

Page Two AI Policy

Last updated: August 2026

Last updated by: Gabrielle Narsted

I. Preamble

Page Two recognizes the transformative potential of artificial intelligence (AI) in the publishing industry, while also acknowledging the ethical, legal, and practical considerations it raises. This policy aims to guide the responsible and strategic use of AI across all aspects of our business, ensuring alignment with our core values of care, curiosity, creativity, excellence, integrity, and partnership.

AI tools assisted in identifying gaps in this policy and in drafting the definitions in Section VII. All content has been reviewed, refined, and approved by Page Two.

II. Scope

This policy applies to the use of AI including generative AI, large language models, predictive AI, image AI, transcription AI, and analytics. It applies to all departments (including employees and contractors) within the organization, along with our work with external partners and clients. It applies to the access and use of AI tools and technologies for Page Two purposes, or otherwise in the course of employment or engagement with Page Two, both during and outside of normal hours of work, whether on a Page Two-owned or personal device, and whether at Page Two's office or from a remote work location. See the definitions section below for more details.

Page Two permits the use of approved AI tools and use cases for work-related activities only in accordance with this policy. Approval conditions for AI platforms and use cases are determined by Page Two's chief operating officer. Please note that these conditions are updated from time to time. It is each user's responsibility to consult this policy regularly to ensure that AI use is approved at the time of use.

This policy has two layers. This document sets out Page Two's AI principles and firm commitments, which can change over time. Alongside it, the chief operating officer maintains a living set of operational controls, including the approved-tools list, prohibited inputs, approved and grey-zone use cases, and review roles. Those controls are reviewed regularly, updated more often than this policy, and are extensions of it.

Freelancers and contractors are not granted access to Page Two's approved AI tools. Freelancers and contractors must not give their personal or any third-party AI accounts access to Page Two manuscripts or other book materials. These requirements are set out in the applicable freelancer



or contractor agreement, which Page Two may ask existing freelancers to re-sign to acknowledge this clause.

This policy supplements and should be consulted in connection with Page Two's other policies and procedures. Page Two may make available detailed guidelines and department-specific policies to address specific AI issues. Those additional guidelines, policies, and procedures are extensions to this policy and must be complied with.

III. Policy

Care

Page Two considers the human and environmental impact of our use of AI, and humans oversee all our engagement with AI. We use AI as a tool to augment, not replace, human creativity, authorship, skill, judgment, and expertise. Critical decisions regarding content, editorial, marketing, and business strategy always involve human review, to verify accuracy and source integrity, and approval.

To ensure our use of AI is appropriate and ethically sound, on any approved AI system or platform we use, we do the following:

- We safeguard against promoting inaccurate information, misleading representations, harmful interactions (including but not limited to bullying and harassment), and discriminatory content.
- We educate ourselves and make informed decisions about the impact of AI use, including the environmental, ethical, and social costs of generative AI, and the capital- and energy-intensive infrastructure required to power it.
- Through our selection of which platforms we will use, we ensure our use promotes diversity, non-discrimination, fairness, and accessibility.
- We understand how user data and copyrighted material is stored, used, and shared, and protect against breach of privacy, confidentiality, and misuse of copyright.

Curiosity

We regularly review our AI policy and strategy. We stay informed about the latest developments in AI and adapt our policies and strategies as needed. We provide ongoing training and education for our employees on the responsible use of AI.

We consider new AI tools for approval and conduct due diligence. To have an AI platform considered for approval by Page Two, employees should inquire with their manager.



Creativity and Excellence

All creative work we publish—writing, design, graphics, illustrations, and so on—is created by humans, not AI-generated. All creative work requires the use of human skill and judgment for copyright protection. We do not use AI-generated material in our published work. We comply with copyright law regarding the use of the intellectual property of others, and respect confidentiality.

We use AI as a companion to our creative processes, not as a replacement for them. We are committed to maintaining the highest standards of accuracy and quality in our publications. AI-generated content used in support of our work is subject to the same rigorous editorial review process as human-created content. A human reviewer is accountable for reviewing and approving any AI-assisted content before it is used in public-facing materials.

We do not use AI-generated images in the interiors or on the covers of our books.

AI is not used to make, or substantially inform, acquisition or manuscript-evaluation decisions about submissions without explicit human governance. Because these decisions affect authors and carry fairness, accountability, and transparency risks, they remain human-led, and any AI assistance is disclosed and reviewed.

Page Two authors must fact-check any AI-assisted output used in their work. They are responsible for the accuracy of their content and to ensure their content does not infringe third party intellectual property or otherwise breach applicable laws or Page Two policies.

Integrity

We comply with all applicable data privacy regulations and best practices when collecting, storing, and using data with AI applications. We protect the privacy of our authors, employees, and readers.

Page Two works within a closed, company-owned AI environment. Our approved tools are Claude (Team plan) and Google Workspace (Gemini), both configured so that Page Two's inputs are not used to train the AI. Because Page Two's files are stored in Google Drive, which is overlaid with Gemini and reached by our approved tools, this approved environment is the sanctioned, secure place for Page Two information to be used. Page Two information, including confidential, proprietary, financial, legal, and personally identifiable information, may be used only within our approved, company-owned AI tools. Page Two information must never be entered into unapproved, personal, or consumer AI accounts outside that environment.



Page Two may use author manuscripts and related book materials within the approved, closed AI environment described above. Author permission for this use is obtained at the time of contracting: signing a Page Two publishing agreement constitutes the author's consent to Page Two's use of their manuscript within our approved AI tools. This applies to manuscripts at every stage, both unpublished and published, because all book files are stored in Google Drive. In all cases, we do not allow AI to train on material that we upload, and we do not enter author or third-party content into unapproved or consumer AI tools.

Our approved AI tools (Claude and Gemini) may be given standing access to the Page Two systems needed to do our work, including email, Slack, calendar, Google Drive, and AI meeting transcripts where relevant, and similar systems. Both tools are governed in the same way under this policy. AI meeting transcripts and recordings are stored in each employee's private "Claude" folder in Google Drive. That folder is private to the individual employee and to the Workspace administrator account.

We respect copyright and intellectual property rights in the use of AI. We ensure that AI-generated content does not infringe on existing copyrights and that we have the necessary rights to use any AI tools or datasets (with that being confirmed prior to any use).

We regularly audit and test approved AI systems for discrimination and biased outcomes, involving a diverse team for their development and review. We regularly review these systems for fairness and transparency, and compliance with applicable law, including but not limited to privacy, employment, human rights, and occupational health and safety law.

Partnership

We are transparent about the use of AI in our processes.

When submitting their manuscripts, and throughout the production process, Page Two clients must disclose to their publisher, editor, or project manager any use of AI in the creation of their manuscript, including which tools they used, what tasks they assigned to an AI platform, and when. This information is used to review and approve content, and to flag any potential copyright or other legal concerns.

We are also transparent with our authors about how Page Two itself uses AI. Page Two discloses its use of approved AI tools and obtains author consent to that use at the time of contracting, as described under Integrity above. We will update our author-facing agreements and guidelines to reflect this.



AI may be used as a tool to assist authors in their work, provided the author's (human) skill and judgment is still used to create the work. Where AI has been used in the development of a book, an AI disclosure statement appears on the copyright page.

Page Two clients must follow Page Two's writing guidelines for engaging with AI, provided by Page Two at the start of the project.

If AI is used to create marketing and promotional content, then a consumer-facing disclosure statement appears with that content.

IV. Employee Training and Guidelines

We regularly train all employees on this AI policy and its implications for their work, including in the development of strategies and tactics that employ AI. This training is mandatory before using any approved AI for work purposes.

We encourage employees to explore and experiment with approved AI tools as long as they follow the guidelines outlined in this policy and comply with all applicable laws in such use. However, it is important to remember that AI tools and technologies must be used only to supplement, not to replace or substitute, the proper performance of the employee's duties and responsibilities to Page Two.

To help everyone apply this policy in practice, AI uses fall into three tiers:

- **Allowed by default:** AI use based on public, non-sensitive material (for example, approved catalogue copy or already-published reviews), with human review. Explore these freely.
- **Grey zone:** AI use involving mixed or uncertain sensitivity, rights, consent, or public-facing impact. Treat these as medium-to-high risk until reviewed. When in doubt, ask your manager. Grey-zone uses are approved by the relevant department head.
- **Prohibited unless approved:** entering any Page Two material into unapproved, personal, or consumer AI tools; and any use of clearly confidential, contractual, financial, rights-related, or substantial unpublished creative material outside our approved AI environment.

Employees are responsible for ensuring the ethical and responsible use of AI in their work and appropriate human oversight and control, and for exercising accountability, good judgment, and critical thinking when using AI. Where AI has made a significant contribution to work intended for publication, employees should retain enough information to support human authorship and copyright, including the nature and extent of their human creative contribution.



Employees should report any concerns about the use of AI (their own or others') to their manager. AI cannot be used to (i) conduct or solicit illegal activities; (ii) interfere with job performance; (iii) generate content to impersonate, discriminate, or harass another person; or (iv) damage Page Two's business or reputation. Personal and non-work-related use of AI tools and technologies on Page Two's IT and communications systems (including Page Two-owned electronic devices) is also not permitted.

Page Two can monitor all content (including but not limited to any prompts or outputs) on any AI tool used on Page Two's IT and communications systems for compliance with this policy. Employees should have no expectation of privacy in this content.

Any employee who becomes aware of an actual or potential violation of this policy must promptly report the relevant details to their manager. Any violation of this policy may result in disciplinary action, up to and including termination of employment or engagement with Page Two.

V. Policy Review and Updates

This AI policy will be reviewed and updated at least annually or more frequently as needed to reflect advancements in AI technology, law, and changes in best practices. Page Two will promptly advise relevant parties when the policy is amended. Users of AI tools are expected to regularly consult this policy and any amendments to ensure their continued compliance.

VI. Contact

For any questions or concerns regarding this AI policy, please contact Chief Operating Officer Gabrielle Narsted.

VII. Definitions

Artificial Intelligence (AI): A broad category of technologies that enable computers to perform tasks that typically require human intelligence, including pattern recognition, language understanding, decision-making, and problem-solving. AI includes both generative and non-generative systems.

Generative Artificial Intelligence (Generative AI): AI systems that create new content—such as text, images, audio, code, or video—based on patterns learned from large datasets. These systems, including large language models, can produce outputs that appear original but may be influenced by the data on which they were trained. Claude, ChatGPT, DALL-E, and Gemini are examples of generative AI tools/systems.



Large Language Model (LLM): A type of AI system trained on vast quantities of text to generate and analyze human-like language. LLMs can summarize information, provide suggestions, assist with drafting, or answer questions, but they may also produce inaccurate or fabricated content (typically referred to as AI hallucinations). Claude and ChatGPT are examples of LLMs.

Predictive AI: AI systems that analyze existing data to forecast outcomes or identify patterns. Examples include tools that evaluate audience behaviors, assess reading levels, or flag inconsistencies in text. Predictive AI does not generate creative content. Microsoft's Azure Machine Learning is an example of a predictive AI service.

AI-Assisted Content: Content that is created by a human with the support of AI tools for tasks such as idea generation, outlining, grammar checking, summarizing, or research assistance. In AI-assisted content, a human remains the primary creator and is responsible for verifying the accuracy and originality of all material.

AI-Generated Content: Content created wholly or primarily by an AI system, including drafted passages, images, illustrations, graphics, or datasets. Because this material typically does not qualify for copyright protection, it cannot be published by Page Two as original creative work and must undergo significant human review and adaptation.

Significant AI Contribution: Any use of AI that influences the structure, meaning, argument, or creative expression of content. This includes generating draft text, outlining chapters, or creating images. All significant AI contributions must be disclosed to Page Two and are subject to editorial review. Page Two does not use AI-generated images in the interior or covers of its books.

Content Intended for Publication: Any material—textual, visual, or audio—that is expected to appear in a published book, article, marketing asset, or other publicly released content. AI-generated content cannot be used as content intended for publication unless it is significantly transformed by a human creator and approved by Page Two.

Training Data: The body of information—text, images, recordings, or other content—used to teach AI systems how to perform tasks. Training datasets may contain copyrighted, personal, or proprietary material. Page Two does not permit its approved AI tools to use Page Two or third-party content for training, and does not enter author or third-party content into unapproved or consumer AI tools. Author manuscripts may be used within Page Two's approved, no-training AI environment, with author consent obtained at contracting (see Section III, Integrity).

Sensitive Content: Material that includes private, personal, confidential, or culturally significant information; material subject to legal, regulatory, or ethical considerations; or content that may cause harm if misrepresented or incorrectly generated by AI. Sensitive content must be handled



only within Page Two's approved, company-owned AI tools, and must never be entered into unapproved, personal, or consumer AI accounts.

Copyrightable Authorship / Human Authorship: The legal principle that copyright protection applies only to original expression produced by a human's skill and judgement. AI-generated text or images must be substantially revised, curated, or transformed through meaningful human creative input.