

Integrated Marketing Communications in the AI Era

Integrated Marketing Communications in the AI Era

*Running the AI-Augmented Agency
from Insight to Pitch*

Paul S. Richardson

Copyright © 2026 Cambria Press

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means without written permission from the publisher.

About this book: This book was created the way it teaches. It was drafted in collaboration with artificial intelligence (Anthropic's Claude); every element was directed, verified, edited, and approved by the author, who — per the accountable-author standard taught in Chapter 1 — signs it.

ISBN 978-1-63857-443-9 (PDF)

ISBN 978-1-63857-454-5 (paperback)

Printed in the United States of America

Contents

Preface	1
Introduction — The Reinvention: AI and the New Marketing Communications	3
1 IMC Foundations and the AI-Augmented Agency	11
2 Brand Strategy, Positioning, and Voice	25
3 Consumer Insight and Segmentation	37
4 Creative Strategy and the Big Idea	47
5 Advertising Copywriting Across Channels	57
6 Visual Identity and Generative Creative	67
7 Media Planning and Budgeting	77
9 Digital, Social, and Content Marketing	87
10 Email and Marketing Automation	97
11 PR, Earned, and Influencer Marketing	107
12 Sales Promotion, Direct Response, and Retail Media	117
13 Measurement, Testing, and Optimization	127
14 Ethics, Law, and the Future of AI in Marketing	135
15 Integration and the Pitch	143
A The Course Roadmap: How the Chapters, the Labs, and the Campaign Book Fit Together	151

B	The Campaign Book and the Pitch: A Specification	155
C	The Four Clients: Extended Case Narratives	159
D	The Prompt Pattern Library	165
E	Careers in AI-Era Marketing: A Field Guide	171
F	The Classical Foundations: Persuasion Theory, Consumer Psychology, and Budgeting Doctrine	175
G	Consolidated Course Glossary	197
	About the Author	223

Preface

In the summer of 2026, the *Wall Street Journal* reported on a new kind of company: the one-person firm scaling past one million dollars in revenue, with some clearing ten million. Their founders are not superhuman. They are well-staffed — by artificial intelligence that answers email, writes and debugs code, fields customer requests, and signs up subscribers. Stripe, the payment processor that watches the revenue of millions of businesses, finds the trend accelerating: twice as many solo companies crossed the million-dollar mark in 2025 as in 2023, and solo founders now account for nearly two-thirds of new incorporations. The one-person company is no longer a curiosity. It is a business model.

Advertising and marketing communications are living the same story at industry scale. When *Ad Age* surveyed one hundred industry leaders on the year ahead, artificial intelligence was no longer a trend on the list — it was the premise beneath the list, the collaborative engine now running content, campaigns, creative, and analytics alike. The industry's own structure is reorganizing around that fact: the holding companies are merging, rebuilding, and openly shifting investment from human headcount to machine capability. Strip away the headlines and a single pattern remains — the professional unit of marketing work is becoming one person directing a staff of machine collaborators. Which is to say: the industry is converging on exactly the organizational form the *Journal* found in those million-dollar one-person firms.

This is the world our students will graduate into, and it changes what an IMC course owes them. AI is no longer an ad hoc tool bolted onto marketing communications; it is in the workflow itself, drafting the copy, mocking the visuals, modeling the budget. What students need now is not another survey of channels with an AI sidebar, but practiced judgment in the three roles the machine cannot fill: directing the work, verifying its claims, and editing it until it deserves a signature.

That is why this book teaches IMC by having each student run one of those one-person firms. In this course, students charter an AI-augmented agency and build a complete campaign for a realistic client — positioning through creative, media through measurement — by completing a series of studio labs alongside an AI coach on the dedicated course platform at www.imcstudio.ai. Week by week, the deliverables accumulate into a bound Campaign Book, audited for ethics and

disclosure and defended in a live pitch. Students finish the semester not having read about how IMC is practiced in the AI era, but having practiced it.

For professors considering the textbook and studio approach, adoption resources include narrated lecture decks for all fifteen weeks, a complete Instructor Companion with fully worked lab sessions from both the student's and instructor's side, a sample syllabus, quizzes, grading rubrics for the Campaign Book and pitch, and a fully graded exemplar campaign. The Studio also lightens the grading load: for every lab submission it prepares a first-pass rubric assessment — scores, comments, and evidence drawn from the student's own work and session transcript — which the professor reviews, edits, and releases under their own name. Student AI access is included with the textbook code: each student receives a semester allowance of AI usage sized for the full lab sequence, with no separate AI accounts or subscriptions required, and instructors can extend any student's allowance in one click. Each course runs in its own private space, and no AI expertise is required to teach with it — the Companion assumes none.

Welcome to *Integrated Marketing Communications in the AI Era*.

This volume is one of the *Marketing in the AI Era* series — four companion courses, one shared universe, each with its own studio. The complete series is at <https://www.marketingintheaiera.ai>.

Paul S. Richardson, PhD
Professor of Marketing
Holzschuh College of Business Administration
Niagara University

Sources

Te-Ping Chen, "How AI Is Helping One-Person Companies Scale to \$1 Million — and Beyond," Wall Street Journal, July 29, 2026.

"100 Ad Leaders Predict 2026 Marketing Trends," Ad Age, January 5, 2026.

eMarketer, "Omnicom's Latest Report Signals an Agency Shift from Human Talent to AI," 2026.

The Reinvention: AI and the New Marketing Communications

Two Christmases in Atlanta

In November 2024, Coca-Cola — owner of perhaps the most beloved holiday advertising tradition in the world — released a remake of its classic “Holidays Are Coming” campaign produced substantially with generative AI. The reaction was swift and brutal. Critics called it soulless; creators accused the company of trading craft for cost savings; the word “slop” — the internet’s term for low-effort machine-generated content — attached itself to one of the most valuable brands on earth. For a company whose Christmas trucks had been rolling through television sets since 1995, it was a remarkable stumble: not a failure of budget or reach, but a failure of *meaning*.

Here is the interesting part. One year later, in November 2025, Coca-Cola did it again — deliberately. A second AI-generated holiday ad, technically far more polished, arrived with the company openly defending the approach; the backlash arrived on schedule anyway, and this time the industry’s reaction split. Some saw a brand cheapening its most emotional asset two years running. Others saw the future being normalized in real time by the one advertiser big enough to absorb the criticism — noting that production costs had fallen by an order of magnitude, that the ad had been made in a fraction of the traditional timeline, and that most viewers scrolling past it on a phone neither noticed nor cared how it was made.

Hold both Christmases in your head at once, because together they are this book’s subject in miniature. The technology worked — better the second year than the first, and it will be better still by the time you read this. The *economics* worked — spectacularly. What remained genuinely contested was everything this book is about: whether the work was any good, whether audiences felt respected or shortchanged, whether the brand’s long-term meaning was being in-

vested or spent, and who — in a production process where machines did most of the making — was actually responsible for the answer.

The technology will not wait for you to finish this course. In the single semester you spend with this book, models will improve, tools will launch, and at least one brand will step on exactly the rake Coca-Cola stepped on twice. That is why this book is not a manual for today's tools. It is a course in the judgment that survives the tools changing.

The Numbers: How Fast, How Big

It is difficult to overstate how quickly artificial intelligence moved from novelty to infrastructure in marketing. A few load-bearing figures, current as of this writing:

Adoption is effectively universal at the organizational level. McKinsey's *State of AI* survey — conducted in mid-2025 across 1,993 organizations in 105 countries — found that **88 percent of organizations now use AI in at least one business function**, up from 78 percent a year earlier, with marketing and sales consistently among the functions reporting both the heaviest use and the clearest revenue benefits. Whatever firm hires you after graduation, the probability that AI is already inside its marketing workflow rounds to certainty.

The money is following. The Interactive Advertising Bureau's 2026 outlook projects U.S. advertising spend growing **9.5 percent in 2026**, with buyers explicitly citing the accelerating adoption of *agentic* AI — systems that don't just generate content but take actions, optimize campaigns, and transact — as a driver of where those dollars go.

The platforms are automating the work itself. Meta — whose advertising machine funds one of the largest businesses in history — has stated its ambition to let advertisers **fully automate ad creation with AI**: a business states an objective and a budget, and the system generates the creative, picks the audiences, and optimizes the spend, end to end. Eleven new AI advertising tools launched at Cannes Lions 2025 alone were steps toward that vision. Whether this thrills or unsettles you probably depends on whether you were planning to be the person who made those ads.

The industry is consolidating around the transition. In November 2025, Omnicom completed its acquisition of Interpublic Group, creating the largest advertising holding company in the world — a merger both companies framed explicitly as a response to the AI era: the scale to build proprietary data and AI platforms, positioned for what their announcement called “intelligent growth.” When the two of the biggest names in an industry combine and give *that* as the reason, students of that industry should take note.

And the audience is watching with narrowed eyes. Getty Images' consumer research finds that **nearly 90 percent of consumers want transparency about**

whether images are AI-generated. When *Vogue*'s August 2025 issue carried a Guess campaign featuring a strikingly beautiful model who did not exist — created by an AI studio, disclosed only in small print — the reader backlash made international news, models' unions raised alarms, and the industry got a preview of a question it will be answering for a decade: *what does an audience deserve to know about how the persuasion aimed at them was made?*

One more set of numbers matters to you personally, and we will return to it at the end of this introduction: the research on what AI actually does to the people who work with it.

What Actually Changed

Strip away the headlines and the change underneath is surprisingly easy to state: **the cost of producing competent persuasive content collapsed.** Not declined — collapsed. A task that consumed a junior copywriter's afternoon — ten headline options, a first-draft press release, a set of social captions — now takes a model seconds, and the model's version is, by blind evaluation, roughly as good as the junior's. The same is increasingly true of images, video, media plans, and analytical narratives.

Economists have a name for what happens next: when the price of a good falls toward zero, the value migrates to its *complements* — the things you still need that did not get cheap. This book is organized around that migration. If competent drafts are nearly free, what remains scarce?

Strategy remains scarce. The machine can write ten positionings in a minute; it cannot know which one your brand can defend, fund, and live up to — that requires knowing things about the client, the market, and the future that are not in any training set.

Distinctiveness remains scarce — and got scarcer. When everyone's first draft comes from similar models trained on similar data, everyone's first draft looks alike. Fluent sameness is now the default state of marketing content; escaping it is deliberate work, and this book spends several chapters teaching exactly how.

Verification became precious. Models state falsehoods with the same confidence as facts. In a discipline whose legal and ethical foundation is truthful claims, the ability to catch the confident error — before it ships, before a regulator or a Reddit thread catches it for you — may be the single most employable skill this book teaches.

Judgment became the job. Someone still decides what the campaign believes, which of the forty options to ship, what the terms and conditions promise, and what the brand will not say under pressure. The junior tasks that once trained that judgment are precisely the tasks that automated — which means your generation will have to acquire judgment *deliberately*, through structured practice,

rather than absorbing it across years of grunt work. That is not a small curricular adjustment. It is the reason this course looks the way it does.

A useful way to hold all this: marketing communication has lived through scarcity regimes before. When production was scarce (a century of print and broadcast), power belonged to those who could make and place the ads. When distribution went infinite (the internet), power moved to those who could win attention. Now that competent *content* is infinite, power is moving again — toward those who can be trusted, verified, and distinguished. Every chapter in this book is, one way or another, about operating in that third regime.

What the Research Actually Shows

Amid the hype and the panic, a serious body of evidence has accumulated about what generative AI does to knowledge work, and three findings matter enough to preview here — you will meet all three again, in depth, later in the book.

First: AI raises the floor dramatically. In controlled experiments, giving professionals access to generative AI made writing tasks substantially faster *and* better, with the largest gains going to the less-skilled performers — compressing the gap between the weakest and strongest workers (Noy & Zhang’s *Science* experiment). Field evidence from thousands of customer-support agents found the same shape: large productivity gains, concentrated among novices, as the AI effectively distributed the top performers’ patterns to everyone (Brynjolfsson, Li & Raymond). For you, this finding cuts two ways. The floor you must clear to be employable just rose — competent is now free. But the tools will also make you competent at things you have never done, faster than any generation before you.

Second: the frontier is jagged. The largest field experiment on AI-assisted professional work — 758 consultants at a major firm — found that AI collaboration dramatically improved performance on tasks within the technology’s competence, and *degraded* the performance of skilled professionals on tasks just outside it, because the model’s fluency concealed its failures and people trusted polished wrong answers (Dell’Acqua and colleagues’ “jagged technological frontier”). The professionals who thrived were those who knew where the edge ran. Teaching you the location of that edge — empirically, by letting you fall over it in safe conditions — is this course’s method, and by the end of the semester you will hold a map of it that most working professionals do not yet have.

Third: audiences penalize the machine less than you’d think — and reward honesty more. In careful studies, people rate AI-generated persuasive content as equal to or better than human work *until told how it was made* — and even then, the penalty is modest, better described as a preference for human-labeled work than punishment of machine-labeled work (Zhang & Gosline). The practical implication threads through this whole book: the disclosure question is strategic,

not existential, and brands that handle it deliberately — as a documented decision rather than a hope nobody asks — are the ones that keep their footing when the *Vogue* moment arrives.

The Industry You Are Entering

What does all this mean for the career you are preparing for? Four honest observations.

The entry-level job is being redefined under your feet. The tasks that once filled a coordinator's or junior copywriter's first two years — first drafts, versioning, basic decks, media math — are exactly what the tools now do. The roles are not disappearing, but they are being re-scoped upward: employers increasingly expect a first-year hire to *direct and verify* machine output rather than produce raw drafts. The uncomfortable truth is that “I can write ten decent headlines” is no longer a differentiator; “I can get fifty out of the machine, pick the three that fit the strategy, catch the claim we can't substantiate, and defend the choice to a client” very much is.

The org charts are in motion. Agencies are consolidating for AI scale (the Omnicom–Interpublic merger being the loudest example), brands continue pulling work in-house now that production is cheap, and the platforms are automating the media-buying layer entirely. Nobody knows precisely where the industry lands. But every plausible configuration needs people who can do what machines cannot: hold the strategy, own the judgment, sign the work.

Verification is becoming a profession. Claims substantiation, provenance tracking, AI-content disclosure, bias auditing — functions that barely existed as jobs five years ago are appearing in agency org charts and job postings, driven by regulators, platforms, and burned brands. This book teaches those practices not as compliance chores but as craft, because the people who can do them credibly are about to be very employable.

Trust is the compounding asset. Every synthetic model in a magazine, every AI holiday ad, every fake review farm draws down a shared reservoir of audience belief. The brands — and the professionals — who operate visibly on the honest side of that ledger are accumulating the one advantage that gets *more* valuable as fakery gets cheaper. It is not an accident that this book ends with a signature.

How This Book Works

This is not a book *about* AI in marketing, to be read at arm's length. It is a book in which you *run* an AI-augmented marketing operation, for real, all semester — and the structure follows from that.

You will operate a one-person agency. In Week 1 you will charter it and meet your AI teammate — the Studio Coach, a senior-mentor AI that accumulates your work as the semester builds. In Week 2 you will take on a client: one of four fictional brands with rich, realistic briefs. Every week after that, you will produce one real component of an integrated campaign — positioning, personas, creative, copy, art direction, media plan, content, email, PR, promotion, measurement — each one building on the last, until Week 13 assembles them into a complete **Campaign Book**, Week 14 audits it, and Week 15 pitches it. Nothing you make is disposable homework; the dependency map in your syllabus shows exactly how each week’s work feeds the next.

The chapters share one rhythm. Each opens with a scenario from Harbor & Vine — a fictional agency whose weekly adventures dramatize the chapter’s core problem. The middle sections teach the marketing discipline itself: the frameworks, the evidence, the craft, exactly as a traditional IMC text would. Then an **AI practice layer** teaches how that week’s work actually gets done with a machine collaborator — including how it goes wrong. An **ethics sidebar** confronts the week’s hardest honest question. And a **lab briefing** hands the whole thing to you to do.

Your AI teammate will sometimes be wrong on purpose. This deserves plain statement here, because it is the course’s most unusual design choice: in several labs, the Coach is *scripted* to err — an arithmetic slip in a budget table, a confident “industry benchmark” with no source, an unconfirmed fact smuggled into a crisis statement drafted under deadline pressure. Real tools fail in exactly these ways, and a course that only ever showed you a flawless machine would graduate students who trust machines. Catching the errors is graded. You have now been warned, which — you will discover — helps less than you think.

The disciplines are the curriculum. Along the way you will acquire a set of working habits with names you will get tired of and then, we suspect, use for the rest of your career: the *claim log* (every fact verified, excluded, or flagged), the *arithmetic rule* (every machine number recomputed by hand), the *voice spec* (brand voice made precise enough to govern a machine), the *swap test*, the *claim ladder*, the *won’t-say list*, and a dozen more. Each was designed for one purpose: to let you move at machine speed without surrendering human judgment.

And at the end, you sign. The Campaign Book closes with an Ethics & Disclosure Statement over your actual name: what the machine made, what you verified, what you disclosed, what risks you flagged. Because the last thing this book believes — and the first thing worth telling you — is that a profession, in the end, is a group of people who sign their work.

The tools in this book’s examples will be outdated before your diploma is printed. The judgment will not. Welcome to *Integrated Marketing Communications in the AI Era*.

Key Terms

- **Agentic AI** — AI systems that go beyond generating content to taking actions: optimizing campaigns, allocating budgets, transacting — the industry’s current frontier.
- **AI-augmented agency** — This course’s operating model: one human strategist directing, verifying, editing, and signing the work of AI collaborators.
- **Campaign Book** — The semester’s capstone artifact: the complete, integrated, audited campaign you assemble week by week and defend in the final pitch.
- **Complements of free content** — What gains value when content generation approaches zero cost: strategy, distinctiveness, verification, and judgment.
- **Floor-raising effect** — The replicated research finding that generative AI improves the least-skilled performers most, compressing skill gaps and raising the baseline of “competent.”
- **Generative AI** — AI systems that produce novel content — text, images, video, audio — from natural-language instruction; the technology whose cost-collapse in content production this book is organized around.
- **Human favoritism** — The evidence that audiences rate AI-made persuasive content equal or better until told its origin, and then show a modest preference for human-labeled work rather than punishment of the machine’s.
- **Jagged frontier** — The research finding that AI collaboration helps dramatically inside the technology’s competence and quietly degrades expert judgment just outside it; knowing where the edge runs is the operating skill.
- **Scarcity regimes** — The historical framing of marketing power: from scarce production, to scarce attention, to the present regime of scarce trust and distinctiveness amid infinite competent content.
- **“Slop”** — The internet’s dismissive term for low-effort, high-volume machine-generated content; what brands risk shipping when generation outruns judgment.

References and Further Reading

Scholarly foundations

- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. — *The controlled experiment behind the floor-raising effect in professional writing.*

- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *Quarterly Journal of Economics*, 140(2), 889–942. — *Field evidence from thousands of support agents: large gains, concentrated among novices.*
- Dell’Acqua, F., McFowland, E., Mollick, E. R., et al. (2023). Navigating the jagged technological frontier. Harvard Business School Working Paper 24-013 (subsequently published in *Organization Science*, 2025). — *The 758-consultant experiment; the frontier map this course teaches you to draw.*
- Zhang, Y., & Gosline, R. (2023). Human favoritism, not AI aversion. *Judgment and Decision Making*, 18, e41. — *The disclosure-penalty evidence.*

Industry and current sources

- McKinsey & Company (2025). *The State of AI: Global Survey*. — *The adoption figures cited in this introduction (88% of organizations; survey of 1,993 respondents in 105 countries, mid-2025).*
- Interactive Advertising Bureau (2026). *2026 Outlook Study*. — *The 9.5% U.S. ad-spend growth forecast and the agentic-AI adoption driver.*
- Getty Images (2024–2025). *Building Trust in the Age of AI* consumer research. — *The ~90%-want-transparency finding.*
- Omnicom Group (2025, November). “Omnicom Completes Acquisition of Interpublic.” — *The merger creating the world’s largest advertising holding company, framed around AI-era scale.*
- Contemporary press coverage of the Coca-Cola AI holiday campaigns (2024, 2025), the Guess/Vogue AI-model controversy (August 2025), and Meta’s AI advertising automation roadmap — *the opening examples of this introduction; readily searchable, and worth reading in the original as exercises in how the trade and popular press narrate this transition.*

IMC Foundations and the AI-Augmented Agency

Opening Scenario: Monday Morning at Harbor & Vine

At 9:04 on a Monday morning, the four people who make up Harbor & Vine — a small marketing agency in a converted warehouse office — learn they have won the pitch. Their new client is a regional craft beverage company with a modest budget, an ambitious growth target, and a launch window twenty-six weeks away. In 2019, an agency this size would have politely declined the assignment. A full integrated campaign — brand positioning, audience research, creative development, copy for a dozen channels, a media plan, an email program, a press strategy, and a measurement framework — was the work of a staff of twenty.

Harbor & Vine has four people and an AI model.

By Wednesday, the strategist has interviewed a dozen synthetic customer personas built from the client's survey data, pressure-tested three positioning territories, and killed two of them for good reasons. The copywriter has generated forty headline directions, discarded thirty-six, and rewritten the surviving four in the client's voice — a voice she defined precisely enough that the model can now imitate it. The media planner has modeled three budget scenarios. The account lead has drafted, edited, and fact-checked a creative brief that would have taken a week to circulate for approval.

Notice what the model did *not* do. It did not decide who the customer is. It did not choose which positioning was strategically defensible. It did not know — until the copywriter caught it — that one generated headline made a claim the client could not legally substantiate. Every decision that mattered was made by a person; every hour of production drudgery that used to stand between decisions was compressed to minutes.

This is the agency you will run for the next fifteen weeks. The theory in this book tells you what decisions matter and why. The Studio labs teach you to make them with an AI teammate — fast, professionally, and with your judgment always in charge.

What Integrated Marketing Communications Is — and Why It Exists

Integrated marketing communications (IMC) is the disciplined process of planning, creating, delivering, and evaluating all of a brand's messages so that every contact a customer has with the brand — an ad, a package, a social post, a support email, an influencer's mention — works together to build one coherent meaning and move people toward action.

The definition contains three commitments worth pausing on.

First, IMC is a *process*, not a set of deliverables. Campaigns fail less often from bad ads than from bad process: objectives nobody agreed on, audiences defined by wishful thinking, channels chosen by habit. The planning sequence you will learn in this course — and execute, week by week, in your Campaign Book — exists to prevent those failures.

Second, IMC concerns *all* messages, not just the ones marketing intends to send. A brand that advertises warmth and then bills customers with hostile fine print is communicating — just not what it meant to. Modern IMC thinking treats every brand contact as media.

Third, integration means *one meaning*, not one message. A campaign says different things in a six-second video, a podcast read, and a 2,000-word founder letter. Integration is achieved when those different executions are recognizably the same brand making the same core promise. Strategists sometimes summarize this as speaking with **one voice, many volumes**.

Why integration became the discipline's central problem

For most of the twentieth century, marketing communication was organized in silos: an advertising agency made ads, a PR firm courted journalists, a promotions house ran sweepstakes, and nobody was responsible for the whole. The silo model survived because media were few and audiences were captive. When three television networks reached most households, “the campaign” and “the TV campaign” were nearly the same thing.

Two forces broke that world. **Audience fragmentation** scattered attention across hundreds of channels, platforms, and micro-communities, so no single medium could carry a brand's meaning alone. And **message skepticism** rose as consumers gained tools to skip, block, mute, and fact-check; a claim in an ad is now instantly checkable against the brand's reviews, its labor record, and what a stranger on a forum said last Tuesday. In a fragmented, skeptical environment, inconsistency is not merely inefficient — it is evidence against you. Integration became a survival requirement.

A third force is the reason this course exists: **generative AI has collapsed the cost of producing messages**. When a passable ad, post, or email can be

generated in seconds, the volume of marketing communication explodes and the scarce resources become strategy, taste, and trust. We return to this in Section 1.5.

The Chapter 1 vocabulary of value: brands and positioning

IMC serves a brand, so we need a working definition. A **brand** is the network of meanings, memories, and expectations a name evokes — an asset that lives in customers' heads, not in the company's files. **Brand equity** is the commercial value of that network: the price premium, preference, and forgiveness a strong brand earns.

Positioning is the deliberate act of choosing which meaning the brand should own in the customer's mind relative to alternatives — the single strategic decision from which all communication flows. You will write a formal positioning statement in Week 2, and every deliverable after that will be graded partly on whether it honors it. Position first; communicate second. Reverse the order and you get campaigns that are merely loud.

The Toolbox: Promotional Mix and the PESO Channel Model

IMC coordinates a set of communication disciplines that developed separately. You will practice each one in this course; here is the map.

Advertising — paid, identified, non-personal messages placed in media (Weeks 5–7). Its strengths are reach, control, and repetition; its weakness is the skepticism paid messages invite.

Public relations and earned media — managing relationships with publics, especially the press, to earn coverage the brand does not pay for or fully control (Week 11). Earned messages carry credibility precisely because they are not ads.

Sales promotion and direct response — incentives (discounts, sampling, contests) and offers engineered to produce measurable action now (Week 12). Powerful and dangerously habit-forming: promotion moves product this quarter and, used carelessly, teaches customers never to pay full price again.

Digital, social, and content marketing — the owned-and-shared world of websites, social platforms, search, and the ongoing publishing operation modern brands run (Week 9).

Direct and lifecycle marketing — individually addressable communication, today dominated by email and marketing automation (Week 10).

Personal selling — human-to-human persuasion. It appears in this course chiefly in Week 15, when you pitch your campaign — because presenting and defending strategy is a selling task.

The classic mix describes *disciplines*. Contemporary planners also need a map of *channel ownership*, and the standard one is **PESO**:

Channel type	You...	Examples	Currency
Paid	buy access	ads, sponsored posts, paid search, influencer fees	money
Earned	persuade gatekeepers	press coverage, reviews, organic word of mouth	newsworthiness & trust
Owned	control outright	website, email list, packaging, app	audience you've built
Shared	participate	social platforms, communities, user-generated content	attention & participation

A well-built campaign moves people deliberately across PESO — paid media generates the first exposure, shared media multiplies it, earned media legitimizes it, and owned media captures the relationship so the brand is not renting its audience forever. When you build your media plan in Week 7, PESO balance will be one of the things your rubric rewards.

How Communication Actually Works

Every discipline in the mix runs on the same underlying machinery, and when campaigns fail, a planner should be able to say *where* in the machinery they failed.

The classical model describes a **sender** who **encodes** an intended meaning into a **message**, transmits it through a **channel**, where a **receiver decodes** it — all amid **noise**, with **feedback** closing the loop. Three implications matter enormously in practice.

Meaning is made by the receiver. Decoding happens inside the audience's frame of reference, not the sender's. A luxury cue in one community is a status joke in another; an idiom that lands with 45-year-olds baffles 19-year-olds. This is why audience insight (Week 3) precedes creative work (Week 4) in every professional process — and why the **field of shared experience** between brand and audience is the ground every campaign is built on.

Noise is now the dominant condition. The average consumer encounters thousands of commercial messages daily and processes almost none of them deeply. Attention is the first conversion. Creative quality — distinctiveness,

emotion, craft — is not decoration; it is the price of admission to the rest of the persuasion process.

Effects are staged. Audiences typically move through some version of a **hierarchy of effects** — from awareness to knowledge to liking to preference to conviction to action — though not always in that tidy order. The practical value of hierarchy models is diagnostic: a brand nobody has heard of has an awareness problem, and no coupon fixes it; a brand everyone knows and nobody prefers has a meaning problem, and no reach buy fixes that. Setting **communication objectives** means naming the stage you intend to move, for whom, by how much, by when. You will do exactly this for your client in your Campaign Book, and your measurement plan in Week 13 will test whether it happened.

One more layer of the machinery has joined the model in your professional lifetime: increasingly, the encoding is done *with a machine's help*, and machine-shaped messages are decoded by audiences — and filtered, ranked, and summarized by other machines — at scale. That does not change the physics of meaning-making. It changes who does the drafting, how many variants exist, and how fast feedback arrives. Which brings us to the second half of this chapter.

The IMC Planning Process — Your Map for the Semester

Professional campaigns follow a planning discipline that this course makes literal: each stage is a week of your life and a section of your Campaign Book.

1. **Situation analysis** — What is true? Client, category, competitors, customers. (Your brand brief is the raw material; Labs 2–3 interrogate it.)
2. **Objectives** — What must change, in whose head or behavior, measured how? (Set in your brand platform; finalized in Week 13's measurement framework.)
3. **Strategy** — Positioning, audience choice, and the creative idea big enough to organize everything (Weeks 2–4).
4. **Tactics** — The executions: copy, design direction, media plan, content calendar, email program, PR kit, promotion (Weeks 5–12).
5. **Measurement and optimization** — How you will know it worked, and how the campaign learns while running (Week 13).

Two habits distinguish professionals at every stage. They can always answer “*what decision does this work product feed?*” — analysis that feeds no decision is decoration. And they revise: the Campaign Book you submit in Week 15 will contain Week 2 work you have rewritten three times, because that is what real campaign development looks like.

The AI-Augmented Agency

Now the second subject of this course. Generative AI models — systems trained on vast text corpora to produce fluent language on demand — have entered every stage of the process you just read. It is tempting to file this under “new tools,” like desktop publishing or programmatic ad buying. The better analogy is structural: AI changes *the economics of the work itself*.

What actually changed

The cost of a draft collapsed. For a century, producing communication — writing the copy, comping the layout, cutting the video — consumed most of an agency’s hours and payroll. When drafts are nearly free and instant, volume stops being the constraint. A strategist can see forty headline territories instead of six; a planner can model twelve budget scenarios instead of two.

Judgment became the scarce resource. When anyone can generate competent-sounding copy, competent-sounding copy is worthless — the market price of “fine” fell to zero. What still commands a premium: knowing which of the forty headlines is strategically right, hearing that a draft is off-voice, catching the claim that is subtly false, and having the taste to push past the model’s first, safest, most average idea. Every rubric in this course allocates explicit points to exactly these acts, because they are what employers now hire for.

Verification became a production step. Language models are fluent guessers, not databases. They produce confident, plausible falsehoods — **hallucinations** — at unpredictable moments. Professional AI-augmented workflow therefore *ends* with human verification the way traditional workflow ended with proofreading. In this course that habit is graded: your transcripts are reviewed for evidence that you checked the machine.

The floor rose and the ceiling didn’t. AI reliably lifts weak output to average. It does not, unattended, produce excellent — excellence still requires a human with a point of view directing the machine through many rounds. This is why the Studio Coach in your labs will refuse to simply hand you finished deliverables: average is free now, and this course is not training you to produce it.

The three roles you will practice

In the AI-augmented agency, the junior marketer’s job description compresses into three recurring roles, and every lab in this course exercises all three:

- **Director** — you supply intent: the strategy, the audience, the constraints, the brief. The single strongest predictor of AI output quality is the quality of the brief it was given — a fact that should sound familiar, because it has always been true of human creative teams too.

- **Editor** — you impose taste: selecting, cutting, reworking, and voice-checking machine drafts against the brand platform *you* wrote.
- **Verifier** — you guarantee truth: every fact, claim, and statistic in work you submit is yours to defend, whoever drafted it.

Hold onto the through-line, because it is the thesis of this entire book: **the machine drafts; you decide**. The strategy chapters of this book teach you what to decide. The labs teach you to direct, edit, and verify at a professional standard.

The AI Practice Layer: Briefing the Machine

Each chapter of this book pairs the marketing theory you just read with the professional AI practice that goes with it. This week's practice is the foundation all the others build on: getting good work out of a language model on purpose rather than by luck.

How the machine works — the sixty-second version

A large language model (LLM) is a statistical system trained on enormous quantities of text to predict what words plausibly come next. From that one trick emerges something remarkable: give it a description of a task and it produces fluent, often skillful attempts at the task. Three practical consequences follow.

The model has read more marketing than any human alive — every framework, every formula, every register of voice — so it is a genuinely strong pattern-matcher for this field. But it optimizes for *plausible*, not *true*, so it will state falsehoods with the same confident fluency as facts. And it has no access to your strategy, your client, or your intent except what you put in the conversation. It is, in short, a brilliant, tireless, occasionally overconfident junior colleague with no context — *until you brief it*.

The Creative Brief Method

Marketers have a superpower here that most professions lack: our field already invented the technology for getting good work out of a talented collaborator who lacks context. It is called a **creative brief**, and prompting a model well is simply briefing it well. The course's standard prompt structure has five parts — **Role, Task, Context, Constraints, Format (R-T-C-C-F)**:

- **Role** — who the model should be: “*You are a senior copywriter known for warm, unpretentious premium-brand voice.*” Role framing activates the right region of the model's learned patterns.
- **Task** — the specific deliverable and its purpose: “*Draft 10 headline options for a paid social campaign whose objective is trial among first-time buyers.*”

- **Context** — the strategy the work must serve: positioning, audience, insight, the campaign's stage in the hierarchy of effects. This is the section beginners omit and professionals load.
- **Constraints** — hard boundaries: banned claims, mandatory disclosures, voice rules, length, reading level, what competitors already own.
- **Format** — what the output should look like, so you can evaluate it fast: a table, three options with rationales, a 200-word maximum.

A worked example

Here is the difference the method makes, using our Harbor & Vine scenario. First, the prompt most beginners write:

Weak prompt: “Write a slogan for a craft beverage company.”

The model will comply instantly — and produce the world's most average slogan, because average is exactly what an unconstrained prediction machine returns. Now the briefed version:

Briefed prompt: “You are a senior brand copywriter. **Task:** draft 10 tagline options for Harbor Craft Sodas, for use across packaging and paid social. **Context:** positioning is ‘the grown-up soda for people who outgrew soda but not joy’; audience is 28–40, urban, health-aware but pleasure-positive; the strategic insight is that quitting sugary soda feels like losing a small daily ritual. **Constraints:** no health claims of any kind (regulatory), never use the words ‘guilt’ or ‘guilty,’ 6 words maximum, wit welcome but no puns on ‘pop.’ **Format:** numbered list, each with a one-line rationale tied to the insight.”

Same model, same second of effort — an entirely different class of output, because the second prompt transmitted *strategy*. Notice that writing it required real marketing knowledge: a positioning, an insight, a regulatory constraint. This is the course in miniature. **Prompt quality is strategy made visible** — which is why your lab transcripts are part of your grade.

Iteration: where the real work happens

Professionals almost never accept a first draft, human or machine. The craft is in the follow-up moves, and you should learn these four as reflexes:

1. **Push for range** — “Give me 10 more, but make five of them risky enough that a cautious client would reject them.” First drafts cluster at the safe center; the good ideas live at the edges.
2. **Critique against strategy** — “Score each option against the positioning and flag any that drift off-voice.” The model is a surprisingly capable critic of its own work when handed criteria.

3. **Invert** — “Argue against your best option. What would a skeptical customer say?” Red-teaming your own creative is cheap insurance.
4. **Refine by direction, not by re-roll** — “Take #7: keep the warmth, cut the cleverness, make it declarative.” Aimless regeneration is slot-machine prompting; direction is editing.

The four failure modes — learn to see them

You will meet all four of these this semester, and your Editorial Judgment rubric score depends on catching them:

Failure mode	What it looks like	Your defense
Hallucination	Invented statistics, fake sources, confident wrong facts	Verify every checkable claim against the brief or a real source before it enters a deliverable
Genericness	Fluent, forgettable, could-be-any-brand output	Sharper context and constraints; push past draft one; add the point of view only you have
Voice drift	Output slowly stops sounding like the brand	Keep the voice guide in the conversation; run explicit voice checks
Sycophancy	The model praises and agrees with whatever you propose	Ask for critique explicitly; invite disagreement; never use the model’s enthusiasm as evidence

Ethics Sidebar: The Accountable Author

Every chapter closes its theory with the ethical issue that lives closest to that week’s practice. This week’s is the one that governs all the others.

When work product passes through a machine, responsibility does not. The professional norm this course holds you to — and the one emerging across the industry — is **the accountable author standard**: whoever submits the work owns every word, claim, and implication in it, exactly as if they had typed it themselves. “The AI wrote that part” has the same professional standing as “my intern wrote that part” — which is to say, none. You signed it; you own it.

Two corollaries follow. First, *verification is not optional*: submitting an unchecked machine-generated statistic is not a technology accident, it is a profes-

sional failure. Second, *disclosure is context-dependent but honesty is not*: whether audiences must be told content is AI-assisted varies by platform, jurisdiction, and situation (Chapter 14 covers the regulatory landscape), but deceiving audiences about authorship where authorship matters — fake reviews, fabricated testimonials, synthetic “customers” — is always over the line. In this course, the disclosure rule is simple: your instructor always knows, because your transcript is part of your submission. In your career, the rule is: when in doubt, disclose.

Studio Lab 1 Briefing: Meet Your Teammate

Open in IMC Studio → Lab 1. Due before Session 2.

This week’s lab has one purpose: turn the model from a curiosity into a colleague, and leave a written record showing you can direct it.

Part A — Calibration drills (in your Studio workspace). Three short exercises: (1) give the model a deliberately weak prompt, then rewrite it with the full R-T-C-C-F structure and compare outputs in a sentence or two; (2) take one output and improve it through at least three *directed* iterations — range, critique, refine; (3) hallucination hunt: ask the model a factual marketing question you can check, verify its answer, and document what you found. These drills are graded from your transcript.

Part B — Your agency charter (the deliverable). Found your agency: a one-page charter containing your agency name, a three-sentence philosophy of AI-augmented marketing (what the machine is for, what only you do), your working principles as director-editor-verifier, and one professional goal for the semester. Draft it with the model; edit it until it sounds like you and not like it.

Rubric emphasis: AI Collaboration Skill 40% · Editorial Judgment 30% · Craft Quality 20% · Strategic Fit 10%.

Review Questions

1. Define IMC and explain why “one voice, many volumes” is a better description of integration than “one message everywhere.”
2. Audience fragmentation and message skepticism both made integration more valuable. Explain the mechanism by which each does so.
3. Classify each of the following as paid, earned, owned, or shared — and justify the tricky ones: a brand’s TikTok account; a customer’s unboxing video; a sponsored newsletter placement; the brand’s packaging; a journalist’s product review.
4. A regional brand has 85% aided awareness in its home market but flat sales and weak preference scores. Using the hierarchy of effects, diagnose the

- communication problem and state what kind of objective the next campaign needs.
5. “The cost of a draft collapsed, so judgment became the scarce resource.” Trace this argument step by step. What, specifically, still commands a premium, and why do those things resist automation?
 6. Rewrite this prompt using the full R-T-C-C-F structure, inventing reasonable strategic details: “Write an Instagram caption for our new running shoe.”
 7. Name the four AI failure modes and give a realistic example of the damage each could cause in a live campaign.
 8. A teammate says, “The model gave me this churn statistic, and it sounds right, so I put it in the client deck.” Using the accountable author standard, explain what should have happened — and who owns the consequences if the number is wrong.

Discussion Prompts

1. Some argue AI-generated marketing will flood every channel with average content, making distinctive human creativity *more* valuable; others argue it will erode the wage premium for creative work entirely. Which force wins, and for whom?
2. Should audiences always be told when campaign content is substantially AI-generated? Construct the strongest case for each side before choosing.
3. Harbor & Vine does with four people what once took twenty. Is that story about job destruction, job transformation, or agency democratization? What evidence would change your mind?

Key Terms

- **Accountable author standard** — The course’s rule that the student is the author of record for everything submitted: verified, edited, and signed, regardless of what the machine drafted.
- **Advertising** — Paid, identified, one-to-many brand messaging placed in media the brand does not own.
- **Brand equity** — The commercial value of those associations: the price premium, loyalty, and forgiveness a strong brand earns over an unbranded equivalent.
- **Brand** — The network of meanings, memories, and associations a name evokes in people’s minds — an asset that exists in the audience, not in the logo file.

- **Communication objectives** — What a message must change in the audience's head or behavior, stated measurably; the standard against which creative and media are judged.
- **Direct response** — Communications built to produce a measurable action from the recipient, with every dollar judged by the actions it produced.
- **Directed iteration** — Working with AI in deliberate rounds — brief, review, redirect — rather than accepting first drafts.
- **Encoding and decoding** — The communication-model idea that a sender translates meaning into a message (encoding) and the receiver reconstructs meaning from it (decoding) — not always the meaning intended.
- **Field of shared experience** — The overlap in knowledge, culture, and reference points between sender and receiver that makes decoding possible; messages outside it misfire.
- **Hallucination** — An AI failure mode: the confident generation of false facts, figures, or sources; the reason the course's verification discipline exists.
- **Hierarchy of effects** — The staged model of persuasion (awareness → knowledge → liking → preference → conviction → action) used to match objectives to a customer's stage.
- **Integrated marketing communications (IMC)** — The strategic coordination of all brand messages and touchpoints so the audience experiences one coherent brand argument rather than disconnected promotions.
- **Large language model (LLM)** — The AI technology behind tools like the Studio Coach: a system trained on text to produce fluent language — and the course's central object of both use and skepticism.
- **Noise** — Anything that interferes between message and receiver — clutter, distraction, competing messages.
- **Personal selling** — Person-to-person persuasion and relationship management, the oldest channel in the mix.
- **PESO model** — A map of modern channels by ownership: Paid (bought), Earned (coverage), Shared (social), and Owned (your properties).
- **Positioning** — The deliberate choice of what a brand should stand for, for whom, against which alternatives.
- **Promotional mix** — The classic toolkit of marketing communications: advertising, public relations, sales promotion, direct response, and personal selling.
- **Public relations** — Earning coverage, credibility, and goodwill through relationships with media and publics rather than purchased space.

- **R-T-C-C-F prompt structure** — The course’s briefing format for AI work: Role, Task, Context, Constraints, Format.
- **Sales promotion** — Incentives designed to change behavior now — trial, repeat purchase, loading, switching.
- **Situation analysis** — The disciplined review of brand, customer, competitor, and context that precedes strategy.
- **Sycophancy** — An AI failure mode: agreeing with whatever the user seems to want to hear, which makes uncritical prompting self-confirming.

References and Further Reading

Scholarly foundations

Batra, R., & Keller, K. L. (2016). Integrating marketing communications: New findings, new lessons, and new ideas. *Journal of Marketing*, 80(6), 122–145. — *The modern research synthesis on why and how integration works; the intellectual backbone of Section 1.1.*

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. — *Large-scale field evidence that AI assistance lifts less-experienced workers most: the “floor rose” argument of Section 1.5.*

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114–126. — *Why humans over- and under-trust machine judgment — context for the verifier role.*

Keller, K. L. (1993). Conceptualizing, measuring, and managing customer-based brand equity. *Journal of Marketing*, 57(1), 1–22. — *The canonical account of brand equity as a network in customers’ minds (Section 1.1).*

Kliatchko, J. (2008). Revisiting the IMC construct: A revised definition and four pillars. *International Journal of Advertising*, 27(1), 133–160. — *A careful tour of how the field has defined and debated IMC.*

Lavidge, R. J., & Steiner, G. A. (1961). A model for predictive measurements of advertising effectiveness. *Journal of Marketing*, 25(6), 59–62. — *The original hierarchy-of-effects paper behind Section 1.3.*

Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. — *The controlled experiment on generative AI in professional writing tasks cited throughout Section 1.5.*

Ries, A., & Trout, J. (1981). *Positioning: The battle for your mind*. McGraw-Hill. — *The classic statement of positioning; short, opinionated, still worth an evening.*

Schultz, D. E., Tannenbaum, S. I., & Lauterborn, R. F. (1993). *Integrated marketing communications: Putting it together and making it work*. NTC Business Books. — *The book that named the discipline*.

Shannon, C. E., & Weaver, W. (1949). *The mathematical theory of communication*. University of Illinois Press. — *The source of the sender–channel–receiver model every marketing adaptation descends from*.

Industry and professional sources

Dietrich, G. (2014). *Spin sucks: Communication and reputation management in the digital age*. Que Publishing. — *Origin of the PESO model used in Section 1.2*.

Duke University Fuqua School of Business / Deloitte / American Marketing Association. *The CMO Survey* (semiannual). — *Ongoing data on marketing budgets, channel spending, and AI adoption in marketing organizations; useful for current figures each term*.

Edelman. *Edelman Trust Barometer* (annual). — *The standard industry reference on institutional and brand trust, relevant to Section 1.1's skepticism argument*.

Federal Trade Commission. (2023). *Guides concerning the use of endorsements and testimonials in advertising* (16 CFR Part 255, revised). — *The U.S. regulatory baseline for honesty in endorsements, reviews, and disclosure — previewed in this chapter's ethics sidebar and treated fully in Chapter 14*.

Interactive Advertising Bureau / PwC. *IAB Internet Advertising Revenue Report* (annual). — *The standard industry accounting of digital ad spending and channel mix*.

McKinsey & Company. (2023). *The economic potential of generative AI: The next productivity frontier*. — *The widely cited industry sizing of generative AI's impact, with marketing and sales among the most-affected functions*.

A note on currency: the AI practice layer of this field moves quickly. Statements in this book about model capabilities and industry adoption reflect the state of practice at the time of writing; your instructor will flag material updates in the IMC Studio.

Next chapter: Brand strategy, positioning, and voice — and Lab 2, where you write the brand platform your entire campaign will be built on.