

Introduction to Hierarchical Continuous Time Dynamic Modelling With `ctsem`

Charles C. Driver

Max Planck Institute for Human Development

Manuel C. Voelkle

Humboldt University Berlin

Max Planck Institute for Human Development

Abstract

`ctsem` allows for specification and fitting of a range of continuous and discrete time dynamic models with stochastic system noise. The models may include multiple indicators (dynamic factor analysis), multiple, interrelated, potentially higher order processes, and time dependent (varying within subject) and time independent (not varying within subject) covariates. Classic longitudinal models like latent growth curves and latent change score models are also possible. Version 1 of `ctsem` provided structural equation model based functionality by linking to the OpenMx software, allowing mixed effects models (random means but fixed regression and variance parameters) for multiple subjects. For version 2 of the R package **`ctsem`**, we include a Bayesian specification and fitting routine that uses the **Stan** probabilistic programming language, via the **rstan** package in R. This allows for all parameters of the dynamic model to individually vary, using an estimated population mean and variance, and any time independent covariate effects, as a prior. Frequentist approaches with `ctsem` are documented in a forthcoming JSS publication (Driver, Voelkle, Oud, in press), and in R vignette form at <https://cran.r-project.org/package=ctsem/vignettes/ctsem.pdf>. The Bayesian approach is discussed more fully and conceptually in a preprint article available at https://www.researchgate.net/publication/310747801_Hierarchical_Bayesian_Continuous_Time_Dynamic_Modeling, but here we provide more specifics on the software for getting started with the Bayesian approach included in version 2.

Keywords: hierarchical time series, Bayesian, longitudinal, panel data, state space, structural equation, continuous time, stochastic differential equation, dynamic models, Kalman filter, R.

1. Overview

1.1. Subject Level Latent Dynamic model

This section describes the fundamental subject level model, and where appropriate, the name of the `ctModel` argument used to specify specific matrices. The description of the full model, including subject level likelihood and population model, is provided at the end of this document. Although we do not describe it explicitly, the corresponding discrete time autoregressive / moving average models can be specified and use the same set of parameter matrices we describe, although the meaning is of course somewhat different.

1.2. Subject level latent dynamic model

The subject level dynamics are described by the following stochastic differential equation:

$$d\boldsymbol{\eta}(t) = \left(\mathbf{A}\boldsymbol{\eta}(t) + \mathbf{b} + \mathbf{M}\boldsymbol{\chi}(t) \right) dt + \mathbf{G}d\mathbf{W}(t) \quad (1)$$

Vector $\boldsymbol{\eta}(t) \in \mathbb{R}^v$ represents the state of the latent processes at time t . The matrix $\mathbf{A} \in \mathbb{R}^{v \times v}$ (DRIFT) represents the so-called drift matrix, with auto effects on the diagonal and cross effects on the off-diagonals characterizing the temporal dynamics of the processes.

The continuous time intercept vector $\mathbf{b} \in \mathbb{R}^v$ (CINT), in combination with \mathbf{A} , determines the long-term level at which the processes fluctuate around.

Time dependent predictors $\boldsymbol{\chi}(t)$ represent inputs to the system that vary over time and are independent of fluctuations in the system. Equation 4 shows a generalized form for time dependent predictors, that could be treated a variety of ways dependent on the assumed time course (or shape) of time dependent predictors. We use a simple impulse form shown in Equation 5, in which the predictors are treated as impacting the processes only at the instant of an observation occasion u . When necessary, the evolution over time can be modeled by extending the state matrices, for an example see [Driver, Oud, and Voelkle \(In Press\)](#).

$$\boldsymbol{\chi}(t) = \sum_{u \in \mathbf{U}} \mathbf{x}_u \delta(t - t_u) \quad (2)$$

Here, time dependent predictors $\mathbf{x}_u \in \mathbb{R}^l$ (tdpreds) are observed at measurement occasions $u \in \mathbf{U}$. The Dirac delta function $\delta(t - t_u)$ is a generalized function that is ∞ at 0 and 0 elsewhere, yet has an integral of 1 (when 0 is in the range of integration). It is useful to model an impulse to a system, and here is scaled by the vector of time dependent predictors \mathbf{x}_u . The effect of these impulses on processes $\boldsymbol{\eta}(t)$ is then $\mathbf{M} \in \mathbb{R}^{v \times l}$ (TDPREDEFFECT).

$\mathbf{W}(t) \in \mathbb{R}^v$ (DIFFUSION) represents independent Wiener processes, with a Wiener process being a random-walk in continuous time. $d\mathbf{W}(t)$ is meaningful in the context of stochastic differential equations, and represents the stochastic error term, an infinitesimally small increment of the Wiener process. Lower triangular matrix $\mathbf{G} \in \mathbb{R}^{v \times v}$ represents the effect of this noise on the change in $\boldsymbol{\eta}(t)$. \mathbf{Q} , where $\mathbf{Q} = \mathbf{G}\mathbf{G}^\top$, represents the variance-covariance matrix of the diffusion process in continuous time.

1.3. Subject level measurement model

The latent process vector $\boldsymbol{\eta}(t)$ has measurement model:

$$\mathbf{y}(t) = \mathbf{\Lambda}\boldsymbol{\eta}(t) + \boldsymbol{\tau} + \boldsymbol{\epsilon}(t) \quad \text{where } \boldsymbol{\epsilon}(t) \sim \mathbf{N}(\mathbf{0}_c, \boldsymbol{\Theta}) \quad (3)$$

$\mathbf{y}(t) \in \mathbb{R}^c$ is the vector of manifest variables, $\mathbf{\Lambda} \in \mathbb{R}^{c \times v}$ (LAMBDA) represents the factor loadings, and $\boldsymbol{\tau} \in \mathbb{R}^c$ (MANIFESTMEANS) the manifest intercepts. The residual vector $\boldsymbol{\epsilon} \in \mathbb{R}^c$ has covariance matrix $\boldsymbol{\Theta} \in \mathbb{R}^{c \times c}$ (MANIFESTVAR).

1.4. Overview of hierarchical model

Parameters for each subject are first drawn from a simultaneously estimated higher level distribution over an unconstrained space, then a set of parameter specific transformations

are applied so that a) each parameter conforms to necessary bounds and b) is subject to the desired prior, then a range of matrix transformations are applied to generate the continuous time matrices described, as well as all relevant discrete time instantiations (More variability in measurement time intervals thus means more computations). The higher level distribution has a multivariate normal prior. A more comprehensive description is found at the end of this document.

1.5. Install software and prepare data

Install ctsem software:

```
install.packages("ctsem")
```

Prepare data in long format, each row containing one time point of data for one subject. We need a subject id column containing numbers from 1 to total subjects, rising incrementally with each subject going down the data structure. This is to ensure coherence with the internal structure of the Stan model. The column is named by default "id", though this can be changed in the model specification. We also need a time column "time", containing numeric values for time, columns for manifest variables (the names of which must be given in the next step using ctModel), columns for time dependent predictors (these vary over time but have no model estimated and are assumed to impact latent processes instantly), and columns for time independent predictors (which predict the subject level parameters, that are themselves time invariant – thus the values for a particular time independent predictor must be the same across all observations of a particular subject).

	id	time	Y1	Y2	TD1	TI1	TI2	TI3
[1,]	1	-0.05	-5.62	-11.52	0.471	0.529	-0.992	0.837
[2,]	1	1.25	-4.21	-10.39	0.824	0.529	-0.992	0.837
[3,]	1	2.10	-4.65	-10.36	0.124	0.529	-0.992	0.837
[4,]	1	3.03	-4.35	-9.92	-1.115	0.529	-0.992	0.837
[5,]	1	3.98	-5.24	-10.53	-0.063	0.529	-0.992	0.837
[6,]	2	36.06	-1.01	-17.12	0.377	1.138	0.323	0.147
[7,]	2	36.88	-2.10	-17.65	-0.159	1.138	0.323	0.147
[8,]	2	37.87	-2.53	-17.69	0.907	1.138	0.323	0.147

At present, missingness is fine on manifest indicators, but not allowed elsewhere.

1.6. Model specification

Specify model using `ctModel(type="stanct", ...)`. "stanct" specifies a continuous time model in Stan format, "standt" specifies discrete time, while "omx" is the classic **ctsem** behaviour and prepares an **OpenMx** model. Other arguments to ctModel proceed as normal, although some matrices used for type 'omx' are not relevant for the Stan formats, either because the between subject matrices have been removed, or because time dependent and independent predictors are now treated as fixed regressors and only require effect (or design) matrices. These differences are documented in the help for ctModel.

```
model<-ctModel(type='stanct',
  n.latent=2, latentNames=c('eta1','eta2'),
```

Argument	Sign	Default	Meaning
n.manifest	c		Number of manifest indicators per individual at each measurement occasion.
n.latent	v		Number of latent processes.
LAMBDA	Λ		$n.manifest \times n.latent$ loading matrix relating latent to manifest variables.
manifestNames		Y1, Y2, etc	$n.manifest$ length character vector of manifest names.
latentNames		eta1, eta2, etc	$n.latent$ length character vector of latent names.
T0VAR	Q_1^*	free	lower tri $n.latent \times n.latent$ matrix of latent process initial covariance, specified with standard deviations on diagonal and (partial) correlations on lower triangle.
T0MEANS	η_1	free	$n.latent \times 1$ matrix of latent process means at first time point, T0.
MANIFESTMEANS	τ	free	$n.manifest \times 1$ matrix of manifest means.
MANIFESTVAR	Θ	free diag	lower triangular matrix of var / cov between manifests, specified with standard deviations on diagonal and (partial) correlations on lower triangle.
DRIFT	A	free	$n.latent \times n.latent$ matrix of continuous auto and cross effects.
CINT	b	0	$n.latent \times 1$ matrix of continuous intercepts.
DIFFUSION	Q	free	lower triangular $n.latent \times n.latent$ matrix containing standard deviations of latent process on diagonal, and (partial) correlations on lower off-diagonals.
n.TDpred	l	0	Number of time dependent predictors in the dataset.
TDpredNames		TD1, TD2, etc	$n.TDpred$ length character vector of time dependent predictor names.
TDPREDEFFECT	M	free	$n.latent \times n.TDpred$ matrix of effects from time dependent predictors to latent processes.
n.TIpred	p	0	Number of time independent predictors.
TIpredNames		TI1, TI2, etc	$n.TIpred$ length character vector of time independent predictor names.

```

n.manifest=2, manifestNames=c('Y1','Y2'),
n.TDpred=1, TDpredNames='TD1',
n.TIpred=3, TIpredNames=c('TI1','TI2','TI3'),
LAMBDA=diag(2))

```

This generates a first order bivariate latent process model, with each process measured by a single, potentially noisy, manifest variable. A single time dependent predictor is included in the model, and three time independent predictors. Additional complexity or restrictions may be added, the table below shows the basic arguments one may consider and their link to the dynamic model parameters. For more details see the `ctsem` help files or papers. Note that for the Stan implementation, `ctModel` requires variance covariance matrices (`DIFFUSION`, `T0VAR`, `MANIFESTVAR`) to be specified with standard deviations on the diagonal, correlations (partial, if > 2 latent processes) the lower off diagonal, and zeroes on the upper off diagonal.

These matrices may all be specified using a combination of character strings to name free parameters, or numeric values to represent fixed parameters.

The parameters subobject of the created model object shows the parameter specification that will go into Stan, including both fixed and free parameters, whether the parameters vary across individuals, how the parameter is transformed from a standard normal distribution (thus setting both priors and bounds), and whether that parameter is regressed on the time

independent predictors.

```
head(model$pars,8)
```

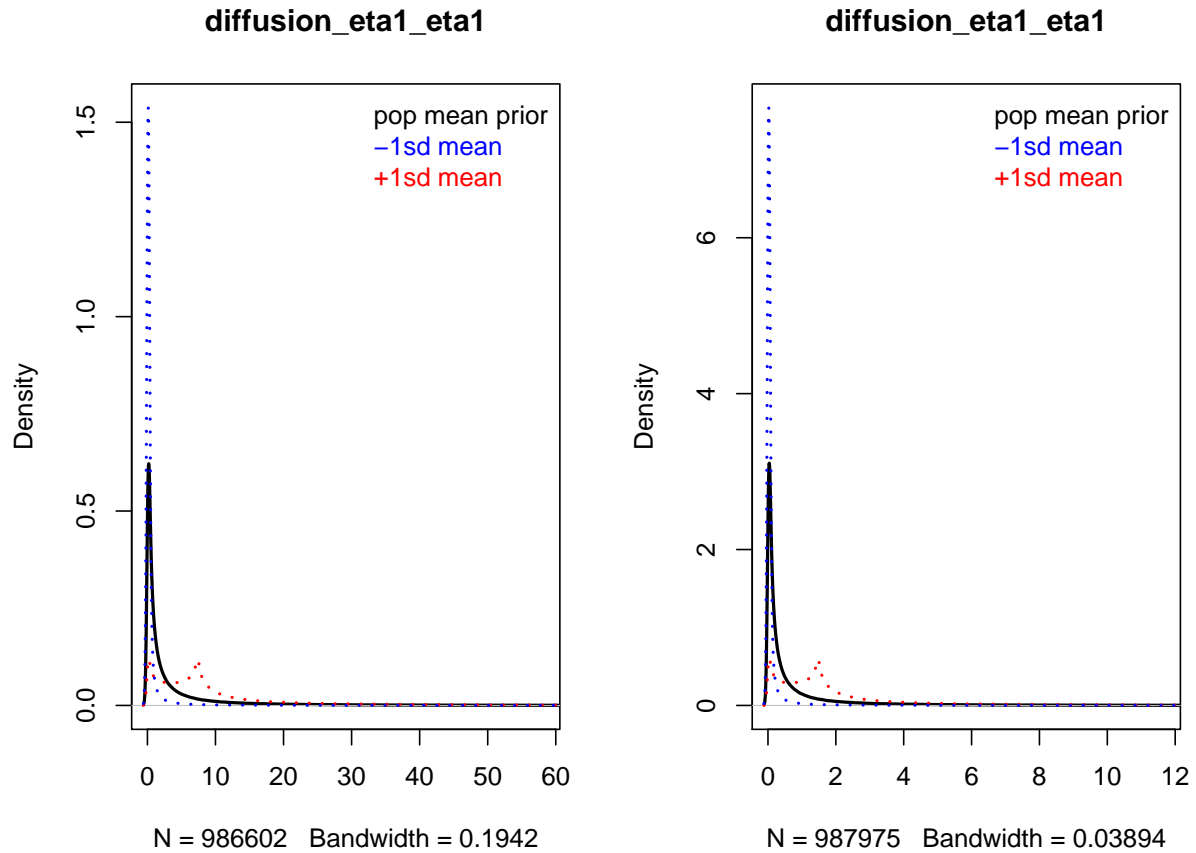
	matrix	row	col	param	value	transform	indvarying	sdscale
1	TOMEANS	1	1	T0mean_eta1	NA	(param) * 10	TRUE	1
2	TOMEANS	2	1	T0mean_eta2	NA	(param) * 10	TRUE	1
3	LAMBDA	1	1	<NA>	1	<NA>	FALSE	1
4	LAMBDA	1	2	<NA>	0	<NA>	FALSE	1
5	LAMBDA	2	1	<NA>	0	<NA>	FALSE	1
6	LAMBDA	2	2	<NA>	1	<NA>	FALSE	1
7	DRIFT	1	1	drift_eta1_eta1	NA	-log(exp(-param*1.5)+1)-.00001	TRUE	1
8	DRIFT	1	2	drift_eta1_eta2	NA	(param)*.5	TRUE	1
	TI1_effect	TI2_effect	TI3_effect					
1	TRUE	TRUE	TRUE					
2	TRUE	TRUE	TRUE					
3	FALSE	FALSE	FALSE					
4	FALSE	FALSE	FALSE					
5	FALSE	FALSE	FALSE					
6	FALSE	FALSE	FALSE					
7	TRUE	TRUE	TRUE					
8	TRUE	TRUE	TRUE					

One may modify the output model to either restrict between subject differences (set some parameters to not vary over individuals), alter the transformation used to determine the prior / bounds, or restrict which effects of time independent predictors to estimate. Plotting the original prior, making a change, and plotting the resulting prior, are shown here – in this case we believe the stochastic latent process innovation for our first latent process, captured by row 1 and column 1 of the DIFFUSION matrix, to be small, so scale our prior accordingly to both speed and improve sampling. Rather than simply scaling by 0.2 as shown here, one could also construct a new form of prior, so long as the resulting distribution was within the bounds required for the specific parameter. Note that the resulting distribution is a result of applying the specified transformation to a standard normal distribution, with mean of 0 and standard deviation of 1. To change the underlying standard normal, one would need to edit the resulting Stan code directly.

```
par(mfrow=c(1,2))
plot(model,rows=11)
print(model$pars$transform[11])

[1] "exp(param*2) +.00001"

model$pars$transform[11]<- "(exp(param*2) +.0001)*.2"
plot(model,rows=11)
```



The plots show the prior distribution for the population mean of DIFFUSION[1,1] in black, as well as two possible priors for the subject level parameters. The blue prior results from assuming the population mean is one standard deviation lower than the mean of the prior, and marginalising over the prior for the standard deviation of the population distribution. This latter prior can be scaled using the `sdscale` column of the parameters subobject, but is by default a truncated normal distribution with mean 0 and SD 1.

Restrict between subject effects as desired. Unnecessary between subject effects will slow sampling and hinder appropriate regularization, but be aware of the many parameter dependencies in these models – restricting one parameter may lead to genuine variation in the restricted parameter expressing itself elsewhere. The prior scale for between subject variance may need to be restricted when limited data (in either the number of time points or number of subjects) is available, to ensure adequate regularisation. Here we restrict MANIFESTVAR effects between subjects, and set all prior scales for the standard deviation of the population distribution to 0.2, from the default of 1.0. A rough interpretation of this change in `sdscale` is simply that we expect lower values for the population standard deviation, but to better interpret the effect of this latter change, see the section on standard deviation transformations.

```
model$pars[c(15,18),]$indvarying<-FALSE
model$pars$sdscale[1:28] <- .5
```

Also restrict which parameters to include time independent predictor effects for in a similar way, for similar reasons. In this case, the only adverse effects of restriction are that the relationship between the predictor and variables will not be estimated, but the subject level

parameters themselves should not be very different, as they are still freely estimated. Note that such effects are only estimated for individually varying parameters anyway – so after the above change there is no need to set the `tipreffect` to `FALSE` for the `T0VAR` variables, it is assumed. Instead, we restrict the `tipreffects` on all parameters, and free them only for the manifest intercept parameters.

```
model$pars[,c('TI1_effect', 'TI2_effect', 'TI3_effect')]<-FALSE
model$pars[c(19,20),c('TI1_effect', 'TI2_effect', 'TI3_effect')]<-TRUE
```

1.7. Model fitting

Once model specification is complete, the model is fit to the data using the `ctStanFit` function as follows – depending on the data, model, and number of iterations requested, this can take anywhere from a few minutes to days. Current experience suggests 300 iterations is often enough to get an idea of what is going on, but more may be necessary for robust inference. This will of course vary massively depending on model and data. For the sake of speed for this example we only sample for 200 iterations with a lowered `max_treedepth` - this latter parameter controls the maximum number of steps the Hamiltonian sampler is allowed to take per iteration, with each increase of 1 doubling the maximum. With these settings the fit should take only a few minutes (but will not be adequate for inference!). Those that wish to try out the functions without waiting, can simply use the already existing `ctstantestfit` object instead of creating the fit object (and adjust the code in following sections as needed!).

```
fit<-ctStanFit(datalong = ctstantestdat, ctstanmodel = model, iter=200,
  chains=2, plot=FALSE, control=list(max_treedepth = 6))
```

```
In file included from C:/Users/driver/Documents/R/win-library/3.3/BH/include/boost/config.hpp:39:0,
                 from C:/Users/driver/Documents/R/win-library/3.3/BH/include/boost/math/tools/confi
                 from C:/Users/driver/Documents/R/win-library/3.3/StanHeaders/include/stan/math/rev
                 from C:/Users/driver/Documents/R/win-library/3.3/StanHeaders/include/stan/math/rev
                 from C:/Users/driver/Documents/R/win-library/3.3/StanHeaders/include/stan/math/rev
                 from C:/Users/driver/Documents/R/win-library/3.3/StanHeaders/include/stan/math/rev
                 from C:/Users/driver/Documents/R/win-library/3.3/StanHeaders/include/stan/math.hpp
                 from C:/Users/driver/Documents/R/win-library/3.3/StanHeaders/include/src/stan/mode
                 from file6ac4a746f09.cpp:8:
C:/Users/driver/Documents/R/win-library/3.3/BH/include/boost/config/compiler/gcc.hpp:186:0: warning
#  define BOOST_NO_CXX11_RVALUE_REFERENCES
^
<command-line>:0:0: note: this is the location of the previous definition
```

The `plot` argument allows for plotting of sampling chains in real time, which is useful for slow models to ensure that sampling is proceeding in a functional manner. Models with many parameters (e.g., many subjects and all parameters varying over subject) may be too taxing for the function to handle smoothly - we have had success with up to around 4000 parameters.

1.8. Summary

After fitting, the summary function may be used on the fit object, which returns details

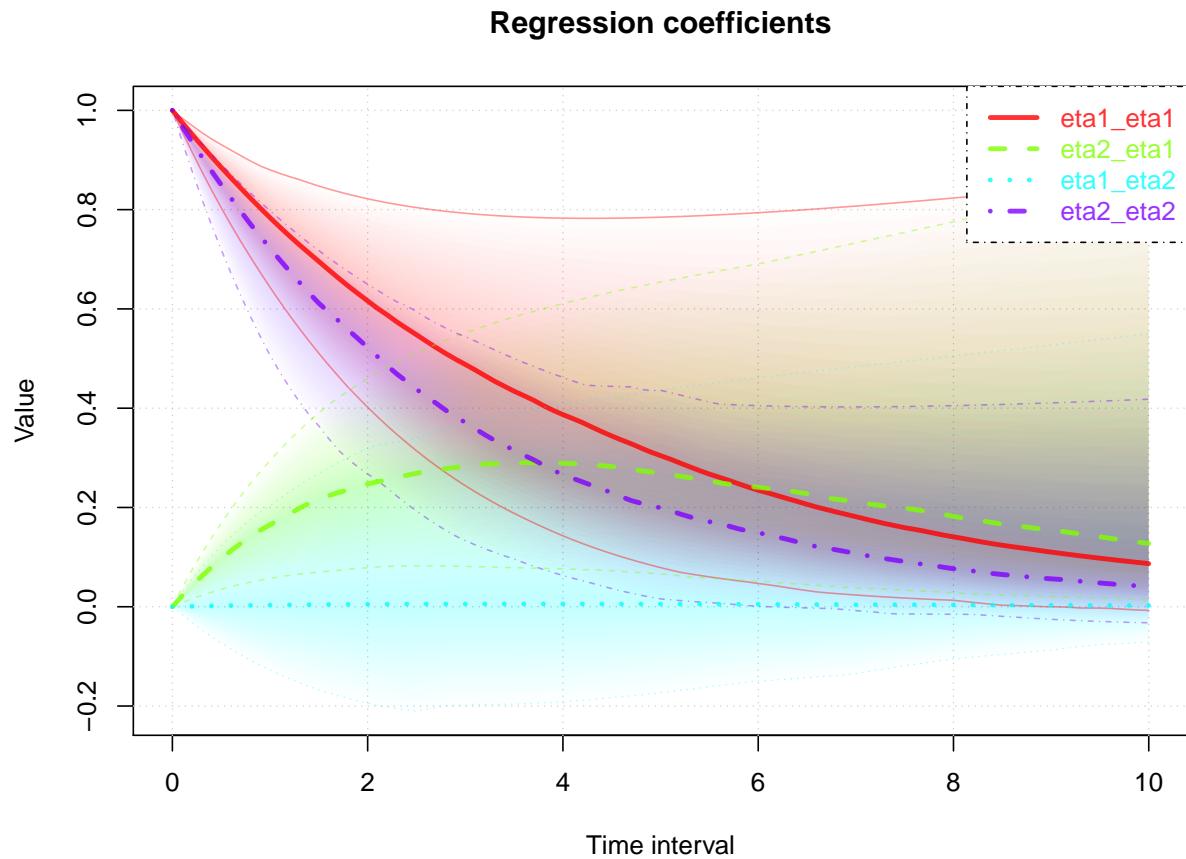
regarding the population mean parameters, population standard deviation parameters, population correlations, and the effect parameters of time independent predictors. Additionally, summary outputs a range of matrices regarding correlations between subject level parameters. `hypercorr_means` reports the posterior mean of the correlation between raw (not yet transformed from the standard normal scale) parameters. `hypercorr_sd` reports the standard deviation of these parameters. `hypercovcor_transformedmean` reports the correlation between transformed parameters on the lower triangle, the variance of these parameters on the diagonal, and the covariance on the upper triangle. To view the posterior median of the continuous time parameter matrices, the `ctStanContinuousPars` function can be used.

```
summary(fit)
```

In the summary output, the population parameters are returned in the same form that they are input to `ctModel` - that is, standard deviations and partial correlations, which can be more difficult to interpret when more than two processes are modeled. To return the full set of continuous time parameter matrices for a specified subject or collection of subjects, the `ctStanContinuousPars` function can be used, with the `calcfunc` argument set appropriately depending on whether one wants the median, mean, sd, or a specific quantile (which would require changing the `probs` argument also).

1.9. Plotting

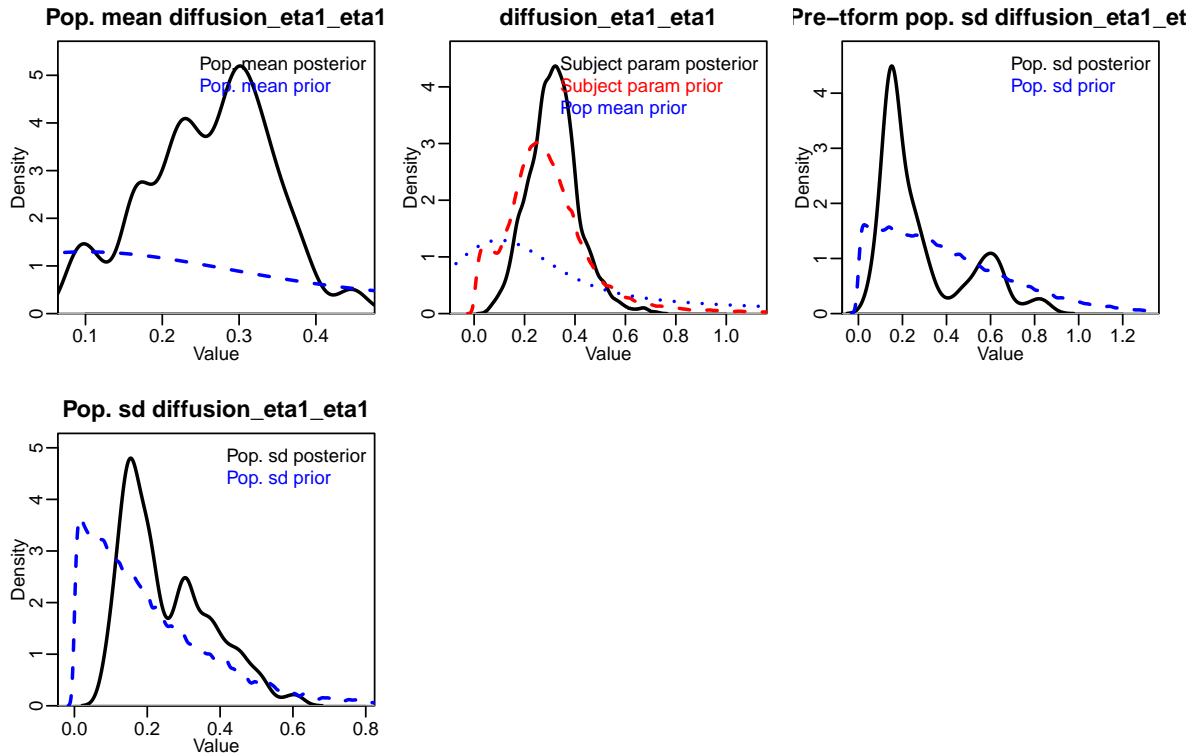
The plot function outputs a sequence of plots, all generated by specific functions. The name of the specific function appears in a message in the R console, checking the help for each specific function and running them separately will allow more customization. Some of the plots, such as the trace, density, and interval, are generated by the relevant rstan function and hopefully self explanatory. The plots specific to the heirarchical continuous time dynamic model are as follows:



The above plot shows the dynamic regression coefficients (between latent states at different time points) that are implied by the model for particular time intervals, as well as the uncertainty of these coefficients.

The relation between posteriors and priors for variables of interest can also be plotted as follows:

```
ctStanPlotPost(ctstanfitobj = fit, rows=11)
```



Shown are approximate density plots based on the post-warmup samples drawn. For each parameter four plots are shown – the first displays the posterior distribution of subject level parameters, the subject level prior (generated from repeated sampling of the hyper parameters), and the prior for the population mean.

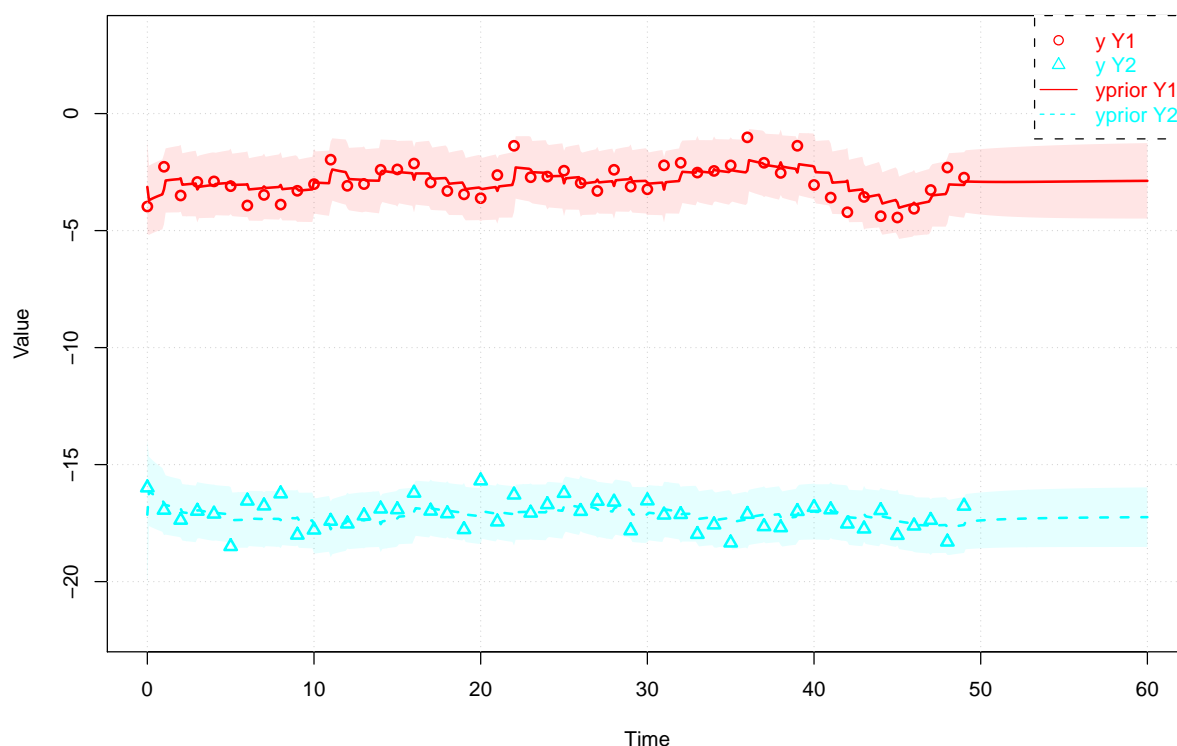
1.10. Stationarity

When it is reasonable to assume that the prior for long term expectation and variance of the latent states is the same as (or very similar to) the prior for initial expectations and variances, setting the argument `stationary=TRUE` to `ctStanFit` can reduce uncertainty and aid in fitting. This argument ignores any `T0VAR` and `T0MEANS` matrices in the input, instead replacing them with asymptotic expectations based on the `DRIFT`, `DIFFUSION`, and `CINT` matrices. Alternatively, a prior can be placed on the stationarity of the dynamic models, calculated as the difference between the `T0MEANS` and the long run asymptotes of the expected value of the process, as well as the difference between the diagonals of the `T0VAR` covariance matrix and the long run asymptotes of the covariance of the processes. Such a prior encourages a minimisation of these differences, and can help to ensure that sensible, non-explosive models are estimated, and also help the sampler get past difficult regions of relative flatness in the parameter space due to colinearities between the within and between subject parameters. However if such a prior is too strong it can also induce difficult dependencies in model parameters, and there are a range of models where one may not wish to have such a prior. To place such a prior, the `model$stationarymeanprior` and `model$stationaryvarprior` slots can be changed from the default of `NA` to a numeric vector, representing the normal standard deviation of the deviations from stationarity. The number of elements in the vector correspond to the number of latent processes.

1.11. Individual level analyses

Individual level results can also be considered, as `ctsem` includes functionality to output prior (based on all prior observations), updated (based on all prior and current observations), and smoothed (based on all observations) expectations and covariances from the Kalman filter, based on specific subjects models. For ease of comparison, expected manifest indicator scores conditional on prior, updated and smoothed states are also included. This approach allows for: predictions regarding individuals states at any point in time, given any values on the time dependent predictors (external inputs such as interventions or events); residual analysis to check for unmodeled dependencies in the data; or simply as a means of visualization, for comprehension and model sanity checking purposes. An example of such is depicted below, where we see observed and estimated scores for a selected subject from our sample. If we wanted to predict unobserved states in the future, we would need only to specify the appropriate timerange (Prediction into earlier times is possible but makes little sense unless the model is restricted to stationarity).

```
ctStanKalman(fit, subjects=2, timerange=c(0,60), timestep=.1, plot=TRUE)
```



1.12. Accessing Stan model code

For diagnosing problems or modifying the model in ways not achievable via the `ctsem` model specification, one can use `ctsem` to generate the Stan code and then work directly with that, simply by specifying the argument `fit=FALSE` to the `ctStanFit` function. Any altered code can be passed back into `ctStanFit` by using the `stanmodeltext` argument, which can be convenient for setting up the data in particular.

1.13. Using Rstan functions

The standard rstan output functions such as `summary` and `extract` are also available, and the `shinystan` package provides an excellent browser based interface. The stan fit object is stored under the `$stanfit` subobject from the `ctStanFit` output. The parameters which are likely to be of most interest in the output are prefixed by `"hmean_"` for hyper (population) mean, `"hsd_"` for hyper standard deviation, and `"tipred_"` for time independent predictor. Any hmean parameters are returned in the form used for input - so correlations and standard deviations for any of the covariance related parameters. Subject specific parameters are denoted by the matrix they are from, then the first index represents the subject id, followed by standard matrix notation. For example, the 2nd row and 1st column of the DRIFT matrix for subject 8 is `"DRIFT[8,2,1]"`. Parameters in such matrices are returned in the form used for internal calculations - that is, variance covariance matrices are returned as such, rather than the lower-triangular standard deviation and correlation matrices required for input. The exception to this are the time independent predictor effects, prefixed with `"tipred_"`, for which a linear effect of a change of 1 on the predictor is approximated. So although `"tipred_TI1"` is only truly linear with respect to internal parameterisations, we approximate the linear effect by averaging the effect of a score of +1 or -1 on the predictor, on the population mean. For any subject that substantially differs from the mean, or simply when precise absolute values of the effects are required (as opposed to general directions), they will need to be calculated manually.

1.14. Oscillating, single subject example - sunspots data

In the following example we fit the sunspots data available within R, which has previously been fit by various authors including [Tómasson \(2013\)](#). We have used the same CARMA(2,1) model and obtained similar estimates - some differences are due to the contrast between Bayes and maximum likelihood, though if desired one could adjust the code to fit using maximum likelihood, as here we have only one subject. There are usually some divergent transitions (indicating a difficulty in the sampling chain and a potential threat to inference) generated in this fit - alternate parameterisations or an increase in the `adapt_delta` control argument to Stan (which defaults to 0.9 in `ctsem`, with a maximum of 1, though 1 is not recommended...) may help.

```
#get data
sunspots<-sunspot.year
sunspots<-sunspots[50: (length(sunspots) - (1988-1924))]
id <- 1
time <- 1749:1924
datalong <- cbind(id, time, sunspots)

#setup model
model <- ctModel(type='stanct', n.latent=2, n.manifest=1,
  manifestNames='sunspots',
  latentNames=c('ss_level', 'ss_velocity'),
  LAMBDA=matrix(c( 1, 'ma1' ), nrow=1, ncol=2),
  DRIFT=matrix(c(0, 1, 'a21', 'a22'), nrow=2, ncol=2, byrow=TRUE),
  MANIFESTMEANS=matrix(c('m1'), nrow=1, ncol=1),
  # MANIFESTVAR=matrix(0, nrow=1, ncol=1),
  CINT=matrix(c(0, 0), nrow=2, ncol=1),
  DIFFUSION=matrix(c(
```

1.15. Population standard deviations - understanding the transforms

Internally, we sample parameters that we will refer to here as the ‘raw’ parameters – these parameters have no bounds and are drawn from normal distributions. Both population mean (internally: `hypermeans`) and subject level (internally: `indparamsbase`) raw parameters are drawn from a $\text{normal}(0, 1)$ distribution. Depending on the specific parameter, various transformations may be applied to set appropriate bounds and priors. The population standard deviation (`hypersd`) for these raw parameters is sampled (by default) from a truncated $\text{normal}(0, 1)$ distribution. This distribution can be scaled on a per parameter basis by the `sdscale` multiplier in the model specification, which defaults to 1. The following script shows a didactic sequence of sampling and transformation for a model with a single parameter, the `auto`

effect of the drift matrix, and 50 subjects. Although we sample the priors here, this is merely to reflect the prior and enable understanding and plotting.

```
#population mean and subject level deviations (pre-transformation)

hypermeans_prior <- rnorm(99999, 0, 1)
hypermeans_post <- -2 #hypothetical sample

indparamsbase_prior <- rnorm(99999, 0, 1)
indparamsbase_post <- rnorm(50, 0, 1) #hypothetical sample

#population standard deviation prior

hypersd_prior <- rnorm(99999, 0, 1)
hypersd_prior <- hypersd_prior[hypersd_prior > 0]

#population standard deviation posterior

hypersd_post <- .4 #hypothetical

#population cholesky correlation matrix
#lower triangle sampled from uniform(-1, 1),
#upper triangle fixed to 0,
#diagonal calculated according to hypersd.
hypercorrchol_post <- 1 #because only 1 parameter here...

#population cholesky covariance matrix
#here based on mean of hypersd_post, for convenience...
#in reality would have multiple samples.
hypercovchol <- diag(hypercorrchol_post,1) %%%
  diag(hypersd_post,1) %%% diag(hypercorrchol_post,1)

#subject level parameters
#first compute pre transformation parameters
#then transform appropriately (here according to drift auto effect)
indparams <- hypercovchol %%% indparamsbase_post + hypermeans_post
indparams <- -log(exp(-1.5 * indparams) + 1)

#post transformation population standard deviation
hsd_ourparameter <- abs( #via delta approximation
  (-log(exp(-1.5 * (hypermeans_post + hypersd_post)) + 1) -
    -log(exp(-1.5 * (hypermeans_post - hypersd_post)) + 1) ) / 2)
```

2. The model

There are three main elements to our hierarchical continuous time dynamic model. There is a subject level latent dynamic model, a subject level measurement model, and a population level model for the subject level parameters. Note that while various elements in the model depend on time, the fundamental parameters of the model as described here are time-invariant. Note also that we ignore subject specific subscripts when discussing the subject level model.

2.1. Subject level latent dynamic model

The subject level dynamics are described by the following stochastic differential equation:

$$d\boldsymbol{\eta}(t) = \left(\mathbf{A}\boldsymbol{\eta}(t) + \mathbf{b} + \mathbf{M}\boldsymbol{\chi}(t) \right) dt + \mathbf{G}d\mathbf{W}(t) \quad (4)$$

Vector $\boldsymbol{\eta}(t) \in \mathbb{R}^v$ represents the state of the latent processes at time t . The matrix $\mathbf{A} \in \mathbb{R}^{v \times v}$ represents the so-called drift matrix, with auto effects on the diagonal and cross effects on the off-diagonals characterizing the temporal dynamics of the processes.

The continuous time intercept vector $\mathbf{b} \in \mathbb{R}^v$, in combination with \mathbf{A} , determines the long-term level at which the processes fluctuate around.

Time dependent predictors $\boldsymbol{\chi}(t)$ represent inputs to the system that vary over time and are independent of fluctuations in the system. Equation 4 shows a generalized form for time dependent predictors, that could be treated a variety of ways dependent on the assumed time course (or shape) of time dependent predictors. We use a simple impulse form shown in Equation 5, in which the predictors are treated as impacting the processes only at the instant of an observation occasion u . When necessary, the evolution over time can be modeled by extending the state matrices, for an example see [Driver et al. \(In Press\)](#).

$$\boldsymbol{\chi}(t) = \sum_{u \in \mathbf{U}} \mathbf{x}_u \delta(t - t_u) \quad (5)$$

Here, time dependent predictors $\mathbf{x}_u \in \mathbb{R}^l$ are observed at measurement occasions $u \in \mathbf{U}$, where \mathbf{U} is the set of measurement occasions from 1 to the number of measurement occasions, with $u = 1$ treated as occurring at $t = 0$. The Dirac delta function $\delta(t - t_u)$ is a generalized function that is ∞ at 0 and 0 elsewhere, yet has an integral of 1 (when 0 is in the range of integration). It is useful to model an impulse to a system, and here is scaled by the vector of time dependent predictors \mathbf{x}_u . The effect of these impulses on processes $\boldsymbol{\eta}(t)$ is then $\mathbf{M} \in \mathbb{R}^{v \times l}$.

$\mathbf{W}(t) \in \mathbb{R}^v$ represents independent Wiener processes, with a Wiener process being a random-walk in continuous time. $d\mathbf{W}(t)$ is meaningful in the context of stochastic differential equations, and represents the stochastic error term, an infinitesimally small increment of the Wiener process. Lower triangular matrix $\mathbf{G} \in \mathbb{R}^{v \times v}$ represents the effect of this noise on the change in $\boldsymbol{\eta}(t)$. \mathbf{Q} , where $\mathbf{Q} = \mathbf{G}\mathbf{G}^\top$, represents the variance-covariance matrix of the diffusion process in continuous time.

2.2. Subject level dynamic model — discrete time solution

The stochastic differential Equation 4 may be solved and translated to a discrete time representation, for any observation $u \in \mathbf{U}$:

$$\boldsymbol{\eta}_u = \mathbf{A}_u^* \boldsymbol{\eta}_{u-1} + \mathbf{b}_u^* + \mathbf{M}\mathbf{x}_u + \boldsymbol{\zeta}_u^* \quad \boldsymbol{\zeta}_u^* \sim \mathbf{N}(0_v, \mathbf{Q}_u^*) \quad (6)$$

The $*$ notation is used to indicate a term that is the discrete time equivalent of the original. \mathbf{A}_u^* then contains the appropriate auto and cross regressions for the effect of latent processes $\boldsymbol{\eta}$ at measurement occasion $u - 1$ on $\boldsymbol{\eta}$ at measurement occasion u . \mathbf{b}_u^* represents the discrete time intercept for measurement occasion u . Since \mathbf{M} is conceptualized as the effect of instantaneous impulses \mathbf{x} , (which only occur at occasions \mathbf{U} and are not continuously present as for

the processes $\boldsymbol{\eta}$), the discrete and continuous time forms are equivalent. $\boldsymbol{\zeta}_u$ is the zero mean random error term for the processes at occasion u . \mathbf{Q}_u^* represents the multivariate normal disturbance at occasion u . The recursive nature of the solution means that at the first measurement occasion $u = 1$, the system must be initialized in some way, with $\mathbf{A}_u^* \boldsymbol{\eta}_{u-1}$ replaced by $\boldsymbol{\eta}_1$, and \mathbf{Q}_u^* replaced by \mathbf{Q}_1^* . These initial states and covariances are later referred to as TOMEANS and T0VAR respectively.

Unlike in a purely discrete time model, where the various effect matrices described above would be unchanging, in a continuous time model the discrete time matrices all depend on some function of the continuous time parameters and the time interval between observations u and $u - 1$, these functions look as follows:

$$\mathbf{A}_u^* = e^{\mathbf{A}(t_u - t_{u-1})} \quad (7)$$

$$\mathbf{b}_u^* = \mathbf{A}^{-1}(\mathbf{A}_u^* - \mathbf{I})\mathbf{b} \quad (8)$$

$$\mathbf{Q}_u^* = \mathbf{Q}_\infty - \mathbf{A}_u^* \mathbf{Q}_\infty (\mathbf{A}_u^*)^\top \quad (9)$$

Where $\mathbf{A}_\# = \mathbf{A} \otimes \mathbf{I} + \mathbf{I} \otimes \mathbf{A}$, with \otimes denoting the Kronecker-product, the asymptotic diffusion $\mathbf{Q}_\infty = \text{irow}(-\mathbf{A}_\#^{-1} \text{row}(\mathbf{Q}))$, row is an operation that takes elements of a matrix row wise and puts them in a column vector, and irow is the inverse of the row operation. The covariance update shown in Equation 9 has not been described in the psychological literature so far as we are aware, but is a more computationally efficient form used in [Tómasson \(2013\)](#).

2.3. Subject level measurement model

The latent process vector $\boldsymbol{\eta}(t)$ has measurement model:

$$\mathbf{y}(t) = \boldsymbol{\Lambda} \boldsymbol{\eta}(t) + \boldsymbol{\tau} + \boldsymbol{\epsilon}(t) \quad \text{where } \boldsymbol{\epsilon}(t) \sim \mathcal{N}(\mathbf{0}_c, \boldsymbol{\Theta}) \quad (10)$$

$\mathbf{y}(t) \in \mathbb{R}^c$ is the vector of manifest variables, $\boldsymbol{\Lambda} \in \mathbb{R}^{c \times v}$ represents the factor loadings, and $\boldsymbol{\tau} \in \mathbb{R}^c$ the manifest intercepts. The residual vector $\boldsymbol{\epsilon} \in \mathbb{R}^c$ has covariance matrix $\boldsymbol{\Phi} \in \mathbb{R}^{c \times c}$.

2.4. Subject level likelihood

The subject level likelihood, conditional on time dependent predictors \mathbf{x} and subject level parameters $\boldsymbol{\Phi}$, is as follows:

$$p(\mathbf{y}|\boldsymbol{\Phi}, \mathbf{x}) = \prod_{u=1}^U p(\mathbf{y}_u | \mathbf{y}_{(u-1)}, \mathbf{x}_u, \boldsymbol{\Phi}) \quad (11)$$

To avoid the large increase in parameters that comes with sampling or optimizing latent states, we use a continuous-discrete (or hybrid) Kalman filter (?) to analytically compute subject level likelihoods, conditional on subject parameters. For more on filtering see [Jazwinski \(2007\)](#) and [Särkkä \(2013\)](#). The filter operates with a prediction step, in which the expectation $\hat{\boldsymbol{\eta}}_{u|u-1}$ and covariance $\hat{\mathbf{P}}_{u|u-1}$ of the latent states are predicted by:

$$\hat{\boldsymbol{\eta}}_{u|u-1} = \mathbf{A}_u^* \hat{\boldsymbol{\eta}}_{u-1|u-1} + \mathbf{b}_u^* + \mathbf{M} \mathbf{x}_u \quad (12)$$

$$\hat{P}_{u|u-1} = A_u^* \hat{P}_{u-1|u-1} (A_u^*)^\top + Q_u^* \quad (13)$$

For the first measurement occasion $u = 1$, the priors $\hat{\eta}_{u|u-1}$ and $\hat{P}_{u|u-1}$ must be provided to the filter. These parameters may in some cases be freely estimated, but in other cases need to be fixed or constrained, either to specific values or by enforcing a dependency to other parameters in the model, such as an assumption of stationarity.

Prediction steps are followed by an update step, wherein rows and columns of matrices are filtered as necessary depending on missingness of the measurements y :

$$\hat{y}_{u|u-1} = \Lambda \hat{\eta}_{u|u-1} + \tau \quad (14)$$

$$\hat{V}_u = \Lambda \hat{P}_{u|u-1} \Lambda^\top + \Theta \quad (15)$$

$$\hat{K}_u = \hat{P}_{u|u-1} \Lambda^\top \hat{V}_u^{-1} \quad (16)$$

$$\hat{\eta}_{u|u} = \hat{\eta}_{u|u-1} + \hat{K}_u (y_u - \hat{y}_{u|u-1}) \quad (17)$$

$$\hat{P}_{u|u} = (I - \hat{K}_u \Lambda) \hat{P}_{u|u-1} \quad (18)$$

The log likelihood (ll) for each subject, conditional on subject level parameters, is typically¹ then (Genz and Bretz 2009):

$$ll = \sum_{u=1}^U \left(-1/2(n \ln(2\pi) + \ln |V_u| + \right. \\ \left. (\hat{y}_{(u|u-1)} - y_u) V_u^{-1} (\hat{y}_{(u|u-1)} - y_u)^\top) \right) \quad (20)$$

Where n is the number of non-missing observations at measurement occasion u .

2.5. Population level model

Rather than assume complete independence or dependence across subjects, we assume subject level parameters are drawn from a population distribution, for which we also estimate

¹For computational reasons we use an alternate but equivalent form of the log likelihood. We scale the prediction errors across all variables to a standard normal distribution, drop constant terms, calculate the log likelihood of the transformed prediction error vector, and appropriately update the log likelihood for the change in scale, as follows:

$$ll = \sum_{u=1}^U \left(\ln(\text{tr}(V_u^{-1/2})) + \sum 1/2(V_u^{-1/2}(\hat{y}_{(u|u-1)} - y_u)) \right) \quad (19)$$

Where tr indicates the trace of a matrix, and $V^{-1/2}$ is the inverse of the Cholesky decomposition of V . The Stan software manual discusses such a *change of variables* (Team 2016).

parameters and apply some prior. This results in a joint-posterior distribution of:

$$p(\Phi, \mu, R, \beta | Y, Z) = \frac{p(Y|\Phi)p(\Phi|\mu, R, \beta, Z)p(\mu, R, \beta)}{p(Y)} \quad (21)$$

Where subject specific parameters Φ_i are determined in the following manner:

$$\Phi_i = \text{tform}\left(\mu + R h_i + \beta z_i\right) \quad (22)$$

$$h_i \sim N(0, 1) \quad (23)$$

$$\mu \sim N(0, 1) \quad (24)$$

$$\beta \sim N(0, 1) \quad (25)$$

$\Phi_i \in \mathbb{R}^s$ represents all parameters for the dynamic and measurement models of subject i . $\mu \in \mathbb{R}^s$ parameterizes the population means of the distribution of subject level parameters. $R \in \mathbb{R}^{s \times s}$ is the Cholesky factor of the population covariance matrix, parameterizing the effect of subject specific deviations $h_i \in \mathbb{R}^s$ on Φ_i . $\beta \in \mathbb{R}^{s \times w}$ is the effect of time independent predictors $z_i \in \mathbb{R}^w$ on Φ_i . Y_i contains all the data for subject i used in the dynamic model – y (process related measurements) and x (time dependent predictors). Z_i contains time independent predictors data for subject i . tform is an operator that applies a transform to each value of the vector it is applied to. The specific transform depends on which subject level parameter matrix the value belongs to, and the position in that matrix — these transforms and rationale are described below, but are in general necessary because many parameters require some bounded distribution, making a purely linear approach untenable.

Besides the tform operator, Equation 22 looks like a relatively standard hierarchical approach, with subject parameters dependent on a population mean and covariance, and observed covariates. Subject specific parameters h_i are in standardised deviation form to effect a non-centered parameterization, which appears to improve sampling efficiency in this model. See [Bernardo, Bayarri, Berger, Dawid, Heckerman, Smith, and West \(2003\)](#) and [Betancourt and Girolami \(2013\)](#) for discussion of non-centered parameterizations.

2.6. Parameter transformations and priors

A range of considerations must be made when deciding how to parameterize the model. Such considerations include: parameter bounds, distributional assumptions, fixed parameters and ease of interpretation, and sampling efficiency. Boundaries apply to many subject level parameters, such as for instance standard deviations which must be greater than 0. These boundaries also imply that the subject level parameters are unlikely to be normally, or symmetrically, distributed, particularly as population means approach the boundaries. There is a need to be able to fix parameters to specific, understandable values, as for instance with elements of the diffusion matrix Q , which for higher order models will generally require a number of elements fixed to 0. This possibility can be lost under certain transformations. Sampling efficiency is reduced when parameters are correlated (with respect to the sampling procedure), because a random change in one parameter requires a corresponding non-random

change in another to compensate, complicating efficient exploration of the parameter space. While the use of modern sampling approaches like Hamiltonian Monte Carlo (Betancourt and Girolami 2013) and the no U-turn sampler (Homan and Gelman 2014) mitigate these issues to some extent, minimizing correlations between parameters through transformations still substantially improves performance. A further efficiency consideration is the inclusion of sufficient prior information to guide the sampler away from regions where the likelihood of the data approaches zero.

To account for the stated considerations, while allowing for straightforward estimation of the covariance between parameters, we first parameterize the population means according to the standard normal distribution. Subject specific deviations from each mean are distributed according to the R matrix. The R matrix accounts for parameter correlations such as would be found when, for example, subjects that typically score highly on measurements of one process are also likely to exhibit stronger auto-effects on another. The R matrix has a truncated standard normal distribution for the prior of standard deviations (diagonals), and a prior on the Cholesky factor covariance which is the result of multiplying a uniform prior on the space of correlation matrices, with the standard deviations on the diagonal. Discussion of this approach may be found in the Stan software manual (Team 2016), and full details in Lewandowski, Kurowicka, and Joe (2009), but it in general has the virtue that the prior for correlations is independent of the scale. At this point, all the subject level parameters have a multivariate-normal prior, but we then transform these parameters according to the necessary constraints and desired prior distribution. The precise forms and resulting prior densities are specified in Appendix 3, but are intuitively described along with implications below:

Subject level standard deviation parameters are obtained by exponentiating a multiple of the unconstrained parameter, which results in a prior similar to the Jeffreys or reference scale prior (Bernardo 1979), but is lightly regularized away from the low or high extremes to ensure a proper posterior. This form, wherein mass reduces to 0 at any boundaries, is used for all parameters where boundaries exist, because typically at such boundaries other parameters of the model become empirically non-identified.

Parameters for off-diagonals of covariance-related matrices are first transformed to a partial correlation parameter (between -1 and 1), and then combined with the relevant standard deviation parameters to determine the covariance parameter, as for the R matrix earlier. This approach ensures that the off-diagonal parameters are independent of the diagonals, in terms of both prior expectations and during sampling.

Because intercept and regression type parameters need not be bounded, for these we simply scale the standard normal to the desired range (i.e., level of informativeness) by multiplication. We could of course also add some value if we wanted a non-zero mean.

Diagonals of the drift matrix A – the temporal auto effects – are transformed to be negative, with probability mass relatively uniformly distributed for discrete time autoregressive effects between 0 and 1, given a time interval of 1, but declining to 0 at the extremes.

Off diagonals of the drift matrix A – the temporal cross effects – need to be specified in a problem dependent manner. For simple first order processes, they can all be left as multiplications of the standard normal distribution, but for higher order processes, the cross effects between a lower and higher order component will in general need to be parameterized similarly to the auto effects, ensuring negative values. This is so that optimization or sampling does not get stuck at a local minimum just above zero.

3. Parameter transformations and resulting priors

Because all sampled parameters are drawn from a standard normal distribution, we need a variety of transformations that can be applied depending on what the parameter represents. The transformation serves to provide both the upper and lower bounds (if they exist) for the parameter of interest, as well as the prior density. While for many model parameters a simple multiplication of the standard normal is adequate, standard deviations and correlations require a bounded distribution, while we have opted to use a bounded distribution on the drift auto effects for pragmatic reasons – values greater than 0 represent explosive, non-stationary processes that are in most cases not theoretically plausible and occur only due to model misspecification. Further, positive values can result in additional local minima, causing problems for optimization or sampling. While allowing for such values may point to misspecification more readily, the constrained form results in what we believe is a computationally simpler and sensible prior distribution for genuine effects. Estimated values close to 0 when using this constrained form may point to the need to consider the model specification, perhaps by including higher order terms. The transformations we use for the various parameter types are as follows:

$$\text{Standard deviation: } e^{4x} \tag{26}$$

$$\text{Drift auto effect (diagonal): } -\log(e^{-1.5x} + 1) \tag{27}$$

$$\text{Partial correlation: } 2/(1 + e^{-x}) - 1 \tag{28}$$

Covariance matrices are handled in a two step procedure, to ensure that priors on correlation parameters are independent of scale parameters – a long recognised problem with various common parameterizations of covariance matrices, see for instance [Huang and Wand \(2013\)](#) and [Gelman \(2006\)](#). Off-diagonals are first transformed via the inverse logit function and scaled to the range -1 to 1, and used in the lower triangle of a Cholesky decomposed correlation matrix. The diagonal of the Cholesky correlation matrix is determined based on both the standard deviations and correlations. Then, the covariance parameters are calculated by pre-multiplying the Cholesky correlation matrix by the scale matrix, which is simply the diagonal matrix containing the standard deviations. This approach ensures that the off-diagonal parameters are independent of the diagonals, in terms of both prior expectations and during sampling. Discussion of this general approach to handling covariance matrices may be found in the Stan software manual ([Team 2016](#)), and full details in [Lewandowski et al. \(2009\)](#).

Figure 1 plots the resulting densities when using the described transformations. Note that of course the density for a variance is directly related to the standard deviation, and the density plot for an autoregression assumes that the time interval is 1 with no cross effects involved. For the sake of completeness we include a prior density for all other parameters, such as the drift cross effects, intercepts, and regression type parameters. These use a simple multiplication of the standard normal.

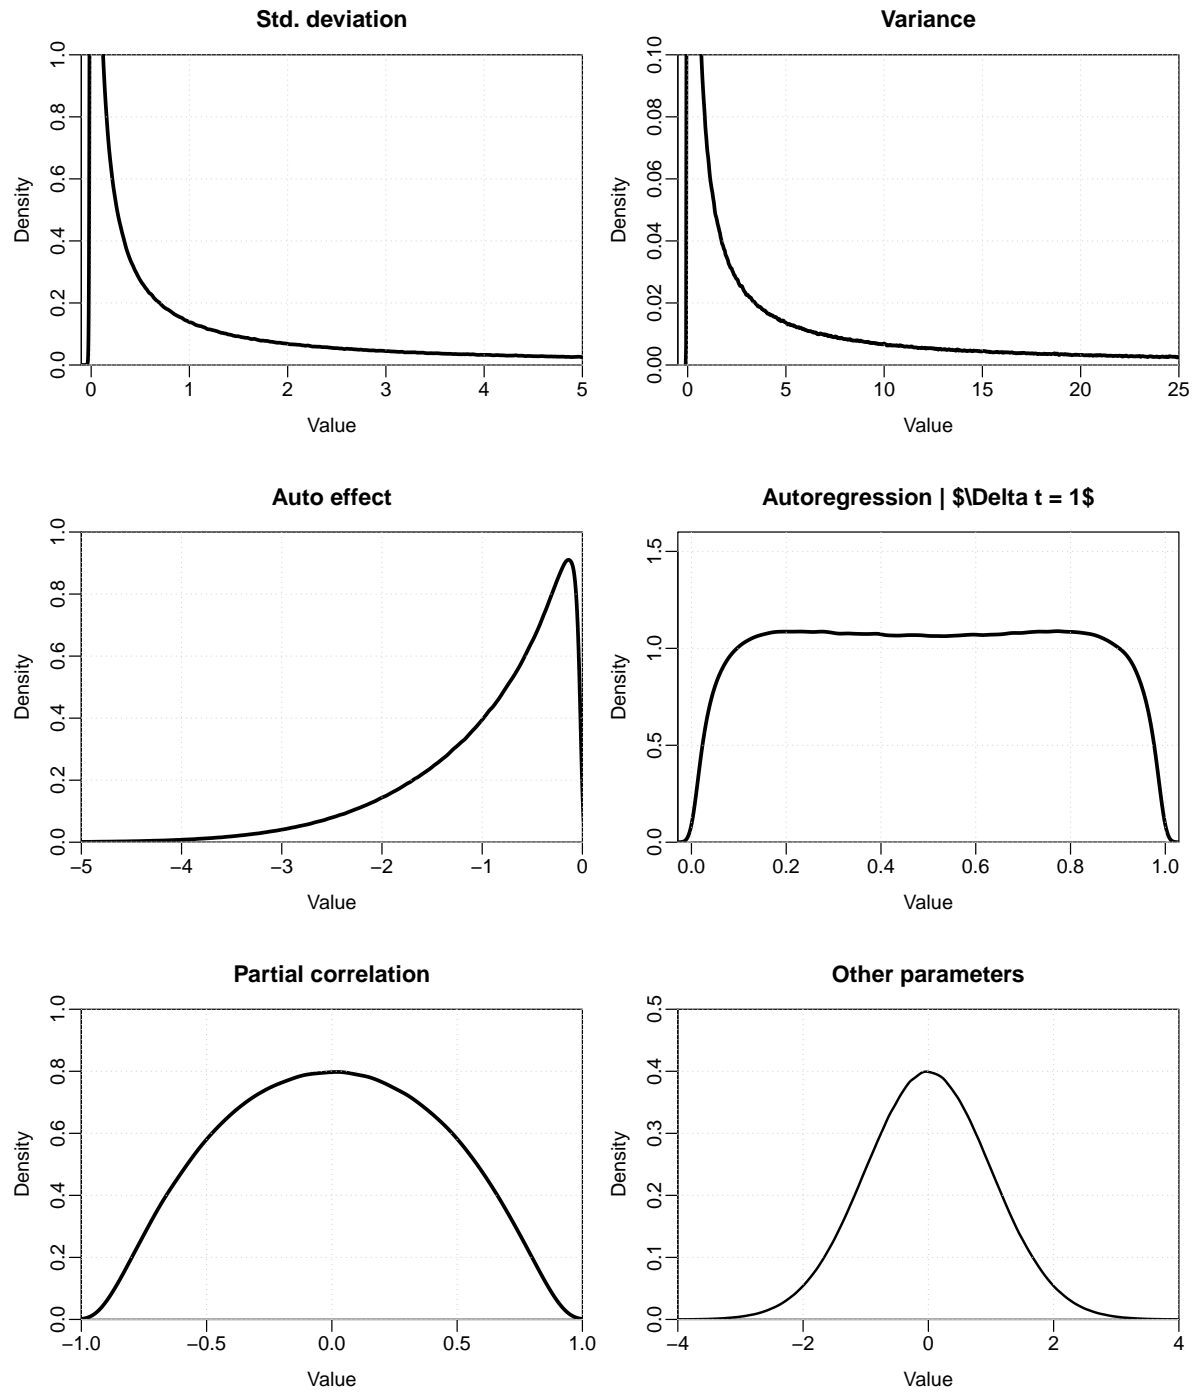


Figure 1: Priors for population mean parameters.

References

- Bernardo JM (1979). “Reference Posterior Distributions for Bayesian Inference.” *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 113–147.
- Bernardo JM, Bayarri MJ, Berger JO, Dawid AP, Heckerman D, Smith AFM, West M (2003). “Non-Centered Parameterisations for Hierarchical Models and Data Augmentation.” *Bayesian Statistics 7: Proceedings of the Seventh Valencia International Meeting*, p. 307.
- Betancourt MJ, Girolami M (2013). “Hamiltonian Monte Carlo for Hierarchical Models.” *arXiv:1312.0906 [stat]*. [1312.0906](#).
- Driver C, Oud J, Voelkle M (In Press). “Continuous Time Structural Equation Modelling with R Package Ctsem.” *Journal of Statistical Software*.
- Gelman A (2006). “Prior Distributions for Variance Parameters in Hierarchical Models (Comment on Article by Browne and Draper).” *Bayesian Analysis*, **1**(3), 515–534. ISSN 1936-0975, 1931-6690. [doi:10.1214/06-BA117A](#).
- Genz A, Bretz F (2009). *Computation of Multivariate Normal and T Probabilities*, volume 195 of *Lecture Notes in Statistics*. Springer Berlin Heidelberg, Berlin, Heidelberg. ISBN 978-3-642-01688-2 978-3-642-01689-9.
- Homan MD, Gelman A (2014). “The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo.” *J. Mach. Learn. Res.*, **15**(1), 1593–1623. ISSN 1532-4435.
- Huang A, Wand MP (2013). “Simple Marginally Noninformative Prior Distributions for Covariance Matrices.” *Bayesian Analysis*, **8**(2), 439–452. ISSN 1936-0975, 1931-6690. [doi:10.1214/13-BA815](#).
- Jazwinski AH (2007). *Stochastic Processes and Filtering Theory*. Courier Corporation. ISBN 978-0-486-46274-5.
- Lewandowski D, Kurowicka D, Joe H (2009). “Generating Random Correlation Matrices Based on Vines and Extended Onion Method.” *Journal of Multivariate Analysis*, **100**(9), 1989–2001. ISSN 0047-259X. [doi:10.1016/j.jmva.2009.04.008](#).
- Särkkä S (2013). *Bayesian Filtering and Smoothing*. Cambridge University Press. ISBN 978-1-107-03065-7. Google-Books-ID: 5VlsAAAAQBAJ.
- Team SD (2016). “Stan Modeling Language Users Guide and Reference Manual, Version 2.9.0.”
- Tómasson H (2013). “Some Computational Aspects of Gaussian CARMA Modelling.” *Statistics and Computing*, **25**(2), 375–387. ISSN 0960-3174, 1573-1375. [doi:10.1007/s11222-013-9438-9](#).

Affiliation:

Charles Driver

Center for Lifespan Psychology

Max Planck Institute for Human Development

Lentzeallee 94, 14195 Berlin

Telephone: +49 30 82406-367 E-mail: driver@mpib-berlin.mpg.de

URL: <http://www.mpib-berlin.mpg.de/en/staff/charles-driver>