

UNITED WISDOM

What can be delegated to artificial intelligence

Freedom and Action

Andrey Zarinsh & Maria Evilkut

A39 · EN

@andrey.zarinsh @ilkutki

What can be delegated to artificial intelligence

Epigraph

> *There are things which are within our power, and there are things which are beyond our power.*

Epictetus, Enchiridion, I, the complete opening sentence in Thomas W. Higginson’s translation, Project Gutenberg 45109 electronic edition. The Stoic distinction concerns one’s own judgement and external things; it does not describe contemporary AI’s technical capabilities.

What this article is about

Delegating to artificial intelligence begins before a useful answer appears. A person first chooses the operation, the information to hand over and the grounds for accepting its result. In a hypothetical story, a small ceramics studio prepares a visitor leaflet. It has verified information about opening times, booking classes and using the premises. AI can suggest a structure, headings or clearer wording for an approved paragraph. Attractive presentation nevertheless creates no facts about materials and permits no new promises to clients. This article examines an assignment’s boundaries: drafts, checkable operations, data access and action on someone’s behalf. Current official NIST materials provide limited risk context; Stoic texts ask a different question about one’s own judgement. We select no best provider, guarantee no model quality and establish no universal legal responsibility. Our practical aim is preserving intelligible grounds for choice where an answer’s convenience can easily be mistaken for its suitability.

I. Delegate an operation with an intelligible result

“Make a good leaflet” leaves almost everything uncertain: who needs it, what visitors should understand and which information may be used. The studio can first name a limited operation: prepare a draft page about class times, booking and access from approved notes. Its result can then be compared with the source information. This guarantees no correct execution, but makes an error more available for examination. The assignment concerns particular work rather than an undefined demand to replace the entire process of thinking. If organisers still disagree about visiting conditions, a polished answer does not remove the need to settle that disagreement first.

@andrey.zarinsh @ilkutki

Different operations need different signs of a suitable result. Sorting a list differs from explaining a rule; shortening a paragraph differs from finding grounds for a claim. When AI turns a table into text, each substantial value and condition needs checking. When it offers headings, someone assesses their suitability for the intended reader. When it proposes a plan, its points need examining: which serve the task, and which appeared merely because plans usually look that way? “Checked” alone conveys none of these differences. Identify examination’s subject before an answer’s ready availability makes acceptance look almost like a formality.

An assignment need not be as small as possible. A whole draft sometimes reveals a structure that does not emerge sentence by sentence. A large result nevertheless needs someone able to understand its components and examine significant transitions. A person who can assess only the beginning and presentation does not yet have sufficient grounds to accept everything. Options include limiting the task, obtaining appropriate expertise, postponing use or choosing another method. This is an editorial choice, not a universal rule against AI. The question is whether the result’s scope corresponds to available examination and the consequences of applying it.

II. Source information remains a separate subject

Rewording a leaflet does not require handing over everything accumulated in the studio. Approved opening times may be needed; participants’ private correspondence may not. Being technically able to insert a file does not establish whether doing so is permissible or necessary. Actual conditions of the selected service, available settings and applicable restrictions for particular data need establishing. This article claims no identical storage or data-use practices across providers. Such conditions change and need their own verification. Our more enduring distinction is that possessing material is not equivalent to having grounds to pass it to another system for a new use.

A draft page may need only prepared general facts without names or private conversations. Removing a name does not guarantee that someone is unrecognisable: an unusual story, date or combination of circumstances may still identify them. Data reduction therefore needs assessment through content rather than a single completed operation. An unnecessary passage can be excluded, a personal example replaced by a hypothetical one, or the service not used for that part at all. These choices promise no complete protection. They connect transmission with the task’s actual need and preserve questions that a simple “personal data removed” tick cannot honestly settle.

Roles within the studio also differ. An administrator may see bookings without being authorised to pass every message over for training an editorial assistant. A leaflet’s author can receive approved facts without receiving the whole archive. Missing information can prompt a limited

request for clarification instead of expanded access to everything. An AI assignment creates no new rights for its giver. If earlier access was tied to a purpose, the new operation needs separate consideration. Data remains part of the choice, rather than invisible technical preparation supposedly justified already by a result’s future usefulness.

III. A plausible answer differs from grounds for a claim

NIST’s Generative Artificial Intelligence Profile, 2.2 describes confidently presented false content, including fabricated reasoning and citations. This is official risk context, rather than proof that every answer from every model is false. The distinction becomes concrete when the leaflet draft adds a completion date absent from the studio’s notes. The phrase sounds natural and usefully answers a visitor’s question. Fluency creates no grounds for it, however. Check the date’s source, remove the addition or mark the condition as unsettled, rather than silently accepting good wording as a new fact.

A citation also needs separate reading. It can exist while failing to support a claim, concern another material or have a different scope. Asking the same model to explain why it is right produces another answer rather than independent confirmation. An explanation may help identify a checkable hypothesis, but it must meet the actual source and task. Technical studio conditions require real information from the responsible specialist and documentation of the particular process. This article supplies no firing or safety instructions. It preserves the difference between seeking a formulation and obtaining sufficient grounds to communicate a technical condition to a visitor.

Preserving uncertainty can itself belong to the result. If the studio does not yet know a date, an honest leaflet identifies where to ask instead of offering attractive unsupported approximation. AI can be asked to express that boundary more clearly, but not to create certainty that is absent. Two contradictory source documents similarly need human clarification rather than selection of the smoothest version. A model can propose a discrepancy list, which also needs checking against the documents. Its usefulness lies in making questions available for further work. It grants no permission for the system to decide which unsupported condition will now bind the studio.

IV. Checking is more than asking for agreement again

A convincing draft tempts a short question: “Is everything correct?” An equally convincing confirmation can follow without examining any original value. More substantive checking concerns particular potential losses: a changed date, a missing booking condition, an expanded accessibility promise or preserved exceptions. Ordinary comparison tools sometimes help; other cases require attentive reading or a competent person. The subject matters more than procedural solemnity. Three repeated assurances acquire no greater evidential force merely because they arrive in separate messages or appear as a list of successful checks.

Suppose the original note permits changing a class date only by agreement. The draft says visitors may choose another date. Familiar words conceal a changed right. The checker needs to notice the relationship between condition and action, rather than matching names alone. If the model restores the qualification after an observation, the changed text needs rereading at that point. A small correction with understood consequences need not restart the entire project. Retaining an old confirmation after a substantial change can nevertheless conceal that another version was examined. Grounds for acceptance concern the actual version people intend to use.

Independence of examination takes different forms with different limits. Another person can err; automatic comparison does not understand every qualification; another model is not necessarily independent in how it fails. The checker’s name cannot replace a question about what they examined and on what grounds. An ordinary leaflet may need comparison with approved facts and a complete reading before release. More consequential use needs other expertise and conditions; this article does not prescribe them for its reader. Examination’s scope is chosen in relation to the result and consequences, with honesty about what remains unexamined and cannot count as confirmed.

V. Preparing a proposal differs from taking action

AI can draft a reply to a visitor, but prepared text is not yet a sent message. When a system connects to email or another working environment, permission to send and its conditions need separate definition. A draft reply about an available place must not become a promise before actual availability is checked. This distinction concerns more than wording: calendar entries, file changes, material orders or deleting an earlier version can affect others. The ability to perform an action is a separate part of an assignment. It does not follow from the system having helped write a suitable proposal already.

A limited task may permit preparing alternatives while leaving application to a person. Some repeated operations may suit a predefined automation with clear scope and stopping

conditions. A new circumstance must nevertheless not silently receive permission under an old general name. A changed class outside agreed dates needs another choice. Conflicting data should return a question rather than conceal it through selection. This describes a working arrangement without promising that any programme can respect such boundaries. The selected system’s actual behaviour needs checking before an assignment’s wording can be considered an operative constraint.

Being able to reverse an action affects permissible scope without making consequences nonexistent. A mistaken email can be clarified after its recipient has changed plans. A deleted file may be recoverable, or intermediate information may disappear with it. Reversibility should therefore be examined under actual conditions rather than treating every electronic action as harmless. A necessary earlier version can be preserved, a particular result shown before application, and access limited to the task. These measures remove no risk entirely. They make the moment of choice visible and allow authorised action to be compared with what will actually change.

VI. Keeping one’s judgement does not require doing everything by hand

In Enchiridion, I, Epictetus distinguishes one’s own judgement and external things within Stoic doctrine. Contemporary tool choices cannot be mechanically derived from it: access, data and technical conditions concern another task. A limited philosophical connection asks which choice a person recognises as their own. Using supplied wording does not necessarily abandon thought. Someone can understand it, examine its grounds and consciously accept it. Conversely, entirely personally written text can contain an unchecked claim. How words were typed establishes neither judgement’s quality nor the honesty of presenting it to someone else.

The author can ask AI for alternatives explaining class bookings. They choose one because it preserves conditions and addresses readers intelligibly, rather than because the system called it best. That participation differs from demanding that every sentence be newly composed to earn an author’s name. Legal authorship depends on applicable rules and is not settled here. Our subject is the work’s actual arrangement and the possibility of explaining a choice. A person who does not understand a phrase cannot yet honestly present it as an explanation they have examined, even when it sounds more professional than their usual language.

Keeping judgement does not deny reliance on another’s expertise. The leaflet’s author uses a maker’s information and may be unable to verify the entire technical process independently. They can accurately identify what a competent participant established, what was rewritten and what was checked editorially. That map of grounds is more useful than declaring complete independence. It also does not permit assigning every question to a model labelled a universal

expert. This article does not rank people’s and products’ competence on one scale. A particular claim needs available grounds and suitable examination; personal choice needs understanding of present reliance and where one’s own consideration ends.

VII. Distinguish the origins of others’ ideas and generated options

AI can help express a chosen thought or suggest an idea its user did not previously have. These are different cases despite a uniform-looking final paragraph. Information about substantial source materials and methods should remain where it affects the result’s presentation. A quotation from someone else’s text does not become one’s own merely after a system paraphrases it. Nor should every generated formulation be treated as an independently verified idea without provenance. What to reveal and how depends on the work, promises to readers and actual rules. This article supplies no single mandatory disclosure wording for all cases.

Suppose the model offers an attractive story about a first class in the studio. If that event never occurred, the story may be an explicitly hypothetical example, but not a real visitor’s testimonial. A person selects its status and is responsible for accurate presentation in an editorial sense. Generated text creates no witness, portrayed person’s consent or factual event. If a real account lies behind it, use conditions and preservation of substantial details need checking. Changing a name does not automatically settle recognisability, meaning or permission. These boundaries concern the material, rather than its attractive final style alone.

In Letter 33.7–11, Seneca distinguishes repeating others’ words from understanding while retaining predecessors as guides. This argument predicts no AI and forbids no assistance. Its limited connection asks what a person actually understood and chose from an offered version. The studio can retain a fact’s source, mark a hypothetical example and explain its choice without publishing every intermediate request. Such accuracy guarantees neither originality nor legal acceptability of any material. It avoids replacing the work’s history with a convenient legend of wholly independent origins for something that actually came from several places.

VIII. Saving writing time does not necessarily save work

A draft appears quickly, but factual comparison, selection, correction and application remain. When these exceed available time, generation’s speed may produce more unchecked material than the studio can use. Delegation’s value concerns the whole operation rather than the instant text appears. This article promises no particular reduction in hours or costs. Sometimes a supplied version reveals structure; sometimes correcting it demands more attention than writing a short explanation independently. That distinction is established through a particular task, rather than general faith in automation or general dislike of it.

A limited beginning is possible: one booking paragraph and an intelligible way to compare it with approved conditions. Afterwards, examine what accepting the version actually required. Besides attractiveness and the number of corrected words, consider hidden changes of meaning, missed questions and additional access needed. One successful case establishes no suitability for every text, and one failure proves no universal uselessness. Changed tasks make earlier experience a limited ground. The studio can expand use as actual understanding of conditions develops, without converting early success into final evidence of safety for every future assignment.

Delegation can also transfer work to someone else. An author receives polished text and gives a maker lengthy checking of technical additions absent before. Without agreement, time saved by one becomes another’s invisible burden. Generation can be limited to approved information, additions marked and examination’s scope discussed beforehand. Even this requires actual comparison: asking the system not to invent does not prove additions are absent. A tool’s benefit concerns the shared process. It matters who gained a new obligation and whether that person could accept or change it.

IX. A failure should clarify the assignment’s boundary

If the model added an incorrect date, correcting one number and forgetting the error’s origin is insufficient. Examine available information: was the date in source materials, did documents contradict one another, or did the assignment expand into unsupported advice? The cause may remain undetermined, which needs stating honestly. One case cannot diagnose an entire system. It can nevertheless reveal a particular boundary: this leaflet cannot accept unsupported dates from a freely drafted version. Correcting text and changing conditions for the next assignment are separate results, each needing its own examination before being called completed.

Failure does not require endlessly repeating requests until a pleasing option appears. That part’s use can stop while needed facts are obtained from a responsible person and a narrower task prepared. If the problem concerns access or action, what already happened matters

alongside the latest answer's content. A mistaken message needs clarification through an appropriate process; a correct new draft does not undo its earlier sending. This offers no universal incident protocol. It distinguishes restoring an actual situation from obtaining another text conveniently called a correction while consequences remain outside it.

Conditions for declining use can be named before error: absent grounds for a substantial fact, unavailable examination, unsuitable data transmission conditions or action outside permission. These concern particular work rather than a person's caution or courage. Sometimes refusal affects only a passage, leaving other material usable after suitable checking. Sometimes the task needs an entirely different arrangement. Reconsideration preserves accurate history without erasing successful parts or concealing failures. A clear application boundary is more valuable than a ritual promise to be more attentive when it remains unclear which action and ground will change next time.

X. Delegation keeps a person at the moment of acceptance

A final leaflet reaches visitors as the studio's material. They may not see which tool rearranged its paragraphs but expect intelligible conditions and honest information. Acceptance cannot dissolve into a long chain where everyone supposedly merely passed along another version. Identify who checked applicable facts, who selected the edition and who authorised its use. This describes responsibility for one's actions within work rather than universally allocating legal blame. Particular fields have specific rules and expertise; the article replaces none and derives none from AI use alone.

In AI RMF 1.0, Appendix C, NIST considers differentiated human roles and varying AI interaction configurations. Human presence is no automatic correctness guarantee, and oversight differs between operations. Accepting our leaflet includes a complete thoughtful reading of its final text and comparison of substantial conditions. If only opening times were checked, the whole material cannot be called technically verified. Saving text does not mean sending or publishing it. Each readiness claim should concern an actual completed action and the result now available to others.

After this choice, the tool remains an assistant for particular operations, and a person remains a participant in particular decisions. This requires no heroic independence and promises no absence of mistakes. It permits an account: these facts are approved, this version chosen, these boundaries checked, and this question remains open. Much can be delegated when grounds for use are intelligible and suitable examination is available. A feeling that the system sounds cleverer and more confident cannot supply those grounds. The hypothetical studio's good leaflet emerges when text, data, permitted action and actual acceptance connect into intelligible work explainable to a visitor and open to revision, rather than with the first fluent answer.

What this means in practice

Choose one operation to assign to AI and identify its result: a draft, comparison, transformation or alternatives. Separately record necessary information and grounds for passing it to the selected system. Check actual service conditions where they matter; assumptions about all products cannot replace them. The hypothetical studio can begin with a booking paragraph using approved general conditions rather than private correspondence. Identify additions, lost qualifications and unsupported facts in advance. This gives subsequent examination a subject without guaranteeing that the system will observe every request constraint.

Compare the result with grounds existing outside its confident tone. Check substantial values and qualifications, read sources where needed and preserve uncertainty when information is missing. Repeated model assurance is not independent confirmation. Separate prepared text from permission to apply it: sending a message, changing a record or presenting a leaflet. Automated action requires checking the actual device’s scope and stopping conditions. Examination depends on the task and consequences. Honestly identify which parts were considered, by whom and on what grounds, and which remain outside completed work.

After correction, read the actual final version to the extent appropriate to the change. Preserve substantial provenance and do not present a hypothetical scene as a real person’s evidence. On failure, distinguish corrected text, consequences of completed action and conditions for another assignment. Unavailable examination may call for limited use or refusal. This practice establishes neither universal prohibition nor universal permission for AI. It preserves four subjects: operation, data, grounds for the result and permission to act. Assistance’s convenience can then enter intelligible work, while human acceptance becomes more than pressing a button after a confident answer supplied by something else.

Key quotation

> *Truth lies open for all;*

Seneca, Moral Letters to Lucilius, 33.11, the first clause before the sentence’s continuing qualification, Richard M. Gummere translation, Wikisource. The selected context retains predecessors as guides and calls for independent inquiry; this is not a historical statement about AI.

Connections between traditions

Question	Bearing	Editorial connection	Boundary
What I choose myself	Epictetus, I	Acceptance has grounds	Stoic control is not technical access
Understanding and repetition	Seneca, 33.7–11	Supplied words require thoughtful selection	Others' assistance is not rejected
Human roles	NIST AI RMF, Appendix C	Checking differs from application	A voluntary framework is no unified legal rule
Confident false answers	NIST AI 600-1, 2.2	Tone does not establish a fact	Not all answers are declared false
A shared boundary question	Stoic texts and a contemporary framework	Comparison across differing subjects	A technical document is not a spiritual tradition

The map compares Stoic philosophy with contemporary technical context while preserving different tasks. It attributes neither AI prediction to ancient authors nor doctrinal unity with Stoicism to NIST.

Primary sources

1. Epictetus. The Enchiridion, I. Thomas W. Higginson translation, Project Gutenberg 45109. Complete I and the electronic edition's note were read, displaying Liberal Arts Press, copyright 1948 and the 1955 second edition. The catalogue reports public domain status in the USA, release in 2014 and update on 24 October 2024; non-US rights are independently unverified. Albert Salomon's introduction was not used as evidence for Epictetus's words.

2. Seneca. Moral Letters to Lucilius, Letter 33. Richard M. Gummere translation, Wikisource; continuous 33.7–11 was read. Collection title metadata displays Heinemann/Putnam and MCMXXV; the letter's particular print copy was not separately collated. The site's licence notice is not evidence of worldwide rights for every edition.

3. NIST. Artificial Intelligence Risk Management Framework, 1.0. NIST AI 100-1, January 2023; Appendix C, printed pages 40–41, was read. The current official page reports revision of 1.0 at the time checked. The framework is voluntary; this is no assertion of a completed new edition.

4. NIST. Generative Artificial Intelligence Profile. NIST AI 600-1, 2024; the continuous main section 2.2, printed page 6, and selected actions MP-2.3-001/MS-2.5-003 about checking

output and sources were read. No direct NIST quotations appear in this article. The documents establish neither a particular provider’s quality nor the proposed practice’s efficacy.

Tags

artificial intelligence, delegation, examination, human judgement, draft, source information, data, access, permission, action, responsibility, authorship, provenance, sources, qualifications, uncertainty, application, version, facts, plausibility, automation, limitations, editorial work, studio, Epictetus, Seneca, Stoicism, NIST, AI RMF, shared work

About the authors

Andrey Zarinsh is the author and founder of United Wisdom. He explores inner transformation, the unity of traditions, intercultural dialogue, and contemporary spiritual discipline. Instagram: @andrey.zarinsh.

Maria Evilkut is a co-author of United Wisdom. Alongside her research, she seeks to test these ideas in practice and observe how principles of unity and mutual understanding work in real life between people and cultures. Instagram: @ilkutki.