




RESEARCH · PLATFORM INTEGRITY

Bot-Like Amplification on X: Findings from CAM's Study of the @MarioNawfal Amplifier Network

October 13, 2026 · Combat Antisemitism Movement (CAM)

40.5%

of the accounts amplifying Mario Nawfal's most-shared posts show strong signals of being bots, fake profiles, or inauthentic accounts

Between August 29 and October 12, 2026, CAM's Antisemitism Research Center analyzed 10,015 unique X accounts that amplified 230 of the most-shared posts by Mario Nawfal (@MarioNawfal), an Australian crypto entrepreneur who runs one of the largest news-aggregation accounts on X, with roughly 3.95 million followers. After auditing, CAM's detection model identified 4,054 of these accounts (40.5%) as showing strong signals of being bots, fake profiles, or inauthentic accounts.

In practical terms, roughly two in every five accounts promoting Mario Nawfal's content behave more like automated or fake profiles than like ordinary users. Reposts are one of the signals X uses to decide what to show more widely, so amplification at this scale can make content appear more popular, and travel further, than human interest alone would carry it. For readers, repost counts on these posts are therefore a weaker measure of genuine public engagement than they appear.


Nawfal's team does not vet posts for accuracy, and since 2025 he has given long-form, largely unchallenged platforms to the Russian foreign minister and the president of Belarus, among other heads of state. When a channel with that editorial record draws a large share of its amplification from bot-like accounts, the effect is a ready-made distribution system that can push unverified claims and state-aligned narratives, including against Israel, to millions of users while appearing to reflect genuine public interest.

Cross-account tests found a dense core of repeat amplifiers but no evidence of a centrally operated network: no copy-paste text across accounts, no sustained sub-minute synchronized reposting, and no clusters of accounts created in batches. The pattern most consistent with the data is a loyal, heavy-reposting audience in which a large share of accounts behave like automated ones.

HOW THE BOT SCORE WORKS

The score comes from CAM's own model (X model v1.4.0), which detects bots, fake profiles, and inauthentic accounts. It is built on CAM's framework of key signs that an account may be a bot, with each signal's weight grounded in published research from DFRLab, Indiana University's Observatory on Social Media, and the Cresci studies. Each account is measured against 32 signals: posting rate and bursts, odd-hour activity, duplicate text, account age, default avatars and digit-suffixed handles, follower patterns, and generic or repetitive replies. No single signal counts on its own; only consistent combinations raise an account's 0 to 100 score.

The model is built to under-count. Post content and ideology are not inputs, so an account is never scored as a bot for what it says. Any one measurable human behavior, such as original varied text, real conversations, or a daily sleep gap, pulls a borderline account out of the bot-like group. The 40.5% figure is what remains after those protections, which makes it a conservative measure rather than an inflated one.



Who is Mario Nawfal


Mario Nawfal is a Lebanese-Australian entrepreneur, business influencer, podcast host, and political commentator. He is the co-founder of NFT Technologies and the CEO and founder of IBC Group. He runs one of the largest news-aggregation accounts on X, with roughly 3.95 million followers as of October 13, 2026. His team repackages other outlets' reporting around the clock without linking to it and, by Nawfal's own account, without vetting it for accuracy. Elon Musk interacts with the account almost daily, regularly pushing its posts to millions of views (NPR, April 17, 2025).

His content leans pro-Russian and hostile to Israel. He gives Kremlin-aligned leaders an unchallenged platform: in his interviews, the Belarusian president and the Russian foreign minister repeated Kremlin talking points without pushback. He has also argued that Western media treat Russian and Chinese leaders unfairly (NPR, April 17, 2025). Conspiracy Watch describes him as spreading false and hateful narratives about Ukraine, Jews and Israel (Conspiracy Watch, March 15, 2025).

What we studied

CAM studied amplifiers: accounts that repost, quote, or reply to a post, pushing it into their own followers' feeds. CAM selected 230 of Nawfal's most-reposted original posts published between August 29 and October 12, 2026, and collected the public profiles of their most recent reposters, using only public, read-only data from the official X API. Counting each account once left 10,015 unique accounts.

The study examines this amplifier population, not Nawfal himself, and makes no claim that he or his team operates or purchased any of these accounts.



CONTEXT: THE YANDEX STRIKES

Between October 8 and 11, 2026, Ukrainian drones struck three Yandex data centers in Russia, and Cloudflare Radar showed automated traffic from Yandex's network falling roughly 94% (Kyiv Post, October 12, 2026). Israeli observers reported that accounts posting aggressive political comments in Russian and Hebrew fell silent, ahead of Israel's October 27 election and against a documented record of Russian bot operations targeting Jewish and Israeli audiences.


Key findings

After the final audit on October 13, 2026, the 10,015 accounts were classified as follows.

CLASSIFICATION OF 10,015 AMPLIFIER ACCOUNTS		
CLASSIFICATION	ACCOUNTS	SHARE
Strong bot-like signals	4,054	40.5%
Moderate bot-like signals (profile-only)	244	2.4%
Some inauthentic signals, without bot-like behavior or reposting patterns	3,460	34.5%
Consistent with human activity and X human profiles	2,240	22.4%
Insufficient data	17	0.2%

1 Bot-like signals rise with loyalty.

Accounts that amplified three or more of the sampled posts showed strong bot-like signals at 54.6%, against 37.6% for one-post amplifiers. The most engaged slice of the amplifier base is also the most bot-like.




2 The model is deliberately conservative.

Any one measurable human behavior, such as original varied text, real conversations, or a daily sleep gap, moves a borderline account out of the bot-like group. The 40.5% headline figure already reflects this benefit of the doubt.



What these accounts look like

Six accounts from the amplifier set, shown as they appear on X. Each combines several of the profile signals the model weighs: a default or placeholder avatar, a digit-suffixed handle, a recent creation date, and a posting volume far beyond typical human use.



CAM Bots Detection Model

The model (X model v1.4.0) detects bots, fake profiles, and inauthentic accounts. It operationalizes CAM's framework, "Key signs that a social media account may be a bot," whose governing principle is that no single sign proves automation; only combinations count. It measures 32 signals across the framework's four categories: account behavior (posting rate, bursts, duplicate text, odd hours), profile (account age, default avatars, digit-suffixed handles, follower patterns), coordination and manipulation (shared hashtags and links, template handles), and language and content (generic replies, repetitive phrasing). Each signal's weight is grounded in published bot-detection research, including DFRLab, Indiana University's Observatory on Social Media, and the Cresci spam-bot studies.

Signals combine into a 0 to 100 score through a logistic model, and every score comes with plain-language reasons, the model version, and the data it was based on. Three design choices keep the model honest:

Content-neutral.

Ideology and post content, including antisemitic content, are not scoring inputs. Hate and automation are different questions and are assessed separately.

Human-indicator protection.

Seven measurable human behaviors (real identity, balanced network, account age with steady posting, organizational verification, original varied text, genuine conversation, a daily sleep gap) each override a borderline bot score.

Three categories kept apart.

Automation (software-driven activity), inauthenticity (a false identity, possibly human-run), and declared bots (self-labeled automation that X permits) are distinguished; declared bots receive a strong score reduction and are not treated as malicious.

X's rules and their limits

X's own rules already prohibit what this study measures. Its authenticity policy states that the platform does not allow "any activity that attempts to manipulate our platform or disrupt our services through inauthentic accounts, behaviors or content." Its platform manipulation policy names artificial amplification, making an account or idea seem more popular than it is through inauthentic reposts and engagement, as a violation.

Yet 40.5% of the accounts amplifying one of X's most-boosted news accounts show strong signals of being bots, fake profiles, or inauthentic accounts. The rules exist; enforcement is not keeping pace.

CAM calls on X to apply its own platform manipulation policy to the amplifier networks around high-reach accounts, and to give independent researchers the data access needed to verify findings like these.