The score calculator, feed simulation and follower goal planner are interactive and are not included in this PDF. Use them at www.howxranks.com.

A reply is worth 5. A report, –234.

Those are X's own weights, from the open-source code behind the **For You** feed. Run your own post through them and see the score it would get.

Created by **@SidneyCur** · Based on the repository at commit b220448 (2026-11-03)

SPECIMEN — ONE VIEWER, ONE POST — FIG. 0

EXAMPLE

+2.00

A reply on your original post, from a mutual follow

The reply weight is 5, plus a 15 boost for mutual follows on original posts: 20 in total. A 10% predicted chance of a reply adds 20 × 0.10.

`20 × 0.10 = +2.00`

REV What's new

X updated the open-source ranking code. Five scoring weights moved and four new signals appeared. Here is each change, what it actually means, and what it does to a post's score.

~~b412112 · 3 Oct 2026~~ → b220448 · 3 Nov 2026 20 commits · 5 weights changed

At a glance

The five weights that moved. Every number below is read from the repository; the chances in the examples are illustrative.

REVISION LEDGER — FIVE WEIGHTS CHANGED — COMMIT B412112 → B220448

#	ACTION	WAS	NOW	CHANGE	
01	**Favorite** FavoriteWeight	~~0.5~~	0.66	+32%	Explain ↓
02	**Not interested** NotInterestedWeight	~~-47.52~~	−52.272	10% harsher	Explain ↓
03	**Video open** VideoOpenWeight	~~0.07~~	0.02	−71%	Explain ↓
04	**Dwell time** ContDwellTimeWeight	~~0.004~~	0.006	+50%	Explain ↓
05	**Unexplored post** PostUnexploredWeight	~~0.02~~	0.015	−25%	Explain ↓

What each change means

Each weight multiplies the *predicted chance* of an action for one viewer, not a count of actions. A weight of 0.66 on favorites means a 10% predicted chance of a like adds 0.066 points — not 0.66, and not one point per like.

01

A like counts for a third more

FavoriteWeight

~~0.5~~ → 0.66 +32%

WHAT IT IS

Every time the ranker scores your post for one viewer, it predicts how likely that person is to like it. That chance is multiplied by the favorite weight and added to the score.

WHAT IT CHANGES FOR YOU

Likes were the cheapest positive signal on the board and they still are — but the gap narrowed. At the chances this site uses, the copy-link term was about 20 times the favorite term; it is now about 15 times. Posts that earn a lot of likes and little else gained a little ground, though shares and replies still dominate.

EXAMPLE

WAS	0.5 × 0.10	+0.050
NOW	0.66 × 0.10	+0.066

A post a viewer is 10% likely to like adds 0.016 more than it did. Small per viewer, repeated across every viewer the post is scored for.

02

“Not interested” costs 10% more

NotInterestedWeight

~~−47.52~~ → −52.272 10% harsher

WHAT IT IS

When a viewer taps “Not interested” on a post, the ranker treats it as a strong negative. This is the penalty most ordinary posts actually risk — reports, blocks and mutes are far rarer.

WHAT IT CHANGES FOR YOU

The absolute cost went up: the same 1% chance now removes 0.048 more from the score. Relative to a like it actually eased slightly — the penalty term is about 8 times the favorite term, down from about 9.5 — because favorites rose faster. The practical read is unchanged: posting off-topic for your audience, or hammering one theme until people tune out, is the most expensive ordinary mistake on this list.

EXAMPLE

WAS	−47.52 × 0.01	−0.475
NOW	−52.272 × 0.01	−0.523

At a 1% chance of the viewer tapping “Not interested”, the post loses 0.048 more than before — which is roughly three-quarters of what a 10% like chance adds back.

03

Opening a video barely counts now

VideoOpenWeight

~~0.07~~ → 0.02 −71%

WHAT IT IS

The chance a viewer taps to open your video. It is a separate signal from watching it — video quality views have their own weight.

WHAT IT CHANGES FOR YOU

This is the change with the sharpest edge. Video opens dropped to almost nothing, video quality views are weighted `0.0`, and the four new watch-time signals are also `0.0`. In this code, essentially nothing about *playing* a video raises its score. What still pays on a video post is what the video makes people do: reply, quote, share the link, follow you — and stay on the post, which is the one video-adjacent thing that went up.

EXAMPLE

WAS	0.07 × 0.10	+0.007
NOW	0.02 × 0.10	+0.002

A video a viewer is 10% likely to open contributes 0.002. One copy-link share at a 5% chance contributes 1.00 — five hundred times more.

04

Time spent reading matters more

ContDwellTimeWeight

~~0.004~~ → 0.006 +50%

WHAT IT IS

Unlike the rest of the table, this weight does not multiply a probability. It multiplies the *predicted number of seconds* a viewer will spend on the post, so it grows with how long you hold someone.

WHAT IT CHANGES FOR YOU

Two upstream changes point the same way. This weight rose by half, and on 8 October X shipped a ranking model that predicts time spent much better — so the signal is both worth more and measured more accurately. Posts built to be read to the end gain: a thread opener that earns the tap, a post dense enough to re-read, a chart worth studying. One-line posts that are consumed in half a second gain nothing here.

EXAMPLE

WAS	0.004 × 10 units	+0.040
NOW	0.006 × 10 units	+0.060

For a post predicted to hold a viewer for 10 model units of time. The repository does not publish what one unit equals in seconds, so read the direction, not the absolute figure.

05

Smaller nudge for followers new to your topic

`PostUnexploredWeight`

~~0.02~~ → 0.015 −25%

WHAT IT IS

A small bonus applied when the post is something the viewer has not engaged with much before. It only applies to people who already follow you — though a new switch, `PostUnexploredIncludeOutOfNetwork`, can extend it to strangers. It is off by default.

WHAT IT CHANGES FOR YOU

Honestly: almost nothing. The bonus was tiny and is now a quarter smaller. It is also not a lever you can pull — it depends on what the viewer has seen, not on what you write. Worth knowing it exists so you do not mistake it for a reward for posting about new topics.

EXAMPLE

WAS	0.02 × 0.30	+0.0060
NOW	0.015 × 0.30	+0.0045

A difference of 0.0015 — about a fortieth of what a single like contributes.

06

A shared list of popular posts now feeds For You

`EnablePopularPostsSource`

~~false~~ → true

WHAT IT IS

A new source of candidate posts, switched on by default. Alongside the personalised retrieval that finds posts for you specifically, For You now also draws from one shared list of popular recent posts — seeded, in the repository, from the most-followed active accounts.

WHAT IT CHANGES FOR YOU

Read this one carefully before celebrating. It is a further route into feeds of people who do not follow you — but because the list starts from the most-followed accounts, for most accounts it means *more competition* in the feed rather than a new door. The routes you can actually influence are unchanged: original posts reach strangers through retrieval; replies and reposts are still filtered out for people who do not follow you.

WHAT IT LOOKS LIKE

> A stranger's For You feed can now contain your post by three routes: Phoenix retrieval matched you to their interests, SimClusters found you through who they engage with, or you landed on the shared popular list. The third is the one you are least likely to reach.

07

Four new signals, all switched off

```
VideoContinuationWeight
UserVideoContinuationWeight
ProfileVisitSecsWeight
PostVideoWeight
```

`new → 0.0`

WHAT IT IS

Four measurements are now wired into the scoring function: how many seconds of your video a viewer watches, how long that viewer watches video generally, how many seconds they spend on your profile after the post, and whether the post carries video. All four are multiplied by a weight of `0.0`.

WHAT IT CHANGES FOR YOU

Today, nothing — a weight of zero contributes zero. What matters is that the plumbing now exists. These weights are feature switches, so any of them can be turned on in production without a code change, and the public repository would not necessarily show it. Three of the four concern watch time and profile attention, which is a fair signal of where ranking is heading.

WHAT IT LOOKS LIKE

NOW	0.0 × anything	0.000

A 40-second video watched to the end still adds nothing through any watch-time term. Treat this as a thing to watch, not a thing to act on.

08

Rules can now attach a notice

visibility-filtering/rules/engine.rs

3 → 4 outcomes

WHAT IT IS

When a visibility rule matched a post, it could do one of three things: show it, filter it from a feed, or drop it. There is now a fourth — attach a notice to it.

WHAT IT CHANGES FOR YOU

A middle path exists where there used to be a cliff. A post that trips a rule can stay in the feed, undownranked, carrying a notice, instead of quietly vanishing. The catch for anyone auditing this: nothing in the published repository draws the notice or says what it reads, so the visible effect cannot be traced from this code.

WHAT IT LOOKS LIKE

show — nothing happens · **filter** — gone from this feed · **drop** — gone everywhere · **notice** — stays, with a label attached

---

09

Smaller changes worth knowing

constants.py
Brazil2026ElectionFilter
recsys_two_tower_model.py

REPLY-SPAM EXEMPTION GOT STRICTER

Accounts above a PageRank score skip the risky-reply label. The bar rose from **62 to 66**, so slightly fewer accounts are exempt. The grey-badge exemption is unchanged.

BRAZIL ELECTION FILTER REMOVED

With the 2026 Brazilian election over, `Brazil2026ElectionFilter` is gone from For You. Posts from the accounts it covered are no longer filtered because of it.

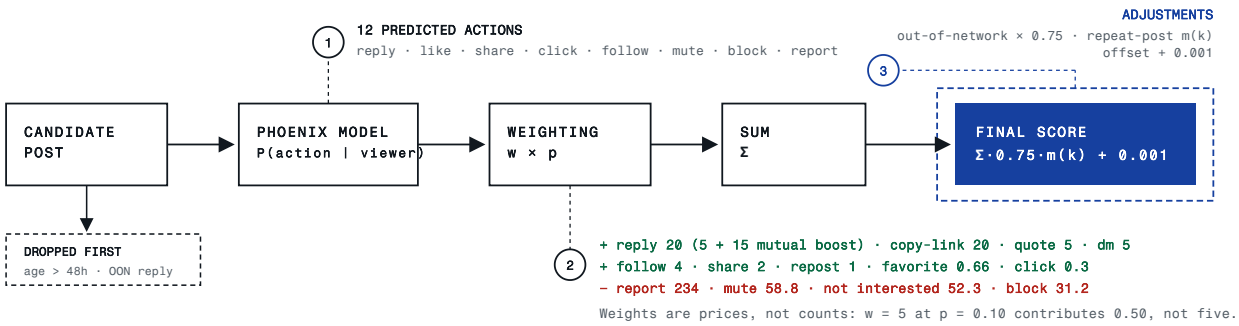
RETRIEVAL FAVOURS POSTS PEOPLE FINISH

From 13 October, retrieval is also trained to surface posts that people read all the way through — the same direction as the dwell-time weight rising. Posts that hold attention to the last line are favoured twice over.

A How a score is built

The ranker does not count likes. It predicts how likely each action is for each viewer, weights each prediction, and adds the results. The four steps below are the same ones the calculator runs.

FIG. 1 — SCORING PIPELINE, ONE POST, ONE VIEWER — SCALE: NOT TO SCALE



01 Predict

The Phoenix model estimates the chance a viewer will reply, like, share, follow, mute, or report a post.

02 Weight

Each chance is multiplied by a fixed weight. Copy-link shares carry 20, a report carries −234.

03 Sum

The weighted values are added. Positive actions push the score up and negative ones pull it down.

04 Adjust

An offset keeps scores non-negative. Out-of-network posts are scaled by 0.75, and each extra post from one author is scaled down.

ⓘ Weights are not counts. A weight of 5 on replies means a 10% chance of a reply adds 0.5 points, not five.

D Worked examples

Each row is one calculation from the weights above, with illustrative chances. Read the result first, then the calculation and what it means for your posts.

DETAIL CALLOUTS — EIGHT WORKED CASES				ILLUSTRATIVE CHANCES
#	CASE	RESULT	CALCULATION	WHAT IT MEANS
01	**Boosted reply on your original post** Largest gain	+2.00	`20 × 0.10 = +2.00`	A mutual follow is 10% likely to reply to your original post. Only original posts from mutual follows get the boost.
02	**Copy-link share** Large gain	+1.00	`20 × 0.05 = +1.00`	Five in a hundred viewers copy the link to your post. Copy-link is the largest non-reply gain, so posts people want to pass on are worth designing for.
03	**One report, at a 1% chance** Largest loss	−2.34	`−234 × 0.01 = −2.34`	A report weighs −234. A 1% chance of reporting cancels several good actions, which is why report rates matter most.
04	**A 2% mute chance against a 30% favorite chance** Mute vs favorite	−1.18	`−58.8 × 0.02 = −1.18 · 0.66 × 0.30 = +0.20`	A mute weighs −58.8, so a 2% chance costs about six favorites' worth of gain (+0.20 at 30%).
05	**The same boosted reply, for a viewer who doesn't follow you** Out of network	1.50	`2.00 × 0.75 = 1.50`	An out-of-network post keeps 75% of its score. Replies and reposts from accounts a viewer doesn't follow are filtered by default.
06	**Your fourth post in one viewer's feed** Fourth post	0.69	`2.00 × 0.344 = 0.69`	Each later post from you is scaled down: the 4th keeps 34% of the starting score. Posting less often reduces how many of your posts compete in one feed.

#	CASE	RESULT	CALCULATION	WHAT IT MEANS
07	**The tenth post and beyond** Floor	0.25	`0.25 + 0.75 × 0.5⁹ = 0.2515`	The multiplier approaches 0.25 but never reaches zero, so a late post still scores a quarter of the starting value.
08	**Report chance that cancels one boosted reply** Break-even	0.85%	`2.00 ÷ 234 ≈ 0.85%`	At a 10% reply chance, a boosted reply is worth +2.00. About one viewer in 117 reporting cancels it out.

Do and don't

Each item links to a weight or a rule in the repository. Points are the change to a post's score. Labels are enforcement flags that last 30 days.

Raises the score

Post original posts to mutual follows

Replies to them carry 20 instead of 5. `reply + boost`

Earn copy-link shares

Worth 20 per unit of chance, the largest non-reply gain. `ShareViaCopyLinkWeight`

Earn quotes and DM shares

Each worth 5. `QuoteWeight`, `ShareViaDmWeight`

Earn follows

Worth 4 per unit of chance. Clear posts that show what you know help. `FollowAuthorWeight`

Reply to followers and mutuals

Their replies to you stay in the feed. Replies to strangers are filtered by default. `OONRetweetReplyFilter`

Write posts worth putting on a popular list

A shared list of popular recent posts now feeds For You by default, a new route to people who do not follow you. `EnablePopularPostsSource`

Post regularly, within 48 hours

Posts older than 48 hours are removed before scoring. `MAX_POST_AGE`

⚠ Lowers the score or gets labelled

Trigger reports, mutes, blocks

Reports –234, mutes –58.8, blocks –31.2 per unit of chance. `ReportWeight`

Automate or burst replies

Flagged as fast-reply spam: a 30-day label that drops the post for people who don't follow you.

`fast_reply_spam_post`

Post LLM-written or gibberish replies

Flagged as LLM text or gibberish. Gibberish gets a 30-day spam label. `llm_slop_post`, `gibberish_post`

Repeat near-identical replies

The multi-step classifier flags them across threads. `classifier_multi_step_reply_spam`

Count on video views, watch seconds or profile clicks

All carry zero weight in this code, including the new watch-seconds and profile-visit terms. `VqvWeight`, `ProfileClickWeight`, `VideoContinuationWeight`, `ProfileVisitSecsWeight`

Post a fourth or later item in one feed

Scaled to about a third of its score. `EnableAuthorDiversity`

Questions

What is howxranks?

howxranks is an independent reference for X's For You feed. It explains the public ranking code, turns its scoring weights into a calculator and simulation, and lists the enforcement rules, with sources.

Is howxranks affiliated with X?

No. howxranks is independent and is not affiliated with, endorsed by, or sponsored by X Corp.

Where do the numbers come from?

Every weight comes from the public repository x-algorithm, commit b220448 (2026-11-03). The repository's parameters are mirrored from feature switches, so production values may differ. The probabilities in the calculator are illustrative.

How is a post's score calculated?

The ranker predicts how likely each action is, multiplies each prediction by its weight, and adds the results. An offset keeps scores non-negative. Out-of-network posts are scaled by 0.75, and each extra post from one author is scaled down.

Why do replies to strangers reach fewer people?

In the repository, out-of-network replies and reposts are filtered by default, so replies mostly reach your followers. Original posts are the route to new people.

Sources and limits

Weights and fusion: `home-mixer/params/param.rs`, `vm-ranker/params.rs`, `xai-value-model/scoring.rs`.

Filters: `home-mixer/filters/`. Enforcement: `abuse-enforcement-service/service-lib/rules/`, `visibility-filtering/rules/`, `grox/flows/reply_spam/`.

The repository's parameters are mirrored from feature switches, so production values may differ. The probabilities are illustrative, not measured for any account. Classifier thresholds are not published.

Design direction: an engineering drawing — graph-paper sheet, titleblock, numbered callouts and figure frames. Dark mode is a cyanotype. No assets were copied from anywhere.

REVISION LOG

REV	DATE	CHANGE	SOURCE
C	2026-11-03	Popular-posts source on by default; weights re-read	commit b220448
B	2026-10-06	Repeat-post floor confirmed at 0.25	home-mixer
A	2026-09-12	First plate	x-algorithm

Created by **@SidneyCur**. Built from the x-algorithm-source repository, commit b220448 (2026-11-03). Not affiliated with X.