

THE 30-MINUTE PROTOCOL

0-5 | Reference

State the instrument, outcome, no-policy counterfactual, binding constraint, non-negotiable limits and failure condition. If these are unclear, stop and rewrite.

5-10 | Incidence

Map cash, real resources and balance sheets. Who receives, pays, administers or is excluded? Which cost moved off-budget, into the future or onto another institution?

10-15 | Adaptation

For households, firms, officials, financiers and foreign partners ask: what becomes rewarding; what will be avoided or gamed; what investment changes; who will preserve the policy?

15-20 | Propagation

Draw O1 direct, O2 adaptation, O3 propagation and O4 inherited state. Trace households, firms, government, external account, finance, institutions/trust and natural/infrastructure assets.

20-25 | Uncertainty

Use low, central and high conditional probabilities. Separate expected system value from critical tail exposure. Shared nodes require joint scenarios, not independent probabilities.

25-30 | Learning and exit

Choose Approve, Pilot, Sequence, Redesign, Hedge or Pause. Name the owner, data, review date, trigger, automatic response and stop rule.

DANGEROUS BLIND SPOTS

Blind spot	What it hides
Headline fixation	Output is not outcome.
Policy-chain neglect	Stopping before response and feedback.
Static actors	Assuming behaviour will not change.
Fiscal illusion	Missing tax expenditure, guarantees, maintenance and utility losses.
Average-case blindness	Ignoring correlated shocks and concentration.
Present/electoral bias	Ignoring costs outside this Budget or term.
Capacity optimism	New duties displace existing work.
Goodhart's trap	The target becomes a gamed substitute.
External-sovereignty illusion	Foreign rules are assumed fixed.
Sunk-cost escalation	Past spending blocks a rational stop.

FIVE QUESTIONS TO ASK ALOUD

1. What will people do because this policy exists?
2. Where will the cost migrate?
3. Through which shared node will the effect return?
4. Which external actor can rewrite the payoff?
5. What evidence would make us change course?

Two screens - never one average

Expected System Value compares probable effects. Critical Tail Exposure uses the upper plausible probability, severity, irreversibility, coupling and weak control so catastrophic harm cannot disappear inside an average.

FEEDBACK-LOOP TYPOLOGY

Loop	Pattern	Example
Accumulation	surplus -> capability -> surplus	Mauritius exports
Cost spiral	falling scale -> higher unit cost -> exit	banana shipping
Fiscal ratchet	windfall -> sticky commitments	volatile revenue
Credibility	weak controls -> external action	visa access
Congestion	policy surge -> scarce capacity -> delay	hotel projects
Scarcity	growth -> input cost -> slower growth	food/wages
Rebound	lower cost -> more use / shifted peak	energy demand
Learning/trust	delivery -> compliance -> finance	water service

SYSTEMIC ADJUSTMENT OR CRITICAL TAIL?

Treat an effect as baseline when at least three are true: a physical/accounting/legal constraint makes it hard to avoid; an actor has incentive and capability; it appears in a relevant reference class; its sign survives optimistic, central and adverse scenarios.

Probability	Consequence	Response
High lower bound	Moderate	Budget in baseline
High lower bound	Severe	Redesign
Low centre; real upper	Moderate	Monitor / insure
Low centre; real upper	Catastrophic / irreversible	Prevent / stage / hedge / stop

Small-State Coupling check

Circle every path that returns to the same sovereign balance sheet, utility, bank, port, airport, tourism market, scarce technical team, foreign-exchange pool or external gatekeeper. Shared nodes make risks correlated.

RIPPLE-4 EFFECT CARD

Decision		Responsible minister		Implementing owner	
Evidence vintage		Review date		Decision deadline	
Intended outcome		No-policy counterfactual		Binding constraint	
Non-negotiable safeguards					

Order	Effect and causal mechanism	Actor / ledger	Sign	Delay	Probability L/C/H	Magnitude range	Evidence confidence
O1 direct							
O2 adaptation							
O3 propagation							
O4 inherited state							

Critical return paths / shared nodes	
Reinforcing and balancing loops	
Thresholds, lock-in and distribution	
External gatekeepers	

Leading indicator	Data source / frequency	Trigger threshold	Automatic response	Accountable owner

Expected System Value		Critical Tail Exposure	
Posture	Approve / Pilot / Sequence / Redesign / Hedge / Pause	Next review	
Reason and pathway that changed the decision			
Revision or stop rule			