Inference With Cross-Lagged Effects – Problems in Time and New Interpretations

Charles C. Driver
Center for Lifespan Psychology, Max Planck Institute for Human Development
Institute of Education, University of Zurich
Institute for Educational Evaluation, Affiliated Institute at the University of Zurich

The interpretation of cross-effects from vector autoregressive models to infer structure and causality amongst constructs is widespread and sometimes problematic. I first explain how hypothesis testing and regularization are invalidated when processes that are thought to fluctuate continuously in time are, as is typically done, modeled as changing only in discrete steps. I then describe an alternative interpretation of cross-effect parameters that incorporates correlated random changes for a potentially more realistic view of how process are temporally coupled. Using an example based on wellbeing data, I demonstrate how some classical concerns such as sign flipping and counter intuitive effect directions can disappear when using this combined deterministic / stochastic interpretation. Models that treat processes as continuously interacting offer both a resolution to the hypothesis testing problem, and the possibility of the combined stochastic / deterministic interpretation.

Keywords: Continuous time, dynamic model, cross-lagged panel, lagged effects, causal inference, multilevel VAR, dual-change score, network analysis, stochastic differential equation

Predicting the future based on a set of past observations is, in at least some sense, a straightforward problem. We can of course be wrong to varying degrees, but the basic tools of statistics and computation have evolved to the point where basic recipes for the use of past data to predict future data can often give quite reasonable results. In contrast, predicting how the future will be *different* if we intervene in a specific way, based on the same set of past observations, is, in most cases, far more difficult. Therein lies both the promise and the pain of much of statistical modeling. Discussing mere correlation is easy, causal interpretation under the assumption of having fit the 'true' model is easy, but unfortunately most science resides somewhere in the grey zone where causality is of interest and models plus data are only approximations of a more complex reality. In classical experimental settings such as a randomized controlled trial, the statistical model can sometimes be quite simple, and the causal interpretation seemingly quite clear – the prime inferential difficulty in such cases is usually mapping understanding of one or more very specific, controlled contexts, to the broader and messier reality we're usually interested in. The use of observational data approaches typically aims to reduce some of these generalization issues (i.e. data better reflects the circumstances of interest), but from this arises what can be thought of as an inverse generalization issue – that of going from 'what is likely to happen to y when x changes for any reason' to 'what is likely to happen to y if x is *manipulated*'.

Cross-lagged effects in longitudinal models have for long offered a tantalizing vision of causal inference in observational data, but their interpretation has for just as long been problematic and controversial, as interpretations often go far beyond what can reasonably be said given the data. In this paper I discuss two important issues when interpreting cross-lagged coefficients. I first discuss a very common form of model misinterpretation, wherein researchers assume that a model structure involving discrete jumps between time points (e.g. typical structural equation or vector autoregressive models) can be used for sensible inference about processes that do not exhibit such a discrete jump structure. Although not a problem in every situation, in moderate to high dimensional scenarios inferences become invalid, and this does not appear to be well understood in the field. This misspecification is in a sense a 'simple' problem, in that although it may involve some mathematical complexity, the so-

Charles C Driver, Institute of Education, University of Zurich. Thanks to Dr. Manuel Voelkle and Dr. Martin Tomasik for comments and critique. Analysis scripts and generated data are available in the supplementary material. Correspondence concerning this article should be addressed to Charles C Driver, Institute of Education, University of Zurich, Freiestrasse 36, CH-8032 Zurich. E-mail: charles.driver2@uzh.ch

lution is known and approaches to remedy the issue are clear. The far trickier problem is what we can understand from estimated parameters when we drop the usually massively incorrect assumption that the specified model structure accurately reflects reality. To this end, I offer no panacea, but a novel approach to visualizing and thinking about cross-lagged effects that considers both deterministic and stochastic temporal effects in a combined manner, rather than only the deterministic parameters (e.g. cross regressions) alone as is typically done. With an empirical example using long-term data on satisfaction, I demonstrate how some confusions and overly strong claims can disappear when considered from this alternative view.

Cross-lagged models

Most models that address coupled processes over time, such as vector autoregressive, cross-lagged panel, and multivariate change-score models, contain, or can be re-written to contain, (e.g. Voelkle & Oud, 2015) equation components for modeling the processes that look something like this:

$$y_u = Ay_{u-1} + b + Gz_u \quad z \sim N(0, 1) \tag{1}$$

Where y is a d length vector of process values (possibly but not necessarily data), u is an index of measurement occasion, A is a matrix of temporal regression coefficients, b is an intercept, and G is the effect matrix of the d length system noise vector z, where z contains independent and identically distributed deviations with zero mean. Many extensions and variations are in use, non-linear forms may be written more generically as functions, but the issues discussed herein, while broadly applicable, are simplest to understand using the basic form shown. For a recent overview contrasting different approaches to cross-lagged models and software, see Ruissen et al. (2021).

Part 1 - Misspecified Temporal Structure

Hypothesis testing and structure discovery via regularization are both regularly misunderstood and misapplied in the context of cross-lagged panel models, dual change scores, vector auto-regressive models, and similar approaches that implicitly assume processes evolve in discrete jumps in time, when in reality researchers tend to think of the underlying processes as continually in flux. While the need for continuous-time systems for inference with stochastic processes has already been discussed somewhat, (see for example Aalen, 1987; Aalen et al., 2016; Deboeck & Preacher, 2016; Kuiper & Ryan, 2018; Ryan & Hamaker, 2021), discrete-time cross-regressive models are still the dominant paradigm for inference in such settings. To explicitly highlight the problem, in this section I provide some simple examples, in the domains of hypothesis testing, and structural discovery via regularization. The take home message will be that even in relatively sparse causal systems, the temporal regression matrix between observations at different times is very likely to be dense, or at least less sparse – though in some cases can also show similar sparsity to the causal system but in *different locations*. The fundamental problem is that this discrepancy correspondingly invalidates attempts to infer structure or causation via the specification of, or regularization to, zeroes in the matrix of dynamic regression coefficients. This issue arises due to a difference in how the systems are typically expected to behave, which is continuously (or very frequently) interacting, and how they are modeled, which is interacting only when measured. This discrepancy can be resolved by directly estimating a continuous-time temporal effects matrix, and specifying or regularizing this to zero – the corresponding discrete-time regression weights and covariances can then be computed based on the continuous time temporal effects, or 'drift', matrix, and the time interval between observations. Important to note, is that this all applies whether or not time intervals vary between observations or subjects.

When a cross-lagged model is used to represent a hypothesis where it is believed that a) the processes actually evolve in discrete jumps, and b) the measurement occasions u actually capture *each and every one of* these discrete jumps, then there is no problem. To consider such a system, we could imagine 3 people exchanging information by post, and the post is delivered daily. Person x writes to person y, and person y in turn passes this information on to person z. If we know that person x has received new information, we can expect y to become aware of this the following day, while it will take an additional day before z knows it. In this case, we can precisely represent the causal pathways using a temporal regression matrix (or path diagram) that represents change over a time interval of one day or one postal delivery, as per Figure 1. y is entirely dependent on x for information, and z is dependent on y.

$$\begin{array}{c} \\ x \\ y \\ z \end{array} \begin{array}{c} \begin{array}{ccc} x & y & z \end{array} \\ \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \end{array}$$

Using such a representation, our model will perform optimally in terms of prediction, where we can use some knowledge of what information each individual has on a particular day to predict future days. Furthermore, we can also use the structure to understand what will happen if the structure *changes* – if for instance the link between y and z is removed, we know that z will not receive any information from x.

What happens to our representation if we instead shift to an observation time interval of two days or two deliveries though? As soon as we deviate from the 'true' time step over which information flow occurs, our representation no longer matches the data generating process, or causality, inherent to the system. For an interval of two days, the temporal regres-

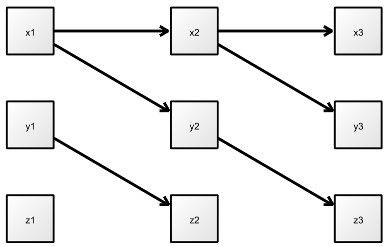

Figure 1

Representation of daily packets of information flow between persons x, y, and z, across 3 days.

sion matrix would look as follows, and the path diagram as per Figure 2. When the new time interval is simply an integer τ multiple of the original, the new temporal regression matrix can be obtained by matrix multiplying the original matrix by itself $\tau - 1$ times.

$$\begin{array}{cc} & \begin{array}{ccc} x & y & z \end{array} \\ \begin{array}{c} x \\ y \\ z \end{array} & \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} \end{array}$$

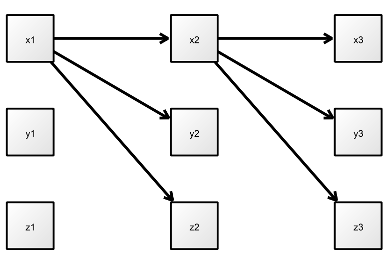

Figure 2

Representation of packets of information flow between persons x, y, and z, across 6 days, when the time interval is two days.

In this case, we have accurately represented how information flows through the system at time intervals of two days, in that once a change in x occurs, 2 days later we see changes in both y and z. So, we can still use the representation to make accurate predictions on a 2 day interval. However, we clearly have *not* represented the actual causality in the system! One might be tempted to infer from this representation that person y is simply unnecessary for information flow between x and z, but in fact person y is required.

This discrepancy between the generating model, and information flow given an arbitrary unit of time, captures a problem inherent to many longitudinal modeling endeavors. Such a discrepancy is precisely why discrete-time cross regression parameters cannot be used as the basis of hypothesis tests regarding causality, or as indicators of some generative structure when achieved via regularization. The causal conclusions that result are not valid, even when the true model is some form of vector autoregression and plenty of data is available – unless every possible occasion of change (whether observed or not) is included in the model.

As we can see from the example, problems related to hypothesis tests and structure determination with discrete time models occur *even when the true data generating function is a discrete time function*, whenever more than one step of change has occurred between observations. Given that most psychological processes tend to be thought of as continuously existing and interacting with each other, rather than interacting only when we decide to measure them, this poses something of a problem, as regular longitudinal structural equation modeling and network modeling approaches can't represent the data generating processes that researchers typically hypothesize. While this is not inherently problematic when models are used only for predictive purposes – in the postal example, we could still make accurate two-day forward predictions using the two day temporal regressions – it *is* a problem whenever the underlying causality (or 'structure') is of interest.

Continuous Time

Instead of modeling processes as a sequence of discrete jumps, differential equations – the mathematics of continuously changing processes – have been available and widely used in the sciences for hundreds of years. While the fields of psychology have recently been showing more awareness of differential equation approaches, the full ramifications with respect to understanding and testing causal structure still appear little understood. While some of the mathematical computations required for software to solve the equations and fit continuous time models can look cruel and unusual, an intuitive understanding can be gained quite easily, and this often resolves a range of time-related confusions that can also occur with more familiar regression oriented approaches (Voelkle et al., 2019, For more on such confusions, see).

In discrete-time systems, the matrix of temporal effects represents regression strengths between two points in time, as in the day or two day examples discussed in the earlier letter writing example. Continuous-time approaches can be intuitively thought of in much the same way, but simply compressing the time interval to a 'very small' value. This allows for usage of a conceptually similar temporal effects matrix, the only difference being that the matrix represents the influence of the current state of the system on the *change* in the system – which is actually analogous to the approach used in change-score structural equation models, where an autoregression of 1.00 is always applied in addition to any estimated self and cross feedback parameters. A continuous-time form of the vector autoregression discussed is:

$$\mathrm{d}y(t) = (Ay(t) + b)\mathrm{d}t + G\mathrm{d}W(t) \tag{2}$$

This looks similar to the discrete-time form, but instead of telling us the new value of y given one step forward in time, it tells us how y is changing *at the moment*. Some mathematical complications due to the nature of stochastic differential equations are present: the dt on the right hand side, which can simply be thought of as a very small step in time, and the d$W(t)$, which represents white noise in continuous time. To compute y at some time point given an earlier value, one needs to solve the system. This is numerically involving and described in detail in Voelkle et al. (2012), however simple approximations using linear extrapolation can be done by hand, and are quite helpful – both for numerical simulation and to follow the basic logic.

$$y_u = y_{u-1} + \Delta t(Ay_{u-1} + b) + Gz_u \quad z \sim N(0, \Delta t)) \tag{3}$$

Essentially then, one computes the deterministic rate of change using A and b and the earlier value of y, then multiplies this by the length of time step Δt desired. The variance of the white noise element z is then also Δt, resulting in a system noise (co)variance of $GG^\top \Delta t$. Shorter steps are (up to a point) more accurate, but require more computations. Many refinements are possible – this simply represents the basic idea, which is called the Euler-Maruyama method in the context of stochastic differential equations. The linear system shown also allows an 'exact' (close, anyway) one-step solution. While this is described in more detail elsewhere (e.g. Driver & Voelkle, 2018a, contains the equations and code for plotting), the basic component involves the matrix exponential, with temporal regression coefficients for particular time intervals given by $e^{A\Delta t}$.

As shown earlier, when using a discrete-time specification to model a system that is assumed to be continuously interacting, zeroes in the matrices do not imply zero causal effect. We can understand this by computing the implied discrete-time coefficients of a continuously interacting system, using $e^{A\Delta t}$. To illustrate this we can consider a similar system to the postal example, with A affecting y which in turn affects z, but where interactions occur continuously. The continuous-time temporal matrix is:

$$\begin{array}{c} \\ x \\ y \\ z \end{array} \begin{array}{c} \begin{array}{ccc} x & y & z \end{array} \\ \begin{bmatrix} -1 & 0 & 0 \\ 1 & -1 & 0 \\ 0 & 1 & -1 \end{bmatrix} \end{array}$$

The negative coefficients on the diagonal tell us that increases in any of the variables exert a downwards pressure on that same variable in the future, while the positive cross-effects show where change in one variable (determined by the column) leads to positive change in another (determined by the row). Figure 3 illustrates the flow of influence through the system, given some change in variable x. Such a plot is interpretable both as an impulse response function (given an impulse of magnitude 1.00), and as the implied discrete-time coefficients. The value of x at time zero (i.e. when the change in x occurs) becomes less predictive of later values of x, so the coefficient reduces. Initially, the time zero value of x is not predictive of either y or z, but as time goes by the change in x exerts influence on y, and then via y on z, so the regression coefficients rise. At some point, once the initial change in x has worn off to a sufficient degree, the coefficient for y reaches a peak and starts to reduce, and similarly for z. Important to note is that any such peaks do not represent a period of 'peak effect' or the like – the cross-effect is constant over time, it is simply that the accumulation of any effect becomes most obvious somewhere between the time of the initial change, and the time when the initial change has mostly dissipated.

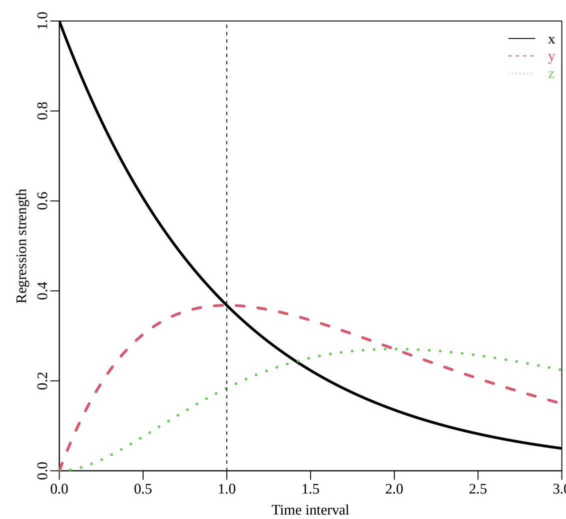

Figure 3

Temporal regression coefficients between x at an earlier point in time and x, y, z. Note that the coefficient between x and z is never zero, even though there is no causal effect. A time interval of 1 unit is marked – the same temporal matrix structure in a discrete-time approach would still show 0.00 for the influence on z at this point.

So, the zeroes in a discrete-time matrix set the coefficients at some particular point in time to zero, but this does not readily map to causality in a continuously interacting system. In contrast, modeling the continuous-time system directly allows one to turn off certain causal paths of the temporal effects A or system noise effects G. This point is easy to miss because there are substantial equivalences between discrete and continuous time models *when temporal effects are freely estimated*, in that with free parameters, equivalent predictions are (in many cases) arrived at. The distinction comes about when restrictions are imposed on the system – zeroes in discrete-time temporal effects do not necessarily translate to zeroes in discrete-time effects given a different sampling rate, or zeroes in continuous-time temporal effects.

For a tangible way to think about the problem, consider a system of motivation to exercise (x), daily exercise levels (y), and fitness (z). A not unreasonable causal hypothesis would look like the structure of Figure 1, where motivation leads to exercise which leads to fitness, or x -> y -> z. Consider then a yearly panel study and a typical discrete-time modeling approach. Naively mapping the underlying causal structure we believe in to the yearly temporal coefficients results in the situation that changes in motivation to exercise take a year to result in changes in exercise, and a further year for this to affect fitness. While Figure 3 is no doubt still inaccurate (because humans are highly complex and it is a made up example using a linear system), it gets much closer to what we would expect – an initial increase in motivation (x, black) starts to increase daily exercise levels (y, red), which in turn start to increase fitness levels (z, green).

To highlight the scope of potentially invalid inferences in these directions, I simulated continuously interacting systems with a range of dimensions and levels of sparsity, observed at time intervals of 1.00. Self-feedback effects (i.e. the diagonals) were fixed to -1, with cross effects randomly set to either 0.5, 0, or -0.5. The correct discrete-time temporal regression matrices were computed for each of these, and then they were checked for inferential consistency. Perfect consistency is found when the pattern of zeroes in the discrete-time matrix exactly matches the pattern of the generative system. To reflect the fact that values near zero are often assumed to be, or regularized to, zero, an arbitrary threshold of 0.05 was used – discrete-time coefficients between -0.05 and 0.05 were assumed zero. All choices in this simulation are necessarily somewhat arbitrary, but the point is found in the general patterns, not the absolute values. The code provided in supplementary material allows for easy changes to these decisions if desired. The top plot of Figure 4 shows the proportion of correctly inferred zeroes (i.e. the correct absence of a causal path) using discrete-time matrices. The lower plot shows the proportion of incorrectly inferred zeroes (i.e. the incorrect presence of a causal path). We see that for three dimensional systems, inferring from freely estimated discrete-time matrices about zeroes in continuously interacting systems appears viable, since the proportion of incorrectly inferred zeroes was 0.00. Inferring the presence of a non-zero path is however, still problematic. Kuiper and Ryan (2018) address some aspects of this in more detail. As system dimensions increase, the proportion of zeroes in the discrete-time matrices is too low, except for systems with either mostly zeroes, or no zeroes. This means that when using discrete-time models a) more parameters are necessary to estimate the same causal system, and b) cross-effect parameters may be needed between variables that have no direct causal path in the generating system. Also as dimensions increase, the proportion of incorrectly inferred zeroes increases, with higher proportions in dense, highly connected systems.

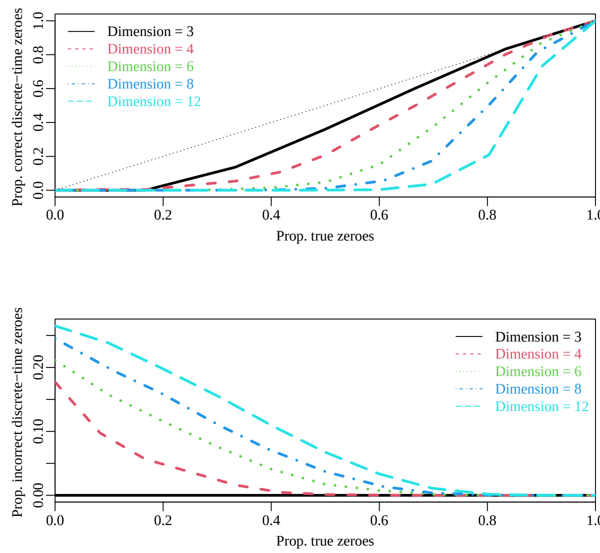

Figure 4

Simulation results, showing how the use of discrete-time models for continuously interacting systems results in invalid conclusions, across different dimensions and sparsity. On top, the expected proportion of correct zeroes in the discrete-time matrix is shown by the thin dotted line. On bottom, the proportion of incorrect zeroes in the discrete-time matrix should always be zero, yet this only holds when dimensions are three or less.

A variety of software is available for modeling continuous-time systems with stochastic elements, as described here. The R software packages ctsem (Driver et al., 2017) and dynr (Ou et al., 2019) are targeted more at social science fields.

Regularization

Having highlighted the problems of relying on discrete-time matrices as a basis for inference when systems evolve continuously (or faster than measured), it is hopefully clear that for confirmatory modeling and hypothesis testing, one needs instead to parameterize the continuous-time temporal effects matrix. From that point, regular tools of hypothesis testing and model comparison can be used as normal. Moving away from confirmatory modeling however, it has become popular within some fields to use regularization to zero (e.g. using the lasso of Tibshirani, 1996) as the basis for inferring causal structure of a system. While such a task is, I believe, difficult and fraught with problems in the psychological domain, it is not necessarily more troublesome than confirmatory modeling, and in both cases we should use the most appropriate tools available.

For the case where structural discovery via regularization is to be attempted, just as with hypothesis testing it is likely that the continuous-time temporal effects are more representative of typical hypothesis and beliefs about psychological systems. The ctsem software (Driver et al., 2017; Driver & Voelkle, 2017) offers the ability to specify Bayesian priors on the system parameters, which, when estimated using the default optimization approach, can also be interpreted as a form of penalized likelihood. The default prior (when enabled) is however a Gaussian distribution, which can be interpreted as L2, or ridge, regression. L2 penalization does not lead to sparse matrices, because coefficients will not, in general, be pushed all the way to zero. A simple fix for this is to replace the Gaussian prior with a Laplace distribution (interpretable as L1 or lasso regression), ensuring that small coefficients will be pushed to zero.

A comparison between L2 and L1 penalization approaches, applied to a simulated 4 dimensional system with a mixture of true parameter values, is shown in Figures 5 and 6, while code for the simulation and plots is in the supplementary material. The system has 4 interacting processes that change continuously, with deterministic effects from the system processes and random effects from uncorrelated noise. The system is stationary and observed without error at intervals of 1 time unit. The plots show how the estimated temporal coefficients and log likelihoods of the system changes as the penalty changes. From the continuous-time model, both the continuous-time and implied discrete-time (ContinuousDiscrete) coefficients are shown, and coefficients from a comparable (i.e. first order) discrete time model are also included for comparison. Note that because the scale of the discrete and continuous coefficients are not necessarily the same, the ideal penalty for out of sample prediction will also not be the same, although the general shape should be – notwithstanding cases of misspecification, such as ignored measurement error.

From Figure 5, we see the expected regularization to zero of some continuous and discrete coefficients before others, when using an L1 penalty approach. When using the L2 approach all coefficients converge at the same point, as the prior scale (inverse of penalty) approaches zero, which is also expected. The main point to notice is that while the discrete-time coefficients implied by the continuous-time model are similar to the discrete-time parameters, they are not the same, and do not necessarily drop to zero when the corresponding continuous parameter does.

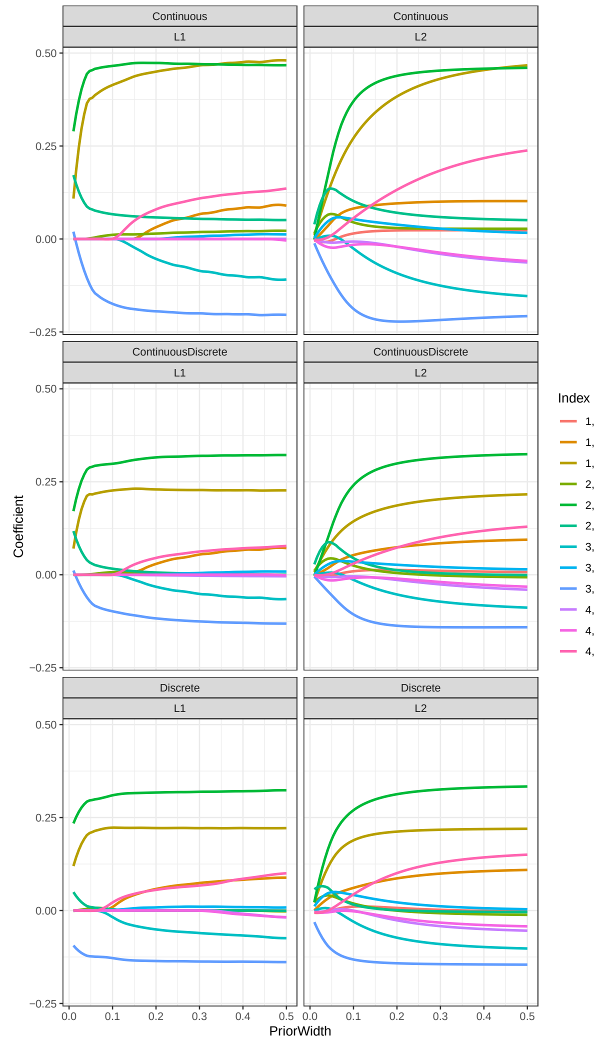

Figure 5

Estimated temporal effect coefficients in a 4 dimensional system, using both L1 (lasso / Laplace prior) and L2 (ridge / Gaussian prior) penalization with different prior scales. Continuous and discrete time models were estimated, and the discrete time coefficients implied by the continuous time model are also shown.

Non-directed relations

Throughout the paper so far I have focused on directed relationships. In many circumstances researchers may instead be interested in covariance and correlation parameters, and the same general pattern of results already described holds. That is, zeroes in a discrete-time system noise covariance

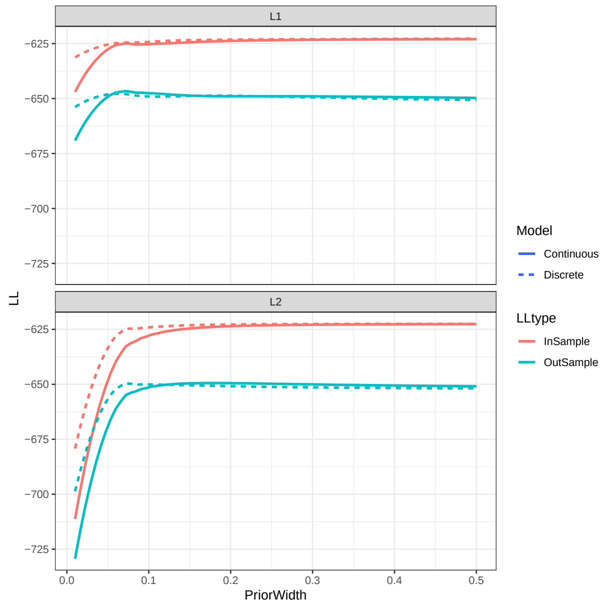

Figure 6

Estimated log likelihoods (in sample as well as out of sample) in the same 4 dimensional system as Figure 5, using both L1 (lasso / Laplace prior) and L2 (ridge / Gaussian prior) penalization with different prior scales. Both continuous and discrete time models were estimated.

matrix can only be interpreted as zeroes in an underlying generative structure of interest, when the real system evolves at (or slower than), the rate at which it has been modeled – which typically corresponds to the frequency of observation.

Part 2 – Interpreting temporal coefficients

Of course, a simple linear model will generally be a vast simplification, whether or not one treats time continuously, or assumes that changes happen in sudden packets. This likely means that any strong claims to causal inference based on the consideration of cross-effect parameters are invalid or at least over-confident, as they are based on the fiction that the true model has been fit. If we take seriously the view that our models are always simplifications however, additional insight into the system dynamics may be gained by considering how the stochastic portion of the model *combines* with the deterministic portion, and may lend at least some additional insight into questions of causality and temporal precedence. The influence of this stochastic portion is represented by the G in the system (latent process) equation, and is, importantly, **not** synonymous with 'the residual', as there may be a measurement layer between the latent process and observations. The stochastic elements should instead be understood as the genuine change in the process that could not be accounted for by the deterministic model elements. Both discrete (Eq. (4)) and continuous (Eq.)(5)) time forms of the system equation are repeated here for convenience:

$$y_u = Ay_{u-1} + b + Gz_u \quad z \sim N(0,1) \tag{4}$$

$$\mathrm{d}y(t) = (Ay(t) + b)\mathrm{d}t + G\mathrm{d}W(t) \tag{5}$$

Typically, when researchers attempt to understand system dynamics by means of cross-effect parameters, the stochastic portion of the system process is either completely ignored, or considered as a separate entity. When, as in a typical simulation study, the true model is used to fit some data, this is at least partially justifiable as the parameters can be interpreted relatively independently. In such a case, the deterministic, cross-effect parameters define how the components of the system influence each other, and the stochastic elements defines how the system components tend to be affected by external, unknown sources. In this convenient but very fictional simulation-land, we can do thought experiments such as 'What happens to the system if we change one component independently of the others?' and obtain the answers very easily, as they are determined by the model design. Such 'answers' can, for example, take the shape of plots like that of Figure 3, which shows how three elements of the system respond to an upwards shock to one element – in simulated data we can be certain that this deterministic response does not vary under different types of shock. Here, shock should be understood as a sudden impulse at a single moment, occurring at time zero in the plots.

To the extent that we believe a model fit to real-world data represents the true data generating process, we can of course obtain answers to such thought experiments in just the same way. However, consider a construct like 'depression', and suppose we measure it every day for one subject. Once we have some data, we can fit a basic autoregressive style model, and obtain parameters telling us approximately how we can expect the system to behave from day to day. Our temporal coefficient and system noise variance could be perfectly precise (imagining infinite data), but they will only ever represent some kind of average *conditional on the various external shocks the subject underwent*. We don't have any specific info on the external shocks the subject underwent – these could be for instance a bad day at work, an illness, or the loss of a loved one – so the best we can do is characterize them via some distribution, but it's important to remember that underlying the distribution is some set of potentially quite different events with different time courses (see Driver & Voelkle, 2018b, for approaches to explicitly incorporate known events with differing time courses). If we also measure a second, related process, such as 'anxiety', we will no doubt find some unique changes, but we are likely to find the disturbance patterns of the two processes have some similarity, i.e., the system noise is correlated. Where does such similarity arise from? Some of the events that change depression may also tend to change anxiety, some of the events that change depression may tend to happen at the same time as events that cause anxiety, or both. Now, imagine that the events which trigger both anxiety and depression tend to have shorter term effects than the 'average' depression triggering event, or in other words, the depression component of the system returns to baseline faster after such a disturbance. This will likely show up in our regular VAR type model as a negative cross-effect of anxiety on depression, but is most definitely *not* because inducing anxiety reduces depression!

While such logic may be unusual thinking, it can resolve some common concerns with such models across a range of fields. Where previously a change in cross-effect sign between alternative models or study designs was a major cause for concern and uncertainty, in many of these cases it may not indicate that there is a truly different direction of effect, but rather just that there is a difference in the estimated speed of changes. The easiest way to understand this combined interpretation of stochastic and deterministic effects is to use a similar visualization as used for the deterministic effects only, but using correlated rather than independent shocks. That is, instead of using some hypothetical intervention that modifies one component alone to plot the expected time course of changes, we can use the estimated correlation in the system noise to generate appropriate shocks to each component of the system, in line with what we would expect to observe empirically. If a shock of magnitude 1.00 is observed in one component, we can expect the other components to also exhibit a non-zero shock (with expectation equal to the correlation), and can then plot the predicted trajectories conditional on this expectation. Such a plot provides answers to the question "Conditional on *observing* a change in one component, what do we expect to happen to other components of the system?". The distinction between the two types of plots is essentially what Gische et al. (2021) terms "forecasting causal effects of interventions versus predicting future outcomes", though here it is handled in continuous-time as opposed to a regression-based framework. Moreover, I will argue that the magnitude of discrepancy between the two forms should probably influence how likely we are to accept the causal interpretation.

To try and unpack all this a little more, in the following I present a brief empirical example of the dynamic relations between overall satisfaction, health satisfaction, and job satisfaction, over long time scales.

Overall, job, and health satisfaction – correlated shock interpretation

Using data from version 33.1 of the German socioeconomic panel (GSOEP) (Goebel et al., 2019), I examine the relations between overall life satisfaction, work satisfaction, and health satisfaction. These constructs are all measured approximately yearly, via single item self-report questions. 2000 participants with 20 or more years of complete data (maximum 33) were randomly selected from the dataset, for the purposes of this demonstration.

The ctsem (Driver & Voelkle, 2021) software for R is used to specify and fit a hierarchical continuous-time state-space model. This has a similar fundamental structure to a simple first-order VAR model, but is based on continuous rather than discrete time, and includes random-subject effects and a basic Gaussian measurement model. Allowing for measurement error ensures that imperfect measurements do not generate spurious cross-effects (Schuurman & Hamaker, 2019). To ensure that the system parameters representing dynamics actually reflect dynamics and not stable individual differences, starting points and intercept terms are allowed to (co)vary across subjects using a random-effects approach. For simplicity of exposition, other system parameters are assumed constant over subjects, though this is not necessary (For the hierarchical Bayes formulation see Driver & Voelkle, 2018a, though note that use of non-linear filtering in modern versions of ctsem also allows for regular maximum likelihood estimation with complex heterogeneity). The R script used for fitting is provided in the supplementary material, as is randomly generated data from the fitted model, since it is not possible to provide the GSOEP data.

After fitting the model using maximum likelihood estimation (the ctsem default), using typical approaches to interpretation one may be tempted to look directly at the temporal coefficients matrix. In this case, examination of the temporal effects parameters leads to similar conclusions regardless of whether ctsem is used to fit a continuous or a discrete time model. Essentially, both life and job satisfaction show negative effects on the other two constructs (though life on work is non significant at the .05 level), while health satisfaction shows positive effects. Such a pattern of results is on face-value somewhat confusing – for example, why should increases in life satisfaction predict future *decreases* in health satisfaction? Moreover, the results conflict with the pattern of between-person effects, which show the expected strong positive correlations between baseline levels of the three satisfaction constructs. Such non-concordance of between and within effect directions is certainly possible, but in the absence of plausible theory, should probably lead to further interrogation of the model! When we consider not only the deterministic changes, but the stochastic changes of the system as well, the picture starts to make more sense – sources of random change in the three constructs are highly positively correlated, such that if we observe a positive change in one of the constructs, years later we *still* expect the other constructs to show positive changes, *even after accounting for the negative temporal coefficients*. So, perhaps rather than generating change in the reverse direction, we could rather interpret the negative coefficients as *faster dissipation*. One can understand this by considering life satisfaction, which is a rather high level construct and likely to be affected by events with very different time courses. Some events might have a large impact but dissipate rapidly, while others could have a smaller but persistent impact. In this case, it seems that changes to life satisfaction that are more work related may tend to dissipate faster than changes which are more health related. These sorts of patterns are however much clearer with visualizations as in Figure 7.

Figure 7 is created using functions from ctsem that compute the impulse response over time of the system[1]. In the regular case, the impulse is assumed independent and only affects one construct at one moment in time, raising it by 1.00. So for example, in the top left of Figure 7 we see that health satisfaction experienced a hypothetical shock in which the other constructs are (initially) unaffected, and health satisfaction is raised by 1.00. As time progresses, the change to health satisfaction dissipates towards zero, but in the meantime the rise in health satisfaction leads to future rises in life and work satisfaction. Considering the proposed interpretation which includes the stochastic component, we can look to the lower left plot and see that whenever a change of 1.00 in health satisfaction is observed, we should also expect to see substantial gains in life and work satisfaction – instead of starting with a change of 0.00 these now start with a change above 0.50. In the context of such an initial change, we no longer see a rise in life or work satisfaction at later time points. Instead, the cross-effects from health satisfaction cause a slower dissipation of the initial change than would be expected if we had considered only the auto-effects of life or work satisfaction.

A similar story occurs if we look at the plots on the right of Figure 7. In this case, if we interpret the parameters as though we could directly intervene and adjust work satisfaction only, the top-right plot will lead the conclusion that an increase in work satisfaction predicts a decrease in both health and life satisfaction. However, after accounting for the fact that the random changes shown by work satisfaction are typically positively correlated with health and life satisfaction, we instead see that when something causes work satisfaction to increase, health and life satisfaction also increase, and the change dissipates more quickly than is otherwise typical for health and life satisfaction processes.

Which interpretation is correct?

Both, and likely neither! The interpretation focusing only on deterministic effects has been popular because it is relatively simple and offers a hopeful vision for inferring causality from longitudinal observational data. It really would be great if from such results we could be confident that decreasing work satisfaction was the key to a happy life! Unfortunately people are right to be skeptical of such simplistic interpretations. In the satisfaction example, the temporal coefficients are estimated in the context of highly correlated random changes in the different processes. Although the linear models in regular use assume it, there is no obvious reason for us to be confident that the same temporal coefficients apply in all circumstances. When for example both work and life satisfaction have been increased by some event, it may be that the impact on life satisfaction declines quicker than for typical events which impact life satisfaction. That by no means implies that when the (apparently) rare event occurs that only directly impacts work satisfaction, that life satisfaction is pushed down – indeed this seems on the face of it somewhat implausible. If we *did* have sophisticated control over the different constructs, one could imagine a scenario in which the random changes to each are induced so they are uncorrelated, and then we collect data and fit a model. This would bring our observed system much closer to answering the question of what happens if we directly intervene on one component only.

Discrete-time interpretation of correlated shocks

In the empirical example shown, a continuous-time approach made more sense because the satisfaction constructs are assumed to function and interact continuously (granularity of space-time concerns notwithstanding). For systems that **do** change and interact only in discrete jumps, the same approach can be used, in which the initial shock is determined by the correlation of the system noise. The ctsem software, despite the continuous-time orientation, will also produce such plots and output when a discrete-time model is fit. Importantly however, just as with the deterministic coefficients, one cannot leverage a stochastic effect matrix (i.e. G in Eq. (4)) from a discrete-time model to apply this combined stochastic / deterministic interpretation when the underlying process operates continuously. This is because the stochastic effect matrix in discrete-time represents the covariance **after** 1 step of time, rather than at time zero.

Conclusion

Advances in data collection technologies have led to an explosion in longitudinal studies and interest in modeling approaches. Cross-lagged effect parameters have been of substantial interest and provided substantial controversy for many years. In this work I have described two issues that relate to the interpretation of cross effects, both of which can be resolved and leveraged by applying more appropriate models such as offered by the ctsem (Driver & Voelkle, 2017) and dynr (Ou et al., 2019) software packages. The first issue was the known but not widely understood issue

[1]For continuous-time systems, the differential equation is solved for each time-interval of interest, while for discrete-time, the temporal effects matrix is simply raised to the power of the number of time steps taken.

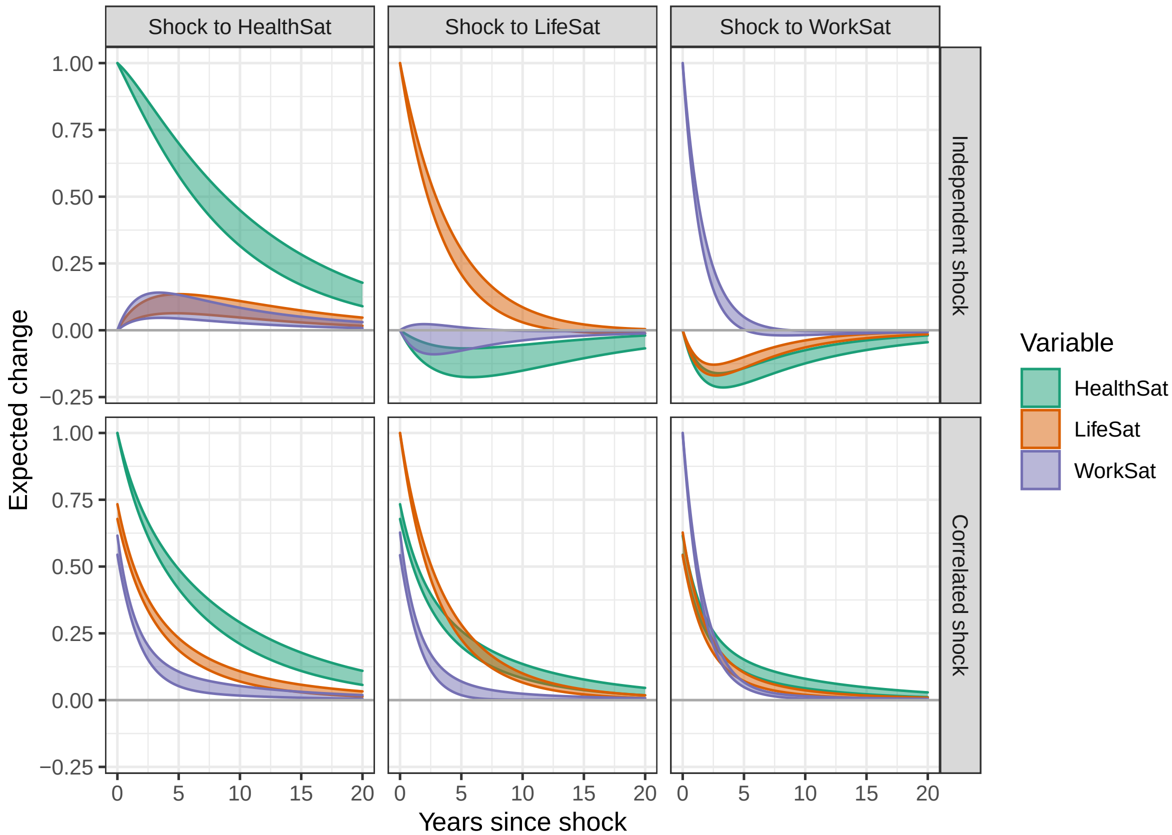

Figure 7

On the top row a standard interpretation of temporal dynamics effects is shown, with a hypothetical independent shock of magnitude 1.00 to each construct. In each column a different construct receives the shock. On the lower row there is an interpretation that focuses on what is typically observed – when there is a shock of 1.00 to a construct, the other constructs receive a shock in accordance with the estimated system noise correlation. In the top, classical interpretation, there are negative effects that are difficult to make sense of, while the combined stochastic and deterministic interpretation suggests these negative effects are more likely just faster dissipation to zero from an initial positively correlated change.

of temporal misspecification, wherein zeroes in the temporal effects matrix of some discrete-time system, with likely quite arbitrary time intervals, are taken to represent no direct causal connection between the variables. This is in most cases an invalid assumption, and is easily rectified by either shifting to a continuous-time perspective or leveraging the discrete-time approach for prediction only. Moving beyond simply highlighting past understanding, I have also described an alternative approach for understanding cross-lagged effects, that also accounts for the correlated random changes occurring in the system. Visualizing the expected trajectories of the system conditional on some initial shock is not uncommon in time-series analysis, but considering this initial shock in correlated rather than independent form is (as far as I am aware) either new or quite uncommon. Taking such an approach can, in at least some cases, resolve seemingly quite confusing problems of unexpected effect directions and concerns over sign-flipping. In addition, this interpretation can also demonstrate that when system noise is correlated, temporal effects do not necessarily imply a particularly interesting change. In the empirical example of satisfaction dynamics, we saw that the apparently positive-effect of a change in health satisfaction on life satisfaction, could also be quite reasonably interpreted as a slower dissipation of some initial correlated change to both. Not uninteresting, but also perhaps not as headline grabbing as change in one variable causing *later* change in another. From at least some perspectives, a little more nuance and less headline grabbing may not be a bad thing when it comes to interpretations of multivariate longitudinal models...

References

Aalen, O. O. (1987). Dynamic modelling and causality. *Scandinavian Actuarial Journal*, *1987*(3-4), 177–190.

Aalen, O. O., Røysland, K., Gran, J. M., Kouyos, R., & Lange, T. (2016). Can we believe the DAGs? A comment on the relationship between causal DAGs and mechanisms. *Statistical methods in medical research*, *25*(5), 2294–2314.

Deboeck, P. R., & Preacher, K. J. (2016). No need to be discrete: A method for continuous time mediation analysis. *Structural Equation Modeling: A Multidisciplinary Journal*, *23*(1), 61–75.

Driver, C. C., Oud, J. H. L., & Voelkle, M. C. (2017). Continuous time structural equation modeling with r package ctsem. *Journal of Statistical Software, Articles*, *77*(5), 1–35. https://doi.org/10.18637/jss.v077.i05

Driver, C. C., & Voelkle, M. C. (2017). Introduction to hierarchical continuous time dynamic modelling with ctsem. *R package Vignette. Available online at: https://cran. r-project. org/web/packages/ctsem/index. html*.

Driver, C. C., & Voelkle, M. C. (2018a). Hierarchical Bayesian continuous time dynamic modeling. *Psychological Methods*, *23*(4), 774.

Driver, C. C., & Voelkle, M. C. (2018b). Understanding the time course of interventions with continuous time dynamic models. In K. van Montfort, J. H. L. Oud, & M. C. Voelkle (Eds.), *Continuous time modeling in the behavioral and related sciences*. Springer International Publishing.

Driver, C. C., & Voelkle, M. C. (2021). Hierarchical continuous time modeling. In J. F. Rauthmann (Ed.), *The Handbook of Personality Dynamics and Processes* (pp. 887–908). Academic Press. https://doi.org/10.1016/B978-0-12-813995-0.00034-0

Gische, C., West, S. G., & Voelkle, M. C. (2021). Forecasting causal effects of interventions versus predicting future outcomes. *Structural Equation Modeling: A Multidisciplinary Journal*, *28*(3), 475–492.

Goebel, J., Grabka, M. M., Liebig, S., Kroh, M., Richter, D., Schröder, C., & Schupp, J. (2019). The German Socio-Economic Panel (SOEP). *Jahrbücher für Nationalökonomie und Statistik*, *239*(2), 345–360. https://doi.org/10.1515/jbnst-2018-0022

Kuiper, R. M., & Ryan, O. (2018). Drawing conclusions from cross-lagged relationships: Re-considering the role of the time-interval. *Structural Equation Modeling: A Multidisciplinary Journal*, *25*(5), 809–823.

Ou, L., Hunter, M. D., & Chow, S.-M. (2019). What's for dynr: A package for linear and nonlinear dynamic modeling in R. *The R journal*, *11*(1), 91.

Ruissen, G. R., Zumbo, B. D., Rhodes, R. E., Puterman, E., & Beauchamp, M. R. (2021). Analysis of dynamic psychological processes to understand and promote physical activity behaviour using intensive longitudinal methods: A primer. *Health Psychology Review*, 1–34.

Ryan, O., & Hamaker, E. L. (2021). Time to Intervene: A Continuous-Time Approach to Network Analysis and Centrality. *Psychometrika*. https://doi.org/10.1007/s11336-021-09767-0

Schuurman, N. K., & Hamaker, E. L. (2019). Measurement error and person-specific reliability in multilevel autoregressive modeling. *Psychological methods*, *24*(1), 70.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, *58*(1), 267–288.

Voelkle, M. C., Gische, C., Driver, C. C., & Lindenberger, U. (2019). The Role of Time in the Quest for Understanding Psychological Mechanisms. *Multivariate Behavioral Research*, 1–24. https://doi.org/10.1080/00273171.2018.1496813

Voelkle, M. C., Oud, J. H. L., Davidov, E., & Schmidt, P. (2012). An SEM approach to continuous time modeling of panel data: Relating authoritarianism and anomia. *Psychological Methods*, *17*(2), 176–192. https://doi.org/10.1037/a0027543

Voelkle, M. C., & Oud, J. H. (2015). Relating latent change score and continuous time models. *Structural Equation Modeling: A Multidisciplinary Journal*, *22*(3), 366–381.