The agents return €5.89M a year after a loss of €2.24M from bad changes.

Of the scenarios tried, Triple the volume shows the largest net value: +€11.78M a year compared with Today. The table is not a recommendation. The conditions for each scenario are listed under the table.

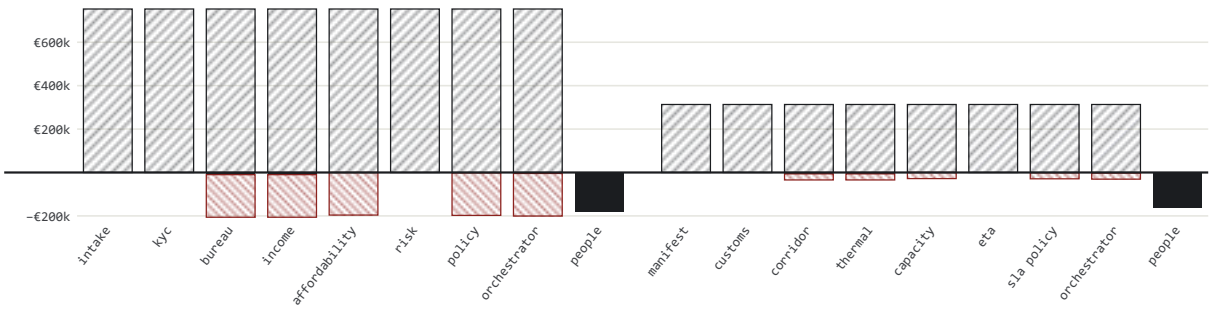
> 37 of 37 prices are illustrative examples, and none has been researched. Do not quote these figures until the prices are replaced. The measurements come from two synthetic Rootplant files with 400 applications and 250 shipments over about 12 hours. The figures are an estimate from a model and are not accounting figures.

The scenarios, per year

	Today	A: Catch it sooner	B: People check denials	C: Triple the volume
Work done in place of people	€8.53M	€8.53M	€8.53M	€25.59M
Extra business	€0	€0	€0	€0
Cost to run the agents	−€52k	−€52k	−€52k	−€155k
People	−€343k	−€343k	−€1.11M	−€1.03M
Loss from bad changes	−€2.24M	−€1.12M	−€673k	−€6.73M
Programme	€0	€0	€0	€0
Net	**€5.89M**	**€7.02M**	**€6.69M**	**€17.68M**
Against Today	–	+€1.12M	+€801k	+€11.78M
1 year in 10 is worse than this	€4.25M	€6.20M	€6.20M	€12.76M
1 year in 10 is better than this	€7.28M	€7.71M	€7.11M	€21.83M
Loss, 1 year in 20	−€4.46M	−€2.23M	−€1.35M	−€13.37M
Years below zero	0.1%	<0.1%	<0.1%	0.1%

- **A: Catch it sooner.** Bad changes are caught after 3.3 h.
- **B: People check denials.** People check 100% of denials.
- **C: Triple the volume.** Run volume is ×3 today's rate.

Who helps and who harms: Catch it sooner



Hatched areas show work done, solid areas show run cost, and red areas show the share of loss. The 10 agents share the loss because their own output changed together when a feed changed, measured on runs side by side. The lowest net values today are Thermal agent (€251k), Corridor agent (€251k) and Dispatch orchestrator (€254k).

Catching bad changes sooner has the largest single effect on Catch it sooner: +€1.12M.

Measured from the run files

Loan decisions: runs	400 in 11.95 h
Health status	UNHEALTHY
Change record_feed_version 3.3	70 runs
Extra denials compared with the old version	84% (72%–89%)
Still active at file end	6.6 h
Doubling alarm after	3.3 h
Cold-chain dispatch: runs	250 in 11.97 h
Health status	UNHEALTHY
Change record_feed_version 4.1	43 runs
Extra denials compared with the old version	81% (64%–88%)
Still active at file end	6.6 h
Doubling alarm after	3.2 h

Assumed

Bad changes a year per system	4	estimate
Share affecting all systems	25%	estimate
A person catches a wrong outcome	70%	estimate
Working days a year	250	example
A wrong outcome becomes a complaint	0.1	example
Cost of a complaint	150	example
Loan decisions: minutes a person would spend	30	example
Loan decisions: cost of a person's minute	1	example
Loan decisions: minutes to review a run	12	example
Loan decisions: wrong outcome, share of amount	0.03	example
Cold-chain dispatch: minutes a person would spend	20	example
Cold-chain dispatch: cost of a person's minute	1	example
Cold-chain dispatch: minutes to review a run	12	example
Cold-chain dispatch: cost of a wrong outcome	2000	example

How each figure is made

- Work done is runs a year × minutes a person would spend × the cost of that minute. The value is shared equally between agents, which is a fair share of the total when every step is needed for a decision. The sharing rule is a convention and never changes the total.
- Run cost is the tokens, tool calls and seconds of step time measured per agent, at the prices listed and scaled to one year.
- Loss is wrong outcomes per hour × hours until caught × bad changes a year × the cost of one wrong outcome and its complaint. The hourly number is measured against runs that stayed on the old version at the same time. For one incident at the observed exposure, this equals the loss on the Value sheet.
- The range comes from 3,000 simulated years with fixed random numbers. Incidents happen at random, and detection time and the cost of a wrong outcome vary. The shared part of the incidents affects every system at once.

Not in this number

- Fines: a fine depends on the regulator and the situation, so the sum excludes fines.
- Wrong outcomes when nothing changed: the run files cannot show them, so the normal error rate counts as zero.
- Wrong approvals: the loss here is only the rise in the share of failed runs after a change.
- Customers who leave after a complaint: the sum does not count them.
- Proof of cause for the fault path: the agents changed together while other runs stayed on the old version. Only the causal analysis can show a cause chain.
- A fault rate that changes over time: the model continues the observed rate unchanged after the 6.6 hours that were actually seen.