# AI 2040: Plan A for Book-Lovers

Abundance, Diversity, Growth, and Innovation

Frederick Zimmerman

July 11, 2026

## A Recommendation, Not a Prediction

Books are not doomed. They are at a fork, and the good branch — abundance, diversity, growth, and innovation — is buildable with tools that already exist. That is the claim of this essay, and it is argued, not assumed.

The framing is borrowed, with credit. On July 9, 2026, the AI Futures Project — Kokotajlo, Larsen, Lifland, Dean, and colleagues — published "AI 2040: Plan A" (blog.aifutures.org / ai-2040.com), a document they took care to label "a recommendation, not a prediction": not what will happen, but "the least bad plan we currently know of," argued for on the merits. Their method is stated with equal care: "work things out step by step rather than rely on high-level abstractions like doom or utopia." Their subject is superintelligence governance — a coordinated US–China slowdown that delays superintelligence to roughly 2040, followed by world GDP growing on the order of 200x and a Citizen's Dividend. This essay does not adjudicate that plan. It asks a narrower question that sits underneath whichever macro-plan prevails: what do books become? And it answers in the same register — a recommendation, not a prediction, worked out step by step from mechanism and evidence rather than from vibes, doom, or utopia.

The evidence comes from an unusual vantage point: inside a working algorithmic publishing laboratory. I run Nimble Books and its production system, descended from the PageKicker software of the 2010s, and I have just finished a 45-chapter book about the state of that art — *State of the Art in Algorithmic Book Publishing* (SOTAPUB) — every chapter of which was produced through the machinery it describes. When this essay cites a validation gate or a production run, it is citing operations, not speculation.

One more credential matters, because optimists should be audited. In 2021 I circulated a Longform Prospectus that made a falsifiable technical prediction. GPT-3 was then capped at 2048 tokens; from the quadratic scaling of self-attention I estimated that book-length generation would require roughly three orders of magnitude more effective capacity, and predicted it would arrive "in a timeframe of a few to several years rather than multiple decades." It did. Long-form generation is now a commodity; the full passage, quoted in Section 3.1, reads today like a description of the present. The optimist in this essay has a verified forecast on the record. That does not make the next forecast right, but it earns the argument a hearing.

The structure follows the fork. Section 2 lays out the ledger as it stands in 2026 — what authors, readers, publishers, and underserved languages are already gaining, each item labeled by how well it is verified. Section 3 argues Plan A step by step: abundance (3.1), the breaking of the quality ceiling (3.2), the escape from the many-buyers-few-dollars collision (3.3), the move beyond coarse categories (3.4), genuinely new kinds of books (3.5), and what remains irreducibly human (3.6). Section 4 then takes Plan B — the dystopia — seriously on its own terms, argues it as rigorously as Plan A, and audits the assumptions it requires. Section 5 states what each actor does that selects one branch over the other.

Plan A is not guaranteed. It is chosen — by authors, readers, publishers, model companies, and educators, in decisions that are being made now. The AI Futures authors were honest enough to call their plan the least bad one they knew of rather than the destined one; the same honesty applies here. The purpose of what follows is to show that choosing Plan A is neither naive nor expensive, because most of its load-bearing components already run in production.

## The Ledger Today

Before arguing about 2040, take inventory of 2026. The doomer premise is that nothing good is happening to books under AI — that the technology's entire effect is flood, theft, and displacement. The ledger says otherwise. What follows is drawn from a 45-phase survey of the vendor landscape conducted in July 2026, plus targeted verification; each item carries its confidence label — *verified reporting* for independently confirmed facts, *market signal* for vendor-asserted claims worth noting but not audited. The discipline matters: an honest optimism keeps its receipts sorted.

**What authors are gaining.** The path from manuscript to market has collapsed in cost. Deterministic formatting tools (Vellum at $199–249 one-time, Atticus at $147) plus AI drafting and continuity aids have compressed what

once required a small production team into a laptop workflow — tool existence and pricing verified, throughput claims market signal. Continuity support that used to require a paid developmental editor now ships as structured data: Sudowrite's Story Bible and Novelcrafter's Codex hold characters, locations, and plot threads that the model consults on every generation call, catching the chapter-20-contradicts-chapter-1 failure that sinks first novels (architecture verified, effectiveness market signal). Authors who cannot afford a $2,000 developmental edit can get a structural gut-check from multi-model beta readers, one of which runs Claude, GPT-4o, and Gemini in parallel and reports cross-model consensus (market signal on quality; the consensus design itself is verifiable). And the slush pile is being read: tools like Storywise triage the 94–97% of submissions that traditionally never got any systematic read at all, with Watkins Media's deputy managing director Vicky Hartley publicly confirming the practice (adoption verified, scoring quality market signal).

New money is appearing where none existed. HarperCollins's Microsoft licensing deal pays consenting backlist nonfiction authors a 50/50 split — roughly $5,000 per title as reported — and O'Reilly pays a "generative AI royalty" to authors whose specific passages are used to construct a cited answer, one of the few concrete compensation mechanisms for AI-mediated reading on record (verified reporting via Publishers Weekly, the Authors Guild, and O'Reilly's own materials; exact contract terms not independently audited). Authors who want to signal human craft now have a market-recognized instrument: the Authors Guild's Human Authored certification, opened to all US authors by March 2026, had certified more than 3,000 authors across roughly 5,000 titles — and it explicitly permits AI for research and outlining while requiring human-written prose (verified reporting). Amazon's Kindle Translate carries English to and from Spanish, German, French, Italian, and Portuguese, with thousands of authors using it and translated output averaging 4+ star ratings (verified via company announcement and Publishers Weekly). A Gotham Ghostwriters survey of 1,481 writers found 61% using AI tools, reporting an average 31% productivity gain concentrated in drafting friction, not voice replacement (market signal — self-reported).

**What readers are gaining.** The European Accessibility Act, in force since June 2025, requires ebooks sold in the EU to be born accessible — screen-reader compatible, braille-transcribable, with DRM that cannot block assistive technology — and InDesign now generates alt text and ARIA-tagged EPUB export by default (regulatory floor verified; the InDesign rollout drew honest craft criticism, which is worth stating rather than smoothing over). AI narration has reached near-human quality for single-narrator work, and HarperCollins uses ElevenLabs for non-English audiobooks explicitly positioned as complementary to hu-

man narration — meaning backlist and niche titles get audio editions the old economics would never have funded (partnership verified, quality claims market signal). NotebookLM lets a reader load a whole book and interrogate it with answers grounded in the uploaded text; O'Reilly's Answers does the same for its technical catalog with citations back to source pages (features verified). Rebind pairs public-domain classics with 20+ hours of recorded commentary from named experts — Margaret Atwood on *A Tale of Two Cities*, Roxane Gay on *The Age of Innocence*, Marlon James on *Adventures of Huckleberry Finn* — using AI only to route reader questions to the expert's own recorded answers, never to improvise new ones; it was a *Time* Invention of the Year (verified reporting; scale figures market signal).

The reading commons is growing, not shrinking. Nancy Drew, Miss Marple, and Sam Spade entered the public domain on January 1, 2026, alongside thousands of 1930 works, many long out of print through ownership uncertainty (legal fact, verified). The StoryGraph passed five million signups in January 2026 and became the first reading-community app integrated with a hardware e-reader; Fable grew to 100,000+ book clubs and three million users (verified reporting). Most telling: the 2026 State of Reading Report found that personal recommendations from people readers actually know have overtaken platforms, social media, and AI tools as the top source of book discovery — evidence that AI tooling is enlarging the ecosystem around reading rather than hollowing out its social core (verified survey, precise percentages market signal).

**What publishers and companies are building.** Ingram eliminated its title setup fee in February 2026 and shipped a metadata assistant and cover generator — verifiable reductions in the professionalism barrier that separated small presses from Big Five polish (verified via Ingram's own blog and Shelf Awareness). OCLC added AI-suggested classification and subject headings to the world's largest shared bibliographic database, explicitly framed as augmenting cataloger judgment, not replacing it (verified). Springer Nature built its Curie editing tool because non-native-English-speaking scientists need an estimated 51% more time to write a publishable paper, and gives it away free — a rare AI tool with an equity rationale stated up front rather than retrofitted as PR (verified reporting, scale figures company-reported). O'Reilly's in-house summary of its attribution pipeline — "the R in RAG stands for Royalties" — is the clearest existing proof that citation-bearing AI reading and author compensation are not mutually exclusive (verified from O'Reilly's published materials). Trust infrastructure is being built across competitors: C2PA content provenance, ratified as ISO/IEC 22144 in 2025, is now signed onto output by OpenAI and Google alike (standards publication verified). Perplexity's Publishers' Program pays a revenue

share when content is cited in an AI answer, with roughly 2,400 partners enrolled — no book publisher among them yet, but the mechanism is exactly what publishing needs built before, not after, answer-engines dominate discovery (program verified, figures market signal). Even the gatekeeping is real: Draft2Digital reportedly declines up to 70% of monthly submissions and added fees specifically to curb low-quality AI volume, and Amazon caps KDP uploads at three titles per day (policy actions verified, the percentage market signal). These are distribution intermediaries spending money to keep signal above noise.

**What underserved cultures and languages are gaining.** Kindle Translate and Taylor & Francis's program to translate scholarship into English from more than 30 languages both attack the same structural gap: languages whose audience size never justified a human translation budget (both verified company announcements). Google's WAXAL dataset spans 21 African languages; the Masakhane research community has published translation tools covering more than 48; Lelapa AI's InkubaLM serves five African languages — Swahili, Yoruba, isiXhosa, Hausa, isiZulu — covering more than 350 million speakers (verified reporting; model-quality claims market signal, and independent reporting is candid that African languages remain nearly absent from training data — this is progress against a hard problem, not a solved one). A study of two million biomedical papers found AI-assisted-writing adoption grew 400% in non-English-speaking countries after ChatGPT versus 183% in English-speaking ones — consistent with the tools disproportionately helping researchers who paid the largest time tax to participate in English-dominant scholarship (study finding verified; the equity interpretation reasonable but not adjudicated). PublishDrive reports 80%+ year-over-year growth in non-English audiobook titles (market signal). Indigenous-language efforts pairing AI translation with native-speaker review are early-stage and mixed — a documented South African project performed poorly on Zulu educational text so far — worth tracking, neither dismissed nor oversold.

This is one year's ledger, labeled honestly. It does not prove Plan A. It defeats the premise that Plan A has nothing to build on.

## Plan A, Argued Step by Step

Plan A rests on six arguments, each made from mechanism or evidence; the first two — that scarcity is now optional, and that quality is an engineering discipline rather than a casualty — carry the rest.

## The Abundance Thesis

The historical condition of books is scarcity, and it is worth saying plainly that the scarcity was never romantic. In 2021, before long-form generation existed, I put the argument this way in the Longform Prospectus:

> The immersive deep reading of high-quality books must rank among the most beneficial and broadly distributed technologies ever invented (see inter alia McLuhan, 1962; McDermott, 2006; Boorstin, 1992; UNESCO 2019). Yet this vital experience remains tightly constrained by scarcity: of classics, of representation, of accessibility, of sustainability, of time. It's time to flip the script to a world of abundance.

Each scarcity in that list is a real ration. Scarcity of classics: most of what deserves to be in print is not. Scarcity of representation: most cultures and languages cannot staff a curriculum from their own shelf. Scarcity of accessibility: large print, audio, and adapted reading levels are afterthoughts funded only where mass markets exist. Scarcity of sustainability: the economics of traditional lists kill most worthy books before they are written. Scarcity of time: authors ration their own ambitions against a finite writing life. None of these scarcities is a law of nature. All of them are artifacts of production cost — and production cost is exactly what has collapsed.

What abundance looks like was also specified in 2021, and I quote it in full because it remains the target:

> A world where the dominant dopamine drip is not the addictive rage spiral of social media but the transformative insight and empathy of great books. A world where every culture and language can be represented in curricula with equal ease. A world where every author has equally powerful co-creators: new and globally available Muses. A world where every reader has an infinite bookshelf. A world where there are a thousand new Agatha Christies. A world where a helpful AI Muse saves George R.R. Martin from his tendency to be easily distracted. A world where our office learnings are instantly translated into immersive reading artifacts for our teams and our customers. A world where a group of Nigerian teenagers narrates their 4C story in a shared universe of novels that are read by millions.

Note what that litany is not. It is not a promise of infinite content; social media already delivers that, and the first sentence names it as the adversary. It is a promise of abundant *books* — structured, immersive, high-quality long forms — for readerships the old economics never reached. The Nigerian teenagers are

the point: not more supply for saturated markets, but first supply for unserved ones.

Why believe the person making this promise? Because the same document made a technical forecast that could have failed, and did not:

> This isn't possible just yet. OpenAI's GPT-3 model has an almost magical capability to generate endless streams of plausible, interesting text, but it is currently limited to no more than 2048 tokens per request (1500 words). As I understand it, the computing resources required for the key "self-attention" step currently scale on a (mostly) quadratic basis, which means that doubling the size of the request equals quadrupling the response time. So–check my work, please, OpenAI!–creating a "long-form" completion the size of an 80,000-word book– 53x larger–would require 53^2 (2809x) more resources. Increasing speed by three orders of magnitude is a big deal, but the history of computing technology suggests this will happen eventually, and probably in a timeframe of a few to several years rather than multiple decades. So an abundance of long-form narrative is coming; the only question is when, who and how.

That was written from an engineering analysis — quadratic attention, three orders of magnitude, the historical cadence of computing — and the window closed on schedule. Long-context models now hold book-length material as a matter of routine, and "generation is cheap enough that a book-shaped draft is a commodity" (SOTAPUB, ch. 1). The 2021 document ended with four words: "This is within reach." The abundance thesis of Plan A is that sentence, promoted from forecast to observation. What remains to argue is that abundance need not mean slop — which is the next step.

## Breaking the Quality Ceiling

Slop is a choice, not a destiny. The doomer syllogism — generation is cheap, therefore quality collapses — skips its middle term: it assumes no one builds the machinery that separates cheap generation from published work. That machinery exists, runs in production, and is the subject of roughly a quarter of SOTAPUB.

Begin with the reframe that dissolves most of the argument. "AI writes slop" is a category error about drafts: "A zeroth draft may contain brilliance, but nobody has yet agreed to be answerable for it… A first draft is the earliest version of a text for which someone, human or system, has accepted that answerability" (SOTAPUB, ch. 18). Slop is unpromoted zeroth-draft material shipped without

a promotion gate. Quality control is the gate — and the gates are software, which means they are testable, versionable, and cheap to run on every title rather than only the lead titles.

What the gates look like in practice: deterministic defect batteries that run before any model's self-report is trusted, because a real failure taught the house that "A first-draft gate trusted without a deterministic backstop is not a gate; it is a hope with a log file" (SOTAPUB, ch. 18). Synthetic reader panels run in deliberately plural configurations — focus group, engagement curve, sensitivity read, senior sanity check — because a single aggregate score hides failures a plural read catches (SOTAPUB, ch. 19). The panels are themselves audited: the judging harness is separated from the generating agent by construction, because "a system that can rewrite its own judge is not a judge" (SOTAPUB, ch. 5). And the terminal gate is semantic, not cosmetic: "Final approval is a semantic checkpoint, not a box-ticking exercise over spelling and file validity" (SOTAPUB, ch. 23) — a deliverable must satisfy the whole contract it was commissioned under, with a deliberate subset of checks reserved for human sign-off. None of this machinery pretends to replace human readers; its stated purpose is "to widen judgment, not counterfeit it" (SOTAPUB, ch. 22).

Does the machinery actually hold at scale? The exhibit is a production run: 48 distributor-ready books, 10,988 pages, produced in three weeks for $42 of total compute, with 100% compilation success and under 5% rework. The operator's own accounting of why: "100% success didn't happen by accident. It happened because of paranoid validation at every checkpoint" (SOTAPUB, ch. 27, Exhibit 27.1). The 45-chapter book this essay accompanies is itself a second exhibit — produced through the gates it documents.

There is a historical arc here, and it favors optimism. Early print was a flood of corrupt texts, unauthorized editions, and outright piracy; the librarians of Alexandria were marking suspect lines with critical signs twenty-three centuries before anyone worried about model output (SOTAPUB, ch. 20). Craft standards emerged because readers and publishers demanded them, and the same maturation is now underway for AI-assisted books — but faster, because this time the standards are executable. A stylesheet had to be enforced by a human; a validation gate enforces itself on every build. Diligence, likewise, becomes a competitive moat rather than an overhead: "both will have ancestors; only one will be able to produce the family records" (SOTAPUB, ch. 21).

The economics point the same direction as the ethics. "Fluency is now free; structure is not" (SOTAPUB, ch. 19) — which means the market value migrates to exactly the layer slop-merchants skip. And rigor, at abundance scale, is not gatekeeping snobbery; it is reader protection: "The flood of machine-generated

prose makes severity kinder, because it prevents whole books from hardening around avoidable weakness. Severity here is not cruelty. It is mercy toward the reader" (SOTAPUB, ch. 19). The quality ceiling was never a ceiling on what machines can draft. It was a ceiling on what unaccountable processes can publish — and accountability, it turns out, compiles.

## Escaping the Many:Many Collision

The strongest argument against abundance is arithmetic. If anyone can generate a book, supply approaches infinity; buyers are finite; income per title approaches zero. The arithmetic is sound. The premise is not. It assumes the demand side of publishing is fixed — that everyone who could be a reader already is one, already served in their own language, at their own reading level, on their own subjects. None of that is true, and the distance between assumed and true is where Plan A operates.

The 2021 Prospectus named the unserved demand directly: "A world where every culture and language can be represented in curricula with equal ease." And, more specifically: "A world where a group of Nigerian teenagers narrates their 4C story in a shared universe of novels that are read by millions." These are not hypothetical readers conjured to rescue a business model. They are people for whom the old economics never closed: translation budgets that required a projected audience no small language could supply, narration budgets reserved for lead titles, acquisition thresholds that made a community's own story commercially invisible.

Those economics are closing now, and the ledger of Section 2 is the receipts: Kindle Translate carrying authors into five new language markets, Taylor & Francis translating scholarship out of more than thirty languages that never justified a human translator's time, the Masakhane and Lelapa work pushing machine capability into African languages spoken by hundreds of millions, the European Accessibility Act making born-accessible the legal floor, AI narration giving audio editions to backlist titles no narration budget would ever reach, large print engineered as a standard product line rather than a charity case. Each of these is the same move: a reader who existed all along, finally addressable.

My own shop applies the move as a design discipline rather than a hope. The house evaluation rubric asks of every manuscript: "Can this find 1,000+ genuine readers who will treasure it?" (SOTAPUB, ch. 15). The validation doctrine behind it holds that "one does not need the broadest possible audience; one often needs the right 5,000 readers rather than an imaginary 50,000" (SOTAPUB, ch. 9). A book aimed at a thousand genuine readers does not collide with anything. It is not competing for the fixed attention of the same exhausted buyer; it

is meeting a demand that no previous supply ever reached. Supply of this kind creates its market rather than dividing an existing one.

The honest caveat is that new supply meets new demand only if the two can find each other, and here the record includes a self-administered lesson. An audit of my own catalog found the publisher-level machine-readable index serving correctly while "per-title agent discoverability for the remaining 655 titles was effectively zero" (SOTAPUB, ch. 36, Exhibit 36.2) — 655 of 656 titles dark to the AI agents that increasingly mediate discovery, an infrastructure defect, not a demand defect, and fixable as infrastructure. The stakes of fixing it are stated plainly in the same book: "When readers' questions are answered by systems that retrieve over structured book data, the publisher whose subject codes, qualifiers, and ONIX records are clean is the publisher whose books exist" (SOTAPUB, ch. 29). Or, compressed to a single design brief: "The blurb was invented to sell a book to someone who couldn't browse it. Its successor must sell the book to something that reads everything and browses nothing" (SOTAPUB, ch. 33).

The many:many collision, in other words, is not a law of nature. It is a failure of matching — and matching is precisely what this generation of machines is best at.

## Beyond BISAC: Millions of Living Categories

Roughly five thousand BISAC subject codes currently ration the attention of the entire American book trade. The rationing is explicit: BISG's own guidance is to "choose the most specific codes possible, generally no more than three when three fully describe the book" (SOTAPUB, ch. 29). Three coarse labels per book, drawn from a fixed list, is the compression through which every store shelf, retail algorithm, and library feed must pass. The category system, not reader interest, is the bottleneck — and the category system is the part we can now rebuild.

Librarians have known the fixed list was a fiction for a century. S. R. Ranganathan's faceted classification decomposed a subject into independent, recombinable facets precisely because, as SOTAPUB puts it, "A book is almost always about more than one thing, and a controlled vocabulary can only pretend otherwise by fiat" (SOTAPUB, ch. 29). The trade adopted the fiat because enumerative lists were the only thing a purchase order could process. That constraint is gone.

The 2015 PageKicker Manifesto planted the seed of the alternative: "The PageKicker catalog includes at least one book for every node in the BISAC nonfiction hierarchy." The ambition was comprehensive coverage of the existing map — with a quality floor attached: books that "pass the duck test: they look like a book

and quack like a book. And sell like a book." A decade later the extension is obvious: not every node of a fixed taxonomy, but a category system as fine-grained as readers' actual interests. The working version already exists in my shop. A production module embeds every BISAC code into a shared vector space: "Treat the taxonomy as a space, not a list", exposing "nearest neighbors, farthest members, greedy walks between codes" (SOTAPUB, ch. 30). In a space, a category is not a slot assigned by committee; it is a neighborhood, discoverable at any resolution the readership warrants. The neighborhoods are real, measurable objects: Ted Underwood's corpus research found that "some genre boundaries — detective fiction — stay machine-separable for a century while others blur" (SOTAPUB, ch. 7). Categories live and die like species. A classification system should be able to watch them do it.

The production side has caught up with the map. In the current ideation program, "a gridded generator crosses the full BISAC subject taxonomy with twenty-one production archetypes — definitive reference, data atlas, casebook, cross-domain survey, and so on — producing on the order of a hundred thousand concept stubs in seconds, for nothing, with no model call at all" (SOTAPUB, ch. 5). A hundred thousand candidate categories, each a potential shelf, evaluated by selection machinery rather than rationed by a code list.

The economics close the argument. When a documented production run delivers 48 validated books for $42 in direct cost (SOTAPUB, ch. 27), a category with two hundred enthralled readers is a viable publishing category — not a loss leader, not a vanity project, a going concern. The long tail has been a demand-side observation since Brynjolfsson, Hu, and Smith measured the consumer surplus of Amazon's long-tail variety in 2003 (SOTAPUB, ch. 9). What it never had was a production side. Now it does. Millions of living categories, each with its small thriving readership, is not a metaphor; it is a catalog architecture with the tooling already in beta.

## New Kinds of Books

Abundance and finer categories improve books we already know how to make. Innovation, the fourth word in this essay's subtitle, means reading experiences that have not existed before. Three are already past the prototype stage, and each can be stated as a mechanism rather than a marvel.

The first is the computable book. "A computable book exposes structure, claims, citations, and data so machines can assist readers without replacing reading" (SOTAPUB, ch. 40). The mechanism is addressability: a book whose claims, apparatus, and boundaries are machine-legible can be verified, cited, indexed, and recombined without a human re-keying anything. Done correctly it costs

nothing at press time: "The result is a book that is simultaneously a printed object and an agent-addressable resource, with no second authoring pass and no opportunity for the two representations to disagree" (SOTAPUB, ch. 40). The codex stays a codex; it also becomes infrastructure.

The second is the conversational book. The mechanism is interrogation with receipts: a reader asks the book a question and receives an answer grounded in — and cited to — the actual text. The discipline that keeps this honest is refusal: "A book that can answer is only useful if it can also point, delimit, and sometimes refuse" (SOTAPUB, ch. 41). And the economics are no longer hypothetical: "O'Reilly built forensic attribution into the pipeline so it can pay a 'generative AI royalty' to the authors whose passages are actually used to construct each answer — described in-house as 'the R in RAG stands for Royalties.' " (SOTAPUB, ch. 41). That is verified from O'Reilly's own published materials, and it matters beyond one vendor: it is an existence proof that machine-mediated reading and author compensation can be designed into the same pipeline.

The third is the adaptive and agentic book. SOTAPUB draws the distinction cleanly: "An *adaptive* book changes what a given reader encounters: reading paths, glossary depth, worked examples, remediation. An *agentic* book is queried and acted upon by software" (SOTAPUB, ch. 42). The design law governing both comes from the market's own verdict: "The products surviving contact with readers (retrieval shells, expert-voice matching, cited answers) all hold the text fixed; the expensive failures (autonomous courseware, mandated national adaptivity) all promised to change the content itself faster than anyone could audit it" (SOTAPUB, ch. 42). Rebind, described in Section 2, is the exemplar: "Core text fixed; expert's voice fixed; only the retrieval path dynamic" (SOTAPUB, ch. 42). Not every entry survives scrutiny — Amazon's "Ask this Book" is verified as a feature, but its net effect on reading is contested, not settled — and that unevenness is what a real innovation frontier looks like.

What these forms share is what they refuse. Each adds a capability at the book's edge while holding its core — fixed text, named authorship, checkable citations — invariant. The book as an institution survives these innovations for the same reason it survived the paperback: the new vessel carries the old promise.

## What Remains Human

Every argument in this Part fails if quality and accountability are rhetorical. Optimism about abundance is credible only if someone remains answerable for what ships — and answerability, to count, must be implemented, not professed.

SOTAPUB's final chapters state the principle as a specification rather than a sentiment: " 'What should remain human' is not a residue — the leftover tasks the machine cannot yet do — but a positive specification: the set of decisions for which a named person must answer, whatever the tooling on either side" (SOTAPUB, ch. 44). In my own production system that specification is executable. A frozen set of human-gated task types is enforced in the orchestrator itself; the distribution validator marks specific checks MANUAL by design; and "No automation is permitted to submit a book to the printer" (SOTAPUB, ch. 12). The rule compresses to four words: "Automate the preparation, never the signature" (SOTAPUB, ch. 44).

This is not a new covenant invented to soothe AI anxiety. It is the oldest fact in the trade: "From the beginning, the human in the chain was the person who could be sued" (SOTAPUB, ch. 44) — the Stationers' charter of 1557 existed so that someone was prosecutable for every printed page. A publisher who cannot name which decisions carry a human signature, SOTAPUB argues, does not have an AI policy; it has a mood. The gates are how the mood becomes a policy.

Nor do the instruments replace judgment; they are built to extend it. The purpose Section 3.2 quoted — to widen judgment, not counterfeit it — is the honest promise of the whole apparatus — more manuscripts seriously read, more defects caught early, more small readerships found — with the deciding mind still human and still named. The instruments scale the preparation; they never inherit the verdict. That division of labor is what makes the scale of Part A's ambitions compatible with the standards of the craft.

The last word of Plan A's affirmative case belongs to the boundary itself: "The validator can enforce that a record is consistent; only an editor can decide that a book deserves to exist. The infrastructure exists to make that judgment cheaper to exercise and harder to fake — not to replace it" (SOTAPUB, ch. 45). That is the shape of the future worth choosing: abundance with a signature on it. Growth, with conscience — engineered in, gated in code, and answerable to a person with a name.

## Plan B: The Dystopia, Taken Seriously

An argument for Plan A that cannot state Plan B at full strength is advocacy, not analysis. The AI Futures Project earns the right to recommend by working the dark branches honestly, and this essay owes its readers the same discipline: mechanism before conclusion, evidence before verdict. So here is Plan B, argued as carefully as I know how — not the doomer caricature, but the serious case.

Every step below is underway somewhere, observable today. The question Plan B poses is whether the steps compose into a trajectory. Take each in turn.

**The supply cataclysm.** It begins with volume. Bowker reported registered U.S. title output rising 32.5% year over year in 2025 to more than four million titles, self-published titles up 38.7% to over 3.5 million — and as SOTAPUB observes, "no human metadata team scales to four million annual titles" (SOTAPUB, ch. 33). In Plan B, nothing else scales to it either. Curation loses the arms race with generation permanently: every gate a retailer builds, the flood learns to pass. Provenance collapses next — with no reliable way to establish who or what made a book, every title carries the discount of suspicion, and honest publishers pay the fraudulent publisher's tax. Then retailer trust fails. The intermediaries respond with blunt instruments — upload caps, account fees, mass delistings — that punish the innocent alongside the guilty, and readers, burned enough times by confident-looking garbage, stop trusting the store shelf altogether. The cautionary microcosm already exists: Spines, the AI-native publisher that drew industry backlash for automating distribution mechanics without a corresponding verification gate, reproducing what SOTAPUB calls "vanity-press economics under an AI label" (SOTAPUB, ch. 32). In Plan B, Spines is not an outlier that failed; it is the median firm that succeeded. Trust is a stock depleted by a flow of bad experiences, and the flow has no ceiling.

**Readers as rare as operagoers.** Demand decays after supply corrodes it. In Plan B, the general reader — the person who buys eight books a year on no particular schedule — concludes that finding a good book is no longer worth the search cost, and drifts to media where the floor is higher because the format is shallower. Reading long-form books becomes what opera became: a genuine art, magnificently maintained, publicly subsidized, demographically narrow, and culturally beside the point. The book's prestige dies with its scarcity. For five centuries, "author" named an achievement precisely because a book was hard to make; when anyone can generate one in an afternoon, the title confers nothing, and the cultural cachet that made people want to be seen reading, gifting, and shelving books evaporates. The book as an object of aspiration — the thing a young person wanted to have written — is gone, and with it the tributary that refilled the reading public each generation.

**Books get crummy, and authors bifurcate.** Race-to-the-bottom economics does the rest. When fluent text costs nothing, price competition drives the market toward whatever is cheapest to make that still clears the (falling) bar of reader tolerance. The midlist dies first — it always does — and the authors who remain sort into two castes. The slop-meisters optimize for volume: the indie operator running hundreds of titles across dozens of pen names becomes not an

outlier but the business model, craft subordinated entirely to throughput. The craftspeople keep their standards and lose their income, subsidizing their writing with other work, their books indistinguishable at the point of sale from the flood around them. Neil Clarke's slush pile at Clarkesworld — flooded with machine submissions "nowhere near the standards we expect" (SOTAPUB, ch. 44) — is the leading indicator: in Plan B, every acquisition channel becomes that slush pile, and unlike Clarke, most gatekeepers give up.

**The extraction scenario.** Above the wreckage of the trade sits the layer that captured its economics. In Plan B, the model companies absorb the margin of every content industry, publishing included, the way platforms absorbed the margins of news. The training corpora were taken first and paid for later, if at all: the Anthropic settlement — "roughly $3,000 per work across some 500,000 books with a claim rate above ninety percent… the largest copyright settlement in U.S. history" (SOTAPUB, ch. 21) — reads, in this scenario, not as accountability but as liquidation pricing: a one-time payment, roughly the price of a used car, for the perpetual capability distilled from a lifetime of craft. The HarperCollins–Microsoft license — a reported $5,000 per title, split 50/50 with the author, output capped at 200 words or 5% of a book (SOTAPUB, ch. 42) — becomes the industry's reference price for the same trade: modest fixed sums exchanged for the capacity to satisfy, forever, much of the demand the books themselves once served. Downstream, publishing persists as a thin user interface over foundation models that keep the economics — imprint logos skinning an API. SOTAPUB's own warning names the terminal condition: a house without real apparatus "is just another spigot on an infinite reservoir" (SOTAPUB, ch. 45). In Plan B, every house is a spigot, and the reservoir's owner collects the rent.

**Institutional retreat.** The deepest damage is quiet and takes a generation. Colleges stop requiring whole books; campus reporting in the mid-2020s already described elite students arriving without ever having been assigned one in high school. In Plan B, the accommodation becomes doctrine: the assigned book — the shared, complete, demanding text as a unit of education — dies without a vote ever being taken, replaced by excerpts, summaries, and model-mediated digests of texts nobody opens. The institutional appetite to reverse course is spent by expensive failures: South Korea's national AI-textbook program, roughly $850 million committed and then stripped of official status by the National Assembly after about two semesters and a 56,000-signature parents' petition (SOTAPUB, ch. 42), teaches ministries a double lesson — the technology disappointed, and the political cost of ambition was severe — so they retreat not to books but to caution. Long-form attention then decays generationally by simple recursion: cohorts assigned fewer books become teachers who assign fewer

still, and the capacity for sustained argument — the cognitive skill books uniquely train — becomes as exotic as Latin.

**The terminal image.** Follow Plan B to its end and you do not find zero books. You find something stranger and sadder: books and readers as the core of a Butlerian-jihad samizdat — a culture that, having been deceived at scale by machine text, trusts only the provably human book. And there is exactly one medium that a model cannot fake at scale: the handpress. Letterpress printing undergoes a genuine renaissance in this world, and a book person cannot pretend not to love it — type set by hand, ink and impression into cotton paper, every copy carrying physical proof of human hours, provenance you can feel with a fingertip. It is the one beautiful thing in Plan B, and it deserves to be said lovingly: the handpress book is the incorruptible book. But name the tragedy plainly. A hand press serves thousands, not billions. The immersive deep reading that the 2021 Prospectus called "among the most beneficial and broadly distributed technologies ever invented" ends up neither beneficial at scale nor broadly distributed — it ends up curated, venerated, and irrelevant to how civilization actually thinks. Plan B's endpoint is not the death of the book. It is the book as relic.

**The load-bearing assumptions.** That is the serious case, and it deserves a serious dissection. Plan B is a conjunction: it requires all five of the following to hold, simultaneously, for fifteen years.

(i) *Quality gates never mature or never get adopted.* Contestable, and already being falsified: the entire editorial machinery of Section 3.2 exists and ships books, distributors are screening at the gate (Draft2Digital reportedly declining up to 70% of monthly submissions, Amazon capping daily uploads), and the market documented in Section 2's ledger is investing in verification, not just generation.

(ii) *Demand is fixed and saturated.* Contestable: the ledger records new readerships being created now — machine translation into markets that never had translation budgets, accessibility mandated into the default, audiobooks for titles that never justified narration. Supply that creates demand is not dividing a fixed pie.

(iii) *Provenance infrastructure fails.* Contestable: C2PA content credentials are now an ISO standard with cross-competitor adoption; ONIX 3.1 carries a provenance composite; the Authors Guild's Human Authored certification covered roughly 5,000 titles by 2026. The infrastructure Plan B requires to fail is being built and adopted on the record.

(iv) *Rents flow entirely to model owners, with no countervailing force.* Contestable: the Anthropic settlement — read straight rather than darkly — estab-

lished that taking corpora has a price; HarperCollins negotiated terms rather than suffering extraction; O'Reilly pays passage-level royalties for AI-mediated reading; citation-triggered payment programs already exist. Courts, contracts, and competition are all countervailing forces with observable outputs.

(v) *Education abandons long-form permanently.* The strongest of the five, and still contestable: the same mid-2020s reporting that documented the retreat provoked institutional counter-reaction, and South Korea's reversal cuts both ways — it shows publics willing to fight, loudly and successfully, over what their children read.

Each assumption is contestable; several are already being falsified by evidence catalogued in Section 2. A scenario that requires a five-way conjunction, several conjuncts of which are visibly failing, is worth taking seriously — that is what this section has done — but it is not a forecast. Plan B is a scenario. It is not a trend line.

## The Fork

Plan A is not a prediction and not a drift; it is a choice, distributed across five kinds of actors, each of whom selects a branch through ordinary decisions made this decade.

**Authors** select Plan A by treating the machine as instrument rather than replacement and by making the distinction legible: certify human authorship where it is true (three thousand authors already have), license training rights on negotiated terms rather than defaulting into extraction, and use the new leverage — continuity tooling, synthetic first readers, translation into markets no agent ever sold — to write books that were previously impossible, not merely faster versions of books that already existed.

**Readers** select Plan A with money and attention: buy from houses that disclose how their books are made, favor retailers that police their shelves, review honestly so the trust signal stays informative, join the reading communities that are demonstrably growing — StoryGraph's five million, Fable's hundred thousand book clubs — rather than lamenting ones that shrank, and keep doing the one thing the 2026 State of Reading Report found still beats every algorithm: pressing a book on a friend by name.

**Publishers** select Plan A by building the apparatus this essay's companion volume documents: validation gates that make "release" a promotion state rather than a mood, provenance records that can survive an audit, metadata clean

enough that agent-mediated discovery finds their books, prices and formats engineered to the markets the new tooling opens, and a signature — a named human answerable for every title — that no spigot can counterfeit. Ingram dropped its setup fees; the cost floor is gone, and with it the last excuse for a dark catalog.

**Model companies** select Plan A by paying at the meter instead of at the settlement table: licensed corpora on HarperCollins-style negotiated terms, passage-level attribution and royalties on the O'Reilly pattern, citation-triggered revenue sharing extended from news to books, and content credentials signed by default. The alternative was priced in court at three thousand dollars a work, plus the distrust.

**Educators** select Plan A by holding the line that Korea's parents held: the whole assigned book, read in common, as a non-negotiable unit of education — while using the new instruments (grounded question-answering on the actual text, study apparatus generated against the assigned edition, versions matched to reading level and home language) to bring more students through more whole books, not fewer. The retreat from the assigned book was an accommodation, never a pedagogy; it can be reversed by syllabus, one course at a time.

None of this requires heroism. It requires ordinary professional choices, made by people who can see the fork. In 2021, before any of the present tooling existed, the Longform Prospectus closed its argument for abundance with a forecast: "This is within reach." The prediction it rested on came true on schedule. In 2026, with the gates built, the ledger filling, and the fork visible — read it again, not as a forecast but as a status report. This is within reach.

## Sources


- AI Futures Project (Kokotajlo, Larsen, Lifland, Dean, et al.). "AI 2040: Plan A." July 2026. blog.aifutures.org / ai-2040.com.
- Zimmerman, Frederick. *Longform Prospectus.* 2021. (Quadratic-attention forecast; "This is within reach.")
- Zimmerman, Frederick. *PageKicker Manifesto.* 2015.
- Zimmerman, Frederick. *State of the Art in Algorithmic Book Publishing* (SOTAPUB). Nimble Books, 2026. Cited by chapter: chs. 21, 32, 33, 42, 44, 45.
- Bowker. U.S. ISBN registration statistics, 2025 (via SOTAPUB, ch. 33).
- Publishers Weekly / Authors Guild reporting: Anthropic settlement; HarperCollins–Microsoft license; Human Authored certification (2025–2026).

- Everand/Fable. *2026 State of Reading Report.* (Market signal on discovery percentages.)
- Horowitch, Rose. "The Elite College Students Who Can't Read Books." *The Atlantic,* October 2024.
- Landscape items and confidence labels per the companion evidence file "The Generous Ledger" (2026); vendor-asserted figures remain market signal.