

BayesPharma Innovator

Quick User Guide

A five-page guide to moving from a governed discovery campaign to an evidence-backed candidate decision.

The one rule to remember

BayesPharma connects predictions, calculations, retrieved evidence, experiments and human decisions on one evidence spine. A computed result is never automatically presented as a measured result.

Start in four steps

- 01** Sign in and open Discovery Campaign. Choose an existing programme or create a focused campaign with target, indication, objective and success criteria.
- 02** Review target readiness and approve only the compute you intend to run. Generation, docking and downstream qualification remain separate evidence stages.
- 03** Use Programme Command Center to see blockers, candidates, evidence gates and the next governed action. Keep the selected **programme_id** in the URL.
- 04** Promote only reviewable leads. In Candidate OS, inspect inherited evidence, contradictions and the v3 nomination matrix before recording a human decision.

Recommended first session

- Open a known example or a small programme; do not start with an unbounded screen.
- Confirm the target structure, binding site, reference ligand and scientific assumptions.
- Save the rationale for every approval, hold, rejection or promotion.

Online guide: bayespharma.com/pharma/user-guide.html

The six-stage lifecycle

The top workflow navigation groups related workspaces into six stages. Use the same programme and candidate identifiers as you move between them.

Stage	Primary tools	Use it for
1. Campaign	Discovery Campaign	Define target and objective; generate and qualify molecules.
2. Programme	Command Center; Programme Ops	See state, blockers, next experiments and governed programme actions.
3. Experiments	Study Execution; Partner Handoff; Experimental Review	Plan assays, exchange CRO/lab packages, QC results and approve measured evidence.
4. Evidence	Candidate Passport	Inspect identity, provenance, versions, limitations and append-only history.
5. Optimize	Lead Optimization	Run Design-Make-Test-Learn cycles; compare series and balance exploit/explore/control.
6. Candidate	Candidate OS; Decision Center	Characterize the lead, review v3 gates and record nominate/conditional/hold/kill.

Identifiers that keep the record connected

Identifier	Meaning
programme_id	The governed therapeutic programme. Required across Command, experiments, optimization and candidate decisions.
candidate_id	The same molecule/candidate identity from Discovery through Candidate OS. Do not create a second authoritative record.
series_id	A related chemical series used for SAR and DMTL prioritization.
evidence_id + record_hash	Append-only lineage identifiers used to verify the evidence record and its history.

Navigation tip

If a page opens without context, use its programme or candidate picker. Avoid manually editing IDs unless you are troubleshooting a known record.

Who approves what?

- The software may rank, calculate, flag or recommend.
- A qualified human approves compute, experimental QC, promotion and nomination.
- A CRO or laboratory result is not nomination-grade until its method, source, units and QC state are traceable.

Campaign to programme command

1. Discovery Campaign

- Define the target, indication, objective and acceptance criteria before generation.
- Check structure readiness: chain, binding site, co-factors/metals, reference ligand, protonation and docking box.
- Run the intended funnel: generation and chemistry filters, fast/accurate/deep docking, pose checks, ADMET, novelty and synthesis assessment.
- Treat docking score as a prioritization signal, not proof of binding or activity.

Recommended funnel

Use broad generation, then progressively more expensive qualification. Preserve all failures and exclusions so the final shortlist is reproducible.

2. Programme Command Center

- Select a programme from the account picker. The page writes the selection to the URL.
- Read the current stage, candidate board, SAR series, DMTL cycles and negative-data intelligence.
- Resolve deterministic blockers before following a model recommendation.
- Use next-best experiments as transparent prioritization; human approval is still required.

3. Evidence and provenance

Class	Definition	Examples	Correct use
COMPUTED	Generated by a declared algorithm/model	Docking, PBPK, ADMET model output	Decision support within stated applicability
RETRIEVED	Imported from a traceable external source	PDB, ChEMBL, PubChem, literature	Context; verify relevance and licensing
MEASURED	Observed in a traceable assay/study	CRO or internal laboratory result	May satisfy a gate after QC and method review
REVIEWED	Human interpretation or governed decision	Target review, promotion, nomination	Records accountability; not a measurement

Never relabel evidence

Do not copy a predicted or simulated value into a MEASURED field. If the source is a QA fixture, label it QA_ONLY / SIMULATED and disable production export.

Experiments to candidate decision

Experimental workflow

- Create the study plan with assay, protocol, species/matrix, units, controls and acceptance criteria.
- Use Experimental Partner Handoff for internal lab, CRO, CDMO, contract lab or academic partner packages.
- On return, validate file identity, units, QC flags, protocol deviations and candidate mapping.
- Approve measured evidence only after review; retain rejected and contradictory results.

Lead Optimization

- Compare molecules within a series; inspect potency, selectivity, developability and synthetic feasibility together.
- Use DMTL decisions explicitly: exploit known SAR, explore uncertainty, and retain a control.
- Promote a lead only when identity and evidence lineage remain intact.

Candidate OS

- Inspect inherited Discovery evidence, evidence fitness and unresolved contradictions.
- Register traceable DMPK, PK/TK, exposure-response, PBPK, safety and CMC evidence.
- Review every required v3 nomination gate. Screening-only or context-dependent evidence cannot silently satisfy a required gate.

Decision	When to use
Nominate	All required gates satisfied; evidence current; contradictions resolved; human rationale recorded.
Conditional	A bounded decision with explicit conditions, unresolved risks and follow-up work.
Hold	Evidence is insufficient, stale, contradictory or an experiment is still pending.
Kill	The candidate no longer meets the programme hypothesis or risk/benefit profile; preserve the negative learning.

Final nomination is human
BayesPharma can determine whether the configured evidence contract is satisfied. It cannot replace the accountable scientific and development review that records the final decision.

Reading results and getting unstuck

Result-reading checklist

Question	Check
What is the claim?	Binding priority, predicted property, measured activity, synthesis route or human decision?
What is the evidence class?	COMPUTED, RETRIEVED, MEASURED or REVIEWED?
What produced it?	Method, model/software version, parameters, structure/protocol and source.
How reliable is it here?	Applicability domain, uncertainty, limitations, contradictions and staleness.
What changes next?	A decision, a confirmatory experiment, another model run, or no action?

Common problems

Symptom	Action
Page asks for programme_id	Use the Programme Command picker or reopen the page from the programme workspace.
No promoted leads	Complete lead promotion with the same candidate identity; Candidate OS intentionally excludes unpromoted records.
Nominate is disabled	Open the v3 matrix; resolve missing/stale gates and contradictions. Do not bypass the gate.
Evidence is missing	Check the Candidate Passport and append-only ledger; confirm candidate_id and source record.
Prediction conflicts with assay	Record the contradiction. Measured evidence does not automatically erase computed evidence; investigate method and context.
A service is unavailable	Preserve the failed state, retry only when appropriate, and never replace a missing result with a default value.

Useful links

[Discovery Campaign](#) | [Programme Command Center](#) | [Candidate OS](#)

Before using a result externally

Export the evidence package, confirm source and version metadata, verify the candidate identity, and obtain the required qualified human approval. Product outputs are decision support, not clinical or regulatory authorization.

Contact: contact@bayespharma.com