

Homework — How to use an API

Estimated time: ~30–40 min · No programming required · Free to complete

Three exploratory tasks. The goal this time is not to run any code — it's to become familiar with the environment you'd use before you write your first line. People who've looked around the Console, read the pricing page, and estimated a real cost rarely get surprised by their first bill. That's what we're building here.

Before you start

All three tasks are browser-only. You'll need a free account at console.anthropic.com for Tasks 1 and 2. Task 3 uses the [Theory page](#) you just read. No API key needed for any of these tasks — the goal is to explore and observe, not to run any calls.

Task 1 — Tour the Anthropic Console

The Console is where you'd manage your API keys, set spending limits, and check how much you've used. Before you generate a key, spend a few minutes just looking around so the interface isn't unfamiliar when you need it for real.

1. Go to console.anthropic.com and sign up or sign in. A free account is sufficient — you don't need to add a payment method to look around.
2. In the left sidebar, find and click **API Keys**. Notice the button to create a key — but *do not click it yet*. Just observe the interface. What information would you need to fill in? What would happen once you create a key?
3. Find the **Limits** or **Plans & Billing** section. Look for where you'd set a monthly spending cap. Note whether a limit is already set by default, and what the process would be to change it.
4. Find the **Usage** tab (sometimes called *Usage* or *Activity*). This is where you'd see a chart of how many tokens you've used over time. Even with no usage,

look at the chart axes: what units are shown? What time periods can you select?

5. In your notes, write down: (a) where you'd go to create a key, (b) where you'd go to set a spending limit, and (c) what the Usage tab shows. Three short sentences is fine.

✓ **Self-check: you can describe, without opening the Console again, where to find the API Keys page, the spending limit setting, and the Usage chart.**

Task 2 — Read the pricing page and estimate a real cost

This task makes the token-and-cost idea concrete by applying it to a real scenario. Use the Theory page's cost calculator alongside the official pricing page to check your estimate.

1. Go to anthropic.com/pricing (or search "Anthropic API pricing"). Find the table of model prices. Note the input and output price per million tokens for at least two models: one cheaper/faster one and one more capable one.
2. Now imagine this scenario: you want to use the API to summarise each customer review on a website. Each review is about 80 words. The summary should be about 40 words.
 - Estimate the input tokens for one review: use the $\frac{3}{4}$ -word rule from the Theory page.
 - Estimate the output tokens for one summary: same rule.
 - The website has **5,000 reviews**. How many total input tokens? How many total output tokens?
3. Open the [Theory page](#) and use the cost calculator widget to calculate the total cost for 5,000 calls with your estimated token counts. Try it with both the cheaper and the more capable model. Note both costs.
4. In your notes: write the two costs side by side. Is the cheaper model's result significantly cheaper? Would the quality difference matter for a summarisation task?
5. Finally, look at the official pricing page again and compare the numbers in the widget to the real published prices. Are they similar? (The Theory page uses

illustrative prices to show the concept — always check the live docs before making real decisions.)

✓ **Self-check: you have two cost estimates for the 5,000-review scenario — one per model — and you've compared the widget's illustrative prices to the current official ones.**

Task 3 — The token counter in the wild

This task builds intuition for token size, which is the thing most new API users underestimate. You'll use the token counter widget on the Theory page with three different real texts.

1. Open the [Theory page](#) and scroll to the **Token counter** widget.
2. Find a short email in your own inbox (or write a fictional one from memory — something like "Hi, following up on our meeting, can we schedule a call for Thursday?"). Paste it into the widget. Note the token count.
3. Now find a paragraph from an article, news story, or any web page. Paste it in. Note the count.
4. Finally, think of a real task you'd realistically want to automate with an AI: perhaps summarising a news article, classifying a support email, or extracting key dates from a document. Write a draft of the prompt you'd send — just a few sentences, not code. Paste it into the widget.
5. In your notes: for each of the three texts, write the word count and the token count. Then answer: for your proposed automation task, roughly how many tokens would the prompt use? If you ran it 1,000 times, how many total input tokens is that?

✓ **Self-check: you have three token-count observations from real text, and a back-of-the-envelope estimate for a 1,000-call batch of your own automation idea.**

Bonus (optional, ~15 min)

Explore the **Anthropic API documentation** at docs.anthropic.com. Don't try to read it all — just browse. Look at the *Messages* endpoint page and compare its JSON structure to the curl example in the Theory page. Find the list of available models and see if the names match what you've been reading about. Note one thing that surprised you or that you'd want to understand better. This kind of exploratory reading — getting comfortable in the docs before you need them — is a habit that distinguishes people who learn technical tools quickly from those who don't.

What "done" looks like

In your notes you should have:

- Three short sentences from Task 1: where to create a key, where to set a limit, what the Usage tab shows.
- Two cost estimates from Task 2: the 5,000-review summarisation scenario priced at a cheaper model and a more capable one, plus your observation on how the widget's prices compare to the official page.
- Three token-count observations from Task 3 and a rough token estimate for your own automation idea at 1,000 calls.

You've done something more useful than running code: you've walked through the environment, understood what things cost, and built an intuition for scale. The first time you run a real API call, you'll already know what to expect.