ucm — Unobserved-components model

Description

Unobserved-components models (UCMs) decompose a time series into trend, seasonal, cyclical, and idiosyncratic components and allow for exogenous variables. ucm estimates the parameters of UCMs by maximum likelihood.

All the components are optional. The trend component may be first-order deterministic or it may be first-order or second-order stochastic. The seasonal component is stochastic; the seasonal effects at each time period sum to a zero-mean finite-variance random variable. The cyclical component is modeled by the stochastic-cycle model derived by Harvey (1989).

Quick start

Random-walk model for y using tsset data

    ucm y

Add a cyclical component of order 2

    ucm y, cycle(2)

Add a seasonal component arising every 3 periods

    ucm y, cycle(2) seasonal(3)

Random-walk model for y with a drift component and a cyclical component of order 1

    ucm y, model(rwdrift) cycle(1)

Smooth-trend model for y with cyclical and seasonal components of order 2

    ucm y, model(strend) cycle(2) seasonal(2)

Menu

Statistics > Time series > Unobserved-components model

Syntax

> ucm *depvar* [*indepvars*] [*if*] [*in*] [, *options*]

options	Description
Model	
model(*model*)	specify trend and idiosyncratic components
seasonal(*#*)	include a seasonal component with a period of *#* time units
cycle(*#* [, **f**requency(*#*$_f$)])	include a cycle component of order *#* and optionally set initial frequency to *#*$_f$, $0 < \#_f < \pi$; cycle() may be specified up to three times
constraints(*constraints*)	apply specified linear constraints
SE/Robust	
vce(*vcetype*)	*vcetype* may be oim or **r**obust
Reporting	
level(*#*)	set confidence level; default is level(95)
nocnsreport	do not display constraints
display_options	control columns and column formats, row spacing, display of omitted variables and base and empty cells, and factor-variable labeling
Maximization	
maximize_options	control the maximization process
collinear	keep collinear variables
coeflegend	display legend instead of statistics

model	Description
rwalk	random-walk model; the default
none	no trend or idiosyncratic component
ntrend	no trend component but include idiosyncratic component
dconstant	deterministic constant with idiosyncratic component
llevel	local-level model
dtrend	deterministic-trend model with idiosyncratic component
lldtrend	local-level model with deterministic trend
rwdrift	random-walk-with-drift model
lltrend	local-linear-trend model
strend	smooth-trend model
rtrend	random-trend model

You must tsset your data before using ucm; see [TS] **tsset**.

indepvars may contain factor variables; see [U] **11.4.3 Factor variables**.

indepvars and *depvar* may contain time-series operators; see [U] **11.4.4 Time-series varlists**.

by, collect, fp, rolling, and statsby are allowed; see [U] **11.1.10 Prefix commands**.

collinear and coeflegend do not appear in the dialog box.

See [U] **20 Estimation and postestimation commands** for more capabilities of estimation commands.

Options

Model

model(*model*) specifies the trend and idiosyncratic components. The default is model(rwalk). The available *model*s are listed in *Syntax* and discussed in detail in *Models for the trend and idiosyncratic components* under *Remarks and examples* below.

seasonal(#) adds a stochastic-seasonal component to the model. # is the period of the season, that is, the number of time-series observations required for the period to complete.

cycle(#) adds a stochastic-cycle component of order # to the model. The order # must be 1, 2, or 3. Multiple cycles are added by repeating the cycle(#) option with up to three cycles allowed.

cycle(#, frequency($#_f$)) specifies $#_f$ as the initial value for the central-frequency parameter in the stochastic-cycle component of order #. $#_f$ must be in the interval $(0, \pi)$.

constraints(*constraints*); see [R] **Estimation options**.

SE/Robust

vce(*vcetype*) specifies the estimator for the variance–covariance matrix of the estimator.

vce(oim), the default, causes ucm to use the observed information matrix estimator.

vce(robust) causes ucm to use the Huber/White/sandwich estimator.

Reporting

level(#), nocnsreport; see [R] **Estimation options**.

display_options: noci, nopvalues, noomitted, vsquish, noemptycells, baselevels, allbaselevels, nofvlabel, fvwrap(#), fvwrapon(*style*), cformat(*%fmt*), pformat(*%fmt*), and sformat(*%fmt*); see [R] **Estimation options**.

Maximization

maximize_options: difficult, technique(*algorithm_spec*), iterate(#), [no]log, trace, gradient, showstep, hessian, showtolerance, tolerance(#), ltolerance(#), nrtolerance(#), and from(*matname*); see [R] **Maximize** for all options except from(), and see below for information on from().

from(*matname*) specifies initial values for the maximization process. from(b0) causes ucm to begin the maximization algorithm with the values in b0. b0 must be a row vector; the number of columns must equal the number of parameters in the model; and the values in b0 must be in the same order as the parameters in e(b).

If your model fails to converge, try using the difficult option. Also see the technical note below example 5.

The following options are available with ucm but are not shown in the dialog box:

collinear, coeflegend; see [R] **Estimation options**.

Remarks and examples

Remarks are presented under the following headings:

An introduction to UCMs

UCMs decompose a time series into trend, seasonal, cyclical, and idiosyncratic components and allow for exogenous variables. Formally, UCMs can be written as

$$y_t = \tau_t + \gamma_t + \psi_t + \boldsymbol{\beta}\mathbf{x}_t + \epsilon_t \tag{1}$$

where y_t is the dependent variable, τ_t is the trend component, γ_t is the seasonal component, ψ_t is the cyclical component, $\boldsymbol{\beta}$ is a vector of fixed parameters, $\mathbf{x}_t$ is a vector of exogenous variables, and ϵ_t is the idiosyncratic component.

By placing restrictions on τ_t and ϵ_t, Harvey (1989) derived a series of models for the trend and the idiosyncratic components. These models are briefly described in *Syntax* and are further discussed in *Models for the trend and idiosyncratic components*. To these models, Harvey (1989) added models for the seasonal and cyclical components, and he also allowed for the presence of exogenous variables.

It is rare that a UCM contains all the allowed components. For instance, the seasonal component is rarely needed when modeling deseasonalized data.

Harvey (1989) and Durbin and Koopman (2012) show that UCMs can be written as state-space models that allow the parameters of a UCM to be estimated by maximum likelihood. In fact, ucm uses sspace (see [TS] **sspace**) to perform the estimation calculations; see *Methods and formulas* for details.

After estimating the parameters, predict can produce in-sample predictions or out-of-sample forecasts; see [TS] **ucm postestimation**. After estimating the parameters of a UCM that contains a cyclical component, estat period converts the estimated central frequency to an estimated central period and psdensity estimates the spectral density implied by the model; see [TS] **ucm postestimation** and the examples below.

We illustrate the basic approach of analyzing data with UCMs, and then we discuss the details of the different trend models in *Models for the trend and idiosyncratic components*.

Although the methods implemented in ucm have been widely applied by economists, they are general time-series techniques and may be of interest to researchers from other disciplines. In example 8, we analyze monthly data on the reported cases of mumps in New York City.

A random-walk model example

▷ Example 1

We begin by plotting monthly data on the US civilian unemployment rate.

```
. use https://www.stata-press.com/data/r19/unrate
. tsline unrate, name(unrate)
```


This series looks like it might be well approximated by a random-walk model. Formally, a random-walk model is given by

$$
\begin{aligned}
y_t &= \mu_t \\
\mu_t &= \mu_{t-1} + \eta_t
\end{aligned}
$$

The random-walk is so frequently applied, at least as a starting model, that it is the default model for ucm. In the output below, we fit the random-walk model to the unemployment data.

```
. ucm unrate
searching for initial values ..........
(setting technique to bhhh)
Iteration 0:  Log likelihood =  84.272992
Iteration 1:  Log likelihood =  84.394942
Iteration 2:  Log likelihood =  84.400923
Iteration 3:  Log likelihood =  84.401282
Iteration 4:  Log likelihood =  84.401305
(switching technique to nr)
Iteration 5:  Log likelihood =  84.401306
Refining estimates:
Iteration 0:  Log likelihood =  84.401306
Iteration 1:  Log likelihood =  84.401307

Unobserved-components model
Components: random walk
Sample: 1948m1 thru 2011m1                          Number of obs = 757
Log likelihood = 84.401307
```

unrate	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
var(level)	.0467196	.002403	19.44	0.000	.0420098	.0514294

```
Note: Model is not stationary.
Note: Tests of variances against zero are one sided, and the two-sided
      confidence intervals are truncated at zero.
```

The output indicates that the model is nonstationary, as all random-walk models are.

We consider a richer model in the next example.

◁

▷ Example 2

We suspect that there should be a stationary cyclical component that produces serially correlated shocks around the random-walk trend. Harvey (1989) derived a stochastic-cycle model for these stationary cyclical components.

The stochastic-cycle model has three parameters: the frequency at which the random components are centered, a damping factor that parameterizes the dispersion of the random components around the central frequency, and the variance of the stochastic-cycle process that acts as a scale factor.

Fitting this model to unemployment data yields

```
. ucm unrate, cycle(1)
searching for initial values ...................
(setting technique to bhhh)
Iteration 0:  Log likelihood =  84.273579
Iteration 1:  Log likelihood =  87.852115
Iteration 2:  Log likelihood =  88.253422
Iteration 3:  Log likelihood =  89.191311
Iteration 4:  Log likelihood =  94.675898
(switching technique to nr)
Iteration 5:  Log likelihood =  98.394691  (not concave)
Iteration 6:  Log likelihood =  98.983093
Iteration 7:  Log likelihood =  99.983625
Iteration 8:  Log likelihood =  104.83113
Iteration 9:  Log likelihood =   114.2697
Iteration 10: Log likelihood =  116.47453
Iteration 11: Log likelihood =  118.45857
Iteration 12: Log likelihood =  118.88057
Iteration 13: Log likelihood =  118.88421
Iteration 14: Log likelihood =  118.88421
Refining estimates:
Iteration 0:  Log likelihood =  118.88421
Iteration 1:  Log likelihood =  118.88421

Unobserved-components model
Components: random walk, order 1 cycle
Sample: 1948m1 thru 2011m1                     Number of obs =       757
                                              Wald chi2(2)  = 26650.81
Log likelihood = 118.88421                    Prob > chi2   =   0.0000
```

unrate	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
frequency	.0933466	.0103609	9.01	0.000	.0730397	.1136535
damping	.9820003	.0061121	160.66	0.000	.9700207	.9939798
var(level)	.0143786	.0051392	2.80	0.003	.004306	.0244511
var(cycle1)	.0270339	.0054343	4.97	0.000	.0163829	.0376848

```
Note: Model is not stationary.
Note: Tests of variances against zero are one sided, and the two-sided
      confidence intervals are truncated at zero.
```

The estimated central frequency for the cyclical component is small, implying that the cyclical component is centered on low-frequency components. The high-damping factor indicates that all the components from this cyclical component are close to the estimated central frequency. The estimated variance of the stochastic-cycle process is small but significant.

We use estat period to convert the estimate of the central frequency to an estimated central period.

```
. estat period
```

cycle1	Coefficient	Std. err.	[95% conf. interval]	
period	67.31029	7.471004	52.66739	81.95319
frequency	.0933466	.0103609	.0730397	.1136535
damping	.9820003	.0061121	.9700207	.9939798

```
Note: Cycle time unit is monthly.
```

Because we have monthly data, the estimated central period of 67.31 implies that the cyclical component is composed of random components that occur around a central periodicity of about 5.61 years. This estimate falls within the conventional Burns and Mitchell (1946) definition of business-cycle shocks occurring between 1.5 and 8 years.

We can convert the estimated parameters of the cyclical component to an estimated spectral density of the cyclical component, as described by Harvey (1989). The spectral density of the cyclical component describes the relative importance of the random components at different frequencies; see *Frequency-domain concepts used in the stochastic-cycle model* for details. We use psdensity (see [TS] **psdensity**) to obtain the spectral density of the cyclical component implied by the estimated parameters, and we use twoway line (see [G-2] **graph twoway line**) to plot the estimated spectral density.

```
. psdensity sdensity omega
. line sdensity omega
```


The estimated spectral density shows that the cyclical component is composed of random components that are tightly distributed at the low-frequency peak.

◁

Frequency-domain concepts used in the stochastic-cycle model

The parameters of the stochastic-cycle model are easiest to interpret in the frequency domain. We now provide a review of the useful concepts from the frequency domain. Crucial to understanding the stochastic-cycle model is the frequency-domain concept that a stationary process can be decomposed into random components that occur at the frequencies in the interval $[0, \pi]$.

We need some concepts from the frequency-domain approach to interpret the parameters in the stochastic-cycle model of the cyclical component. Here we provide a simple, intuitive explanation. More technical presentations can be found in Priestley (1981), Harvey (1989, 1993), Hamilton (1994), Fuller (1996), and Wei (2006).

As with much time-series analysis, the basic results are for covariance-stationary processes with additional results handling some nonstationary cases. We present some useful results for covariance-stationary processes. These results provide what we need to interpret the stochastic-cycle model for the stationary cyclical component.

The autocovariances γ_j, $j \in \{0, 1, \ldots, \infty\}$, of a covariance-stationary process y_t specify its variance and dependence structure. In the frequency-domain approach to time-series analysis, the spectral density describes the importance of the random components that occur at frequency ω relative to the components that occur at other frequencies.

The frequency-domain approach focuses on the relative contributions of random components that occur at the frequencies $[0, \pi]$.

The spectral density can be written as a weighted average of the autocorrelations of y_t. Like autocorrelations, the spectral density is normalized by γ_0, the variance of y_t. Multiplying the spectral density by γ_0 yields the power-spectrum of y_t.

In an independent and identically distributed (i.i.d.) process, the components at all frequencies are equally important, so the spectral density is a flat line.

In common parlance, we speak of high-frequency noise making a series look more jagged and of low-frequency components causing smoother plots. More formally, we say that a process composed primarily of high-frequency components will have fewer runs above or below the mean than an i.i.d. process and that a process composed primarily of low-frequency components will have more runs above or below the mean than an i.i.d. process.

To further formalize these ideas, consider the first-order autoregressive (AR(1)) process given by

$$y_t = \phi y_{t-1} + \epsilon_t$$

where ϵ_t is a zero-mean, covariance-stationary process with finite variance σ^2, and $|\phi| < 1$ so that y_t is covariance stationary. The first-order autocorrelation of this AR(1) process is ϕ.

Below are plots of simulated data when ϕ is set to 0, -0.8, and 0.8. When $\phi = 0$, the data are i.i.d. When $\phi = -0.8$, the value today is strongly negatively correlated with the value yesterday, so this case should be a prototypical high-frequency noise example. When $\phi = 0.8$, the value today is strongly positively correlated with the value yesterday, so this case should be a prototypical low-frequency shock example.

The plots above confirm our conjectures. The plot when $\phi = -0.8$ contains fewer runs above or below the mean, and it is more jagged than the i.i.d. plot. The plot when $\phi = 0.8$ contains more runs above or below the mean, and it is smoother than the i.i.d. plot.

Below we plot the spectral densities for the AR(1) model with $\phi = 0$, $\phi = -0.8$, and $\phi = 0.8$.

The high-frequency components are much more important to the AR(1) process with $\phi = -0.8$ than to the i.i.d. process with $\phi = 0$. The low-frequency components are much more important to the AR(1) process with $\phi = 0.8$ than to the i.i.d. process.

❑ Technical note

Autoregressive moving-average (ARMA) models parameterize the autocorrelation in a time series by allowing today's value to be a weighted average of past values and a weighted average of past i.i.d. shocks; see Hamilton (1994), Wei (2006), and [TS] **arima** for introductions and a Stata implementation. The intuitive ARMA parameterization has many nice features, including that one can easily rewrite the ARMA model as a weighted average of past i.i.d. shocks to trace how a shock feeds through the system.

Although it is easy to obtain the spectral density of an ARMA process, the parameters themselves provide limited information about the underlying spectral density.

In contrast, the parameters of the stochastic-cycle parameterization of autocorrelation in a time series directly provide information about the underlying spectral density. The parameter ω_0 is the central frequency at which the random components are clustered. If ω_0 is small, then the model is centered on low-frequency components. If ω_0 is close to π, then the model is centered on high-frequency components. The parameter ρ is the damping factor that indicates how tightly clustered the random components are at the central frequency ω_0. If ρ is close to 0, there is no clustering of the random components. If ρ is close to 1, the random components are tightly distributed at the central frequency ω_0.

In the graph below, we draw the spectral densities implied by stochastic-cycle models with four sets of parameters: $\omega_0 = \pi/4$, $\rho = 0.8$; $\omega_0 = \pi/4$, $\rho = 0.9$; $\omega_0 = 4\pi/5$, $\rho = 0.8$; and $\omega_0 = 4\pi/5$, $\rho = 0.9$. The graph below illustrates that ω_0 is the central frequency at which the other important random components are distributed. It also illustrates that the damping parameter ρ controls the dispersion of the important components at the central frequency.

❑

Another random-walk model example

▷ Example 3

Now let's reconsider example 2. Although we might be happy with how our model has identified a stationary cyclical component that we could interpret in business-cycle terms, we suspect that there should also be a high-frequency cyclical component. It is difficult to estimate the parameters of a UCM with two or more stochastic-cycle models. Providing starting values for the central frequencies can be a crucial help to the optimization procedure. Below we estimate a UCM with two cyclical components. We use the frequency() suboption to provide starting values for the central frequencies; we specified the values below because we suspect one model will pick up the low-frequency components and the other will pick up the high-frequency components. We specified the low-frequency model to be order 2 to make it less peaked for any given damping factor. (Trimbur [2006] provides a nice introduction and some formal results for higher-order stochastic-cycle models.)

```
. ucm unrate, cycle(1, frequency(2.9)) cycle(2, frequency(.09))
searching for initial values ...................
(setting technique to bhhh)
Iteration 0:  Log likelihood =  115.98563
Iteration 1:  Log likelihood =  125.04043
Iteration 2:  Log likelihood =  127.69387
Iteration 3:  Log likelihood =  134.50864
Iteration 4:  Log likelihood =  136.91353
(switching technique to nr)
Iteration 5:  Log likelihood =   138.5091
Iteration 6:  Log likelihood =  146.09273
Iteration 7:  Log likelihood =  146.28132
Iteration 8:  Log likelihood =  146.28326
Iteration 9:  Log likelihood =  146.28326
Refining estimates:
Iteration 0:  Log likelihood =  146.28326
Iteration 1:  Log likelihood =  146.28326
```

```
Unobserved-components model
Components: random walk, 2 cycles of order 1 2
Sample: 1948m1 thru 2011m1                          Number of obs  =      757
                                                    Wald chi2(4)   = 7681.33
Log likelihood = 146.28326                          Prob > chi2    =  0.0000
```

| unrate	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
cycle1						
frequency	2.882382	.0668017	43.15	0.000	2.751453	3.013311
damping	.7004295	.125157	5.60	0.000	.4551262	.9457328
cycle2						
frequency	.0667929	.0206848	3.23	0.001	.0262514	.1073345
damping	.9074708	.0142273	63.78	0.000	.8795858	.9353559
var(level)	.0207704	.0039669	5.24	0.000	.0129953	.0285454
var(cycle1)	.0027886	.0014363	1.94	0.026	0	.0056037
var(cycle2)	.002714	.001028	2.64	0.004	.0006991	.0047289

Note: Model is not stationary.
Note: Tests of variances against zero are one sided, and the two-sided
      confidence intervals are truncated at zero.

The output provides some support for the existence of a second, high-frequency cycle. The high-frequency components are centered at 2.88, whereas the low-frequency components are centered at 0.067. That the estimated damping factor is 0.70 for the high-frequency cycle whereas the estimated damping factor for the low-frequency cycle is 0.91 indicates that the high-frequency components are more diffusely distributed at 2.88 than the low-frequency components are at 0.067.

We obtain and plot the estimated spectral densities to get another look at these results.

```
. psdensity sdensity2a omega2a
. psdensity sdensity2b omega2b, cycle(2)
. line sdensity2a sdensity2b omega2a
```
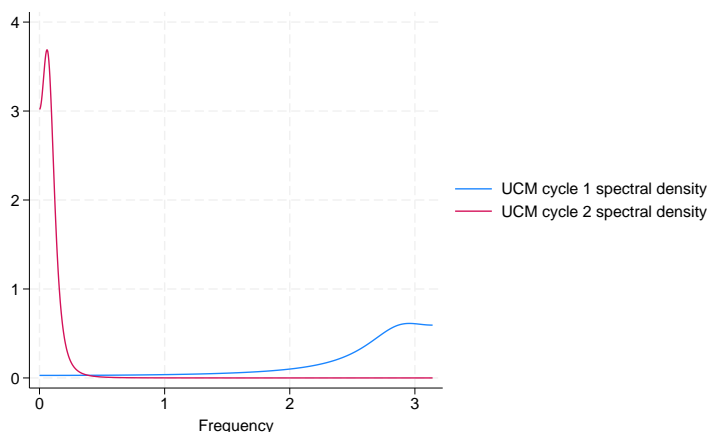

The estimated spectral densities indicate that we have found two distinct cyclical components.

It does not matter whether we specify omega2a or omega2b to be the x-axis variable, because they are equal to each other.

◁

❑ Technical note

That the estimated spectral densities in the previous example do not overlap is important for parameter identification. Although the parameters are identified in large-sample theory, we have found it difficult to estimate the parameters of two cyclical components when the spectral densities overlap. When the spectral densities of two cyclical components overlap, the parameters may not be well identified and the optimization procedure may not converge.

❑

Comparing UCM and ARIMA

▷ Example 4

This example provides some insight for readers familiar with autoregressive integrated moving-average (ARIMA) models but not with UCMs. If you are not familiar with ARIMA models, you may wish to skip this example. See [TS] **arima** for an introduction to ARIMA models in Stata.

UCMs provide an alternative to ARIMA models implemented in [TS] **arima**. Neither set of models is nested within the other, but there are some cases in which instructive comparisons can be made.

The random-walk model corresponds to an ARIMA model that is first-order integrated and has an i.i.d. error term. In other words, the random-walk UCM and the ARIMA(0,1,0) are asymptotically equivalent. Thus

```
ucm unrate
```

and

```
arima unrate, arima(0,1,0) noconstant
```

produce asymptotically equivalent results.

The stochastic-cycle model for the stationary cyclical component is an alternative functional form for stationary processes to stationary autoregressive moving-average (ARMA) models. Which model is preferred depends on the application and which parameters a researchers wants to interpret. Both the functional forms and the parameter interpretations differ between the stochastic-cycle model and the ARMA model. See Trimbur (2006, eq. 25) for some formal comparisons of the two models.

That both models can be used to estimate the stationary cyclical components for the random-walk model implies that we can compare the results in this case by comparing their estimated spectral densities. Below we estimate the parameters of an ARIMA(2,1,1) model and plot the estimated spectral density of the stationary component.

```
. arima unrate, noconstant arima(2,1,1)

(setting optimization to BHHH)
Iteration 0:  Log likelihood =   129.8801
Iteration 1:  Log likelihood =  134.61953
Iteration 2:  Log likelihood =  137.04909
Iteration 3:  Log likelihood =  137.71386
Iteration 4:  Log likelihood =  138.25255
(switching optimization to BFGS)
Iteration 5:  Log likelihood =  138.51924
Iteration 6:  Log likelihood =  138.81638
Iteration 7:  Log likelihood =  138.83615
Iteration 8:  Log likelihood =   138.8364
Iteration 9:  Log likelihood =  138.83642
Iteration 10: Log likelihood =  138.83642

ARIMA regression
```

Sample: 1948m2 thru 2011m1

			Number of obs	=	756
			Wald chi2(3)	=	683.34
Log likelihood = 138.8364			Prob > chi2	=	0.0000

D.unrate	Coefficient	OPG std. err.	z	P>|z|	[95% conf. interval]
ARMA					
ar					
L1.	.5398016	.0586304	9.21	0.000	.4248882 .6547151
L2.	.2468148	.0359396	6.87	0.000	.1763744 .3172551
ma					
L1.	-.5146506	.0632838	-8.13	0.000	-.6386845 -.3906167
/sigma	.2013332	.0032644	61.68	0.000	.1949351 .2077313

Note: The test of the variance against zero is one sided, and the two-sided
      confidence interval is truncated at zero.

```
. psdensity sdensity_arma omega_arma
. line sdensity_arma omega_arma
```


The estimated spectral density from the ARIMA(2,1,1) has a similar shape to the plot obtained by combining the two spectral densities estimated from the stochastic-cycle model in example 3. For this particular application, the estimated central frequencies of the two cyclical components from the stochastic-

cycle model provide information about the business-cycle component and the high-frequency component that is not easily obtained from the ARIMA(2,1,1) model. On the other hand, it is easier to work out the impulse–response function for the ARMA model than for the stochastic-cycle model, implying that the ARMA model is easier to use when tracing the effect of a shock feeding through the system.

◁

A local-level model example

We now consider the weekly series of initial claims for unemployment insurance in the United States, which is plotted below.

▷ Example 5

```
. use https://www.stata-press.com/data/r19/icsa1, clear
. tsline icsa
```


This series looks like it was generated by a random walk with extra noise, so we want to use a random-walk model that includes an additional random term. This structure causes the model to be occasionally known as the random-walk-plus-noise model, but it is more commonly known as the local-level model in the UCM literature.

The local-level model models the trend as a random walk and models the idiosyncratic components as independent and identically distributed components. Formally, the local-level model specifies the observed time-series y_t, for $t = 1, \ldots, T$, as

$$y_t = \mu_t + \epsilon_t$$
$$\mu_t = \mu_{t-1} + \eta_t$$

where $\epsilon_t \sim$ i.i.d. $N(0, \sigma_\epsilon^2)$ and $\eta_t \sim$ i.i.d. $N(0, \sigma_\eta^2)$ and are mutually independent.

We fit the local-level model in the output below:

```
. ucm icsa, model(llevel)
searching for initial values ..........
(setting technique to bhhh)
Iteration 0:   Log likelihood = -9954.8223
Iteration 1:   Log likelihood =  -9917.406
Iteration 2:   Log likelihood = -9905.6679
Iteration 3:   Log likelihood = -9897.7588
Iteration 4:   Log likelihood = -9894.2015
(switching technique to nr)
Iteration 5:   Log likelihood = -9893.4337
Iteration 6:   Log likelihood = -9893.2469
Iteration 7:   Log likelihood = -9893.2469
Refining estimates:
Iteration 0:   Log likelihood = -9893.2469
Iteration 1:   Log likelihood = -9893.2469

Unobserved-components model
Components: local level
Sample: 07jan1967 thru 19feb2011                    Number of obs = 2,303
Log likelihood = -9893.2469
```

icsa	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
var(level)	116.558	8.806587	13.24	0.000	99.29745	133.8186
var(icsa)	124.2715	7.615506	16.32	0.000	109.3454	139.1976

```
Note: Model is not stationary.
Note: Tests of variances against zero are one sided, and the two-sided
      confidence intervals are truncated at zero.
Note: Time units are in 7 days.
```

The output indicates that both components are statistically significant.

◁

❑ Technical note

The estimation procedure will not always converge when estimating the parameters of the local-level model. If the series does not vary enough in the random level, modeled by the random walk, and in the stationary shocks around the random level, the estimation procedure will not converge because it will be unable to set the variance of one of the two components to 0.

Take another look at the graphs of `unrate` and `icsa`. The extra noise around the random level that can be seen in the graph of `icsa` allows us to estimate both variances.

A closely related point is that it is difficult to estimate the parameters of a local-level model with a stochastic-cycle component because the series must have enough variation to identify the variance of the random-walk component, the variance of the idiosyncratic term, and the parameters of the stochastic-cycle component. In some cases, series that look like candidates for the local-level model are best modeled as random-walk models with stochastic-cycle components.

In fact, convergence can be a problem for most of the models in `ucm`. Convergence problems occur most often when there is insufficient variation to estimate the variances of the components in the model. When there is insufficient variation to estimate the variances of the components in the model, the optimization routine will fail to converge as it attempts to set the variance equal to 0. This usually shows up in the iteration log when the log likelihood gets stuck at a particular value and the message (`not concave`) or (`backed up`) is displayed repeatedly. When this happens, use the `iterate()` option

to limit the number of iterations, look to see which of the variances is being driven to 0, and drop that component from the model. (This technique is a method to obtain convergence to interpretable estimates, not a model-selection method.)

❑

▷ Example 6

We might suspect that there is some serial correlation in the idiosyncratic shock. Alternatively, we could include a cyclical component to model the stationary time-dependence in the series. In the example below, we add a stochastic-cycle model for the stationary cyclical process, but we drop the idiosyncratic term and use a random-walk model instead of the local-level model. We change the model because it is difficult to estimate the variance of the idiosyncratic term along with the parameters of a stationary cyclical component.

```
. ucm icsa, model(rwalk) cycle(1)
searching for initial values ...................
(setting technique to bhhh)
Iteration 0:  Log likelihood = -10055.453
Iteration 1:  Log likelihood = -10047.163
Iteration 2:  Log likelihood = -10047.146  (backed up)
Iteration 3:  Log likelihood = -10047.146  (backed up)
Iteration 4:  Log likelihood = -10047.145  (backed up)
(switching technique to nr)
Iteration 5:  Log likelihood = -10047.142  (not concave)
Iteration 6:  Log likelihood = -9889.8038
Iteration 7:  Log likelihood =  -9883.967
Iteration 8:  Log likelihood = -9883.3818  (not concave)
Iteration 9:  Log likelihood = -9883.3817  (not concave)
Iteration 10: Log likelihood = -9883.3815  (not concave)
Iteration 11: Log likelihood = -9883.3789  (not concave)
Iteration 12: Log likelihood =  -9883.376  (not concave)
Iteration 13: Log likelihood = -9883.3684  (not concave)
Iteration 14: Log likelihood = -9882.0677  (not concave)
Iteration 15: Log likelihood = -9881.6608
Iteration 16: Log likelihood = -9881.4451
Iteration 17: Log likelihood = -9881.4441
Refining estimates:
Iteration 0:  Log likelihood = -9881.4441
Iteration 1:  Log likelihood = -9881.4441

Unobserved-components model
Components: random walk, order 1 cycle
Sample: 07jan1967 thru 19feb2011                  Number of obs =   2,303
                                                  Wald chi2(2)  =   23.04
Log likelihood = -9881.4441                       Prob > chi2   =  0.0000
```

| icsa	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
frequency	1.469634	.3855654	3.81	0.000	.7139395	2.225328
damping	.1644576	.0349537	4.71	0.000	.0959495	.2329656
var(level)	97.90982	8.320047	11.77	0.000	81.60282	114.2168
var(cycle1)	149.7323	9.980797	15.00	0.000	130.1703	169.2943

Note: Model is not stationary.
Note: Tests of variances against zero are one sided, and the two-sided
      confidence intervals are truncated at zero.
Note: Time units are in 7 days.

Although the output indicates that the model fits well, the small estimate of the damping parameter indicates that the random components will be widely distributed at the central frequency. To get a better idea of the dispersion of the components, we look at the estimated spectral density of the stationary cyclical component.

```
. psdensity sdensity3 omega3
. line sdensity3 omega3
```


The graph shows that the random components that make up the cyclical component are diffusely distributed at a central frequency.

◁

Comparing UCM and ARIMA, revisited

▷ Example 7

Including lags of the dependent variable is an alternative method for modeling serially correlated errors. The estimated coefficients on the lags of the dependent variable estimate the coefficients in an autoregressive model for the stationary cyclical component; see Harvey (1989, 47–48) for a discussion. Including lags of the dependent variable should be viewed as an alternative to the stochastic-cycle model for the stationary cyclical component. In this example, we use the large-sample equivalence of the random-walk model with pth order autoregressive errors and an ARIMA$(p, 1, 0)$ to illustrate this point.

In the output below, we include 2 lags of the dependent variable in the random-walk UCM.

```
. ucm icsa L(1/2).icsa, model(rwalk)
searching for initial values ..........
(setting technique to bhhh)
Iteration 0:  Log likelihood = -10026.649
Iteration 1:  Log likelihood = -9947.9671
Iteration 2:  Log likelihood = -9896.4778
Iteration 3:  Log likelihood = -9890.8199
Iteration 4:  Log likelihood = -9890.3202
(switching technique to nr)
Iteration 5:  Log likelihood = -9890.1546
Iteration 6:  Log likelihood =  -9889.561
Iteration 7:  Log likelihood = -9889.5608
Refining estimates:
Iteration 0:  Log likelihood = -9889.5608
Iteration 1:  Log likelihood = -9889.5608

Unobserved-components model
Components: random walk
Sample: 21jan1967 thru 19feb2011                 Number of obs =   2,301
                                                 Wald chi2(2)  = 271.88
Log likelihood = -9889.5608                      Prob > chi2   = 0.0000
```

icsa	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
icsa						
L1.	-.3250633	.0205148	-15.85	0.000	-.3652715	-.2848551
L2.	-.1794686	.0205246	-8.74	0.000	-.2196961	-.1392411
var(level)	317.6474	9.36691	33.91	0.000	299.2886	336.0062

Note: Model is not stationary.
Note: Tests of variances against zero are one sided, and the two-sided
      confidence intervals are truncated at zero.
Note: Time units are in 7 days.

Now we use arima to estimate the parameters of an asymptotically equivalent ARIMA(2,1,0) model. (We specify the technique(nr) option so that arima will compute the observed information matrix standard errors that ucm computes.) We use nlcom to compute a point estimate and a standard error for the variance, which is directly comparable to the one produced by ucm.

```
. arima icsa, noconstant arima(2,1,0) technique(nr)

Iteration 0:   Log likelihood = -9896.4584
Iteration 1:   Log likelihood =  -9896.458

ARIMA regression
Sample: 14jan1967 thru 19feb2011            Number of obs    =       2302
                                            Wald chi2(2)     =     271.95
Log likelihood = -9896.458                  Prob > chi2      =     0.0000
```

		OIM				
D.icsa	Coefficient	std. err.	z	P>|z|	[95% conf.	interval]
ARMA						
ar						
L1.	-.3249383	.0205036	-15.85	0.000	-.3651246	-.284752
L2.	-.1793353	.0205088	-8.74	0.000	-.2195317	-.1391388
/sigma	17.81606	.2625695	67.85	0.000	17.30143	18.33068

```
Note: The test of the variance against zero is one sided, and the two-sided
      confidence interval is truncated at zero.

. nlcom _b[sigma:_cons]^2

      _nl_1: _b[sigma:_cons]^2
```

D.icsa	Coefficient	Std. err.	z	P>|z|	[95% conf.	interval]
_nl_1	317.4119	9.355904	33.93	0.000	299.0746	335.7491

It is no accident that the parameter estimates and the standard errors from the two estimators are so close. As the sample size grows the differences in the parameter estimates and the estimated standard errors will go to 0, because the two estimators are equivalent in large samples.

◁

Models for the trend and idiosyncratic components

A general model that allows for fixed or stochastic trends in τ_t is given by

$$\tau_t = \tau_{t-1} + \beta_{t-1} + \eta_t \tag{2}$$
$$\beta_t = \beta_{t-1} + \xi_t \tag{3}$$

Following Harvey (1989), we define 11 flexible models for y_t that specify both τ_t and ϵ_t in (1). These models place restrictions on the general model specified in (2) and (3) and on ϵ_t in (1). In other words, these models jointly specify τ_t and ϵ_t.

To any of these models, a cyclical component, a seasonal component, or exogenous variables may be added.

Table 1. Models for the trend and idiosyncratic components

Model name	Syntax option	Model
No trend or idiosyncratic component	model(none)	
No trend	model(ntrend)	$y_t = \epsilon_t$
Deterministic constant	model(dconstant)	$y_t = \mu + \epsilon_t$ $\mu = \mu$
Local level	model(llevel)	$y_t = \mu_t + \epsilon_t$ $\mu_t = \mu_{t-1} + \eta_t$
Random walk	model(rwalk)	$y_t = \mu_t$ $\mu_t = \mu_{t-1} + \eta_t$
Deterministic trend	model(dtrend)	$y_t = \mu_t + \epsilon_t$ $\mu_t = \mu_{t-1} + \beta$ $\beta = \beta$
Local level with deterministic trend	model(lldtrend)	$y_t = \mu_t + \epsilon_t$ $\mu_t = \mu_{t-1} + \beta + \eta_t$ $\beta = \beta$
Random walk with drift	model(rwdrift)	$y_t = \mu_t$ $\mu_t = \mu_{t-1} + \beta + \eta_t$ $\beta = \beta$
Local linear trend	model(lltrend)	$y_t = \mu_t + \epsilon_t$ $\mu_t = \mu_{t-1} + \beta_{t-1} + \eta_t$ $\beta_t = \beta_{t-1} + \xi_t$
Smooth trend	model(strend)	$y_t = \mu_t + \epsilon_t$ $\mu_t = \mu_{t-1} + \beta_{t-1}$ $\beta_t = \beta_{t-1} + \xi_t$
Random trend	model(rtrend)	$y_t = \mu_t$ $\mu_t = \mu_{t-1} + \beta_{t-1}$ $\beta_t = \beta_{t-1} + \xi_t$

The majority of the models available in ucm are designed for nonstationary time series. The deterministic-trend model incorporates a first-order deterministic time-trend in the model. The local-level, random-walk, local-level-with-deterministic-trend, and random-walk-with-drift models are for modeling series with first-order stochastic trends. A series with a dth-order stochastic trend must be differenced d times to be stationary. The local-linear-trend, smooth-trend, and random-trend models are for modeling series with second-order stochastic trends.

The no-trend-or-idiosyncratic-component model is useful for using ucm to model stationary series with cyclical components or seasonal components and perhaps exogenous variables. The no-trend and the deterministic-constant models are useful for using ucm to model stationary series with seasonal components or exogenous variables.

Seasonal component

A seasonal component models cyclical behavior in a time series that occurs at known seasonal periodicities. A seasonal component is modeled in the time domain; the period of the cycle is specified as the number of time periods required for the cycle to complete.

▷ Example 8

Let's begin by considering a series that displays a seasonal effect. Below we plot a monthly series containing the number of new cases of mumps in New York City between January 1928 and December 1972. (See Hipel and McLeod [1994] for the source and further discussion of this dataset.)

```
. use https://www.stata-press.com/data/r19/mumps, clear
(Hipel and Mcleod (1994), http://robjhyndman.com/tsdldata/epi/mumps.dat)

. tsline mumps
```


The graph reveals recurring spikes at regular intervals, which we suspect to be seasonal effects. The series may or may not be stationary; the graph evidence is not definitive.

Deterministic seasonal effects are a standard method of incorporating seasonality into a model. In a model with a constant term, the s deterministic seasonal effects are modeled as s parameters subject to the constraint that they sum to zero; formally, $\gamma_t + \gamma_{t-1} + \cdots + \gamma_{t-(s-1)} = 0$. A stochastic-seasonal model is a more flexible alternative that allows the seasonal effects at time t to sum to ζ_t, a zero-mean, finite-variance, i.i.d. random variable; formally, $\gamma_t + \gamma_{t-1} + \cdots + \gamma_{t-(s-1)} = \zeta_t$.

In the output below, we model the seasonal effects by a stochastic-seasonal model, we allow for the series to follow a random walk, and we include a stationary cyclical component.

```
. ucm mumps, seasonal(12) cycle(1)
searching for initial values ...................
(setting technique to bhhh)
Iteration 0:  Log likelihood = -3270.0908
Iteration 1:  Log likelihood = -3257.6792
Iteration 2:  Log likelihood = -3257.3017
Iteration 3:  Log likelihood = -3249.8741
Iteration 4:  Log likelihood = -3249.4774
(switching technique to nr)
Iteration 5:  Log likelihood = -3248.9154
Iteration 6:  Log likelihood = -3248.7233
Iteration 7:  Log likelihood = -3248.7138
Iteration 8:  Log likelihood = -3248.7138
Refining estimates:
Iteration 0:  Log likelihood = -3248.7138
Iteration 1:  Log likelihood = -3248.7138

Unobserved-components model
Components: random walk, seasonal(12), order 1 cycle
```

Sample: 1928m1 thru 1972m6

```
Log likelihood = -3248.7138
```

Number of obs =   534
Wald chi2(2) = 2141.69
Prob > chi2   = 0.0000

mumps	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
frequency	.3863607	.0282037	13.70	0.000	.3310824	.4416389
damping	.8405622	.0197933	42.47	0.000	.8017681	.8793563
var(level)	221.2131	140.5179	1.57	0.058	0	496.6231
var(seasonal)	4.151639	4.383442	0.95	0.172	0	12.74303
var(cycle1)	12228.17	813.8394	15.03	0.000	10633.08	13823.27

```
Note: Model is not stationary.
Note: Tests of variances against zero are one sided, and the two-sided
      confidence intervals are truncated at zero.
```

The output indicates that the trend and seasonal variances may not be necessary. When the variance of the seasonal component is zero, the seasonal component becomes deterministic. Below we estimate the parameters of a model that includes deterministic seasonal effects and a stationary cyclical component.

```
. ucm mumps ibn.month, model(none) cycle(1)
searching for initial values .......
(setting technique to bhhh)
Iteration 0:  Log likelihood = -4138.3988
Iteration 1:  Log likelihood = -3706.3937
Iteration 2:  Log likelihood =  -3533.649
Iteration 3:  Log likelihood = -3426.8902
Iteration 4:  Log likelihood = -3392.0083
(switching technique to nr)
Iteration 5:  Log likelihood = -3374.3692
Iteration 6:  Log likelihood = -3284.0435
Iteration 7:  Log likelihood = -3283.0343
Iteration 8:  Log likelihood = -3283.0284
Iteration 9:  Log likelihood = -3283.0284
Refining estimates:
Iteration 0:  Log likelihood = -3283.0284
Iteration 1:  Log likelihood = -3283.0284
```

```
Unobserved-components model
Components: order 1 cycle
Sample: 1928m1 thru 1972m6                          Number of obs =      534
                                                    Wald chi2(14) = 3404.29
Log likelihood = -3283.0284                         Prob > chi2   =  0.0000
```

mumps	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
cycle1						
frequency	.3272754	.0262922	12.45	0.000	.2757436	.3788071
damping	.844874	.0184994	45.67	0.000	.8086157	.8811322
mumps						
month						
1	480.5095	32.67128	14.71	0.000	416.475	544.544
2	561.9174	32.66999	17.20	0.000	497.8854	625.9494
3	832.8666	32.67696	25.49	0.000	768.8209	896.9122
4	894.0747	32.64568	27.39	0.000	830.0904	958.0591
5	869.6568	32.56282	26.71	0.000	805.8348	933.4787
6	770.1562	32.48587	23.71	0.000	706.4851	833.8274
7	433.839	32.50165	13.35	0.000	370.1369	497.541
8	218.2394	32.56712	6.70	0.000	154.409	282.0698
9	140.686	32.64138	4.31	0.000	76.7101	204.662
10	148.5876	32.69067	4.55	0.000	84.51508	212.6601
11	215.0958	32.70311	6.58	0.000	150.9989	279.1927
12	330.2232	32.68906	10.10	0.000	266.1538	394.2926
var(cycle1)	13031.53	798.2719	16.32	0.000	11466.95	14596.11

Note: Tests of variances against zero are one sided, and the two-sided
      confidence intervals are truncated at zero.

The output indicates that each of these components is statistically significant.

◁

❑ Technical note

In a stochastic model for the seasonal component, the seasonal effects sum to the random variable $\zeta_t \sim$ i.i.d. $N(0, \sigma_\zeta^2)$:

$$\gamma_t = -\sum_{j=1}^{s-1} \gamma_{t-j} + \zeta_t$$

❑

Stored results

Because ucm is estimated using sspace, most of the sspace stored results appear after ucm. Not all of these results are relevant for ucm; programmers wishing to treat ucm results as sspace results should see *Stored results* of [TS] **sspace**. See *Methods and formulas* for the state-space representation of UCMs, and see [TS] **sspace** for more documentation that relates to all the stored results.

ucm stores the following in e():

Scalars
e(N)	number of observations
e(k)	number of parameters
e(k_eq)	number of equations in e(b)
e(k_dv)	number of dependent variables
e(k_cycles)	number of stochastic cycles
e(df_m)	model degrees of freedom
e(ll)	log likelihood
e(chi2)	χ^2
e(p)	p-value for model test
e(tmin)	minimum time in sample
e(tmax)	maximum time in sample
e(stationary)	1 if the estimated parameters indicate a stationary model, 0 otherwise
e(rank)	rank of VCE
e(ic)	number of iterations
e(rc)	return code
e(converged)	1 if converged, 0 otherwise

Macros
e(cmd)	ucm
e(cmdline)	command as typed
e(depvar)	unoperated names of dependent variables in observation equations
e(covariates)	list of covariates
e(tvar)	variable denoting time within groups
e(eqnames)	names of equations
e(model)	type of model
e(title)	title in estimation output
e(tmins)	formatted minimum time
e(tmaxs)	formatted maximum time
e(chi2type)	Wald; type of model χ^2 test
e(vce)	*vcetype* specified in vce()
e(vcetype)	title used to label Std. err.
e(opt)	type of optimization
e(initial_values)	type of initial values
e(technique)	maximization technique
e(tech_steps)	iterations taken in maximization technique
e(properties)	b V
e(estat_cmd)	program used to implement estat
e(predict)	program used to implement predict
e(marginsnotok)	predictions disallowed by margins
e(asbalanced)	factor variables fvset as asbalanced
e(asobserved)	factor variables fvset as asobserved

Matrices
e(b)	parameter vector
e(Cns)	constraints matrix
e(ilog)	iteration log (up to 20 iterations)
e(gradient)	gradient vector
e(V)	variance–covariance matrix of the estimators
e(V_modelbased)	model-based variance

Functions
e(sample)	marks estimation sample

In addition to the above, the following is stored in r():

Matrices
r(table)	matrix containing the coefficients with their standard errors, test statistics, p-values, and confidence intervals

Note that results stored in r() are updated when the command is replayed and will be replaced when any r-class command is run after the estimation command.

Methods and formulas

Methods and formulas are presented under the following headings:

Introduction

The general form of UCMs can be expressed as

$$y_t = \tau_t + \gamma_t + \psi_t + \mathbf{x}_t \boldsymbol{\beta} + \epsilon_t$$

where τ_t is the trend, γ_t is the seasonal component, ψ_t is the cycle, $\boldsymbol{\beta}$ is the regression coefficients for regressors $\mathbf{x}_t$, and ϵ_t is the idiosyncratic error with variance σ_ϵ^2.

We can decompose the trend as

$$\tau_t = \mu_t$$
$$\mu_t = \mu_{t-1} + \alpha_{t-1} + \eta_t$$
$$\alpha_t = \alpha_{t-1} + \xi_t$$

where μ_t is the local level, α_t is the local slope, and η_t and ξ_t are i.i.d. normal errors with mean 0 and variance σ_η^2 and σ_ξ^2, respectively.

Next consider the seasonal component, γ_t, with a period of s time units. Ignoring a seasonal disturbance term, the seasonal effects will sum to zero, $\sum_{j=0}^{s-1} \gamma_{t-j} = 0$. Adding a normal error term, ω_t, with mean 0 and variance σ_ω^2, we express the seasonal component as

$$\gamma_t = -\sum_{j=1}^{s-1} \gamma_{t-j} + \omega_t$$

Finally, the cyclical component, ψ_t, is a function of the frequency λ, in radians, and a unit-less scaling variable ρ, termed the damping effect, $0 < \rho < 1$. We require two equations to express the cycle:

$$\psi_t = \psi_{t-1}\rho\cos\lambda + \tilde{\psi}_{t-1}\rho\sin\lambda + \kappa_t$$
$$\tilde{\psi}_t = -\psi_{t-1}\rho\sin\lambda + \tilde{\psi}_{t-1}\rho\cos\lambda + \tilde{\kappa}_t$$

where the κ_t and $\tilde{\kappa}_t$ disturbances are normally distributed with mean 0 and variance σ_κ^2.

The disturbance terms ϵ_t, η_t, ξ_t, ω_t, κ_t, and $\tilde{\kappa}_t$ are independent.

State-space formulation

ucm is an easy-to-use implementation of the state-space command sspace, with special modifications, where the local linear trend components, seasonal components, and cyclical components are states of the state-space model. The state-space model can be expressed in matrix form as

$$\mathbf{y}_t = \mathbf{Dz}_t + \mathbf{Fx}_t + \epsilon_t$$
$$\mathbf{z}_t = \mathbf{Az}_{t-1} + \mathbf{C}\boldsymbol{\zeta}_t$$

where y_t, $t = 1, \ldots, T$, are the observations and $\mathbf{z}_t$ are the unobserved states. The number of states, m, depends on the model specified. The $k \times 1$ vector $\mathbf{x}_t$ contains the exogenous variables specified as *indepvars*, and the $1 \times k$ vector $\mathbf{F}$ contains the regression coefficients to be estimated. ϵ_t is the observation equation disturbance, and the $m_0 \times 1$ vector $\boldsymbol{\zeta}_t$ contains the state equation disturbances, where $m_0 \leq m$. Finally, $\mathbf{C}$ is a $m \times m_0$ matrix of zeros and ones. These recursive equations are evaluated using the diffuse Kalman filter of De Jong (1991).

Below we give the state-space matrix structures for a local linear trend with a stochastic seasonal component, with a period of 4 time units, and an order-2 cycle. The state vector, $\mathbf{z}_t$, and its transition matrix, $\mathbf{A}$, have the structure

$$\mathbf{A} = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & -1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \rho\cos\lambda & \rho\sin\lambda & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & -\rho\sin\lambda & \rho\cos\lambda & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \rho\cos\lambda & \rho\sin\lambda \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\rho\sin\lambda & \rho\cos\lambda \end{pmatrix} \quad \mathbf{z}_t = \begin{pmatrix} \mu_t \\ \alpha_t \\ \gamma_t \\ \gamma_{t-1} \\ \gamma_{t-2} \\ \psi_{t,1} \\ \psi_{t,1} \\ \psi_{t,2} \\ \psi_{t,2} \end{pmatrix}$$

$$\mathbf{C} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} \quad \boldsymbol{\zeta}_t = \begin{pmatrix} \eta_t \\ \xi_t \\ \omega_t \\ \kappa_t \\ \tilde{\kappa}_t \end{pmatrix}$$

$$\mathbf{D} = (1 \quad 0 \quad 1 \quad 0 \quad 0 \quad 1 \quad 0 \quad 0 \quad 0)$$

Cyclical component extensions

Recall that the stochastic cyclical model is given by

$$\psi_t = \rho(\psi_{t-1}\cos\lambda_c + \psi_{t-1}^*\sin\lambda_c) + \kappa_{t,1}$$
$$\psi_t^* = \rho(-\psi_{t-1}\sin\lambda_c + \psi_{t-1}^*\cos\lambda_c) + \kappa_{t,2}$$

where $\kappa_{t,j} \sim$ i.i.d. $N(0, \sigma_\kappa^2)$ and $0 < \rho < 1$ is a damping effect. The cycle is variance-stationary when $\rho < 1$ because $\text{Var}(\psi_t) = \sigma_\kappa^2/(1 - \rho)$. We will express a UCM with a cyclical component added to a trend as

$$y_t = \mu_t + \psi_t + \epsilon_t$$

where μ_t can be any of the trend parameterizations discussed earlier.

Higher-order cycles, $k = 2$ or $k = 3$, are defined as

$$\psi_{t,j} = \rho(\psi_{t-1,j} \cos \lambda_c + \psi_{t-1,j}^* \sin \lambda_c) + \psi_{t-1,j+1}$$
$$\psi_{t,j}^* = \rho(-\psi_{t-1,j} \sin \lambda_c + \psi_{t-1,j}^* \cos \lambda_c) + \psi_{t-1,j+1}^*$$

for $j < k$, and

$$\psi_{t,k} = \rho(\psi_{t-1,k} \cos \lambda_c + \psi_{t-1,k}^* \sin \lambda_c) + \kappa_{t,1}$$
$$\psi_{t,k}^* = \rho(-\psi_{t-1,k} \sin \lambda_c + \psi_{t-1,k}^* \cos \lambda_c) + \kappa_{t,2}$$

Harvey and Trimbur (2003) discuss the properties of this model and its state-space formulation.

Andrew Charles Harvey (1947–) is a British econometrician. After receiving degrees in economics and statistics from the University of York and the London School of Economics and working for a period in Kenya, he has worked as a teacher and researcher at the University of Kent, the London School of Economics, and now the University of Cambridge. Harvey's interests are centered on time series, especially state-space models, signal extraction, volatility, and changes in quantiles.

References

Baum, C. F., and S. Hurn. 2021. *Environmental Econometrics Using Stata*. College Station, TX: Stata Press.

Burns, A. F., and W. C. Mitchell. 1946. *Measuring Business Cycles*. New York: National Bureau of Economic Research.

De Jong, P. 1991. The diffuse Kalman filter. *Annals of Statistics* 19: 1073–1083. https://doi.org/10.1214/aos/1176348139.

Durbin, J., and S. J. Koopman. 2012. *Time Series Analysis by State Space Methods*. 2nd ed. Oxford: Oxford University Press.

Fuller, W. A. 1996. *Introduction to Statistical Time Series*. 2nd ed. New York: Wiley.

Hamilton, J. D. 1994. *Time Series Analysis*. Princeton, NJ: Princeton University Press. https://doi.org/10.2307/j.ctv14jx6sm.

Harvey, A. C. 1989. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge: Cambridge University Press.

———. 1993. *Time Series Models*. 2nd ed. Cambridge, MA: MIT Press.

Harvey, A. C., and T. M. Trimbur. 2003. General model-based filters for extracting cycles and trends in economic time series. *Review of Economics and Statistics* 85: 244–255. https://doi.org/10.1162/003465303765299774.

Hipel, K. W., and A. I. McLeod. 1994. *Time Series Modelling of Water Resources and Environmental Systems*. Amsterdam: Elsevier.

Priestley, M. B. 1981. *Spectral Analysis and Time Series*. London: Academic Press.

Trimbur, T. M. 2006. Properties of higher order stochastic cycles. *Journal of Time Series Analysis* 27: 1–17. https://doi.org/10.1111/j.0143-9782.2005.00462.x.

Wei, W. W. S. 2006. *Time Series Analysis: Univariate and Multivariate Methods*. 2nd ed. Boston: Pearson.

Also see