

Chapter 9

Betting on Uncertain Demand: The Newsvendor Model¹

Matching supply and demand is particularly challenging when supply must be chosen before observing demand and demand is stochastic (uncertain). To illustrate this point, suppose you are the owner of a simple business: selling newspapers. Each morning you purchase a stack of papers with the intention of selling them at your newsstand at the corner of a busy street. Even though you have some idea regarding how many newspapers you can sell on any given day, you never can predict demand for sure. Some days you sell all of your papers, while other days end with unsold newspapers to be recycled. As the newsvendor you must decide how many papers to buy at the start of each day. Because you must decide how many newspapers to buy before demand occurs, unless you are very lucky, you will not be able to match supply to demand. A decision tool is needed to make the best out of this difficult situation. The *newsvendor model* is such a tool.

You will be happy to learn that the newsvendor model applies in many more settings than just the newsstand business. The essential issue is that you must take a firm bet (how much inventory to order) before some random event occurs (demand) and then you learn that you either bet too much (demand was less than your order) or you bet too little (demand exceeded your order). This trade-off between “doing too much” and “doing too little” occurs in other settings. Consider a technology product with a long lead time to source components and only a short life before better technology becomes available. Purchase too many components and you risk having to sell off obsolete technology. Purchase too few and you may forgo sizable profits. A notable example of the latter occurred to IBM in the early 1990s. IBM’s laptops were more popular than they initially imagined. But due to shortages in components, IBM estimated that they were unable to satisfy \$100 million in demand, a sizeable amount even for IBM.

This chapter begins with a description of the production challenge faced by O’Neill Inc., a sports apparel manufacturer. O’Neill’s decision also closely resembles the newsvendor’s task. We then describe the newsvendor model in detail and apply it to O’Neill’s problem. We also show how to use the newsvendor model to forecast a number of performance measures relevant to O’Neill.

¹ Data in this chapter have been disguised to protect confidential information.