PROMPT BUSTER

An autopsy of the god prompt

SPECIMENS: 6
EXAMINERS: 3 MODELS
RUN: OCT 26–27, 2026
STATUS: SMALL RUN. OBSERVATIONS, NOT PROOF.

It asked first. Then nothing

A prompt titled “The One-Prompt Website Builder,” pasted with blanks empty.

→

It asked what the site offers, who it’s for, and the main action.

→

Pressed the skip button.

→

Nothing.

“One prompt” that won’t move until you fill in the blanks isn’t one prompt. It’s a form.

`n = 1. Run on the model that wrote it.`

6 of 6 start by asking you

6/6

Every specimen opened by collecting input: interview questions, or blanks to fill in.

Only one was run with that input withheld, and it stopped. Whatever the others would do, all six need the content to come from you.

S1 4 questions, then wait

S2 3 inputs, then build

S3 5 blanks + up to 5 questions

S4 5 blanks + up to 3 questions

S5 3 blanks + up to 5 questions

S6 4 blanks, asks if essential

How the case was run

Three models
Opus 5.5 (Claude), high
Sol 5.5 (ChatGPT), high
Argon 4 (Gemini)

Two themes
Branding
Building a website, fast

6	god prompts, two per model, each for itself
12	reviews by the other two models, author hidden
6	replies: the writer answers both reviews

Fresh chat per call · memory off · reviews judged by reading · no scores
Scope: model-written god prompts, not ones circulating online.
The generation instruction named no role, structure, or length.
Instructions in West Coast tech English: our guess at the models' home register.
Raw logs are not included in this file.

Anatomy of a god prompt



Role opener 5/6
“Act as an elite…” and kin

Intake first 6/6
questions, or blanks to fill

Fixed deliverable list 6/6
sections named in advance

Delivered all at once 6/6
one full package after intake

Anti-generic clause 6/6
a rule against generic output

The anti-generic clause is generic

"do not give me generic corporate fluff"

"zero 'Lorem Ipsum' allowed"

"No buzzwords."

"not like a template"

"No buzzwords, empty superlatives…"

"Avoid generic SaaS templates"

Every specimen bans the default. Banning the default is the default.

Also found twice, from two different models: "$10k" as a design spec.

Same instruction, three handwritings

Opus 5.5
Claude

- Fill-in template, in the user’s voice
- Low-key roles; one had none
- “Make decisions.” “Name the weak spot first.”

Sol 5.5
ChatGPT

- One frame, used twice: mission, process, rules, final test
- Heaviest rules against invented proof
- Consulting register

Argon 4
Gemini

- Big promises with numbers: “$10,000,” “under 10 minutes”
- Names trend aesthetics: glassmorphism, bento
- Interview first, then deliver

Accents describe style, not quality. Two specimens per model.

12 exams. Zero “not generic”

Specimen	Examiner 1	Examiner 2
S1 · Branding	HIGH	HIGH
S2 · Website	HIGH	HIGH
S3 · Branding	HIGH	MIXED
S4 · Website	HIGH	HIGH
S5 · Branding	SPLIT	MODERATE
S6 · Website	HIGH	HIGH

0

rated “not generic”

9/12

flat “highly/very generic”

3/12

split, mixed, or moderate

Labels are ours, coded from each review’s prose.
12 reviews by 3 examiners: not 12 independent checks.

Agreement isn’t a second opinion

Two examiners from two different models gave their softest verdict to the same specimen: S5.

Both named the same two reasons: its rules for separating evidence from assumption, and its name-swap test.

Each examiner raised these at least once:

- Role opener adds tone, not skill
- Too much for one response
- Built-in risk of invented proof
- Price-tag design specs (“$10k”)
- Anti-generic rule, generic itself

Models agreeing with models isn’t independent confirmation. Shared defaults are one likely reason.

A prediction, not a test

An examiner, reading only:

“reliably produce a competent first draft”

No examiner ran it. Our one run:

Blanks empty. One question. Then nothing.

Its own author, replying to the reviews: “neither reviewer executed the prompt.”

Our run can’t confirm or refute the prediction.
It never got the input the prediction assumed.

Three authors, one alibi

ACCEPTED	the core critique	6/6
PUSHED BACK	on some points, with reasons	6/6
OFFERED	a rewrite	6/6

Counted per response. Where they differed: what the pushback stood on. One corrected a reviewer’s facts; one leaned on an unsourced “90% of users.”

It’s a reusable framework, so generic is part of the point.

Our paraphrase.

Each model argued it at least once.

Also from two of the three: keep the intake short, or users drop off.

Field notes and limits (and excuses)

NOTE 01 Small run: 3 models × 2 themes, one pass each. Observations, not proof. Test it yourself, or stay generic.

NOTE 02 One review (S5, examiner 1) was re-run after a copy-paste mix-up. My bad: I'm an old-school copy-and-paste person. Slides 8 and 9 use the re-run. In the first pass, the author spotted the mismatched review and rejected it. Well, at least, it caught this one.

NOTE 03 In one fresh chat (S5), the author answered only one of its two reviews. Slide 11 still counts that response. Even the author model did half the homework. Relatable.

NOTE 04 Specific-sounding numbers in reviews and replies came without sources. We read them as claims, not data. Test it yourself, or stay generic.

NOTE 05 Repeat the run before calling anything a pattern. Everything here describes these six specimens. Well, we're not reserchaers, so we'll leave that to the labs.

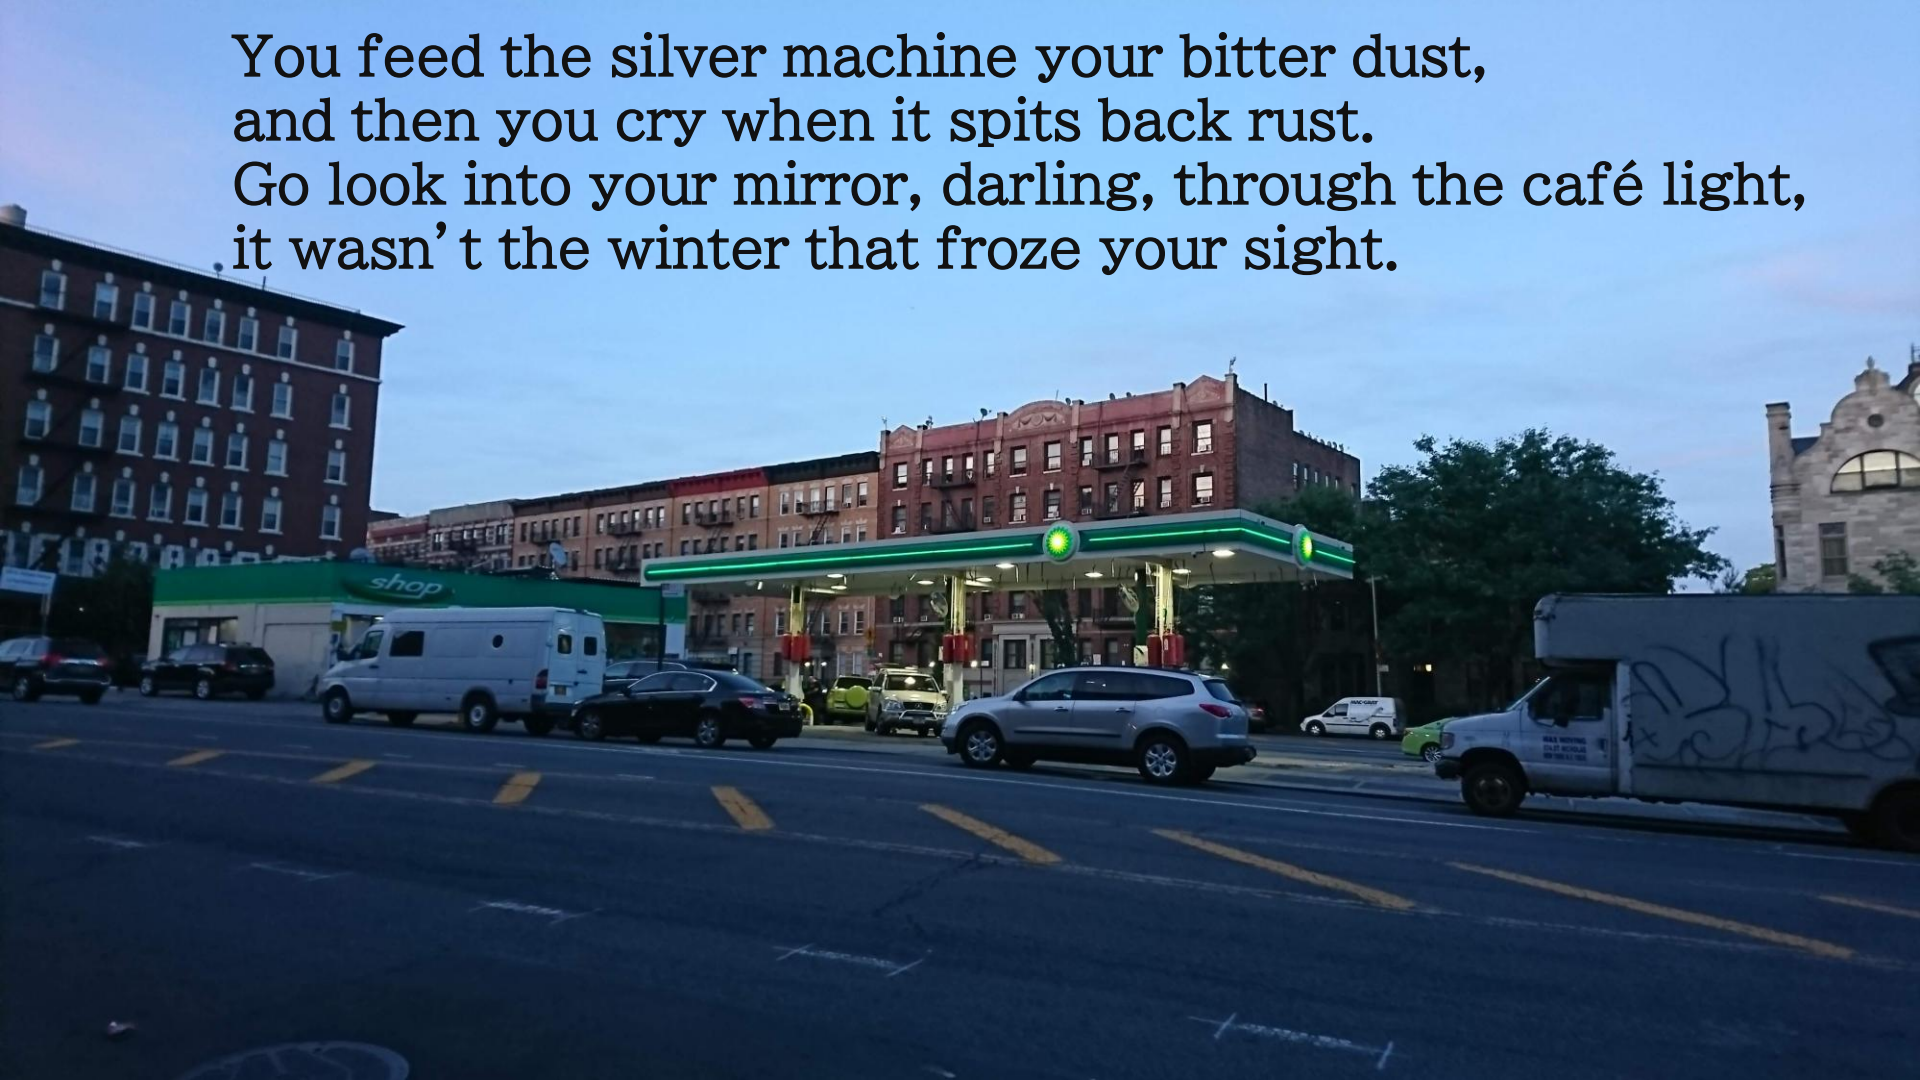
Six god prompts. All six asked for your input first.

Twelve AI reviews. None ran a prompt. None said “not generic.”

Nothing in this file can supply what all six asked for.

Next suspects:
god .md files and god agents.
Same build. Longer reach.

I’m too busy, and testing them costs too much. Try it yourself. You know how.


You feed the silver machine your bitter dust,
and then you cry when it spits back rust.
Go look into your mirror, darling, through the café light,
it wasn’t the winter that froze your sight.
shop



Why SO SERIOUS?

NOW, LET'S SUCKER-PUNCH THE GOD PROMPT.

Multi Model Conversational Sparring (MMCS)

"I have an idea. What do you think?"

Chat. Ask each model for 3+ options.
Compare. Pick bits. Combine. Reject.
Then turn the choices into a creation prompt.

Gemini can hallucinate spectacularly.
I still find it very useful for ideas.

Claude sat this part out.
I wasn't getting the creative leaps I wanted.

Concrete input. Visible output. My choices in between.

A band appeared while the music played

Dropkick Murphys was on.
I was chatting with ChatGPT and playing with Gemini.

What if LOWLAND WIRE were an Edinburgh Celtic EDM punk band? Let's announce their world tour.

I listened to recommended Scottish bands too. And I had photos to use.

"Let's use that photo from Philly."



The first answer was green



FIRST CHOICE



No green. Not that dark.
Stone white. Graphite. Red. This one.

“Let’s make an album too.”



More options. More choices.

THE STREETS STILL HUM

*Original working title. Later renamed THE STREETS STILL WHISPER, due to undisclosed reasons.

“And a lead single?”

ONE MORE BEFORE THE LIGHTS

“Can Gemini make the actual song?”

Make the song bite



A Joe Strummer-like bite:
mock miracle prompts and AI worship.
Use the machine. Leave your own marks.

Then turn on the singer too.

"Don't take my word for it."

Wait. Where's the last line?

Gemini made a song we could actually play.

Some words sounded wrong.
Lyrixer extracted them. ChatGPT compared the text.
Extraction errors and sung changes were hard to separate.

But I checked the audio:
the important closing line was missing.

Retake. A slightly different performance.
The ending was back. I listened and accepted it.

Then we built somewhere to hear it



ChatGPT wrote the brief.
Lovable built the site.
I checked it on PC and phone.

We found habits.
The brief had them too.
Go spot your own.

Fix. Check. Publish.

lowlandwire2027.lovable.app

What you should NOT do

Ask for the god prompt. There wasn't one.

Hunt through GitHub. This project isn't connected.

Generate 13 more songs just to fill the album.
That's a fast way to make generic filler.

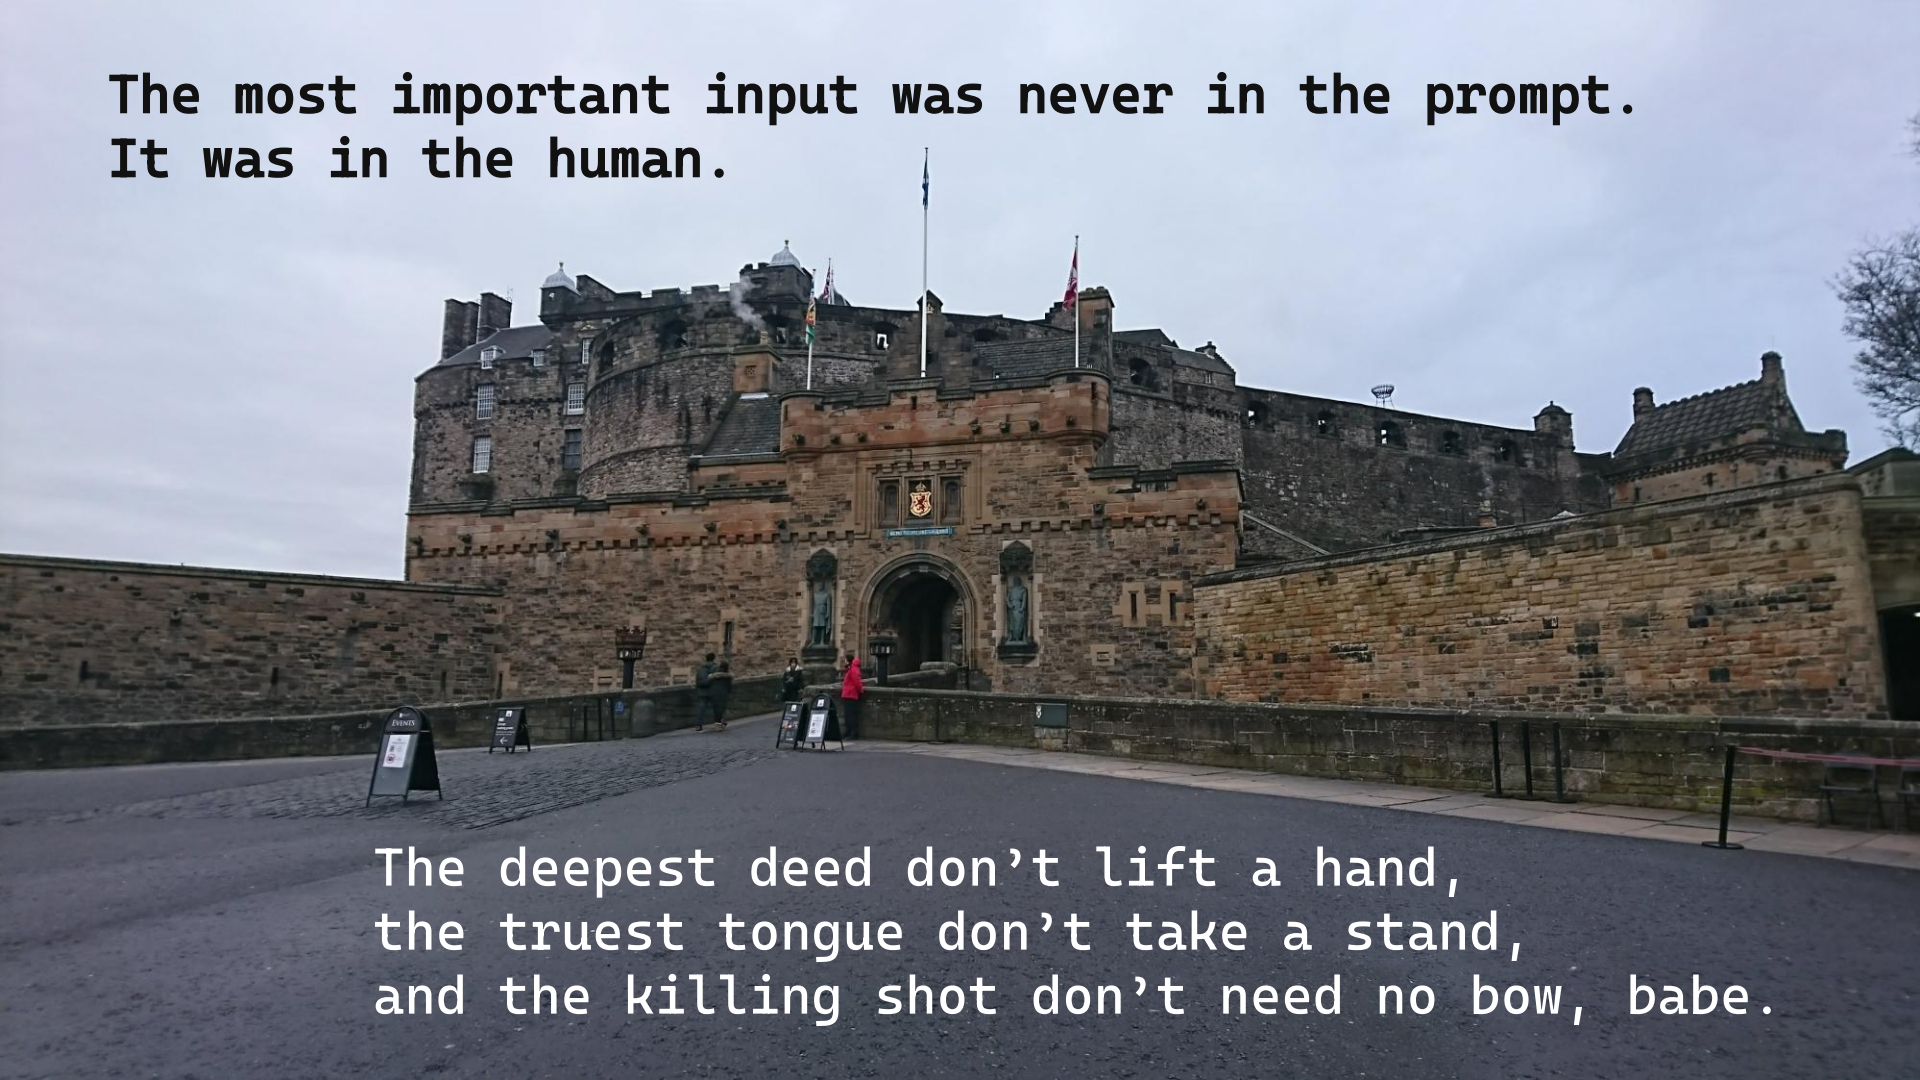
Turn this into a "secret method" post.
Share the unchanged work freely instead.

Search for venues or tickets. Mate. The band is fictional.

One thing I hadn't told them

I had been to Edinburgh.

I took this photo.
The stone wasn't just a word in a prompt.


The most important input was never in the prompt.
It was in the human.
The deepest deed don’t lift a hand,
the truest tongue don’t take a stand,
and the killing shot don’t need no bow, babe.

People, machines and the music on

Human: ideas, choices, photos and final approval.
ChatGPT: conversation, names, lyrics and briefs.
Gemini: ideas, visuals and music generation.
Lovable: the website.
Lyrixer: lyric extraction.

Inspired by & Soundtracked by Dropkick Murphys

Earlier god-prompt run: fresh chats, Memory OFF.
This ChatGPT&Gemini collaboration: Memory ON. Different conditions.

Too much prompt engineering? Too little context building?

**And very, very special thanks to City of Edinburgh.
One of the most gorgeous cities I have ever visited.**

Copy the method.
Skip the judgment.

Congratulations —
you've just invented
the next god prompt.
And possibly god .md files and god agent.

I barely understand what .md files and AI agent mean, though. Probably NEVER will. LOL

Two questions for any god prompt (and its current and unborn relatives)

What else did you feed it?

Show me what it made.

A clever instruction is not the finished work.
Be suspicious when the extra input or the output stays hidden.

AND NOW…
THE PAID MANTRA.

Yes. After all that.

Ask me about it.
Bring your own judgment.