`software, no person`

01 **Every basket**
Item-level EBT data from every store. Volume drives computing cost only.
`millions / day`

02 **Item price vs reference price**
Each line checked against the regional reference for its barcode and store class.
`every line`

03 **Store-month scores**
Excess price share, lines above range, basket plausibility, supply consistency.
`1 per store / month`

04 **Ranked by expected loss**
Highest expected loss first, not highest sales.
`probability x excess $`

05a

Automatic actions

`most stores above normal`

`warning letter` `payout delay` `volume cap`

Reversible, no investigator time. Scores keep updating; a store that keeps rising enters the queue.

05b

Top N stores to investigators

`N set by budget`

A person reviews a case only when its expected benefit tops the cost of one review.

↱ **Feedback to step 03.** Every closed case, confirmed or cleared, becomes a label. The scoring model retrains on both, so it learns which patterns are benign.

COMPUTING COST	HUMAN COST
grows with transaction volume, and stays small	grows with investigator headcount, which the budget fixes first