

Pricing a Good That Costs Nothing to Copy

Discount dynamics and demand on Steam

Preliminary and incomplete — comments welcome

Aleksandr Shneider¹ Arseniy Stolyarov²

October 2026

¹University of Minnesota ²New Economic School

An unusually clean pricing problem

A released PC game is close to a pure **information good**:

- marginal cost of a copy = 0; no inventory, no depreciation, infinite shelf life
- **no capacity constraint**: any number of copies can be sold at any moment — no stock-outs, no production lead times, no reason to clear inventory
- **durable**: one purchase per consumer, consumed repeatedly, no resale
- competes for a scarce input — the player's *time and attention* — with new releases *and* with the games the consumer already owns
- the only “cost” of a sale is the revenue lost on consumers who would have paid more

⇒ the posted price is a pure **allocation and screening device**: intertemporal price discrimination (Stokey 1979; Bulow 1982), the Coase problem for a durable-good seller, skimming with forward-looking consumers (Nair 2007; Conlon 2012)

And the platform shapes the game: a 30-day cooldown between discounts, a 90% cap, no list-price increase right before a sale, and four announced platform-wide *seasonal sales* — a coordinated, anticipated event structure.

Questions

1. **How do sellers price such a good over its life cycle?**

How often, how deep, how long, and when — and what role do the platform's rules and calendar play?

2. **How does demand respond?**

How elastic is the purchase flow to a temporary price cut? Is the response a price effect or an *attention* effect? Do consumers wait, and do firms cannibalise later sales?

3. **Who buys when?**

At the user level: do consumers time purchases to sales, do they anticipate predictable sales, and are there persistent types?

4. **[Next step] Equilibrium discount policy and the welfare role of coordinated sales and the cooldown rule.**


Today: new facts on (1), reduced-form evidence on (2) and (3), a model sketch for (4).

Data: prices, activity, products, users

Block	N	Units	Period
US posted-price history (IsThereAnyDeal, Steam store)	5.60m price events	167,507 apps	2020-01 – 2026-06
Discount spells, main sample, known start	1.43m spells	71,542 games	2020-01 – 2026-06
Weekly price panel, main sample	15.6m game-weeks	79,672 games	2020-01 – 2026-05
Dated, deduplicated reviews	154.2m reviews	146,022 apps	2010 – 2026-04
Balanced weekly activity panel (reviews, Steam-purchase reviews, mean CCU)	0.88m game-weeks	8,424 games	2024 – 2025
Daily peak concurrent players (SteamDB)	5.0m game-weeks	21,298 games	2020-01 – 2026-05
Sampled users: profiles, libraries, all their reviews	6.6m reviews, 30.3m library rows	139,966 users	May 2026
Sampled users' reviews matched to the price at the time of the review	3.5m reviews	100,333 users	2020-01 – 2026-05
External sales estimates (Gamalytic, VG Insights)	18.5k games		2026 snapshot
Product metadata (release, genres, tags, publisher)	218,981 apps	140,017 paid games	March 2026

- **Main sample** = paid games (free-to-play excluded) minus 9.6k NSFW titles and 14.0k titles currently in early access; depth regressions also drop 13.7k games that ever raised their list price (early-access graduations, repricing). “Budget” below means a *paid* game under \$10.
- A discount *spell*: a maximal run of states with cut > 0 ; depth = cut at start; duration = time to the first zero-cut state.
- What we do *not* observe: transactions, wishlist histories, update/patch dates. Reviews are dated, require ownership, carry a Steam-purchase flag and playtime at review; one review $\approx$ 30–40 sales (median reviews/sales 2.5% Gamalytic, 3.3% VG Insights, flat across tiers).

Price paths by budget tier




List price dashed, posted price solid; shaded: the four seasonal sales. AAA titles discount later and more shallowly at first; AA and indie titles ratchet down fast.

Fact 1: discounts are the price instrument


- 2.8 discount starts per game-year (4.6 among games fully observed in 2025) vs. 0.10 list-price changes per game-year
- 65% of the within-game variance of the log posted price comes from discounts
- list-price changes are more often cuts (29.2k) than increases (19.5k)
- the typical discounted game is on sale 14% of the time; 30% of listed games are not discounted in a given year
- 53% of new releases carry a launch discount (median 15%); 73% are discounted within a year

How often is a game discounted? (n = 61,616 games fully observed in 2025)





Fact 2: timing follows the platform calendar ...

Daily number of discount starts, paid Steam games (US store). Shaded: 4 seasonal sales (orange), other Valve-wide sales (yellow)



43% of all discount starts fall inside the four seasonal sales, 5% inside other Valve-wide sales, **51% are publisher-timed**. In the week a seasonal sale starts, 57% of eligible games start a discount.


... and the 30-day cooldown rule



Median gap between consecutive publisher-timed discounts: 32 days; 63% restart within 35 days. The weekly hazard of a new discount is 0.4% for never-discounted games, 22% in the week the cooldown expires, then declines monotonically. (The 7% inside the window are mostly smaller Valve fests, which waive the rule, and the pre-2023 six-week rule.)

Full sample: 11.8m game-weeks outside Valve-wide sales, game not on sale at the start of the week. `tab_hazard_by_cooldown`.

Fact 3: depth is deep, discrete, and the same inside and outside sales




tab_depth_by_timing, tab_duration_by_timing, tab_modal_depths.

	publisher-timed	seasonal
spells	733k	617k
mean depth (%)	55.6	55.4
median depth (%)	51	51
share $\geq 50\%$	0.69	0.70
median duration (days)	7.0	13.8

Modes: 50% (16% of spells), 80, 75, 51, 90, 60, 30, 40.

Durations cluster at 7 and 14 days (publisher-timed) and at the sale length (seasonal).

Fact 4: a life-cycle ratchet



Depth at start: 20% at launch → 36% in months 0–3 → 47% at 6–12 months → 54% at 1–2 years → 62% at 3–5 years → 71% after 8 years.

Within game, relative to the previous discount: 64% same depth, 23% deeper, 13% shallower. 60% of spells are at or above the game's running maximum.

The posted price path is a *sawtooth with a falling floor*: the list price almost never moves, the depth ratchets down, each cut is temporary.

tab_depth_by_age, facts_descriptives.json.

What determines the depth of a discount?

Dependent variable: depth at start (pp). Main sample *excluding* early-access titles and games that ever raised their list price; 2020–2026; $0 < \text{depth} < 100$. SE clustered by publisher.

	(1) all + publisher FE	(2) all game FE	(3) publisher-timed + publisher FE	(4) seasonal only + publisher FE	(5) + games that raised list price, + publisher FE
log list price	6.34*** (0.41)		6.70*** (0.46)	6.02*** (0.37)	7.09*** (0.39)
log(1 + months since release)	10.57*** (0.37)	11.48*** (0.31)	10.57*** (0.40)	10.19*** (0.38)	10.47*** (0.32)
log(1 + prior discounts)	4.37*** (0.34)	2.91*** (0.34)	4.44*** (0.37)	4.82*** (0.35)	3.69*** (0.32)
log(1 + cumulative reviews)	-2.00*** (0.15)	-0.99*** (0.37)	-1.99*** (0.16)	-2.02*** (0.14)	-2.15*** (0.14)
share of positive reviews	-5.55*** (1.08)	-4.26*** (1.55)	-5.29*** (1.20)	-5.75*** (1.04)	-3.34*** (1.02)
game ever raised list price					-1.75*** (0.44)
start in a seasonal sale	0.54*** (0.08)	0.43*** (0.08)			0.56*** (0.08)
genre dummies, year-month FE	yes	yes	yes	yes	yes
N	527,383	524,770	266,706	223,630	661,385
R ² within	0.376	0.220	0.396	0.375	0.342

A game twice as expensive discounts 4.4 pp deeper; one year after release the depth is ≈ 27 pp deeper than at launch; popular and well-rated games discount a little less. Adding the previous depth (coef. 0.69) cuts the age gradient to 2.8 pp: **most of the ratchet is persistence**. Early-access titles are excluded because their list price rises at 1.0 (e.g. Titan Quest 2); games that raised their price discount 1.8 pp less, and nothing else moves.

When does a publisher start a discount?

Linear probability model, weekly, in pp. Game not on sale at the start of the week. Full main sample; SE clustered by game.

	(1) all weeks, game FE	(2) non-seasonal, game + week FE	(3) non-seasonal, + activity
cooldown <28 d (vs. never discounted)	-0.8***	-4.5***	-4.9***
cooldown 28-35 d	14.0***	16.5***	16.9***
cooldown 36-49 d	7.6***	6.3***	8.5***
cooldown 50-70 d	3.1***	1.9***	4.0***
cooldown 71-98 d	3.8***	1.0***	3.4***
cooldown >1 year	-2.9***	-2.4***	-3.4***
seasonal sale starts this week	48.0***		
Δ log players, last 4 wks vs. previous 4			-4.05*** (0.11)
log players, last 4 wks			-0.43*** (0.07)
N (game-weeks)	13,459,837	11,761,044	2,783,602
games	79,304	79,303	16,916
mean of dep. var.	9.6	5.4	8.4

Publishers discount when the rule allows them to and **when attention is declining** (a halving of concurrent players raises the weekly start probability by ≈ 2.8 pp on a base of 5). Price changes are not exogenous: this matters for everything that follows.

tables/reg_hazard_long; two-way FE by alternating projections (scripts/twfe.py). A 1/8 sample in the previous version gave the same estimates.

From prices to demand: what we can measure

- We do not observe transactions. We observe **dated reviews** (a review requires ownership; 154m of them), the **Steam-purchase flag**, the **playtime at review**, and **concurrent players**.
- **Steam-purchase reviews** = reviews whose author bought the game in the Steam store (`steam_purchase`) and did not receive it for free. The rest activated a key bought elsewhere (Humble/Fanatical bundles, third-party stores) or got a gift/review copy — purchases that do not respond to the Steam price. In the random user sample 79% of fresh reviews are Steam purchases.
- **Reviews per sale**: median 2.5% (Gamalytic) / 3.3% (VG Insights) across 16k games, flat across budget tiers $\Rightarrow$ one review $\approx$ 30–40 sales. The maintained assumption for an elasticity is that this ratio does not depend on the price paid.
- Mean playtime at review *falls* when the price falls (elasticity +0.15, game and month FE): sale-period reviews come from fresh owners.
- Free-to-play titles are excluded everywhere; price tiers below are tiers of *paid* games.
- Two threats: (i) firms discount into declining demand (bias toward zero); (ii) a discount is also *advertising*: Steam e-mails every wishlister and features discounted games.


Monthly review-flow elasticity, broad panel

1.80m game-months, 35,671 games aged ≥ 6 months, 2020-01 – 2026-03. Game and month FE; SE clustered by game. Price = time-weighted posted price.

	coefficient on log price	N
(1) OLS $\log(1 + \text{reviews})$	−0.145 (0.004)	1,797,591
(2) PPML reviews	−0.329 (0.045)	1,791,404
(3) PPML reviews on discount exposure (share of month $\times$ depth)	+1.013 (0.062)	1,791,404
(4) PPML: log list price / discount exposure	+0.115 (0.072) / +1.008 (0.062)	1,791,404
(5) PPML, seasonal share $\times$ log price	−0.331; interaction +0.020 (0.042)	1,791,404
(6) PPML, t / $t-1$ / $t-2$ / lead $t+1$	−0.502 / +0.101 / +0.070 / +0.208	1,683,761
(7) PPML Steam-purchase reviews	−0.358 (0.047)	1,771,139
(8) PPML, months without any Valve-wide sale	−0.292 (0.057)	810,346
(9) PPML, excluding games that ever raised their list price	−0.479 (0.049)	1,462,421

- The response comes from discounts, not from list-price changes — which are endogenous to updates and 1.0 releases and bias the pooled elasticity toward zero (−0.33 vs. −0.48 without the price-raisers).
- A one-week 50% discount (≈ 0.12 exposure) raises the month's reviews by $\approx 12\%$.
- Same elasticity inside and outside seasonal sales; the positive lead mixes consumer anticipation with the firms' response to weak months.

Who is price-sensitive? Young games and the \$10+ tiers



PPML elasticity of monthly reviews to the posted price, game and month FE, 95% CIs. Paid games only; "budget" = list price below \$10, not free-to-play.

[tables/reg_demand_monthly_heterogeneity.](#)

Weekly dynamics around a discount start



Game and week FE, 34,654 publisher-timed and 26,494 seasonal starts with four clean pre-weeks, reference week -1 . Reviews: $+18\%$ in the start week, $+42\%$ the week after, back to baseline by week 5, **no dip in weeks 5–8**. Seasonal starts look small because the week FE absorb the platform-wide surge. **Updates:** we have no patch/DLC feed (SteamDB annotations in the data are bundle events only), so the slightly rising pre-trend may reflect updates timed with discounts; concurrent players move like reviews, and patch notes (Steam news API) are the next data task.

The discount event matters more than its depth

PPML on weekly counts, publisher-timed starts, game and week FE; effect in week 1 relative to week -1.

Depth bin	mean depth	events	reviews, week 1	Steam-purchase reviews	mean CCU	arc elasticity (reviews)
$\leq 33\%$	25%	2,912	+63% (0.09)	+70%	+45%	-1.67
34-50%	47%	7,806	+65% (0.06)	+72%	+30%	-0.79
51-66%	59%	5,473	+62% (0.04)	+69%	+30%	-0.53
67-99%	79%	18,463	+79% (0.03)	+90%	+44%	-0.38

- A 25% cut and a 79% cut produce almost the same proportional jump. Depth is chosen (deeper for older, less popular games), so this is not a causal dose-response — but it is hard to square with a pure price response.
- Consistent with a **notification / visibility channel**: wishlisters are e-mailed when the game goes on sale, discounted games are featured. The discount is partly advertising to a stock of waiting consumers.
- Weeks 4-8 after the start: reviews +2 to +8% above baseline, not below — no visible cannibalisation of later demand on this horizon.

tables/dose_response_weekly_ppml (standard errors in parentheses).

Weekly elasticities and engagement

Game and week FE; SE clustered by game. A: balanced panel, 8,424 games $\times$ 104 weeks (2024–25). B: broad players panel, 16,436 games, 3.6m game-weeks (2020–26).

panel	outcome	OLS $\log(1+y)$	PPML
A	reviews	−0.248 (0.004)	−0.419 (0.018)
A	Steam-purchase reviews	−0.269 (0.004)	−0.453 (0.019)
A	mean concurrent users	−0.215 (0.004)	−0.225 (0.028)
A	PPML reviews, log price at t / $t-1$ / $t-2$ / $t-3$	−0.36 / −0.20 / −0.07 / −0.02	
A	cumulative $t \dots t-4$		−0.68
B	$\log(1 + \text{weekly peak players})$	−0.247 (0.004)	
B	weeks outside Valve-wide sales	−0.275 (0.005)	

- Reviews lag purchases by about a week: the contemporaneous coefficient understates the response; the cumulative review-flow elasticity is about -0.7 .
- Engagement (concurrent players) responds with an elasticity of about -0.25 — a discount brings in players, not only reviewers.

Users: what we see, and how we date a purchase


139,966 sampled reviewers: 30k most prolific + 110k uniform random (of 47.5m reviewer IDs). Profile (country, account age), public library for 69k (30m owned games with playtime), and *all* 6.6m of their reviews with metadata.

Not observed: purchase dates, wishlist additions/removals, prices paid.

The dating trick: a review written with ≤ 2 h of playtime (“fresh”) follows the purchase within days. We match every review to the posted US price *at the second it was written* and to the game’s last discount start: 3.5m reviews, 100k users, 59k games; 892k fresh reviews (15k by random users).


scripts/09_micro_build.py, 10_micro_analysis.py; tables/facts_micro.json.

Reviews with little playtime cluster right after a discount starts (purchase-timing resolution)



15.6% of fresh reviews arrive within two days of a discount start, 36.5% within a week (6.4% / 19.3% for reviews after 50+ hours of play).

Users time purchases to sales — in every segment



Random users, 14,944 fresh purchases:

- 42.5% made while the game is on sale, vs. 21.4% of time on sale for the same games
- average discount at purchase 20.4 pp vs. a time-average of 11.2 pp; 48% conditional on buying on sale
- the same for prolific users (41.6% vs. 22.0%)
- AAA titles least (35%), indie most (47%)

Consumers extract roughly twice the average posted discount by timing — the aggregate review spike is made of these purchases.

Inside a sale: day one. Before a sale: no anticipation

Timing inside a spell (on-sale fresh purchases, spells ≤ 21 days):

	observed	uniform over spell
first day	20.7%	10.4%
first two days	34.3%	20.7%
last two days	16.4%	20.7%

A first-day spike and no end-of-sale rush: the user-level counterpart of the aggregate notification effect.

Do full-price buyers see the next sale coming?

Share of full-price fresh purchases made shortly before the game's next discount start (483k purchases):

	observed	random timing
within 7 days	14.5%	15.9%
within 14 days	27.8%	30.3%
within 30 days	54.3%	55.8%

Full-price buyers barely avoid the days before a sale, although the cooldown clock makes the next start highly predictable. The buyers who could wait are not in this group — heterogeneous patience/attention.

Random-timing benchmark: share of off-sale time of the same games lying within 7/14/30 days of a start. tables/facts_micro.json, facts_micro_benchmarks.json.


Who buys on sale? Persistent types, and the early adopters pay full price

LPM: $100 \times [\text{game on sale at the moment of the fresh review}]$, game and month FE, SE clustered by game, 613,537 fresh reviews, 24,725 users.

	pp
log games owned (library size at review)	-1.64*** (0.08)
log(1 + account age, years)	-1.56*** (0.14)
log(1 + reviews written)	+0.53*** (0.08)
Russia / Turkey / UK (vs. US)	+5.6 / +5.7 / +5.0***
Brazil / China (vs. US)	+2.8 / +2.8***
France / Poland (vs. US)	-1.5** / -1.0**

Purchase lag since release (game FE): -0.31 log points per log unit of games owned, +0.28 per log unit of account age.

Large-library users buy early and at full price; sale buying is the budget-constrained behaviour (lower-income markets, heavy reviewers).



Users with ≥ 10 fresh purchases: on-sale share p10 / median / p90 = 19 / 40 / 64%; split-half correlation of a user's share 0.39; user FE explain 15% of the variance, game FE 23%, both 30%.

A demand-like estimate from the users: the purchase-timing hazard

For each game we know *when* every sampled eventual buyer bought $\Rightarrow$ a discrete-time proportional hazard on game $\times$ week cells, $\mathbb{E}[\text{buys}_{gt}] = R_{gt} \exp(\beta \log p_{gt} + \text{age bins} + \alpha_g + \delta_t)$, R_{gt} = eventual buyers not yet bought (offset); PPML, SE by game. 2.08m game-weeks, 9,358 games, 587k purchases. Raw weekly hazard: 1.31% in sale weeks, 0.64% otherwise.

	estimate
hazard elasticity to the posted price	-0.55 (0.01)
weeks outside Valve-wide sales	-0.63 (0.02)
hazard ratio, week on sale vs. not	2.15
by depth 1-33 / 34-50 / 51-66 / 67-99%	1.90 / 2.38 / 2.47 / 2.62
timing elasticity (hazard ratio) by user group	
random stratum	-0.82 (2.24)
prolific stratum	-0.55 (2.15)
library 100-499 games	-0.60 (2.28)
library 500+ games	-0.49 (2.00)
high-income countries	-0.52 (1.98)
middle-income countries	-0.59 (2.30)

What it is: the pure **re-timing margin** — how much eventual buyers move their purchase into sale weeks — estimated without observing non-buyers. A sale week **doubles** the weekly purchase probability. The aggregate review-flow elasticities (-0.33 to -0.48) mix this with new demand: most of the short-run response is re-timing by people who buy anyway (median release-to-purchase lag 4-5 months).

Do sale buyers value the game less? No. Within game and month the discount at purchase is unrelated to playtime a year later (+0.002 log points per 10 pp; -0.005 with user FE) and slightly *positively* related to playing ≥ 10 h and to recommending. Discounts exploit heterogeneity in **patience, attention and budgets**, not in valuations.

What the facts say

1. The price of a zero-marginal-cost durable good is a **sawtooth with a falling floor**: list prices are sticky, discounts are frequent, short, deep, discrete, and ratchet down with age.
2. Timing is governed by **platform rules and events** (cooldown, seasonal sales) and by **declining attention**: firms discount to reach consumers who are no longer looking.
3. Demand responds strongly but **almost independently of depth**; no post-discount dip within two months.
4. At the user level: **42% of purchases in 21% of time**, a first-day spike, persistent sale-buyer types, early adopters who pay full price, full-price buyers who do not anticipate predictable sales; a sale week **doubles the purchase hazard** of eventual buyers (timing elasticity -0.55), and sale buyers play and recommend as much as full-price buyers.

A model has to deliver: (a) frequent temporary cuts rather than permanent ones, (b) a depth ratchet, (c) a notification channel, (d) heterogeneous patience/attention types rather than valuation sorting, (e) a role for the platform calendar.

Next step: a structural model — primitives

Good. Game g with list price P_g , zero marginal cost; posted price $p_{gt} = (1 - d_{gt})P_g$, $d_{gt} \in \{0, 0.1, \dots, 0.9\}$, weekly.

Consumers. Potential buyers arrive over time (release spike, then a flow driven by attention). Consumer i has valuation $v_i \sim F_g$, a persistent **patience/attention type** θ_i (the user-level trait), interest that decays at rate λ , and sees the price only when she visits the page or is *notified*. Wishlisting = choosing to be notified. Each period a non-owner buys, waits, or exits. Waiting pays because the price process is predictable: a Markov chain over (time since last discount, last depth, calendar) — exactly what Sections 2–3 estimate.

Publisher. Chooses d_{gt} to maximise the present value of revenue, subject to the cooldown (no new spell within 30 days of the previous end unless a Valve event is running), $d \leq 0.9$, and a cost of changing the depth (the persistence we see). State: stock of waiting consumers W_{gt} , time since last discount, last depth, age, calendar.

Platform. Seasonal sales are anticipated common shocks to arrivals/attention that waive the cooldown; the platform chooses the calendar and the rules.

Structural model — demand side and identification

Purchase flow: $Q_{gt} = W_{gt} \cdot h(p_{gt}, \mathbb{E}_t[p_{g,t+1}, \dots], \text{notified}_{gt}; \theta)$.

- **Re-timing margin** directly from the user-level purchase hazard (-0.55 ; hazard ratio 2.15 in sale weeks, rising slowly with depth); **price sensitivity** h_p from depth variation *within* the notified set, from spell duration, and from seasonal sales where the platform, not the publisher, creates the traffic.
- **Notification effect**: the common jump at any start (flat dose-response) and the **first-day spike** in user-level purchases.
- **Types** θ : the persistent user-level on-sale share (p10–p90: 19–64%), the early adopters with large libraries, and the full-price buyers' lack of anticipation (an inattentive/impatient type) pin down the mixture of patient-attentive and impatient-inattentive consumers.
- **Stock of waiting consumers** W_{gt} : needs wishlist/follower *histories* (SteamDB) — the first data task.
- **Reviews** $\rightarrow$ **purchases**: the calibrated 2.5–3.3% ratio (flat across tiers) plus the library cross-section.
- Estimation by indirect inference on the event-study, dose-response, no-dip and user-level timing moments, separately for young and mature games.

Structural model — supply, platform, counterfactuals

Supply. A dynamic discrete-choice problem (Rust): choice-specific values from the estimated start-hazard and depth policies (CCP approach). The ratchet is the equilibrium skimming path: as high-valuation waiting consumers are exhausted, the optimal depth rises. The flat dose-response plus the cooldown explain *short, frequent, deep* discounts rather than permanent cuts: each discount is a periodic notification to an accumulating stock, and the platform caps the frequency.

Objects recovered: the depth-adjustment cost, the publisher's discount factor, the distribution of valuations and of patience/attention types, the decay of interest, the value of a notification.

Counterfactuals.

- no cooldown rule; no seasonal sales; a longer or shorter sale calendar
- a ban on discounts (the Coasian durable-good seller with commitment)
- notification without a price cut (pure advertising) vs. price cut without notification
- entry of similar games (release pressure) and the discount response

Welfare split between sale buyers and early adopters, small and large publishers, and the platform.

What we do next

- **Data:** wishlist/follower histories (SteamDB); Steam news / patch notes per game (ISteamNews) to separate updates and DLC from discounts; fest and Daily-Deal participation; regional prices (FR, GB parsed); the broad weekly review panel for all games (script ready); repeated library snapshots to observe acquisitions directly.
- **Identification clean-ups:** stacked event design with matched clean controls (as in our earlier matched analysis: +31.7 reviews over weeks 0–3, 95% CI 23.5–39.9); the cooldown expiry as a timing instrument; user-level event studies around notifications.
- **Model:** estimate the price process by publisher type; the consumer hazard with notification and types; the publisher's dynamic problem; counterfactuals on the calendar and the cooldown rule.

Thank you.

Backup: definitions and samples


- **Weeks** run Thursday–Wednesday UTC (a seasonal sale starting on a Thursday falls entirely in one week).
- **Known start**: spells already running at the first observation (synthetic 2020-01-01 state for 49,099 apps) are excluded from timing and depth analyses. 100% cuts (giveaways) are excluded from depth analyses.
- **Main sample**: Steam store type “game”, not free-to-play, listed in March 2026, not NSFW (ITAD mature flag or Hentai / NSFW / Sexual Content tags or Steam Sexual Content / Nudity descriptors: 9,639 games), not currently in early access (ITAD flag or Steam genre: 14,027 games). 116,808 paid games remain. Depth regressions also exclude 13,730 games that raised their US list price by $\geq 10\%$ at any time in 2020–2026.
- **Budget tiers** by maximum list price: budget $< \$10$, indie $\$10\text{--}29.99$, AA $\$30\text{--}49.99$, AAA $\geq \$49.99$.
- **Seasonal windows** (start, end): Summer 2020 06-25/07-10, Autumn 2020 11-25/12-02, Winter 2020 12-22/01-06, Summer 2021 06-24/07-09, Autumn 2021 11-24/12-02, Winter 2021 12-22/01-06, Summer 2022 06-23/07-08, Autumn 2022 11-22/11-30, Winter 2022 12-22/01-06, Spring 2023 03-16/03-24, Summer 2023 06-29/07-14, Autumn 2023 11-21/11-29, Winter 2023 12-21/01-04, Spring 2024 03-14/03-21, Summer 2024 06-27/07-11, Autumn 2024 11-27/12-04, Winter 2024 12-19/01-02, Spring 2025 03-13/03-20, Summer 2025 06-26/07-10, Autumn 2025 09-29/10-06, Winter 2025 12-18/01-05, Spring 2026 03-19/03-26. Lunar New Year 2020–22 and Halloween 2020–21 are “other Valve-wide” events.
- **Monthly demand sample**: main-sample games with ≥ 20 reviews in total, age ≥ 6 months, full price coverage in the month. **Weekly**: the balanced panel of the earlier project (complete price, review and mean-CCU coverage in 2024–2025), main sample.
- **Event sample**: a start in week 0 with no discount in weeks $-4..-1$; publisher-timed = no Valve-wide sale in week 0.
- **User sample**: reviews by the sampled users written 2020-01-02 – 2026-05-27 on main-sample games with a price state at the time of the review; *fresh* = playtime at review ≤ 120 minutes; library size = the author’s number of owned games at the time of the review (hidden if the profile is private).

Backup: reviews per sale, and depth distribution / duration


Median ratio of total reviews to third-party sales estimates (main-sample games with >1,000 estimated sales and >50 reviews):

tier	games	Gamalytic	VG Insights
budget < \$10	5,173	2.45%	3.13%
indie \$10–29.99	8,755	2.54%	3.32%
AA \$30–49.99	1,631	2.59%	3.34%
AAA \$50+	663	2.86%	3.34%
all	16,245	2.53%	3.32%

IQR of the Gamalytic ratio: 1.8–3.6%.



Backup: dose-response, $\log(1+y)$ version



Mean effect in weeks 0-1 relative to week -1, by depth bin of publisher-timed starts. `tables/dose_response_weekly`.