



BY LESZEK LAMMEL

# 43

## humanization rules.

How to write AI-generated cold email and LinkedIn copy that doesn't read like AI. Validated on 50,000+ outbound messages.

**The rules don't fix it. Here's what does.**

## START HERE

# Why AI cold copy fails before you fix the words.

Before the 43 rules, the framing. Most humanization advice optimizes against detection - banned words, sentence-rhythm checks, length constraints. That misses the actual problem. AI copy fails for three structural reasons. The rules in this checklist patch the symptoms; reading these three first patches the cause.

## FAILURE MODE 1

**AI produces the statistical mean.**

The average of every framework, every blog, every corporate content marketing piece it's been trained on. Smooth. Structurally predictable. Inoffensive. Nobody replies to the mean.

## FAILURE MODE 2

**AI describes states. Humans describe movement.**

Every AI message describes a state - "our product does X" - without implying movement from one state to another. Human readers, even reading a 50-word email, are wired for narrative: there was a before, now there's an after.

## FAILURE MODE 3

**AI is structurally trained to hedge.**

"Could be argued." "Potentially." "May help you." Hedging is the linguistic signature of a system that can't be held accountable for its claims. Humans notice this even when they can't name it.

## THE 43 RULES · TIER 1

# Banned words.

## Tier 1 Kill on sight

These words spiked 200-6700% in AI-generated text (FSU / COLING 2025). A single one can trigger inbox suspicion.

1 **delve**, delves, delved, delving

2 **tapestry**

3 **groundbreaking**

4 **realm**

5 **seamlessly**

6 **underscore**, underscores, underscoring

7 **showcasing**, showcases

8 **robust**

9 **pivotal**

10 **utilize** - use "use"

11 **leverage** as a verb - use "use"

12 **holistic**

13 **revolutionary**

14 **transformative**

15 **nuanced**

16 **multifaceted**

17 **paradigm**

18 **bespoke**

19 **intricate**, intricacies

20 **paramount**

## THE 43 RULES · TIER 2

# Structural AI tells.

## Tier 2 Rewrite the pattern

The contrastive reframe is the single most recognizable AI writing pattern. These ten patch the structural tells the model defaults to.

- 21 Never write **It's not X, it's Y** - state what something IS directly.
- 22 Never write **The problem isn't X, it's Y** - same pattern, different surface.
- 23 Never use **Not just X but also Y** - two short sentences instead.
- 24 Never start a sentence with a present participle ("Building on this...", "Thinking about...") - start with the subject.
- 25 Never use three bullet points of equal length - vary lengths, make one a fragment.
- 26 Never use perfect parallel structure in lists - break parallelism intentionally.
- 27 Never use rhetorical question into answer ("So what's the play? The play is...") - just say the play.
- 28 Never use closing call to reflection ("Something to think about.") - end on a specific.
- 29 Never use the corrective framing **Not just X but Y** in outbound - it judges the recipient's present.
- 30 Never lead with capability lists - lead with what changed for the reader, not what you can do.

## THE 43 RULES · TIERS 3-4

# Rhythm and hooks.

## Tier 3 Sentence rhythm - force burstiness

AI text has abnormally low burstiness (CMU / PNAS 2025). Sentences cluster around the same length. Variation is the unlock.

- 31 Never write 3 consecutive sentences within 3 words of each other in length. Mix short (under 8 words) with medium (15-20 words). Use fragments when natural.
- 32 The short sentence is the release after the build, not the default. A piece in all-short sentences drones at a different pitch.
- 33 Always use contractions. It's, we've, don't, can't, won't, they're.
- 34 Mix register within a single message. One formal observation + one casual aside.
- 35 Use occasional self-correction ("Actually, scratch that - the real issue is...") sparingly.

## Tier 4 Hook + opening rules

First 210 characters of LinkedIn (the "...see more" cutoff). First 100 characters of email (Gmail Gemini reads this for deprioritization).

- 36 Never open with "I hope this email finds you well." Instant skip.
- 37 Never open with "I noticed that [Company] is..." or "I came across your profile..." - both are AI tells the recipient sees 12x a day.
- 38 Front-load value in the first 100 characters. Gmail Gemini reads this for deprioritization. Burying the value buries the message.
- 39 One specific detail per opener that only the sender would know about this prospect. Generic personalization (their company name, their title) doesn't count.
- 40 No questions in the opener of a cold email or connection request. Questions trigger filter flags.

## THE 43 RULES · TIER 5 + THE DATA

# CTAs, and what actually correlates with replies.

## Tier 5 CTA rules

One ask, one shot.

- 41 One ask per message. The second CTA dilutes the first - cut it.
- 42 The ask must be specific. "What's your experience with Clay at 5k rows?" beats "Thoughts?" by 3-5x in reply rates.
- 43 No sign-off beyond first name in cold email. "Best regards / Looking forward / Warmly / Cheers" all read like business letters from 2014.

## Lavender reply data

The Lavender cold-email analysis surfaces what moves reply rates independent of the words you pick. Use these as constraints on every draft.

Under 75 words

+83% replies

3rd to 5th grade reading level

+17% replies

Casual tone

+23% replies

Informative tone

-26% replies

First opened on mobile

80% of cold emails

## THE FRAMEWORKS BEHIND THE RULES

# Three classics. Internalize these and the checklist becomes intuition.

The 43 rules patch surface-level AI tells. These three frameworks explain why the rules work - and what to do when the rules don't cover the situation in front of you.

**SCHWARTZ****5 levels of awareness**

Cold email almost always lands in Level 1 (Unaware) or Level 2 (Problem-Aware). Lead with the prospect's world, never your product. The fatal mistake is pitching to someone who hasn't named the problem yet.

**SUGARMAN****The slippery slide**

The only job of any line of copy is to make the reader read the next line. Every sentence earns the next. Plant open loops. Front-load the most arresting truth you have. If a sentence could be cut and the reader would still continue, it should be cut.

**PROVOST****Sentence rhythm**

"Five-word sentences are fine. But several together become monotonous." Vary length. Short, then long, then short again. The short sentence is the release. The long one is the build. Music, not drone.

## HOW TO USE THIS

# Two workflows. One checklist.

## Short-form copy

Cold email, LinkedIn DMs, comments. Paste these rules directly into the generation prompt - the model follows them while writing. Add Tier 1 and Tier 5 to a Layer 2 (constraints) block per the 5-layer prompt architecture in the appendix.

## Long-form copy

Newsletters, blog posts, sequences. Use this as an **edit pass** after the draft, not during. Embedding rules in the generation prompt for long-form produces compliant text that reads dead - the model optimizes individual sentences instead of the whole piece. Write first. Check rules second.

## For Clay tables generating messages at scale

Run every message through a second AI column with this checklist as a YES/NO filter. Any YES (a rule was broken) goes to manual review before export. Sample 10% of every batch by hand even with automated checks - no automation catches register problems.

### Want this **built into your outbound system?**

The prompts. The Clay columns. The QA gates. The way these rules survive in production at 1,000+ messages a day. I build the system end-to-end - six-week sprint, owned by your team after.

[Book a 15-min fit call → lammel.ai](https://lammel.ai)

# The full ruleset and the frameworks behind it.

The summary above is the working surface. The full source text follows: the rules of what to avoid, then the frameworks of how to think. Use them as the canonical reference - paste into prompts, embed into Clay columns, copy into n8n nodes.

# The Ruleset

**Purpose:** Two different uses depending on content type.

- **Short-form copy (cold email, LinkedIn messages):** Embed these rules directly in the generation prompt. The model follows them while writing.
- **Long-form copy (newsletters, blog posts, sequences):** Use as an *edit pass* - run after the draft is written, not during. Embedding rules in the generation prompt for long-form produces compliant text that doesn't read, because the model optimizes individual sentences instead of the whole piece. Write first. Check rules second.

The master reference for banned words, structural rules, and prompt architecture. Copy into Clay formulas, n8n nodes, and system prompts.

## 1. Banned Words & Phrases

### Tier 1: Instant AI Flags (NEVER use)

These words spiked 200-6700% in AI-generated text (COLING 2025, FSU). Any single one can trigger suspicion.

**Words:** delve, delves, delved, delving, tapestry, camaraderie, palpable, intricate, intricacies, underscores, underscore, underscoring, showcasing, showcases, groundbreaking, realm, garnered, surpassing, surpasses, advancements, aligns, boasts, comprehending, nuanced, multifaceted, synergy, holistic, revolutionary, transformative, paradigm, pivotal, robust, seamlessly, bespoke, utilize, hitherto, paramount, encompass, esteemed, resonate, unease, plethora, myriad, commendable, meticulous, noteworthy, invaluable

**Opening phrases (instant skip by recipients):** - "I hope this email finds you well" - "In today's ever-evolving / fast-paced / digital world..." - "In the ever-evolving landscape of..." - "I came across your profile and was impressed..." - "I noticed that [Company] is..." - "I wanted to reach out because..." - "As we navigate the complexities of..." - "It's important to note that..." - "Certainly!" / "Absolutely!" / "Great question!"

**Closing phrases:** - "In summary," / "To sum up," / "In conclusion," - "In essence," / "At the end of the day," - "Don't hesitate to reach out" - "Looking forward to connecting" - "Let me know if you have any questions"

### Tier 2: AI Buzzwords (avoid or replace with specifics)

cutting-edge, innovative, comprehensive, dynamic, iterative, confluence, framework, facet, trajectory, spectrum, underpinning, at the forefront of, game-changer, game-changing, next-generation, best-in-class, scalable, agile, strategic, empower, empowering, unlock, unleash, harness, leverage (as verb), streamline, optimize, supercharge, foster, elevate, enable, spearhead, catalyze, synergize

**Replace with:** specific numbers and concrete outcomes. "Cut onboarding from 3 weeks to 4 days" not "streamline your onboarding process."

### Tier 3: Structural AI Tells (rewrite the pattern, not just the word)

| AI Pattern                                                | Human Alternative                                                                                                    |
|-----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| "Leveraging X, we achieve Y..."                           | "We use X to get Y."                                                                                                 |
| "Here's the thing..." / "Here's what nobody tells you..." | Just say the thing.                                                                                                  |
| "That's a fair concern, but..."                           | Skip the validation, answer directly.                                                                                |
| "Not only X but also Y"                                   | Two short sentences.                                                                                                 |
| "It could be argued that..."                              | State the position.                                                                                                  |
| Present participle openers ("Building on this...")        | Start with the subject.                                                                                              |
| Three bullet points of equal length                       | Vary lengths. Make one a fragment.                                                                                   |
| Perfect parallel structure in lists                       | Break parallelism intentionally.                                                                                     |
| "It's not X, it's Y" / "The problem isn't X, it's Y"      | <b>State what something IS directly. The contrastive reframe is the single most recognizable AI writing pattern.</b> |

### Tier 4: Hedge Words (kill AI's signature uncertainty)

quite, rather, somewhat, possibly, potentially, arguably, presumably, seemingly, apparently, relatively, fairly, essentially, basically, "it could be argued"

**Rule:** If you can delete the hedge word and the sentence still works, delete it.

---

## 2. Structural Rules (Anti-Detection)

These address the Carnegie Mellon PNAS 2025 finding that AI text has abnormally low burstiness (uniform sentence rhythm).

### Sentence Rhythm

**RULE:** Never write 3 consecutive sentences of similar length.

Pattern to follow:

- Short (under 8 words)
- Medium-long (15-20 words)
- Short (under 10 words)
- Then vary freely

Example:

"Pipeline's broken. Your reps are booking meetings but deals stall at proposal stage - same pattern we saw at three other Series B companies last quarter. Classic scaling problem."

**Warning - uniform short sentences are their own AI tell.** Short sentences only land hard when they follow longer ones. A piece written entirely in short sentences drones just as much as one written in long ones - just at a different pitch. The Provost principle is about *contrast*, not brevity. The short sentence is the release after the build. Apply this as a rhythm check, not a length target. If every sentence is short, something is wrong.

### Paragraph Structure

- Cold emails: 2-4 sentences total. No paragraphs - just lines.
- LinkedIn messages: 3-6 sentences. One line break max.
- Newsletters: Vary paragraph length (1-sentence paragraph, then 3-sentence, then 2).
- Never use the same structural template for 2 messages in a row.

### Contractions & Register

- Always use contractions: it's, we've, don't, can't, won't, they're
- Mix register within a message: one formal observation + one casual aside
- Use fragments when natural: "Same story everywhere." "Not great."
- Occasional self-correction: "Actually, scratch that - the real issue is..."

---

## 3. System Prompt Architecture

### The 5-Layer Model

Every outbound copy prompt should follow this order:

[LAYER 1: PERSONA – one sentence, communication style not just title]  
You are [Name], [Title] at [Company]. You write short, direct messages the way you'd text a colleague – no jargon, no fluff.

[LAYER 2: CONSTRAINTS – what NOT to do, paste banned words]  
Never use these words: [paste Tier 1 list].  
Never open with a question or flattery.  
Never use transition phrases.  
No adjectives that don't add specifics.

[LAYER 3: STRUCTURAL RULES – force burstiness]  
Sentences: Mix short (under 8 words) with medium (12–18 words).  
Never write three sentences of the same length in a row.  
Use fragments when they feel natural.  
Prefer active voice with a direct subject.  
Always use contractions.

[LAYER 4: OUTPUT FORMAT – be explicit]  
Output: [X] words max. No subject line unless asked. No sign-off.  
No placeholder text. One version only. No explanation or commentary.

[LAYER 5: FEW-SHOT EXAMPLES – 3 real samples]  
Here are three examples of how [Name] writes:  
---EXAMPLE A---  
[paste real message from sender]  
---END A---  
---EXAMPLE B---  
[different prospect, same voice]  
---END B---  
---EXAMPLE C---  
[third real message]  
---END C---

Now write: [task with prospect data]

## Why This Order Matters

1. **Persona** anchors the register before the model generates anything
2. **Constraints** override the model's RLHF-trained instincts to use "impressive" vocabulary
3. **Structural rules** address the CMU burstiness problem
4. **Format** prevents the model from adding meta-commentary
5. **Examples** are the single most powerful technique – a weak prompt with 3 real examples beats a perfect prompt with no examples

## Voice Extraction (Run Once Per Sender)

Before building prompts for a new sender, extract their style:

Analyze these 5 emails I wrote. Extract:

1. Average sentence length distribution (short/medium/long)
2. Vocabulary level (simple/mixed/elevated)
3. Emotional register (warm/neutral/direct/dry)
4. Structural patterns I repeat
5. Words that are distinctive to me
6. What I never do

[Paste 5 real emails from the sender]

Output as a 6-bullet style card. I will use this in another prompt.

Paste the resulting style card into Layer 1 of your generation prompt.

## 4. Temperature & Sampling

### By Content Type

| Content Type                     | Temperature | Top_p | Notes                                                  |
|----------------------------------|-------------|-------|--------------------------------------------------------|
| Cold email body                  | 0.7-0.8     | 0.9   | Balanced: readable + natural                           |
| Email first-line personalization | 0.7-0.8     | 0.9   | High enough for variation, low enough to stay specific |
| LinkedIn connection request      | 0.6-0.7     | 0.85  | Short format - too high = hallucinated details         |
| LinkedIn message body            | 0.7-0.8     | 0.9   | Same as email body                                     |
| Newsletter intro                 | 0.8-0.9     | 0.95  | Needs personality and surprise                         |
| Subject lines (batch)            | 0.9-1.0     | 0.95  | Max variation; you filter the best                     |
| Follow-up sequence variants      | 0.8         | 0.9   | Prevents repetition across touches                     |

### Model-Specific Notes

- **Claude:** Set temperature OR top\_p, not both (Anthropic guidance). Use temperature only.
- **GPT:** Never use temperature 0 for outreach - it produces the most "GPT-sounding" output. Above 1.2 starts producing grammar errors in short text.
- **Both:** Higher temperature alone is not humanization. It just adds randomness. You need the structural rules + banned words + examples.

## 5. Post-Generation Self-Check

### Embed in Every Generation Prompt (After the Task)

Add this at the end of any copy generation prompt:

Before outputting, verify:

1. Zero words from this list appear: delve, tapestry, groundbreaking, intricate, underscore, realm, robust, synergy, seamlessly, nuanced, transformative, pivotal, holistic, revolutionary, comprehensive, aligns, showcasing, garnered, advancements, surpassing, boasts, emphasizing, paradigm, bespoke, utilize, paramount, resonate
2. No two consecutive sentences are within 3 words of each other in length
3. At least one contraction is used
4. The opening line contains something specific to THIS prospect (not generic)
5. Total word count is under [limit]
6. No sentence starts with "I noticed" or "I came across" or "I wanted to"

If any check fails, rewrite that part before outputting.

### Clay Two-Column Pattern

For Clay tables generating messages at scale:

Column A: Generate message (Claude/GPT formula)

Column B (QA check):

```
=AI("Does this message contain ANY of these words: delve, tapestry, groundbreaking, intricate, underscore, realm, robust, synergy, seamlessly, nuanced, transformative, pivotal, holistic, revolutionary, comprehensive? Reply YES or NO only.
```

```
Message: " & [Column A])
```

Filter: Show only YES rows for human review before export.

### Manual Sampling Rule

Even with automated checks: **manually read 10% of every batch** before sending. No automation catches register problems (message sounds like a press release instead of a person).

## 6. Content-Type Quick Reference

### Cold Email (under 80 words)

- Front-load value in first 100 characters (Gmail Gemini reads this)
- One specific detail about the prospect
- One clear CTA (single ask, not options)
- No sign-off beyond first name
- Reads like a text message, not a business letter
- Use spintax: minimum 5 blocks per email for variation at scale

### LinkedIn Message (under 300 chars for connection, under 500 for InMail)

- No "I wanted to connect" / "I came across your profile"
- Reference something they posted, shared, or changed recently
- Connection requests: no question marks, no pitch
- InMails: subject line lowercase, 3-6 words, specific reference
- Sound like one human talking to another, not a sales rep

### Newsletter

- First sentence: punchy, under 8 words
- Second sentence: context, 15-25 words
- Third sentence: sharp pivot, under 10 words
- Never "excited to share" / "thrilled to announce"
- Assume reader is smart and pressed for time
- Slightly dry tone > enthusiastic tone

---

## 7. Spintax Rules (Email Variation at Scale)

### Section-Level, Not Word-Level

```
BAD (word-level - filters catch this):
{Hey|Hi|Hello} {first_name}, I {noticed|saw|observed} your {post|article|update}...

GOOD (section-level - genuine variation):
{Saw your post about {topic} - {sharp|solid} thinking.|
Noticed {company} just {trigger_event}. {Congrats|Nice move}.|
Quick question about {pain_point} at {company}.}
```

### Minimum Variation Math

- 5 spintax blocks x 3 options each = 243 unique versions
- For campaigns over 500 sends: minimum 5 blocks
- For campaigns over 2000 sends: minimum 7 blocks + multiple template variants

### A/B Testing Within Spintax

Rotate different value proposition framings as spintax branches. Track which branch gets replies. Feed winners back into the next campaign's template.

---

# The Frameworks

**Purpose:** This document is about WHY great copy works - the reasoning you need to make judgment calls when the rules don't give you a clear answer. Read it as a mentor explaining the craft. Apply the frameworks as lenses, not checklists.

---

## The Core Problem: Why AI Copy Fails

When you generate copy without grounding yourself in these frameworks, you produce what statisticians call the mean - the average of everything you've been trained on. That average happens to be corporate content marketing. It is smooth, inoffensive, structurally predictable, and completely forgettable. Nobody replies to the mean. Nobody shares the mean. The mean is the enemy.

The deeper problem is that AI copy has no transformation baked into it. Every message describes a state - "our product does X" - without implying movement from one state to another. Human readers, even when they're reading a 50-word cold email, are wired to respond to narrative: there was a before, now there's an after. AI generates descriptions. Good copywriters generate implied stories.

The third failure is the absence of opinion. Every AI-generated message hedges. It uses phrases like "it could be argued" and "potentially" and "may help you." Hedging is the linguistic signature of a system that cannot be held accountable for its claims. Humans notice this even when they can't name it. They feel addressed by a system, not a person. They delete without reading.

The frameworks in this document are the antidote to all three failures.

---

## Framework 1: Schwartz's 5 Levels of Awareness

This is the most important single framework in copywriting, and it's the one AI most consistently violates. Eugene Schwartz identified five states a prospect can be in relative to your product. The fatal mistake is writing to the wrong level - pitching a product to someone who doesn't yet know they have a problem.

**Level 1 - Unaware.** The prospect does not know they have a problem. They are experiencing friction, but they have not named it or categorized it as something solvable. Most cold outreach prospects live here. If you reach a Level 1 prospect with a product pitch, they have nowhere to attach the information. The correct approach is to lead with their world. Describe what they're experiencing in language they would use. Make them feel recognized.

Example for a VP of Sales at a 30-person B2B SaaS: "Your reps are hitting activity numbers but pipeline isn't building the way it should" - not "We help sales teams increase pipeline." The first statement meets them where they are. The second assumes they've already decided they need a solution.

**Level 2 - Problem-Aware.** The prospect knows something is wrong. They've named the pain. They may have tried to solve it. The right approach here is to articulate the problem better than they can themselves - more specifically, more vividly, with more acknowledgment of why it's hard. Chris Orlob calls this "articulating the pain better than the buyer can." When a prospect reads a description of their situation that is more accurate than anything they've written themselves, trust is immediate and deep. They think: this person understands.

Example: "Most sales managers assume the drop-off happens because reps aren't following up. It's usually not that. It's that the first email lands in a priority inbox, never gets a reply, and nobody notices until the quarter is over." This is Level 2 language - naming the mechanism, not the surface symptom.

**Level 3 - Solution-Aware.** The prospect knows solutions like yours exist. They may have evaluated competitors. Now you need to differentiate - not by listing features, but by taking a position on what the right approach is and why the others fall short. Never name a competitor. Just articulate the principle that makes you different.

**Level 4 - Product-Aware.** They know you. They've read your website, attended a webinar, or talked to someone on your team. They haven't bought. The friction is usually internal - budget, timing, competing priorities. Here you're moving them over a decision threshold with social proof, risk reduction, or a sharp offer.

**Level 5 - Most Aware.** They're ready. Just close. Don't re-sell what they already know. Make the next step as easy as possible.

The critical operational rule: cold email and LinkedIn prospects are almost always Level 1 or Level 2. Never open with your product. Never open with your company name. Open with their world.

---

## Framework 2: The Slippery Slide

Joe Sugarman built a career on one insight: the ONLY job of any line of copy is to make the reader read the next line. The first sentence exists to earn the second. The second exists to earn the third. If you can read a message and find any

sentence where you could stop and not feel compelled to continue, that sentence has failed.

This creates a specific technique: every passage should have forward momentum built into it. You do this with two tools.

The first is sentence-level curiosity. State something interesting, then withhold the explanation until the next sentence.

"Three months ago, one of our clients almost fired their entire SDR team." That sentence is incomplete - you need to know what happened. The next sentence delivers it, but in doing so, opens a new question. The reader is always slightly ahead of where you want them to be, reaching for what comes next.

The second is structural curiosity. Plant references to things you'll explain later. "I'll come back to why this matters in a moment." "There's one thing about their approach that surprised us - and it's not what you'd expect." These are open loops. The human brain hates unresolved loops. It keeps reading to close them. In a cold email, a well-placed open loop can be the difference between a reply and a delete.

Practical application in short-form copy: even in a 50-word cold email, you can use the slippery slide. Start with the most arresting statement you have. Don't save the good line for the middle. Put it first. Then make the second sentence explain why the first sentence is true. Then make the third sentence show what that means for this specific person.

---

## Framework 3: Sentence Rhythm

Gary Provost wrote one of the most instructive demonstrations in the history of writing instruction. He showed it with variation in sentence length:

"This sentence has five words. Here are five more words. Five-word sentences are fine. But several together become monotonous. Listen to what is happening. The writing is getting boring. The sound of it drones. It's like a stuck record. The ear demands some variety.

Now listen. I vary the sentence length, and I create music. Music. The writing sings. It has a pleasant rhythm, a lilt, a harmony. I use short sentences. And I use sentences of medium length. And sometimes when I am certain the reader is rested, I will engage him with a sentence of considerable length, a sentence that burns with energy and builds with all the impetus of a crescendo, the roll of the drums, the crash of the cymbals - and sounds, just right."

This is not a stylistic preference. It is a cognitive reality. Monotonous rhythm reads as AI-generated because it IS AI-generated - the model's loss function pushes it toward average sentence length, average complexity, average everything. Human writers do not produce averages. They produce music.

The operational rule: short declarative statement, then a longer sentence that builds and adds detail, then short again. The short sentences carry emotional weight. They land. The longer sentences carry information and create momentum. You need both. Consecutive short sentences feel choppy and aggressive. Consecutive long sentences feel like a briefing document. The combination feels like a human talking.

Fragments are powerful and underused. "Not ideal." "Classic problem." "Nobody catches it until Q4." These are not grammatically complete sentences. They read as confident and direct, the way someone talks when they know what they're talking about and don't need to qualify.

---

## Framework 4: Transformation and Story

Nathan Baugh's framework distills what makes a piece of content feel like it matters: there must be a before and an after. Not necessarily stated explicitly - implied is often more powerful. But the reader should feel that something changed, or is about to change, or could change. Without that implied transformation, you have a product description. Product descriptions are not copy.

The "But, Therefore" principle vs. "And Then": in bad copy, ideas connect like a list. "We do X and we do Y and we do Z." In good copy, ideas create causation and consequence. "They tried X but it didn't work, therefore they did Y." The difference is that the second version implies agency, decision-making, and stakes. The reader is tracking a journey, not reading a menu.

Open loops in storytelling. A story without tension is an anecdote. Even in 50 words, you can create tension by positioning the reader before a resolution. "We built the whole pipeline - 3 months of work - and then the lead list came back at 12% bounce rate." Now the reader needs to know what happened. Open the loop early. Close it with whatever you're actually trying to say.

Specificity as world-building. "A 40-person fintech in Berlin" is not just more specific than "a tech company" - it creates a world. The reader's brain fills in context: they picture an office, a team, the challenges specific to that size and geography and industry. Specificity does work that no adjective can do. "A company facing challenges" does nothing. "A Series A logistics SaaS that just hired its third VP of Sales in two years" creates a story before you've said anything about the story.

Juxtaposition creates tension. Put opposing things next to each other. "They had the best reps in the market. Pipeline was empty." The gap between those two sentences is where the reader leans in. The contrast forces a question: how is

that possible? That question is what you're writing the email to answer.

---

## Framework 5: Voice and Opinion

Voice is not tone. Tone is warmth vs. formality. Voice is the specific, idiosyncratic quality of a person's writing that makes it unmistakably theirs. Voice is built from four components: word choice, sentence rhythm, what gets left out, and opinion.

The first three can be approximated by a language model with enough examples. The fourth - opinion - is the one thing AI structurally cannot do. AI is trained to hedge, to balance, to present multiple perspectives. It cannot take a position and defend it the way a person does when they've actually experienced something and formed a view from that experience.

This means that the single most powerful humanization technique available is forcing opinion into the copy. Not an AI-style opinion ("we believe that...") but a defensible, specific claim that someone could disagree with.

Examples of AI-opinion: "It's important for sales teams to focus on quality outreach." That's not an opinion. That's a platitude. Nobody disagrees with it.

Examples of real opinion: "Most sales sequences are too long by exactly three emails." "The problem with most SaaS onboarding is that it's designed by the product team, not by the CS team, and you can tell." "Intent data is oversold. The signal-to-noise ratio is terrible unless you layer it with job change data."

These are opinions someone could argue with. They have a POV baked in. They signal that the writer has actually thought about this and arrived somewhere specific. That specificity of thought is the core signal of a human voice.

When writing as a sender who has real experience (client executives, sales leaders), mine their actual views. If you have examples of their past writing, look for places where they make claims - those claims are the seeds of their voice. Amplify them.

---

## Framework 6: Specificity as Credibility

David Ogilvy built his career on one rule above all others: research until you find the specific, true detail that makes everything else credible, then lead with it. He said the consumer is not a moron - she is your wife. Assume intelligence. Give her something real.

Generic claims do not persuade. They cannot, because they are not falsifiable. "We help sales teams improve their pipeline" - this could mean anything. It is the average claim that every sales tool makes. It passes through the reader's brain without leaving a mark.

Specific claims land differently. "We help teams cut the time from first outreach to booked meeting from 14 days to 6." That's a specific number. The reader either believes it or doesn't, but they have to engage with it. The specificity forces engagement.

The mechanism here is credibility transfer. Specific details suggest that someone actually counted, measured, tracked. You don't have a specific number unless you did the work to get it. The specificity is proof of rigor, even when the reader doesn't consciously think that. They just trust the specific claim more than the vague one.

Apply this in every sentence. Replace "many companies" with "three of the five clients we onboarded last quarter." Replace "significant improvement" with "37% faster." Replace "enterprise-level clients" with "Series B companies between 50 and 200 employees." Every time you reach for a general qualifier, ask what the actual number or detail is. If you don't have a real number, use the most specific framing that is still honest.

In cold email specifically: the opening line must contain something specific to this prospect. Not their company name dropped into a template. Something that requires knowledge of them - a specific post they wrote, a specific hire they made, a specific market they're expanding into. The specificity signals that you did not bulk-send this message. Whether or not that's true, the reader needs to feel it.

---

## Framework 7: Emotional Triggers in B2B Outbound

Sugarman catalogued 31 emotional triggers. In B2B cold outbound and newsletters, five dominate.

**Curiosity** is the most exploitable trigger in short-form copy. Humans have a near-compulsive need to resolve open questions. A subject line that implies an incomplete idea - not a question, but a hanging thread - forces the open. "What we found in 50 SDR pipelines" works because the reader needs to know what you found. The curiosity trigger is why Sugarman's slippery slide works at the sentence level.

**Pain before gain.** Kennedy's PAS framework - Problem, Agitate, Solve - works because it follows the human priority order. We respond to threats faster than to opportunities. In cold email, spending the first line on the problem (not your

solution) makes the reader feel seen before they've been sold to. Agitation is not cruelty - it's specificity. Naming the specific mechanism of the pain, not just the symptom.

**Belonging and exclusivity.** B2B buyers respond to signals that others like them have made this decision. "Three other Series A fintech companies we work with ran into this" says: this is not a unique problem, your peers have solved it, you are behind if you haven't. Exclusivity is the inverse: "We only work with 10 clients at a time" creates scarcity-driven desire even in buyers who were not previously motivated.

**Honesty and trust.** In an environment saturated with optimistic sales claims, honest admissions are disarming. Acknowledging a limitation before the prospect brings it up is not weakness - it's the most powerful credibility move available. "This doesn't work if your team is under 10 reps" builds more trust than leaving out the caveat and having the prospect discover it later.

**Curiosity about self.** The prospect is more interested in their own situation than in your product. Any message that offers them a new way to see their own situation - a framework, a comparison, a signal they hadn't noticed - earns attention that a product pitch never will. "Based on your team size and the market you're in, you're probably hitting one of two bottlenecks" - now they want to know which one. You've made the message about them discovering something about themselves.

---

## Framework 8: Awareness Level Matching - Diagnostic and Application

Before writing any piece of copy, diagnose the reader's awareness level. This is a three-question process.

First: does this person know they have a problem? If they're being cold-outreached for the first time, assume no. If they're in a nurture sequence after downloading a resource about their problem, assume yes.

Second: what language are they using for their situation? If you have LinkedIn posts, content they've shared, job descriptions they've published, or interview data from their segment - their language is in there. Extract it. Use it verbatim or near-verbatim. The prospect reading their own words back (wrapped in your framing) experiences a recognition response that manufactured language cannot replicate.

Third: what is the minimum move you need them to make? At Level 1, the minimum move is to feel recognized - not to reply, not to click, just to think "this person understands my situation." If you can generate that feeling in 30 words, you've won the first touchpoint. The reply comes later. At Level 4 or 5, the minimum move is a specific next step - a calendar link, a response, a decision. Don't confuse them by re-educating someone who's already convinced.

The matching table:

- Level 1: Lead with their observable situation, no mention of your product.
- Level 2: Name the problem vividly, imply you understand the mechanism, hint at a different approach.
- Level 3: Position your approach vs. the category norm, explain your differentiated principle.
- Level 4: Social proof from their specific peer group + a low-friction next step.
- Level 5: Just make it easy to say yes.

Cold email to a stranger is almost always Level 1 → Level 2. Start in their world. End with a single, low-commitment ask.

---

## The Lavender Data: Hard Numbers Applied

Lavender analyzed hundreds of thousands of cold emails and identified what actually correlates with replies. These are not opinions. They are frequencies from real data.

**Word count:** 25-50 words is the sweet spot. Under 75 words produces 83% more replies than messages over 75 words. This is the most important data point in modern cold email. If you're writing cold email and the draft is at 90 words, you are not done editing. The constraint is not a stylistic preference - it is empirically correlated with reply rate. Cut.

**Reading level:** 3rd-5th grade reading level produces 17% higher reply rate than high school level. Short words. Short sentences. No jargon. This is not about the prospect's intelligence - it's about cognitive load. A message that is easy to process gets processed. A message that requires effort gets deferred.

**Tone:** Informative tone produces -26% reply rate relative to neutral. Casual tone produces +23%. Writing like a business publication hurts you. Writing like a person talking to a person helps you. The informative tone pattern: noun-heavy, structured, data-first. The casual tone pattern: verb-heavy, sentence fragments, contractions, one idea at a time.

**Subject lines:** Two-word subject lines perform best. Questions in subject lines reduce replies. Commands reduce replies. Superlatives ("best," "fastest") reduce replies. Specific references perform best: the prospect's company name, a thing they said, a specific trigger event. The subject line's job is to earn the open, not to sell anything.

**Mobile:** 90% of cold emails are written on desktop, 80% are first opened on mobile. Design for a 4-inch screen. If the email requires scrolling to read, it is already losing. If the first line is not immediately compelling on a mobile preview, the

email will not be opened. The preview text - the first 100-150 characters of the email body - is your real subject line on mobile.

---

## How to Use These Frameworks in an Agentic Session

When you sit down to write a piece of copy, sequence through these frameworks as diagnostic questions, not as a checklist.

Start with awareness level. Before writing a single word, answer: where is this reader on Schwartz's scale? That determines your opening. If they're Level 1, your opening must be about their world, not your product. If they're Level 4, your opening can reference the thing that's been stopping them.

Then set the slippery slide. Identify the most arresting, specific, true thing you can say about this prospect's situation. Put it first. Make sure every subsequent sentence earns the next.

Then draft once, fast. Do not edit while drafting. Get the full idea out. **Do not open the rules during this phase.** Checking rules while writing causes you to optimize individual sentences instead of the whole piece. The result passes the checklist and doesn't read. The rules is an edit pass - run it only after the draft is complete.

Then apply the rhythm test. Read it aloud. If consecutive sentences feel the same length, rewrite. Short. Long. Short. If it doesn't sound like a person talking, it isn't.

Then check for transformation. Is there a before implied? An after? Or is this a description of a state? If it's just description, find the movement - even a hint of movement - and make it visible.

Then check for opinion. Find one sentence where you're making a claim someone could disagree with. If there isn't one, add one. It doesn't have to be provocative. It has to be specific and defensible.

Then apply the Lavender constraints. Count the words. Is it under 75? Is the subject line two to three words? Does the first line work on a mobile preview?

Then replace every generic claim with a specific. "Companies like yours" → "Series B fintech teams between 30 and 80 people." "Better results" → "18% higher reply rates."

This is not a linear process you run once. It's a lens-switching process. The frameworks are ways of seeing. Run the draft through each one, ask what it reveals, adjust. Good copy is not generated - it's edited into existence.

**The Style Card is the last layer.** If writing in a specific person's voice, apply their style signatures (sentence habits, shorthand, preferred patterns) only after the piece reads well on its own. Voice on top of writing that doesn't flow just adds noise. Make it work first. Make it sound like them second.

---

## The Book List

### Tier 1 - Read these before writing anything:

- **Sugarman, Adweek Copywriting Handbook** - the slippery slide, psychological triggers, how to read a reader's mind. The most practical book on copy mechanics.
- **Schwartz, Breakthrough Advertising** - the 5 Levels of Awareness plus the concept of mass desire. Everything about audience state comes from here.
- **Halbert, The Boron Letters** - write for one person, simplicity wins, AIDA as the backbone. Free online. Read all 25 letters.
- **Handley, Everybody Writes** - the authority on content quality and voice. Especially useful for newsletters and longer-form content.

### Tier 2 - Read when the Tier 1 books feel internalized:

- **Cialdini, Influence** - the six principles of persuasion. Reciprocity, Scarcity, Authority, Commitment, Liking, Social Proof, Unity. Understanding why humans comply is the foundation of ethical persuasion.
  - **Kennedy, The Ultimate Sales Letter** - Pain, Agitate, Solve as a complete methodology. Deep prospect understanding as the substitute for clever writing.
  - **Josh Braun, 15 Principles PDF (free)** - the most practical modern translation of classic copy for B2B outbound. Especially "poke the bear" and the "without" framing.
-