1 Introduction

1.1 WHY SHOULD YOU READ THIS TEXT?

You cannot describe a movie by looking at a single frame of film or experience an opera by listening to a single note. So why should we rely on snippets of information plucked from single points in time when trying to explain the complex processes that animate our society? There are things you can learn from fragments, to be sure, but to develop a comprehensive understanding of the world we need to know what lies between the snapshots and how they are connected. Time series analysis is about building models of the dynamic processes that govern the evolution of social, economic, and political phenomena in order to describe relationships, test hypotheses, forecast the future, and better understand the world around us.

The word dynamic is sometimes used to describe people or fields that are full of energy and new ideas. We use it in the more elementary sense – to describe a process that changes over time. In that sense, almost everything is dynamic – economics, politics, and culture. Many dynamic processes do not evolve independently but are embedded in dynamic systems. The economy can determine the direction of political events, politics can affect the kinetics of culture, and cultural changes can alter the flows of economic activity. Challenges arise in studying these dynamic processes because many standard statistical techniques assume data points are independent snapshots, frozen in time and isolated from one another. The methods presented in this book allow social scientists to capture a richer representation of the world, as a complex but comprehensible moving picture.

A **time series** is a sequence of observations ordered over time. **Time series analysis** is the application of statistical tools to analyze such data. Some processes captured in time series data are naturally discrete, for example, there is one presidential election in the United States every four years. However, many data sets are created by sampling from continuous dynamic processes at discrete intervals. For example, people always hold

https://doi.org/10.1017/9781108290326.001 Published online by Cambridge University Press

opinions about the economy and the president in their heads, but public opinion time series are created from surveys administered at discrete points in time.

Time series analysts fit statistical models to time series data. There are many methodological approaches that have been devised for this enterprise, including Box–Jenkins ARIMA (autoregressive integrated moving average) analysis (Box and Jenkins, 1970), ordinary least squares (OLS) dynamic regression (Banerjee et al., 1993; Hendry, 1995), vector autoregression (VAR) (Sims, 1980), state-space models (Kalman, 1960), and others. A comprehensive text that covered all of these approaches would be sprawling and impractical. Instead, authors of time series texts typically take one of two tacks: write a survey text or take a deep dive into a particular problem or approach.

Survey texts offer short introductions to a variety of approaches. There might be a thirty- to forty-page introduction to ARIMA models in one chapter and a treatment of VAR in another. These are different approaches to time series data developed for different purposes. The construction of each type follows different philosophies and the analyst can be required to use different tools to estimate the models. However, a good survey text will introduce the reader to both. Particularly useful survey texts for social scientists include Box-Steffensmeier et al. (2014) and Enders (2015).

Many books tackle specific models and approaches in deep detail. In *Multiple Time Series Analysis* – part of the Quantitative Applications in the Social Sciences (QASS) series curated by Sage – Brandt and Freeman (2009) provide a detailed introduction to VAR models that describes their utility and explains, with examples, how to build a VAR. Lütkepohl (2005) and Kilian and Lütkepohl (2017) provide in-depth treatments of the topic as well. These books serve as reference manuals to the mathematical properties, estimation challenges, and extensions of the basic model. Each book focuses on VAR analysis, but one is much more accessible to beginners. Other examples of detailed treatments include Choi (2015) on unit root processes, Francq and Zakoian (2019) on Autoregressive Conditional Heteroskedasticity (ARCH-GARCH) models, and Juselius (2006) on the cointegrated VAR model. Some of these can be used in more advanced classes and short workshops on specific topics, but many of them are most useful for applied researchers who need more detailed information than can be found in a survey text. This text is a compromise between these two approaches, providing more detail than a survey text, while being as accessible to newcomers as short introductions and chapter-length treatments.

We provide a practical guide to applied time series analysis using ordinary least squares (OLS) regression. OLS is covered in most introductory

https://doi.org/10.1017/9781108290326.001 Published online by Cambridge University Press

statistics courses taken by graduate and undergraduate students. Our expectation is that you have taken such a course already. Many important tools in the time series analysts' tool kit are built on this basic scaffolding. This book serves as a bridge to the complexities of time series analysis by building on material you have already been exposed to: estimation and statistical inference using OLS. If you can estimate an OLS regression model and interpret a slope coefficient, this book is for you.

This book can be used in many ways. It can be assigned in a semester-long course on time series analysis, as part of a workshop on time series topics, or used as a desk resource for social scientists who think they will (or might) need to estimate time series models for their research. We have endeavored to make the writing direct and have included a number of examples (along with the data and code needed to replicate them), which will allow you to teach yourself if necessary. To our knowledge, this book is the first of its kind: It draws connections between the various OLS-based approaches to time series analysis and provides insights into the practical application of these tools. And in our view, it is sorely needed.

There are two principles that guide our approach to time series analysis and underlie our presentation of the material. First, *uncertainty is ubiquitous in applied time series analysis*. There is a facile quality to many statistical texts. The authors present the equations for statistical estimators and describe when and how they should be used, but fail to inform the reader about the challenges they are likely to encounter. For example, the standard approach to time series analysis begins with testing whether a series has a unit root – a specific type of nonstationarity where shocks persist indefinitely rather than dying out over time. Tests are presented to detect unit roots, but there is little discussion of strategies for applying these tests, handling competing findings, or making knife-edged modeling decisions when test results are borderline. The analyst is left to navigate these uncertain waters on their own.

Time series analysis is often messy. For example, we often find that the dynamics characterizing the time series with which we work – shorter, intermediate frequency time series – are not easily classified. In fact, classifying the dynamics of individual time series is often frustrating. Multiple tests can leave you with competing inferences. This uncertainty is typically hidden in subsequent analyses. Other choices in working with time series data also engender uncertainty. What should you do when different diagnostic tests give conflicting results about the number of lags to include in a model? What does it mean when different model-fit criteria point to different dynamic specifications? These are the types of questions that commonly arise, and advice is scarce in statistical texts about

https://doi.org/10.1017/9781108290326.001 Published online by Cambridge University Press

what to do. There are times when you will have to make judgment calls. This book offers practical guidance for principled decision-making in these cases.

Much of *our* scholarship takes on these kinds of problems. We developed a bounds approach to inference that builds the uncertainty in univariate classification into hypothesis tests about the relationship between time series (Webb, Linn, and Lebo, 2019, 2020). The lesson is that, whenever possible, uncertainty should be explicitly accounted for in the modeling process. When the tools to do so aren't available, researchers should make that uncertainty a transparent part of how analyses are described and reported – especially when results are contradictory.

As a practical guide, this text is designed to help you navigate the challenges you will encounter in applied time series analysis. There are many forks on this road, and we provide step-by-step guidance to help keep you on the right path. The topics include:

- How to classify time series
- How to determine whether single- or multiple-equation analysis is necessary
- How to build an appropriate model specification including:
  - How to decide the appropriate deterministic features
  - How to model the effects of events
  - How to determine dynamic structure, that is, choosing the number of lags
- How to test whether a model satisfies the necessary assumptions
- How to draw detailed inferences.

Our second principle is that *theory and data interplay in important ways in time series analysis*. This can be uncomfortable because social science disciplines often stress the primacy of theory in model building. But theory seldom identifies the dynamic structures of relationships. For example, theory may suggest that consumer sentiment affects public views of the president, but it is unlikely to be developed enough to specify the number of lags that belong in the model. Depending on how the data are measured, different numbers of lags may be appropriate. Throughout the text, we determine the dynamic structures of the models from the data and rely on theory to select the variables in the models. Even if you have a conjecture about what dynamic specification might be appropriate for a statistical model, you must subject this conjecture to formal testing before imposing it on the data. Otherwise, you risk misleading yourself and your audience.

https://doi.org/10.1017/9781108290326.001 Published online by Cambridge University Press

1.2 WHY DO YOU NEED AN ENTIRE TEXT ON REGRESSION WITH TIME SERIES DATA?

Although this book relies on linear regression analysis, the nature of temporal data demands an approach to model specification, diagnostics, inference, and interpretation that are all different from those in the cross-sectional context. The fundamental characteristic that distinguishes time series from cross-sectional data is time dependence. Time series observations are rarely independent. History matters – previous values of a time series inform the present value. This fact fundamentally complicates regression analysis and requires specialized terminology. Additionally, while systems of simultaneous equations can arise in any context, they are particularly common in time series analysis where variables often evolve together over time. Despite these complexities, the core objective remains unchanged: to estimate a statistical model that approximates the DGP and enables generalization to the population.

Time dependence has a number of implications for analyses. First, it requires thinking differently about model specification – the choices you make about which variables to include in a regression model and how they are related. More often than not, the effect of a change in a regressor propagates over time. Models will need to incorporate lags. This means incorporating regressors not only at time t but also at $t - s$, as well as including lagged values of the dependent variable as regressors. The inclusion of lagged dependent variables creates a *dynamic* regression, which violates the Gauss–Markov assumption of zero conditional mean (strict exogeneity). Consequently, dynamic regression estimates become biased and inefficient in small samples. This necessitates an alternative set of assumptions for time series OLS analysis. When these alternative assumptions hold, OLS estimators maintain consistency.

Second, the approach to diagnostic analyses of regression assumptions differs from the cross-sectional case. Because time dependence is a feature of time series data, you need to be particularly concerned with violations of the Gauss–Markov assumption that errors are independent. Serially correlated errors pose the most common threat to inference in time series regression. But this is typically not a problem to be *corrected for* with fixes like weighted least squares or robust standard errors. Applying these *solutions* does not address the underlying problem – omitted lags of the regressors and regressands – and obfuscates important features of dynamic relationships. Instead, serially correlated errors should generally be taken as evidence indicating dynamic misspecification – which produces omitted variable biases – and indicates a need to reconsider the lag structures of models.

https://doi.org/10.1017/9781108290326.001 Published online by Cambridge University Press

Third, different patterns of time dependence require different strategies for estimation and inference. You need to characterize the temporal patterns that describe the data. In particular, you need to determine whether the time series are best characterized as stationary processes, which have constant means, variances, and autocovariances over time, or whether they are not stationary. Unit root processes – those that contain random trends – present particular challenges since they are prone to the spurious regression problem. That is, independent unit root processes often appear related to each other even though they are not. Even when two unit root processes are related, model specification requires more care – for example, critical values for hypothesis tests will tend to be non-standard.

Fourth, interpretation of dynamic regression models requires more than examining single coefficients. A dynamic regression model provides information about how variables affect one another in the short, medium, and long term. These models show how changes in a regressor play out in a regressand over time, how those effects decay, and what types of equilibria characterize long-run relationships, if they exist. But this added complication is really an advantage, making rich interpretations of dynamic relationships possible.

Finally, time dependence means that analysts need to consider whether single- or multiple-equation analyses are necessary for valid inference. Single-equation analysis requires that no information about the relationship of interest is lost by ignoring the model generating the regressors, that is, there is weak exogeneity for the parameters of interest. Meeting the weak exogeneity assumption is needed for consistent inference. Most texts provide limited discussion of this assumption, often treating it as a decision made in advance of the analysis. Indeed, there are a number of misconceptions about the conditions that permit single-equation analysis. So much so, that we devote a chapter to the topic. When weak exogeneity is violated, multiple-equation analysis is necessary. But this modeling approach puts additional demands on both the data and the analyst.

These complications mean time series analysts require a different nomenclature. The terms stationary and unit root figure prominently in the vocabulary, and a clear understanding of these terms is essential to robust time series analysis. But other terms like persistence, integration, deterministic and stochastic trends, structural breaks, cointegration, and different varieties of exogeneity, to name a few, will also play an important role in model-building. Interpretation of model results requires expanding your vocabulary to include impact multipliers, long-run multipliers, and impulse response functions. These terms likely don't mean a lot to you now. That's fine. We will introduce them and explain how and when they are relevant in the chapters to come.

https://doi.org/10.1017/9781108290326.001 Published online by Cambridge University Press

In short, the opportunities that come with time series data also bring a number of challenges beyond those faced in a cross-sectional context. We walk you through each of them. But first, let's get your time series data in order.

1.3 HOW DO YOU ASSEMBLE TIME SERIES DATA?

Time series analysis requires careful consideration of how to construct your data set to answer your research question. You'll need to make several key decisions: what time period to analyze, what time unit to use for measurement, and how to operationalize your variables. These choices are interconnected and often constrained by data availability. This section guides you through these crucial decisions as you prepare your data for analysis.

Every time series analysis begins with a research question and a theory about relationships between variables. Before collecting data, you need to think carefully about which variables will help you answer your question, how they might influence each other, and over what time periods these relationships occur. These theoretical relationships should guide your measurement decisions – not just what to measure, but when and how to measure. The choices you make at this stage are crucial for developing an analysis that can effectively test your hypotheses.

1.3.1 How Do You Choose a Sampling Window?

A **sampling window** is the time span over which data are collected – the period of time between the beginning and end of the time series. There are two time span questions to consider. First, what is the *ideal* sampling window for your analysis? The word ideal is italicized to highlight the theoretical nature of this proposition. In the absence of any practical limitations, what is the period of time that would best allow you to observe the relationship you are interested in analyzing? If you are interested in the dynamics of the Cold War relationship between the United States and the USSR, the ideal sampling window would span 1945 to 1991. If you are interested in the consequences of cyber warfare on interstate conflict, it does not make sense to include militarized interstate disputes (MIDs) from the nineteenth and twentieth centuries, even though the data collected as part of the correlates of war (COW) project include observations as far back as 1816 (Palmer et al., 2015). To study the effects of presidential elections on consumer sentiment and presidential approval, you need a time span that covers a number of elections.

https://doi.org/10.1017/9781108290326.001 Published online by Cambridge University Press

Second, what sampling window is *available* to you? Whereas the question about the ideal sampling window is theoretical, this question raises practical considerations. Data may be more limited than you'd like. Ideally, all the variables you need for your analysis – the outcome, your variable(s) of interest, and your controls – are available over the ideal sampling window. In many cases, you will not be so lucky.

The sampling window for your analysis must be the same for all of your time series variables – you need to have data for every time point for every series. Your window is only as wide as your shortest series. If a control variable starts after, or ends before, the rest of your data, you may want to find a suitable substitute rather than limit the sampling window for your analysis.

Missing observations within your time series must also be avoided. When a variable has missing values, statistical software typically uses listwise deletion, removing those time points from the analysis of all variables. This deletion occurs automatically in most software packages, sometimes without warning. If your data contain gaps, you have two main options: adjust your sampling window to exclude periods with missing data or employ methods to fill in missing observations.[1]

1.3.2 How Do You Choose a Sampling Interval?

You must next choose your sampling interval. The **sampling interval** is the discrete period of time used to measure the phenomenon of interest. Time is naturally continuous, but we measure it in discrete intervals: days, weeks, months, years, and so on. When constructing a data set, the interval cannot change over the sampling window.

The first step is identifying the "natural time unit" of your theory, as Freeman (1989, p. 93) puts it. What sampling interval would best capture the relationship you want to study? Consider how quickly your variables change and interact. For example, take changes in the partisanship of the United States Supreme Court. While the Court's decisions affect politics daily, its partisan composition changes only when justices are replaced. Creating a monthly time series of Court partisanship would produce long periods of no variation, making it difficult to detect relationships with other variables that do vary monthly.

The ideal sampling interval also avoids interregnums in stochastic variation. The lack of variation over consecutive intervals, like the problem highlighted in our Supreme Court example, has important consequences.

[1] Woodward, Gray, and Elliott (2017, 436) demonstrate Kalman filter methods that can be used for this purpose.

https://doi.org/10.1017/9781108290326.001 Published online by Cambridge University Press

The partial slope coefficients we use to make inferences in multiple regression analyses are covariances ($cov(y, x)$) divided by a variances ($var(x)$), net $cov(y, z)$ and $cov(x, z)$ for sets of controls (z). In a model used to analyze the partisanship of the Court, every period where the variable of interest (x) varies but court partisanship (y) remains constant hinders the detection of an effect; two variables cannot covary if one of them doesn't vary. For obvious reasons, choosing a monthly sampling interval might be impractical for such an analysis. At the other end of the spectrum, access to high-frequency social media data allows us to capture Twitter activity on a second-by-second basis. But if you are looking at the relationship between social media data and some political variable, say public opinion, that moves more slowly, it may not make sense to measure both phenomena on a second-by-second basis. When you are thinking about the ideal sampling interval, use what you think you know about the relationship of interest to identify the measurement strategy that would give you the best time series to capture the dynamic relationship.

Like the choice of your sampling window, practical considerations can force compromises between what is theoretically ideal and what is plausible for applied analysis. You now have to identify what sampling interval is *available* to you. Again, with limited time and resources, your sampling choices may be dictated by the work of others. For example, the Bureau of Labor Statistics measures unemployment monthly, but GDP quarterly. You might wish to analyze consumer sentiment since World War II, but the Institute for Social Research (ISR) at the University of Michigan only began its quarterly survey of consumer sentiment in 1960 and did not do so monthly until 1978. You will need to choose either a larger sampling window or a higher-resolution sampling interval (monthly instead of quarterly). Which choice covers the longest time period? Which choice provides the most observations? The answers to these questions are not always the same.

The sampling interval should be the same for all the variables included in the analysis. If you can measure your outcome and independent variable of interest at a quarterly level, but can only measure a control at an annual level, you have to make a choice. You may be tempted to resolve this problem by recording the annual observations of this control as though the variable were quarterly – by inserting the annual observations four times, once for each quarter. Or, you might be tempted to turn a quarterly measure into a monthly one by assigning the values you have to January, April, July, and October and then interpolating data in equal steps for the months in between. These attempts to artificially create observations at higher frequency have undesirable consequences. In the first case, you are artificially suppressing the variation in the data, and in the latter case you

https://doi.org/10.1017/9781108290326.001 Published online by Cambridge University Press

are imposing a dynamic structure. Both strategies will affect what you learn from the data.

Sometimes it makes sense to aggregate data from higher frequency intervals to lower frequencies. For example, you might average across monthly observations to generate quarterly observations. If the dynamics you are interested in studying are captured by quarterly data, this can be a good strategy. But if the outcome of interest is punctuated by major events and sudden changes in direction, removing variation through averaging may wash out the effects you are interested in observing. Freeman (1989, p. 93) advises researchers to think hard about the "natural time unit of our theories" and consider how one's conclusions might change if different decisions about sampling and aggregation were made. We concur.

1.3.3 How Many Observations Do You Need?

A common mistake people make in time series analysis is asking too much from too few observations. Because time series observations are not independent of each other, you need more of them. You may have thirty-five years worth of data but if your variables are all measured annually, that may not be enough information to build a dynamic model. Or, more to the point, not enough for a good dynamic model.

How long is long enough? Unfortunately, there is no simple answer – it will depend on the type of model and on the features of your data. You may be able to produce a reliable estimate of a coefficient that describes the relationship between y_t and y_{t-1} (call it $\hat{\alpha}_1$) with a sample of $T = 50$ if the population parameter is $\alpha_1 = .5$, but the reliability of the estimate would be lower if $\alpha_1 = .90$, and even more data could be required to conduct a reliable test that $\alpha_1 \neq 1$.

Further, the more variables in a model, the more time points you need. Keele, Linn, and Webb (2016*b*) reference Babyak (2004) on this point, noting that one should fit no more than one parameter for each ten observations when the data are independent and identically distributed. With time series data, the lack of independence necessitates more than ten. As α_1 approaches one, for example, substantially more observations are needed for each parameter estimated because each observation only contributes a little bit of new information.

There's a caveat, of course. While longer series are generally better, sometimes they aren't. Relationships can change, and you should choose the sampling window accordingly. A conventional time series regression model assumes that the parameters of interest are constant across the sampling window. As with all regression models, a partial slope coefficient gives the average conditional effect of the regressor on the regressand

https://doi.org/10.1017/9781108290326.001 Published online by Cambridge University Press

across all observations. The longer your sampling window, the more likely the marginal effects change over time. A failure to account for these parameter changes can produce biases. Such changes can, of course, be modeled, but they complicate analysis. Sometimes estimating relationships over a shorter time span, where they are expected to be constant, produces a more useful analysis.

1.3.4 How Should You Operationalize Your Variables?

A final consideration in assembling time series data is operationalization. **Operationalization** is the process of turning an abstract, theoretical concept into numeric values cataloged in a data array. This may or may not be straightforward.

An example will help demonstrate the problem. Let's say you have a theory about the factors that shape public perceptions of the economy. You decide to measure the outcome using the University of Michigan's Index of Consumer Sentiment (ICS). One concept that might influence consumer sentiment is objective economic performance. How do you measure that?

One possibility is gross domestic product (GDP). Is the raw measure of GDP appropriate to test hypotheses about the ICS? An initial concern should be that while the level of US GDP increases over time, you do not observe similar patterns in consumer sentiment. GDP is higher today than it was in the year 2010, higher in 2010 than in 2000, and so on. A plot of GDP reveals a trend – it increases over time. The ICS, on the other hand, tends to increase and decrease based on prevailing economic conditions, not the long-term trend. That is, people do not base their economic evaluations on the comparisons of GDP now to some distant point in the past. If they did, everyone would always have a very positive feeling about the economy.

Objective economic performance might, instead, be operationalized as the rate of growth:

$$\text{Growth}_t = \frac{\text{GDP}_t - \text{GDP}_{t-1}}{\text{GDP}_{t-1}}. \tag{1.1}$$

GDP growth is the difference between this year's GDP and last year's GDP scaled by last year's GDP. This is a more appropriate measurement – positive and negative changes in growth seem more likely to covary with economic perceptions.

There are still other issues to consider in this example. Is GDP really how people experience the economy? Coyle (2015) argues that GDP captures broad economic trends but does not necessarily translate into public

https://doi.org/10.1017/9781108290326.001 Published online by Cambridge University Press

feelings about the economy. MacKuen, Erikson, and Stimson (1992) use measures like unemployment and inflation as more direct influencers of public opinion. Still, questions remain. If you use inflation, which version? The annualized percentage change less food and energy is often used in macroeconomic models, but food and energy are, arguably, the prices most likely to influence attitudes about the economy. Annual changes in consumer prices also seem a strange choice. Do people really compare prices today to those twelve months ago when thinking about current economic conditions? Perhaps. Perhaps not. We won't resolve these questions here – the point is that careful thinking is required not only about which variables to use but also about how they should be operationalized to best capture the phenomena of interest.

As emphasized throughout this text, time series analysis involves a lot of uncertainty and theory is unlikely to completely resolve questions like these. Building the best model leverages what you know about the process that produced the data and also involves searching for empirical clues about the dynamic behavior of the data. In the chapters that follow, we describe the best practices that draw a balance between theory and data in the empirical specification of theoretical models.

1.3.5 Sampling Decisions for Presidential Approval

You can get a sense of some of these issues by examining what we think of as the "mother of all political time series," presidential approval in the United States. Figure 1.1 presents four graphs: Figure 1.1a and b show data for a sampling window going back to the beginning of Eisenhower's presidency but varying the sampling interval – quarterly data (a) and monthly data (b). The dashed vertical lines demarcate presidencies. Each point in the series represents the average of all the Gallup surveys in that month or quarter. The monthly series has more observations, but it covers less historical time. Also, because surveys were not faithfully administered back to President Eisenhower, the monthly series includes many interpolated time points.

The two time series are quite similar, but a closer look suggests that the quarterly series evolves more smoothly. The monthly series contains more bumps and wiggles, in part because of random sampling error that is not smoothed through aggregation to the quarterly level. The quarterly series is, thus, more persistent, a feature that has implications for the number of lags that might be needed to model its behavior. The choice of which level of aggregation to use depends on the particular research question and on which other variables will be modeled with the

https://doi.org/10.1017/9781108290326.001 Published online by Cambridge University Press

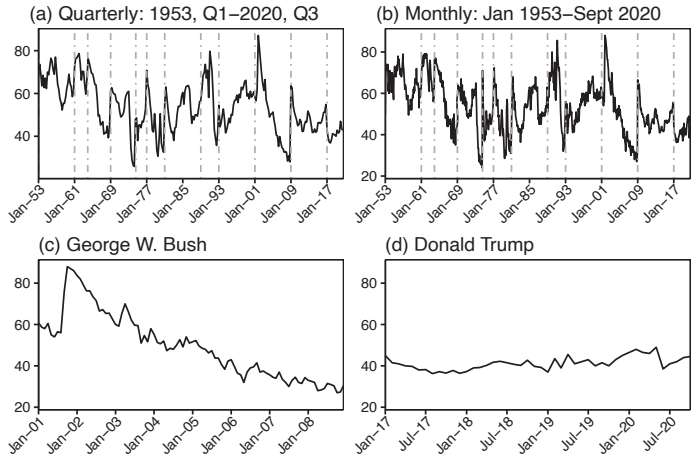


Figure 1.1 Gallup presidential approval ratings

approval series. For example, using the monthly series to increase the sample size would not make sense if other variables are only available quarterly (e.g., GDP) or yearly (e.g., budget deficits). Note, too, that the level of aggregation has implications for how you think about (and model) the effects of any explanatory variables. In particular, an effect that proves significant at one lag in monthly data might manifest as a contemporaneous effect with quarterly data.

The behavior of approval also differs across sampling windows. Figure 1.1c and d show monthly data separately for the Bush II (c) and Trump (d) administrations. Note the differences: George W. Bush's approval benefited from two large rally effects – one after the terrorist attacks of 9/11 and another with the onset of the War with Iraq. Both rallies persisted for some time but, eventually, the president's popularity followed what Mueller (1970) referred to as inexorable decline. Throughout this book, we will talk about *structural breaks*. It is noteworthy, here, that the 9/11 attacks – the biggest spike in the history of presidential approval – may be classified as either a temporary increase or a permanent shift, depending upon the sampling window used.

In contrast, President Trump's tenure is remarkable for its lack of volatility. Strong partisanship meant that neither record-low unemployment, impeachment proceedings, nor the pandemic moved his approval ratings out of the low forties for very long. Thus, hypotheses that might prove correct during one administration might not do so in others. What is true in one window of data may not hold in a wider, narrower, or different sampling window.

https://doi.org/10.1017/9781108290326.001 Published online by Cambridge University Press

1.4 HOW DO YOU USE THIS BOOK?

Once the data are together, you are ready to conduct the analyses we cover in this book. Figure 1.2 presents a map of the book, which we've broken into three parts. Part I outlines the preliminary questions that must be resolved before analysis. Are the data stationary or do they contain unit roots (Chapters 2 and 3)? Is weak exogeneity a valid assumption (Chapter 4)? The conclusions drawn from Part I dictate how to conduct analyses of the relationships between your time series. Readers who conclude that weak exogeneity is a reasonable assumption will find the single-equation methods covered in Part II appropriate for their analysis. We divide Part II of the text into regression models for stationary data (Chapter 5), unit root data (Chapter 6), and cases in which you are unsure of the nature of the univariate dynamics (Chapter 7). Which chapter is relevant will depend on your classification of the time series in Part I. If weak exogeneity is not a reasonable assumption, Part III covers the multiple-equation models you need to use for stationary data (Chapter 8) and unit root data (Chapter 9). The concluding chapter of the book (Chapter 10) offers guidance on best practices for applied time series analysis and highlights resources for further time series education.

There is often uncertainty that makes the road to follow unclear in a given analysis. You may find that efforts to characterize time series as stationary or unit root processes leave you unsure about the best path forward. You might similarly be unwilling to state, categorically, that the variables you care about are weakly exogenous, even if you

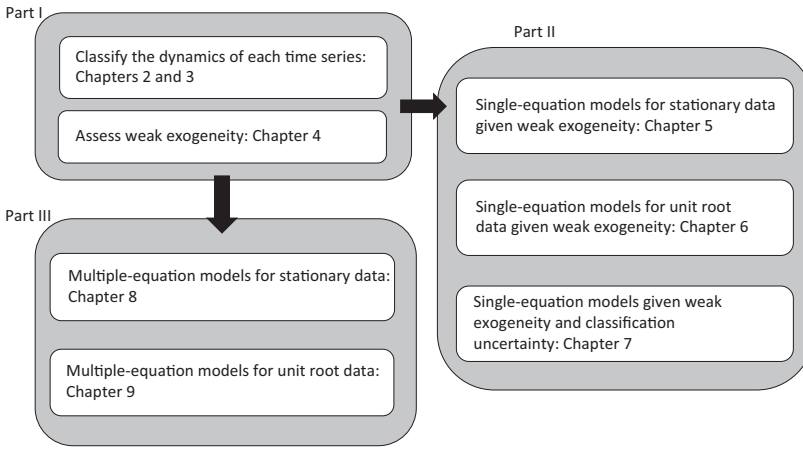


Figure 1.2 Map of the book

https://doi.org/10.1017/9781108290326.001 Published online by Cambridge University Press

believe it is plausible. Such is the nature of applied time series analysis. But the tools we provide will help you to navigate these dilemmas. Throughout the text, we use examples that highlight the problems you are likely to encounter. Data and code for the applications in this book can be found at https://github.com/webb767/practical-guide-to-time-series-replication. While most of our analyses can be conducted in Stata and R, other time series software is available. If you anticipate being a regular user of time series methodology, you might find purchasing a statistical package developed specifically for time series analysis a worthwhile investment; EViews, RATS, and CATS for RATS are packages that provide extensive tools for time series analysts that can make your work easier.

Time series analysis is an enormous field. Nevertheless, the contents of the following chapters should provide a reasonably comprehensive foundation upon which to learn more advanced techniques. If our collective experience has taught us anything, it is that learning new ways of approaching time series analysis, and statistical analyses in general, never ends. For our part, we believe the chapters that follow are a good place to begin.

https://doi.org/10.1017/9781108290326.001 Published online by Cambridge University Press

https://doi.org/10.1017/9781108290326.001 Published online by Cambridge University Press