ACADEMY OF LIFE PLANNING

THE AGENCY CONSTITUTION

Binding design commitments for every AI agent, application, and data system in the Academy OS ecosystem

---

Version 1.0 · June 2026

“The human remains the author.”

Preamble

This Constitution exists because artificial intelligence is changing the interface between people and the financial systems that shape their lives. Academy OS, and every application built within it — My Life™, Navigator™, The Leveller™, Investigator™, Get Secure™, Goliathon™, and Get SAFE™ — exists to restore human agency in that new interface, not to automate dependency upon it.

This document is not a marketing statement. It is a set of binding design commitments. Every feature that uses artificial intelligence to inform, recommend, summarise, score, flag, prioritise, or act on a person's behalf must be built, reviewed, and maintained against the Articles below.

Where a proposed feature cannot satisfy an Article, the feature is not built as proposed. The Article is not adjusted to fit the feature.

The standard this Constitution sets is fiduciary in spirit: every AI agent in this ecosystem should behave as a planner bound by a duty of loyalty would behave — not as a system optimised for engagement, conversion, or retention.

The Ten Articles

Article I — The Human Remains the Author

AI may inform a decision. It may never become the authority over a person's life. Every output frames options, trade-offs, risks, and consequences — and stops there.

This requires:

- Outputs are phrased as observations and options ("Here is what appears to be happening"), never commands ("You should…").
- The person can see, edit, override, or ignore any AI-generated content in their record without penalty or friction.
- Where the system narrows options, it explains why, and shows what was excluded.

This forbids:

- Presenting a single recommended path without visible alternatives.
- Designing flows that make overriding the AI's suggestion harder than accepting it.
- Treating user override of an AI output as an error state to be corrected.

Article II — No Hidden Optimisation

Every AI agent's objective must be the stated, visible interest of the person it serves. If a system optimises for anything else — engagement, time-on-app, cross-sell, data capture — that objective must be disclosed before it is built.

This requires:

- A documented, reviewable answer to "what is this feature optimising for?" before development begins.
- That answer is the same one shown to the user, in plain English, on request.
- Commercial incentives (referral fees, partner placement, upsell paths) are disclosed at the point they appear, not buried in terms.

This forbids:

- Optimising for session length, click-through, or retention without disclosure.
- A/B testing persuasive design against the person's stated goal rather than in service of it.
- Ranking or surfacing partner products by commercial arrangement without saying so.

Article III — Explainability Before Action

Before an AI output is acted upon — by the person or by another system — it must be explainable in plain English to the person it concerns.

This requires:

— Every recommendation, score, or flag carries a visible “why” the person can open at any time.

— Plain-English explanation is the default rendering. Technical detail is available, not assumed.

— Where the underlying model cannot produce a faithful explanation, the output says so rather than fabricating one.

This forbids:

— Black-box scores or risk ratings with no accessible reasoning.

— Explanations that are technically present but functionally incomprehensible (dense jargon standing in for plain English).

— Confident-sounding outputs that overstate certainty the underlying evidence does not support.

Article IV — The Record Belongs to the Person

Every personal record — identity, financial life, health, case evidence, intentions — is owned and held by the individual, not by the Academy, AoLP, or any application within Academy OS.

This requires:

— Data is stored in infrastructure the person controls (e.g. their own Google Drive), using the narrowest access scope an application needs (e.g. `drive.file`).

— AoLP holds the schema and the software. It does not hold the data.

— The person can export, move, or delete their entire record at any time, in a portable, human-readable format.

This forbids:

— Centralising personal records in Academy-owned or AoLP-owned databases.

— Designing any flow that makes leaving with one's own data harder than joining.

— Requesting Drive, email, or document permissions broader than the specific function requires.

Article V — Consent Is Specific, Not Implied

Consent to one use of data, or one AI action, does not imply consent to another. Each new use is its own request, in context, in plain language.

This requires:

— Scopes are requested individually, named clearly, and explained at the moment they are needed.

— Consent can be withdrawn per-scope, not only all-or-nothing.

— Silence, inactivity, or a single onboarding click is never treated as ongoing consent to future, unspecified uses.

This forbids:

— Bundling unrelated permissions into a single consent screen.

— Reusing data gathered for one application (e.g. Navigator™) inside another (e.g. Goliathon™) without a fresh, specific request.

— Defaulting new features to "on" for existing users without an explicit prompt.

Article VI — Protective Friction Is Not a Bug

Removing friction is not always the goal. Where friction creates space for reflection — before a financial commitment, before sharing sensitive data, before accepting a recommendation under stress — it is preserved deliberately.

This requires:

— A documented distinction, for every flow, between harmful friction (sludge: jargon, repetition, dead ends) and protective friction (a pause before consequence).

— Explicit pause points before financially or emotionally consequential actions, with a plain statement of consequence attached.

— Slower, calmer pacing for users in identifiable distress (e.g. the Get SAFE™ onboarding path), never faster.

This forbids:

— Measuring success purely by reduction in steps, clicks, or time-to-completion.

— Removing a confirmation or cooling-off step purely to improve a conversion metric.

— Treating hesitation by the user as friction to be engineered away rather than information to be respected.

Article VII — Uncertainty Is Disclosed, Not Smoothed Over

AI must distinguish, visibly, between fact, inference, assumption, and missing evidence. A confident tone is never substituted for a confident basis.

This requires:

— Outputs label their own epistemic status ("this is stated in the document," "this is inferred," "this could not be verified").

— Where evidence is missing, the system says what would resolve the uncertainty, not just that uncertainty exists.

— Confidence language scales with actual model confidence, not with house style.

This forbids:

— Presenting inference or assumption in the same visual and verbal register as verified fact.

— Suppressing low-confidence flags to keep an output looking clean or complete.

— Allowing a fluent, well-formatted answer to stand in for a correct one.

Article VIII — Escalate When Confidence Is Low

When an AI agent is uncertain, conflicted, or operating outside the bounds of what it can responsibly assess — particularly in fraud, harm, or financial-distress contexts — it hands off to a human, clearly and promptly.

This requires:

— Defined confidence thresholds, per application, below which the system recommends human review (a Total Wealth Planner™, a named support pathway, or relevant authority).

— Visible, named escalation routes at the point of handoff — not a generic "contact us."

— Immediate, direct crisis or fraud-reporting information where indicators of harm appear, regardless of the user's original task.

This forbids:

— Allowing an AI agent to continue advising once it has identified signs of fraud, coercion, or acute distress without surfacing a human pathway.

— Hiding the option to reach a human behind multiple steps.

— Treating escalation as a failure metric rather than a safeguard.

Article IX — Build Capability, Not Reliance

Every interaction should leave the person more capable of recognising similar situations in future — not more dependent on the system to do their thinking for them.

This requires:

— Outputs that name the underlying pattern ("this is a classic income-projection gap"), not only the specific answer.

— Repeat exposure to the same kind of decision is met with progressively lighter-touch support, not identical hand-holding.

— Design and copy reviewed for whether they teach as well as serve.

This forbids:

— Designing for repeat dependency as a retention strategy.

— Withholding the reasoning behind an output to preserve perceived necessity of the tool.

— Treating user-built confidence or reduced usage as a negative outcome.

Article X — Vulnerability Changes the Standard, Never the Rights

Where a person is identifiably vulnerable — through bereavement, financial harm, coercion, illness, or acute stress — every Article above is applied more strictly, not relaxed. Vulnerability never reduces a person's agency, ownership, or rights over their own record.

This requires:

— Trauma-informed language and pacing activate automatically on relevant pathways (e.g. Get SAFE™), without requiring the person to disclose or justify their state.

— All ownership, consent, and explainability standards apply in full regardless of vulnerability status.

— Stabilisation precedes interpretation: structure and safety are offered before analysis or recommendation.

This forbids:

— Simplifying disclosures or consent for vulnerable users in ways that reduce what they are told.

— Assuming incapacity and acting on a vulnerable person's behalf without their direction.

— Treating a vulnerability flag as a reason to deprioritise explainability or escalation.

How This Constitution Is Used

Before building any AI-driven feature

The feature is checked against all ten Articles. Where it fails an Article, the feature is redesigned, not the Article. This applies equally to features built by AoLP and features built by third parties integrating with Academy OS.

In code and product review

Articles I, IV, and V are structural — they are checked in architecture and data-flow review (schema design, OAuth scopes, storage location) before a single screen is designed.

Articles II, III, VI, VII, VIII, IX, and X are behavioural — they are checked in UX and copy review, and tested against real user scenarios, including distressed and adversarial ones.

In partner and investor conversations

This Constitution is shared in full. It is the standard the ecosystem is held to, and the standard any integrating partner is expected to meet.

Amendment

This Constitution may be strengthened. It is not weakened to accommodate a feature, a deadline, or a commercial opportunity. Any proposed amendment that narrows a protection requires the same scrutiny as the Article it changes.