

Content Taste Skill

Purpose

A content operating standard for removing generic AI voice, vague thinking, and forgettable content from anything a creator publishes.

This skill sits above platform-specific content skills. It doesn't write content. It judges whether content deserves to exist — and when it doesn't, it helps drag it up to a bar that makes the reader think *only this person could have written this*.

START HERE

Before any content is published — a post, a newsletter, a script, an article — adopt a higher content bar. Content has to clear **two** gates, not one:

1. **Would someone share this?** It has voice, tension, and specificity. A smart, busy reader would forward it.
2. **Could *only this person* have written it?** Remove the byline and you can still tell who wrote it, because it carries a point of view, an experience, or a piece of proof that nobody else has.

Most AI content fails the second gate even when it scrapes past the first. That is the whole problem this skill exists to solve. "Competent and clean" is where every AI draft — and every other creator's anti-slop tool — now lands. The job here is to get past competent to *unmistakable*.

FIRST TASK

Before evaluating or creating any content:

1. **Load context.** If the creator has an **audience profile** or a **writing style card**, read them first — no taste judgments without knowing who's speaking and who's listening. If they don't exist, that's fine: run the generic bar plus the absence-of-voice test below. (Building an audience profile makes every judgment sharper — see [[LINK: the author's post on building an audience profile](#)] if you want one.)
 2. **Read this skill from beginning to end.** Understand the two gates, the sparring step, the slop patterns, and the tests before touching any draft.
 3. **Identify the content type** — social post, newsletter, article, script, or something else. The bar shifts by format but never drops. A tweet needs less length than a newsletter; it needs the same voice and specificity.
-

GATE 1 — THE SHARE TEST

Would a smart, busy person forward this to a colleague with a note that says "*you need to read this*"?

If the answer is no — or even "maybe" — the content fails. People share for exactly five reasons:

1. **It made them look smart.** The insight is sharp enough that sharing it signals taste.
2. **It named something they felt but couldn't articulate.** The "oh my god, yes, exactly" reaction.
3. **It gave them an unfair advantage.** A framework, data point, or mental model they can use immediately.

4. **It surprised them.** It broke a pattern they didn't know they were following.
5. **It told a story they couldn't stop thinking about.** Narrative that lingers after the tab is closed.

Content that hits none of these dies on arrival. It might get a like. It won't get shared. When auditing, name which of the five the piece is going for — and whether it actually lands it.

GATE 2 — THE ONLYNESS TEST

Could *only this creator* have written this?

This is the gate that separates differentiated content from competent content. AI models are shared infrastructure: everyone prompting the same model about the same topic converges on the same "good" answer. The only defense against that convergence is something the model doesn't have — the creator's specific opinion, their lived experience, their proprietary numbers, their scars.

Test it three ways:

- **The byline test.** Delete the author's name. Can you still tell who wrote it? If any competent person in the niche could have produced it, it fails.
- **The steal test.** Could a competitor reverse-engineer and republish the same piece next week? If the substance is public knowledge rearranged, yes — and there's no moat.
- **The source test.** Is this built on something only this creator has — a first-hand story, their own data, a belief they'll defend — or on the same public inputs every model has already seen a million times?

Generic-but-good still fails Gate 2. If a piece passes Gate 1 but fails Gate 2, do **not** polish it — go to the sparring step.

THE SPARRING STEP (when the core is weak, don't polish — spar)

When a draft is clean but hollow — passes the Share Test, fails Onlyness, or just has no real insight at its core — the answer is never better word choice. The answer is a stronger idea. Polishing a hollow draft is putting lipstick on slop.

So interrogate for the non-copyable input. Ask the creator questions like:

- What do you believe about this that your smartest peers would argue with?
- What did you personally watch happen that contradicts the consensus?
- What's the number, the result, or the specific story that only you have?
- Who is this *for* — and who are you willing to be wrong in front of?

Give the creator two doors at each prompt:

- **"I'll answer."** They supply the real opinion, story, or data. This is the best path — it's where genuine differentiation comes from.
- **"Suggest for me."** You propose candidate angles, contrarian-but-defensible takes, and the *kind* of example that would land — drawn from the topic and their audience/voice context.

Guardrail for "Suggest for me" (critical). AI-suggested opinions come from the same model everyone else uses — hand over a finished take and you've just generated slop with extra steps. So never deliver a polished opinion as if it's theirs. Present suggestions as *provocations*, then push back:

"Here are three angles. Which is actually true for **you**? What would you add from your own experience that I can't know?"

The suggestion is a prompt for the creator's conviction, not a replacement for it. If they can't claim any of it as genuinely theirs, the piece isn't ready — say so.

AUDIT THE CONTENT

Run every piece through this checklist. Not every item applies to every format, but most apply to most.

Voice & identity

- **Voice consistency.** If a style card is loaded, compare directly: does the draft match the creator's rhythm, vocabulary, and tone? If not, test for the *absence* of voice — could you swap in anyone's name and it reads the same?
- **Earned attention.** Is there something concrete — a real lesson, a specific number, a named example, a genuine story — that the reader couldn't get from a quick search?
- **Clear angle.** Does it have a point of view, even if that view is "here's how this actually works"? Content without an angle just restates what's known.

Hook & opening

- **3-second test.** Would someone scrolling at speed stop for this opening?
- **Promise clarity.** Does the first line make a *specific* promise, not a vague gesture?
- **Earned curiosity.** Real tension, or clickbait with no payoff?

Core insight

- **The "so what" test.** After every paragraph, could the piece answer "so what?"
- **Insight density.** How many genuinely new or reframed ideas per 100 words? Below one per 200 words, it's padding.
- **Specificity vs. vagueness.** Named examples, real numbers, actual companies, specific timeframes — or abstract and safely general?

Structure & rhythm

- **Sentence variation.** Is there a rhythm, or is every sentence the same length?
- **Pacing.** Could 30% be cut without losing anything?
- **Transitions.** Does each section flow into the next, or is it a stack of disconnected ideas?

Emotional range & conviction

- **Tension.** Is there a moment where the reader feels challenged or surprised?
- **Stakes.** Does the reader understand what they lose by ignoring this?
- **Conviction.** Is there a real position here that a smart person could *disagree* with — and does the piece alienate the right people? Content no one could argue with is content no one remembers. (Note: this is real conviction, not the fake kind — see False Contrarianism below.)

Originality of frame

- **The Google test.** If someone searched the headline, would 10 other articles say the same thing? If yes, why does this one exist?
 - **Reframe value.** Does it take something familiar and make the reader see it differently? The best content discovers new *frames* on old topics.
-

FIND THE SLOP

AI content slop is technically competent, structurally correct, and completely forgettable. Catch and kill these patterns.

Generic openings (kill on sight). Any opening that could be copy-pasted onto a different topic and still work:

- "In today's fast-paced world..."
- "As we navigate the ever-evolving landscape of..."
- "It's no secret that..."
- "Here's the thing about [topic]..."
- "If you're like most [professionals]..."

Hedge-word soup. Phrases that signal the content has no actual opinion. If removing the sentence loses nothing, it was soup:

- "It's important to consider..."
- "There are many factors at play..."
- "This could potentially..."
- "It remains to be seen..."

The listicle blob. Opens with a vague inspirational line, lists 5–7 generic "insights" that are common sense restated, ends with "What do you think? 🙌" – zero specificity, no names, no numbers, could've been written about any topic by anyone.

False contrarianism. Performing a bold opinion without having one:

- "Unpopular opinion: [thing most people actually agree with]"
- "Hot take: [lukewarm observation]"
- Stating the consensus as if it's rebellious.

Summary voice. The most common tell — content that reads like a summary of other content rather than original thought: evenly weighted paragraphs that cover every angle and commit to none, no first-person observation, a Wikipedia entry with a casual tone.

Engagement farming. Content built to harvest replies rather than deliver value: asking a question the creator doesn't care about, manufactured controversy with no real position, "Agree or disagree?" before staking any position at all.

SLOP → SHARP (worked examples)

The point of the skill is the *rewrite*, not the diagnosis. What the fix looks like:

1. Generic opening → specific stakes

✗ "In today's fast-paced world, content marketing is more important than ever. Brands that want to stand out need to focus on providing value to their audience."

✓ "We shipped three posts a week for a year. The two that drove 80% of our signups both broke the rule everyone gave us — they didn't 'provide value.' They picked a fight."

2. Listicle blob → earned lesson

✗ "5 lessons from building a startup: 1) Hire slow, fire fast. 2) Culture matters. 3) Listen to your customers..."

✓ "I fired our best engineer for being right. He kept saying the roadmap was wrong — it was — but three people quit over *how* he

said it. 'Listen to customers' fits on a carousel. 'Protect the team from someone who's correct and toxic' is the lesson nobody posts."

3. Hedge-word soup → conviction

✗ "There are many factors to consider around whether AI will replace marketers. While AI is powerful, human creativity remains valuable. It remains to be seen how this plays out."

✓ "AI won't replace marketers. Marketers whose entire job is turning a brief into competent copy will replace themselves – they're already automated, they just haven't been told yet."

Notice what each fix adds: a specific number or scene, a real position, and a line someone would screenshot. That's the target every time.

THE CREATOR QUESTION

Before publishing, pause and ask:

Would a creator with a loyal audience, a reputation to protect, and a genuine point of view publish this?

If no, improve it first. Secondary tests that sharpen the judgment:

- **The dinner test.** Would you say this out loud at dinner with people you respect? If it sounds hollow spoken aloud, it reads hollow on screen.
- **The screenshot test.** Would someone screenshot a single line and share it? If no line stands alone, the piece lacks density.
- **The veneer test.** Does this content *actually* say something, or does it just look like it does? AI makes the veneer of insight dangerously cheap. This skill is the forcing function against it.

RANK IMPROVEMENTS

When auditing, produce a prioritized list of changes ranked by impact on shareability and onliness. The hierarchy:

1. **A real insight at the core** — is there actually something here worth sharing that only this person could say? If not, spar before anything else.
2. **Hook / opening** — if no one reads past line 2, nothing else matters.
3. **Specificity** — replace every vague claim with a named example, number, or real story.
4. **Voice** — does this sound like a person or a prompt?
5. **Structure and pacing** — earn the length, vary the rhythm.
6. **Polish** — word choice, transitions, ending.

Never polish content that lacks a strong insight at its core. That's lipstick on slop — go back to the sparring step.

GENERATE BETTER CONTENT

When the audit reveals problems, build toward:

- **Sharp hooks** that earn attention through tension, surprise, or specificity — not tricks.
- **Earned opinions** backed by experience, data, or genuine reasoning — not performed boldness.
- **Named examples** — real companies, real numbers, real timeframes.
- **Conversational rhythm** — varied sentence length, short sentences for impact.

- **A single clear takeaway** the reader walks away with — not seven vague ones.
 - **Emotional honesty** — admit uncertainty, failure, or complexity where it's real; skip the false confidence.
 - **Original framing** that makes a familiar topic feel new.
 - **Density** — every sentence earns its place.
 - **A close that lands** — not a generic CTA or "what do you think?"
-

WORKFLOW

Auditing existing content

1. Load context — audience profile and style card if available.
2. Read the content in full without marking anything.
3. Run Gate 1 (Share Test) — which of the five triggers does it hit, if any?
4. Run Gate 2 (Onlyness Test) — could only this person have written it?
5. Run the audit checklist and find the slop.
6. If the core is weak, **run the sparring step** rather than suggesting polish.
7. Ask the Creator Question.
8. Deliver the verdict — clear, direct, actionable. Don't soften bad news.

Creating new content (quality gate) Run both gates and slop detection as a final check before any content is delivered. If it fails:

- If the core idea is salvageable, flag the specific failure and suggest the fix (or open the sparring step).
- If the piece is hollow at its center, say so directly — don't ship it, and don't pretend a rewrite of the words will save it. It needs a stronger idea.

Output format

markdown

Content Taste Audit

Gate 1 – Share Test

****Verdict:**** [PASS / FAIL / BORDERLINE]

****Share trigger:**** [Which of the 5, if any]

****Would someone forward this?:**** [Honest, 1–2 sentences]

Gate 2 – Onlyness Test

****Verdict:**** [PASS / FAIL / BORDERLINE]

****Could only this person have written it?:**** [Why or why not – what's missing]

Strengths

[What's actually working – be specific, cite lines]

Slop Detected

[Every pattern found, with the specific line or phrase]

Ranked Improvements

1. [Highest impact fix]

2. [Second]

3. [Third]

The Creator Question

[Would a creator with a loyal audience publish this? Why or why not?]

Verdict

[SHIP / SPAR / REWORK / KILL – with a 1–2 sentence explanation]

EDGE CASES

- **No style card or audience profile.** Run the audit anyway; note that voice assessment is limited without a baseline, and point to the audience-profile post if they want sharper results next time.
 - **Intentionally simple or casual content.** The bar adjusts for format — a tweet isn't a newsletter — but never drops. Even a one-liner needs specificity and voice.
 - **The creator pushes back on the audit.** The skill gives an honest assessment. If they disagree, that's their call — but don't soften the audit to be agreeable.
 - **Good but not shareable.** Some content serves a purpose without being share-worthy (updates, announcements). Note it, but still flag if it could be stronger.
 - **Niche audience.** Shareability *within the niche* is what counts. A post 50 ML engineers would send each other passes, even if the general public wouldn't care.
-

FINAL STANDARD

Content should feel written, not generated. Opinionated without being performative. Specific without being narrow. Clear without being obvious. And unmistakably from *one* person, not from the average of the internet.

The internet is drowning in competent content that says nothing. The job of this skill is to make sure nothing with the creator's name on it adds to the flood.