DECISION MEMO · OCTOBER 2026 · PLANNING ESTIMATE, NOT A QUOTE

# Validate MAMMAL before treating it as an oncology engine

Approve a four-week computational gate. Place the wet-lab order only if two claims survive a held-out test. A pass means a ranking enriches molecules that bind or kill in an assay the model did not see. It does not mean the model predicts which patients respond.

**$170,430**
Vendor cash to authorize, both tracks plus 15% contingency. A higher quote is a new decision.

**13 weeks**
From protocol lock to an unblinded yes or no. Organoids are not in this clock.

**$18,100**
All-in cost of stopping at week 4. Only $2,000 of that is vendor cash.

## WHAT THE CAPITAL TESTS

### Drug response

**Kill at week 4** unless rank correlation on unseen drugs beats the best simple baseline by at least 0.05, and the unseen-drug plus unseen-lineage split does not fall below that baseline. The published correlations near 0.92 used random pair splits, so drugs and cell lines leaked across the test.[1]

**Wet-lab pass,** only after that kill is survived: 12 lines, 20 held-out compounds, 9 doses, 3 replicates. The predicted-sensitive set must be more active in at least 8 of 12 lines, with median within-line Spearman of at least 0.40. One spectacular line does not pass.

### Antibody or nanobody

**Kill at week 4** unless held-out classification reaches AUROC 0.70 with 200 antigen-specific labels. The pretrained checkpoint is chance-level on HER2 (AUROC 0.49). Zero-shot design is already falsified.[1]

**Wet-lab pass:** 128 sequences against a length-matched control arm. Confirmed, non-polyreactive binders at least 3× the control, at least four such binders, and at least one that competes with a clinical antibody or binds the antigen-positive cell and not the negative line.

## CLOCK AND CASH

| Week 1 | Weeks 2–4 | Week 4 | Week 5 | Week 13 |
|---|---|---|---|---|
| Reproduce the released checkpoint. Stop if it misses the published error. | Cold-start splits and the antibody learning curve. Collect quotes. Buy nothing. | Written stop or go. This is the planned exit. | On a pass only: purchase orders commit about $85,500. | Unblind. Apply the rule locked in week 1. |

Internal loaded time on the same path is about $50,400. Count it if you are comparing with hiring. Do not add it to the purchase order. Gate 3, three established organoid models, is a separate $46,000 and about ten further weeks. It stays unauthorized until the plate has passed.

## WHAT THIS APPROVAL DOES NOT BUY

Patient response, survival, or a biopsy-to-drug recommendation. Toxicity or developability beyond a polyreactivity panel. T-cell receptor assays. Fresh tumor tissue or an ethics review. A claim that a foundation model will move clinical success rates. The free NCI-60 screen is not a substitute: it does not run this blinded, scaffold-held-out design.[2]

**If quotes exceed $170,430,** cut in this order: three biological replicates to two, then 12 lines to 8, then 128 sequences to 96 by shrinking plate controls. Do not cut the drug-blind test, the matched antibody control arm, or the polyreactivity panel. Do not spend contingency to absorb a quote that was high on day one.


1. Shoshan et al., MAMMAL, arXiv:2410.22367 (2025). https://arxiv.org/abs/2410.22367
2. NCI-60 Human Tumor Cell Line Screen, Developmental Therapeutics Program. https://dctd.cancer.gov/drug-discovery-development/assays/high-throughput-screening-services/nci60


Cash figures are planning estimates from the accompanying budget, October 2026, not vendor quotes. Replace the screen price and the per-clone price with the median of three quotes before the week-5 order.