



Queensland University of Technology
Brisbane Australia

This may be the author's version of a work that was submitted/accepted for publication in the following source:

[Graham, Timothy](#) & Andrejevic, Mark
(2024)

A computational analysis of potential algorithmic bias on platform X during the 2024 US election.

[Working Paper]
(Unpublished)

This file was downloaded from: <https://eprints.qut.edu.au/253211/>

© Consult author(s) regarding copyright matters

This work is covered by copyright. Unless the document is being made available under a Creative Commons Licence, you must assume that re-use is limited to personal use and that permission from the copyright owner must be obtained for all other uses. If the document is available under a Creative Commons License (or other specified license) then refer to the Licence for details of permitted re-use. It is a condition of access that users recognise and abide by the legal requirements associated with these rights. If you believe that this work infringes copyright please provide details by email to qut.copyright@qut.edu.au

License: Creative Commons: Attribution-Noncommercial-No Derivative Works 4.0

Notice: *Please note that this document may not be the Version of Record (i.e. published version) of the work. Author manuscript versions (as Submitted for peer review or as Accepted for publication after peer review) can be identified by an absence of publisher branding and/or typeset appearance. If there is any doubt, please refer to the published source.*

A computational analysis of potential algorithmic bias on platform X during the 2024 US election

Timothy Graham (Queensland University of Technology)

Timothy.Graham@qut.edu.au

Mark Andrejevic (Monash University)

Mark.Andrejevic@monash.edu

Summary

This technical report presents findings from a two-phase analysis investigating potential algorithmic bias in engagement metrics on X (formerly Twitter) by examining Elon Musk's account against a group of prominent users and subsequently comparing Republican-leaning versus Democrat-leaning accounts. The analysis reveals a structural engagement shift around mid-July 2024, suggesting platform-level changes that influenced engagement metrics for all accounts under examination. The date at which the structural break (spike) in engagement occurs coincides with Elon Musk's formal endorsement of Donald Trump on 13th July 2024.

In Phase One, focused on Elon Musk's account, the analysis identified a marked differential uplift across all engagement metrics (view counts, retweet counts, and favourite counts) following the detected change point. Musk's account not only started with a higher baseline compared to the other accounts in the analysis but also received a significant *additional* boost post-change, indicating a potential algorithmic adjustment that preferentially enhanced visibility and interaction for Musk's posts.

In Phase Two, comparing Republican-leaning and Democrat-leaning accounts, we again observed an engagement shift around the same date, affecting all metrics. However, only view counts showed evidence of a group-specific boost, with Republican-leaning accounts exhibiting a significant post-change increase relative to Democrat-leaning accounts. This finding suggests a possible recommendation bias favouring Republican content in terms of visibility, potentially via recommendation mechanisms such as the "For You" feed. Conversely, retweet and favourite counts did not display the same group-specific boost, indicating a more balanced distribution of engagement across political alignments.

Overall, the results imply that while some aspects of engagement on the platform appear to have been enhanced broadly, specific visibility advantages may have been selectively applied, raising important questions about the potential impact of algorithmic adjustments on public discourse and the 'neutrality' of social media platforms as information carriers.

Data collection

For both Phase One and Phase Two of the study, data were collected using the Zeeschuimer¹ open source tool by searching for posts by the target set of accounts via the X advanced search web interface. This was conducted over a period of two days, ensuring that all posts sent by each account between 1st January 2024 and 25 October 2024 were collected. This produced a dataset of 56,184 posts.

The accounts selected for this analysis are divided into Republican-leaning and Democrat-leaning categories, each consisting of five accounts.

Republican-leaning accounts:

- Jack Posobiec (conservative commentator known for his strong support of Republican causes)
- Tucker Carlson (political commentator and former Fox News host)
- Ben Shapiro (founder of The Daily Wire and prominent conservative voice)
- Donald Trump Jr. (political activist and son of former President Trump)
- Libs of TikTok (a controversial account sharing liberal TikTok content, often for conservative critique).

Democrat-leaning accounts:

- Alexandria Ocasio-Cortez (AOC) (progressive congresswoman from New York)
- Bernie Sanders (Senator from Vermont known for his democratic socialist stance)
- Cory Booker (Senator from New Jersey with a strong focus on social justice)
- Kamala Harris (current Vice President of the United States)
- Stacey Abrams (political leader and voting rights advocate).

The Republican-aligned accounts tended to post much more than the Democrat-aligned accounts. The number of posts authored by each account within the time range is as follows:

- JackPosobiec: 25165 posts
- elonmusk: 17595 posts
- libsoftiktok: 8792 posts
- DonaldJTrumpJr: 1624 posts
- KamalaHarris: 1515 posts
- AOC: 502 posts
- BernieSanders: 438 posts
- TuckerCarlson: 283 posts
- CoryBooker: 263 posts
- staceyabrams: 213 posts

All analysis is conducted in the Python language using open source packages for data wrangling and statistical analysis. Code and instructions to reproduce this analysis is provided on GitHub: <https://github.com/timothyjgraham/AlgorithmicBiasX>.

Limitations

¹ <https://github.com/digitalmethodsinitiative/zeeschuimer>

The study is primarily limited by the comprehensiveness of the data. While careful efforts have been made to collect as many posts from the timelines of the accounts as possible, there is no guarantee that 100% of posts have been collected. The scraping process used the Advanced Search tool on the web-based version of X, which scrolls backwards in time to collect the posts. We do not find any evidence of missing posts, as the plot of post counts over time for all accounts do not show any patterns of missing data. However, given the difficulty in collecting data on X after the Academic API was removed, we note that the study is limited because of the relatively small amount of data we could collect and the possibility of missing posts. Future work should analyse this at scale using more comprehensive datasets, ideally collected using the Premium API.

Phase 1: Analysis of Potential Algorithmic Boost to Elon Musk's Engagement Metrics

Methods

To assess whether Elon Musk's account received differential engagement boosts on X compared to other prominent accounts, we examined three engagement metrics: view counts, retweet counts, and favourite counts. We aimed to detect potential algorithmic bias favouring Musk's account, using statistical models to rigorously compare engagement patterns before and after a key change point. Data included timestamped engagement counts across multiple accounts, allowing us to analyse patterns in daily engagement for Musk's account relative to other accounts on the platform.

Change Point Detection Using CUSUM

To identify key shifts in engagement, we applied the Cumulative Sum (CUSUM) method to each metric. CUSUM captures deviations from the historical average, allowing us to detect sudden, significant increases or decreases in engagement. For each metric:

1. CUSUM Computation: Daily engagement values were resampled, calculating cumulative deviations from the mean. This approach highlighted days with substantial shifts in engagement.
2. Change Point Identification: The maximum deviation in each CUSUM series marked the "change point" for Musk's account and for other accounts. We then selected the latter of these two points as a unified change point for each metric, ensuring consistent analysis and facilitating post/pre-change comparison.

Difference-in-Differences (DiD) Analysis

With the unified change point established for each metric, we implemented a Difference-in-Differences (DiD) analysis to estimate the effect of the change point on Musk's engagement compared to other accounts:

1. Group Assignment: A binary variable (Group) identified Musk's account (coded as 1) versus other accounts (coded as 0).
2. Post-Change Indicator: We added a binary indicator (Post) marking observations after the unified change point.

3. **Interaction Term:** An interaction term (Post_Group) between Group and Post measured if the engagement shift post-change was specific to Musk's account. A significant interaction term would suggest a unique engagement boost for Musk after the change point.
4. **OLS Regression:** We modelled each engagement metric (view count, retweet count, favourite count) as the dependent variable, with Post, Group, and Post_Group as independent variables. The resulting Ordinary Least Squares (OLS) regression model enabled us to isolate the specific effect of the change on Musk's engagement patterns.

Mann-Whitney U Test

To validate changes in engagement distributions for Musk's account and other accounts, we conducted a Mann-Whitney U test on each metric:

1. **Pre- and Post-Change Distributions:** For each metric, we calculated engagement distributions separately for the periods before and after the unified change point.
2. **Non-Parametric Comparison:** Given the skewed nature of social media engagement data, the Mann-Whitney U test (a non-parametric test) was employed to determine if shifts in engagement distribution were statistically significant.


This multi-faceted approach - combining CUSUM for change detection, DiD for comparative analysis, and Mann-Whitney U for distributional shifts - enables a robust examination of potential algorithmic bias. By comparing Musk's engagement patterns to those of other accounts, this analysis provides insights into whether algorithmic bias may influence engagement on social media platforms.

Results

Across all metrics (view counts, retweet counts, and likes counts), the Bayesian change point detection method identified a structural break for Musk's metrics around July 13, 2024. Detailed statistics are provided in the Appendix A and we provide summary findings here for each metric under analysis.

View Count Analysis

Change Point: A unified change point for view counts was identified around July 13, 2024, suggesting a significant shift in viewing patterns around this date.



OLS Regression Analysis: The regression analysis for view counts revealed several noteworthy findings:

- **Constant (Baseline View Count):** The baseline view count for other accounts was estimated at approximately 392,100 ($p < 0.001$), suggesting a considerable baseline level of visibility before the change point.
- **Post (Overall Change):** The Post coefficient ($2.215e+05$, $p = 0.038$) indicates a significant increase in view counts across all accounts after the change point, pointing to a general uplift in visibility after July 12, 2024.
- **Group (Elon Musk's Account):** The Group coefficient ($4.428e+06$, $p < 0.001$) indicates that Elon Musk's account had substantially higher baseline view counts compared to other accounts, with an average difference of about 4.43 million views prior to the change point. This stark baseline difference aligns with Musk's high-profile status on the platform.
- **Post_Group (Interaction Effect):** The Post_Group interaction term ($6.398e+06$, $p < 0.001$) shows that Elon Musk's account experienced an additional significant increase in view counts post-change. This finding suggests that, following the change point, Musk's posts received an incremental boost of around 6.4 million views on average, independent of the general uplift seen across accounts.

Mann-Whitney U Test:

- **Elon Musk's Account:** The Mann-Whitney U test for Musk's account shows a highly significant change in view counts after the unified change point ($p \approx 2.27e-95$), reinforcing the presence of a substantial post-change increase.
- **Other Accounts:** The Mann-Whitney U test for other accounts also reveals a significant shift in view counts post-change ($p \approx 1.00e-113$), albeit to a lesser degree than that observed for Musk's account.

Percentage difference in view counts


Following the identified structural break on July 13, 2024, Elon Musk's account showed a substantial increase in view counts, rising by approximately 138.27% compared to his average view count before the change. In contrast, other accounts experienced a more moderate increase of 56.93% in view counts over the same period. This marked difference suggests that, post-change, Musk's content gained significantly more visibility than that of other accounts, further supporting the potential influence of algorithmic bias amplifying his reach disproportionately.

Summary of results for view counts

These results suggest that Elon Musk's account not only started from a higher baseline of views but also benefited from an enhanced increase in visibility post-change relative to other accounts. The significant Post_Group interaction term points to a unique post-change boost for Musk's account, potentially indicating algorithmic adjustments that favoured Musk's content in terms of visibility or recommendation. This differential uplift supports the hypothesis of a visibility bias that specifically enhanced Musk's reach beyond the general trend.

Retweet Count Analysis

Change Point: A unified change point for retweet counts was detected around July 13, 2024, signalling a shift in retweet patterns on and around this date.



OLS Regression Analysis: The regression analysis on retweet counts yielded the following key insights:

- **Constant (Baseline Retweet Count):** The baseline retweet count for other accounts was estimated at approximately 1,066 retweets ($p < 0.001$), providing a reference point for retweet activity before the change.
- **Post (Overall Change):** The Post coefficient (1,601.37, $p < 0.001$) indicates a significant increase in retweets across all accounts after the change point, reflecting a general uptick in retweet engagement after July 12, 2024.
- **Group (Elon Musk's Account):** The Group coefficient (2,070.91, $p < 0.001$) suggests that Elon Musk's account had a significantly higher baseline retweet count relative to other accounts prior to the change point, with an average difference of approximately 2,071 retweets.
- **Post_Group (Interaction Effect):** The Post_Group interaction term (5,861.18, $p < 0.001$) shows that Musk's account experienced a unique additional increase in retweets post-change. This result implies that, following the change point, Musk's tweets garnered an average of approximately 5,861 more retweets than those of other accounts, over and above the general post-change increase.

Mann-Whitney U Test:

- **Elon Musk's Account:** The Mann-Whitney U test for Musk's account indicates a highly significant change in retweet counts after the change point ($p \approx 8.23e-261$), underscoring a substantial increase in retweet activity.
- **Other Accounts:** The Mann-Whitney U test for other accounts also yielded a significant result ($p \approx 0$), indicating a notable shift in retweet counts, though the magnitude was less pronounced than for Musk's account.

Percentage difference in retweet counts


For retweets, the change point coincided with a sharp increase in Musk's engagement, with his average retweet count rising by 237.94% after July 13, 2024. Other accounts also saw a notable increase in retweet counts, averaging a 152.15% gain, yet this uplift was considerably smaller than that observed for Musk. This disparity implies that Musk's tweets not only benefited from a high baseline of retweet activity but also received a uniquely heightened amplification post-change, underscoring a potential platform-driven engagement boost favouring his content.

Summary of results for retweet counts

These findings suggest that Elon Musk's account had an elevated baseline of retweet engagement, with a marked additional increase in retweets following the detected change point. The significant Post_Group interaction term points to a specific post-change uplift in retweets for Musk's account, which likely reflects algorithmic adjustments favouring Musk's content in terms of engagement. This observed differential increase in retweets suggests that Musk's tweets may have been more actively promoted or featured, potentially contributing to a distinct visibility advantage post-change.

Favourite (Like) Count Analysis

Change Point: A unified change point for favourite counts was detected around July 13, 2024, indicating a shift in user interaction patterns on this date.



OLS Regression Analysis: The analysis of favourite counts revealed the following significant factors:

- **Constant (Baseline Favourite Count):** The baseline favourite count for other accounts was estimated at around 4,628 favourites ($p < 0.001$), representing the average engagement before the change point.
- **Post (Overall Change):** The Post coefficient (5,947.47, $p < 0.001$) indicates a significant increase in favourites across all accounts after the change point, pointing to an overall boost in likes after July 12, 2024.
- **Group (Elon Musk's Account):** The Group coefficient (18,920, $p < 0.001$) shows that Musk's account had a significantly higher baseline in favourite counts than other accounts before the change point, with an average difference of around 18,920 favourites.
- **Post_Group (Interaction Effect):** The Post_Group interaction term (37,540, $p < 0.001$) demonstrates that Musk's account experienced a unique, additional increase in favourite counts post-change. This result suggests that, following the change point, Musk's tweets received an average of about 37,540 more favourites than those of other accounts, beyond the general post-change increase.

Mann-Whitney U Test:

- **Elon Musk's Account:** The Mann-Whitney U test for Musk's account showed a highly significant change in favourite counts after the change point ($p \approx 1.10e-180$), indicating a considerable increase in likes.
- **Other Accounts:** The Mann-Whitney U test for other accounts was also highly significant ($p \approx 1.62e-210$), suggesting a significant increase in favourites post-change, although less substantial than Musk's account.

Percentage difference in favourites counts

For retweets, the change point coincided with a sharp increase in Musk's engagement, with his average retweet count rising by 237.94% after July 13, 2024. Other accounts also saw a notable increase in retweet counts, averaging a 152.15% gain, yet this uplift was considerably smaller than that observed for Musk. This disparity implies that Musk's tweets not only benefited from a high baseline of retweet activity but also received a uniquely heightened amplification post-change, underscoring a potential platform-driven engagement boost favouring his content.

Summary of results for favourites counts

These findings reveal that Elon Musk's account not only maintained a higher baseline of favourite counts compared to other accounts but also experienced a pronounced increase following the change point. The significant Post_Group interaction term points to a specific post-change boost for Musk's tweets, likely reflecting increased visibility or algorithmic emphasis on his content. This differential boost in likes suggests that Musk's tweets may have been further promoted, leading to a substantial visibility advantage and higher engagement post-change.

Conclusion for Phase One


The analysis of Elon Musk's engagement metrics on X (formerly Twitter) reveals that Musk's account exhibited distinct and elevated engagement patterns compared to other accounts, particularly around a key structural change on July 13, 2024. Across view counts, retweet counts, and favourite counts, Musk's content not only began with a higher baseline of engagement but also benefited from an additional, significant increase following the change point. This amplified post-change boost, highlighted by significant Post_Group interaction terms across all metrics, suggests a potential platform-level adjustment that disproportionately enhanced Musk's visibility and engagement relative to other accounts.

Specifically, the regression and Mann-Whitney U test results indicate that Musk's posts received a marked increase in visibility (view counts), amplification (retweet counts), and user interaction (favourite counts) that outpaced the general engagement trends observed across the platform. The unique post-change uplift for Musk's account suggests the possibility of algorithmic prioritisation or bias, positioning Musk's content favourably in terms of platform visibility and user engagement.

These findings underscore a distinct pattern that may indicate an algorithmic shift that disproportionately favoured Musk's account, contributing to a considerable engagement advantage. This visibility bias, if linked to platform algorithm adjustments, highlights the impact of such structural changes on engagement dynamics and the potential for differential treatment among users. This outcome has broader implications for understanding how platform algorithms may shape public discourse by influencing which voices and content receive heightened visibility and interaction.

Appendix A

Detailed results for view count analysis



OLS Regression Results for view_count:

OLS Regression Results						
Dep. Variable:	view_count		R-squared:	0.110		
Model:	OLS		Adj. R-squared:	0.110		
Method:	Least Squares		F-statistic:	2323.		
Date:	Thu, 31 Oct 2024		Prob (F-statistic):	0.00		
Time:	12:02:39		Log-Likelihood:	-9.8655e+05		
No. Observations:	56184		AIC:	1.973e+06		
Df Residuals:	56180		BIC:	1.973e+06		
Df Model:	3					
Covariance Type:	nonrobust					
	coef	std err	t	P> t	[0.025	0.975]
const	3.921e+05	6.64e+04	5.903	0.000	2.62e+05	5.22e+05
Post	2.215e+05	1.07e+05	2.071	0.038	1.19e+04	4.31e+05
Group	4.428e+06	1.17e+05	37.767	0.000	4.2e+06	4.66e+06
Post_Group	6.398e+06	1.93e+05	33.200	0.000	6.02e+06	6.78e+06
Omnibus:	66578.675		Durbin-Watson:	1.827		
Prob(Omnibus):	0.000		Jarque-Bera (JB):	9336908.691		
Skew:	6.304		Prob(JB):	0.00		
Kurtosis:	64.883		Cond. No.	5.87		


Elon Musk's Account - Mann-Whitney U Test for view_count:

Stat=29081256.0, P-value=2.2687672435022933e-95

Other Accounts - Mann-Whitney U Test for view_count: Stat=152285518.0,

P-value=1.004693197870457e-113

Detailed results for retweet count analysis




OLS Regression Results for retweet_count:

OLS Regression Results						
Dep. Variable:	retweet_count		R-squared:	0.077		
Model:	OLS		Adj. R-squared:	0.077		
Method:	Least Squares		F-statistic:	1556.		
Date:	Thu, 31 Oct 2024		Prob (F-statistic):	0.00		
Time:	12:02:40		Log-Likelihood:	-5.9676e+05		
No. Observations:	56184		AIC:	1.194e+06		
Df Residuals:	56180		BIC:	1.194e+06		
Df Model:	3					
Covariance Type:	nonrobust					
	coef	std err	t	P> t	[0.025	0.975]
const	1066.0661	64.448	16.542	0.000	939.748	1192.384
Post	1601.3686	103.766	15.432	0.000	1397.986	1804.751
Group	2070.9056	113.755	18.205	0.000	1847.945	2293.866
Post_Group	5861.1837	186.987	31.345	0.000	5494.687	6227.680
Omnibus:	83917.835		Durbin-Watson:	1.811		
Prob(Omnibus):	0.000		Jarque-Bera (JB):	67985190.774		
Skew:	8.995		Prob(JB):	0.00		
Kurtosis:	172.462		Cond. No.	5.87		

Elon Musk's Account - Mann-Whitney U Test for retweet_count: Stat=24617179.0, P-value=8.227240892960878e-261

Other Accounts - Mann-Whitney U Test for retweet_count: Stat=129256985.5, P-value=0.0

Detailed results for favourite (likes) count analysis



OLS Regression Results for favorite_count:

OLS Regression Results						
Dep. Variable:	favorite_count	R-squared:	0.080			
Model:	OLS	Adj. R-squared:	0.080			
Method:	Least Squares	F-statistic:	1624.			
Date:	Thu, 31 Oct 2024	Prob (F-statistic):	0.00			
Time:	12:02:42	Log-Likelihood:	-7.0255e+05			
No. Observations:	56184	AIC:	1.405e+06			
Df Residuals:	56180	BIC:	1.405e+06			
Df Model:	3					
Covariance Type:	nonrobust					
	coef	std err	t	P> t	[0.025	0.975]
const	4628.2738	423.632	10.925	0.000	3797.952	5458.596
Post	5947.4661	682.083	8.720	0.000	4610.578	7284.354
Group	1.892e+04	747.742	25.309	0.000	1.75e+04	2.04e+04
Post_Group	3.754e+04	1229.116	30.545	0.000	3.51e+04	4e+04
Omnibus:	96410.815	Durbin-Watson:	1.868			
Prob(Omnibus):	0.000	Jarque-Bera (JB):	192676371.161			
Skew:	11.699	Prob(JB):	0.00			
Kurtosis:	288.933	Cond. No.	5.87			

Elon Musk's Account - Mann-Whitney U Test for favorite_count: Stat=26507926.0, P-value=1.1001526625562207e-180

Other Accounts - Mann-Whitney U Test for favorite_count: Stat=143444419.5, P-value=1.616655138095897e-210

Phase 2: Is X algorithmically boosting the reach of Republican-aligned accounts?

Methods

This second phase of the study investigates whether there is a measurable difference in engagement metrics - view counts, retweet counts, and favourite (likes) counts - between selected Republican-leaning and Democrat-leaning accounts on X (formerly Twitter). The dataset used in this study is the same as in Phase One, reported above. As with the Phase One methodology, the analysis considers three primary dimensions: identifying potential change points in engagement over time, measuring differences in engagement across the two account groups, and assessing whether any observed shifts align with algorithmic changes or preferential boosts. The methods are therefore the same as in Phase One, however we highlight them again here for readability.

CUSUM Analysis for Change Point Detection

To identify significant shifts in engagement over time, we applied a Cumulative Sum (CUSUM) analysis. This approach detects deviations from the average value over time, allowing us to pinpoint moments where a change in engagement behaviour might indicate external factors, such as algorithmic adjustments. Engagement metrics for each group were aggregated on a daily basis using a time-resampling approach. The daily values were analysed to calculate the cumulative sum of deviations from the mean, plotted to visually identify change points. This approach helps highlight days where engagement levels sharply diverged from the norm, suggesting a potential alteration in algorithmic prioritisation or user interest patterns.

For each metric - view count, retweet count, and favourite count - a separate change point was detected using the CUSUM method. Interestingly, while these change points were detected independently, they aligned across all metrics, pointing to July 14, 2024, as a unified change point. This alignment strengthens the likelihood that this date represents a platform-wide alteration affecting all engagement types, rather than isolated shifts within specific metrics. Thus, this unified change point became the standard point of reference for subsequent analyses.

Rationale for the Unified Change Point

The alignment of individual change points across metrics provides a compelling basis for adopting a unified change point in this analysis. A single change point facilitates consistent comparisons across engagement metrics and reduces metric-specific noise, which can lead to more robust inferences. Given that platform-wide algorithmic changes would likely impact multiple engagement metrics simultaneously, interpreting this unified change point as a systemic shift offers a more coherent framework for analysing potential biases. Consequently, this unified change point serves as a reliable anchor for examining the effects of possible platform adjustments on Republican and Democrat-leaning accounts alike.

Difference-in-Differences (DiD) Analysis

With the unified change point established, we conducted a Difference-in-Differences (DiD) analysis to evaluate the impact of these changes on engagement metrics. DiD is a quasi-experimental statistical method that estimates the causal impact of a treatment - in this case, the unified change point - on an outcome by comparing pre- and post-change engagement for Republican- and Democrat-leaning accounts. This model enables us to separate the effect of time from that of group affiliation, as well as their interaction.

We introduced a binary variable, *Post*, to indicate whether a given observation occurred after the unified change point and a *Group* variable to distinguish between Republican and Democrat accounts. An interaction term, *Post_Group*, was also added to capture any combined effect of group and change period. Using an Ordinary Least Squares (OLS) regression model with *Post*, *Group*, and *Post_Group* as predictors, we analysed each metric separately to assess whether the change in engagement varied significantly between the two groups. A significant interaction term (*Post_Group*) would suggest that the engagement pattern changed differently for one group, potentially indicating algorithmic bias.

Pre/Post Comparison Using Mann-Whitney U Test

To further understand changes in engagement patterns before and after the unified change point, we performed a Mann-Whitney U test for each metric within both groups. This non-parametric test evaluates whether the distributions of engagement counts before and after the change point differ significantly. Unlike parametric tests, the Mann-Whitney U test does not assume normality, making it a robust choice for social media engagement data, which often exhibit non-normal distributions due to viral spikes or other irregularities.

By applying these three complementary analyses - CUSUM for change detection, DiD for assessing differential effects between groups, and Mann-Whitney U for testing distributional shifts - we aimed to assess whether Republican- and Democrat-leaning accounts received differential boosts in engagement metrics, potentially attributable to platform algorithmic changes. The results from these analyses are presented below.


Findings

Overall, we do not find as strong evidence of a potential algorithmic bias as we did in Phase One, where we looked at Elon Musk's account versus the rest of the accounts. The findings suggest that there is likely some kind of pro-Republican bias for *view counts* after the detected structural break on July 12. This is not the case for the structural break we detect for retweet and favourites counts, which occurs on July 14 for both engagement metrics and appears to affect both groups equally.

These change points - July 12th for views and July 14th for retweets and favourites - suggests a likely systemic adjustment in engagement patterns across these metrics, even if the differential engagement after the break is only clearly observed for one of the metrics (i.e., view counts are higher for pro-Republican accounts).

Below, we summarise the findings for each metric, focusing on statistically significant changes and notable interaction effects.

View Count



Change Point Detection: The unified change point for view counts was detected on July 12, 2024, indicating a significant shift in engagement behaviour on the platform for both Republican and Democrat accounts.

OLS Regression Results:

- The **Post coefficient** ($1.199\text{e}+06$, $p = 0.003$) indicates a statistically significant increase in view counts across both groups following the change point. This suggests a platform-wide engagement boost, translating to an approximate increase of 1.2 million views post-change, which affected accounts across the political spectrum.
- The **Group coefficient** ($1.322\text{e}+06$, $p < 0.001$) shows that Republican accounts consistently had higher baseline view counts, with around 1.3 million more views than Democrat accounts before the change. This disparity highlights an underlying engagement gap favouring Republican accounts, independent of the change point.
- The **Post_Group interaction term** ($9.523\text{e}+05$, $p = 0.021$) is significant and large, suggesting that Republican accounts experienced an additional increase in view counts specifically after the change point. This effect, amounting to approximately 952,300 additional views, could reflect an amplified visibility for Republican accounts, beyond the general boost experienced by both groups.


Mann-Whitney U Test:

- The Mann-Whitney U test results reinforce the regression findings, showing a highly significant shift in view counts after the change point for both groups. Republican accounts, in particular, exhibit a pronounced increase in post-change view distributions ($p \approx 4.62\text{e}-177$ for Republicans; $p \approx 7.14\text{e}-47$ for Democrats).

Interpretation: The significant Post_Group interaction term, combined with the substantial increase in views for Republican accounts post-change, suggests a possible

recommendation or visibility bias favouring these accounts. This finding implies that Republican-leaning content may have been algorithmically promoted to a broader audience after the change point - potentially through features like the “For You” feed - which would increase the visibility of these accounts and drive higher view counts. The results underscore a noteworthy platform-driven amplification of engagement, with a marked advantage for Republican accounts in the post-change period.

Retweet Count



Change Point Detection: The unified change point for retweet counts was detected on July 14, 2024, marking a significant shift in retweet engagement across Republican and Democrat accounts.

OLS Regression Results:


- The **Post coefficient** (3297.73, $p < 0.001$) indicates a statistically significant increase in retweet counts for both Republican and Democrat accounts following the change point. This suggests an approximate boost of 3,298 retweets for each account after the change, pointing to an overall uplift in engagement across the platform.
- The **Group coefficient** (165.42, $p = 0.543$) shows no significant baseline difference in retweet counts between Republican and Democrat accounts prior to the change point. This suggests that both groups had similar levels of retweet engagement leading up to the change.
- The **Post_Group interaction term** (8.53, $p = 0.983$) is not significant, indicating that the post-change increase in retweet engagement was evenly distributed across both groups without a specific benefit to either Republican or Democrat accounts.

Mann-Whitney U Test:

- Results from the Mann-Whitney U test reveal significant shifts in retweet distributions for both Republican ($p \approx 0$) and Democrat ($p \approx 6.73e-97$) accounts after the change point. This further underscores the substantial overall increase in retweet counts, but does not show a marked divergence between the two groups.

Interpretation: The lack of a significant Post_Group interaction term, alongside the similar baseline retweet counts between Republican and Democrat accounts, suggests that the increase in retweet engagement post-change was likely driven by a platform-wide adjustment rather than preferential treatment for one group. The detected increase in retweet counts across both groups could reflect a general algorithmic shift in how content was promoted or surfaced, enhancing visibility and shareability equally. Unlike the view count findings, these results suggest that retweet engagement gains were broadly consistent across political affiliations, indicating no evident bias in retweet amplification between Republican and Democrat accounts in the post-change period.

Favourite Count



Change Point Detection: A unified change point for favourite counts was detected on 12 July 2024, indicating a distinct shift in how favourite engagements evolved for both Republican and Democrat accounts.

OLS Regression Results:

- The **Post coefficient** ($1.942e+04$, $p < 0.001$) reveals a statistically significant increase in favourite counts following the change point for both groups, suggesting that all accounts saw an average increase of nearly 19,420 favourites post-change.
- The **Group coefficient** (4689.18, $p = 0.010$) shows that Republican accounts had a significantly higher baseline of favourite counts compared to Democrat accounts prior to the change point. This indicates a pre-existing disparity in favourites, favouring Republican-leaning accounts.
- The **Post_Group interaction term** (-2640.43, $p = 0.307$) is not statistically significant, indicating that the increase in favourite counts post-change point was distributed across both groups without one group benefiting disproportionately over the other.

Mann-Whitney U Test:

- The Mann-Whitney U test results indicate highly significant differences in favourite counts post-change for both Republican ($p \approx 0$) and Democrat accounts ($p \approx 1.31e-74$). These results further affirm that both groups experienced a substantial post-change uplift in favourite counts.

Interpretation: The lack of significance in the Post_Group interaction term, combined with the higher baseline of favourites for Republican accounts, suggests that the post-change increase in favourite counts was applied broadly, affecting both groups similarly. Unlike view counts, where Republican accounts appeared to gain a unique post-change boost, favourite counts rose evenly between the two groups post-change, suggesting a general increase in engagement without apparent platform bias. This uniformity indicates that the observed uplift may have resulted from a platform-wide algorithmic shift promoting content engagement through increased visibility for all accounts rather than any selective amplification favouring one political affiliation.

Conclusion for Phase Two

The change points detected for the structural breaks in engagement are quite close in time (12 July for views and 14 July for retweets and favourites), suggesting that system wide changes occurred around those dates. For the view count results, the significant Post_Group interaction term suggests that *Republican accounts benefited from increased visibility or recommendation bias post-change*. This could imply an algorithmic adjustment, potentially in the For You feed, designed to place certain content in front of more users, thus driving up views selectively for these accounts.

In contrast, *retweet and favourite counts showed no such group-specific boost*, indicating that the overall engagement boost observed in these metrics was distributed fairly across both Republican and Democrat accounts. This consistency in change points without group-specific boosts in retweets and favourites suggests that the observed increase was likely the result of a general engagement shift rather than a targeted amplification.

We speculate that the view count findings suggest possible recommendation algorithm bias for visibility that privileges pro-Republican accounts, warranting further investigation to


OLS Regression Results						
=====						
Dep. Variable:	view_count		R-squared:	0.010		
Model:	OLS		Adj. R-squared:	0.010		
Method:	Least Squares		F-statistic:	186.2		
Date:	Thu, 31 Oct 2024		Prob (F-statistic):	3.76e-120		
Time:	11:37:04		Log-Likelihood:	-9.8956e+05		
No. Observations:	56184		AIC:	1.979e+06		
Df Residuals:	56180		BIC:	1.979e+06		
Df Model:	3					
Covariance Type:	nonrobust					
=====						
	coef	std err	t	P> t	[0.025	0.975]

const	5.457e+05	2.84e+05	1.923	0.055	-1.06e+04	1.1e+06
Post	1.199e+06	4.01e+05	2.987	0.003	4.12e+05	1.99e+06
Group	1.322e+06	2.9e+05	4.562	0.000	7.54e+05	1.89e+06
Post_Group	9.523e+05	4.13e+05	2.307	0.021	1.43e+05	1.76e+06
=====						
Omnibus:	68795.720		Durbin-Watson:	1.643		
Prob(Omnibus):	0.000		Jarque-Bera (JB):	9348091.430		
Skew:	6.726		Prob(JB):	0.00		
Kurtosis:	64.743		Cond. No.	21.6		
=====						

Republican Accounts - Mann-Whitney U Test for view_count:
Stat=283090267.5, P-value=4.622362594560296e-177

Democrat Accounts - Mann-Whitney U Test for view_count: Stat=721540.5,
P-value=7.135371717953282e-47

Detailed results for retweet count analysis



OLS Regression Results for retweet_count:


OLS Regression Results						
=====						
Dep. Variable:	retweet_count		R-squared:	0.024		
Model:	OLS		Adj. R-squared:	0.024		
Method:	Least Squares		F-statistic:	457.9		
Date:	Thu, 31 Oct 2024		Prob (F-statistic):	5.40e-294		
Time:	11:37:06		Log-Likelihood:	-5.9832e+05		
No. Observations:	56184		AIC:	1.197e+06		
Df Residuals:	56180		BIC:	1.197e+06		
Df Model:	3					
Covariance Type:	nonrobust					
=====						
	coef	std err	t	P> t	[0.025	0.975]

const	1603.5126	266.189	6.024	0.000	1081.780	2125.245
Post	3297.7316	379.614	8.687	0.000	2553.686	4041.778
Group	165.4196	271.902	0.608	0.543	-367.510	698.349
Post_Group	8.5302	390.560	0.022	0.983	-756.969	774.030
=====						
Omnibus:	85387.126		Durbin-Watson:	1.715		
Prob(Omnibus):	0.000		Jarque-Bera (JB):	67972104.791		
Skew:	9.343		Prob(JB):	0.00		
Kurtosis:	172.370		Cond. No.	21.4		
=====						

Republican Accounts - Mann-Whitney U Test for retweet_count: Stat=238548439.0,
P-value=0.0

Democrat Accounts - Mann-Whitney U Test for retweet_count: Stat=575277.5,
P-value=6.734339153650122e-97

Detailed results for favourite (likes) count analysis



OLS Regression Results for favorite_count:

OLS Regression Results						
Dep. Variable:	favorite_count	R-squared:	0.015			
Model:	OLS	Adj. R-squared:	0.015			
Method:	Least Squares	F-statistic:	276.6			
Date:	Thu, 31 Oct 2024	Prob (F-statistic):	2.95e-178			
Time:	11:37:08	Log-Likelihood:	-7.0448e+05			
No. Observations:	56184	AIC:	1.409e+06			
Df Residuals:	56180	BIC:	1.409e+06			
Df Model:	3					
Covariance Type:	nonrobust					
=====						
	coef	std err	t	P> t	[0.025	0.975]

const	6207.4623	1775.574	3.496	0.000	2727.325	9687.599
Post	1.942e+04	2511.041	7.732	0.000	1.45e+04	2.43e+04
Group	4689.1766	1813.505	2.586	0.010	1134.696	8243.657
Post_Group	-2640.4298	2582.866	-1.022	0.307	-7702.863	2422.003
=====						
Omnibus:	96423.379	Durbin-Watson:	1.745			
Prob(Omnibus):	0.000	Jarque-Bera (JB):	172955357.354			
Skew:	11.755	Prob(JB):	0.00			
Kurtosis:	273.792	Cond. No.	21.6			

Republican Accounts - Mann-Whitney U Test for favorite_count: Stat=269981297.0,
P-value=6.485712072189484e-284

Democrat Accounts - Mann-Whitney U Test for favorite_count: Stat=634673.5,
P-value=1.3084276945000918e-74