

THE GAME ODDS SERIES

How to win Professional Catan

*The 36-outcome lattice, placement expected value, the robber's price list,
and a seeded simulation campaign of 265,000 games*

Winfield

Library of Alexandria · An Odisena Imprint
Living Edition 1.0 · July 2026
libraryofalexandria.store

How to win Professional Catan

© 2026 Odisena. All rights reserved.

Written by Winfield.

Published by the Library of Alexandria, an Odisena imprint.

libraryofalexandria.store · libraryofalexandria.shop

ISBN 979-8-951779-42-7 (ebook, PDF)

ISBN 979-8-951779-43-4 (ebook, EPUB)

This is a Living Edition. The mathematics in this book is computed by an open solver and re-verified on every revision. Field results, corrections, and new simulation runs are folded back into later editions, and every edition records what changed. Nothing here is a promise about your next game; everything here is reproducible.

Exact figures in this book are computed in closed form. Simulated figures come from seeded Monte Carlo runs and are labelled as such, with the seed recorded in the appendix. Where no honest computation is possible — most of all in negotiation — the text says so and calls the advice doctrine.

Edition changelog

Living Edition 1.0 — first publication, 2026-07-24.


Contents

A Living Edition strategy monograph from the Library of Alexandria, an Odisena imprint. Written by Winfield. Catan is the first game in this series that can actually be beaten — no house edge, no commission, no rake, just four people and a lattice of 36 dice outcomes. Every probability in this book was computed in an open solver, every simulation was run at a fixed seed, and every claim that cannot be computed is labelled as doctrine rather than dressed up as arithmetic.

- 1 The Pip Map
- 2 Scoring Every Intersection Before the First Roll
- 3 Income Is Not Points
- 4 The Second Settlement
- 5 Harbours and the Trade-Surplus Ledger
- 6 The Robber: Denial Mathematics and Table Politics
- 7 The Twenty-Five-Card Deck
- 8 Tempo: Racing the Ten-Point Clock
- 9 When Eight Never Rolls
- 10 Trading With Humans
- 11 The Field Test
- 12 Appendix: The Engine and the Tables

Chapter 1 — The Pip Map

Catan prints its own probability distribution on the board and most players never read it. The dots under each number token are not decoration; they are the number of ways two dice can produce that total. Six ways make a seven, five ways make a six or an eight, one way makes a two or a twelve. The token is a price tag, and the pips are the price (CATAN base rulebook¹).



The solver behind this book enumerates the lattice exactly, in fractions, never in decimals: $P(7) = 6/36 = 1/6$, $P(6) = P(8) = 5/36$, $P(5) = P(9) = 4/36$, $P(4) = P(10) = 3/36$, $P(3) = P(11) = 2/36$, $P(2) = P(12) = 1/36$. Those eleven fractions sum to one, and the test suite refuses to let the simulator run until they do. This is the anchor benchmark of the whole book, the Catan equivalent of the house edge cross-checks that opened the earlier volumes in this series: if the engine cannot reproduce the numbers printed on the cardboard, nothing else it says can be trusted.

Now count the island. The standard three-to-four-player board has 19 hexes — four forest, four pasture, four fields, three hills, three mountains, and one desert — and 18 number tokens: one 2, one 12, and two each of 3, 4, 5, 6, 8, 9, 10, and 11 (CATAN base rulebook¹). Add the pips and the island carries exactly **58 pips of production**, of which **20 sit on the four red tokens** (the two sixes and the two eights). More than a third of all production on the board is printed in red ink, which is why the red tokens are the most contested cardboard in the hobby.

The geometry is equally finite. The solver builds the board from first principles rather than hard-coding it — the 19 hexes are the axial hexagon of radius two, corners are deduplicated vertices, edges are corner pairs one side apart — and it recovers the printed counts: **19 hexes, 54 intersections, 72 edges**, of which 24 intersections are interior (touching three hexes) and 30 are coastal (touching one or two). Every opening decision in Catan is a choice among at most 54 objects, and only 24 of them can ever touch three

¹ CATAN base rulebook. https://www.catan.com/sites/default/files/2021-06/catan_base_rules_2020_200707.pdf

producing hexes.

One last exact figure belongs in this chapter, because the rest of the book keeps colliding with it. A seven arrives on average once every six rolls, and it does two things: nothing produces, and every player holding more than seven resource cards discards half, rounded down (CATAN base rulebook¹). The robber is not an event that happens to other people. Over a 70-roll game the lattice predicts about 11.7 sevens, and the discard rule turns every large hand into a taxed position — a tax computed to the card in Chapter 6.

¹ CATAN base rulebook. https://www.catan.com/sites/default/files/2021-06/catan_base_rules_2020_200707.pdf

Chapter 2 — Scoring Every Intersection Before the First Roll

An intersection's production is the sum of its neighbours' pips, and its expected yield is that sum over 36. The doctrinal example on BoardGameGeek is the 6-5-9 corner: five pips plus four plus four is thirteen, so the spot produces on 13 of 36 rolls — a little over a third of them (Rempton Games, BoardGameGeek¹). A settlement takes one card per matching hex; a city takes two, so a city on the same corner expects 13/18 of a card per roll. That is the entire mathematics of placement EV, and it is exact.

Community doctrine has converged on the same arithmetic expressed as a checklist: count pips rather than numbers, treat 10 pips as the minimum for a serious spot and 12 or more as top-tier, cover as many of the five resources as possible across the two openers, and weight grain and ore because the late game runs on cities and development cards (Ore You Serious, starting-settlement guide²). Tournament-level guidance adds board-reading heuristics on top — where the desert sits determines whether the strong plays are the inland three-hex corners or the coastal two-hex ones (Catan placement walkthrough³).

The solver's answer to how much a spot is worth is blunt. Across 100,000 randomly dealt standard boards at seed 20260724, a drafter alone on the island who simply takes the two highest-pip legal intersections averages **25.67 pips**, which is 0.713 cards per roll, or **49.9 cards over a 70-roll game** (simulated, ± 7.6 cards of standard deviation). That is the reference income against which every other doctrine in this book is measured.

Two structural facts fall out of the same run and neither is obvious at the table. First, the pip-maximiser is not resource-poor by accident: its two settlements touch **3.79 distinct resources** on average, because high-pip corners tend to be interior and interior corners touch three different hexes. Second, harbours are essentially never worth taking on pips alone — the pip-max drafter lands on a harbour intersection in fewer than one game in two thousand. A harbour corner is coastal, a coastal corner touches at most two hexes, and two hexes cannot outbid three. Any port plan is therefore a deliberate purchase, priced in Chapter 5.

¹ Rempton Games, BoardGameGeek. <https://boardgamegeek.com/blogpost/121267/the-ultimate-catan-strategy-guide-top-tips-to-win>

² Ore You Serious, starting-settlement guide. <https://oreyouserious.com/blog/starting-settlement-strategy>

³ Catan placement walkthrough. <https://www.youtube.com/watch?v=wHhvztIH7A>

Chapter 3 — Income Is Not Points

The cost table is the reason raw pips mislead. A road costs one brick and one lumber; a settlement one brick, one lumber, one wool, and one grain; a city two grain and three ore; a development card one wool, one grain, and one ore (CATAN base rulebook¹). Grain appears in three of the four recipes. Ore appears in two, in quantity. Wool appears in two, once each.

Take a canonical route from the two opening settlements to ten victory points: three more settlements, three cities, six roads, and four development cards. The solver adds up the recipes and gets a shopping list of **51 resource cards**: 13 grain, 13 ore, 9 brick, 9 lumber, and **7 wool**. Then compare that demand with what the island prints — 18 producing hexes split four lumber, four wool, four grain, three brick, three ore — and divide demand share by supply share. Normalised to a mean of one, the scarcity weights come out as **ore 1.49, grain 1.11, brick 1.03, lumber 0.77, wool 0.60**.

A pip of ore is worth roughly two and a half pips of wool to a player who intends to build cities. That is not a probability; it is a modelling choice, and this book says so plainly — change the build plan and the weights move. But the direction is not in doubt, and it is exactly the direction community doctrine already points: prioritise grain and ore, treat wool as a trading commodity rather than a build resource (Ore You Serious²).

Now run the same 100,000 boards with a drafter who maximises scarcity-weighted income instead of raw pips. It ends up with **24.20 pips** — a pip and a half poorer in cards, **47.1 cards per game** against the pip-maximiser's 49.9 — but its weighted income is **52.7 against 48.5**, the largest weighted income of any doctrine tested. It out-earns the pip-maximiser in weighted terms in 21.4 percent of games outright while collecting fewer physical cards. The resource mix, not the card count, is what the cost table actually buys.

Chapter 8 will show that this weighted advantage does not automatically become victories. Holding both facts at once is the discipline this book is trying to teach.

¹ CATAN base rulebook. https://www.catan.com/sites/default/files/2021-06/catan_base_rules_2020_200707.pdf

² Ore You Serious. <https://oreyouserious.com/blog/starting-settlement-strategy>

Chapter 4 — The Second Settlement

Catan's opening is a snake draft: seats pick in order 1-2-3-4, then 4-3-2-1, so the last seat picks twice in a row and the first seat waits the longest for its second spot (Ore You Serious¹). Tournament rules formalise the clock around it: five minutes of open board study and discussion, then two minutes per placement and three minutes for a double placement (CATAN Championship Tournament Rules, revised July 2025²).

The doctrinal advice for the second pick is to fix what the first one missed rather than repeat it. The solver can price that advice. A **diversity-max** drafter — pips plus a two-pip-equivalent premium for every resource type not yet held — gives up almost nothing in production and buys a materially wider base: **25.18 pips against 25.67, 49.0 cards per game against 49.9, and 4.46 distinct resources against 3.79**. Half a card per game per pip of production, in exchange for two-thirds of an extra resource type. In 17.2 percent of the 100,000 games the diversity drafter out-earned the pip-maximiser outright even in raw cards, and in another 51.8 percent the two doctrines earned identically because they had picked the same corners.

Why the wide base matters is a build-tree argument, not an income argument. Every recipe in the game needs at least two resources and three of the four need three or more. A drafter missing a resource entirely pays the bank's exchange rate on every unit of it forever, which Chapter 5 prices at one to two extra cards per acquisition, and a drafter missing grain or ore is locked out of cities and development cards until it fixes the hole. The pip lead is worth less than the tax it incurs.

The engine test is in Chapter 8, and it is emphatic: in 20,000 four-player games, diversity-max won **44.4 percent** of its games against three rival doctrines, against pip-max's 25.2 percent. Same dice, same build policy, same island. The only difference was which two corners the doctrine chose and how many resources they covered.

¹ Ore You Serious. <https://oreyouserious.com/blog/starting-settlement-strategy>

² CATAN Championship Tournament Rules, revised July 2025. <https://catanevents.com/hosts/tournament-rules/>


Chapter 5 — Harbours and the Trade-Surplus Ledger

Harbours are the most misunderstood real estate on the island. The standard board carries nine: four generic 3:1 harbours and five 2:1 harbours, one per resource (CATAN base rulebook¹). Note that tournament play does not use port tokens at all — the printed harbours on the frame are the only trade rates on the table (CATAN Championship Tournament Rules²).

The exact part of the ledger is trivial. Upgrading from the bank's 4:1 rate to a generic 3:1 saves one card per acquired card; upgrading to a matched 2:1 saves two. Buying a city entirely through the bank at 4:1 costs 20 surplus cards; through a matched 2:1 harbour it costs 10.

The interesting part is what that saving does to the clock. The solver ran 40,000 seeded games in which a pip-max placement produced from an empty hand until a target became affordable, with identical boards and identical dice across three trade regimes. Mean rolls to the first affordable build (simulated, seed 20260726):

Target	Bank only (4:1)	Generic 3:1	Matched 2:1
Settlement	15.7	12.8	12.0
City	21.2	17.4	16.5
Development card	13.3	10.8	9.9



A generic harbour pulls three to four rolls out of every build cycle; a matched 2:1 pulls four to five. Over a game that is real tempo, and it is the same structural finding an independent browser-based Monte Carlo model reports — that harbour access measurably compresses the distribution of turns to a city (Catan

¹ CATAN base rulebook. https://www.catan.com/sites/default/files/2021-06/catan_base_rules_2020_200707.pdf

² CATAN Championship Tournament Rules. <https://catanevents.com/hosts/tournament-rules/>

³ Catan Monte Carlo, Columbia. <https://www.columbia.edu/~jsl2239/catan.html>

Monte Carlo, Columbia³). That model explicitly excludes the robber, development cards, and player trading, so it is used here as a directional check, not a numeric benchmark.

Then comes the bill. A port plan is not free, because harbour corners are coastal. The solver's **port-oriented** drafter, which is required to take one harbour intersection and then plays pips, averages **21.69 pips and 42.2 cards per game — 7.7 cards per game less than the pip-maximiser**, a 15 percent income cut, and only 3.46 resource types. It beats the pip-maximiser on raw income in 5.8 percent of games. In the tournament it won **14.2 percent** of its games, the worst of the four doctrines tested.

The honest conclusion is the one the community reached by feel: ports break ties, they do not override pips (Ore You Serious¹). A harbour that matches a resource you already produce heavily is a genuine accelerator; a harbour bought with three pips of production is a fifteen-percent pay cut for a discount card you may not need.

¹ Ore You Serious. <https://oreyouserious.com/blog/starting-settlement-strategy>

Chapter 6 — The Robber: Denial Mathematics and Table Politics

The robber has three separable prices, and confusing them is how most tables mis-play it.

Denial. Parking the robber on a hex stops that hex's production for every building on it. The rate is exact: pips over 36 per building per roll, doubled for a city. A city on an eight loses 10/36 of a card per roll — **1.67 cards over the six rolls a robber typically sits before the next seven**. That is the real currency of the robber, and it dwarfs the theatrical part.


The steal. One random card from one player. One card. It is the part everyone remembers and the smallest number in the chapter.

The discard tax. A seven arrives once per six rolls, and every hand above seven cards sheds half, rounded down (CATAN base rulebook¹). The solver tabulates the tax exactly: a hand of 8 or 9 cards loses **0.67 cards per roll** in expectation, a hand of 12 or 13 loses **1.00**, a hand of 16 loses **1.33**. A player sitting on twelve cards is paying a card per roll for the privilege. Seven cards is free. The rule is a wealth tax with a published rate, and the correct response is to spend down before the roll rather than to admire the hand.

Targeting is where mathematics ends and politics begins. The solver sampled 200,000 random table states — three opponents with a pip total, a building count, and a hand size — and compared four doctrines. Over a six-roll block, mean cards denied plus stolen (simulated, seed 20260733):

Doctrine	Cards denied	Times it hit the leader
Maximum denial	3.27	33.1%
Leader first (1.5-card premium)	2.97	76.7%
Grudge (same target regardless)	2.39	33.4%
Random	2.39	33.3%

¹ CATAN base rulebook. https://www.catan.com/sites/default/files/2021-06/catan_base_rules_2020_200707.pdf



The robber's price list

Three readings. First, playing the robber for maximum denial is worth about **0.87 cards per placement** over random or grudge play — small per event, roughly ten cards over a full game's sevens, which is a fifth of a placement's entire income. Second, leader-targeting is cheap: insisting on the leader costs only 0.30 cards of denial against pure denial-maximising, and it buys a 77 percent hit rate on the person who is actually about to win. Denial is a card metric; winning is a points metric, and paying a third of a card to slow the leader is usually correct. Third, the grudge is indistinguishable from random. Robbing the same person out of spite is precisely as productive as not thinking at all — a result that matches the standing community advice to rob the player who acts before you when there is no clear leader, and to avoid revenge cycles (Rempton Games, BoardGameGeek¹).

One structural defence deserves emphasis. Do not build a settlement whose value is concentrated in a single hex; a robber on that hex switches you off. Corners that share hexes with opponents are better protected, because the robber cannot sit there without hurting someone who can move it (Rempton Games, BoardGameGeek¹).

¹ Rempton Games, BoardGameGeek. <https://boardgamegeek.com/blogpost/121267/the-ultimate-catan-strategy-guide-top-tips-to-win>

Chapter 7 — The Twenty-Five-Card Deck

The development deck is completely knowable, so there is no excuse for playing it on vibes. It holds 25 cards: **14 knights, 6 progress cards (2 road building, 2 year of plenty, 2 monopoly), and 5 victory-point cards** (CATAN base rulebook¹; World of Catan wiki²; BoardGameArena rules summary³).

Exact consequences, all hypergeometric because drawing is without replacement:

- The first card off a full deck is a knight with probability $14/25 = 56.0\%$, a victory point with probability $1/5 = 20.0\%$, and a progress card with probability 24.0% .
- Three or more knights — Largest Army's requirement — appear in the first three buys with probability 15.8%, in six buys **79.1%**, in eight buys 95.7%.
- The expected number of buys to reach the third knight is exactly **5.2**.
- Four buys yield **0.80 victory-point cards** and, through Largest Army, **1.59 expected points in total** for 12 resource cards spent: about **0.13 points per card**. Six buys yield 2.78 points for 18 cards, or 0.155 points per card — the best rate the deck offers.

Set that against the alternative uses of the same cards. A city costs five cards for one net point and doubles a corner's production forever. A settlement costs four cards for a point and opens new production. On raw points per card, the deck at six buys (0.155) is competitive with a city (0.20 points per card, plus the income), and it is the only source of points that cannot be blocked by the board running out of legal corners.

The practical doctrine follows from the numbers rather than from taste. Buy development cards when your production includes wool, grain and ore — the deck's recipe is exactly the surplus a grain-and-ore position generates. Buy them when the board is closed and there is nowhere left to settle. Do not buy them as a substitute for a city you can afford, because the city pays twice. And keep count of what others have bought: five of the 25 cards are silent points, so an opponent with three unrevealed cards may already be at nine (Rempton Games, BoardGameGeek⁴).

¹ CATAN base rulebook. https://www.catan.com/sites/default/files/2021-06/catan_base_rules_2020_200707.pdf

² World of Catan wiki. https://catan.fandom.com/wiki/Development_card

³ BoardGameArena rules summary. <https://en.doc.boardgamearena.com/Gamehelpcatan>

⁴ Rempton Games, BoardGameGeek. <https://boardgamegeek.com/blogpost/121267/the-ultimate-catan-strategy-guide-top-tips-to-win>


Chapter 8 — Tempo: Racing the Ten-Point Clock

Ten points ends the game on the winner's own turn (CATAN base rulebook¹), which makes Catan a race, and races reward tempo rather than wealth. The solver's tournament is the cleanest evidence in this book.

Twenty thousand four-player games, seed 20260728, randomised standard boards. Four agents share one identical build policy — city first, then settlement, then development card, then road, with bank and harbour trades — and differ **only** in opening-placement doctrine. Seat assignments rotate every game so seating order cancels: the observed win share by seat ran 24.0, 24.8, 25.3, 25.9 percent, within noise of the 25 percent par that a fair four-way field must produce.

Doctrine	Win share	Mean final VP	Mean cards/game (income run)
diversity-max	44.4%	7.93	49.0
pip-max	25.2%	6.63	49.9
scarcity-max	16.2%	6.23	47.1
port-oriented	14.2%	5.86	42.2

Income is not victory: four placement doctrines



Income is not victory

The ranking is not the income ranking. The pip-maximiser earns the most cards and finishes second. The scarcity-maximiser has the highest weighted income of all four and finishes third. The diversity drafter, third in raw cards, wins nearly twice as often as anyone else. The reason is tempo: a wide resource base converts income into buildings *now*, without waiting for a bank trade or a partner, and buildings are what the clock counts. Concentrated income is money that arrives in the wrong denominations.

Two honest caveats, both structural. The engine does not allow player-to-player trading, which is the single largest simplification in this book; a specialised, resource-poor position is much stronger at a table

¹ CATAN base rulebook. https://www.catan.com/sites/default/files/2021-06/catan_base_rules_2020_200707.pdf

where you can trade than at a table where you cannot, so the port and scarcity doctrines are penalised here relative to real play. And Largest Army is credited on knights bought rather than played, while Longest Road is approximated by road count. The direction of the result — breadth beats depth for tempo — is robust across seeds; the exact percentages belong to this engine and no other.

For calibration against the wider literature: published Monte-Carlo tree-search agents reach roughly parity with the JSettlers heuristic bots at 1,000 simulations per move (27 percent win rate against a 25 percent par) and clearly beat them at 10,000 (49 percent), while their authors report that a strong human still beat the agent — partly because the agent over-favoured the ore-and-wheat plan and under-built settlements (Szita, Chaslot & Spronck, Monte-Carlo Tree Search in Settlers of Catan, ACG 2012¹). An expensive search engine landing on the same critique that this solver's placement tournament produces — build more settlements, do not over-invest in one narrow plan — is the closest thing to independent confirmation available in the published literature.

¹ Szita, Chaslot & Spronck, Monte-Carlo Tree Search in Settlers of Catan, ACG 2012. <https://spronck.net/pubs/ACG12Szita.pdf>

Chapter 9 — When Eight Never Rolls

Every Catan player has a story about the game where the eight never came. The lattice has an answer, and it is less comforting than "bad luck" and more useful.

The solver dealt 100,000 boards at seed 20260725, drafted a pip-max pair and a deliberately weaker pair at least four pips worse on each board, then dealt both the same 70 rolls. Results (simulated):

- The strong pair earned 49.9 cards on average against the weak pair's 42.2 — a 7.7-card edge over four pips, or roughly 1.9 cards per pip per game.
- The strong pair nevertheless **lost the game's income race in 6.4 percent** of games and tied in another 2.2 percent. A four-pip advantage is beaten by the dice about one game in fifteen.
- The strong pair finished **below its own expectation in 50.4 percent** of games — as it must, since the median of a near-symmetric distribution sits at the mean — **more than 10 percent below expectation in 22.9 percent**, and **more than 20 percent below in 7.1 percent**. One game in fourteen, the best board on the island produces a fifth less than it promised.
- The coefficient of variation of a 70-roll income is **0.152**. Income noise of fifteen percent is the normal weather of a Catan game, not a scandal.
- An eight is rolled fewer times than a ten in **13.1 percent** of 70-roll games, and produces fewer than six times — against an expectation of 9.7 — in **6.5 percent**.

That last figure is the tilt chapter in one line. Roughly one game in fifteen, the red eight underperforms a mediocre ten. It is not the board's fault, it is not the dice's fault, and it does not indicate that pips were the wrong criterion. It is the 36-outcome lattice behaving exactly as specified over a sample of seventy.

The practical consequences are three. First, do not re-plan strategy from a ten-roll sample; the sample is too small to contain information. Second, since fifteen percent income noise is the standing condition, prefer positions that are robust to a single number going quiet — which is another argument for breadth, and the same argument the robber chapter made from a different direction. Third, when a game is lost to a dry number, log it. The field-test protocol in Chapter 11 exists precisely so that the reader's own dice can be checked against the lattice rather than against memory, which keeps the fistfuls and forgets the drip.

Chapter 10 — Trading With Humans

Everything up to here is computable. This chapter is not, and pretending otherwise would violate the standard this series is built on. The solver's engine has no player-to-player trading at all, because there is no honest way to simulate a negotiation between four people who can see each other's faces. What follows is doctrine, labelled as such, informed by the mathematics rather than derived from it.

Four principles survive contact with the numbers.

Price by scarcity weights, not by count. The cost table says ore and grain are the expensive resources (weights 1.49 and 1.11) and wool is the cheap one (0.60). A one-for-one wool-for-ore trade is not a fair trade; it is a subsidy. If you must overpay, overpay in wool.

Never trade a leader into a build. A trade that hands the leader the last card of a city has cost you more than the card. The robber chapter measured the value of slowing the leader at about 0.30 cards of denial for a 77 percent hit rate; a trade refusal is free by comparison, and community doctrine agrees that the endgame is where trades should dry up (Rempton Games, BoardGameGeek¹).

Trade on your own turn. Cards acquired on your turn get spent immediately; cards acquired on someone else's turn sit in your hand until the next seven taxes them at the rate computed in Chapter 6 (Rempton Games, BoardGameGeek¹).

Reputation is a resource with a real exchange rate. A player nobody will trade with is playing the solver's engine — the no-trading version — and the solver's engine takes about 105 player-turns to reach ten points. Being tradeable is worth more than any single trade you can win.

The kingmaker problem deserves its own paragraph, because tournaments have institutionalised it. Under official rules, preliminary rounds are ranked by number of victories, with total victory points and then Victory Point Percentage — your score divided by the total points at your table — as tiebreakers (CATAN Championship Tournament Rules²). That structure means a player who cannot win still has something to play for, which is exactly the incentive that keeps a table honest. Home games have no such mechanism, and the field-test protocol in this book substitutes the mildest possible one: the winner is exempt from cleanup, and nobody is playing for anything larger.

For scale on what the professional version of this game looks like: the CATAN World Championship runs every two years, sending English-language national champions to the international final (CATAN Championship programme³). The 2025 edition in Stuttgart ran four preliminary games on the Saturday, then semifinals and a four-player final on the Sunday, and was won by Kohulan Narendran of Canada ahead of Guido Grossi of Italy and Raúl Jiménez of Mexico (CATAN news⁴; event page⁵), with roughly 90 players from 60 countries in the field (Ringbote report⁶). The 2026 United States National Championship is scheduled for Chicago, 25–27 September (CATAN US Nationals 2026⁷), fed by local qualifiers and regional preliminaries into a 16-player national final (US championship structure⁸). Every one of those tables is a negotiation, on a five-minute study clock, with a two-minute limit per placement (CATAN Championship Tournament Rules²).

¹ Rempton Games, BoardGameGeek. <https://boardgamegeek.com/blogpost/121267/the-ultimate-catan-strategy-guide-top-tips-to-win>

² CATAN Championship Tournament Rules. <https://catanevents.com/hosts/tournament-rules/>

³ CATAN Championship programme. <https://catanevents.com/players/champ/>

⁴ CATAN news. <https://www.catan.com/catan-fans/news/catan-worldchampionship-2025-stuttgart>

⁵ event page. <https://www.catan.com/championships/catan-world-championship-2025>

⁶ Ringbote report. <https://www.ringbote.de/aktuelles/catan-wm-2025-in-stuttgart>

⁷ CATAN US Nationals 2026. <https://www.catan.com/championships/catan-us-national-championship-2026>

⁸ US championship structure. <https://catanevents.com/players/champ/cnc-usa/>

Chapter 11 — The Field Test


The Living Edition mechanism of this series is the field test: the solver predicts observable frequencies before play, the log sheet records what actually happened, and the next revision reports the comparison. Catan's version is the cheapest field test in the series, because the stakes are structural rather than financial.

Protocol. Ten games, four players, standard base-game board, home table. Randomised board each game. No expansions. Stakes: the winner is exempt from cleanup. Nothing else is wagered.

What gets logged. Every roll, tallied by number — roughly 70 rolls per game, about 700 across the protocol. Each player's two opening placements with their pip totals and resource types. Every robber placement, with the target and the hex's pips. Every trade: what for what, with whom, on whose turn. Development cards bought per player. Final victory points, and how they were composed.

Three predictions, computed before the first roll.

1. **Roll frequencies.** Over 700 rolls the lattice predicts 116.7 sevens with a standard deviation of 9.86, so a 95 percent interval of **97 to 136 sevens**. Sixes and eights: 97.2 each. Twos and twelves: 19.4 each. Any observation inside those intervals is the dice behaving; anything outside is worth re-counting before it is worth a theory. 2. **Income by pip total.** A settlement pair totalling 25 pips should produce **0.69 cards per roll**, about **48.6 cards per 70-roll game** — measured as cards received, before robbery and discards. Note that in a four-player game every roll produces for everyone, so a player's cards accumulate on other people's turns too. 3. **Opening pips versus final points.** This is the prediction most likely to embarrass its author. In the engine, with all four seats using identical doctrine and identical build policy, the correlation between opening pip total and final victory points is only $r = 0.053$ across 20,000 seats — statistically real, practically tiny. But the win share by opening pip total climbs steadily: **18.5 percent at 18 pips, 22.6 at 20, 25.6 at 21, 26.7 at 22, and 29.6 percent at 24 pips**. The prediction for the field test is therefore deliberately modest: pip totals will barely correlate with final scores in ten games, and the pip-rich seats will win a little more than a quarter of the time. Ten games cannot resolve a difference that small, which is itself the lesson.



Debrief questions. Were the seven counts inside the interval? Did the observed cards-per-roll match the pip prediction within the fifteen-percent noise band of Chapter 9? Did the highest-pip opening win more often than the widest-resource opening — and if not, does the tournament result in Chapter 8 explain why? Which robber placements were made for denial and which for grudge, and what did the grudges cost? Every answer feeds the next Living Edition revision.

Chapter 12 — Appendix: The Engine and the Tables

The engine. ``gamestats.catan`` is a Python module in the open series codebase, with a pytest suite that must pass before any simulation output is used. It contains: ``dice`` (the exact 36-outcome lattice in Fractions), ``board`` (geometry built from first principles — 19 hexes, 54 intersections, 72 edges — plus randomised standard boards from the printed terrain and token multisets), ``econ`` (the printed cost table, the canonical build plan, and the scarcity weights), ``devdeck`` (the 25-card deck, hypergeometric draws), ``robber`` (denial rates, the discard tax, target scoring), ``placement`` (four placement doctrines and the snake draft), and ``simulate`` (the income campaign, the variance study, the build-time study, the robber study, and the four-player game engine).

The benchmark gate. Before any simulation was trusted, the suite verified against published, printed figures: $P(7) = 6/36$; $P(6) = P(8) = 5/36$; $P(2) = P(12) = 1/36$; the lattice summing to 36 ways; the 18-token multiset with one 2 and one 12; the 58-pip board with 20 red pips; the 19/54/72 geometry; the printed cost table; the 25-card deck as 14 knights, 6 progress, 5 victory points; the first-draw knight probability of 14/25; and the discard threshold of more than seven cards. All 46 Catan tests pass, and so do the 80 pre-existing tests of the other five game modules.

Table A — the lattice (exact).

Total	Ways	Probability	One in
2	1	2.78%	36.0
3	2	5.56%	18.0
4	3	8.33%	12.0
5	4	11.11%	9.0
6	5	13.89%	7.2
7	6	16.67%	6.0
8	5	13.89%	7.2
9	4	11.11%	9.0
10	3	8.33%	12.0
11	2	5.56%	18.0
12	1	2.78%	36.0

Table B — scarcity weights (exact, given the canonical plan).

Resource	Cards demanded	Demand share	Supply share	Weight
Ore	13	25.5%	16.7%	1.49
Grain	13	25.5%	22.2%	1.11

Resource	Cards demanded	Demand share	Supply share	Weight
Brick	9	17.6%	16.7%	1.03
Lumber	9	17.6%	22.2%	0.77
Wool	7	13.7%	22.2%	0.60

Table C — placement doctrines (simulated; 100,000 boards, seed 20260724; 20,000 games, seed 20260728).

Doctrine	Pips	Cards/game	Weighted income	Resource types	Win share
pip-max	25.67	49.90	48.47	3.79	25.2%
diversity-max	25.18	48.96	48.00	4.46	44.4%
scarcity-max	24.20	47.05	52.70	3.66	16.2%
port-oriented	21.69	42.16	40.97	3.46	14.2%

Table D — the development deck (exact).

Quantity	Value
Deck size	25 cards
Knights	14 (56.0% of the deck)
Progress cards	6 (24.0%)
Victory-point cards	5 (20.0%)
P(3+ knights in 6 buys)	79.1%
Expected buys to third knight	5.2
Points from 6 buys	2.78 (0.155 per card spent)

Table E — the robber (exact rates, simulated policies).

Quantity	Value
P(seven) per roll	1/6
Denial, city on an 8, six rolls	1.67 cards
Discard tax, hand of 8–9	0.67 cards/roll
Discard tax, hand of 12–13	1.00 cards/roll
Max-denial policy, six-roll block	3.27 cards
Leader-first policy	2.97 cards, 76.7% leader hit rate
Grudge or random policy	2.39 cards

Reproduction. From the `gamestats` root: `PYTHONPATH=src pytest -q` runs the suite; `PYTHONPATH=src python3 scripts/run\_catan\_campaign.py` reproduces every simulated figure in this book at the seeds named above and writes them to `results/catan\_\*.json`; `PYTHONPATH=src python3 scripts/make\_catan\_charts.py` redraws the charts.

What this book does not claim. It does not claim a win rate for any doctrine at a real table, because the engine cannot trade with humans and humans are most of Catan. It does not claim that the scarcity weights are canonical — they follow from one stated build plan. It does not report tournament win rates by placement, because no public dataset of tournament placements and outcomes exists. And it does not claim the pip map can be beaten. It cannot. It can only be read, and most of the table is not reading it.

Sources

CATAN base game rulebook (2020) —

https://www.catan.com/sites/default/files/2021-06/catan_base_rules_2020_200707.pdf

CATAN — The Game rulebook (2025 printing) — [https://www.catan.com/sites/default/files/2025-03/CN3081%20CATAN%E2%80%93The%20Game%20Rulebook%20secure%20\(1\).pdf](https://www.catan.com/sites/default/files/2025-03/CN3081%20CATAN%E2%80%93The%20Game%20Rulebook%20secure%20(1).pdf)

CATAN Championship Tournament Rules, revised July 2025 — <https://catanevents.com/hosts/tournament-rules/>

CATAN Championship programme — <https://catanevents.com/players/champ/>

CATAN Championship, United States — <https://catanevents.com/players/champ/cnc-usa/>

CATAN US National Championship 2026 —

<https://www.catan.com/championships/catan-us-national-championship-2026>

CATAN World Championship 2025, Stuttgart — <https://www.catan.com/championships/catan-world-championship-2025>

CATAN news, 2025 World Championship result —

<https://www.catan.com/catan-fans/news/catan-worldchampionship-2025-stuttgart>

Ringbote, CATAN WM 2025 report — <https://www.ringbote.de/aktuelles/catan-wm-2025-in-stuttgart>

Rempton Games, "The Ultimate Catan Strategy Guide" (BoardGameGeek) —

<https://boardgamegeek.com/blogpost/121267/the-ultimate-catan-strategy-guide-top-tips-to-win>

Ore You Serious, "How to pick your starting settlements in Catan" —

<https://oreyouserious.com/blog/starting-settlement-strategy>

Catan placement walkthrough (tournament-level heuristics) — <https://www.youtube.com/watch?v=wHhvztiHM7A>

World of Catan wiki, development cards — https://catan.fandom.com/wiki/Development_card

BoardGameArena rules summary — <https://en.doc.boardgamearena.com/Gamehelpcatan>

Szita, Chaslot & Spronck, "Monte-Carlo Tree Search in Settlers of Catan" (ACG 2012) —

<https://spronck.net/pubs/ACG12Szita.pdf>

Lee, "Catan Monte Carlo" (Columbia) — <https://www.columbia.edu/~jsl2239/catan.html>

Field-Test Log Sheet

One sheet per game. Ten games, four players, standard base-game board, home table, winner exempt from cleanup. Photocopy or redraw freely; the point is that the tallies exist before the theories do.

Game number _____ Date _____ Seats (in order) _____

Roll tally

Total	2	3	4	5	6	7	8	9	10	11	12
Tally											
Count											
Expected in 70 rolls	1.9	3.9	5.8	7.8	9.7	11.7	9.7	7.8	5.8	3.9	1.9

Openings and outcome

Player	Opening 1 (pips / hexes)	Opening 2 (pips / hexes)	Total pips	Resource types	Dev cards	Final VP

Robber placements and trades

#	Robber hex / pips	Target	Reason (denial / leader / grudge)	Trade made (what for what, with whom)

Debrief: sevens observed _____ (95% interval across ten games: 97–136) · cards per roll at 25 pips, predicted 0.69 _____ · did the widest-resource opening beat the highest-pip opening? _____

About the Living Edition

A Living Edition is a book that expects to be wrong about something and builds the correction into its own machinery. Every number in this volume is either computed in closed form from the printed rules of the game or produced by a seeded simulation whose code, seeds, and output files are named in the appendix. Anyone with the repository can rerun the campaign and get the same figures to the last decimal.

That is also the limit of the guarantee. This book does not know what your table will do. It cannot simulate a negotiation, it does not trade with humans, and it says so in Chapter 10 rather than hiding the gap behind a percentage. Where the engine simplifies — no player-to-player trading, Largest Army credited on knights bought, Longest Road approximated by road count, harbours randomised across coastal edges — the simplification is stated in the text and in the module docstring.

The field test in Chapter 11 is how the next edition gets written. Ten logged games at a home table cannot settle a three-percentage-point question, and the book predicts in advance that it will not. What ten games can do is check the roll tallies against the lattice, check income against the pip prediction, and catch any place where the solver and the cardboard disagree. If they disagree, the solver is wrong, and Living Edition 1.1 will say where.

Series: The Game Odds Series from the Library of Alexandria — Blackjack, Poker, Roulette, Craps, Baccarat, Rummikub, Catan.

Corrections and field logs: libraryofalexandria.store

Edition record

Edition	Date	Change
Living Edition 1.0	2026-07-24	First publication. Exact lattice, deck, cost-table and geometry benchmarks verified; 265,000-game seeded campaign.