W1: Token, cost, and temperature worksheet

How Large Language Models Work: Foundations for Choosing and Applying LLMs at Work (FAST-14) · Lab 1: Tokens, probabilities, and temperature

Name: _____ Date: _____

Item	Your result
FAQ words / tokens / tokens per word	
String that split into the most tokens, and why you think so	
Tokens per request (FAQ + instructions + question)	
Cost per question (show the formula)	
Monthly cost, 3,000 questions; with caching	
Top 5 next tokens for "every ..." and their probabilities	
Temperature 0 vs 1.0: how much did answers vary for welcome vs check-out?	
Setting you would choose for the FAQ bot, with one-sentence reason	

Completion criteria

- Token counts are recorded and the cost is calculated with the formula shown.
- The answer separates input tokens from output tokens.
- The temperature choice is justified by the type of task (factual versus creative).

7formations · AI Training 2026 · Instructor: Wadood Chaudhary