

# A Shared-Switch Model for Protein XOR

## Architecture and limits

Vasanth Sreeram | Computational working paper | 17 October 2026

Revision: shared-switch-2026-10-17-v1. Not peer reviewed. No protein gate was built or tested.

### Abstract

This working paper examines an equilibrium model in which two inputs A and B each bind a shared switch S, while A:B pairing suppresses switch occupancy when both inputs are present. The architectural distinction from Chen et al.'s XNOR gate is the assumed output: occupancy of one shared switch replaces association of an output protein pair. Mutual input sequestration is already present in that prior work. At the reference concentrations, the modeled switch is 15.57% occupied with one input and 3.65% with both. The resulting truth table meets a specified relative criterion, but depends on balanced inputs, excluded ternary complexes and an unimplemented physical readout. We report conditional affinity-noise sensitivity and a preliminary, separately evaluated mutation-effect predictor. The study does not establish a new XOR principle, a complete three-protein device, a calibrated design-success probability or a performance advantage over existing gates.

## 1. Scope and architectural distinction

This study asks whether a small, specified binding network can generate an XOR-like equilibrium truth table and what engineering requirements that network leaves unresolved. A and B are input proteins; S is a shared switch. A:S and B:S are both counted as active output, while A:B is an inactive input-pairing complex. With one input present, the switch can be occupied. With both inputs present at similar amounts, sufficiently tight A:B binding suppresses free input and switch occupancy.

The direct comparison is Chen et al.'s XNOR gate, Fig. 2F [1]. Their inputs already sequester each other. With neither input present, an output protein pair associates; one input disrupts it, and both inputs bind each other so the output pair can associate again. Here the assumed observable changes to occupancy of one shared switch. The architectural distinction is the arrangement and readout; mutual sequestration is not introduced by this work.

| Input state | Chen XNOR: output-pair association | Shared-switch model           |
|-------------|------------------------------------|-------------------------------|
| Neither     | ON                                 | OFF                           |
| A only      | OFF                                | ON                            |
| B only      | OFF                                | ON                            |
| Both        | ON                                 | LOW under suitable conditions |


These logical complements do not establish interchangeable physical reporters. A real readout must recognise either A:S or B:S and remain low for free S and A:B. Physically inverting a demonstrated XNOR reporter is itself an engineering operation.

Protein XOR is already established in receptor-state theory [2] and experimentally through proteolysis/coiled-coil circuits [3]. Competitive-dimerization theory includes a six-monomer XOR with a single output dimer [4]. This model instead counts a sum of two complexes. Therefore,

three core components versus six monomers is not an established saving in complete device complexity. No claim is made of the smallest protein XOR, a new XOR principle, or superior cost, speed or robustness.

## 2. Equilibrium model and readout

The model contains three components and six chemical species: free A, B and S, plus A:S, B:S and A:B. It permits three reversible reactions:



The dissociation constants obey:

$$[A][S] = Kd1[A:S] \quad [B][S] = Kd2[B:S] \quad [A][B] = Kd3[A:B]$$

Component totals are conserved:

$$\begin{aligned} A_{\text{total}} &= [A] + [A:S] + [A:B] \\ B_{\text{total}} &= [B] + [B:S] + [A:B] \\ S_{\text{total}} &= [S] + [A:S] + [B:S] \end{aligned}$$

The plotted observable is  $f = ([A:S] + [B:S]) / S_{\text{total}}$ . Exclusion of ternary complexes, homodimers, larger assemblies and nonspecific binding is a physical assumption. In particular, the earlier design scaffold describes non-overlapping binding surfaces, which do not by themselves establish the exclusion imposed by these equations. A future sequence-level implementation must reconcile that difference.

An earlier reporter proposal fuses LgBiT to S and supplies free SmBiT. The code multiplies occupied switch by the constant  $\text{SmBiT} / (Kd_{\text{NanoBiT}} + \text{SmBiT})$ . At fixed SmBiT, this is linear in occupancy. It does not engineer an input-dependent conformational gate, establish a background level or model reporter back-action. The NanoBiT literature supplies an assay concept [6], not validation of this proposed switch. The readout proposal also adds a protein species beyond the three-component binding core.

For the principal relative score, each simulated truth table defines  $T = 0.30$  times its own mean one-input signal. Both ON states must exceed  $T$  and both OFF states must remain below  $T$ . The threshold is recomputed for each affinity-noise draw. This is an analysis criterion, not a calibrated assay detection threshold. The original solver demonstration used different thresholds, so its scores are reported separately.

## 3. Reference results and concentration sensitivity

When present, each input is at 1  $\mu\text{M}$ ; switch concentration is 0.5  $\mu\text{M}$ . The nominal affinities are  $Kd1 = Kd2 = 5 \mu\text{M}$  and  $Kd3 = 10 \text{ nM}$ . These are stipulated parameters rather than predictions for completed protein sequences.

| Quantity                               | Modeled value |
|----------------------------------------|---------------|
| Occupancy, neither input               | 0%            |
| Occupancy, A only or B only            | 15.57%        |
| Occupancy, both inputs                 | 3.65%         |
| Both-input / mean one-input occupancy  | 23.43%        |
| Direct one-input / both-input contrast | 4.27-fold     |

The table passes the 30% relative rule. It fails the original demonstration's 10% OFF rule. The historical mean-ON/mean-OFF metric is 8.54-fold because the two OFF states are averaged and

one is zero; it is not a measured reporter dynamic range.




Figure 1. Equilibrium calculations with the parameters above. Left: four input states. Right: B titration with A fixed at  $1 \mu\text{M}$ . Curves are model calculations, not measured kinetics or biological experiments. At  $B = 2 \mu\text{M}$ , residual B restores an occupancy close to the A-only level. The lowest output occurs near matched inputs.

Tightening the pair  $K_d$  from 10 nM to 1 nM to 0.1 nM reduces the leak ratio from 23.43% to 7.89% to 2.54% under equal inputs and equal switch affinities. When most input is paired and switch occupancy is low, free input is approximately proportional to the square root of pair  $K_d$  times total input. This approximation has a restricted regime and is not a new general binding law.

Switch affinity creates a brightness-leakage tradeoff. At switch  $K_d$  100 nM and pair  $K_d$  1 nM, a single input occupies about 85% of S but both-input leak is about 44% of the one-input level. At switch  $K_d$   $5 \mu\text{M}$  with the same pair, leak falls to about 8%, while one-input occupancy falls to about 16%. No experimental noise floor establishes whether either output is detectable.

## 4. Conditional affinity-noise sensitivity

The illustrative noise model specifies independent, zero-mean Gaussian errors in binding energy, with standard deviation  $\sigma_G$  in kcal/mol. For each interaction,  $\log(K_d/K_{d\_nominal})$  has standard deviation  $\sigma_G/RT$ , using  $RT = 0.592$  kcal/mol. The geometric standard deviation of the multiplicative affinity error is  $\exp(\sigma_G/RT)$ . The parameter is an assumed standard deviation, not a measured predictor MAE.

| Nominal pair $K_d$ | Pass at $\sigma_G = 0.50$ | Pass at $\sigma_G = 1.10$ |
|--------------------|---------------------------|---------------------------|
| 10 nM              | 63%                       | 33%                       |
| 1 nM               | 94%                       | 56%                       |
| 0.1 nM             | 94%                       | 63%                       |

These rounded rates use 6,000 draws, seed 42, and the relative threshold described above. They are conditional sensitivity results, not probabilities that synthesized proteins will work.

The old scoring pipeline substituted mutation-prediction MAE for Gaussian standard deviation. For zero-mean Gaussian errors,  $\sigma_G = \text{MAE} \times \sqrt{\pi/2}$ . A separate 100,000-draw check, seed 42, calibrated to an illustrative MAE of 1.10 gives 26.33%, 44.03% and 52.75% passing at the same three pair  $K_d$  settings. Monte Carlo sampling uncertainty is approximately  $\pm 0.3$  percentage points at 95%; this excludes uncertainty in the assumed distribution and model.

Other assumptions materially affect the result. With perfectly correlated switch-affinity errors and independent pair-affinity error, the calibrated-MAE reference rate becomes 50.71%. Under the original independent  $\sigma_G = 1.10$  setting, declaring a fixed threshold from the nominal gate reduces its pass rate from 32.95% to 18.87%. These scenarios illustrate sensitivity; none is a fitted model of actual gate-design errors. Tighter pairing alone does not repair unequal ON branches, imbalanced inputs or an unresolved reporter.

## 5. Numerical checks and a separate predictor benchmark

The browser uses a one-dimensional positive bisection solver with algebraic elimination of the other free concentrations. The reference table was independently reproduced by reducing the symmetric equations to one free-input unknown and using a scalar root solve in micromolar units. The maximum reference conservation residual was  $3.3\text{e-}16$   $\mu\text{M}$ . A separate log-variable positive least-squares solution agreed on 150 asymmetric parameter sets to within  $8.8\text{e-}14$  times the largest total concentration. These checks support the numerical calculation; they do not validate the physical reaction network.

The original repository root finder ignores convergence status and clamps negative concentrations. Repeating the archived 3,000-point extreme-parameter sweep found 11 disagreements above 0.001 times switch total against the conserving solver. Its outputs are not universally reliable. Reference results in this note use the checked calculation.

A separate MLP benchmark uses SKEMPI 2.0 mutation-affinity records [5]. Two saved baseline runs achieved best validation MAEs of 1.1025 and 1.1100 kcal/mol, compared with 1.1951 for predicting the training median. These records support preliminary benchmark results, not a new affinity-prediction method or an autonomous improvement campaign.

The preprocessing retains 6,723 rows: 6,072 training and 651 validation, grouped by PDB ID. Related binding sites still cross the split: 480 validation records have an explicit SKEMPI hold-out PDB link to a training structure, spanning 25 of 49 validation PDB IDs. In addition, 186 retained observations contain affinity bounds treated as exact values, including 16 validation rows. Repeats, experimental temperatures and the magnitude filter need deliberate handling. Validation selects the best epoch; there is no untouched final test, and two unseeded runs do not establish variability.

A mutation effect converts a known starting affinity through  $Kd_{mut} = Kd_{wt} \times \exp(\Delta\Delta G/RT)$ . It does not supply an absolute affinity for a newly designed pair without a suitable baseline. The model has no protein-sequence or three-dimensional interface context, and no candidate's three affinities have been predicted and validated through the proposed workflow. Thus, the predictor benchmark and equilibrium sensitivity analysis remain separate exploratory pieces.

## 6. Weaknesses, comparison requirements and next experiments

The principal unresolved weaknesses are physical exclusion of omitted complexes; activation and background of the readout; sensitivity to matched input amounts; the tradeoff between desired occupancy and unwanted leak; and the absence of kinetic, expression, degradation and cellular resource models. A high-affinity pair is not evidence of fast response, reversibility or lack of kinetic trapping.

Chen et al. engineered cooperativity to reduce sensitivity to certain component imbalances [1]. No matched concentration-range comparison with their architecture has been performed. The present four-point truth table cannot support a superiority claim. Equally, experimental proteolysis-based XOR [3] and receptor-state theory [2] address different mechanisms; they are not interchangeable benchmarks solely on component count.

A fair comparison should require a physically specified readout and count all required components. It should use the same input concentration ranges, fixed detection thresholds, ON-state brightness requirements and calibrated affinity uncertainties. It should report background, undesired states and both ON branches separately. If response time is compared, association/dissociation rates and observation time must also be specified.

The immediate next step is to reconcile the intended binding surfaces and design an observable that responds to either bound-switch state. A repaired, family-aware mutation benchmark is relevant if pursuing sequence design, but a better MLP cannot repair an incorrect reaction network. Baseline affinities and all four input states would ultimately need measurement for a concrete implementation.

## 7. Conclusion

The shared-switch model produces an XOR-like truth table under stated equilibrium conditions. Its specific architectural distinction is the use of one shared occupied switch as the assumed output, compared with association of an output protein pair in Chen's XNOR gate. The input-pairing mechanism is already present in prior work. The evidence presently supports an explicit model and an account of its failure modes. A complete three-protein device, a calibrated design-success probability and an advantage over existing architectures remain unestablished.

## Reproducibility and provenance

The original repository snapshot is commit 4706bbfd180d0bf09d8ae6a08e90ad6194925354 of the author's private protein-comp repository. This revision preserves the baseline logs and the audit results under paper/evidence; the interactive model is available at [vas.sh](https://vas.sh). No new neural-network training or wet-lab experiment was performed for this revision. The note was revised with AI-assisted drafting and code checking; the author remains responsible for its claims. The software and data snapshots document what was run, not independent biological validation.

## References

1. Chen et al. (2020). De novo design of protein logic gates. *Science* 368, 78-84. <https://doi.org/10.1126/science.aay2790>
2. de Ronde, ten Wolde and Mugler (2012). Protein logic: a statistical mechanical study of signal integration at the single-molecule level. *Biophysical Journal* 103, 1097-1107. <https://doi.org/10.1016/j.bpj.2012.07.040>
3. Fink et al. (2019). Design of fast proteolysis-based signaling and logic circuits in mammalian cells. *Nature Chemical Biology* 15, 115-122. <https://doi.org/10.1038/s41589-018-0181-6>
4. Parres-Gold et al. (2025). Contextual computation by competitive protein dimerization networks. *Cell* 188, 1984-2002.e17. <https://doi.org/10.1016/j.cell.2025.01.036>
5. Jankauskaitė et al. (2019). SKEMPI 2.0: an updated benchmark of changes in protein-protein binding energy, kinetics and thermodynamics upon mutation. *Bioinformatics* 35, 462-469. <https://doi.org/10.1093/bioinformatics/bty635>
6. Dixon et al. (2016). NanoLuc complementation reporter optimized for accurate measurement of protein interactions in cells. *ACS Chemical Biology* 11, 400-408. <https://doi.org/10.1021/acscchembio.5b00753>
7. Parres-Gold et al. (2026), correction. Erratum for: Contextual computation by competitive protein dimerization networks. *Cell* 189, 2194. <https://doi.org/10.1016/j.cell.2026.03.015>