


COGNITIVE-TERRAIN

PRACTITIONER'S MANUAL • RUN IT • BUILD IT • KNOW WHY



RUN IT WITH A GROUP


Session plans from a lunch break to a multi-organisation exercise, and how to lead the debrief.

BUILD YOUR OWN CAMPAIGN

The campaign and technology editors, balance checks and sharing.

WHY IT WORKS THIS WAY

Design notes: what each mechanic models, and how it was arrived at.



FOR WHOEVER RUNS THE ROOM

CONTENTS

03	1.0 INTRODUCTION <i>Who this is for, and what the game is and is not</i>
04	2.0 RUNNING A SESSION <i>Formats, timings, the room, privacy</i>
06	3.0 CAMPAIGNS <i>All eight campaigns, difficulty and choosing one</i>
09	4.0 THE WHITE CELL <i>A facilitator who steers a two-player game</i>
13	5.0 THE DEBRIEF <i>The three verdicts, the replay and questions</i>
14	6.0 THE MODEL IN NUMBERS <i>Every formula and every rule, with worked examples</i>
18	7.0 BUILDING CAMPAIGNS <i>The editor, crises, technologies, balance</i>
21	8.0 EXERCISES ACROSS ORGANISATIONS <i>Turn codes without a shared network</i>
22	9.0 DESIGNER'S NOTES <i>Why each part of the game works the way it does</i>
27	10.0 SOURCES AND LIMITS <i>What comes from research, what has been checked</i>
29	APPENDICES <i>A Checklist · B Debrief questions · C Reference tables</i>

HOW TO READ THIS MANUAL	
4.5	A section of this manual.
CAPITALS	Buttons and labels you will see on screen, such as END TURN.
DESIGN NOTE	Why a part of the game works the way it does, and where it comes from.
EXAMPLE	A rule worked through with real numbers.

FIRST TIME RUNNING IT?

Play **Typhoon Ines** on **Standard** once yourself, end to end, before you run it for anyone else. It takes about forty minutes. Then read Appendix A.

1.0 INTRODUCTION

1.1 WHO THIS IS FOR

Instructors, exercise designers and facilitators who run Cognitive Terrain for other people, build campaigns for their own audience, or need to explain why the game models cognitive warfare the way it does.

Players need only the **Field Manual**, which is also the guide inside the game. This manual assumes you have read it or played a campaign.

1.2 WHAT THE GAME TRIES TO GET ACROSS

- The target of cognitive warfare is the quality of the other side's decisions, not territory or attention.
- Reach is not effect: a viral post on the wrong area changes nothing that matters.
- Attack and defence happen at once, and open and closed societies are vulnerable in different ways.

1.3 THREE DOCUMENTS

The **Field Manual** teaches players the game on screen, with a worked turn; it is also the Guide (?) inside the game. The **Quick Reference** is a two-sided sheet to keep beside the keyboard. This manual covers what a facilitator needs on top: sessions, the campaigns, the White Cell, the debrief, every rule and formula (6.0), the editor, and the reasoning and research behind each part (9.0, 10.0).

1.4 WHAT THE GAME IS, AND IS NOT

Cognitive Terrain is an **independent hobby project**, made for learning and discussion. It is free.

It is **not doctrine**, and not an official, approved or endorsed product of any organisation. The authors of the research it draws on have not reviewed it.

It has **not been independently assured or validated**. Its numbers are game rules, not measurements or predictions of real-world effects. Section 10.0 says which rules come from research, which are the designer's assumptions, and what has and has not been checked.

Every campaign, town, outlet, person and event is **fiction**.

1.5 USING IT WITH A GROUP

Treat the game as a prompt for discussion, not as evidence. The most useful moments are usually the gaps: a side that wins on points but loses the debrief's verdicts, or a viral move that changed nothing [5.0].

2.0 RUNNING A SESSION

The game runs offline on any computer with a browser, so the format is set by your time and your room, not by the technology. Timings are suggestions; adjust them to your group.

2.1 FORMATS

SOLO

ONE PLAYER • 30–45 MIN

- 1 Send the link or the HTML file with the Field Manual.
- 2 Ask for **Typhoon Ines** on **Recruit** with the tutorial first.
- 3 Ask them to bring a screenshot of the debrief and one turning point from the replay.

CLASSROOM

GROUPS OF 2–3 • 50–90 MIN

- 1 10 min: the decision loop on the projector (the Field Manual's opening pages).
- 2 30–60 min: each group plays one campaign on one screen, one driver, others advising.
- 3 15–20 min: debrief as a room (5.0).

PAIRED SEMINAR

BLUE vs RED • 60–90 MIN

- 1 Pair students; one plays Blue, one Red, on two computers in the same room.
- 2 They swap turn codes by chat (8.0). Sixteen weeks takes about an hour.
- 3 Debrief each pair, then compare how Red was played across pairs.

DISTRIBUTED EXERCISE

SEVERAL ORGANISATIONS • DAYS

- 1 Build or choose a campaign and send the file to every cell first (7.1). Add a White Cell to steer it (4.0).
- 2 Cells play one week per day or per session, codes by email.
- 3 Collect debriefs and replays for a joint hot-wash.

2.2 DECIDING BEFOREHAND

WHAT TO DECIDE BEFOREHAND

Plain or military terms. Plain words suit mixed audiences; military terms (OODA, DQ, Seed Channels) suit staff courses. Players switch in the top bar.

Campaign and difficulty. See 3.0.

Week in review on or off. On for learning; off for a faster, harder exercise.

KEEP IT MOVING

Set Red's turn speed to **FAST** once players know what Red's moves look like, and result pop-ups to **KEY ONLY**. Both are in Settings and change nothing in the rules.

EXAMPLE

A 90-minute classroom session. Groups of two or three, one screen each, Typhoon Ines on Recruit.



Keep the last twenty minutes for the debrief (5.0), even if the campaign is unfinished.

2.3 SETTING UP THE ROOM

2.3 GETTING IT ONTO MACHINES

OPTION	BEST WHEN
Play in the browser (playtest page)	Internet is available and nothing can be installed.
One HTML file	Locked-down or offline machines: copy it over and double-click.

2.4 ON A PROJECTOR

Settings → **Chart Room** (the light scheme) reads best on washed-out projectors. Set **Interface size** to LARGER or LARGEST for the back row, and turn on **Colour-safe mode** so the sides differ by shape as well as colour.

2.5 PRIVACY

The game has no accounts, no analytics and no server; its only network request is for web fonts, and it works without them. Settings, saved two-player games, custom campaigns and a loaded voice pack stay in the browser on that machine. Clearing the browser's site data removes them.




FIG. 2.4 Settings. Chart Room, the light scheme, reads best on a washed-out projector; interface and text size are set separately.

2.6 IMMERSIVE BRIEFINGS IN A ROOM

With the optional voice pack, ► **LISTEN** in the week in review reads the week aloud as the broadcaster of the side ahead. It never starts by itself. In a room it works best on the presenter's machine with speakers.

- 1 Download **CognitiveTerrain-Voice-Emma.zip** (about 111 MB) once, before the session.
- 2 Settings → Immersive Briefings → **LOAD VOICE PACK...**, in the browser the session will use.
- 3 Press **HEAR A SAMPLE** to check the speakers.

It runs offline, on that machine only. The voice is synthetic speech from an open-source model; tell your audience so. Slower machines take a few seconds before the first line.

2.7 · ACCESSIBILITY CHECKLIST

Text size and interface size are separate settings. All text in every colour scheme passes 4.5:1 contrast. Reduced motion stops pulsing and drifting. Number keys pick moves; Esc cancels. Captions show whenever the narrator speaks.

3.0 CAMPAIGNS

Every campaign uses the same rules. What changes is the setting, the opponent, the twist at the start and the four crises. Every campaign's crises follow the decision loop: week 4 tests Notice, week 8 Interpret, week 12 Decide and week 16 Act, which counts double. Pick by what you want the room to argue about.

3.1 CHOOSING ONE

CAMPAIGN	VS	TWIST AT THE START	CRISES (WEEKS 4 · 8 · 12 · 16)	GOOD FOR
Typhoon Ines	Olvana	None	Relief Flotilla · Blackout · False Alarm · Reconstruction Compact	First games; the core loop
The Valmora Vote	Olvana	Both TRUST areas weakened	Ballot Rumor · Leaked Audio · Count Night · Transfer of Power	Elections, polarisation
Collision at Varrin Reef	Olvana	Both officials' entry points weakened	Two Videos · Missing Crew · Drills · Tribunal Ruling	Crisis diplomacy, competing videos
Fever in Port Sabre	Olvana	Rumours spread faster for both; both NOTICE areas strained	Patient Zero · Bad Batch · Forged Order · Health Accord	Infodemics, public health
Winter Grid	Donovia	Blue's FOLLOW-THROUGH strained; 4 dormant narratives	First Blackout · Price of Heat · Substation Fire · Winter Compact	Energy coercion
The Lake Referendum	Donovia	Donovia already in Blue's INTERPRET; TRUST weakened	Ten Thousand Neighbours · Engineers' Report · Debate Night · Referendum Day	Referendums, planted premises
The Silent Survey Ship	Donovia	Both DECIDE areas strained	Radio Silence · Three Stories · Snap Drills · The Inquiry	Ambiguity, alliance choices
The Vell Valley	Donovia	5 dormant narratives; Blue's NOTICE and TRUST strained	Protection Claim · Men Without Insignia · People's Vote · New Map	Annexation and its cover story

3.2 DIFFICULTY

LEVEL	WHAT CHANGES
Recruit	Red is slower, makes more mistakes and earns 2 fewer points a week. For first games and short sessions.
Standard	A steady Red that guards its weak spots, wins back what Blue takes and saves its extra moves for crisis weeks. It moves to another area when three weeks on one have gained it nothing, and defends harder once it falls behind on health. Campaigns are balanced at this level.
Veteran	Red plans two moves ahead, expects Blue's answer, saves for crises, goes after Blue's income and hides its aim. No extra budget, no hidden information.
Elite	Veteran's play under four harder rules (6.11–6.12). The only level that changes the rules; never in two-player games.

3.3 · MATCH THE OPPONENT TO YOUR TOPIC

Olvana is designed around patience against nerve: a closed state that works one area at a time from inside.

Donovia is designed to show that the aim can be doubt and inaction rather than persuasion: floods, planted premises, shows of force.

Section 6.10 and the start screen set out each one's rules and sources.

SAME MAP, FAIR COMPARISON

Groups playing the same campaign get different maps unless they start from the same seed. For a like-for-like comparison, run it as a two-player game from one Blue code (8.0).

3.4 THE OLVANA CAMPAIGNS

Same rules, different pressure. Each campaign has its own story, crises and starting conditions; every one tests Notice, Interpret, Decide, then Act. How *Olvana* fights is on the start screen and in 6.10.

TYPHOON INES

DISASTER RELIEF · THE KALVARA ISLANDS

A typhoon flattens a small island republic that trades with both sides. Over sixteen weeks of relief and rebuilding, Kalvara chooses whom to trust.

TWIST None. The best first campaign.

W04	👁	The Relief Flotilla	NOTICE
W08	🔍	The Blackout	INTERPRET
W12	⚡	The False Alarm	DECIDE
W16	➔	The Reconstruction Compact	ACT ×2

THE VALMORA VOTE

A SNAP ELECTION · VALMORA

A partner democracy holds a snap election. *Olvana* quietly favours the candidate who would cancel the Blue trade pact.

TWIST Both TRUST areas start weakened (Blue's more), and AI Synthetic Media arrives first.

W04	👁	The Ballot Rumor	NOTICE
W08	🔍	The Leaked Audio	INTERPRET
W12	⚡	Count Night	DECIDE
W16	➔	Transfer of Power	ACT ×2

COLLISION AT VARRIN REEF

MARITIME INCIDENT · TELUGA

An *Olvana* cutter and a fishing boat from Teluga, a small Blue ally, collide at a disputed reef. Two fishermen are missing.

TWIST Both sides' Government Contacts start weakened, so the fight over envoys starts on day one.

W04	👁	Two Videos	NOTICE
W08	🔍	The Missing Crew	INTERPRET
W12	⚡	The Drills	DECIDE
W16	➔	The Tribunal Ruling	ACT ×2

FEVER IN PORT SABRE

AN OUTBREAK · PORT SABRE

A fast-spreading virus hits a trading city-state that depends on both sides. Rumours move faster than the virus.

TWIST Infodemic. Held nodes push harder on their neighbours and both NOTICE areas start strained.

W04	👁	Patient Zero	NOTICE
W08	🔍	The Bad Batch	INTERPRET
W12	⚡	The Forged Order	DECIDE
W16	➔	The Health Accord	ACT ×2

◆ AGAINST OLVANA
Closed, centralized state

Typhoon Ines
the Kalvara Islands
START HERE

The Valmora Vote
Valmora

Collision at Varrin Reef
Teluga

Fever in Port Sabre
Port Sabre

the Kalvara Islands · against Olvana

TYPHOON INES

Typhoon Ines has just devastated the Kalvara Islands, a small republic that trades with both Blue and Olvana. Over the next sixteen weeks of relief and rebuilding, Kalvara will choose whom to trust.

Twist No special conditions. The best first campaign.

●

◆

●

◆

●

◆

●

◆

NOTICE
Week 4

INTERPRET
Week 8

DECIDE
Week 12

ACT
Week 16 · double

FIG. 9 The start screen. Left: campaigns grouped by opponent, the chosen one highlighted. Right: its story, its twist, and the crisis calendar with the area each crisis tests.

Replay value: every game draws a fresh map and its headlines from a large story bank, so stories rarely repeat. **Try this:** play the same campaign on Standard twice, attacking every turn and then defending first, and compare the two debriefs.

COGNITIVE TERRAIN · PRACTITIONER'S MANUAL

07

3.5 THE DONOVIA CAMPAIGNS

Donovia fights by its own rules (6.10). Its four campaigns are set in fictional countries from the U.S. Army's training environment; every town, outlet, person and event in them is invented.

WINTER GRID

A COLD SNAP · BOTHNIA

A record cold snap strains the grid of a small northern democracy that buys a third of its winter power from **Donovia**. **Donovia** offers to keep the lights on, for a price.

TWIST Blue's ACT area starts strained; 4 dormant narratives.

W04	👁	The First Blackout	NOTICE
W08	🔍	The Price of Heat	INTERPRET
W12	⚡	The Substation Fire	DECIDE
W16	➔	The Winter Compact	ACT ×2

THE LAKE REFERENDUM

A REFERENDUM · FRAMLAND

Framland votes on rebuilding a dam on the lake it shares with **Donovia**. **Donovia** wants the joint project; many Framlanders do not.

TWIST **Donovia** starts with a foothold in Blue's INTERPRET; Blue's TRUST starts weakened.

W04	👁	Ten Thousand Neighbours	NOTICE
W08	🔍	The Engineers' Report	INTERPRET
W12	⚡	Debate Night	DECIDE
W16	➔	Referendum Day	ACT ×2

THE SILENT SURVEY SHIP

A SHIP GOES QUIET · ARNLAND

An Arnland research ship stops transmitting near **Donovian** waters, while Arnland decides whether to join Blue's alliance.

TWIST Both DECIDE areas start strained.

W04	👁	Radio Silence	NOTICE
W08	🔍	Three Stories	INTERPRET
W12	⚡	Snap Drills	DECIDE
W16	➔	The Inquiry	ACT ×2

THE VELL VALLEY

ANNEXATION · GORGAS

Donovia has said for years that **Donovian** speakers in a Gorgas valley need its protection. Now armed men without insignia are on the valley roads.

TWIST 5 dormant narratives; Blue's NOTICE and TRUST start strained.

W04	👁	The Protection Claim	NOTICE
W08	🔍	Men Without Insignia	INTERPRET
W12	⚡	The People's Vote	DECIDE
W16	➔	The New Map	ACT ×2

FICTION, WITH REAL METHODS

No campaign retells a real event. What comes from research is **how Donovia operates**: flooding the zone, planting premises, dormant narratives, shows of force, denial and staged legitimacy. Each campaign names its sources on the start screen (**How Donovia fights**). **Donovia**, Bothnia, Framland, Arnland and Gorgas are countries in the U.S. Army TRADOC Decisive Action Training Environment (DATE); they are fictional training countries.

4.0 THE WHITE CELL

In a structured session a facilitator joins as the **White Cell**. It sees both minds, and once a week, after Red and before Blue, it can adjust the game: run a story, change a budget, move a crisis, strike a move. It is the game's version of the control cell in a staff exercise.

4.1 ORDER OF PLAY




FIG. 3 The White Cell console before week 2: both sides' health, budget and next crisis; what is hidden from the players; last week's moves with STRIKE buttons; this week's changes and their effect on the right.

ONLY WHEN YOU CHOOSE IT

Ordinary two-player games are unchanged. The White Cell exists only when Blue starts a **NEW STRUCTURED SESSION · WITH WHITE CELL**.

4.2 SETTING ONE UP

- 1 **Blue starts it.** TWO PLAYERS → NEW STRUCTURED SESSION · WITH WHITE CELL. Blue plays week 1 and sends the code to Red.
- 2 **Red joins** with Blue's code, as in any two-player game, and sends its code to the White Cell, not to Blue.
- 3 **The White Cell joins** with Red's first code: TWO PLAYERS → JOIN AS WHITE CELL. Both week 1 turns replay and the console opens.
- 4 **Release.** Queue changes, or none, and press RELEASE TO BLUE. Send the code to Blue. Week 2 begins.

4.3 WHAT EACH PLAYER SEES

WHAT THE WHITE CELL SEES

Both minds with no fog, both budgets, every move as it replays, Red's secret technology choices and Blue's dormant narratives. It plays no moves of its own.

The **WHITE CELL** button (where END TURN usually is) reopens the console; **LOOK AT THE BOARD** closes it to study the map.

WHAT THE PLAYERS SEE

Nothing is secret. Every story, ruling and change appears in a White Cell notice (4.7): Blue sees it before its turn, Red when Blue's code arrives. It is also in the log, the debrief and the replay.

4.4 THE CONSOLE

Eight tabs, one queue, almost all of it buttons. Each change is listed on the right as you add it, with a preview of both sides' health and the next crisis. Remove any item before you release; nothing happens until you do.

SITUATION Health, budgets, the next crisis, hidden facts, and what the computer would play next. STRIKE a move here.	BOARD Pick a node on the map, then one click to shift control, set or clear a status, or plant a dormant narrative.	RESOURCES One click to add or take budget points, or give either side a move more or fewer this week.	CRISES Move, retarget or double a coming crisis, add a new one, or schedule a new technology.
STORY A breaking news item, with an optional picture. Blue reads it before its turn, Red with Blue's next code.	RULING An adjudication the players must follow, with your reason.	NOTE A teaching note kept out of play; it appears in the debrief and replay.	LIMITS Your own caps: largest shift and budget change; switch moves, crises or strikes off.

4.5 CHANGES AND THEIR LIMITS		
CHANGE	LIMIT THE GAME KEEPS	YOUR DEFAULT LIMIT
Shift a node's control	Stays within -100 to +100	Up to 25 per change
Watched, prebunked, framed	0 to 4 weeks; 0 clears it	—
Dormant narrative	Blue's mind only	—
Budget	Stays within 0 to 18	Up to 3 per change
Moves this week	One more or one fewer; Blue's at once, Red's when its turn starts	Can be switched off
Crisis	Not in the past, not in a week that already has one	Can be switched off
New technology	A future week without one	Can be switched off

Set your own limits in the **LIMITS** tab; they change only what the console offers.

4.6 STRIKING A MOVE
A strike removes one of the past week's moves. Every copy rebuilds the week without it, and both sides are told why. Strikes can be switched off in **LIMITS**.

4.7 WHO SEES WHAT, AND WHEN
Every change, story and ruling is shown to both sides. Blue sees the White Cell notice before its turn; Red sees it when Blue's code arrives. Teaching notes are the exception: they appear only in the debrief and the replay.

WORKED EXAMPLE
The next page follows one real week at the console, from finding the pressure on the map to the notice Blue reads.




FIG. 3 The **SITUATION** tab: health, budgets, the next crisis, what is hidden from the players, and last week's moves with **STRIKE** buttons.

SECTION 4 • EXAMPLE • ONE WEEK AT THE CONSOLE

A WEEK AT THE CONSOLE

Typhoon Ines, before week 2. Olvana spent its first turn hammering Blue's Social Media. The facilitator wants the room to feel an outside shock and to see how Blue responds to pressure on resources. Every picture is the real game.




THE PLAN

- 1 **Story:** "Port strike spreads to three cities", with a picture.
- 2 **Board:** shift Blue's Social Media 10 toward Olvana, as the clip spreads online.
- 3 **Resources:** Blue +2 budget, to see where it spends.
- 4 **Ruling:** attacks on medical services are out of bounds.

FACILITATOR'S NOTE Pair every board change with a story, so players know why it happened.

STEP 1 • FIND THE PRESSURE The White Cell sees Blue's mind with no fog. Social Media is down to Blue 12, contested, after Olvana's Plant Stories. **PICK A NODE ON THE MAP** closes the console; the bar at the bottom says what to do. Click Social Media.



STEP 2 • QUEUE THE CHANGES The console reopens on Social Media. One click on **10 toward Olvana** queues the shift. The list on the right holds all four changes; the preview shows both sides still SOUND for the week 4 crisis. Remove any item with x. Then **RELEASE TO BLUE**.



STEP 3 • WHAT THE PLAYERS SEE Blue loads the code and reads this before its turn: the story and picture, the ruling, and every change in plain words. Red reads the same notice when Blue's week 2 code arrives (4.7).

4.8 STEERING THE EXERCISE

4.8 STORIES AND PICTURES

A story is a short breaking-news item with an optional picture. Drop, paste or browse for an image: the game shrinks it to fit inside the turn code, so it travels by email or chat like any other code. A picture that cannot be made small enough is refused with a message.

4.9 RULINGS AND NOTES

A **ruling** is an adjudication the players must follow, with your reason. A **note** is a teaching note kept out of play for the debrief. Neither changes the board.

4.10 ENDING THE SESSION

When Red's last turn ends the game, Red's final code comes to you; release it and send the last code to Blue so every copy ends together. The debrief carries the White Cell's record and notes.

POINT AND CLICK

Click any node on the map and the console opens on it, ready for a shift or a status. **LOOK AT THE BOARD** closes the console; the **WHITE CELL** button brings it back.

4.11 SUGGESTED INJECTS

IDEA	TRY
Events outrun preparation	Bring the next crisis forward a week, or make it test the area a side has neglected.
An outside shock	A story plus a small control shift on the matching node, e.g. a viral clip and a nudge on Social Media.
Resource pressure	Cut a side's budget by 2 before a crisis and watch what it gives up.
Rules of engagement	A ruling, then strike any move that breaks it.
A losing side recovers	One extra move for a week; ask in the debrief what it did with it.

FACILITATOR'S NOTE Say in the debrief why you changed what you changed.

5.0 THE DEBRIEF

The game ends with a neutral debrief and a replay of both minds week by week. Most of the learning happens here, so leave time for it.





FIG. 1 The debrief: crisis points and health week by week, followed by three verdicts, turning points and lessons further down.

EXAMPLE

An illustration of a split result. Blue wins the campaign 7–5 on crisis points. But Olvana's average Decision Health across the sixteen weeks was higher, and it moved more influence for each point it spent. Two of the three verdicts go to the side that lost.

Ask the room: what did Blue do that won the crises, and what did it cost?



5.3 QUESTIONS

QUESTIONS THAT WORK

- Which crisis did you lose, and when did you first know it was coming?
- Find a move with a great spread score that changed nothing. Why?
- When did defending your own mind matter more than attacking theirs?
- What did Red seem to be trying to do? What told you?
- Which of Blue's advantages as an open society did you use, and which did you lose?
- What would you tell a decision-maker in the real version of this crisis?


6.0 THE MODEL IN NUMBERS

Every number the game uses to score a campaign, in plain words first and then as a formula, with a worked example. Nothing here is hidden from players: the side panel and the week in review show the same sums.

6.1 AREA HEALTH

Each area's health is how much of it the owner still controls, from 0 to 1. Nodes that are slipping get part credit, so the score moves smoothly rather than in jumps. A node fully yours counts 1; a neutral node about a third; a node the enemy holds hard counts nothing.

$$\text{area} = \text{average over its nodes of clamp((lean + 50) \div 150, 0, 1)}$$



6.3 WHAT MOVES THE NEEDLE

How much each change is worth, to Decision Health and to the next crisis:

CHANGE	DECISION HEALTH	NEXT CRISIS
Tested area +0.10	+0.01 to +0.04	+0.05 to +0.06
Any other area +0.10	+0.01 to +0.04	+0.01 to +0.02
TRUST +0.10	about +0.02	about +0.01 (more if it is tested)
Lose an area	-0.06 and more	-0.04; -0.12 and more if tested

In short: the area the next crisis tests is worth three to five times any other. That is why the side panel names it in the two weeks before a crisis.

6.2 DECISION HEALTH

Decision Health adds up the four loop areas, weighted by how much each side relies on them. The total is then scaled by trust, and 0.06 comes off for every area lost.

$$\begin{aligned} \text{areas} &= w \cdot \text{NOTICE} + w \cdot \text{INTERPRET} + w \cdot \text{DECIDE} + w \cdot \text{ACT} \\ \text{DQ} &= \text{areas} \times (0.75 + 0.25 \cdot \text{TRUST}) \\ &\quad - 0.06 \times \text{areas lost} - \text{flood} \end{aligned}$$

WEIGHT	NOT	INT	DEC	ACT
BLUE	0.20	0.25	0.30	0.25
OLVANA · DONOVIA	0.15	0.20	0.40	0.25

Trust can take a quarter off everything; with trust gone, the most a side can score is 0.69. **Flood** (Donovia only): -0.02 for each of its attacks that lands in a week beyond the first two (one more is free per Lessons Learned Blue has played, up to 3), at most -0.06, until Donovia's next turn.

EXAMPLE

Blue's areas: NOTICE 0.9, INTERPRET 0.8, DECIDE 1.0, ACT 0.9; TRUST 0.7; nothing lost. Each area adds its share; trust then trims the total.

NOT	INT	DEC	ACT	
.18	.20	.30	.225	0.905

× TRUST 0.925 (0.75 + 0.25 × 0.7)

0.84

HEALTH 0.84 · A CRISIS ALSO WEIGHS ONE AREA (6.4)

6.4 CRISES AND WINNING

6.4 A CRISIS

A crisis mixes a side's general health (60%) with the health of the one area it tests (40%). Losing that area costs another 0.08. The result falls in one of three bands.

$$\text{crisis} = 0.6 \cdot \text{DQ} + 0.4 \cdot \text{tested area} - 0.08 \text{ if the tested area is lost}$$

ERROR below 0.72 · 0 pts	DELAYED 0.72+ · 1 pt	SOUND 0.86+ · 3 pts
------------------------------------	--------------------------------	-------------------------------

Plus 1 to a side whenever the other side scores **ERROR**. A crisis that counts double (week 16 in the built-in campaigns) doubles everything, bonus included.

The "decided in ... hours" shown in a crisis result is $6 \div (\text{Decision Health, or the crisis score at a crisis})^2$, rounded, and never more than 267. It is a way of showing health, not a separate rule.

6.5 WINNING

Most crisis points after week 16 wins. **Tied?** The side whose Decision Health is more than 0.03 higher wins; otherwise it is a draw. **Knockout**, earlier: the other side has lost at least 2 areas, and at least 2 more than you, at the end of 2 weeks running.

FACILITATOR'S NOTE Ask players to predict a crisis result before it lands; then compare.

EXAMPLE

The same Blue, one crisis, two outcomes. Blue's Decision Health is 0.84 (6.2) and the crisis tests **INTERPRET**, which is at 0.8:

$0.6 \times 0.84 + 0.4 \times 0.80 = 0.82 \rightarrow$ **DELAYED**, 1 point

Had Blue spent the two weeks before lifting **INTERPRET** to 0.9, its Decision Health would be 0.86, and:

$0.6 \times 0.86 + 0.4 \times 0.90 = 0.88 \rightarrow$ **SOUND**, 3 points

One area, raised by a tenth, is worth 2 points here. That is the lesson of 6.3 in a single crisis.

Week 8 · crisis 2 of 4 · tests **INTERPRET**

THE BLACKOUT

The undersea cable that carries most of Kalvara's internet traffic goes dark. Islanders are cut off from family abroad and from relief updates. Both Blue and Olvana offer to restore service, and both are accused of cutting the line.

This crisis leans on each side's **INTERPRET** area: it counts for 40% of the result, and losing it costs extra.

BLUE	OLVANA
DELAYED +1 point	SOUND +3 points
Crisis score 0.85 · health 0.79 · INTERPRET 95% · decided in 8 hours	Crisis score 0.96 · health 0.93 · INTERPRET 100% · decided in 7 hours
Blue's offer stalled in interagency review while Olvana's technicians were already on the beach.	Olvana offered repair ships with no conditions attached and stayed out of the blame game.

Each side's result comes from its Decision Health: how much of its **NOTICE**, **INTERPRET**, **DECIDE**, **ACT** and **TRUST** areas it still controls. Sound earns 3 points, delayed 1, a mistake 0, and a mistake also hands the other side 1 point.

CONTINUE

FIG. 6.4 How players see a crisis result: each side's band, the points, and what each side actually did.

6.6 MOVES AND ODDS

How hard a move hits, how often it fails and what comes in each week. The Field Manual shows each one in play; here is the arithmetic in one place.

6.6 SPREAD SCORE AND GRADE

Every attack except Show of Force rolls an audience, scores it with the research's engagement formula, and turns the score into a grade that multiplies the push.

$$\text{score} = (0.1 \cdot \text{views} + 0.3 \cdot \text{likes} + 0.7 \cdot \text{comments} + 1 \cdot \text{shares}) \div \text{posts}$$

GRADE	SCORE	PUSH ×
F · E · D	under 3 · 3 to under 4 · 4 to under 10	0.25 · 0.45 · 0.7
C · B	10+ · 100+	1.0 · 1.2
A · A+	1,000+ · 10,000+	1.4 · 1.6

Reach is rolled on a ten-times scale, so from grade C upward ×10 is one grade. ×10: their TRUST lost; each technology you used. ÷10: a Watch Team on the target. **Smaller:** entry point ×2.5; their News Media held ×3.2; Red into Blue ×1.25; Blue into Olvana ÷1.4 (Donovia ÷1.25). Every one of these is already counted in the preview a player sees.

6.7 PUSH

The push is the card's power, made bigger or smaller by the area, the audience and the node's state.

$$\begin{aligned} \text{push} &= \text{card power} \times \text{area fit} \times \text{grade} \\ &\times 0.5 \text{ if prebunked (0.45 for Donovia's attacks)} \\ &\times 1.4 \text{ if framed} \cdot \times 2 \text{ if it wakes a dormant narrative} \\ &\times 1.1 \text{ for Blue's attacks on Donovia's TRUST} \end{aligned}$$

Card power: Plant Stories 20 (×1.5 on an entry point), Frame 16, Bot Swarm 12 plus 7 on every link (12 if you hold their Social Media, or used Social Bots or Autonomous AI Agents and the other side did not prepare), Repost 10, Show of Force 16 with no grade. Official Statement restores 14 (Blue ×1.6 while TRUST holds; ×1.5 on a watched node; ×0.5, with no Blue bonus, once TRUST is lost). Worked example: 6.7 below.

6.8 ODDS OF FAILURE

Chances add up. A failed attack pushes back toward the defender.

MOVE FAILS IF...	CHANCE
Any attack	8%
+ target watched	+15%
+ target prebunked	+8%
+ Blue posting into Olvana / Donovia	+2% / +1%
Show of Force	+4%
Official Statement ignored	6% (+20% if TRUST lost)

A failed attack pushes the node back 6 toward the defender (9 when Donovia's fail); a failed Show of Force also pushes each neighbour back 4. With ACT lost, 1 move in 4 misfires before any of this is rolled.

6.9 EACH WEEK

Budget: 6 + 1 per enemy node held (Top Leadership 2) + 2 for their Government Contacts; at least 2 a week; at most 18 held. Recruit Red −2; Donovia +1. Extra move: Blue 2, Red 5, once a turn after the last move, only if 2 would be left.

Recovery: each node +1 toward its owner (+1.5 per Lessons Learned level on that node, half for Red), halved with TRUST lost. Blue's TRUST nodes +2 more; Red's entry points ×1.2 (Donovia's ×1.1); Red's nodes leaning Blue ×0.35. None while framed; once held by the enemy, only the Lessons Learned part, at half rate. Any neighbour leaning 50+ the other way pushes 5 (6 in Fever in Port Sabre; half if prebunked), and +1 more toward the owner if any neighbour leans 50+ its way.

6.10 RULES IN FULL

The Field Manual teaches the game screen by screen. These are the remaining rules a facilitator gets asked about, with the exact numbers.

6.10 HOLDING NODES AND LOSING AREAS

LEAN TOWARD A SIDE	MEANS
10 or more	Foothold: that side's attacks reach every linked node.
30 or more	Held: earns income in the enemy mind; no recovery for the owner.
Half an area's nodes held	Area lost (rounded up: 3 of 5, 3 of 6). DECIDE also falls while its Top Leadership is held. Won back at once when that stops.

AREA LOST	PENALTY FOR ITS OWNER
NOTICE	The other mind is hidden (?? on the map and bars).
INTERPRET	Every move costs 1 more.
DECIDE	2 moves a turn instead of 3.
ACT	1 move in 4 misfires; it is still paid for.
TRUST	Official Statements halved with no Blue bonus and ignored 26% of the time; recovery halved; enemy posts reach $\times 10$.

Each lost area also takes 0.06 off Decision Health (6.2).

6.11 STATUSES

STATUS	SET BY LASTS	EFFECT
WATCHED	Watch Team on a node and its links · 3 wk	Enemy reach $\div 10$, fail +15%; your statements $\times 1.5$; framing cleared; a dormant narrative there removed.
PREBUNKED	Prebunk on a node and its links · 3 wk (2 when preparing for a technology)	Enemy push $\times 0.5$ ($\times 0.45$ v Donovan), fail +8%; spread from neighbours halved.
FRAMED	Enemy Frame the Story · 2 wk (3 by Donovan)	No recovery; every enemy attack $\times 1.4$. Cleared by Watch Team or Lessons Learned.
TRAINED	Lessons Learned on a node and its links · rest of game, up to 3 levels	Recovery +1.5 a level (Red +0.75); once held by the enemy, half that still pulls it back.

A status counts its first week and loses one at each upkeep. Setting it again restarts the count. Repost & Rebrand strips Watched and Prebunked.

6.12 HOW THE OPPONENTS DIFFER

RULE	BLUE	OLVANA	DONOVIA
Start in Blue's Social Media	—	Blue 80	Blue 80
Bot Swarm cost	3	2	2
Extra move	2	5	5
Health weight on DECIDE	30%	40%	40%
Lessons Learned	+6, recovery +1.5	+3, +0.75	+3, +0.75
Blue's reach into it	—	$\div 1.4$, fail +2%	$\div 1.25$, fail +1%
Its entry points recover	—	$\times 1.2$	$\times 1.1$
Budget each week	—	—	+1
Framing lasts	2 wk	2 wk	3 wk
Failed attack backfires	6	6	9

Blue: TRUST nodes recover 2 more a week; Official Statements push 22 while TRUST holds; Red nodes leaning Blue recover at about a third of the rate.

Donovia only. Flooding: each attack landing beyond the second in a week takes 0.02 off Blue's health until Donovia's next turn (up to 0.06; each Lessons Learned Blue plays, up to 3, absorbs one more). Dormant narratives: 3 hidden in Blue's NOTICE, INTERPRET and TRUST (4 or 5 in some campaigns); the first attack on one hits $\times 2$ and a Watch Team removes it. Show of Force: a ninth card, 3 points, any ACT or DECIDE node with no foothold and no audience roll, push 16, fail chance +4%; Blue's next move costs 1 more; if it fails it backfires 9 on the node and 4 on each neighbour. Prebunk works better against it ($\times 0.45$), and Blue's attacks on its TRUST hit 10% harder.

Elite only. A card Blue has used on Red 3 weeks running: $\times 0.7$ until skipped a week. Red's Watch Teams and Prebunks stay hidden until Blue attacks a covered node; then the whole group shows. A node Red wins back from Blue is DUG IN: $\times 0.7$ until Blue holds it. A failed Blue attack on Red gives Red 1 point.

EXAMPLE

Upkeep on one node. An Olvana node at -20 recovers 1 toward Olvana (-21); a linked node Blue holds at $+60$ spreads 5 toward Blue (-16). Every change is worked out from the board as it stood, then applied at once.

7.0 BUILDING CAMPAIGNS

Set the scene for your own audience, place the crises and new technologies on the calendar, tune how strong each side starts, then test it, check its balance and share it. No code involved.




FIG. 5 The campaign editor, editing a copy of Typhoon Ines with a new crisis in week 10. Click a crisis or technology on the calendar to edit it on the right; + adds one in that week.

7.1 MAKE ONE

- 1 On the start screen press **CAMPAIGNS** (or ⚙️ → **CAMPAIGN EDITOR**).
- 2 + **NEW CAMPAIGN**, or **COPY & EDIT** a built-in one. Copying is the fastest way to a balanced start.
- 3 Choose the opponent (Olvana or Donovan), write the setting and the headlines' places and outlets, and set any twist at the start.
- 4 Place 1 to 6 crises and up to 5 new technologies (7.2).
- 5 **CHECK BALANCE**, then **TEST PLAY**.

SHARING

Save a `.campaign.json` file or copy a code that starts CTC1 - . Campaigns are plain text and cannot run code, so they are safe to pass around. Campaigns live in the browser that made them: keep the file.




FIG. 6 The campaign library: play a campaign, or copy one to edit.

7.2 CRISES AND TECHNOLOGIES

7.2 WRITING A CRISIS

A crisis is a moment that tests one area of both minds. For each one you write what happens, choose the area it tests, and write how each side handles it at each result: SOUND, DELAYED and ERROR. Players read your text at the moment of truth, so keep each line to one or two sentences about a decision, not an outcome.

Pick the area by the decision. A flood of rumours tests NOTICE; a contested video tests INTERPRET; a deadline tests DECIDE; a delivery that must happen tests ACT; a scandal tests TRUST.

In the built-in campaigns the last crisis counts double; in your own, tick **Counts double** on any crisis. Spread the tested areas so no single defence wins the campaign.

CHECKING BALANCE

CHECK BALANCE plays your campaign 40 times, computer against computer, and reports how often each side wins. 7.4 explains how to read the result.

ORIGINAL FICTION

The built-in campaigns invent every place, outlet, person and event. Doing the same keeps a session about methods rather than about a real event, and keeps it usable in any room.

YOUR TECHNOLOGY · WEEK 02

NAME

Social Bots

YEAR

1999

WHEN

PAST

PRESENT

FUTURE

ATTRIBUTE SCORES

S = 9

FLAGGED · S = 9

Pathways

LOW

MED

HIGH

Data collection

LOW

MED

HIGH

Targeting

LOW

MED

HIGH

Social networks

LOW

MED

HIGH

MOVES IT BOOSTS

Plant Stories

Frame the Story

Bot Swarm

Repost & Rebrand

HOW MUCH

+1 SPREAD GRADE

+2 SPREAD GRADES

☐ Plant Stories reaches NOTICE and INTERPRET nodes with no foothold

☒ Bot Swarm hits linked nodes at full strength

WHAT IT IS IN REAL LIFE

Automated and semi-automated accounts imitate human behavior, letting a small team simulate a crowd.

In the game: Bot Swarm posts gain +1 spread grade. Bot Swarm hits links at full strength.

DONE

[Remove this technology](#)

FIG. 7 Building a technology: name, year, era, the four attribute scores and their total S, and which moves it boosts.

7.3 TECHNOLOGIES

Pick from 18 preset technologies, past, present and speculative, or **BUILD YOUR OWN**. Each is scored 1 to 3 on up to four attributes from Rushing, Umeokolo & Xu (2025): new pathways, data collection, targeting and social networks. A total of 9 or more is the paper's mark of a disruptive innovation. Where the paper gives no date or score, the editor marks the figure as illustrative or the designer's estimate.

7.4 READING A BALANCE CHECK

CHECK BALANCE plays your campaign 40 times, the computer against itself, and reports how often Blue wins. Forty games is enough to spot a lopsided campaign, but not enough to tell 45% from 55%.

7.4 WHAT THE NUMBER MEANS

The editor calls 25–65% Blue wins a **fair match**. The built-in campaigns sit between about 35% and 60%.


Any count of 40 games carries luck. A result is only good to about ± 15 points either way: a campaign that shows 35% might really be close to 50%. The chart shows the likely true range for three results.

7.5 GETTING A FIRMER ANSWER

Run **CHECK BALANCE** again and add up the wins. Two runs (80 games) narrow the range to about ± 11 points; four runs (160 games) to about ± 8 .

7.6 IF IT IS LOPSIDED


- Soften the twist at the start, or give it to both sides.
- Move a crisis to an area the stronger side has not prepared.
- Spread the tested areas, so no single defence wins every crisis.
- Change one thing at a time, then check again.



FACILITATOR'S NOTE Balance is tested computer against computer; people will play it differently.

8.0 ACROSS ORGANISATIONS

Two-player games need nothing but a way to send a line of text. That makes the game usable between cells, units or agencies that cannot share a network or install software.



One week of a structured session between organisations. Without a White Cell, Red's code goes straight back to Blue.




FIG. 8 A turn code with a COPY button, and a box for the reply.

8.1 • WHAT A CODE CARRIES

The game and its map seed, the campaign, the week, the side and that side's moves. Each copy replays the other's moves and checks the board matches; if it ever does not, the log says so. Codes contain no personal data.

8.2 CHECKLIST

- 1 **Same version.** Codes work only between builds with the same minor version (0.19.x with 0.19.x), and the game cannot check this for you: a mismatch shows only as a failed board check. Send everyone the same download.
- 2 **Same campaign.** For a custom campaign, Red and any White Cell import the file before their first code (7.1). The code checks it.
- 3 **Assign sides.** Blue starts the game; Red joins with Blue's first code. The campaign decides whether Red is Olvana or Donovia. To steer it yourself, start a structured session and join as the White Cell (4.0).
- 4 **Set a rhythm.** One week per session or per day keeps a sixteen-week campaign moving.
- 5 **Collect the debriefs.** Each side sees the same debrief at the end; Red sends one final code so Blue's game ends too (in a structured session, the White Cell sends it).

ANY CHANNEL WORKS

Email, chat, Teams, a text file on removable media. Extra text around a pasted code is ignored, so players can add a note.

Each player's game saves after every code and appears under **RESUME**, so a game can pause for days.


9.0 DESIGNER'S NOTES

Each note takes one part of the game. It says what the rule models, how I arrived at it, and where it comes from.

FROM RESEARCH A published source shaped this rule.

DESIGNER'S ASSUMPTION My own design choice, made so the game plays well.

THE BOARD: FIVE AREAS OF ONE MIND



NOTE 1 THE MIND AS TERRAIN

MODELS A nation's decision-making as the ground being fought over.

Two brains of about thirty connected nodes. Every node is a real kind of person or institution: a newsroom, an analysts' cell, a cabinet, a crisis team, the public's trust in its own government.

HOW I GOT THERE Cognitive warfare is hard to teach because you can't see it, and people picture it as "fake news". I wanted the real target on the table: whether the other side can still decide well under pressure.

RUSHING, HERSCH & XU (2026)

NOTE 2 LOSING AN AREA

MODELS What it costs to lose a capability.

An area falls when the other side holds at least half its nodes (DECIDE also falls if Top Leadership is held). Each loss has one plain effect (the strip above), so players feel it the way a staff would.

HOW I GOT THERE The five areas follow the research's cognitive OODA loop with trust underneath. The effects are deliberately simple so they can be read off the map.

OODA FRAMING: RUSHING ET AL.

EFFECTS: DESIGNER'S ASSUMPTION

NOTE 3 DECISION HEALTH

MODELS How well a side can still decide, 0 to 1.

The score is not views or nodes held. It is the health of the four loop areas, scaled by trust (between $\times 0.75$ and $\times 1$), minus a penalty for each area lost. The arithmetic is in 6.2.

HOW I GOT THERE Spreading far is not the same as changing minds. A centralized state leans 40% on DECIDE (Blue 30%): when everything runs through the top, the top is the target.

RUSHING ET AL. · WEIGHTS AFTER SHIN (NIU)

WEIGHTS: DESIGNER'S READING

NOTE 4 ENTRY POINTS AND Footholds

MODELS Messages need a way in before they can travel.

Each brain has three entry points: News Media, Social Media and Government Contacts. Attacks start there and spread along links from nodes you have a foothold in.

HOW I GOT THERE It makes players plan an approach instead of hitting the most valuable node first, which is how I believe influence work has to be sequenced.

DESIGNER'S ASSUMPTION

"SPREADING FAR IS NOT THE SAME AS CHANGING MINDS."

MOVES & METRICS

Eight moves, one engagement metric, and odds that change when a side prepares.

-] Plant Stories
- ≡ Frame the Story
- ⦿ Bot Swarm
- ↻ Repost & Rebrand
- 👁 Watch Team
- 🛡 Prebunk
- 🔊 Official Statement
- 📖 Lessons Learned

→] NOTE 5 EIGHT MOVES, EIGHT CAPABILITIES

MODELS The attacker's and defender's capability families. Four attack cards and four defence cards, one for each capability family in the research. Military mode keeps doctrinal names. Donovan adds a ninth, Show of Force.

HOW I GOT THERE *I did not invent a list of tricks. Mapping one card to each family keeps the game honest to the framework and easy to cite in a debrief.*

RUSHING, HERSCH & XU (2026) §2.2.2, §2.3.2

SHOW OF FORCE: ISW, IISS

⦿ NOTE 6 THE SPREAD SCORE

MODELS How engaging a campaign of posts actually was.

Every attack except Show of Force rolls an audience and is graded with the research's engagement metric. The grade then multiplies the move's push:



HOW I GOT THERE *I kept the formula and grade scale because they make the point that interaction outweighs views. The multipliers are mine. A viral post on the wrong area still loses the crisis.*

METRIC AND GRADES: RUSHING & XU (2025)

MULTIPLIERS: DESIGNER'S ASSUMPTION

👁 NOTE 7 FAILURE AND BACKFIRE

MODELS Platforms, watchers and prepared audiences.

Any attack can be taken down or fall flat. Watching and prebunking raise those odds, and a failed attack nudges the node back toward the defender (6.8 has the numbers).

HOW I GOT THERE *I wanted defence to change the odds, not just add hit points, so players feel why detection and inoculation matter before a campaign lands.*

DESIGNER'S ASSUMPTION

📖 NOTE 8 RECOVERY, FRAMING AND SPREAD

MODELS Minds drifting back, and ideas settling in.

Each week every node recovers a little toward its own side; strongly held nodes push on their neighbours. Framing stops recovery for two weeks (three when Donovan frames), and every attack on the node lands 40% harder meanwhile: the game's version of priming.

HOW I GOT THERE *Tuned so that a single hit fades but a sustained, sequenced campaign holds.*

DESIGNER'S ASSUMPTION

"DEFENCE SHOULD CHANGE THE ODDS, NOT JUST ADD HIT POINTS."

THE MOMENTS THAT COUNT

Sixteen weeks so it fits a lunch break; four crises so preparation pays.

01	02 ♦ TECH	03	04 CRISIS	05	06 ♦ TECH	07	08 CRISIS	09	10 ♦ TECH	11	12 CRISIS	13	14 ♦ TECH	15	16 CRISIS ×2
----	--------------	----	--------------	----	--------------	----	--------------	----	--------------	----	--------------	----	--------------	----	--------------------

NOTE 9 FOUR CRISES

MODELS Decisions under pressure at the moments that matter.

In weeks 4, 8, 12 and 16 a crisis tests one area of both minds; in the built-in campaigns the last counts double. The result depends on overall health and, heavily, on the tested area (6.4).

HOW I GOT THERE *Weighting the tested area makes players read the calendar and prepare, which is the skill I wanted to reward.*

OODA FRAMING: RUSHING ET AL.

SCORING: DESIGNER'S ASSUMPTION

NOTE 10 COGNITIVE SUPERIORITY

MODELS One side's decision-making clearly dominating.

If one side has lost at least two more areas than the other for two weeks running, the campaign ends early. The race shows in the side panel as AREAS TAKEN: a bar from the other side +2 to you +2.

HOW I GOT THERE *The research defines cognitive superiority as a sustained advantage in the decision loop; the knockout gives that idea a hard edge.*

RUSHING, HERSCH & XU (2026)

THRESHOLDS: DESIGNER'S ASSUMPTION

NOTE 11 NEW TECHNOLOGIES

MODELS Disruptive innovations arriving mid-campaign.

Both sides choose in secret: use it for an attack bonus, or prepare and deny the other side's. Using it against a prepared opponent gains nothing.

HOW I GOT THERE *A dilemma, because that is the real problem with new tools: you don't know whether the other side is ready. The scoring and the threshold of 9 are the paper's.*

RUSHING, UMEOKOLO & XU (2025)

SOME SCORES: DESIGNER'S ESTIMATE

NOTE 12 BUDGET AND MOVES

MODELS Finite attention and resources.

Three moves a week from a budget that grows with what you hold, capped at 18, with one extra move a turn after your last, if you can pay for it and still have 2 left. Access to officials pays more than any post.

HOW I GOT THERE *The cap stops hoarding; the extra move rewards spending.*

DESIGNER'S ASSUMPTION

"READ THE CALENDAR, AND PREPARE THE AREA THE NEXT CRISIS TESTS."

RED IS NEVER A MIRROR

	BLUE	OLVANA	DONOVIA
IS	An open democracy	A closed, one-party state	An oligarchic state at permanent information war
STRENGTH	Free press, truth that sticks, cheap extra move	Censored internet, state-funded content, already inside	Floods, planted premises, dormant narratives, shows of force
WEAKNESS	Its open internet lets Red's posts spread	Runs from the top; slow to learn	Predictable; threats can blow back
SHOULD FEEL	Under pressure, keeping its head	A fortress with one gate	Loud

NOTE 13 OPEN AGAINST CLOSED

MODELS Why open and closed societies are vulnerable in different ways.

Blue's posts into a censored internet spread less; Red's into Blue's open internet spread further. Blue gets trust that recovers, statements that hit harder while trust holds, and truth that sticks.

HOW I GOT THERE *The asymmetry is the point: you fight a closed system differently from how it fights you.*

ATP 3-13.1 · SHIN (NIU)

SOME EFFECTS: DESIGNER'S ASSUMPTION

NOTE 14 TWO OPPONENTS, KEPT FICTIONAL

MODELS Two doctrines of cognitive warfare.

Red's countries come from the Army's training world, so they are familiar and clearly fictional. Every rule names its source or is labelled as my assumption (10.1).

HOW I GOT THERE *Both share a core and differ in temperament, so each feels different with a person playing it.*

ODIN DATE · NIU · EADS ET AL.

NATO C2COE · ISW · IISS · CARNEGIE

NOTE 15 AN OPPONENT THAT PLAYS FAIR

MODELS A competent adversary at four levels of skill.

Recruit is noisier and earns 2 fewer points a week. Veteran plans two moves ahead on a scratch copy of the game and expects Blue's best answer: no extra budget, no hidden information, never the game's own dice. Elite is the one level that changes the rules, and the start screen says how.

HOW I GOT THERE *Meant to be hard to beat, without cheating. Each of its moves also shows the likely reason, judged only from what the player could see.*

DESIGNER'S ASSUMPTION

NOTE 16 BUILT TO BE USED

MODELS A teaching game that runs anywhere.

It explains itself at every moment that matters. Two players need only a turn code; campaigns are plain text that cannot run code; Immersive Briefings are optional; all text in every scheme passes 4.5:1 contrast.

HOW I GOT THERE *Keep the original game: new features go around the core loop, never instead of it. The editor calls a campaign a fair match when Blue wins 25–65% of computer games; the built-in ones sit between about 35% and 60%.*

DESIGNER'S ASSUMPTION

"MEANT TO BE HARD TO BEAT, WITHOUT CHEATING."

HOW IT SHOULD FEEL

Rules decide what happens; these notes are about what a player should feel while it happens. They are my intentions, not findings.

NOTE 17 OLVANA: A FORTRESS WITH ONE GATE

HOW IT SHOULD FEEL *Olvana should feel like a fortress with one gate. Against a human Olvana, Blue feels a slow squeeze: no fireworks, just patient pressure on one area at a time from someone who was already inside. Playing Olvana, you feel safe behind the firewall but know everything hangs on DECIDE, and that once it cracks you learn too slowly to fix it. It is patience against nerve.*

SHIN (NIU)

TEMPERAMENT: DESIGNER'S READING

NOTE 18 DONOVIA: LOUD

HOW IT SHOULD FEEL *Donovia should feel loud. Against a human Donovia, Blue feels rushed and off balance: too many stories at once, old narratives going off like mines, a show of force just as you planned your turn. The test is whether you keep your head and defend the premises instead of chasing every claim. Playing Donovia, you feel the thrill of cheap, fast aggression, and the risk when your threats blow back.*

ISW · NATO C2COE · IISS · CARNEGIE

TEMPERAMENT: DESIGNER'S READING

NOTE 19 IMMERSIVE BRIEFINGS: THE SECOND TELLING

WHY IT EXISTS *Every week in this game happens twice: once on the board, and once in how each side tells it. Immersive Briefings let you hear the second version, and hearing the other side's broadcaster tell your public that its leaders cannot decide is the point. It is optional because the base game must stay small, quiet and offline.*

DESIGNER'S CHOICE

"EVERY WEEK IN THIS GAME HAPPENS TWICE."

10.0 SOURCES AND LIMITS

A model is only as useful as your sense of where it is solid. This section is the honest map.

10.1 WHAT COMES FROM WHERE

RULE	FIELD MANUAL	BASIS
Five areas and the decision loop	9.0	From research: Rushing, Hersch & Xu (2026).
What losing each area does	6.10	Designer's assumption.
Decision Health weights	6.2	Designer's reading of Shin (2026) on centralized command.
The eight cards	6.10	From research: the capability families in Rushing, Hersch & Xu (2026). Their numbers are assumptions.
Spread score and grades	6.6	From research: Rushing & Xu (2025). The push multipliers are assumptions.
Technology scores and threshold	7.3	From research: Rushing, Umeokolo & Xu (2025). 16 of the 18 score sets are the designer's estimates.
Crisis scoring and bands	6.4	Designer's assumption.
Knockout	6.5	Idea from Rushing, Hersch & Xu (2026); thresholds are assumptions.
Budget, costs, failure, recovery	6.8, 6.9	Designer's assumption.
Opponent rules	6.10	Each rule names its source in the game, or is marked as an assumption.
Elite: learns your playbook	6.12	From research: defenders adapt and learn, and repeated exposure to a technique builds resistance (Rushing, Hersch & Xu 2026). How much softer is an assumption.
Elite: hidden defences, dug in	6.12	Designer's reading of the same paper: detection works unseen until it bites, and an audience that has seen a campaign resists it better.
Elite: blowback	6.12	Designer's assumption.

10.2 WHAT HAS BEEN CHECKED

- The rules and numbers in both manuals were checked against the game's code for this version.
- Balance was tested by the computer playing itself: on Standard the built-in campaigns land between about 35% and 60% Blue wins [7.4]. Against a Blue that also thinks ahead, Blue wins 25–38% on Veteran and 5–23% on Elite.
- Automated tests check that every copy of a two-player or White Cell game stays in step.

10.3 WHAT HAS NOT

- No independent review, assurance or validation of the game, its numbers or its value for teaching.
- No study of what players learn from it.
- The cited authors and their institutions have not reviewed it.
- The computer opponent is a game opponent, not a model of any real actor's decisions.

10.4 REFERENCES

- Rushing, B., Hersch, W., & Xu, S. (2026). *Cognitive Warfare: Definition, Framework, and Case Study*. arXiv:2603.05222.
- Rushing, B., & Xu, S. (2025). *Quantifying the Engagement Effectiveness of Cyber Cognitive Attacks: A Behavioral Metric for Disinformation Campaigns*. arXiv:2510.15805.
- Rushing, B., Umeokolo, M.-R., & Xu, S. (2025). *Towards Proactive Defense Against Cyber Cognitive Attacks*. arXiv:2510.15801.
- Headquarters, Department of the Army (2018). *ATP 3-13.1, The Conduct of Information Operations*.
- NATO Science & Technology Organization. *Cognitive Warfare: NATO Chief Scientist Research Report*.
- Shin, D. W. (2026). *Viewing China's Cognitive Warfare Through the Lens of Its Global Strategy*. National Intelligence University.
- Eads, L. J., Clarke, R., Lin, X. S., & McCreight, R. (2023). *Warfare in the Cognitive Age: NeuroStrike and the PLA's Advanced Psychological Weapons & Tactics*.
- Iida, M. (2024, May 5). *China's Chilling Cognitive Warfare Plans*. The Diplomat (cited in Shin 2026).
- DISARM Foundation. *DISARM Framework* (red: incident tactics and techniques; blue: countermeasures). github.com/DISARMFoundation/DISARMframeworks
- NATO Command and Control Centre of Excellence (2026). *Information Operations in Russian Command and Control* (short read).
- Bugayova, N., & Stepanenko, K. (2025). *A Primer on Russian Cognitive Warfare*. Institute for the Study of War.
- Voo, J., & Singh, V. V. (2025). *Russia's Information Confrontation Doctrine in Practice (2014–Present): Intent, Evolution and Implications*. International Institute for Strategic Studies.
- Wilde, G., & Sherman, J. (2023). *No Water's Edge: Russia's Information War and Regime Security*. Carnegie Endowment for International Peace.
- U.S. Army TRADOC G-2. *Decisive Action Training Environment (DATE)*, Operational Environment Data Integration Network (ODIN): Olvana, Donovia, Bothnia, Framland, Arnland, Gorgas. odin.tradoc.army.mil/DATE
- Roozenbeek, J., & van der Linden, S. (2020). *Breaking Harmony Square: A game that inoculates against political misinformation* (cited in Rushing et al.).

ATTRIBUTION AND LICENSE

This game and manual adapt concepts, the capability taxonomy, the engagement-effectiveness metric (formula, interaction weights and grade scale) and two sets of innovation attribute scores from the three Rushing et al. papers above. Each paper is licensed under Creative Commons Attribution 4.0 International (CC BY 4.0): creativecommons.org/licenses/by/4.0. Papers: arxiv.org/abs/2603.05222 · arxiv.org/abs/2510.15805 · arxiv.org/abs/2510.15801.

Changes made: the material was adapted into game mechanics and rewritten for play. Engagement grades are mapped to effect multipliers invented for this game; the OODA effects, capability families and cognitive superiority definition are expressed as game rules; 16 of the 18 technology score sets are the designer's own estimates (the paper scores AI Synthetic Media and Humanoid Robots); all scenario, narrative and tutorial text is original.

No endorsement: the authors and their institutions have not reviewed or endorsed this game.

Immersive Briefings voice pack (optional, separate download): Kokoro-82M by hexgrad, kokoro-js and transformers.js (Apache 2.0); ONNX Runtime (MIT); eSpeak NG (GPL-3.0-or-later), whose source is published with the pack.

HOW THIS GAME WAS MADE

Design and direction: Kevin Williamson. Cognitive Terrain is my design: I conceived the game, built its rules from the research, chose every source, defined both opponents and the campaign briefs, set the balance targets and accessibility standards, and reviewed every rule, campaign and page.

Tools: I used AI-assisted development tools to help implement the code, draft text and run balance tests, working from my specifications; I reviewed, revised or rejected everything they produced. The published sources, not the tools, are the authority for any rule that models real-world behaviour. The art, icons and music are drawn and synthesized by the game's own code; no AI image or music generators were used, and the base game runs no AI while you play. The optional Immersive Briefings voice is synthetic speech from an open-source model that runs on your own computer (Settings > Briefings).

Olvana and Donovia, and the countries Bothnia, Framland, Arnland and Gorgas, are fictional training countries from the US Army's ODIN DATE environment; they are used here without endorsement. All eight scenarios, their towns, outlets, people and every event in them are original fiction. Any resemblance to real places or events is coincidental. Donovia's rules follow the NATO, ISW, IISS and Carnegie research above; game effects are abstractions of broad, publicly reported patterns, not claims about any specific operation.

FACILITATOR'S CHECKLIST

BEFORE

- ☐ Play the campaign yourself, end to end [3.0].
- ☐ Pick campaign, difficulty and plain or military terms [2.2].
- ☐ Put the same version on every machine [8.0].
- ☐ Test the projector scheme and text size [2.3].
- ☐ Load and test the voice pack, if used.
- ☐ Print the Quick Reference, one per player or pair.
- ☐ White Cell: plan two or three injects [4.11].
- ☐ Decide: week in review on or off.

DURING

- ☐ Open with the decision loop in five minutes.
- ☐ First game: tutorial on.
- ☐ Point players at the top bar's next crisis.
- ☐ When the other side pulls ahead, read its edge panel together.
- ☐ Note turning points as they happen.
- ☐ Two players: check health bars match after each code.
- ☐ White Cell: say why, in the ruling or note.
- ☐ Keep time; leave a third for the debrief.

AFTER

- ☐ Open the debrief together [5.0].
- ☐ Compare points with the three verdicts.
- ☐ Step the replay to two or three turning points.
- ☐ Ask two questions from Appendix B.
- ☐ Ask what would differ in the real world, and why the model may be wrong.
- ☐ Keep a screenshot of each debrief.

TIMING GUIDE

A solo campaign takes 30–45 minutes; a paired two-player game about an hour. Plan the debrief at a third of the session [2.1].

DEBRIEF QUESTIONS

Two starters for each campaign, tied to its twist and crises. Use them with the general questions in 5.3.

CAMPAIGN	QUESTIONS
Typhoon Ines	The Relief Flotilla tested NOTICE. What did each side need people to see first, and who got there? Who was trusted with the rebuilding by week 16, and what earned it?
The Valmora Vote	Both TRUST areas started weak. What did that do to Official Statements, and how did you adapt? The Leaked Audio tested INTERPRET. What would have helped people judge it?
Collision at Varrin Reef	Both sides' Government Contacts started weak. Who reached the officials first, and what did it pay? Two Videos tested NOTICE. How did each side compete over which clip people saw?
Fever in Port Sabre	Rumours spread faster here. When did a cluster of held nodes start doing the work for one side? The Bad Batch tested TRUST. What protects trust in health advice before a scandal, not after?
Winter Grid	Four dormant narratives were hidden in Blue's mind. When did you find the first, and how? The Price of Heat tested DECIDE. What made the decision hard to take?
The Lake Referendum	Donovia began with a foothold in Blue's INTERPRET. What premise had it planted, and did Blue answer the premise or the stories? Debate Night tested DECIDE. Who set the terms of the debate?
The Silent Survey Ship	Both DECIDE areas started strained. Who acted first under uncertainty, and was acting first right? Three Stories tested INTERPRET. How did competing explanations change the decision?
The Vell Valley	The Protection Claim tested INTERPRET. How did Donovia's framing shape what the valley's events meant? Did a Show of Force work, or blow back? What decided it?

FOR ANY CAMPAIGN

- Which rule in this game would you change to make it more like the real world? What would that change in how it plays?
- Where did the game's numbers feel wrong to you? (Section 10.0 says which are assumptions.)

REFERENCE TABLES

Every player-facing number in one place.

CARDS

CARD	COST	PUSH	COVERS
PLANT STORIES	2	20 (×1.5 on entry)	Target
FRAME THE STORY	2	16; framed	Target
BOT SWARM	3*	12; 7 on links	Target + links
REPOST & REBRAND	2	10; strips watched, prebunked	Target; strips on links too
SHOW OF FORCE	3	16 (no grade)	Any ACT/DECIDE
WATCH TEAM	2	Watched 3 wk	Target + links
PREBUNK	2	Prebunked 3 wk	Target + links
OFFICIAL STATEMENT	2	14 back (Blue 22)	Target
LESSONS LEARNED	3	+6 (Red +3); trains	Target + links

* 2 for Red. Add 1 to every cost while INTERPRET is lost.

FAILURE

Attack	8%, +15% watched, +8% prebunked
Blue into Olvana / Donovia	+2% / +1%
Show of Force	+4%
Official Statement	6%; 26% if TRUST lost
Any move, ACT lost	1 in 4 misfires

GRADES

SCORE	GRADE	PUSH ×
under 3	F	0.25
3 to under 4	E	0.45
4 to under 10	D	0.7
10+	C	1
100+	B	1.2
1,000+	A	1.4
10,000+	A+	1.6

CRISES

Score = $0.6 \times \text{Decision Health} + 0.4 \times \text{tested area}$, -0.08 if that area is lost

SOUND	0.86 or above: 3 points
DELAYED	0.72 to under 0.86: 1 point
ERROR	under 0.72: 0, other side +1

WEEKLY RECOVERY

Recovery	1 + 1.5 a Lessons level (Red 0.75)
Lost node	Half the Lessons part only
Support	+1 if a link leans 50+ your way
Spread	5 per link leaning 50+ the other way

THRESHOLDS

Foothold	10	Held	30
Area lost	half its nodes	Budget cap	18