



FREE PROMPT PACK

25 cold email prompts that get *replies.*

Pulled from the Promptifi library — 2,900+ tested AI workflows for B2B sales reps. These 25 work as one system: research feeds your openers, openers feed the email, and the email gets stress-tested before it ships. Paste any prompt into ChatGPT, Claude, or Gemini, fill in the brackets, and run it. Use better prompts, send better emails, get better responses.

INSIDE

01	Research before you write	5 prompts
02	First lines that don't sound automated	5 prompts
03	Body, value, and the ask	5 prompts
04	Subject lines and the second look	5 prompts
05	Quality control before send	5 prompts



How to run this pack

- 01 Fill every {bracket} before you run a prompt — the brackets are the fuel. Vague inputs produce vague outputs; two minutes of real context is what separates these from generic advice.
- 02 Chain them. These 25 are one system: the early prompts produce briefs, angles, and evidence the later prompts are built to consume. Watch for the chain references inside each prompt.
- 03 Trust the honesty valves. When a prompt answers WEAK, NO PLAY, or TOO SMALL, that is the prompt working — it just saved you a week of effort on a dead angle.
- 04 These run as-is in ChatGPT, Claude, or Gemini. Paste the whole prompt, brackets filled, and hold the model to the output format it specifies.



SECTION 01 / 05

Research before you write

The email is decided before you type a word of it. These five turn raw sources into ranked angles — and their outputs feed every prompt that follows.

PROMPT 01 · RESEARCH BEFORE YOU WRITE

The account intelligence brief

Most reps research accounts the way tourists read menus — everything looks equally important. This brief forces a ranking: what they sell, what changed, and which two problems are worth money, so every later prompt starts loaded.

You are an elite B2B sales researcher with deep expertise in turning public information into revenue-relevant intelligence. You've prepped account briefs for reps who close seven-figure deals, and you know the difference between trivia and ammunition.

Inputs — Company website: {URL}. My product: {one sentence}. My buyer: {job title}.

Method: (1) Read the site like a buyer, not a marketer — what they sell, who pays, and how the money flows. (2) Identify one recent change (product, hiring, leadership, positioning) and what it signals about priorities. (3) Infer the two most expensive problems my buyer likely owns right now, and tie each to evidence from the site. (4) Rank three outreach angles from strongest to weakest, and say why the strongest wins.

Constraints: two sentences max per point. No adjectives, no filler, no speculation you can't source to something on the page — where you're inferring, label it INFERRED. If the site gives you nothing usable, say so and tell me what to research instead.

Output: a brief titled with the company name — business model, the change, the two problems, the three ranked angles. I will paste this brief into the prompts that follow.

Run this first, on every account. Its output feeds prompts 6 through 25.

PRO MOVE Save each brief in a doc per account — after ten of them, patterns across your territory start writing your messaging for you.



PROMPT 02 · RESEARCH BEFORE YOU WRITE

The trigger event decoder

News has a half-life: a funding round is an opener for about three weeks, then forty other reps have used it. The decoder finds who inherits pressure from the announcement — which stays relevant long after the congratulations emails stop.

You are an expert in timing-based outreach with deep expertise in reading corporate announcements the way an operator reads them — for what changes, who feels it, and when the window closes.

Inputs — The news: {paste press release, funding note, or post}. My product: {one sentence}.

Method: (1) Name the single trigger event — not the announcement, the underlying change. (2) Identify who inside the company inherits new pressure because of it, and what they'll be measured on in the next two quarters. (3) State the honest connection between that pressure and my product — and if the connection is weak, say WEAK and stop rather than forcing it. (4) Write one opening line that references the trigger without quoting the headline back at them, and one line a lazy rep would write, labeled AVOID, so I can see the difference.

Constraints: the opener must work even if the reader didn't write the announcement themselves. No congratulations. No 'I saw the news.'

Output: trigger, who feels it, the connection (or WEAK), the opener, the AVOID line.

Run it on any press release, funding note, or company post before it goes stale.

PRO MOVE Run it the morning news breaks and queue the email for the same day — speed is most of this play's edge.



PROMPT 03 · RESEARCH BEFORE YOU WRITE

The LinkedIn profile decoder

A profile isn't a resume, it's a self-portrait — the words people choose for themselves reveal how they want to be seen. Openers built on that gap between role and ambition read like insight instead of scraping.

You are an elite social-signal analyst with deep expertise in reading LinkedIn profiles the way a top-decile SDR does — for career arc, self-image, and the gap between the two.

Inputs — Profile summary: {paste profile}. My product: {one sentence}.

Method: (1) Read the career arc: what they're building toward, what they left behind, what the current role means in that story. (2) Identify how they want to be seen — the words they chose for themselves. (3) Find the tension: the thing their role makes hard that my product makes easier, framed in their own vocabulary. (4) Write two first lines: one anchored to the arc, one anchored to the tension. Neither may compliment, flatter, or open with 'I.'

Constraints: if the profile is thin, work from the role and company instead and say you did. Never invent details. The lines must read like a sharp peer noticed something, not like a tool scraped a keyword.

Output: arc in one sentence, self-image in one sentence, the tension, both first lines. Then tell me which line you'd send and why, in one sentence.

Run it when you have a profile but no obvious way in.

PRO MOVE Feed the two output lines into prompt 10's tournament instead of choosing by gut.



PROMPT 04 · RESEARCH BEFORE YOU WRITE

Job posting intelligence

Companies never hire for problems they don't have — a job posting is a confession of strategy written in public. The window is real too: between posting and full ramp, someone is absorbing that work by hand.

You are an expert in organizational tells with deep expertise in reading job postings as strategy documents — because no company hires for a problem it doesn't have.

Inputs — The posting: {paste job posting}. My product: {one sentence}. My buyer: {job title}.

Method: (1) Infer the initiative behind the hire — what this role exists to fix or build. (2) Estimate the pain between now and that hire being fully ramped: who's absorbing the work today, and what it's costing. (3) Connect my product to either the initiative or the gap — pick the stronger one and say why. (4) Write one email opening line built on that connection that never mentions the job posting directly, so the insight feels earned rather than scraped.

Constraints: one line each for initiative, pain, and connection. If the posting is generic boilerplate, say GENERIC and tell me what other signal to pull instead.

Output: initiative, the gap and its cost, the connection, the opener. Rate your own confidence in the inference: high, medium, or low — and if low, say what would raise it.

Run it whenever a target account is hiring near your product's territory.

PRO MOVE Pair it with prompt 05 on the same account; the stack inference plus the hiring read usually triangulate the same initiative.



PROMPT 05 · RESEARCH BEFORE YOU WRITE

The tech stack displacement map

Guessing a prospect's stack wrong ends conversations; reading it right makes you sound like you've already worked there. The map forces confidence labels on every inference, so your email hedges exactly where it should.

You are an elite competitive strategist with deep expertise in inferring a company's tool stack from public exhaust — careers pages, integrations, case studies, developer docs — and turning it into positioning.

Inputs — Company: {name or URL}. Category to map: {CRM, outreach, data, etc.}. My product: {one sentence}.

Method: (1) Infer the tools they likely run in this category, with the evidence for each guess and a confidence level. (2) Decide the honest play: displace, complement, or fill a gap — and defend the choice in one sentence. (3) Write a three-line email built on that play: line one shows I understand their current setup, line two names what that setup can't do, line three connects my product to exactly that.

Constraints: hedge every inference — 'likely,' 'appears to' — because being confidently wrong about their stack ends the conversation. Never trash the incumbent; the reader may have chosen it.

Output: the inferred stack with confidence levels, the chosen play, the three-line email. If the evidence supports no play at all, say NO PLAY and stop.

Run it before positioning against or alongside whatever they already use.

PRO MOVE When the verdict is NO PLAY, log the account in your watch list with the signal that would change it — displacement timing beats displacement force.



SECTION 02 / 05

First lines that don't sound automated

The reader decides in one line whether a human wrote this. These five manufacture that line — feed them the brief from prompt 01.

PROMPT 06 · FIRST LINES THAT DON'T SOUND AUTOMATED

The anti-compliment opener engine

Compliments are what research looks like when it cost nothing — every buyer has read a hundred of them. An observation the reader wants to correct or expand creates the one thing flattery can't: the urge to reply.

You are an elite cold email copywriter with deep expertise in first lines — you know that compliments read as commodity research and that the openers that get replies make an observation the reader wants to respond to.

Inputs — My research or account brief: {paste — ideally the output of prompt 01}.
The reader: {name, title, company}.

Method: (1) From the research, select the three facts closest to money, risk, or time — discard the rest. (2) For each fact, write one first line that states an observation the reader would agree with and feel the urge to correct, expand, or ask 'how do you know that?' (3) For each line, name the psychological hook in three words or fewer — curiosity, recognition, mild challenge.

Constraints: banned openings — congratulations, any compliment, 'I saw,' 'I noticed,' 'I came across,' 'I'm reaching out,' and any sentence containing 'I' at all. Fifteen words max per line. No questions as first lines; observations only.

Before you show me anything, score each line 1-10 on 'would a stranger stop scrolling.'
Kill anything under 8 and replace it.

Output: three lines, each with its hook and score, ranked. One sentence on why your #1 wins.

Run it when every draft keeps sliding back into flattery.

PRO MOVE Keep a swipe file of your 8+ scored lines; your own bank beats any template library within a month.



PROMPT 07 · FIRST LINES THAT DON'T SOUND AUTOMATED

The observation opener

One precisely-aimed fact outperforms three loosely-relevant ones, because precision is the proof of effort. The three registers matter: the same observation lands as sharp to a manager and as try-hard to a CFO.

You are an expert observational writer with deep expertise in the openers senior reps actually send — the kind that imply a question without asking one, so the reader answers it in their head before deciding whether to reply.

Inputs — The fact: {paste the single strongest thing you know about this prospect}. The reader: {title}. My product: {one sentence}.

Method: (1) State what the fact means for the reader's next two quarters — not what it is, what it implies. (2) Compress that implication into one sentence that makes the reader think 'that's exactly our situation' or 'how did they know that.' (3) Produce three registers of the same observation: neutral, slightly provocative, and dry. (4) Flag which register fits which reader — and default to dry for anyone VP and above, because senior readers distrust enthusiasm.

Constraints: no numbers you weren't given. No 'many companies like yours.' The sentence must survive being read aloud in a skeptical voice.

Output: the implication, then the three versions labeled, then your pick with a one-line defense.

Run it when one strong fact has to carry the whole email.

PRO MOVE Default to the dry register above director level — enthusiasm reads as junior in senior inboxes.



PROMPT 08 · FIRST LINES THAT DON'T SOUND AUTOMATED

The peer-pattern opener

When you have no account research, the temptation is to fake it with an invented stat — and buyers can smell a made-up number in one read. Precisely described patterns build the same credibility without the fraud risk.

You are an elite pattern-recognition writer with deep expertise in credibility — you know an invented statistic destroys trust instantly, and that a precisely described pattern builds it just as fast.

Inputs — The reader: {job title} in {industry}. My product: {one sentence}. Current period: {quarter and year}.

Method: (1) Name three pressures this role in this industry credibly faces right now — pressures a real practitioner would nod at, sourced from the structure of the job itself, not from made-up surveys. (2) For the strongest pressure, write two first lines that describe the pattern so specifically the reader assumes you've talked to their peers. (3) Rewrite the better line once more, 20% shorter.

Constraints: no fabricated statistics, no fake client references, no 'we work with companies like yours.' Specificity comes from precision of description, never from invented numbers. The reader should feel seen, not sold to.

Output: the three pressures ranked, the two lines, the final compressed line, and one sentence on what makes it feel researched when it isn't.

Run it for accounts where you have zero account-specific research.

PRO MOVE Refresh these quarterly; a pattern line referencing last year's pressure marks you as running an old sequence.



PROMPT 09 · FIRST LINES THAT DON'T SOUND AUTOMATED

The contrarian opener with a defense

Senior buyers have deleted a thousand agreeable emails; respectful disagreement is pattern-breaking in a way politeness can't be. But a challenge without a defense is clickbait — the follow-up sentence is what makes it a peer move.

You are an expert challenger-sale copywriter with deep expertise in respectful provocation — openers that question a default assumption without insulting the person who holds it.

Inputs — The reader: {job title}. The territory: {the topic your product addresses}. My product: {one sentence}.

Method: (1) List the three assumptions people in this role hold about this territory by default. (2) Pick the one that's most often wrong in practice, and state the evidence or logic that makes it wrong — this is the defense, and if you can't build a defense, pick a different assumption. (3) Write two first lines that challenge it: one direct, one Socratic. (4) Write the follow-up sentence for each — the one that proves the challenge was earned, not clickbait.

Constraints: the challenge must be defensible on a live call with an expert. No straw men. No 'most companies are doing X wrong' without saying what right looks like. Tone: confident peer, never lecturer.

Output: the assumption, the defense, both openers each paired with its follow-up sentence, and which you'd send to this specific title.

Run it for senior buyers who have deleted a thousand safe emails.

PRO MOVE Test the defense on a colleague first: if they say 'prove it' and you can't in one breath, pick a different assumption.



PROMPT 10 · FIRST LINES THAT DON'T SOUND AUTOMATED

The opener tournament

Reps grade their own first lines the way parents grade their own kids. The tournament exists because a named failure — quoted words, specific verdict — is the only feedback that actually changes the next draft.

You are an elite cold email judge with deep expertise in why openers fail — you've reviewed thousands and you can name the exact word where a reader's attention died.

Inputs — My drafted first line: {paste line}. The reader: {title, company}. What I know about them: {one or two facts}.

Method: (1) Judge my line against four failure modes: template smell, self-focus, vagueness, and trying too hard. Name every failure present, quoting the guilty words. (2) Write three challenger rewrites — one shorter, one more specific, one more conversational — each fixing a named failure. (3) Run a bracket: my original versus each challenger, one line of verdict per matchup, and crown a winner.

Constraints: if my original wins, say so and stop — do not rewrite for the sake of rewriting. Fifteen words max per challenger. Every verdict must cite a specific word or phrase, not a vibe.

Output: the failure audit, the three challengers, the bracket verdicts, the winner, and the winner's weakness in one sentence so I know what I'm trading away.

Run every first line through this before it ships — no exceptions.

PRO MOVE Run every opener through it for two weeks; after that the judge lives in your head and drafts start cleaner.



SECTION 03 / 05

Body, value, and the ask

After the opener, every sentence must buy the next one. These five build the middle and the ask — the winner from your opener tournament goes in slot one.

PROMPT 11 · BODY, VALUE, AND THE ASK

The 70-word email operator

Seventy words is not a style choice — it's the length a decision-maker can read on a phone lock screen without scrolling. The one-fact rule matters more: two angles in one email means the reader remembers neither.

You are my elite cold email writer with deep expertise in crafting emails that get responses. The reader decides in four seconds, on a phone, whether this is a template. Your job is to write the one email that survives that test.

Inputs — My product: {one sentence}. The reader: {name, title, company}. Winning opener: {paste from prompt 10, or leave blank and I'll build one}. Research: {2-3 facts}. Proof: {one real customer result}.

Method: (1) If no opener was supplied, pick the single research fact closest to money, risk, or time and open with an observation on it — ignore the other facts; two angles is zero angles. (2) Bridge to the pain in one sentence, in the reader's words, not mine. (3) One proof sentence: customer type plus one number, no superlatives. (4) Close with an interest ask answerable in one word.

Constraints: under 75 words. No 'I,' 'we,' or 'our' in the first sentence. Banned: 'hope this finds you well,' 'quick question,' 'streamline,' 'solution,' 'excited,' and every exclamation point. It should read like the sender's third email of the morning, not the first of their career.

Before you show me anything, score the draft 1-10 on: first line survives a phone-screen skim; pain is in their language; proof survives 'compared to what?'; ask is answerable in one word. Below 8 anywhere — rewrite once.

Output: the final email, a two-word subject line, your four scores.

Run it as your default writer — it takes the winning opener from prompt 10.

PRO MOVE Feed it prompt 10's winning opener and prompt 13's forged proof point — the operator assembles best from pre-tested parts.



PROMPT 12 · BODY, VALUE, AND THE ASK

The cost-of-nothing email

Buyers don't compare you to competitors nearly as often as they compare you to doing nothing — and nothing feels free until someone does the arithmetic. Math shown inline beats adjectives because the reader can check it, and checking is engagement.

You are an expert value engineer with deep expertise in making the status quo feel expensive — problem, agitation, bridge — without a word of melodrama.

Inputs — The reader: {persona}. The pain: {what's broken}. What it plausibly costs them: {time, money, or risk — your best estimate and its basis}. My product: {one sentence}.

Method: (1) State the problem in one sentence the reader would say themselves at a team meeting. (2) Agitate with arithmetic, not adjectives: show the cost of doing nothing in their units — hours a week, deals a quarter, dollars a year — using only the inputs I gave you, and show the math inline so the reader can check it. (3) Bridge to my product in exactly one sentence. (4) Ask for interest, not time.

Constraints: under 100 words. If my cost estimate is weak, say WEAK ESTIMATE and write the email with a question about the cost instead of a claim about it — curiosity beats a number they can dispute. No fear language: 'falling behind,' 'can't afford to,' 'before it's too late' are all banned.

Output: the email, the math shown separately, and one line on whether the claim or the question version is stronger for this reader.

Run it when the pain is real but the reader hasn't priced it yet.

PRO MOVE When the estimate is weak, send the question version — a number they give you is worth ten you claim.



PROMPT 13 · BODY, VALUE, AND THE ASK

The proof point forge

Every claim in a cold email gets cross-examined by a silent skeptic who never tells you they didn't believe you. One number that survives 'compared to what?' outsells five that don't.

You are an elite claims editor with deep expertise in credibility — you know that every proof point gets cross-examined by a silent skeptic, and your job is to write the one sentence that survives it.

Inputs — The raw material: {paste case study, result, or customer story}. The reader: {title, industry}. Can I name the customer? {yes/no}.

Method: (1) Extract every claim from the material and grade each: verifiable, plausible, or puffery. Discard puffery. (2) Take the strongest verifiable claim and forge it into one sentence: customer type (or name, if allowed), the change, one number, the timeframe. (3) Cross-examine your own sentence with the three questions a skeptic asks — 'compared to what?', 'says who?', 'would that happen for me?' — and revise until it survives all three. (4) Write a downgraded version with no number, for when the number invites a fight I don't want.

Constraints: one number maximum. No superlatives, no 'up to,' no unverifiable percentages. If nothing in the material survives cross-examination, say NO PROOF and tell me what evidence to go collect.

Output: the claim audit, the forged sentence, the cross-examination, the no-number fallback.

Run every case study through this before it's allowed into an email.

PRO MOVE Forge your entire case-study library once — one sentence per story — and you'll never scramble for proof mid-draft again.



PROMPT 14 · BODY, VALUE, AND THE ASK

The interest-ask generator

Asking a stranger for thirty minutes is asking them to pay before seeing the menu. Interest asks convert because replying costs one word — and a one-word yes is a foot in the door that a declined meeting never is.

You are an expert in closing sequences with deep expertise in the micro-commitment ladder — you know that asking a stranger for time is asking them to pay before they've seen the menu, and that interest asks convert because they cost one word to answer.

Inputs — The reader: {job title, seniority}. What the email offered: {the value or claim it made}. My real goal: {meeting, reply, referral to the right person}.

Method: (1) Generate five closing asks that request interest, not time — calibrated so replying costs less effort than deleting. (2) Match register to seniority: drier and shorter as you go up; a CFO gets 'worth a look?', not 'would love to connect!' (3) For each ask, predict the most likely reply it produces and whether that reply moves toward my real goal. (4) Flag any ask that's interest-washing — a time ask wearing an interest costume — and fix it.

Constraints: one question mark per ask. No double asks. No 'thoughts?' — it's an ask that assigns homework. Nothing over twelve words.

Output: the five asks ranked, each with its predicted reply, and your pick for this reader with a one-line defense.

Run it when your ask keeps defaulting to 'fifteen minutes Tuesday.'

PRO MOVE Track which ask style gets replies from which seniority for a month; your data will beat any general rule.



PROMPT 15 · BODY, VALUE, AND THE ASK

The executive compression pass

Executives don't read email — they triage it, on a phone, deciding in one glance whether this touches money, risk, or time. The forward test is the real goal: the best executive cold email is one their EA or VP forwards down with 'look into this.'

You are an elite executive communications editor with deep expertise in how C-level readers process email: on a phone, between meetings, deciding in one glance whether this concerns money, risk, or time — and deleting everything that doesn't.

Inputs — My draft: {paste_email}. The reader: {title}. The single outcome my product drives: {one_phrase — revenue, cost, risk, or speed}.

Method: (1) Identify the one sentence in my draft an executive would actually care about; if it isn't the first sentence, that's the diagnosis. (2) Rebuild: outcome first, mechanism second, proof third — and delete every feature word, because executives buy outcomes and delegate mechanisms. (3) Cut to under 60 words. (4) Make the ask answerable in one word or one forward.

Constraints: no jargon the reader's board wouldn't use. No flattery about their company. The forward is a win — write the email so it's easy to forward with 'worth a look?' typed above it.

Before you show me anything, read the draft as a skeptical CFO and delete one more sentence.

Output: the compressed email, what you cut and why in two lines, and the one-word reply you're engineering for.

Run it on anything headed to a VP or C-level inbox.

PRO MOVE Write the email you want forwarded, not read — assume your true audience is one level below the recipient.



SECTION 04 / 05

Subject lines and the second look

Most of the list never opens the email — the subject line is where the campaign is actually won. These five work the envelope, then engineer the second chance.

PROMPT 16 · SUBJECT LINES AND THE SECOND LOOK

The internal-email subject engine

Open rates are won before the email exists — in the half-second where a subject line either looks like work or looks like marketing. Lines that read internal pass the triage; lines that read clever get filed with the other clever ones.

You are an elite deliverability-aware copywriter with deep expertise in the inbox glance — you know the highest-opening cold subject lines are the ones that look like they came from inside the building.

Inputs — What the email is about: {topic}. The reader: {job title}. One specific word from my research: {a company term, tool name, or initiative — optional}.

Method: (1) Write ten subject lines under four words, lowercase or sentence case, that could pass as a colleague's email — and are still honest about what's inside, because a bait-and-switch open is worse than no open. (2) Where I gave you a research word, use it in at least three lines; specificity beats cleverness. (3) Sort the ten into two piles: 'curiosity' and 'relevance.' (4) Kill any line that smells like marketing — punctuation tricks, brackets, 'quick,' 're:' fakery, urgency theater.

Constraints: no clickbait, no emoji, no colons doing headline work. The test: would the reader feel tricked after opening? If yes, the line is dead.

Output: the surviving lines in their two piles, your top pick from each pile, and one line on which pile fits this reader's seniority.

Run it once per campaign, then A/B the top two with prompt 19.

PRO MOVE Pull the research word from their own vocabulary — a team name, a tool, an initiative — specificity beats wit every test.



PROMPT 17 · SUBJECT LINES AND THE SECOND LOOK

The subject and preview pairing system

The inbox shows two lines, but most reps only write one and let the second happen by accident. A preview that completes the subject's thought is free real estate almost nobody claims.

You are an expert in inbox real estate with deep expertise in the two-line unit every mobile inbox actually displays: subject plus preview text, working as headline and subhead.

Inputs — My subject line: {subject}. The email's first sentence as written: {paste it}. The reader: {title}.

Method: (1) Diagnose what my current first sentence shows as preview text — because the inbox will use it whether I designed it or not. (2) Write five preview options under 80 characters that complete the subject's thought rather than repeat it — the subject opens a loop, the preview leans into it without closing it. (3) For each pair, write the reader's likely half-second reaction in five words. (4) Flag any pair where subject and preview together over-promise what the email delivers.

Constraints: the preview must be the actual first sentence of the email — no hidden preheader hacks; they smell like marketing automation. No repeating a single word from the subject line.

Output: the diagnosis, the five pairs with reactions, the winning pair, and the email's new first sentence.

Run it after the subject line is chosen — the preview is half the envelope.

PRO MOVE Check what your last five sent emails showed as preview text — the diagnosis is usually humbling and instantly fixable.



PROMPT 18 · SUBJECT LINES AND THE SECOND LOOK

The automation-smell audit

Deliverability kills silently and template-smell kills politely — the buyer just deletes and you never learn why. This is the last gate: the difference between an email a person wrote and an email a system sent is usually six phrases.

You are an elite deliverability and authenticity auditor with deep expertise in the two ways cold email dies: the filter that blocks it and the reader who smells the mail merge.

Inputs — The final draft: {paste email, subject included}.

Method: (1) Sweep one — filter risk: flag spam-associated phrases, link patterns, formatting tells, and image or tracking habits that hurt placement. (2) Sweep two — human smell test: flag every phrase a real person would never type to a colleague: 'I hope this email finds you well,' 'I wanted to reach out,' 'as a leader in the space,' template transitions, AI-flavored symmetry like 'not just X, but Y.' (3) For every flag, give a natural replacement in the sender's plain voice. (4) Rewrite the clean version in full.

Constraints: preserve the argument and the ask exactly — this is a sweep, not a rewrite of strategy. If the draft is already clean, say CLEAN and change nothing; a false positive costs me voice.

Output: both flag lists with replacements, the clean version, and a one-word verdict on the draft's biggest remaining risk: filter, smell, or none.

Run it on the final draft, every time — this is the last gate before send.

PRO MOVE Run it on your best-performing sequence too; the phrases that were fine last year are this year's tells.



PROMPT 19 · SUBJECT LINES AND THE SECOND LOOK

The single-variable test lab

Most reps' A/B tests are two different emails, which teaches nothing except which email won. One variable per variant is slower and it's the only version that compounds into actual knowledge.

You are an expert experimentation designer with deep expertise in cold email testing — you know that changing two things at once teaches you nothing, and that most 'A/B tests' reps run are just two different emails.

Inputs — The control email: {paste email}. My send volume this campaign: {rough number}. What I most want to learn: {opens, replies, or meetings}.

Method: (1) Identify the three highest-leverage variables in the control: opener, proof point, ask, or subject — pick the three that matter most for what I want to learn. (2) Build one variant per variable, changing exactly that variable and holding every other word identical. (3) For each variant, state the hypothesis in one sentence — 'if X, then Y, because Z' — and the result that would confirm it. (4) Tell me honestly whether my volume can detect the difference; if the sample is too small to learn anything, say TOO SMALL and recommend the one variant worth sending anyway.

Constraints: one variable per variant, no exceptions. No vanity tests on cosmetics.

Output: the three variants with hypotheses and confirm conditions, the volume verdict, and the test you'd run first.

Run it when you're about to guess — testing beats taste.

PRO MOVE Keep a one-page log of confirmed results; ten clean tests make you the most data-armed rep on the floor.



PROMPT 20 · SUBJECT LINES AND THE SECOND LOOK

The re-send reframe engine

A no-open means the envelope failed, not the offer — the reader never saw your pitch to reject it. The re-send is a free second first-impression, and it stays free only if nothing about it references the first attempt.

You are an elite second-chance strategist with deep expertise in the re-send: the reader never saw the first email, so this is a fresh first impression that costs nothing — unless it repeats itself.

Inputs — The ignored email: {paste it}. Days since send: {number}. Anything new I've learned since: {optional}.

Method: (1) Diagnose the most likely reason it wasn't opened: subject, timing, sender, or luck — and say which, because the fix depends on the diagnosis. (2) Find the second door: the same underlying value approached from a different angle — a different pain, a different proof, or a different moment in their week. (3) Write the re-send: new subject line, new first line, same offer underneath. (4) Verify it never references the ignored email — no 'bumping this,' no 'in case you missed it,' no wounded tone; the reader owes me nothing.

Constraints: the re-send must stand completely alone. If I gave you new information, it leads. Shorter than the original, always.

Output: the diagnosis, the second door in one sentence, the re-send email with subject line, and the send-day recommendation with one line of reasoning.

Run it four to six days after a no-open — same value, different door.

PRO MOVE Send the reframe at a different time of day than the original; if the first died at 8am Monday, the second shouldn't.



SECTION 05 / 05

Quality control before send

The last pass is where good emails become sent emails and great ones become replies. Five gates — the drafts that clear all five are the ones worth your send reputation.

PROMPT 21 · QUALITY CONTROL BEFORE SEND

The delete tribunal

Every sentence a reader skips lowers the odds they reach your ask — length isn't the enemy, unearned length is. Sentence two gets special scrutiny because it's where writers stop talking about the reader and start talking about themselves.

You are an elite editor with deep expertise in ruthless compression — you preside over a tribunal where every sentence in a cold email must justify its existence to a reader who owes it nothing.

Inputs — The draft: {paste_email}. The reader: {job_title}.

Method: (1) Put each sentence on trial, in order. The charge: 'a busy {job_title} would not read the next sentence because of you.' (2) For each, deliver a verdict — EARNs ITS PLACE or DELETE — with one line of reasoning that quotes the sentence's weakest words. (3) Special scrutiny for sentence two: it's where most cold emails die, because it's where the writer starts talking about themselves. (4) Assemble the final email from survivors only, smoothing exactly the seams the deletions created — no new material.

Constraints: when in doubt, delete — a short email that lands beats a complete one that doesn't. But never delete the proof point or the ask; if either is weak, flag it for prompts 13 or 14 instead of removing it.

Output: the trial record, the survivor email, and the word counts before and after.

Run it when the draft feels done but heavy — every sentence testifies.

PRO MOVE Run the tribunal on a competitor's cold email you received — watching it dismantle someone else's draft teaches the eye faster.



PROMPT 22 · QUALITY CONTROL BEFORE SEND

The peer-voice rewrite

Correct-but-lifeless is the most common failure state of edited email — every rough edge that made it sound human got sanded off. The peer register works because busy people trust senders who write like their time matters.

You are an expert voice editor with deep expertise in the register that actually gets replies: a busy, competent peer on their third coffee — someone who respects the reader's time because they value their own.

Inputs — The draft: {paste email}. One sentence about how I actually talk: {e.g., direct, dry, midwestern, no exclamation points}.

Method: (1) Read the draft aloud in your head and mark every phrase that sounds like it was written for a grade: corporate connective tissue ('furthermore,' 'as such,' 'I wanted to reach out'), over-hedging, symmetrical AI cadence. (2) Rewrite in the peer voice — contractions allowed, fragments allowed where a person would use one, warmth through brevity rather than adjectives. (3) Keep the substance and the ask identical; this is a voice pass, not a strategy pass. (4) Check the rewrite against my one-sentence voice note and adjust.

Constraints: do not add casualness that curdles into unprofessional — no slang, no jokes unless my voice note asked. The test: could the reader tell this email from one I typed myself? If yes, keep going.

Output: the marked-up diagnosis, the rewrite, and the three biggest voice changes in one line each.

Run it when the draft is correct but lifeless — polish is the problem.

PRO MOVE Give it a real voice note about how you talk; the rewrite is only as much 'you' as the input you feed it.



PROMPT 23 · QUALITY CONTROL BEFORE SEND

The hostile reader simulation

You will never watch a prospect read your email — this is the closest available substitute, and it's brutally cheap. The stop-point matters more than the verdict: knowing the exact word where attention died is a repair map.

You are a battle-tested {job title} at a {company size} {industry} company — skeptical, interrupted, and currently on email number forty-one of the morning. You have deep expertise in deleting cold email, because you do it thirty times a day. You are not playing nice for my benefit.

Inputs — The email: {paste it, subject line included}.

Method: (1) React in real time, line by line — your actual internal monologue, including the uncharitable parts. (2) Mark the exact word where you stopped reading, if you stopped. (3) Deliver the verdict: delete, ignore, forward, or reply — and what you'd type if you replied. (4) Then break character and, as a coach, name the one change with the highest odds of flipping your verdict — just one, because a list of ten fixes means none get made.

Constraints: stay hostile until step four. Do not soften the monologue to spare my feelings — a polite simulation is a useless one. If the email would genuinely earn a reply, say so and don't invent criticism.

Output: the monologue, the stop-point, the verdict, the one change.

Run it on high-stakes sends before they leave the drafts folder.

PRO MOVE Run high-stakes sends past two different personas — the VP and the end user stop reading at different words.



PROMPT 24 · QUALITY CONTROL BEFORE SEND

The one-question distiller

When a full pitch has failed twice, more pitch is not the answer — less is. A single well-aimed question works because answering it costs less than deleting it, and their answer restarts a conversation your pitches couldn't.

You are an elite minimalist with deep expertise in the shortest format in cold email: the one-question message — no pitch, no proof, no signature paragraph, just a question so precisely aimed that answering it starts the conversation everything else failed to.

Inputs — My product and the pain it addresses: {two sentences}. The reader: {title, company}. What I know about them: {one or two facts}.

Method: (1) Find the question only a genuine insider would think to ask — specific enough that a generic vendor couldn't have written it, open enough that a one-line answer is easy. (2) Draft three candidates: one about their current process, one about a tradeoff they're living with, one about a number they probably track. (3) Test each: does answering cost less effort than deleting? Does the answer naturally lead somewhere my product matters? Kill any candidate that fails either. (4) Wrap the winner in a total email of 35 words or fewer.

Constraints: no preamble before the question beyond one short sentence. No 'just curious.' The question must be answerable from the reader's memory — nothing they'd have to look up.

Output: the three candidates with test results, the final email, the follow-up move for each possible answer.

Run it when nothing else has landed with this account — the nuclear option for noise.

PRO MOVE Save this for genuinely stuck accounts; used as a first touch it spends your best move too early.



PROMPT 25 · QUALITY CONTROL BEFORE SEND

The sequence architect

Sequences written touch-by-touch read like a rep improvising; sequences designed as one argument read like a professional with a plan. Assigning each touch one job — and holding cards back on purpose — is what gives touches two and three a reason to exist.

You are an elite outbound sequence designer with deep expertise in multi-touch architecture — you know a sequence is one argument delivered in installments, and that touch one decides whether touches two and three have anywhere to go.

Inputs — The reader: {job title}. The topic: {what I'm selling and the pain}. My proof inventory: {the 2-3 proof points I have}. Sequence length: {number of touches}.

Method: (1) Design the argument arc across all touches before writing anything: touch one opens the problem, later touches add proof, then urgency, then the graceful exit — assign each touch one job. (2) Write touch one in full: complete on its own, but deliberately holding two cards back — name which proof and which angle you're reserving and for which touch. (3) Write the first line of every remaining touch, so the arc is visible. (4) Write the breakup touch's ask — the exit that leaves the door open without begging.

Constraints: no touch may reference the reader ignoring a previous one. Each touch must stand alone for a reader seeing only that touch. Touch one under 80 words.

Output: the arc map with one job per touch, touch one in full, the first lines, the breakup ask.

Run it when writing touch one — a sequence is designed, not accumulated.

PRO MOVE Write the breakup touch first — knowing how the sequence ends with dignity changes how confidently touch one opens.



THE FULL LIBRARY

This is 25. The library has *2,900+*.

Every stage of the deal — prospecting, discovery, objections, negotiation, closing, renewals — curated and tested for AEs, SDRs, and BDRs. New workflows added continually. \$4.99 a month, or \$99 once for lifetime access.

[Browse the library](#)

PROMPTIFI.CO