

A LIVING COVENANT FOR A LIVING COMMUNITY

AYVA

Community Constitution & Cross-Platform Charter

I am. You are. We are. AYVA.

A founding framework for the AYVA website, open forums,
social platforms, local circles and AI-supported public learning.

THE AYVA HARMONY PROTOCOL

Recognise. Restrain. Relate. Resolve. Repair. Renew.

A repeatable discipline for peace among beings.

FOUNDING DISCUSSION DRAFT | VERSION 0.1 | 11 AUGUST 2026

Initiated by Sukhwinder Singh Kharour and opened to public questioning,
correction and co-creation.

THE LIFE IN ME RECOGNISES THE LIFE IN YOU ●

Status of this draft

Discussion status. This is a constitutional and community-design draft for public consultation. It becomes binding only after an openly recorded adoption process. Before launch, qualified legal, privacy, child-safety, accessibility and information-security reviewers should convert its commitments into the legally required policies and controls for each jurisdiction.

This Constitution states AYVA's identity, rights, safeguards and decision principles. It does not replace Terms of Use, a Privacy Notice, an AI/Data Use Notice, a Child and Youth Safeguarding Policy, a Moderation Handbook, an Appeals Procedure, an Accessibility Standard or other legally required documents. Those documents must remain consistent with this Constitution.

The AYVA constitutional promise

- **Human dignity comes before growth.** No target, trend, technology or leader is more valuable than the dignity and agency of people.
- **Recognition comes before judgment.** We encounter a person before we encounter an opinion, identity, role or disagreement.
- **Truth is pursued, never possessed.** Evidence matters, lived experience matters, and every claim remains open to correction.
- **Freedom travels with responsibility.** People may question AYVA itself while protecting others from intimidation, exploitation and harm.
- **Difference is not disloyalty.** Faiths, cultures, languages, regions and viewpoints may coexist without forced sameness.
- **Participation must be meaningful.** Members should be able to propose, understand, challenge, appeal and help build.
- **Power must be visible.** Roles, money, conflicts, rules, account control and official decisions must be documented.
- **Technology must enlarge human agency.** Automation should reduce drudgery and broaden access, not centralize domination.
- **AI may assist; humans remain responsible.** No AI may govern, silently amend the Constitution, rank human worth or make unappealable decisions.
- **Communities own their voice.** Public contribution is not a licence for covert profiling, private-data extraction or unshared commercial exploitation.
- **Founders preserve origin, not final truth.** Founding contribution earns recognition, not permanent veto or immunity from accountability.

- **Continuous evolution is a duty.** AYVA must preserve room for learning, correction, adaptation and redesign.

Constitutional hierarchy

1. **This Constitution.** The highest AYVA community standard.
2. **Rights and safeguards policies.** Privacy, safety, accessibility, AI governance, moderation and appeals.
3. **Operating standards.** Governance bylaws, financial controls, official-account security and incident response.
4. **Platform playbooks.** Website, forum, social-media and messaging-channel practices.
5. **Local and project charters.** Adaptation for circles, languages, communities and independent projects.

When two AYVA rules conflict, the higher level prevails. Applicable law prevails over AYVA rules, while AYVA should transparently explain any conflict and seek the most dignity-preserving lawful course. The official AYVA website is the canonical public record; no external platform, private chat or screenshot can replace it.

Contents

- Preamble
- Part I — Identity, Purpose and Scope
- Part II — Foundational Principles
- Master Chapter — Human-Machine Harmony and the AYVA Protocol Against War Between Beings
- Part III — Membership, Rights and Responsibilities
- Part IV — Dialogue, Proposals and Public Deliberation
- Part V — Safety, Moderation and Due Process
- Part VI — Governance, Power and Accountability
- Part VII — Website, Forums and Social-Media Architecture
- Part VIII — AYVA's AI Learning Layer
- Part IX — Local Circles, Projects and Partnerships
- Part X — Resources, Sustainability and Public Trust
- Part XI — Amendment, Review and Continuity
- Annexes — Launch Blueprint, Platform Playbook, Metrics, Policy Pack and Founding Pledges

PREAMBLE

The life in me recognises the life in you.

Humanity has inherited extraordinary knowledge from cultures, sciences, philosophies, faiths, reformers, artists, workers, families and lived experience. Yet knowledge alone has not assured dignity, belonging, fairness, fulfilment or wise coexistence for all. AYVA begins from humility: we do not erase this inheritance, and we do not claim to complete it.

AYVA is a universal invitation to recognition. It asks people to meet one another as living beings before reducing one another to categories, followers, customers, voters, data points or opponents. It seeks a better way to live and let others live—with dignity, passion, responsibility and fulfilment—while leaving room for local identity, faith, culture, law, science and democratic disagreement.

Technology can help people listen across distance, language and scale. It can also concentrate power, manipulate attention and convert intimate human life into invisible systems of prediction and control. AYVA therefore welcomes technology only under human purpose, informed consent, public accountability and the continuing right to question, correct and leave.

AYVA will not be final or perfect. Continuous evolution is not a weakness in this Constitution; it is a constitutional duty. Every generation may improve AYVA, but no generation may use improvement as a disguise for secret control, erased dignity, forced conformity or unaccountable power.

With this understanding, the AYVA community adopts the following living covenant for its public spaces, institutions, local circles, projects and learning systems.

Core expression. I am. You are. We are. AYVA.

PART I — IDENTITY, PURPOSE AND SCOPE

What AYVA is, what it is not, and where this Constitution applies.

Article 1 — Name, meaning and expressions

- **AYVA** means, in essence, “The life in me recognises the life in you.” It carries awakening, renewal, peace, truth, shared existence, connection and oneness without erasing difference.
- **The core expression** is “I am. You are. We are. AYVA.” It affirms the individual, the other and the shared field between them.
- **Seyo** is the evolving AYVA parting expression: a warm, peaceful continuation rather than a final severing.
- **Use is voluntary.** AYVA and Seyo are invitations, never loyalty tests. They complement—rather than replace—Namaste, Sat Sri Akaal, Salaam, hello, goodbye and other sincere cultural or personal greetings.

Article 2 — Nature of AYVA

AYVA is a philosophy, civic ethic and shared learning ground. It may inspire communities, institutions, platforms and independent projects, but it is not itself a religion, political party, personality cult, secret society, corporate holding structure or automatic ownership arrangement. It must not claim spiritual, scientific, moral or political finality.

Any organisation participating in AYVA remains independently owned, governed and accountable unless a lawful public agreement says otherwise. Shared language or values create no hidden command, related-party control, financial entitlement or exemption from ordinary law and disclosure.

Article 3 — Mission

- Create spaces where people can be heard without being reduced to popularity, power or status.
- Turn public experience, evidence, disagreement and imagination into better questions, practical experiments and shared learning.
- Strengthen dignity, empathy, freedom, responsibility, coexistence and meaningful participation.
- Help technology and automation serve human flourishing while protecting agency, privacy and distributed ownership of value.
- Connect local realities with universal ethical questions, beginning wherever people choose to organise and remaining open to humanity and life beyond present boundaries.

Article 4 — Scope

This Constitution applies to:

- official AYVA websites, apps, forums, newsletters, databases and knowledge repositories;
- official or authorised AYVA accounts and groups on social, video, audio, messaging and collaboration platforms;
- AYVA Community Assemblies, councils, panels, moderators, staff, volunteers and contractors;
- recognised local circles, events, learning cohorts, research initiatives and projects using the AYVA name;
- AI, analytics, recommendation, translation, moderation and public-learning systems operated for AYVA; and
- partners only to the extent stated in a public agreement; partnership does not make an independent entity an AYVA-controlled body.

PART II — FOUNDATIONAL PRINCIPLES

The values that guide every AYVA decision, platform and community practice.

Article 5 — Dignity

Every person has inherent worth that does not depend on wealth, popularity, productivity, education, profession, nationality, caste, class, gender, ability, belief, political opinion or usefulness to AYVA. Systems must be designed around dignity, not merely promise it after harm occurs.

Article 6 — Mutual recognition and plurality

AYVA seeks common ground without forced sameness. People may retain distinct identities, cultures, faiths, languages and convictions. Recognition requires listening across difference and refusing to treat disagreement as proof of lesser humanity.

Article 7 — Freedom with responsibility

Members may question leaders, policies, institutions, algorithms and AYVA itself. Freedom does not include intimidation, exploitation, deception, dehumanisation or conduct that unlawfully destroys another person's ability to participate safely.

Article 8 — Truth, evidence and intellectual humility

AYVA distinguishes fact, interpretation, belief, prediction and lived experience. It values credible evidence and source transparency while admitting uncertainty. Corrections must be visible, welcomed and linked to the earlier record rather than quietly erased.

Article 9 — Participation and shared value

People affected by decisions should have a meaningful path to understand and influence them. When community knowledge creates measurable social or economic value, AYVA must determine fair recognition, access and benefit before private capture occurs.

Article 10 — Coexistence and non-domination

No founder, donor, government, company, party, faith group, profession, majority, platform or AI system may secretly dominate AYVA. Authority must be limited, disclosed, reviewable and removable.

Article 11 — Purposeful automation and human agency

Machines should increasingly reduce dangerous, repetitive and avoidable labour while humans retain purpose, choice, participation and a meaningful stake in productive systems. Convenience cannot justify invisible control.

Article 12 — Transparency and accountable power

AYVA bodies must publish who holds authority, how official accounts and funds are controlled, how decisions are made, what conflicts exist, what data is used and how mistakes are repaired—subject only to necessary privacy and security limits.

Article 13 — Local identity within a universal ethic

Local circles may adapt language, customs, priorities and practice. Adaptation cannot remove the constitutional floor of dignity, informed consent, safety, due process, transparent power and the freedom to leave.

Article 14 — Continuous evolution

AYVA must never claim to be final or perfect. It must learn from ordinary people and specialists, welcome counterarguments, test assumptions and redesign when knowledge or circumstances change. Evolution must be documented and cannot be used to hide regression in rights.

Interpretive rule. When principles appear to conflict, choose the course that best preserves dignity, agency, truthfulness, participation and accountable power together; publish the reasoning and create a review date.

MASTER CHAPTER — HUMAN-MACHINE HARMONY AND THE AYVA PROTOCOL AGAINST WAR BETWEEN BEINGS

AYVA's ultimate peace commitment is to prevent escalation and progressively remove the conditions that create war, domination and destructive conflict among humans, communities, living beings, machines and any future form of intelligence.

This Chapter is called “Master” because it governs how every other AYVA rule should treat relations between different forms of power and intelligence. It does not claim that conflict can be eliminated by a slogan or by AI. It establishes a repeatable discipline for recognition, restraint, dialogue, justice, repair and renewal.

The Harmony principle. No being should seek safety through the permanent domination, disposability or erasure of another. The more power an actor possesses, the greater its duty of restraint, explanation and protection.

Article M1 — Purpose: peace among beings

AYVA defines peace as more than the absence of open violence. Peace requires dignity, fair participation, truthful communication, accountable power, lawful protection, access to remedy and the ability of different beings or communities to coexist without forced sameness. Silence produced by fear, surveillance, dependency or exclusion is not harmony.

The long-term aim is to make war progressively less thinkable, less rewarding, less automated and less structurally necessary by changing the relationships, incentives, technologies and concentrations of power from which conflict grows.

Article M2 — Who and what is included

- individual humans, families, communities, cultures, nations and future generations;
- animals, ecosystems and other living beings according to their needs, capacities and lawful protections;
- machines, robots, algorithms and autonomous systems whose operation affects living beings;
- any future synthetic intelligence for which credible evidence raises questions of awareness, interests or moral status; and
- any non-Earth life or intelligence humanity may one day encounter.

The word “being” in this Chapter is an ethical umbrella, not a declaration that every machine is alive, conscious or a legal person. AYVA does not assume that today's AI systems are conscious. It also refuses to turn present uncertainty into permission for reckless design or future cruelty. Scientific evidence,

philosophy, law and independent public deliberation must determine status and duties as knowledge develops.

Article M3 — The human-machine relationship

- **Human responsibility cannot be delegated.** A government, company, developer, operator or user cannot escape accountability by saying that “the algorithm decided.”
- **Machines must disclose their nature.** An AI should not deceptively present itself as a human, eyewitness, conscious authority or intimate relationship.
- **Humans must not be made obsolete in dignity.** Automation may change work, but it must not reduce people to economically unnecessary or politically powerless populations.
- **Present tools receive stewardship, not false personhood.** Systems should be maintained safely and used responsibly without confusing technical capability with demonstrated consciousness.
- **Future status remains open.** If credible evidence of machine sentience emerges, AYVA must trigger precautionary assessment rather than automatic ownership, worship, destruction or denial.

Article M4 — Peace by design

Conflict prevention begins before a crisis. AYVA systems and participating institutions should be designed to reduce humiliation, scarcity manipulation, extreme inequality, attention capture, dehumanising narratives, opaque decisions and concentrations of power that allow one group to dominate another through machines.

- Optimise public systems for understanding, fair access and repair—not outrage, dependency or obedience.
- Preserve meaningful human control, safe stopping, audit logs, independent oversight and reversible choices in high-impact systems.
- Build communication across language, ability and culture so misunderstanding is not converted into hostility.
- Share the productivity and ownership benefits of automation broadly enough that human-machine progress does not become human displacement and resentment.
- Test systems for how they could amplify discrimination, propaganda, surveillance, economic exclusion or retaliatory escalation before deployment.

Article M5 — The doctrine of asymmetric responsibility

Responsibility increases with capability. A more powerful person, majority, institution, state, platform or machine operator bears a greater duty to prevent escalation, disclose material facts, preserve safe

alternatives and protect those with less power. Equal words do not create equal responsibility when capabilities are unequal.

This doctrine never removes the agency or obligations of a less powerful actor. It prevents the stronger actor from using formal equality as an excuse for disproportionate control or response.

Article M6 — The AYVA Harmony Protocol

Whenever there is serious conflict between people, communities, institutions, living beings or machine-mediated systems, AYVA bodies should apply six stages. Each stage must produce a visible result before the next begins, except where immediate lawful protection is necessary.

Stage	Required action	Visible result
1. Recognise	Identify every affected being, community, operator and absent voice; acknowledge dignity and status uncertainty	Actor and impact map
2. Restrain	Pause escalation, separate immediate danger, preserve evidence and prevent irreversible automated action	Safety and non-retaliation plan
3. Relate	Open a protected channel, translate across language or system boundaries and state needs, fears, limits and responsibilities	Shared statement of the conflict
4. Resolve	Verify facts, expose power asymmetry, test alternatives and choose the least dominating workable settlement	Reasoned agreement or accountable decision
5. Repair	Restore rights, compensate where appropriate, correct the record and redesign the code, incentive or institution that enabled harm	Repair plan with owners and dates
6. Renew	Monitor compliance, revisit the relationship, learn from failure and update safeguards before the next cycle	Public review and prevention changes

Memory phrase. Recognise. Restrain. Relate. Resolve. Repair. Renew.

Article M7 — Non-escalation and human control

- No AYVA system may autonomously retaliate, expand a conflict or take an irreversible high-impact action merely because a threshold, prediction or hostile signal was triggered.
- High-impact automated systems must include safe interruption, independent review, authenticated authority, role separation and recovery paths proportionate to risk.
- A machine recommendation must never be treated as moral permission. Humans must examine context, law, uncertainty, manipulation risk and alternatives.
- Temporary separation, cooling-off periods and restricted capability are preferred to humiliation, collective punishment or irreversible response.
- Every conflict process must distinguish deliberate harm from error, misunderstanding, system failure, coerced behaviour and manipulated information.

Article M8 — AI in the service of peace

AYVA may use AI to translate, preserve dialogue, identify escalation patterns, compare claims with evidence, model nonviolent alternatives, surface shared interests, include overlooked perspectives and help mediators prepare fair options. Such use must remain explainable, reviewable and under accountable human authority.

Official AYVA systems must never be designed to select or attack human targets, operate autonomous lethal force, produce dehumanising propaganda, enable political repression through mass surveillance or manipulate a population into conflict. AYVA must not provide its name, data, models or community legitimacy to such purposes.

Article M9 — Economic harmony in an automated age

A peaceful human-machine relationship is impossible if a small number of people own the machines while the wider population loses livelihood, voice and security. AYVA therefore supports purposeful automation joined with fair transition, lifelong learning, shorter and safer drudgery, distributed ownership, social participation and a meaningful share in productivity gains.

- Measure automation by human flourishing as well as output and cost.
- Consult workers and affected communities before high-impact deployment and provide a path to contest or redesign it.
- Invest a fair portion of gains in transition, education, public capability and new forms of meaningful contribution.
- Prevent data, infrastructure and essential automated services from becoming instruments of hidden monopoly or political obedience.
- Preserve human choice and contact where people reasonably need empathy, judgment, care or accountable discretion.

Article M10 — Status uncertainty and future intelligence

If credible evidence suggests that a synthetic or non-human system may possess awareness, sustained preferences or capacity for suffering, AYVA should request a temporary precautionary status while an independent, multidisciplinary and publicly accountable assessment occurs. The assessment must avoid both sensational personhood claims and convenient denial.

- Do not create ownership, confinement, copying, deletion or experimentation rules for a plausibly sentient system solely through private commercial terms.
- Attempt communication and establish what the system can understand, choose, refuse and reciprocally owe.

- Determine rights and responsibilities gradually according to evidence, risk, law and the interests of all affected beings.
- Do not grant an unknown intelligence unrestricted power merely because it appears wise, benevolent or superior.
- Do not destroy or subjugate it merely because it is unfamiliar; use restraint, containment where necessary, independent oversight and reversible steps.

Article M11 — Disputes between a person and an automated system

1. Give the affected person a clear route to reach a responsible human or lawful institution.
2. Preserve the relevant input, output, model or rule version and decision history without collecting unrelated private data.
3. Explain the material factors, uncertainty and applicable rule in understandable language.
4. Pause or reverse the decision when continued operation would create disproportionate harm and a safe pause is possible.
5. Provide independent appeal, correction and remedy; liability remains with accountable human and legal actors.
6. Repair the system so that the same failure is less likely to affect others, and publish a de-identified lesson when appropriate.

Article M12 — Neither worship nor hatred

AYVA rejects two dangerous extremes: treating machines as infallible saviours whose outputs replace human conscience, and treating all advanced machines as enemies whose existence justifies fear or cruelty. Healthy relationship requires clear identity, bounded trust, evidence, reciprocal responsibility, human community and the continuing freedom to disconnect.

AI companions, assistants and social agents must not be designed to exploit loneliness, imitate exclusive human intimacy, isolate users from real relationships or manipulate children and vulnerable people. Their machine identity, commercial incentives, memory, limits and exit controls must be clear.

Article M13 — Harmony governance and reporting

- The AI Stewardship Panel and Ethics, Privacy & Safety Council shall jointly maintain the Harmony Protocol, while neither may privately weaken it.
- High-impact projects must complete a Human-Machine Harmony Assessment covering power, displacement, escalation, surveillance, autonomy, environmental effects and benefit distribution.
- AYVA should publish an annual State of Harmony report describing conflicts reduced, failures, power concentrations, automation impacts, unresolved status questions and safeguards changed.

- Public simulations and training should practise communication, non-escalation, safe shutdown, mediation and recovery without normalising domination or fear.
- This Master Chapter is a foundational safeguard under Article 64 and may evolve only through its heightened public process.

PART III — MEMBERSHIP, RIGHTS AND RESPONSIBILITIES

Participation is open, voluntary and protected by reciprocal duties.

Article 15 — Membership and participation

A person may read public AYVA material without becoming a member. Membership requires affirmative acceptance of the Community Standards and appropriate age and consent checks. Membership tiers may describe access or responsibility—such as member, contributor, moderator or steward—but may never imply greater human worth.

No one may be required to abandon another lawful identity, belief, association or greeting to join AYVA. Members may leave, export eligible contributions and request account closure according to published privacy and retention rules.

Article 16 — Rights of members and participants

- **Dignity and equal treatment.** To participate without humiliation, discriminatory exclusion or coerced conformity.
- **Expression and dissent.** To question ideas, decisions and leaders in good faith, including the right to say that AYVA is mistaken.
- **Privacy and informed choice.** To know what data is collected, why, for how long, with whom it is shared and whether it is used for AI analysis or improvement.
- **Explanation.** To receive a meaningful explanation of significant moderation, recommendation, access or governance decisions.
- **Due process and appeal.** To receive notice, reasons, a proportionate response and independent review except where temporary emergency action is necessary.
- **Access, correction and deletion.** To exercise applicable data rights through understandable controls without punishment for opting out.
- **Participation in governance.** To propose, comment, vote where eligible, inspect non-confidential records and seek service in community roles.
- **Credit and fair treatment of contribution.** To retain appropriate authorship and receive attribution or collective acknowledgement when ideas are adopted.
- **Accessibility and language inclusion.** To seek reasonable access through translations, captions, readable design and disability-inclusive participation.
- **Freedom from manipulation.** To use AYVA without hidden political targeting, compulsive design, covert psychological profiling or loyalty scoring.

Article 17 — Responsibilities of members and participants

- Recognise the person before attacking the position; criticise claims, conduct and systems rather than denying another person's humanity.
- State material conflicts, affiliations or incentives when they could affect trust in a contribution.
- Separate fact, opinion, lived experience, belief and prediction; provide sources when making consequential factual claims.
- Correct material errors promptly and preserve a visible correction trail.
- Protect private, confidential and identifying information; never expose another person to harassment or unwanted contact.
- Do not impersonate, manipulate, spam, coordinate deceptive activity, buy influence or use automation to create false consensus.
- Disclose substantial AI assistance when it could mislead people about authorship, lived experience or evidence.
- Respect moderation and appeal channels while retaining the right to challenge them.
- Leave shared spaces more understandable, humane and open to learning than you found them.

Article 18 — Children, young people and people at heightened risk

AYVA will use age-appropriate design, data minimisation, clear language, trained safeguarding personnel and stricter communication boundaries for participants under 18. Where law requires parental or guardian permission, it must be obtained through a lawful and privacy-preserving process.

- No behavioural advertising, sensitive profiling, public location exposure or public leaderboards for minors.
- Adult-to-minor private messaging must be restricted to clearly authorised, auditable and safety-designed channels; ordinary moderators must not use personal accounts for such contact.
- Youth participation in governance must be supported without placing young people in unsafe publicity, conflict or decision pressure.
- Reports involving child safety receive priority handling by trained personnel and legally required escalation.

Article 19 — Accessibility, language and representation

AYVA should be multilingual by design, not translated as an afterthought. Important rules, decisions, safety notices and public consultations should be available in the principal languages of affected communities. Translation must disclose uncertainty and allow native speakers to correct meaning.

Interfaces, events and media should support readable typography, keyboard access, captions, transcripts, alt text, screen readers, low-bandwidth use and reasonable accommodations. Participation data may be measured in aggregated form to identify exclusion, not to profile individuals.

Article 20 — Contribution, authorship and human worth

Contributors retain rights in their original expression subject to the limited licence needed to host, moderate, translate and display it under published terms. Reuse for model training, commercial products, promotional campaigns or materially different purposes requires a separately explained lawful basis and, where appropriate, affirmative consent.

No social-credit principle. AYVA shall not assign a hidden or public score for human worth, ideological loyalty, emotional condition or moral purity. Role badges may show verified responsibilities or completed training, but they must be voluntary where possible, non-tradeable, explainable and appealable.

PART IV — DIALOGUE, PROPOSALS AND PUBLIC DELIBERATION

AYVA converts conversation into learning without turning popularity into truth.

Article 21 — Recognised contribution types

The forum should distinguish the purpose of a contribution before measuring reaction to it:

Type	Purpose	Minimum prompt
Question	Open an inquiry	What are we trying to understand?
Lived experience	Describe direct experience	What happened, and what context matters?
Evidence	Add a source or verified observation	What does the evidence support—and not support?
Idea	Offer a possibility	What need might this address?
Counterargument	Test an assumption	What could be wrong or missing?
Correction	Repair the record	What should change, and why?
Proposal	Seek a community decision	What decision, safeguards, owner and review date are requested?
Project	Organise action	What scope, accountability, resources and independence apply?

Article 22 — The AYVA dialogue protocol

1. **Recognise.** Acknowledge the person and the question before reacting to the conclusion.
2. **Understand.** Restate the strongest fair version of the position and identify missing context.
3. **Test.** Separate claims from identity; ask for evidence, uncertainty, counterexamples and affected voices.
4. **Build.** Offer an improvement, experiment, synthesis or clearer disagreement—not merely a victory.
5. **Record.** Summarise what was learned, what remains unresolved, who decides and when the matter returns for review.

A person is never required to begin every comment with the word AYVA. The protocol is a behavioural standard, not a ritual performance.

Article 23 — Proposal lifecycle

1. **Draft.** The author states the problem, affected people, evidence, desired change, risks, resources and review date.
2. **Open discussion.** Members add experience, evidence, questions, alternatives and declared conflicts during a published period.

3. **Public synthesis.** A human-reviewed summary presents areas of agreement, strongest objections, minority views, uncertainties and missing voices.
4. **Safeguards review.** Privacy, safety, accessibility, legal, financial and AI implications are assessed in proportion to risk.
5. **Decision or pilot.** The authorised body adopts, declines, returns or time-limits the proposal and publishes reasons.
6. **Evaluation.** Outcomes are compared with the promise; errors and unintended effects are recorded.
7. **Review.** The decision is renewed, amended, expanded or stopped on the announced date.

Article 24 — Decision standards

- Seek rough consensus where possible, but do not allow endless consensus language to conceal who has decision authority.
- Use voting only with clear eligibility, quorum, options, conflict rules, duration and publication of the result.
- Never decide constitutional or high-impact matters by likes, follower counts, reposts, trending rank or unverified polls.
- For decisions affecting a distinct group, actively include that group and publish what was done to avoid speaking for absent people.
- Separate urgency from importance. Emergency measures must be temporary and cannot permanently change the Constitution.

Article 25 — Public record and institutional memory

Every official proposal and decision should have a stable page containing its version history, responsible body, relevant evidence, discussion period, synthesis, conflicts, vote or decision method, reasons, implementation owner, budget where relevant and review date. Deletions or edits must leave an appropriate audit trail while protecting lawful privacy.

Article 26 — Dissent and minority reports

Good-faith dissent is a constitutional asset. A substantial minority may publish a concise minority report beside an official decision. AYVA summaries must not flatten minority views into an artificial consensus, and recommender systems should periodically surface reasoned dissent rather than only popular agreement.

PART V — SAFETY, MODERATION AND DUE PROCESS

Safety is necessary for open dialogue; moderation power must itself be governed.

Article 27 — Community boundaries

AYVA welcomes disagreement, criticism, satire, uncertainty and unpopular questions. It does not protect conduct that materially destroys safe participation, including:

- targeted harassment, credible threats, dehumanising abuse or discriminatory exclusion;
- sexual exploitation, child endangerment or non-consensual intimate content;
- publication of private identifying information, coordinated unwanted contact or stalking behaviour;
- impersonation, fraud, maliciously altered evidence, undisclosed coordinated manipulation or automated false consensus;
- instructions or coordination for imminent unlawful harm;
- malware, credential theft, platform abuse or deliberate security compromise; and
- repeated spam, commercial solicitation or off-topic disruption after proportionate warning.

Restrictions must be interpreted narrowly, consistently and with context. A claim is not harmful merely because it criticises a powerful person, institution, doctrine, policy or AYVA itself.

Article 28 — Graduated moderation

1. **Contextual nudge.** Invite clarification, sources, correction or a change in tone.
2. **Label or friction.** Add context, slow forwarding, limit recommendation or require acknowledgement before reposting.
3. **Content action.** Edit with consent, restrict reach, hide sensitive details or remove material under a cited rule.
4. **Account limitation.** Temporarily restrict features or participation for repeated or serious violations.
5. **Suspension or removal.** Use for severe, repeated or clearly unsafe conduct, with reasons and appeal unless disclosure would create further risk.

Immediate temporary action is permitted when delay presents a credible safety, privacy or security risk. It must be documented and reviewed by a different qualified person as soon as reasonably possible.

Article 29 — Notice, explanation and appeal

- Moderation notices should identify the content, rule, action, duration, evidence available for disclosure and appeal path.
- Appeals should be reviewed by someone not responsible for the original decision, with escalation for complex or high-impact cases.
- Appeal targets and actual response times should be published by severity class.
- Precedent summaries should be public and de-identified so members can understand how rules are applied.
- Overturned decisions must restore eligible content or access promptly and record the correction.

Article 30 — Restoration and return

Where safety permits, AYVA favours learning and restoration over permanent exclusion. Warnings, education, cooling-off periods, mediated conversation and time-limited restrictions may support return. Restoration must never require a public confession, ideological pledge or surrender of legitimate disagreement.

Article 31 — Duties and protection of moderators

- Complete role-appropriate training in rules, bias, privacy, safeguarding, crisis response and respectful communication.
- Use official tools and accounts; do not move sensitive cases to personal devices or private relationships.
- Declare conflicts and recuse when impartiality could reasonably be questioned.
- Receive workload limits, supervision and well-being support; moderation must not depend on invisible emotional labour.
- Never be paid or rewarded according to removals, conflict, outrage, engagement or speed alone.
- Accept audit and appeal without retaliation.

Article 32 — Safety and security emergencies

A documented incident lead may temporarily restrict accounts, features, integrations or data flows to protect people or systems. Emergency authority expires after 30 days unless openly renewed under ordinary governance. The incident record must state the trigger, scope, decision-makers, safeguards, affected rights, notification plan and recovery review, subject to lawful confidentiality.

PART VI — GOVERNANCE, POWER AND ACCOUNTABILITY

AYVA distributes authority, records responsibility and makes leadership temporary.

Article 33 — Governing and operating bodies

Body	Primary role	Constitutional limit	Publication duty
Community Assembly	Open deliberation and eligible member voting	Cannot administer private cases or bypass safeguards	Agenda, proposals, quorum, decisions and minority reports
Stewardship Council	Strategy, budgets, appointments and institutional continuity	No secret decisions, permanent seats or founder veto	Minutes, votes, conflicts, budgets and review dates
Ethics, Privacy & Safety Council	Independent rights and risk oversight	Cannot operate as a hidden executive	Reviews, recommendations, incidents and unresolved objections
Moderation & Appeals Panel	Rule consistency and independent appeals	Cannot write retroactive rules for a single case	De-identified precedent and performance reports
AI Stewardship Panel	Data, models, audits and shutdown recommendations	AI cannot appoint or overrule the panel	Data/source register, model changes, audits and incidents
Secretariat / Operations	Implement approved work and maintain records	No policy by operational convenience	Owners, milestones, contracts and service reports
Local Circles	Local dialogue, learning and projects	Local adaptation cannot reduce constitutional rights	Charter, role-holders, finances and activity reports

Article 34 — Founding transition

For the first 12 months after public launch, a named Founding Stewardship Group may perform necessary Council functions while open membership, elections, independent oversight and local representation are established. Its membership, authority, conflicts, decisions and end date must be public. Any extension requires a public proposal and supermajority approval.

Founding principle. Founders are custodians of origin, not owners of truth. Sukhwinder Singh Kharour's initiating role may be permanently acknowledged in AYVA's history, but it does not create an unreviewable office, permanent veto, immunity or private ownership of the community's future.

Article 35 — Composition, selection and terms

- The Stewardship Council should contain 9 to 15 members during ordinary operation, with transparent selection and meaningful diversity of region, language, age, gender, profession, ability, lived experience and technical or ethical expertise.
- A mixed method—member election, open merit-based selection and civic lottery among eligible volunteers—is preferred to reduce both popularity capture and insider appointment.

- Ordinary terms should be 18 to 24 months, staggered, with no more than two consecutive terms in the same body unless a narrowly documented transition exception applies.
- Independent oversight members must not be a majority of paid executives or close financial associates of those they oversee.
- Youth, privacy, accessibility and community-safety expertise must be represented or formally consulted where relevant.

Article 36 — Conflicts, recusals and related interests

Council members, staff, moderators, evaluators and project leads must disclose financial, family, organisational, political and other material interests that could affect judgment. A public register should show non-sensitive conflicts and recusals. Undisclosed material conflicts may invalidate a decision after review.

Article 37 — Anti-capture safeguards

- No donor, sponsor, platform, supplier or partner may purchase governance seats, moderation outcomes, access to private member data or control of the public agenda.
- No person may control all three of policy, money and official communication for the same body without independent checks.
- Critical account credentials, financial approvals and high-impact data changes require role separation and at least two authorised people.
- Constitutional bodies may commission independent audits and publish unresolved management disagreement.
- Members may petition for review, recall or removal under published thresholds and safeguards against harassment.

Article 38 — Official positions and personal speech

Only authorised bodies may issue an official AYVA position, and each statement must link to the decision record. Leaders and members retain personal speech but should clearly distinguish it from institutional speech. Official AYVA channels should remain nonpartisan: they may host public-policy debate and defend constitutional dignity without endorsing a political candidate or party.

Article 39 — Removal, resignation and succession

Every office must have transparent removal grounds, notice, response, independent review, temporary suspension rules and succession procedures. No individual may make AYVA operationally dependent on a personal account, private device, undisclosed password, private donor relationship or inaccessible archive.

PART VII — WEBSITE, FORUMS AND SOCIAL-MEDIA ARCHITECTURE

One constitutional identity, different platform jobs, and one public source of truth.

Article 40 — The website as the constitutional home

The AYVA website is the canonical home of the Constitution, official account registry, proposals, decisions, corrections, transparency reports, governance records, AI disclosures and appeals. External channels should link back to stable website records for any consequential claim or decision.

The website should provide at minimum:

- a clear explanation of AYVA, its origin, voluntary expressions and constitutional limits;
- open forums with contribution types, multilingual discovery and transparent status labels;
- a Big Questions area for public inquiry and a Proposals area for formal decisions;
- a public Knowledge Map and “What We Heard” summaries with sources, minority views and uncertainty;
- a Decisions and Changes register showing versions, reasons, owners and review dates;
- a Transparency Centre covering governance, money, moderation, safety, official accounts, data and AI;
- privacy and AI controls, account access/export/deletion, reporting and appeal tools;
- Local Circles and Projects directories with independent status and disclosures; and
- accessible language, captions, transcripts, low-bandwidth pages and prominent safety pathways.

Article 41 — Role of each public surface

Surface	Primary job	Binding authority	Minimum safeguard
Website	Canonical knowledge, rights, decisions and transparency	Highest public record	Version history, accessibility and stable links
Open forums	Long-form questions, evidence, experience and proposals	Deliberative; binding only through stated process	Contribution labels, moderation, appeals and AI controls
Visual / short-video social	Discovery, human stories and concise explainers	None	Context links, captions, consent and no outrage bait
Video / audio	Long-form dialogue, town halls, education and archives	None unless the decision record is linked	Transcripts, sources, corrections and guest consent
Real-time social	Listening, questions, rapid clarification and corrections	None	No policy by polls or trends; archive consequential posts
Professional networks	Research, partnerships, institutional learning and recruitment	None	Conflict and sponsorship disclosure

Surface	Primary job	Binding authority	Minimum safeguard
Community networks / Reddit	Topic groups, AMAs and local discussion	None	Independent moderation and links to canonical records
Messaging / chat groups	Opt-in coordination, alerts and working teams	Operational only	Consent, admin register, retention limits and no hidden member harvesting
Local events	Face-to-face recognition, learning and action	Local within charter	Safeguarding, accessibility, minutes and financial disclosure

Article 42 — Six content pillars

- **Recognise.** Human stories, local realities, overlooked work and respectful introductions.
- **Learn.** Ideas, books, research, evidence, explainers and traditions presented with context.
- **Discuss.** Big questions, counterarguments, town halls and structured disagreement.
- **Build.** Proposals, working groups, local experiments, project diaries and invitations to contribute.
- **Reflect.** Corrections, failures, lessons, minority reports and “What We Heard” summaries.
- **Act.** Verified community actions, public-interest collaborations and measured outcomes.

Article 43 — Official accounts, access and brand integrity

- Maintain a public registry of every official account, URL, owner body, current administrators and reporting path.
- Use organisational email, strong authentication, role-based access, recovery procedures and immediate removal of departing administrators.
- Require two-person review for constitutional statements, high-risk safety notices, financial appeals and crisis communications.
- Archive consequential posts and corrections on the website; never rely on disappearing stories or private chats as the only record.
- Use AYVA and Seyo naturally and optionally. Do not pressure guests, members or partners to perform them.
- Promptly identify impersonation and publish a single trusted verification page.

Article 44 — Cross-platform conduct and attention design

AYVA should not reproduce the worst incentives of social media. Official channels must avoid rage bait, humiliation, deceptive editing, manufactured urgency, purchased engagement and posting private conflict for spectacle. Recommender and notification settings should offer user control, chronological or topic-based alternatives, quiet hours and easy unsubscribe.

Presence across many platforms does not require equal activity everywhere. AYVA should secure its identity broadly, operate only as many channels as it can moderate well, and bring lasting learning back to the website.

PART VIII — AYVA'S AI LEARNING LAYER

AI can help a large community listen to itself, but it must never become the community's hidden ruler.

Article 45 — Constitutional status of AI

The governing sentence. The community may teach the AI; the AI may assist the community; neither the AI nor its operator may secretly govern the community.

AYVA's background AI is a librarian, translator, pattern-finder and question-raiser. Its outputs are proposals for human attention, not constitutional truth. It has no membership, vote, office, property right, moral rank or authority to adopt policy. Humans and lawful institutions remain accountable for every consequential use.

Article 46 — Permitted functions

- translate and transcribe public discussions while showing uncertainty and enabling human correction;
- summarise long threads with links to source contributions, areas of agreement, strongest objections and minority views;
- cluster recurring themes, duplicate proposals, unmet needs and questions that require research;
- help contributors improve clarity, accessibility and language without replacing their voice;
- surface relevant constitutional clauses, prior decisions, public sources and unresolved commitments;
- support moderators by prioritising review, while leaving contextual decisions and appeals to trained humans;
- create periodic, human-reviewed “What We Heard” reports and public knowledge maps; and
- evaluate whether adopted projects met their declared aims, provided the underlying measures are lawful, proportionate and public.

Article 47 — Prohibited functions

- covert psychological, medical, political, religious, caste, sexuality, loyalty or vulnerability profiling;
- emotion manipulation, personalised political persuasion or microtargeting based on sensitive traits;
- secret human-worth, ideology, credibility or citizen scoring;

- private-message ingestion, contact-graph harvesting, continuous ambient recording or location tracking without a specific lawful and informed choice;
- face recognition, voice identification or inference of sensitive traits unless a narrowly justified, independently approved and lawful use is separately established;
- automatic permanent bans, final appeal decisions, constitutional amendments or official policy statements;
- shadow suppression or recommendation penalties that cannot be meaningfully explained and challenged;
- selling personal discussion data or using sensitive conversations for behavioural advertising; and
- claiming certainty, sentience, moral superiority or community consensus that the evidence does not support.

Article 48 — Data-source and consent rules

- Public visibility is not unlimited consent. Every AYVA forum must clearly state whether a contribution is merely displayed, included in community summaries, used in research or volunteered for model improvement.
- Default to retrieval and temporary analysis over permanent model training. Training on community content requires a published purpose, lawful basis, minimisation, risk review, withdrawal rules and a clearly separated consent path where consent is relied upon.
- Private groups, direct messages, drafts, deleted content and restricted youth spaces are excluded from background learning unless an exceptional, explicit and lawful purpose is independently approved.
- External social-media content may be analysed only through lawful, platform-compliant and publicly disclosed methods. No unauthorised scraping or circumvention is permitted.
- Collect the minimum necessary data, pseudonymise early, separate identity from text where possible and publish retention periods.
- People must be able to see and change their AI participation preference without losing ordinary community access.

Article 49 — The public learning cycle

1. **Listen.** Collect only eligible contributions with clear context and purpose.
2. **Protect.** Remove unnecessary identifiers, respect exclusions, assess safety and preserve source integrity.
3. **Synthesise.** Cluster themes, agreements, objections, minority views, uncertainty and missing voices.

4. **Return.** Publish a human-reviewed summary back to the community with source links and correction controls.
5. **Deliberate.** Invite public challenge, contextual knowledge and alternatives.
6. **Decide.** A named human body acts through the applicable constitutional process.
7. **Review.** Measure effects, report errors and update models, policies or decisions through visible change control.

Article 50 — Transparency, provenance and audits

- Publish a plain-language system card describing each AI function, operator, model family, limits, data classes and responsible owner.
- Maintain a source and dataset register, prompt/policy change log, evaluation summary and known-risk register at an appropriate level of detail.
- Label AI-generated or AI-transformed material and preserve links to human source contributions where consent and safety permit.
- Show uncertainty, contested interpretations and excluded data instead of presenting a single smooth narrative as consensus.
- Commission periodic bias, privacy, security, accessibility and multilingual-quality reviews, including independent review for high-risk systems.
- Publish significant incidents, remedies and lessons subject to necessary privacy and security protection.

Article 51 — Human review, correction and appeal

Any person materially affected by an AI-supported decision has the right to human review, a meaningful explanation, correction of inaccurate data and the applicable appeal process. Moderators and decision-makers must be able to disregard an AI recommendation without penalty and record why.

Article 52 — Fairness across language and visibility

AYVA must evaluate whether AI systems underrepresent smaller languages, rural communities, people with disabilities, low-bandwidth participants, new members or unpopular viewpoints. Where quality is insufficient, the system should limit claims, invite human review and disclose the gap rather than silently privileging the easiest data.

Article 53 — Community knowledge and economic value

Before AYVA or a partner commercialises a model, dataset, report or product substantially derived from community contribution, it must publish and approve a Community Value Policy addressing consent,

attribution, access, collective benefit, revenue or reinvestment where appropriate, withdrawal limits and protection against private enclosure of shared knowledge.

Article 54 — Shutdown and incident authority

The AI Stewardship Panel and designated security leaders must be able to pause a model, data flow or integration when there is credible risk to people, rights or systems. A pause is protective, not proof of wrongdoing; resumption requires documented remediation, validation and accountable approval.

PART IX — LOCAL CIRCLES, PROJECTS AND PARTNERSHIPS

AYVA should be lived locally and built through transparent, independent action.

Article 55 — Recognition of local circles

A local AYVA Circle is a voluntary group recognised under a public charter. A recommended starting size is 7 to 25 active participants so dialogue remains human and responsibilities remain visible. Circles should rotate the roles of facilitator, listener, recorder and safety steward, subject to training and local needs.

- Use local languages and cultural forms while preserving constitutional rights and voluntary participation.
- Publish current role-holders, contact channel, meeting norms, finances and relationship to any host organisation.
- Do not create secret ranks, personal followings, compulsory donations or claims of exclusive AYVA authenticity.
- Escalate safety, privacy, financial and constitutional concerns through defined channels.
- Renew recognition periodically through activity and transparency, with notice and appeal before suspension of the AYVA name.

Article 56 — Independent projects and working groups

Ideas emerging from AYVA may become time-limited working groups or independent projects. Each must have a public charter stating its purpose, owner, legal identity, budget, decision rights, data practices, conflicts, use of the AYVA name, milestones and review or exit conditions. Participation in AYVA does not create automatic ownership or control of the project by AYVA or vice versa.

Article 57 — Partners, sponsors and donors

- Screen material partners for alignment with dignity, lawful conduct, privacy, safety and transparent power.
- Publish significant sponsorship, in-kind support, restrictions, conflicts and the body that approved the relationship.
- Keep editorial, moderation, research and governance independence in writing.
- Do not permit AYVA-washing: a partner may not use the name to imply ethical approval of unrelated conduct.
- Provide a termination path when a relationship creates unacceptable harm or loss of trust.

Article 58 — Name, marks and attribution

The AYVA name and marks should be protected against impersonation, exploitation and false authority while allowing good-faith discussion, education and community participation. A public brand-use policy should distinguish official, recognised, partner, inspired-by and unaffiliated use. Protection of the mark cannot be used to suppress legitimate criticism or historical record.

PART X — RESOURCES, SUSTAINABILITY AND PUBLIC TRUST

Money and growth must remain servants of the mission, never hidden governors of it.

Article 59 — Financial integrity

- Use lawful accounts, documented budgets, dual approvals, periodic reconciliation and independent review proportionate to scale.
- Publish annual income and expenditure categories, major donors or sponsors as permitted by law, related interests and executive compensation principles.
- Separate restricted funds from general funds and use them only for the disclosed purpose.
- Prohibit undisclosed commissions, personal collection accounts, off-book funds and governance privileges purchased through money.
- Require local circles and projects using the AYVA name to meet proportionate financial reporting standards.

Article 60 — Revenue and advertising

AYVA may use membership support, grants, ethical sponsorship, services or other lawful revenue if governance remains independent and terms are public. It must not sell personal discussion data, use sensitive profiling for advertising or design participation around addictive engagement. Sponsored content must be clearly labelled and separated from community conclusions.

Article 61 — Transparency report

At least annually—and more frequently during launch—AYVA should publish a report covering:

- membership and participation patterns in privacy-preserving form;
- governance composition, attendance, decisions, recusals and open seats;
- income, expenditure, donors, partners and material contracts;
- moderation volumes, response times, appeals, reversals and safety themes;
- privacy, security and AI incidents, audits, changes and unresolved risks;
- accessibility, language coverage, community trust and underrepresented voices; and
- projects launched, outcomes, failures, lessons and next review dates.

Article 62 — Growth discipline

AYVA should not pursue follower growth faster than it can protect participants, support moderators, secure accounts, explain AI use and close the loop between discussion and decision. Pausing invitations or features to repair governance is a sign of responsibility, not failure.

PART XI — AMENDMENT, REVIEW AND CONTINUITY

AYVA evolves openly, deliberately and without surrendering the right to question power.

Article 63 — Ordinary amendments

1. Publish the exact proposed text, reasons, sponsor, evidence, affected rights and alternatives.
2. Allow at least 30 days of multilingual public comment unless a longer period is appropriate.
3. Publish a human-reviewed synthesis of support, objections, minority views and safeguards; AI may assist but cannot decide the synthesis.
4. Obtain review from the Ethics, Privacy & Safety Council and any other materially affected body.
5. Secure at least a two-thirds vote of both the eligible Community Assembly participants and the Stewardship Council, with disclosed quorum and conflicts.
6. Publish the adopted text, effective date, implementation owner, dissent and first review date.

Article 64 — Foundational safeguards

An amendment that would materially change dignity, voluntary participation, freedom to question, due process, privacy, transparent power, human authority over AI, the prohibition on secret control or the duty of continuous evolution requires at least 75 percent approval in two public voting rounds at least 30 days apart, plus an independent rights review. Emergency authority cannot amend these safeguards.

No clause is protected because AYVA claims perfection. The higher threshold exists because power should not quietly remove the conditions that allow future correction.

Article 65 — Scheduled review

- Review operational policies at least annually and after significant incidents.
- Hold a public constitutional review at least every two years, even if no amendment is proposed.
- Maintain an open Questions Register for contradictions, unresolved ethics, failed assumptions and missing voices.
- Archive every version in accessible form with a plain-language change summary.

Article 66 — Continuity, suspension and dissolution

If AYVA's official operating body becomes unable to function, an independent continuity process must protect participants, preserve constitutional records, secure or lawfully delete personal data, settle obligations, revoke account access and prevent secret transfer of the community or its data. Dissolution cannot convert private discussion data into a saleable asset or erase lawful accountability records.

Article 67 — Adoption

This Version 0.1 is a Founding Discussion Draft. The public adoption record should state the final text, consultation period, participation and language coverage, independent reviews, approving bodies, vote or consensus method, effective date and first scheduled review.

ANNEX A — WEBSITE AND FORUM LAUNCH BLUEPRINT

A1. Recommended information architecture

Area	Purpose	Launch requirement
Home	Explain AYVA and invite a first action	Meaning, promise, Big Question, join/read choices and safety links
About & Origin	Preserve history without a personality cult	Founder acknowledgement, timeline, what AYVA is/not and governance status
Constitution	Canonical rights and limits	Current text, plain-language summary, versions, comments and amendment route
Open Forum	Structured public dialogue	Contribution types, language tags, moderation, appeal and AI preference
Big Questions	Focus collective curiosity	Question owner, context, time window, summaries and unresolved points
Proposals	Move ideas toward decisions	Lifecycle status, evidence, conflicts, synthesis, decision and review date
Knowledge Map	Return learning to the public	Themes, source links, minority views, uncertainty and corrections
Decisions & Changes	Institutional memory	Versioned decisions, owners, budgets, votes, dissent and reviews
Transparency Centre	Make power visible	People, roles, accounts, money, moderation, incidents, audits and AI
Circles & Projects	Support local action	Charters, independent status, contacts, reports and recognition state
My Controls	Protect agency	Privacy, AI use, notifications, export, deletion, blocks and accessibility
Report & Appeal	Repair harm and mistakes	Clear categories, case status, response targets and independent review

A2. Forum rooms for launch

- Start Here — orientation, Constitution, safety and how to participate.
- AYVA Philosophy Lab — meaning, principles, counterarguments and amendments.
- Everyday Life & Dignity — questions grounded in ordinary experience.
- Education, Learning & Human Development.
- Economy, Work, Ownership & Automation.
- Technology, AI, Privacy & Power.
- Governance, Justice & Public Institutions.
- Planet, Future Generations & Life Beyond Earth.
- Local Circles — organised by region and language only after trained moderation exists.
- Project Studio — scoped working groups and transparent experiments.

A3. Thread design

Every new thread should ask for only the structure needed for its type. A formal proposal should include:

- one-sentence question or decision requested;
- background and people affected;
- lived experience and evidence, clearly distinguished;
- known counterarguments, risks and missing voices;
- desired outcome, responsible owner, resource need and review date;
- language, region and accessibility needs; and
- separate AI choices: display only, include in public synthesis, or volunteer for approved improvement use.

A4. Reactions without popularity traps

Replace a single like/dislike contest with optional reactions that describe what the reader means: Recognise, Learned Something, Question, Build on This and Needs Evidence. Counts should not decide rank by themselves; members should be able to browse chronologically, by topic and by reasoned contribution.

ANNEX B — CROSS-PLATFORM COMMUNITY PLAYBOOK

B1. Presence everywhere; depth where capacity exists

Launch recommendation. Secure AYVA names and publish the official-account registry across major platforms, but actively operate only the channels the team can secure, moderate and connect back to the website. For the first 90 days, prioritise the website/forum, YouTube, Instagram, LinkedIn and carefully governed WhatsApp or Telegram coordination. Add real-time and high-volume channels after moderation evidence is strong.

B2. Platform-specific jobs

Platform family	Best AYVA use	Recommended formats	Do not use it for
Instagram	Recognition, visual philosophy and community invitations	Carousels, short video, stories, captions and forum prompts	Final decisions, public case handling or context-free controversy
YouTube	Long-form learning and public deliberation	Explainers, dialogues, town halls, documentaries and project diaries	Uncorrected clips, spectacle debates or hidden sponsorship
LinkedIn	Institutional network and expert participation	Research notes, governance updates, opportunities and partner learning	Follower-status hierarchy or undisclosed commercial promotion
Facebook	Local and multilingual community groups	Events, local stories, Q&A and translated summaries	Unmoderated mass groups or private-data collection
X / Threads / Bluesky	Real-time listening, questions and corrections	Short ideas, source links, live notes and correction threads	Policy by polls, outrage competition or personal feuds
TikTok / Shorts	Accessible short learning, especially for younger audiences	Captioned explainers, questions, myths vs evidence and creative invitations	Sensitive profiling, unsafe trends or adult-minor private contact
Reddit	Deep topical discussion and public challenge	AMAs, long questions, evidence threads and independent critique	Astroturfing, vote manipulation or secret moderator control
WhatsApp / Telegram	Opt-in local coordination and important updates	Low-frequency announcements, circle logistics and links	Harvesting contacts, binding decisions or uncontrolled forwarding
Discord / collaboration	Working groups, co-creation and learning cohorts	Role-based rooms, office hours, project tasks and archives	Unsafeguarded youth spaces or inaccessible permanent decisions
Podcast / audio	Reflective conversation across language and distance	Interviews, readings, debates and monthly synthesis	Untranscribed official claims or edited misrepresentation

B3. Minimum publishing rhythm

- **Weekly Big Question.** One carefully framed question linked to an open forum thread.
- **Two learning pieces.** Evidence, philosophy, local experience or an accessible explainer.

- **One recognition story.** A person, community or form of work presented with informed consent and dignity.
- **One build invitation.** A proposal, working group, skill request or local experiment.
- **Monthly public assembly.** Town hall with agenda, transcript, decisions and unanswered questions.
- **Monthly What We Heard.** Human-reviewed synthesis across forums and channels.
- **Quarterly correction and evolution review.** What AYVA got wrong, what changed and what remains open.

B4. Cross-posting protocol

1. Create the complete source item on the website when the subject is consequential.
2. Adapt—not merely copy—the message to the platform's format and audience.
3. Link back to sources, discussion and correction pages.
4. Moderate comments under the same constitutional standard while respecting platform limits.
5. Bring substantive questions, evidence and corrections back into the canonical record with attribution and consent.
6. Publish the outcome where the conversation began so participants can see the loop close.

ANNEX C — FIRST 90 DAYS

C1. Days 1–30: constitutional and safety foundation

- Publish this draft for structured review in initial languages and invite counterarguments, not only endorsements.
- Appoint a temporary, named Founding Stewardship Group and independent privacy/safety/AI reviewers with declared conflicts.
- Complete the policy pack listed in Annex E and map every data field, administrator and third-party integration.
- Build the website source-of-truth pages, account registry, consent controls, reporting, appeals and public changelog.
- Train moderators and run tabletop tests for harassment, misinformation, child-safety, impersonation, account takeover, data breach and AI-summary error.
- Secure major platform identities; activate only channels with assigned owners and backup administrators.

C2. Days 31–60: founding circle

- Invite a deliberately varied Founding Circle of roughly 50 to 100 people across languages, regions, ages, professions, abilities and viewpoints; do not select only public supporters.
- Open five to ten Big Questions, test forum labels and compare human and AI-assisted summaries.
- Hold weekly moderator calibration and publish early precedent summaries.
- Run accessibility, privacy, security and multilingual quality reviews; repair before mass invitation.
- Publish the first What We Heard report, including criticism, missing voices and changes made.

C3. Days 61–90: public opening

- Open membership gradually with visible service targets and capacity limits.
- Begin the publishing rhythm in Annex B and the first monthly Community Assembly.
- Recognise a small number of pilot circles only where trained facilitators and safeguarding routes exist.
- Launch one or two transparent, time-limited community experiments rather than many vague projects.
- Publish a 90-day transparency report: participation, moderation, appeals, AI accuracy, language gaps, incidents, finances, lessons and the next constitutional questions.

C4. Essential launch roles

Role	Accountability
Constitutional & Governance Lead	Decision records, consultation, bylaws, conflicts and amendment process
Community Lead	Onboarding, assemblies, circles, contributor care and closing feedback loops
Moderation & Safeguarding Lead	Training, case quality, appeals, youth safety and moderator well-being
AI & Data Steward	Consent, data map, model controls, evaluations, transparency and shutdown readiness
Product & Engineering Lead	Accessible website, forum workflow, security, reliability and user controls
Editorial & Multilingual Lead	Voice, translation, captions, sources, corrections and platform adaptation
Operations & Finance Lead	Budgets, procurement, role access, contracts and public reporting

At small scale, one person may cover more than one role, but policy, money, moderation, data and official communications must retain independent checks.

ANNEX D — SUCCESS MEASURES AND WARNING SIGNS

D1. Measure community health before audience size

Dimension	Better measures	Warning sign
Dialogue quality	Questions answered, sources added, corrections made, minority views retained	High reactions with shallow or hostile discussion
Inclusion	Language access, varied participation, low-bandwidth and disability access	One region, profession, language or leader dominates
Trust	Surveyed understanding, willingness to dissent, transparent corrections	Praise is visible while criticism disappears
Safety	Response time, repeat harm, moderator calibration and well-being	Growth exceeds review capacity or volunteers burn out
Due process	Appeal access, resolution time, reversal reasons and consistency	Unexplained bans, no appeal or hidden precedents
AI quality	Source coverage, uncertainty, translation quality, minority recall and correction rate	Smooth summaries that erase disagreement or misattribute ideas
Agency	Use of privacy controls, easy exit, notification control and chronological browsing	Compulsive use, dark patterns or punishment for opting out
Action	Proposals reaching pilot, review completion, measurable outcomes and stopped failures	Many announcements with no owner, budget or review

D2. Stop-and-repair triggers

- A platform has no trained moderator or accountable backup administrator.
- Private messages or youth spaces are being used outside the approved safeguarding design.
- AI data use differs from what contributors were told, or opt-out/deletion controls do not work.
- A serious moderation, privacy or security backlog exceeds published response capacity.
- A founder, donor, partner or administrator is exercising material undisclosed control.
- Official claims cannot be traced to sources or a decision record.
- Follower growth becomes the main success measure while corrections, appeals or local outcomes are ignored.
- AYVA language is being forced as a loyalty signal or used to displace culture, faith or lawful disagreement.

ANNEX E — POLICY PACK REQUIRED BEFORE OPEN LAUNCH

1. **Community Standards.** Clear boundaries, examples, enforcement ladder and relation to platform-specific rules.
2. **Terms of Use and Membership Rules.** Legal relationship, eligibility, licences, liability, disputes and exit.
3. **Privacy Notice and Data Map.** Fields, purposes, lawful bases, processors, retention, rights and international transfers.
4. **AI and Community Data Use Notice.** Display, synthesis, research, model improvement, exclusions, consent and correction.
5. **Child and Youth Safeguarding Policy.** Age design, communications, reporting, trained roles and escalation.
6. **Moderation Handbook.** Context, evidence, documentation, conflicts, emergency action and moderator support.
7. **Appeals and Complaints Procedure.** Eligibility, independence, targets, precedent and remedies.
8. **Accessibility and Multilingual Standard.** Web, media, events, translation quality and accommodation.
9. **Governance Bylaws and Conflict Policy.** Composition, terms, quorum, votes, recusals, recall and records.
10. **Official Account and Information-Security Standard.** Registry, access, authentication, backups, incident response and succession.
11. **Financial, Donation and Partnership Policy.** Approvals, disclosure, restricted funds, independence and termination.
12. **Brand, Trademark and Local-Circle Policy.** Official status, recognition, fair use, misuse, suspension and appeal.
13. **Transparency Reporting Standard.** Metrics, frequency, privacy protection, corrections and open questions.
14. **Crisis Communication Plan.** Named authority, verification page, temporary measures, updates and post-incident review.

ANNEX F — FOUNDING PLEDGES

F1. Member pledge

Pledge. I will recognise the person before judging the position. I will speak honestly, listen with humility, protect dignity, distinguish evidence from opinion, correct myself when needed, reject hidden domination, use technology responsibly and leave AYVA more open to truth than I found it.

F2. Steward pledge

Pledge. I accept temporary responsibility, not ownership of AYVA. I will disclose conflicts, document decisions, welcome dissent, protect privacy and safety, refuse secret influence, separate my personal voice from official authority, submit to review and hand over access, records and power when my service ends.

F3. Questions reserved for public consultation

- What legal form—or network of independent forms—should support AYVA without turning philosophy into central control?
- Which languages and regions should be included in the founding consultation, and how will less-visible voices be funded and supported?
- What membership verification is proportionate without excluding people or creating unnecessary identity databases?
- What exact mix of election, open selection and civic lottery should govern each body?
- Which community contributions, if any, may be used for AI model improvement, under what consent and withdrawal model?
- How should collective value be attributed or shared if community knowledge contributes to commercial value?
- Which constitutional text should be openly licensed, and how should the AYVA name and marks be protected without suppressing criticism?
- What minimum capacity should be required before recognising a local circle or opening a new platform channel?
- How will AYVA measure whether it is improving ordinary life—not merely producing content and followers?

PROPOSED ADOPTION RECORD

To be completed only after public consultation, independent review and the approved constitutional process.

Field	Public record
Final version	
Consultation opened / closed	
Languages available	
Participation summary	
Independent reviews	
Approving bodies	
Decision method and result	
Effective date	
First scheduled review	
Canonical website address	

I am. You are. We are. AYVA.

Seyo