4 Time Series Analysis

Traditional and Contemporary Approaches

Itzhak Yanovitzky

Arthur VanLear

One commonly hears that communication is a process, but most communication research fails to exploit or live up to that axiom. Whatever the full implications of viewing communication as a process might be, it is clear that it implies that communication is dynamically situated in a temporal context, such that time is a central dimension of communication (Berlo, 1977; VanLear, 1996). We believe that there are several reasons for the gap between our axiomatic ideal of communication as process and the realization of that ideal in actual communication research. Incorporating time into communication research is difficult because of the time, effort, resources, and knowledge necessary. Temporal data not only offers great opportunity and advantage, but it also comes with its own set of practical problems and issues (Menard, 2002; Taris, 2000). The paucity of process research exists not only because of the greater effort and difficulty in time series data collection but also because the exploitation and analysis of time series data call for knowledge and expertise beyond what is typically taught in our communication methods sequences in graduate school, and mastering these techniques by self-teaching is difficult for many. This chapter is designed as the first step in developing the knowledge and skills necessary to do communication process research.

Authors' Note: We wish to thank Stacy Renfro Powers and Christian Rauh for use of the interaction data used in the example of frequency domain time series, and we wish to thank Jeff Kotz and Mark Cistulli for unitizing and coding that data.

A time series is a set of observations obtained by measuring a single variable regularly over a period of time. The number of news stories about breast cancer appearing each month in the *New York Times* between 2000 and 2005 or the amount of time spent smiling in every 2-second interval over a 10-minute conversation are both examples of time series.[1] Time series data are different from cross-sectional data in that observations have a temporal order, and the analysis of such data leads to new and unique problems in statistical modeling and inference. Characteristic properties of a time series are that the data are not independently sampled, their dispersion varies in time, and they are often governed by a trend and cyclic components. In particular, the correlation introduced by the sampling of adjacent points in time can severely restrict the applicability of the many conventional statistical methods traditionally dependent on the assumption that observations are independent and that errors are therefore uncorrelated. Statistical procedures that suppose independent and identically distributed data are, therefore, excluded from the analysis of time series. The systematic approach for the statistical modeling of such data is commonly referred to as time series analysis.

The main objective of this chapter is to offer readers a reasonably broad and nontechnical exposition of traditional and contemporary time series analysis methods as they apply to communication research. We intentionally leave out any technical details related to the statistical application of these methods. Time series analysis is quite complex, and there are several excellent texts on this topic that we would recommend to readers who are interested in applying time series analysis methods to their data (see Chatfield, 1989; Cromwell, 1994; Cromwell, Labys, & Terraza, 1994; Gottman, 1981; McCleary & Hay, 1980; McDowall, 1980; Ostrom, 1990; Sayrs, 1989; Shumway & Stoffer, 2000; StatSoft, 2003; Watt & VanLear, 1996). Our goal here is to provide a conceptual introduction to time series analysis, one that (a) illustrates to readers the benefits of incorporating time series analysis into the existing repertoire of communication research methods, (b) describes the common application of time series analysis and the potential weaknesses of this approach, and (c) introduces a set of standards that communication scholars should use when reporting on or evaluating studies that employ time series analysis methods. Thus, we begin with an overview of the potential application of time series analysis methods in communication research and provide examples of how these methods have been used by communication scholars to date. Our discussion here focuses on the kinds of research questions that can be addressed through this family of methods and the proper use of time series analysis in communication research. Next, we review the basic terminology and critical assumptions of time series analysis, distinguish between the time domain approach and the frequency domain approach to time series analysis, and outline the traditional approach to the analysis of time series data. Here, our primary interest is to point out a number of problems in the application of these approaches to communication research and suggest some alternatives.

Time Series Analysis Methods in Communication Research

The impact of time series analysis on scientific applications within the field of communication can be partially documented by listing the kinds of communication research to which time series methods have been applied. In the area of mass communication research, time series analysis methods have been most commonly applied to the study of agenda-setting processes in a variety of contexts including AIDS policy (Rogers et al., 1991), breast cancer screening by women 40 years and older (Yanovitzky & Blitz, 2000), marijuana use among adolescents (Stryker, 2003), drunk driving policy and behavior (Yanovitzky & Bennett, 1999), global warming (Trumbo, 1995), consumer confidence (Blood & Phillips, 1995; Fan & Cook, 2003), and political judgments (Gonzenbach, 1996; Shah, Watts, Domke, & Fan, 2002; Shah, Watts, Domke, Fan, & Fibison, 1999), to name a few. For example, Yanovitzky and Blitz (2000) employed time series regression analysis to compare the contribution of news coverage of mammography screening and physician advice to the utilization of mammography by women 40 years and older in the United States between 1989 and 1991. Data on mammography-related national media attention between January 1989 and December 1991 were generated by analyzing the content of seven nationally and regionally prominent newspapers (the *New York Times, Washington Post, Los Angeles Times, Chicago Tribune, Boston Globe, St. Petersburg Times,* and *USA Today*). All relevant news stories appearing in these newspapers in a course of each month during the research period ($N = 36$ months) were aggregated to represent the volume of media attention to this issue in that particular month. Comparable national-level data on mammography utilization by women ages 40 and older and prevalence of physicians' advice to have a mammogram were compiled from the Behavioral Risk Factor Surveillance System (BRFSS) that is administered each month by the Centers for Disease Control and Prevention (CDC) to a representative cross-section of noninstitutionalized adults nationwide. The proportion of women 40 years and older in each month who had a mammogram in the year preceding the interview served as the dependent variable in the analysis. To estimate the prevalence of physician advice to have a mammogram in each month, the proportion of women 40 years and older indicating that having a mammogram in the past year was their physician's idea was used. Using time series analysis to examine the direction of influence between these three variables controlling for potential confounding variables, the researchers found that both channels of communication (news coverage and physician advice) accounted together for 51% of the variance over time in mammography-seeking behavior by women 40 years and older. Moreover, they found that physician advice was particularly influential for women who had regular access to a physician, while news coverage of mammography was more influential among women who did not have regular access to a physician (mainly due to lack of health insurance).

In this and other similar studies, the typical approach taken by the researchers was to correlate national news coverage of issues over time with outcomes related to public opinion or public policy on these issues during the same time period (Dearing & Rogers, 1996). In virtually all cases, some form of aggregated data was used and most studies were limited to the investigation of the relationship between two time series. However, the time series methods employed in these studies vary considerably, ranging from trend analysis (Brosius & Kepplinger, 1992; Funkhouser, 1973; Smith, 1980; Tedrow & Mahoney, 1979) and cross-correlation methods (Brosius & Kepplinger, 1992; Winter & Eyal, 1991), to time series regression and traditional ARIMA methods (Gonzenbach, 1996; Shoemaker, Wanta, & Leggett, 1989; Trumbo, 1995; Yanovitzky & Bennett, 1999), and to nonlinear methods (Fan, 1988; Fan & Cook, 2003; Yanovitzky, 2002a).

There is little doubt that time series analysis methodology has enriched agenda-setting research in a number of important ways, including the ability to describe and analyze the agenda-setting process and to correlate it with a host of hypothesized outcomes over time (for a complete review, see Gonzenbach & McGavin, 1997). For example, by applying these methods, researchers were able to estimate lagged effects of news coverage on individuals, groups, and institutions (Yanovitzky, 2002b) and calculate the rate in which these effects decay for different issues (Fan, 1988; Watt, Mazza, & Snyder, 1993). They were also able to compare media effects across issues and populations (e.g., McCombs & Zhu, 1995; Yanovitzky & Blitz, 2000) as well as to examine indirect (or mediated) effects between the news, the policy agenda, and personal behavior over time (Yanovitzky & Bennett, 1999). Perhaps more importantly, the use of these methods allows more rigorous tests of agenda-setting theory (Gonzenbach & McGavin, 1997) and facilitates multilevel theorizing and research (Pan & McLeod, 1991; Slater, Snyder, & Hayes, 2006).

In the interpersonal domain, time series analysis has greatly enhanced our understanding of the interaction patterns used by relational partners including reciprocity and compensation, conversational control and coordination among adult dyads and mother-infant dyads (Cappella, 1981, 1996; Street & Cappella, 1989; VanLear et al., 2005), relationship emergence and development (Huston & Vangelisti, 1991; VanLear, 1987, 1991), and decision emergence in small groups (Poole, 1981; Poole & Roth, 1989; VanLear & Mabry, 1999).

One of the best examples in the literature is the research program carried out by Joseph Cappella during the 1980s and 1990s on reciprocity and compensation as forms of mutual adaptation in dyadic conversations (Cappella, 1981, 1996; Street & Cappella, 1989). For example, Cappella (1996) and his colleagues coded conversational behaviors (e.g., vocalizing, smiles, body orientation, eye gaze, illustrators) at every 0.3 second along with an identification of which partner "held the floor" (as defined by Jaffe & Feldstein, 1970). Each data point was an individual's frequency of that behavior in a 3-second window over 30 minutes of conversation (Cappella, 1996). A cluster of behaviors (vocalizing, illustrator gestures, and averting gaze) were highly correlated and clearly associated with actually "holding the floor." Cappella fit each dyadic partner's time series for the composite variable ("turn index") using time domain time series analyses (ARIMA). The ARIMA diagnostics showed that most of the series displayed a first-order autoregressive process, with a few showing trends, first-order moving averages, or some combination (see our later discussion). He then ran individual time series analyses to correlate the time series of speaker A with that of speaker B covarying out the effects of actually "holding the floor." Finally, treating each dyad's analysis as a separate study, he used meta-analysis to aggregate the effects (mutual adaptation scores) across dyads (see later discussion of aggregating postanalysis). The fact that the turn index (along with actually holding the floor) exhibited strongly complementary alternations between dyadic partners is hardly surprising and almost trivial. However, by covarying out the effects of actually "holding the floor," Cappella (1996) was able to show that there is still some mutual adaptation in the form of compensation between the "turn-taking" complex of behaviors, especially among low-expressive dyads and especially between newly acquainted dyads. Cappella (1996) interprets this compensation as "exaggerated politeness" in which partners respond to these behaviors by their partner to avoid "stepping into each other's conversational space" (p. 384).

IMPORTANT ADVANTAGES AND LIMITATIONS OF TIME SERIES ANALYSIS

There are a number of obvious reasons to collect and analyze time series data when studying communication-related phenomena. Among these are the desire to describe variation in variables of interest over time, to gain a better understanding (or explanation) of the data-generating mechanism, to be able to predict future values of a time series, and to allow for the optimal monitoring and control of a system's performance over time (Chatfield, 1989). More importantly, however, time series analysis can greatly enhance our ability to study human communication as a set of dynamic phenomena and to devise more rigorous empirical tests of theoretical propositions about communication-related processes, their determinants, and their effects. Many communication-related phenomena are by definition time-bound processes, and many of the theories that guide research in the field, such as cultivation (Gerbner, Gross, Morgan, & Signorielli, 1986), diffusion of innovations (Rogers & Shoemaker, 1973), structuration (Poole, Seibold, & McPhee, 1996), and relational pragmatics (Fisher, 1978), treat communication, inherently, as a process. Yet, empirical investigations of communication-related phenomena are seldom process oriented, with most being limited to the investigation of simultaneous or short-term relationships between communication variables of interest (Poole, 2000; Watt, 1994).

Time series analysis methods can be a powerful instrument for studying communication processes (Watt & VanLear, 1996). For one, they allow researchers to avoid the pitfalls of studying a communication-related phenomenon in isolation from its past and future. Most communication variables, being realizations of underlying communication processes, are ever evolving (VanLear, 1996). For example, individuals often modify their verbal and nonverbal communication many times in the course of a single discussion (Cappella, 1996), and news coverage of a particular issue, such as the AIDS epidemic, can greatly change in volume and content over a period of a decade (Rogers, Dearing, & Chang, 1991). Some of the changes observed in these variables over time may be random, but many tend to be systematic or deterministic. For instance, certain variables, such as the use of the Internet to search for health information, may trend upward over time (Rice & Katz, 2001), others, such as ambiguity in group decision making, will follow a curvilinear pattern over the course of deliberations (VanLear & Mabry, 1999), while still others, such as television viewing, may follow regular seasonal patterns (Barnett, Chang, Fink, & Richards, 1991) or shorter cycles of attention (Meadowcroft, 1996). Such trends and cycles, in turn, may explain why communication-related phenomena vary across units of analysis at a given point in time or for the same unit of analysis at different points in time. Even stability over time (or inertia) can be quite consequential in this respect. For example, suppose we would like to explain, or even predict, the degree to which a certain person depends on newspapers alone to get the news. We could probably come up with several competing explanations, but our task would be less complicated if we find out that this person has been relying primarily on newspapers to get the news for the past 20 years. We could reasonably propose, then, that this person's current preference for newspapers as the source of news can be explained, to a great degree, by this old habit. Importantly, this habit can also help to explain why this person's current preference for newspapers is similar to or different from those of another person: If they share the same habit, we would predict similarity in current behavior; if they do not share this habit, we would predict a difference. The bottom line is that studying variables in relation to their past can greatly enhance researchers' ability to more fully understand and to predict communication-related phenomena.

Other than facilitating one's capacity to model and to predict communication-related processes, time series analysis methods also have some desirable properties in terms of enhancing causal inference about the relationships between phenomena of interest. One important advantage in this respect is the ability to establish the temporal ordering of variables as a way of delineating which of two variables may be the likely cause of the other. Establishing temporal order between variables may not be an issue when the researcher controls the timing of introducing the independent variable, as is the case when experimental or quasi-experimental methodology is used,[2] but is crucial (though rarely sufficient) for drawing causal inference in the context of cross-sectional or nonexperimental research where researchers cannot have such control. In these cases, the ability to establish temporal ordering can greatly enhance researchers' ability to draw causal inference from their data. For example, employing time series analysis to the relationship between news coverage of mammography and the observed increase in mammography utilization by women 40 years and older between 1989 and 1991, Yanovitzky and Blitz (2000) found convincing evidence that news coverage preceded the observed behavior change, thus supporting the argument that exposure to news coverage about mammography contributed positively to mammography utilization during this period. Similarly, VanLear, Brown, and Anderson (2003) found a series of complex relationships between social support and emotional quality of life among recovering alcoholics such that the supportiveness of an AA sponsor predicted long-term improvement in emotional quality of life, but the emotional quality of life predicted the perceived quality of the relationship with the alcoholic's significant other. In both cases, establishing the temporal order between variables of interest afforded the opportunity to gain better insight into the nature of the relationships that exist between these variables.

One other notable benefit of employing time series analysis when studying the relationships between two or more time-bound variables is the ability to model and test hypotheses about lagged effects. For example, a recent study using time series analysis showed an almost instantaneous effect of news coverage of drunk driving on policymakers' attention to the issue but a delayed effect (of about three months) on policymakers' legislative behavior (Yanovitzky, 2002b). Similarly, in online support groups, the level of self-presentation of a speaker affects not only the self-presentation and other-orientation of the next speaker but of speakers at subsequent lags as well (VanLear, Sheehan, Withers, & Walker, 2005). Standard methods, including those that are seemingly sensitive to the temporal ordering of variables such as repeated-measures ANOVA, are not well equipped to differentiate instantaneous from delayed effects either because measures of all variables are taken at one particular point in time (as is the case when cross-sectional or posttest-only experimental data are used) or because repeated measurements of variables are frequently limited (typically, 2–3 time points at most) and taken at time intervals that are either too close or too distant to capture lagged effects. In contrast, time series designs allow for the collection of time-sequenced data over multiple and equally sequenced time intervals and are better equipped to detect both instantaneous and delayed effects, providing that decisions about the frequency in which data are collected are grounded in strong theoretical or empirical rationale about the expected or hypothesized timing of effects.

Time series analysis may also be useful when comparing two or more communication processes or estimating the effect that these different processes have on a particular outcome. For example, Yanovitzky and Stryker (2001) used time series analysis to estimate the extent to which adolescents' exposure to information about other peers' use of alcohol occurred through the mass media or interpersonal channels. By comparing two different hypothesized processes of exposure (direct exposure to media content vs. diffusion of information within peer networks), they were able to determine that exposure to mass communication channels had an independent contribution to adolescents' perception of alcohol use by peers (see also Zhu, Watt, Snyder, Yan, & Jiang, 1993).

Finally, time series analysis has the advantage of modeling both linear and nonlinear relationships between variables over time. Most standard data analysis methods used in communication research such as ANOVA and OLS regression assume linear association between variables of interest. Frequently, however, nonlinear functions provide better approximation of the true relationship between communication-related variables (Brosius & Kepplinger, 1992; Poole, 2000). Thus, whereas this chapter focuses mainly on linear time series methods with the goal of helping novice users of time series analysis to acquire the methodological foundations of this approach, readers should be aware of recent developments regarding the application of nonlinear time series methods in communication and related disciplines (Fan, 1988; Heath, 2000; Poole, 2000).

On the other hand, many of the well-known problems of collecting and analyzing longitudinal data are relevant to time series analysis. For example, measuring participants repeatedly can influence their behavior and perceptions over time in addition to the impact that independent variables of interest may have on these changes. Similarly, when subject attrition is systematic, trends may reflect the changing nature of the sample rather than the dynamics of the phenomena under investigation. Other confounding influences may be created by historical events and changes in measures or recording practices of variables over time. It is also often difficult to disentangle cohort effects from true temporal trends (Menard, 2002; Taris, 2000). One of the most significant of the temporal problems is "regression toward the mean" where people with extreme scores in their first measurement tend to score closer to the mean on subsequent measures (Campbell & Kenny, 1999). One way of dealing with many of these problems is to use a revolving panel design in which measurements are repeated on a subgroup four times with different participants being rotated into the sample at different waves of the study (Menard, 2002). Subsamples can even be rotated in for several waves, rotated out for several waves, and rotated back into the sample (Mansur & Shoemaker, 1999). Generally, differences can be observed between groups rotated into the sample for the first time and those groups that have been in the sample for some time, and this has been referred to as "time-in-sample bias" (Mansur & Shoemaker, 1999). In some cases, these effects can be statistically controlled for or appropriate transformations may be used to debias the data if a rotating panel design is used.

SUMMARY

Given the importance of the potential contributions to achieving progress in communication research, it is surprising that only a handful of studies have applied this methodology to research problems relevant to communication. There may be several reasons for this, including the cost of collecting time series data (Tabachnick & Fidell, 2001), the inherent complexity of this data analysis method (Shumway, 1988), and the fact that communication-related variables are rarely included in longitudinal survey systems that collect time series data such as the Monitoring the Future Project and the General Social Survey (though such sources could be augmented with comparable communication time series data such as aggregated measures of news coverage of a particular issue). It is also worth noting that time series analysis is not always the best approach to analyzing longitudinal data. Time series approaches are generally appropriate for answering research questions regarding systematic and random patterns of change in a series over time, the association between two or more time series over time, and the effect of interventions (also known as interrupted time series analysis). When longitudinal designs feature large numbers of cases (say, hundreds) but small numbers of repeated observations (e.g., less than 30), then other methods are often utilized such as hierarchical linear modeling and latent growth models (see Chapter 3 of this volume). The strength of the time series models discussed here is in their ability to detect, analyze, and explain complex temporal processes. Their main weakness is in their lack of a model for generalizability across cases. We will discuss the issue of aggregating time series across multiple cases later in this chapter.

The Basics of Time Series Analysis

There are two separate but not necessarily mutually exclusive approaches to time series analysis: the time domain approach and the frequency domain approach. The *time domain* approach is motivated by the presumption that correlation between adjacent points in time is best explained in terms of the dependence of current values on past values of the same series. This approach focuses on modeling some future values of a time series as a linear function of current and past values. The most popular time series analysis techniques that follow this approach—autoregressive integrated moving average (ARIMA) models—receive special attention in this chapter. The *frequency domain* approach, on the other hand, assumes that the primary characteristics of interest in time series analysis involve the periodic or systematic sinusoidal variations found naturally in most data. These periodic variations are often caused by external or environmental factors of interest or may be an intrinsic feature of the phenomenon (e.g., biological rhythms). The partition of the various kinds of periodic variations is typically achieved through Fourier or spectral analysis, which we discuss in detail later on.

BASIC TERMINOLOGY

Before we move to discuss the underlying logic of these two approaches, it is useful to introduce the basic terminology of time series analysis. A time series' most basic unit of analysis is a *point in time* (i.e., a second, an hour, a day, a week, a month, a year, etc.). For each point in time we record the value of a certain variable (e.g., media salience, domestic violence incidents, eye gaze, etc.). The temporal distance between two time periods is a *lag* and is quantified by the number of time units that are included in this time interval. A *first-order series* is a series in which one lag is separating two correlated observations, while a *second-order series* is a series in which two lags are separating two correlated observations, and so on. A time series is said to be *continuous* when observations are made continuously without interruption in time (e.g., heart beats as recorded by EKG trace). A time series is said to be *discrete* when observations are taken only at specific time intervals (usually equally spaced). An *aggregated* time series is a discrete time series in which values are aggregated over equal intervals of time (e.g., the total number of news stories on the economy published each day in the *Washington Post* during the month of January). A *pooled* time series of cross sections contains measures of a particular variable taken from a relatively large number of units of observations (such as individuals, countries, or organizations) over a relatively large number of time points (also called "cross-sectional time series").

One major consideration in analyzing a process captured by a time series involves assumptions about the relationship between observations of the same variable at different time points. Many standard data analysis methods such as regression analysis assume that observations are independent of each other and, therefore, the errors are uncorrelated. Time series analysis, on the other hand, is based on the assumption that successive observations are usually dependent and that we must account for the time order of the observations in our analysis of the time series data.[3] This assumption is formalized in the notion of autocorrelation (also called serial correlation), namely, the correlation of a variable with itself over successive time intervals. Autocorrelation has one important consequence regarding time series data: When successive observations are dependent, we can predict future values from past observations of the same variable and not exclusively by exogenous variables. If a time series future value can be predicted from past values, the time series is said to be *deterministic*.

This characteristic of time series data forces us to identify the internal mechanism that is capable of having produced the set of observed values of a variable over time. With cross-sectional data, this mechanism is the covariance of two or more variables at the same point in time. In contrast, with time series data, this mechanism is assumed to be a *stochastic process*. In order to provide a statistical setting for describing the character of data that seemingly fluctuate in a random fashion over time, we assume that a time series can be defined as a collection of random variables indexed according to the order in which they are obtained in time. Put differently, a stochastic process is a random function that varies in time. For this reason, the future values of a time series (being a realization of a stochastic process) can be predicted with only a certain probability of being correct. This assumption does not mean that the process behaves in a completely unpredictable manner, only that its behavior is partially governed by a random mechanism. In fact, the Wold decomposition thereom holds that a time series can be thought of as a combination of a trend, a deterministic cycle, and a stochastic process (Gottman, 1981). VanLear (1996; VanLear & Li, 2005) has used this theorem to suggest that communication processes can be decomposed into scheduled or programmed processes (i.e., deterministic) and unscheduled stochastic adaptations as they evolve over time.

A second common assumption in many time series techniques (e.g., time domain time series) is that the data are *stationary*. A stationary process has the property that the mean, variance, and autocorrelation structure do not change over time. Stationarity can be defined in precise mathematical terms, but for our purpose we mean a series free of a trend and periodic fluctuations (seasonality). In reality, however, the behavior of most time series is determined by two basic classes of systematic components: trend and seasonality. *Trend* represents a general systematic linear or (most often) nonlinear component that changes over time and that does not repeat in our data (e.g., a monotonic increase or decrease in the level of the series over time). *Seasonality* may have a similar nature, but it tends to repeat itself in systematic intervals over time. Those two general classes of time series components may coexist in real-life data.

Figure 4.1 illustrates these phenomena through a hypothetical example: local news coverage of outdoor events held in a particular community between 1987 and 1995. Just by inspecting the time series visually, we can detect a slight upward trend in the number of outdoors events covered by the local news media from year to year. In addition, as one may expect, news

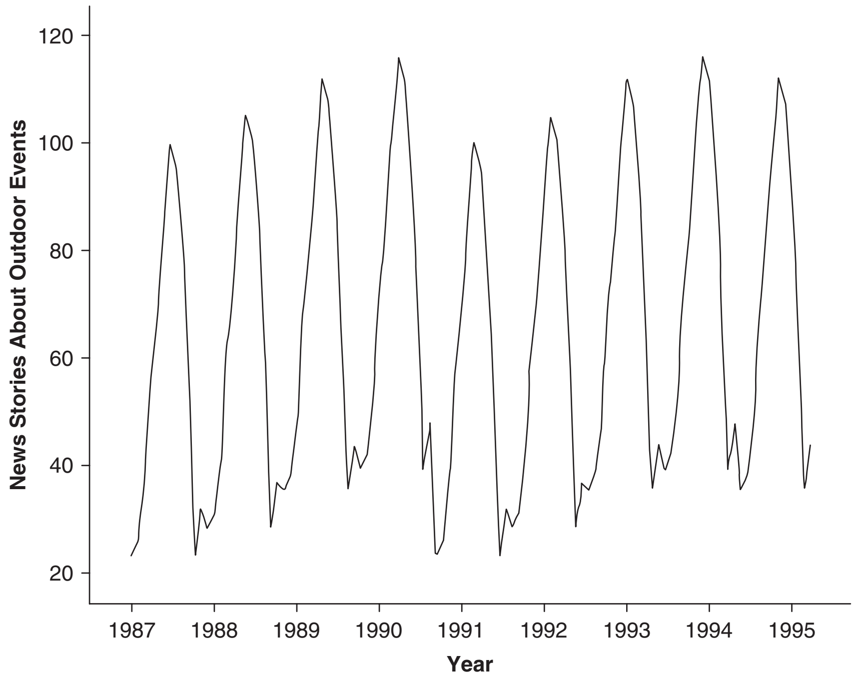

Figure 4.1 Hypothetical News Coverage of Outdoor Events, 1987–1995

coverage in each year peaks during the spring and summer months and declines during the winter months. These "cycles" are responsible for seasonal effects on the level of the series. The problem with using standard methods (e.g., regression analysis) to model the process underlying a time series is that such deterministic components can introduce large systematic errors into the estimation procedure that likely violate critical assumptions such as the assumption that models' residuals are uncorrelated or that they are equally distributed among the different categories of the independent variable (homoscedasticity). Thus, before we can use a time series as a "conventional" variable, we must make sure it is stationary or free from any systematic and deterministic effects over time. If it is not a stationary time series, it must be transformed into one (a procedure that is often referred to as *prewhitening*) before standard statistical procedures could be employed.

The Time Domain Approach to Modeling Time Series Data

This basic rationale of time series analysis guides the common approach to analyzing time series data. In this respect, traditional time domain approaches to time series analysis are best understood as techniques for adapting standard regression methods to the problematic nature of time series data. As noted above, the classical regression model was developed for the static case, namely, when a dependent variable is allowed to be influenced by current values of the independent variable. In the time series case, it is desirable to allow the dependent variable to be influenced by the past values of the independent variables and possibly by its own past values. The need to incorporate these lagged relationships into the explanation of a variable's variance over time led to the development of the *autoregressive integrated moving average* (ARIMA) model that was popularized by Box and Jenkins (1976) and that seeks to uncover persistent patterns in the behavior of time series, often so that unbiased estimates of standard deviations can be calculated and that accurate forecasts of future values can be generated.

It is worth noting that because of its power and flexibility, ARIMA is a rather complex technique—it is not easy to use, it requires a great deal of experience, and although it often produces satisfactory results, those results depend on the researcher's level of expertise. Readers who are interested in applying this approach could benefit from reading texts dedicated to these methods (e.g., McCleary & Hay, 1980) as well as books, monographs, and book chapters that discuss the application of ARIMA methods in agenda-setting research (e.g., Gonzenbach, 1996; Gonzenbach & McGavin, 1997; Trumbo, 1995). Our discussion below leaves out much of the technical details of fitting ARIMA models and focuses instead on the gist of this approach, which can be summarized in a few basic steps: (a) plotting the data against time, (b) possibly transforming the data, (c) identifying the dependence order of the model (identification), (d) estimating the ARIMA parameters (estimation), and (e) evaluating the estimated model's goodness of fit (diagnostics).

ARIMA METHODOLOGY: A PRIMER

The first step to take in any time series analysis is to plot the time series against time. In many cases, the researchers can detect the presence of possible deterministic components just by inspecting the behavior of the series of time, as is the case in Figure 4.1. Depending on the degree of fluctuations in a time series behavior over time, one may also find it useful to use some form of a *smoothing* procedure. Smoothing techniques are used to reduce irregularities (random fluctuations) in time series data and thus provide a clearer view of the true underlying behavior of the series. Two common smoothing procedures are moving average and natural log transformation, but other options exist (for more options, see Chatfield, 1989). However, the visual display may be misleading at times, particularly when no systematic change in the series' level can be detected. In these cases, researchers are strongly cautioned against concluding that the time series is stationary before more formal statistical tests are performed (e.g., the Box-Ljung Q statistic).

Two additional tools that most common statistical packages (e.g., SPSS, SAS, STATA) offer are the autocorrelation function (ACF) and the partial autocorrelation function (PACF) *correlograms.* Autocorrelation correlograms are a commonly used tool for checking randomness in a series' behavior over time. This randomness is ascertained by computing correlations for data values at varying time lags. If random, such autocorrelations should be near zero for any and all time-lag separations. If nonrandom, then one or more of the autocorrelations will be significantly nonzero. Partial autocorrelations are the autocorrelations between two time points separated by a certain lag controlling for any dependence on the intermediate time points within this lag. If a lag of 1 is specified (i.e., there are no intermediate elements within the lag), then the partial autocorrelation is equivalent to autocorrelation. In a sense, the partial autocorrelation provides a "cleaner" picture of serial dependencies for individual lags (not confounded by other serial dependencies). Data can be assumed to be stationary if no partial autocorrelation is found to be statistically significant. Autoregressive models are created with the idea that the present value of the series can be explained as a linear function of past values on this series.

Figure 4.2 shows the ACF correlogram for a hypothetical discrete series measuring exposure to media messages about domestic violence per 100,000 viewers in a finite population over a period of 145 weeks ($N = 145$ equally spaced time points). Figure 4.3 presents the corresponding PACF correlogram. The bars in each diagram represent the estimated autocorrelation coefficients and are bounded by 95% confidence intervals to detect statistically significant autocorrelations. It is apparent from Figure 4.2 that a significant degree of autocorrelation exists within the series (as indicated by the presence of a statistically significant correlation in each of the lags), which suggests a nonstationary series. The PACF plot in Figure 4.3 indicates that the autocorrelation in the series is of first order (namely, that the strongest autocorrelations exist between observations separated by a single time lag).

When a time series is determined to be nonstationary, the ACF and PACF correlograms have an important role in the identification of deterministic components using the ARIMA approach. ARIMA stands for the three types of mathematical processes that can be employed to generate a stationary process: (1) autoregression (AR), (2) trend or integrated series (I), and (3) moving average (MA). The *autoregressive* component in the ARIMA model accounts for the autocorrelation or the magnitude of the dependency between adjacent observations. These dependencies can be removed by regressing the present value of the series on the linear function of past values of the same series at k lags and replacing the original

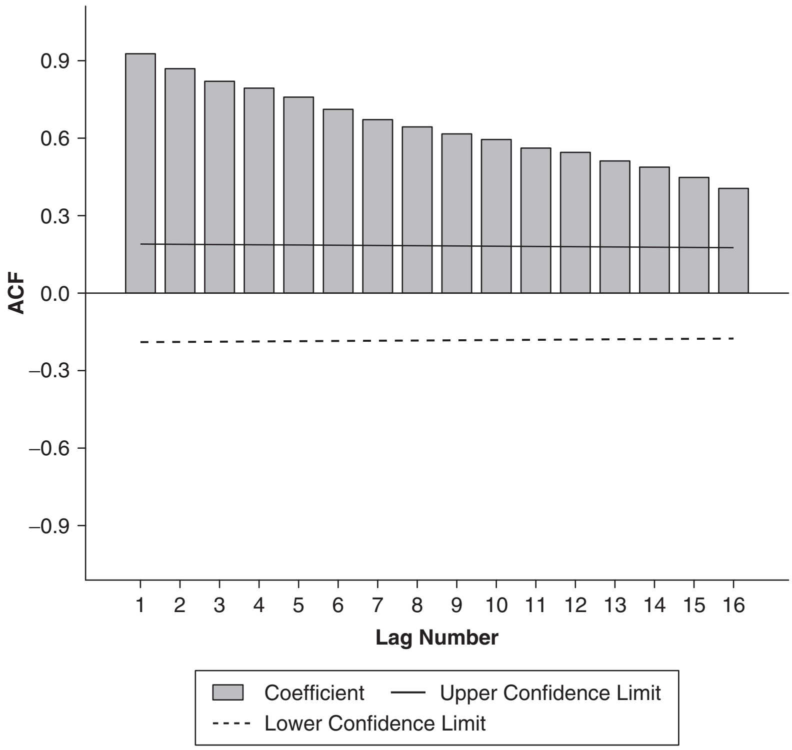

Figure 4.2 Autocorrelation Function (ACF) Correlogram

series with the newly created residual series. The *integrated* component in the ARIMA model addresses the issue of stationarity in the average level of the series over time (note, however, that it does not address the issue of stationarity in variance over time). Many, if not most, time series can be made stationary by *differencing.* The method of differencing replaces each time series observation with the difference of the current observation and its adjacent observation *k* steps backward in time. The moving average component is a bit less intuitive. It addresses the persistence of a random shock (or a past error that cannot be accounted for by an autoregressive process) from one observation to the next. A shock is an external event that takes place at a particular point in the series and whose impact is not contained to the point at which it occurs. The method of moving averages dampens fluctuations in a time series by first taking successive averages of groups of observations and then replacing each successive overlapping sequence of *k* observations in the series with the mean of that sequence.

ACF and PACF correlograms are typically used to identify the three ARIMA model's parameters that correspond to each of these three components. The autoregressive parameter (*p*) represents the number of time lags that separate two correlated observations. Thus, a first-order series,

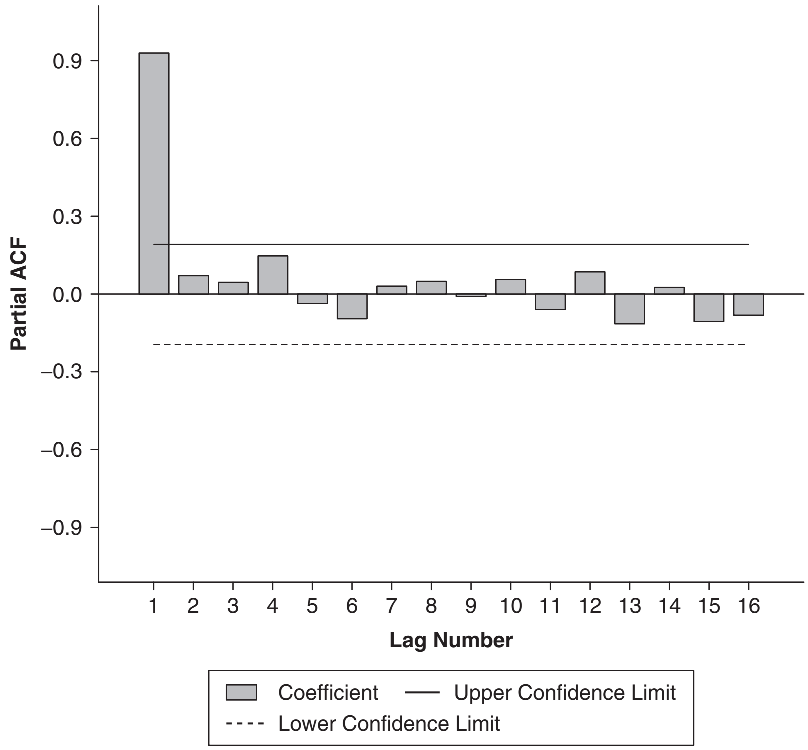

Figure 4.3 Partial Autocorrelation Function (PACF) Correlogram

abbreviated $AR_{(1)}$, is a series in which one lag is separating two correlated observations, while a second-order series, $AR_{(2),}$ is a series in which two lags are separating two correlated observations, and so on. The integrated parameter (*d*) represents the shape of the trend that exists within the time series: A first-order *d* parameter ($d = 1$) indicates a linear trend, a second-order *d* parameter indicates a quadratic trend, and a third-order *d* parameter indicates a cubic trend. The moving average parameter (*q*) represents the number of time lags (or window) over which the effect of a random shock persists in the series. A first-order moving average parameter ($q = 1$) means that current observations are correlated with shocks at lag 1, a second-order *q* parameter means they are correlated with shocks at lag 2, and so on. By convention, these parameters are represented as (*p,d,q*) and denote the ARIMA term. Thus, an ARIMA (1,0,0) indicates a series characterized by a first-order autoregressive component with no trend or moving average components. Particular ARIMA models tend to be associated with a particular output of the ACF and PACF correlograms. For example, an ACF output that demonstrates an exponential decay of serial correlations of the type shown in Figure 4.2 and a PACF output demonstrating a spike at lag 1 with no serial correlations for other lags, as is the case in

Figure 4.3, suggest the presence of a first-order autoregressive process (1,0,0). Note that the boundary lines around the functions in these figures are the 95% confidence bounds. If the bar representing an autocorrelation at some lag crosses the boundary lines, the interpretation is that this autocorrelation is significantly different from zero and that it should be included in the ARIMA model. ACF and PACF representations of other ARIMA processes are included in most basic texts about time series (e.g., Box & Jenkins, 1976; McCleary & Hay, 1980) as well as manuals of commonly used statistical packages (for example, the SPSS Trends manual). At times, a series may have seasonal components (seasonal autocorrelations) or structural dependency among observation separated by one period or cycle, such as an annual cycle or 6-month cycles. Seasonal components are evident in ACF and PACF plots that wear the shape of cycles with strong positive ACFs equal to the length of the cycle and negative ACFs equal to one half the period of the cycle. A seasonal ARIMA model takes the seasonal components into account while using the same components of a regular ARIMA (with seasonal parameters denoted by uppercase letters: P,D,Q).

Once the ARIMA model's parameters have been identified, they can be used to estimate the ARIMA model that best fits the data. ARIMA models use a maximum likelihood estimation procedure that is designed to maximize the likelihood of the observed series, given the estimated parameter values. This can be done using any standard statistical package that includes a time series module. The next step is the evaluation (diagnosis) of the estimated model's fit to the observed series. Here, a combination of three strategies is recommended. The first is to verify that each ARIMA parameter in the model is statistically significant using standard hypothesis testing procedures (e.g., effect/standard error). If not significant, the respective parameter can in most cases be dropped from the model without affecting substantially the overall fit of the model. A second straightforward test involves the accuracy of the estimated model's forecast of future values of the series. Typically, this procedure entails estimating the ARIMA model based on partial data (e.g., the first two thirds of the observations in a series) and using it to predict the remaining observations, which are then compared with the known (original) observations. However, a good model should not only provide sufficiently accurate forecasts, it should also produce small, random, and statistically independent residuals. The patterns of ARIMA model–generated residuals are typically inspected through the use of ACF and PACF correlograms. If no serial dependencies are detected (namely, no remaining autocorrelations or partial autocorrelations at various lags appear to cross the 95% boundary lines in the ACFs' and PACFs' plots), the series is said to be stationary. If serial dependencies are detected, the ARIMA model would need to be reestimated using a different combination of parameters until stationarity is achieved.

The procedure described thus far pertains to a time series that consists of single observations recorded sequentially over equal time increments (or a *univariate time series*). To conduct a *multivariate time series analysis,* several additional steps are necessary. An important requirement of the ARIMA approach is that all time series involved in the analysis of multivariate relationships will be made stationary (or prewhitened) before standard correlation or regression methods are used. However, there is a disagreement among researchers about the most appropriate prewhitening approach. Box and Jenkins recommended differencing (ARIMA 0,1,0) as the preferred method of removing deterministic components from each nonstationary series (Box & Jenkins, 1976). However, some (e.g., Cappella, 1996) proposed that a better approach is to use the same ARIMA model used to prewhiten the independent series to prewhiten the dependent series, while others (e.g., Granger, 1969) prescribe that using an autoregressive model (ARIMA 1,0,0) would suffice in most cases given that most time series studied in the social sciences seem to be governed by a first-order autoregressive process (McCleary & Hay, 1980). A reasonable approach has been proposed by Watt (1994), who expressed the concern that while first-order differencing may produce a stationary series, it also destroys all information about the absolute value and the trend in the original data that are theoretically important for explaining the variance in the series over time. Watt proposed instead that variables be made stationary by fitting a least squares regression line to the data and creating a transformed time series based on the regression residuals. This technique was used successfully in a number of studies (e.g., Yanovitzky & Bennett, 1999; Yanovitzky & Blitz, 2000), though the other approaches work as well. The disadvantage of this approach is that either linear or polynomial trends can lead to unrealistic forecasts of values far beyond the temporal horizon of the study. Therefore, this approach should be used with caution for forecasting.

Next, a combination of cross-correlation analysis and a Granger causality test typically allows researchers to sort out the causal direction between two or more variables. The *cross-correlation function* is a measure of the degree of the linear relationship between two time series as a function of the time lag between the two. Conceptually, it is similar to the autocorrelation function except that it compares values in two different time series instead of comparing different values within the same series. In cross-correlation analysis, the correlation of one time series with a time-lagged version of a second time series is examined, as illustrated in Figure 4.4, where the association between exposure to media messages about domestic violence and volume of calls received at a domestic violence hotline within a week's period is examined. A statistically significant correlation at a certain lag indicates the time lag required for the independent (or leading) series to affect the dependent series. A cross-correlation matrix is computed for each relationship of interest where in each case the independent series in Step 1 becomes the dependent series in Step 2. A time series will be considered as leading another time series if statistically significant correlations are observed only in the lags that represent its effect on the second time series. In this example, the only statistically significant cross-correlation appears to be at lag +1, indicating that exposure to domestic violence messages leads the behavior of calling a domestic violence hotline such that messages received influence behavior one week later. On the other hand, the nonsignificant effect at lag –1 indicates that we cannot say that the number of calls to the hotline necessarily influences media messages about domestic violence.

However, when working with a set of variables and a range of lags, a large set of correlations is produced and spuriousness becomes a concern. If a nonrandom distribution of positive or negative correlations is observed, the spuriousness hypothesis is rejected. If the only meaningful and significant correlation is to be found at unity (time t_0), the cross-correlation analysis indicates the existence of covariation between the two series. On the other hand, if no significant correlation is found in any of the lags, one may conclude that the two time series are independent. The information gathered from the cross-correlation analysis is used in the computation of *transfer function models.* These are essentially multiple

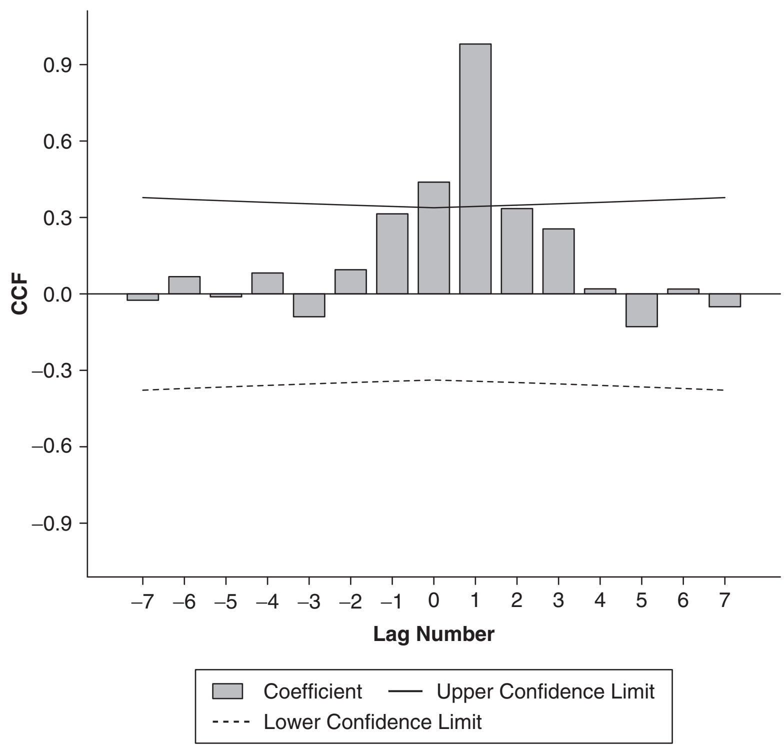

Figure 4.4 Cross-Correlation Function (CCF)

regression models in which the dependent variable is the given series' residuals and the independent variables are created based on all the cross-lagged correlations that were statistically significant.

The *Granger causality test* (Granger, 1969) compares the effect of one time series on another in order to verify the causal direction between the two time-bound variables. The basic logic of this test is that time series *X* may be considered a cause of time series *Y* when *X* predicts *Y* significantly better than *Y* predicts itself. To test this proposition, both series are prewhitened first, and then *Y* is regressed on its previous (lagged) values alone. In the next step, *Y* is regressed on both its lagged values and the lagged values of *X*. After being transformed to white noise (i.e., made stationary), *Y* should have a very limited ability to predict itself based on its previous value. The R^2 for the initial and second models are then compared with an *F*-test to determine if any predictive improvement due to the effect of *X* is actually significant. The procedure is then repeated for the effect of *Y* on *X*. In the final step, the regression coefficients representing the effect of each series on the other controlling for previous values of the dependent series are compared to determine which series (if any) leads the other.

LIMITATIONS OF THE ARIMA APPROACH AND SOME USEFUL ALTERNATIVES

The ARIMA approach, then, is an iterative model-building procedure through which any deterministic component of the time series is identified and removed in order to make the data stationary before standard data analysis methods could be used with time series data. Although powerful and flexible, this approach has some known limitations (Shumway, 1988). For instance, the ARIMA method performs best when the number of sequential observations available for analysis is 50 or greater, the unit of time is consistent among all variables measured, and each time series is uniform and unbroken. These may not be reasonable requirements for much of the data collected or used by communication researchers, particularly when researchers have no control over data collection such as when secondary time series data are used. The greatest difficulty involves the analysis of unequally spaced time series observations. In data analysis practice, such a characteristic of time series data is often ignored and standard analyses that treat data as equally spaced are used. This practice can clearly introduce a significant bias into estimates of ARIMA parameters leading to incorrect predictions. It is therefore necessary to use continuous time series models for serial correlations (see Jones & Ackerson, 1990), which can be quite complicated. In addition, ARIMA assumes that the values of the estimated parameters are constant throughout the life of a series (i.e., stationarity of process), which may be a false assumption in some cases. Furthermore, ARIMA can model only linear relationships between time series, whereas relationships between communication-related variables often take a nonlinear form. Finally, although much of time series analysis focuses on analyzing the mean behavior of a time series, there has been increasing attention to the study of volatility or variability of a time series. ARIMA models assume constant variance and, therefore, cannot be used in these cases. Autoregressive conditionally heteroscedastic (ARCH) models (Engle, 1982) should be used instead.

More importantly, though, ARIMA is appropriate only for a time series that is stationary (i.e., its mean, variance, and autocorrelation should be approximately constant through time). This may not be a problem if one is interested in forecasting future values of a series, a common objective of time series analysis in economic research. However, communication scholars are often more interested in studying theoretically meaningful communication processes than in forecasting. That is, they are far more interested in explaining systematic changes in the behavior of variables over time (such as trends and cycles in news coverage of issues) than removing them statistically so that they may produce more precise prediction of future values on these variables. In this sense, as suggested by Watt (1994), the powerful filters employed by ARIMA models remove much or all of the variance in a time series that communication scholars may seek to explain.

A number of statistical alternatives to the ARIMA approach have emerged in recent years (Brockwell & Davis, 2002), two of which are *distributed lag models* (Ostrom, 1990) and *differential equation models* (Fan, 1988; Zhu, 1992). Differential equation models can be used to model both linear and nonlinear relationships in time, but a discussion of this approach is beyond the scope of this chapter (for an overview, see Fan & Cook, 2003). However, distributed lag models are briefly discussed here.

Distributed lag analysis is a specialized technique for examining the relationships between variables that involve some delay. This technique relies on a simple structural equation model that can be estimated by an ordinary least squares (OLS) regression and is mathematically expressed as follows:

$$Y_t = b_0 + b_1 Y_{t-1} + b_2 X_{t-1} + e_t$$

where Y_t is the dependent series at time t, Y_{t-1} is the dependent series lagged by a single time point, X_{t-1} is the independent series that is also first-order lagged, and e_t is the error in estimation of Y_t. The model estimates three parameters: the constant or intercept (b_0) and two partial time series regression coefficients (b_1 and b_2). Similar to the logic of a Granger causality test, X is said to cause Y when lagged values of X are significantly related to Y after controlling for the previous history of Y (i.e., lagged values of Y). To ensure that the standard regression assumptions are not violated when estimating the model, three statistical tests are employed. The first, the Durbin-Watson test of correlated errors (serial correlation), is designed to detect first-order autocorrelations. When there is no serial correlation, the expected value of the Durbin-Watson statistic is approximately 2, whereas a value under 1.5 indicates a positive serial correlation and a value above 2.5 a negative serial correlation. The second, tolerance, estimates the amount of variation in a single predictor that is not explained by its association to other predictors in a multiple regression model. Tolerance values range from 0 (perfect collinearity) to 1 (no collinearity). Finally, the autoregressive conditional heteroscedasticity (ARCH) test (Engle, 1982) is used to test the null hypothesis of homoscedasticity in the errors. This statistic has a chi-square distribution with 1 degree of freedom, where a nonsignificant result (i.e., a value of 7.8 or lower) indicates that the errors are homoscedastic.

The Frequency Domain Approach to Modeling Time Series Data

The logic of frequency domain time series analysis is somewhat different from that of time domain time series analysis. Time domain time series generally treat temporal patterns (e.g., trends, seasonal or cyclical fluctuations) as potential confounds that can create "spurious" relationships between two time series and, thus, obscure the true influence of one series on another. For example, if the level of self-disclosure of one person is correlated with the level of self-disclosure of his/her relational partner over time, it could be because they are both following the same developmental norm of incremental increase rather than either person's disclosure level influencing the disclosure level of the other (i.e., reciprocity). Likewise, if TV advertising and TV programming follow the same seasonal pattern, they may be correlated because of that common pattern rather than because programming variation leads to advertising variation. Hence the prewhitening process seeks to eliminate these confounding effects from the data so that the effects of one variable on another can be observed. Frequency domain time series, on the other hand, begins by looking for hidden temporal patterns in the data (especially cyclical or periodic patterns) and models them. Bivariate spectral analysis then attempts to correlate two time series by identifying and correlating the common patterns. For example, Chapple (1970) argued that people's physiological rhythms entrain their behavioral rhythms, which in turn influence the rhythms of their social interaction. In a series of studies, Warner (1996) utilized frequency domain time series to demonstrate correlations between people's physiological rhythms and their communicative behavior as well as the behavioral rhythms of their relational partners. Altman, Vinsel, and Brown (1981) argued that as relationships evolve, relational partners experience a dialectic tension between openness and closeness, and this leads to a cyclical pattern of openness behaviors over the course of a relationship; in successful relationships, partners match and time their cycles of openness to create a synchronous pattern. Utilizing frequency domain time series, VanLear (1991) found that (a) relational partners did evidence cyclical patterns of openness using both behavioral and perceptual data, (b) relational partners generally matched and timed their behavioral cycles in a synchronous pattern, and (c) the nature of the cycles of openness were associated with communication and relationship satisfaction. Despite these differences in perspective, the Wiener-Khintchine theorem shows that mathematically, the time domain and frequency domain are just two sides of the same coin. They are mathematically the same (Gottman, 1981).

Frequency domain time series begins with the recognition that a*ny* time series, from an extremely patterned process to a random, white-noise process, can be represented by a series of weighted, orthogonal, sinusoidal (i.e., sine and cosine) functions. These functions involve several parameters. The height of the cycle from zenith to baseline is the *Amplitude R*. The *period* is the time it takes to complete a single cycle. Instead of the period, one could choose to use the frequency ω, the number of full repetitions the function makes in a single unit of time (usually measured in radians per unit of time). Finally, we identify the phase angle φ (in radians) with respect to the time of origin.

Given that any time series can be described by a set of weights representing the sinusoids' amplitudes and frequencies, that set of weights is given by the Discrete Fourier Transform (DFT). This set is chosen because they are orthogonal, and given the sum of these $N - 1$ sine and cosine coefficients, they can completely represent the data. Of course, that number of functions offers no more parsimonious representation than the raw data, so the goal is to find a small number of functions that, when inverted, will adequately reconstruct the data with little error. Watt (1994; VanLear & Watt, 1996) suggests a stepwise procedure that will identify the major patterns in the data (Bloomfield, 1976). First, the amplitudes (or amplitudes squared) of the Fourier coefficients are plotted against the Fourier frequencies in what is called a *periodogram* (see Figure 4.6). The amplitude is an indication of the strength of a given pattern in the data. Peaks in the periodogram at various frequencies indicate that those frequencies are particularly strong patterns within the data. If the strongest pattern in the data is not at one of the Fourier frequencies, the frequencies closest to the one that actually represents the pattern will show a peak. When the strongest function is found, a sinusoidal function representing that frequency is fit to the data using least squares minimization, much as a linear trend is fit to the data in ordinary regression. The frequency can be adjusted to find the optimal frequency of the strongest component, even when it doesn't fall on one of the Fourier frequencies (Watt, 1996). If this model explains a substantial portion of the variance of the series, then the residuals around that function can be extracted, a DFT can be calculated for the residual series, a new periodogram can be plotted, and the next-strongest frequency can be identified and fitted to the data. This iterative process is continued until the series is adequately explained and no more major functions are identified (VanLear & Watt, 1996). Once all meaningful functions are identified, the series is a white-noise process. In a white-noise process, the periodogram values will display an "exponential distribution" (StatSoft, 2003). Watt (1998) has developed a program (FATS, Fourier Analysis of Time Series) that will perform the stepwise procedure and also can fit a priori cyclical models to a time series.

There are some cautions and limitations involved in the stepwise method. It can be shown that while this procedure is very sensitive to the identification of actual hidden patterns in the data, it can yield "false positives" (i.e., false evidence of a pattern that is only a random variation) when examining both significance and amount of variance explained by a function. This bias is most pronounced with short time series and diminishes as the length of the series increases (VanLear & Li, 2005). Therefore, we recommend that several precautions be used. First, statistical significance should not be used as the primary decision rule even though the program (FATS) generates such tests. Second, this procedure should ideally be used on very long time series data sets (hundreds of time points).[4] Third, the researcher should assess what size of effect could be obtained by a random process. This can be determined by running a number of random series of the same length as the actual data in the study and adjusting the decision rule to substantially exceed that baseline value. It can be shown that this problem is mainly due to fitting such a large number of potential functions in a stepwise procedure, such that this bias does not pose a major problem when fitting an a priori model to the data (VanLear & Li, 2005). Finally, it is wise to limit the extraction of sinusoidal components to a small number of very strong components and to be wary of patterns that do not have a clear theoretical interpretation. Nevertheless, because this procedure can be shown to detect real patterns in data obscured by random noise, it should not be completely shunned as an inductive method, especially for uncovering hidden periodicity. We like to use this procedure in conjunction with the examination of ACFs and PACFs for detecting hidden patterns in time series data (VanLear & Li, 2005).[5] Sometimes researchers will "smooth" their time series data and apply these methods to the smoothed series. This is typically referred to as spectral analysis (Gottman, 1981). When analyzed, a smoothed series will not tend to show the random spikes in the periodogram. One looks for the frequencies with the greatest "spectral densities," which are the frequency regions with many adjacent frequencies that account for most of the overall periodic behavior in the series.

Figure 4.5 provides a graphic view of this procedure on real data. A dyad interacted over a mediated channel for about 10 minutes regarding a controversial political issue. Each member was coded for smiling (smiles, doesn't smile) every half second, and the number of smiling units in every 2 seconds of interaction comprised the measure. Each member of the dyad in this analysis is represented by a time series of 248 repeated measures.

The data for person A is graphed in Figure 4.5 along with the first three sinusoidal components extracted using the stepwise procedure. The periodogram for this series is presented in Figure 4.6. The first component extracted has a period of 64.7 (32.3 seconds). This component explains 20.3% of the variance in the series, which well exceeds what would be expected if this were a random series. After the first component was analyzed, it was removed and the residual series was analyzed. The second sinusoidal component has a period of 49 units (24.5 seconds) and accounts for an additional 10% of the variance of the series (total R^2 = .305, ΔR^2 = .102), which also exceeds what would be expected in a random series of this length. The third component in the stepwise procedure has a

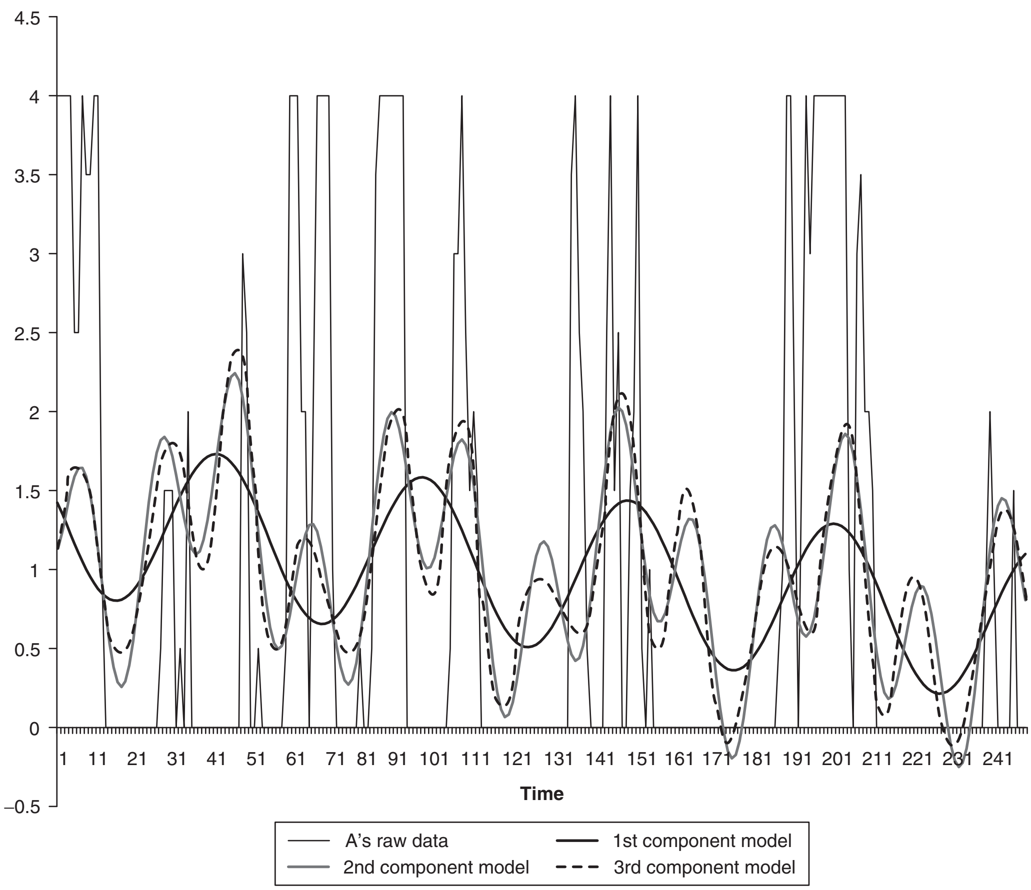

Figure 4.5 A Three-Component Stepwise Fourier Analysis of Smiling for Speaker

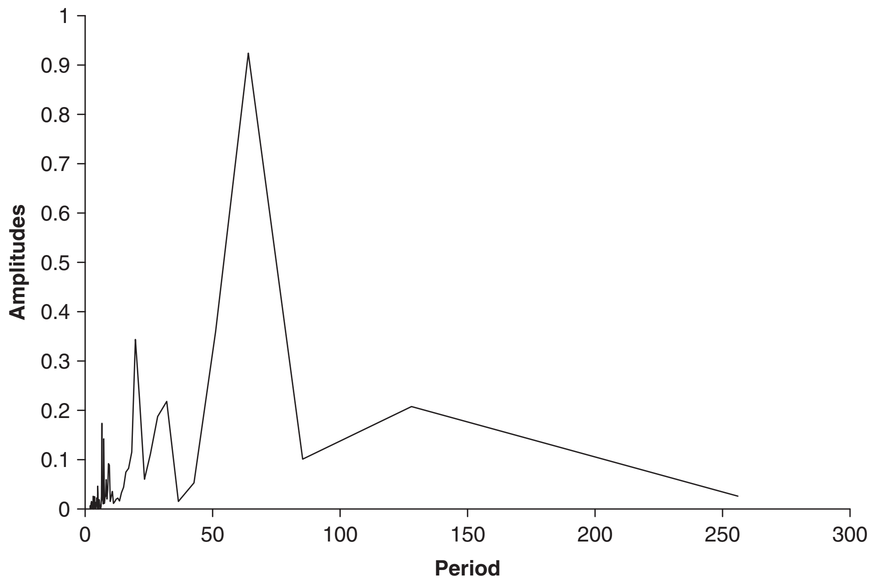

Figure 4.6 The Periodogram for Speaker A's Time Series of Smiling

period of 19.3 units (9.65 minutes) and explains an additional 6.7% of the variance of the series (total $R^2 = .371$, $\Delta R^2 = .066$). Although this exceeds the average amount of explained variance for random series of this length, it did so by a small amount. Therefore, the third component must be interpreted with more caution. Figure 4.6 shows what the periodic function formed from the summation of these three components looks like as plotted against the raw data. Figure 4.7 shows the ACFs for speaker A on the same series. The significant negative ACF (–.27) at lag 31 and the corresponding positive ACFs at lags 58, 59, and 60 (r_s = .18, .21, and .18, respectively) correspond closely to the first component revealed through the stepwise Fourier analysis. This evidence suggests that speaker A's smiling behavior follows a periodic cyclical pattern.

Suppose that we have a theoretical reason to expect a cycle of a given frequency/period. A sinusoidal function representing that frequency can be fit to the data and assessed for the amount of variance explained. Because we are fitting one function to the data, this procedure does not contain the same potential for false positives as the stepwise procedure (VanLear & Li, 2005). For example, one might predict that alcohol advertising will follow the same seasonal cycle as TV programming. Or if one has daily measures of mood or emotion, one might predict a weekly cycle with lows on Mondays and highs on Fridays and that the same cycle might be evidenced in certain spontaneous nonverbal behaviors known to communicate emotion. One could fit the cycle found in the emotion series to the contemporaneous behavioral series.

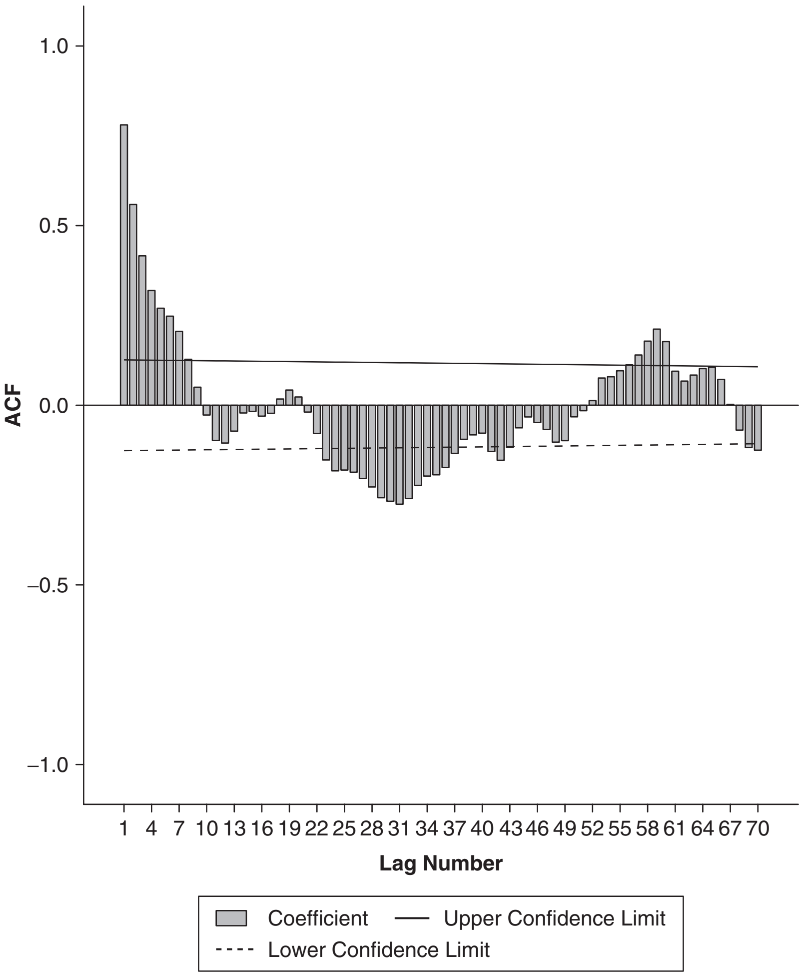

Figure 4.7 ACFs for Speaker A's Smiling Series

VanLear and Watt (1996) presents procedures in which these techniques can be used in an experimental design to detect how an experimental stimulus affects not just the level of a variable but also the nature of cycles that a whole time series displays (e.g., frequency, amplitude). Likewise, models built on one part of a time series can be used to forecast the future values of the series. In a related design (interrupted time series), the technique can be used to model changes in processes occurring after a significant event, whether the event was manipulated by the researcher or was a naturally occurring phenomenon.

Finally, one of the most common ways to use frequency domain time series is to examine the relationship between two concurrent time series. This can be accomplished through *cross-spectral analysis.* In this approach, the DFTs are calculated for each variable's series, and the values of the two periodograms can be multiplied to produce a "cross-periodogram" (or if the series are smoothed, a "cross-spectrum"). A cross-periodogram will show large peaks at the frequencies that the two series have in common. Cross-spectral analysis analyzes the cross-amplitude and the relative phase of the two series at each frequency. A standardized measure analogous to the square of the correlation at each frequency is given by the squared "coefficient of coherence" (Gottman, 1981; StatSoft, 2003).[6] Each series has a gain value for each frequency in a cross-spectral analysis. The gain values for each frequency are interpreted like the standardized betas in a regression for that series at that frequency. The "phase shift" estimates are measures of the extent that one series "leads" the other at each frequency. SPSS is capable of conducting cross-spectral analysis of time series data.

For example, we can take the sinusoidal function that best fit the data for person A in the above example and fit one-, two-, and three-component models to the time series representing person B's smiling. Figure 4.8 graphs the fit of these functions. The first sinusoidal component from speaker A's model (period = 64.67) explained 4.9% of the variation in speaker B's time series. The two-component model from A's series explained 13.8% (ΔR^2 = .09) of the variation in B's smiling over time. The three-component model from A's series explains 23.2% (ΔR^2 = .094) of the variance of B's series. The fit is both significant and meaningful. Speaker B's smiling appears to match

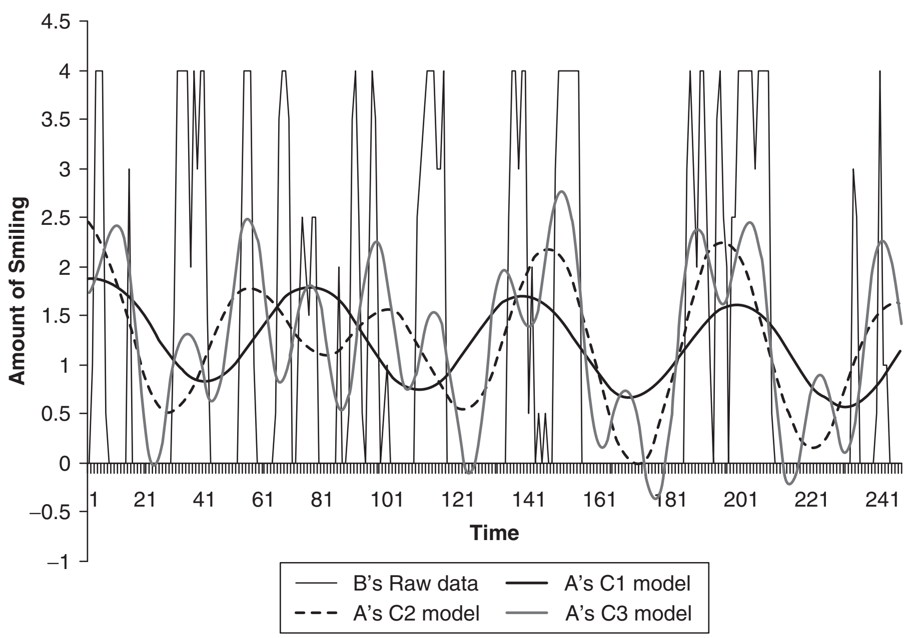

Figure 4.8 Speaker A's Fourier Model's Fit to Speaker B's Data

the rhythm of speaker A's smiling. We also matched speaker B's model to speaker A's data as well. The first component of speaker B's model explained 5.4% of A's variance, the second component explained 12.8% of A's variance, and the third component explained 13.8% of A's variance. Speaker A seems to match speaker B's pattern, especially the second component (a cycle with a period of 19), although the fit is not quite as good as when matching speaker A's model to speaker B's data.

Figure 4.9 displays the cross-periodogram for speaker A's and speaker B's smiling from the above example. This figure shows the largest peaks (cross amp = .24) at the adjacent Fourier periods of 64 (ω = .098) and 51.2 (ω = .123) and a smaller peak at period 19.7 (ω = .319, cross amp = .18). These generally correspond to the peaks at the frequencies identified in modeling the individual series. We can conclude that these two speakers tend to synchronize their smiling behavior by matching their periodic behaviors.

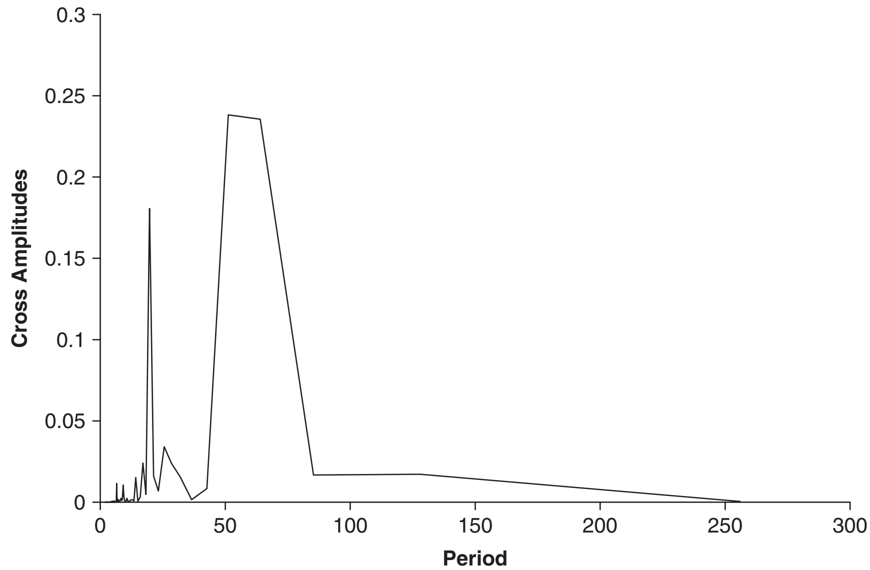

Figure 4.9 The Cross-Periodogram Between Speakers A's and B's Time Series of Smiling Behavior

Pooling and Aggregation of Time Series

The strength of time series analyses is assessing complex patterns and relationships over time and forecasting future values, not generalizing to a population of similar cases. If one has a large number of cases and a small number of replications, then one can use either structural equation modeling, where each time point is a different set of variables, or multilevel modeling (see Chapter 3, this volume). The sophisticated types of temporal patterns modeled by individual time series are not easily assessed with these methods. Time series analyses are often useful when one has a large number of repeated measures on a single or small number of cases. An organizational communication consultant may have measures on a single client company over an extended period of time. A marital therapist may have interaction data on a single client couple. A media consultant may have data on the client station or network over an extensive period of time. Or we may have a research question that we wish to explore on a small number of cases (e.g., the major networks) over an extensive number of time periods. Many time series are conducted on data that are already pooled or aggregated across cases (e.g., Nielsen ratings, public opinion poll results, economic indices, and crime statistics). However, often communication scholars wish to gather data over many time points across many cases and conduct sophisticated analyses of temporal processes. In such cases, the data must be aggregated or pooled.

There are two ways in which data can be pooled or aggregated: preanalysis or postanalysis. Pooling time series from different cases before the analysis has the great advantage that only one set of time series analyses needs to be conducted. If there are a large number of cases (e.g., hundreds), this is a tremendous advantage. Imagine conducting a whole set of analyses like those presented in the above examples hundreds of times. However, in order to pool the series before analysis, certain conditions should be met. First, the various series representing different cases must be exactly contemporaneous and in sync. For example, if hundreds of TV viewers are measured every 3 seconds as they watch the same half-hour program and each measure in time is synchronized to the exact same point in the program, then the data meet this first condition. However, if people do daily monitoring of their communication with their significant other over a 3-month period, each person's time series is unlikely to be synchronized with other participants if the relationships have been in existence for different periods of time and different participants begin their data recording on different days. The second condition is that the different time series must be homogeneous with regard to process across cases. If some series trend, some are first-order autoregressive processes, some are first-order moving average processes, and some show seasonal cycles that others do not, then aggregating before analysis will lose important between-case variance and may be inappropriate. If one can meet the first condition, but is unsure about the second, then a random subsample of the cases could be assessed using individual time series as explained here. If all or nearly all of the cases have the same kind of pattern of serial dependency, then the data could be pooled and the pooled data assessed for that pattern.

When the conditions for preanalysis aggregation cannot be met, then postanalysis aggregation can be employed. In this situation, each case is analyzed separately as if it were a study unto itself, and then the results are aggregated using meta-analysis if the results are homogeneous. If the results are not homogeneous, they can be input as scores for each case and their variation analyzed using traditional statistics to search for moderating variables. For example, Street and Cappella (1989) used this approach in the time domain to analyze the adaptation of children to adults' speech characteristics, and VanLear (1991) used this approach in the frequency domain to analyze dialectic cycles of openness in relationships.

Conclusion

Gathering and analyzing time series data presents communication researchers with a unique set of challenges. Meeting these challenges requires time, effort, training, and skill. Nevertheless, we believe that this extra training and effort are worth the payoff. Communication is a dynamic process, and to be true to this axiom, we must be willing to model communication over time. The present chapter is an introduction. Scholars interested in the dynamic modeling of communication processes are advised to study these techniques in depth.

Notes

1. Sometimes time series analyses are distinguished from repeated measures with time series analysis referring to methods used to analyze data consisting of large numbers of replications (at least 30 or so) and repeated measures referring to data with smaller numbers of replications across a large number of cases. We will use the term "time series" in the broader sense and use the term "pooled time series" or "cross-sectional time series" for the latter group of analyses. The analytical methods we focus on are individual time series.

2. However, even experimental designs are not well suited for identifying reciprocally or mutually causal processes involving feedback loops, which are usually held to be central to viewing communication as a process (Berlo, 1960, 1977; VanLear, 1996). Likewise, some variables cannot be easily or ethically manipulated experimentally while retaining ecological validity.

3. Observations in a time series are not independently sampled even if there is no autocorrelation or when the statistical dependency is removed. As a result, inferential tests do not provide evidence of generalizability to other cases, though they may be used to forecast values of future observations of the cases analyzed.

4. If significance tests are used, the problem actually gets worse with longer time series, whereas if effect size is used (R^2), the problem diminishes as the length of the series increases.

5. SPSS uses a default of 16 lags for ACFs and PACFs, which is usually adequate for detecting AR and MA processes but is often too short to detect long cycles. Because of the large number of ACFs possible in a series, ACFs can also

lead to the detection of false positives if significance of one lag is a sufficient criterion for the presence of a pattern.

6. One should be careful in interpreting these values by themselves, because one can obtain large values for coherency when the spectral density values of both series (the divisor when coherency is computed) are both small, indicating no strong periodic components in the data.

References

Altman, I., Vinsel, A., & Brown, B. (1981). Dialectic conceptions in social psychology: An application to social penetration and privacy regulation. In L. Berkowitz (Ed.), *Advances in experimental social psychology: Vol. 14* (pp. 76–100). New York: Academic Press.

Barnett, G. A., Chang, H. J., Fink, E. L., & Richards, W. D. (1991). Seasonality in television viewing: A mathematical model of cultural processes. *Communication Research, 18,* 755–772.

Berlo, D. K. (1960). *The process of communication.* New York: Holt, Rinehart & Winston.

Berlo, D. K. (1977). Communication as process: Review and commentary. In B. D. Ruben (Ed.), *Communication yearbook I* (pp. 11–27). New Brunswick, NJ: Transaction Books.

Blood, D. J., & Phillips, P. C. B. (1995). Recession headline news, consumer sentiment, the state of the economy and presidential popularity—a time-series analysis 1989–1993. *International Journal of Public Opinion Research, 7,* 1–22.

Bloomfield, P. (1976). *Fourier analysis of time series: An introduction.* New York: John Wiley.

Box, G. E., & Jenkins, G. M. (1976). *Time series analysis: Forecasting and control.* San Francisco: Holden-Day.

Brockwell, P. J., & Davis, R. A. (2002). *Introduction to time series and forecasting* (2nd ed.). New York: Springer.

Brosius, H. B., & Kepplinger, H. M. (1992). Linear and nonlinear models of agenda-setting in television. *Journal of Broadcasting & Electronic Media, 36,* 5–23.

Campbell, D. T., & Kenny, D. A. (1999). *A primer on regression artifacts.* New York: Guilford.

Cappella, J. N. (1981). Mutual influence in expressive behavior: Adult-adult and infant-adult dyadic interaction. *Psychological Bulletin, 89,* 101–132.

Cappella, J. N. (1996). Dynamic coordination of vocal and kinesic behavior in dyadic interaction: Methods, problems, and interpersonal outcomes. In J. H. Watt & C. A. VanLear (Eds.), *Dynamic patterns in communication processes* (pp. 353–386). Thousand Oaks, CA: Sage.

Chapple, E. D. (1970). *Culture and biological man: Explorations in behavioral anthropology.* New York: Holt, Rinehart & Winston.

Chatfield, C. (1989). *The analysis of time series: An introduction* (4th ed.). London; New York: Chapman and Hall.

Cromwell, J. B. (1994). *Multivariate tests for time series models.* Thousand Oaks, CA: Sage.

Cromwell, J. B., Labys, W. C., & Terraza, M. (1994). *Univariate tests for time series models.* Thousand Oaks, CA: Sage.

Dearing, J. W., & Rogers, E. M. (1996). *Agenda-setting.* Thousand Oaks, CA: Sage.

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

Fan, D. P. (1988). *Prediction of public opinion from the mass media: Computer content analysis and mathematical modeling.* New York: Greenwood Press.

Fan, D. P., & Cook, R. D. (2003). A differential equation model for predicting public opinions and behaviors from persuasive information: Application to the index of Consumer Sentiment. *Journal of Mathematical Sociology, 27,* 29–51.

Fisher, B. A. (1978). *Perspectives on human communication.* New York: Macmillan.

Funkhouser, G. R. (1973). The issues of the sixties: An exploratory study in the dynamics of public opinion. *Public Opinion Quarterly, 37,* 62–75.

Gerbner, G., Gross, L., Morgan, M., & Signorielli, N. (1986). Living with television: The dynamics of the cultivation process. In J. Bryant & D. Zillman (Eds.), *Perspectives on media effects* (pp. 17–40). Hillsdale, NJ: Lawrence Erlbaum Associates.

Gonzenbach, W. J. (1996). *The media, the president, and public opinion: A longitudinal analysis of the drug issue, 1984–1991.* Mahwah, NJ: Lawrence Erlbaum Associates.

Gonzenbach, W. J., & McGavin, L. (1997). A brief history of time: A methodological analysis of agenda setting. In M. McCombs, D. L. Shaw, & D. Weaver (Eds.), *Communication and democracy: Exploring the intellectual frontiers in agenda-setting theory* (pp. 115–136). Mahwah, NJ: Lawrence Erlbaum Associates.

Gottman, J. M. (1981). *Time-series analysis: A comprehensive introduction for social scientists.* Cambridge, UK: Cambridge University Press.

Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica, 37,* 424–438.

Heath, R. A. (2000). *Nonlinear dynamics: Techniques and applications in psychology.* Mahwah, NJ; London: Lawrence Erlbaum Associates.

Huston, T. L., & Vangelisti, A. L. (1991). Socioemotional behavior and satisfaction in marital relationships: A longitudinal study. *Journal of Personality and Social Psychology, 41,* 721–733.

Jaffe, J., & Feldstein, S. (1970). *Rhythms of dialogue.* New York: Academic Press.

Jones, R. H., & Ackerson, L. M. (1990). Serial correlation in unequally spaced longitudinal data. *Biometrika, 77,* 271–731.

Mansur, K. A., & Shoemaker, H. H. (1999). The impact of changes in the current population survey on time-in-sample bias and correlations between rotation groups. *Proceedings of the survey methods section of the American Statistical Association* (pp. 180–185). Retrieved January 23, 2006, from http://www.amstat.org/sections/srms/Proceedings/papers/1999_028.pdf

McCleary, R., & Hay, R. (1980). *Applied time series analysis for the social sciences.* Beverly Hills, CA: Sage.

McCombs, M., & Zhu, J. H. (1995). Capacity, diversity, and volatility of the public agenda: Trends from 1954 to 1994. *Public Opinion Quarterly, 59,* 495–525.

McDowall, D. (1980). *Interrupted time series analysis.* Beverly Hills, CA: Sage.

Meadowcroft, J. M. (1996). Attention span cycles. In J. H. Watt & C. A. VanLear (Eds.), *Dynamic patterns in communication processes* (pp. 255–276). Thousand Oaks, CA: Sage.

Menard, S. (2002). *Longitudinal research* (2nd ed.). Thousand Oaks, CA: Sage

Ostrom, C. W. (1990). *Time series analysis: Regression techniques* (2nd ed.). Newbury Park, CA: Sage.

Pan, Z., & McLeod, J. M. (1991). Multilevel analysis in mass communication research. *Communication Research, 18,* 140–173.

Poole, M. S. (1981). Decision development in small groups: A comparison of two models. *Communication Monographs, 50,* 206–232.

Poole, M. S. (2000). *Organizational change and innovation processes: Theory and methods for research.* Oxford, UK; New York: Oxford University Press.

Poole, M. S., & Roth, J. (1989). Decision development in small groups V: Test of a contingency model. *Human Communication Research, 15,* 549–589.

Poole, M. S., Seibold, D. R., & McPhee, R. D. (1996). A structurational approach to theory-building in group decision-making research. In R. Y. Hirokawa & M. S. Poole (Eds.), *Communication and group decision-making* (2nd ed., pp. 114–146). Thousand Oaks, CA: Sage.

Rice, R. E., & Katz, J. E. (2001). *The internet and health communication: Experiences and expectations.* Thousand Oaks, CA: Sage.

Rogers, E. M., Dearing, J. W., & Chang, S. (1991). AIDS in the 1980s: The agenda setting process for a public issue. *Journalism Monographs, 126.*

Rogers, E. M., & Shoemaker, F. (1973). *Communication of innovations.* Glencoe, IL: Free Press.

Sayrs, L. W. (1989). *Pooled time series analysis.* Newbury Park, CA: Sage.

Shah, D. V., Watts, M. D., Domke, D., & Fan, D. P. (2002). News framing and cueing of issue regimes—explaining Clinton's public approval in spite of scandal. *Public Opinion Quarterly, 66*(3), 339–370.

Shah, D. V., Watts, M. D., Domke, D., Fan, D. P., & Fibison, M. (1999). News coverage, economic cues, and the public's presidential preferences, 1984–1996. *Journal of Politics, 61,* 914–943.

Shoemaker, P. J., Wanta, W., & Leggett, D. (1989). Drug coverage and public opinion, 1972–1986. In P. J. Shoemaker (Ed.), *Communication campaigns about drugs: Government, media, and the public* (pp. 67–80). Hillsdale, NJ: Lawrence Erlbaum Associates.

Shumway, R. H. (1988). *Applied statistical time series analysis.* Englewood Cliffs, NJ: Prentice-Hall.

Shumway, R. H., & Stoffer, D. S. (2000). *Time series analysis and its applications.* New York: Springer.

Slater, M. D., Snyder, L. B., & Hayes, A. F. (2006). Thinking and modeling at multiple levels: The potential contribution of multilevel modeling to communication theory and research. *Human Communication Research, 32,* 375–384.

Smith, T. (1980). America's most important problems—a trend analysis, 1946–1976. *Public Opinion Quarterly, 44,* 164–180.

StatSoft, Inc. (2003). *Time series analysis.* Retrieved August 8, 2005, from www.statsoft.com.

Street, R. L., & Cappella, J. N. (1989). Social and linguistic factors influencing adaptation in children's speech. *Journal of Psycholinguistic Research, 18,* 497–519.

Stryker, J. E. (2003). Media and marijuana: A longitudinal analysis of news media effects on adolescents' marijuana use and related outcomes, 1977–1999. *Journal of Health Communication, 8,* 305–328.

Tabachnick, B. G., & Fidell, L. S. (2001). *Using multivariate statistics* (4th ed.). Needham Heights, MA: Allyn & Bacon.

Taris, T. W. (2000). *A primer in longitudinal data analysis.* London: Sage.

Tedrow, L. M., & Mahoney, E. R. (1979). Trends in attitudes toward abortion: 1972–1976. *Public Opinion Quarterly, 43,* 181–189.

Trumbo, C. (1995). Longitudinal modeling of public issues: An application of the agenda-setting process to the issue of global warming. *Journalism and Mass Communication Monographs, 152.*

VanLear, C. A. (1987). The formation of social relationships: A longitudinal study of social penetration. *Human Communication Research, 13,* 299–322.

VanLear, C. A. (1991). Testing a cyclical model of communicative openness in relationship development: Two longitudinal studies. *Communication Monographs, 58,* 337–361.

VanLear, C. A. (1996). Communication process approaches and models: Patterns, cycles, and dynamic coordination. In J. H. Watt & C. A. VanLear (Eds.), *Dynamic patterns in communication processes* (pp. 35–70). Thousand Oaks, CA: Sage.

VanLear, C. A., Brown, M., & Anderson, E. (2003, May). *Communication, social support, and emotional quality of life in the twelve-step sobriety maintenance process: Three studies.* A paper presented at the annual meeting of the International Communication Association, San Diego, CA.

VanLear, C. A., & Li, S. (2005, May). *Dynamic modeling of scheduled and nonscheduled communication processes.* A paper presented at the annual meeting of the International Communication Association, New York, NY.

VanLear, C. A., & Mabry, E. A. (1999). Testing contrasting interaction models for discriminating between consensual and dissentient decision-making groups. *Small Group Research, 30,* 29–58.

VanLear, C. A., Sheehan, M., Withers, L., & Walker, R. (2005). AA online: The enactment of supportive computer mediated communication. *Western Journal of Communication, 69,* 5–26.

VanLear, C. A., & Watt, J. H. (1996). A partial map to a wide territory. In J. H. Watt & C. A. VanLear (Eds.), *Dynamic patterns in communication processes* (pp. 3–34). Thousand Oaks, CA: Sage.

Warner, R. M. (1996). Coordinated cycles in behavior and physiology during face-to-face social interactions. In J. H. Watt & C. A. VanLear (Eds.), *Dynamic patterns in communication processes* (pp. 327–351). Thousand Oaks, CA: Sage.

Watt, J. H. (1994). Detection and modeling of time-sequenced processes. In A. Lang (Ed.), *Measuring psychological responses to media messages* (pp. 181–207). Hillsdale, NJ: Lawrence Erlbaum Associates.

Watt, J. H. (1998). *FATS: Fourier analysis of time series program user's guide.* Troy, NY: Rensselaer Polytechnic Institute.

Watt, J. H., Mazza, M., & Snyder, L. (1993). Agenda-setting effects of television news coverage and the effects decay curve. *Communication Research, 20,* 408–435.

Watt, J. H., & VanLear, C. A. (1996). *Dynamic patterns in communication processes.* Thousand Oaks, CA: Sage.

Winter, J. P., & Eyal, C. (1991). Agenda-setting for the civil rights issue. In D. L. Protess & M. E. McCombs (Eds.), *Agenda setting: Readings on media, public opinion, and policymaking. Communication textbook series* (pp. 101–107). Hillsdale, NJ: Lawrence Erlbaum Associates.

Yanovitzky, I. (2002a). Effect of news coverage on the prevalence of drunk-driving behavior: Evidence from a longitudinal study. *Journal of Studies on Alcohol, 63*(3), 342–351.

Yanovitzky, I. (2002b). Effects of news coverage on policy attention and actions—a closer look into the media-policy connection. *Communication Research, 29,* 422–451.

Yanovitzky, I., & Bennett, C. (1999). Media attention, institutional response, and health behavior change—the case of drunk driving, 1978–1996. *Communication Research, 26,* 429–453.

Yanovitzky, I., & Blitz, C. L. (2000). Effect of media coverage and physician advice on utilization of breast cancer screening by women 40 years and older. *Journal of Health Communication, 5,* 117–134.

Yanovitzky, I., & Stryker, J. (2001). Mass media, social norms, and health promotion efforts: A longitudinal study of media effects on youth binge drinking. *Communication Research, 28,* 208–239.

Zhu, J. (1992). Issue competition and attention distraction: A zero-sum theory of agenda-setting. *Journalism & Mass Communication Quarterly, 68,* 825–836.

Zhu, J., Watt, J. H., Snyder, L. B., Yan, J., & Jiang, Y. (1993). Public issue priority formation: Media agenda-setting and social interaction. *Journal of Communication, 43,* 8–29.