

CHAPTER OVERVIEW

14: Comparing Several Means (One-way ANOVA)

This chapter introduces one of the most widely used tools in statistics, known as “the analysis of variance”, which is usually referred to as ANOVA. The basic technique was developed by Sir Ronald Fisher in the early 20th century, and it is to him that we owe the rather unfortunate terminology. The term ANOVA is a little misleading, in two respects. Firstly, although the name of the technique refers to variances, ANOVA is concerned with investigating differences in means. Secondly, there are several different things out there that are all referred to as ANOVAs, some of which have only a very tenuous connection to one another. Later on in the book we’ll encounter a range of different ANOVA methods that apply in quite different situations, but for the purposes of this chapter we’ll only consider the simplest form of ANOVA, in which we have several different groups of observations, and we’re interested in finding out whether those groups differ in terms of some outcome variable of interest. This is the question that is addressed by a **one-way ANOVA**.

The structure of this chapter is as follows: In Section 14.1 I’ll introduce a fictitious data set that we’ll use as a running example throughout the chapter. After introducing the data, I’ll describe the mechanics of how a one-way ANOVA actually works (Section 14.2) and then focus on how you can run one in R (Section 14.3). These two sections are the core of the chapter. The remainder of the chapter discusses a range of important topics that inevitably arise when running an ANOVA, namely how to calculate effect sizes (Section 14.4), post hoc tests and corrections for multiple comparisons (Section 14.5) and the assumptions that ANOVA relies upon (Section 14.6). We’ll also talk about how to check those assumptions and some of the things you can do if the assumptions are violated (Sections 14.7 to 14.10). At the end of the chapter we’ll talk a little about the relationship between ANOVA and other statistical tools (Section 14.11).

- [14.1: Summary](#)
- [14.2: An Illustrative Data Set](#)
- [14.3: How ANOVA Works](#)
- [14.4: Running an ANOVA in R](#)
- [14.5: Effect Size](#)
- [14.6: Multiple Comparisons and Post Hoc Tests](#)
- [14.7: Assumptions of One-way ANOVA](#)
- [14.8: Checking the Homogeneity of Variance Assumption](#)
- [14.9: Removing the Homogeneity of Variance Assumption](#)
- [14.10: Checking the Normality Assumption](#)
- [14.11: Removing the Normality Assumption](#)
- [14.12: On the Relationship Between ANOVA and the Student t Test](#)

This page titled [14: Comparing Several Means \(One-way ANOVA\)](#) is shared under a [CC BY-SA 4.0](#) license and was authored, remixed, and/or curated by [Danielle Navarro](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.