Business Operations Analytics

Business Operations Analytics

MELISSA CHRISTENSEN



Business Operations Analytics Copyright © by Melissa Christensen is licensed under a Creative Commons Attribution 4.0 International License, except where otherwise noted.

Preface

This text was written for students at Plymouth State University enrolled in MGM 3190: Business Operations Analytics with the intention of providing students an alternative to commercial textbooks while ensuring the content is relevant to their work in the course. It was developed to provide a course-specific resource that connects operations concepts with applied analytics using Excel, Power BI, and R.

Although designed with MGM 3190 in mind, the material extends beyond a single class or institution. The emphasis on applied learning, case studies, and tool integration makes it a resource that can be adapted for other courses in operations and analytics.

I would like to thank Evan J. Christensen for his contributions to several of the graphics that illustrate and clarify the concepts presented in this text. I am also grateful to Plymouth State University and the New England Board of Higher Education for their support in the development of this OER.

This resource is designed to grow over time. Students in MGM 3190 will have opportunities to shape it by contributing new cases, examples, or refinements for future students.

Licensed under CC By 4.0

Case Study Overview – Pemi Coffee Roasters



Pemi Coffee Roasters is the running case study for this course. It is a fictional, mid-sized specialty coffee company designed to illustrate real-world operational challenges. While it does not exist in reality, it models the complexities of small but growing companies with local personality and operational interdependencies.

Students will use Pemi Coffee Roasters to examine supply chains, workforce management, technology use, and business operations analytics. The scenarios are realistic, continuous, and applicable to both manufacturing and service contexts.

Company Background

- **Location:** Plymouth, New Hampshire, a town shaped by mountains, lakes, and a college campus. Seasonal tourism

and a vibrant arts scene influence local demand.

- **Name Inspiration:** The Pemigewasset River flows through Plymouth, representing steady flow and seasonal change. The company's branding and marketing reflect this connection to place.
- **Origins:** Founded by two college friends who started with a small café, a compact roaster, a few wholesale accounts, and a loyal local customer base.
- **Growth:** Expanded to two café locations, a wholesale roasting arm, and an online store shipping across New England and beyond. Marketing includes seasonal blends, local imagery, and community-focused messaging.

Operational Overview

Pemi Coffee Roasters' operations encompass both manufacturing (coffee roasting, packaging, inventory management) and service (cafés, customer experience, online order fulfillment). Key operational considerations include:

- **Supply Chain Management:** Sourcing coffee beans globally, accounting for harvest cycles, quality grades, and ethical sourcing requirements.
- **Production Scheduling:** Balancing freshness, efficiency, and cost management.
- **Inventory Management:** Timing, turnover, and waste prevention.
- **Human Resources:** Training baristas and wholesale account managers, maintaining brand standards in customer interactions.
- **Logistics:** Coordination of roasting, packaging, shipping, and café operations.

Challenges reflect realistic business pressures: seasonal demand fluctuations, global commodity price swings, supply chain delays, and competitive pressures from local and national players.

Learning Applications

Using Pemi Coffee Roasters, students will:

1. **Apply analytics to operations:** Examine production schedules, inventory turnover, and staffing efficiency.
2. **Explore both manufacturing and service operations:** Analyze cost-effective delivery routes and barista scheduling during peak hours.
3. **Develop a course portfolio:** Build analyses, process maps, and operational improvement plans throughout the course.
4. **Understand real-world challenges:** Consider how shifts in demand, supply chain disruptions, and customer preferences create operational opportunities and constraints.

Role in the Course

- Pemi Coffee Roasters will be referenced throughout the textbook via example boxes and applied exercises, allowing students to focus on operations concepts without reintroducing the company's backstory each time.

- The case study integrates business operations analytics, showing how data-driven decisions support efficiency, quality, and customer satisfaction.
- By following Pemi Coffee Roasters, you will gain practical experience applicable to real organizations.

SECTION I

FOUNDATIONS OF BUSINESS OPERATIONS & ANALYTICS

1. Introduction to Business Operations and Analytics



Learning Objectives

- Define business operations and explain its role in transforming strategy into actionable results.
- Differentiate between manufacturing and

service operations, recognizing the unique challenges of each context.

- Identify the scope of operations within an organization, including process flows, quality control, technology, and compliance.
- Explain how operations contribute to efficiency, quality, customer satisfaction, and strategic alignment.
- Distinguish between business analytics and business operations analytics, and describe how each supports organizational decision-making.
- Apply abstract operations concepts to a real-world scenario using example case studies, such as Pemi Coffee Roasters, to analyze workflow, scheduling, and resource management.

1.1 What Is Business Operations

Business operations is the everyday reality of how an organization gets work done. It's not a slogan or a vision statement—it's the people, the tools, the schedules, the problem-solving that actually make things happen. You can write the smartest strategy in the world, but without operations, it stays stuck on paper.

It looks different in different settings. A clothing manufacturer might be chasing down a delayed shipment of fabric while figuring out how to keep the sewing line busy. A healthcare clinic might be adjusting patient appointments on the fly when two nurses call out. The details change, but the

purpose doesn't: get the right things to the right people at the right time, without breaking the budget.

Example: Operational Juggling at Pemi Coffee Roasters

Pemi Coffee Roasters balances daily café service with wholesale orders and online shipments. A late green-bean delivery affects roasting schedules, which in turn can delay café supply and online orders. Staff must adjust roasting, packaging, and delivery plans quickly to keep everything on track, showing how operations decisions ripple across the organization.

The work isn't just about doing tasks. It's about deciding how those tasks should be done in the first place. Who's the best person for this step? Which system will keep the process moving? Where are we wasting time or resources without realizing it? People who manage operations tend to have a knack for seeing the whole picture—like a chess player who's already thinking three moves ahead—but they also know when to improvise.

And the ground is always shifting. Supply chains tangle. Software updates roll out whether you're ready or not. Customer expectations creep higher each year. A manufacturer might bring in automation to keep pace; a service provider might adjust its offerings to focus on what it does best. Surprises happen, and operations is often where the scramble to adapt plays out.

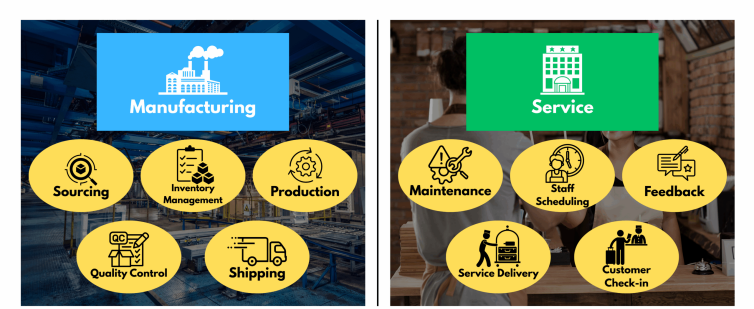
Over time, patterns emerge. The smart operators pay attention to those patterns and look for ways to make things

smoother. Sometimes the fix is small—rearranging a workspace so people stop tripping over each other. Sometimes it's a bigger change, like overhauling the way orders are processed. Bit by bit, those adjustments determine whether the organization runs smoothly… or just runs in circles.

Example: Workflow Adjustments in Action

At Pemi, staff noticed bottlenecks during the morning rush when both café and wholesale roasting tasks coincided. By shifting packaging to the afternoon and reallocating baristas temporarily, the company smoothed workflow and reduced wait times, showing how observing patterns leads to practical operational solutions.

Whether you end up in a warehouse, behind a desk, or managing a team in a busy hotel, knowing how operations work changes how you see things. You start to notice the chain of events that takes an idea, a box of materials, or a set of instructions and turns it into something useful. You also notice how small decisions—some you might barely remember making—can shift costs, slow things down, or make customers happier. In this course, we'll get into the methods and tools that keep operations on track, but we'll also talk about the messier side: the judgment calls, trade-offs, and quick thinking that show up when the real world doesn't go by the book.



Operations vary by context: manufacturing emphasizes production flow, while services focus on customer interaction. Both involve coordination, timing, and problem-solving.

1.2 The Scope of Business Operations

Business operations is a wide-reaching field that connects more parts of an organization than most people think about. It's not a single room full of people doing the same task—it's the thread that runs through almost everything a company touches. Making something or delivering a service is only the surface layer; beneath that is a network of decisions, handoffs, and fixes that keep it all from stalling.

In manufacturing, you can see the physical flow. A bicycle maker might place orders for components months in advance, track them across oceans, assemble them on the factory floor, and send finished bikes to stores before the season changes. If a shipment of pedals shows up late or a step in assembly gets held up, the whole rhythm shifts. Operations isn't just "make the bike." It's making sure every part of that chain stays in sync.

Services can be harder to picture, but the same rules apply. A hotel on a busy weekend is a good example—housekeeping turning over rooms, maintenance trying to stay ahead of small repairs, the front desk managing early arrivals while the lobby

fills. One delay in cleaning can back up check-ins, which means the front desk staff starts the day explaining instead of welcoming. That's operations too.

Example: Service Flow at Pemi

Pemi cafés must manage a mix of regular customers and tourist traffic. During peak leaf-peeping season, baristas and front-of-house staff adjust schedules, streamline drink prep, and monitor inventory in real-time. These small operational decisions ensure smooth service, illustrating that operations principles apply equally in service contexts.

It's not all front-line work either. Quality checks, safety protocols, compliance paperwork, environmental commitments—they live inside operations as well. When a company decides to cut waste, switch to a greener supplier, or change packaging, those choices land here first. They can improve efficiency, but they also change how the business is seen by customers and partners.

Technology changes the game further. Small companies and large enterprises alike use software to track inventory, plan production runs, manage shipments, and predict maintenance needs. The tools are different, but the goal is the same: keep the work moving and the surprises to a minimum.

And nothing stays still for long. A change in the market, a shift in customer preferences, or something unexpected—weather, a supplier going under, even a news event—can throw things off quickly. The organizations that

handle it best treat operations as something alive. They check in on it, tweak it, and sometimes overhaul it before they have to. That habit isn't just about surviving problems—it's often how they get ahead.



From sourcing to feedback, each stage affects efficiency, quality, and customer satisfaction. Operations analytics helps optimize this flow.

1.3 The Role of Operations in Organizational Success

Operations is more than a support function; it is a central driver of organizational success. An organization can have a brilliant marketing strategy and an innovative product, but without strong operations, it will struggle to deliver on promises and maintain customer trust.



Operations directly influence efficiency, quality, customer satisfaction, and strategic alignment. Improvements in one area often reinforce others, creating a cycle of organizational success.

One of the most significant contributions operations makes is to **efficiency**. Streamlined processes reduce waste, save time, and minimize costs. This efficiency can translate into competitive pricing, higher profit margins, or reinvestment into areas like research and development.

Example: Efficiency Gains at Pemi

By coordinating roasting schedules with café demand and online orders, Pemi ensures beans are roasted just-in-time. This reduces storage costs, keeps coffee fresh, and minimizes labor spent managing inventory. Small operational efficiencies like this can have a large impact on overall performance.

Operations also play a critical role in **quality**. Consistency is a hallmark of successful organizations. Customers expect the same level of quality each time they interact with a company, whether they are ordering a product online, visiting a store, or booking a service. In manufacturing, quality assurance processes verify that products meet specifications before they leave the factory. In services, training programs and standard procedures help ensure that customers receive the same level of care and attention at every location or touchpoint.

Customer satisfaction is another area where operations has a direct impact. Reliable delivery times, responsive service, and a seamless experience contribute to customer loyalty. When operations falter, customers notice. A late delivery, a defective product, or a poorly executed service can quickly erode trust and damage the brand's reputation.

Strategic alignment is an often-overlooked aspect of operations. The best-run operations are those that support and reinforce the organization's overall strategy. If a company differentiates itself based on rapid delivery, its operations must be designed for speed, from supplier relationships to warehouse layout to last-mile logistics. If the company emphasizes sustainability, operations must focus on sourcing environmentally responsible materials, reducing energy use, and minimizing waste.

In addition, operations are a source of valuable data. By

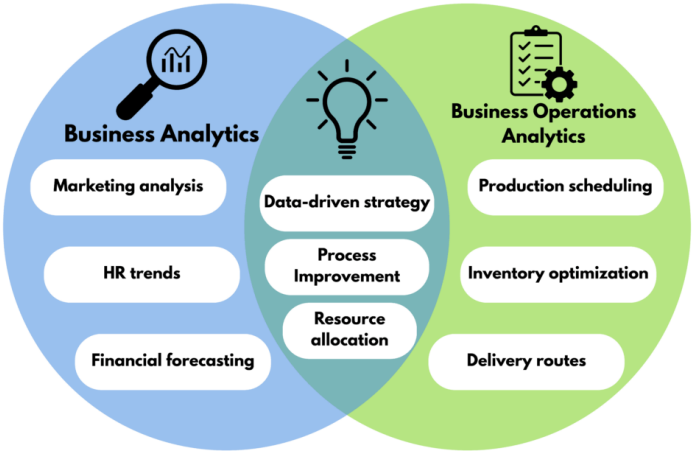
tracking production times, defect rates, inventory levels, and customer feedback, organizations can identify trends and make informed decisions. This is where business operations analytics becomes essential, turning raw data into actionable insights that guide process improvements and strategic planning.

Ultimately, operations is not just about keeping the business running; it is about enabling the business to thrive. When operations are aligned with strategy, managed efficiently, and **continuously improved**, they become a powerful engine for growth and competitive advantage.

Business operations and business operations analytics are not just academic fields. They represent large and growing areas of employment. According to the U.S. Bureau of Labor Statistics, management and operations related roles remain among the strongest employment sectors, with steady demand projected (U.S. Bureau of Labor Statistics, 2024). This outlook reinforces why learning the fundamentals of operations and operations analytics is valuable for managers preparing to enter the workforce.

1.4 The Difference Between Business Analytics and Business Operations Analytics

The terms business analytics and business operations analytics are closely related but not interchangeable. Understanding the distinction between them is important because it shapes how we approach problem-solving in an organization and what kinds of decisions we aim to improve.



Business analytics examines performance across all areas of a company. Business operations analytics zooms in on operational efficiency and process improvement. Both use data to make better decisions, and their insights often overlap.

Business analytics is the broader discipline. It encompasses the collection, analysis, and interpretation of data to support decision-making across all areas of an organization. This might include marketing analytics to understand customer behavior, financial analytics to assess investment performance, or human resources analytics to evaluate employee retention trends. The focus is on using data to guide strategy and improve outcomes in any business function.

Business operations analytics is a specialized subset of business analytics. It applies the principles of analytics specifically to the processes and systems that produce and deliver an organization's products or services. The emphasis is on operational efficiency, process improvement, and the performance of the **value chain**. Instead of asking, "How can

we improve our overall business performance?" operations analytics might ask, "How can we reduce production cycle time by 10 percent without sacrificing quality?" or "Which delivery routes minimize both time and cost?"

Example: Operations Analytics at Pemi

Pemi uses sales and inventory data to adjust roasting schedules and optimize delivery routes. By analyzing historical demand patterns, the company reduces stockouts, avoids overproduction, and cuts shipping costs. This demonstrates how operations analytics translates data into concrete, actionable improvements.

The distinction can be illustrated by looking at the same organization from two perspectives. Consider a large hotel chain. A business analytics team might analyze guest demographics and booking trends to determine where to open new locations. Meanwhile, the operations analytics team would focus on optimizing housekeeping schedules, improving check-in processes, or managing inventory levels for linens and supplies to keep operations running smoothly.

The relationship between the two is complementary. Business analytics provides a wide-angle view of the organization's performance and strategic opportunities, while business operations analytics zooms in on the day-to-day functions that make strategy possible. Strong organizations recognize the value of both perspectives and ensure they are integrated.

In this course, we will explore business operations analytics in depth, using analytical methods and tools to address challenges in production, service delivery, supply chain management, and process optimization. We will also draw connections to broader business analytics concepts so you can see how operational insights fit into the bigger strategic picture.

Key Takeaways

- Operations turn strategy into action.
- Operations are interconnected and dynamic across manufacturing and services.
- Efficiency, quality, and customer satisfaction are core outcomes of strong operations.
- Data and analytics make operations decisions faster and better.
- Theory + practice: case examples (e.g., Pemi Coffee) translate concepts into action.
- Continuous improvement keeps operations aligned with goals as conditions change

Chapter One References

Davenport, T. H., & Harris, J. G. (2017). *Competing on analytics: The new science of winning*. Harvard Business Review Press.

Heizer, J., Render, B., & Munson, C. (2020). *Operations management: Sustainability and supply chain management* (13th ed.). Pearson.

OpenStax. (2023). *Introduction to business*. OpenStax, Rice University. https://openstax.org/details/books/introduction-business

Slack, N., Brandon-Jones, A., & Burgess, N. (2022). *Operations management* (10th ed.). Pearson.

U.S. Bureau of Labor Statistics. (2024). *Business employment dynamics*. https://www.bls.gov/bdm/

SECTION II

OPERATIONS STRATEGY & ANALYTICS IN EXCEL

2. Operations Strategy and Data Quality



Learning Objectives

- **Explain** the role of operations strategy in linking long-term vision to daily decision-making.
- **Identify** and **describe** the four competing priorities of operations strategy (cost, quality, speed, flexibility).

- **Discuss** the relationship between data quality and operational success.
- **Recognize** the five dimensions of data quality and their operational implications.
- **Describe** common causes of poor data quality and strategies to address them.
- **Apply** a step-by-step data cleaning process to real-world datasets.
- **Integrate** data quality management into an organization's continuous improvement strategy.
- **Compare** operational challenges between goods-producing and service-oriented businesses.

2.1 Understanding Operations Strategy

Operations strategy connects a company's long-term vision to the everyday decisions that make that vision a reality. It is not just a theoretical framework; it shapes how resources are allocated, how processes are designed, and how performance is evaluated. When well executed, it aligns every operational choice with the organization's competitive priorities—whether that priority is to be the fastest, the most cost-efficient, the highest quality, or the most flexible provider in the market.

Example: Retail vs. Wholesale Strategy

In Pemi's cafés, operations strategy shapes menu variety, pricing, service speed, and the customer experience. In its wholesale roasting arm, it drives decisions on bean sourcing, production scheduling, inventory management, and distribution partnerships.

While manufacturing firms produce physical goods, service companies also rely heavily on operations strategy. In a consulting firm, for example, the "production" is the delivery of expertise. Timelines, staff utilization, and project workflows take the place of assembly lines and production schedules. In both goods and services, operations strategy determines how the organization responds to fluctuations in demand, how it ensures quality, and how it manages costs.

An effective strategy balances four dimensions: cost, quality, speed, and flexibility. These dimensions are interdependent. A focus on cost efficiency may limit flexibility. Prioritizing speed may require extra labor or technology investments. The trade-offs a company chooses must match its competitive positioning.

Example: Strategic Trade-offs

Pemi cannot compete on price with mass-market coffee brands, so it prioritizes quality and brand experience while keeping operations lean enough to remain profitable.

2.2 Linking Operations Strategy to Data Quality

Even the most well-crafted operations strategy depends on reliable, timely, and accurate information. Without quality data, managers cannot make informed decisions about scheduling, inventory, staffing, or resource allocation. Poor data quality can turn a sound strategy into an operational liability.

Example: Data Accuracy Impact

In Pemi's wholesale operation, misrecorded orders or incorrect delivery addresses can delay shipments and damage customer relationships. In cafés, inaccurate sales data may cause stockouts and lost revenue during peak demand.

In a service industry example, consider a marketing agency tracking campaign performance. If the analytics tool misclassifies website visits, the agency might misinterpret the

effectiveness of different channels and shift its budget in the wrong direction.

The connection is simple: strategy guides the "what" and "why" of operations, but data quality drives the "how" and "how well."

2.3 Dimensions of Data Quality

There are five basic dimensions of data quality: accuracy, completeness, consistency, timeliness, and relevance. Each one matters for daily operations.

Accuracy ensures that the data reflects reality.

> If Pemi's point-of-sale (POS) system mis-records a cappuccino as a latte, sales reports misrepresent customer preferences, potentially skewing menu planning and inventory purchasing.

Completeness means that all necessary fields are filled in and no critical information is missing.

> A wholesale order without a delivery date or customer contact information creates bottlenecks and extra administrative work.

Consistency ensures that data follows uniform formats and naming conventions.

Variations in recording state names (“NH” vs. “N.H.”) can create mismatches across systems.

Timeliness ensures data is available when needed. Outdated information can lead to misaligned staffing schedules or production runs.

Relevance keeps the focus on collecting only what is useful for the decision at hand. Storing unnecessary details adds clutter and distracts from key metrics.

2.4 Common Causes of Poor Data Quality

Poor data quality can come from human mistakes, flawed processes, or technical limitations. Manual entry errors are among the most common issues. At Pemi’s café, a barista might accidentally record a large coffee as a medium, throwing off inventory counts. In a consulting firm, employees might misclassify billable hours, leading to incorrect invoices.

System issues can also introduce errors. If a customer order management platform does not integrate with the inventory system, discrepancies can arise when orders are fulfilled. Outdated technology may lack validation checks, making it easier for mistakes to slip through.

Process flaws often go unnoticed until they cause recurring problems. If a business collects information on paper forms before transcribing it into a digital system, every transcription step becomes a potential error point.

Culture plays a role as well. When employees do not understand why data quality matters, they may view it as a

low-priority administrative task rather than as part of delivering good service.

2.5 The Role of Data Cleaning in Operations

Data cleaning is the deliberate process of finding and fixing errors, filling in gaps, and standardizing information so that it is usable for decision-making. It is not just a technical chore—it is an operational safeguard. Without clean data, forecasts become unreliable, schedules inefficient, and quality control ineffective.

In Pemi's wholesale operation, data cleaning takes place before production schedules are finalized. Orders are checked for missing quantities, incorrect addresses, or unverified delivery dates. Catching and correcting errors before the beans are roasted prevents costly waste.

In service industries, the stakes are similar. A law firm that sends invoices with incorrect billing hours risks client disputes and payment delays. Regular data cleaning ensures the accuracy of records and preserves professional credibility.

2.6 A Detailed Walkthrough of the Data Cleaning Process

The cleaning process is sequential, but in practice, steps may loop back as new issues are found. Here is how a typical cycle might look when preparing Pemi's wholesale order data for a weekly roasting plan.

1. Import the dataset into a tool that supports sorting,

filtering, and validation—often Excel, Google Sheets, or a database platform.

2. Scan for duplicate entries. At Pemi, a duplicate might occur if a wholesale customer calls and emails the same order, and both are entered separately. Duplicates can lead to double production and unnecessary shipping.
3. Identify structural errors. This includes inconsistent date formats, misaligned columns, or misspelled product names. A typo turning “Dark Roast” into “Dak Roast” could prevent the order from being matched with inventory records.
4. Address missing values. This may mean reaching out to the customer for confirmation, using historical averages to fill gaps, or removing incomplete entries if they cannot be verified.
5. Cross-verify against a reliable source. For addresses, this could be a postal database; for order quantities, the original order confirmation. This step ensures accuracy before the data is acted upon.
6. Standardize formats. All product names, abbreviations, and customer identifiers should follow the same conventions. Consistency helps prevent future mismatches between systems.
7. Document the changes made. Keeping a log of adjustments builds an audit trail and supports training efforts, showing employees where errors occur most often and how they were resolved.

This same process applies in service industries. For example, a digital marketing agency cleaning campaign analytics would import raw reports, remove duplicates from multiple tracking platforms, ensure consistent date and campaign naming conventions, address missing conversion data, verify numbers against client systems, and record all adjustments before analysis.

2.7 Embedding Data Quality into Operations Strategy

For data quality to last, it has to be part of daily operations. It cannot be treated as a one-time project. Businesses need to build quality checks into routine workflows, train employees, and review results on a regular basis. This makes data quality part of the culture rather than an occasional fix.

At Pemi Coffee Roasters, data entry rules are included in new employee training. Baristas, sales staff, and managers are shown how mistakes at their step affect production, delivery, and customer satisfaction. A wrong entry may seem small, but it can lead to missing inventory, late orders, or unhappy customers.

In a service business, project managers may review data at specific checkpoints. This makes sure the information passed to the next stage is complete and accurate. A missing client approval or an incomplete project log can cause delays later in the process.

Monitoring is also important. Leaders can track error rates, correction times, and the costs linked to mistakes. These numbers show whether the business is making progress or if problems continue. Regular reviews help find repeating issues, adjust training, and improve processes over time.

2.8 Continuous Improvement and Data Quality

Data quality is not a one-time achievement. It requires constant attention. Many organizations use continuous improvement methods such as PDCA (Plan–Do–Check–Act) or Six Sigma to keep data quality connected to their overall goals for operational excellence.

At Pemi Coffee Roasters, this might include quarterly audits of both café and wholesale systems. These audits can show patterns in errors and point to areas where training or processes need to change. For example, if delivery address mistakes become common, the order form could be redesigned to include address checks.

In service businesses, improvement often means making data collection easier and reducing manual entry. A consulting firm, for instance, could switch to one integrated system for time tracking and billing. This step reduces the chance of differences between separate platforms and helps keep records consistent.

2.9 Case Integration: Service Operations and Pemi Coffee Roasters

Although Pemi produces a tangible product, many of its operational challenges mirror those of service organizations. Both must align strategy with quality data to manage resources effectively.

Consider a busy Saturday at Pemi's café alongside a high-demand day for a spa. In both cases, reservations (or orders) need to be accurate, inventory (whether coffee beans or massage supplies) must be sufficient, and staffing must be aligned with customer flow. Data quality in booking systems ensures smooth execution.

By studying operations in both sectors, students can see how the same principles apply whether the "product" is a latte or a legal consultation. The skills learned here—especially around cleaning and validating data—transfer directly between goods and service industries.

2.10 Summary and Assignment

Operations strategy sets the framework for how a business delivers value, and data quality makes sure that framework rests on solid ground. If the data is not accurate, complete, consistent, timely, and relevant, even the strongest strategies are likely to fall apart during execution.

Pemi Coffee Roasters shows how these concepts come together in practice. From café sales to wholesale roasting, decisions rely on clean data. The same applies in service settings, where customer satisfaction and efficiency hinge on reliable information.

Assignment: For this week's applied activity, you will receive a dataset with intentional errors. Your task is to clean the data using the step-by-step process outlined in this chapter. Document each change, explain why it was necessary, and ensure that the final dataset meets all five dimensions of data quality.

Key Takeaways

- Ops strategy links vision to daily trade-offs (cost, quality, speed, flexibility).
- Good data makes strategy executable; focus on accuracy, completeness, consistency, timeliness, relevance.
- Most data issues come from people/process/systems—use validation, integration, and training.

- Clean with a repeatable workflow (dedupe → fix structure/missing → verify → standardize → document).
- Make data quality routine (checkpoints + PDCA) across goods and services.

Chapter Two References

CARL Open Education Working Group. (n.d.). *Data cleaning*. In *OER data collection toolkit*. Pressbooks. https://pressbooks.openedmb.ca/oerdata/chapter/data-cleaning/ pressbooks.openedmb.ca

OpenStax. (2018, September 19). *Production and operations management*. In *Introduction to business*. https://openstax.org/books/introduction-business/pages/10-introduction

3. Process Design and Improvement



3.1 Understanding Process Design

Process design is the foundation of operational efficiency. It involves creating a structured sequence of activities that transforms inputs into outputs in a way that maximizes value while minimizing waste. A well-designed process provides clarity for employees, sets expectations for quality, and establishes a consistent rhythm for production or service delivery. Whether in manufacturing or service industries, process design must consider the flow of materials, information, and people.

In the manufacturing sector, process design often focuses on the movement of raw materials through various stages of production, such as assembly, inspection, and packaging. This involves precise planning for the location of workstations, the timing of operations, and the integration of quality checks. In service industries, the focus shifts to the flow of customer interactions, digital information, and task sequences. For example, in a coffee shop, the process might include order-taking, beverage preparation, and customer pickup. Each step must be timed and organized to prevent bottlenecks and ensure a smooth customer experience.

Good process design also anticipates variability. Businesses must account for fluctuations in demand, supply chain interruptions, and differences in employee performance. A resilient process design allows for adjustments without breaking the entire system. This often requires a balance between standardization, which ensures consistency, and flexibility, which enables adaptation.

[Placeholder for Visual: Diagram showing parallel examples of a manufacturing process flow and a service process flow]

3.2 Key Principles of Effective Processes

There are several guiding principles that underpin effective process design. The first is alignment with organizational strategy. Processes must directly support the goals and objectives of the business. A process that is efficient but misaligned with the company's value proposition may save time yet fail to deliver what customers truly want.

The second principle is simplicity. Complex processes often lead to confusion, errors, and unnecessary delays. This does not mean oversimplifying to the point of inefficiency, but rather removing redundant steps, unclear handoffs, and outdated procedures. In practice, simplicity often reveals itself through

ease of training. If a process is too complicated to explain to a new hire in a reasonable amount of time, it likely needs refinement.

Third, processes should be designed for measurement. Without clear performance indicators, it is impossible to determine whether a process is successful or in need of improvement. Metrics might include cycle time, error rate, customer satisfaction, or cost per unit. These measures create a feedback loop for ongoing improvement efforts.

Finally, effective processes are customer-focused. Every step should add value from the perspective of the customer. In a service context, this might mean reducing wait times or personalizing interactions. In manufacturing, it could involve improving durability, aesthetics, or ease of use.

[Placeholder for Visual: Chart illustrating process principles mapped to manufacturing and service examples]

3.3 Mapping and Documenting Processes

Before improving a process, it must be fully understood. Process mapping is the act of creating a visual representation of the steps, decision points, and flows involved. This is often the first step in both designing new processes and diagnosing problems in existing ones.

Common tools for process mapping include flowcharts, swimlane diagrams, and value stream maps. Flowcharts provide a straightforward step-by-step visualization, while swimlane diagrams add layers of responsibility, showing who is responsible for each part of the process. Value stream mapping, a tool from lean manufacturing, highlights both value-added and non-value-added steps, making it easier to identify waste.

Documentation should go beyond just the diagram. Each process should have a written description outlining its purpose,

scope, inputs, outputs, and performance measures. This ensures that even if key employees leave, the organization retains knowledge of how the process works. It also promotes consistency across teams and locations.

At Pemi Coffee Roasters, for example, mapping the roasting process revealed inefficiencies in how beans were moved between roasting and packaging stations. The visual map made it clear that equipment placement caused extra walking time, which slowed production during peak order periods. By reorganizing the workspace, Pemi reduced time per batch without compromising quality.

[Placeholder for Visual: Process map of Pemi Coffee Roasters' roasting-to-packaging workflow]

3.4 Analyzing Processes for Improvement

Once a process is mapped and documented, the next step is analysis. Analysis involves looking at each step critically to determine whether it adds value, creates waste, or introduces risk. Common frameworks include the lean concept of eliminating the seven types of waste and Six Sigma's focus on reducing variability.

One useful method is to walk through the process as though you were the customer. This "customer journey" approach helps uncover delays, confusion, or dissatisfaction that may not be visible from an internal perspective. For instance, a customer calling Pemi Coffee Roasters' customer service might experience a long hold time. While internally this delay might be due to staff attending to in-store customers, from the customer's perspective it feels like neglect.

Another approach is using data-driven analysis. By measuring process performance over time, patterns emerge that indicate where improvements could yield the greatest benefit. If a roasting machine at Pemi has a higher defect rate

on certain settings, analysis of production data can lead to targeted adjustments.

Process analysis also benefits from cross-functional input. People from different roles often see different problems and opportunities. A warehouse worker may notice delays in receiving raw beans, while a marketing manager may see gaps in the timing of product launches. Together, these insights lead to more comprehensive improvements.

3.5 Lean Thinking in Process Design

Lean thinking is a philosophy that seeks to maximize value while minimizing waste. It originated in manufacturing, particularly in the Toyota Production System, but has since been widely applied to services, healthcare, logistics, and technology. In lean, "waste" refers to any activity that consumes resources but does not create value for the customer.

There are traditionally seven categories of waste: overproduction, waiting, unnecessary transportation, over-processing, excess inventory, unnecessary motion, and defects. In service industries, these wastes manifest differently. For example, unnecessary motion in a café could mean baristas walking back and forth to retrieve ingredients that could be kept closer to the work area. In manufacturing, it might involve moving parts between distant workstations when a more efficient layout could bring them closer together.

For Pemi Coffee Roasters, lean thinking has been applied in both production and fulfillment. One improvement came from reducing the amount of roasted coffee stored before packaging. By packaging immediately after roasting, they reduced inventory holding time, preserved freshness, and minimized the risk of unsold stock going stale.

[Placeholder for Visual: Side-by-side example of lean waste types in manufacturing vs. service at Pemi Coffee Roasters]

3.6 Continuous Improvement and the PDCA Cycle

Continuous improvement, or kaizen, is the idea that processes should never be considered "finished." Even when a process is performing well, there is always room for refinement. The Plan-Do-Check-Act (PDCA) cycle provides a framework for this ongoing work.

In the "Plan" phase, the improvement team identifies a specific problem and proposes a solution. The "Do" phase involves implementing the solution on a small scale or as a pilot project. In the "Check" phase, the results are measured against the expected outcome. Finally, in the "Act" phase, the solution is either fully implemented, modified, or abandoned based on the findings.

Pemi Coffee Roasters uses PDCA to test new production schedules during the holiday rush. By piloting a schedule that staggered roasting shifts and packaging times, they discovered they could handle a 20% increase in orders without overtime pay. However, they also learned that some shifts led to bottlenecks at quality control, prompting adjustments before full-scale rollout.

Continuous improvement thrives in a culture that encourages employees to suggest changes. Small improvements, like adjusting the position of tools or revising a form to make it easier to complete, can add up to significant gains over time.

[Placeholder for Visual: PDCA cycle diagram with a Pemi-specific example in each phase]

3.7 Tools for Process Analysis and Redesign

Several tools support process improvement efforts. One is the cause-and-effect diagram, also known as a fishbone diagram, which helps teams identify potential causes of a problem. This encourages thinking beyond obvious explanations to uncover root causes. Another is the Pareto chart, which ranks problems by frequency or impact, allowing teams to focus on the issues that will yield the greatest benefit.

Simulation modeling is also valuable, especially in complex operations where changes could have multiple ripple effects. For instance, using simulation software, Pemi Coffee Roasters can test how altering packaging line speed will affect overall throughput and staffing needs without physically making the change first.

Workflow management software provides digital visibility into process steps, particularly useful in service or knowledge work. It allows managers to see where tasks are stalled, track deadlines, and measure turnaround times. When paired with process mapping, these tools create a powerful improvement toolkit.

The selection of tools should align with the complexity of the problem. While high-tech simulations may be necessary for major layout changes, a simple flowchart and discussion may be enough to resolve a bottleneck in customer service response times.

[Placeholder for Visual: Fishbone diagram example for a Pemi order fulfillment delay]

3.8 Reducing Process Variability

Process variability refers to the inconsistency in how a process performs over time. High variability leads to unpredictable outcomes, reduced customer satisfaction, and wasted resources. While some variability is inevitable due to factors like changing customer demand, much of it can be reduced through standardization and training.

At Pemi Coffee Roasters, roasting time and temperature must be carefully controlled to ensure consistent flavor. Early in their operations, variability arose when different employees adjusted roaster settings based on personal judgment rather than following a standardized profile. By introducing detailed roasting profiles and training staff to follow them precisely, Pemi significantly reduced flavor inconsistencies.

In service processes, variability often stems from differences in how employees handle customer interactions. Standard operating procedures (SOPs) help create consistency without removing all flexibility. For example, customer service representatives can be given guidelines on tone, response times, and problem resolution steps, while still being encouraged to add personal touches.

Reducing variability also involves preventive maintenance on equipment, regular audits, and a commitment to documenting changes to processes so that improvements are consistently applied.

3.9 Implementing and Sustaining Process Improvements

Designing and improving processes is only half the challenge—ensuring that improvements last is equally

important. Implementation requires clear communication, employee training, and a plan for monitoring performance after changes are made.

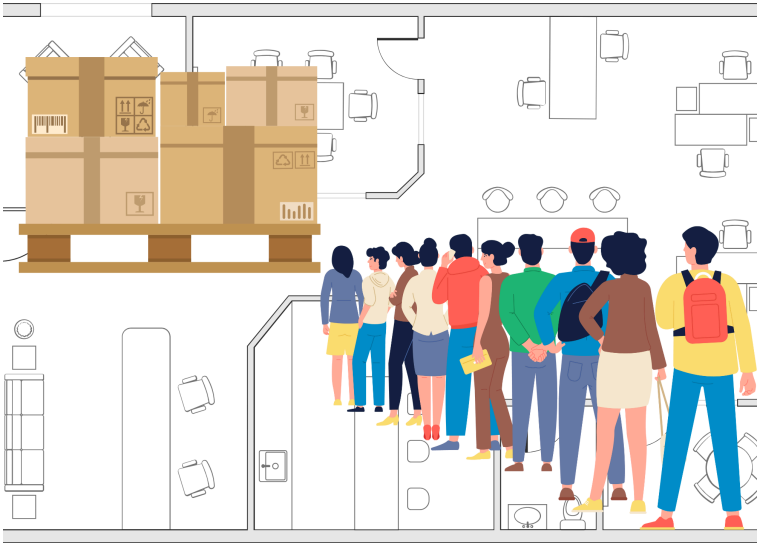
Resistance to change is common. Employees may be comfortable with existing processes or skeptical about new ones. Involving staff early in the redesign process helps reduce resistance and increases buy-in. When employees understand the reasons for change and see their feedback incorporated, they are more likely to support the new approach.

Sustaining improvements depends on accountability and measurement. Assigning process owners ensures that someone is responsible for monitoring performance and initiating new improvement cycles when needed. This prevents processes from gradually reverting to old habits.

At Pemi Coffee Roasters, after reorganizing their packaging area for greater efficiency, the operations manager conducted monthly reviews for six months to confirm that the new layout was being used correctly. Minor adjustments were made based on staff input, keeping the improvements both relevant and sustainable.

[Placeholder for Visual: Before-and-after layout comparison of Pemi's packaging area]

4. Capacity Planning and Facility Layout



4.1 Understanding Capacity in Business Operations

Capacity is the maximum output a business can produce within a given period under normal working conditions. It is a measure of how much a company can deliver without overstretching its resources, and it forms the foundation for making informed operational decisions. Whether the organization is a manufacturing plant, a service provider, or a hybrid of both, capacity defines the upper limit of productivity before performance begins to degrade.

In manufacturing, capacity is often expressed in units of output per day or week, while in service settings it might be measured in customers served, appointments completed, or transactions processed. For example, a coffee roasting facility may quantify capacity by pounds of beans roasted daily, while a café might measure it by the number of beverages served during peak hours. The distinction between these metrics is crucial because service operations often rely heavily on variable human interaction, making capacity less predictable than in manufacturing.

For Pemi Coffee Roasters in Plymouth, New Hampshire, understanding capacity is about balancing two distinct but interconnected sides of the business. On one side, the roasting operation produces coffee beans for wholesale distribution to regional cafés and retailers. On the other, the café and tasting room serve a steady flow of customers who expect not just a product but an experience. Capacity planning here must account for the output of the roasters, the seating available, and the flow of orders through the espresso machines, ensuring that neither side suffers from delays or bottlenecks.

Effective capacity planning requires consideration of both short-term and long-term perspectives. Short-term capacity decisions often involve managing current resources—adjusting staff schedules, extending operating hours, or prioritizing certain orders. Long-term capacity planning involves decisions about facility expansion, purchasing new equipment, or implementing technology that can automate processes. Businesses that fail to align capacity with demand risk losing customers, overburdening employees, or wasting resources on underused infrastructure.

[Visual Placeholder: Diagram showing the relationship between design capacity, effective capacity, and actual output in a service-manufacturing hybrid like Pemi Coffee Roasters.]

4.2 Measuring and Managing Capacity

Measuring capacity is not a one-size-fits-all process. The metrics chosen depend heavily on the type of business, its products or services, and the variability of demand. For manufacturing operations, managers might rely on production rates, machine utilization, or throughput time. Service organizations may measure appointment slots, transaction times, or the number of customers served per employee. Hybrid businesses like Pemi Coffee Roasters must consider both physical production and service delivery metrics.

One of the most common approaches to managing capacity is to compare actual output against effective capacity. Effective capacity is the maximum possible output after accounting for downtime, maintenance, staffing limitations, and inefficiencies. This differs from design capacity, which assumes optimal conditions with no interruptions—a rarely achievable scenario in real-world operations. At Pemi, the roasters might have a design capacity of 1,200 pounds per day, but effective capacity may be closer to 950 pounds after accounting for cleaning, calibration, and scheduled breaks.

Capacity management is an ongoing process that requires continuous monitoring and adjustment. For example, during tourist season in the Lakes Region, Pemi may experience a significant spike in café customers. To meet this demand, they could extend weekend hours, increase staff on the espresso bar, or pre-roast larger batches to free up production time. Each decision must weigh the potential revenue gain against the operational cost, ensuring that the business maintains both profitability and service quality.

Capacity planning is not static. It must evolve alongside the business and its market. Changes in customer preferences, technology, and competitive pressures all influence capacity

needs. By integrating regular capacity reviews into operational strategy, businesses can remain agile and avoid crises where demand outpaces capability or where overinvestment in unused capacity erodes profitability.

[Visual Placeholder: Graph showing seasonal capacity fluctuations at Pemi Coffee Roasters, with peaks during summer tourist season and holiday months.]

4.3 Short-Term vs. Long-Term Capacity Planning

Short-term capacity planning focuses on making immediate adjustments to meet fluctuating demand without making major capital investments. This may involve adding temporary staff, outsourcing certain tasks, or adjusting shift schedules. At Pemi Coffee Roasters, a short-term adjustment could be bringing in part-time baristas during a popular community event or renting an additional espresso machine for a seasonal festival. These changes are flexible and reversible, allowing the business to respond quickly to unexpected surges or drops in demand.

Long-term capacity planning, on the other hand, addresses fundamental changes to the business's infrastructure and capabilities. These decisions often require significant financial investment and strategic forecasting. Examples include building an additional roasting room, purchasing high-capacity grinders, or redesigning the café layout to accommodate more customers. For Pemi, a long-term plan might also involve expanding into a second location to capture a larger market share in New Hampshire.

The key difference between short- and long-term planning is the level of commitment and risk. Short-term changes are relatively low-risk because they can be reversed or scaled back

if demand shifts unexpectedly. Long-term changes carry greater risk but also the potential for substantial reward if executed at the right time. Businesses must evaluate these decisions using both market data and operational analytics to avoid costly missteps.

Balancing the two planning horizons is essential. An overemphasis on short-term fixes can lead to inefficiencies and missed growth opportunities, while overly aggressive long-term investments can result in underutilized assets. By aligning both approaches with demand forecasts and business objectives, organizations like Pemi Coffee Roasters can ensure sustainable growth and customer satisfaction.

4.4 Understanding Facility Layout

Facility layout is the arrangement of physical spaces, equipment, and resources to optimize workflow and efficiency. It directly affects production speed, employee productivity, customer experience, and safety. A poor layout can lead to wasted motion, congestion, and even increased operational costs. A well-designed layout streamlines processes, reduces unnecessary movement, and supports both capacity and quality goals.

There are several types of facility layouts, including product layouts, process layouts, cellular layouts, and fixed-position layouts. Each type serves different operational needs. A product layout arranges resources according to the sequence of production steps, making it efficient for high-volume, standardized products. A process layout groups similar tasks together, offering flexibility for varied products. Cellular layouts arrange workstations to handle families of similar products, while fixed-position layouts are used for large projects where the product remains stationary.

At Pemi Coffee Roasters, the roasting facility might use a product layout, where green coffee beans move through roasting, cooling, grinding, and packaging in a linear flow. In the café, however, a process layout might be more appropriate, with stations for espresso preparation, pastry display, and customer payment organized to allow multiple orders to be processed simultaneously. This hybrid approach reflects the dual nature of their operations.

The layout must also consider customer experience. In the café, customers should be able to order, pay, and receive their beverages with minimal confusion. The seating area should balance comfort with turnover rates, ensuring a pleasant environment that still supports the café's capacity goals. In the roasting area, safety, ventilation, and workflow efficiency take priority, ensuring employees can work effectively without hazards or bottlenecks.

[Visual Placeholder: Floor plan sketch showing both the roasting facility and café areas of Pemi Coffee Roasters, highlighting workflow patterns.]

4.5 Factors Influencing Capacity Decisions

Capacity decisions are rarely made in isolation. They are shaped by a complex interplay of market demand, operational costs, workforce availability, and technological capabilities. Seasonal fluctuations can dramatically influence capacity requirements, especially in businesses with strong ties to tourism or holiday cycles. For Pemi Coffee Roasters, the summer months in New Hampshire bring in an influx of vacationers, which increases both wholesale orders from local cafés and foot traffic in the café itself. This surge means more roasted beans, more pastries, and more labor hours must be planned in advance.

Economic conditions also play a role. During times of economic growth, demand for premium coffee and café experiences may rise, justifying an expansion of roasting or seating capacity. Conversely, during economic downturns, consumer spending may shift toward lower-priced options, prompting businesses to adjust capacity downward to avoid excess inventory or idle staff. These decisions require not only data analysis but also an understanding of customer behavior patterns.

Technology is another major factor. Upgrading to faster roasting equipment, implementing automated ordering systems, or introducing advanced scheduling software can expand capacity without requiring more physical space. Pemi Coffee Roasters could, for example, adopt a system that lets customers order and pay from their phones before arriving, reducing congestion at the counter and improving throughput. Such investments, however, must be weighed against their cost and expected return.

Finally, the availability and skill level of labor directly influence capacity. A shortage of trained baristas or skilled roasters can limit output even if equipment and facilities are capable of more. Cross-training employees to handle multiple roles can provide flexibility and resilience, allowing Pemi to maintain consistent service levels even during staffing challenges.

4.6 Capacity Strategies

Organizations typically adopt one of three broad capacity strategies: lead, lag, or match capacity. A lead strategy involves adding capacity in anticipation of future demand, which can position a company to capture market share quickly but carries the risk of underutilization if demand does not materialize. For Pemi Coffee Roasters, this might mean purchasing an

additional roaster before wholesale orders have fully justified the investment, betting on continued growth in the specialty coffee market.

A lag strategy waits until demand has clearly increased before adding capacity. This conservative approach reduces the risk of excess capacity but can result in lost sales or service delays during periods of high demand. If Pemi were to wait until the café was consistently overcrowded before expanding seating, they might lose customers to competitors in the meantime.

The match strategy takes a middle-ground approach, adding capacity in small increments that align closely with demand trends. For Pemi, this could mean gradually expanding roasting shifts, hiring one or two additional staff during busy months, or rearranging the café layout to add a few more tables without a major renovation. This approach balances risk and responsiveness, but it requires constant monitoring of performance metrics to ensure the timing of changes is optimal.

Each strategy has its place, and businesses often blend them depending on the situation. For example, Pemi might take a lead approach in preparing for the holiday season by increasing roasted bean inventory ahead of time, while using a match approach for café staffing during less predictable shoulder seasons.

4.7 Facility Layout Design Considerations

Designing a facility layout is as much about anticipating movement and interaction as it is about placing equipment and furniture. Every step taken by an employee, every turn in a workflow, and every point of customer interaction should

be intentional. Poor layout decisions can create choke points, redundant tasks, or unsafe working conditions.

The physical characteristics of the building often dictate some layout constraints, such as the location of structural walls, windows, and utility hookups. At Pemi Coffee Roasters, the roasting machines must be placed where ventilation and exhaust systems can operate effectively, which may limit flexibility in that area. The café, however, may allow more freedom to adjust counters, seating arrangements, and traffic flow.

Workflow analysis is critical. Observing how employees move through their tasks can reveal inefficiencies or unnecessary steps. For example, if baristas must cross paths frequently while preparing drinks, rearranging the espresso machines and milk stations might reduce collisions and speed up service. In the roasting area, having packaging materials stored closer to the cooling bins could save minutes on each batch, which adds up over time.

Aesthetics and branding also influence layout decisions. The café space should not only function efficiently but also reflect Pemi's identity as a welcoming, artisanal coffee brand. Sight lines to the roasting equipment can engage customers by letting them see and smell the roasting process, turning part of the production area into an experience rather than something hidden away. This blending of operational efficiency and brand storytelling can create a competitive advantage.

4.8 Integrating Capacity Planning and Layout for Business Success

Capacity planning and facility layout are deeply interconnected. An efficient layout can enhance capacity without the need for additional staff or equipment, while poor

design can undermine even the best capacity strategy. At Pemi Coffee Roasters, a well-optimized roasting and café layout ensures that every step of production and service supports the business's ability to meet demand.

Consider a scenario where Pemi has invested in a second roaster to double production. If the layout does not allow beans to move easily from roasting to cooling to packaging, the full potential of that capacity increase will not be realized. Likewise, if the café's seating area is not arranged to accommodate the increased customer volume that comes with more wholesale recognition, the customer experience could suffer.

The integration of these two elements begins with data. Demand forecasts inform capacity requirements, which in turn guide layout decisions. If data shows that weekend traffic is growing faster than weekday traffic, the café layout might be adjusted to accommodate more customers during those peak hours, perhaps with flexible seating that can be reconfigured quickly. In the roasting area, workflow analysis might lead to rearranging equipment so that more batches can be processed with less downtime.

Regular review is essential. Both capacity needs and optimal layouts can change over time due to shifts in market trends, technology, or customer behavior. Pemi Coffee Roasters might find that a layout designed for in-store service is less effective if takeout orders suddenly increase, prompting a redesign of counter space and order pickup areas. By treating capacity planning and layout as dynamic, interconnected processes, businesses can remain agile and responsive in a competitive market.

Key Takeaways

- Capacity sets the maximum sustainable output and guides both short- and long-term decisions.
- Effective capacity is always less than design capacity due to real-world constraints.
- Short-term planning relies on flexible adjustments, while long-term planning requires larger investments.
- Facility layout shapes efficiency, productivity, and customer experience.
- Lead, lag, and match are the main capacity strategies, each balancing risk differently.
- Capacity planning and layout must be integrated to align resources, workflows, and demand.

Chapter Four References

Lean Enterprise Institute. (n.d.). *What is lean?* Retrieved from https://www.lean.org/explore-lean/what-is-lean

National Institute of Standards and Technology. (2024, October 29). *Revision of Baldrige framework to increase its value as improvement guide*. U.S. Department of Commerce. https://www.nist.gov/news-events/news/2024/10/revision-baldrige-framework-increase-its-value-improvement-guide

OpenStax. (2019). *Principles of management*. OpenStax CNX. Retrieved from https://openstax.org/books/principles-management/pages/6-introduction

Stevenson, W. J. (2021). *Operations management* (14th ed.). McGraw-Hill Education.

Jacobs, F. R., & Chase, R. B. (2021). *Operations and supply chain management* (16th ed.). McGraw-Hill Education.

5. Inventory Management



5.1 Introduction to Inventory Management

Inventory management is the process of overseeing, controlling, and optimizing the stock of goods and materials an organization uses or sells. It plays a critical role in ensuring that a company can meet customer demand without holding excess inventory that ties up capital. Effective inventory management strikes a balance between too much stock, which can lead to waste and higher carrying costs, and too little stock, which can cause lost sales and dissatisfied customers.

Whether a business operates in manufacturing or service sectors, understanding how to control inventory is essential to operational efficiency and profitability.

The concept applies equally to a company like Pemi Coffee Roasters and to a large automotive manufacturer. For Pemi Coffee Roasters, inventory includes green coffee beans, roasted beans, packaging materials, and even the cups and sleeves used in their cafes. For a manufacturer, inventory may include raw materials, work-in-progress goods, and finished products waiting for distribution. In both cases, inventory is not just a physical asset—it is a strategic resource that directly affects customer satisfaction, cash flow, and competitive advantage.

The service sector often views inventory differently than manufacturing, focusing on consumables, supplies, and sometimes capacity as forms of inventory. A hotel, for example, may think of available room nights as a type of perishable inventory, while a hospital manages its inventory of medical supplies with the same precision a manufacturer applies to raw materials. Regardless of the context, the principles of managing inventory—forecasting needs, controlling costs, and aligning with demand—are universal.

[Visual Placeholder: Diagram showing the relationship between overstock, understock, and optimal inventory level]

5.2 Types of Inventory

Businesses categorize inventory into types to manage it more effectively. In manufacturing, three main categories dominate: raw materials, work-in-progress (WIP), and finished goods. Raw

materials are the unprocessed inputs, like green coffee beans for Pemi Coffee Roasters or sheet metal for an auto plant. WIP inventory includes goods that are partially assembled, such as coffee beans in the roasting stage or vehicles on the assembly line. Finished goods are the final, sale-ready products that are ready for customers.

In service industries, inventory often takes a different form. It may be spare parts for maintenance work, linens and cleaning supplies for hotels, or pre-prepared meals for catering businesses. The perishable nature of many service-sector goods introduces added urgency to inventory management. In restaurants, for instance, unsold meals cannot be stored indefinitely, creating a direct link between demand forecasting and waste reduction.

Pemi Coffee Roasters manages multiple inventory types simultaneously. Their raw coffee beans must be stored under controlled conditions to maintain freshness. Their roasted beans, a finished product, have a limited shelf life and require careful turnover to ensure quality. Even their branded mugs and merchandise form part of their finished goods inventory, serving both functional and marketing purposes.

Some organizations also classify inventory into maintenance, repair, and operations (MRO) supplies, which are not directly part of production but are essential for keeping operations running smoothly. These can include cleaning supplies, machine parts, or even office stationery. MRO inventory management ensures that breakdowns or shortages in support materials do not interrupt the main production or service delivery.

[Visual Placeholder: Chart illustrating different inventory types for manufacturing vs. service industries]

5.3 Inventory Functions in Operations

Inventory serves several critical functions within operations. It acts as a buffer against fluctuations in demand and supply. In manufacturing, this buffer ensures that production lines keep running even if suppliers are late or demand spikes unexpectedly. For a coffee roaster like Pemi, maintaining a buffer of green beans means they can continue roasting even if a shipment from a coffee-growing region is delayed by weather or transportation issues.

Another key function of inventory is decoupling different stages of production. In complex manufacturing, having work-in-progress inventory allows each stage of production to operate independently without waiting for other stages to catch up. In service settings, inventory can help separate preparation from delivery. For instance, a catering company may prepare certain food items in advance, storing them until the event begins.

Inventory also plays a strategic role in hedging against price fluctuations. When raw material prices are expected to rise, companies may choose to purchase and store more inventory in advance. Pemi Coffee Roasters might buy larger quantities of beans if they anticipate higher prices due to a poor harvest in a key coffee-growing region. This approach can protect margins and stabilize pricing for customers.

Additionally, inventory supports seasonal and promotional strategies. Retailers often build up inventory ahead of holiday seasons, and manufacturers prepare for peak demand periods. Service businesses, such as tour operators, may stock up on promotional materials or equipment before the busy season. In every case, the goal is to align inventory levels with operational and marketing plans to maximize profitability.

5.4 The Role of Inventory in Supply Chain Performance

Inventory is not just a passive element in the supply chain; it actively shapes the speed, cost, and reliability of the entire system. Effective inventory management ensures that the right products are available at the right time and in the right place, minimizing delays and customer dissatisfaction. In global supply chains, inventory acts as a shock absorber, smoothing out disruptions caused by transportation delays, customs issues, or supplier problems.

For Pemi Coffee Roasters, the role of inventory in the supply chain is particularly visible in their multi-location operations. Their central roasting facility must coordinate inventory levels with retail outlets to ensure each location has the correct mix of coffee varieties and merchandise. Too much inventory at a store increases holding costs and risks spoilage; too little results in missed sales opportunities.

In manufacturing, the role of inventory in supply chain performance is closely tied to lead times. Shorter lead times generally require less safety stock, while longer lead times demand higher inventory levels to maintain service levels. The challenge is finding the balance that supports customer expectations without tying up excessive capital in stock. This balance can be influenced by technology, supplier relationships, and demand forecasting accuracy.

Inventory visibility is another crucial factor. Without clear, real-time insight into inventory levels across the supply chain, decision-making becomes reactive and prone to error.

Advances in inventory tracking systems, such as RFID and integrated ERP solutions, have greatly improved this visibility. For service businesses like Pemi's cafes, point-of-sale systems linked to inventory management software can trigger automatic replenishment orders, ensuring popular items never run out unexpectedly.

5.5 Costs Associated with Inventory

Inventory management is not simply about counting and storing items; it involves managing the costs that come with holding and replenishing stock. The primary costs include ordering costs, holding costs, and shortage costs. Ordering costs are incurred each time a company places an order, whether for raw materials, finished goods, or consumables. These costs include administrative work, transportation, and receiving processes. For Pemi Coffee Roasters, ordering costs may involve negotiating with suppliers, coordinating shipping for green beans, and scheduling deliveries to ensure freshness.

Holding costs, also known as carrying costs, refer to expenses related to storing inventory over time. This includes warehousing, insurance, depreciation, spoilage, and obsolescence. For example, roasted coffee has a limited shelf life, so holding it too long not only ties up capital but can also reduce quality, leading to customer dissatisfaction. In manufacturing, holding costs can be even higher due to the space and equipment needed to maintain inventory in good condition.

Shortage costs occur when a company does not have enough inventory to meet demand. These can include lost sales, rushed shipping fees, and even damage to the brand's reputation. For a cafe, a shortage of popular coffee blends during peak hours

may frustrate loyal customers and lead them to competitors. In manufacturing, shortages can halt production entirely, incurring both lost revenue and increased costs to restart operations.

Balancing these costs requires careful analysis. The goal is to minimize total inventory-related expenses while maintaining service levels. Businesses often turn to economic order quantity (EOQ) models, just-in-time (JIT) approaches, or safety stock strategies to achieve this balance. The right mix depends on the nature of the products, demand variability, and the company's overall operations strategy.

5.6 Inventory Control Systems

An inventory control system is the framework a business uses to monitor stock levels, manage replenishment, and ensure accuracy. These systems range from manual counts and spreadsheets to fully automated, integrated solutions that link point-of-sale data directly to procurement. For Pemi Coffee Roasters, a cloud-based inventory management platform allows them to track bean usage at each cafe, monitor freshness dates, and trigger roasting schedules based on demand trends.

Two primary approaches to inventory control are periodic and perpetual systems. A periodic system updates inventory records at set intervals, such as weekly or monthly, often using physical counts. This method is simpler but provides less real-time accuracy. A perpetual system, on the other hand, updates inventory data continuously as transactions occur. This approach offers greater accuracy and enables quicker responses to demand changes but requires more sophisticated technology.

In manufacturing environments, inventory control systems often integrate with production planning tools. This ensures that materials are available when needed without overstocking. For example, a factory using just-in-time production needs precise, real-time inventory data to avoid costly delays. In service industries, inventory control systems can help manage perishable goods, schedule replenishments, and even optimize storage space.

Technology plays a critical role in modern inventory control. Barcoding, RFID tagging, and IoT-enabled sensors provide accurate, real-time data on stock levels and locations. For Pemi, using barcode scanning for coffee bags not only improves tracking but also reduces human error during order fulfillment. These technological tools enhance visibility, accuracy, and efficiency across the supply chain.

5.7 Demand Forecasting for Inventory Planning

Accurate demand forecasting is essential to effective inventory management. Forecasting predicts future customer demand so businesses can stock the right amount of inventory at the right time. Inaccurate forecasts can lead to overstocking, tying up capital in unsold goods, or understocking, resulting in missed sales. Pemi Coffee Roasters uses historical sales data, seasonal trends, and promotional calendars to forecast demand for different coffee blends, ensuring they roast the correct quantities.

There are two main forecasting approaches: qualitative and quantitative. Qualitative forecasting relies on expert judgment, market research, and customer feedback. This method is

particularly useful when launching new products or operating in markets with little historical data. Quantitative forecasting uses statistical models and historical sales patterns to predict future demand. These models may account for seasonality, trends, and cyclical patterns.

In manufacturing, demand forecasting aligns production schedules with material procurement, minimizing waste and reducing holding costs. In service industries, forecasting helps anticipate demand for consumables, labor, and capacity. For example, a hotel may use forecasting to determine how much laundry detergent to order in preparation for peak tourist season.

Integrating forecasting with inventory management systems creates a dynamic feedback loop. As actual sales data comes in, forecasts can be adjusted to improve accuracy. For Pemi Coffee Roasters, real-time sales tracking allows for quick adjustments during promotions or unexpected demand spikes, helping avoid both shortages and overstock.

[Visual Placeholder: Chart comparing forecasted demand with actual sales over time]

5.8 Inventory Metrics and Performance Indicators

Measuring inventory performance is crucial for identifying improvement opportunities. Common metrics include inventory turnover, days sales of inventory (DSI), and stockout rate. Inventory turnover measures how many times a company sells and replaces its inventory over a given period. Higher turnover generally indicates efficient inventory management.

For Pemi Coffee Roasters, a high turnover rate for their flagship blend suggests strong demand and efficient restocking.

Days sales of inventory calculates the average number of days it takes to sell the inventory on hand. Lower DSI indicates quicker sales, reducing holding costs. Stockout rate measures how often products are unavailable when customers want them. A high stockout rate signals potential problems in forecasting, replenishment, or supplier reliability.

Other useful metrics include carrying cost percentage, order accuracy rate, and fill rate. Carrying cost percentage shows the proportion of inventory value tied up in holding costs. Order accuracy rate reflects how often customer orders are fulfilled without errors. Fill rate measures the percentage of customer demand met without backorders or delays.

Regularly monitoring these metrics enables businesses to make data-driven decisions. For example, if Pemi notices an increasing DSI for a seasonal blend, they can adjust production schedules or run promotions to move the stock faster. In manufacturing, monitoring turnover and fill rates can highlight bottlenecks and inefficiencies in the supply chain.

5.9 Inventory Management in Manufacturing vs. Services

While the principles of inventory management apply broadly, the strategies can differ significantly between manufacturing and service industries. In manufacturing, inventory often represents a significant portion of a company's assets, and managing it effectively can directly impact profitability. Manufacturers must coordinate raw material procurement,

production schedules, and finished goods distribution to meet customer demand efficiently.

In service industries, inventory may be less central but still critical to operations. For example, a catering company manages perishable food inventory with a focus on minimizing waste while meeting event requirements. Pemi Coffee Roasters operates in a hybrid space, producing roasted coffee (manufacturing) while also running cafes (service). This dual nature means they must manage both manufacturing-style inventories like raw beans and service-oriented inventories like pastries and milk.

Lead times also differ between sectors. Manufacturing often involves longer lead times for raw materials and production processes, requiring careful planning. Service industries typically deal with shorter lead times but must be highly responsive to fluctuating customer needs. For Pemi, green bean procurement may require months of planning, while replenishing milk or syrups for cafes can be handled within days.

Ultimately, effective inventory management in any sector requires understanding the specific nature of demand, the perishability of goods, and the operational constraints unique to the business model.

5.10 Future Trends in Inventory Management

Inventory management is evolving rapidly, driven by advances in technology, shifts in consumer expectations, and global supply chain changes. One major trend is the increased use of artificial intelligence and machine learning to improve

demand forecasting, optimize stock levels, and automate replenishment decisions. For Pemi Coffee Roasters, AI could analyze sales data across all locations, weather patterns, and even social media trends to predict demand for specific coffee blends.

Another trend is the integration of sustainability into inventory management practices. Businesses are seeking ways to reduce waste, minimize overproduction, and source materials responsibly. For a coffee roaster, this might involve working with suppliers who use sustainable farming practices and reducing packaging waste.

Real-time inventory tracking is becoming standard, enabled by IoT devices, cloud computing, and mobile technology. These tools provide visibility across the supply chain, allowing for quicker decision-making and better coordination. Omnichannel retailing is also reshaping inventory strategies, as companies must manage stock across physical and online channels seamlessly.

Finally, the shift toward localized and resilient supply chains is influencing inventory practices. Companies are diversifying suppliers, increasing safety stocks for critical items, and using nearshoring strategies to reduce dependency on distant sources. For Pemi, sourcing beans from multiple regions can help mitigate the risks of crop failure or political instability in a single country.

Key Takeaways

- nventory management balances demand with cost by avoiding overstock and stockouts
- Types include raw materials, work-in-progress, finished goods, and MRO supplies
- Inventory functions as a buffer, decoupling tool, and hedge against price shifts
- Costs include ordering, holding, and shortage costs
- Control systems can be periodic or perpetual, with technology improving accuracy
- Forecasting demand is essential to align stock with customer needs
- Metrics like turnover, DSI, and stockout rate measure performance
- Manufacturing and service contexts differ, but both require alignment with demand
- Trends include AI forecasting, sustainability, and resilient supply chains

Chapter Five References

American Society for Quality. (n.d.). *Process mapping*. ASQ. https://asq.org/quality-resources/process-mapping

Jacobs, F. R., & Chase, R. B. (2021). *Operations and supply chain management* (16th ed.). McGraw-Hill Education.

Lean Enterprise Institute. (n.d.). *What is lean?* Lean Enterprise Institute. https://www.lean.org/explore-lean/what-is-lean

National Institute of Standards and Technology. (2024, October 29). *Revision of Baldrige framework to increase its value as improvement guide*. U.S. Department of Commerce. https://www.nist.gov/news-events/news/2024/10/revision-baldrige-framework-increase-its-value-improvement-guide

OpenStax. (2019). *Principles of management*. OpenStax, Rice University. https://openstax.org/books/principles-management/pages/6-introduction

Stevenson, W. J. (2021). *Operations management* (14th ed.). McGraw-Hill Education.

SECTION III

BUSINESS INTELLIGENCE WITH POWER BI

6. Business Intelligence Foundations



6.1 Introduction to Business Intelligence (BI)

Business Intelligence (BI) is the process of collecting, integrating, analyzing, and presenting business data to support decision-making. It combines technology, processes, and people to transform raw information into meaningful insights. BI helps organizations understand their performance, identify trends, and make better strategic choices. Whether

in a manufacturing plant tracking production metrics or a service-based business monitoring customer satisfaction scores, BI serves as the bridge between operational data and actionable strategy.

In today's competitive marketplace, BI is no longer optional—it is essential. Companies that embrace BI can identify inefficiencies faster, predict shifts in customer behavior, and adapt before competitors. For Pemi Coffee Roasters, BI tools provide visibility into sales patterns, roasting efficiency, and supply chain performance, allowing them to anticipate seasonal demand and optimize staffing levels. Without BI, these insights might remain buried in spreadsheets or siloed across departments.

The roots of BI can be traced back to traditional management information systems, but today's BI is far more dynamic. Cloud computing, real-time analytics, and user-friendly visualization platforms have made it accessible to businesses of all sizes. A small regional coffee roaster can now leverage the same analytical capabilities as a multinational manufacturer, leveling the competitive playing field.

[Visual Placeholder: Diagram showing the flow from raw data to actionable insights through BI processes]

6.2 Core Components of BI

At its core, BI consists of several interconnected components: data sources, data integration, data storage, analysis tools, and presentation layers. Each plays a role in converting raw data into decision-ready information. Understanding these components is key to building a strong BI foundation.

The first step is identifying data sources. These can include transactional systems, CRM platforms, ERP systems, IoT sensors, and even external data such as market trends or

weather patterns. Pemi Coffee Roasters, for instance, pulls data from its point-of-sale system, supplier databases, and production logs to get a holistic view of operations.

Data integration involves consolidating these disparate sources into a single, cohesive dataset. This step ensures that decision-makers have a unified view of the organization. In manufacturing, integration might mean linking machine performance data with quality control results. In service industries, it could involve combining customer feedback with operational metrics to spot service gaps.

The data storage component often involves a data warehouse or data lake. A data warehouse is structured and optimized for querying, while a data lake can store unstructured or semi-structured data for more flexible analysis. Finally, analysis tools and presentation layers—such as dashboards, reports, and visualization platforms—bring the data to life, making patterns and trends visible and actionable.

6.3 Data Sources for BI

The effectiveness of BI hinges on the quality and breadth of its data sources. Internal data is often the starting point, including sales transactions, inventory counts, production logs, and employee performance records. External data, such as market research, economic indicators, and social media sentiment, can enrich the internal perspective.

For Pemi Coffee Roasters, internal data sources include daily sales from each cafe, roasting batch quality reports, and supply chain lead times for green beans. External data might include global coffee price indexes, competitor pricing trends, and seasonal weather forecasts that affect crop yields. By combining both types of data, Pemi can make more accurate forecasts and adapt procurement strategies.

Manufacturers may rely on sensor data from production lines, supplier performance records, and real-time order tracking. Service industries often pull data from appointment scheduling systems, feedback surveys, and operational workflows. Each sector's BI approach is shaped by the data most critical to its success.

A challenge in using diverse data sources is ensuring compatibility and accuracy. Inconsistent formats, incomplete fields, and outdated information can undermine BI efforts. Data governance—setting rules for data entry, maintenance, and quality—is essential to building a reliable BI foundation.

6.4 Data Warehousing and Storage

Data warehousing is the backbone of BI, providing a centralized repository for integrated, historical data. A data warehouse stores data in a structured format optimized for querying and reporting, while newer data lake technologies store raw, unstructured data for flexible exploration. The choice between them—or a hybrid approach—depends on business needs and analytical goals.

Pemi Coffee Roasters might use a cloud-based data warehouse to store sales, production, and inventory data, enabling quick access for analysis. For more advanced insights, they could maintain a data lake with unstructured data such as customer reviews, social media posts, and supplier communications. This combination would allow them to merge structured transactional data with qualitative insights for richer analysis.

In manufacturing, a data warehouse may consolidate data from multiple plants, providing leadership with performance benchmarks and quality trend analysis. In service industries,

warehousing supports customer segmentation, loyalty program analysis, and service quality monitoring.

The design of a data warehouse involves choosing an appropriate schema—such as star, snowflake, or galaxy schema—to balance performance with flexibility. Regular updates, whether in batch or real time, ensure that decision-makers are working with current information.

6.5 Data Integration and ETL Processes

Data rarely arrives in a ready-to-use format. Integration involves extracting data from various sources, transforming it into a standardized structure, and loading it into the target system—a process known as ETL (Extract, Transform, Load). This step ensures that all data follows consistent definitions and formats, making analysis more accurate and efficient.

For Pemi Coffee Roasters, ETL might involve pulling daily sales data from POS systems, transforming it to match the warehouse schema, and loading it alongside production and supply chain data. In manufacturing, ETL processes might merge quality control data with machine sensor readings to identify performance issues.

Modern BI systems also use ELT (Extract, Load, Transform) processes, where raw data is loaded first and transformed later within the warehouse or lake. This approach can be more flexible and better suited for large, diverse datasets.

Integration is not just technical—it also requires aligning business definitions. For example, defining “order completion” consistently across departments prevents confusion in reporting. Data integration tools, whether standalone software or built-in features of BI platforms, streamline this process and help maintain accuracy.

6.6 BI Tools and Platforms

The BI landscape is rich with tools that cater to different needs, from self-service visualization platforms to enterprise-level analytics suites. Popular tools include Microsoft Power BI, Tableau, Qlik Sense, and Looker. These platforms allow users to create interactive dashboards, run ad-hoc queries, and share insights across the organization.

Pemi Coffee Roasters could use Power BI to monitor sales performance by cafe location, track roasting yields, and visualize supply chain bottlenecks. Manufacturing companies might use Tableau to analyze production efficiency across plants, while service businesses could leverage Qlik Sense for customer satisfaction trend analysis.

The choice of BI platform depends on several factors: integration capabilities, ease of use, scalability, cost, and the specific analytical needs of the business. Some organizations prefer cloud-based solutions for their flexibility and lower upfront costs, while others opt for on-premises systems to maintain tighter control over sensitive data.

Self-service BI is a growing trend, empowering non-technical users to create their own reports without relying on IT departments. While this democratizes data access, it also requires strong data governance to prevent inconsistent definitions and duplicate reports.

6.7 BI in Manufacturing and Services

BI applications vary across industries, reflecting different priorities and metrics. In manufacturing, BI often focuses on production efficiency, quality control, supply chain optimization, and predictive maintenance. For example,

analyzing machine downtime trends can reveal patterns that help schedule preventive maintenance, reducing costly breakdowns.

In service industries, BI is more likely to focus on customer behavior, service quality, and operational efficiency. A hotel chain might use BI to monitor booking trends, optimize staffing levels, and track guest satisfaction scores in real time.

Pemi Coffee Roasters bridges both worlds. On the manufacturing side, BI tracks roasting efficiency, defect rates, and green bean sourcing costs. On the service side, it monitors cafe foot traffic, transaction values, and customer loyalty program engagement. This hybrid approach requires a BI system flexible enough to handle both manufacturing and service data without losing analytical depth.

6.8 BI Implementation Process

Implementing BI is not just a matter of buying software—it's a structured process that involves strategic planning, technical integration, and cultural adoption. Organizations often start by defining clear objectives: what decisions do they want to improve, and what data is needed to support them? This step ensures that BI efforts are tied to business goals, not just technology trends.

The next phase is assessing current data infrastructure. For a manufacturer, this might involve mapping production data flows from sensors to the ERP system. For a service company, it could mean reviewing CRM data quality and accessibility. Pemi Coffee Roasters might evaluate how well their sales, roasting, and supply chain systems connect and where gaps in reporting occur. Without this assessment, BI projects risk being built on incomplete or unreliable data.

Once the baseline is understood, organizations select tools

and platforms that match their needs. This decision should weigh technical features, scalability, integration capabilities, and user-friendliness. After tool selection, data integration and warehouse setup follow, ensuring that all necessary data is accessible and standardized.

The final stage is user adoption. Even the best BI tools are useless if employees do not use them effectively. Training sessions, clear documentation, and ongoing support are critical. Some companies launch with a pilot group to refine processes before a full rollout. Pemi Coffee Roasters might start with a few cafe managers testing the dashboard, providing feedback on usability and the types of reports most helpful for daily operations.

6.9 Data Governance in BI

Data governance ensures that the information feeding BI systems is accurate, consistent, secure, and compliant with regulations. This is not just an IT responsibility—it's an organizational discipline involving policies, procedures, and cultural norms around data use.

At its core, governance defines who can access what data, how it should be stored, and how long it should be retained. For manufacturers, governance might involve strict protocols for production data due to regulatory requirements. In service industries, it often covers customer privacy, consent management, and compliance with laws like GDPR or CCPA.

Pemi Coffee Roasters could establish a data governance council to oversee BI data quality, ensuring definitions like "completed sale" or "batch yield" are applied consistently across locations. This prevents confusion when comparing metrics between cafes or roasting facilities.

Good governance also includes data lineage—tracking

where data comes from, how it has been transformed, and where it is used. This is critical for audit trails and for diagnosing anomalies in BI reports. Without governance, organizations risk "data chaos," where multiple versions of the truth undermine trust in analytics.

6.10 Key BI Metrics and KPIs

Metrics and Key Performance Indicators (KPIs) are the lifeblood of BI. While metrics are simply measured values, KPIs are strategic indicators tied to organizational goals. Choosing the right KPIs ensures BI delivers insights that drive meaningful action.

In manufacturing, common KPIs include Overall Equipment Effectiveness (OEE), defect rates, on-time delivery, and production cycle time. Service industries might track customer satisfaction scores, average response times, or client retention rates. Pemi Coffee Roasters could track KPIs like average daily sales per cafe, roasting efficiency percentage, and supplier lead time accuracy.

The key is not to overload BI dashboards with every possible metric. Instead, focus on a balanced set that covers financial performance, operational efficiency, customer experience, and future readiness. Metrics should be reviewed periodically to ensure they remain relevant as business priorities change.

A strong BI system not only tracks KPIs but also supports drill-down analysis—allowing managers to explore the underlying data behind a trend. For example, if Pemi sees a drop in average transaction size, BI tools can reveal whether it's due to fewer high-margin product sales, seasonal changes, or increased discounting.

6.11 Real-Time vs. Historical BI

BI can operate on different time horizons. Historical BI looks at past data to understand trends, measure performance, and inform strategic planning. Real-time BI, on the other hand, processes data as it is generated, enabling immediate responses to changing conditions.

Manufacturers benefit from real-time BI by monitoring machine health, detecting defects, and adjusting production on the fly. Service companies use it to track customer interactions in progress, enabling immediate service recovery if needed. Pemi Coffee Roasters might use real-time BI to monitor cafe sales during a promotional event, adjusting staffing levels or restocking inventory in response to demand.

Historical BI remains important for identifying long-term patterns. For instance, reviewing several years of sales and production data could help Pemi forecast seasonal peaks, optimize roasting schedules, and negotiate better supplier contracts.

The choice between real-time and historical BI depends on the decision-making context. Many organizations adopt a hybrid approach, combining the immediacy of real-time monitoring with the depth of historical analysis for comprehensive insight.

6.12 Challenges and Best Practices in BI

Implementing BI comes with challenges, from technical hurdles to cultural resistance. Data silos, inconsistent definitions, and poor data quality can undermine BI projects before they deliver value. User adoption is another common

obstacle—employees may distrust new systems or prefer old reporting methods.

One best practice is starting small with high-impact use cases. For example, Pemi Coffee Roasters might first implement BI for inventory management, reducing waste and stockouts. A manufacturing company might begin with downtime tracking to increase productivity. Early wins help build momentum and buy-in.

Another best practice is involving end users early in the design process. This ensures dashboards and reports are intuitive and relevant to daily work. Providing ongoing training is essential to maintain engagement and proficiency.

Finally, BI success requires continuous improvement. As business goals evolve, data sources expand, and technology advances, BI strategies must adapt. Organizations that treat BI as a static project risk falling behind competitors who embrace it as a dynamic, evolving capability.

6.13 Summary

Business Intelligence provides the tools, processes, and cultural foundation for turning raw data into actionable insight. From manufacturing plants optimizing equipment uptime to service companies improving customer experience, BI transforms how decisions are made.

Pemi Coffee Roasters demonstrates the power of BI in a hybrid manufacturing-service context—roasting beans efficiently, managing supplier relationships, and enhancing the cafe customer experience. By combining high-quality data, robust tools, and clear governance, they create a BI environment that supports both day-to-day operations and long-term strategy.

Ultimately, BI is most effective when aligned with business

goals, supported by strong governance, and embraced across the organization. With these foundations in place, companies can navigate change, seize opportunities, and maintain a competitive edge.

Key Takeaways

- Business Intelligence (BI) transforms raw data into actionable insights that support decision-making.
- Core components include data sources, integration, storage, analysis tools, and presentation layers.
- Effective BI depends on data quality, governance, and integration across systems.
- BI tools such as Power BI, Tableau, and Qlik provide visualization and self-service analytics.
- Applications vary by sector: manufacturing emphasizes efficiency and quality, while services focus on customer behavior and experience.
- Implementation requires clear goals, technical integration, and strong user adoption.
- Data governance ensures accuracy, consistency, and compliance.
- Metrics and KPIs must align with organizational

strategy, balancing real-time and historical insights.

- Common challenges include data silos, adoption resistance, and evolving business needs; best practices stress early wins, user involvement, and continuous improvement.

Chapter Six References

Gartner. (2023). *Business intelligence (BI): An overview.* Gartner. https://www.gartner.com/en/information-technology/glossary/business-intelligence-bi

Kimball, R., & Ross, M. (2013). *The data warehouse toolkit: The definitive guide to dimensional modeling* (3rd ed.). Wiley.

Microsoft. (n.d.). *What is Power BI?* Microsoft. https://learn.microsoft.com/en-us/power-bi/fundamentals/power-bi-overview

OpenStax. (2019). *Principles of management*. OpenStax, Rice University. https://openstax.org/books/principles-management/pages/6-introduction

Tableau. (n.d.). *Business intelligence and analytics: What it is and why it matters*. Tableau. https://www.tableau.com/learn/articles/business-intelligence

7. Data Visualization for Operations Insights



Learning Objectives

By the end of this chapter, students will be able to:

- Match an operational question to the right visual and audience.
- Choose chart types for time series, categories, distributions, relationships, and maps.

- Apply clear design (hierarchy, labels, honest scales) and avoid misleading visuals.
- Build a compact, usable dashboard in Excel/ Power BI with filters/slicers.
- Communicate insight and a concrete action using Pemi Coffee Roasters data.

7.1 Introduction to Data Visualization in Operations

Data visualization has become a critical bridge between raw operational data and the strategic decisions that stem from it. While spreadsheets and databases can store enormous quantities of information, the human mind is not well-suited to quickly interpret columns of figures without visual context. In operations, where managers are often balancing tight production schedules, service commitments, and fluctuating demand, visualization offers clarity. Well-crafted visuals allow decision-makers to quickly identify patterns, detect anomalies, and track performance over time. A line graph showing a steady increase in on-time deliveries, for instance, communicates a story far more intuitively than a dense table of dates and percentages.

In both manufacturing and service settings, visualization transforms complex, multidimensional datasets into actionable insights. A factory producing custom furniture might use a dashboard to track real-time production efficiency by workstation, while a hotel could use heat maps to show which times of year experience the highest booking rates. These representations allow managers to quickly see where

bottlenecks are forming or where additional resources could be allocated. They also enable cross-functional teams to speak the same language—visual evidence that supports operational decisions.

At Pemi Coffee Roasters, visualization is a daily practice rather than an occasional reporting tool. The company relies on dashboards that integrate order fulfillment data, roasting schedules, delivery timelines, and quality control metrics. For example, a bar chart comparing weekly production volumes across different coffee blends helps Pemi quickly adjust roasting priorities when a seasonal blend begins outselling a year-round product. In this way, visual tools not only display performance but also shape the day-to-day adjustments that keep operations smooth.

Shape

7.2 Principles of Effective Visualization

Effective visualizations in operations follow a few core principles that ensure the message is clear, accurate, and actionable. First, the visualization must be aligned with its purpose. A chart meant to monitor production quality should highlight trends and anomalies, not overwhelm the viewer with unnecessary detail. Second, clarity is paramount. Managers should be able to interpret the message within seconds, without having to decode complex legends or navigate ambiguous scales. Third, accuracy in representation prevents misleading conclusions. Poorly chosen chart types or manipulated scales can distort the reality of the data, potentially leading to flawed decisions.

In the manufacturing world, this might mean using a line chart to track machine downtime over the last year, ensuring

the vertical axis starts at zero to avoid exaggerating small differences. In service industries, a similar principle applies—a call center might use a stacked bar chart to show call resolution rates by agent, ensuring each segment is proportionally represented. Clarity and honesty in data presentation build trust among stakeholders, which is critical when operational adjustments may affect jobs, budgets, or customer satisfaction.

Pemi Coffee Roasters applies these principles in every visual it produces. The company's weekly operations dashboard is carefully structured so that the most important indicators—roasting efficiency, order fulfillment rates, and customer satisfaction—are positioned at the top in clear, color-coded charts. Supporting metrics, such as energy use and packaging waste, are placed below in simpler visuals that require less frequent monitoring. By prioritizing the hierarchy of visuals, Pemi ensures that decision-makers focus on the most urgent performance signals before drilling down into supplementary data.

Shape

7.3 Choosing the Right Visualization for the Data

One of the most common mistakes in operations analytics is selecting a visualization type that does not match the nature of the data or the question being asked. Choosing the right format requires understanding the data's structure—whether it is categorical, numerical, time-based, or a combination—and the audience's needs. A time series of production output should typically be displayed as a line chart, while a categorical comparison, such as order accuracy by warehouse, may be better suited to a bar chart. Scatter plots are excellent for

identifying relationships, such as the connection between staff hours worked and the number of customer complaints.

In manufacturing, capacity utilization over time often benefits from a simple line or area chart that shows peaks and valleys. Service operations might use Gantt charts to visualize project timelines or flowcharts to map out process bottlenecks. Choosing the right visualization type is not just about aesthetics—it is about ensuring the visual answers the operational question as directly as possible.

At Pemi Coffee Roasters, this principle comes to life in the way the team tracks delivery performance. When the company expanded to new distribution hubs, the operations manager chose a map-based visualization to show delivery times across regions. This choice was deliberate—maps made it instantly clear that certain rural areas were consistently experiencing longer delivery times, prompting a review of routing strategies. Without the correct visualization type, these patterns might have been buried in a spreadsheet.

[Placeholder for Visual: Example comparison of delivery times shown as a map versus a table, demonstrating clarity gained from the visual representation.]

Shape

7.4 Role of Dashboards in Operational Decision-Making

Dashboards bring multiple visualizations together into a single, coherent view of operations. They allow managers to see the entire operational landscape at a glance while still enabling them to drill down into specific areas when needed. A well-designed dashboard is not merely a collection of charts—it is an interactive tool that prioritizes relevant metrics and updates in near real-time. In manufacturing, dashboards can display live

machine performance data, quality checks, and supply chain status all in one place. For service industries, they can combine customer satisfaction metrics, staffing levels, and transaction volumes.

The power of dashboards lies in their ability to unify disparate datasets into a single decision-making environment. Instead of jumping between reports in different systems, managers can view the latest operational status in one screen, reducing delays in response time. This real-time accessibility is particularly critical when small shifts in demand or production can have significant downstream effects.

For Pemi Coffee Roasters, the operations dashboard is the nerve center of decision-making. It integrates data from roasting equipment, order management systems, and shipping carriers into a unified platform. During peak holiday seasons, managers can monitor whether roasting is keeping pace with incoming orders, whether packaging is falling behind, and whether delivery schedules remain on track. This integrated view allows the company to anticipate bottlenecks before they impact customers, demonstrating how effective dashboards can serve as a competitive advantage in operations.

7.5 Designing for Interpretation & Accessibility

Good visuals help people see the answer quickly and correctly. In operations, that means emphasizing the signal, minimizing friction, and making the display usable for everyone who needs it.

Purpose and hierarchy

- Lead with the point: use a title that states the takeaway ("On-time delivery below 95% for second week").

- Put the most important KPI(s) top-left; supporting charts flow right/down. Keep 3–5 primary visuals per screen.

Encodings and chart choices

- Favor position and length (bars/lines) over area or angle (pies, bubbles). Avoid 3D and heavy gradients.
- Bars: start at zero; sort bars by value or business order.
- Lines: consistent time intervals; don't exaggerate with broken axes. Use small multiples instead of dual axes.
- Part-to-whole: consider 100% stacked bars for composition over time; avoid many thin slices.
- Variation: show distributions with box/violin/histogram rather than averages alone.

Labels, scales, and annotation

- Label units in titles or axes (orders/day, % on-time). Use direct labels on key lines instead of crowded legends.
- Add reference lines/bands for targets, SLAs, or control limits; annotate notable events (promo, outage, weather).

Color and emphasis

- Use a restrained palette; highlight exceptions with one accent color.
- Ensure color-blind safety (avoid red/green alone; pair color with line style or marker shape).
- Maintain contrast and legible font sizes (≥ 11–12 pt on screen).

Interactivity and accessibility

- Keep filters/slicers to the few a manager actually uses; set sensible defaults.
- If the platform supports it, add alt text, keyboard focus order, and tooltip summaries.
- Test on the device people will use (laptop, projector, phone) and in print/export.

For Pemi Coffee Roasters, this means a top row with roasting throughput, on-time delivery, and defect rate (each with target lines), followed by compact small multiples by blend or region, with a single date slicer and one filter for channel (wholesale/retail/online).

7.6 From Visuals to Decisions

A chart is only useful if it changes what we do. Move from display to action with a simple, repeatable flow.

Question → evidence → implication → decision

- Question: what operational decision are we informing (staffing, routing, roasting schedule)?
- Evidence: 1–2 visuals that directly answer it (trend + breakdown).
- Implication: what the pattern means for service, cost, or quality.
- Decision: the action, owner, and when it will happen.

Thresholds and triggers

- Define targets and action rules before looking at the data (e.g., on-time delivery target 95%).
- Example triggers Pemi can use:
  - On-time < 95% for 2 of 3 days → review carrier mix and packing shift coverage.
  - Defect rate > 2% for 3 consecutive shifts → halt-and-fix on affected line.
  - Backorders > 30 units for any blend for 2 days → advance roasting of that blend by one slot.

Cadence and accountability

- Use the dashboard in a short daily huddle: what's off target, why, what will we do today.
- Assign an owner and due date for each action; track follow-up on the same screen (simple "open/closed" status).

One-page memo template (attach to the dashboard or export)

- Decision statement (one line): "Increase Friday roasting by 15% for two weeks."
- Evidence (one chart, one sentence): "Orders spike Fri/Sat; last two Fridays missed SLA."
- Action, owner, deadline: "Roastery lead adds 1 PM shift; starts this week."
- Follow-up metric: "On-time delivery ≥ 95% next two Fridays; review Monday."

For Pemi Coffee Roasters, a weekly review might focus on the delivery map and the on-time trend, confirm triggers, and log a single, concrete action (reroute two rural zones; add one afternoon packer on Thursdays during the fall).

7.7 Dashboard Build Lab (Power BI, Pemi Coffee Roasters)

This lab turns the ideas in the chapter into a single, decision-ready dashboard. The purpose is simple: give a manager one page they can glance at in a morning huddle and decide what to do today. You will use the provided Pemi Coffee Roasters datasets (Orders, Production, Quality, Calendar). The Calendar table should control time across the page; Orders, Production, and Quality relate to it on Date.

Begin with the data model. Load the four tables, open Model view, and relate each operational table to Calendar[Date] (many-to-one, single direction). Mark Calendar as the date table so relative date filtering and time intelligence behave predictably. This structure keeps the dashboard stable as you add visuals.

Create three metrics that appear every week: throughput, on-time delivery, and defect rate. In Power BI you can do this without writing formulas by using built-in aggregation and Quick measures. For throughput, place a Card visual and drop in Production units (it will sum automatically). For on-time delivery, use Quick measure → Division: numerator is the count/sum of on-time orders, denominator is total orders; format as a percentage. For defect rate, do the same with defects over inspected units. Rename the quick measures with clear business names (On-Time %, Defect Rate) so the Fields pane reads like a dashboard spec, not a data dictionary.

Lay out the page with intent. The top row holds three Card

visuals—Throughput Units, On-Time %, and Defect Rate. Use conditional formatting on the data labels (green when at/above target, red when below). In the Analytics pane for the trend and bar charts, add constant lines for targets (for example, On-Time target at 95%, Defect target at 2%). Titles should read like conclusions ("On-time below target in week 42"), not captions.

Explain the pattern beneath the KPIs. Add a Line chart showing daily On-Time % for the last four weeks. Add a Clustered bar for throughput by blend (or region) and another visual—small multiples or a bar—showing defect rate by line or shift. Keep the color palette restrained and use one accent color for exceptions. Place two filters on the page: a Relative date slicer (default to last 28 days) and one business slicer (Blend or Region). Test the page at 100% zoom and in Mobile layout; if labels crowd, simplify rather than shrinking text.

Finish by making the dashboard actionable. Add a small textbox titled Action with three prompts—Decision, Owner, Date—so the team records one commitment in the review. Add a lightweight tooltip page (optional): when users hover over a trend point, they see date, region, On-Time %, and a note for last promo day. These touches convert a report into a decision surface.

7.8 Applied Assignment: Weekly Operations Dashboard & Memo (Power BI)

Deliver two items: a one-page Power BI dashboard and a one-page memo. Use the provided Pemi datasets. The dashboard should make it immediately clear whether throughput, on-time delivery, and defects are on target this month and this

week; the memo should turn that evidence into one concrete action.

Build the dashboard first. Confirm relationships to the Calendar table in Model view. Create the three top-row indicators using built-in aggregation and Quick measures (no typing formulas). Add target constant lines in the Analytics pane and conditional formatting on KPI cards so status is visible at a glance. Use a Relative date slicer set to the last 28 days and one business slicer (Blend or Region). Place a daily On-Time % line, throughput by blend bars, and defect rate by line or shift. Titles should state the takeaway; axes show units and percent; legends stay minimal.

Write the memo second. In five or six sentences:

- State the operating question for the week (for example, “Are we meeting targets?”).
- Cite each visual with one sentence of evidence (trend, breakdown, exception).
- Explain the operational implication (service, cost, or quality).
- Name a specific action with owner and timing. If off target, include one plausible cause and one mitigation.

Submission: a single-page **.pbix** and a one-page memo (PDF/DOCX). Grading emphasizes (1) analytic correctness and honest design, (2) clarity and usefulness of the layout and slicers, and (3) actionability—the memo names a decision someone will take this week.

Key Takeaways

- Data visualization translates complex operational data into clear, actionable insights.
- Effective visuals are purpose-driven, simple, accurate, and aligned with decision needs.
- Choosing the right chart type depends on the data structure and the operational question.
- Dashboards unify multiple metrics into a single, decision-ready environment.
- Good design emphasizes hierarchy, clarity, accessibility, and interactivity.
- Visuals must drive decisions by linking evidence to actions with clear thresholds and accountability.
- Tools like Power BI and Tableau make visualization accessible for both manufacturing and service operations.

Chapter Seven References

Gartner. (2023). *Business intelligence (BI): An overview.* Gartner. https://www.gartner.com/en/information-technology/glossary/business-intelligence-bi

Microsoft. (n.d.). *What is Power BI?* Microsoft. https://learn.microsoft.com/en-us/power-bi/fundamentals/power-bi-overview

OpenStax. (2019). *Principles of management*. OpenStax, Rice

University. https://openstax.org/books/principles-management/pages/6-introduction

Tableau. (n.d.). *Business intelligence and analytics: What it is and why it matters*. Tableau. https://www.tableau.com/learn/articles/business-intelligence

Tufte, E. R. (2001). *The visual display of quantitative information* (2nd ed.). Graphics Press.

8. Advanced Performance Analysis



Learning Objectives

By the end of this chapter, students will be able to:

- Interpret a throughput and on-time trend against explicit targets and call out meaningful patterns.
- Compare periods (week-over-week or month-

over-month) to avoid mistaking seasonality or day-of-week effects for true change.

- Build and read a Pareto view that identifies the few categories contributing most to defects or late units.
- Translate findings into one specific operational action (what, where, when, who) for the next review.

This chapter sharpens the dashboard you built in Chapter 7 into a decision tool. The aim isn't more charts; it's better judgment. We'll read the *pulse* of throughput and service against targets, compare performance across periods so we don't chase noise, and use a Pareto to focus on the few issues that cause most of the pain. You'll see how to frame a finding ("what changed, by how much, and where?"), and how to turn it into a concrete action for the coming week. Everything runs in Power BI using the datasets you already loaded.

8.1 What "advanced" means here

Advanced performance work is not about more sophisticated formulas; it's about separating signal from noise so managers intervene where it matters. In operations, random day-to-day bumps are normal. We care about *runs* (several days below target), *shifts* (a new level), and *where* the misses concentrate (by region, route, line, or blend). The rest of this chapter teaches you to see those patterns quickly and defend your interpretation with simple, honest visuals.

8.2 The pulse: throughput and on-time vs. target

Start with the pulse: two compact visuals that answer “Are we on pace?” and “Are we meeting service?” Place On-Time % as a daily line filtered to the last 28 days, and a companion visual for Throughput Units over the same window. Add the target as a constant line (Analytics pane) so the eye knows immediately whether we’re where we should be.

How to read the pulse:

- Look for *runs* below the target (e.g., two or more consecutive days). That’s an operational event, not random noise.
- Scan for day-of-week rhythm. If Fridays and Saturdays dip, that’s a capacity or sequencing issue during a predictable peak.
- Pair the story: if throughput spikes on promo days while on-time drops, the likely cause is load, not a carrier failure.

Make the visuals carry the story. Titles state the takeaway (“On-time below 95% for two of the last five days”), axes include units/percent, and one short annotation calls out an exception (“Promo 8/15”). If labels crowd, reduce detail rather than shrinking text—managers should be able to read this on a laptop or projector without squinting.

(Clicks only: add constant lines from Analytics; use a Relative Date slicer set to “Last 28 days.”)

8.3 Rolling comparisons: week-over-week and month-over-month

A pulse can look alarming until you compare it with the prior period. Rolling comparisons keep us from overreacting to seasonality.

Use week-over-week for daily dashboards and month-over-month for monthly reviews. In Power BI, you can duplicate the pulse visuals and switch the page's date range to the *previous* period, or add a paired card showing the change (e.g., On-Time % Δ vs last week) using Quick Measures → Time intelligence (percentage difference from previous period). The important part is interpretation: are we trending up or down relative to a comparable window, and is that movement operationally meaningful?

Reading the comparison:

- If On-Time % is down vs last week and the deficit clusters on the same weekdays, the fix is likely staffing or sequencing rather than a random blip.
- If throughput is higher but service held steady, celebrate: that's evidence your process absorbed a surge.
- If both worsened, link back to the slice that explains it (region, blend, route, or line).

Keep the display simple: one Δ card per KPI, green when better than last period, red when worse. You're teaching managers to expect context, not just a snapshot.

(Clicks only: Quick Measure → Time intelligence → "Month over month change" or "Previous period" requires marking Calendar as the Date table, which you already did in Ch. 7.)

8.4 Pareto: focus on the biggest few

Once you know performance moved, the next question is *where to intervene*. The Pareto principle says a small number of categories often account for most of the problem. Build a sorted bar chart (descending) of your chosen driver—defect type, region, route, or line—using the adverse metric (Defects, or Late Units). Add a Top N filter (e.g., Top 5) so the eye lands on the levers first. If you'd like to emphasize proportion, show data labels as percentages of the total.

How to read it:

- Name the concentration: "Two routes account for 57% of late units this month."
- Ask for plausibility: does this match what supervisors see on the floor or in the field?
- Propose the smallest action that could move the needle (reroute, add a shift, recheck a line).

The aim isn't a perfect 80/20 curve; it's a ranked list that credibly points to where action will help most.

(Clicks only: sort the bar chart by value descending; Visual level filters → Top N; add data labels showing percent of grand total.)

8.5 Stability check: simple rules without full SPC

You don't need a full control chart to avoid knee-jerk reactions. Two simple heuristics catch most meaningful shifts:

1. Two-day rule: two or more consecutive days below target merits a look.

2. Four-of-five rule: if four of five recent days are below target—even if not consecutive—flag it.

Use these as *review triggers*, not conclusions. Triggers send you back to the slice or process step that explains the miss. On the page, a small text box labeled "Triggers this week" lists which rule fired and on which dates. That makes your thresholding explicit and repeatable.

(Clicks only: you can record triggers manually in a text box or use a small table visual filtered to the last 7 days where On-Time % < target.)

8.6 Applied assignment — Performance Pulse (Power BI)

What you submit (lightweight):

- One Power BI page that extends your Chapter 7 dashboard with:

    1. a 28-day On-Time % pulse (target line, one annotation),

    2. a Throughput Units pulse aligned to the same dates, and

    3. a Pareto (sorted bars with Top N) for the driver you choose (defect type, region, route, or line).
    Add two small Δ cards (week-over-week or month-over-month) for On-Time % and Throughput.

- One paragraph memo (5–6 sentences) that states the question, summarizes the pulse and the comparison,

names the top contributor from the Pareto, and commits to one action with owner and timing.

How it's graded:

- The page makes the status clear (targets visible, Δ cards readable, labels meaningful).
- The Pareto is sorted and focused (Top N).
- The memo proposes a realistic action for the next week.

Key Takeaways

- Focus on separating signal from noise rather than reacting to daily variation.
- Pulse visuals track throughput and on-time performance against targets.
- Rolling comparisons (week-over-week, month-over-month) add essential context.
- Pareto charts reveal the small number of drivers causing most issues.
- Simple stability rules (two-day, four-of-five) flag meaningful shifts.
- Dashboards are decision tools when visuals link directly to actions.

Chapter 8 References

American Society for Quality. (n.d.). *What is a Pareto chart? Analysis & diagram*. ASQ. https://asq.org/quality-resources/pareto

Microsoft. (2024, March 15). *Use quick measures for common and powerful calculations*. Microsoft Learn. https://learn.microsoft.com/en-us/power-bi/transform-model/desktop-quick-measures

Microsoft. (n.d.). *What is Power BI?* Microsoft. https://learn.microsoft.com/en-us/power-bi/fundamentals/power-bi-overview

OpenStax. (2019). *Principles of management*. OpenStax, Rice University. https://openstax.org/books/principles-management/pages/6-introduction

SECTION IV

OPERATIONS ANALYTICS USING R

9. Introduction to Statistical Analysis for Operations



Learning Objectives

By the end of this chapter, students will be able to:

1. Summarize operational data (mean, median, standard deviation) in Excel and R.

2. Choose and apply Normal, Binomial, and Poisson models; compute probabilities/quantiles.

3. Run a two-sample t-test; report the difference in means, confidence interval, and p-value.

4. Analyze association with correlation and simple linear regression; interpret slope and R^2.

5. Perform quick checks for outliers, variance, and independence; note limits of the analysis.

6. Convert results into a concise operations recommendation.

9.1 Basic Statistical Terms (Operations Context)

In earlier chapters you organized data, improved quality, and visualized performance. Here we formalize those summaries with a few core statistics you already know from your prerequisite course, focusing on how to apply and interpret them in operations. We'll compute each measure in Excel and mirror it in R, using two *Pemi Coffee Roasters* datasets: Dataset A (fulfillment times, days) and Dataset B (items per order and fulfillment time).

Mean and Median

The mean summarizes average performance; the median is the middle value and is less sensitive to a few slow orders.

- **Excel** (Dataset A in B2:B25)
  =AVERAGE(B2:B25) → 5.875
  =MEDIAN(B2:B25) → 5

- **R**

```
dfA <-
read.csv("Chapter9_Dataset_A_FulfillmentTimes.cs
v")
mean(dfA$FulfillmentTimeDays) # [1] 5.875
median(dfA$FulfillmentTimeDays) # [1] 5
```

Variability: Range, Variance, Standard Deviation

Consistency matters. The range is max − min. Variance and standard deviation quantify spread around the mean.

- **Excel**
  =MAX(B2:B25)-MIN(B2:B25) → 8
  =VAR.S(B2:B25) → 3.592
  =STDEV.S(B2:B25) → 1.895

- **R**

```
max(dfA$FulfillmentTimeDays) -
min(dfA$FulfillmentTimeDays) # [1] 8
var(dfA$FulfillmentTimeDays) # [1] 3.592389
sd(dfA$FulfillmentTimeDays) # [1]
1.895360
```

Association (Preview): Correlation

Correlation indicates whether two variables move together—for example, whether larger orders take longer to fulfill.

- **Excel** (Dataset B: items in `B2:B31`, time in `C2:C31`)
  `=CORREL(B2:B31, C2:C31) → 0.991`
- **R**

```
dfB <-
read.csv("Chapter9_Dataset_B_OrderSize_vs_Fulfil
lment.csv")
cor(dfB$ItemsInOrder, dfB$FulfillmentTimeDays) #
[1] 0.9911393
```

Applied Example: Fulfillment Times at Pemi (Dataset A)**Excel**

Mean: `=AVERAGE(B2:B25) → 5.875`
Median: `=MEDIAN(B2:B25) → 5`
Std. dev.: `=STDEV.S(B2:B25) → 1.895`
Optional histogram: Data → Data Analysis → Histogram
R

```
dfA <- read.csv("Chapter9_Dataset_A_Fulfillment
mean(dfA$FulfillmentTimeDays) # [1] 5.875
median(dfA$FulfillmentTimeDays) # [1] 5
sd(dfA$FulfillmentTimeDays) # [1] 1.895360
hist(dfA$FulfillmentTimeDays)
```

Interpretation

The mean exceeds the median, indicating a right tail (a few slow orders). A standard deviation near 1.9 suggests noticeable variability. Next steps: isolate slow orders and investigate causes (capacity, supplier delays, address issues).

9.3 Hypothesis testing: two-sample t-test (small vs. large orders)

Question: do larger orders take longer to fulfill?

Define two groups in Dataset B:

- Small orders: items ≤ 3
- Large orders: items > 3

Set up in **Excel**: add a helper column `size_group` with `=IF(B2<=3,"Small","Large")`

Compare mean fulfillment time by group using Welch's two-sample t-test (unequal variances).

Excel

1. Compute group means and standard deviations with `=AVERAGEIFS()` and `=STDEV.S(IF(...))` or with filters/pivots.
2. Data → Data Analysis → t-Test: Two-Sample Assuming Unequal Variances.
    - Variable 1 Range: times for Small
    - Variable 2 Range: times for Large
    - Hypothesized mean difference: 0

- Output to a new range

R

```
dfB <-
read.csv("Chapter9_Dataset_B_OrderSize_vs_Fulfillmen
t.csv")
dfB$size_group <- ifelse(dfB$ItemsInOrder <= 3,
"Small", "Large")
t.test(FulfillmentTimeDays ~ size_group, data = dfB)
# Welch by default
```

What you should see with the provided data:

- Difference in means ≈ 1.54 days (Large − Small)
- 95% CI ≈ [1.22, 1.86]
- $p < 0.001$

Interpretation: in this sample, larger orders take longer by about a day and a half on average. The effect is practically and statistically meaningful.

Transition: correlation and a simple linear model quantify the items → time relationship directly.

9.4 Correlation and simple linear regression

Use correlation to summarize association and a simple linear regression to estimate how much time increases per added item.

Excel

- Correlation: `=CORREL(B2:B31, C2:C31)`

- Scatterplot: Insert → Scatter. Add a trendline → display equation and R^2.
- Functions: `=SLOPE(C2:C31, B2:B31)`, `=INTERCEPT(C2:C31, B2:B31)`, `=RSQ(C2:C31, B2:B31)`

R

```
cor(dfB$ItemsInOrder,
dfB$FulfillmentTimeDays) # ≈ 0.991

m <- lm(FulfillmentTimeDays ~
ItemsInOrder, data = dfB)

summary(m)
```

With the provided data, you should see results close to:

- Slope ≈ 0.430 days per additional item
- Intercept ≈ 4.232 days
- R^2 ≈ 0.982

Interpretation: each additional item adds roughly 0.43 days on average. The fit is strong in this sample. Do not extrapolate beyond the observed range without a check.

Transition: before relying on these results, confirm basic assumptions and scan for obvious issues.

9.5 Assumptions and quick diagnostics

Hypothesis tests and linear models rest on simple conditions. A few quick checks go a long way.

t-test (Small vs. Large)

- Independence: orders should be independent observations. Segment by shift/promo if needed.
- Shape: with n ≈ 15 per group, the t-test is robust, but heavy skew/outliers can distort results. Use group histograms or boxplots.
- Variances: Welch's test (used above) does not assume equal variances. If you suspect very different spreads, keep using Welch.

Excel (practical checks)

- Build small and large group histograms or boxplots (via charts).
- Compute group variances with `=VAR.S(...)`.
- If you ran the regression, add predicted values `=intercept + slope*Items` and create a residuals column `=Actual - Predicted`. Plot residuals vs. predicted; look for no pattern.

R (quick commands)

```
boxplot(FulfillmentTimeDays ~
size_group, data = dfB)

by(dfB$FulfillmentTimeDays,
dfB$size_group, sd) # group spreads

# Regression diagnostics

plot(m, which = 1) # Residuals vs
Fitted
```

```
plot(m, which = 2) # Normal Q-Q
```

Regression

- Linearity: residuals vs. fitted should look like noise, not a curve.
- Constant variance: avoid funnel shapes; consider a variance-stabilizing approach or segmenting if needed.
- Influence: in small samples, a single unusual order can sway the fit; inspect points with large residuals.

Transition: with checks in place, synthesize findings in a brief memo that a manager can act on.

9.6 Applied assignment: Pemi memo

Goal: use the chapter workflow to answer a specific operations question and communicate a clear recommendation.

Datasets

- Chapter9_Dataset_A_FulfillmentTimes.csv
- Chapter9_Dataset_B_OrderSize_vs_Fulfillment.csv

Tasks

1. Excel
    - Dataset A: mean, median, standard deviation; histogram of fulfillment times.
    - Dataset B: correlation (items vs. time); two-sample t-test (Small vs. Large).

- Optional: scatterplot with trendline, show equation and R^2.

2. R

   - Replicate Excel results with `mean()`, `median()`, `sd()`, `hist()`, `cor()`, `t.test()`, and `lm()`.
   - Confirm that Excel and R agree within rounding.

3. Managerial memo (1–2 paragraphs)

   - Is the difference in mean time by size group meaningful?
   - About how much time does each additional item add?
   - One concrete action you recommend (e.g., staffing during peaks, batching rules, target setting).

Deliverables

- One Excel file with clearly labeled sheets (A_Summary, B_ttest, B_Regression).
- One R script (.R) or R Markdown (.Rmd) with commands and printed outputs.
- A PDF or DOCX memo.

Key Takeaways

- Mean and median summarize central tendency; the median is less sensitive to outliers.
- Variance and standard deviation measure variability around the mean, indicating consistency in performance.
- Correlation quantifies the strength of association between two variables.
- A two-sample t-test compares group means and can show whether differences are statistically significant.
- Linear regression estimates the relationship between predictors (e.g., order size) and outcomes (e.g., fulfillment time).
- Diagnostic checks (shape, variance, residuals) ensure statistical results are reliable for decision-making.
- Excel and R can be used together to calculate, visualize, and cross-verify these statistics in an operations context.

Chapter Nine References

Microsoft. (n.d.). *Statistical functions (reference)*. Microsoft. https://support.microsoft.com/en-us/excel

R Core Team. (2023). *R: A language and environment for*

statistical computing. R Foundation for Statistical Computing. https://www.r-project.org

Wickham, H., & Grolemund, G. (2017). *R for data science: Import, tidy, transform, visualize, and model data*. O'Reilly Media. https://r4ds.had.co.nz

10. Forecasting and Predictive Analytics



Learning Objectives

By the end of this chapter, students will:

1. Prepare operational time series in R and split into training/test (and rolling-origin).
2. Build and compare baseline forecasts (naïve, seasonal naïve, moving average).

3. Fit exponential smoothing (ETS) and ARIMA models; interpret components and diagnostics.

4. Evaluate accuracy with MAE, RMSE, and MAPE; select a model for deployment.

5. Produce a short operations recommendation grounded in forecast results.

10.1 Framing and data prep (R)

Forecasting is about estimating future values in a time series — here, daily orders. Predictive analytics goes a step further by adding outside drivers, such as promotions or weather, that may explain or improve those forecasts.

Before modeling, clarify the business question: *"How many orders should we expect in the next 28 days, so we can plan staffing and inventory?"* Then:

- Choose the right frequency (daily orders with a weekly pattern).
- Hold out the last 28 days of data as a test window. This gives us an honest way to judge model accuracy.
- Prepare candidate external drivers (e.g., promo days, average daily temperature) for models that support them.

R

```
library(readr); library(dplyr); library(lubridate);
library(ggplot2); library(forecast)
```

```
  df <- read_csv("Chapter10_DailyOrders_PlymouthNH.csv")
%>% mutate(date = as.Date(date))
  # hold out the last 28 days
h <- 28
df_train <- df |> slice_head(n = n() – h)
df_test <- df |> slice_tail(n = h)
  # weekly seasonality -> frequency = 7
y_train <- ts(df_train$orders, frequency = 7)
y_full <- ts(df$orders, frequency = 7)
  # candidate external drivers (optional in some models)
x_train <- as.matrix(df_train[, c("promo_flag","avg_temp_f")])
x_test <- as.matrix(df_test[, c("promo_flag","avg_temp_f")])
  ggplot(df, aes(date, orders)) + geom_line()
```

10.2 Baselines: naïve, seasonal naïve, moving average

Baselines are simple "no-frills" forecasts. They set the minimum that more complex models must beat.

- Naïve: tomorrow equals today. Good when series wander slowly.
- Seasonal naïve: next Monday equals last Monday. Strong when day-of-week patterns dominate.
- Moving average: smooths noise to show the local level but does not project seasonality by itself.

R

```
fc_naive <- naive(y_train, h = h)
fc_snaive <- snaive(y_train, h = h) # respects
weekly patternautoplot(y_full) +
autolayer(fc_snaive, series = "Seasonal naive",
```

PI = FALSE) +
autolayer(fc_naive, series = “Naive”, PI = FALSE)

10.3 Exponential smoothing (ETS)

ETS breaks the series into level, trend, and seasonality. It adapts quickly when things change and is a common default for operational data. R can automatically pick the best combination of additive or multiplicative forms. After fitting, check residuals to confirm no leftover pattern.

R

```
fit_ets <- ets(y_train) # auto-selects additive/
multiplicative forms
fc_ets <- forecast(fit_ets, h = h)autoplot(fc_ets)
checkresiduals(fit_ets) # want residuals ~
white noise
```

10.4 ARIMA and ARIMAX (with temperature and promos)

ARIMA uses past values and past errors to explain the series. Seasonal ARIMA adds repeating patterns (like weekly cycles).

- **ARIMA**: relies only on the series’ own history.
- **ARIMAX**: adds external predictors (e.g., promos, temperature). This can sharpen forecasts when drivers explain part of demand.

Residual checks confirm whether the model has captured the structure, or whether systematic patterns remain.

R

```
# ARIMA without drivers
fit_arima <- auto.arima(y_train, seasonal = TRUE)
fc_arima <- forecast(fit_arima, h = h)
```
ARIMAX with promo + temperature
fit_arimax <- auto.arima(y_train, xreg = x_train, seasonal = TRUE)
fc_arimax <- forecast(fit_arimax, xreg = x_test, h = h)checkresiduals(fit_arima)
checkresiduals(fit_arimax)

Notes: Be careful not to "peek" at the future when preparing drivers (xreg for the test window must be known at forecast time). If residuals still show weekly structure or long runs, revisit the specification.

10.5 Accuracy comparison and rolling-origin

We compare competing models on the same holdout window using:

- MAE (average absolute error): easy to interpret in original units.
- RMSE (root mean square error): penalizes large misses more heavily.
- MAPE (percentage error): intuitive, but avoid when actuals can be near zero.

A single split can be lucky. Rolling-origin cross-validation re-forecasts many times from expanding windows and averages the error, giving a more stable view of performance.

R

```
library(Metrics)
  y_test <- df_test$orders
compare <- dplyr::bind_rows(
tibble::tibble(model="naive",
MAE=mae(y_test, fc_naive$mean),
RMSE=rmse(y_test, fc_naive$mean),
MAPE=Metrics::mape(y_test, fc_naive$mean)),
tibble::tibble(model="snaive",
MAE=mae(y_test, fc_snaive$mean),
RMSE=rmse(y_test, fc_snaive$mean),
MAPE=Metrics::mape(y_test, fc_snaive$mean)),
tibble::tibble(model="ETS",
MAE=mae(y_test, fc_ets$mean),
RMSE=rmse(y_test, fc_ets$mean),
MAPE=Metrics::mape(y_test, fc_ets$mean)),
tibble::tibble(model="ARIMA",
MAE=mae(y_test, fc_arima$mean),
RMSE=rmse(y_test, fc_arima$mean),
MAPE=Metrics::mape(y_test, fc_arima$mean)),
tibble::tibble(model="ARIMAX",
MAE=mae(y_test, fc_arimax$mean),
RMSE=rmse(y_test, fc_arimax$mean),
MAPE=Metrics::mape(y_test, fc_arimax$mean))
) |> arrange(RMSE)
compare
```

Rolling-origin example (7-day horizon):

```
e_snaive <- tsCV(y_full, function(y, h)
forecast(snaive(y), h=h), h = 7)
e_arima <- tsCV(y_full, function(y, h)
forecast(auto.arima(y), h=h), h =
7)sqrt(mean(e_snaive^2, na.rm=TRUE))
sqrt(mean(e_arima^2, na.rm=TRUE))
```

Choose the simplest model that meets your accuracy target and leaves residuals looking like noise.

10.6 Applied example: 28-day staffing plan (R)

Forecasts matter when they drive action. For Pemi Coffee Roasters, we use the chosen model (e.g., ETS or ARIMAX) to forecast 28 days of orders. Orders are converted into staff requirements using a simple rule (one staff per 40 orders).

The deliverable is a short memo with:

- The forecast horizon (28 days).
- The chosen model and why it was selected.
- A staffing table by day.
- Any flagged risks (e.g., promotions or heat waves) and one suggested mitigation.

R

```
best <- fit_arimax # or fit_ets / fit_arima based on compare + residuals
fc <- forecast(best, xreg = x_test, h = h)staff_per_day <- ceiling(as.numeric(fc$mean) / 40)
out <- tibble::tibble(
date = tail(df$date, 1) + 1:h,
forecast_orders = as.numeric(fc$mean),
staff_needed = staff_per_day
)
print(out)
```

Memo pointers: state the horizon, the chosen model and why, attach the daily staffing table, flag any operational risks (e.g., a promo cluster) and one mitigation.

Key Takeaways

- Forecasting projects future demand; predictive analytics adds drivers like promos or weather.
- Define the question, set frequency, and hold out data for testing.
- Baselines (naïve, seasonal naïve, moving average) are the benchmark.
- ETS, ARIMA, and ARIMAX capture trend, seasonality, and external drivers.
- Accuracy checks use MAE, RMSE, and MAPE, plus rolling-origin validation.
- Forecasts should guide concrete actions, such as staffing plans.

Chapter 10 References

Microsoft. (n.d.). *Forecasting functions in Excel*. Microsoft. https://support.microsoft.com/excel

R Core Team. (2023). *R: A language and environment for*

statistical computing. R Foundation for Statistical Computing. https://www.r-project.org

Hyndman, R. J., & Athanasopoulos, G. (2021). *Forecasting: Principles and practice* (3rd ed.). OTexts. https://otexts.com/fpp3/

Wickham, H., & Grolemund, G. (2017). *R for data science: Import, tidy, transform, visualize, and model data*. O'Reilly Media. https://r4ds.had.co.nz

11. Scenario Modeling and Decision Support



11.1 Introduction to Scenario Modeling

Operations rarely happen in steady conditions. Customer demand can spike or fall, suppliers can delay shipments, and costs can shift with little warning. Scenario modeling gives managers a way to prepare by asking “what if” and then seeing how those answers play out.

At Pemi Coffee Roasters, leadership often wonders what will happen if wholesale orders grow faster than expected or if café sales soften. A simple model in Excel lets them adjust demand up or down and watch the effect on roasting schedules, staffing, and delivery. A 15 percent jump in wholesale demand

may show the need for an extra roasting shift and added drivers. A 10 percent drop highlights unused capacity and rising unit costs. In both cases, the team can talk through trade-offs before the situation actually happens.

The goal is not to predict the future exactly but to understand how different paths would affect the business. Scenario modeling makes those conversations concrete. It shows risks and opportunities clearly, and it helps leaders explain why one course of action makes more sense than another.

11.2 Building Scenario Models

Once the concept of scenarios is introduced, the next step is to design models that can capture them. A scenario model is essentially a framework that allows managers to ask “what if?” and receive structured, data-based answers. At Pemi Coffee Roasters, this might involve using Excel to estimate how changes in wholesale demand affect both café operations and warehouse capacity. By building a flexible model, leaders can test not only their expectations but also unexpected shifts that may arise in the market.

The process usually begins with identifying the variables that matter most. For a coffee business, these could include the volume of beans purchased, café sales, labor hours, and delivery costs. Each variable is then linked through formulas so that a change in one input automatically adjusts the outputs. If bean prices rise by 15 percent, the model recalculates margins; if wholesale orders fall, the impact on inventory and staffing becomes clear. This structure allows managers to test different scenarios without having to rebuild their calculations each time.

Equally important is the choice of assumptions. Models are only as reliable as the estimates they are built on. Using

historical sales patterns, industry benchmarks, or supplier quotes provides a stronger foundation than rough guesses. Students should practice challenging their assumptions by asking how sensitive their results are to small changes. A forecast that collapses when assumptions shift slightly is not robust enough for real-world use.

Scenario modeling is most effective when the outcomes are presented clearly. Instead of a long spreadsheet full of numbers, a simple chart or table that shows best-case, middle-ground, and worst-case results side by side makes the trade-offs obvious. Managers can quickly see how a decision might play out if demand grows, levels off, or declines, and they can discuss actions for each situation. In practice, scenario models serve as both a decision-making aid and a communication tool, helping leadership teams compare alternatives and align on strategies under uncertainty.

11.3 Building What-If Scenarios in Excel

Operations often involve decisions made with incomplete information. Managers may not know if costs will rise, whether customers will buy more or less than expected, or if supplies will arrive on time. What-if analysis in Excel gives them a way to test possible changes before those changes happen in real life.

Scenario Manager is one of the tools that makes this possible. It allows several sets of assumptions to be saved and compared side by side. At Pemi Coffee Roasters, one scenario might assume steady demand, another a sharp increase in sales, and a third a slower market. Switching between these versions shows how each situation affects production schedules, staffing, or inventory levels.

Data Tables are another option. They work well when you want to test one or two variables repeatedly. If Pemi Coffee is considering expanding roasting capacity, a Data Table could

show profits at different combinations of sales volume and cost per pound. This type of grid makes it easier to see where growth is realistic and where expenses outweigh benefits.

Goal Seek takes the process in a different direction. Instead of asking "what happens if," it answers "what input is needed to hit a goal." If the target is a two-day order lead time, Goal Seek can identify the maximum number of orders the system can manage while still meeting that commitment. This provides a clear limit for managers.

Used together, these tools help operations teams think ahead without pretending they can predict the future. The value comes from seeing possible outcomes in advance and planning for how to handle them.

11.4 Scenario Planning with Power BI

Excel is often the first stop for scenario planning, but Power BI changes how those scenarios are shared and understood. Instead of scrolling through spreadsheets, managers can see outcomes directly on a dashboard. Filters let them compare regions, product lines, or time periods in real time, and the charts update as assumptions shift.

At Pemi Coffee Roasters, a Power BI dashboard could show profits under several sales forecasts. Leaders might look at one view where wholesale sales grow quickly, then switch to another where café traffic slows. Having those pictures side by side makes it easier to decide whether to hire more staff, delay a store opening, or adjust pricing.

Another benefit is accessibility. People who are not building the models themselves can still interact with them. A regional manager can click through the dashboard, ask "what if" questions, and see the answers without waiting for a new spreadsheet. This keeps scenario planning active and part of

regular conversations rather than something that happens in the background.

11.5 Using R for Advanced Scenarios

R brings another layer to scenario analysis by allowing more advanced statistical and predictive modeling. Instead of only testing predefined changes, R can uncover relationships in the data and simulate a range of outcomes.

Take customer demand. In Excel, Pemi Coffee Roasters might assume demand rises or falls by 10%. With R, managers could run a regression to see how demand changes with season, weather, or promotions. They could also cluster customers into groups based on ordering habits. These approaches reveal patterns that simple "increase or decrease" assumptions can miss.

R also makes it easier to automate scenario runs. A script can generate dozens of versions—different cost structures, delivery times, or growth rates—without manually adjusting cells. The output can then be linked back into Excel or Power BI for visualization. This keeps the analysis flexible while grounded in data.

11.6 Bringing the Tools Together

Each tool—Excel, Power BI, and R—adds a different strength. Excel provides an easy entry point for exploring scenarios. Power BI turns scenarios into interactive visuals for quick interpretation. R allows for deeper, more technical exploration of the data. Together, they give a fuller picture of what could happen in an uncertain environment.

For Pemi Coffee Roasters, using all three tools means

leadership can explore “what if demand spikes,” see the financial impact clearly, and check the results against patterns in past data. The company moves from reacting to surprises toward preparing for multiple possibilities.

Scenario planning is not about predicting the future with certainty. It is about being ready for several possible futures and choosing the best path when the time comes.

Key Takeaways

- Scenario modeling tests “what if” situations to prepare for uncertainty.
- Models link key variables (demand, labor, costs) so one change updates outcomes.
- Excel tools (Scenario Manager, Data Tables, Goal Seek) make scenarios easy to build.
- Power BI adds interactive, shareable visuals for quick comparison.
- R enables deeper, data-driven simulations and automation.
- Together, Excel, Power BI, and R provide a fuller view of possible futures.

Chapter 11 References

Camm, J. D., Cochran, J. J., Fry, M. J., Ohlmann, J. W., & Anderson, D. R. (2021). *Business analytics* (4th ed.). Cengage Learning.

Delen, D. (2020). Prescriptive analytics: The final frontier for evidence-based management and optimal decision making. *Decision Support Systems, 131,* 113246. https://doi.org/10.1016/j.dss.2020.113246

Hyndman, R. J., & Athanasopoulos, G. (2021). *Forecasting: Principles and practice* (3rd ed.). OTexts. https://otexts.com/fpp3/

Microsoft. (2023). *Introduction to Power BI*. Microsoft Learn. https://learn.microsoft.com/en-us/power-bi/fundamentals/power-bi-overview

OpenStax. (2023). *Introductory business statistics*. OpenStax. https://openstax.org/books/introductory-business-statistics

SECTION V

INTEGRATION & FINAL COMPREHENSIVE PROJECT

12. Integrating Operations and Analytics for Decision-Making



Chapter 12: Integrating Operations and Analytics for Decision-Making

12.1 Introduction

Over the course of this class, we have walked step by step through the world of business operations, learning not just the concepts but also how to use tools to make sense of them. We started with Excel, moved into Power BI, and finally worked with R. Each tool was introduced with a purpose, not just as software, but as a way to help you answer the kinds of questions businesses face every day. By now, you should see how operations strategy, capacity planning, process design, inventory management, and analytics all connect. This final chapter is about pulling those threads together into one clear picture.

12.2 From Concepts to Practice

Operations strategy is not only theoretical. It is about aligning what a company does every day with its broader goals. You saw this when we examined how process design determines efficiency, how capacity shapes customer satisfaction, and how inventory management affects financial stability. At each step, you tested these ideas with data. Whether it was building an Excel model to simulate demand, using Power BI to create a dashboard for decision makers, or running an R script to analyze trends, the point was the same. Operations must be both managed and measured.

12.3 Tools in Context

Excel gave you a way to handle structured problems such as calculations, forecasting, and what if analysis. Power BI introduced visualization and the ability to communicate findings quickly to different audiences. R added another dimension with flexibility for statistical analysis and reproducibility. Taken together, the tools mirror how professionals actually work. Rarely does one tool do everything. Instead, the strength lies in knowing how to choose the right one for the situation and how to integrate them into a workflow.

12.4 The Role of Data in Decisions

One of the themes running through this course is that decisions need to be grounded in evidence. In operations, there is always pressure to move quickly, but rushing without data often leads to bigger problems later. The examples we worked through, such as balancing trailer utilization in shipping or analyzing bottlenecks in a service process, show that better data leads to better questions, and better questions lead to stronger outcomes. The more comfortable you are moving between raw data and strategic choices, the more effective you will be in practice.

12.5 Portfolio as a Story

By now your portfolio should not feel like a collection of disconnected assignments. It should read as a story of how you grew in your ability to apply both concepts and tools. Employers or colleagues who look at it should see evidence of

critical thinking, practical analysis, and communication skills. They should also see that you can handle real problems, not just textbook exercises. This final project will be the anchor of that portfolio, demonstrating your ability to synthesize everything you have learned.

12.6 Final Comprehensive Project

For your final project, you will return to the fictional company we have been analyzing all semester. Your task is to create a complete operations analysis that integrates strategy, process, capacity, and inventory, supported by data.

Deliverables:

1. **Excel Component**. Build a model that addresses one operational decision such as forecasting demand, evaluating capacity, or testing an inventory policy. The model should allow for adjustments so decision makers can see how different choices play out.

2. **Power BI Component**. Create a dashboard that communicates key insights visually. This should include metrics that would matter to leadership, such as efficiency, cost, and customer impact, and it should be designed with clarity in mind.

3. **R Component.** Write and run an R script that provides a deeper level of analysis. This might include regression, clustering, or trend analysis depending on your focus. The code should be annotated so that someone else could follow your steps.

4. **Executive Presentation.** Deliver a professional presentation to the class as if you are addressing the executive team at Pemi Coffee Roasters. This presentation

should highlight your findings, explain their significance, and make clear recommendations. Visual clarity, concise messaging, and executive presence will be emphasized.

5. **Portfolio Submission.** Compile the three tool-based deliverables and the written report into your portfolio. The portfolio is the final product, and it should present your work as a cohesive demonstration of business operations analytics.

Evaluation criteria:

- Integration of concepts with practice.
- Technical execution of Excel, Power BI, and R work.
- Clarity of communication in the dashboard, report, and portfolio.
- Critical thinking that moves beyond surface level answers to examine trade offs, risks, and outcomes.

12.7 Conclusion

This course began with basic building blocks of operations and ended with a complete toolkit for analysis and decision making. The final project is not just an assignment. It is a simulation of the work you will be expected to do in professional settings. The ability to move from data to insight to action is the outcome of the course, and your portfolio is the proof.

Key Takeaways

- Operations strategy connects daily actions to organizational goals.
- Excel, Power BI, and R complement each other in analysis and decision-making.
- Strong decisions are grounded in evidence, not intuition.
- A portfolio should show applied skills, critical thinking, and communication.
- The final project integrates all course concepts into a professional deliverable.

Chapter 12 References

Microsoft. (n.d.). *Excel for Microsoft 365 functions*. Microsoft. https://support.microsoft.com/excel

Microsoft. (n.d.). *Power BI documentation*. Microsoft. https://learn.microsoft.com/power-bi

R Core Team. (2023). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. https://www.r-project.org

Wickham, H., & Grolemund, G. (2017). *R for data science: Import, tidy, transform, visualize, and model data*. O'Reilly Media. https://r4ds.had.co.nz

This is where you can add appendices or other back matter.