

Balancing Contact Front!

Monte Carlo and the 38% Solution. On simulation, difficulty, and the mathematics of meaningful loss.

b1tr0n1n // 2026

THE PROBLEM

Contact Front! started as a feel. A WW2-themed solo card game played on a standard 54-card deck, where every draw is a tactical decision and every round could be your last. The mechanics felt right in my hands — the doctrine system gave strategic depth, the scar mechanics created narrative weight, and the campaign mode turned individual missions into something you cared about losing.

But feel is not balance. I could play fifty games and think the difficulty was fair, but I'd be testing against one brain — mine. I'd learned the patterns. I'd internalized the optimal lines. What I needed was a way to test the system against ten thousand brains that played every possible way, from brilliant to reckless.

I needed Monte Carlo simulation.

...

WHAT MONTE CARLO ACTUALLY DOES

Monte Carlo simulation is simple in concept: run the game thousands of times with randomized decisions and measure the outcomes. No AI, no optimization — just brute statistical force applied to your system. You feed it the rules, the deck composition, and a decision model, and it tells you what happens when probability plays itself out across a large enough sample.

For Contact Front!, I built a Python simulation that modeled the full game loop: draw phase, tactical deployment, combat resolution, morale checks, and doctrine activation. Each simulated player made decisions using weighted randomness — sometimes optimal, sometimes suboptimal, roughly approximating the decision quality of a real human learning the game.

Then I ran it ten thousand times.

...

FINDING THE NUMBER

The first pass came back at 52% win rate. Too easy. A game you win more than half the time doesn't create the tension that makes solo gaming worth the table space. You need to feel like the system is hunting you.

I started adjusting: enemy scaling, draw limits, morale thresholds, doctrine costs. Each change rippled through the simulation. Tighten one knob and win rates drop, but so does the feeling that your decisions matter — if the game kills you regardless of what you do, that's not difficulty, that's nihilism.

The target emerged through iteration: 38%. Just over one in three. High enough that a skilled player feels rewarded for good decisions. Low enough that every win feels earned and every loss feels like it could have gone differently. The gap between 38% and 50% is where all the interesting decisions live.

A game that's too easy teaches you nothing. A game that's too hard teaches you to quit. The sweet spot is where you lose often enough to respect the system and win just enough to believe you can beat it.

...

WHAT THE DATA REVEALED

The simulation didn't just find the win rate — it exposed the architecture of difficulty. I could see which cards were dead draws, which doctrine combinations were dominant, and where the decision tree collapsed into obvious choices. Data like that is invisible from inside the game. You need distance.

Three specific findings reshaped the final design. First, the original morale system was too binary — you were either fine or dead, with no middle ground. I added degradation states that created a slow bleed rather than a cliff. Second, two doctrine cards were accounting for a disproportionate share of wins. I didn't neef them — I increased their cost, turning them into high-risk commitments rather than free advantages. Third, late-game draws were statistically irrelevant because the outcome was usually decided by turn four. I restructured the draw economy to keep late turns meaningful.

None of these changes came from intuition. They came from letting the numbers speak.

...

THE LESSON

Game design is systems design, and systems design requires testing at scale. Your gut tells you what feels right. Simulation tells you what is right. The two don't always agree, and when they don't, the data wins.

Contact Front! ships at 38% not because that number sounds good, but because ten thousand simulated games proved it's the point where difficulty becomes meaningful. Where losing teaches you something and winning means you actually learned it.

That's the only kind of balance worth shipping.

...

b1tr0n1n — Builder. Thinker. Maker. — b1tr0n1n.github.io
