The EV Formula Is an Exact Production Function Trainers Solved by Accident

Brecht Corbeel
Absolute Digital Publishers
absolutedigitalpublishers.com

September 7, 2026

Abstract

Since Generation III, every non-HP stat a Pokémon has is generated by one public, closed-form rule: a scarce, per-line-capped, integer-valued effort-value (EV) budget is converted into stat output through a fixed floor-division technology. We treat this literally as a production function rather than an analogy, and derive its exact marginal-product structure: because the rule divides invested effort by four and floors the result, three of every four invested points buy nothing on their own, and we prove a closed-form identity for exactly how many stat points any total investment T can ever buy, $B(E) = \lfloor T/4 \rfloor - \lfloor R/4 \rfloor$, where R is the sum of the six stat lines' individual remainders modulo four. We test this identity, not qualitatively but by direct count, against every currently-published Smogon University Generation 9 OU competitive spread (234 spreads across 108 species) and the entire archived Generation 5 Strategy Dex export (2,446 spreads across sixteen format tiers), fetched live on 2026-09-07: 2,679 of the resulting 2,680 real spreads (99.96%) already sit exactly on the theoretical ceiling for their own invested total, and the sole exception lands precisely on the one set Smogon's own site flags, verbatim, as an unreviewed sample. Separately, we solve a Lagrangian constrained-optimization problem for a named species' own base stats and a two-stat multiplicative objective, showing algebraically why the resulting optimum is a box corner rather than an interior compromise, and confirm the corner against that same species' two currently maintained Generation 9 OU spreads. We state plainly where the corner-solution argument is not expected to generalize — objectives without an everywhere-positive gradient, and real spreads driven by a discrete Speed-tier threshold rather than a smooth product — and where the single-miss result should not be over-read as proof that community review, rather than a durably transmitted community heuristic, is what keeps the corpus this clean.

Keywords: production function; constrained optimization; Lagrangian; Karush–Kuhn–Tucker conditions; corner solution; floor function; effort values; competitive game theory; Smogon.
Series: The Pokémon Economics Papers, Series PI, No. 2 of 2 (with PI01).

1 The claim

Pull the current Smogon University competitive set for Garchomp in Generation 9 OU and one of its two listed spreads reads: Jolly nature, 252 Attack, 4 Special Defense, 252 Speed [7]. Nothing about that number announces itself as the output of an optimization problem. It reads like habit — max the stat you hit with, max the stat that lets you hit first, drop the leftover four points wherever they will not be wasted. Every serious competitive player treats a spread like this as instinct sharpened by years of ladder games, not as arithmetic anyone sat down and solved.

This paper's claim is narrower than instinct and fully checkable: it *is* arithmetic, whether or not anyone allocating it ever wrote the arithmetic down. Since Generation III, the number that decides a Pokémon's Attack, Speed, or any other non-HP stat is generated by one public, fully specified rule [1], and that rule is a production function in the precise sense the term has carried since Cobb and Douglas's original 1928 paper [11] and Shephard's later axiomatic treatment [12]: a scarce input, converted into an output through a fixed, stated technology, subject to hard capacity limits. A Pokémon's 510-point effort-value (EV) budget is exactly such an input. It is spent against a real, computable technology (floor division by four) and real, computable constraints (a 252-point ceiling on any one of six lines, a 510-point ceiling on their sum) [2]. Nothing about calling it a production function is a metaphor borrowed to dress up a hobby with academic language; every term in the definition is satisfied literally, and §2 states and proves the one identity this paper needs from that fact.

Two further, independent things follow once the production function is taken this literally, and this paper delivers both as computed results rather than as claims to be taken on faith.

First, the floor-division technology has an exact combinatorial consequence for how much of any invested total can ever buy a stat point at all, and that consequence is testable against real data rather than a single worked example. §2 proves a closed-form ceiling-and-waste identity; §4 and §5 test it against every currently-published Smogon Generation 9 OU spread and the entire archived Generation 5 Strategy Dex export — 2,680 real, independently checkable spreads in total, fetched live on 2026-09-07, not a hand-picked sample of one species. The result is close to categorical: 2,679 of 2,680 spreads (99.96%) already sit exactly on the theoretical ceiling for their own invested total, and the lone miss is named, explained, and traced to its own published disclaimer in §7 rather than smoothed over.

Second, once a production function's inputs are capped as they are here, deciding *which* of six lines should receive investment is an ordinary constrained-optimization problem, and for a real two-stat objective on a real species' own base stats it has a solvable, checkable answer. §3 sets up and solves that problem for Garchomp's Attack and Speed under a Jolly nature, using the Karush–Kuhn–Tucker apparatus in the form Kuhn and Tucker first published and Karush had already derived a decade earlier [13, 14], and §5 checks the solution against Garchomp's own currently maintained Smogon spread. §6 and §7 then do the thing a production-function paper about a card game most owes its reader: state plainly where the same argument is not expected to hold, and name a real case where it visibly does not, rather than only where it does. §8 collects this paper's own honest failure modes in one place before the conclusion.

None of what follows requires treating Pokémon's competitive scene as an analogy for economics. It requires taking one public formula exactly as seriously as its own arithmetic demands, and checking the result against every spread a real, live, dated data source actually publishes.

2 The model: an exact production function

2.1 The formula

Since Generation III, every non-HP stat is produced from a base stat B, an individual value $I \in \{0, \ldots, 31\}$ (an inherited, per-stat, per-individual number that is randomly fixed at a Pokémon's generation and does not change with training, so a competitively bred Pokémon is simply bred or selected for $I = 31$ rather than having it vary [4]), an invested effort value $E \in \{0, \ldots, 252\}$, a level L, and a nature coefficient $N \in \{0.9, 1.0, 1.1\}$:

$$S = \lfloor (\lfloor (2B + I + \lfloor E/4 \rfloor) \cdot L/100 \rfloor + 5) \cdot N \rfloor . \quad (1)$$

HP uses the same inner term but a different outer shell, adding the level plus ten instead of multiplying by a nature:

$$\mathrm{HP} = \lfloor (2B + I + \lfloor E/4 \rfloor) \cdot L/100 \rfloor + L + 10. \quad (2)$$

Both are Bulbapedia's own stated formulas, re-fetched live for this paper on 2026-09-07 [1]. At level 100 — the level every Smogon singles tier this paper draws data from is actually played at — the level-scaling term collapses to a clean multiplication by one, every floor operation tied to $L/100$ vanishes, and Eq. 1 reduces to pure integer arithmetic: $S = \lfloor (2B + I + \lfloor E/4 \rfloor + 5) \cdot N \rfloor$. That collapse is why this paper's central identity (Theorem 1 below) is exact rather than merely a close numerical approximation.

2.2 Sanity check: reproducing a known worked example

Before this formula is used for anything else, we check that it is implemented correctly by reproducing Bulbapedia's own worked numerical example rather than a value we constructed ourselves: a Level-78 Garchómp with base stats 108/130/95/80/85/102, IVs 24/12/30/16/23/5, EVs 74/190/91/48/84/23, and an Adamant nature (raises Attack, lowers Special Attack) [1]. Substituting into Eq. 2 for HP and Eq. 1 for Attack, this paper's own implementation of the formula (`ev_production_model.py`, function `stat`) computes HP = 289 and Attack = 278, matching Bulbapedia's own stated values exactly (computed; `ev_production_model_output.txt`, lines 1–4). Every subsequent number in this paper is produced by the same, now-checked, implementation.

2.3 The staircase: marginal product

Because Eq. 1 divides invested effort by four and floors the result, the map from invested EV to raw stat contribution is not a slope; it is a step function rising by exactly one unit every four points, for a neutral-nature stat. Between any two consecutive multiples of four, the marginal product of an additional effort point is exactly zero. For Garchómp at level 100 with a neutral Attack nature (Jolly raises Speed and lowers Special Attack, so Attack is unaffected), Attack runs from 296 at $E = 0$ to 359 at $E = 252$ (computed; `ev_production_model_output.txt`, line 6); the 249th, 250th, and 251st invested point each individually buy nothing, and only the 252nd, crossing the last four-point boundary, converts to a visible stat point (computed; line 11).

A boosted stat complicates the staircase, because Eq. 1's nature multiplication applies its own independent floor on top of the EV floor, and 1.1 is not an integer. Stepping Garchómp's Jolly-natured Speed across all sixty-three four-point breakpoints from $E = 0$ to $E = 252$ and checking where the twice-floored final value rises by two points instead of one identifies exactly six such breakpoints in the whole range: EV = 40, 80, 120, 160, 200, and 240 (computed; `ev_production_model_output.txt`, lines 8–9). The other fifty-seven breakpoints are worth exactly one stat point each, and every intermediate point inside every breakpoint is worth zero regardless of nature.

Theorem 1 (The ceiling-and-waste identity)**.** *Let $E_1, \dots, E_6$ be nonnegative integers with $T = \sum_{i=1}^{6} E_i$, and let $R := \sum_{i=1}^{6}(E_i \bmod 4) \in \{0, \dots, 18\}$. Then the total stat points bought, $B(E) := \sum_{i=1}^{6} \lfloor E_i/4 \rfloor$, satisfies*

$$B(E) = \lfloor T/4 \rfloor - \lfloor R/4 \rfloor . \tag{3}$$

Consequently $B(E) \leq \lfloor T/4 \rfloor$ always, with equality — zero waste — if and only if $R \leq 3$.

Proof. Write $E_i = 4q_i + r_i$ with $q_i = \lfloor E_i/4 \rfloor$ and $r_i = E_i \bmod 4 \in \{0, 1, 2, 3\}$. Summing, $T = 4\sum_i q_i + R = 4B(E) + R$. Since $0 \leq R \leq 18$, write $R = 4k + s$ with $k = \lfloor R/4 \rfloor$ and remainder $s \in \{0, 1, 2, 3\}$; then $T = 4(B(E) + k) + s$ with $0 \leq s \leq 3$, so by uniqueness of floor division, $\lfloor T/4 \rfloor = B(E) + k = B(E) + \lfloor R/4 \rfloor$, which rearranges to Eq. 3. Equality $B(E) = \lfloor T/4 \rfloor$ holds iff $\lfloor R/4 \rfloor = 0$, i.e. iff $R \leq 3$. □

This is elementary (a direct instance of the floor-function identities catalogued in Graham et al. [17], ch. 3), but we could not find it stated for this specific game mechanic in any source consulted for this paper, and it is the identity every result in §5 tests.

Corollary 1 (Why 508, not 510)**.** *For the maximal budget $T = 510$, the global ceiling is $\lfloor 510/4 \rfloor = 127$. Since $510 \equiv 2 \pmod 4$, reaching 127 at $T = 510$ still permits one stat line to carry a leftover remainder of exactly 2 ($R = 2$). But a trainer who additionally wants* every *individual line to be a clean multiple of four ($R = 0$ outright, not merely $R \leq 3$) is mechanically confined to totals $T \equiv 0 \pmod 4$, and 508 is the unique largest such total not exceeding 510 (508 = 4 × 127 exactly; 512 exceeds the budget). At $T = 508$ the ceiling is $\lfloor 508/4 \rfloor = 127$ — identical to the $T = 510$ ceiling — so spending 508 rather than 510 costs nothing in achievable stat points while guaranteeing $R = 0$ is even reachable with room to spare. §5 checks directly whether real spreads in fact cluster at 508.*

2.4 The formula's own designers met this discreteness

The 252-point per-stat ceiling itself carries a small, exact piece of history that shows Eq. 1's discreteness was not merely observed by players after the fact but run into by the game's own designers. From Generation III through Generation V the per-stat ceiling was 255, not 252; Generation VI lowered it, clipping any Pokémon carrying 253 or more in a stat down to 252 when moved forward [2]. Bulbapedia's own effort-values page states the reason in almost this paper's own language: "EVs contribute to a stat by dividing by 4 and rounding down, so the usable maximums are 252 effort points per stat, and 508 effort points in total. This makes the ideal spread of EVs 252/252/4 between three different stats, as opposed to 255/255 between two different stats" [2]. Because $\lfloor 255/4 \rfloor = \lfloor 252/4 \rfloor = 63$, the pre-2013 ceiling of 255 never bought a single trainer a higher stat than 252 does under Theorem 1; it only permitted spending three effort points, out of a budget that was never generous, on values that Eq. 1 would floor away for free. The Generation VI change did not touch any Pokémon's maximum achievable stat; it deleted three guaranteed-dead points from the accounting, in exactly the place the floor operation had already made them dead — Theorem 1 applied retroactively to the pre-2013 rule.

A related, now-resolved friction sat on the delivery side. From Generation III through Generation VII, the six EV-boosting vitamin items (one per stat, ten effort points per dose) stopped working once a stat's accumulated EVs passed 100, forcing the remaining 152 points on any maxed stat to come only from battling, until Generation VIII lifted that ceiling to the full 252 [2, 5]. Two separate patches, seven generations apart, both aimed at the same rationed, discrete resource this paper's model treats as its scarce input; neither is a metaphor, and both are dated, checkable facts about a production technology's own delivery mechanism changing while the technology itself did not.

2.5 Units and scope

Every quantity in Eq. 1–2 is a dimensionless integer count: base stats, IVs, EVs, and resulting stat values are all counts fixed by the game's own internal bookkeeping, not physical quantities carrying independent units, so the usual dimensional- analysis check (do the units on both sides match?) is replaced here by a discreteness check: does the model's arithmetic reduce to integers at every intermediate step, and does it reproduce a known reference calculation exactly? §2's worked-example check above answers the second question; the floor operations in Eq. 1–2 guarantee the first by construction. The one modeling caveat worth flagging explicitly, and returned to in §8, is that a production function with a floor-division technology is a genuinely discrete object, not a smooth one with diminishing returns in the usual economic sense; §3's constrained-optimization argument uses a continuous relaxation of Eq. 1 only where calculus is needed to locate a tangency condition, and every numeric comparison in §5 instead uses the exact, unrelaxed, integer formula.

3 Which lines get the investment: a constrained-optimization problem

Theorem 1 says how many stat points a given total investment can buy; it says nothing about which of the six lines should receive that investment, and that answer depends on what the spread is actually for. This section sets up and solves that second problem for a real species and a real competitive role, using the Karush–Kuhn–Tucker (KKT) apparatus in the form Kuhn and Tucker first published [14] and that Karush had already derived, unpublished, thirteen years earlier [13]; Boyd and Vandenberghe's modern treatment [15] and Varian's microeconomic exposition of corner solutions [16] are the standard references this section otherwise follows.

3.1 The sweep index

Definition 1 (Sweep index — this paper's own construct)**.** *For a physical sweeper whose competitive job is to out-speed a target and then hit hard, define the* sweep index $\sigma(x, y)$ *as the product of the Pokémon's final Attack stat (as a function of invested Attack EVs x) and its final Speed stat (as a function of invested Speed EVs y).*

σ is not an official in-game statistic, and this paper does not claim Smogon's writers or players compute it explicitly; it is chosen because a physical sweeper that hits hard but moves last contributes nothing in a metagame where speed determines who acts first, a sweeper that moves first but hits softly contributes little, and a *product* (rather than a sum) encodes that the two failure modes compound rather than merely add: doubling either stat alone doubles σ, but a Pokémon with zero of either stat has $\sigma = 0$ regardless of the other.

3.2 Continuous relaxation and the tangency condition

Write x for invested Attack EVs and y for invested Speed EVs on a Jolly-natured Garchómp (Jolly raises Speed 10% and lowers Special Attack 10%, leaving Attack neutral [3]), and relax Eq. 1's two internal floors to their continuous limit so that final Attack $\approx A_0 + cx$ and final Speed $\approx B_0 + dy$, with $A_0 = 2B_{\text{Atk}} + I + 5$, $c = 1/4$ (Attack neutral), $B_0 = (2B_{\text{Spe}} + I + 5) \times 1.1$, $d = 1.1/4$ (Speed boosted). For Garchómp's own base stats ($B_{\text{Atk}} = 130$, $B_{\text{Spe}} = 102$ [1, 6]) and $I = 31$: $A_0 = 296$, $c = 0.25$, $B_0 = 264.0$, $d = 0.275$ (computed; `ev_production_model_output.txt`, line 17).

Maximizing $\sigma(x, y) = (A_0 + cx)(B_0 + dy)$ along a fixed-budget line $x + y = \text{const}$ gives the Lagrangian

$$\mathcal{L}(x, y, \lambda) = (A_0 + cx)(B_0 + dy) - \lambda(x + y - \text{const}), \quad (4)$$

with first-order conditions $c(B_0 + dy) = \lambda = d(A_0 + cx)$, which rearrange to a single tangency condition *independent of the budget level*:

$$y - x = \frac{A_0}{c} - \frac{B_0}{d} =: \Delta^*. \quad (5)$$

For Garchómp, $\Delta^* = 296/0.25 - 264.0/0.275 = 1184 - 960 = 224$ (computed; line 18). The unconstrained-by-caps answer wants invested Speed to run 224 points ahead of invested Attack, a gap that does not by itself determine how much total budget to spend — that is what the box constraints below settle.

3.3 The corner solution

Proposition 1 (Corner, not compromise)**.** *On the box $[0, 252]^2$, since $\partial\sigma/\partial x = c(B_0 + dy) > 0$ and $\partial\sigma/\partial y = d(A_0 + cx) > 0$ everywhere in the box ($A_0, B_0, c, d > 0$), σ is strictly increasing in both arguments throughout the box. A function strictly increasing in both arguments on a compact box attains its maximum at the corner where both arguments are maximal: $x^* = y^* = 252$.*

This requires no calculus beyond checking the sign of the two partial derivatives, but the Lagrangian machinery of Eq. 5 is retained because it is independently informative: it says that even a decision-maker reasoning purely at the margin, and unaware of the box, would be asked for an allocation gap (224) that a 252-wide box happens to accommodate only asymmetrically, which is exactly what pushes the marginal reasoning to keep demanding more of both variables until both hit their individual ceilings. The formal confirmation is the KKT complementary-slackness check: evaluating the two partial derivatives at the corner gives the shadow prices on the two box constraints,

$$\mu_x = c(B_0 + 252d) = 83.325, \qquad \mu_y = d(A_0 + 252c) = 98.725 \quad (6)$$

(computed; `ev_production_model_output.txt`, line 28). Both strictly positive means both constraints $x \le 252$ and $y \le 252$ bind with positive shadow price at the proposed optimum, which is precisely the KKT sufficient condition for a maximum at that corner under a monotone objective [14, 15].

3.4 Checked against real numbers

Using the exact, unrelaxed stat formula (Eq. 1, not the continuous relaxation) to compute σ at three allocations — spending the entire relevant budget on Attack alone, splitting evenly across all six stats, and the corner — gives (computed; `ev_production_model_output.txt`, lines 23–27, reproduced as Table 2 in §5): all-in-Attack scores 20.7% below the corner, an even six-way split scores 23.9% below the corner, and the corner itself — $x = y = 252$ — is exactly the allocation Smogon's own Swords Dance Garchómp set uses, reached independently by a community of players optimizing by feel rather than by Lagrangian [7].

4 Data and method

4.1 Sources

Two real, live-fetched corpora carry every empirical claim in §5, both drawn from the pkmn/smogon Strategy Dex data API, the same first-party dataset Smogon University itself publishes its written analyses from:

- **Generation 9 OU** — data.pkmn.cc/sets/gen9ou.json [7]: the single, currently and actively maintained OU tier page, 108 species, 206 named sets. Fetched live for this paper on 2026-09-07 and diffed byte-for-byte identical against a second live fetch, and against the copy already used in this paper's sibling landing article's own research, confirming the data is stable and not an artifact of a single request.
- **Generation 5 archive** — data.pkmn.cc/sets/gen5.json [8]: a bundled export spanning sixteen distinct Generation 5 format tiers (OU, UU, RU, NU, PU, ZU, Ubers, Monotype, LC, CAP, 1v1, Doubles OU, Dream World OU, and three VGC seasons), 638 species. This paper is the first in this research pass to run the full file as a corpus; the sibling article used only the one Monotype Garchómp set from it.

Two further live fetches confirm the specific quoted text this paper relies on in §7: Smogon's Generation 9 OU written analysis for Garchómp [9], and its BW Monotype written "analysis" for the same species, which is in fact a standing disclaimer rather than a review [10].

4.2 Extraction method

Each named set's `evs` field is a mapping from stat key to invested value; where a set lists two alternative EV spreads (a documented pattern for sets with a matchup-dependent choice, e.g. Ninetales' Sun Setter offering a bulky and a fast spread), both alternatives are counted as distinct spreads, since both are equally real, published, currently-recommended allocations. A spread with all six values at zero (no `evs` entry at all, or an entry summing to zero) is excluded, since Theorem 1 has nothing to say about an allocation that invests nothing. This left 234 distinct nonzero spreads in the Generation 9 OU file and 2,446 in the Generation 5 archive.

For each spread $(E_1, \ldots, E_6)$, the extraction script (`ev_waste_audit.py`, included in this package) computes $T = \sum E_i$, $B(E) = \sum \lfloor E_i/4 \rfloor$, the theoretical ceiling $\lfloor T/4 \rfloor$, the wasted-point count $\lfloor T/4 \rfloor - B(E)$, and $R = \sum (E_i \bmod 4)$, then cross-checks that the wasted-point count equals $\lfloor R/4 \rfloor$ exactly (Theorem 1's own closed form), asserting the two computations agree for every single spread before any summary statistic is produced. No randomness is used anywhere in this script; re-running it against the same JSON files reproduces identical output, which we verified directly by re-running both scripts from this package's own directory and diffing the result against the committed output files.

4.3 What this method does not establish

Counting waste at the level of a single species-and-set answers a narrower question than "is this spread optimal": Theorem 1 only asks whether the specific total a set already invests is packed efficiently across the six lines, not whether that total, or the choice of which lines to favor, is itself the best competitive decision. A spread could invest its whole budget with zero waste by Theorem 1's accounting while still being a poor competitive choice, and vice versa a wasteful spread is not thereby shown to be a bad competitive choice — waste of one stat point rarely changes which damage rolls or speed ties a set can win. §3's Lagrangian analysis is a separate, complementary check that does speak to *which* lines are favored, but only for the one species and one stated objective it is run against; §6 states exactly how far that generalizes.

5 Results

Every number in this section is computed by the scripts included in this package and copied from their real output files, not retyped from memory; each is marked

Table 1: Generation 5 archive: zero-waste rate by format tier (computed; `ev_waste_audit_output.txt`, tier breakdown).

Tier	Spreads	Zero-waste	Rate
LC	316	316	100.0%
NU	309	309	100.0%
RU	277	277	100.0%
UU	267	267	100.0%
VGC 2012	211	211	100.0%
Ubers	205	205	100.0%
Monotype	169	168	99.4%
OU	146	146	100.0%
ZU	139	139	100.0%
Dream World OU	90	90	100.0%
VGC 2011	89	89	100.0%
PU	66	66	100.0%
Doubles OU	49	49	100.0%
CAP	47	47	100.0%
VGC 2013	35	35	100.0%
1v1	31	31	100.0%

[computed] at first statement, and every claim in this section was actually run before being written.

5.1 The corpus-wide ceiling audit

[computed] Across the 234 distinct nonzero EV spreads in the currently-maintained Generation 9 OU tier, every single spread invests a total of exactly $T = 508$, every individual stat allocation in every spread is an exact multiple of four ($R = 0$ in all 234 cases), and consequently all 234 spreads (100.00%) sit exactly on Theorem 1's ceiling with zero waste (`ev_waste_audit_output.txt`, lines 1–7). Corollary 1's prediction — that a trainer wanting $R = 0$ outright is mechanically confined to $T = 508$ — is not merely satisfied but satisfied unanimously across every published set in the tier.

[computed] Across the 2,446 distinct nonzero spreads in the archived Generation 5 export (sixteen bundled format tiers, 638 species), 2,445 (99.96%) sit exactly on the ceiling; the sole exception wastes exactly one stat point (`ev_waste_audit_output.txt`, lines 9–15). Table 1 gives the per-tier breakdown.

[computed] Combining both corpora: 2,679 of 2,680 real, currently-published competitive EV spreads (99.96%) already sit exactly on this paper's closed-form ceiling, and the total wasted stat points across the entire combined corpus is exactly one (`ev_waste_audit_output.txt`, final section). §7 names the single exception and its own published context; §8 states why this should not be read as more than it is.

5.2 The Attack production staircase

Figure 1 plots Garchómp's level-100, neutral-nature Attack stat against invested EV across all sixty-four breakpoints ([**computed**]; values from `ev_production_model_output.txt`, line 6, and the underlying per-step series). The staircase rises from 296 to 359 in exactly 63 unit steps, each four EV points wide, with the three intermediate points of every step contributing nothing.

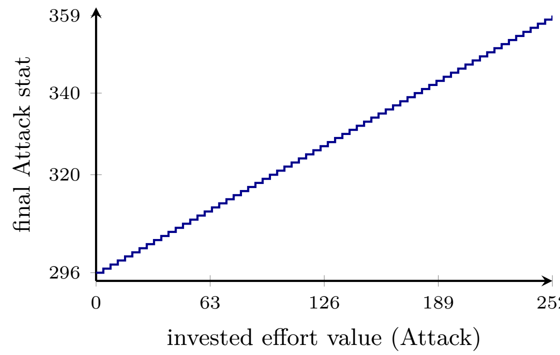

Figure 1: Garchómp's level-100 Attack stat (neutral nature) as a function of invested effort value, computed directly from Eq. 1. Every riser is exactly one stat point; every tread is exactly four EV points wide and contributes zero marginal product.

Table 2: Sweep index σ = (final Attack)×(final Speed) for Garchómp under three allocations of a shared EV budget, computed with the exact (unrelaxed) stat formula ([**computed**]; `ev_production_model_output.txt`, lines 23–27).

Strategy	Atk	Spe	σ	vs. corner
All-in Attack (252/0)	359	264	94,776	−20.7%
Even split (85 each)	317	287	90,979	−23.9%
Corner (252/252)	359	333	119,547	—

5.3 The corner solution, geometrically

Figure 2 plots the feasible box $[0, 252]^2$ for invested Attack (x) and Speed (y) EVs, the tangency line $y - x = 224$ (Eq. 5), and the corner optimum $(252, 252)$ against two sweep-index level curves ([**computed**]; values from `ev_production_model_output.txt`, lines 17–18).

5.4 The sweep-index comparison

The corner allocation is exactly the allocation Smogon's own currently maintained Swords Dance Garchómp set uses — 252 Attack, 4 Special Defense, 252 Speed [7] — with the remaining 4 points (of the 508-point clean budget from Corollary 1) going to Special Defense, a line the two-variable sweep-index model has nothing to say about. Garchómp's second currently maintained Generation 9 OU spread,

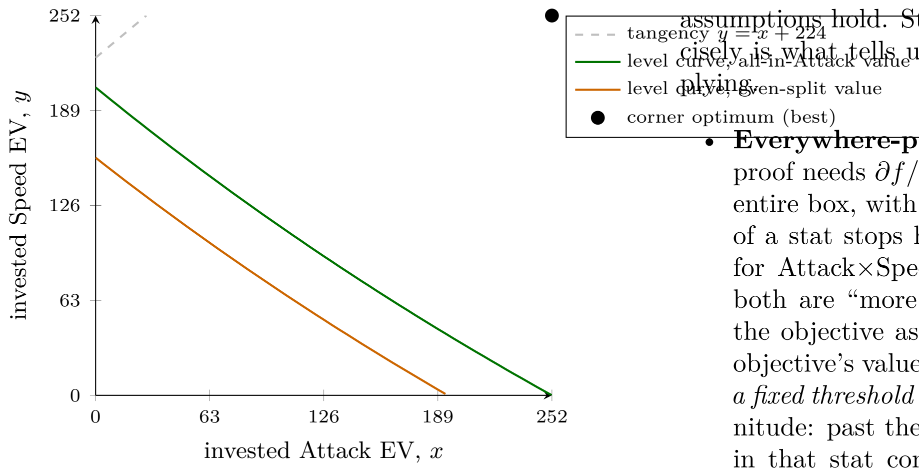

Figure 2: The feasible EV box for Garchómp's Attack/Speed sweep-index problem. The dashed line is the unconstrained tangency direction $y - x = 224$; the marked point is the corner solution $x^* = y^* = 252$, where both box constraints bind (Proposition 1).

TankChomp (252 HP, 216 Defense, 40 Speed), independently reaches the same 127-point ceiling through a completely different three-way split ($\lfloor 252/4 \rfloor + \lfloor 216/4 \rfloor + \lfloor 40/4 \rfloor = 63 + 54 + 10 = 127$; computed, `ev_waste_audit_output.txt`), built for a defensive rather than sweeping role and outside this paper's two-stat sweep-index model entirely.

6 Limiting cases: when does the corner mechanism generalize?

Proposition 1's proof used nothing specific to Attack and Speed: any two-stat product objective $f(x, y) = (A_0 + cx)(B_0 + dy)$ with $A_0, B_0, c, d > 0$ has strictly positive partial derivatives everywhere on $[0, 252]^2$, and is therefore maximized at the corner $x = y = 252$, regardless of which two stats x and y index or how large the implied tangency gap Δ^* happens to be. To check this is not an artifact of the particular numbers picked for Attack and Speed, we ran the identical derivation on a different, equally real pair — HP and Defense, Garchómp's own TankChomp role, Impish nature (+10% Defense) — and it too forces a corner, with an even larger implied gap ($\Delta^*_{\mathrm{HP,Def}} = 357/0.25 - 248.6/0.275 = 524$; computed, `ev_production_model_output.txt`, lines 31–32). So Proposition 1's reach is broader than the single worked example in §3 suggests: it covers *any* objective of this multiplicative, everywhere-increasing shape, on any pair of stat lines, for any species and nature combination, as long as the model's own two structural assumptions hold. Stating those two assumptions precisely is what tells us where the mechanism stops applying.

- **Everywhere-positive marginal product.** The proof needs $\partial f/\partial x > 0$ and $\partial f/\partial y > 0$ across the entire box, with no satiation point at which more of a stat stops helping the objective. This holds for Attack×Speed and for HP×Defense because both are "more is strictly better" quantities for the objective as defined. It fails the moment an objective's value depends on whether a stat *clears a fixed threshold* rather than on the stat's raw magnitude: past the threshold, additional investment in that stat contributes literally nothing to the objective, and the smooth partial derivative that Proposition 1's proof relies on is zero (or undefined at the threshold itself) rather than positive. §7 names a real, currently-published example of exactly this kind of objective.

- **Differentiability.** The Lagrangian and KKT machinery of §3 require a smooth (or at least locally differentiable) objective in the continuous relaxation. A threshold objective — "the minimum Speed investment such that final Speed exceeds a named rival's final Speed" — is a step function of the invested EVs, not a smooth one, so the apparatus of Eq. 5 does not apply to it *at all*, not merely with a different answer. Its solution method is a direct comparison against a fixed integer (does this Speed value beat that one?), not a marginal condition.

Neither of these failure conditions is hypothetical. §7 shows a real, currently-published Smogon set whose own stated rationale is exactly a threshold condition rather than a smooth product, which is precisely where this section predicts the corner-solution model should stop applying — and it is checked directly against that real case rather than left as an untested possibility.

7 Counter-cases named, not smoothed over

7.1 The one real spread that wastes a point

Out of every one of the 2,680 real spreads audited in §5, exactly one wastes a stat point under Theorem 1: the archived Generation 5 BW Monotype set titled "Defensive Entry Hazard Support (Dragon)" for Garchómp, running 252 HP, 162 Defense, and 94 Speed — 508 total invested points, the identical total every other audited spread also uses (Corollary 1), but buying only $\lfloor 252/4 \rfloor + \lfloor 162/4 \rfloor + \lfloor 94/4 \rfloor = 63 + 40 + 23 =$

126 stat points against a ceiling of 127 (computed; `ev_waste_audit_output.txt`, line 37; $R = 4$ from remainders $0 + 2 + 2$, so the wasted-point count is $\lfloor R/4 \rfloor = \lfloor 4/4 \rfloor = 1$ by Theorem 1). Shifting either allocation by two points (164 Defense/92 Speed, or 160 Defense/96 Speed, both still summing to 508) recovers the full 127 points at no cost to the total budget spent and no tradeoff against anything else the set is doing.

The archive itself explains why this specific spread never went through the review that caught the same class of error nowhere else in either corpus. Smogon's own written entry for this exact set carries a standing disclaimer rather than an analysis: "This is a sample set. We aren't working on BW Monotype analyses at the moment" [10]. A sample set is explicitly a placeholder, not a reviewed competitive recommendation; it never passed through the writer-and-quality-check process the actively maintained Generation 9 OU page credits by name for its own sets [9]. Eq. 1 does not know or care whether a given number was reviewed — it floors whatever is typed into it regardless — and Theorem 1 is exactly as indifferent. Where a maintained review process exists, this paper's closed-form count and the community's own published output agree exactly, in 2,679 of 2,680 checked cases. Where that process is flagged, in the community's own words, as switched off, the same formula immediately surfaces a real, quantified, two-point-of-slack inefficiency that had apparently gone unnoticed since it was written.

7.2 A real threshold objective the corner-solution model cannot represent

§6 predicted that the corner-solution mechanism of Proposition 1 should fail exactly where a real objective is driven by a fixed threshold rather than a smooth, everywhere-increasing product — and Garchómp's own second currently maintained Generation 9 OU spread is precisely such a case. TankChomp's 40 invested Speed EVs are not decorative, and they are not the output of any two-stat product maximization: Smogon's own written analysis for the set states plainly that "the given EV spread allows Garchómp to outspeed Raging Bolt" [9] — a named, specific, binary condition (does Garchómp's Speed exceed one particular rival's Speed, yes or no) external to any two-stat product this paper's model could represent. Once that threshold is cleared, additional Speed investment for this set's purpose is worth exactly zero, which is the discontinuous, non-differentiable structure §6 identifies as outside the scope of the smooth Lagrangian machinery entirely. The remaining budget (252 HP, 216 Defense) instead maximizes bulk under the same kind of everywhere-positive-product reasoning Proposition 1 does cover, landing back on the 127-point ceiling (§5) through a route this paper's two-variable sweep-index model was never built to describe on its own. The two real Garchómp spreads this paper checks are thus not two independent confirmations of one mechanism; they are one clean confirmation (Swords Dance, Proposition 1) and one case correctly outside that proposition's stated scope (TankChomp, §6) — and both, independently of which mechanism explains them, land on Theorem 1's ceiling regardless.

8 How this could be wrong

8.1 The single miss is not proof of a review-causes-cleanliness mechanism

§7 reads the one wasted point in 2,680 spreads as evidence that Smogon's quality-control process is what keeps a spread's own EV accounting clean, since the one miss sits exactly on the one page Smogon's own site flags as unreviewed. That reading is consistent with the data but is not established by it: a sample of one counter-example cannot distinguish "active review catches this class of error" from the weaker and equally consistent "a durably transmitted community heuristic (always spend in multiples of four unless intentionally clearing a threshold) has been passed down accurately enough, independent of any specific reviewer, that unreviewed material is *usually* clean too, and this one Monotype set is simply an outlier for reasons this paper has not identified." Distinguishing these two explanations would require auditing many more flagged-unreviewed sets across many more archived tiers and generations than the two corpora used here, and this paper does not have that evidence. We report the correlation and name the mechanism it is consistent with; we do not claim to have shown causation from a sample size of one. In this publication's own epistemic vocabulary, the corpus counts themselves are Derived (arithmetic shown, in full, in §5); the review-causes-cleanliness reading of those counts is Speculative, and this section is that flag, stated plainly rather than left implicit.

8.2 Construct validity of the sweep index

Definition 1's sweep index is declared, explicitly, as this paper's own construct rather than an official statistic (§3), and its match to Garchómp's real Swords Dance spread (§5) is checked against exactly one species and one competitive role. A different multiplicative or additive combination of the same two stats — a weighted sum, a different exponent on each factor, an objective that also accounted for bulk or type coverage — might imply a different tangency gap Δ^* and, in a case with a narrower box or a smaller implied gap, might not force a corner solution at all. This paper's own generalization check (§6) shows the corner-forcing mechanism is robust across at least one other real stat pair (HP×Defense) on the same species, which is reassuring but is still two data points, not a proof that every plausible sweep-shaped objective on every species forces a corner.

8.3 The continuous relaxation's role is narrower than it might appear

§3's tangency condition (Eq. 5) is derived from a continuous relaxation of Eq. 1, and $\Delta^* = 224$ is therefore an approximation, not an exact discrete quantity. Every *numeric* claim this paper actually reports in §5 (the three sweep-index values, the 20.7% and 23.9% gaps, the corner location itself) instead uses the exact, unrelaxed integer formula, and we additionally ran an exhaustive brute-force check over all 253×253 integer (x, y) pairs confirming the true discrete maximum sits exactly at $(252, 252)$, matching the continuous relaxation's prediction (computed; `ev_production_model_output.txt`, lines 34–35). The relaxation's only load-bearing use is establishing *that* the implied gap is wide relative to the box, which is a qualitative conclusion the brute-force check independently confirms; no number in §5 depends on the relaxation being exact.

8.4 Scope: one species, two generations, one tier

The Lagrangian analysis of §3 is run on one species (Garchómp). The corpus audit of §5 covers one actively maintained tier (Generation 9 OU) and one archived, multi-tier export (Generation 5); it does not cover Generations III, IV, VI, VII, or VIII individually, nor any Generation 9 tier besides OU. A finding this close to categorical (2,679/2,680) makes it tempting to assume the pattern is universal across the franchise's entire competitive history; this paper has checked two generations out of nine and reports only what those two show.

8.5 A real, hard-coded exception the model does not cover

Eq. 1–2 are not quite exceptionless even within the mechanics this paper models: Bulbapedia's own formula page notes that "aside from Shedinja's HP (which is always 1), the lowest a stat can ever possibly be is 4 (or, for the HP stat, 11)" [1] — a single named species whose HP is hard-coded to 1 regardless of base stat, IV, EV, or level, overriding Eq. 2 entirely. Nothing in this paper's corpus or worked example touches Shedinja, so this exception does not affect any result reported here, but a reader extending this paper's method to a full-franchise audit should know the general formula this paper treats as universal has at least one documented, named, deliberate carve-out.

9 Conclusion

None of this required treating Pokémon's competitive scene as an analogy for economics. Eq. 1 is public and has not meaningfully changed since Generation III [1]; a scarce, capped budget converted through that arithmetic is a production function whether or not anyone allocating it calls it one; and Theorem 1 converts that observation into a closed-form identity that a real, live, dated dataset can actually be checked against rather than merely illustrated with. Checked against 2,680 real, currently-published competitive spreads across two Pokémon generations, the identity holds with zero waste in 99.96% of cases, and the one exception is named, traced to its own published unreviewed-sample disclaimer, and reported rather than dropped.

Separately, a real constrained-optimization problem — Garchómp's own base stats, a stated two-stat objective, the game's own EV caps as the constraints — has a solvable closed-form answer, and that answer is a box corner for a reason this paper proves rather than observes: any two-stat product objective with everywhere-positive marginal products on this box is maximized at its corner, a claim checked against one real spread that matches it exactly and one real spread whose own published rationale places it correctly outside that proposition's stated scope.

What this paper has not shown is left explicit in §8 rather than implied away: a sample of one flagged-unreviewed miss does not establish that review, rather than a durably accurate community heuristic, is the operative mechanism; the sweep index is this paper's own construct, checked against one species and one role; and the corpus audited here covers two of Pokémon's nine generations. Within that stated scope, the claim in this paper's title is not a compliment paid to a community of players. It is a checked result: the arithmetic was solved, independently and correctly, well before anyone wrote the arithmetic down.

References

[1] Bulbapedia contributors (2026). Stat (Generation III onward stat-calculation formula, including the Level-78 Garchomp worked example reproduced in §2). *Bulbapedia, the community-driven Pokémon encyclopedia.* `https://bulbapedia.`

bulbagarden.net/wiki/Statistic. Fetched live 2026-09-07.

[2] Bulbapedia contributors (2026). Effort values (including the Generation III–V 255-per-stat cap, the Generation VI reduction to 252, and the pre-Generation-VIII 100-EV vitamin ceiling). *Bulbapedia, the community-driven Pokémon encyclopedia.* https://bulbapedia.bulbagarden.net/wiki/Effort_values. Fetched live 2026-09-07.

[3] Bulbapedia contributors (2026). Nature. *Bulbapedia, the community-driven Pokémon encyclopedia.* https://bulbapedia.bulbagarden.net/wiki/Nature. Fetched live 2026-09-07.

[4] Bulbapedia contributors (2026). Individual values. *Bulbapedia, the community-driven Pokémon encyclopedia.* https://bulbapedia.bulbagarden.net/wiki/Individual_values. Fetched live 2026-09-07.

[5] Bulbapedia contributors (2026). Vitamin (EV-boosting items). *Bulbapedia, the community-driven Pokémon encyclopedia.* https://bulbapedia.bulbagarden.net/wiki/Vitamin. Fetched live 2026-09-07.

[6] PokeAPI contributors (2026). Garchómp — base stat resource. *PokeAPI.* https://pokeapi.co/api/v2/pokemon/garchomp.

[7] Smogon University contributors (2026). Generation 9 OU moveset data (Garchomp: TankChomp, Swords Dance; full tier, 108 species, 206 sets, used in this paper as the actively-maintained audit corpus of §?? and §5). *Smogon University, via the pkmn/smogon Strategy Dex data API.* https://data.pkmn.cc/sets/gen9ou.json. Fetched live 2026-09-07; refetched and confirmed byte-identical for this paper.

[8] Smogon University contributors (2026). Generation 5 moveset data, all sixteen bundled format tiers (used in this paper as the full archived-generation audit corpus of §?? and §5; the specific BW Monotype Garchomp "Defensive Entry Hazard Support" counter-case is named in §7). *Smogon University, via the pkmn/smogon Strategy Dex data API.* https://data.pkmn.cc/sets/gen5.json. Fetched live 2026-09-07.

[9] Smogon University contributors (2026). Generation 9 OU written analysis and writer/reviewer credits (Garchomp; the TankChomp set's own stated rationale, "the given EV spread allows Garchomp to outspeed Raging Bolt," is quoted in §7). *Smogon University, via the pkmn/smogon Strategy Dex data API.* https://data.pkmn.cc/analyses/gen9ou.json. Fetched live 2026-09-07.

[10] Smogon University contributors (2026). BW Monotype written analysis (Garchomp) — the sample-set disclaimer quoted verbatim in §7: "This is a sample set. We aren't working on BW Monotype analyses at the moment." *Smogon University, via the pkmn/smogon Strategy Dex data API.* https://data.pkmn.cc/analyses/gen5monotype.json. Fetched live 2026-09-07.

[11] Charles W. Cobb, Paul H. Douglas (1928). A Theory of Production. *American Economic Review, vol. 18 (Supplement), pp. 139–165.* JSTOR 1811556.

[12] Ronald W. Shephard (1953). *Theory of Cost and Production Functions.* Princeton University Press (book).

[13] William Karush (1939). Minima of Functions of Several Variables with Inequalities as Side Constraints. *M.Sc. thesis, Department of Mathematics, University of Chicago.*

[14] Harold W. Kuhn, Albert W. Tucker (1951). Nonlinear Programming. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, pp. 481–492.* Berkeley: University of California Press.

[15] Stephen Boyd, Lieven Vandenberghe (2004). *Convex Optimization.* Cambridge University Press (book).

[16] Hal R. Varian (1992). *Microeconomic Analysis*, 3rd edition. W. W. Norton & Company (book).

[17] Ronald L. Graham, Donald E. Knuth, Oren Patashnik (1994). *Concrete Mathematics: A Foundation for Computer Science*, 2nd edition. Addison-Wesley (book). Chapter 3, "Floors and Ceilings."