

Artificial intelligence principles and practice

How models are governed, overseen and made accountable across the institution.

DERIVED FROM	POLICY	PAGES CONSOLIDATED
Ethical Technology Charter	IGC-DT-005	4

Governing artificial intelligence

The group uses learning systems in 26 approved ways and in no others. Each use is entered in a register before it runs, sorted into one of four classes, approved by the Ethics & Technology Council for a stated scope and a period of not more than twelve months, and withdrawn when it drifts outside either. The fourth class, autonomous action, is prohibited.

The register comes first

No learning or generative system is used for any group purpose until it is entered in the artificial intelligence register and approved. The entry names the class, the purpose, the accountable executive, the system owner, the datasets used for training, tuning and prompting with their origin, the population the system acts on, the inferences it is designed to draw and the inferences it is forbidden to draw.

The entry also states how those prohibitions will be tested and the tier the system carries under the Model Risk Management policy. A learning system is a model as well as a system, so it takes an owner, a tier and a validation cycle in addition to everything asked of it here. Nothing in the artificial intelligence rules relaxes a model requirement.

The register carried 26 approved uses at 30 June 2026: eleven of Class 1, nine of Class 2, six of Class 3 and none of Class 4. Six proposals were declined during the twelve months to that date. A use that is not re-approved within twelve months lapses and the system is switched off by the platform team without further notice.

Four classes of use

Every use is placed in a class according to what the system is permitted to do with its output, not according to how the system is built. Class 1 covers drafting, summarising, translating, searching and code assistance, where a named person reads, edits and issues everything produced. Class 2 covers pattern detection, anomaly flagging, screening prioritisation and research signal generation.

A Class 2 output is an input to a person and is never shown to a client as a conclusion. Class 3 covers systems that put a recommendation and the evidence for it in front of a named approver, who may accept, amend or reject it and whose decision is recorded with reasons. Class 3 is permitted only where the approver has the time, the information and the authority to decide differently.

- Class 1, assistive: every output is read, edited and issued by a named person. Eleven uses.
- Class 2, analytical: the output is an input to human analysis and never a client-facing conclusion. Nine uses.
- Class 3, decision support: a recommendation with its evidence goes to a named approver who records the decision. Six uses.

- Class 4, autonomous action: prohibited without an amendment to the Charter. No uses.

The prohibition that holds the rest together

Class 4 is any system that would take an action affecting a client, an employee, a counterparty, a position or a regulatory filing without a human decision in between. It is prohibited. The Council cannot grant an exception, nor can the Executive Committee, nor can the Board without first amending the Charter itself. That is a deliberately high bar and it has not been approached.

The effect is visible in the ordinary run of the business. A signal proposes orders and a named person authorises their transmission. A model produces a valuation and the Risk & Valuation Committee adopts the figure. A screening tool surfaces a match and a compliance officer reads it. No client, in seven years of operation, has had a decision taken about them by a machine.

Public generative services are blocked at the group boundary. Approved systems run inside the group environment or inside a contracted vendor environment bound by Charter terms: no retention of group prompts or outputs, no training on group data, no onward disclosure and audit rights for the group. A vendor that will not accept those terms is not approved, whatever its system can do.

How a use is stopped

An incident is any output that breaches a prohibition, any use outside its approved scope, any client complaint about a system-assisted decision and any failure of a Class 3 approver to record a decision with reasons. Incidents are reported to the Head of Technology and Ethics within one business day. Waiting to establish the full facts is not permitted, and a first report may be incomplete.

Two incidents were recorded in the twelve months to 30 June 2026. Both were scope breaches by Class 1 tools, in which staff used an approved drafting system for a purpose the approval did not cover. Both uses were suspended until the teams concerned had been retrained and the scope restated in the user documentation.

A system found in use to produce a prohibited inference is withdrawn on the day of discovery by the system owner or by the Head of Technology and Ethics, without waiting for a committee. The withdrawal is reported to the Council and to the Board. Restoring the use requires a fresh approval, not a resumption of the old one.

What the Board is told

The Ethics & Technology Council reviews the register, the incident log, the approval rates of every Class 3 system and every request for a human decision at each quarterly sitting in Singapore. A Class 3 system whose recommendations are almost never rejected is re-examined, because an approver who always agrees is not exercising a judgement.

The Council reports to the Board once a year with the full register, the proposals it declined and its own assessment of whether the seven principles were kept. Internal Audit examines the register and a sample of approvals on a two-year cycle. The Head of Technology and Ethics presents the report, and the independent director with a data-governance background reads it before the Board does.

Commitments: Governing artificial intelligence

- No learning or generative system runs on group work before it is registered and approved.
- No system takes an action affecting a client, an employee, a counterparty, a position or a filing.
- Every approval is for a stated scope and lapses after twelve months unless renewed.
- Every incident reaches the Head of Technology and Ethics within one business day.

- The full register goes to the Board once a year, including the proposals that were refused.

Accountability for automated work

Every automated estimate the group makes has a name attached to it. The 71 models in the inventory have 71 named owners, each an individual. Every approved use of a learning system has an accountable executive who sits on the Executive Committee. When automated work goes wrong, the consequences reach the people who owned it, including their remuneration.

No orphan systems

A model without a named owner cannot be entered in the inventory, and a model that is not in the inventory cannot be used for any decision. The rule is absolute and it is enforced at the platform level: the systems on which models run refuse to schedule a job for an unregistered model. There is no route by which an unowned system reaches production work.

The owner answers for the model being fit for its purpose, for its inputs being current and lawfully held, for its outputs being used only inside the approved scope, and for stopping it when it fails. Ownership sits with a person because a team can dissolve, a vendor can be replaced and a committee can disagree with itself, while a name on a register cannot do any of those things.

When an owner leaves the group or changes role, a successor is named before the change takes effect, or the model is suspended on the day of the change. Eleven ownership transfers were recorded in the twelve months to 30 June 2026 and none required a suspension. The inventory carries the date of every transfer and the name of the person who approved it.

The accountability record

For each model the group keeps a single record that answers four questions: who built it, who validated it, who owns it and who authorised its use. For each approved use of a learning system the record adds a fifth: which member of the Executive Committee answers to the Board for it. The five answers are names, with dates.

The accountable executive is expected to be able to explain what the system does and what it cannot do, in plain language, without the team that built it in the room. That expectation is not a formality. Two proposals in the twelve months to 30 June 2026 were deferred by the Council because the executive putting them forward could not describe the limits of the system.

The record can be produced to a client whose mandate the system touched, to the supervisor of a licensed entity that relies on it and to Internal Audit. It is retained for ten years after the model is retired, alongside the validation reports and the monitoring history, in the George Town archive.

When the work is wrong

A model that fails does so quietly, which is why the group treats detection as an accountability question rather than a technical one. The owner reports a threshold breach the day it is seen. The validation unit grades the finding. A grade one finding stops the model until it is closed. A grade two finding permits use under a written restriction with a closure date not more than 90 days out.

Where a client was affected, the group tells the client what happened, what it has done and what the client can do, in the next report or sooner if the matter will not wait. Where a valuation was wrong, the Risk & Valuation Committee restates it. Where an allocation was wrong, the position is corrected at the group cost rather than the client cost.

Consequences that reach people

Model ownership and Charter compliance are conduct and control objectives for every owner, for the heads of the Quantitative Strategies and Information & Data divisions and for the accountable executives. The Remuneration Committee applies malus where a breach was known and not escalated, where a model was used outside its scope, or where an incident was reported late.

That link is the reason the register is accurate. A control that costs nobody anything is a description rather than a control. Failure to submit a new system for review before committing resources to it is a breach of the Code of Conduct and is treated as one, whatever the quality of the work that followed.

Owners attest once a year that their inventory entries are complete and current, that the inputs remain lawfully held and that the model has been used only inside its scope. Attestation was 100 per cent of owners at 30 June 2026. Internal Audit tests a sample of attestations against the monitoring records on a two-year cycle.

Commitments: Accountability for automated work

- No model runs without a named individual owner recorded before its first use.
- A successor is named before an owner leaves, or the model is suspended that day.
- The accountable executive can explain the system without the builders present.
- Clients affected by a model failure are told what happened and what was done.
- Remuneration carries malus where a breach was known and not escalated.

How models are overseen

The group ran 71 models at 30 June 2026, tiered by the cost of being wrong rather than by their complexity. Twelve were Tier 1, 26 were Tier 2 and 33 were Tier 3. A validation unit of six people in Zurich and Singapore, with no reporting line to any developer or user, decides whether each may be used and for what.

Tiering by the cost of error

A model is any method that applies statistical, mathematical or algorithmic technique to input data to produce an estimate or a ranking the group may act on. A spreadsheet that does that is a model. The tier follows the consequence of the output, so a simple calculation that sets a reported valuation ranks above an elaborate system that produces internal management information.

Tier 1 outputs determine an order, a valuation reported to a client or an investor, a capital or liquidity figure reported to the Board, or a screening outcome for a client. The twelve Tier 1 models include the four signal families in Quantitative Strategies, the unlisted valuation approaches used by Private Equity, the group risk aggregation engine and the liquidity forecast.

An owner proposes the tier and the Chief Risk Officer sets it. Where a use changes so that an output would carry further, the tier is raised before the new use begins, not after. Using a model outside its validated scope makes it a new model, and a new model cannot be used until the extension has been validated.

- Tier 1: the output determines an order, a reported valuation, a capital figure or a screening outcome.
- Tier 2: the output informs a human decision without determining it.
- Tier 3: the output is internal management information with no external or position effect.

Validation by people with nothing at stake

The validation unit sits inside the Risk function, reports to the Chief Risk Officer and has no reporting line to any model developer or any desk that uses a model. Six people work in it, three in Zurich and three in Singapore. Their work is not charged to the desks whose models they examine, because a validator paid by

the business is a consultant rather than a check.

Validation tests the soundness of the approach, the quality and origin of the inputs, the implementation against the documented design, the behaviour of the model at the limits of its inputs and the accuracy of its outputs against realised outcomes. Tier 1 models are validated before first use and every year after. Tier 2 models are validated every two years. Tier 3 models are self-assessed annually and validated every three years.

Findings are graded. The 2025 cycle raised 44 findings across the inventory. Nine were open at 30 June 2026, all grade two or grade three, and none older than 90 days. Internal Audit reviews the work of the validation unit on a three-year cycle, so that the validators are themselves validated by someone else.

Monitoring and the stop rule

Every Tier 1 model is monitored against thresholds fixed at validation, covering prediction error, input data quality, output stability and, for trading signals, realised return against expectation. The owner and the validation unit review the results monthly. The Risk & Valuation Committee reviews them quarterly. Tier 2 models are monitored quarterly. A breach of a threshold is a finding and follows the finding timetable.

Any Tier 1 model can be stopped within fifteen minutes by its owner, by the Chief Risk Officer or by the head of the desk that uses it, with no further approval required. Each has a documented fallback, either an earlier validated version or a manual process, and the fallback is tested at every revalidation. A stopped model restarts only on the written approval of the Chief Risk Officer.

The output is not the decision

A model output is information. It becomes a decision when a person decides to act on it, and the group has placed that person in every chain. For Tier 1 trading signals, a named person authorises the transmission of the orders the signal proposes. For unlisted positions, the Risk & Valuation Committee adopts the figure the approach produces.

For screening, a compliance officer reads every match the model surfaces and records the conclusion. For deal screening, the ranking orders the queue and does not shorten it. This is why the group can say without qualification that it takes no automated decision about any client, counterparty or employee.

What is reported and to whom

The Chief Risk Officer reports to the Risk & Valuation Committee each quarter on the inventory, the validation calendar, open findings, threshold breaches and any model stopped in the period. The Board receives an annual model risk report and approves the list of Tier 1 models itself, because the Tier 1 list is a statement of where the group has accepted the cost of automated error.

Client reports name the figures and decisions that a model assisted, in the same fixed form used for learning systems. The assistance statement applies to a fifteen-year-old valuation method as readily as to a system trained last quarter, since a client asking whether a machine was involved is asking about the decision rather than about the technique.

Commitments: How models are overseen

- Every model is tiered by the cost of its being wrong, and the Board approves the Tier 1 list.
- Validation is done by people with no reporting line to the developer or the user.
- A model carrying an open grade one finding does not run until that finding is closed.
- Any Tier 1 model can be stopped in fifteen minutes and has a tested fallback.
- A person authorises every action taken on a model output.

The person who signs

Oversight means a signature that could have gone the other way. The group fixes the point at which a person must sign, the conditions that make the signature meaningful and the measurements that show whether it is being exercised. Approvers amended or rejected 18 per cent of decision-support recommendations in the twelve months to 30 June 2026.

Where a signature is required

The Charter requires a human decision between any system and any action affecting a client, an employee, a counterparty, a position or a regulatory filing. That principle becomes operational through a schedule of sign-off points maintained by the Chief Operating Officer and approved by the Executive Committee each year.

The schedule states who signs, at what threshold, on what evidence and within what time. It covers automated and manual work alike, because the risk of an unconsidered decision does not depend on whether a machine proposed it. Sixty-one sign-off points were in force at 30 June 2026, of which nineteen sit in the Quantitative Strategies division.

- A dealer named on the desk rota releases each basket of orders, and the release is timestamped.
- Orders above US\$2 million require a second authoriser recorded before release.
- Unlisted valuations above US\$10 million are adopted by the Risk & Valuation Committee.
- Every screening match is read and concluded by a compliance officer, whatever the score.
- Every client-facing recommendation is approved by a person who can be asked to explain it.
- Every decision about an employee is taken by a person, with no system assistance at all.

What makes a signature real

A signature means something when four conditions hold. The approver can see the evidence. The approver has enough time to read it. The approver understands the work well enough to disagree with it. The approver has the authority to decide differently without seeking permission from the person whose proposal it is.

The third and fourth conditions are the ones that fail quietly. An approver who does not understand a recommendation will approve it, and an approver whose rejections are routinely overturned will stop rejecting. The group tests both by reading the record, not by asking people how they feel about the work they do.

Time is protected by a cap. A decision-support queue is limited to 25 recommendations for one approver in a day. Where the volume exceeds the cap, the queue waits or a second approver is rostered. A queue never clears itself by lowering what an approval requires.

Measuring whether oversight happens

Across the six decision-support uses, approvers amended or rejected 18 per cent of recommendations in the twelve months to 30 June 2026. The rate ranged from 9 per cent on a screening queue to 31 per cent on a deal-ranking tool. The Ethics & Technology Council reads the range as closely as it reads the average.

A use whose recommendations are amended in fewer than 5 per cent of cases across two consecutive quarters is re-examined. One was re-examined in the year to 30 June 2026. The examination found the tool proposing the only lawful outcome available, and its scope was narrowed so that the approver was deciding something.

The platform records the time between a recommendation being opened and a decision being entered. Decisions taken faster than the evidence can be read are flagged for the accountable executive. Forty-one

were flagged in the year: thirty-four had a sound explanation and seven led to retraining on the desk concerned.

Disagreement is part of the design

An approver who rejects a recommendation records the reason, and the reason enters the monitoring record for the system. Rejections are evidence about the model rather than friction in a process. A pattern of rejections for the same cause is a finding for the validation unit and is graded like any other.

A rejection may be reversed only by the accountable executive, in writing, with reasons that go to the Ethics & Technology Council. Two reversals were recorded in the twelve months to 30 June 2026. Neither concerned a client-facing decision, and both were reported at the following quarterly sitting.

No approver is measured on the rate at which they agree with a system, and no remuneration in the group depends on the throughput of a queue. The Remuneration Committee tests this when it reviews the conduct objectives of the divisions that run decision-support systems.

When a person is not enough

Some work cannot be overseen at the volume proposed. The group then declines to automate it rather than accept a signature that means nothing. The test is stated plainly in the review template: if no approver could read the evidence in the time the process allows, the process is not approved.

One of the two proposals declined by the Ethics office in the twelve months to 30 June 2026 was an automated limit-setting tool for Commercial Finance. The volume implied 300 approvals in a day for one credit officer. The division returned with a higher referral threshold and a narrower scope, and that proposal was cleared with conditions.

Commitments: The person who signs

- No recommendation becomes a decision until a named person records that decision with reasons.
- Approvers hold the evidence, the time, the competence and the authority to decide differently.
- A decision-support queue waits when volume exceeds the daily cap; the standard does not fall.
- A use whose recommendations are almost never amended is re-examined by the Council.
- Only the accountable executive may reverse a rejection, in writing and on the record.
- Where no approver could do the work properly, the automation is not adopted.

Colophon

IGUAKO Capital is a work of institutional fiction by the Iguako Institute for Applied Unreality (iguako.tech). This document is part of that work. It creates no rights or obligations, describes no licensed entity and is not an offer, a contract, a client record or advice.