

Figure 1: AI-Assisted Power BI for Business Analytics

Case Study 3: Customer Support Root Cause Analysis

Microsoft Global Customer Support

Module Alignment: Module 3 (Discovering Insights with AI Visuals) **Discussion Time:** 30 minutes

Protagonist: Priya Sharma, VP of Customer Support Operations

The Story

Chapter 1: The Number No One Could Explain

The quarterly review had barely begun when Priya Sharma felt her stomach drop.

The VP of Customer Support Operations at Microsoft stared at the slide: **CSAT 3.96**. Down from 4.31 the previous quarter. An 8% decline in customer satisfaction in ninety days.

“Priya,” the CVP of Customer Experience said, leaning forward, “what happened?”

She opened her mouth to answer, but the truth was brutal: she didn’t know. Her organization handled millions of customer interactions annually. Somewhere in that ocean of tickets, calls, chats, and emails, something had gone catastrophically wrong. And she had exactly zero insight into what.

“We’re investigating,” she said. “We’ll have answers within -”

“Within when?” the CVP interrupted. “The board meets in six weeks. They’re going to ask about this. What do I tell them?”

Priya’s analytics team had already estimated the investigation timeline: two analysts, three to four weeks of work, multiple Excel models, no guarantee of finding the root cause.

She didn’t have four weeks.

Chapter 2: The Conference Room Full of Theories

That afternoon, Priya convened her leadership team in a windowless conference room.

“We need hypotheses,” she said. “What’s driving the drop?”

The responses came fast - and contradictory.

Head of Operations: “It’s staffing. We lost 15% of our agents to summer vacations. Response times went up.”

Head of Quality: “It’s training. We onboarded 200 new agents in Q2. They’re not performing at the level we need.”

Head of Channel Strategy: “It’s the shift from phone to chat. Customers hate chat. We pushed too hard.”

Head of Product Escalations: “It’s the Windows update. We got flooded after the July release. Support volume spiked.”

Priya listened as each leader defended their domain. No one had data. Everyone had theories. And each theory conveniently exonerated its proponent’s team.

“We can’t investigate all of these,” she said. “And even if we could, it would take a month. We need a different approach.”

Chapter 3: The Analyst Who Didn't Believe in AI

Raj Krishnamurthy had been a support analyst at Microsoft for eleven years. He'd seen every trend, every crisis, every flavor-of-the-month analytics initiative. When Priya's director asked him to try "those new AI visuals in Power BI," he was skeptical.

"I've been doing root cause analysis since before these tools existed," he said. "I know how to build a regression model. I don't need AI to tell me what's important."

But the timeline forced his hand. Traditional analysis would take weeks. The AI visuals promised answers in days. Reluctantly, Raj opened Power BI and started configuring Key Influencers.

What he set up: - **Analyze:** CSAT Score - **Explain by:** Channel, Support Plan, Agent Skill Level, Product Family, Resolution Time, Reopen Count, Day of Week

He clicked "Analyze" and waited for what he assumed would be garbage.

What appeared changed his career.

Chapter 4: The Factor No One Had Mentioned

The Key Influencers visual rendered in seconds. At the top of the list, with a clear margin over everything else:

#1: When Resolution Time Hours is greater than 24, CSAT decreases by 0.89 points

Raj leaned back. Resolution time. Not channel. Not staffing. Not training. Not the product release.

He drilled deeper. The data was unambiguous:

Factor	Impact on CSAT
Resolution Time > 24 hours	-0.89 points
Reopen Count > 0	-0.67 points
Agent Skill Level = Tier 1	-0.34 points
Channel = Email	-0.28 points
Day = Monday	-0.21 points

Resolution time was nearly three times more impactful than channel - the factor the Channel Strategy team had been certain was the culprit.

Key Influencers Visual *Figure: Power BI Key Influencers showing the factors that decreased CSAT. Resolution time (-0.89 pts) was 3x more impactful than channel (-0.28 pts).*

But the real revelation came from "Top Segments." Key Influencers had identified the deadliest combination:

Segment: Cases where Resolution Time > 24 hours AND Reopen Count > 0 AND Channel = Email

Average CSAT: 2.41 (vs. 3.96 overall) **Volume:** 12,847 cases **Impact:** This segment - just 8.5% of total cases - was responsible for 30% of the CSAT decline.

Raj stared at the screen. In eleven years, he'd never isolated a problem this precisely, this fast.

Chapter 5: The Pattern in the Chaos

Key Influencers had answered the "what." But Priya needed the "why."

Why had resolution times spiked? What had caused the backlog that pushed cases past 24 hours?

Raj turned to Anomaly Detection. He plotted daily case volume over the quarter, let the algorithm find unusual spikes.

Three dates lit up red:

Date	Expected Volume	Actual Volume	Variance
July 15	4,200	7,891	+88%
August 3	4,100	6,234	+52%
August 22	3,900	5,678	+46%

Raj cross-referenced with the company's product release calendar:

- **July 15:** Windows Update release
- **August 3:** Azure service outage
- **August 22:** M365 service degradation

Three major product events. Three volume spikes. But that wasn't the insight. The insight was what happened *after*.

He plotted CSAT over time with Anomaly Detection. New red markers appeared - two days after each volume spike:

Date	Expected CSAT	Actual CSAT	Lag
July 17	4.18	3.42	2 days after volume spike
August 5	4.15	3.51	2 days after volume spike
August 24	4.12	3.67	2 days after volume spike

The pattern was clear: volume surge → backlog builds → resolution time exceeds 24 hours → CSAT craters. The damage appeared precisely two days after each triggering event - the time it took for the backlog to impact customers.

```
%%{init: {'theme': 'base', 'themeVariables': {
  'edgeLabelBackground': '#ffffff',
  'critBkgColor': '#fce8ed',
  'critBorderColor': '#f0d0dc',
  'doneBkgColor': '#e8f6ea',
  'doneBorderColor': '#c2dbc8',
  'activeTextColor': '#334155',
  'sectionBkgColor': '#f8f9fa',
  'altSectionBkgColor': '#f0f7ff'
}}}%
gantt
  title Anomaly Detection: Volume Spikes & CSAT Impact
  dateFormat YYYY-MM-DD
  axisFormat %b %d

  section Product Events
  Windows Update      :milestone, event1, 2024-07-15, 0d
  Azure Outage        :milestone, event2, 2024-08-03, 0d
  M365 Degradation    :milestone, event3, 2024-08-22, 0d

  section Volume Spikes
  +88% Volume (7,891) :crit, vol1, 2024-07-15, 2d
  +52% Volume (6,234) :crit, vol2, 2024-08-03, 2d
  +46% Volume (5,678) :crit, vol3, 2024-08-22, 2d
```

```

section CSAT Impact (2-day lag)
CSAT 3.42 (-0.76)      :done, csat1, 2024-07-17, 3d
CSAT 3.51 (-0.64)      :done, csat2, 2024-08-05, 3d
CSAT 3.67 (-0.45)      :done, csat3, 2024-08-24, 3d

```

Figure: The 2-day lag pattern between product events and CSAT drops - invisible to traditional analysis but surfaced by Anomaly Detection.

Chapter 6: The Root Cause Map

By day three, Raj had assembled the complete picture. He presented it to Priya using a diagram that traced the chain of causation:

```

%%{init: {'theme': 'base', 'themeVariables': { 'edgeLabelBackground': '#ffffff'}}}%
flowchart TD
    A[" Product Events<br/><i>Windows Update, Azure outage, M365 issues</i><br/>3 major incidents"]
    B[" Case Volume Spikes<br/><i>+50-90% above normal</i><br/>12,000+ cases/day"]
    C[" Resolution Time Exceeded<br/><i>18 hrs → 29 hrs average</i><br/>Over 24-hour SLA"]
    D[" Cases Reopened<br/><i>4.2% → 7.8%</i><br/>Due to rushed fixes"]
    E[" CSAT Dropped<br/><i>4.31 → 3.96</i><br/>-8.1% quarter"]

    A -->|"Volume surge"| B
    B -->|"Backlog builds"| C
    C -->|"Pressure to close"| D
    D -->|"Customer frustration"| E

    style A fill:#fef3f7,stroke:#f0d0dc
    style B fill:#fefcf0,stroke:#e8dbb8
    style C fill:#f0f7ff,stroke:#c7d8ed
    style D fill:#fefcf0,stroke:#e8dbb8
    style E fill:#fef3f7,stroke:#f0d0dc

```

“This is it,” he said. “Product events caused volume spikes. We lacked surge capacity, so resolution times blew past 24 hours. Under pressure, agents rushed fixes - which led to reopens. Customers got slow answers *and* wrong answers. That’s why satisfaction cratered.”

Priya studied the diagram. Three days earlier, she’d had nothing but theories. Now she had a causal chain, backed by 150,000 data points, surfaced by AI in a fraction of the time traditional analysis would have taken.

“What do we do about it?”

Chapter 7: The 60-Day Recovery

The response came in two phases.

```

%%{init: {'theme': 'base', 'themeVariables': { 'edgeLabelBackground': '#ffffff'}}}%
flowchart LR
    subgraph Week1["Week 1: Stop the Bleeding"]
        A[" Activate<br/>Contractors"] --> B[" Prioritize<br/>24+ hr Cases"]
        B --> C[" Pause Quality<br/>Sampling"]
        C --> D[" Daily CSAT<br/>Standup"]
    end

    subgraph Weeks24["Weeks 2-4: Fix the System"]
        E[" Surge Staffing<br/>Protocol"] --> F[" Pre-release<br/>Briefings"]
        F --> G[" Reopen RCA<br/>Analysis"]
        G --> H[" Email+Chat<br/>Migration"]
    end

```

Week1 --> Weeks24

```
style A fill:#fef3f7,stroke:#f0d0dc
style B fill:#fef3f7,stroke:#f0d0dc
style C fill:#fef3f7,stroke:#f0d0dc
style D fill:#fef3f7,stroke:#f0d0dc
style E fill:#edf6ee,stroke:#c2dbc8
style F fill:#edf6ee,stroke:#c2dbc8
style G fill:#edf6ee,stroke:#c2dbc8
style H fill:#edf6ee,stroke:#c2dbc8
```

Week 1: Stop the Bleeding

Action	Rationale	Owner
Activate overflow contractors	Reduce backlog immediately	Operations
Prioritize 24+ hour cases	Clear oldest tickets first	Queue Management
Pause quality sampling	Let agents focus on volume	Quality Team
Daily CSAT standup	Track recovery in real-time	Leadership

Weeks 2-4: Fix the System

Change	Problem Addressed	Expected Impact
Surge staffing protocol	No plan for volume spikes	-50% backlog buildup
Pre-release support briefings	Agents unprepared for new issues	+15% FCR during releases
Reopen root cause analysis	Rushed resolutions causing reopens	-40% reopen rate
Email-to-chat migration	Email channel underperforming	+0.28 CSAT improvement

The changes cascaded through the organization. Operations stood up a rapid-response team that could activate within 24 hours of a product event. Quality began analyzing every reopen to identify training gaps. Channel Strategy - who had been so certain the problem was chat - quietly redirected their energy toward improving email workflows.

60 Days Later:

Metric	Q3 (Crisis)	Q4 (Recovery)	Change
CSAT	3.96	4.24	+7.1% ↑
FCR	71%	79%	+8 pts ↑
Reopen Rate	7.8%	4.0%	-3.8 pts ↓
Resolution Time	29 hrs	19 hrs	-34% ↓

CSAT didn't just recover - it surpassed the pre-crisis baseline.

Chapter 8: The Skeptic's Conversion

Six months later, Raj Krishnamurthy gave a presentation at an internal analytics conference. His topic: "How AI Visuals Saved My Quarter."

"I've been doing root cause analysis for eleven years. I thought I'd seen every approach. When someone told me to try Key Influencers, I rolled my eyes. I figured it would give me obvious insights I already knew.

Instead, it showed me that resolution time was three times more important than channel - the opposite of what most of the leadership team believed. It found a segment representing 8.5% of our cases that was causing 30% of our damage. And Anomaly Detection revealed a two-day lag between volume spikes and CSAT drops that none of us had ever noticed.

AI didn't replace my judgment. It accelerated it. What would have taken me four weeks took four days. And the insights were better - because the AI had no political stake in the answer."

In the audience, Priya smiled. The skeptic had become an evangelist.

The Analysis Behind the Story

What AI Found That Humans Missed

Insight	Why Humans Missed It
Resolution time was 3x more important than channel	Teams were defending their domains, not following data
The "Monday effect" was real but minor	Intuition overweighted visible patterns
8.5% of cases caused 30% of damage	Aggregate views hide toxic segments
2-day lag between spikes and CSAT drops	Time-series patterns require anomaly detection

When to Trust AI Insights

Trust Level	Criteria
High	Finding aligns with business logic and passes sanity check
Medium	Finding is surprising but data supports it
Verify	Finding contradicts domain expertise - investigate further
Skeptical	Finding seems too convenient or sample size is small

In this case, all AI findings passed the "makes sense" test: - Of course slow resolutions hurt satisfaction - Of course reopens indicate failed first attempts - Of course volume spikes overwhelm capacity

AI didn't discover magic - it quantified what intuition suspected and revealed what data showed.

Configuration Tips

Key Influencers Setup: - Analyze CSAT as continuous (not just "satisfied/unsatisfied") - Use categorical factors, not raw IDs - Exclude highly correlated metrics (like NPS) - Ensure sufficient volume (150K cases provided statistical power)

Anomaly Detection Setup: - Sensitivity: 80% (high - wanted to catch significant deviations) - Seasonality: Weekly (support has day-of-week patterns) - Expected range: Statistical (let algorithm determine bounds)

Discussion Questions

For Class Discussion

1. **AI vs. Human Analysis:** Could traditional analysis have found the same insights? What would have been missed or delayed?
2. **Key Influencers Interpretation:** The #1 factor was resolution time. How actionable is this insight? How do you turn correlation into causation?

3. **Anomaly Detection Value:** The 2-day lag was a key discovery. Why might humans have missed this pattern?
4. **Organizational Response:** The immediate actions included “pause quality sampling.” Is this a good trade-off?
5. **Prevention:** How could this crisis have been prevented? What monitoring would have provided earlier warning?

Reflection Prompts

- Have you experienced a “data but no insights” situation?
- How might AI visuals have helped in that scenario?
- What root cause investigations does your organization struggle with?

Connection to Course Labs

Lab	Connection
Lab 3	You’ll build Key Influencers to analyze CSAT in SupportInsights data
Lab 3	You’ll add Anomaly Detection to find case volume spikes
Lab 3	You’ll write an executive summary synthesizing AI findings

Note: The SupportInsights dataset is designed to reveal similar patterns. You’ll discover embedded insights about resolution time, channel performance, and agent effectiveness.

Additional Resources

- Key Influencers Visual Documentation
- Anomaly Detection in Power BI
- Root Cause Analysis Best Practices

Case materials based on real Microsoft support operations. Details modified for educational purposes.