STAT 0219: Time Series Analysis

Spring 2026

Instructor:	Christian Stratton	**Time:**	TR 9:45 – 11:00
Email:	cstratton@middlebury.edu	**Place:**	Warner 101
Office:	Warner 203	**Office hours:**	TBD Also by appointment

Course description: An introduction to statistical methods for time series analysis for students with a background in statistics. Topics include time series regression, auto-regressive models, moving average models, and ARIMA models, with an emphasis on estimation and forecasting with real data applications. Students will develop skills visualizing and summarizing serially correlated data structures and fitting time series models in various statistical software packages, including R and Julia.

Correspondence: My goal is to maximize my availability for help and discussion throughout the semester. Office hours will be determined via poll during the first week of class, but please feel free to contact me via email at anytime. Additionally, I am happy to meet outside of office hours by appointment.

Meeting format: Class will generally be used to learn new statistical concepts through a mixture of lecture and in-class activities. Most class periods will feature a short lecture introducing a new concept, followed by an in-class guided activity to be worked on in small groups. You will need to have access to a laptop during class. See more details below.

Learning objectives: Through this course, students will:

- Visualize and summarize time series data using statistical software
- Understand and implement stochastic models for forecasting
- Implement time series techniques in regression to account for temporal correlation

Textbook and materials: There is nothing that need be purchased for this class; all materials are free.

- The website for this course is on Middlebury Canvas. Please check Canvas often for assignments, deadlines, resources, and announcements.
- Students must have access to a laptop with the statistical computing language R, which can be downloaded for free at https://cran.rstudio.com/. Additionally, I recommend using RStudio as an integrated development environment (IDE) for interfacing with R. RStudio may be downloaded for free at https://posit.co/download/rstudio-desktop/.
  - Laptops with R/RStudio pre-installed are available to borrow from the Davis Family Library, which are a good option for those without access to a laptop or those experiencing short-term issues with your laptop. Please talk to me or the front desk of the Davis Library for more info.
- We will use two free online textbook: *Forecasting: Principles and Practice* by Rob J Hyndman and George Athanasopoulos and *Introductory Time Series with R* by Andrew V. Metcalfe and Paul S.P. Cowpertwait. These books may be downloaded at https://otexts.com/fpp3/ and https://link.springer.com/book/10.1007/978-0-387-88698-5, respectively.

Academic integrity: You are bound by Middlebury College's honor code, including its policies on plagiarism and cheating. Violation of these rules is ground for failure. To avoid charges of plagiarism, cite all the

sources used to complete your assignments/homework, including any peers with whom you collaborated. I encourage you to seek help in understanding the concepts and problems in your assignments from various sources, including peers, instructors, peer tutors, class notes, textbooks, and online sources.

Use of LLM and generative AI: Large language models (LLM) and generative AI, such as ChatGPT, are powerful tools enabled by statistics and data science techniques that may be used to enhance your learning of statistics and coding languages. As such, the use of large language models (LLM) and generative AI, such as ChatGPT, is permitted in this class and may be used on all assignments, unless explicitly prohibited by the assignment. However, **you may not copy responses verbatim from these tools, nor may you use these tools to generate complete responses or assignments.** Additionally, if content from generative AI is used on an assignment, **you must provide appropriate citation.** To clarify this policy, examples of acceptable and unacceptable prompts for ChatGPT are provided below.

Acceptable:

- Please provide an example of how to fit a Holt-Winters model in R.
- What is an ACF plot?
- The following code keeps giving me an error: ...

Unacceptable:

- Please analyze the attached data set.
- Answer the following question: *copy-paste from assignment*

Disclaimer: I am compelled to note that while generative AI can be a powerful tool, it is not infallible; we will explore some of the shortcomings of AI on the first day of class. It is possible that generative AI will provide you with incorrect information, and it is your responsibility to use generative AI critically. "ChatGPT said so," is not sufficient justification for an answer, and I am unlikely to be sympathetic to such comments on assignments.

One final note on AI use and grades. I reserve the right to ask you to explain any responses submitted in this class. Your grade may be reassessed based on the outcome of this conversation, and up to all points on the assignment are at stake. Generative AI should be a tool that facilitates your learning, in the same way your peers, professor, textbook, and internet assist your learning. Please do not use it to replace your learning!

Late policy: Consistent engagement with the course material is essential for your learning and academic growth. However, I understand that unforeseen circumstances may occasionally arise:

- When you become aware that you won't be able to make a deadline, please notify me and inform me of what day in the next week you anticipate completion of the assignment. You do not need to disclose why you are missing the deadline. So long as you communicate to me **before** the deadline, no late penalty will be applied. There may be some exceptions to this policy by assignment and category, but these exceptions will be communicated to you.
- **If you do not communicate with me before the deadline, late submissions will receive no credit.**

Course assessment: Your grade will be determined by daily activities, statistical reports, exams, and participation in an event called DataFest. Each category is loosely defined as follows:

10%	**Activities**	Most class days will conclude with an activity that demonstrates the day's concepts; the activity must be completed by the next class day and submitted on canvas. **Graded on effort**. The late policy does not apply to this assignment category; activities must be submitted before class.
40%	**Exams**	There will be two exams in this class: the midterm and the final. Both exams will consist of an in-class and take-home component. To be completed **on your own**.
45%	**Statistical reports**	Throughout the semester, you will be asked to write statistical reports. Full details will be provided as each report arises, but these reports must be **completed with two other students.**
5%	**DataFest participation**	You are required to attend and participate in an event hosted by the Department of Mathematics and Statistics called DataFest. More details will be provided throughout the semester, but be aware that the event spans the weekend of Friday, April 17^{th} to Sunday, April 19^{th}. You must be present for the opening ceremony on Friday afternoon and the closing ceremony on Sunday to receive full credit.

Letter grades: Letter grades will be assigned according to the following scale. Note that I may adjust thresholds at the end of the semester, but they will only ever be adjusted *down*.

F	**D**	**C-**	**C**	**C+**	**B-**	**B**	**B+**	**A-**	**A**
$[0, 60)$	$[60, 70)$	$[70, 74)$	$[74, 77)$	$[77, 80)$	$[80, 84)$	$[84, 87)$	$[87, 90)$	$[90, 94)$	$[94, 100)$

Accommodations for disabilities: Students who have Letters of Accommodation in this class are encouraged to contact me early in the semester to ensure that such accommodations are implemented in a timely fashion. For those without Letters of Accommodation, assistance is available to eligible students through the Disability Resource Center (DRC).

Disclaimer: This syllabus is a dynamic document and may change throughout the semester as a result of our conversations and decisions as a class.

Tuesday	Thursday
Feb 10th **1** Generative AI lecture and time series data	12th **2** FPP Ch. 2 - Visualizing time series
17th **3** FPP Ch. 3 - Decomposition of time series	19th **4** FPP Ch. 3 - Decomposition of time series STL decomposition
24th **5** ITSR Ch. 2 - Correlation and Stationarity	26th **6** FPP Ch. 8 - Forecasting and Smoothing Simple exponential smoothing and Holt's linear method
Mar 3rd **7** FPP Ch. 8 - Forecasting and Smoothing Holt-Winters	5th **8** FPP Ch. 8 - Forecasting and Smoothing Innovations state space models
10th **9** FPP Ch. 8 - Forecasting and Smoothing Innovations state space models	12th **10** FPP Ch. 9 - Autoregressive models and characteristic polynomials
17th **11** FPP Ch. 9 - Moving Averages and ARMA models	19th **12** FPP Ch. 9 - Non-seasonal ARIMA Models
24th Spring break	26th Spring break
31st **13** FPP Ch. 9 - Non-seasonal ARIMA Models	Apr 2nd **14** FPP Ch. 9 - Seasonal ARIMA Models
7th **15** FPP Ch. 9 - Seasonal ARIMA Models	9th **16** FPP Ch. 10 - Time series regression Deterministic trends
14th **17** FPP Ch. 10 - Time series regression GLS and Residual series	16th **18** FPP Ch. 10 - Time series regression Harmonic regression
21st **19** FPP Ch. 10 - Time series regression Regression with ARIMA errors	23rd **20** FPP Ch. 10 - Time series regression Case-study: comparing stochastic and deterministic trends
28th **21** FPP Ch. 10 - Time series regression Longitudinal models	30th **22** FPP Ch. 10 - Time series regression Longitudinal models cont.
May 5th **23** Flex day/special topics	7th **24** Flex day/special topics