

The calc-block map

Every apartment acquisition underwrite is seven blocks. Each one has a quiet lever — an input that changes the answer without announcing itself. Read down the flow; the levers are what you check.

01	The rent roll In-place vs market rent · month to roll · renewal probability · free rent · vacancy days · releasing cost	QUIET LEVER One blank cell in a <i>named</i> row corrupts the aggregation across all six monthly engines. Outputs don't error — they move. In-place = market on a value-add deal silently zeroes the entire thesis.
02	Expenses + renovation Year-one operating costs · management fee · the capex plan	QUIET LEVER Sign convention: this model wants positive year-one dollars and handles the sign itself. On a value-add deal the capex plan <i>is</i> the thesis — check it carries what the story claims.
03	NOI aggregation Six monthly engines roll up to the annual cash flow	QUIET LEVER Every upstream error concentrates here, which is exactly why NOI always looks clean. Clean is not the same as right.
04	The pricing block Manual input · present value at a target return · cap on Year-1 NOI · replacement cost	QUIET LEVER One dropdown decides which of the four is live. A price typed into the manual cell while the method is set elsewhere is simply ignored — the single most common configuration mistake on this model.
05	The debt block Senior loan · fixed rate · optional interest-only · amortisation	QUIET LEVER Is DSCR proven at the events the loan cares about — post-development, full lease-up, refinance — or rebuilt as a side calculation whenever the lender asks? Name which constraint binds: DSCR, debt yield, LTV or LTC.
06	Property returns The deal unlevered, then levered	QUIET LEVER Cap-rate growth is often entered in <i>basis points</i> , not percent. "Grow it 0.25% a year" typed as 0.25 is a hundred-times error that still looks plausible.
07	The waterfall Preferred return · promote hurdles · the sponsor / LP split	QUIET LEVER Hurdles must ascend — the model does not check, and will produce a distribution anyway. Sponsor equity at 100% flips it to GP-only and every LP return reads N/A.

The property to respect: these models compute. They do not validate intent. Nothing on this page is exotic — each one is an input doing exactly what it was told, in a place nobody was looking.


The silent-failure checklist

Run this before you trust any AI-populated workbook. Seven habits, one per failure mode. None of them needs a budget line, and none of them slows the AI down — they are what lets you run it at deal speed.

- ☐ **01 Prove the recalc**
Never accept a number the file computed *before* the last edit. Code-driven AI edits don't recalculate formulas — the values it reads back can describe a workbook that no longer exists.
- ☐ **02 Check the switch, not the cell**
Confirm which input the live dropdown or toggle actually reads. A number entered into an inert cell is invisible in every output.
- ☐ **03 Read the achieved value, not the requested target**
Solvers stop close, not exact. If you asked for 8.50% and the file landed at 8.28%, say the miss out loud.
- ☐ **04 Validate completeness**
Every named row full; every model check cell reading zero. Refuse to report outputs until it does.
- ☐ **05 Say the units**
Basis points or percent, positive or negative — state it in the prompt, then read it back from the cell.
- ☐ **06 Trace three numbers end to end**
Summary → cell → source document, in one motion. A number that can't make that trip doesn't go in the memo.
- ☐ **07 Demand the source, not just the justification**
Every AI-entered input names its document or carries an explicit "guess" flag. A justified guess is still a guess.

Where the skill's job ends

Three boundaries — scope, not defect. The record doesn't persist beyond the chat session. Every session starts fresh and never learns your firm. The checking is advisory: nothing enforces it.



Excel for the sandbox — the engine for the numbers you sign. Everything on the checklist above, Excel relies on a person to remember. That is the honest boundary of the tool, and the part a governed underwriting engine adds.