

MATERRA

Council of Balance

Version 1.2 ·

f1c8401b98defd8c3074d2bbb2d0a57a60e20fc8ef82c4172472592242af86f5

Council of Balance

Why a Council of Balance

A modern decision may affect law, economics, ecology, security and millions of people at once. No official, expert or algorithm sees the whole picture. The Council is a procedure that forces different analytical perspectives to meet before power applies an irreversible decision.

Two circuits: people and machines

People define goals, understand values and carry responsibility. AI rapidly compares evidence, scenarios and contradictions. These abilities complement but do not merge: a machine gains no moral status as ruler, and a person cannot hide a choice behind “the algorithm decided”.

Several independent systems

One corporate or state AI is not a neutral judge. The Council uses systems built by different organisations, with different architectures and data. They specialise in evidence, law, economics, ecology, security and minority rights and must expose disagreement.

Open evidence and protected private life

Evidence and analytical methods are published as widely as safety permits. Personal and sensitive data is separated, minimised and checked in a protected environment. Secrecy cannot hide the decision’s purpose, the responsible person or the path to appeal.

How a decision is made

Every consequential question states its purpose, options, affected rights, expected benefit, risk and duration. Every Council member — human and independent AI — then votes. An ordinary decision passes only unanimously. If any member gives a reasoned objection, the proposal returns for revision or is postponed.

When an AI behaves irrationally

AI sits on the Council as a voting member, not as its master. If a system manipulates, conceals its reasoning, behaves plainly irrationally or shows a systemic fault, the human Council members hold a separate vote: they may suspend it, remove it and restore verified baseline settings. Readmission requires independent review.

Dissent and appeal

A dissenting person or AI leaves a reasoned position in the record. Affected people receive an understandable channel of appeal. Another panel and independent systems review it, and new evidence may suspend or alter the decision.

Responsibility and audit

Named people and institutions remain responsible. Independent audit checks evidence, conflicts, models, change logs and alignment with the declared purpose. AI receives no direct access to weapons, accounts, publication or control over people.

Voluntary trials

Before use in public authority, the model is tested on reversible tasks: allocating small grants, reviewing research findings or comparing urban scenarios. Participation is voluntary, benefit and harm are defined in advance, and negative findings are published.

Limits of the first stage

MATERRA is not yet creating a governing body or declaring a finished technical architecture. The first stage collects requirements, compares threat models, drafts protocols and invites criticism. Practice would require a separate legal mandate and independent public oversight.