

A Comprehensive Method of Single-Case Data Analysis: Interrupted Time-Series Simulation (ITSSIM)

Kevin R. Tarlow and Daniel F. Brossart
Texas A&M University

Single-case experimental methods are used across a range of educational and psychological research. Single-case data are analyzed with a variety of methods, but no statistic has demonstrated clear superiority over other methods. The time-series nature of single-case designs requires special consideration for baseline trend and autocorrelation when estimating intervention effect size. However, standard correction methods are limited because they assume precise statistical estimation of trend and autocorrelation. Unlike standard correction methods, Monte Carlo simulation methods can address the poor precision of single-case effect size indices. This paper presents the rationale for a new simulation method, Interrupted Time-Series Simulation (ITSSIM). A small field test was also conducted, and ITSSIM performed similarly to sophisticated multilevel methods for single-case research. ITSSIM is accessible as a free software application that requires no prior knowledge of statistical computing or syntax. ITSSIM may be used to estimate the effect size of a single interrupted time-series (AB design), and multiple ITSSIM effect size estimates may be combined via meta-analysis.

Impact and Implications

A new method of single-case data analysis, ITSSIM, uses computer simulation to test the effectiveness of school/psychological interventions at the individual subject level. ITSSIM gave results similar to sophisticated multilevel modeling methods when tested on a small example data set, suggesting ITSSIM is a viable tool for single-case research. This new method of analysis is easy for all investigators to use because ITSSIM software is free to download and requires no previous statistical computing experience.

Keywords: effect size, meta-analysis, Monte Carlo simulation, single-case, time-series

Supplemental materials: <http://dx.doi.org/10.1037/spq0000273.supp>

Single-case experimental designs have for decades been an important feature of educational and behavioral research, and they continue to gain popularity across all areas of psychology (Smith, 2012). Single-case designs can demonstrate causal intervention effects, similar to randomized controlled trials, and they can support the development of evidence-based interventions (APA Presidential Task Force on Evidence-Based Practice, 2006; Barlow & Hersen, 1984). Single-case methods are also accessible to a broad range of investigators because they do not require the considerable resources of large-sample studies (Barlow & Nock, 2009; Morgan & Morgan, 2001).

An unresolved issue in single-case research is the lack of consensus about the analysis of experimental data. Most single-case

studies evaluate data with visual analysis—a well-established approach with limited reliability (Brossart, Parker, Olson, & Mahadevan, 2006; Danov & Symons, 2008; DeProspero & Cohen, 1979; Harbst, Ottenbacher, & Harris, 1991; Lieberman, Yoder, Reichow, & Wolery, 2010; Park, Marascuilo, & Gaylord-Ross, 1990; Ximenes, Manolov, Solanas, & Quera, 2009). A meta-analysis of visual analysts' interrater agreement found interrater reliabilities were better than previously thought, but on average interrater agreement fell short of “minimally acceptable” levels (Ninci, Vannest, Willson, & Zhang, 2015). Statistical analysis complements visual analysis when it clarifies or confirms the findings of visual raters. Effect size statistics can also aid the quantitative synthesis of findings across studies. However, there is no consensus on how best to quantify the data patterns identified by visual inspection. This is a problem for investigators because different single-case statistics may lead to different effect size estimates for any given data set (Brossart et al., 2006; Parker & Brossart, 2003; Parker & Vannest, 2012).

An ideal single-case statistical method should meet several criteria. First, a single-case statistic should minimally yield an easy-to-interpret effect size estimate for comparing the magnitudes of intervention outcomes (Maggin & Odom, 2014). Second, a

This article was published Online First October 4, 2018.

Kevin R. Tarlow, School of Public Health, Texas A&M University; Daniel F. Brossart, Department of Educational Psychology, Texas A&M University.

Correspondence concerning this article should be addressed to Kevin R. Tarlow, School of Public Health, Texas A&M University, 4225 TAMU, College Station, TX 77843-4225. E-mail: krtarlow@gmail.com

single-case statistic should quantify the precision of the effect size estimate; for example, a small but highly reliable intervention effect may be more clinically meaningful than a large but unreliable intervention effect. Third, an ideal single-case statistic should model interrupted time-series data accurately; it should incorporate the data parameters which are known to characterize participant responses. For decades researchers have discussed which parameters should be included in single-case data analysis. The most debated parameters involve processes which affect behavior over time, such as baseline trend and autocorrelation (Manolov, 2017; Tarlow, 2017; Wampold, 1988). Fourth, analytic methods should be accessible to as wide a range of investigators as possible (Parker & Vannest, 2012; Shadish, 2014b). This includes practitioners and applied researchers who may not possess advanced statistical training. Accessibility is enhanced when investigators can understand and apply statistical methods to their own data and interpret results appropriately in order to answer their questions. Fifth, given the growing interest in single-case meta-analysis, effect size statistics should facilitate the quantitative synthesis of experimental results across cases and studies. Meta-analysis is one area where statistical methods are expected to be superior to visual analysis (Beretvas & Chung, 2008; Busk & Serlin, 1992; Maggin & Odom, 2014; Scruggs & Mastropieri, 2001).

Though no current methods satisfy all criteria perfectly, statistical methods for single-case research tend to fall into one of two heterogeneous groups. The first includes relatively simple effect size indices well-suited for use by clinicians. These methods are easily calculated from a single interrupted time-series data set, and their interpretation is fairly intuitive. For example, many single-case statistics quantify the degree of data nonoverlap between baseline and intervention phases, with greater nonoverlap suggesting an effective intervention (Parker, Vannest, & Davis, 2011). Other straightforward effect size statistics do not estimate nonoverlap, but instead fit trend lines to time-series via regression or other methods to better estimate the effect of intervention when stable baselines are not attainable (Tarlow, 2017). These methods are often critiqued for being overly simplistic, and failing to achieve the third criterion described above—the accurate modeling of brief time-series data. These methods often capture one important feature of data (e.g., data nonoverlap) while ignoring another (e.g., autocorrelation, trend). As a result, popular methods may yield different effect size estimates for the same data set (Smith, 2012).

The second group of single-case statistics includes sophisticated methods that utilize multilevel modeling. These statistics allow researchers to pool data from across cases and estimate complex time-series data patterns with a high degree of precision (Shadish, Rindskopf, & Hedges, 2008; Ugille, Moeyaert, Beretvas, Ferron, & Van den Noortgate, 2012; Van den Noortgate & Onghena, 2003). Multilevel single-case methods have great potential for organizing the emerging field of single-case meta-analysis, as they are well-suited for synthesizing results of many single-case experiments. However, these methods are critiqued for being difficult to implement and interpret without advanced statistical training—that is, they fail the fourth criterion of accessibility (Parker & Vannest, 2012). Single-case research is a historically pragmatic discipline, and new methods which are hard to adapt to clinical settings may offer little practical value to investigators despite their statistical sophistication (Schlosser, Lee, & Wendt, 2008).

This paper introduces a new method of single-case data analysis called Interrupted Time-Series Simulation (ITSSIM). ITSSIM was developed to address the five criteria outlined above. Essential to ITSSIM is the assumption that *one observed single-case data set can be explained by many plausible intervention effects*. ITSSIM uses Monte Carlo simulation modeling to determine what effect size is most likely (and how reliable that effect is) based on the many possible conditions which could plausibly yield the observed data. In this way ITSSIM is somewhat like the first group of simpler statistics, in that it is easily applied and interpreted. Yet it is also like the second group of sophisticated multilevel methods, because it estimates intervention effects by pooling information across many simulated cases.

Outline of Study

This paper proceeds as follows: (a) limitations of standard single-case statistical control methods (for trend and autocorrelation) are discussed, and Monte Carlo simulation methods are presented as one method of addressing those limitations; (b) ITSSIM effect size calculation and interpretation are discussed, and limitations of ITSSIM are addressed; (c) a small field test of ITSSIM is conducted where data from a published single-case study are reanalyzed, and a meta-analysis of ITSSIM effect sizes is compared to five multilevel methods; (d) the results of the field test are reviewed and recommendations for ITSSIM's use are presented.

Limitations of Statistical Control in Single-Case Data Analysis

Before ITSSIM is formally introduced, some limitations of conventional single-case statistical control methods are illustrated below. These limitations will establish the rationale for simulation-based effect size statistics like ITSSIM. Two control problems—baseline trend and autocorrelation—involve analyzing data that has some time-dependent process which confounds a straightforward preintervention/postintervention (or A/B) comparison.

Baseline Trend

For the baseline trend problem, consider the dilemma of the participant who is already improving before the experimental intervention is introduced. If the participant is improving during the baseline phase, and then the participant fully recovers during the intervention phase, one cannot infer the recovery was due solely (or even partly) to the intervention. Perhaps the participant would have recovered without treatment. Or perhaps the participant would have recovered more quickly were it not for some unexpected noxious effect of the treatment.

Single-case investigators have proposed a number of statistical trend control methods to account for baseline trend (Manolov, 2017). These methods include regression models, the Extended Celeration Line (ECL), the Mean Phase Difference (MPD), Baseline Corrected Tau, Tau-U, and others (Tarlow, 2017). Most methods control trend by estimating a baseline trend parameter and statistically removing the effect of the estimated trend from both A and B phases. In theory, the corrected baseline phase data no longer contain observable trend, and thus whatever intervention

effects remain in the corrected B phase data are assumed to exist beyond the influence of baseline trend. The effect size estimated from trend-corrected data is therefore considered the “true” intervention effect. There is a general consensus about this approach among single-case statistics that account for baseline trend, though the methods of trend estimation and correction vary from statistic to statistic. Most (but not all) trend correction methods assume linear trend patterns; trend may be estimated with parametric or nonparametric statistical models; it may be assumed a priori that all data sets are trended and in need of correction, or it may be assumed that data are stable and not trended unless proven otherwise.

Unfortunately, despite the variety of correction methods, few single-case statistics explicitly account for the reliability or unreliability of baseline trend estimates. This practice raises several concerns. It is unclear whether trend should be corrected in all cases, or only in cases where some criterion is met—such as statistical significance of trend, magnitude of trend, or some combination of both. It is plausible that, given the brief time-series common in single-case designs, observed trends may be explained by the chance variability of the small sample data, and not to true preexisting baseline trends. Parker, Vannest, Davis, and Sauber (2011) recommended that investigators control trend when an observed baseline trend coefficient falls above a given cutoff point. Tarlow (2017) recommended trend correction only when observed baseline trends were statistically significant; essentially, a null hypothesis of no trend must be rejected before the investigator alters observed data. Other approaches disregard hypothesis testing completely, and incorporate a baseline trend correction automatically and regardless of the trend’s magnitude or statistical signif-

icance (Allison & Gorman, 1993; Huitema & McKean, 2000b; Manolov & Solanas, 2009, 2013; White & Haring, 1980).

Trend correction is risky when the precision of the trend estimate is unknown. To illustrate, consider the hypothetical time-series in Figure 1. This time-series of 100 data points was generated from random normally distributed data points with unit variance ($s^2 = 1.00$) and an increasing linear trend with slope $\beta = 0.20$. Suppose an investigator estimates the trend in a baseline of 10 points (assume the other 90 points are not observed). How accurately would that estimate reflect the true trend parameter of $\beta = 0.20$?

The Figure 1 time-series is partitioned into brief time-series of 10 points each. Suppose each sample of 10 points is corrected for baseline trend via ordinary least squares (OLS) regression—but the correction is based only on the estimated trend in the 10-point sample, not the “unknown” true parameter of $\beta = 0.20$. In some cases, the estimated trend will be less than the true trend, and in other cases it will be greater than the true trend. Two corrected time-series are illustrated in the bottom graph of Figure 1. In one corrected series, baseline trend was controlled with the trend estimates of each 10-point sample; in the other, trend was controlled with the true trend.

Although the two corrected time-series in Figure 1 appear visually similar, there is one important difference. The variance of the time-series corrected with the true slope parameter is unchanged from the original data, $s^2 = 1.00$. However, the variance of the time-series corrected with the sampled slope estimates is reduced by over 20%, $s^2 = 0.78$. Because parameter estimates are never perfectly precise, 20% of the data set’s information was mistaken for trend and discarded. Any inferences made from the




Figure 1. Hypothetical time-series with corrections for trend ($n = 100$). Original data is displayed in the top graph. In the bottom graph, trend was removed using the known slope parameter, $\beta = 0.20$ (dark line), and removed using an estimated slope from each interval of 10 data points (light line), resulting in an unwanted reduction of variance.

overcorrected data should take into account this reduction in variance. Failure to do so would distort conclusions about intervention effectiveness.

Here is the crucial point of this heuristic example: *nearly all baseline trend control methods in single-case research assume trend estimates are perfectly precise*. In most single-case statistics, a trend coefficient is estimated and then data are corrected and analyzed. The unreliability of trend estimates is rarely considered. Perhaps more importantly, the precision of trend estimates decreases as the number of baseline phase data points decreases. For example, if the [Figure 1](#) data were partitioned into samples of five points instead of 10 points, the corresponding reduction in variance from trend correction would increase from 20% to nearly 40% ($s^2 = 0.63$). This result should be of great concern to single-case investigators, who may be drawing conclusions from “trend controlled” data that have little resemblance to reality.

Autocorrelation

Autocorrelation describes the degree of serial dependence in time-series data. In autocorrelated data, past fluctuations in the time-series predict future fluctuations. This violates the assumption of independence shared by most methods. Autocorrelation distorts visual analysis ([Jones, Vaught, & Weinrott, 1977](#); [Matyas & Greenwood, 1990](#)) and statistical analysis of single-case data ([Brossart et al., 2006](#); [Manolov & Solanas, 2008](#); [Tarlow, 2017](#)).

Single-case methods of autocorrelation correction rarely account for the lack of precision in autocorrelation estimation, just as

with trend correction. Some have recommended “cleansing” autocorrelation from data (also called “back-casting” or “pre-whitening”) before an effect size is calculated (e.g., [Parker, Cryer, & Byrns, 2006, 2011](#)). The process is similar to trend correction: an autocorrelation model is estimated from observed data, then data are “corrected” to remove the influence of the estimated autocorrelation, and then an effect size is calculated. Unfortunately, autocorrelation estimation requires a large number of data points to yield reliable estimates, far more than are typically available to single-case investigators ([Box & Jenkins, 1970](#); [Glass, Willson, & Gottman, 1975](#)). Back-casting does remove the autocorrelation observed in a sample of data points—just as trend control removes trend in an observed baseline—but the procedure may lead to misleading conclusions about intervention effects. Statistical removal of autocorrelation (or baseline trend) from data does not guarantee results are reliable or valid.

The hypothetical [Figure 2](#) data illustrates the hazard of cleansing autocorrelation from brief time-series data with ARIMA model fitting (i.e., back-casting). Similar to the [Figure 1](#) example, the [Figure 2](#) data was generated with a lag-1 autoregressive parameter $\phi_1 = 0.20$ and unit variance ($s^2 = 1.00$). Two corrected time-series are presented. In one, autocorrelation is cleansed using the true ϕ_1 parameter; the variance of the time-series is nearly unchanged, $s^2 = 0.96$. In the other corrected time-series, autocorrelation is cleansed within each 10-point sample using an autocorrelation estimate; the result is a reduction in variance to $s^2 = 0.79$, or again roughly 20%. The results from [Figure 1](#) and [Figure 2](#) beg the




Figure 2. Hypothetical time-series with corrections for autocorrelation ($n = 100$). Original data is displayed in the top graph. In the bottom graph, trend was removed using the known autocorrelation parameter, $\phi_1 = 0.20$ (dark line) and removed using an estimated slope from each interval of 10 data points (light line), resulting in an unwanted reduction of variance.

question: If a large portion of information contained in the observed data is erroneously discarded during the statistical correction process, how useful are the effect size estimates calculated from corrected data?

Simulation Methods Do Not Assume Perfectly Precise Control

As the Figure 1 and Figure 2 examples demonstrate, statistical control of brief time-series data is limited by the precision of baseline trend and autocorrelation estimates. However, many single-case statistics assume correction methods are perfectly precise—though we know this assumption is false, particularly for autocorrelation but also for baseline trend. In most statistical control methods, a baseline trend and/or autocorrelation coefficient is calculated and used to adjust observed data for effect size estimation, but those final results are almost never reported within the context of the reliability, or unreliability, of the correction parameter estimates.

Computer-aided simulation methods have the potential to address this limitation. Rather than calculate one corrected data set and one effect size estimate, many plausible effect sizes may be calculated from many plausible corrected data sets. Rather than assume correction methods are perfectly precise, a simulation approach assumes many intervention effects (and trend parameters and autocorrelation parameters) could have plausibly produced the observed data. It may in fact be quite meaningful for investigators to consider the range of plausible intervention effects that could explain their observed data. A simulation approach therefore offers context for interpretation absent in most single-case statistical analyses.

Interrupted Time-Series Simulation: ITSSIM

Interrupted Time-Series Simulation (ITSSIM) follows a three-stage process to yield an effect size estimate for an interrupted time-series (AB) design. The three stages are (a) parameter estimation, (b) time-series simulation, and (c) effect size calculation. In the first stage, estimates and standard errors are calculated for seven data parameters: A phase level, B phase level, A phase trend, B phase trend, A phase error variance, B phase error variance, and across-phase autocorrelation. These parameter estimates are used to construct two models. A *null effect model* (based on the A phase data) describes the participant's behavior prior to intervention. An *experimental effect model* (based on the B phase data) describes the participant's behavior during/after intervention. In the second stage of ITSSIM, hypothetical data sets are simulated from the null effect model and experimental effect model. This stage yields thousands of artificial time-series that represent the null effect and intervention effect within a range of plausible parameter values. Essentially, the simulation stage yields two distributions of time-series, one representing the range of plausible outcomes without intervention, and one representing the range of plausible outcomes with intervention. In the third stage, the null effect distribution and the experimental effect distribution are compared to calculate a standardized mean difference effect size, d . This effect size statistic is suitable for meta-analysis.

ITSSIM was developed to estimate effect sizes for AB single-case designs, as in a multiple baseline design. However, more

complex designs, such as an ABAB reversal, are also appropriate for ITSSIM analysis. When a design includes multiple AB phase contrasts, a separate ITSSIM effect size is calculated for each AB phase contrast (therefore an ABAB design will yield two sequential effect size values). These effects may be reported separately or combined via meta-analysis. The synthesis of multiple ITSSIM effect sizes is demonstrated later in this paper, following the explanation of ITSSIM methodology.

How Is ITSSIM Different From Other Single-Case Statistics?

It is useful to keep in mind three differences between ITSSIM and other single-case statistics while reviewing the ITSSIM procedure below. First, the ITSSIM effect size is not calculated from exact estimates of level, trend, and so forth, nor is it calculated directly from observed data. Instead, the ITSSIM effect size is calculated from many simulated time-series that are based on a range of conditions which could plausibly account for the observed data. For example, rather than fit a single estimate of baseline trend to a small sample of data points, ITSSIM models a range of many baseline trend values that could fit the data. If observed data lead to very precise parameter estimates with small standard errors (e.g., "From the twenty baseline data points, I am confident that the true baseline trend coefficient lies somewhere close to $\beta = 0.20$ "), then simulated data sets will reflect that confidence. In this way, unlike many other methods, ITSSIM incorporates the reliability of the parameter estimates into effect size estimation. ITSSIM thus "rewards" investigators for time-series that are stable (i.e., less trend, variance) and long (i.e., larger n), though it can be implemented with time-series of any length. (When extended/stable baselines are not feasible for practical or ethical reasons, the meta-analysis of multiple ITSSIM effect sizes can similarly provide more precise estimates, as discussed below.)

Second, ITSSIM effect size estimates are calculated from predicted data rather than corrected data. Many single-case statistics model trend or autocorrelation by correcting observed data points to make comparisons between A and B phases more tenable. This approach is used because A phase and B phase observations are recorded at different time intervals, and time confounds direct A-to-B data comparisons when data points are changing deterministically over time. ITSSIM instead simulates B phase data from both the null effect model and the experimental effect model. These artificial B phase time-series, which are simulated predictions, are directly comparable without correction because they are simultaneous predictions for the same time interval. This strategy is similar to single-case effect size statistics proposed by Gorsuch (1983), Allison and Gorman (1993), and Manolov and Solanas (2013), where B phase observations are compared with data predicted by baseline observations.

Third, the ITSSIM effect size is not calculated by a comparison of A phase and B phase data. Instead, an ITSSIM effect size is calculated from *distributions* of simulated time-series. The null effect distribution (based on A phase data) and the experimental effect distribution (based on B phase data) are both composed of simulated time-series—predictions—for the B phase time interval. These distributions do not directly represent A phase and B phase data, but they do represent a range of plausible data parameters based on the A and B phases, respectively.

The three stages of the ITSSIM procedure are discussed further below, and the [online supplemental materials](#) demonstrates ITS-SIM calculations with a hypothetical data set. However, [Figure 3](#) presents a simplified illustration of the three ITSSIM stages.

Stage I: Parameter Estimation

A statistical model represents an implicit model of reality (Thompson, 2006). Models based on unrealistic parameters—or models based on an insufficient number of realistic parameters—will do a poor job of describing reality. There has been considerable discussion about which parameters should be included in the analysis of single-case data. Many agree the most fundamental parameters in an interrupted time-series (AB) design are level and trend (e.g., Kazdin, 1982). Clinical treatments are often expected to affect either the level (i.e., magnitude) of the outcome variable, the trend (i.e., slope) of the outcome variable over time, or both.

Level and trend alone may be insufficient to comprehensively model single-case data. As discussed earlier, autocorrelation is often cited as an important parameter in single-case data analysis (Campbell & Herzinger, 2010; Hartmann et al., 1980; Jones et al., 1977; Smith, 2012; Shadish, 2014b). The serial dependence of sequential data points may mask or exaggerate underlying intervention effects (Brossart et al., 2006; Jones, Weinrott, & Vaught, 1978; Matyas & Greenwood, 1990). Variability, that is, the stability, or amount of “bounce” in time-series data points, is also recognized as an important parameter in single-case analysis (Hayes, 1981; Parsonson & Baer, 1992; Kazdin, 1982; Kratochwill et al., 2010). Generally, data that are more stable offer more compelling evidence of intervention effect.

Four parameters were selected for the ITSSIM model: level, trend, variability, and autocorrelation. In an ITSSIM analysis, coefficients for level, trend, and variability are estimated separately for the A and B phase; autocorrelation is estimated for the whole series after controlling for the other three variables.¹ This procedure leads to an estimation of seven total coefficients: A phase level, A phase trend, A phase variability, B phase level, B phase trend, B phase variability, and across-phase autocorrelation.

The order in which parameters are estimated is important. For example, the presence of trend can greatly distort estimates of autocorrelation (Huitema & McKean, 1998; Yue, Pilon, Phinney, & Cavadias, 2002). Trend should be estimated and then removed from the time-series in order to accurately measure the degree of autocorrelation. Within-phase variability should also be controlled before estimating cross-phase autocorrelation, because heteroscedasticity within a time-series could distort the estimated autocorrelation coefficient. Given these constraints, ITSSIM parameter estimation is conducted with the following steps.

Estimate level and slope parameters: Theil-Sen robust regression. Regression methods are widely used in single-case analysis and simulation models (Huitema & McKean, 2000b). Regression is useful in single-case data analysis because it models both trend and level changes. However, ordinary least squares (OLS) regression is limited by its assumptions and sensitivity to outliers (Brossart, Parker, & Castillo, 2011). Nonparametric robust regression is a promising alternative to OLS regression in single-case research because it makes fewer distributional assumptions and is less sensitive to outliers. Theil-Sen nonparametric regression (Sen, 1968; Theil, 1950) is used in ITSSIM analysis because

it is robust, yields a relatively small standard error, and is efficient in small samples—or in this case, brief time-series (Wilcox, 1998, 2001). Theil-Sen regression has been recommended for single-case data analysis (Tarlow, 2017; Vannest, Parker, Davis, Soares, & Smith, 2012) and has been used in other fields to analyze autocorrelated, non-normal, monotonically trended time-series data (e.g., Yue et al., 2002).

Theil-Sen level and trend (i.e., intercept and slope) coefficients are therefore estimated for observed A and B phase data separately. Standard errors for these coefficients are then calculated using a bootstrap procedure (Wilcox, 2001).

Estimate variability parameters: Error variance. Variability within each phase is calculated as the variance of the Theil-Sen regression residuals. Both the phase error variance (s^2) and its standard error are calculated using OLS methods.

Estimate lag-1 autocorrelation coefficient: Unbiased r_1 . An unbiased r_1 estimator (Ferron, 2002; Huitema & McKean, 1998, 2000a) is calculated across both A and B phase data after standardizing the Theil-Sen residuals (i.e., dividing residuals by their within-phase standard deviation). There are several ways to estimate the standard error of the r_1 statistic, with Bartlett’s (1946) method being the most popular. However, Huitema and McKean (1998) demonstrated that Bartlett’s method was biased in small samples; their modified $\text{var}(r_1)$ estimator is used in the ITSSIM model.

At the end of the parameter estimation stage, the ITSSIM analysis will yield seven coefficients and their corresponding standard errors. These parameter estimates comprise the *null effect* and *experimental effect* models. [Table 1](#) presents hypothetical null effect and experimental effect models.

Stage II: Time-Series Simulation

It is useful to note that, following the parameter estimation stage, the investigator is left with a time-series that has unit variance and an estimated autocorrelation value. It would be possible at this point to work backward through the parameter estimation stage, restoring the within-phase error variance and then blending the residuals with the Theil-Sen regression coefficients, to return these standardized residuals to the original data set. In fact, using only the parameter estimates and standard errors from the first stage (e.g., [Table 1](#)), it would be possible to recreate a time-series very similar to the original data set, with differences attributable only to random variability in the initial standardized residuals. This process of “working backwards” is essentially how the time-series simulation stage of ITSSIM is conducted.

Randomly sample simulation parameters. A random time-series is created with a preset degree of autocorrelation, and then within-phase variability, trend, and level are added. But instead of using the parameter estimates of the original data set (e.g., [Table 1](#)), new coefficients are randomly drawn from the estimated parameters’ sampling distributions. These new level, trend, variability, and autocorrelation coefficients—though not the exact values calculated from the observed data—plausibly represent the true

¹ Autocorrelation is estimated across both phases because of the low power of autocorrelation estimators in brief time-series (see discussion above, e.g., [Figure 2](#)). The ITSSIM model assumes autocorrelation is relatively stable from phase to phase.

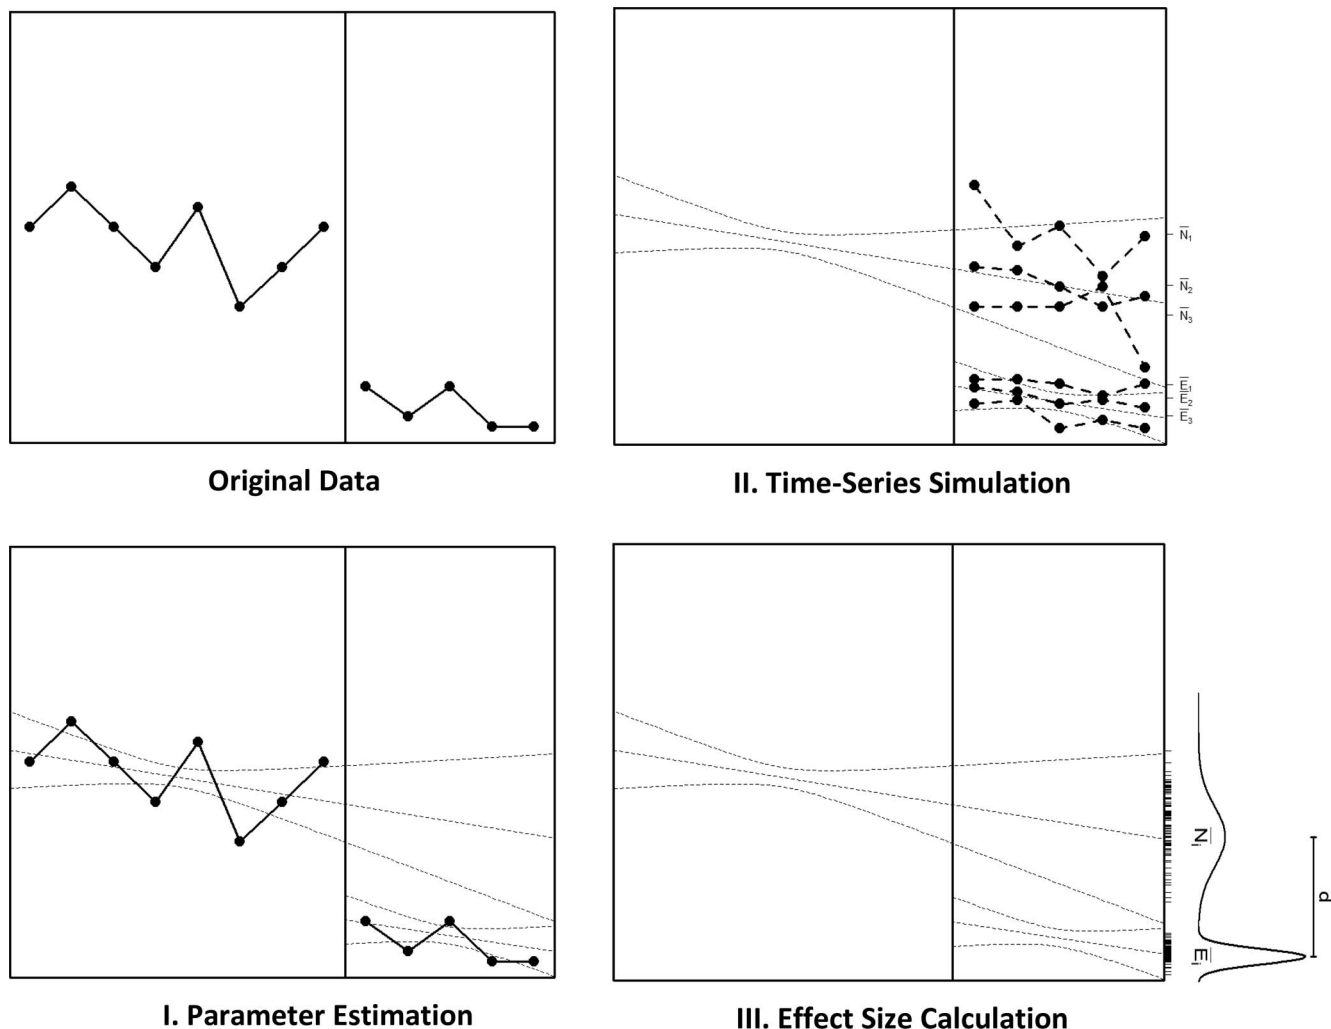


Figure 3. Simplified illustration of ITSSIM stages.

nature of the participant’s response pattern given the known precision of the original parameter estimates. This procedure is arithmetically simplified because the sampling distributions of Theil-Sen, variance, and lag-1 autocorrelation coefficients are normal (Anderson, 1942; Cox, 1966; Mann, 1945; Sen, 1968). Thus, about

two thirds of the randomly sampled coefficients will fall within $\pm SE$ from the parameter estimate ($SE =$ standard error).

This step is iterated 100,000 times each with the null model and experimental model estimates, yielding two sets of intercept, slope, error variance, and autocorrelation coefficients.² The first set of 100,000 coefficients represents the plausible patterns of participant response in the absence of intervention, that is, the null effect. The second set of 100,000 randomly sampled coefficients represents the plausible patterns of participant response under the effect of intervention, that is, the experimental treatment effect. Using each set of sampled coefficients, one could create 100,000 AB time-series with the specified level, trend, variability, and autocorrelation. However, instead of creating AB data sets (like the original

Table 1
Hypothetical Parameter Estimates in an ITSSIM Analysis

Model	Coefficient	SE
Null model		
Intercept	5.375	(2.752)
Slope	-0.083	(0.540)
s	2.092	(0.501)
Exp. model		
Intercept	0.375	(1.604)
Slope	-0.125	(0.489)
s	0.845	(0.295)
Autocorrelation		
r_1	0.133	(0.244)

² Different iteration procedures were tested during the development of ITSSIM. It was determined that sampling 10,000 time series from each model yielded acceptable results that did not vary with the randomization seed. However, 100,000 iterations yielded very stable results and did not significantly increase software execution time. Therefore, the more intensive procedure was used, though 10,000 iterations is considered sufficient.

observed data set), the ITSSIM method uses these simulated coefficients to create simultaneous B phase time-series.

Simulate B phase time-series. Recall from earlier that ITSSIM does not control for time-dependent data patterns like trend or autocorrelation. Instead, ITSSIM makes predictions using estimates of level, trend, variability, and autocorrelation. Predictions based on the A phase data are compared with predictions made from the B phase. To make the two sets of predictions directly comparable, they are simulated for the same interval of time. After coefficients are randomly drawn from the parameters' sampling distributions, a corresponding B phase time-series is generated from each of the 100,000 simulated null effect models and simulated experimental effect models. The result is 100,000 null effect time-series and 100,000 experimental effect time-series, all corresponding to the same interval of time—the time interval of the B phase.

Stage III: Effect Size Calculation

To find a treatment effect size, the mean of each simulated time series is first calculated. This procedure yields two distributions of means—one for the means of the simulated null effect time-series and one for the means of simulated experimental effect time-series. The distributions of means are approximately normal because of the Central Limit Theorem. Thus, as with any large sample study of control group and experimental group data, the ITSSIM null effect distribution and the experimental effect distribution may be compared to find a standardized mean difference, d (see Figure 1). The effect size, d , may also be converted to an equivalent r or R^2 effect size. The unstandardized mean difference, D , may also be reported as an effect size.

Put another way, the simulation process in Stage II generates a range of plausible mean scores based on the data in each phase (i.e., the null effect distribution and the experimental effect distribution). If A phase predictions are close to B phase predictions, that is, if there is overlap between the two distributions, then the effect size is considered to be relatively small. If A phase predictions are far apart, resulting in less overlap between distributions, then the effect size will be larger.

Assumptions and Limitations

ITSSIM has several statistical assumptions and limitations which investigators should consider before analyzing their data. The violations of some assumptions, such as the expectation of stable autocorrelation across phases, may not greatly affect results (the degree of autocorrelation observed in brief time series is relatively small, so cross-phase differences are also expected to be minimal). The violation of other assumptions, like the assumption of linear trends, may lead to less accurate results. ITSSIM assumptions include the following:

- A linear trend model is assumed. However, the Theil-Sen estimator finds a best-fit line for any monotonic (i.e., linear or nonlinear) trend pattern, so the error introduced by nonlinear trends are expected to be less severe than with OLS regression.
- Lag-1 autoregressive error is assumed to be unaffected by treatment. This assumption is consistent with the general view that autocorrelated behavior is a function of the

participant and design rather than an effect of treatment (e.g., Baer, 1988).

- Only lag-1 autoregressive error structures (ϕ_1) are modeled. Moving average, integrated, and higher-order error structures are not considered. Other single-case statistics which model autocorrelation often assume a lag-1 autoregressive error structure (e.g., Moeyaert, Ferron, Beretvas, & Van den Noortgate, 2014; Shadish, Hedges, et al., 2014).
- The outcome variable is assumed to be continuous. Investigators analyzing time-series with count data (or other noncontinuous data) should exercise caution in interpreting results. For example, “percentage of time on task” would be considered a continuous variable and would therefore be appropriate for ITSSIM analysis; however, “number of mands” would be considered a count variable and may give inaccurate results. Statistics that assume continuous outcome variables are frequently applied to count data, though the effects of this violation assumption are not clear. Future updates of ITSSIM software should include options for noncontinuous data modeling.
- Normally distributed data are assumed in the current version of ITSSIM software. Future software releases should include options for modeling non-normal data distributions (e.g., Poisson, binomial, etc.).
- Level, trend, error variance, and autocorrelation parameters are assumed to be independent of each other. There is little published research to suggest whether this assumption is tenable. It is expected that relationships between these parameters vary from study to study, though more exploration is needed to determine whether and how violations to this assumption would impact statistical results.
- The current version of ITSSIM software (1.0) does not permit missing data—that is, all observations are assumed to be recorded at equally spaced time intervals. This assumption is shared by many nonoverlap and rank-based methods, whereas regression-based methods tend to accommodate missing data better. Data sets with missing data may still be analyzed with ITSSIM; however, results may be increasingly inaccurate as the percentage of missing data points increases. Further study is needed to better assess the effect of missing data on ITSSIM, and future software updates should accommodate incomplete data sets.

ITSSIM Effect Size Interpretation

As noted above, ITSSIM effect sizes may be converted into different metrics depending on the preferences and needs of the investigator. ITSSIM software outputs the effect sizes frequently used in clinical practice, intervention outcomes research, and meta-analysis. A guide for reporting and interpreting these effect size statistics is presented below, although investigators may find other equally valid interpretations more suited to their needs.

D. The ITSSIM unstandardized mean difference, D , indicates the average improvement during treatment, accounting for level, trend, variability, and autocorrelation. D is reported in the original metric of the outcome variable, for example, “ D fewer intervals of disruptive behavior” or “ D percent increase in time on task.” D is

not bounded (values have no set range) and may be positive or negative. D may be particularly relevant to practitioners who wish to report their treatment results in a clinically relevant way, though it is less useful for synthesizing results across studies with different outcome measures.

d. The ITSSIM standardized mean difference, d , indicates the average improvement during treatment, accounting for level, trend, variability, and autocorrelation. d is a standardized metric (D /standard deviation of simulation distributions), and therefore d values may be compared across studies that use different outcome measures. d is not bounded (though most values are expected to fall within ± 5) and may be positive or negative.

r. The ITSSIM correlation coefficient, r , indicates the association between intervention and outcome, accounting for level, trend, variability, and autocorrelation. r is a standardized metric, and therefore r values may be compared across studies that use different outcome measures. r is bounded between -1 and $+1$. A positive r indicates the outcome variable increased during intervention after accounting for level, trend, variability, and autocorrelation (i.e., a positive correlation between intervention and outcome); a negative r indicates the outcome variable decreased during intervention (i.e., a negative correlation between intervention and outcome).

R^2 . The ITSSIM R^2 effect size indicates the proportion of outcome variance explained by the intervention effect, accounting for level, trend, variability, and autocorrelation. R^2 is a standardized metric, and therefore R^2 values may be compared across studies that use different outcome measures. R^2 , which may be reported as a percentage, is bounded between 0 and $+1$. As R^2 decreases toward 0, less of the outcome is explained by the intervention; as R^2 approaches $+1$, more of the outcome is explained by the intervention.

A note about effect size conventions: “small,” “medium,” and “large.” It is expected that ITSSIM effect sizes will vary with the design, outcome measure, and domain of any given single-case experiment. For example, a “small” effect size in an applied behavior analysis (ABA) study might be considered quite large in a single-case study of psychotherapy outcomes. In addition, many single-case statistics yield results that do not conform to Cohen’s (1988) popular conventions for “small,” “medium,” and “large” effect size (Parker et al., 2006). Therefore, when assessing ITSSIM effect size magnitude, two interpretive strategies are recommended in lieu of relying on standard cut-off scores (an operation which Cohen described as “fraught with many dangers”). First, single-case investigators, who are presumably experts in their respective fields of research, can establish their own interpretive benchmarks by applying specialized knowledge of participant outcomes and visual analysis of data. These investigators are ultimately in the best position to determine what a “large” change is in a clinically meaningful sense. Second, statistical methods like ITSSIM should be used not only to determine whether the effect of Intervention X is “large,” but also (and perhaps more importantly) to determine whether the effect of Intervention X is *larger* than the effect of Intervention Y. Identifying and replicating the largest intervention effects is one goal that statistical methods are well suited to address, in addition to quantitative synthesis of research findings.

ITSSIM Software

ITSSIM is accessible to single-case investigators via free software download at <http://ktarlow.com/stats> (Tarlow, 2018a). This is a free standalone program that does not require additional statistical computing software. When opened, the program prompts the user to enter A phase and B phase data. After the observed data are inputted, the program estimates the model parameters and performs the simulations. All coefficients, standard errors, and effect size statistics are reported in the software output. Use of ITSSIM software requires no computing code or syntax, and should be easy to implement for applied researchers and practitioners without statistical computing experience. An example of ITSSIM software input and output is presented in Figure 4.

Meta-Analysis With ITSSIM

Researchers are increasingly interested in the meta-analysis of single-case designs. Synthesizing results across cases is one area where statistical methods are expected to be superior to visual analysis (Beretvas & Chung, 2008; Busk & Serlin, 1992; Maggin & Odom, 2014; Scruggs & Mastropieri, 2001). Many single-case statistics do not possess the properties necessary for meta-analysis; however, ITSSIM is ideal for aggregating results across studies.

Within-case meta-analysis. There are three ways in which single-case effect sizes may be combined. First, effects may be combined within one individual. For example, consider an ABAB design. An ITSSIM analysis will yield one effect size for the first AB phase contrast and one effect size for the second AB phase contrast. These within-case effects could be combined to yield one overall effect size. However, the problem of autocorrelation in time-series data described above also applies to the analysis of multiple sequential effects sizes within one individual. Standard meta-analytic methods assume combined effects are independent.

Across-case meta-analysis. The second way meta-analysis may be applied to single-case research is the synthesis of effect sizes across multiple cases, as in a multiple baseline design (the most common single-case experimental design; Smith, 2012). For example, consider a multiple baseline study with three AB time-series for three participants; an effect size could be calculated for each AB phase contrast, and then meta-analytically averaged across all three cases. This type of single-case meta-analysis is less likely to include violations of statistical assumptions, and standard meta-analytic methods may be applied to combine effect sizes across cases. The following section of this paper includes a demonstration of this second type of single-case meta-analysis with ITSSIM. In this field test, an ITSSIM effect size was calculated for each AB phase contrast and then combined across cases.

Between-groups meta-analysis. Third, groups of cases may be synthesized via meta-analysis. Between-groups meta-analyses are also described as “subgroup analyses” (Borenstein, Hedges, Higgins, & Rothstein, 2009). For example, an investigator may wish to aggregate the findings across several studies, where each study includes several cases. In this scenario, it is reasonable to assume cases within the same study may be more similar to each other than to cases from other studies. Put another way, the effect sizes in the meta-analysis are not statistically independent, because within-study and between-study variance both contribute to the distribution of effects. Subgroup analyses can accurately model

```

ITSSIM: Interrupted Time-Series Simulation
Version 1.0 (2018)
ktarlow.com

Copyright (c) 2018 Kevin R. Tarlow

Press ENTER to input data...

Enter the number of A phase (baseline) points: 8

    Enter the A phase points one at a time...

    A[ 1]: 5
    A[ 2]: 8
    A[ 3]: 5
    A[ 4]: 4
    A[ 5]: 6
    A[ 6]: 3
    A[ 7]: 4
    A[ 8]: 9

Enter the number of B phase points: 5

    Enter the B phase points one at a time...

    B[ 9]: 2
    B[10]: 0
    B[11]: 1
    B[12]: 0
    B[13]: 0

ESTIMATING PARAMETERS...

                est      (SE)
Null:  int      5.375 (2.752)
       slope  -0.083 (0.540)
       s       2.092 (0.501)

Exp:   int      1.500 (1.232)
       slope  -0.500 (0.385)
       s       0.652 (0.251)

Auto:  r1       0.224 (0.244)

SIMULATING CASES...

Null Effect mean 4.453 (3.424)
Exp Effect mean  0.000 (0.862)

CALCULATING EFFECT SIZE...

D (unstd)      -4.453
d (std)        -1.602 (0.651)
r              -0.615
R-squared      0.378 (decrease)
p              0.192

```

Figure 4. Example of ITSSIM console with input and output. Software is available for download at <http://ktarlow.com/stats> (Tarlow, 2018a).

these nested relationships (effect sizes within groups). Multilevel models can also address autocorrelation in the meta-analysis of within-subject effect sizes (the first type of meta-analysis described above). Further exploration of multilevel modeling exceeds the scope of this paper; however, Van den Noortgate and

Onghena (2003, 2008), Jenson, Clark, Kircher, and Kristjansson (2007), and Moeyaert et al. (2014) offered detailed presentations of single-case multilevel solutions to these problems. Meta-analysis models that account for moderator relationships could also be applied to the synthesis of ITSSIM effect sizes across

groups of cases, assuming cases within each group are independent (Borenstein et al., 2009).

Meta-analysis software. Many statistical software packages allow for quick aggregation of effect size statistics (e.g., RevMan, Comprehensive Meta-Analysis, Open Meta-Analysis). For this study, R code developed by Tarlow (2018b; <http://ktarlow.com/stats>) and adapted from Borenstein et al. (2009) was used. To use this R code, the user inputs a vector of effect size and variance values (which in this case are found in the ITSSIM software output)

```
d <- c(1.63, 3.06, 1.90, 2.28, 2.76, 4.48, 4.66, 1.60, 2.77)
v <- c(0.42, 0.67, 0.53, 0.56, 0.50, 1.21, 1.14, 0.38, 0.50)
```

After loading the code, effect size vector, and variance vector into the R console, the user then calls the random-effects meta-analysis function `remeta()` with the syntax `remeta(d,v)`.

The results include a random-effects mean, as well as a p value, confidence intervals, and heterogeneity statistics (Q , I^2).

Field Test of ITSSIM

The primary goal of this paper is expository, to introduce the rationale for computer-intensive simulation methods in single-case research and explain ITSSIM's use and interpretation. However, an initial field test of ITSSIM was also conducted to determine whether it performed similarly to other sophisticated single-case analytic methods. The *Journal of School Psychology* published a special issue (Shadish, 2014a) in which five research teams independently analyzed the same single-case data from a study by Lambert, Cartledge, Heward, and Lo (2006) using five different multilevel methods. This data set was reanalyzed with ITSSIM and an across-case meta-analysis, and results were compared with the five previously published methods.

Summarizing the special issue, Shadish (2014a) found the results of the five comparison methods "reasonably consistent with each other" (p. 112), and similar ITSSIM effect size estimates would suggest ITSSIM is suitable for use in single-case research. Like ITSSIM, all five previously published methods produced unstandardized effect sizes; however, not all of the comparison methods gave an equivalent standardized effect. Two approaches (Shadish, Hedges, et al., 2014; Swaminathan, Rogers, & Horner, 2014) gave standardized mean differences like the ITSSIM d statistic, which are straightforward to compare across studies. The remaining three methods (Moeyaert et al., 2014; Rindskopf, 2014; Shadish et al., 2014) gave log-odds ratios as the standardized outcome metric, which are less intuitive when describing treatment effects. Only two of the five methods (Moeyaert et al., 2014; Swaminathan et al., 2014) modeled both trend and autocorrelation. One method (Shadish et al., 2014) used generalized additive models (GAMs) to fit nonlinear trends. Two methods (Rindskopf, 2014; Swaminathan et al., 2014) used Bayesian estimation because Bayesian methods are well-suited for analyzing small samples. Two of the approaches (Moeyaert et al., 2014; Shadish et al., 2014) involved testing multiple analytic models for goodness of fit and/or agreement.

Results of the Field Test

Data from the nine students in Lambert et al.'s (2006) study were analyzed with ITSSIM simulation software (Tarlow, 2018a).

Table 2
ITSSIM Standardized Effect Sizes for Lambert et al. (2006) Data

Participant	A1B1 Effect		A2B2 Effect	
	d	SE	d	SE
Student A1	1.63	0.65	3.71	0.80
Student A2	3.06	0.82	5.68	1.81
Student A3	1.90	0.73	3.72	0.85
Student A4	2.28	0.75	1.83	0.64
Student B1	2.76	0.71	3.20	0.72
Student B2	4.48	1.10	1.16	0.57
Student B3	4.66	1.07	6.89	1.32
Student B4	1.60	0.62	-0.39 ^{ns}	0.52
Student B5	2.77	0.71	0.83 ^{ns}	0.54
Random-effects mean	2.56	0.32	2.63	0.65

^{ns} Not statistically significant at $p < .05$.

The study used a reversal ABAB design for each student, so a total of 18 AB phase contrasts were analyzed. Essentially, two across-case meta-analyses were performed, one with all nine A1B1 effect size estimates, and one with all nine A2B2 effect size estimates. A random-effects meta-analysis model was used to aggregate the A1B1 and A2B2 effect size estimates. The R code by Tarlow (2018b) described above was used.³ Results of the ITSSIM analyses are presented in Table 2. The A1B1 standardized effects ranged from $d = 1.60$ to 4.66 ; all A1B1 effects were statistically significant at the $p < .05$ level. The A2B2 standardized effects ranged from $d = -0.39$ to 6.89 ; seven of the nine A2B2 effects were statistically significant. The random-effects mean of the A1B1 treatment effects was $d = 2.56$, $p < .001$, 95% CI [1.93, 3.19] ($Q = 12.70$, $df = 8$, $p = .122$); the random-effects mean of the A2B2 treatment effects was $d = 2.63$, $p < .001$, 95% CI [1.35, 3.90] ($Q = 56.23$, $df = 8$, $p < .001$).

An unstandardized treatment effect, D , was also calculated for each of the 18 AB phase contrasts. The A1B1 unstandardized mean effect was $D = 7.51$ fewer intervals of disruptive behavior; the A2B2 unstandardized mean effect was $D = 8.43$ fewer intervals of disruptive behavior. The mean unstandardized effects for ITSSIM and the other five methods is presented in Table 3, in addition to the mean standardized effects (for the methods that yielded an interpretable d -type effect size).

The results in Table 3 suggest ITSSIM yields slightly larger mean effect sizes than the multilevel comparison methods. The ITSSIM standardized mean effect sizes, d , were 0.05 to 0.29 standard deviations larger than the other reported effect sizes. ITSSIM unstandardized D results were larger than the comparison methods by about one to three intervals of disruptive behavior.

Discussion

The purpose of this study was to introduce a computer simulation method, ITSSIM, for the analysis of single-case experimental designs. ITSSIM incorporates baseline trend, level- and slope-change, variability, and autocorrelation parameters, making it more comprehensive than most single-case statistical models. In-

³ For demonstration purposes, the example R code provided for inputting effect size and variance vectors was based on the A1B1 effect size estimates in Table 2 (standard errors were converted to variances).

Table 3
Six Analyses of Lambert et al. (2006) Data

Study	Unstandardized effect size	Standardized effect size
Shadish, Hedges, and Pustejovsky (2014) ^a	5.46	2.51
Shadish, Zuur, and Sullivan (2014) ^b	6.70	—
Rindskopf (2014) ^{a,b}	5.70	—
Moeyaert et al. (2014): <i>continuous outcome</i>	5.10–5.76 (A1B1)	—
	4.92–5.77 (A2B2)	
Moeyaert et al. (2014): <i>logistic outcome</i> ^{a,b}	5.61	—
Swaminathan et al. (2014)	5.38 (A1B1)	2.47 (A1B1)
	5.03 (A2B2)	2.34 (A2B2)
ITSSIM	7.51 (A1B1)	2.56 (A1B1)
	8.43 (A2B2)	2.63 (A2B2)

^a Trend not modeled. ^b Autocorrelation not modeled.

investigators may use ITSSIM to analyze a single case, and standard meta-analysis methods may be used to combine multiple ITSSIM effect sizes. In the present study, single-case data from Lambert et al.'s (2006) study were reanalyzed and ITSSIM meta-analysis results were compared to results from five other multilevel methods.

ITSSIM and Comparison Methods

Several of the comparison multilevel methods did not account for trend and/or autocorrelation, and this is expected to account for some of the discrepancy with ITSSIM effect sizes. This is true in particular for the A2B2 phase contrast, where ITSSIM yielded the largest effects and the comparison methods gave relatively smaller ones. The Lambert et al. (2006) design included an ABAB reversal. An increasing trend is visually apparent in many of the A2 phases as the participants return to the higher baseline level of disruptive behavior. The return to baseline in the A2 reversal phase often produces an inverted “V” shaped time-series in the A2B2 phase contrast, as the participant changes from an increasing to a decreasing pattern of response—indeed, this inverted “V” is visually apparent in many of the Lambert et al. (2006) participants. Analyses that do not account for this trend will tend to underestimate effect sizes, because, unlike ITSSIM, their models do not assume that the participants' responses would have continued to increase in the absence of treatment. Ma (2006), Allison and Gorman (1993), and Schlosser et al. (2008) explored this very issue as a limitation of some single-case effect size statistics that do not model trend. Similarly, the failure to model autocorrelation can attenuate some single-case effect size statistics (Tarlow, 2017); however, the effect of autocorrelation on multilevel methods is not well understood (Ugille et al., 2012).

Another possible reason ITSSIM yielded slightly larger effect size estimates is the presence of missing data points in the Lambert et al. (2006) data sets. The current version of ITSSIM software (1.0) does not recognize missing data, so observed changes are interpreted as occurring over a shorter period of time because missing “gaps” in data are ignored; this may lead to steeper trend estimates, and therefore larger effect size estimates. Additional field testing of ITSSIM may reveal whether this overestimation is consistent in other data sets, and whether the missing data limitation is the reason.

Conclusions

The results of this small field test suggest ITSSIM merits further use and study. Despite small but notable discrepancies, the simulation method performed similarly to the comparison multilevel modeling methods. ITSSIM software is also relatively easy to use and interpret and may be a welcome tool for investigators who seek a comprehensive tool for analysis of single-case data. ITSSIM may also be a useful tool for visual raters because—in addition to providing an overall effect size estimate—it isolates and outputs level, trend, variance, and autocorrelation estimates, which are noted in the visual analysis process (Kratochwill et al., 2010).

Three strategies are recommended for future evaluation of ITSSIM. First, future studies should explore how well ITSSIM agrees with visual analysis. Second, additional data sets should be analyzed with ITSSIM, and results compared with other methods (similar to the field test in this study). ITSSIM's agreement with visual raters and other established statistical methods is ultimately the standard by which its viability should be evaluated. Third, simulation methods have been used to empirically evaluate the limitations other single-case statistics (e.g., Manolov & Solanas, 2008; Tarlow, 2017). This approach could be applied to ITSSIM (i.e., “simulations of simulations”), though it is expected to be computationally challenging.

While describing “the shape of things to come” in single-case research, Shadish (2014b) suggested the ideal single-case statistical analysis would model trend and autocorrelation, and would yield standardized effect sizes suitable for meta-analysis. Shadish also stated that a good method would be accessible via easy-to-use software: “The importance of the last criterion [accessibility] cannot be overstated. Many clinical scientists will understandably use simple statistical programs even if they are not state of the art” (p. 144). ITSSIM was developed in accordance with that request, and while ITSSIM software is simple to use, its design was intended to place comprehensive computer-intensive simulation methods into the hands of any single-case investigator.

References

- Allison, D. B., & Gorman, B. S. (1993). Calculating effect sizes for meta-analysis: The case of the single case. *Behaviour Research and Therapy*, 31, 621–631. [http://dx.doi.org/10.1016/0005-7967\(93\)90115-B](http://dx.doi.org/10.1016/0005-7967(93)90115-B)

- Anderson, R. L. (1942). Distribution of the serial correlation coefficient. *Annals of Mathematical Statistics, 13*, 1–13. <http://dx.doi.org/10.1214/aoms/1177731638>
- APA Presidential Task Force on Evidence-Based Practice. (2006). Evidence-based practice in psychology. *American Psychologist, 61*, 271–285. <http://dx.doi.org/10.1037/0003-066X.61.4.271>
- Baer, D. M. (1988). An autocorrelated commentary on the need for a different debate. *Behavioral Assessment, 10*, 295–297.
- Barlow, D. H., & Hersen, M. (1984). *Single-case experimental designs: Strategies for studying behavior change* (2nd ed.). Elmsford, NY: Pergamon Press.
- Barlow, D. H., & Nock, M. K. (2009). Why can't we be more idiographic in our research? *Perspectives on Psychological Science, 4*, 19–21. <http://dx.doi.org/10.1111/j.1745-6924.2009.01088.x>
- Bartlett, M. S. (1946). On the theoretical specification and sampling properties of autocorrelated time-series. *Journal of the Royal Statistical Society, 8*, 27–41.
- Beretvas, S. N., & Chung, H. (2008). A review of meta-analyses of single-subject experimental designs: Methodological issues and practice. *Evidence-Based Communication Assessment and Intervention, 2*, 129–141. <http://dx.doi.org/10.1080/17489530802446302>
- Borenstein, M., Hedges, L. V., Higgins, J. P. T., & Rothstein, H. R. (2009). *Introduction to meta-analysis*. Hoboken, NJ: Wiley. <http://dx.doi.org/10.1002/9780470743386>
- Box, G. E. P., & Jenkins, G. M. (1970). *Time series analysis: Forecasting and control*. San Francisco, CA: Holden-Day.
- Brossart, D. F., Parker, R. I., & Castillo, L. G. (2011). Robust regression for single-case data analysis: How can it help? *Behavior Research Methods, 43*, 710–719. <http://dx.doi.org/10.3758/s13428-011-0079-7>
- Brossart, D. F., Parker, R. I., Olson, E. A., & Mahadevan, L. (2006). The relationship between visual analysis and five statistical analyses in a simple AB single-case research design. *Behavior Modification, 30*, 531–563. <http://dx.doi.org/10.1177/0145445503261167>
- Busk, P. L., & Serlin, R. C. (1992). Meta-analysis for single-case research. In T. R. Kratochwill & J. R. Levin (Eds.), *Single-case research design and analysis: New directions for psychology and education* (pp. 187–212). Hillsdale, NJ: Erlbaum.
- Campbell, J. M., & Herzinger, C. V. (2010). Statistics and single-subject research methodology. In D. Gast (Ed.), *Single-subject research methodology in behavioral science* (pp. 417–453). New York, NY: Routledge.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Mahwah, NJ: Erlbaum.
- Cox, D. R. (1966). The null distribution of the first serial correlation coefficient. *Biometrika, 53*(3/4), 623–626. <http://dx.doi.org/10.2307/2333676>
- Danov, S. E., & Symons, F. J. (2008). A survey evaluation of the reliability of visual inspection and functional analysis graphs. *Behavior Modification, 32*, 828–839. <http://dx.doi.org/10.1177/0145445508318606>
- DeProspero, A., & Cohen, S. (1979). Inconsistent visual analyses of intrasubject data. *Journal of Applied Behavior Analysis, 12*, 573–579. <http://dx.doi.org/10.1901/jaba.1979.12-573>
- Ferron, J. (2002). Reconsidering the use of the general linear model with single-case data. *Behavior Research Methods, Instruments, & Computers, 34*, 324–331. <http://dx.doi.org/10.3758/BF03195459>
- Glass, G. V., Willson, V. L., & Gottman, J. M. (1975). *Design and analysis of time-series experiments*. Boulder, CO: Colorado Associated University Press.
- Gorsuch, R. L. (1983). Three methods for analyzing limited time-series (N of 1) data. *Behavioral Assessment, 7*, 141–154.
- Harbst, K. B., Ottenbacher, K. J., & Harris, S. R. (1991). Interrater reliability of therapists' judgements of graphed data. *Physical Therapy, 71*, 107–115. <http://dx.doi.org/10.1093/ptj/71.2.107>
- Hartmann, D. P., Gottman, J. M., Jones, R. R., Gardner, W., Kazdin, A. E., & Vaught, R. S. (1980). Interrupted time-series analysis and its application to behavioral data. *Journal of Applied Behavior Analysis, 13*, 543–559. <http://dx.doi.org/10.1901/jaba.1980.13-543>
- Hayes, S. C. (1981). Single case experimental design and empirical clinical practice. *Journal of Consulting and Clinical Psychology, 49*, 193–211. <http://dx.doi.org/10.1037/0022-006X.49.2.193>
- Huitema, B. E., & McKean, J. W. (1998). Irrelevant autocorrelation in least-squares intervention models. *Psychological Methods, 3*, 104–116. <http://dx.doi.org/10.1037/1082-989X.3.1.104>
- Huitema, B. E., & McKean, J. W. (2000a). A simple and powerful test for autocorrelated errors in OLS intervention models. *Psychological Reports, 87*, 3–20. <http://dx.doi.org/10.2466/pr0.2000.87.1.3>
- Huitema, B. E., & McKean, J. W. (2000b). Design specification issues in time-series intervention models. *Educational and Psychological Measurement, 60*, 38–58. <http://dx.doi.org/10.1177/00131640021970358>
- Jenson, W. R., Clark, E., Kircher, J. C., & Kristjansson, S. D. (2007). Statistical reform: Evidence-based practice, meta-analyses, and single subject designs. *Psychology in the Schools, 44*, 483–493. <http://dx.doi.org/10.1002/pits.20240>
- Jones, R. R., Vaught, R. S., & Weinrott, M. (1977). Time-series analysis in operant research. *Journal of Applied Behavior Analysis, 10*, 151–166. <http://dx.doi.org/10.1901/jaba.1977.10-151>
- Jones, R. R., Weinrott, M. R., & Vaught, R. S. (1978). Effects of serial dependency on the agreement between visual and statistical inference. *Journal of Applied Behavior Analysis, 11*, 277–283. <http://dx.doi.org/10.1901/jaba.1978.11-277>
- Kazdin, A. E. (1982). *Single-case research designs: Methods for clinical and applied settings*. New York, NY: Oxford University Press.
- Kratochwill, T. R., Hitchcock, J., Horner, R. H., Levin, J. R., Odom, S. L., Rindskopf, D. M., & Shadish, W. R. (2010). *Single-case design technical documentation*. Washington, DC: What Works Clearinghouse.
- Lambert, M. C., Cartledge, G., Heward, W. L., & Lo, Y. (2006). Effects of response cards on disruptive behavior and academic responding during math lessons by fourth-grade urban students. *Journal of Positive Behavior Interventions, 8*, 88–99.
- Lieberman, R. G., Yoder, P. J., Reichow, B., & Wolery, M. (2010). Visual analysis of multiple baseline across participants graphs when change is delayed. *School Psychology Quarterly, 25*, 28–44. <http://dx.doi.org/10.1037/a0018600>
- Ma, H. H. (2006). An alternative method for quantitative synthesis of single-subject researches: Percentage of data points exceeding the median. *Behavior Modification, 30*, 598–617. <http://dx.doi.org/10.1177/0145445504272974>
- Maggin, D. M., & Odom, S. L. (2014). Evaluating single-case research data for systematic review: A commentary for the special issue. *Journal of School Psychology, 52*, 237–241. <http://dx.doi.org/10.1016/j.jsp.2014.01.002>
- Mann, H. B. (1945). Nonparametric tests against trend. *Econometrica, 13*, 245–259. <http://dx.doi.org/10.2307/1907187>
- Manolov, R. (2017). Linear trend in single-case visual and quantitative analyses. *Behavior Modification*. Advance online publication. <http://dx.doi.org/10.1177/0145445517726301>
- Manolov, R., & Solanas, A. (2008). Comparing $N = 1$ effect size indices in presence of autocorrelation. *Behavior Modification, 32*, 860–875. <http://dx.doi.org/10.1177/0145445508318866>
- Manolov, R., & Solanas, A. (2009). Percentage of nonoverlapping corrected data. *Behavior Research Methods, 41*, 1262–1271. <http://dx.doi.org/10.3758/BRM.41.4.1262>
- Manolov, R., & Solanas, A. (2013). A comparison of mean phase difference and generalized least squares for analyzing single-case data. *Journal of School Psychology, 51*, 201–215. <http://dx.doi.org/10.1016/j.jsp.2012.12.005>

- Matyas, T. A., & Greenwood, K. M. (1990). Visual analysis of single-case time series: Effects of variability, serial dependence, and magnitude of intervention effects. *Journal of Applied Behavior Analysis*, 23, 341–351. <http://dx.doi.org/10.1901/jaba.1990.23-341>
- Moeyaert, M., Ferron, J. M., Beretvas, S. N., & Van den Noortgate, W. (2014). From a single-level analysis to a multilevel analysis of single-case experimental designs. *Journal of School Psychology*, 52, 191–211. <http://dx.doi.org/10.1016/j.jsp.2013.11.003>
- Morgan, D. L., & Morgan, R. K. (2001). Single-participant research design. Bringing science to managed care. *American Psychologist*, 56, 119–127. <http://dx.doi.org/10.1037/0003-066X.56.2.119>
- Ninci, J., Vannest, K. J., Willson, V., & Zhang, N. (2015). Interrater agreement between visual analysts of single-case data: A meta-analysis. *Behavior Modification*, 39, 510–541. <http://dx.doi.org/10.1177/0145445515581327>
- Park, H., Marascuilo, L., & Gaylord-Ross, R. (1990). Visual inspection and statistical analysis of single-case designs. *Journal of Experimental Education*, 58, 311–320. <http://dx.doi.org/10.1080/00220973.1990.10806545>
- Parker, R. I., & Brossart, D. F. (2003). Evaluating single-case research data: A comparison of seven statistical methods. *Behavior Therapy*, 34, 189–211. [http://dx.doi.org/10.1016/S0005-7894\(03\)80013-8](http://dx.doi.org/10.1016/S0005-7894(03)80013-8)
- Parker, R. I., Cryer, J., & Byrns, G. (2006). Controlling baseline trend in single-case research. *School Psychology Quarterly*, 21, 418–444. <http://dx.doi.org/10.1037/h0084131>
- Parker, R. I., & Vannest, K. J. (2012). Bottom-up analysis of single-case research designs. *Journal of Behavioral Education*, 21, 254–265. <http://dx.doi.org/10.1007/s10864-012-9153-1>
- Parker, R. I., Vannest, K. J., & Davis, J. L. (2011). Effect size in single-case research: A review of nine nonoverlap techniques. *Behavior Modification*, 35, 303–322. <http://dx.doi.org/10.1177/0145445511399147>
- Parker, R. I., Vannest, K. J., Davis, J. L., & Sauber, S. B. (2011). Combining nonoverlap and trend for single-case research: Tau-U. *Behavior Therapy*, 42, 284–299. <http://dx.doi.org/10.1016/j.beth.2010.08.006>
- Parsonson, B. S., & Baer, D. M. (1992). The visual analysis of data, and current research into the stimuli controlling it. In T. R. Kratochwill & J. R. Levin (Eds.), *Single-case research design and analysis* (pp. 15–40). Hillsdale, NJ: Erlbaum.
- Rindskopf, D. (2014). Nonlinear Bayesian analysis for single case designs. *Journal of School Psychology*, 52, 179–189. <http://dx.doi.org/10.1016/j.jsp.2013.12.003>
- Schlosser, R. W., Lee, D. L., & Wendt, O. (2008). Application of the percentage of non-overlapping data (PND) in systematic reviews and meta-analyses: A systematic review of reporting characteristics. *Evidence-Based Communication Assessment and Intervention*, 2, 163–187. <http://dx.doi.org/10.1080/17489530802505412>
- Scruggs, T. E., & Mastropieri, M. A. (2001). How to summarize single-participant research: Ideas and applications. *Exceptionality: A Special Education Journal*, 9, 227–244.
- Sen, P. K. (1968). Estimates of the regression coefficient based on Kendall's Tau. *Journal of the American Statistical Association*, 63, 1379–1389. <http://dx.doi.org/10.1080/01621459.1968.10480934>
- Shadish, W. R. (2014a). Analysis and meta-analysis of single-case designs: An introduction. *Journal of School Psychology*, 52, 109–122. <http://dx.doi.org/10.1016/j.jsp.2013.11.009>
- Shadish, W. R. (2014b). Statistical analyses of single-case designs: The shape of things to come. *Current Directions in Psychological Science*, 23, 139–146. <http://dx.doi.org/10.1177/0963721414524773>
- Shadish, W. R., Hedges, L. V., Pustejovsky, J. E., Boyajian, J. G., Sullivan, K. J., Andrade, A., & Barrientos, J. L. (2014). A *d*-statistic for single-case designs that is equivalent to the usual between-groups *d*-statistic. *Neuropsychological Rehabilitation*, 24(3–4), 528–553. <http://dx.doi.org/10.1080/09602011.2013.819021>
- Shadish, W. R., Rindskopf, D. M., & Hedges, L. V. (2008). The state of science in the meta-analysis of single-case experimental designs. *Evidence-Based Communication Assessment and Intervention*, 2, 188–196. <http://dx.doi.org/10.1080/17489530802581603>
- Shadish, W. R., Zuur, A. F., & Sullivan, K. J. (2014). Using generalized additive (mixed) models to analyze single case designs. *Journal of School Psychology*, 52, 149–178. <http://dx.doi.org/10.1016/j.jsp.2013.11.004>
- Smith, J. D. (2012). Single-case experimental designs: A systematic review of published research and current standards. *Psychological Methods*, 17, 510–550.
- Swaminathan, H., Rogers, H. J., & Horner, R. H. (2014). An effect size measure and Bayesian analysis of single-case designs. *Journal of School Psychology*, 52, 213–230. <http://dx.doi.org/10.1016/j.jsp.2013.12.002>
- Tarlow, K. R. (2017). An improved rank correlation effect size statistic for single-case designs: Baseline Corrected Tau. *Behavior Modification*, 41, 427–467. <http://dx.doi.org/10.1177/0145445516676750>
- Tarlow, K. R. (2018a). *ITSSIM: Interrupted time-series simulation, version 1.0*. College Station, TX: Author. Retrieved from <http://ktarlow.com/stats>
- Tarlow, K. R. (2018b, January). *R functions for meta-analysis (R code)*. Retrieved from <http://ktarlow.com/stats>
- Theil, H. (1950). A rank-invariant method of linear and polynomial regression analysis, I, II, and III. *Proceedings of the Royal Netherlands Academy of Sciences*, 53, 386–392, 521–525, and 1397–1412.
- Thompson, B. (2006). *Foundations of behavioral statistics: An insight-based approach*. New York, NY: Guilford Press.
- Ugille, M., Moeyaert, M., Beretvas, S. N., Ferron, J., & Van den Noortgate, W. (2012). Multilevel meta-analysis of single-subject experimental designs: A simulation study. *Behavior Research Methods*, 44, 1244–1254. <http://dx.doi.org/10.3758/s13428-012-0213-1>
- Van den Noortgate, W., & Onghena, P. (2003). Combining single-case experimental data using hierarchical linear models. *School Psychology Quarterly*, 18, 325–346. <http://dx.doi.org/10.1521/scpq.18.3.325.22577>
- Van den Noortgate, W., & Onghena, P. (2008). A multilevel meta-analysis of single-subject experimental design studies. *Evidence-Based Communication Assessment and Intervention*, 2, 142–151. <http://dx.doi.org/10.1080/17489530802505362>
- Vannest, K. J., Parker, R. I., Davis, J. L., Soares, D. A., & Smith, S. L. (2012). The Theil-Sen slope for high-stakes decisions from progress monitoring. *Behavioral Disorders*, 37, 271–280. <http://dx.doi.org/10.1177/019874291203700406>
- Wampold, B. E. (1988). Introduction. *Behavioral Assessment*, 10, 227–228.
- White, O. R., & Haring, N. G. (1980). *Exceptional teaching: A multimedia training package*. Columbus, OH: Merrill.
- Wilcox, R. R. (1998). A note on the Theil-Sen regression estimator when the regressor is random and the error term is heteroscedastic. *Biometrical Journal Biometrische Zeitschrift*, 40, 261–268. [http://dx.doi.org/10.1002/\(SICI\)1521-4036\(199807\)40:3<261::AID-BIMJ261>3.0.CO;2-V](http://dx.doi.org/10.1002/(SICI)1521-4036(199807)40:3<261::AID-BIMJ261>3.0.CO;2-V)
- Wilcox, R. R. (2001). *Fundamentals of modern statistical methods: Substantially improving power and accuracy*. New York, NY: Springer-Verlag. <http://dx.doi.org/10.1007/978-1-4757-3522-2>
- Ximenes, V. M., Manolov, R., Solanas, A., & Quera, V. (2009). Factors affecting visual inference in single-case designs. *The Spanish Journal of Psychology*, 12, 823–832. <http://dx.doi.org/10.1017/S113874160002195>
- Yue, S., Pilon, P., Phinney, B., & Cavadias, G. (2002). The influence of autocorrelation on the ability to detect trend in hydrological series. *Hydrological Processes*, 16, 1807–1829. <http://dx.doi.org/10.1002/hyp.1095>

Received September 28, 2017

Revision received June 6, 2018

Accepted June 7, 2018 ■