochastic simulation results using EFECT. Modeler Workflow: A modeler (*i.e.*, the producer of stochastic results) iteratively samples a stochastic model, or parameter distributions of a model, until the aggregated sample is sufficiently reproducible. After showing reproducibility of results, the modeler reports minimal supporting data along with their publication that facilitates reproducing results by others. Curator Workflow: A curator (*i.e.*, the reproducer of stochastic results) reproduces published results by implementing the published model, generating a sample, and testing for equality in distribution of their results according to the supporting data that accompanies the model.

A Modeler iteratively adds simulation results to a sample until results are sufficiently reproducible. The Modeler evaluates reproducibility by randomly generating two equally sized subsamples of their current sample and calculating the EFECT error when comparing ECFs of the subsamples. Evaluations of the EFECT error are collected as a random variable and the process of randomly selecting and evaluating subsamples is repeated until the absolute relative change in the mean of EFECT error evaluations converges below a threshold (here tested with a threshold of 0.1%). Reproducibility is determined according to a threshold for the statistics of the EFECT error. Here we assume that the modeler, curator, and others in their field have defined such a threshold, an *EFECT convergence point*, through consensus (here accepted when the mean plus three standard deviations is below 0.075). When a sample is not reproducible, the Modeler generates more results, adds them to the sample, and repeats the process. When a sample is reproducible, we recommend that the Modeler generates minimal supporting data, called an *EFECT report*, to allow others to reproduce a subsample. Excluding relevant information specific to a model, simulation, or data formatting or than what our workflow requires, the contents of an EFECT report is as follows:

- *Stochastic variable names*: A list of all named stochastic variables in reported results.



- *Simulation times*: A list of all simulation times at which results are reported.
- *EFECT error statistics*. The accepted mean and standard deviation of the EFECT error during the test for reproducibility.
- *EFECT sample size*: The size of an accepted sample when testing for reproducibility.
- *ECF evaluations*: Evaluations of ECFs generated from a subsample of an accepted sample for each named stochastic variable at each simulation time.
- *ECF domains*: Maximum transform variable value for all reported ECF evaluations according to [ 5 ].
- *Input sampling*: Specification of distributions from which model inputs were sampled (if any) according to the ProbOnto ontology (M. J. Swat et al., 2016).
- *Significant figures*: The significant figures of results produced by the employed simulator(s) when testing for reproducibility.

A Curator develops their own implementation of a published model and generates results distributions according to information provided by the EFECT report accompanying the model. Specifically, the sample generated by the Curator will be of a size equal to the reported sample size and record results for all named stochastic variables at all reported simulation times. The Curator performs the test for reproducibility and generates their own EFECT error statistics. Finally, the Curator randomly selects a subsample, computes the EFECT error for comparing the generated ECFs to those reported in the EFECT report, and computes the probability that the EFECT error when comparing ECFs was taken from the (unknown) distribution of the EFECT error during the Curator’s test for reproducibility. The Curator determines an outcome according to whether the computed probability is above (reproduced) or below (not reproduced) a significance level.

To test whether our SSR workflow and EFECT report sufficiently support a lifecycle of model development that includes reproducing stochastic simulation results, we simulated the processes defined for a Modeler and Curator and exchange of information between them for an ODE model of viral infection (Kermack & McKendrick, 1927) defined in SBML (Figure 4). In these simulations, a Modeler reports results from an ODE model with random parameter sampling as generated by libRoadRunner (Welsh et al., 2023), and a Curator reproduces results using COPASI (Hoops et al., 2006). In our simulation of the processes performed by the Modeler (*i.e.*, the “Modeler Workflow” in Figure 3), the Modeler begins by iteratively increasing the size of their sample and testing for reproducibility until a power law could be fit to EFECT error evaluations. The Modeler rounds simulation results to six significant figures for compatibility with COPASI default output (neglecting this step results in failure to reproduce, Source Code S3.26). The fit is used to estimate the required sample size to achieve an EFECT convergence point, here defined *ad hoc* as an EFECT error with mean plus three standard deviations below 0.075, when testing for reproducibility (EFECT error of  $0.0537 \pm 6.45 \times 10^{-3}$ , Figure 4A, B, C), the final size of which was 10,120. An EFECT report is then generated and stored to file (Source Code S3.23). In our simulation of the processes performed by the Curator (*i.e.*, the “Curator Workflow” in Figure 3), the Curator implements the same model in COPASI, generates a sample of the reported size for all reported stochastic model variables and simulation times, and performs the test for reproducibility (error metric of  $0.0535 \pm 6.59 \times 10^{-3}$ ). The Curator then compares the ECFs generated from a randomly selected subsample to those reported in the EFECT report. Our simulation showed that the Curator accepts that results are

reproduced with EFECT error 0.0568 and p-value 1 (Figure 4D, Source Code S3.24) and can readily detect differences in model parameter distributions (Source Code S3.25).

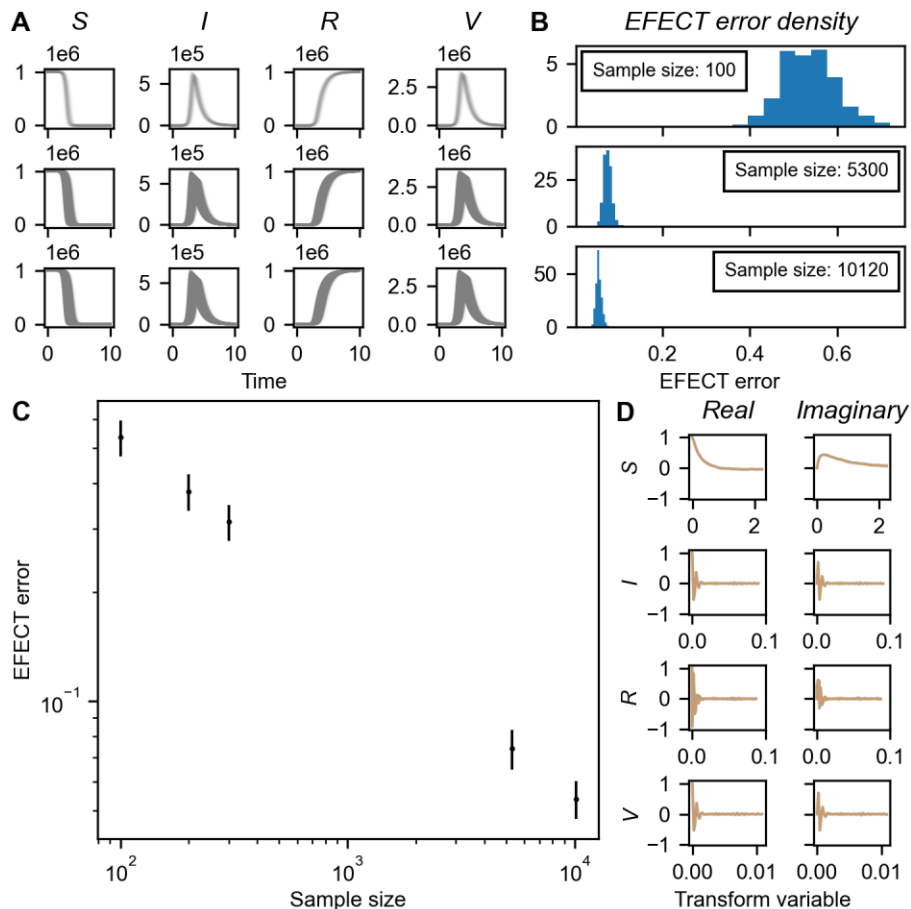


Figure 4. Simulation of the workflow for reproducing stochastic simulation results using different simulators. A, B, C: Modeler workflow using libRoadRunner. Trajectories (panel A) and corresponding EFECT error sampling (panel B) per select sample size are shown when testing for reproducibility (panel C). A sample was considered reproducible when the EFECT error plus three standard deviations was below a threshold (here 0.075). Error bars show three standard deviations. D: Comparison of empirical characteristic functions during the Curator workflow. Curator results were generated using COPASI to generate empirical characteristic functions (orange) for comparison to those reported from the Modeler workflow (blue). The two sets of empirical characteristic functions cannot be individually distinguished because they are nearly identical.

## Discussion

EFECT is agnostic to modeling methodology, broadly supports analysis of stochastic simulation results, and is readily computable. Its methodological constraints are that evaluated results must be formattable as vectors of evaluations of bounded, real-valued random variables. The interpretability of the EFECT error is independent of modeling methodology. Hence, outcomes of testing reproducibility of results are comparable irrespective of model or modeling methodology. Our workflow produces a minimal data set supporting reproducible simulation without requiring storage of all simulation data, a major advantage for promoting reproducible stochastic simulation as a broadly accepted hallmark of credible stochastic modeling. Our SSR workflow with EFECT enables testing for reproduced simulation results according to a significance level, which could be determined *a priori* through consensus in various modeling

communities with different expectations concerning tradeoffs in accuracy and computational cost.

The EFECT report can be incorporated into the SED-ML standard. For example, in reference to SED-ML Level 1 Version 4, data for each model variable can be taken from DataSet entries specified in SEDBase::Output::Report::ListOfDataSets, where variable objects of corresponding data generators are scalar and output data includes simulation times. Sample size can be specified through a RepeatedTask. Parameter sampling (if any) can be specified through a RepeatedTask (*i.e.*, via SetValue) for all distributions supported by SED-ML (*i.e.*, those supported by MathML). Thus, specifying a SED-ML script for data generation during a Curator's workflow naturally follows and could be automated using SED-ML library implementations (e.g., libSEDML in C++, Java, Python, and other programming languages).

The computational cost of EFECT can be non-negligible when testing for reproducibility, depending on sample size and variance of results. With the prototype implementation in Python used in this work, tests for reproducibility in all test cases ranged over orders of seconds to hours. Furthermore, the sample sizes required to produce reproducible results may be prohibitive for computationally expensive simulations, though various modeling communities can decide what level of accuracy (*i.e.*, a significance level) is appropriate for the typical computational cost of producing simulation results. Future work should develop infrastructure supporting usage of EFECT through our SSR workflow, including a markup language, publicly distributed libraries in popular scientific programming languages (e.g., C++, Python, R, Julia), hardware-accelerated algorithms, and online web services. Related strategies have been previously demonstrated for accelerating ECF computations (O'Brien et al., 2014). To this end, we have started such a software project, called "libSSR" ([github.com/tjsego/libSSR](https://github.com/tjsego/libSSR)), which is intended to provide such infrastructure in anticipation of future standardization efforts on stochastic simulation reproducibility. Current initial work is providing support in Python. Future work should also define the computations and minimal necessary data to recover the probability distribution functions encoded in ECFs according to our standard data, which could provide an efficient means to share statistical model results using an enhanced form of the EFECT Report.

Lastly, future work should promote community adoption of EFECT by simulation tools, public repositories, journals, modeling communities, and regulatory agencies. Simulation tools could provide built-in features that perform our Modeler's workflow. Modeling communities could develop standard test cases like those performed here to develop consensus on EFECT convergence points that define sufficiently reproduced results. Communities that use expensive or exceptionally stochastic simulations may be inclined to adopt more forgiving criteria for reproducibility. Public repositories could curate stochastic models according to EFECT convergence points and report EFECT errors from testing for reproduced results. Journals could promote or even require including our standard data as supplementary information accompanying publication of stochastic models (Mendes, 2018; Peng, 2011). Regulatory agencies could define what EFECT convergence point must be met, or proven, for stochastic simulation results to be considered valid evidence supporting applications and guidance on policymaking.

## Materials and methods

We test for equality in distribution of two stochastic processes  $X(t)$  and  $Y(t)$  by testing whether the two stochastic processes have the same unknown probability distribution  $f(x; t)$  (*i.e.*,

$X(t) \stackrel{D}{=} Y(t)$  if and only if  $X, Y \sim f$ ). Our test is performed over an arbitrary set of simulation times on samples produced from repeated execution of stochastic simulation without regard for determining, or performing any characterization of, the underlying distribution(s). Rather, our method relies on the limit behavior of empirical distributions. Specifically, our method relies on characteristic functions (CFs). For a stochastic process  $X(t)$ , the CF  $\phi_X(\tau; t)$  is defined as

$$\phi_X(\tau; t) = \mathbb{E}[e^{iX(t)\tau}]. \quad [ 1 ]$$

Here  $\tau$  is a transform variable and  $i$  is the imaginary unit. We built our method around CFs because of specific useful properties of a CF: 1. The magnitude of all CFs is in  $[0,1]$ , which allows assessment of model variables without case-specific decision-making to handle their statistical features (e.g., means at different orders of magnitude); 2. A CF supports all bounded, real-valued random variables, whether discrete or continuous, whereas the probability distribution function only supports continuous random variables; and 3. The empirical characteristic function (ECF) is naturally computed from a sample of stochastic simulation results (Lukacs, 1972). For a sample  $\{X_i(t)\}_{i=1}^n$  of size  $n$  of  $X(t)$  at time  $t$ , the ECF  $\phi_{X,n}(\tau; t)$  is defined as

$$\phi_{X,n}(\tau; t) = \frac{1}{n} \sum_{1 \leq j \leq n} e^{iX_j(t)\tau} \quad [ 2 ]$$

We use Lévy's Continuity Theorem to test for equality in distribution of simulation results: for samples  $\{X_i(t)\}_{i=1}^n$  of  $X(t)$  and  $\{Y_i(t)\}_{i=1}^n$  of  $Y(t)$ ,  $\{X_i(t)\}_{i=1}^n \xrightarrow{D} X(t)$  if and only if  $\lim_{n \rightarrow \infty} \phi_{X_n}(t) = \phi_X(t)$ , and likewise for  $\{Y_i(t)\}_{i=1}^n$  and  $Y(t)$ . For two sampled random processes  $X(t)$  and  $Y(t)$  over  $t \in \mathcal{T}$ , we test for pointwise convergence of the characteristic functions  $\phi_X(\tau; t)$  and  $\phi_Y(\tau; t)$  over  $\tau \in \mathcal{t}$  to the same unknown distribution using the EFECT error  $\delta$ ,

$$\delta(\phi_{X,n}(\tau; t), \phi_{Y,n}(\tau; t)) = \sup\{|\phi_{X,n}(\tau; t) - \phi_{Y,n}(\tau; t)| : \tau \in \mathcal{t}, t \in \mathcal{T}\}. \quad [ 3 ]$$

It naturally follows that  $X(t) \stackrel{D}{=} Y(t)$  if and only if  $\delta(\phi_{X,n}(\tau; t), \phi_{Y,n}(\tau; t)) \rightarrow 0$  as  $n \rightarrow \infty$ . When comparing two samples of a model, we consider the maximum EFECT error of all stochastic variables of the model as the EFECT error of the two samples.

### *Methodological Details*

Our test for reproducibility evaluates the likelihood that a sample would be assessed as reproduced by another sample of the same model according to comparison of generated ECFs. A sample of size  $2N$  is tested for reproducibility by sampling the EFECT error [ 3 ] via random selection and evaluation of two  $N$  – combinations. If the sample produces statistics of the EFECT error that satisfy some *a priori* criterion, an EFECT convergence point, then the sample is considered sufficiently reproducible. While we intend in this work to enable consensus building as to what criterion should be considered broadly acceptable, generally an EFECT error closer to zero produces more reliably reproducible results. Tests thus far suggest that requiring an EFECT error mean below 0.05 may be unnecessarily demanding to detect differences in samples that may not be significant to some modeling communities (Figure 1).

The CF has support for the transform variable (*i.e.*,  $\tau$  in [ 1 ] and [ 2 ]) over the real line. Since the intent of the EFECT report is to encode sufficient information to test whether the sample represented by the EFECT report is sufficiently reproduced by another sample, we define a parameterization of the domain of ECFs to produce representative evaluations to test for equality. Specifically, choosing an excessively small domain produces evaluations that are all approximately equal to one, while choosing an excessively large domain can produce evaluations that are equal to either one (*i.e.*, when  $\tau = 0$ ) or zero elsewhere (since, by Riemann-Lebesgue theorem,  $\phi_X(\tau; t) \rightarrow 0$  as  $|\tau| \rightarrow \infty$  if  $X$  has a density, of which  $\phi_X(\tau; t)$  is the Fourier dual). We rewrite [ 1 ] to produce a form for parameterizing the domain over which to evaluate an ECF,

$$\phi_X(\tau; t) = \theta(\tau; t) \sqrt{\frac{\left(\mathbb{E}(\cos(Z(t)\sigma(X(t))\tau)\right)^2 + \left(\mathbb{E}(\sin(Z(t)\sigma(X(t))\tau)\right)^2}{\sigma(X(t))Z(t) = X(t) - \mu(X(t))}}, \quad [ 4 ]$$

Here  $\theta(\tau; t)$  is the phase of  $\phi_X(\tau; t)$ , the details of which are unimportant here except to note that  $|\theta(\tau; t)| = 1$ , and  $Z(t)$  is standardized  $X(t)$  with mean  $\mu(X(t))$  and standard deviation  $\sigma(X(t))$ . Oscillations in the amplitude of  $\phi_X(\tau; t)$  are expected to occur on the scale of periods inversely proportional to the standard deviation of  $X(t)$  (since  $Z(t)$  is standardized), hence we evaluate an ECF over  $m$  periods inversely proportional to the standard deviation,

$$\tau \in \left[ 0, \frac{2\pi m}{\sigma(X(t))} \right], \quad m \geq 1. \quad [ 5 ]$$

In application, we evaluate ECFs over a uniform distribution of the transform variable according to [ 5 ] for a given value of  $m$ . All results in this work were performed with a value of  $m$  equal to 3 and 100 uniformly distributed values of the transform variable, which were determined *ad hoc* by inspection of results. Generally, greater values of  $m$  for the same number of evaluations risks loss of sufficient encoding, and greater numbers of evaluations increases the storage requirements of the EFECT report. We make no formal recommendation here as to what values for these parameters should be broadly acceptable except to note that preliminary tests suggest that a value of  $m$  equal to one may be sufficient for a number of evaluations on the order of 100, whereas increasing the number of evaluations nearer to 1,000 may provide marginal improvements to the quality of encoding for at least  $m$  equal to three (Source Code S3.29, S3.30).

Testing for reproduced results is performed using the null hypothesis that the EFECT error from comparing Modeler and Curator ECFs was taken from the same distribution as the EFECT error from the Curator's test for reproducibility. Assessment of whether results have been reproduced is performed according to an *a priori* significance level using Chebyshev's inequality for unknown population mean and variance (Kabán, 2012). For Curator's and Modeler's stochastic processes  $X(t)$  and  $Y(t)$ , respectively, a Curator performs the test for reproducibility with their model implementation and generates a sample of EFECT error  $\delta_{XX}$  with sample size  $N$ , mean  $\bar{\delta}_{XX}$ , and variance  $\tilde{\delta}_{XX}$ . The Curator then generates an EFECT error  $\delta_{XY}$  by comparing the ECF(s) of their sample to that(those) reported by the Modeler and calculates the probability that  $\delta_{XY}$  was taken from the same distribution as  $\delta_{XX}$ ,

$$\Pr(|\delta_{XX} - \bar{\delta}_{XX}| \geq |\delta_{XY} - \bar{\delta}_{XX}|) = \frac{1}{N+1} \left[ \frac{N+1}{N} \left( \frac{N^2-1}{N} \frac{\tilde{\delta}_{XX}}{(\delta_{XY} - \bar{\delta}_{XX})^2} + 1 \right) \right]. \quad [6]$$

Hence, the Curator concludes whether they have reproduced the Modeler's results via testing of the null hypothesis for a significance level  $\alpha$ ,

$$X \stackrel{D}{=} Y \leftrightarrow \Pr(|\delta_{XX} - \bar{\delta}_{XX}| \geq |\delta_{XY} - \bar{\delta}_{XX}|) \geq \alpha. \quad [7]$$

### *Implementation Details*

ODEs were executed with libRoadRunner v2.4.0, except those described as being executed in COPASI v4.42. Stochastic ODE simulation with libRoadRunner was performed using the libRoadRunner implementation of the Gillespie stochastic simulation algorithm. ODEs executed in libRoadRunner were specified in Antimony v2.13.4 (L. P. Smith et al., 2009) to generate SBML specification, except those imported from BioModels. ODEs executed in COPASI were specified using generated SBML from Antimony specification. Parameter sampling for ODE simulation with libRoadRunner was performed using methods provided in SciPy v1.11.3 (Virtanen et al., 2020). The GSW ABM NetLogo implementation was executed in NetLogo v6.1.1 (September 26, 2019; Source Code S4.2); the MATLAB implementation was executed in MATLAB 2020b (Source Code S4.1); the Python implementation was executed in Python v3.10.13 (Source Code S4.3). All ECF generation and analysis was performed using library functions available in NumPy v1.26.0 (Harris et al., 2020). The stochastic Boolean network model was executed in the MaBoSS deployment in CompuCell3D v4.5.0 (M. H. Swat et al., 2012). Source code to perform all tests and analyses in Python and simulations with libRoadRunner and MaBoSS are available in Source Code S2. Source code to perform data generation with COPASI is available in Source Code S5. Source code to install all Python dependencies is available in Source Code S1. All project source code for this work is available at [https://github.com/tjsego/ssr\\_project\\_2024](https://github.com/tjsego/ssr_project_2024). Instructions to install Python dependencies and execute Python source code and Jupyter Notebooks are available in Instruction S1.

### **Conclusion**

This work delivers EFECT, a method for reproducible stochastic simulation results, broadly defined. EFECT capability is neither limited by scientific discipline nor specific to modeling methodology. Implementation of such capability has been demonstrated to readily support reproducible stochastic simulation in an existing digital ecosystem that enables standardized computational research. Current efforts continue to promote widespread adoption of reproducible stochastic simulation as common practice by working to deliver open-source, publicly available software libraries, and a publicly maintained standard specification, among other necessary infrastructure. Our hope is to reduce the crisis of reproducibility in computational research by providing a general, unambiguous means by which peers can rigorously reproduce the stochastic simulation results of others.

### **Acknowledgments and funding sources**

TJS acknowledges funding from National Institutes of Health grant number U24EB028887 and National Science Foundation grant number 2000281. LLF acknowledges funding from National Institutes of Health grants number R01 GM127909 and R01 AI135128, and Defense Advanced

Research Projects Agency grant number HR00112220038. ACK acknowledges funding from National Science Foundation grant number 2325776. KT, HH and RMS acknowledge EMBL Core funding. HMS acknowledges funding from National Institutes of Health grant numbers U24EB028887 and P41EB023912. JAG acknowledges funding from National Institutes of Health grant number U24EB028887. RCL acknowledges funding from National Institutes of Health grants number R01 GM127909, R01 AI135128, and R01 HL169974, and Defense Advanced Research Projects Agency grant number HR00112220038.

The authors would like to thank Frank Bergmann, Randy Heiland, Jürgen Pahle, Dilan Pathirana, Sven Sahle, and Lucian Smith for their insightful feedback during reproducibility workshop discussions at the COMBINE 2023 (University of Connecticut) and HARMONY 2024 (University College London) events.

### Competing Interests

The authors declare no competing interests.

### Supplementary Materials

- Figure S1. Testing sensitivity of the error metric to number of model variables. A: Oscillator models with increasing number of variables of 2, 3, 4, 5, and 10 were sampled while randomly sampling the frequency of oscillatory variables. B: Box plot of error metrics while sampling for self-similarity for 100, 1k, and 10k sample sizes and 2 (Test 1), 3 (Test 2), 4 (Test 3), 5 (Test 4), and 10 (Test 5) model variables.
- Figure S2. Testing sensitivity of the error metric to differences in orders of magnitude in model variables. A: A biphasic variable (left) controls the frequency of an oscillatory variable (right) with amplitude 1 (Test 1), 1E-3 (Test 2), or 1E-6 (Test 3). B: Box plot of error metrics while sampling for self-similarity for 100, 1k, and 10k sample sizes and oscillatory variable amplitude 1 (Test 1), 1E-3 (Test 2), or 1E-6 (Test 3).
- Instruction S1. Instructions for software installation and execution.
- Source Code S1. `setup/`: Source code to install all software.
- Source Code S2. `code/`: Source code implementing tests, simulations, tested models, and method equations.
  - Source Code S2.1. `cc3d_cell_fate.py`: Data generation for testing the approach on results from a stochastic Boolean network model.
  - Source Code S2.2. `perf_compare_size_1.py`: Runtime test for computational cost of the approach with increasing number of model variables.
  - Source Code S2.3. `sim_2_modeler_py.py`: Simulation of the Modeler's workflow. Generates data for use in simulation of the Curator's workflow for cases where data with six and nine significant figures was used.
  - Source Code S2.4. `sim_2_curator_py.py`: Simulation of the Curator's workflow. Uses generated data from simulation of the Modeler's workflow and externally generated data from COPASI. Tests cases where results are successfully reproduced, unsuccessfully reproduced due to differences in parameter distributions, and unsuccessfully reproduced due to significant figures of data from different simulators.
  - Source Code S2.5. `test_comparison.py`: Tests the likelihood of detecting parameter differences in a viral infection model as a function of sample size and scale of difference in distribution mean and standard deviation.

- Source Code S3. notebooks/: Jupyter Notebooks demonstrating various test cases and simulation of the reproducible simulation workflow.
  - Source Code S3.1. proto\_paramvar.ipynb: Demonstration of basic approach with a viral/bacterial coinfection model while sampling one parameter with a normal distribution.
  - Source Code S3.2. proto\_2paramvar.ipynb: Demonstration of basic approach with a viral/bacterial coinfection model while sampling two parameters with a normal distribution.
  - Source Code S3.3. proto\_10paramvar.ipynb: Demonstration of basic approach with a viral/bacterial coinfection model while sampling ten parameters with a normal distribution.
  - Source Code S3.4. proto\_10paramvar2.ipynb: Demonstration of basic approach with a viral/bacterial coinfection model while sampling ten parameters with a uniform distribution.
  - Source Code S3.5. proto\_bistable.ipynb: Demonstration of basic approach with a two-variable biphasic/oscillatory model while sampling one initial condition with a uniform distribution.
  - Source Code S3.6. proto\_bistable2.ipynb: Demonstration of basic approach with a two-variable biphasic/oscillatory model while sampling one parameter with a normal distribution.
  - Source Code S3.7. proto\_banerjee2008.ipynb: Demonstration of basic approach with BioModels entry MODEL2001300001.
  - Source Code S3.8. proto\_liu2012.ipynb: Demonstration of basic approach with BioModels entry MODEL2004140002.
  - Source Code S3.9. proto\_lo2005.ipynb: Demonstration of basic approach with BioModels entry MODEL1805160001.
  - Source Code S3.10. proto\_lorenz.ipynb: Demonstration of basic approach with a Lorenz system model using a stochastic simulation algorithm.
  - Source Code S3.11. proto\_pendulum.ipynb: Demonstration of basic approach with a nonlinear pendulum model with two cases of sampling an initial condition with a uniform distribution, and one case of sampling a parameter with a normal distribution.
  - Source Code S3.12. proto\_pulse.ipynb: Demonstration of basic approach with a pulse-like oscillatory model while sampling one parameter with a normal distribution.
  - Source Code S3.13. proto\_seir.ipynb: Demonstration of basic approach with a viral infection model using a stochastic simulation algorithm.
  - Source Code S3.14. proto\_compare\_1.ipynb: Demonstration of detecting structural differences in models, one of which is subjected to parameter sampling, and the other of which is simulated with a stochastic simulation algorithm.
  - Source Code S3.15. proto\_compare\_2.ipynb: Demonstration of detecting parametric differences for a viral infection model subject to parameter sampling. Parametric differences are implemented as proportionally scaled differences in parameter distribution mean and standard deviation.
  - Source Code S3.16. proto\_compare\_3.ipynb: Demonstration of detecting parametric differences in constant-valued models. Parametric differences are implemented as differences in line intercept.
  - Source Code S3.17. proto\_compare\_4.ipynb: Demonstration of detecting parametric differences for a two-variable biphasic oscillatory model subject to parameter sampling. Parametric differences are implemented as differences in parameter distribution mean.



- Source Code S3.18. proto\_compare\_5.ipynb: Demonstration of detecting parametric differences for a reaction-kinetics model subject to parameter sampling. Parametric differences are implemented as proportionally scaled differences in parameter distribution mean and standard deviation.
- Source Code S3.19. proto\_compare\_5a.ipynb: Demonstration of detecting parametric differences for a reaction-kinetics model simulated with a stochastic simulation algorithm.
- Source Code S3.20. proto\_compare\_6.ipynb: Demonstration of detecting parametric differences for a viral infection model subject to parameter sampling. Parametric differences are implemented as differences in parameter distribution standard deviation.
- Source Code S3.21. proto\_compare\_7.ipynb: Demonstration of summary statistics incorrectly assessing samples with different initial conditions as equal.
- Source Code S3.22. gsw.ipynb: Demonstration of basic approach with three implementations of a GSW ABM.
- Source Code S3.23. sim\_modeler\_2.ipynb: Simulation of the Modeler's workflow when generating reproducible simulation results.
- Source Code S3.24. sim\_curator\_2\_pass.ipynb: Simulation of the Curator's workflow when successfully reproducing simulation results using supporting data generated from the Modeler's workflow.
- Source Code S3.25. sim\_curator\_2\_fail\_params.ipynb: Simulation of the Curator's workflow when unsuccessfully reproducing simulation results using supporting data generated from the Modeler's workflow. Results are not reproduced due to differences in a model parameter distribution.
- Source Code S3.26. sim\_curator\_2\_fail\_sigfig.ipynb: Simulation of the Curator's workflow when unsuccessfully reproducing simulation results using supporting data generated from the Modeler's workflow. Results are not reproduced due to differences in significant figures of data output by the simulators used in the Modeler's and Curator's workflows.
- Source Code S3.27. proto\_compare\_boolean.ipynb: Demonstration of basic approach with a stochastic Boolean network model.
- Source Code S3.28. proto\_compare\_mag\_1.ipynb: Test for sensitivity of the method to differences in magnitude of model variables.
- Source Code S3.29. proto\_compare\_periods\_1.ipynb: Test for sensitivity of the method to different periods of evaluation of ECFs.
- Source Code S3.30. proto\_compare\_res\_1.ipynb: Test for sensitivity of the method to different discretization resolutions of the transform variable when evaluating ECFs.
- Source Code S3.31. proto\_compare\_size\_1.ipynb: Test for sensitivity of the method to different numbers of model variables.
- Source Code S3.32. proto\_compare\_var\_2.ipynb: Test for sensitivity of the method to different orders of magnitude of the standard deviation of a parameter distribution.
- Source Code S4. gsw/: Source code for GSW ABM implementations.
  - Source Code S4.1. GSW\_MatLab\_source: MATLAB implementation of the GSW ABM.
  - Source Code S4.2. GSW\_NetLogo\_source: NetLogo implementation of the GSW ABM.
  - Source Code S4.3. GSW\_Python\_source: Python implementation of the GSW ABM.
- Source Code S5. copasi/: Source code generating simulation data for the Curator's workflow.

- Source Code S5.1. curator\_results\_diff/: Source code to generate results with COPASI when simulating the Curator's workflow with no model differences.
- Source Code S5.2. curator\_results\_same/: Source code to generate results with COPASI when simulating the Curator's workflow with parametric differences.

## References

- Baker, M. (2016). 1,500 scientists lift the lid on reproducibility. *Nature*, 533(7604), 452–454.  
<https://doi.org/10.1038/533452a>
- Bergmann, F. T., & Sauro, H. M. (2008). Comparing simulation results of SBML capable simulators. *Bioinformatics*, 24(17), 1963–1965. <https://doi.org/10.1093/bioinformatics/btn319>
- Calzone, L., Tournier, L., Fourquet, S., Thieffry, D., Zhivotovsky, B., Barillot, E., & Zinovyev, A. (2010). Mathematical modelling of cell-fate decision in response to death receptor engagement. *PLoS Computational Biology*, 6(3), e1000702.  
<https://doi.org/10.1371/journal.pcbi.1000702>
- Casdagli, M. (1992). Chaos and Deterministic Versus Stochastic Non-Linear Modelling. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 54(2), 303–328.  
<https://doi.org/10.1111/j.2517-6161.1992.tb01884.x>
- Courtot, M., Juty, N., Knüpfer, C., Waltemath, D., Zhukova, A., Dräger, A., Dumontier, M., Finney, A., Golebiewski, M., Hastings, J., Hoops, S., Keating, S., Kell, D. B., Kerrien, S., Lawson, J., Lister, A., Lu, J., Machne, R., Mendes, P., ... Le Novère, N. (2011). Controlled vocabularies and semantics in systems biology. *Molecular Systems Biology*, 7, 543.  
<https://doi.org/10.1038/msb.2011.77>
- Evans, T. W., Gillespie, C. S., & Wilkinson, D. J. (2008). The SBML discrete stochastic models test suite. *Bioinformatics*, 24(2), 285–286. <https://doi.org/10.1093/bioinformatics/btm566>

- Friedrich, R., Peinke, J., Sahimi, M., & Reza Rahimi Tabar, M. (2011). Approaching complexity by stochastic methods: From biological systems to turbulence. *Physics Reports*, 506(5), 87–162. <https://doi.org/10.1016/j.physrep.2011.05.003>
- Fuller, A., Fan, Z., Day, C., & Barlow, C. (2020). Digital twin: enabling technologies, challenges and open research. *IEEE Access: Practical Innovations, Open Solutions*, 8, 108952–108971. <https://doi.org/10.1109/ACCESS.2020.2998358>
- Ghosh, S., Matsuoka, Y., Asai, Y., Hsin, K.-Y., & Kitano, H. (2011). Software for systems biology: from tools to integrated platforms. *Nature Reviews. Genetics*, 12(12), 821–832. <https://doi.org/10.1038/nrg3096>
- Gillespie, D. T. (1976). A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, 22(4), 403–434. [https://doi.org/10.1016/0021-9991\(76\)90041-3](https://doi.org/10.1016/0021-9991(76)90041-3)
- Goldenfeld, N., & Kadanoff, L. P. (1999). Simple lessons from complexity. *Science*, 284(5411), 87–89. <https://doi.org/10.1126/science.284.5411.87>
- Goodman, S. N., Fanelli, D., & Ioannidis, J. P. (2016). What does research reproducibility mean? *Science Translational Medicine*, 8(341), 1–6. <https://doi.org/10.1126/scitranslmed.aaf5027>
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., Del Río, J. F., Wiebe, M., Peterson, P., ... Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, 585(7825), 357–362. <https://doi.org/10.1038/s41586-020-2649-2>

- Hoops, S., Sahle, S., Gauges, R., Lee, C., Pahle, J., Simus, N., Singhal, M., Xu, L., Mendes, P., & Kummer, U. (2006). COPASI--a COmplex PATHway Simulator. *Bioinformatics*, *22*(24), 3067–3074. <https://doi.org/10.1093/bioinformatics/btl485>
- Kabán, A. (2012). Non-parametric detection of meaningless distances in high dimensional data. *Statistics and Computing*, *22*(2), 375–385. <https://doi.org/10.1007/s11222-011-9229-0>
- Kareva, I., Berezovskaya, F., & Castillo-Chavez, C. (2010). Myeloid cells in tumour-immune interactions. *Journal of Biological Dynamics*, *4*(4), 315–327. <https://doi.org/10.1080/17513750903261281>
- Keating, S. M., Waltemath, D., König, M., Zhang, F., Dräger, A., Chaouiya, C., Bergmann, F. T., Finney, A., Gillespie, C. S., Helikar, T., Hoops, S., Malik-Sheriff, R. S., Moodie, S. L., Moraru, I. I., Myers, C. J., Naldi, A., Olivier, B. G., Sahle, S., Schaff, J. C., ... SBML Level 3 Community members. (2020). SBML Level 3: an extensible format for the exchange and reuse of biological models. *Molecular Systems Biology*, *16*(8), e9110. <https://doi.org/10.15252/msb.20199110>
- Kermack, W. O., & McKendrick, A. G. (1927). A Contribution to the Mathematical Theory of Epidemics. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *115*(772), 700–721. <https://doi.org/10.1098/rspa.1927.0118>
- Kirouac, D. C., Cicali, B., & Schmidt, S. (2019). Reproducibility of quantitative systems pharmacology models: current challenges and future opportunities. *CPT: Pharmacometrics & Systems Pharmacology*, *8*(4), 205–210. <https://doi.org/10.1002/psp4.12390>
- Liu, Z., Pu, Y., Li, F., Shaffer, C. A., Hoops, S., Tyson, J. J., & Cao, Y. (2012). Hybrid modeling and simulation of stochastic effects on progression through the eukaryotic cell cycle. *The Journal of Chemical Physics*, *136*(3), 034105. <https://doi.org/10.1063/1.3677190>

- Lo, K., Denney, W. S., & Diamond, S. L. (2005). Stochastic modeling of blood coagulation initiation. *Pathophysiology of Haemostasis and Thrombosis*, 34(2–3), 80–90.  
<https://doi.org/10.1159/000089929>
- Lukacs, E. (1972). A survey of the theory of characteristic functions. *Advances in Applied Probability*, 4(1), 1–37. <https://doi.org/10.2307/1425805>
- Malik-Sheriff, R. S., Glont, M., Nguyen, T. V. N., Tiwari, K., Roberts, M. G., Xavier, A., Vu, M. T., Men, J., Maire, M., Kananathan, S., Fairbanks, E. L., Meyer, J. P., Arankalle, C., Varusai, T. M., Knight-Schrijver, V., Li, L., Dueñas-Roca, C., Dass, G., Keating, S. M., ... Hermjakob, H. (2020). BioModels-15 years of sharing computational models in life science. *Nucleic Acids Research*, 48(D1), D407–D415. <https://doi.org/10.1093/nar/gkz1055>
- McDougal, R. A., Bulanova, A. S., & Lytton, W. W. (2016). Reproducibility in computational neuroscience models and simulations. *IEEE Transactions on Bio-Medical Engineering*, 63(10), 2021–2035. <https://doi.org/10.1109/TBME.2016.2539602>
- Medley, J. K., Goldberg, A. P., & Karr, J. R. (2016). Guidelines for reproducibly building and simulating systems biology models. *IEEE Transactions on Bio-Medical Engineering*, 63(10), 2015–2020. <https://doi.org/10.1109/TBME.2016.2591960>
- Mendes, P. (2018). Reproducible research using biomodels. *Bulletin of Mathematical Biology*, 80(12), 3081–3087. <https://doi.org/10.1007/s11538-018-0498-z>
- O'Brien, T. A., Collins, W. D., Rauscher, S. A., & Ringler, T. D. (2014). Reducing the computational cost of the ECF using a nuFFT: A fast and objective probability density estimation method. *Computational Statistics & Data Analysis*, 79, 222–234.  
<https://doi.org/10.1016/j.csda.2014.06.002>
- Peng, R. D. (2011). Reproducible research in computational science. *Science*, 334(6060), 1226–1227. <https://doi.org/10.1126/science.1213847>

Porubsky, V. L., Goldberg, A. P., Rampadarath, A. K., Nickerson, D. P., Karr, J. R., & Sauro, H. M. (2020). Best practices for making reproducible biochemical models. *Cell Systems*, 11(2), 109–120. <https://doi.org/10.1016/j.cels.2020.06.012>

Porubsky, V. L., & Sauro, H. M. (2023). A practical guide to reproducible modeling for biochemical networks. *Methods in Molecular Biology*, 2634, 107–138. [https://doi.org/10.1007/978-1-0716-3008-2\\_5](https://doi.org/10.1007/978-1-0716-3008-2_5)

Read, M. N., Alden, K., Timmis, J., & Andrews, P. S. (2020). Strategies for calibrating models of biology. *Briefings in Bioinformatics*, 21(1), 24–35. <https://doi.org/10.1093/bib/bby092>

Reiczigel, J., Zakariás, I., & Rózsa, L. (2005). A bootstrap test of stochastic equality of two populations. *The American Statistician*, 59(2), 156–161. <https://doi.org/10.1198/000313005X23526>

Sandve, G. K., Nekrutenko, A., Taylor, J., & Hovig, E. (2013). Ten simple rules for reproducible computational research. *PLoS Computational Biology*, 9(10), e1003285. <https://doi.org/10.1371/journal.pcbi.1003285>

Schaduangrat, N., Lampa, S., Simeon, S., Gleeson, M. P., Spjuth, O., & Nantasenamat, C. (2020). Towards reproducible computational drug discovery. *Journal of Cheminformatics*, 12(1), 9. <https://doi.org/10.1186/s13321-020-0408-x>

Shaikh, B., Smith, L. P., Vasilescu, D., Marupilla, G., Wilson, M., Agmon, E., Agnew, H., Andrews, S. S., Anwar, A., Beber, M. E., Bergmann, F. T., Brooks, D., Bruschi, L., Calzone, L., Choi, K., Cooper, J., Detloff, J., Drawert, B., Dumontier, M., ... Karr, J. R. (2022). BioSimulators: a central registry of simulation engines and services for recommending specific tools. *Nucleic Acids Research*, 50(W1), W108–W114. <https://doi.org/10.1093/nar/gkac331>

- Shin, J., Porubsky, V., Carothers, J., & Sauro, H. M. (2023). Standards, dissemination, and best practices in systems biology. *Current Opinion in Biotechnology*, 81, 102922. <https://doi.org/10.1016/j.copbio.2023.102922>
- Smith, A. M., & Smith, A. P. (2016). A critical, nonlinear threshold dictates bacterial invasion and initial kinetics during influenza. *Scientific Reports*, 6, 38703. <https://doi.org/10.1038/srep38703>
- Smith, L. P., Bergmann, F. T., Chandran, D., & Sauro, H. M. (2009). Antimony: a modular model definition language. *Bioinformatics*, 25(18), 2452–2454. <https://doi.org/10.1093/bioinformatics/btp401>
- Sordo Vieira, L., & Laubenbacher, R. C. (2022). Computational models in systems biology: standards, dissemination, and best practices. *Current Opinion in Biotechnology*, 75, 102702. <https://doi.org/10.1016/j.copbio.2022.102702>
- Stainforth, D. A., Aina, T., Christensen, C., Collins, M., Faull, N., Frame, D. J., Kettleborough, J. A., Knight, S., Martin, A., Murphy, J. M., Piani, C., Sexton, D., Smith, L. A., Spicer, R. A., Thorpe, A. J., & Allen, M. R. (2005). Uncertainty in predictions of the climate response to rising levels of greenhouse gases. *Nature*, 433(7024), 403–406. <https://doi.org/10.1038/nature03301>
- Stodden, V., Seiler, J., & Ma, Z. (2018). An empirical analysis of journal policy effectiveness for computational reproducibility. *Proceedings of the National Academy of Sciences of the United States of America*, 115(11), 2584–2589. <https://doi.org/10.1073/pnas.1708290115>
- Stoll, G., Caron, B., Viara, E., Dugourd, A., Zinovyev, A., Naldi, A., Kroemer, G., Barillot, E., & Calzone, L. (2017). MaBoSS 2.0: an environment for stochastic Boolean modeling. *Bioinformatics*, 33(14), 2226–2228. <https://doi.org/10.1093/bioinformatics/btx123>

- Swat, M. H., Thomas, G. L., Belmonte, J. M., Shirinifard, A., Hmeljak, D., & Glazier, J. A. (2012). Multi-scale modeling of tissues using CompuCell3D. *Methods in Cell Biology*, 110, 325–366. <https://doi.org/10.1016/B978-0-12-388403-9.00013-8>
- Swat, M. J., Grenon, P., & Wimalaratne, S. (2016). ProbOnto: ontology and knowledge base of probability distributions. *Bioinformatics*, 32(17), 2719–2721. <https://doi.org/10.1093/bioinformatics/btw170>
- Tiwari, K., Kananathan, S., Roberts, M. G., Meyer, J. P., Sharif Shohan, M. U., Xavier, A., Maire, M., Zyoud, A., Men, J., Ng, S., Nguyen, T. V. N., Glont, M., Hermjakob, H., & Malik-Sheriff, R. S. (2021). Reproducibility in systems biology modelling. *Molecular Systems Biology*, 17(2), e9982. <https://doi.org/10.15252/msb.20209982>
- Vargha, A., & Delaney, H. D. (2000). A Critique and Improvement of the CL Common Language Effect Size Statistics of McGraw and Wong. *Journal of Educational and Behavioral Statistics*, 25(2), 101–132. <https://doi.org/10.3102/10769986025002101>
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., ... SciPy 1.0 Contributors. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- Waltemath, D., Adams, R., Bergmann, F. T., Hucka, M., Kolpakov, F., Miller, A. K., Moraru, I. I., Nickerson, D., Sahle, S., Snoep, J. L., & Le Novère, N. (2011). Reproducible computational biology experiments with SED-ML--the Simulation Experiment Description Markup Language. *BMC Systems Biology*, 5, 198. <https://doi.org/10.1186/1752-0509-5-198>



Wearing, H. J., Rohani, P., & Keeling, M. J. (2005). Appropriate models for the management of infectious diseases. *PLoS Medicine*, 2(7), e174.

<https://doi.org/10.1371/journal.pmed.0020174>

Welsh, C., Xu, J., Smith, L., König, M., Choi, K., & Sauro, H. M. (2023). libRoadRunner 2.0: a high performance SBML simulation and analysis library. *Bioinformatics*, 39(1).

<https://doi.org/10.1093/bioinformatics/btac770>

Wilensky, U., & Reisman, K. (2006). Thinking Like a Wolf, a Sheep, or a Firefly: Learning Biology Through Constructing and Testing Computational Theories—An Embodied Modeling Approach. *Cognition and Instruction*, 24(2), 171–209.

[https://doi.org/10.1207/s1532690xci2402\\_1](https://doi.org/10.1207/s1532690xci2402_1)