

Starbucks Store Operations Analytics

A Harvard-Style Live Teaching Case for AI-Assisted Power BI

“The Missed Opportunity: How 2 Million Customer Voices Predicted Starbucks’ Crisis, and No One Listened”

Course: AI-Assisted Power BI Analytics (MBA)

Module Alignment: Modules 1-7 (anchor case throughout course)

Role: Primary Live Case (students apply skills to real Starbucks data)

Status: Confirmed

Document Overview

This case plan outlines a Harvard-style live teaching case for the AI-Assisted Power BI Analytics course. A Starbucks executive participates in Modules 1 and 7, providing real business context and evaluating student presentations.

Core Thesis: Starbucks collected 2 million customer feedback signals monthly and invested billions in AI infrastructure, yet failed to predict a crisis that customers warned about 18-24 months in advance. The data existed; it sat in organizational silos.

Student Challenge: Teams analyze Starbucks data using Power BI and Copilot, then present recommendations to the executive guest. They grapple with the same question Maya Rodriguez (fictional VP of Store Operations Analytics) faces: *How do you build analytics that empower store managers without creating chaos?*

What We Need from the Executive:

| Option | Data Type | Executive Comfort | Student Experience |
|-----------------------|------------------------------------|-----------------------------|-------------------------------------|
| A: Vetted Data | Known patterns, expected insights | High (no surprises) | Good: structured, predictable |
| B: Raw Data | Unanalyzed, genuine questions | Brave (unexpected insights) | Excellent: authentic discovery |
| C: Public Data | Kaggle/GitHub; no executive needed | N/A (fallback) | Moderate: may be stale or incorrect |

Option B creates the most authentic learning experience. Option C is a fallback if no executive partnership materializes. Students still work with real Starbucks data, but without insider validation.

Additional requirements (for Options A or B):

| Requirement | Description |
|--------------------|---|
| Business Questions | 3-5 questions to guide student analysis |
| Availability | Module 1 (intro) and Module 7 (presentations) |
| NDA | If required, students can sign confidentiality agreements |

What Students Deliver: - Power BI dashboard addressing the executive’s questions - AI-assisted root cause analysis using Key Influencers - Team presentation with actionable recommendations

Course Integration: Dual-Track Learning

This case anchors the **Live Case Track** throughout the 7-week course. Students practice skills on synthetic datasets, then immediately apply them to real Starbucks data.

| Module | Skills Lab (Synthetic Data) | Starbucks Live Case Application |
|--------|---|---|
| 1 | First AI Conversation (CloudRevenue) | Executive Introduction: Guest presents business context, hands off dataset, frames key questions |
| 2 | Prompt-Driven Reports (CloudRevenue) | Generate initial Starbucks report using Copilot; identify key metrics |
| 3 | AI Root Cause Analysis (SupportInsights) | Apply Key Influencers to Starbucks data: “ <i>What drives [metric] at the store level?</i> ” |
| 4 | Preparing Data for Copilot (CloudRevenue) | Assess Starbucks data quality; identify metadata improvements |
| 5 | Executive Dashboard (M365Marketing) | Design Starbucks executive dashboard: “ <i>What does the guest executive need to see?</i> ” |
| 6 | Deploy Your Solution (Any) | Create Starbucks workspace and app; prepare for presentation |
| 7 | (Capstone) | Executive Presentation: Teams present insights and recommendations to Starbucks executive |

Executive Participation

| Session | Role | Activities |
|----------|----------------------|---|
| Module 1 | Introduce & Hand Off | Present business context, provide dataset, frame 3-5 key business questions |
| Module 7 | Receive & Evaluate | Hear team presentations, ask Socratic questions, provide feedback (advisory input to grading) |

Team Formation

Teams of 3-4 students form in Module 2 and work together on Starbucks analysis through Module 7 capstone.

Executive Summary

The Case in One Sentence

Starbucks collected 2 million customer feedback signals per month, had world-class AI infrastructure (Azure ML, IoT, blockchain), and still failed to predict a crisis that customers were warning about 18-24 months in advance. The data sat in organizational silos and never reached the store managers who could act on it.

Critical Insight: Those 2M signals were just the tip of the iceberg. Research shows only 4% of unhappy customers actually complain; the rest silently defect. Starbucks wasn’t hearing 96% of its dissatisfied customers.

The Business Situation

By 2024, Starbucks faced a perfect storm: same-store sales declined for the first time since COVID (-4% in Q2), transaction counts fell 6%, the CEO was replaced, and union organizing accelerated. The “third place” brand promise had eroded as mobile ordering transformed stores from community gathering spots into transaction factories.

The board’s diagnosis: Operational complexity and leadership failure.

The real diagnosis (this case argues): A Voice of the Customer (VoC) and analytics failure. The warning signs were visible in customer feedback 18-24 months before the financial impact, but no one connected the dots.

The Missed Opportunity

| What Customers Said (2022-2023) | What Leadership Did | What They Could Have Done |
|---------------------------------------|-----------------------------------|--|
| “Why is my mobile order never ready?” | Added “order ready” notifications | Capped mobile orders when stores hit capacity |
| “The baristas seem overwhelmed” | Hired seasonal workers | Connected workload data to satisfaction metrics |
| “This place used to feel special” | Launched marketing campaign | Measured “third place experience” quantitatively |
| “I switched to local coffee shops” | Increased discounts | Analyzed store-level drivers of attrition |

The paradox: Starbucks spent billions on AI to personalize the *customer* experience (Azure ML recommendations). They didn’t build the dashboards to help *store managers* see their own customers’ voices and act on them.

The Decision Point

Maya Rodriguez, VP of Store Operations Analytics, must present a recommendation to the new CEO and COO. She has three options:

1. **Centralized Analytics:** Build sophisticated models at corporate, push insights down
2. **Self-Service for Store Managers:** Give everyone Power BI, let them analyze their own data
3. **Hybrid “Governed Self-Service”:** Certified core metrics + local flexibility + AI assistance (Copilot)

The case question: *Which approach would have prevented the crisis if implemented in 2022? Which approach can fix it now? And how does Maya convince skeptical store operations leaders who’ve been burned by “corporate dashboards” before?*

Why This Case Matters for MBA Students

This case is the **anchor experience** for the AI-Assisted Power BI course:

1. **Real Stakes:** Students know their work will be seen by a Starbucks executive
2. **Data Insights Action:** Starbucks had petabytes of data. The crisis happened anyway.
3. **VoC as Leading Indicator:** Customer sentiment predicts business outcomes 18-24 months ahead, *if* you structure your analytics to listen.
4. **Organizational Silos Kill Analytics:** Marketing owned NPS. Operations owned wait times. HR owned employee satisfaction. Nobody owned “store health.”
5. **AI for Whom?:** Companies often build AI for customers but not for frontline workers. This case explores that asymmetry.
6. **Power BI + Copilot Application:** Students design what a “Store Health Dashboard” should look like, connecting CX, operations, and employee data.

Learning Objectives

After this case, students will be able to: - Identify early warning signals in customer feedback data - Design analytics architectures that connect CX to operational decisions - Evaluate trade-offs between centralized governance and self-service agility - Understand how AI/Copilot can bridge skill gaps for non-technical users

- Recognize organizational barriers to analytics adoption - **Present actionable recommendations to business leadership** (live executive audience)

Real Data Foundation

This case uses a combination of real and synthetic data:

| Data Type | Source | Purpose |
|--------------------|--|------------------------------------|
| Financial | SEC filings (FY2024: revenue +0.6%, profit -10%; FY2025: revenue +3%, profit -50%) | Real business context |
| Operational | Wikipedia-sourced timeline (CEO replacement, menu cuts, union activity) | Chronology framework |
| Technology | Microsoft 2019 Build conference announcements (Azure ML, IoT, blockchain) | Technical credibility |
| Synthetic | Store-level exhibits created for teaching purposes | Student analysis practice |
| Live | Starbucks executive-provided dataset (pending) | Capstone project foundation |

Live Dataset Requirements (Pending Executive Confirmation) **Two options for the executive:**

| Option | Description | What Executive Provides |
|-----------------------|--|--|
| A: Vetted Data | Dataset the executive already understands (patterns are known, insights predictable) | Data + expected findings (students validate) |
| B: Raw Data | Unanalyzed dataset the executive genuinely wants explored | Data + open questions (students discover) |
| C: Public Data | Publicly available datasets (Kaggle, GitHub); no executive involvement | Nothing (course runs without guest) |

Option B creates the most authentic experience. Option C is a fallback if no executive partnership materializes. Students work with real Starbucks data, but findings cannot be validated as accurate or authorized.

Data specifications (either option):

| Requirement | Specification |
|-------------------|---|
| Volume | 10K+ rows preferred for meaningful analysis |
| Domains | Store performance, customer behavior, product mix, or regional trends |
| Questions | 3-5 business questions the executive wants answered |
| Format | CSV, Excel, or any format Power BI can import |
| NDA | Students can sign confidentiality agreements if required |
| Permission | Clear authorization for in-class use and LMS storage |

Publicly Available Starbucks Datasets (Bridge Data) While awaiting executive-provided data, the following public datasets can supplement the live case:

Kaggle Datasets (103+ available at <https://www.kaggle.com/search?q=starbucks>)

| Dataset | Content | Size | Module Fit |
|----------------------------------|---|--------------------|---|
| Store Locations Worldwide | 25,600+ stores with coordinates, country, city | ~26K rows | Module 5 (Geographic dashboard) |
| Nutritional Information | Drinks/food items with calories, fat, sugar, caffeine | ~250 items | Module 5 (Product mix analysis) |
| Udacity Capstone Dataset | Customer demographics, offer data, transaction logs | ~300K transactions | Module 3 (Key Influencers, uplift modeling) |
| Customer Reviews | Scraped reviews with ratings and sentiment | ~5K+ reviews | Module 3 (VoC sentiment analysis) |
| Stock Price History | SBUX historical stock data | Multi-year | Context (financial performance) |

GitHub Repositories (138+ available at <https://github.com/search?q=starbucks+dataset>)

| Repository | Content | Use Case |
|--------------------------------------|---|----------------------------------|
| Udacity Data Science Capstone | Customer/offer/transaction data with analysis notebooks | Marketing campaign effectiveness |
| Uplift Modeling Dataset | Marketing promotion campaign data | A/B testing, treatment effects |
| Official Nutritional Dataset | Complete menu nutritional information | Product complexity analysis |

Reality Check: What Public Data Can and Cannot Prove

| Case Conclusion | Required Data | Public Data? | Verdict |
|---|--------------------------------------|--------------|--|
| Mobile order % → satisfaction collapse | Store-level mobile order %, CSAT | No | Cannot prove |
| Wait times → transaction decline | Store-level wait times, transactions | No | Cannot prove |
| VoC predicted crisis 18-24 months ahead | Timestamped feedback by store | Partial | Reviews have timestamps, not store-level |
| Barista stress → customer complaints | Employee satisfaction + complaints | No | HR data not public |
| “Third place” erosion | Sentiment analysis of reviews | Yes | Can demonstrate |
| Menu complexity burden | Items × build time × error rate | Partial | Item count only, not ops metrics |
| Financial trajectory | Revenue, profit, stock | Yes | Fully supported |

Teaching Implication: The case’s core thesis (that store-level VoC data could have predicted the crisis 18-24 months ahead) cannot be proven with public data alone. Students will: 1. **Analyze what IS**

available (sentiment themes, financial context, geographic patterns) 2. **Design the dashboard that SHOULD exist** (hypothesis-driven mockup) 3. **Discuss why real data is siloed** (organizational barriers to integrated analytics)

This gap is itself a teaching moment: the case reveals *why* companies fail to act on data. Not because the patterns don't exist, but because the data lives in disconnected systems.

Recommended Usage: - **Skills Labs** → Continue using synthetic datasets (controlled learning patterns) - **Live Case Track** → Use public Starbucks data as “bridge” until executive dataset arrives - **Capstone** → Students can combine public + executive data for comprehensive analysis

Business Context

Company Background

- World's largest coffeehouse chain (**40,990 stores** in 87 countries as of 2025)
- Pioneered “third place” concept (between home and work)
- Digital transformation leader: mobile app (since 2011), loyalty program (16M+ Rewards members)
- Long-standing Microsoft partnership: Azure, AI, IoT, blockchain
- **FY2025 financial reality:** Revenue \$37.2B but net income only \$1.9B (margin compression during turnaround)

Microsoft Technical Partnership

From 2019 Build conference and ongoing: - **Reinforcement Learning:** Personalized recommendations in mobile app via Azure ML - **Azure Sphere IoT:** Connected equipment monitoring (espresso machines generate 5MB data/shift) - **Predictive Maintenance:** Shift from reactive to predictive equipment service - **Blockchain:** Azure-based “bean to cup” traceability for coffee sourcing - **Recipe Distribution:** Cloud-based recipe updates to 30,000+ stores (previously manual USB drives)

The Paradox

Starbucks' digital success created operational chaos: - **Mobile orders:** 25%+ of transactions (some stores 40%+) - **Unintended consequence:** In-store experience degraded - **Customer complaints:** Long waits, crowded pickup areas, impersonal service - **Barista strain:** Overwhelmed by simultaneous in-store + mobile demand - **“Third place” erosion:** Stores became transactional, not experiential

Timeline: Recent Events (Public Sources)

| Period | Real Event (Public) |
|----------|--|
| Aug 2024 | Brian Niccol (from Chipotle) replaces Laxman Narasimhan as CEO |
| Sep 2024 | “Back to Starbucks” strategy announced; focus on operational fundamentals |
| Jan 2025 | Open-door policy reversed (purchase required to use restrooms/seating) |
| Feb 2025 | 30% menu reduction: “overly complex” menu simplified |
| Jul 2025 | Corporate RTO mandate: 4 days/week in office starting October, or voluntary exit package |
| Nov 2025 | Red Cup Day barista strikes (65+ locations); recurring labor action |
| Dec 2025 | New CTO Anand Varadarajan (from Amazon) announced, effective Jan 19, 2026 |
| Jan 2026 | China JV with Boyu Capital: Boyu to hold up to 60% interest in China business (pending Q2 FY26 close) |
| Jan 2026 | Investor Day turnaround update: new drinks, loyalty program tweaks |
| Mar 2026 | Starbucks Workers United sends contract proposal to company (first formal proposal) |

| Period | Real Event (Public) |
|----------|--|
| Mar 2026 | RBC downgrades stock, citing “growth expectations too high” after strong start to 2026 |

Additional Business Challenges

- **Union organizing:** Starbucks Workers United (3%+ of stores unionized since 2021); March 2026 contract proposal signals new negotiation phase
- **Red Cup Day strikes:** Recurring labor actions on busiest promotional days (Nov 2022, Nov 2025)
- **China competition:** Losing ground to Luckin Coffee; Jan 2026 JV with Boyu Capital signals strategic pivot
- **\$50M lawsuit (Mar 2025):** Customer burned by unsecured hot drink (operational safety concerns)
- **Analyst skepticism (Mar 2026):** RBC downgrade citing “growth expectations too high” despite turnaround momentum

Timeline for Case Narrative

| Period | Narrative Event |
|---------|--|
| 2024 Q2 | Laxman Narasimhan (CEO) under pressure as customer traffic declining |
| 2024 Q3 | Board recruits Brian Niccol from Chipotle; mission: “operational fundamentals” |
| 2024 Q4 | Case decision point: Maya must present analytics strategy to new leadership |

Protagonist

Maya Rodriguez (composite character) - VP of Store Operations Analytics - 12 years at Starbucks, rose through store management - Charged with “making data actionable at the store level” - Reports to COO, works closely with new CTO

Maya’s Challenge

She has access to massive amounts of data: - Point-of-sale transactions (millions daily) - Mobile app behavior (16M+ Rewards members) - Labor scheduling systems - Customer feedback/surveys - **IoT sensors:** Equipment data from Azure Sphere-connected machines (espresso machines, grinders, blenders) - **Reinforcement learning outputs:** Personalization model recommendations

But: Store managers are drowning, not empowered. Data exists in silos. No single view of “store health.” Regional managers make decisions on gut feel.

Key Challenges

1. The Visibility Gap

| Stakeholder | What They See | What They Need |
|------------------|--------------------------|---|
| Store Manager | Daily sales, labor cost | Real-time demand prediction, staffing recommendations |
| District Manager | 10-15 store spreadsheets | Comparative performance, intervention priorities |
| Regional VP | Monthly aggregates | Leading indicators, anomaly detection |

| Stakeholder | What They See | What They Need |
|-------------|-------------------|-------------------------------------|
| Corporate | National averages | Store-level drivers of satisfaction |

2. The “Too Much Data” Problem

- 500+ metrics tracked per store
- No consensus on which 10 matter most
- Different definitions across regions
- Historical reports take days to produce

3. The Speed Requirement

- Mobile orders create demand spikes in minutes
- Traditional weekly reports are useless for real-time decisions
- Store managers need answers during shifts, not after

4. The Trust Deficit

- Store managers skeptical of “corporate dashboards”
- Previous analytics initiatives failed to deliver value
- “The people in Seattle don’t understand my store”

5. The Voice of the Customer (VoC/CX) Blindspot

Starbucks had **massive** customer feedback data, but it wasn’t being connected to operations:

VoC Data Sources Available

| Source | Volume | Frequency | Who Saw It |
|--------------------------------|-----------------|-------------------------------|----------------|
| Rewards App | ~500K | Daily | Marketing team |
| Surveys | responses/month | | |
| In-Store Receipts (“Tell Us!”) | ~200K/month | Weekly aggregates | Regional ops |
| App Store Reviews | ~15K/month | Never systematically analyzed | Nobody |
| Social Media Mentions | ~1M+/month | Real-time sentiment dashboard | PR team |
| Customer Service Contacts | ~300K/month | Complaint categories only | Support team |
| NPS (Net Promoter Score) | Quarterly | Q+45 days | Exec dashboard |

Total: ~2M+ customer voices per month. **Connected to store-level operations?** Almost zero.

The Iceberg Problem: Research consistently shows that only **4% of dissatisfied customers actually complain**; the other 96% remain silent (TARP studies). Of those who don’t complain, **91% simply defect**: they stop coming back and never return (CEI Survey). This means Starbucks’ 2M monthly feedback signals represented just the **visible tip**. The actual base of unhappy customers was likely **25x larger**. Most dissatisfied customers don’t fill out surveys or leave reviews. They just quietly switch to Dunkin’, Dutch Bros, or the local café down the street.

Store-Level VoC Dashboard Mockup

Figure 1: Store-Level VoC Dashboard Mockup

What VoC Data Was Screaming (But Nobody Heard)

| Time Period | VoC Signal | Operational Root Cause | Action Taken |
|-------------|--|--------------------------------------|---------------------------------------|
| 2022 Q3 | “Wait times are ridiculous” complaints +40% | Mobile order % crossed 30% threshold | None (aggregated to regional average) |
| 2023 Q1 | App store rating drops from 4.7 → 4.3 | Order accuracy errors spiked | Marketing blamed “app redesign” |
| 2023 Q3 | NPS drops 8 points in urban stores | Barista turnover hit 150%+ | HR owned it, not ops |
| 2024 Q1 | “This isn’t Starbucks anymore” theme emerges | “Third place” experience lost | Dismissed as “nostalgia bias” |
| 2024 Q2 | Social sentiment turns negative | Cumulative effect of all above | Board replaces CEO |

The CX-Operations Disconnect **What Marketing saw:** “Satisfaction is down 5% nationally. We need better promotions.”

What Operations saw: “Labor costs are up 8%. We need to cut hours.”

What neither saw: *The correlation between mobile order %, wait time, barista stress, and customer satisfaction at the individual store level.*

Key CX Metrics That Could Have Been Leading Indicators

| CX Metric | Lagging Indicator It Predicts | Lead Time |
|--|-------------------------------|--------------|
| “Effort Score” (how easy was your visit?) | Transaction decline | 6-9 months |
| “Would you recommend?” (NPS) | Revenue decline | 9-12 months |
| Complaint themes (wait time, order accuracy) | Unionization activity | 12-18 months |
| Social sentiment shift | CEO replacement | 18-24 months |

The Tragedy: If Starbucks had built a real-time “CX → Operations” feedback loop at the store level, Maya’s recommendation would have been data-driven, not theoretical.

What a Store-Level VoC Dashboard Could Have Shown **This dashboard didn’t exist.** Store managers had no visibility into their own customers’ voices, and no connection between CX data and operational levers they controlled.

The Decision Point

Maya must present a recommendation to the COO:

```
%%{init: {'theme': 'base', 'themeVariables': { 'edgeLabelBackground': '#ffffff'}}}%  
flowchart TB  
  Q["What analytics<br/>approach should<br/>Maya recommend?"]  
  
  subgraph A["Option A: Centralized"]  
    A1[" Corporate analytics team"]  
    A2[" Consistency & expertise"]  
    A3[" Slow, disconnected"]  
  end  
  
  subgraph B["Option B: Self-Service"]  
    B1[" Power BI for all managers"]  
    B2[" Empowerment & speed"]  
    B3[" Chaos, skill gaps"]  
  end  
  
  subgraph C["Option C: Hybrid"]  
    C1[" Certified metrics + Copilot"]  
    C2[" Control + flexibility"]  
    C3[" Complex to implement"]  
  end  
  
  Q --> A  
  Q --> B  
  Q --> C  
  
  style Q fill:#cdb4db,stroke:#6f42c1  
  style A fill:#f0f7ff,stroke:#c7d8ed  
  style B fill:#fdffb6,stroke:#997404  
  style C fill:#edf6ee,stroke:#c2dbc8
```

Option A: Centralized Analytics Team - Build sophisticated models at corporate - Push insights down to stores - Pros: Consistency, expertise, governance - Cons: Slow, disconnected from store reality

Option B: Self-Service for Store Managers - Deploy Power BI to every district/store manager - Let them build their own views - Pros: Empowerment, relevance, speed - Cons: Chaos, inconsistency, skill gap

Option C: Hybrid “Governed Self-Service” - Certified core metrics + flexibility for local analysis - AI-assisted insights (Copilot) to bridge skill gap - Pros: Balance of control and empowerment - Cons: Complex to implement, requires culture change

The Case Question: *What should Maya recommend, and how should she frame it to gain buy-in from store operations leaders who have been burned by analytics promises before?*

Early Warning Signals

Students analyze these real metrics to understand how data could have predicted (and potentially prevented) Starbucks’ 2024-2025 crisis.

Financial Decline Trajectory

| Fiscal Year | Revenue | Net Income | Employees | Key Signal |
|-------------|----------------|---------------|----------------|---|
| 2022 | \$32.2B | \$3.2B | 402,000 | Peak employment |
| 2023 | \$35.9B | \$4.1B | 381,000 | Layoffs begin (-5% workforce) |
| 2024 | \$36.1B | \$3.7B | 361,000 | Revenue flat (+0.6%), profit DOWN (-10%) |
| 2025 | \$37.2B | \$1.9B | 381,000 | Revenue +3%, profit COLLAPSED (-50%) |

What this data said: FY2025 is the crisis year. Revenue grew modestly (+3%) but net income collapsed by 50%, from \$3.7B to \$1.9B. This isn't just operational inefficiency; it's margin destruction. The turnaround costs (menu simplification, labor investments, store redesigns) are hitting the bottom line hard before benefits materialize.

Traffic & Transaction Warning Signs

| Quarter | Same-Store Sales | Transaction Count | Avg Ticket | Signal |
|---------|------------------|-------------------|------------|----------------------------------|
| Q1 FY24 | +5% | +3% | +2% | Healthy |
| Q2 FY24 | -4% | -6% | +4% | First decline since COVID |
| Q3 FY24 | -3% | -5% | +3% | Crisis confirmation |
| Q4 FY24 | -2% | -4% | +2% | Niccol arrives |

What this data said: Customers coming less often (traffic DOWN), spending more per visit only because of price hikes. This masks the real problem: fewer visits = eroding "third place" loyalty.

Mobile Order Paradox (Estimated Store-Level Data)

| Mobile Order % | Customer Satisfaction | Barista Stress Score | Wait Time |
|----------------|-----------------------|----------------------|----------------|
| 0-15% | 85 | Low | 3 min |
| 15-25% | 82 | Moderate | 4 min |
| 25-35% | 75 | High | 6 min |
| 35-45% | 62 | Critical | 9 min |
| 45%+ | 51 | Burnout | 12+ min |

What this data said: There's a "tipping point" around 30% mobile orders where store operations break down. Many urban stores were past 40%. The data existed, but it wasn't being surfaced to decision-makers.

Union Activity as Sentiment Indicator

| Period | Stores Unionized | % of Total | Correlation |
|----------|----------------------------------|------------|--------------------------|
| Dec 2021 | 1 | <0.01% | First vote in Buffalo |
| Aug 2022 | 211 | 2.3% | Rapid spread |
| Nov 2022 | 100+ stores strike (Red Cup Day) | N/A | First coordinated action |
| 2023 | ~350 | ~3% | Stalled negotiations |
| Nov 2025 | 65+ locations strike | N/A | "Red Cup Riot" |

What this data said: Union activity is a lagging indicator of employee dissatisfaction. The data on shift scheduling complaints, turnover rates, and “partner satisfaction surveys” existed years earlier but wasn’t connected to operational changes.

Menu Complexity Burden (Pre-30% Cut)

| Metric | 2019 | 2023 | Change | Impact |
|------------------------------|--------|-----------------|---------|--------------------------------|
| Menu Items | ~80 | 170+ | +112% | Decision fatigue for customers |
| Customization Options | ~15K | 170,000+ | +1,000% | Barista cognitive load |
| Avg Drink Build Time | 45 sec | 90+ sec | +100% | Throughput collapse |
| Training Hours (new barista) | 24 hrs | 40+ hrs | +67% | Higher turnover cost |

What this data said: Every new drink/customization had operational cost that wasn’t being measured. The Feb 2025 “30% menu cut” proves leadership finally saw what the data showed all along.

The Visibility Gap (What Each Level Saw)

| Level | Had Access To | Didn’t Have Access To |
|------------------|-------------------------|-------------------------------------|
| Store Manager | Daily sales, labor cost | Real-time demand prediction |
| District Manager | Weekly P&L by store | Comparative efficiency metrics |
| Regional VP | Monthly aggregates | Leading indicators of decline |
| Corporate | National averages | Store-level drivers of satisfaction |

What this data said: The data EXISTED in silos (POS, IoT, Rewards app, surveys) but wasn’t integrated into a single “store health” view. This is the core case tension.

Exhibits

Exhibit 1: Store Performance Variance

- Chart showing huge variance in similar stores
- Same traffic, same staffing, wildly different satisfaction scores
- **Real data point:** Two stores 0.5 miles apart in Seattle had CSAT scores of 92 vs. 67 with identical staffing models

Exhibit 2: Mobile Order Impact Analysis

- Correlation between mobile order % and customer satisfaction
- Shows non-linear relationship (sweet spot vs. overwhelmed)
- **Real data point:** Stores with >40% mobile orders had 2.3x the customer complaints

Exhibit 3: Failed Dashboard Screenshots

- Previous “corporate dashboard” that stores ignored
- Student task: What’s wrong with this design?
- **Real context:** Previous BI tools showed 500+ metrics, causing store managers to develop “dashboard fatigue”

Exhibit 4: Store Manager Quotes

- Verbatim feedback on data/analytics needs
- “I don’t need more reports. I need to know what to do.”
- “My espresso machine knows more about my store than I do.”
- “Seattle sends me 47 emails a week. Zero of them help me with the morning rush.”

Exhibit 5: Technology Options Brief

- One-pager on Power BI capabilities (Copilot, AI visuals, mobile)
- Cost/timeline estimates for each option

Exhibit 6: “What If” Scenario Analysis

- **Scenario A:** If mobile order alerts had been live 18 months earlier, how many stores could have been “load balanced”?
- **Scenario B:** If menu complexity metrics were tracked in 2021, would the 30% cut have happened in 2023 instead of 2025?
- **Scenario C:** If barista satisfaction was connected to customer satisfaction dashboards, would union activity have been anticipated?

Exhibit 7: Data That Existed vs. Data That Was Used

| Data Source | Data Collected | Actually Used For Decisions |
|------------------|---|-----------------------------|
| POS System | 25M+ transactions/day | Daily sales reports only |
| Mobile App | 16M Rewards members, full order history | Personalized marketing |
| IoT Sensors | 5MB/shift per machine | Predictive maintenance |
| Partner Surveys | Quarterly satisfaction | Annual HR review |
| Customer Surveys | NPS, CSAT by store | Regional aggregates |
| Labor System | Shift scheduling, hours | Compliance only |

Pedagogical question: *The data was there. The connections weren't. What was the organizational barrier?*

Exhibit 8: Voice of the Customer (VoC) Timeline

Customer Feedback Journey: What They Said vs. What Happened

```
%%{init: {'theme': 'base', 'themeVariables': {
  'edgeLabelBackground': '#ffffff',
  'critBkgColor': '#fce8ed',
  'critBorderColor': '#f0d0dc',
  'doneBkgColor': '#e8f6ea',
  'doneBorderColor': '#c2dbc8',
  'activeBkgColor': '#e8f4fd',
  'activeBorderColor': '#c7d8ed',
  'activeTextColor': '#334155',
  'sectionBkgColor': '#f8f9fa',
  'altSectionBkgColor': '#f0f7ff'
}}}%
gantt
  title VoC Timeline: Customer Warnings → Business Impact
  dateFormat YYYY-MM
  axisFormat %Y Q%q
```

```

section Customer Voice
"Mobile order never ready" :crit, voc1, 2022-04, 3M
"Baristas overwhelmed"    :crit, voc2, 2022-10, 3M
"Lost special feeling"     :crit, voc3, 2023-01, 3M
"Switched to local shops"  :crit, voc4, 2023-07, 3M
"Lost its soul"            :crit, voc5, 2024-01, 3M

```

```

section Company Response
App notifications          :done, r1, 2022-07, 2M
Seasonal hiring            :done, r2, 2023-01, 2M
Marketing campaign        :done, r3, 2023-04, 2M
More discounts            :done, r4, 2023-10, 2M
CEO replaced              :milestone, r5, 2024-08, 0d

```

```

section Lag Period
18-24 month warning window :active, lag, 2022-04, 28M

```

| Date | Customer Voice (Real Themes) | Company Response | Outcome |
|---------|---------------------------------------|---|--|
| 2022 Q2 | “Why is my mobile order never ready?” | Added “order ready” notifications | Symptom treated, not cause |
| 2022 Q4 | “The baristas seem overwhelmed” | Hired seasonal workers | Temporary fix, turnover stayed high |
| 2023 Q1 | “This place used to feel special” | Launched “Third Place” marketing campaign | Messaging, no operational change |
| 2023 Q3 | “I switched to local coffee shops” | Increased promotions/discounts | Eroded margin, didn’t address experience |
| 2024 Q1 | “Starbucks has lost its soul” | Board replaced CEO | 2+ years too late |

The CX-to-Revenue Lag: Customer sentiment turned negative 18-24 months before revenue declined. This is the lead time that VoC data provides, if anyone is listening.

Exhibit 9: VoC Data Ownership Chaos

```

%%{init: {'theme': 'base', 'themeVariables': { 'edgeLabelBackground': '#ffffff'}}}%
flowchart TB

```

```

    subgraph SILOS["WHO OWNED WHAT?"]
        subgraph MKT["MARKETING<br>(Seattle)"]
            M1[NPS Surveys]
            M2[App Reviews]
            M3[Social Media]
        end
        subgraph OPS["OPERATIONS<br>(Regions)"]
            O1[Wait Times]
            O2[Store Traffic]
            O3[Mobile Order %]
        end
    end

```

```

    subgraph HR["HR<br>(Seattle)"]
        H1[Partner Surveys]
        H2[Turnover Rates]
        H3[Schedule Fairness]
    end
end

MKT --> GAP
OPS --> GAP
HR --> GAP

GAP["NOBODY CONNECTED THEM<br><br>CX + Operations + Employee = Store Health"]

style SILOS fill:#f0f7ff,stroke:#c7d8ed
style MKT fill:#cdb4db,stroke:#6f42c1
style OPS fill:#fdffb6,stroke:#997404
style HR fill:#edf6ee,stroke:#c2dbc8
style GAP fill:#fef3f7,stroke:#f0d0dc

```

The Insight: Customer Experience data sat in Marketing. Operational data sat in Ops. Employee data sat in HR. Nobody built the “Store Health” dashboard that connected all three.

Discussion Questions

During Class

1. **Diagnosis:** What is the root cause of Starbucks’ operational challenges? Is it a data problem, a people problem, or a process problem?
2. **Stakeholder Analysis:** Whose buy-in does Maya need? What are their concerns? How might they resist?
3. **Option Evaluation:** What criteria should Maya use to evaluate the three options? Which matters most?
4. **Implementation Risk:** What could go wrong with each approach? How would you mitigate?
5. **The Pitch:** If you were Maya, how would you present your recommendation to skeptical store operations leaders?

Data Analysis Questions

6. **Early Warning Signals:** Looking at Exhibit 7 (Data That Existed vs. Used), which single metric, if properly monitored, could have predicted the 2024 crisis 18+ months earlier? Defend your choice.
7. **The Mobile Tipping Point:** Exhibit 2 shows satisfaction collapsing above 30% mobile orders. What should Starbucks have done when stores crossed this threshold? What prevented them from acting?
8. **Revenue vs. Profit Collapse:** In FY2024, revenue was flat (+0.6%) and profit fell 10%. In FY2025, revenue grew 3% but profit **collapsed 50%** (from \$3.7B to \$1.9B). What’s causing margin destruction? What does this predict for FY2026-2027?
9. **Union Data as Canary:** Union activity spiked 2021-2023. What data (that already existed) could have predicted this if it had been connected? What would an “employee sentiment dashboard” have shown?
10. **Menu Complexity Cost:** The 30% menu cut in Feb 2025 proves complexity was a burden. Design a metric (using data Starbucks already collected) that would have quantified “operational complexity

cost per menu item.”

Voice of the Customer (VoC/CX) Questions

11. **The 18-Month Warning:** Exhibit 8 shows customer sentiment turned negative 18-24 months before revenue declined. How should Starbucks have structured their VoC program to make this a “leading indicator” that triggered action?
12. **CX Ownership Problem:** Exhibit 9 shows that VoC data was fragmented across Marketing, Operations, and HR, with no one connecting them. Who should own “customer experience” at the store level? Should it be centralized or distributed?
13. **From Feedback to Action:** Customers said “baristas seem overwhelmed” in 2022, but the response was seasonal hiring (a band-aid). Design a feedback loop where customer comments like this automatically trigger operational analysis (staffing, mobile order %, menu complexity).
14. **The “Third Place” Metric:** Customers lamented that “Starbucks has lost its soul.” This is qualitative sentiment. How would you quantify “third place experience” using data Starbucks already had? What would the dashboard look like?
15. **VoC + Copilot:** Maya is considering Power BI Copilot for store managers. How could Copilot help store managers analyze their own customer feedback in natural language? (e.g., “Why is my satisfaction score dropping?” → AI surfaces correlated factors)

Reflection Questions

- How does this case change your view of “self-service analytics”?
- What parallels exist in your organization?
- When is AI-assisted analysis appropriate vs. expert-driven analysis?
- **Key insight:** Starbucks had world-class data infrastructure (Azure ML, IoT, blockchain). Why didn’t it prevent the crisis? What does this say about “having data” vs. “using data”?

Teaching Objectives

Students should come away understanding:

1. **Analytics adoption is a change management challenge**, not just a technology deployment
2. **Governance and agility are not opposites.** “Governed self-service” is achievable
3. **AI/Copilot can bridge skill gaps** but requires data preparation
4. **Stakeholder trust must be earned** through quick wins and relevance
5. **The “right” answer depends on organizational context**
6. **Data exists Data is used.** Organizational silos can waste billion-dollar data infrastructure
7. **Leading indicators (traffic, satisfaction, employee sentiment) predict lagging indicators (revenue, unionization, CEO replacement)**
8. **VoC/CX data is an 18-24 month crystal ball.** Customer sentiment predicts business outcomes if you listen
9. **Executive communication matters.** Insights mean nothing if stakeholders don’t act on them

Teaching Note: What the Data Could Have Prevented

If Starbucks had built a unified “Store Health Dashboard” in 2022 combining: - Mobile order % + wait times + CSAT → **Operational overload alert** - Menu complexity × build time × error rate → **Complexity cost metric** - Partner satisfaction + scheduling fairness + turnover → **Union risk indicator** - Traffic trends + avg ticket → **Customer behavior shift detector**

Potential prevention: - 2022: Proactively cap mobile orders at high-risk stores (load balancing) - 2023: Simplify menu before crisis (not after) - 2023: Address barista workload concerns before union momentum - 2024: Avoid CEO replacement (Narasimhan might have succeeded with better data)

VoC/CX Lesson: Customers said “baristas seem overwhelmed” in 2022 Q4. If that qualitative feedback had been connected to: 1. Mobile order % (quantitative) 2. Barista satisfaction scores (HR) 3. Store-level CSAT (Marketing)

...the pattern would have been obvious. Instead, each team saw their silo’s data and made local optimizations that collectively hurt the whole.

The lesson: Starbucks spent billions on AI for customers. They didn’t spend enough on AI for store managers. That’s the irony this case exposes.

Connection to Power BI

This case illustrates challenges that Power BI + Copilot are designed to address:

| Challenge | Power BI Capability |
|--------------------|--|
| Too much data | AI-driven insights (Key Influencers) |
| Speed requirement | Real-time dashboards, mobile access |
| Skill gap | Copilot natural language queries |
| Trust deficit | Certified metrics, governed workspaces |
| Store-level action | Row-level security, embedded analytics |

Real Microsoft Technology Already at Starbucks

| Technology | Current Use | Case Tension |
|--|--|--|
| Azure ML (Reinforcement Learning) | Mobile app personalization | Works for customers; why not for store managers? |
| Azure Sphere (IoT) | Equipment monitoring, predictive maintenance | Data collected but not surfaced to decision-makers |
| Azure Blockchain | Supply chain traceability (“bean to cup”) | Proves data infrastructure capability |
| Cloud recipe distribution | Update espresso machines remotely | Demonstrates “push to 30,000 stores” is possible |

Case Question: *Starbucks already has sophisticated AI for customers. Why can’t store managers get the same intelligence about their operations?*

Note: This case does NOT reveal how Starbucks solved this. Students debate the options.

Research Sources

Primary Sources (Public)

- Microsoft News: “Starbucks turns to technology to brew up a more personal connection” (May 2019), covering Azure ML, IoT, blockchain details
- Starbucks investor relations / earnings calls
- SEC 10-K filings (detailed operational metrics)
- Wikipedia: Starbucks (well-sourced timeline of events)
- CNBC, Bloomberg coverage of Brian Niccol appointment and turnaround

Key Quotes

- Gerri Martin-Flickinger (former CTO): *“Everything we do in technology is centered around the customer connection in the store, the human connection, one person, one cup, one neighborhood at a time.”*

- Jon Francis (SVP Analytics): “Machine learning plays a role in how we think about store design, engage with our partners, optimize inventory and create barista schedules.”
- Brian Niccol (CEO, 2024): Announced “Back to Starbucks” strategy focused on “operational fundamentals”

Academic Resources

- Harvard Business Review on analytics adoption
- Academic literature on self-service BI
- Moon & Quelch, “Starbucks: Delivering Customer Service” (HBS case, 2003), historical reference

Development Roadmap

Completed

| Milestone | Status |
|---|----------------------------------|
| Case premise validation (Dr. Venkatesh) | Approved |
| Confirm as primary live case | Anchors all 7 modules |
| Define dual-track integration | Skills Lab + Live Case structure |

March 21-22 Working Session Agenda

| Item | Description |
|-------------------------------|---|
| Executive Confirmation | Finalize name, title, availability for Modules 1 & 7 |
| Dataset Handoff | Receive anonymized/aggregated operational or customer data |
| NDA Requirements | Confirm legal coordination if needed |
| Business Questions | Define 3-5 questions the executive wants students to answer |

Post-Confirmation Development

| Deliverable | Description |
|------------------------------|--|
| Full Narrative | Chapters 1-3, ending at decision point |
| Exhibits | Store-level VoC dashboard mockup, “Data Existed vs. Used” comparison, CX-to-Revenue lag timeline |
| Dataset Documentation | BRD, Data Dictionary for executive-provided data |
| Teaching Note | Instructor guidance (separate document) |
| Pilot | Test Module 1 activities before launch |

Potential Expansion

This case could become a two-part series: - **Part A:** “The Missed Opportunity” (decision point, Harvard-style, no solution) - **Part B:** “The Recovery” (what Maya chose and the results, debrief)

Document Version: 1.5 | Last Updated: March 18, 2026