Chapter 9
Idiographic Data Analysis: Quantitative Methods—From Simple to Advanced

Ellen L. Hamaker and Conor V. Dolan

Time series analysis is a technique by which a large number of repeated measures taken from a single case can be modeled. As it requires observations from only one case, this is a useful technique for researchers interested in idiographic data analysis. The most basic time series technique is the well-known autoregressive moving average (ARMA) model (Box & Jenkins, 1970; Chatfield, 2004; Hamilton, 1994). It combines the AR model and the MA model, both of which were separately invented in 1927 to handle the autocorrelation typically observed in time series data (Tong, 2001). Characteristic of the AR model is that the current observation is predicted from previous observations (Box & Jenkins, 1970; Granger & Morris, 1976). The part of an observation that cannot be predicted based on previous observations is called the random shock, residual, or innovation. In contrast, the MA model consists of predicting the current observation from a weighted sum of previous random shocks (Box & Jenkins, 1970; Granger & Morris, 1976). Combining these two models resulted in the ARMA model, which gained widespread popularity through the 1970 book *Time series analysis: Forecasting and control* by Box and Jenkins.

This chapter provides a brief tour of the original ARMA model and some of its most popular extensions, which were developed in econometrics and other fields that rely heavily on time series analysis. To emphasize the potential of ARMA-based modeling for the social sciences, we include references to applications within psychology, sociology and criminology that illustrate the use and interpretation of these models. We do not focus on how to implement these models, nor will we discuss issues related to model estimation and evaluation, but the interested reader is referred to standard introductory texts such as Hamilton (1994), Chatfield (2004), Durbin and Koopman (2001), Harvey (1989), and Fan and Yao (2003).

In the following four sections we present the basic ARMA model and its extensions: Section "ARIMA Models" is on the building blocks of the integrated ARMA

E. L. Hamaker (✉)
Department of Methods and Statistics
Faculty of Social Sciences
Utrecht University
PO Box 80140
3508 TC Utrecht, The Netherlands
e-mail: e.l.hamaker@uu.nl

J. Valsiner et al. (eds.), *Dynamic Process Methodology in the Social and Developmental Sciences*, 191
DOI 10.1007/978-0-387-95922-1_9,

(ARIMA) model; Section “Univariate Extensions of the ARIMA model” includes, ARMA models with deterministic trends and cycles, seasonal ARIMA models, fractionally integrated ARMA models, and impact ARIMA models; Section “Multivariate Extensions of ARMA model” includes the vector ARMA (VARMA) model, VARMA models with exogenous variables, latent VARMA models, and the cointegrated model; and Section “Nonlinear Extensions of the ARMA model” includes the bilinear model, the conditional heteroscedastic model, the threshold AR model, and the Markov-switching AR model. Each section ends with a discussion of applications of these techniques in the social sciences.

For all models discussed in this chapter it is assumed that the data are measured at interval or ratio level, and that observations are made at equal time intervals. However, at the end of this chapter we briefly mention some alternative techniques that are not based on these assumptions. Another important assumption for some of the models discussed in this chapter is stationarity. Assuming Gaussian data, stationarity implies that the mean, variance and autocovariances[1] of the series are independent of time. The basic ARMA model is based on the assumption that the data are stationary, but many of its extensions are nonstationary (e.g., the ARIMA model). In what follows we consistently indicate whether certain processes are stationary or not.

ARIMA Models

In this section we introduce the building blocks of the general ARIMA model, that is: (a) the AR model; (b) the MA model; (c) the mixed ARMA model; and finally (d) the full ARIMA model.

Autoregressive (AR) Model

Let y_t be a univariate observation at occasion t. In the most simple version of the AR model, the AR (1), the observation y_t can be predicted from the previous observation y_{t-1}. This can be represented as

$$y_t = \phi_0 + \phi_1 y_{t-1} + u_t, \tag{1}$$

where ϕ_0 is a constant, ϕ_1 is the AR parameter, that is, it is the regression coefficient in the regression of y_t on y_{t-1}, and u_t is the part of y_t that could not be predicted from y_{t-1}, and which is referred to as the innovation, residual, prediction error, or random shock. As the innovation at occasion t is the part of y_t that is independent of the

[1] The autocovariance is the covariance between y_t and y_{t+k}, that is, $E[(y_t-\mu)(y_{t+k}-\mu)]$, where μ is the mean of the series. The lag k is the distance in time. When $k = 0$, we obtain the variance of the series. The autocorrelation at lag k can be obtained by dividing the autocovariance at lag k by the variance of the series.

observations before occasion t, it is also independent of the innovations u prior to and after occasion t. Such a sequence is referred to as a white noise sequence. The mean of this sequence is zero and its variance is denoted as σ_u^2. To ensure that the AR model in Eq. (1) is stationary, the parameter ϕ_1 has to lie between −1 and 1. If this restriction is violated, the variance of the process will increase over time. It can be shown that if $|\phi_1| < 1$, the mean of the observed series is $\mu_y = \phi_0/(1-\phi_1)$, and its variance is $\sigma_y^2 = \sigma_u^2/(1-\phi_1^2)$ (Chatfield, 2004).

A general expression of the AR model of order p (i.e., AR (p)) is

$$y_t = \phi_0 + \phi_1 y_{t-1} + \phi_2 y_{t-2} + \cdots + \phi_p y_{t-p} + ut, \tag{2}$$

where ϕ_0 is a constant, ϕ_1 to ϕ_p are the AR parameters in the regression of y_t on y_{t-1} to y_{t-p}, and u_t is the innovation. For such higher order AR processes the stationarity restrictions are quite complicated (see Hamilton, 1989, pp. 27–33). If the process is stationary, the mean can be shown to equal $\mu_y = \phi_0/(1-\phi_1-\ldots\phi_p)$. The expression for the variance of a pure AR process in terms of the variance of the innovations and the AR parameters is given in Hamilton (1994, p. 59).

Granger and Morris (1976) indicated that an AR process can be interpreted as a momentum effect in a random variable. To illustrate this, suppose we are driving down the freeway at a speed of 70 miles/h. If we measure the exact speed at different occasions, we will find that the speed is not exactly 70 miles/h every time, but it actually fluctuates around this value. If we measure our speed once every minute, we will probably find no sequential relationship between successive observations. However, if we measure speed every five seconds, there probably will be some sequential dependency, simply because the variable speed needs *time* to change. If the interval between measurement occasions grows smaller, we will find a stronger sequential relationship between successive observations. In general we can state that the sequential dependency of a variable that is continuous over time (such as our speed), depends on the intervals between observations.

Moving Average (MA) Model

If an observation y_t can be predicted by the unpredictable parts at previous occasions, we have an MA process. It implies that the observation y_t is a weighted sum of two or more innovations. An MA process of order one, denoted as an MA (1), can be expressed as

$$y_t = \mu_y + u_t - \theta_1 u_{t-1}, \tag{3}$$

where u_t is a white noise sequence, μ_y is the mean of the observed series, and $-\theta_1$ is the MA parameter by which the innovation of the previous occasion is weighted.[2]

[2] In some texts $-\theta_1$ is replaced by Ψ_1, such that the minus sign is omitted. However, the above notation is more conventional, as it has some important advantages for the expression of particular characteristics of an MA process.

We can also say that $-\theta_1$ is the parameter that is used to regress the observation y_t upon the unpredictable part of the previous observation, that is, u_{t-1}. The variance of the observed series can be shown to equal $\sigma_y^2 = (1 + \theta_1^2)\sigma_u^2$ (Chatfield, 2004).

The general expression for an MA process of order q is

$$y_t = \mu_y + u_t - \theta_1 u_{t-1} - \theta_2 u_{t-2} - \cdots - \theta_q u_{t-q}, \tag{4}$$

where μ_y is the mean of the observed series, $-\theta_1$ to $-\theta_q$ are the parameters by which the previous innovations u_{t-1} to u_{t-q} are weighted. The variance of this process is equal to $(1+\theta_1^2+\cdots+\theta_q^2)\sigma_u^2$ (Chatfield, 2004).

Pure MA processes are by definition stationary. However, there are restrictions necessary to ensure the model is invertible, which implies that it can be rewritten as an AR model (we elaborate on this below). These restriction are analogous to the restriction on the AR parameters to ensure stationarity. For an MA (1), this implies that Q_x must lie between −1 and 1 (see Hamilton, 1994, p. 67, for invertibility restrictions for higher order MA processes).

Granger and Morris (1976) described an MA process as involving a variable in equilibrium, which is buffeted by a sequence of unpredictable events with a delayed or discounted effect. Hence, the innovation u_t is interpreted as being due to events or circumstances that influence the variable under investigation y_t. To illustrate this, suppose we ask an individual repeatedly to answer the question how good (s)he feels today. The score y_t at a certain day is influenced by the circumstances that day u_t (e.g., attending a party, getting some good news, having a disagreement with a good friend), but it may also depend on the events that took place in the recent past, i.e., u_{t-1} to u_{t-q}.

Mixed Autoregressive Moving Average (ARMA) Models

The two processes described above can also be combined, resulting in an ARMA (p, q) process. The general expression for such a process is

$$\begin{aligned} y_t &= \phi_0+\phi_1 y_{t-1}+\cdots+\phi_p y_{t-p}+u_t-\theta_1 u_{t-1}-\cdots-\theta_q u_{t-q} \\ &= \phi_0+\sum_{j=1}^{p}\phi_j y_{t-j}+u_t-\sum_{j=1}^{q}\theta_j u_{t-j}, \end{aligned} \tag{5}$$

where the innovation u_t is a white noise sequence. Hence, y_t is a weighted sum of previous observations, going back to y_{t-p}, and previous innovations, going back to u_{t-q}. To ensure stationarity, the same restrictions apply to the AR parameters as in the pure AR model. Similarly, to ensure invertibility, the same restriction apply to the MA parameters as in the pure MA model. The expression of the mean of the observed series is the same as for a pure AR model, but the expression for the variance is more complicated (see Hamilton, 1994, pp. 61–63).

A mixed ARMA process is difficult to interpret in substantive terms. However, each AR process of finite order can be rewritten as an MA process of infinite order, i.e., an MA (∞). Conversely, each invertible MA process of finite order can be represented as an AR process of infinite order. Moreover, each mixed ARMA process of finite orders p and q, can be rewritten as either a pure AR process of infinite order, or a pure MA process of infinite order. This implies that the differences between pure AR, pure MA, and mixed ARMA models are not absolute, which in turn gives rise to the question how to choose between these different representations.[3]

In the Box and Jenkins approach the aim is forecasting and control, and the interpretation of the parameters is mainly in terms of predictive relations (Box & Jenkins, 1970). Hence, in this context it makes sense to find the model with the minimum number of parameters. For the social scientist often the substantive interpretation is more important than forecasting, and from this perspective pure AR and pure MA models may be preferable over mixed ARMA models.

Yet another interesting relationship between ARMA models was published by Granger and Morris (1976), who showed that mixed ARMA processes can arise from summing independent, stationary processes. For instance, summing two AR (1) processes results in an ARMA (2, 1) process, and adding a white noise sequence to an AR (p) process results in an ARMA (p, p). Although it is not possible to disentangle the original processes that have given rise to a mixed ARMA processes, mixed processes may be interpreted in terms of a summation of pure processes. Specifically, models in which white noise is added to a pure AR process are compatible with the idea of noisy measurements: It would imply that there is both measurement error (i.e., the white noise sequence), and prediction error (i.e., the unpredictable part in the AR process), which are two separate sources of variation.

Integrated Autoregressive Moving Average (ARIMA) Model

All the models discussed above are stationary, meaning that the mean, variance and autocovariances are invariant over time. A special class of nonstationary models is formed by the integrated models. Characteristic of an integrated process is that it becomes stationary after differencing it, meaning the previous observation is subtracted from the current observation. Thus, while y_t is nonstationary, $\Delta y_t = y_t - y_{t-1}$ is stationary. Sometimes, differencing needs to be carried out multiple times to obtain a stationary series. If differencing the data once results in stationarity, the process is said to have a unit root (cf., Hamilton, 1994), and it may be referred to as an I (1) process.

A simple example of an integrated process is an ARIMA (0, 1, 0) model, which is also referred to as a random walk, that is

$$y_t = y_{t-1} + u_t. \tag{6}$$

[3] In practice, a process of infinite order is not appealing, as there will be more parameters to estimate than observations. However, in finite samples, the parameters beyond a certain lag will be insignificant and can be omitted from the model. The important issue is that there are no fundamental differences between these processes.

While y_t is nonstationary, differencing it results in $\Delta y_t = y_t - y_{t-1} = u_t$, which is a stationary process. What is typical for a random walk is that while the mean is independent of time (i.e., $\mathrm{E}[y_t] = 0$ for all t), the variance is ever increasing over time (i.e., $\mathrm{Var}[y_t] \to \infty$ as $t \to \infty$).

A unit root process for which the differenced series obeys a stationary ARMA model is denoted as an ARIMA (p, 1, q), that is

$$y_t = y_{t-1} + z_t \tag{7a}$$

$$z_t = \phi_1 z_{t-1} + \cdots + \phi_p z_{t-p} + u_t - \theta_1 u_{t-1} - \cdots - \theta_q z_{t-q}, \tag{7b}$$

where z_t is the differenced series $\Delta y_t = y_t - y_{t-1}$, and is a stationary ARMA process. Another unit root process which is applied frequently in econometrics is the random walk with drift. This model can be represented as

$$y_t = \delta + y_{t-1} + u_t, \tag{8}$$

where δ is referred to as the drift, or the stochastic trend (as opposed to a deterministic trend which is discussed in the following section). If $\delta > 0$, y_t tends to increase over time, while if $\delta < 0$, y_t tends to decrease.

The interpretation of unit root processes focuses on the difference scores Δy_t, which can be modeled as an ARMA (p, q) process. For instance, if a researcher finds that an ARIMA (1, 1, 0) fits the data, this can be interpreted as meaning that the change from the previous occasion to the current occasion (i.e., Δy_t), can be predicted from the change that took place right before that (Δy_{t-1}).

Applications in the Social Sciences

Many applications of ARMA and ARIMA modeling in the social sciences serve the purpose of prewhitening the data, which implies the data are transformed into a white noise series. The goal of prewhitening in these applications is to determine whether there are indications for causal relationships between two or more series, while controling for autocorrelation due to AR-, I-, and MA-components. It is well known that failing to account for such autocorrelation in the univariate series may result in spurious relationships between the series. While this procedure is still used today, there are multivariate extensions of the ARMA model which allow for the simultaneous modeling of the ARMA relations, and the mutual effects. Moreover, differencing may remove important information about the long run relationship between two or more series (e.g., in the case of cointegration, see below).

Applications of ARIMA modeling as a prewhitening technique in the social sciences have been used relatively often to establish a relationship between aggregate time series, such as the alcohol consumption per capita and suicide or criminal violence rates (e.g., Bye, 2007; Razvodovsky, 2007). Bye (2007) for instance, concluded that there was evidence for a causal effect of alcohol consumption on violence. An example of prewhitening in psychological research is the study done by

Andersson and Yardley (2000), who investigated the relationship between the prewhitened measures of dizziness and physical, mental, and emotional stress. They found evidence for concurrent relations mainly, although two of the ten participants were characterized by an increase in stress (either mental or emotional) prior to increases in dizziness. In another study, Andersson, Hägnebo, and Yardley (1997) used prewhitening to study the relationship between stress and symptoms associated with Meniere disease.

There are some studies in which ARIMA modeling was not used merely as a prewhitening device, but rather as a procedure to unveil the dynamics underlying the observed series. In particular, Fortes, Delintnières, and Ninot (2004) used the ARIMA (0, 1, 1) model as a means to understand the balance between two opposite forces: preservation and adaption. Let $\hat{y}_t = y_t - u_t$ be the expectation (i.e., the predictable part) of y_t. Since $y_t = y_{t-1} + u_t - \theta_1 u_{t-1}$, we can also write

$$\begin{aligned}\hat{y}_t &= y_{t-1} - \theta_1 u_{t-1} \\ &= \hat{y}_{t-1} + u_{t-1} - \theta_1 u_{t-1}.\end{aligned}$$

From the latter expression it becomes clear that if the MA coefficient θ_1 is close to 1, this serves as a restoring mechanism, in which the expectation at occasion t is close to the expectation at occasion $t - 1$. Such a process may be interpreted as a form of preservation, meaning there is resistance to the influence of temporal effects (Fortes et al., 2004). In contrast, an MA coefficient further away from 1 implies the expectation changes, as the expectation at t is inflected by the innovation. The latter is more indicative of adaption to change, in which temporal disturbances tend to leave a persistent trace in the data (Fortes et al., 2004).

Fortes et al. (2004) apply ARIMA modeling to the data obtained from seven individuals on six variables related to self-esteem and physical self, and concluded that for 35 of the 42 series an ARIMA (0, 1, 1) model was the most appropriate model. ARIMA (0, 1, 1) models were also used by Peterson and Leckman (1998), who measured inter-tic interval in patients with Gilles de la Tourette syndrome, and investigated the temporal patterning of tics. They concluded that the tics intervals are nonstationary. In addition, the change in tic intervals oscillates rapidly, with large changes followed by small ones and vice versa. Note however that this application differs from usual ARIMA applications, which are based on observations made at equal intervals.

Univariate Extensions of the ARIMA Model

In this section we discuss several univariate extensions of the ARIMA model, that is: (a) the ARMA model with trends, which are applicable if there is some kind of smooth development over time; (b) the seasonal ARIMA model, referred to as SARIMA model, which can be used if the process has a cyclic component to it; (c) the fractionally integrated ARMA model (denoted as ARFIMA or FARIMA model), which can be used if a process exhibits long-range dependency; and (d) the impact ARIMA model, which can be used if there is a sudden impact of an intervention or another sudden change.

ARMA Model with Trends

A logical extension of the ARMA model is to add a deterministic trend, such that the ARMA model describes the variability around this deterministic trend. The trend may have various functional forms, for instance, linear, quadratic, or cyclic. An ARMA model with a linear trend can be represented as

$$y_t = b_0 + b_1 t + \tilde{y}_t \tag{9a}$$

$$\tilde{y}_t = \phi_1 \tilde{y}_{t-1} + \cdots + \phi_p \tilde{y}_{t-p} + u_t - \theta_1 u_{t-1} - \cdots - \theta_q u_{t-q}, \tag{9b}$$

where b_0 is the intercept and b_1 is the slope by which the observed series are regressed on time. The residual $\tilde{y}_t = y_t - (b_0 + b_1 t)$ is then modeled as an ARMA (p, q) process. Alternatively, the model in (9a) and (9b) may be represented in a single equation as

$$y_t = b_0^* + b_1^* t + \phi_1 y_{t-1} + \cdots + \phi_p y_{t-p} + u_t - \theta_1 u_{t-1} - \cdots - \theta_q u_{t-q}. \tag{10}$$

Note however that in this presentation the parameters b_0^* and b_1^* no longer have the easy interpretation of intercept and slope, which the parameters b_0 and b_1 have in Eq. (9a) (Hamaker, 2005).

Processes which consist of a deterministic trend with ARMA residuals are referred to as trend-stationary (Hamilton, 1994): Although these processes are not stationary themselves, they become stationary once the trend is removed. A trend-stationary process may be difficult to distinguish from an integrated process with drift, as described in the previous section. However, if the process is an integrated process with drift, subtracting a linear trend would remove the time-dependency of the mean, but not of the variance. Thus, the resulting series would be mean-stationary, but not variance-stationary (Hamilton, 1994). Determining whether to subtract a linear trend or to difference the data can be done based on the results of a unit root test (see Hamilton, 1994, pp. 444–447).

Seasonal ARIMA Model

Box and Jenkins (1970) extended the ARIMA model to deal with seasonal effects. The basic idea of adding this seasonal component is to accommodate a cyclic effect. For instance, if we consider monthly data, the observation y_t may depend to some extent on y_{t-12}, which represents a annual effect. Similarly, for daily data the observation y_t may depend on y_{t-7}, representing a weekly effect. To deal with these dependencies, the data may be differenced to remove this seasonality, but one can also specify AR or MA relationships at this seasonal interval. This results in the SARIMA (p, d, q) × (P, D, Q) model, where p, d, and q refer to the ARIMA effects discussed before, and P, D, and Q refer to the ARIMA effects at a seasonal lag.

To represent a SARIMA model, we introduce another series z_t which is obtained from y_t by differencing both seasonally and in the way used for ARIMA models (Chatfield, 2004). Then this differenced series is modeled as a SARMA model, in which ARMA relationships can occur directly, or seasonally. For instance, if we consider a simple SARIMA model with only $D = 1$ for a weekly effect (i.e., SARIMA (0, 0, 0) (0, 1, 0)), z_t can be written as

$$z_t = \Delta^{(7)} y_t = y_t - y_{t-7}. \tag{11}$$

Assuming that $d=D=1$, and that we are dealing with daily measurements for which we want to consider a weekly effect, z_t becomes

$$\begin{aligned} z_t &= \Delta\Delta^{(7)} y_t = \Delta^{(7)} y_t - \Delta^{(7)} y_{t-1} \\ &= y_t - y_{t-7} - y_{t-1} + y_{t-8.} \end{aligned} \tag{12}$$

Finally, if a SARIMA (1, 0, 0) × (0, 1, 0) model is considered with a weekly effect, this can be written as

$$\begin{aligned} z_t &= \Delta^{(7)} y_t \\ &= \phi_1 \Delta^{(7)} y_{t-1} + u_t \\ &= \phi_1 (y_{t-1} - y_{t-8}) + u_t \end{aligned} \tag{13}$$

such that $y_t = y_{t-7} + \phi_1 (y_{t-1} - y_{t-8}) + u_t$.

Clearly, this approach allows for many possibilities. However, the interpretation in substantive terms may be difficult. For instance, it is conceivable that a person's emotional state is subject to a weekly effect pattern, which may be captured with the model in (13). This would mean that the difference between today's score and last week's score can be predicted from the difference between yesterday's score and the score on the same day last week using ϕ_1. It is doubtful whether applied researchers will find such explanations plausible. Alternatively, one may choose to model a seasonal effect as a deterministic cyclic trend (as discussed above), such as a sine wave. Other options for handling seasonal effects are discussed at the end of this section.

Fractionally Integrated ARMA (ARFIMA) Model

The ARFIMA model is a generalization of the ARIMA (p, d, q) model in which the integration parameter d can take on noninteger values. Integrated processes are non-stationary, but become stationary after differencing the data. When d is an integer, it is easy to write down the expression of the stationary series in terms of the original series: For instance, when $d = 1$, the stationary series is $\Delta y_t = y_t - y_{t-1}$, and for

$d = 2$ we can write $\Delta^2 y_t = \Delta y_t - \Delta y_{t-1} = y_t - 2y_{t-1} + y_{t-2}$. But when $d = 0.5$, differencing is fractionally, and cannot be expressed in a simple difference equation (Granger, 1980). However, it implies that $\Delta^{0.5} y_t$ is a stationary series. Such fractional integration can be combined with the usual AR and MA relationships, resulting in the ARFIMA model.

Characteristic of ARFIMA processes is that they exhibit long-term dependencies which becomes clear from a very slowly decaying autocorrelation function. This implies that an innovation at occasion t continues to influence future observations for a long time. For this reason such processes are also referred to as long-memory processes. However, Granger and Ding (1996) point out that many other processes may exhibit long-term memory and that this is not a unique feature of fractionally integrated processes.

When $0 < d < 0.5$, the variance of y_t is finite, while $0.5 \leq d < 1$ results in infinite variance (Granger, 1980). Hamilton showed that for $d < 1$, a fractionally integrated process can be rewritten to a pure MA process of infinite order, in which the MA parameters decay slowly, that is

$$y_t = h_0 u_t + h_1 u_{t-1} + h_2 u_{t-2} + \cdots \tag{14}$$

where $h_j \cong (j+1)^{d-1}$ (Hamilton, 1994, pp. 448–449). Hence, if $d = 0.5$, the MA coefficients would be: $h_0 = 1$, $h_1 = 0.71$, $h_2 = 0.58$, $h_3 = 0.50$, $h_4 = 0.43$, etc. As $d \to -\infty$, the process becomes a white noise sequence.

Granger (1980) also showed that ARFIMA models may arise from aggregating other processes. This implies that, as with mixed ARMA processes, an ARFIMA process can be interpreted as the sum of different processes.

Impact ARIMA Models

McCleary, Hay, Meidinger, McDowall, and Land (1980) present the impact or interrupted ARIMA model which can be used to study the effect of an intervention (or event), while assuming that both before and after the intervention an ARIMA model is appropriate. To model the intervention effect they make use of a transfer function. The simplest version of this is the zero-order transfer function, which results in an abrupt, permanent change. Let I_t be a step function, such that $I_t = 0$ before the intervention, and $I_t = 1$ afterward. Then, the observed series can be represented as

$$y_t = \lambda I_t + \omega_t, \tag{15}$$

where ω_t is an ARIMA (p, d, q) model, or potentially a SARIMA model. The parameter λ represents the effect of the intervention.

To allow for a more gradual impact of the intervention, we can use a first-order transfer function. To this end we define the intervention component as $y_t^* = y_t + \omega_t$, such that $y_t = y_t^* + \omega_t$, where ω_t is as defined above. Then,

$$y_t^* = \psi y_{t-1}^* + \lambda I_t. \tag{16}$$

This implies that prior to the intervention, the intervention component $y_t^* = 0$ so that $y_t = \omega_t$. After the intervention takes place at $t = \tau$, the intervention component can be expressed as

$$y_{\tau+n}^* = \sum_{i=1}^{n} \psi^{i-1}\lambda. \tag{17}$$

From this it follows that if $\psi = 0$, we have the zero-order transfer function such as discussed above, with an immediate and abrupt effect of the intervention; if $\psi = 1$ the intervention component y_t^* continues to grow in a linear fashion with slope λ; and if $0 < \psi < 1$, the intervention has a gradual effect which levels off some time after τ.

While interrupted time series models, such as discussed here, have proved valuable in studying the effects of community interventions (e.g., the effect of safety warnings on antidepressants used among youths, see Olfson, Marcus, & Druss, 2008), these models may be less appropriate for studying interventions in the form of psychotherapy, because there the changes are likely to take place more slowly, typically across the entire course of therapy. Moreover, a patient in psychotherapy may display various degrees of relapse, which may require repeated or revised therapeutic intervention. Another potential limitation of these interrupted ARIMA models is that it is assumed that only the level changes, while the ARIMA process ω_t is unaffected by the intervention. To overcome these limitations, one could decide to model separate trends and ARIMA processes before and after the intervention, and determine whether certain parameters may be constrained across these two phases (e.g., Hamaker, Dolan, & Molenaar, 2003).

Rather than using the step function as represented by I_t, one may consider a pulse function P_t, which is defined as $P_t = 1$ at the time of the intervention, and $P_t = 0$ before and after the intervention. Such an intervention model can be valuable if the effect of the intervention is reversible, for instance, the effect of medication on the hyperactivity behavior of a child diagnosed with attention deficit hyperactivity disorder. Such an intervention model has parallels with what is known as ABA-designs (cf., Hersen & Barlow, 1976).

Applications in the Social Sciences

Hamaker et al. (2003) illustrated ARIMA modeling with deterministic trends using three data sets: concentration of luteinizing hormone in blood samples from a healthy female measured at 10 min intervals during the late follicular phase; annual employment percentages of different populations between 1972 and 1998; and the perceptual speed scores of a patient diagnosed with schizophrenia before and after intervention with medication.

Buck and Morley (2006) used SARIMA modeling to study attentional pain control strategies. Because they obtained three measurements per day, they used seasonal differencing to model the time-of-day effect. However, it seems that the actual SARIMA modeling procedure has not seen many applications in the social sciences, and often an alternative way to account for seasonal effects is employed. An example of this can be found in Ichii (1991), who studied the effect of suicide news on monthly suicide rates in Japan. In order to control for a possible seasonal effect, the current suicidal rate is not only regressed upon last month's suicidal rate (and of the month before that in some models), but also on the suicidal rate 12 months ago. Although this may seem like a SARIMA (2, 0, 0) (0, 1, 0) model, it is not: The latter would result in $y_t = y_{t-12} + \phi_1(y_{t-1} - y_{t-13}) + \phi_2(y_{t-2} - y_{t-14}) + u_t$, while the model used by Ichii is $y_t = \phi_1 y_{t-1} + \phi_2 y_{t-2} + \phi_{12} y_{t-12} + u_t$ (Ichii, 1991).

Fractionally integrated processes have enjoyed an increasing interest in the area of reaction time data, where it has been stated by some that long-range memory processes are omnipresent. However, Wagenmakers, Farrell, and Ratcliff (2004) have shown that about half of the empirical time series they considered could be described better with a stationary ARMA (1, 1) process, than with an ARFIMA (1, *d*, 1) process. Delignières, Fortes, and Ninot (2004) applied fractional models to the repeated measurements of self-esteem and physical self of four participants. They conclude that there is not only a balance between preservation and adaption in the short run (as can be shown with an ARIMA (0, 1, 1) model), but that this balance occurs at multiple time scales in a self-similar way. They indicate that a fractionally integrated process is a compromise between the absolute preservation of the expectation (i.e., $\hat{y}_t = y_t - u_t$), as in a white noise process where the expectation is equal to zero for all occasions, and the absolute adaption to change as in the (non-fractionally) integrated process where the expectation is equal to the last observation.

An example of interrupted time series on aggregate time series can be found in Haker, Lauber, Malti, and Rössler (2004), who studied the effect of the 9/11 attacks and the 9/27 amok in Zug in Switzerland (i.e., there are two interventions), on weekly psychiatric patient admissions. They concluded that, contrary to ordinary belief, external psychosocial factors do not influence the need for hospitalization of patients with severe mental disorders. Another example is the study by Cohan and Cole (2002), who investigated the effect of a natural disaster on major family transitions. Their data consist of annual marriage, birth and divorce rates in South Carolina. They also used a pulse function, i.e., a variable with value zero for the years 1975–1989, value one for the year 1990 to model the effect of Hurricane Hugo in 1989, and value zero for the years 1991–1997. They found that (after controling for the general changes over the 24 year span), birth, marriage, and divorce rates were elevated in 1990, indicating that natural disasters mobilizes people to take action in their personal lives.

Multivariate Extensions of ARMA Model

The multivariate extensions of the ARMA model can be divided into four classes: (a) extensions in which the AR and MA relationships are modeled between the observed variables; (b) an ARMA model in which exogenous variables are included;

(c) extensions which are based on introducing latent variables that are measured by multiple indicators, with an ARMA process at the latent level; and (d) cointegrated models in which the combination of two nonstationary processes is stationary.

Vector ARMA Model

The vector ARMA or VARMA model is a straightforward extension of the univariate ARMA model, which was discussed above. Let $\mathbf{y}_t$ be an M-variate observation at occasion t, which may be predicted from previous observations, and from unpredictable parts of previous observations. The VARMA (p, q) model is denoted as

$$\begin{aligned} \mathbf{y}_t &= \phi_0 + \mathbf{\Phi}_1 \mathbf{y}_{t-1} + \cdots + \mathbf{\Phi}_p \mathbf{y}_{t-p} + \mathbf{u}_t - \mathbf{\Theta}_1 \mathbf{u}_{t-1} - \cdots - \mathbf{\Theta}_q \mathbf{u}_{t-q} \\ &= \phi_0 + \sum_{j=1}^{p} \mathbf{\Phi}_j \mathbf{y}_{t-j} + \mathbf{u}_t - \sum_{j=1}^{q} \mathbf{\Theta}_j \mathbf{u}_{t-j} \end{aligned} \tag{18}$$

where $\boldsymbol{\phi}_0$ is an M-variate vector with constants, and $\mathbf{u}_t$ is an M-variate vector with innovations. Although not strictly necessary, the elements of $\mathbf{u}_t$ are often assumed to be uncorrelated with each other.

The $M \times M$ matrices Φ contain the AR parameters on the main diagonal (i.e., the parameters that are used to regress the series upon itself at an earlier occasion), while the off-diagonal elements represent the cross-regression parameters. Thus, element $\phi_{ij,k}$ is used to regress the series i at occasion t on the series j at $t - k$. These matrices are not necessarily symmetric. For instance, series i may be regressed upon series j at previous occasions ($\phi_{ij,k} \neq 0$), whereas series j is not regressed on series i ($\phi_{ji,k} = 0$). The $M \times M$ matrices Θ contain the MA coefficients on the main diagonal. The off-diagonal elements are the parameters by which the unpredictable part of one series at a particular occasion may be predictive of the observation of another series at a later occasion. The model in Eq. (18) may be further simplified to a VAR model, in which case all Θ matrices are zero matrices (e.g., Hamilton, 1994, p. 291). Such models may be used to determine whether there are indications of causal relationships between two or more variables that were measured repeatedly. Hence, it is a more sophisticated alternative to investigating reciprocal influences than by means of prewhitened series.

At first sight the VAR model may seem useful for modeling all kinds of data for which we assume one of the variables has a causal effect on the other (and possible vice versa). However, a VARMA model represents a stationary model and thus it requires the data to be stationary or to be rendered stationary by a suitable transformation. Suppose a researcher is interested in the effect of the empathy of a therapist on the depressive symptoms of a client. If the latter actually show a decline over time, the raw observations can not be modeled directly according to a VAR process. Rather, the researcher will have to make the series stationary, either by detrending the data or by differencing the data. Both approaches have disadvantages.

On the one hand, detrending the data through subtracting a linear (or another) trend implies that if one uses the therapist's data to predict the detrended client's data, one is merely predicting the deviations to the deterministic trend.[4] However, a beneficial effect of the therapy is represented by a decrease over time in symptomatology, which has been taken out when the data are detrended. Thus, one is not really modeling the part one is interested in when applying the VARMA model to the detrended data.

On the other hand, differencing the data is based on the assumption that the process has properties related to a random walk. Recall that a random walk has an expectation of zero, which would imply that the client could just as well improve as worsen over time, with no structural change in the long run. To ensure a positive change in the long run, we would have to find a negative drift (which indicates a decrease in symptomatology). However, this drift is a constant, and is not modeled as a function of the therapist's behavior. In sum, neither solution allows for modeling the structural change as a function of the therapist's behavior. An alternative that may be more appropriate is the cointegration technique discussed later in this section.

VARMAX Model

A VARMAX model is simply a VARMA model with J observed exogenous variables, denoted as $\mathbf{x}_t$. The VARMAX model can be written as

$$\mathbf{y}_t = \phi_0 + \sum_{j=1}^{p} \mathbf{\Phi}_j \mathbf{y}_{t-j} + \mathbf{u}_t - \sum_{j=1}^{q} \mathbf{\Theta}_j \mathbf{u}_{t-j} + \sum_{j=0}^{r} \mathbf{\Gamma}_j \mathbf{x}_{t-j}, \tag{19}$$

where Γ_j is an M **x** J matrix with regression coefficients by which we predict $\mathbf{y}_t$ from $\mathbf{x}_{t-j}$. Such models are particularly suited if we are interested in modeling $\mathbf{y}_t$ and we know or expect it to depend on the x-variables, and we are not interested in how the y-variables influence $\mathbf{x}_t$, either because this is not our focus, or because it is theoretically impossible for $\mathbf{x}_t$ to be affected by the y-variables. An example of the latter would be the effect of weather (e.g., temperature, amount of sunshine, amount of rain) on mood variables (e.g., positive and negative affect): Then the mood variables are modeled as a VARMA process with exogenous variables in the form of weather aspects. Note that if we have reason to believe the exogenous variables are in fact influenced by the other variables, we should turn to a VARMA model which contains all variables as y-variables.

A special case of the VARMAX model is formed by having $M = 1$, such that the outcome y_t is univariate. Such a model is referred to as an ARMAX model. Moreover, when time t (and/or polynomials of t) are used as the exogenous variable (with $r = 0$), this model becomes a multivariate extensions of the ARMA model with a

[4] One can model the trend and the VARMA relations at the same time using a VARMAX model discussed below, but the point made here remains the same: One is modelling the deviations from the deterministic trend (rather then the trend itself) as a function of another variable.

deterministic trend. Note however that the current presentation of the model corresponds to the representation in (10) rather than that in (9a) and (9b), which makes it difficult to interpret the regression coefficient(s) in Γ_0.

Latent VARMA Model

The VARMA model can be extended to a model with multiple indicators measuring a reduced number of latent variables, which follow a VARMA process. Suppose we have a *K*-factor model with *M* observed variables. The factor loadings are restricted to be equal over time, such that the model can be represented as

$$\mathbf{y}_t = \boldsymbol{\mu} + \boldsymbol{\Lambda}\boldsymbol{\eta}_t + \mathbf{e}_t \tag{20}$$

where $\boldsymbol{\mu}$ is an *M*-variate vector with means, $\boldsymbol{\Lambda}$ is an $M \times K$ matrix with factor loadings which do not depend on time, $\boldsymbol{\eta}_t$ is a *K*-variate vector with latent variables at occasion *t*, and $\mathbf{e}_t$ is an *M*-variate vector with measurement errors at occasion *t*. At the latent level, a VARMA model is specified, such that

$$\boldsymbol{\eta}_t = \boldsymbol{\Phi}_1\boldsymbol{\eta}_{t-1} + \cdots + \boldsymbol{\Phi}_p\boldsymbol{\eta}_{t-p} + \mathbf{u}_t - \boldsymbol{\Theta}_1\mathbf{u}_{t-1} - \cdots - \boldsymbol{\Theta}_q\mathbf{u}_{t-q} \tag{21}$$

where Φ and Θ are now $K \times K$ matrices, and $\mathbf{u}_t$ is a *K*-variate vector with innovations of this latent VARMA process. This model can be recognized as a special version of the more general dynamic factor model as discussed by Molenaar (1985). Moreover, when all the Φ and Θ matrices are zero matrices, this model becomes the P-technique model discussed by Cattell, Cattell, and Rhymer (1947). If $q = 0$, the model in Eqs. (20) and (21) reduces to a latent VAR (*p*) model, which is also known as the direct autoregressive factor score model (Nesselroade, McArdle, Aggen, & Meyers, 2002). This model has been compared to the white noise factor score model (Nesselroade et al., 2002), which is also a special version of the more general dynamic factor model discussed by Molenaar (1985). Although the white noise factor score model can not be conceived of as an extension of the ARMA model (because the lagged relationships are not modeled in an ARMA manner, but by use of lagged factor loadings instead), there are situation in which the the white noise factor score model can be rotated into a direct autoregressive factor score model (Molenaar & Nesselroade, 2001). Moreover the latent VMA (*q*) can be shown to be a special case of the white noise factor score model.

Cointegrated Model

Cointegration (Engle & Granger, 1987) has proved one of the most successful discoveries in econometrics, and has earned its discoverers Robert Engle and Clive Granger the Noble Memorial Prize in 2003. A process is said to be cointegrated if each of the univariate series are nonstationary (but is rendered stationary by differencing), while there is a linear combination of the series, which is stationary (Hamilton, 1994). If a process is cointegrated this implies that even though many developments can cause permanent changes in the univariate elements of $\mathbf{y}_t$, there is some long-run equilibrium which ties the individual components of $\mathbf{y}_t$ together. This long-run equilibrium is represented by $z_t = \mathbf{a}'\mathrm{y}_t$, where z_t is a stationary, univariate process. The M-variate vector $\mathbf{a}$ is referred to as the cointegrating vector. Since there is no unique vector that results in a stationary process (because multiplying all elements of the cointegrating vector with the same constant results in another cointegrating vector), some arbitrary normalization is chosen, such as fixing the first element of $\mathbf{a}$ to one.

An example of a bivariate cointegrated process is given by

$$\begin{aligned} y_{1,t} &= \gamma y_{2,t} + u_{1,t} \\ y_{2,t} &= y_{2,t-1} + u_{2,t}. \end{aligned}$$

Note that $y_{1,t} - \gamma y_{2,t} = u_{1,t}$, which is by definition white noise. Thus, the cointegrating vector for this model is $\mathbf{a}' = \begin{bmatrix} 1 - \gamma \end{bmatrix}$.

In general it can be stated that if there are h series, there are at most $h - 1$ cointegrating vectors. The more cointegrating vectors a system actually has, the more constrained its long term behavior is. An illuminating way to think about cointegration is to consider it from a geometric perspective (Dickey, Jansen, & Thornton, 1991). Suppose our system consists of three variables: The behavior of this system can be thought of as the movement of a point in three dimensional space R^3. If all three processes are stationary, the variability is bounded in all three directions, and the observations center around a point. This point can be thought of as the systems equilibrium, from which it never wanders too far. If all three variables are I (1) processes, but they are not cointegrated, this implies there is no restriction on the variability in any direction. Such a system is not characterized by any kind equilibrium. If there is one cointegrating vector, then the plane that is perpendicular to this vector forms the equilibrium of the system. This implies that the variance in the plane is infinite (i.e., unbounded in two directions), but the variance around the plane is finite (i.e., bounded in one direction). This plane can be thought of as the system's equilibrium. If there are two cointegrating vectors, there are two perpendicular planes. The equilibrium of the system is formed by the line which forms the intersection of the two planes. Again, variance on this line is infinite (now unbounded in one direction), while the variance around the line is finite (now bounded in two directions). This shows that more cointegrating vectors imply more constrained behavior of the system in the long run. An illustration in R^2 is given in Fig. 9.1.

Estimating and interpreting cointegrated models is not an easy task. This may give rise to the question: Why not difference the series (or detrend them by subtracting a deterministic trend), and determine whether there are relations between the residual parts? However, if the process is truly cointegrated, differencing the data would overlook the long-term dependencies (Hamilton, 1994).

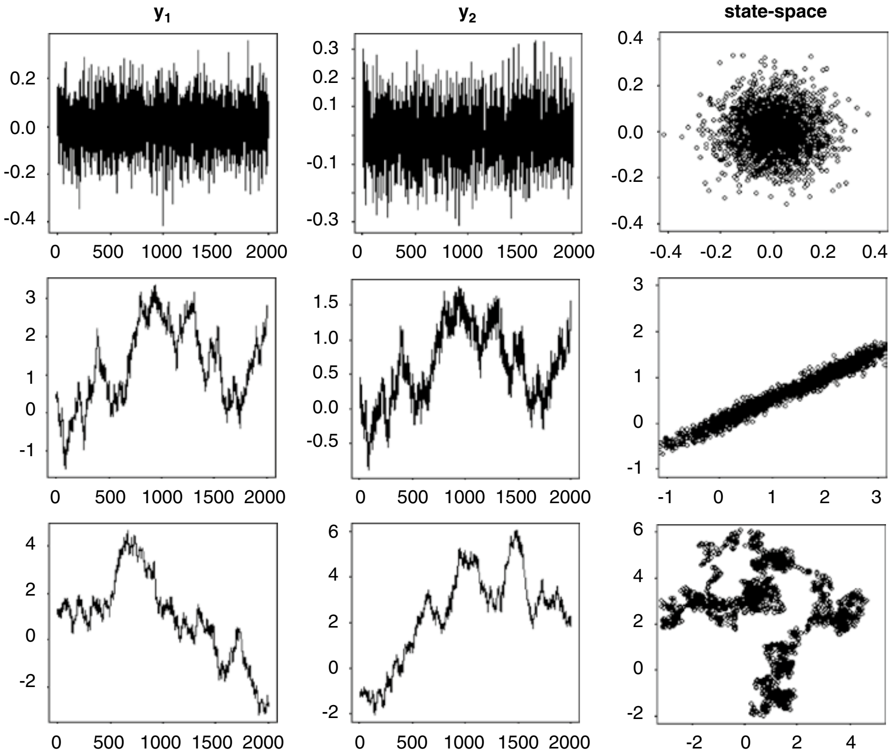

Fig. 9.1 Plots of three bivariate time series: top panel contains two unrelated white noise processes; middle panel contains a cointegrated process; and bottom panel contains two integrated process which are unrelated to each other. Last column contains the behavior of the bivariate series in R^2, where the axes are formed by the two variables. From this it is clear that the first process (which is stationary) has an equilibrium at $\{0, 0\}$. The cointegrated process has an equilibrium formed by the line in the plot on the right. The unrelated nonstatonairy process has no equilibrium

Applications in the Social Sciences

An illustrative application of the VAR (1) model on psychological data can be found in Schmitz and Skinner (1993). They obtained time series data from five children on their effort, performance, subjective evaluation, and control regarding academic tasks in the class room. The authors concluded that the children differed greatly with respect to the relationships between these aspects. For instance, in one child there was no link between effort and performance, while in others this relation was quite strong. Similarly, some children were characterized by a strong link between subjective evaluation on one task and effort on the next, meaning that if they believed they had not performed well, they would try harder the next time (and vice versa), while in other children there was no such relationship.

An application of an ARMAX model (i.e., a VARMAX model with a univariate y_t) can be found in Bollen and Philips (1982), who investigated whether highly publicized suicide stories have an increasing effect on daily suicides. Exogenous variables included dummies for whether a highly publicized suicide had appeared on a particular day (i.e., that day, previous day, and so on up to ten days ago), day of the week, month, year, and certain holidays. Note that, with the exception of the first dummy, these dummies are an alternative way of controling for seasonal effects. Bollen and Philips (1982) concluded that there were two peaks in suicides: at the same day and the next day, and again after six and seven days.

Applications of latent VAR (1) models can be found in Ferrer and Nesselroade (2003), Hamaker, Dolan, and Molenaar (2005), and Hamaker, Nesselroade, and Molenaar (2007). In Hamaker et al. (2005) daily affect measures based on the Five Factor Model (FFM) of personality are analyzed in an exploratory manner. That is, rather than to assume the FFM holds for the variability within individuals as well, it is investigated how many factors are needed for each individual separately. In addition, it is investigated whether there are lagged auto- and/or cross-regressive relationships between the latent variables. It was concluded that individuals differed in both the number and the nature of their intraindividual factors. In addition, some individuals were characterized by lagged relationships at the latent level, while others were not.

Ferrer and Nesselroade (2003) used daily measures of the positive and negative affect from married couple to investigate the reciprocal influences they had on each other. Using a latent VAR (2), they concluded that the wife was influenced by her own affect the preceding day, while the husband was influenced by his own affect the preceding two days. In addition, the wife's affect was influenced by the husband's negative affect at the preceding day, while the husband's positive affect had no effect. The husband was not affected by his wife's affect.

Although cointegration has had many applications in econometrics, only few applications in the social sciences exist. Lin and Brannigan (2003) used cointegration to investigate the relationship between crime and immigration between 1896 and 1940 in Canada. The authors concluded that there was no evidence for a long-term relationship between immigration and crime, with the exception of vagrancy and drunkenness. Stroe-Kunold and Werner (2008) used cointegration to model the interaction between activity, aggressiveness, and depression of a married couple on a day-to-day basis. In addition, the husband's skin symptoms were measured, and the wife's bulimic symptoms. They found some evidence for cointegration of a person's aggressiveness and the spouse's symptoms.

Nonlinear Extensions of the ARMA Model

In this section several nonlinear extensions of the ARMA model are discussed, that is: (a) the bilinear (BL) model; (b) the heteroscedastic autoregressive (ARCH) model. (c) the threshold autoregressive (TAR) model; and (d) the Markov-switching autoregressive (MSAR) model.

Bilinear (BL) Models

The BL model was introduced by Granger and Andersen (1978), and consists of extending the ARMA (p, q) model with product terms between previous observations and previous innovations. Such models are linear in y_t and in u_t, which explains the term bilinear. The BL (p, q, P, Q) model is defined as

$$y_t = \sum_{j=1}^{p} \phi_j y_{t-j} + u_t - \sum_{j=1}^{q} \theta_j u_{t-q} + \sum_{j=1}^{P} \sum_{i=1}^{Q} v_{ji} y_{t-j} u_{t-i}. \tag{22}$$

As with many of the models discussed in this chapter, the BL model was suggested mainly to improve forecasting. Hence, applying this technique in the social sciences, where substantive interpretations may be of greater interest than prediction per se, may result in difficulties as it is not clear how the interactions should be interpreted from a substantive point of view. Moreover, Fan and Yao (2003) state that successful applications (in any field) of the BL model are rare, and they point out diverse unresolved issues regarding estimation and evaluation of the BL model. These issues taken together make it a less attractive candidate for social sciences researchers.

Conditional Heteroscedastic Autoregressive (ARCH) Models

ARCH models were proposed by Engle (1982) to handle volatility, a feature that is often associated with financial data. In contrast to the linear AR model, in which the focus is on predicting the observation y_t based on previous observations, ARCH modeling consists of predicting the *variance* of y_t (i.e., the variance in the prediction error u_t), based on previous observations. Thus the term conditional heteroscedasticity refers to the varying variance which is conditional on previous observations.

Let σ_t be the variance of y_t at occasion t, and let z_t be a white noise sequence with mean zero and variance one. The ARCH (p) model can be defined as

$$y_t = u_t = \sigma_t z_t \tag{23}$$

$$\sigma_t^2 = \phi_0 + \phi_1 y_{t-1}^2 + \cdots + \phi_p y_{t-p}^2. \tag{24}$$

From this it is clear that the uncertainty in predicting y_t depends on y_{t-1} to y_{t-p}. This corresponds well with data characteristics in econometric practice, in which the ability to predict future observations often varies. Another interpretation of this model is that the heteroscedasticity is due to an omitted (i.e., unobserved) variable, in which case the ARCH model is a better approximation of reality than a linear ARMA model (Engle, 1982).

Although the predictive variance σ_t^2 of an ARCH process varies over time, the variance itself is not a function of t. Hence, an ARCH process is stationary.

The ARCH model has been extended with moving average parts to the generalized ARCH (GARCH) model, for which it can be shown that y_t^2 follows an ARMA process (Fan & Yao, 2003, p. 150). GARCH (1, 1) models have shown to be widely applicable in economics, while the ARCH models often require a very large p in order to fit well to empirical data.

These models could be useful in psychological research if for certain data it is known that there is heteroscedasticity over time. For instance, in the study of tics in Gilles de la Tourette discussed earlier (Peterson & Leckman, 1998), instead of measuring the time between the tics, one could also measure the amount of tics per interval: Because of the burst nature of such data, this is likely to result in heteroscedasticity, which could be modeled with a GARCH model.

Threshold Autoregressive (TAR) Models

Threshold models were introduced by Tong and Lim (1980). A TAR process consists of two or more AR processes, which can be thought of as representing separate regimes. The system switches between these regimes when the threshold variable passes a threshold. Suppose there are k regimes, and let z_{t-d} be the threshold variable with delay d, then a TAR(k, p) process is defined as

$$y_t = \sum_{j=1}^{k} \left\{ \phi_0^{(j)} + \phi_1^{(j)} y_{t-1} + \cdots + \phi_p^{(j)} y_{t-p} + \sigma^{(j)} e_t \right\} I(z_{t-d} \in A_j), \tag{25}$$

where $I(\cdot)$ is the indicator function (i.e., it equals one if z_{t-d} falls in A_j, and it is zero otherwise), and the superscript (j) identifies the regimes ($j=1, \ldots, k$). Typically $A_j = (\tau_{j-1}, \tau_j]$, with $-\infty = \tau_0 < \tau_1 < \cdots < \tau_k = \infty$, where τ to τ_{k-1} are the thresholds of interest. Hence, if $z_{t-d} \leqslant \tau_1$, y_t falls in regime 1, if $\tau_1 < z_{t-d} \leqslant \tau_2$, y_t falls in regime 2, and so on. Since the regime-switching is independent of time, the process is stationary.

If y_t serves as its own threshold variable z_t, the model is referred to as a self-exciting TAR (SETAR) model. Such models imply a feedback loop, in which the system corrects itself when its behavior becomes too extreme. If another variable is used as the threshold variable, the model is referred to as an open-loop TAR system (TARSO; Tong & Lim, 1980). This implies that another variable controls the system. If two variables are generated by a TAR model, and each variable serves as the other's threshold parameter, this is referred to as a closed-loop TAR system (TARSC; Tong & Lim, 1980). Such TARSCs were used to model predator-prey data, in which an increase in prey population leads to an increase in predator population until some threshold is reached after which the prey population decreases which in turn leads to a decrease in prey population until another threshold is reached and there is an increase in predator population again (Fan & Yao, 2003).

Extensions of the basic model in Eq. (25) consist of including other (lagged) variables as predictors in the equation, and incorporating moving average terms

(De Gooijer, 1998; Tong, 2003). In addition, multivariate (vector) extensions of the TAR model have been developed (Koop, Pesaran, & Potter, 1996; Tsay, 1998). Further extensions consist of allowing for different orders of the AR processes in each regimes (De Gooijer, 2001).

Markov Switching Autoregressive (MSAR) Models

Hamilton (1989) suggested a nonlinear extension of the AR model that is based on a hidden discrete Markov process. As with the TAR model, it is assumed the system switches between two or more regimes, and each regime is characterized by a different AR process. However, the process that triggers the switching differs between these two models. In TAR modeling the switching occurs when the threshold variable passes a threshold. In contrast, switching in MSAR models is triggered by a hidden discrete Markov process.

Suppose we have k distinct processes, or regimes between which our system switches. Let p_{ij} be the probability of switching to regime j, given that the system is in regime i, that is, $p_{ij} = P[s_t = j | s_{t-1} = i]$, where $i = 1, \ldots, k$ and $j = 1, \ldots, k$. These transition probabilities can be gathered in a matrix,

$$p = \begin{bmatrix} p_{11} & p_{21} & \cdots & p_{k1} \\ p_{12} & p_{22} & \cdots & p_{k2} \\ \cdots & & & \\ p_{1k} & p_{2k} & \cdots & p_{kk} \end{bmatrix}. \tag{26}$$

Note that since $\sum_{j=1}^{k} p_{ij} = 1$, there are only $k \times (k-1)$ non-redundant parameters in this matrix. This matrix governs the Markov switching process s_t, which in turns underlies the regime switching in the MSAR process. The MSAR (k, p) model can be expressed as

$$y_t = \sum_{j=1}^{k} \left\{ \phi_0^{(j)} + \phi_1^{(j)} y_{t-1} + \cdots + \phi_p^{(j)} y_{t-p} + \sigma^{(j)} e_t \right\} I(s_t = j), \tag{27}$$

where $I(s_t = j)$ equals 1 if the system is in regime j at occasion t, and is 0 otherwise. Note that since the parameters in Eq. (26) are independent of time, the MSAR process is stationary.

Kim (1994) extended the work of Hamilton (1989) to the state-space model, such that it can be used for a wide range of time series models. Another useful extension was proposed by Durland and McCurdy (1994), which allows the transition from one regime to another to be duration-dependent. This means that the transition probabilities are not only conditional on the regime the system is in, but also on the amount of time already spent in that regime.

Another model that is related to the MSAR model is the mixture AR (MAR) (Frühwirth-Schnatter, 2006). In this model the observation is supposed to come from a mixture of AR processes. The MAR model can be thought of as a special case of the MSAR model, in which the transition probabilities are equal to the mixing proportions. That is, if π_j is the long-run probability of being in regime *j*, than $P(s_t = j|s_{t-1} = i) = \pi_j$ for all $j = 1, \ldots, k$. Hence, the probability of switching to regime *j* does not depend on the regime the system was in at the precious occasion, but only depends on the long-run probability of making an observation in regime *j*, such that it can be interpreted as the mixing proportion.

Applications in the Social Sciences

The techniques discussed in this section have seen few—if any—applications in the social sciences. To our knowledge, neither the BL nor the ARCH model have been applied in the social sciences. TAR models have been used by Warren (2002; Warren, Hawkins, & Sprott, 2003) to model the behavior of sex offenders and alcohol abusers. They concluded that they could distinguish between periods of recovery versus periods of relapse. Recently, Hamaker, Zhang, and Van der Maas (in press) have shown that the models used by Gottman, Murray, Swanson, Tyson, and Swanson (2002) to model dyadic interaction are in fact TAR-based models. Regarding the MSAR model we are aware of just one application in the social sciences[5], which consists of modeling the daily mood swings in a manic-depressive patient as a two-regime MSAR model (Hamaker, Grasman, & Kamphuis, in press).

Discussion

The majority of studies in the social sciences qualify as nomothetic research, in which a large number of cases were measured on one or a few occasions, and the goal is to find relationships that can be generalized to the population from which the cases were sampled. Exceptions are found in sociology and criminology, where a substantial part of research deals with population aggregates, for instance, unemployment or crime rates, as discussed in this chapter. One could state that in these studies the population itself is dealt with as the single case, and the goal is to understand the process that unfolds at the level of the population.

Despite the dominance of the nomothetic approach in most branches of social science, there is a growing interest in idiographic techniques, as psychologists and

[5] A related technique, which is popular in speech recognition for instance, is the Hidden Markov model (HMM). The difference between the HMM and the MSAR model is that the former requires categorical observations, while the latter requires continuous observations. Moreover, while the MSAR model allows for autoregressive relationships between observations, the sequential dependency in the HMM is modelled exclusively by the hidden Markov process.

other researchers are coming to understand that the standard nomothetic approach presents only one side of the story (e.g., Borsboom, Mellenbergh, & Van Heerden, 2003; Hamaker et al., 2005; Molenaar, 2004; Nesselroade, 2001). The techniques and applications discussed in this chapter illustrate the potential of time series analysis for obtaining a more complete picture of processes that are studied in the social sciences. But even if one is not interested in embracing a fully idiographic approach to the matter, the models presented in this chapter can still be of use: There have been several extensions of ARMA-based models to handle multiple cases, making it compatible with the nomothetic approach. Roughly, we can distinguish between two ways in which the single-case models discussed in this chapter can be extended to handel multiple cases.

First, a straightforward extension consists of fixing the parameters across individuals. We refer to this as the fixed effect approach. Examples of this are the panel version of the VARMA model discussed by Du Toit and Browne (2001), the MI VARMA model discussed by Sivo (2001), and the MSAR model developed by Schmittmann, Dolan, and Van der Maas (2005). Second, a more sophisticated way of extending these models to include multiple cases are the multilevel extensions. For instance, Rovine and Walls (2006) extended the regular AR model in such a way that the AR parameter is random.

As indicated in the introduction, the focus in this chapter was on models for data measured at interval or ratio level, and at regular intervals. Clearly, many measurements in social sciences do not meet these criteria. Recently, Van Rijn (2008) proposed a technique for modeling AR models using ordinal data. Moreover, to model the sequential dependency in both ordinal and nominal data, one can make use of hidden Markov models.

To model measurements obtained at irregular time intervals, one can make use of models based on differential equations. In many diary studies for instance, the intervals between measurements are varied on purpose, to avoid the subject anticipating the next measurement. To model such data of multiple subjects, Oravecz, Tuerlinckx, and Vandekerckhove (in press) developed a multilevel model based on the Ohrnstein-Uhlenbeck process, which is the continuous-time variant of an AR (1) process. Besides having observations at irregular intervals, there are two other reasons for preferring differential equations rather than difference equations. First, Van der Maas and Raijmakers (2000) stressed the fact that while some processes may be understood best in discrete time, others take place in continuous time, warranting a differential equation approach. Second, using differential equations instead of difference equations has the advantage that, while difference equations lead to different results when the intervals change (e.g., daily versus weekly measurements), such arbitrarily evoked differences do not arise when differential equations are used (Oud, 2007). The latter exemplifies that the ARMA-based time series techniques discussed in this chapter form but one approach within idiographic analysis. That is, ARMA-based techniques are a specific branch within time series analysis, which in turn is just one of the possibilities for idiographic research. The aim of the present chapter was to provide an overview of ARMA-based models, and to demonstrate their potential for the social sciences.

Acknowledgements This work was supported by the Netherlands Organization for Scientific Research (NWO), VENI grant 451–05–012 awarded to Ellen L. Hamaker.

References

Andersson, G., Hägnebo, C., & Yardley, L. (1997). Stress and symptoms of Meniere's disease: A time-series analysis. *Journal of Psychosomatic Research*, *43*, 595–603.

Andersson, G., & Yardley, L. (2000). Time-series analysis of the relationship between dizziness and stress. *Scandinavian Journal of Psychology*, *41*, 49–54.

Bollen, K. A., & Philips, D. P. (1982). Imitative suicides: A national study of the effects of television news stories. *American Sociological Review*, *47*, 802–809.

Borsboom, D., Mellenbergh, G. J., & Van Heerden, J. (2003). The theoretical status of latent variables. *Psychological Review*, *110*, 203–219.

Box, G. E. P., & Jenkins, G. M. (1970). *Time series analysis: Forecasting and control*. San Francisco, CA: Holden-Day.

Buck, R., & Morley, S. (2006). A daily process design of attention pain control stategies in the self-management of cancer pain. *European Journal of Pain*, *10*, 385–398.

Bye, E. K. (2007). Alcohol and violence: Use of possible confounders in a time-series analysis. *Addiction*, *102*, 369–376.

Cattell, R. B., Cattell, A. K. S., & Rhymer, R. D. (1947). P-technique demonstrated in determining psycho-physiological source traits in a normal individual. *Psychometrika*, *12*(4), 267–288.

Chatfield, C. (2004). *The analysis of time series: An introduction* (6th ed.). London: Chapman and Hall.

Cohan, C. L., & Cole, S. W. (2002). Life course transitions and natural disaster: Marriaga, birth, and divorce following Hurricane Hugo. *Journal of Family Psychology*, *16*, 16–25.

De Gooijer, J. G. (1998). On threshold moving-average models. *Journal of Time Series Analysis*, *19*, 1–18.

De Gooijer, J. G. (2001). Cross-validation criteria for SETAR model selection. *Journal of Time Series Analysis*, *22*, 267–281.

Delignières, D., Fortes, M., & Ninot, G. (2004). The fractal dynamics of self-esteem and physical self. *Nonlinear Dynamics in Psychology and Life Sciences*, *8*, 479–510.

Dickey, D. A., Jansen, D. W., & Thornton, D. L. (1991). A primer on cointegration with an application to money and income. *Review: Federal Reserve Bank of St. Louis*, 58–78.

Durbin, J., & Koopman, S. J. (2001). *Time series analysis by state space methods*. New York: Oxford University Press.

Durland, J. M., & McCurdy, T. H. (1994). Duration-dependent transitions in a Markov model of U.S. GPN growth. *Journal of Business and Economic Statistics*, *12*, 279–288.

Du Toit, S. H. C., & Browne, M. W. (2001). The covariance structure of a vector ARMA time series. In R. Cudeck, S. Du Toit, & D. Sörbom (Eds.), *Structural equation modeling: Present and future. A festschrift in honor of Karl Jöreskog* (pp. 279–314). Scientific Software International.

Engle, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of the United Kingdom inflation. *Econometrica*, *50*, 987–1007.

Engle, R. F., & Granger, C. W. J. (1987). Co-integration and error correction: Representation, estimation, and testing. *Econometrica*, *55*, 251–276.

Fan, J., & Yao, Q. (2003). *Nonlinear time series: Nonparameteric and parameteric methods*. New York, NY: Springer-Verlag.

Ferrer, E., & Nesselroade, J. (2003). Modeling affective processes in dyadic relations via dynamic factor analysis. *Emotion*, *3*, 344–360.

Fortes, M., Delintnières, D., & Ninot, G. (2004). The dynamics of self-esteem and physical self: Between preservation and adaption. *Quality and Quantity*, *38*, 735–751.

Frühwirth-Schnatter, S. (2006). *Finite mixture and Markov switching models*. New York, NY: Springer.

Gottman, J. M., Murray, J. D., Swanson, C. C., Tyson, R., & Swanson, K. R. (2002). *The mathematics of marriage: Dynamic nonlinear models*. Cambridge, MA: MIT Press.

Granger, C. W. J. (1980). Long memory relationships and the aggregation of dynamic models. *Journal of Econometrics*, *14*, 227–238.

Granger, C. W. J., & Andersen, A. P. (1978). *An introduction to bilinear time series models*. Göttingen: Vandenhoeck und Ruprecht.

Granger, C. W. J., & Ding, Z. (1996). Varieties of long memory models. *Journal of Econometrics*, *73*, 61–77.

Granger, C. W. J., & Morris, M. J. (1976). Time series modelling and interpretation. *Journal of the Royal Statistical Society*, *139*, 246–257.

Haker, H., Lauber, C., Malti, T., & Rössler, W. (2004). Is there an impact of global and local disasters on psychiatric inpatient admissions? *European Archive of Psychiatry and Clinical Neuroscience*, *254*, 330–334.

Hamaker, E. L. (2005). Conditions for the equivalence of the autoregressive latent trajectory model and a latent growth curve model with autoregressive disturbances. *Sociological Methods and Reasearch*, *33*, 404–418.

Hamaker, E. L., Dolan, C. V., & Molenaar, P. C. M. (2003). ARMA-based SEM when the number of time points T exceeds the number of cases N: Raw data maximum likelihood. *Structural Equation Modeling*, *10*, 352–379.

Hamaker, E. L., Dolan, C. V., & Molenaar, P. C. M. (2005). Statistical modeling of the individual: Rationale and application of multivariate time series analysis. *Multivariate Behavioral Research*, *40*, 207–233.

Hamaker, E. L., Grasman, R. P. P. P., & Kamphuis, J.-H. (in press). Regime-switching models to study psychological processes. In P. C. M. Molenaar & K. Newell (Eds.). *Learning and development: Individual pathways of change.* Washington, DC: American Psychological Association.

Hamaker, E. L., Nesselroade, J. R., & Molenaar, P. C. M. (2007). The integrated trait-state model. *Journal of Research in Personality*, *41*, 295–315.

Hamaker, E. L., Zhang, Z., & Van der Maas, H. L. J. (in press). Using threshold autoregressive models to study dyadic interactions. *Psychometrika*.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationairy time series and the business cycle. *Econometrica*, *57*, 357–384.

Hamilton, J. D. (1994). *Time series analysis*. Princeton, NJ: Princeton University Press.

Harvey, A. C. (1989). *Forecasting, structural time series models and the Kalman filter*. Cambridge: University Press.

Hersen, M., & Barlow, D. H. (1976). *Single-case experimental designs: Strategies for studying behavior change*. New York, NY: Pergamon Press.

Ichii, K. (1991). Measuring mutual causation: Effects of suicide news on suicides in Japan. *Social Science Research*, *20*, 188–195.

Kim, C.-J. (1994). Dynamic linear models with Markov-switching. *Journal of Econometrics*, *60*, 1–22.

Koop, G., Pesaran, M. H., & Potter, S. N. (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics*, *74*, 119–147.

Lin, Z., & Brannigan, A. (2003). Advances in the anlysis of non-stationary time series: An illustration of coingration and error correction methods in research on crime and immigration. *Quality and Quantity*, *37*, 151–168.

McCleary, R., Hay, R. A., Meidinger, E. E., McDowall, D., & Land, K. C. (1980). *Applied time series analysis for the social sciences*. Beverly Hills, CA: Sage Publications.

Molenaar, P. C. M. (1985). A dynamic factor model for the analysis of multivariate time series. *Psychometrika*, *50*, 181–202.

Molenaar, P. C. M. (2004). A manifesto on psychology as idiographic science: Bringing the person back into scientific psychology—This time forever. *Measurement: Interdisciplinary Research and Perspectives*, *2*, 201–218.

Molenaar, P. C. M., & Nesselroade, J. R. (2001). Rotation in the dynamic factor modeling of multivariate stationary time series. *Psychometrika*, *66*, 99–107.

Nesselroade, J. R. (2001). Intraindividual variability in development within and between individuals. *European Psychologist*, *6*, 187–193.

Nesselroade, J. R., McArdle, J. J., Aggen, S. H., & Meyers, J. M. (2002). Dynamic factor analysis models for representing process in multivariate time-series. In D. M. Moskowitz & S. L. Hershberger (Eds.), *Modeling intrainidividual variability with repeated measures data: Methods and applications*. Mahwah, NJ: Lawrence Erlbaum Associations.

Olfson, M., Marcus, S. C., & Druss, B. (2008). Effects of food and drug administration warnings on antidepressant use in a national sample. *Archives of General Psychiatry*, *65*, 94–101.

Oravecz, Z., Tuerlinckx, F., & Vandekerckhove, J. (in press). A hierarchical Ornstein-Uhlenbeck model for continuous repeated measurement data. *Psychometrika*.

Oud, J. H. L. (2007). Continuous time modeling of reciprocal relationships in the cross-lagged panel design. In S. M. Boker & M. J. Wenger (Eds.), *Data analytic techniques for dynamic systems in the social and behavioral sciences* (pp. 87–129). Mahwah, NJ: Lawrence Erlbaum Associates.

Peterson, B. S., & Leckman, J. F. (1998). The temporal dynamics of tics in Gilles de la Tourette syndrome. *Biological Psychiatry*, *44*, 1337–1348.

Razvodovsky, Y. E. (2007). Suicide and alcohol psychosis in Belarus 1970–2005. *Crisis: The Journal of Crisis Intervention and Suicide Prevention*, *28*, 61–66.

Rovine, M. J., & Walls, T. A. (2006). Multilevel autoregressive modeling of interindividual differenes in the stability of a process. In T. A. Walls & J. L. Schafer (Eds.), *Models for intensive longitudinal data*. New York, NY: Oxford University Press.

Schmittmann, V. D., Dolan, C. V., & Van der Maas, H. L. J. (2005). Discrete latent Markov models for normally distributed response data. *Multivariate Behavioral Research*, *40*, 461–488.

Schmitz, B., & Skinner, E. (1993). Perceived control, effort, and academic performance: Interindividual, intrainidividual, and multivariate time-series analyses. *Journal of Personality and Social Psychology*, *64*, 1010–1028.

Sivo, S. A. (2001). Multiple indicator stationary time series models. *Structural Equation Modeling*, *8*, 599–612.

Stroe-Kunold, E., & Werner, J. (2008). Modeling human dynamics by means of cointegration methodology. *Methodology*, *4*, 113–131.

Tong, H. (2001). A personal journey through time series in Biometrika. *Biometrika*, *88*, 195–218.

Tong, H. (2003). *Non-linear time series: A dynamic system approach*. Oxford: Oxford Science Publications.

Tong, H., & Lim, K. S. (1980). Threshold autoregression, limit cycles and cyclical data. *Journal of the Royal Statistical Society*, B, *42*, 245–292.

Tsay, R. S. (1998). Testing and modeling multivariate threshold models. *Journal of the Americal Statistical Association*, *93*, 1188–1202.

Van der Maas, H. L. J., & Raijmakers, M. E. J. (2000). A phase transition model for mother-child interaction: Comment on Olthof et al., 2000. *Infant and Child Development*, *9*, 75–83.

Van Rijn, P. W. (2008). *Categorical time series in psychological measurement*. Ph.D. thesis, University of Amsterdam.

Wagenmakers, E.-J., Farrell, S., & Ratcliff, R. (2004). Estimation and interpretation of $1/f^{\alpha}$ noise in human cognition. *Psychonomic Bulletin and Review*, *11*, 579–615.

Warren, K. (2002). Thresholds and the abstinence violation effect: A nonlinear dynamic model of the behaviors of intellectually disabled sex offenders. *Journal of Interpersonal Violence*, *17* (1198–1217), 369–374.

Warren, K., Hawkins, R. C., & Sprott, J. C. (2003). Substance abuse as a dynamical disease: Evidence and clinical implications of nonlinearity in a time series of daily alcohol consumption. *Addictive Behaviors*, *28*, 369–374.