Interpreting parameters of dynamic regression models

Sigert Ariens[1,*], Niels Vanhasbroeck[2], Ginette Lafit[3], and Eva Ceulemans[1]

[1]Quantitative Psychology and Individual Differences, KU Leuven, Leuven, Belgium
[2]Psychological Methods, University of Amsterdam, Amsterdam, the Netherlands
[3]Methodology of Educational Sciences, KU Leuven, Leuven, Belgium
[*]Corresponding author, sigert.ariens@kuleuven.be

Author note

This manuscript has not yet been peer reviewed. A preprint has been made available together with the supplementary materials on the open science framework (https://osf.io/fn7mj/).

Author contributions

Sigert Ariens (Conceptualization, Formal analysis, Software, Funding acquisition, Visualization, Writing—original draft); Niels Vanhasbroeck (Software, Conceptualization, Formal analysis, Visualization, Writing—review and editing); Ginette Lafit (Conceptualization, Formal analysis, Writing—review and editing); Eva Ceulemans (Conceptualization, Formal analysis, Supervision, Visualization, Writing—review and editing)

Conflict of interest: The authors report no conflict of interest.

Funding

This research is supported by a grant from the Fund for Scientific Research - Flanders (FWO; 1278525N), awarded to Sigert Ariens, grants from the Research Council of KU Leuven (STG/23/050, C+/24/002, C16/23/011) awarded to Ginette Lafit and (C14/23/062; iBOF/21/090) awarded to Eva Ceulemans, and by grant EOS 40007528 / G0I2422N from the Research Foundations FWO and F.R.S.-FNRS under the Excellence of Science (EOS) program awarded to E. Ceulemans. Niels Vanhasbroeck is supported by funding from the European Research Council (ERC) under the European Union's Horizon 2020 Excellent Science program (Grant agreement No. 101053880)

Abstract

With the growth of the dynamic paradigm in psychological science, the use of dynamic regression models to answer questions about how specific contexts or experimental conditions influence the psychological process over time has become popular. Although these methods provide unique opportunities to uncover the dynamics which drive psychological processes, the parameters of these dynamic regression models can be difficult to interpret. In this paper, we use the theory of system responses to elucidate how popular dynamic regression models differ in the dynamics they imply. Crucially, this theory allows us to reason about complicated dynamic models qualitatively, without requiring a technical background. We exemplify the opportunities of system response analysis by reanalyzing data from an experience sampling study. Finally, the manuscript introduces the R package impulseR, which provides researchers with a user-friendly toolbox for interpreting dynamic regression models or the results of a dynamic regression analysis.

1 Introduction

The analysis of time series of psychological data has become a popular topic in the recent methodological literature, particularly with the growth of the dynamic paradigm in psychological science. In this paradigm, the study of how psychological variables evolve over time takes center stage. In many domains of inquiry, e.g. affect, decision making (e.g. Lueckmann et al., 2018), learning (Barto et al., 1983), perception, (Ceylan and Pascucci, 2023), attention (Hamker, 2004), psychosomatics (Ariens et al., 2020, Huys et al., 2021), time series of data are gathered and analyzed to arrive at conclusions relevant for our understanding of these constructs.

Researchers often want to investigate how specific contextual variables (X) influence the psychological process (Y) over time. To this end, regression models of various forms are often specified and fit to the time series of data (e.g. Bilacchi et al., 2022). These models contain *parameters*, which are interpreted in a psychologically meaningful way. For example, when fitting a linear regression model:

$$y_t = \alpha + \beta x_t + v_t$$

β is typically interpreted as reflecting the strength of the momentary relationship between x_t and y_t. The parameter β can therefore be interpreted as quantifying how psychologically reactive an individual is towards the covariate time series over the study.

In a recent paper, Ariens et al. (2023) argued that when fitting such regression models to psychological time series, one must carefully consider how the specified model accommodates dynamic dependencies in the data, and whether the model captured these dynamics adequately. For example, specifying a simple linear regression model is a poor strategy if y_t depends on y_{t-1}. The model does not include any term which accommodates such dependencies. The failure of the model to adequately accommodate an effect is broadly termed the problem of *Misspecification*. Misspecification can have severely detrimental consequences for the validity of inferences drawn, as it is a prime source of bias in the estimates of parameters and their standard errors. Such biases can therefore severely distort the results of significance tests, and other inferential procedures based on the parameter estimates.

There is, therefore, good news and bad news. The good news is that how we specify the dynamics in our statistical models *matters*. Since the results of a regression analysis appear to depend strongly on how we specify the dynamics, we can learn about these dynamics through the specification and comparison of regression models. The bad news is that the parameters of these dynamic regression models are notoriously difficult to interpret (e.g. Bulteel, Tuerlinckx, et al., 2018; Eaton et al., 2023). For example, lagged and contemporaneous effects are often interpreted in separate subsections of the manuscript (e.g. Mansueto et al., 2023), despite both contributing to the dynamics. It can be difficult to understand the model-implied behavior based on the values of the individual parameter estimates.

In this manuscript, we show how researchers can interpret the parameters of many commonly specified dynamic regression models through the visualization of the model-implied *impulse response functions*. These functions tell us how a system responds over time once it has been perturbed by an input variable. Psychological scientists typically have some background in experimentalism, where we theorize about how changing some variable influences a dependent variable, keeping other variables constant. Impulse response functions convey this type of information for dynamic regression models: what would happen to y_t *over time* if I change x_t by a unit, controlling for all other model influences? Although straightforward for the linear regression model, this can become difficult when lagged effects are present. Rather than interpreting each parameter (estimate) individually, impulse responses show how the different parameters contribute to the overall shape of the dynamics. Crucially, once they are derived, the impulse responses can be interpreted *visually*, without requiring a technical background. For these reasons, we think that impulse responses could be a powerful tool for psychological researchers interested in studying psychological dynamics. Impulse responses have various useful generalizations which are collectively termed *system responses*. For example, scaled impulse responses can be used to take into account the scale of the input variables when visualizing model-implied dynamics, cumulative responses allow us to visualize how a system behaves when met with an entire time series of input variables, and total responses show us how the system integrates responses to multiple time series of input variables.

To make it feasible for any scientist to calculate and visualize the system responses of their models without requiring a technical background to derive these responses, we introduce the R package impulseR. The package contains user-friendly functions that allow the user to calculate and visualize system responses to aid them in interpreting a model structure or the results of a dynamic regression analysis. The main text provides the necessary conceptual background to understand how to use the materials in practice. A full introduction to the impulseR package, including the package functions, documentation, and tutorial vignettes with example code can be found on Gitlab (https://gitlab.kuleuven.be/u0133721/impulseR).

The manuscript is structured as follows. Section 2 provides a descriptive overview of different types of dynamic regression models that are commonly encountered in psychological applications. We show how these models are related and provide concrete examples of recent applications in psychology. Section 3 provides a conceptual introduction to system responses. We visualize system responses for well known autoregressive models. In Section 4, we apply the theory of system responses to the dynamic regression models we introduced in Section 2. This section offers new, qualitative, and important insights into how popular dynamic regression models differ in their assumptions about the dynamics. Section 5 provides an empirical application, where we fit dynamic regression models to data from an experience sampling (ESM) study. We illustrate how the visualization of system responses can provide much needed assistance when it comes to interpreting the results of a dynamic regression analysis. In the discussion we first reflect on some strengths and limitations of impulse response analyses. We close by discussing extensions which we avoided in the main text to make the materials as accessible as possible.

2 Dynamic regression models

In this section, we introduce those *dynamic regression models* which, in our experience, are most popularly employed in psychological applications. The models we consider encompass many others discussed in the methodological literature (e.g. Dynamic Structural Equation Model [DSEM], Asparouhov et al., 2018; residualized dynamic SEM [RDSEM], Asparouhov and Muthén, 2020). In the following section, we pay special attention to these connections and provide recent examples of works where these models have been applied.

2.1 Notation

We will use t as the discrete time index, $t \in \{0, 1, 2, ...\}$. For technical purposes, we will start the time index at $t = 0$. For example, y_0 is the value of the dependent variable at the first measurement occasion. All parameters and vectors are real valued. Symbols in bold will denote t-dimensional vectors.

Parameters reflecting the effects of the covariate $\boldsymbol{x}$ are represented as β's. The subscript provides further information about which lag the parameter is referring to. For instance, β_x is the contemporaneous effect of the covariate, x_t. β_{Lx} is the parameter for the lagged covariate, x_{t-1}.

2.2 An overview of dynamic regression models

A concise overview of how we define the models is given in the next table. We introduce each in turn, together with a recent application in psychology. Note that the table introduces models without intercepts, and allow only for contemporaneous and lag-1 covariate effects. Both are for didactic purposes, we will discuss the interpretation of the intercept in Section 4.1, and the impulseR package can be used to study the behavior of models with intercepts and arbitrary lag orders. Common to all these models is that they are discrete-time models. Therefore, the measurements are assumed to be interleaved by equal time intervals (Continuous-time models are discussed in, for instance, Driver et al., 2017).

2.2.1 Autoregression (AR)

Autoregressive models are widely used across various scientific domains. They regress the current value of the process, y_t on its p preceding values. In case only a single lag is allowed for, the model is the well-known AR(1) model (see Table 1). The AR(1) model regresses the current y_t on the previous state y_{t-1}. One example where AR(1) models are often used is the field of affect dynamics (Vanhasbroeck et al., 2021). Here, they are often fit to quantify dynamic features of interest. For example, the AR parameter ϕ is often interpreted to reflect *emotional inertia*, (Kuppens and Verduyn, 2017; Koval et al., 2013; Koval et al., 2015; Revol et al., 2025) or the tendency for affect to linger over time.

AR models also include an *innovation term*, v_t. The innovations v_t are interpreted theoretically as latent inputs to the system. At each time point, an innovation *perturbs* the system. In observational settings (such as experience

Name	Abbreviation	Equation				
Autoregression	AR	$y_t =$	ϕy_{t-1}			$+v_t$
Linear Regression	LR	$y_t =$		$+\beta_x x_t$		$+v_t$
Lagged Criterion Variable	LCV	$y_t =$	ϕy_{t-1}	$+\beta_x x_t$		$+v_t$
Common Factor	CF	$y_t =$	ϕy_{t-1}	$+\beta_x x_t$	$-\phi\beta_x x_{t-1}$	$+v_t$
Dead Start	DS	$y_t =$	ϕy_{t-1}		$+\beta_{Lx} x_{t-1}$	$+v_t$
Stimulus Dynamics	SD	$y_t =$		$+\beta_x x_t$	$+\beta_{Lx} x_{t-1}$	$+v_t$
Autoregressive Distributed Lag	ADL	$y_t =$	ϕy_{t-1}	$+\beta_x x_t$	$+\beta_{Lx} x_{t-1}$	$+v_t$

Table 1: An overview of lag-1 dynamic regression models.

sampling settings), v_t is often interpreted to reflect a kind of average of all internal and external forces acting on the system. At some time points, this average will perturb the system positively, while at others the net effect will be negative. In experimental settings, v_t is similarly interpreted as *system noise*, the effects of which feed forward in time due to the system dynamics. Importantly, the temporal effects arise due to the dynamics: the innovations themselves should not contain temporal information. The typical specification $v_t \overset{iid}{\sim} N(0, \sigma_v^2)$, that the inputs are normally distributed around 0, are independent over time, and have constant variance over time, should therefore ideally be interpreted as a set of theoretical assumptions about the system we are trying to model.

The specification $v_t \overset{iid}{\sim} N(0, \sigma_v^2)$ is also "useful" in the sense that it implies that one can use maximum likelihood to yield estimates of the model parameters and their standard errors. For each dynamic regression model, the innovations will be denoted by v_t. When estimating the parameters of dynamic regression models in our empirical application (see also the `irf_empirical()` function in the impulseR package), we will use ordinary least squares and hence assume that the innovations are independent and identically distributed over t, with bounded variance.

2.2.2 Linear Regression (LR)

The LR model is often specified when the goal is to estimate the momentary relationship between a specific covariate and a criterion variable over time (e.g. Strickland et al., 2020; Gilmore et al., 2016; Ferrari et al., 2018). For example, Strickland et al. (2020) investigated performance in a prospective memory task, expecting performance to differ depending on experimental conditions such as the type of stimulus presented to participants. They fit a linear mixed model to the data. For each participant (i.e. at level-1), the model is a multiple linear regression model without lagged effects. A second example comes from Gilmore et al. (2016). These researchers investigated how performance in a dot-comparison task depends on confounding visual cues, such as the convex-hull of the arrays. Their model is a multiple linear regression model regressing accuracy on these non-numerical cues.

The main concern when specifying the LR model on time series variables is that it does not allow for feed-forward behavior of any kind. As such, if there are any temporal dependencies in the data, but a LR model is specified, the parameter estimates are likely to be biased due to misspecification (Ariens et al., 2023). For example it is not inconceivable that information regarding prior non-numerical cues biases current performance.

2.2.3 Lagged criterion variables (LCV)

The LCV model allows for both an effect of a covariate x_t, and an AR effect of the preceding value of the criterion variable, y_{t-1}. The model has also been termed the DSEM model (Asparouhov et al., 2018, McNeish and Hamaker, 2020). These models have been extensively used in affective science. For example, Zhou et al., 2021 suggest a model regressing drinking behavior on drinking behavior at the previous time point, plus a term accommodating the urge to drink at each time point.

2.2.4 Common factor model (CF)

The common factor (CF) model is specified in applied research with great regularity. Although perhaps appearing to be a rather exotic model due to the non-linear parameter restriction, the idea that it "corrects for" serial dependence makes its specification prevalent. The CF model is *always* specified when the residuals of a linear regression model are endowed with an AR(1) structure (Ariens et al., 2023). For instance, whenever researchers specify RDSEM models (Asparouhov and Muthén, 2021) or specify a corAR(1) error structure in the R package lme4 (Bates et al., 2015), they are specifying a CF model. As an example, Blanke et al. (2020) investigated how variability in emotion regulation strategies is related to negative affect at each time point in an experience sampling study. They fit a linear regression model with an AR(1) error structure, predicting negative affect using strategy variability and other covariates of interest.

2.2.5 Dead start model (DS)

The Dead Start (DS) model assumes that a covariate at time t has no immediate effect on the system. Rather, the effect only appears at time $t+1$. It is often specified in applied and methodological works. For example, Takano et al. (2021) fit VAR models with diagonal innovation covariance matrices, implying dead start dynamics. A more applied example comes from Wang et al. (2024), who specified multiple dead start models to assess auto and cross-lagged relationships between various affect variables. We note that the absence of any contemporaneous relationships is rarely argued for theoretically.

2.2.6 Stimulus dynamics (SD)

Models restricting $\phi = 0$ do not include autoregressive parameters, only contemporaneous and lagged covariate parameters. We refer to these models as models of stimulus dynamics (SD). They have been used extensively to model temporal discounting of stimuli. For example, Bilacchi et al., 2022 fit a model predicting a participants temporal error at each trial based on reaction speed at time t and reaction speed at time $t-1$. This is a direct example of the SD model in table 1. Often researchers allow for more than one lag in models of stimulus dynamics. For example, the well known Rutledge computational model of happiness (Rutledge et al., 2014), specifies that the effects of the covariates decay exponentially over time without imposing an AR structure (see also Vanhasbroeck et al., 2024). This comes down to specifying many lagged covariate effects and restricting them to decay exponentially over time.

2.2.7 Autoregressive Distributed Lags (ADL)

At this point we have introduced a lot of different models without showing how they differ in the implied dynamics. The ADL model will help us in this respect. The ADL model derives its name from the fact that lags of both x_t and y_t are allowed for. It therefore allows for more freedom with respect to the dynamics compared to the previous models. It is important to realize that all the previous models are *special cases* of the ADL model. By *restricting* one or more parameters of the ADL model, one of the previous models is returned. For example, the LCV model is a special case of the ADL model in the sense that the LCV model *restricts* $\beta_{Lx} = 0$. The CF model is a special case in that it restricts $\beta_{Lx} = -\phi\beta_x$. The DS model is a special case in that it restricts $\beta_x = 0$. The *SD* model is a special case in that it restricts $\phi = 0$. The LR model restricts $\phi = \beta_{Lx} = 0$, and the AR model restricts $\beta_x = \beta_{Lx} = 0$. In Section 4, we will examine the consequences of imposing these restrictions on the shapes of the impulse response functions. First, we introduce system responses.

3 Introducing system responses

In this section, we provide a conceptual treatment of the tools we will use to interpret parameters of dynamic regression models. The goal of this section is to ensure that researchers understand how the different elements of the visualizations can be interpreted and applied in practice. Since AR(1) models are well known in psychology, we start our exposition of system responses by focusing on the AR(1) model.

3.1 Impulse responses

In section 2.2 we discussed the role of innovations as latent *inputs* to the system. For the AR(1) model, the parameter ϕ determines how the system responds to these innovations over time. It is well known that if $|\phi| < 1$, the influences of innovations decay exponentially over time. This "characteristic behavior" of the AR(1) model is difficult to notice from the equation ($y_t = \phi y_{t-1} + v_t$). It can however be retrieved from the impulse response function of the AR(1) model.

An impulse response function is a function which tells us how the system responds to an input variable *of unit magnitude* as a function of the number of time steps $0 \leq s \leq t$ ($s \in \mathbb{N}$) that have elapsed since the system was met with the input variable. We will denote an impulse response function as $h(s)$. The AR model has only one impulse response function because there is only one input vector, namely the innovations $\boldsymbol{v}$. In later parts of the manuscript, covariates and intercept parameters will also be present. Therefore, we specify which input vector an impulse response function is acting on in subscript. For the AR(1) model, we write the impulse response function as $h_v(s)$. It is given by:

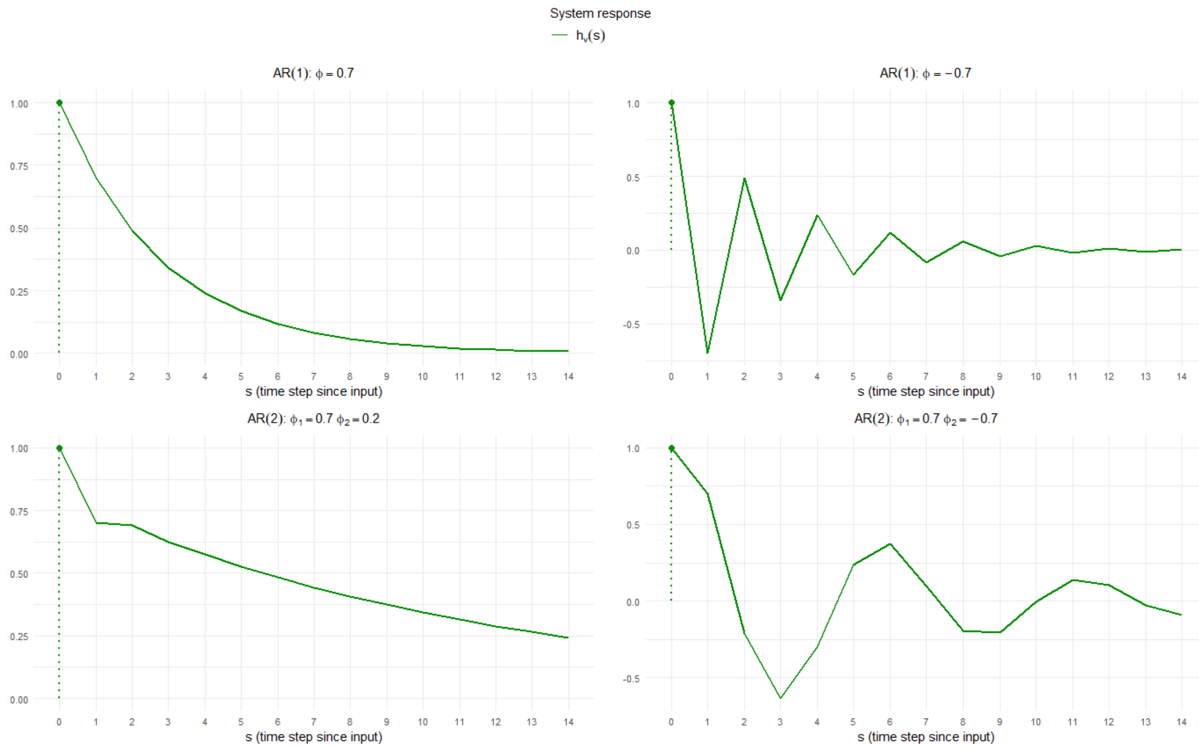

Figure 1: Impulse responses for AR models. The first row of plots displays the impulse response functions $h_v(s)$ for AR(1) models. On the left, $\phi = 0.7$ and on the right $\phi = -0.7$. The second row of plots displays the impulse response functions for AR(2) models, with different settings for the parameters ϕ_1 and ϕ_2. For each plot, an innovation of value 1 is presented at $s = 0$, illustrated with a green lollipop.

$$h_v(s) = \phi^s. \tag{1}$$

The function $h_v(s)$ provides the system response to a unit innovation for arbitrary time steps s. $h_v(0)$ gives us the immediate or contemporaneous effect of the innovation, $h_v(1)$ tells us the system response to this innovation after one time step has occurred, and so on. If $|\phi| < 1$, the impulse response decays exponentially over time (e.g. if $\phi = 0.7$, $1 > 0.7 > 0.7^2 > 0.7^3 > ...$). The closer ϕ is to 1 or -1, the slower this rate of decay becomes. If $|\phi| > 1$, the effect of an innovation will increase rather than decay as more time passes. Throughout the manuscript, we will be concerned with systems which are stable[1] in the sense that the influences of input variables decay rather than increase without bound.

It is clear that *the parameter of the model determines the "shape of decay"*. By seeing how changing the parameters changes the shape of the impulse response functions, we can understand how the parameters determine the system's behavior in a visual, intuitive way. In the first row of Figure 1, we visualize impulse responses for two AR(1) models with different settings for ϕ. We see that changing the value of the parameter ϕ changes the shape of the decay. If ϕ is positive, the decay is exponential. If ϕ is negative, one observes an exponentially decaying "zig-zag", or *sawtooth* pattern.

Although the exponential decay following perturbations is a well-known feature of AR(1) models, the dynamics implied by AR(p) models can be a lot more complicated. Nevertheless, we can use impulse responses to visualize these dynamics too. Consider for example the AR(2) model:

$$y_t = \phi_1 y_{t-1} + \phi_2 y_{t-2} + v_t \tag{2}$$

Impulse responses for two AR(2) models with different settings for the parameters ϕ_1 and ϕ_2 are visualized in the second row of Figure 1. In the first case, both AR effects are positive. In the second case, the first is positive and the second negative.

[1]If there are p AR parameters, the system is stable if all the roots of the lag polynomial $1 - \phi_1 L - \phi_2 L^2 - ... - \phi_p L^p$ are outside the complex unit circle, see Hamilton (1994), p. 34. For dynamic regression models with covariates the same restrictions on the AR parameters guarantee stability.

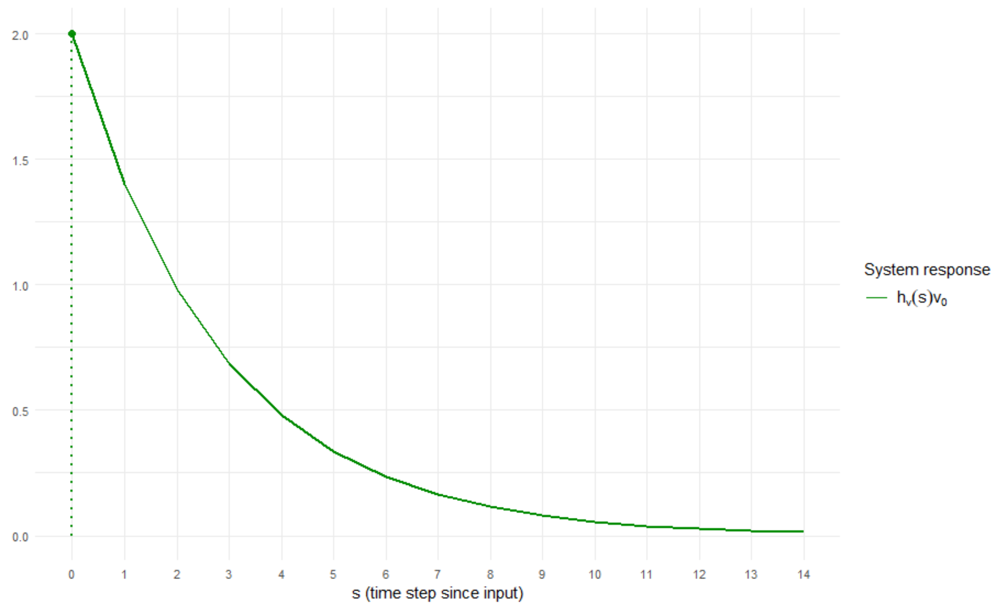

Figure 2: Scaled impulse response for an AR(1) model with $\phi = 0.7$. We observe that the value of the innovation (green dot) is now equal to 2, and the system responds accordingly.

The impulse response function for the AR(1) model had a very elegant form: $h_v(s) = \phi^s$. Although there is no easy way of expressing the impulse response function implied by AR(p) models (or dynamic regression models with covariates, for that matter), Appendix A allows one to calculate the impulse response regardless.

3.2 Scaled impulse responses

Impulse responses show us how the system responds to a unit input variable. However, usually the system will be met with input variables of non-unit magnitude. A first generalization of the impulse response is the *scaled impulse response*. Instead of supplying the system with a unit input, we can supply the system with an arbitrary value of the input variable, and trace its effect over time. For instance, instead of a unit, we may present the system with an innovation of magnitude v_0, where v_0 can in principle be any value. The scaled impulse response is given by:

$$h_v(s)v_0 = \{v_0, \phi v_0, \phi^2 v_0, ..., \phi^s v_0, ...\} \tag{3}$$

We can think about meaningful values of v_0, an example could be the standard deviation of the innovation process, σ_v. To illustrate, after fitting the AR(1) model, we might find that $\hat{\sigma}_v = 2$. Setting $v_0 = 2$ allows us to visualize how the system responds when the innovation at $t = 0$ is a positive standard deviation away from its mean (see Figure 2). Scaled impulse responses can thus be used to visualize system responses while taking into account the scale of the input variables.

3.3 Cumulative responses and total responses

A scaled impulse response tells us how the system responds to an input variable of arbitrary magnitude. Cumulative responses are a further generalization that allow us to assess how the system responds to an *arbitrary time series* of consecutive input variables. This is useful because it is not always evident how the system will behave in more complicated (and realistic) settings, where the inputs vary over time and the responses accumulate due to feed-forward effects.

Although the input variables may vary over time, the impulse response function of the system is assumed to be invariant over time[2]. Each time the system is met with an input variable, the system will respond by multiplying the current value of the input variable with the impulse response function at $s = 0$. However, previous values of the input variable will also still have an effect, to the extent that the impulse response function is non-zero at $s > 0$. As such, current and prior input variables are weighted by the impulse response function. The cumulative response is the sum of these weighted inputs at each time point.

[2]The parameters, e.g. ϕ, are not assumed to vary over the time index t

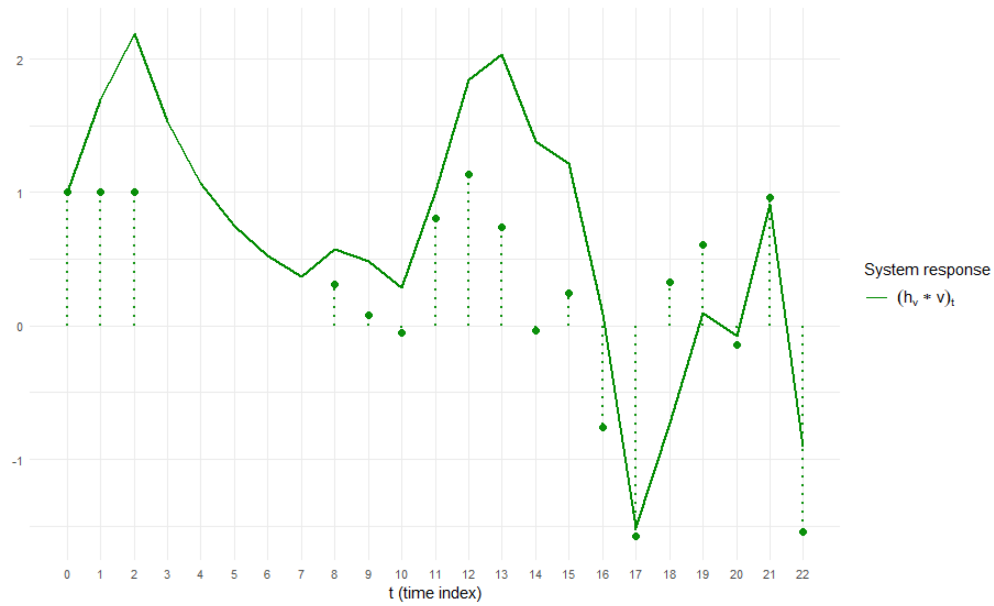

Figure 3: Cumulative response for an AR(1) model with $\phi = 0.7$. We observe that the innovations (green lollipops) now vary over time. The cumulative response towards the innovations $(h_v * v)_t$ is shown as a green line.

To make this concrete, imagine the input vector $\boldsymbol{v} = \{v_0, v_1, v_2, ..., v_t\}$ and the impulse response function $h_v(s) = \phi^s$. Then at the first measurement occasion, the cumulative response is given by:

$$y_0 = h_v(0)v_0 = v_0 \tag{4}$$

At the second measurement occasion, the cumulative response is given by:

$$\begin{aligned} y_1 &= h_v(1)v_0 + h_v(0)v_1 \\ &= \phi v_0 + v_1 \end{aligned} \tag{5}$$

At the third measurement occasion, the cumulative response is given by:

$$\begin{aligned} y_2 &= h_v(2)v_0 + h_v(1)v_1 + h_v(0)v_2 \\ &= \phi^2 v_0 + \phi v_1 + v_2 \end{aligned} \tag{6}$$

and so on. In general, the cumulative response can be written as a *convolution sum*:

$$y_t = \sum_{s=0}^{t} h_v(s)v_{t-s}. \tag{7}$$

It is worth recalling that parametric restrictions on the values of the AR parameters will change the shape of $h_v(s)$ (see Figure 1). Regardless of the precise shape of $h_v(s)$, the output y_t will always be given by a convolution of the impulse response function with the input vector. The output y_t is the *total response* of the system. When covariates are involved, the total response of the system will depend on not only the system response towards the innovations, as in this section, but also the system response towards the covariate.

When visualizing a cumulative response in a figure, we will denote it as $(h_\cdot * \cdot)_t$, where $\cdot$ is a placeholder for an input vector, and the $*$ symbol is a short way of writing a convolution like the right hand side of equation 7. Figure 3 visualizes the cumulative response of an AR(1) system towards a hypothetical time series of innovations. First, 3 innovations of unit magnitude are provided, after which the system is left to relax for 5 time points. Then, 15 innovations drawn from a standard normal distribution are provided. The figure shows how the effects of the innovations accumulate over time.

4 Interpreting dynamic regression models

It is clear that the concepts introduced above can already be used as interpretative tools. Depending on the restrictions imposed on AR parameters, researchers can expect different shapes of decay following innovation perturbations, information that is conveyed by the impulse response. We could fit the AR(1) model to real data, and use the estimate of the innovation variance to assess how the system responds to a "typical" innovation, information which is conveyed by the scaled impulse response. We can also provide time series of input variables to evaluate how the system responses accumulate, information that is conveyed by the cumulative responses.

In many cases, researchers have a particular covariate in mind to tie to the psychological process. In this section, we discuss model interpretation when there is a single time varying covariate x_t included in the model. As Section 2.2 illustrated, the covariate x_t can be included in dynamic regression models in many different ways. Here we will show that *how* the covariate is included changes the model-implied "shape of decay" in meaningful ways.

By including covariates, the model responds to inputs from both the innovations, $\boldsymbol{v}$, and the covariate, $\boldsymbol{x}$. These models, therefore, have 2 impulse response functions. As before, $h_v(s)$ is the impulse response function for the input vector $\boldsymbol{v}$. $h_x(s)$ is the impulse response function for the input vector $\boldsymbol{x}$. Because there are 2 impulse responses at play, the impulse response functions represent partial effects. In other words, $h_x(s)$ tells us how the system responds when it is perturbed by a single (unit) covariate x_t in isolation, keeping all other inputs at all other time points constant.

This leads us to an important question. What happens when we supply the system with an innovation and a covariate at the same time? How does the system integrate the responses to different input vectors? The *impulse response representation* (see Appendix A equation 33) of the model tells us how the system integrates the responses to the different input vectors (e.g. $\boldsymbol{x}$ and $\boldsymbol{v}$) over time to form the total response y_t at each time point. Specifically, the total response at each time point, y_t, is the sum of the individual system responses at each time point.

In Figure 4 we present an example of the impulse responses of each dynamic regression model. These figures will be used to guide the remainder of this section. For each dynamic regression, the model equation is provided on the left, together with the specific parameter values used in this example. On the right, impulse response plots are shown. Here, a unit covariate and unit innovation are supplied at the same time. The impulse response towards the innovation is provided as a green line, the impulse response towards the covariate is drawn in red. The value of the innovation is drawn as a green lollipop, and the value of the covariate is drawn as a red lollipop. In this case, the lollipops overlap since both inputs have a value of 1. The total response to these impulses is provided by a black dashed line. In this case, the total response represents how the system behaves when met with a unit innovation and a unit covariate at the same time.

To further highlight qualitative differences between the models, each figure contains a *trajectory plot* per model. These trajectory plots show us the cumulative responses of the systems towards the same time series of covariate and innovation values. Differences in the plots are thus purely due to the different restrictions imposed by each model. We start off by presenting the systems with 4 unit innovations. We then provide no innovations to the systems for 16 time points to show how the systems relax towards the innovations. We interrupt the systems at $t = 6$ by providing 4 unit covariate inputs, after which the covariate is again set to 0. After the relaxation period, we present the systems with 5 covariate inputs increasing in intensity, and 20 innovations, drawn from a standard normal distribution.

We start by discussing the ADL model, so that we can compare its dynamics with those of a restricted dynamic regression model. This makes it easier to tie parametric restrictions to qualitative behavior. The following conclusions hold irrespective of the specific values of the parameters or input vectors, depending only on how the parameters of the dynamic regression models are restricted.

1. ADL: The ADL allows for freedom with respect to the contemporaneous and lag-1 covariate effects. Varying β_x allows for freedom in how the system responds to the covariate contemporaneously, i.e. at $s = 0$. Varying β_{Lx} allows for freedom in how the system responds to the covariate one time step after it was met with the input variable, i.e. at $s = 1$. After $s = 1$, the system response is qualitatively similar as the system response towards the innovations: Exponential decay with a decay rate governed by $|\phi|$.

2. CF: Under the restriction $\beta_{Lx} = -\phi\beta_x$, the effect of the covariate *does not feed forward or last in time* (Appendix A.2, see also Ariens et al., 2023). The system has forgotten about the influence of the covariate immediately after a perturbation, the overall decay is determined purely by how the system responds to innovations. The system still displays memory towards the innovation process, but not towards the covariate.

3. LCV: Under the restriction $\beta_{Lx} = 0$, the way in which the covariate feeds forward in time is *qualitatively similar* (Appenidx A.3) as the way in which any other innovation would feed forward in time. There is, in a sense, nothing unique about our covariate from a dynamic point of view. By qualitatively similar is meant that the impulse response towards the covariate and the impulse response towards the innovations will have the same shape. In this case since $\beta_x = -1$, $h_x(s)$ has the same shape as $h_v(s)$, just reflected around the s axis.

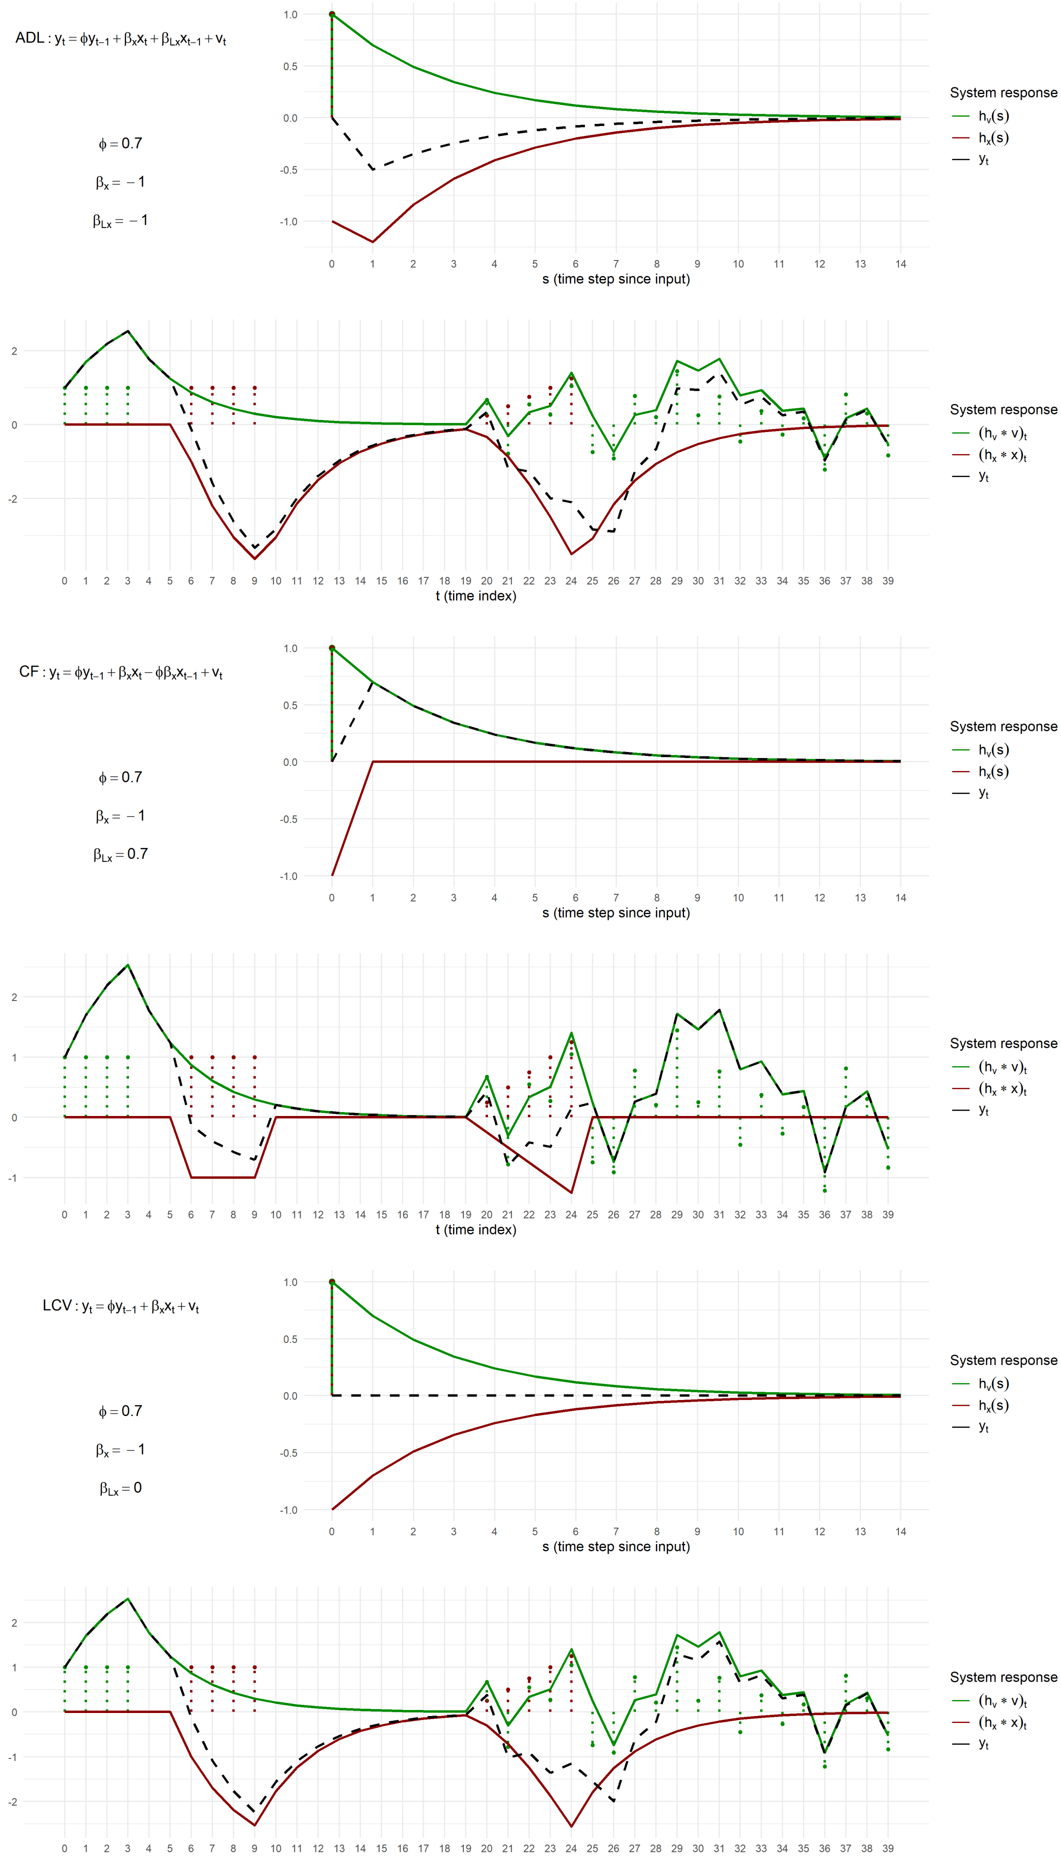

Figure 4: Parameters, impulse response plots, and trajectory plots for dynamic regression models. Details can be found in Section 4

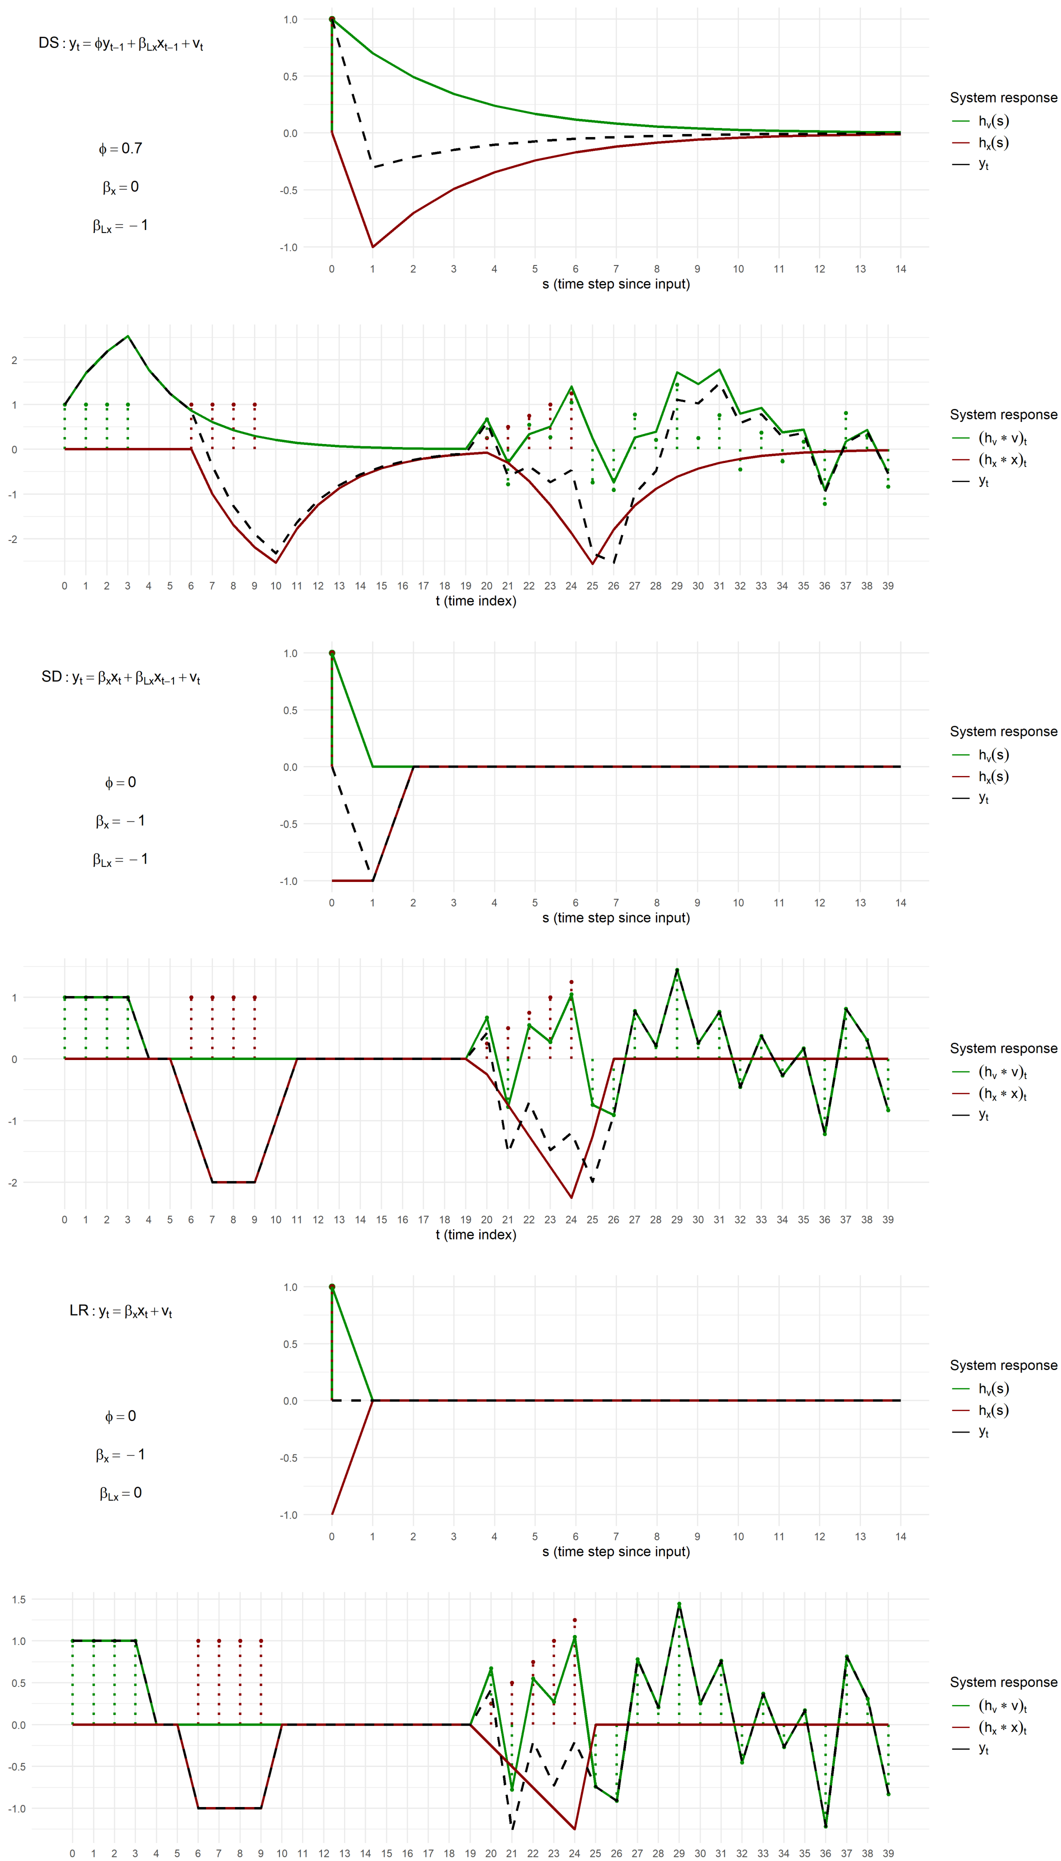

Figure 4: Continued

4. DS: Under the restrictions $\beta_{Lx} \neq 0$ but $\beta_x = 0$, we see that the system does not respond to the covariate immediately. The covariate elicits a purely delayed response. After this delayed response, the system responds as the LCV model responds: the influence of the covariate decays in precisely the same way as the influence of the innovations.

5. SD: Under the restriction $\phi = 0$, there is no autoregressive term in the model. The system only displays contemporaneous and lagged responses towards the covariate. Because there is no autoregressive term, $h_x(s)$ decays immediately after 1 time step has passed.

6. LR: The linear regression model can also be recovered by restricting all temporal parameters to 0. This model does not allow for any feed forward effects of any kind, all inputs are presumed to have strictly transient effects. The system is "memoryless".

From the impulse response plots, we can conclude that by restricting the parameters β_x and β_{Lx} we are changing how we believe the system is responding towards the covariate over time. If ϕ is non-zero, it also plays an important role in determining the shape of $h_x(s)$. Indeed, a zero value for β_{Lx} (a LCV restriction) does not mean that the effect of the covariate does not feed forward in time, it only means that the feed-forward behavior is not qualitatively distinct from the feed forward behavior towards the innovations.

The trajectory plots in Figure 4 further highlight the qualitative differences between these models. In each case, the conclusions from the impulse response plots explain the model-implied behavior even in more complex settings, where the input variables vary over time. Trajectory plots show us how the different cumulative responses are integrated to form the total output at each time point. The black dashed line is the model implied total output, y_t. If a cumulative response is positive, the system is "pushed" upwards. If a cumulative response is negative, the system is "pushed" downwards. The impulseR function `irf_generator()` can be used to calculate the different types of system responses for dynamic regression models (see impulseR, Vignette `irf_generator()`).

4.1 Intercepts and process means

Intercept parameters are straightforward to include in the different dynamic regression models we have discussed. They are also useful, since they allow for non-zero means (equilibrium states) in autoregressive models. To understand the role of the intercept parameter for the system responses, notice that the intercept implicitly implies a covariate that takes on a value of 1 for all time points. We can thus investigate how a system responds when an intercept is included by including the cumulative response to this covariate over time (see Appendix A, equation 33). The cumulative response towards the unit vector is denoted by $\sum_{s=0}^{t} h_1(s)1_{t-s}$, and for an AR(1) model it is given by:

$$\sum_{s=0}^{t} h_1(s)1_{t-s} = \alpha + \phi\alpha + \phi^2\alpha + ... + \phi^t\alpha. \tag{8}$$

If only this intercept is present (no innovations or other covariates), the total response y_t will be given by the cumulative response towards the unit vector. For the first, second, and third measurement occasions:

$$\begin{aligned} y_0 &= \alpha \\ y_1 &= \alpha + \phi\alpha \\ y_2 &= \alpha + \phi\alpha + \phi^2\alpha \end{aligned} \tag{9}$$

As we add more and more time points, this cumulative response converges[3] to the process mean $\frac{\alpha}{1-\phi}$, and changes no longer as the time index increases. The intercept parameter is the point where the cumulative response towards the unit vector starts (at $t = 0$). The model-implied mean or equilibrium state of the process is the point which the cumulative response towards the unit vector *approaches*, as $t \to \infty$. These two points are annotated in Figure 5, which visualizes the cumulative response towards the unit vector for the model $y_t = 5 + 0.5y_{t-1}$ as a gray line. The cumulative response starts at $\alpha = 5$ and approaches $\frac{\alpha}{1-\phi} = \frac{5}{1-0.5} = 10$. By introducing or changing the value of an intercept parameter, the process will not reach a new equilibrium state immediately, but gradually (Lütkepohl, 2005, p. 605).

When simulating data, this "warm-up" behavior is often removed, or *burned in.* Specifically, to obtain a time series of T measurements, one generates a large number (e.g. 1000 + T) of data and only retains the final T observations (e.g. Bulteel, Mestdagh, et al., 2018; Ariens et al., 2022 ;Y. Li et al., 2022). This allows one to study the system as

[3]If there are p AR effects, $\lim_{t\to\infty} \sum_{s=0}^{t} h_1(s)1_{t-s} = \frac{\alpha}{1-\sum_{i=1}^{p}\phi_i}$

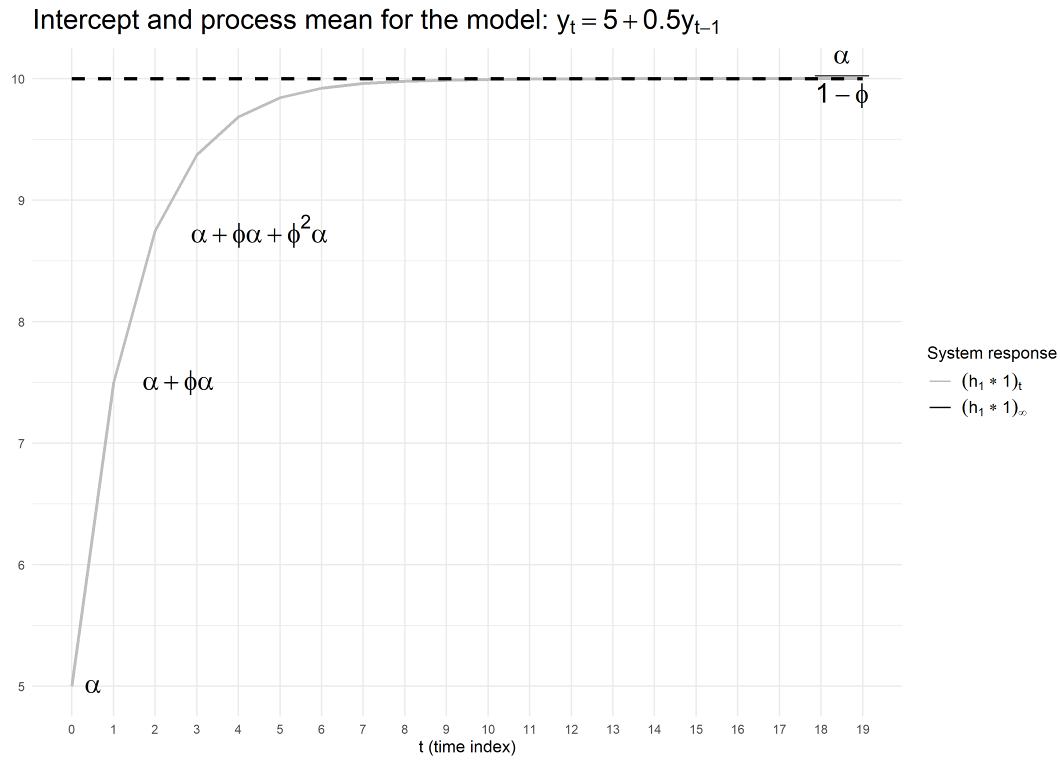

Figure 5: The interpretation of the intercept parameter, and an illustration of "burning in the intercept". Data is generated from the model $y_t = 5 + 0.5y_{t-1}$. Two trajectory plots are shown. In gray, the cumulative response towards the unit vector $(h_1 * 1)_t$ is shown. The cumulative response can be seen to start at $\alpha = 5$ and approach $\frac{\alpha}{1-\phi} = \frac{5}{1-0.5} = 10$. In black, $(h_1 * 1)_t$ is fixed to its model implied equilibrium state, $(h_1 * 1)_\infty = \frac{\alpha}{1-\phi} = 10$ for all time points (i.e., it is burned in).

if it had already approached the mean, or was already in its equilibrium state $\frac{\alpha}{1-\phi}$, at $t = 0$. Similarly, impulseR can conduct a *burn in* analytically when generating data, by fixing the part of the system response due to the intercept to its model-implied equilibrium state (Figure 5, black dashed line). However, when constructing estimated trajectory plots (see Section 5) we will use the information contained in the intercept estimate, and not burn in this part of the response. We discuss impulseR's default treatment of intercept parameters in the impulseR documentation, see Vignettes `irf_generator()` and `irf_empirical()`.

It is important to note that if a covariate x_t is included in the model, then $\frac{\alpha}{1-\phi}$ is no longer the model-implied mean of y_t. Rather, $\frac{\alpha}{1-\phi}$ is the model-implied mean of y_t if the covariate was not present, or fixed to $\mathbf{0}$. Centering x_t such that $x_t^* = x_t - \bar{x}$ prior to fitting a lag-1 dynamic regression model can thus allow one to interpret $\frac{\hat{\alpha}}{1-\hat{\phi}}$ as an estimate of the expected value of y_t given an average value of the covariate.

5 An empirical application

In this section, we illustrate how the R package impulseR can be used to interpret the results of a dynamic regression analysis. We will apply the package functions to an empirical dataset. We reanalyze analyze data from a previously published experience sampling (ESM) study (Dejonckheere et al., 2021) where the affect states of university students were monitored in the period preceding, during, and following the announcement of their exam grades. Ethical approval for the secondary processing of these data was granted by the local ethics committee (PRET KU Leuven, G-2023-7445)

5.1 Data

Over the course of 9 consecutive days, participants were prompted to respond to 10 momentary questionnaires according to a stratified random interval scheme. These momentary questionnaires were presented each day between 10 AM and 10 PM. The ESM protocol started 2 days prior to the students reading their exam grades. At each measurement occasion, participants were presented with the question "please indicate how positive you are feeling right now", measured on a slider scale from 0 to 100. The values were not visible to the participants. This serves as an indicator of positive affect (PA). After this first question, 6 negative affect questions were presented targeting their feelings specifically towards the exam context. For example, one question was "When you think about your exam grades right now, how stressed are you feeling". Dejonckheere et al. (2021) calculated the mean of the exam-related negative affect items at each time point to form a latent indicator of *exam-related negative affect* (exam NA). We will use the reproducible data file, made available by Dejonckheere et al., 2021 at https://osf.io/yte2w/files/osfstorage. This file consists of the data of 101 participants. For further details about the data, see Dejonckheere et al., 2021.

5.2 Procedure

We will fit dynamic regression models to these data aimed at modeling how the PA system (PA, Y) of an individual responds to exam-related negative affect (exam NA, X). Since extending impulse response analyses to $N = many$ cases lies beyond the scope of the paper, we perform the analysis on a single individual. The general PA time series of all individuals were first screened visually, this participant was selected since a clear difference between the general PA before and after seeing the exam results was visible in the time series. We selected participant with identification code 227. The participants had responses at 90 time points, of which 4 contained missing values for PA. The analysis was conducted on observations from 86 time points.

We fit two dynamic regression models to the data to illustrate how the implied dynamics change depending on how we restrict the model parameters. Specifically, we will fit a lag-1 ADL model and a lag-1 LCV model to these data. From Section 4, we would expect that the LCV model implies that the PA system responds to exam related negative affect in a qualitatively similar way as it does to other innovations to the PA system. The ADL model on the other hand allows exam NA to have unique effects on the PA system up to lag-1.

After fitting the models, we interpret the parameter estimates and p-values together with visualizations supplied by the impulseR package. We will rely on *scaled impulse response plots* (which take into account the scale of the covariate and estimated innovations), and *estimated trajectory plots*. These are analogous to the trajectory plots in Section 4, but where instead of providing hypothetical parameter values and input vectors, we use parameter estimates, observed covariate values, and estimated innovations to visualize how the system was estimated to respond over the observation period. The impulseR function `irf_empirical()` provides a user friendly way to fit a dynamic regression model using OLS, and calculate these responses while handling technical details (see Appendix B) automatically.

5.2.1 Analysis

The code used for the analyses is provided as supplementary materials (see https://osf.io/fn7mj/). We load in the data of the participant and install and load the impulseR package. We then call the `irf_empirical()` function to the data. The function uses the lm() function in R to fit the model with ordinary least squares, and supplements the lm() output with the estimated trajectories over the study period. These are used to construct *estimated trajectory plots*. In addition, the fit object contains information such as $\hat{\sigma}_v$, the estimated standard deviation of the innovations and $\bar{X}$, the sample mean of exam NA, which can be used to construct scaled impulse response plots. We will visualize how the PA system is estimated to respond to a standard deviation increase in the innovations, and an increase in the covariate of value $\bar{X}$, where $\bar{X}$ is calculated in the period prior to the participant seeing their exam results. This gives us an indication of how the system responds to a typical value of exam NA during this period, though other choices could have been possible.

The output is shown in Figures 6 and 7, and is interpreted together with other elements of the output of the lm() fit object, such as $p-$values, next. Significance is judged at the .05 level

5.3 Results and interpretation

We start by interpreting the LCV model solution. The top left of Figure 6 shows the parameter estimates, the right shows the scaled impulse response plot. As one might expect, increases in exam NA are associated with contemporaneous decreases in general PA. $\hat{\beta_x}$ is negative, and significantly different from 0: $t(82) = -5.852$, $p < .001$. The estimate $\hat{\phi}$ is not significantly different from 0: $t(82) = 1.867$, $p = .066$. Using the scaled impulse response, we would conclude that an increase in exam NA from 0 to $\bar{X} = 24.28$ will be associated with a change in the PA score of -17.69 points contemporaneously. The system relaxes following this perturbation. Due to the LCV model structure, the system is restricted to decay towards the exam-related negative affect in a qualitatively similar way as it decays towards innovations.

Moving to the trajectory plot, we see that at the start of the study, exam NA (red lollipops) varies and slightly increases over time until the participant sees their exam results. After this, exam NA drops immediately and is close to 0 for the remainder of the study. Some slight nonzero reports of exam NA can be seen after this period, which might indicate that some measurement error (e.g. Schuurman et al., 2015) is present. The estimated latent response to the exam NA is shown with a red line. We see that when the exam NA is "turned off", the PA system rapidly

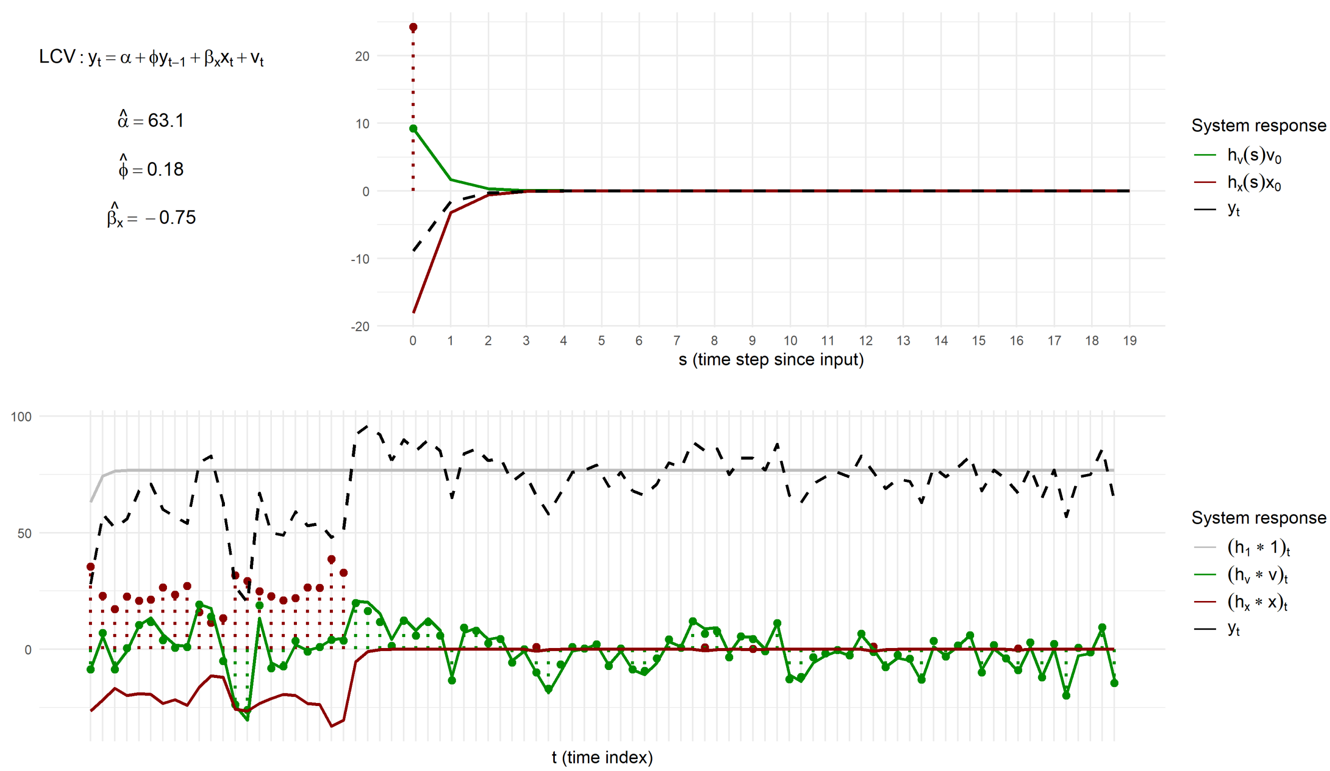

Figure 6: Visualizations of estimated system responses for the LCV model. On the left, the fitting model together with the parameter estimates are displayed. On the right, a scaled impulse response plot shows how the system is estimated to respond when met with an innovation of magnitude $\hat{\sigma}_v = 9.21$ (green lollipop) and a covariate of magnitude $\bar{X} = 24.28$ (red lollipop). These scaled impulse responses are shown by green and red lines respectively. The model implied response of the system towards these inputs at the same time is given by the total response, y_t (black dashed line). Below, an estimated trajectory plot is shown. The observed covariate values x_t are displayed as red lollipops. The estimated innovations $\hat{v}_t$ are displayed as green lollipops. Lines indicate estimated cumulative responses. The implied total response y_t is given as a black dashed line, and will equal the observed values of the criterion variable y_t^{obs} (not shown).

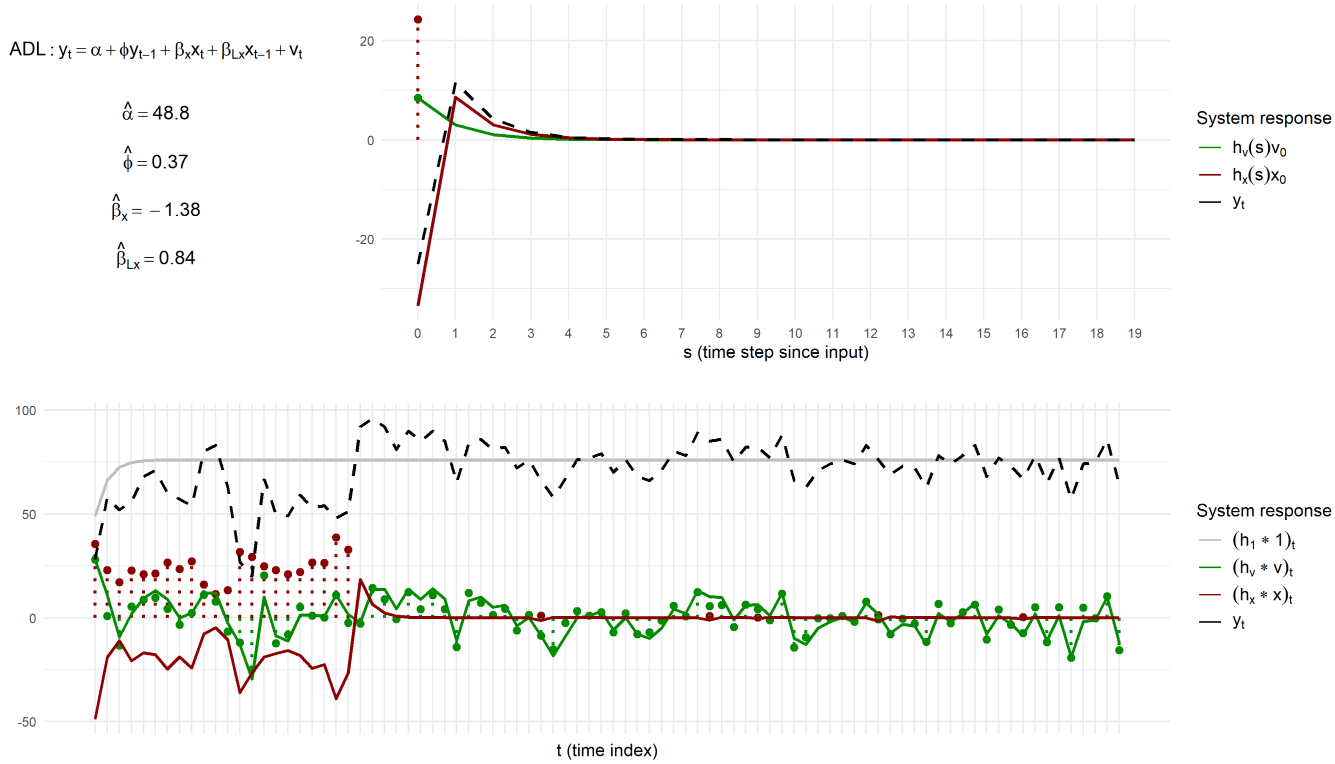

Figure 7: Visualizations of system responses for the ADL model. The figure structure is analogous to that of Figure 6, save for the fact that for the ADL model, $\hat{\sigma}_v = 8.44$.

recovers. The estimated innovations $\hat{v}_t$ under the LCV model are shown as green lollipops, and the cumulative response to these innovations as a green line. The response to the innovations should display autocorrelation, but as per the model assumptions, the estimated innovations themselves should be independent over time. The estimated trajectory plot can be used to gain a first appraisal of the validity of the assumption that the innovations are independent and identically distributed over the time index t. Here, $\hat{v}_t$ appears to display autocorrelation (violating the independence assumption), in particular surrounding the period where the exam NA is "turned off", where $\hat{v}_t$ appears clustered in time. Overall, the autocorrelation of the estimated innovations was positive, though not significantly different from 0: $\hat{\rho} = 0.20$, $t(82) = 1.803$, $p = .075$.

Turning to the ADL, we see some noticeable differences. First, the estimate $\hat{\beta}_x$ remains significantly different from 0: $t(81) = -6.92$, $p < .001$, but is now more negative. Next, $\hat{\phi}$ has become more positive, and is now significantly different from 0: $t(81) = 3.6$, $p < .001$. The lagged covariate effect $\hat{\beta}_{Lx}$ is estimated to be positive, and is significantly different from 0: $t(81) = 3.939$, $p < .001$. Freeing β_{Lx} to be nonzero does seem to have an impact on the conclusions (see Ariens et al., 2023 for a discussion on how estimates change depending on how the dynaimcs are specified). What does this positive estimate tell us about the estimated dynamics? The scaled impulse response plot shows that the PA system will tend to *overshoot* after an exam NA perturbation. An increase in exam NA from 0 to $\bar{X} = 24.28$ is associated with a change in the PA score of -32.70 points on the PA scale contemporaneously. One time step after this exam NA perturbation, the participant responds with a PA score on average 8.39 points higher than their average PA score (all else held constant). The overshoot can also be seen clearly in the trajectory plot: The moment from which exam NA is reported to be close to 0, the PA system overshoots. Such overshooting behaviors may emerge due to relief processes for instance, and are clearly of theoretical relevance for the field of affect dynamics. The estimated innovations surrounding the period where exam NA transitions from being "on" to "off" display less obvious signs of autocorrelation. Overall, The autocorrelation of the estimated innovations was not significantly different from 0, $\hat{\rho} = -0.11$, $t(82) = -0.977$, $p = .33$.

Finally, the gray lines in the plots are the part of the responses due to the intercept parameter α. They approach $\frac{\hat{\alpha}}{1-\hat{\phi}}$, which equals 76.83 for the LCV model, and 75.95 for the ADL model. Both figures illustrate well that these points are estimates of the expected values of PA given that the exam NA equals $\mathbf{0}$ (see section 4.1).

6 Discussion

The empirical illustration demonstrates several advantages of using system responses to aid in interpreting the results of a dynamic regression analysis. First, since the estimated latent system responses are visualized, the behavior of the system can be conveyed and reasoned about visually, through impulse response plots, scaled impulse response plots, and trajectory plots. These make it intuitive to see how a system behaves over time, and how partial responses combine to form the total output. Essentially, we can go from interpreting individual parameter values on a case-by-case basis to interpreting the system responses as a whole. Second, the responses will always be calculated on the same scale as the data. One could transform variables a priori, the system responses will be calculated relative to this new scale. This facilitates interpretation and highlights that the materials can be applied to a great variety of data analytic scenarios. Finally, since we can calculate system responses for arbitrary order dynamic regression models, it becomes conceptually feasible to interpret the behavior of linear systems with more general lag structures than those typically employed in psychological applications.

It is however crucial to keep in mind that system responses can be used to visualize model-implied behavior *regardless of the quality of the model.* For example, if in reality a lag-5 ADL generated the data but one fits a linear regression model, one will still be able to visualize the impulse responses and cumulative responses implied under the fitting model despite the fitting model being misspecified and the parameter estimates untrustworthy.

In an applied data analysis situation, one should perform careful checks of the model assumptions while interpreting the results of a (dynamic) regression analysis (Ariens et al., 2023). For example, if there are time-series dependencies among the estimated innovations such as autocorrelation or changes in variance, one should reconsider the model specification (see Ariens et al., 2023). Further lags may be required (Ivanov and Kilian, 2005), some dynamics may require nonlinear models to be accommodated for (e.g. Vanhasbroeck et al., 2021), or require time-varying parameters to capture adequately (e.g. Adolf et al., 2017). Many assumptions about the innovation process can be checked using classical regression residual diagnostics, though to check for dependence over time, additional tools such as autocorrelation plots are often recommended (for a thorough overview on time series diagnostics, see W. K. Li (2003)). The estimated trajectory plots from impulseR can provide a visual inspection of the estimated innovations of a model over the observation period (green lollipops in the estimated trajectory plots of Figures 6 and 7).

Besides model misspecification, there are many other reasons why a model solution might not be "reliable". There might be a limited signal-to-noise ratio in the data (for instance due to a suboptimal study design, e.g. Tuerlinckx et al., 2025), there may be small sample biases (e.g. Maeshiro, 2000) or uselessly large uncertainty estimates if limited numbers of time points are available, and our measures may be contaminated with measurement error, biasing the parameter estimates if not adequately accounted for in the model specification (Schuurman et al., 2015). In short, the materials allow one to visualize the results of an analysis. The materials do not guarantee that this analysis was appropriate. We close by discussing some extensions of the materials.

6.1 Extensions

6.1.1 Dynamic regression models with more general lag structures

To maximize the accessibility of the materials, we focused mostly on lag-1 dynamic regression models. Indeed, these appear to us to be the most commonly specified dynamic regression models in the psychological literature. However, Appendix A derives impulse response representations of ADL models with arbitrary lag orders. As such, impulseR can be used to calculate and visualize system responses for a very broad class of dynamic regression models. To make this concrete, imagine a very general ADL(p, q) model, where we allow for p AR effects and q lagged covariate effects:

$$y_t = \alpha + \beta_x x_t + \sum_{i=1}^{p} \phi_i y_{t-i} + \sum_{j=1}^{q} \beta_{L^j x} x_{t-j} + v_t. \tag{10}$$

If $p = 1$, the above reduces to

$$y_t = \alpha + \beta_x x_t + \phi y_{t-1} + \sum_{j=1}^{q} \beta_{L^j x} x_{t-j} + v_t. \tag{11}$$

If in addition $q = 1$, the above reduces to the lag-1 ADL provided in Table 1. Crucially we know what the qualitative implications are of different restrictions, and these also hold for models with more general lag structures. It will always be true that for a LCV model (i.e. a model where all lagged covariate effects are 0), the way in which the covariate feeds forward in time is qualitatively similar as the way in which the innovations feed forward in time. CF restrictions guarantee that the effect of the covariate is purely contemporaneous, regardless of how the innovations feed forward in time. These claims are proven in Appendices A.3 and A.2 respectively.

6.1.2 Custom estimators

Although designed to be easy to use, the `irf_empirical()` function we used in the empirical application is limited by the lm() function in R. For example, missing values are handled by listwise deletion, and lm() uses OLS estimation. For researchers seeking more flexibility with respect to estimation, the `irf_generator()` function can be used directly to construct estimated impulse response plots or estimated scaled impulse response plots when other estimators are used. Indeed, all one needs to supply are the parameter estimates and use information returned by your custom fit object. For estimated trajectory plots, some technical details concerning how the estimator treated initial values require clarification, see Appendix B.

6.1.3 Multiple covariates

To keep the materials of the paper contained, we focused on a single covariate and possibly an intercept. From a technical point of view, it is straightforward to extend the materials to allow for multiple covariates. If there are m covariates indexed by d, then the impulse response representation will be given by:

$$y_t = \sum_{s=0}^{t} h_1(s) 1_{t-s} + \sum_{d=1}^{m}\sum_{s=0}^{t} h_{x_d}(s) x_{d\,t-s} + \sum_{s=0}^{t} h_v(s) v_{t-s}. \tag{12}$$

Currently, impulseR only allows for a single covariate and an intercept term. Calculating the values of each impulse response function $h_{x_d}(s)$ can however be done using the techniques of Appendix A.

A Impulse response representations

These appendices use formal power series methods to derive the impulse response representations of ADL(p,q) models. Since the coefficients of the moving average representations correspond to the coefficients of the impulse response functions (e.g. Lütkepohl, 2005, p. 52), we first derive the moving average representation of the ADL(p,q). We start by deriving the MA representation of the ADL(2,2). We discover recurrence relationships which allow us to map the parameters of the ADL models to the values of the different impulse response functions via the moving average representations. Solutions for ADL(p,q) models are then derived by induction. The impulseR package uses these solutions to generate ADL system responses for the user. In Appendices A.2 and A.3, we use these results to show the implications of CF restrictions and LCV restrictions respectively.

A.1 Moving average representations of distributed lag models

The lag-2 ADL is defined as:

$$y_t = \alpha + \phi_1 y_{t-1} + \phi_2 y_{t-2} + \beta_x x_t + \beta_{Lx} x_{t-1} + \beta_{L^2x} x_{t-2} + v_t \tag{13}$$

We introduce the lag operator L, which acts on a variable x_t such that $Lx_t = x_{t-1}$. We can lag a variable k times by applying the operator k times: $L^k x_t = x_{t-k}$ (for further details, see Hamilton, 1994, Chapter 2). It is useful to notice that $L^0 x_t = x_t$. As such, L^0 is an identity operator.

We start off by rewriting the ADL in terms of polynomials in the lag operator:

$$(L^0 - \phi_1 L^1 - \phi_2 L^2) y_t = \alpha L^0 1_t + (\beta_x L^0 + \beta_{Lx} L^1 + \beta_{L^2x} L^2) x_t + L^0 v_t. \tag{14}$$

$1_t = \mathbf{1}$. The moving average representation can be derived by multiplying this equation by $(L^0 - \phi_1 L^1 - \phi_2 L^2)^{-1}$, which we abbreviate with $\phi(L)^{-1}$:

$$\phi(L)^{-1}\phi(L) y_t = \phi(L)^{-1}\alpha L^0 1_t + \phi(L)^{-1}(\beta_x L^0 + \beta_{Lx} L^1 + \beta_{L^2x} L^2) x_t + \phi(L)^{-1} L^0 v_t. \tag{15}$$

We first derive $\phi(L)^{-1}$. We will assume that the roots of $\phi(L)$ are all outside the complex unit circle, ensuring that the process is stable. Under this assumption, the inverse and the following identity are well-defined (Hamilton, 1994, p.34, eq 2.4.6):

$$(L^0 - \phi_1 L^1 - \phi_2 L^2)^{-1}(L^0 - \phi_1 L^1 - \phi_2 L^2) = L^0. \tag{16}$$

We also know that $\phi(L)^{-1}$ can be written as a formal power series in the lag operator (e.g. Loossens, 2021, theorem B.2.1):

$$(L^0 - \phi_1 L^1 - \phi_2 L^2)^{-1} = \sum_{k=0}^{\infty} \theta_k L^k. \tag{17}$$

We derive the θ_k coefficients by multiplying the previous equation by $(L^0 - \phi_1 L^1 - \phi_2 L^2)$:

$$L^0 = (L^0 - \phi_1 L^1 - \phi_2 L^2)(\theta_0 L^0 + \theta_1 L^1 + \theta_2 L^2 + ...). \tag{18}$$

Expanding the right hand side yields:

$$L^0 = L^0(\theta_0 L^0 + \theta_1 L^1 + \theta_2 L^2 + ...) - \phi_1 L^1(\theta_0 L^0 + \theta_1 L^1 + \theta_2 L^2 + ...) - \phi_2 L^2(\theta_0 L^0 + \theta_1 L^1 + \theta_2 L^2 + ...) \tag{19}$$

Factoring the right hand side into powers of L, and using the fact that the identity operator L^0 is idempotent:

$$L^0 = \theta_0 L^0 + (\theta_1 - \phi_1\theta_0)L^1 + (\theta_2 - \phi_1\theta_1 - \phi_2\theta_0)L^2 + (\theta_3 - \phi_1\theta_2 - \phi_2\theta_1)L^3 + ... \tag{20}$$

Knowing that the coefficient for the identity operator L^0 is θ_0, we can fix θ_0 to the real number 1. From this, the equality $L^0 = L^0$ can be made to hold by forcing all terms involving nonzero powers of L to cancel. This yields the following solutions for θ_k

$$\begin{aligned}
&\theta_0 = 1 \\
&\theta_1 = \phi_1 \\
&\theta_2 = \phi_1\theta_1 + \phi_2 \\
&\theta_3 = \phi_1\theta_2 + \phi_2\theta_1 \\
&\quad . \\
&\quad . \\
&\quad . \\
&\theta_k = \phi_1\theta_{k-1} + \phi_2\theta_{k-2}.
\end{aligned} \tag{21}$$

As such, we can rewrite Equation 15 as:

$$y_t = \sum_{k=0}^{\infty} \alpha\theta_k L^k 1_t + \sum_{k=0}^{\infty} \theta_k L^k(\beta_x + \beta_{Lx}L + \beta_{L^2x}L^2)x_t + \sum_{k=0}^{\infty} \theta_k L^k v_t. \tag{22}$$

Now we derive $\sum_{k=0}^{\infty} \theta_k L^k(\beta_x L^0 + \beta_{Lx}L^1 + \beta_{L^2x}L^2)$. We can again write this as a formal power series in the lag operator, which we denote by $\sum_{k=0}^{\infty} \psi_k L^k$:

$$\sum_{k=0}^{\infty} \psi_k L^k = (L^0 - \phi_1 L^1 - \phi_2 L^2)^{-1}(\beta_x L^0 + \beta_{Lx}L^1 + \beta_{L^2x}L^2). \tag{23}$$

By expanding and factoring the right hand side in terms of L, we have:

$$\begin{gathered}
\sum_{k=0}^{\infty} \psi_k L^k \\
= \theta_0 L^0(\beta_x L^0 + \beta_{Lx}L^1 + \beta_{L^2x}L^2) + \theta_1 L^1(\beta_x L^0 + \beta_{Lx}L^1 + \beta_{L^2x}L^2) + \theta_2 L^2(\beta_x L^0 + \beta_{Lx}L^1 + \beta_{L^2x}L^2) + ... \\
= \theta_0\beta_x L^0 + (\theta_1\beta_x + \theta_0\beta_{Lx})L^1 + (\theta_2\beta_x + \theta_1\beta_{Lx} + \theta_0\beta_{L^2x})L^2 + (\theta_3\beta_x + \theta_2\beta_{Lx} + \theta_1\beta_{L^2x})L^3 + ...
\end{gathered} \tag{24}$$

The coefficients ψ_k of $\sum_{k=0}^{\infty} \psi_k L^k$ are thus provided by the recurrence relation:

$$
\begin{aligned}
\psi_0 &= \theta_0\beta_x \\
\psi_1 &= \theta_1\beta_x + \theta_0\beta_{Lx} \\
\psi_2 &= \theta_2\beta_x + \theta_1\beta_{Lx} + \theta_0\beta_{L^2x} \\
\psi_3 &= \theta_3\beta_x + \theta_2\beta_{Lx} + \theta_1\beta_{L^2x} \\
&. \\
&. \\
&. \\
\psi_k &= \theta_k\beta_x + \theta_{k-1}\beta_{Lx} + \theta_{k-2}\beta_{L^2x}.
\end{aligned}
\tag{25}
$$

Combining the results, we have the moving average representation:

$$y_t = \sum_{k=0}^{\infty} \alpha\,\theta_k L^k 1_t + \sum_{k=0}^{\infty} \psi_k L^k x_t + \sum_{k=0}^{\infty} \theta_k L^k v_t. \tag{26}$$

In this way, the moving average representation for AR models with more general lag structures can be derived. Let p denote the AR order, and q denote the number of free lagged covariate effects. In general,

$$\theta_k = \phi_1\theta_{k-1} + \phi_2\theta_{k-2} + ... + \phi_p\theta_{k-p}, \tag{27}$$

and

$$\psi_k = \beta_x\theta_k + \beta_{Lx}\theta_{k-1} + ... + \beta_{L^qx}\theta_{k-q}. \tag{28}$$

A.1.1 Impulse response representation

For a finite time series length, we can rewrite the moving average representation as:

$$y_t = \theta_t y_0 + \sum_{k=0}^{t-1} \alpha\,\theta_k L^k 1_t + \sum_{k=0}^{t-1} \psi_k L^k x_t + \sum_{k=0}^{t-1} \theta_k L^k v_t. \tag{29}$$

Specifying the initial condition as $y_0 = \alpha + \beta_x x_0 + v_0$, we can write the above as:

$$y_t = \sum_{k=0}^{t} \alpha\,\theta_k L^k 1_t + \sum_{k=0}^{t} \psi_k L^k x_t + \sum_{k=0}^{t} \theta_k L^k v_t. \tag{30}$$

To show that the impulse response functions are given by the coefficients of the MA representation, we use the change of variables $s = t - k$:

$$y_t = \sum_{s=0}^{t} \alpha\theta_{t-s} 1_s + \sum_{s=0}^{t} \psi_{t-s} x_s + \sum_{s=0}^{t} \theta_{t-s} v_s. \tag{31}$$

Since convolutions of functions[4] commute (Phillips et al., 2003, eq 10.23), this is equivalent to

$$y_t = \sum_{s=0}^{t} \alpha\theta_s 1_{t-s} + \sum_{s=0}^{t} \psi_s x_{t-s} + \sum_{s=0}^{t} \theta_s v_{t-s}. \tag{32}$$

The impulse response function $h_x(s)$ is thus given by the coefficients ψ_k: $h_x(s) = \psi_k$. In other words, the k'th MA coefficient ψ_k is exactly the value of the impulse response function towards the covariate at time shift s. Likewise, $h_v(s) = \theta_k$ and $h_1(s) = \alpha\theta_k$. As such, we can rewrite the above as:

$$y_t = \sum_{s=0}^{t} h_1(s) 1_{t-s} + \sum_{s=0}^{t} h_x(s) x_{t-s} + \sum_{s=0}^{t} h_v(s) v_{t-s}. \tag{33}$$

The impulse response representation tells us that at each time point the output y_t is given by a sum of the convolutions of the impulse response functions with the input vectors.

[4] Implicit in the above is that $x_s = \delta_{t-s} x_t$, where δ_{t-s} is the shifted unit sample function, Phillips et al. (2003), p.493

A.2 Common factor restrictions

To derive the CF restrictions for the lag-2 case, recall that $h_x(s) = \psi_k$ and $h_v(s) = \theta_k$. If the impulse response $h_x(s)$ is to decay immediately at $s > 0$, we require from equation 25 that $\psi_1 = 0$ and $\psi_2 = 0$:

$$\begin{aligned}\psi_1 = 0 &\implies \beta_{Lx} = -\phi_1\beta_x \\ \psi_2 = 0 &\implies \beta_{L^2x} = -\phi_1\beta_{Lx} - (\phi_1^2 + \phi_2)\beta_x\end{aligned} \tag{34}$$

If both ψ_1 and ψ_2 are equal to 0, we have:

$$\begin{aligned}\psi_1 = \psi_2 = 0 \implies \beta_{Lx} &= -\phi_1\beta_x \\ \beta_{L^2x} &= -\phi_2\beta_x.\end{aligned} \tag{35}$$

It follows that if there are q free lagged covariate effects, the CF restrictions will be satisfied if:

$$\beta_{L^qx} = -\phi_q\beta_x \tag{36}$$

for all $q > 0$. To prove that these restrictions are sufficient, notice that since the roots of $\phi(L)$ are outside the complex unit circle, $|h_x(s)|$ will be strictly decreasing at $s > q$. As such, if $\psi_1 = \psi_2 = ... = \psi_q = 0$, $\psi_k = 0$ for all $k > q$.

A.3 LCV restrictions

LCV restrictions imply that all lagged covariate effects are 0. Equation 25 implies the following:

$$\begin{aligned}\psi_0 &= \theta_0\beta_x \\ \psi_1 &= \theta_1\beta_x + 0 \\ \psi_2 &= \theta_2\beta_x + 0 + 0 \\ &. \\ &. \\ &. \\ \psi_k &= \theta_k\beta_x.\end{aligned} \tag{37}$$

Since $h_x(s) = \psi_k$ and $h_v(s) = \theta_k$, it is clear that:

$$h_x(s) = \beta_x h_v(s). \tag{38}$$

As such, the impulse response towards covariate will differ from the impulse response towards the innovations only by the scalar β_x. In other words, the impulse responses will be *similar* in the sense of having the same shape. For example, if $\beta_x = -1$, $h_x(s)$ will have the same shape as $h_v(s)$, just reflected around the s axis. If $\beta_x = 2$, $h_x(s)$ will have the same shape as $h_v(s)$, only scaled by a factor of 2.

B Estimated trajectories and initial innovations

When an ADL(p,q) model is fit to the data, we can use the estimated parameters, observed values of the covariate, and estimated innovations to construct estimated trajectory plots. The `irf_empirical()` function uses the OLS estimator. As such, the estimator conditions on the first p observed criterion and covariate values. The first p innovations will not be returned by the fit object because the first set of rows where there is a complete set of predictor and criterion values is at $t = p$. For example, when estimating the parameters of an AR(1) model, the first set of complete predictor-criterion pairs occurs at $t = 1$, not $t = 0$. OLS would treat y_0 as known. To construct estimated trajectory plots, we require t estimated innovations as input vector, not $t - p$ estimated innovations.

By using all the information available to the OLS estimator, we have enough information to fix the first p implied innovations uniquely. This requires knowledge of $y_0, y_1, ..., y_p$ and $x_0, x_1, ..., x_p$. These are given in the data frame. For a model with $p = 1$, we can use the equation $y_0 = \alpha + \beta_x x_0 + v_0$. y_0 and x_0 are treated as known by the estimator, and we also know the estimates $\hat{\alpha}$, $\hat{\beta_x}$. We can thus solve for $\hat{v}_0 = y_0 - \hat{\alpha} - \hat{\beta}_x x_0$ straightforwardly. In general if there are p free AR effects, we will need information on the first p values of the criterion variable and covariate, which we can be used to solve for the first p innovations. For a p'th order system, we must fix p initial innovations. Each will be of the form $\hat{v}_p = y_p - \hat{\alpha} - \hat{\phi}_1 y_{p-1} - \hat{\phi}_2 y_{p-2} - ... - \hat{\phi}_p y_0 - \hat{\beta}_x x_p - \hat{\beta}_{Lx} x_{p-1} - \hat{\beta}_{L^2x} x_{p-2} - ... - \hat{\beta}_{L^px} x_0$. The `irf_empirical()` function fixes the initial innovations automatically using the initial values supplied in the data frame.

Estimated trajectory plots can be constructed using the `irf_generator()` function directly if custom estimators are used. If estimators are used which do not condition on initial values, your own fit object should return all $p+(t-p)$ estimated innovations, and as such no fixing procedure is needed. Estimators which condition on initial values should use the same fixing procedure as above prior to calling the `irf_generator()` function.

References

Adolf, J. K., Voelkle, M. C., Brose, A., & Schmiedek, F. (2017). Capturing context-related change in emotional dynamics via fixed moderated time series analysis. *Multivariate behavioral research*, *52*(4), 499–531.

Ariens, S., Adolf, J. K., & Ceulemans, E. (2022). Collinearity issues in autoregressive models with time-varying serially dependent covariates. *Multivariate Behavioral Research*, 1–19.

Ariens, S., Adolf, J. K., & Ceulemans, E. (2023). One does not simply correct for serial dependence. *Structural Equation Modeling: A Multidisciplinary Journal*, 1–15.

Ariens, S., Ceulemans, E., & Adolf, J. K. (2020). Time series analysis of intensive longitudinal data in psychosomatic research: A methodological overview. *Journal of Psychosomatic Research*, *137*, 110191.

Asparouhov, T., Hamaker, E. L., & Muthén, B. (2018). Dynamic structural equation models. *Structural Equation Modeling: A Multidisciplinary Journal*, *25*(3), 359–388.

Asparouhov, T., & Muthén, B. (2020). Comparison of models for the analysis of intensive longitudinal data. *Structural Equation Modeling: A Multidisciplinary Journal*, *27*(2), 275–297.

Asparouhov, T., & Muthén, B. (2021). Residual structural equation models.

Barto, A. G., Sutton, R. S., & Anderson, C. W. (1983). Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE transactions on systems, man, and cybernetics*, (5), 834–846.

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*(1), 1–48. https://doi.org/10.18637/jss.v067.i01

Bilacchi, C. M., Sirius, E. V. P., Cravo, A. M., & de Azevedo Neto, R. M. (2022). Temporal dynamics of implicit memory underlying serial dependence. *Memory & Cognition*, *50*(2), 449–458.

Blanke, E. S., Brose, A., Kalokerinos, E. K., Erbas, Y., Riediger, M., & Kuppens, P. (2020). Mix it to fix it: Emotion regulation variability in daily life. *Emotion*, *20*(3), 473.

Bulteel, K., Mestdagh, M., Tuerlinckx, F., & Ceulemans, E. (2018). Var (1) based models do not always outpredict ar (1) models in typical psychological applications. *Psychological methods*, *23*(4), 740.

Bulteel, K., Tuerlinckx, F., Brose, A., & Ceulemans, E. (2018). Improved insight into and prediction of network dynamics by combining var and dimension reduction. *Multivariate behavioral research*, *53*(6), 853–875.

Ceylan, G., & Pascucci, D. (2023). Attractive and repulsive serial dependence: The role of task relevance, the passage of time, and the number of stimuli. *Journal of Vision*, *23*(6), 8–8.

Dejonckheere, E., Mestdagh, M., Verdonck, S., Lafit, G., Ceulemans, E., Bastian, B., & Kalokerinos, E. K. (2021). The relation between positive and negative affect becomes more negative in response to personally relevant events. *Emotion*, *21*(2), 326.

Driver, C. C., Oud, J. H., & Voelkle, M. C. (2017). Continuous time structural equation modeling with r package ctsem. *Journal of Statistical Software*, *77*, 1–35.

Eaton, N. R., Bringmann, L. F., Elmer, T., Fried, E. I., Forbes, M. K., Greene, A. L., Krueger, R. F., Kotov, R., McGorry, P. D., Mei, C., et al. (2023). A review of approaches and models in psychopathology conceptualization research. *Nature Reviews Psychology*, *2*(10), 622–636.

Ferrari, M., Yap, K., Scott, N., Einstein, D. A., & Ciarrochi, J. (2018). Self-compassion moderates the perfectionism and depression link in both adolescence and adulthood. *PloS one*, *13*(2), e0192022.

Gilmore, C., Cragg, L., Hogan, G., & Inglis, M. (2016). Congruency effects in dot comparison tasks: Convex hull is more important than dot area. *Journal of Cognitive Psychology*, *28*(8), 923–931.

Hamilton, J. D. (1994). *Time series analysis* (Vol. 2). Princeton University Press.

Hamker, F. H. (2004). A dynamic model of how feature cues guide spatial attention. *Vision research*, *44*(5), 501–521.

Huys, Q. J., Browning, M., Paulus, M. P., & Frank, M. J. (2021). Advances in the computational understanding of mental illness. *Neuropsychopharmacology*, *46*(1), 3–19.

Ivanov, V., & Kilian, L. (2005). A practitioner's guide to lag order selection for var impulse response analysis. *Studies in Nonlinear Dynamics & Econometrics*, *9*(1).

Koval, P., Butler, E. A., Hollenstein, T., Lanteigne, D., & Kuppens, P. (2015). Emotion regulation and the temporal dynamics of emotions: Effects of cognitive reappraisal and expressive suppression on emotional inertia. *Cognition and Emotion*, *29*(5), 831–851.

Koval, P., Pe, M. L., Meers, K., & Kuppens, P. (2013). Affect dynamics in relation to depressive symptoms: Variable, unstable or inert? *Emotion*, *13*(6), 1132.

Kuppens, P., & Verduyn, P. (2017). Emotion dynamics. *Current Opinion in Psychology*, *17*, 22–26.

Li, W. K. (2003). *Diagnostic checks in time series* (1st ed.). Chapman; Hall / CRC Monographs on Statistics; Applied Probability.

Li, Y., Wood, J., Ji, L., Chow, S.-M., & Oravecz, Z. (2022). Fitting multilevel vector autoregressive models in stan, jags, and mplus. *Structural equation modeling: a multidisciplinary journal*, *29*(3), 452–475.

Loossens, T. (2021). Toward parsimonious modeling of affect dynamics in daily life: An analysis of continuous time, non-linearities and sudden jumps.

Lueckmann, J.-M., Macke, J. H., & Nienborg, H. (2018). Can serial dependencies in choices and neural activity explain choice probabilities? *Journal of Neuroscience*, *38*(14), 3495–3506.

Lütkepohl, H. (2005). *New introduction to multiple time series analysis* (1st ed.). Springer Science & Business Media.

Maeshiro, A. (2000). An illustration of the bias of ols for yt= λ y t-1+ u t. *The Journal of Economic Education*, *31*(1), 76–80.

Mansueto, A. C., Wiers, R. W., van Weert, J., Schouten, B. C., & Epskamp, S. (2023). Investigating the feasibility of idiographic network models. *Psychological methods*, *28*(5), 1052.

McNeish, D., & Hamaker, E. L. (2020). A primer on two-level dynamic structural equation models for intensive longitudinal data in mplus. *Psychological methods*, *25*(5), 610.

Phillips, C. L., Parr, J. M., Riskin, E. A., et al. (2003). *Signals, systems, and transforms*. Prentice Hall Upper Saddle River.

Revol, J., Ariens, S., Lafit, G., Adolf, J., & Ceulemans, E. (2025). Episode-contingent experience-sampling designs for accurate estimates of autoregressive dynamics. *Psychological Methods*, *Advance online publication*.

Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2014). A computational and neural model of momentary subjective well-being. *Proceedings of the National Academy of Sciences*, *111*(33), 12252–12257.

Schuurman, N. K., Houtveen, J. H., & Hamaker, E. L. (2015). Incorporating measurement error in n= 1 psychological autoregressive modeling. *Frontiers in psychology*, *6*, 1038.

Strickland, L., Loft, S., & Heathcote, A. (2020). Investigating the effects of ongoing-task bias on prospective memory. *Quarterly journal of experimental psychology*, *73*(9), 1495–1513.

Takano, K., Stefanovic, M., Rosenkranz, T., & Ehring, T. (2021). Clustering individuals on limited features of a vector autoregressive model. *Multivariate Behavioral Research*, *56*(5), 768–786.

Tuerlinckx, F., Kuppens, P., Ariens, S., Cloos, L., Dejonckheere, E., Lafit, G., Niemeijer, K., Revol, J., Schat, E., Schreuder, M., et al. (2025). New developments in experience sampling methodology. *British Journal of Mathematical and Statistical Psychology*.

Vanhasbroeck, N., Ariens, S., Tuerlinckx, F., & Loossens, T. (2021). Computational models for affect dynamics. *Affect Dynamics*, 213–260.

Vanhasbroeck, N., Loossens, T., & Tuerlinckx, F. (2024). Two peas in a pod: Discounting models as a special case of the varmax. *Journal of Mathematical Psychology*, *120*, 102856.

Wang, X., Shao, S., Cai, Z., Ma, C., Jia, L., Blain, S. D., & Tan, Y. (2024). Reciprocal effects between negative affect and emotion regulation in daily life. *Behaviour Research and Therapy*, *176*, 104518.

Zhou, L., Wang, M., & Zhang, Z. (2021). Intensive longitudinal data analyses with dynamic structural equation modeling. *Organizational Research Methods*, *24*(2), 219–250.