

The Type Chart Is a Payoff Matrix, and the Metagame Only Half-Solves It

Brecht Corbeel

Absolute Digital Publishers

absolutedigitalpublishers.com

September 7, 2026

Abstract

Pokémon’s 18-type effectiveness table is a fully specified, officially published payoff structure, not a metaphor for one. We restrict it to the eight species Smogon University’s own August 2026 Generation 9 OU usage file (1500-rating cutoff, 730,502 real logged battles) ranks highest, build the induced zero-sum game $U(i, j) = \log_2(E(i, j)/E(j, i))$ from each pair’s best same-type-attack-bonus multiplier, and solve it exactly by linear programming rather than by integrating a dynamical system: the Nash equilibrium is a five-species mix (Great Tusk, Gholdengo, Kingambit, Dragonite, Zamazenta at 18.18%, 22.73%, 13.64%, 31.82%, 13.64%; three species excluded), with value exactly zero, the dimensional-sanity check any symmetric zero-sum game must pass. Tested against the real, renormalized August usage vector, this equilibrium loses badly to both a uniform null and a pure month-to-month persistence null, by Kullback–Leibler divergence and by chi-square, failing our pre-registered kill criterion outright. The more consequential, actually-computed finding is methodological: the replicator equation this same game is conventionally paired with does *not* converge to that equilibrium. A seventeen-seed robustness sweep of the identical ordinary differential equation lands on five different apparent “winners” depending only on the starting mix, because the equilibrium’s own support conceals an intransitive triangle — Great Tusk beats Kingambit beats Dragonite beats Great Tusk — the textbook structural signature of heteroclinic-cycle, non-convergent replicator dynamics [9, 10]. We isolate this triangle as its own three-strategy sub-game and show it does not settle over an 8,000-time-unit integration, and we check the headline equilibrium’s robustness against two alternative zero-sum payoff functionals, finding it stable under one and not under the other — an honestly reported limiting case. A real, dated administrative event, Mega Kangaskhan’s 2013 ban from OverUsed sixty-eight days after its release, supplies the paper’s counter-case: the one correction on record for this system was never com-

puted from any equilibrium at all. We conclude that an exact, officially published payoff matrix is a meaningfully stronger object than an analogy, precisely because it can be solved and checked — and that solving it correctly still requires choosing and stating a solution concept, a choice a companion informal treatment of the same data does not make explicit.

Keywords: zero-sum game; Nash equilibrium; replicator dynamics; evolutionarily stable strategy; heteroclinic cycle; linear programming; competitive metagame; Smogon.

Series: The Pokémon Economics Papers, Series PI, No. 1 of 2 (with PI10).

1 The claim

On October 12, 2013, Pokémon X and Y introduced Mega Evolution, a held stone that lets one Pokémon per team transform mid-battle into a stronger, sometimes retyped, form [35]. Sixty-eight days later, on December 19, a single line change to the file that defines Smogon University’s competitive tier rules added one word to the OverUsed banlist: **Kangaskhanite** [36]. It has not returned to that tier since. Sixty-eight days is not long enough for a population of any realistic size to find, converge on, and exhaust a best response to a new strategy under any of the standard dynamics this paper uses. Whatever correction happened to Mega Kangaskhan, it was not the metagame correcting itself; it was a committee reading the same evidence everyone else had and acting on it faster than any population process plausibly could.

That fact sets the question this paper actually answers, stated precisely rather than as a rhetorical opener: given a real, exact, and completely public payoff structure — the type-effectiveness table — how much of what a real competitive population actually plays can that structure alone explain, once it is solved rather than merely gestured at? Most economics writing that borrows a game from outside economics is

borrowing an analogy: markets behave *like* an ecosystem, a firm competes *like* an organism. The type chart needs no such hedge. Every one of its multipliers — 0, 0.5, 1, or 2 — is a real payoff in the literal game-theoretic sense: it states exactly what one choice returns against another, with no error term and no revision [16]. A payoff matrix a game studio publishes and never quietly revises is a better empirical object than most matrices economists actually get to test a theory against, and this paper’s first commitment is to use it as one.

Using it as one means solving it, not describing it. §2 restricts the full multiplier table to the eight species Smogon’s own real, dated telemetry ranks highest in a named tier and month, builds the induced zero-sum game, and states the solution concept this paper actually uses: the game’s exact Nash equilibrium, found by linear programming [1, 3, 4], together with the dimensional-sanity check that a symmetric zero-sum game’s value must be exactly zero. §3 states the data and method in full: two Smogon usage files, eight independently sourced dual typings, and a type-chart cross-verification pass that itself corrected two omitted cells before anything was computed. §4 reports the equilibrium and tests it against the real ladder and two named nulls; it loses to both, by a wide and honestly reported margin. §5 reports this paper’s central methodological finding: the replicator dynamics conventionally paired with this kind of game [7, 8] do not converge to the equilibrium §4 computed, because an intransitive triangle sits inside the equilibrium’s own support — the same structural mechanism responsible for non-convergent, heteroclinic-cycle dynamics in three-species ecological competition [10, 11] and, empirically, in at least two independently documented biological systems [12, 13]. §6 checks whether this finding survives a change of payoff functional; it does, for one alternative, and does not for another, and both results are reported. §7 returns to Mega Kangaskhan as the paper’s counter-case: the one real correction this system’s history records was never computed from anything this paper builds. §8 collects this paper’s own honest limitations in one place, including a direct statement about a companion informal treatment of the same data. §9 states what changes once a payoff structure this exact is actually solved rather than only invoked.

This paper shares a source table with a sibling paper in the same commission, ‘the-type-chart-has-no-apex-predator’, which treats the full 18-type super-effective relation as an unweighted directed graph and asks order-theoretic and graph-theoretic questions of its static structure — whether it forms a partial order, whether it forms a tournament, how it decomposes under strongly-connected-components analysis. That is a question about the chart’s own topology, answered with no usage data and no game solved. This paper

asks a different question of the same source: restricted to eight real, usage-selected species and read as a population game, what does solving it predict, and does solving it well-posedly even produce a single answer? The two papers do not compete for the same result; §2 and §5 both name the distinction again at the point where a reader might otherwise conflate them.

2 The model

2.1 From a multiplier table to a payoff matrix

Let $M(t, d) \in \{0, 0.5, 1, 2\}$ denote the official type-effectiveness multiplier for an attack of type t against a defending type d , as published by the games themselves and reproduced, cross-checked type by type, from Bulbapedia [16–24]. For two species i, j with official dual typings $\text{types}(i), \text{types}(j)$ (one or two types each), define i ’s best same-type-attack-bonus multiplier into j as

$$E(i, j) = \max_{t \in \text{types}(i)} \prod_{d \in \text{types}(j)} M(t, d). \quad (1)$$

This rewards nothing but typing: no coverage moves outside a species’ own two types, no items, no abilities, no base stats, no speed. It is the one variable the chart alone can price, and restricting the model to exactly that variable, rather than smuggling in anything else, is what lets §4 attribute this paper’s negative result to a specific, named omission rather than to an unstated one.

A matrix of pairwise advantages is not yet a symmetric zero-sum game; it is a real number in each direction, not a signed one. We convert it by a log-ratio,

$$U(i, j) = \log_2 \left(\frac{E(i, j)}{E(j, i)} \right), \quad (2)$$

which is antisymmetric by construction ($U(i, j) = -U(j, i)$) and reads in doublings of relative advantage: $U(i, j) = 2$ means i ’s best attack into j outclasses j ’s best attack into i by a factor of four. This specific functional is a modeling choice, not a fact handed down by the chart; §6 checks two alternatives against it directly rather than asserting the choice is innocuous. Replicator dynamics on a zero-sum payoff of this general kind are not a new tool for this house: prior papers in this publisher’s Evolutionary Economics series apply the same replicator-equation apparatus to real market selection and a matched null-model discipline to test claims of selection against chance [38, 39], and both apply here without needing to be rederived — this paper’s contribution is the empirical system (a transparently telemetered, officially exact payoff table), not the mathematics of the tool itself.

Proposition 1 (Well-defined and zero-sum). *For any two distinct species $i \neq j$ drawn from the 18-type chart, $E(i, j) > 0$, so $U(i, j)$ is always defined, and $U(i, j) = -U(j, i)$ exactly, for every ordered pair.*

This is the paper’s first dimensional-sanity check: every type has at least one neutral ($1 \times$) interaction with every other, so the maximum in Equation 1 is never taken over an empty or all-zero set, and `sim_typechart.py` verifies this numerically for all 56 ordered pairs among the eight species used below (Appendix A reproduces both full matrices), rather than assuming it holds.

2.2 The equilibrium concept: solve the game, do not merely run it

A zero-sum game between symmetric players — both draw from the same eight-strategy set, with antisymmetric payoffs — has a value fixed by construction:

Proposition 2 (A symmetric zero-sum game’s value is exactly zero). *For any antisymmetric U ($U(i, j) = -U(j, i)$, $U(i, i) = 0$) and any mixed strategy $x \in \Delta^7$, $x^\top U x = 0$ identically.*

Appendix B proves this in three lines. It follows from von Neumann’s minimax theorem [1] that an optimal mixed strategy x^* exists satisfying $(x^{*\top} U)_j \geq 0$ for every pure strategy j — equivalently, no pure deviation earns a positive payoff against x^* , the standard Nash condition for a symmetric game [3]. We solve for x^* directly by linear programming (maximize the guaranteed floor $v = \min_j (x^\top U)_j$ subject to $x \in \Delta^7$), using the open-source HiGHS dual-simplex solver [5] as wrapped by SciPy [6] — an exact, reproducible solve, not an approximation and not the outcome of integrating a dynamical system to some finite stopping time.

We emphasize this because the literature this paper draws its dynamics from conventionally pairs a matrix game with the replicator equation [8],

$$\dot{x}_i = x_i(f_i(x) - \phi(x)), \quad f_i(x) = \sum_j x_j U(i, j), \quad \phi(x) = \sum_i x_i f_i(x), \quad (3)$$

and an evolutionarily stable strategy in Maynard Smith and Price’s sense [7] is often treated as though it were simply *what the replicator equation converges to*. For a general zero-sum matrix game with more than two strategies, that identification is not guaranteed, and Hofbauer and Sigmund’s standard treatment states the general condition under which it fails: when the game’s interior equilibrium sits inside an intransitive (rock-paper-scissors-type) cycle among a subset of strategies, the replicator flow need not converge to that equilibrium at all [9], a phenomenon first characterized for three-species ecological competition by May

and Leonard [10] and placed inside evolutionary game theory’s own formal apparatus by Zeeman [11]. §5 shows this is exactly what happens here, and treats it as a real result about this specific game rather than a caveat to file away: the Nash equilibrium (Proposition 3 below) is the well-posed, reproducible object; a single finite-time integration of Equation 3 is not.

2.3 Non-duplication

A sibling paper on this same commission, ‘the-type-chart-has-no-apex-predator’, restricts its own analysis to the unweighted directed graph formed by the chart’s super-effective ($2 \times$) edges alone, asking whether that 51-edge relation forms a strict partial order or a tournament and how it decomposes under Tarjan’s strongly-connected-components algorithm. It uses no usage data, no weighting by multiplier magnitude, and solves no game. Every object in this paper — E ’s full $0/0.5/1/2$ range, the usage-weighted restriction to eight species, the induced zero-sum game, its Nash equilibrium, and the replicator dynamics’ relationship to that equilibrium — is outside that paper’s scope, and nothing here restates or depends on its results.

3 Data and method

3.1 The restricted eight-species game

Smogon publishes usage statistics for every ranked format every month, tallying which species appeared on which teams across every logged ladder battle above a stated rating floor. The August 2026 file for Generation 9 OverUsed at the 1500-rating cutoff — 730,502 recorded battles — ranks Great Tusk first at 32.15476% usage, followed by Gholdengo (25.09918%), Kingambit (22.59541%), Dragonite (17.71140%), Zamazenta (17.68923%), Ogerpon-Wellspring (17.63260%), Iron Valiant (16.21874%), and Raging Bolt (15.55163%) [25]. These eight are this paper’s restricted game: the highest-usage species in a real, named, dated tier, not a hand-picked or hypothetical roster. Smogon publishes four parallel cutoffs (0, 1500, 1695, 1825) for the same tier and month; we use 1500 for the reason Smogon’s own tiering discussion typically gives — high enough to exclude a large mass of unranked and early-placement games that have not yet separated signal from noise, low enough to keep the battle count in the hundreds of thousands. A July 2026 file at the same cutoff [26] supplies this paper’s persistence null (§4).

Each species’ official dual typing is independently confirmed against its own Bulbapedia page, not assumed from memory: Great Tusk is Ground/Fighting [27], Gholdengo is Steel/Ghost [28], Kingambit is

Dark/Steel [29], Dragonite is Dragon/Flying [30], Zamazenta’s Crowned Shield form — the one nearly every competitive set uses — is Fighting/Steel [31], Ogerpon’s Wellspring Mask form is Grass/Water [32], Iron Valiant is Fairy/Fighting [33], and Raging Bolt is Electric/Dragon [34].

3.2 Type-chart verification: a correction made before computing, not after

Every offensive multiplier used in E was checked directly against Bulbapedia’s own dedicated page for the attacking type in question, after an initial broader single-page fetch of the general “Type” article was found, on cross-check, to have silently omitted two real cells: Dragon-type attacks’ $0\times$ multiplier against Fairy, and Grass-type attacks’ $0.5\times$ multiplier against Steel [16, 18, 20, 21, 23]. Both were confirmed independently (the Fairy-type page states plainly that Fairy-type Pokémon are immune to Dragon-type moves; the Steel-type page’s own offensive table lists Grass among its $0.5\times$ targets) and corrected in `sim_typechart.py` before any equilibrium was computed. Neither correction changes this paper’s headline results: Grass only meets a Steel sub-type as a defending type for Ogerpon-Wellspring’s opponents (Gholdengo, Kingambit, Zamazenta), and in each of those three matchups Ogerpon-Wellspring’s Water sub-type already produces an equal or larger multiplier than the corrected Grass value, so E ’s printed entries are identical under both the erroneous and the corrected chart — verified directly by rerunning the script under both versions, not assumed. We report the correction because a paper whose central object is “the chart, read exactly” owes its reader the exact chart, and because catching this kind of omission is itself a methodological point worth making plainly rather than silently fixing.

3.3 What the restricted game deliberately excludes

Equation 1 rewards nothing but typing: no coverage moves outside a species’ own two types, no held items, no abilities, no base stats, no speed tiers, no prediction, and no six-slot team-building constraint against a legal pool of several hundred other ranked species. This is a deliberate restriction, not an oversight, and §4 and §8 return to exactly what the resulting gap between model and ladder is and is not entitled to say once the restriction is named this plainly.

4 The exact equilibrium, and its failure

Proposition 3 (The Nash equilibrium, computed). *Solving Equation 2’s induced game by linear programming (`scipy.optimize.linprog`, `HiGHS`) returns support $\{\text{Great Tusk, Gholdengo, Kingambit, Dragonite, Zamazenta}\}$ at shares 2/11, 5/22, 3/22, 7/22, 3/22 (18.18%, 22.73%, 13.64%, 31.82%, 13.64%); Ogerpon-Wellspring, Iron Valiant, and Raging Bolt are excluded, each with a strictly negative deviation payoff against x^* (-0.4545 , -0.1364 , -0.1818 respectively — verified explicitly, not assumed from exclusion). The solver returns game value $t^* = -0.000000000000$, matching Proposition 2’s prediction of exactly zero to machine precision.*

Two entries make the underlying E concrete. Great Tusk’s Fighting sub-type lands at $2\times$ on Kingambit’s Steel half and $2\times$ again on its Dark half, for a combined $4\times$ — tied for the single strongest matchup among all 56 ordered pairs in this restricted game (Zamazenta and Iron Valiant also reach $4\times$ into Kingambit via their own Fighting sub-type; Appendix A reports the full matrix). Against Dragonite, Great Tusk’s best option is Fighting into Dragon/Flying, and Flying halves Fighting damage while Dragon is neutral to it, so its best multiplier there is only $0.5\times$ — among Great Tusk’s weakest matchups on this list, and, as §5 shows, the other leg of a three-way cycle.

4.1 Testing the equilibrium against the real ladder

Renormalizing the eight species’ real August shares to sum to one gives the observed target: Great Tusk 19.53%, Gholdengo 15.24%, Kingambit 13.72%, Dragonite 10.76%, Zamazenta 10.74%, Ogerpon-Wellspring 10.71%, Iron Valiant 9.85%, Raging Bolt 9.45%. Table 1 and Figure 1 compare this to the Nash equilibrium (Proposition 3), a uniform null (12.5% each — “typing conveys no information”), and a pure persistence null (July 2026’s renormalized shares for the same eight species, with no game-theoretic term at all [26]).

Scoring the fit with Kullback–Leibler divergence [14], $D_{\text{KL}}(p||q) = \sum_i p_i \log_2(p_i/q_i)$, from the observed distribution to each candidate: the Nash equilibrium scores 10.69 bits, the uniform null scores 0.045 bits, and July’s raw persistence scores 0.00046 bits. On Pearson’s chi-square test [15] against the eight species’ derived real August team-appearance counts (each species’ usage fraction times 730,502 battles, summing to 1,202,793 team-appearances across these eight species alone), the Nash equilibrium scores approximately 4.35×10^{19} — a figure this large only because




Figure 1: Real August 2026 usage share against the Nash equilibrium (Proposition 3), a uniform null, and a July-2026 persistence null, for the same eight species.

Table 1: Real August 2026 usage against the Nash equilibrium and two nulls (shares of these 8 species only).

Species	Obs.	Nash	Unif.	Persist.
Great Tusk	19.53	18.18	12.50	19.68
Gholdengo	15.24	22.73	12.50	15.38
Kingambit	13.72	13.64	12.50	13.84
Dragonite	10.76	31.82	12.50	10.37
Zamazenta	10.74	13.64	12.50	10.18
Ogerpon-W.	10.71	0.00	12.50	10.77
Iron Valiant	9.85	0.00	12.50	10.21
Raging Bolt	9.45	0.00	12.50	9.57

the equilibrium assigns three real, frequently played species an expected share of exactly zero, and squaring a large real count against a zero expectation is undefined in the limit and explodes numerically — against 8.09×10^4 for the uniform null and 767 for persistence.

Proposition 4 (The kill criterion fails). *By this paper’s pre-registered kill criterion — the equilibrium should be credited only if the real distribution fits it at least as well as it fits both named nulls — the equilibrium fails, by roughly two orders of magnitude in D_{KL} divergence and by many more in chi-square, against both nulls.*

This is Derived, not Proposed: every number above follows mechanically from the chart, the two usage files, and the stated formulas, computed once by `sim_typechart.py` and reported here exactly as printed (Appendix C), with nothing adjusted after the fact to look better or worse. The equilibrium correctly identifies which single matchup structure is strongest in the field — Great Tusk’s Ground/Fighting typing is close to strictly dominant against this particular eight, neutral or better into six of the other seven — but

Table 2: Replicator-dynamics apparent winner by starting mix (17 seeded runs, $t = 5000$).

Start	Winner	Share
Uniform	Dragonite	33.6%
Dirichlet (8 draws)	Gholdengo×3, Dragonite×2, Kingambit, Great Tusk, Zamazenta	34–65%
Near-corner (8 starts)	Zamazenta, Gholdengo, Dragonite×2, Great Tusk×2, Gholdengo×2	53–87%

the magnitude the exact game assigns to that structural advantage is not the magnitude the real ladder shows.

5 The dynamics do not agree with each other, let alone with the ladder

Proposition 3’s equilibrium is a static object: a mixed strategy no pure deviation can profitably exploit. Nothing in that statement claims a population starting anywhere else will actually arrive there under any particular dynamic. We check this directly rather than assume it, integrating the replicator equation (Equation 3) with SciPy’s adaptive Runge–Kutta solver (`solve_ivp`, method RK45, $\text{rtol} = 10^{-10}$) from seventeen different starting mixes: the uniform mix, eight Dirichlet-random draws under a fixed, declared seed (20260907), and eight “near-corner” starts (98% one species, the rest split evenly), each integrated to $t = 5000$.

Proposition 5 (The dynamics do not converge to the Nash equilibrium; this paper’s central methodological finding). *Across the seventeen runs, five different single-species “winners” appear depending only on*

the starting mix (Table 2): Great Tusk, Gholdengo, Kingambit, Dragonite, and Zamazenta. Every one of these five is exactly the Nash equilibrium’s own support (Proposition 3); none of the three Nash-excluded species (Ogerpon-Wellspring, Iron Valiant, Raging Bolt) ever wins, in any of the seventeen runs.

This is not noise from an under-integrated system settling slowly; it is the correct qualitative behavior of a replicator system containing an intransitive cycle. Restricting attention to Great Tusk, Kingambit, and Dragonite alone — exactly the three of the equilibrium’s five support species with the most one-sided pairwise payoffs among each other — gives

$$U_3 = \begin{pmatrix} 0 & 2 & -2 \\ -2 & 0 & 1 \\ 2 & -1 & 0 \end{pmatrix} \quad (4)$$

in the order (Great Tusk, Kingambit, Dragonite): Great Tusk beats Kingambit (+2), Kingambit beats Dragonite (+1), and Dragonite beats Great Tusk (+2) — a genuine three-cycle, verified directly from Equation 2 rather than asserted. This is the textbook structural signature May and Leonard identified for three-species ecological competition [10] and Zeeman placed inside evolutionary game dynamics directly [11]: Hofbauer and Sigmund’s standard treatment states that a replicator system built on an intransitive cycle of this kind need not converge to its own interior equilibrium at all, producing sustained, non-monotonic orbits instead [9]. Two independently documented real biological systems show the same mechanism operating outside any simulation: side- blotched lizard mating morphs [12] and colicin-mediated competition among *E. coli* strains [13]. We do not claim either biological system behaves identically to this game; both are cited as the phenomenon’s standing empirical precedent, not as evidence about Pokémon.

Figure 2 integrates Equation 4 in isolation to $t = 3000$ from a 0.5/0.3/0.2 start. Great Tusk’s share falls from 50% to under 4% by $t \approx 200$, recovers to nearly 40% by $t \approx 2000$, and falls again; Kingambit and Dragonite trade the same pattern out of phase. Extending the same integration to $t = 8000$ (Appendix C) shows no narrowing of the oscillation’s amplitude, consistent with a genuine heteroclinic cycle rather than slow convergence.

Remark 1. A single finite-time replicator integration therefore does not report a well-posed prediction of this model on its own. Which of the five Nash-support species appears to be “winning” at any chosen stopping time depends on the starting mix and on how long the clock is allowed to run, not on anything the type chart alone determines. A reader who wants one reproducible number from this formal machinery




Figure 2: Isolated three-strategy sub-game (Great Tusk, Kingambit, Dragonite; Eq. 4), replicator trajectory from a 0.5/0.3/0.2 start, $t \in [0, 3000]$. No convergence: each share continues to swing by tens of percentage points late in the integration.

Table 3: Nash equilibrium support under three payoff functionals.

Species	Log-ratio	Linear diff.	Sign-only
Great Tusk	18.18	12.70	9.09
Gholdengo	22.73	30.16	27.27
Kingambit	13.64	7.94	9.09
Dragonite	31.82	41.27	—
Zamazenta	13.64	7.94	—
Ogerpon-W.	—	—	27.27
Iron Valiant	—	—	18.18
Raging Bolt	—	—	9.09

should use Proposition 3’s linear-programming solution, which does not depend on a stopping time or a starting point; §8 returns to this point directly.

6 Does the payoff functional matter?

Equation 2’s log-ratio is one defensible way to turn E into an antisymmetric zero-sum payoff; it is not the only one. We re-solve the identical game’s exact Nash equilibrium under two alternatives, using `sim_sensitivity.py`: a linear difference, $U_{\text{diff}}(i, j) = E(i, j) - E(j, i)$, and a sign-only functional that discards magnitude entirely, $U_{\text{sign}}(i, j) = \text{sign}(E(i, j) - E(j, i))$. Both are exactly antisymmetric, so both induce a symmetric zero-sum game whose value is again exactly zero (Proposition 2 does not depend on which antisymmetric functional is used); only the equilibrium’s support and shares can differ.

Proposition 6 (Payoff-functional sensitivity). *The linear-difference functional reproduces exactly the same*

five-species support as the log-ratio functional used in the main text (shares shift, but Great Tusk, Gholdengo, Kingambit, Dragonite, and Zamazenta remain the equilibrium’s entire support, and no other species enters). The sign-only functional changes the support qualitatively: Dragonite and Zamazenta drop out, Ogerpon-Wellspring and Iron Valiant enter.

This is an honest limiting case, reported because it is real, not because it is comfortable. The five-species finding survives a change from a logarithmic to a linear scaling of advantage — two functionals that agree on which matchups matter, disagreeing only on how much more a $4\times$ matchup counts than a $2\times$ one. It does not survive discarding magnitude altogether: once every matchup is scored purely as win, lose, or tie with no record of *how much*, the official chart’s own $0.5\times/2\times/4\times$ structure — precisely the feature this paper argues makes the chart a real payoff matrix rather than a bare tournament graph (§2) — stops doing any work, and a different five-species set becomes optimal. The chart’s exactness, in other words, is only fully used by a functional that keeps its magnitudes; a sibling paper’s graph-theoretic treatment of the same source table, which keeps only the binary super-effective relation, is by construction closer to this paper’s sign-only sensitivity case than to its own main result, which is one further reason the two papers ask genuinely different questions of the same eighteen rows and columns.

7 When the correction is not the model’s own

Return to Mega Kangaskhan, introduced in §1, because it names a boundary this paper has been building toward rather than illustrating a separate point. If a payoff structure’s own equilibrium logic actually governed outcomes, an over-concentrated strategy should erode on its own: opponents adapt, counter-picks rise, and a system’s own dynamics eventually punish over-concentration — exactly the way this paper’s replicator computation drives three of the eight species toward zero share (§4). Mega Kangaskhan never got the chance to test that story either way. Sixty-eight days after Mega Evolution existed at all, a tiering council removed it by administrative vote, and the same item was independently re-banned from OverUsed’s own Monotype variant five days later and from the 1v1 format twenty-five days into the following month — three separate rulings within the mechanic’s first four months on record, all administrative, none a measured population response [36].

Even the administrative mechanism is not one uniform process, which matters for how much weight “a committee decides” can bear as a tidy alternative to

“the equilibrium decides.” The same Kangaskhanite that took sixty-eight days to be banned from Singles OverUsed took until April 2017 — more than three years — to be banned from Doubles OverUsed [36], a structurally different metagame with its own usage patterns and its own institutional patience for the same held item. Two formats, one payoff-relevant fact about one species, and two entirely different correction timescales, neither computed from anything resembling Equation 3.

It is also worth naming the boundary between the two correction mechanisms Smogon actually runs, because only one of them is the kind of thing this paper’s model could in principle track. One is automatic and continuous: a species whose usage share falls below a stated threshold in its current tier drops to the tier below with no vote at all, a mechanism that responds directly to the same usage numbers §4 compares against the Nash equilibrium [37]. The other is exactly what happened to Kangaskhanite: discretionary, vote-based, and triggered when a strategy’s dominance is judged to distort competitive play regardless of whether its raw usage number alone would have triggered an automatic drop. Everything this paper computes speaks to the first kind of correction; the counter-case that anchors this section is the second kind, and conflating the two would understate how much of what actually happens to a species’ competitive standing is procedural rather than computed from any payoff matrix at all — exact or otherwise.

8 How this could be wrong

The restricted game omits everything the chart cannot price. Equation 1 scores typing alone. Real teams carry coverage moves outside their own two types, items that swap resistances, abilities that void whole attacking types, speed tiers that decide who acts first, prediction, and a six-slot team-building constraint against a legal pool of several hundred other ranked species this restricted game excludes entirely. §4’s finding that the equilibrium loses to both nulls is a measurement of how much of the real ladder’s structure this specific, named omission leaves unexplained — not a claim that typing is irrelevant, since the equilibrium does correctly flag Great Tusk’s typing as the field’s strongest single structural case (§4).

The rating cutoff is a choice. Smogon publishes four cutoffs for the same tier and month; this paper uses only the 1500 file. A different cutoff changes both the eight-species roster and the observed usage vector, and we have not re-run §4 at the other three cutoffs. We expect the qualitative failure to persist — the equilibrium’s concentration mechanism does not depend on the specific cutoff chosen — but this is an expectation,

not a checked result, and is stated as such.

The payoff functional is a choice, checked but not closed. §6 shows the five-species Nash support survives a change from log-ratio to linear-difference scaling and does not survive discarding magnitude altogether. We checked three functionals, not an exhaustive family; a functional between the two that changed the result and the one that did not is possible and unexamined.

The solution-concept ambiguity is this paper’s own central limitation, surfaced rather than hidden. §5’s seventeen-seed sweep shows that a single finite-time integration of the replicator equation is not a reproducible statement about this model: five different species can appear to be “the answer” depending only on where the integration starts and how long it runs. This paper’s own linked landing article reports one such single integration — a run that (using Euler-method integration rather than this paper’s adaptive RK45, and a different random or default starting condition) lands near a Great-Tusk-dominant outcome rather than near any of this paper’s own seventeen sampled outcomes. That number is not wrong in the sense of a miscalculation; it is one point drawn from exactly the non-convergent, starting-point-dependent dynamics Proposition 5 demonstrates, reported without the robustness check this paper performs. We say this plainly rather than let it stand uncorrected: a reader who wants a single, reproducible “the model predicts X ” number from this specific formal machinery should use this paper’s linear-programming Nash equilibrium (Proposition 3), which does not depend on a stopping time, a starting mix, or a choice of integrator, rather than any one run of Equation 3. Reconciling the landing article’s own reported figure with this finding is outside the scope of a paper-authoring pass and is flagged for the corpus’s editorial process rather than resolved here.

Single month, single tier, no claim beyond it. Every number in this paper describes Generation 9 OverUsed at the 1500-rating cutoff for July and August 2026 specifically. Nothing here is a claim about any other tier, generation, or time window, and the two-month persistence comparison in §4 is not evidence that persistence would keep beating a correctly solved equilibrium over a longer horizon or a different metagame regime.

9 What a fully exact, fully public payoff matrix is and is not good for

Treating a real payoff structure as a solvable game, rather than as an analogy for one, is a meaningfully

stronger and more falsifiable claim, and this paper’s own results are the demonstration of both halves of that strength. Solved exactly, the type chart’s restriction to eight real, usage-ranked species yields a genuine, checkable prediction — one this paper computed by linear programming rather than assumed — and that prediction fails a pre-registered kill criterion against real ladder data, plainly and by a wide, honestly reported margin (§4). An analogy cannot fail this way, because an analogy was never solved in the first place; a solved, exact payoff structure can, and did.

The second half is more consequential for how this kind of formal machinery should be used going forward, inside and outside this commission. A payoff matrix’s exactness does not, by itself, make every question asked of it well-posed. Asking for a static equilibrium is well-posed here and answerable by linear programming, with an internal consistency check (Proposition 2) this paper verified to machine precision. Asking what a population dynamic converges to is a different question, and §5 shows it does not have a single answer for this specific game at all — not because the mathematics is hard, but because an intransitive cycle sits inside the equilibrium’s own support, the same structural fact responsible for non-convergent dynamics in three-species ecological competition and in at least two real biological systems. A paper, or an article, that reports one integration’s outcome as “the solved game predicts X ” without checking whether that outcome survives a change of starting point has not fully solved the game; it has sampled one trajectory of a system this paper shows does not settle.

None of this requires treating Pokémon’s competitive scene as standing in for financial markets or for economic competition generally. It requires taking one public, exact payoff structure exactly as seriously as its own mathematics demands: computing what it actually implies, checking that computation’s internal consistency, testing it against the real, dated telemetry a live competitive population actually produces, and reporting the result whichever way it comes out — including, as it did here, a real negative result and a real methodological finding about when a textbook dynamic simply does not apply.

References

- [1] J. von Neumann (1928). Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, vol. 100, no. 1, pp. 295–320. <https://doi.org/10.1007/BF01448847>. DOI: 10.1007/BF01448847.
- [2] John von Neumann, Oskar Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton University Press.

- [3] John Nash (1951). Non-Cooperative Games. *Annals of Mathematics*, vol. 54, no. 2, pp. 286–295. <https://doi.org/10.2307/1969529>. DOI: 10.2307/1969529.
- [4] George B. Dantzig (1963). *Linear Programming and Extensions*. Princeton University Press.
- [5] Q. Huangfu, J. A. J. Hall (2018). Parallelizing the Dual Revised Simplex Method. *Mathematical Programming Computation*, vol. 10, no. 1, pp. 119–142. <https://doi.org/10.1007/s12532-017-0130-5>. DOI: 10.1007/s12532-017-0130-5.
- [6] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, et al. (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, vol. 17, no. 3, pp. 261–272. <https://doi.org/10.1038/s41592-019-0686-2>. DOI: 10.1038/s41592-019-0686-2.
- [7] J. Maynard Smith, G. R. Price (1973). The Logic of Animal Conflict. *Nature*, vol. 246, no. 5427, pp. 15–18. <https://doi.org/10.1038/246015a0>. DOI: 10.1038/246015a0.
- [8] P. D. Taylor, L. B. Jonker (1978). Evolutionary Stable Strategies and Game Dynamics. *Mathematical Biosciences*, vol. 40, no. 1–2, pp. 145–156. [https://doi.org/10.1016/0025-5564\(78\)90077-9](https://doi.org/10.1016/0025-5564(78)90077-9). DOI: 10.1016/0025-5564(78)90077-9.
- [9] Josef Hofbauer, Karl Sigmund (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press. <https://doi.org/10.1017/CBO9781139173179>. DOI: 10.1017/CBO9781139173179.
- [10] R. M. May, W. J. Leonard (1975). Nonlinear Aspects of Competition Between Three Species. *SIAM Journal on Applied Mathematics*, vol. 29, no. 2, pp. 243–253. <https://doi.org/10.1137/0129022>. DOI: 10.1137/0129022.
- [11] E. C. Zeeman (1980). Population Dynamics from Game Theory. In *Global Theory of Dynamical Systems, Lecture Notes in Mathematics*, Springer, pp. 471–497. <https://doi.org/10.1007/BFb0087009>. DOI: 10.1007/BFb0087009.
- [12] B. Sinervo, C. M. Lively (1996). The Rock–Paper–Scissors Game and the Evolution of Alternative Male Strategies. *Nature*, vol. 380, no. 6571, pp. 240–243. <https://doi.org/10.1038/380240a0>. DOI: 10.1038/380240a0.
- [13] Benjamin Kerr, Margaret A. Riley, Marcus W. Feldman, Brendan J. M. Bohannan (2002). Local Dispersal Promotes Biodiversity in a Real-Life Game of Rock–Paper–Scissors. *Nature*, vol. 418, no. 6894, pp. 171–174. <https://doi.org/10.1038/nature00823>. DOI: 10.1038/nature00823.
- [14] S. Kullback, R. A. Leibler (1951). On Information and Sufficiency. *Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86. <https://doi.org/10.1214/aoms/1177729694>. DOI: 10.1214/aoms/1177729694.
- [15] Karl Pearson (1900). On the Criterion that a Given System of Deviations from the Probable in the Case of a Correlated System of Variables is Such that it can be Reasonably Supposed to have Arisen from Random Sampling. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, Series 5*, vol. 50, no. 302, pp. 157–175. <https://doi.org/10.1080/14786440009463897>. DOI: 10.1080/14786440009463897.
- [16] Bulbapedia contributors (2026). Type. *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. <https://bulbapedia.bulbagarden.net/wiki/Type>.
- [17] Bulbapedia contributors (2026). Dark (type). *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. [https://bulbapedia.bulbagarden.net/wiki/Dark_\(type\)](https://bulbapedia.bulbagarden.net/wiki/Dark_(type)).
- [18] Bulbapedia contributors (2026). Steel (type). *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. [https://bulbapedia.bulbagarden.net/wiki/Steel_\(type\)](https://bulbapedia.bulbagarden.net/wiki/Steel_(type)).
- [19] Bulbapedia contributors (2026). Ghost (type). *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. [https://bulbapedia.bulbagarden.net/wiki/Ghost_\(type\)](https://bulbapedia.bulbagarden.net/wiki/Ghost_(type)).
- [20] Bulbapedia contributors (2026). Dragon (type). *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. [https://bulbapedia.bulbagarden.net/wiki/Dragon_\(type\)](https://bulbapedia.bulbagarden.net/wiki/Dragon_(type)).
- [21] Bulbapedia contributors (2026). Fairy (type). *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. [https://bulbapedia.bulbagarden.net/wiki/Fairy_\(type\)](https://bulbapedia.bulbagarden.net/wiki/Fairy_(type)).
- [22] Bulbapedia contributors (2026). Flying (type). *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. [https://bulbapedia.bulbagarden.net/wiki/Flying_\(type\)](https://bulbapedia.bulbagarden.net/wiki/Flying_(type)).

- [23] Bulbapedia contributors (2026). Grass (type). *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. [https://bulbapedia.bulbagarden.net/wiki/Grass_\(type\)](https://bulbapedia.bulbagarden.net/wiki/Grass_(type)).
- [24] Bulbapedia contributors (2026). Electric (type). *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. [https://bulbapedia.bulbagarden.net/wiki/Electric_\(type\)](https://bulbapedia.bulbagarden.net/wiki/Electric_(type)).
- [25] Smogon University Statistics (2026). Gen 9 OU Usage Statistics, August 2026 (1500+ rating). *Smogon.com*. <https://www.smogon.com/stats/2026-08/gen9ou-1500.txt>.
- [26] Smogon University Statistics (2026). Gen 9 OU Usage Statistics, July 2026 (1500+ rating). *Smogon.com*. <https://www.smogon.com/stats/2026-07/gen9ou-1500.txt>.
- [27] Bulbapedia contributors (2026). Great Tusk (Pokémon). *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. [https://bulbapedia.bulbagarden.net/wiki/Great_Tusk_\(Pok%C3%A9mon\)](https://bulbapedia.bulbagarden.net/wiki/Great_Tusk_(Pok%C3%A9mon)).
- [28] Bulbapedia contributors (2026). Gholdengo (Pokémon). *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. [https://bulbapedia.bulbagarden.net/wiki/Gholdengo_\(Pok%C3%A9mon\)](https://bulbapedia.bulbagarden.net/wiki/Gholdengo_(Pok%C3%A9mon)).
- [29] Bulbapedia contributors (2026). Kingambit (Pokémon). *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. [https://bulbapedia.bulbagarden.net/wiki/Kingambit_\(Pok%C3%A9mon\)](https://bulbapedia.bulbagarden.net/wiki/Kingambit_(Pok%C3%A9mon)).
- [30] Bulbapedia contributors (2026). Dragonite (Pokémon). *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. [https://bulbapedia.bulbagarden.net/wiki/Dragonite_\(Pok%C3%A9mon\)](https://bulbapedia.bulbagarden.net/wiki/Dragonite_(Pok%C3%A9mon)).
- [31] Bulbapedia contributors (2026). Zamazenta (Pokémon). *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. [https://bulbapedia.bulbagarden.net/wiki/Zamazenta_\(Pok%C3%A9mon\)](https://bulbapedia.bulbagarden.net/wiki/Zamazenta_(Pok%C3%A9mon)).
- [32] Bulbapedia contributors (2026). Ogerpon (Pokémon). *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. [https://bulbapedia.bulbagarden.net/wiki/Ogerpon_\(Pok%C3%A9mon\)](https://bulbapedia.bulbagarden.net/wiki/Ogerpon_(Pok%C3%A9mon)).
- [33] Bulbapedia contributors (2026). Iron Valiant (Pokémon). *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. [https://bulbapedia.bulbagarden.net/wiki/Iron_Valiant_\(Pok%C3%A9mon\)](https://bulbapedia.bulbagarden.net/wiki/Iron_Valiant_(Pok%C3%A9mon)).
- [34] Bulbapedia contributors (2026). Raging Bolt (Pokémon). *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. [https://bulbapedia.bulbagarden.net/wiki/Raging_Bolt_\(Pok%C3%A9mon\)](https://bulbapedia.bulbagarden.net/wiki/Raging_Bolt_(Pok%C3%A9mon)).
- [35] Bulbapedia contributors (2026). Pokémon X and Y. *Bulbapedia, The Community-Driven Pokémon Encyclopedia*. https://bulbapedia.bulbagarden.net/wiki/Pok%C3%A9mon_X_and_Y.
- [36] Smogon / Pokémon Showdown contributors (2013). “Kangaskhanite has been banned from OU” (config/formats.js commit c48d01d). *GitHub — smogon/pokemon-showdown*. <https://github.com/smogon/pokemon-showdown/commit/c48d01dc81e1e916789b83b5b1000cfece8dd5bd>.
- [37] Wikipedia contributors (2026). Smogon. *Wikipedia*. <https://en.wikipedia.org/wiki/Smogon>.
- [38] Brecht Corbeel (2026). Replicator Markets: When Competition Computes. *Absolute Digital Publishers, The Evolutionary Economics Papers No. 2*. <https://absolutedigitalpublishers.com/papers/evolutionary-economics-02-replicator-markets.pdf>.
- [39] Brecht Corbeel (2026). Drift or Selection: Null Models for Markets. *Absolute Digital Publishers, The Evolutionary Economics Papers No. 10*. <https://absolutedigitalpublishers.com/papers/evolutionary-economics-10-drift-or-selection.pdf>.

A Full computed matrices

Reproduced verbatim from `sim_typechart_output.txt` (species order: Great Tusk (GT), Gholdengo (GH), Kingambit (KG), Dragonite (DN), Zamazenta (ZM), Ogerpon-Wellspring (OG), Iron Valiant (IV), Raging Bolt (RB)).

Table 4: $E(i, j)$: best type-advantage multiplier, attacker row into defender column.

	GT	GH	KG	DN	ZM	OG	IV	RB
GT	—	2	4	0.5	2	1	1	2
GH	1	—	0.5	1	1	1	2	1
KG	1	2	—	1	0.5	1	2	1
DN	2	0.5	0.5	—	1	2	2	2
ZM	1	0.5	4	1	—	1	2	1
OG	2	1	1	0.5	1	—	1	0.5
IV	2	0.5	4	2	2	1	—	2
RB	1	1	1	2	1	1	1	—

Table 5: $U(i, j) = \log_2(E(i, j)/E(j, i))$: the zero-sum payoff, in doublings of relative advantage.

	GT	GH	KG	DN	ZM	OG	IV	RB
GT	0	1	2	-2	1	-1	-1	1
GH	-1	0	-2	1	1	0	2	0
KG	-2	2	0	1	-3	0	-1	0
DN	2	-1	-1	0	0	2	0	0
ZM	-1	-1	3	0	0	0	0	0
OG	1	0	0	-2	0	0	0	-1
IV	1	-2	1	0	0	0	0	1
RB	-1	0	0	0	0	1	-1	0

Both tables are read row-into-column: row i , column j gives species i attacking species j . Every entry of Table 5 is the negative of its transpose entry by construction (Proposition 1), verified numerically for all 56 off-diagonal ordered pairs in `sim_typechart.py` before any equilibrium was solved.

B The linear program, and why the value is exactly zero

B.1 Proof of Proposition 2

For any antisymmetric U ($U(i, j) = -U(j, i)$ for all i, j , hence $U(i, i) = 0$) and any $x \in \Delta^7$,

$$x^\top Ux = \sum_{i,j} x_i x_j U(i, j) = \sum_{i < j} x_i x_j (U(i, j) + U(j, i)) = \sum_{i < j} x_i x_j \cdot 0 = 0,$$

where the middle step pairs each off-diagonal term with its transpose partner (each unordered pair $\{i, j\}$ contributes $x_i x_j U(i, j) + x_j x_i U(j, i) = x_i x_j (U(i, j) + U(j, i))$) and diagonal terms vanish since $U(i, i) = 0$. This holds for *every* $x \in \Delta^7$, not only at an equilibrium — in particular, playing x^* against itself always yields exactly zero, which is why a symmetric zero-sum game’s value cannot be anything other than zero.

B.2 The linear program

Following the standard formulation for a symmetric zero-sum matrix game (a direct consequence of von Neumann’s minimax theorem, 1; see also 2), we solve for $x \in \mathbb{R}^8$ and $t \in \mathbb{R}$:

$$\text{maximize } t \quad \text{subject to} \quad \left(\sum_i x_i U(i, j) \right) \geq t \quad \forall j, \quad \sum_i x_i = 1, \quad x_i \geq 0.$$

At the optimum, t^* is the game’s value (exactly 0, by the appendix’s first result) and x^* is a Nash equilibrium: no pure strategy j earns more than t^* against it. We solve this with `scipy.optimize.linprog(method="highs")` [5, 6]; the solver returns $t^* = -0.000000000000$ and the support reported in Proposition 3.

B.3 Verifying exclusion, not assuming it

A species j is correctly excluded from x^* ’s support only if playing j strictly loses against x^* . Because the game is symmetric zero-sum, the payoff to a pure strategy j played against x^* is $-(x^{*\top}U)_j$ (antisymmetry: the column player’s payoff is minus the row player’s). Table 6 reports this value for all eight species; the three excluded species are strictly negative, and the five in-support species are zero to solver precision, exactly the condition Proposition 3 requires.

Table 6: Payoff to a pure deviation against x^* (must be ≤ 0 ; verified, not assumed).

Species	Payoff	Status
Great Tusk	−0.000000	in support
Gholdengo	−0.000000	in support
Kingambit	−0.000000	in support
Dragonite	+0.000000	in support
Zamazenta	+0.000000	in support
Ogerpon-Wellspring	−0.454545	excluded
Iron Valiant	−0.136364	excluded
Raging Bolt	−0.181818	excluded

C Simulation method and full output

C.1 `sim_typechart.py`

Deterministic except for one declared, fixed seed. Requires `numpy` and `scipy`. Computes, in order: (1) E and U for all 56 ordered pairs among the 8 species (§2); (2) the exact Nash equilibrium by linear programming (Appendix B); (3) a seventeen-run replicator-dynamics robustness sweep (SciPy `solve_ivp`, RK45, `rtol` = 10^{-10} , `atol` = 10^{-12} , t = 5000): the uniform mix; eight Dirichlet-random draws under `numpy.random.default_rng(20260907)` (the only randomness anywhere in this package, and it affects only *which* eight of infinitely many possible starting mixes are sampled, not the equilibrium computation itself); eight near-corner starts (98% one species, 2% split over the rest); (4) the isolated Great Tusk/Kingambit/Dragonite three-strategy sub-game, integrated to t = 8000; (5) D_{KL} divergence and Pearson χ^2 of the Nash equilibrium, the uniform null, and the July-2026 persistence null against the real, renormalized August 2026 usage vector. Run: `python3 sim_typechart.py`. Full std-out below, reproduced verbatim from `sim_typechart_output.txt`; machine-readable results are also written to `sim_typechart_output.json`.

```
E(i,j): best type-advantage multiplier
GT->GH=2 GT->KG=4 GT->DN=0.5 GT->ZM=2
GT->OG=1 GT->IV=1 GT->RB=2
GH->GT=1 GH->KG=0.5 GH->DN=1 GH->ZM=1
GH->OG=1 GH->IV=2 GH->RB=1
KG->GT=1 KG->GH=2 KG->DN=1 KG->ZM=0.5
KG->OG=1 KG->IV=2 KG->RB=1
DN->GT=2 DN->GH=0.5 DN->KG=0.5 DN->ZM=1
DN->OG=2 DN->IV=2 DN->RB=2
ZM->GT=1 ZM->GH=0.5 ZM->KG=4 ZM->DN=1
ZM->OG=1 ZM->IV=2 ZM->RB=1
OG->GT=2 OG->GH=1 OG->KG=1 OG->DN=0.5
OG->ZM=1 OG->IV=1 OG->RB=0.5
IV->GT=2 IV->GH=0.5 IV->KG=4 IV->DN=2
IV->ZM=2 IV->OG=1 IV->RB=2
```

RB->GT=1 RB->GH=1 RB->KG=1 RB->DN=2
RB->ZM=1 RB->OG=1 RB->IV=1

Cross-check vs. article prose: E(Great Tusk->Kingambit)=4.0
(expect 4.0), E(Great Tusk->Dragonite)=0.5 (expect 0.5)

EXACT Nash equilibrium (linear programming):

LP success: True; game value t* = -0.000000000000

Great Tusk: 18.1818%
Gholdengo: 22.7273%
Kingambit: 13.6364%
Dragonite: 31.8182%
Zamazenta: 13.6364%
Ogerpon-Wellspring: 0.0000%
Iron Valiant: 0.0000%
Raging Bolt: 0.0000%

Payoff a pure deviation earns against x* (<=0 required):

Great Tusk: -0.000000 (in support)
Gholdengo: -0.000000 (in support)
Kingambit: -0.000000 (in support)
Dragonite: +0.000000 (in support)
Zamazenta: +0.000000 (in support)
Ogerpon-Wellspring: -0.454545 (excluded)
Iron Valiant: -0.136364 (excluded)
Raging Bolt: -0.181818 (excluded)

Replicator-dynamics robustness sweep (RK45, 17 starts, seed=20260907):

uniform -> Dragonite (33.59%)
random_dirichlet_0 -> Gholdengo (34.03%)
random_dirichlet_1 -> Gholdengo (41.90%)
random_dirichlet_2 -> Kingambit (40.06%)
random_dirichlet_3 -> Dragonite (64.98%)
random_dirichlet_4 -> Dragonite (60.01%)
random_dirichlet_5 -> Gholdengo (55.74%)
random_dirichlet_6 -> Great Tusk (60.23%)
random_dirichlet_7 -> Zamazenta (40.38%)
near_corner_Great Tusk -> Zamazenta (61.94%)
near_corner_Gholdengo -> Gholdengo (79.14%)
near_corner_Kingambit -> Dragonite (86.82%)
near_corner_Dragonite -> Great Tusk (56.35%)
near_corner_Zamazenta -> Dragonite (58.16%)
near_corner_Ogerpon-Wellspring -> Great Tusk (69.70%)
near_corner_Iron Valiant -> Gholdengo (53.31%)
near_corner_Raging Bolt -> Gholdengo (54.68%)

Distinct apparent winners: [Dragonite, Gholdengo, Great Tusk,
Kingambit, Zamazenta]

Nash-equilibrium support: [Dragonite, Gholdengo, Great Tusk,
Kingambit, Zamazenta]

Every apparent winner lies in the Nash support: True

Embedded 3-cycle (Great Tusk, Kingambit, Dragonite):

```
[[ 0.  2. -2.]  
 [-2.  0.  1.]  
 [ 2. -1.  0.]]
```

Trajectory (0.5/0.3/0.2 start), no convergence:

```
t=0      GT=50.00% KG=30.00% DN=20.00%
t=50     GT=41.62% KG=15.25% DN=43.13%
t=200    GT=3.43%  KG=41.92% DN=54.65%
t=500    GT=19.91% KG=14.50% DN=65.60%
t=1000   GT=3.70%  KG=37.51% DN=58.79%
t=2000   GT=39.78% KG=45.41% DN=14.81%
t=4000   GT=25.65% KG=60.50% DN=13.85%
t=8000   GT=8.08%  KG=70.85% DN=21.07%
```

Goodness of fit vs. real August 2026 usage:

```
D_KL(observed||model), bits: Nash=10.692888 uniform=0.045110
  persistence=0.000456
Chi-square: Nash=4.353429e+19 uniform=8.093723e+04
  persistence=7.672983e+02
Kill criterion: FAILS -- kill criterion triggered
```

C.2 `sim_sensitivity.py`

No randomness at all: three deterministic linear-program solves on the identical E under three different antisymmetric payoff functionals (§6). Run: `python3 sim_sensitivity.py`. Full output, verbatim:

```
--- log-ratio (main text) --- value=-0.0000000000
  Great Tusk: 18.1818%
  Gholdengo: 22.7273%
  Kingambit: 13.6364%
  Dragonite: 31.8182%
  Zamazenta: 13.6364%

--- linear difference --- value=-0.0000000000
  Great Tusk: 12.6984%
  Gholdengo: 30.1587%
  Kingambit: 7.9365%
  Dragonite: 41.2698%
  Zamazenta: 7.9365%

--- sign only --- value=-0.0000000000
  Great Tusk: 9.0909%
  Gholdengo: 27.2727%
  Kingambit: 9.0909%
  Ogerpon-Wellspring: 27.2727%
  Iron Valiant: 18.1818%
  Raging Bolt: 9.0909%
```

C.3 `cycle_trajectory_dense.txt`

A 121-row, finer-time-resolution ($t \in [0, 3000]$, step 25) resample of Equation 4's same isolated sub-game, generated purely for Figure 2's plotting resolution; it reports the identical dynamical system as the coarser 8-point trace above, not a different computation.