CAKE MURDER ADVENTURE — Difficulty Caps

2026-04-11

CAKE MURDER ADVENTURE — Difficulty Caps

A record of every difficulty parameter cap applied to `cake-murder-adventure.html` and the Python oracle (`~/git/frank/manifold.py`). Both files are kept in sync — the sim trains on the same physics the game runs.

Why caps exist

The original difficulty curve scaled all parameters linearly from ch1 to ch101 with no upper bounds. At high chapters several parameters crossed into mechanically unbeatable territory — the knife could outrun the cake, or the drop was too fast to react to regardless of player skill. The caps define a "hard but beatable" ceiling rather than a punishment function.

The machine learning bot and the human player face the same physics. If the bot can't beat it, a human can't be expected to either.

Cap timeline

Round 1 — warnTrackSpeed, lateralDrift, warningTime, bladeScale

Date: 2026-04-11
Commit context: difficulty analysis after bot plateaued at 52/101

param	original	cap	hits at
`warnTrackSpeed`	`0.15 + t×0.60` → 0.75	**≤ 0.42**	ch46
`lateralDrift`	`t × 4.5` → 4.5	**≤ 2.40**	ch55
`warningTime`	`0.90 − t×0.50` → 0.40s	**≥ 0.69s**	ch43
`bladeScale`	`1 + t×0.65` → 1.65	**≤ 1.30**	ch47

Why warnTrackSpeed ≤ 0.42:
At 0.42, the knife tracks at `0.42 × 420px = 176px/s`. The cake escapes at 190px/s. Running away always opens a gap — the knife can never close faster than the cake can flee. Above 0.42 the knife outpaces the escape and running toward it becomes the only option, which is mechanically broken.

Why warningTime ≥ 0.69s:
Below 0.69s there is not enough warning time to read the knife spawn position and move to safety before the drop starts.

Why lateralDrift ≤ 2.40 and bladeScale ≤ 1.30:
Hitbox and drop-zone size. With uncapped values the safe corridor between knives narrows to near zero at ch65+.

Round 2 — zSpeedBoost, leadFactor

Date: 2026-04-11
Commit context: bot stuck at 54/101; diagnosed drop time reaching 40ms at high z

param	original	cap	hits at
`zSpeedBoost`	`1 + t×2.2` → 3.2×	**≤ 2.0×**	ch46
`leadFactor`	`t × 1.6` → 1.6	**≤ 0.80**	ch51

Why zSpeedBoost ≤ 2.0:
`zSpeedBoost` scales actual knife drop duration: `dropDur = dropDuration / (1 + (walker.z/22) × zSpeedBoost)`

At 3.2× with a typical walker.z of 7, the knife drops in ~40ms — well below human reaction time and below the 60ms floor. Capping at 2.0 keeps the fastest drops at ~60ms (floor) only for extreme manifold z values.

Why leadFactor ≤ 0.80:
Lead factor controls predictive aiming: `leadX = victim.x + victim.vx × leadFactor × dropDuration`

Above 0.80 the knife aimed well ahead of where the cake was, making moving in any direction feel like running into the aim point.

Round 3 — trackingBias

Date: 2026-04-11
Commit context: bot stuck at 58/101; LEMNISCATE and CARDIOID chapters identified as structural walls regardless of genome

param	original	cap	hits at
`trackingBias`	`0.48 + t×0.44` → 0.92	**≤ 0.69**	ch49

Why trackingBias ≤ 0.69:
`trackingBias` controls how accurately a knife aims: `spread = (manifoldX − dome.cx) × (1 − trackingBias)`

At 0.92, only 8% manifold randomness remains. On LEMNISCATE and CARDIOID manifolds (shapes with cusps and self-intersections), the walker generates clustered high-z spawn positions that the bot cannot escape when knives also track at 92% accuracy.

Cap value 0.69 was chosen over 0.65 to preserve the natural scaling through ch1-48 (trackingBias reaches 0.69 at ch49). Mid-game chapters feel unchanged; only ch49+ are affected.

Current caps (all active as of 2026-04-11)

```
// cake-murder-adventure.html — levelConfig()
warningTime:    Math.max(0.69, 0.90 - t * 0.50)   // floor 0.69s
warnTrackSpeed: Math.min(0.42, 0.15 + t * 0.60)   // cap 0.42
lateralDrift:   Math.min(2.40, t * 4.5)            // cap 2.40
bladeScale:     Math.min(1.30, 1 + t * 0.65)       // cap 1.30
zSpeedBoost:    Math.min(2.00, 1 + t * 2.2)        // cap 2.00
leadFactor:     Math.min(0.80, t * 1.6)            // cap 0.80
trackingBias:   Math.min(0.69, 0.48 + t * 0.44)   // cap 0.69
```

Uncapped (still scaling freely through ch101): - `knifeInterval` — 1.4s → 0.45s (floor already present) - `knifeGoal` — 3 → 21 knives per chapter - `maxSimultaneous` — 1 → 4 simultaneous knives - `zSpeedBoost` effect still varies with `walker.z` per-knife (caps the coefficient, not the physics)

Bot architecture alongside the caps

Two training eras: a look-ahead controller era that evolved alongside cap changes, and a CakeWalker era that replaced the fixed architecture with a systematic sense search.

Look-ahead era (2026-04-11 → 04-12)

date	arch	best total
2026-04-11 (start)	8-gene gradient + look-ahead v1	38/101
2026-04-11	look-ahead v2 (floor=0.20s)	54/101
2026-04-11 + round 1 caps	look-ahead v2	54/101
2026-04-11 + round 2 caps	look-ahead v2	58/101
2026-04-11 + round 3 caps	8-gene, fresh	58/101
2026-04-12	**9-gene + z_anticipate**	**101/101**

Gene 9 — z_anticipate (added 2026-04-11): The bot reads `walker.z` and `walker_to_dome_x(walker)` — the oracle's NEXT spawn position and drop speed — and adds pre-emptive repulsion from that zone before the knife appears. Targets the LEMNISCATE/CARDIOID wall where high-z clustered spawns defeat all 8-gene strategies.

CakeWalker era (2026-04-12 → 04-13)

Replaced the fixed look-ahead architecture with a systematic sense vocabulary search. **98 CakeWalker types** across **23 senses** in four families. CEM trains each independently against the same 101 chapters and the same difficulty caps.

date	type	d	best total	OOS
2026-04-13	`float_warn`	2	95/101	—
2026-04-	`triad_float_phase_wall`	3	101/101	12/13 (ch102

13				fails)
2026-04-13	`float_warn_danger_edge`	4	101/101	12/13 (ch102 fails)
2026-04-13	`float_warn_edge_exit`	4	101/101	12/13 (ch102 fails)
2026-04-13	**`float_warn_sprint_edge`**	**4**	**101/101**	**13/13** ✓

The difficulty caps do not change between eras — both architectures train and evaluate against identical physics. The CakeWalker champion uses 4 weights where the look-ahead champion used 9.

Human high score: **54 consecutive** (fox, 2026-04-11).