Teacher Answer Key

Case Study

1. First failure: When Priya said she thought the charge was an error, the AI responded with payment plan information — a complete mismatch to her actual need. Second failure: When she explicitly asked to 'dispute this charge,' the AI provided a broken PDF link — offering a path that didn't work with no fallback. Third failure: When she asked for a billing specialist, the AI deflected back to generic helpfulness. A human agent would have recognized the dispute intent immediately, accessed her account, and either initiated the dispute process or transferred her to a specialist who could.

2. Priya should call the hospital's billing department directly — saying clearly that she has an insurance dispute and needs a billing specialist. She should have her insurance card, explanation of benefits (EOB) from her insurer, and the bill itself ready. She should not attempt to resolve a $847 insurance dispute through a chatbot that has already demonstrated it cannot handle complexity. She might also call her insurance company directly to confirm what they should have covered and get that information in writing before calling the hospital.

Multiple Choice

1. B
2. C
3. A

Matching

1-B, 2-E, 3-D, 4-C, 5-A

Fill-in-the-Blank

1) 84, 2) hallucination, 3) agent

Short Answer (Sample Responses)

Q1: AI chatbots excel at high-volume, low-complexity tasks: order tracking, password resets, FAQs, appointment confirmations, standard policy questions. What these tasks have in common is that they are structured, repetitive, and have clear, retrievable correct answers. The AI doesn't need judgment or empathy — it needs accurate lookup and consistent communication, which it does well.

Q2: If customers give up before reaching a human agent, the company avoids the cost of a human interaction. Since many customers abandon a service request rather

than fight through multiple barriers, the math works in the company's favor. Consumer advocates find this problematic because it effectively denies customers the service they're entitled to, uses technology as an obstruction rather than a tool, and disproportionately affects less persistent or less technically savvy customers.

Q3: Arguments for national disclosure standard: customers have a right to know they're not talking to a human, especially in high-stakes situations. Confidence in a conversation changes when you know the other party cannot exercise judgment or empathy. Arguments against: companies might argue it slows deployment of beneficial AI tools or that customers don't actually need or want to know. Most consumer advocates and ethicists would support the national standard, given the documented harm of doom loops and hallucinated policy answers.