CERNAL
COMPILER-LIKE ENGINE FOR RNA LOGIC

CERNAL Risk Analysis

How We Assessed Risk

Likelihood and Impact are each scored on a 3-level scale. Priority is based on the same 3×3 matrix used in the Wiki: Low = score 1-2, Moderate = score 3-4, High = score 6-9.

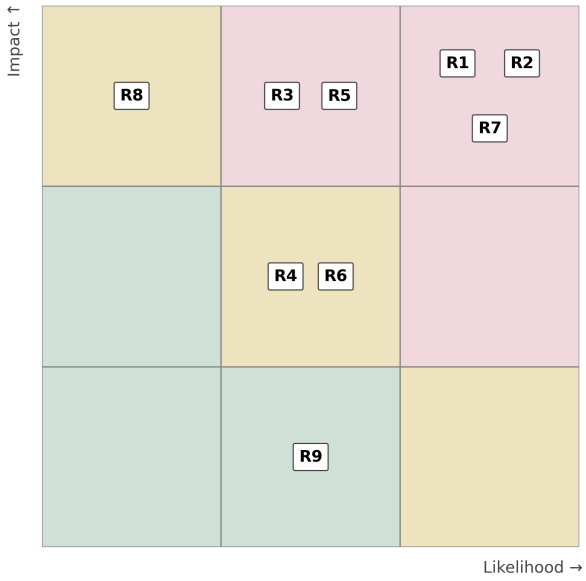

Figure 1. CERNAL Probability × Impact matrix for risks R1-R9.

Detailed Risk Register

#	RISK	TYPE	WHY IT MATTERS	LIKELIHOOD	IMPACT	PRIORITY	DETAILED MITIGATION
R1	Prediction–experiment mismatch	Technical	CERNAL's MVP depends on computational predictions being useful enough to guide experimental work. A design that looks strong in silico may still behave differently inside the cell.	3	3	High	Compare predicted and measured behavior for every design, retain failed designs as learning data, and recalibrate the engine using structured partner results.
R2	Commercial demand & repeat purchasing unvalidated	Market	The MVP will be paid, so initial willingness to pay will be tested directly. However, one paid MVP does not prove recurring demand, repeat purchasing, or a scalable commercial relationship.	3	3	High	Track conversion into the paid MVP, repeat use, follow-on projects, price acceptance, and whether the same buyer returns with additional design needs.
R3	Customers may be unable or unwilling to share RNA-seq data	Data & IP	RNA-seq is a core input to the CERNAL workflow. Companies may treat transcriptomic datasets as confidential or strategically sensitive, which could limit deployment and the learning loop.	2	3	High	Minimize the data required, use clear confidentiality and data-use terms, and define explicitly what information may be retained or reused.
R4	Concierge model may not scale	Operational	The MVP intentionally relies on hands-on intake, manual case execution, reporting and feedback coordination. This is useful for learning, but too much team effort per project could limit growth.	2	2	Moderate	Standardize repeated steps, track team hours by stage, and progressively automate the workflow as patterns become stable.
R5	Cross-cell-type generalization	Technical	The MVP starts with HEK293T and expands one mammalian cell type at a time. Gate behavior and circuit performance may change across CHO, T/NK or other systems.	2	3	High	Validate the gate library one cell type at a time and define a clear supported operating range before expanding to additional systems.

CERNAL
COMPILER-LIKE ENGINE FOR RNA LOGIC

#	RISK	TYPE	WHY IT MATTERS	LIKELIHOOD	IMPACT	PRIORITY	DETAILED MITIGATION
R6	Partner feedback may be inconsistent	Operational	The partner builds and tests the circuit, and CERNAL relies on returned experimental results to improve the engine. Inconsistent assays, missing controls or delayed data could reduce comparability.	2	2	Moderate	Use a fixed Feedback Protocol with agreed controls, replicates, metrics, data formats, timing and interim checkpoints.
R7	Value vs. current workflow unproven	Market / Product	Even if CERNAL works technically, customers need a meaningful improvement over manual design. The core business case depends on reducing time, expert effort or design iterations.	3	3	High	Benchmark CERNAL against the customer's existing workflow and quantify the difference in time, effort, iterations and experimental outcome.
R8	Competitive / IP advantage may narrow	Competition / IP	Existing design platforms or therapeutic-circuit companies could extend toward state-first design, while CERNAL's IP protection is still developing.	1	3	Moderate	Continue FTO and claim-scope review, monitor competitor capabilities, and strengthen defensibility through validated design-outcome data and the integrated state-to-circuit workflow.
R9	No self-serve interface in the MVP	Product / Operational	The MVP is intentionally delivered through a concierge model rather than a polished autonomous interface. This may reduce convenience, but it does not block early validation because the CERNAL team operates the engine.	2	1	Low	Provide guided concierge support during the MVP and develop a self-serve interface only after the workflow and customer requirements are validated.

How Each Risk Will Be Tested

#	WHAT WE NEED TO LEARN	EVIDENCE / METRIC
R1	Can CERNAL's predicted behavior reliably guide experimental design?	Track prediction accuracy, OFF/ON behavior, specificity, leakage and recurring failure modes across design-partner cycles.
R2	Will customers return for additional paid projects after the initial MVP?	Measure paid-MVP conversion, repeat purchase rate, follow-on projects, time to second project, expansion signals and buyer/budget-owner continuity.
R3	Can partners share enough RNA-seq data under acceptable terms?	Test data-sharing requirements during partner intake and document whether private/local processing or restricted retention is needed.
R4	How much manual effort is required per design-partner cycle?	Measure scientist-engineer hours, customer-success hours, cycle time and the proportion of each project that can be standardized or automated.
R5	How much performance transfers when CERNAL expands to a new cell type?	Compare predicted and measured circuit behavior across each newly supported cell type and characterize cell-type-specific failure modes.
R6	Can partner-generated results be made sufficiently comparable?	Track completeness, comparability and turnaround time of the returned datasets across partners.
R7	Does CERNAL create a measurable advantage over manual design?	Record baseline workflow metrics before each paid MVP and compare them with the CERNAL-assisted cycle.
R8	Does CERNAL retain a meaningful technical and IP differentiation?	Review relevant patents and competitor products periodically and track whether CERNAL's differentiators remain unique and valuable.
R9	Which interactions truly need a user interface rather than concierge support?	Track where partners need repeated support and use those interactions to define future interface and automation requirements.