MESAS 2026 // MODELLING AND SIMULATION FOR AUTONOMOUS SYSTEMS

Analysis of Unreliability Factors in the Neutralization of Small Unmanned Aerial Systems

A probabilistic C-UAS kill-chain model with correlated stress and saturation



Daniel Korec[1] · Václav Minařík[1] · Norbert Świętochowski[2]

[1] University of Defence, Brno, Czech Republic
[2] Military University of Land Forces, Wrocław, Poland

PRAGUE // PVA EXPO
23 OCT 2026 // 09:30

[01] THE PROBLEM

Cheap drones, unreliable neutralization

A

Small RCS

Hard to detect and to tell apart from birds and clutter

B

Low and agile

Low-altitude manoeuvre complicates tracking and fire control

C

Cost asymmetry

The defensive round often costs more than the drone

The weak point is the whole process, from first detection to confirmed kill.

“Systemic” unreliability is widely claimed, but rarely stated in a form that can be computed, tested or simulated.

[02] THIS PAPER

Three contributions

01	**Explicit model**	The C-UAS kill chain as a fully specified probabilistic model. Reproducible and re-parameterizable.
02	**Correlated Monte Carlo**	A shared environmental and EM stress factor shows what single-number stage efficiencies miss.
03	**Saturation and autonomy**	An N-target extension links C2 latency to engagement capacity and leakage.

SCOPE: reliability structure of the engagement. Not a technology survey; no ranking of commercial systems.

[03] MODEL

The C-UAS kill chain: six phases

B = BASELINE: 1 sensor · 1 effector · human C2 **F** = FUSION: sensor fusion · integrated C2 · interceptors

PH-1	PH-2	PH-3	PH-4	PH-5 · WEAKEST	PH-6
Detect	**Identify**	**Track**	**C2 decide**	**Engage**	**Assess**
B 0.85	B 0.75	B 0.80	B 0.85	B 0.62	B 0.86
F 0.95	**F 0.88**	**F 0.90**	**F 0.93**	**F 0.68**	**F 0.90**

Output: the probability of a confirmed neutralization of one target.

BDA is part of the chain: a drone may be destroyed without the defender knowing it.

[04] MODEL

From stage efficiencies to correlated failure

USUAL PRACTICE · EQ. (1)

$$P = p1 \times p2 \times p3 \times p4 \times p5 \times p6$$

Exact by the chain rule, if every p is truly conditional on the earlier phases.

In practice each stage efficiency is estimated on its own, under average conditions. Multiplying them quietly assumes that phases fail **independently**.

OUR MODEL · EQS. (2)-(3)

$$p(s) = p \times (1 - \lambda \cdot s)$$

s: one shared stress value per engagement, from 0 to 1, Beta(2, 6), mean 0.25. Severity of weather and EM conditions.

λ: how strongly each phase degrades, 0.20–0.30. Because s is common, bad conditions hit all phases together.

- **Weather, terrain masking and jamming degrade sensing, decision and guidance at once.**

[05] MODEL

Parameters: a transparent baseline

Phase	Baseline	Fusion	Stress λ
Detection	0.85	0.95	0.30
Identification	0.75	0.88	0.28
Track maintenance	0.80	0.90	0.24
C2 decision	0.85	0.93	0.20
Engagement	**0.62**	**0.68**	**0.26**
Assessment (BDA)	0.86	0.90	0.30

Expert estimates

Based on reported performance ranges. Not measurements of a fielded system.

Parameter uncertainty

Each p drawn from a Beta distribution, κ = 60, about ±4–6 pp per phase.

Stress

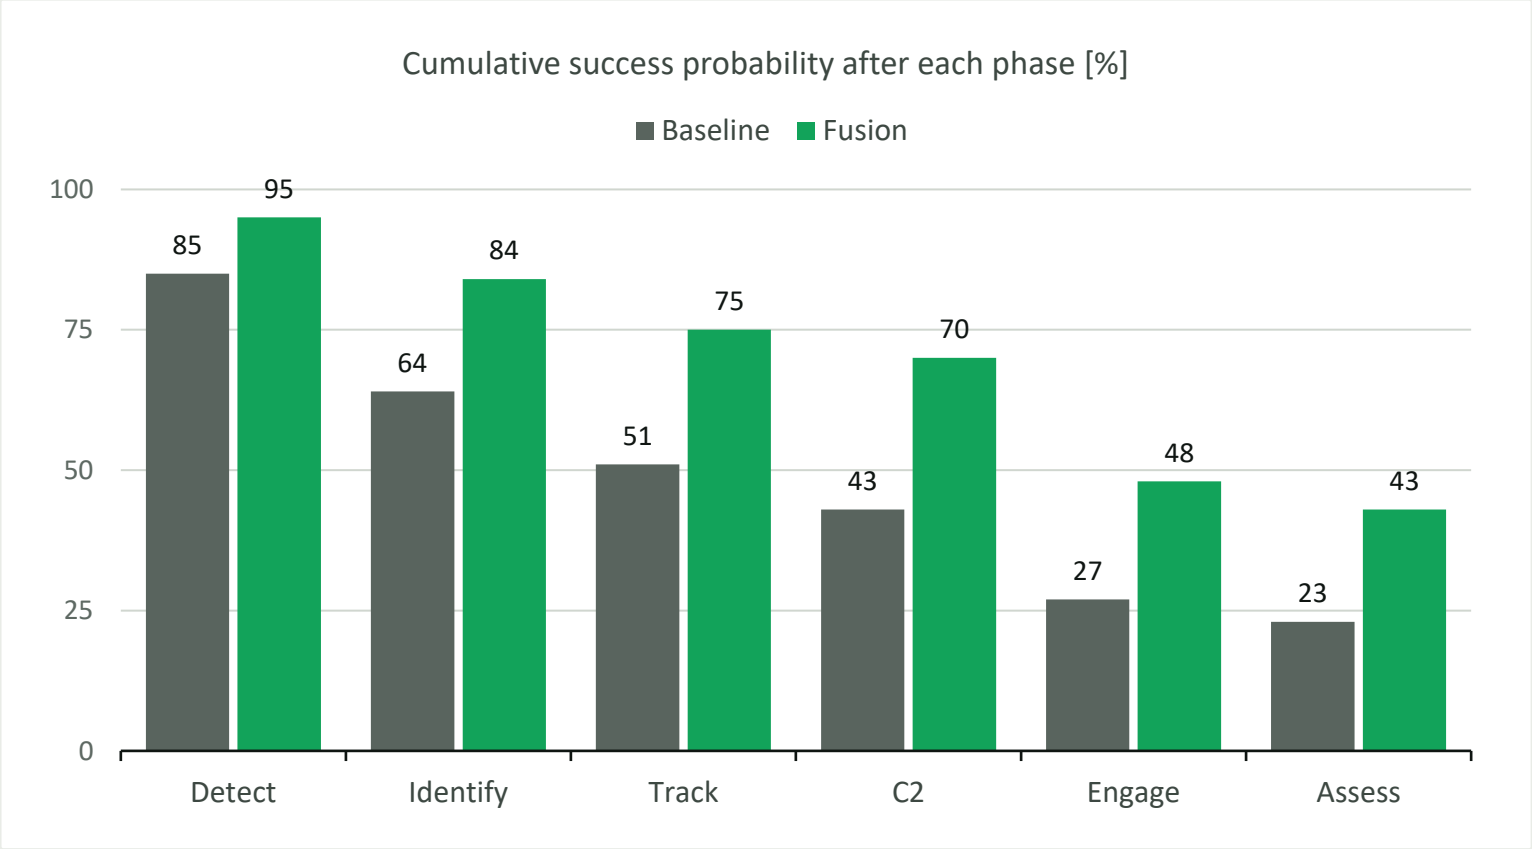
s ~ Beta(2, 6), mean 0.25: most engagements in moderate conditions.

Simulation

100,000 replications per configuration. Intervals = 5th–95th percentile.

[06] RESULT 1 · NOMINAL VALUES

No phase looks alarming. The product does.



23 %

Baseline

43 %

Fusion

Neutralization of one target at nominal values, no stress. Fusion still loses more than half along the chain.

[07] RESULT 2 · MAIN RESULT

The price of ignoring shared stress

	NOMINAL		SHARED STRESS
Baseline	23 %	→	16 %
	17–30 %		9–24 %
Fusion	43 %	→	29 %
	34–52 %		17–42 %

≈ 1.5×

Overestimate when stage efficiencies ignore shared degradation: +47 % relative, in this scenario.

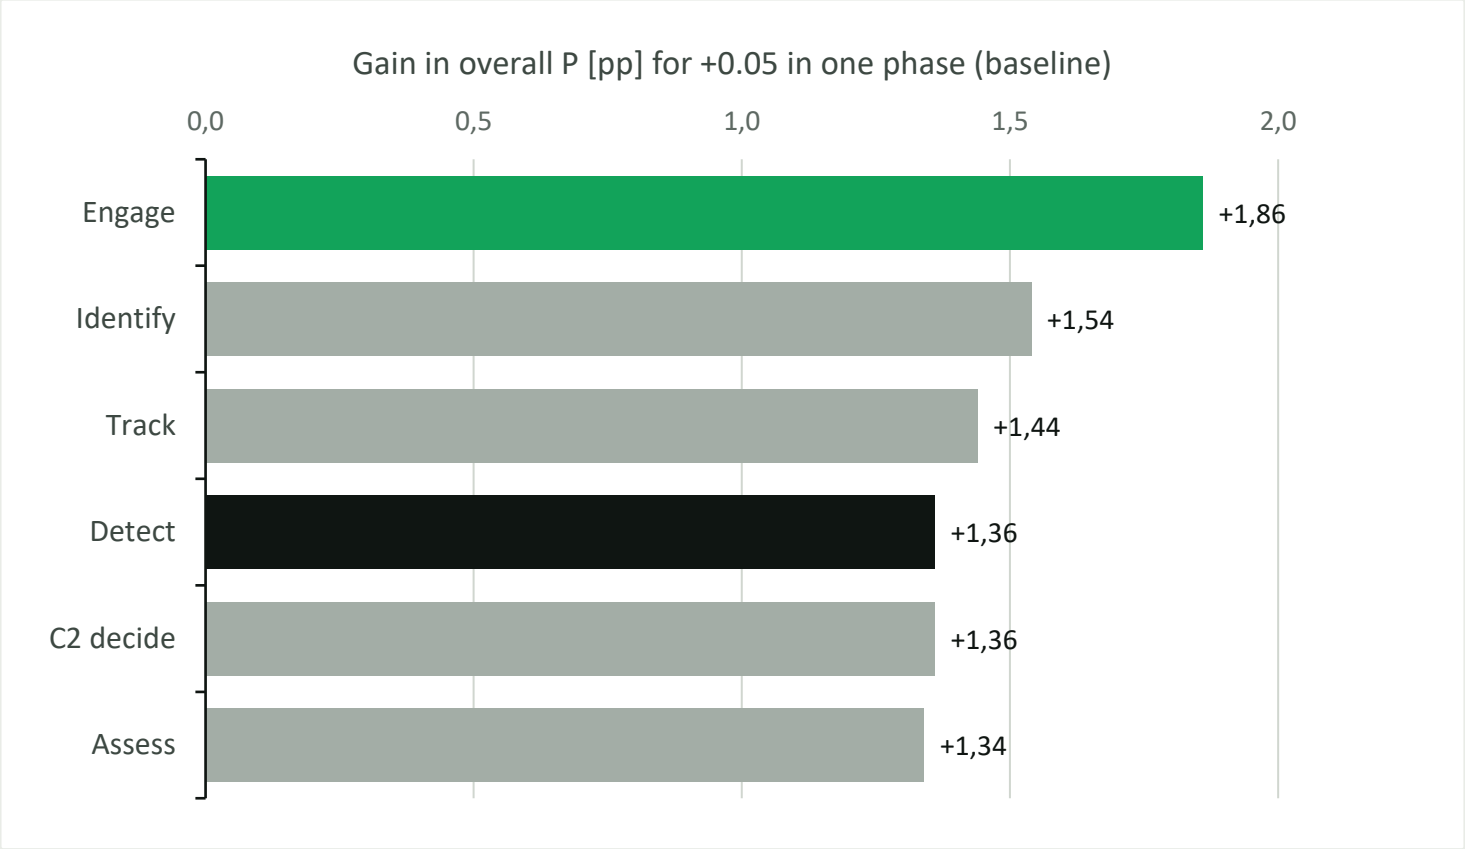
Shared vs. phase-independent stress: same mean, heavier lower tail. Fusion 5th percentile: 17 % instead of 21 %.

- **Shared causes of failure lower the mean and widen the bad-day tail. Plan for the lower tail.**

Ranges: 90 % simulation intervals (5th–95th percentile)

[08] RESULT 3 · SENSITIVITY

The weakest link leads, but only slightly



Engagement first

The lowest phase probability gives the largest marginal gain.

Detection alone is not enough

Most researched, yet among the smallest returns.

No silver bullet

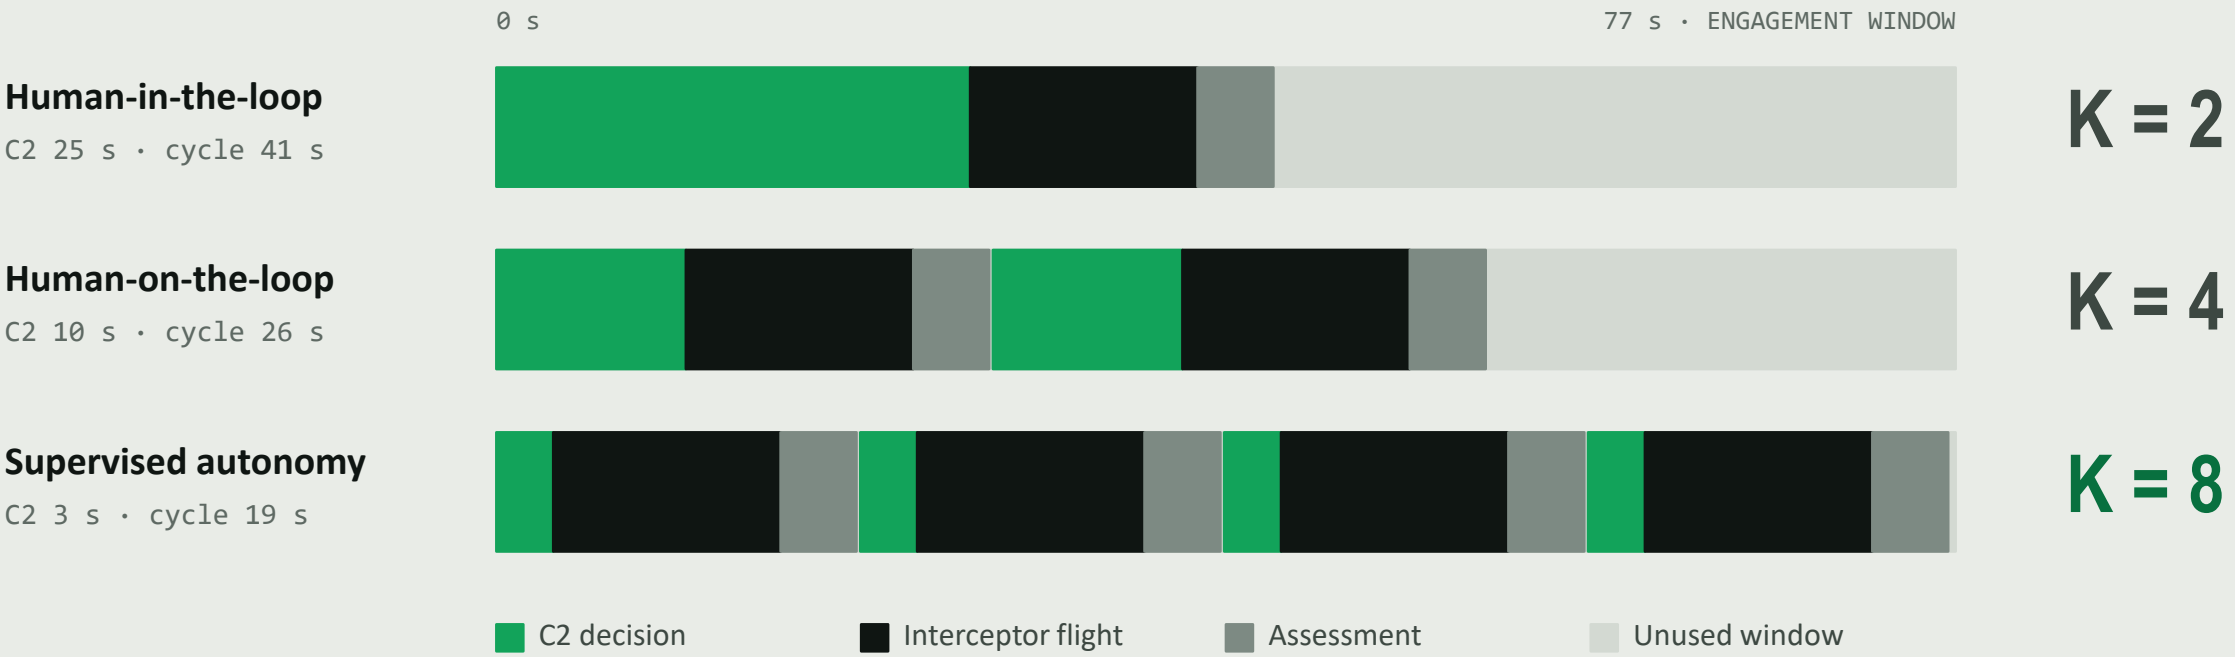
Best vs. worst phase: under 1.4×. Gains must span the whole chain.

pp = percentage points

[09] SATURATION · N TARGETS

Under saturation, time becomes the constraint

Window ≈ 77 s (2 500 m → 200 m at 30 m/s) · 2 channels · magazine 8 · cycle = C2 + flight 12 s + assessment 4 s

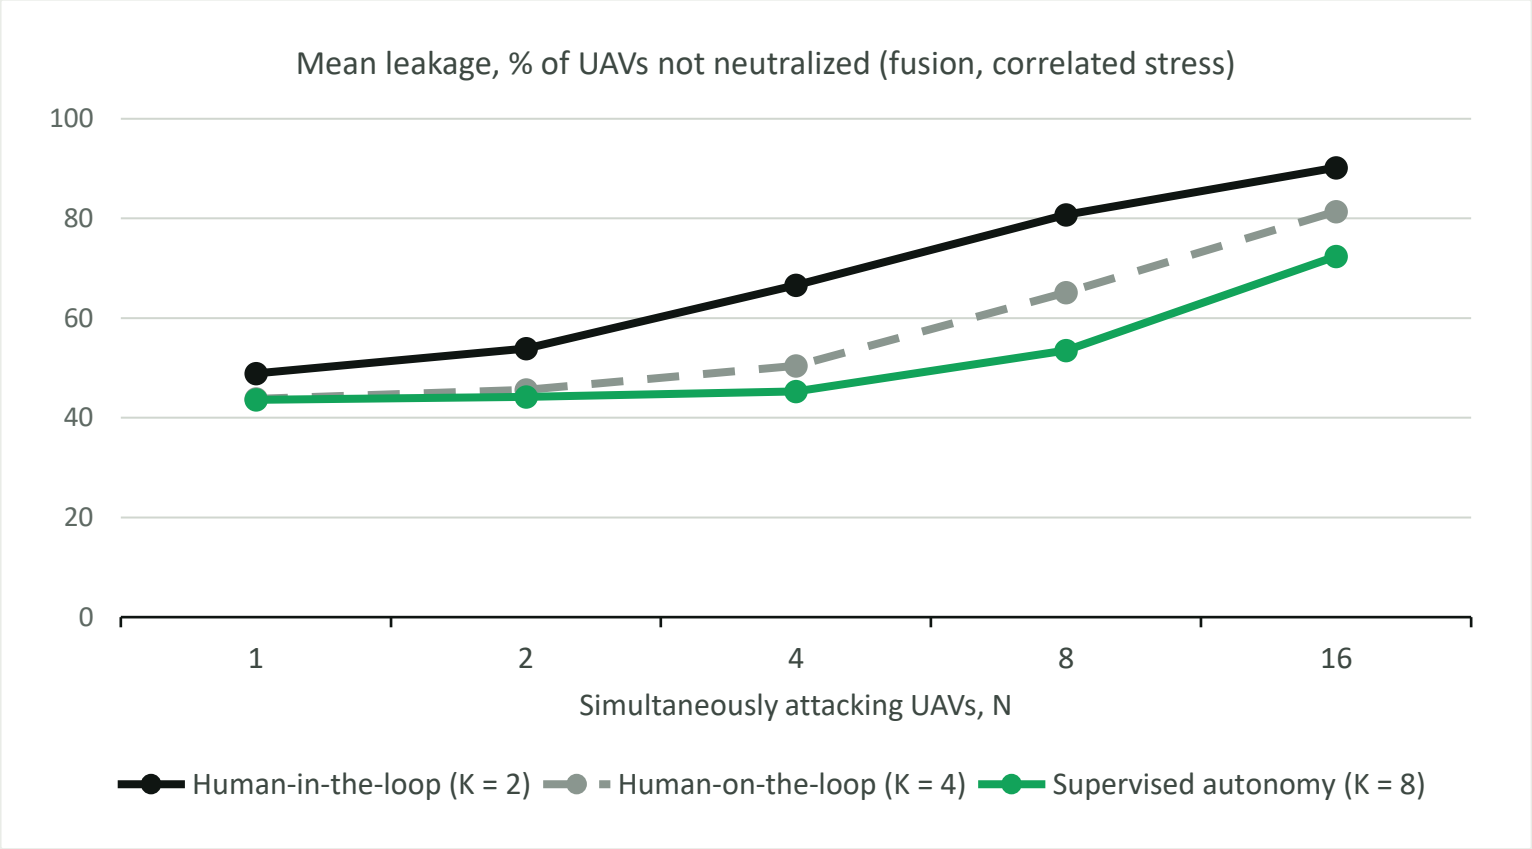


Same sensors, interceptors and magazine: 4× the engagement capacity.

One channel shown; K counts both

[10] SATURATION · LEAKAGE

Autonomy delays the collapse



Whole raid neutralized

SUPERVISED AUTONOMY, BEST CASE

N = 1	56 %
N = 4	≈ 9 %
N = 8	≈ 0 %

Beyond a few targets, magazine depth and the number of channels set the ceiling.

4 000 replications per point

[11] DISCUSSION

What this means for C-UAS capability

01

An optimistic bound

Assessments built on multiplied stage efficiencies overstate performance. Design for the lower tail.

02

The constraint moves

Single intruder: engagement effectiveness, a physical limit. Raid: decision latency, a time limit.

03

Autonomy = a shorter cycle

The same gain can come from procedure: pre-delegated engagement authority in a defined area. ROE matter as much as algorithms.

04

Cost asymmetry

Attackers add targets more cheaply than defenders add engagement capacity. Layered soft- and hard-kill defence stays the rational answer.

[12] LIMITATIONS

Limitations and next steps

Limitation	Next step
Phase probabilities are expert estimates	Validate and re-parameterize with instrumented trial data
One shared stress factor with linear effects	Several correlated stressors: weather, EW, terrain
Raid model: simultaneous arrival, first come first served, no prioritization, no evasion	Heterogeneous, manoeuvring, coordinated raids. Current leakage figures are optimistic.
Autonomy modelled only as shorter decision time	Add misclassification and fratricide risk to the trade-off

[13] CONCLUSIONS

Test C-UAS against raids, not single intruders

1 Specify and test systems against raid sizes, not only against single intruders.

2 Treat decision latency as a primary design parameter, next to detection range and kill probability.

3 Read assessments built on multiplied stage efficiencies as optimistic bounds.

Thank you. Questions?

npor. Ing. Daniel Korec · daniel.korec@unob.cz

Supported by the Czech Ministry of Defence (LANDOPS, DZRO-FVL22-LANDOPS) and MEYS (SV25-FVL-K107-KOR). Thanks to the 13th Artillery Regiment “Jaselský”.