

The Attention Trap: How International Celebrity Intervention Displaces Policy Discourse in Protest Movements

Shreyas Meher

April 14, 2026

Abstract

Does international celebrity attention help or hinder social movements? I analyse 1.08 million tweets from the 2020–2021 Indian farmer protests to measure the effect of celebrity attention shocks on protest discourse. When Rihanna tweeted about the protests on 2 February 2021, daily volume increased 22-fold, but the share of tweets discussing policy issues fell by 10 percentage points while celebrity and meta-commentary rose by 45 percentage points. I call this the *attention trap*: celebrity intervention maximizes visibility while displacing the substantive discourse movements depend on. The 70,000 new participants drawn by the tweet were disproportionately non-Indian, apolitical in profile, and overwhelmingly transient, with two-thirds departing within a single day. Local projection impulse responses show that this discourse displacement is specific to the celebrity channel: domestic events, including state violence at Lakhimpur Kheri, generate policy-focused responses with no comparable displacement. Granger causality confirms that celebrity content predicts subsequent declines in policy discussion, but not the reverse. The cost was strategically exploitable: the BJP government launched a coordinated nationalist counter-campaign within hours, on discursive terrain made favourable by the displacement itself. These findings reframe celebrity attention as a trade-off with a hidden cost that authoritarian regimes can exploit.

Keywords: celebrity activism, discourse displacement, social movements, attention economy, Indian politics

Word count: 6,711

Introduction

On 2 February 2021, the American musician Rihanna posted a single tweet: ‘why aren’t we talking about this?!’ with a link to a CNN article about the Indian farmer protests. Within 24 hours, the hashtag #FarmersProtest had attracted more than 80,000 tweets, an increase of over 2,000 percent from the previous day’s baseline. International figures including Greta Thunberg, Mia Khalifa, and Meena Harris amplified the message, directing a global audience toward a movement that had received comparatively limited international attention. The Indian government responded with a coordinated counter-campaign the following morning, mobilizing Bollywood celebrities to tweet under #IndiaTogether and #IndiaAgainstPropaganda in synchronized waves that investigators traced to the Ministry of External Affairs ([Mishra et al., 2021](#)).

The conventional wisdom in both activist circles and the academic literature on social movements holds that such attention is an unqualified good. Visibility generates sympathy, sympathy generates political pressure, and political pressure generates concessions. This logic undergirds the extensive scholarship on celebrity activism, which has documented how celebrity endorsements can raise the salience of political causes, attract media coverage, and channel public attention toward issues that would otherwise languish in obscurity ([Thrall et al., 2008](#); [Brockington, 2014](#)). It also undergirds the practice of contemporary social movements, which devote considerable organizational resources to courting celebrity support and engineering viral moments on social media platforms. The assumption, rarely interrogated with rigorous evidence, is that attention and substantive political engagement move in the same direction: more of the former produces more of the latter.

In this paper, I show that the relationship between attention and substance is more complicated, and in certain respects, adversarial. The Rihanna tweet dramatically expanded the farmer protest conversation on Twitter while simultaneously transforming its content. Before the tweet, users discussed minimum support prices, the three farm acts, and corporate consolidation. After the tweet, the conversation shifted toward commentary about Rihanna

herself, about whether international celebrities should involve themselves in Indian affairs, and about the ‘toolkit’ conspiracy theory the government promoted in response. The share of tweets containing substantive policy content fell by 10 percentage points; the share containing celebrity and meta-commentary rose by 45 percentage points. I call this pattern the *attention trap*: a trade-off in which external attention surges maximize a movement’s visibility while displacing the substantive discourse through which it builds political pressure. The trap is not that celebrity attention is worthless; it is that the attention comes with a hidden cost, one that is invisible in the volume metrics that movements and scholars typically use to measure success, and one that opponents can exploit.

I study this phenomenon using a dataset of 1.08 million tweets collected over 386 days, spanning the full arc of the farmer protest movement from the initial Dilli Chalo march in November 2020 through Prime Minister Narendra Modi’s surprise announcement repealing the three farm laws in November 2021. The dataset captures the complete life cycle of one of the largest protest movements in world history, including multiple discrete shocks of different types: coordinated domestic mobilization (the December 2020 Bharat Bandh), escalatory confrontation (the Republic Day tractor rally of 26 January 2021), international celebrity attention (the Rihanna and Thunberg tweets of 2 February 2021), state violence (the Lakhimpur Kheri massacre of 3 October 2021, in which the son of a sitting Union Minister drove a vehicle into a crowd of protesting farmers), and government concession (the farm law repeal of 19 November 2021). This variation allows me to estimate not only the celebrity attention effect in isolation, but also to compare it against substantively different events, identifying what is distinctive about the celebrity channel. I bring two econometric frameworks: local projection impulse response estimation following [Jordà \(2005\)](#) and vector autoregressive Granger causality tests following [Kulichkina et al. \(2025\)](#).

Three sets of findings emerge. First, the Rihanna shock generated a massive but transient volume increase, peaking at 83,866 tweets the day after the tweet (a 22-fold increase) and decaying to baseline within eight days. This surge was driven by 70,000 new entrants who ac-

counted for 71.5 percent of post-shock tweets, producing overwhelmingly meta-commentary rather than policy content. Second, the discourse displacement is specific to the celebrity shock: when farmers were killed at Lakhimpur Kheri, the celebrity/meta-commentary share barely moved, and when the government repealed the farm laws, policy content spiked above 50 percent. Third, Granger causality tests confirm that celebrity discourse predicts subsequent declines in policy content ($F = 4.63, p = 0.032$), but not the reverse.

These findings contribute to several literatures. Most directly, they speak to the scholarship on celebrity activism, which has debated the conditions under which celebrity endorsements translate into meaningful political engagement rather than superficial awareness. The attention trap framework offers a mechanism for when and why celebrity attention can be counterproductive: it floods a conversation with participants who are responsive to the celebrity signal rather than the underlying political content, diluting the substantive discourse that movements depend on to articulate demands and build coalitions. This is not an argument against celebrity activism in general, but rather an identification of a specific cost that has not been previously measured.

The findings also contribute to the study of digitally networked social movements. [Tufekci \(2017\)](#) has argued that digital platforms enable ‘tactical freeze,’ in which movements can scale rapidly but struggle to develop the organizational capacities needed for sustained political engagement. The attention trap is a discursive analogue of tactical freeze: movements can achieve visibility rapidly through celebrity amplification, but the participants drawn in by that visibility lack the substantive engagement needed to sustain a policy-oriented conversation. The 70,000 users who entered the farmer protest conversation after Rihanna’s tweet were, in a meaningful sense, talking about a different subject than the 20,000 users who had been discussing farm laws for the preceding months.

Finally, in the specific context of Indian politics, the findings connect to scholarship on the BJP’s use of conspiracy theories and nationalist framing to delegitimize protest movements ([Chacko, 2025](#)). The government’s ability to deploy the toolkit conspiracy theory and

the foreign interference frame depended on the prior shift in discourse terrain. When the conversation was about minimum support prices and the Swaminathan Commission report, the government was on weak ground: the protesters had expertise, legitimacy, and broad public sympathy on these substantive questions. When the conversation was about whether Rihanna should be commenting on Indian affairs, the government was on much stronger ground, able to invoke nationalist sentiment and the ‘anti-national’ label that has been a cornerstone of BJP rhetorical strategy since 2014. The attention trap is not only costly for movements in its own right; it is exploitable by the very regimes that movements seek to challenge.

Theoretical Framework

The starting point is Herbert Simon’s foundational observation that ‘a wealth of information creates a poverty of attention’ (Simon, 1996). In any information-rich environment, the binding constraint on engagement is not information availability but the cognitive capacity to process it. This insight, elaborated by Wu (2017) into a theory of the ‘attention merchants’ who compete for human focus, has profound implications for political discourse on social media. When a movement’s hashtag suddenly attracts millions of new users, the total volume of conversation increases, but the share devoted to the movement’s substantive demands need not. New attention does not simply layer on top of existing discourse; it restructures the conversation’s centre of gravity toward whatever stimulus attracted the new participants.

This restructuring is especially pronounced when the source of new attention is a celebrity. Celebrities command attention not because of what they say but because of who they are. Their communicative power derives from what Driessens (2013) calls ‘celebrity capital,’ a form of accumulated visibility analytically distinct from cultural or political capital in the Bourdieusian sense (Bourdieu, 2018). When a celebrity directs attention toward a political cause, the attention follows the celebrity, not the cause. Kristofferson et al. (2014) find that

initial acts of ‘token support’ (such as sharing a celebrity’s post) do not reliably translate into more substantive engagement and may in fact reduce subsequent prosocial action by satisfying the individual’s need for moral self-regard.¹

The academic literature on celebrity activism has evolved from early optimism toward a more ambivalent assessment. [Thrall et al. \(2008\)](#) found that while celebrities can attract media coverage to neglected issues, the coverage tends to be shallow, personality-driven, and short-lived. [Brockington \(2014\)](#) argued that celebrity advocacy often serves the celebrities’ reputational interests more than the causes they champion. [Wheeler \(2013\)](#) theorized the ‘political celebrity’ as a political actor whose influence derives from visibility rather than expertise or accountability. What this literature lacks, however, is a clear account of the *costs* that celebrity attention can impose on the movements it purports to support. The prevailing framework treats celebrity attention as a windfall that may or may not convert into policy change, but does not consider the possibility that the attention itself might be actively harmful to a movement’s discursive position. The attention trap framework fills this gap by identifying a specific mechanism: celebrity attention undermines movements not by opposing them, but by displacing their substantive discourse with meta-commentary about the celebrity intervention itself.

The mechanism at the heart of the attention trap is discourse displacement, understood through agenda-setting theory ([McCombs and Shaw, 1972](#)) and the framing literature ([Entman et al., 1993](#); [Benford and Snow, 2000](#)). When a celebrity intervenes in a political conversation, the intervention becomes the new item on the agenda, displacing the policy questions around which the movement had organized. This is related to but distinct from ‘counter-framing’ in the social movements literature ([Benford and Snow, 2000](#)). Discourse displacement does not require an opponent’s deliberate intervention; it can occur as a structural byproduct of the attention economy. That said, discourse displacement can be

¹This is the phenomenon commonly referred to as ‘slacktivism’ in both popular and academic commentary ([Morozov, 2012](#)), though the empirical evidence on whether digital engagement crowds out or complements offline participation remains contested.

exploited by opponents who recognize that the shifted terrain is more favourable to their counter-narrative.

The attention trap is particularly relevant to digitally networked protest movements. [Bennett and Segerberg \(2023\)](#) distinguish between ‘collective action,’ which requires shared identity and formal organizational coordination, and ‘connective action,’ which operates through the sharing of personalized content with minimal organizational infrastructure. Connective action can scale rapidly, but because participation requires only sharing a hashtag, the threshold for entry is low enough that celebrity attention shocks can flood the network with participants who have no prior engagement with the movement’s substantive concerns. [Tufekci \(2017\)](#) identified a related vulnerability she calls tactical freeze: movements can mobilize quickly but struggle to develop the organizational capacities needed for sustained engagement. The attention trap is a discursive analogue. Where tactical freeze concerns the organizational dimension, discourse displacement concerns the communicative dimension: movements can attract attention but cannot control what that attention is focused on.²

Recent empirical work supports these dynamics. [Barberá et al. \(2015\)](#) study the ‘critical periphery’ in social protests, finding that peripheral participants who join during attention spikes are structurally different from core participants. [González-Bailón et al. \(2011\)](#) show that recruitment into online protest networks proceeds through waves, with early adopters differing systematically from later entrants in their network position and engagement intensity. [Steinert-Threlkeld \(2017\)](#) demonstrates that peripheral mobilization during the Arab Spring, while important for scaling, was qualitatively different from the core mobilization that initiated protests. My analysis of entry cohorts around the Rihanna event is directly in conversation with this literature, providing evidence that the distinction between core and periphery extends to the *content* of engagement, not only its structural position.

²The farmer protests are an interesting hybrid case. The physical encampments at the Delhi border were organized through traditional collective action structures. The Twitter conversation, by contrast, operated as a decentralized connective action space. The attention trap manifested in the latter, not the former.

The 2020–2021 Indian Farmer Protests

India’s agricultural sector employs roughly half the population but contributes only about 15 percent of GDP. Approximately 86 percent of the country’s 146.5 million farms are smaller than one hectare, and 52 percent of farm households carry substantial debt ([Narayanan, 2021](#)). The primary mechanism through which the state supports small farmers is the minimum support price (MSP) system,³ under which the government purchases certain crops at set prices through regulated market yards known as Agricultural Produce Marketing Committees (APMCs).⁴

In September 2020, the Indian Parliament passed three pieces of legislation that farmers viewed as an existential threat to this system. The Farmers’ Produce Trade and Commerce Act limited APMC regulatory authority and allowed transactions outside the established market yard system. The Essential Commodities (Amendment) Act relaxed stocking limits on private traders, removing constraints that had limited the market power of large agribusiness corporations. The Farmers (Empowerment and Protection) Agreement on Price Assurance and Farm Services Act created a framework for contract farming that allowed corporate ‘sponsors’ to engage directly with farmers outside state regulation. Although framed by the government as liberalizing reform, farmers and agricultural economists argued that the laws were designed to benefit large corporations, particularly Reliance Industries and the Adani Group, at the expense of small producers. The fear was not abstract: farmers pointed to India’s historical experience with contract farming, in which companies like PepsiCo had sued potato farmers in Gujarat for breach of contract, and to the Adani Group’s rapidly expanding presence in agricultural logistics ([Narayanan, 2021](#)).

On 26 November 2020, tens of thousands of farmers, predominantly from Punjab and

³The MSP is an administratively determined floor price at which the government guarantees to purchase specified commodities. In practice, procurement is concentrated in Punjab and Haryana for wheat and paddy rice.

⁴APMCs are state-regulated wholesale markets originally established to protect farmers from exploitative intermediaries. Farmers feared that the alternative, unregulated private markets dominated by large corporations, would be worse.

Haryana but with participation from Uttar Pradesh, Rajasthan, and other states, marched toward Delhi under the banner of ‘Dilli Chalo’ (Onward to Delhi). When police barricaded the major highways with concrete barriers, water cannons, and tear gas, the farmers established semi-permanent encampments at Singhu, Tikri, and Ghazipur that would persist for over a year. These encampments became extraordinary sites of democratic practice: communal kitchens served free meals to anyone regardless of religion, caste, or gender; libraries and art installations appeared alongside tractors and makeshift tents; and a farmer-produced newspaper, *Trolley Times*, published in Hindi and Punjabi, provided independent coverage of the movement (Behl, 2022). The Samyukt Kisan Morcha (SKM), a coalition of more than 40 farmer unions, served as the movement’s coordinating body, maintaining a discipline and organizational coherence that defied the government’s repeated attempts to dismiss the protests as the work of a disorganized fringe.

Several key events punctuated the movement’s trajectory. On 8 December 2020, the SKM called a Bharat Bandh (national strike) that demonstrated the movement’s capacity for coordinated action. On 26 January 2021, Republic Day, a tractor rally turned violent when some protesters breached the Red Fort, and the government ordered Twitter to suspend hundreds of protest-associated accounts (Chacko, 2025). The farmer leader Rakesh Tikait’s emotional refusal to abandon the Ghazipur encampment prevented the camps from dissolving.

Rihanna’s tweet arrived on 2 February 2021, directing her 100 million followers toward a CNN article about internet shutdowns at protest sites. Greta Thunberg shared a toolkit document, which the government seized upon as evidence of an international conspiracy. The BJP responded with a coordinated celebrity counter-campaign: actors and cricketers posted nearly identical messages at roughly twenty-minute intervals under #IndiaTogether (Mishra et al., 2021). The Delhi Police arrested climate activist Disha Ravi for editing the toolkit, charging her with sedition. For the government, the episode provided the discursive raw material for a narrative it had been constructing: that the farmer movement was a

foreign-backed conspiracy involving Khalistani separatists, Western NGOs, and international celebrities (Chacko, 2025).

The most violent episode came on 3 October 2021, when a convoy including a vehicle belonging to Ashish Mishra, son of the Union Minister of State for Home Affairs, drove into protesting farmers at Lakhimpur Kheri, killing four farmers and a journalist. On 19 November 2021, Prime Minister Modi made a surprise announcement repealing all three farm laws, widely interpreted as an electoral calculation ahead of state elections in Punjab and Uttar Pradesh.

Data and Methods

Data

The primary dataset consists of 1,084,452 tweets collected using the Twitter search API with farmer-protest-related hashtags, principally #FarmersProtest and its variants.⁵ The dataset spans 386 days from 1 November 2020 to 21 November 2021. Each tweet record includes the full text, timestamp (UTC), user identifier, and engagement metrics. A companion dataset of 397,718 user profiles provides account-level metadata. Because the data was collected using pro-farmer hashtags, the corpus represents the protest side of the conversation. This is the appropriate frame of analysis: the attention trap operates within the movement’s own discursive space.

I supplement the Twitter data with 1,486,389 articles from the Times of India (TOI). For the overlap period, I identify 6,278 articles as relevant to the farmer protests using keyword matching on titles and content. The daily count of farmer-protest articles serves as a measure of mainstream media attention for the Granger causality analysis.

⁵The tweet and user data were collected and published by Pratham Sharma as the Farmers Protest Tweets Dataset on Kaggle, released under a CC0 Public Domain license.

Table 1: Dataset Overview

Total tweets	1,084,452
Unique users	165,177
Date range	1 November 2020 – 21 November 2021
Days covered	386
Mean daily tweets	2,809
Maximum daily tweets	83,866 (3 February 2021)
Mean likes per tweet	14.1
Mean retweets per tweet	8.8
Mean replies per tweet	0.6

To measure discourse displacement, I classify each tweet along two non-exclusive dimensions: *policy substance* and *celebrity/meta-commentary*. A tweet is classified as containing policy substance if its text matches keywords relating to the farmer movement’s substantive demands, including ‘MSP,’ ‘minimum support price,’ ‘farm law,’ ‘APMC,’ ‘contract farming,’ ‘Adani,’ ‘Ambani,’ ‘privatization,’ ‘Swaminathan,’ and ‘repeal.’ A tweet is classified as containing celebrity or meta-commentary if it matches terms including ‘Rihanna,’ ‘Greta,’ ‘Thunberg,’ ‘toolkit,’ ‘conspiracy,’ ‘foreign,’ ‘propaganda,’ ‘anti-national,’ ‘Khalistani,’ and ‘Bollywood.’ The full keyword list is provided in the supplemental materials. Across the full dataset, 11.5 percent of tweets contain policy substance and 10.6 percent contain celebrity or meta-commentary.

Table 2: Key Events in the Data Window

Date	Event	Type	Peak tweets	Peak multiple
26 Nov 2020	Dilli Chalo march	Mobilization	6,368	—
8 Dec 2020	Bharat Bandh	Coordinated	11,209	2.0×
26 Jan 2021	Republic Day tractor rally	Mobilization	15,368	10.4×
2 Feb 2021	Rihanna/Thunberg tweets	Celebrity	83,866	21.9×
28 Aug 2021	Karnal lathi charge	Violence	4,145	2.3×
3 Oct 2021	Lakhimpur Kheri	Violence	3,674	2.6×
19 Nov 2021	Farm law repeal	Concession	13,404	20.9×

Note: Peak multiple is the ratio of peak daily tweet count to the 14-day pre-event mean. The Dilli Chalo march was the movement’s inception, so no meaningful pre-event baseline exists.

Empirical Strategy

The primary empirical framework is local projection estimation (Jordà, 2005). For each outcome variable y_t and each shock event S , I estimate a separate regression at each horizon h :

$$y_{t+h} = \alpha_h + \beta_h \cdot \mathbf{1}[t = t_S] + \sum_{l=1}^7 \gamma_{h,l} \cdot y_{t-l} + \delta_h \cdot \text{DOW}_t + \varepsilon_{t+h} \quad (1)$$

where $\mathbf{1}[t = t_S]$ is an indicator for the event date, seven lags of the outcome control for serial dependence, and day-of-week fixed effects absorb weekly patterns. The coefficient β_h is the impulse response at horizon h . I estimate for horizons $h = -7, \dots, 0, \dots, 14$. Standard errors are computed using the Newey and West (1987) HAC estimator with bandwidth $\max(|h|, 1) + 1$. I report estimates for three outcomes: log daily tweet count, the celebrity/meta-commentary share, and the policy substance share, estimated separately for

four events: the Rihanna tweet, the Republic Day rally, Lakhimpur Kheri, and the farm law repeal.

I complement this with Granger causality tests within bivariate VAR frameworks ([Kulichkina et al., 2025](#)). The first system tests whether the daily celebrity/meta-commentary share temporally precedes the policy substance share. The second tests the direction of information flow between log daily Twitter volume and log daily TOI farmer-protest article count. Lag order is selected by AIC; stationarity is verified by Augmented Dickey-Fuller tests, with first-differencing applied where non-stationarity is detected.

Results

Volume and Content Displacement

Figure 1 presents the daily time series of tweet volume over the full study period. The Rihanna spike is visually dominant. Before the tweet, the conversation sustained a baseline of approximately 3,800 tweets per day. On the day of the tweet, volume jumped to 28,548; the following day it reached 83,866, a 22-fold increase. The decay was rapid, declining roughly exponentially to baseline by day 8. Other events produced much smaller volume responses: the Republic Day rally peaked at 15,368, the farm law repeal at 13,404, and Lakhimpur Kheri at 3,674.




Figure 1: Daily tweet volume, November 2020 to November 2021. Vertical dashed lines mark key events from Table 2. The dominant spike corresponds to the Rihanna tweet of 2 February 2021.

The volume story, however, is only the surface. Figure 2 reveals what happened to the *content* of the conversation. Before the tweet, the policy share averaged 12.2 percent and the celebrity share 4.5 percent. After the tweet, these positions inverted: the celebrity share spiked to over 45 percent while the policy share collapsed to approximately 5 percent. The two lines cross in the manner of a scissors diagram: at the moment of maximum visibility, the conversation’s substantive content was at its minimum. The inversion persisted for approximately four days before content shares returned toward pre-event levels.

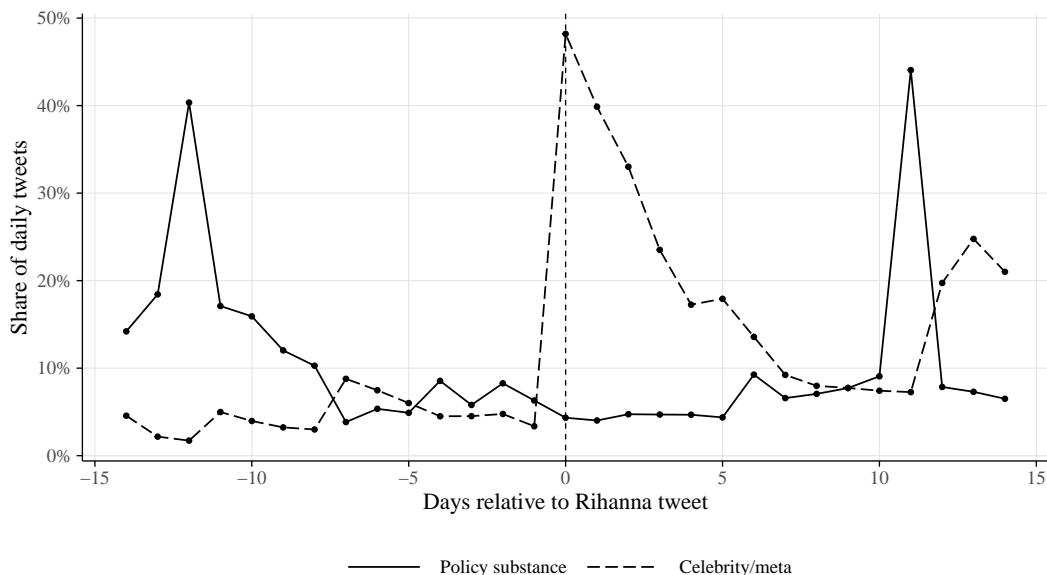


Figure 2: Content composition around the Rihanna tweet (2 February 2021). The solid line shows the daily share of tweets containing policy substance; the dashed line shows celebrity/meta-commentary. Before the tweet, policy content exceeds celebrity content. After the tweet, the lines cross: celebrity dominates while policy collapses.

The contrast with Lakhimpur Kheri is instructive. When farmers were killed on 3 October, the celebrity/meta-commentary share barely moved, rising from 1.9 to 3.4 percent. Domestic violence, even of an extreme and politically consequential nature, did not trigger discourse displacement. The attention trap is specifically a feature of celebrity-driven attention, not of attention shocks in general.

Local Projection Estimates

Figure 3 presents the local projection impulse responses for the Rihanna shock. The top panel shows the effect on log daily tweets: a sharp spike of approximately 2.5 log points at horizons 0 and 1, decaying to zero by horizon 8. Pre-event coefficients at negative horizons are tightly centred on zero, confirming the absence of pre-trends. The middle panel shows a 45-percentage-point spike in the celebrity/meta share, statistically significant at the 95 percent level. The bottom panel shows a statistically significant negative effect of approximately 10 percentage points on the policy substance share.




Figure 3: Local projection impulse responses for the Rihanna shock. Each panel plots $\hat{\beta}_h$ from Equation 1 with 95 percent Newey-West confidence intervals (dotted lines). Top: log daily tweets. Middle: celebrity/meta-commentary share. Bottom: policy substance share. Pre-event coefficients confirm the absence of pre-trends.

Figure 4 presents the cross-event comparison for the celebrity/meta-commentary share. The contrast is stark. For the Rihanna shock, the impulse response is large and statistically significant. For the Republic Day rally, the farm law repeal, and Lakhimpur Kheri, the impulse responses are indistinguishable from zero at all horizons. Discourse displacement is

specific to the celebrity channel.

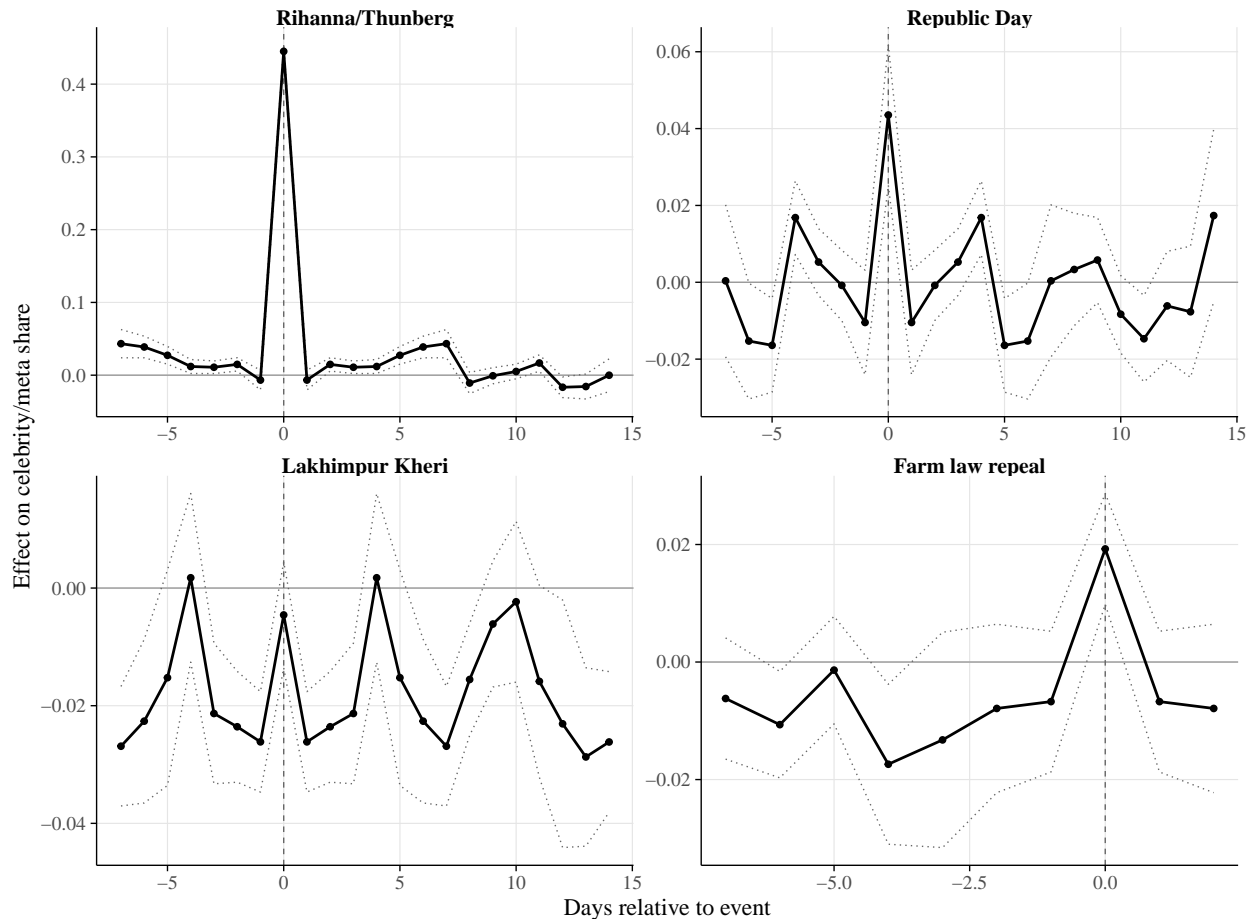


Figure 4: Local projection impulse responses for the celebrity/meta-commentary share across four events, with 95 percent Newey-West confidence intervals. Only the Rihanna shock produces a statistically significant positive response; the other three events produce responses indistinguishable from zero.

Granger Causality

Table 3 reports the Granger causality results. Celebrity/meta-commentary Granger-causes subsequent declines in policy substance ($F = 4.63$, $p = 0.032$), but policy substance does not Granger-cause changes in celebrity commentary ($F = 0.002$, $p = 0.963$). The displacement is unidirectional: celebrity content crowds out policy content, but policy content does not displace or prevent celebrity content. Separately, Twitter activity Granger-causes subsequent increases in TOI coverage ($F = 3.56$, $p = 0.002$), while newspaper coverage does not predict

Twitter activity ($F = 0.13$, $p = 0.993$). Social media led the mainstream media narrative. Because Twitter led the media, the displacement of policy content on Twitter may have propagated into newspaper framing of the protests.

Figure 5 presents VAR impulse response functions that confirm these dynamics. A shock to celebrity/meta-commentary produces a statistically significant negative response in policy content, peaking at approximately day 2 and dissipating by day 10 (top left). The reverse direction shows no significant effect (top right), consistent with the unidirectional Granger causality. A shock to Twitter volume produces a positive response in TOI coverage the following day (bottom left), while TOI shocks do not predict Twitter activity (bottom right).




Figure 5: VAR impulse response functions with 95 percent bootstrapped confidence intervals. Top row: celebrity/meta-commentary and policy substance system. A shock to celebrity content depresses policy content (left), but not the reverse (right). Bottom row: Twitter volume and TOI article count system. Twitter activity predicts subsequent newspaper coverage (left), but not the reverse (right).

Table 3: Granger Causality Tests

System	Direction	Lags	F -statistic	p -value
Policy \leftrightarrow Celebrity	Celebrity \rightarrow Policy	1	4.628	0.032
	Policy \rightarrow Celebrity	1	0.002	0.963
Twitter \leftrightarrow TOI	TOI \rightarrow Twitter	6	0.129	0.993
	Twitter \rightarrow TOI	6	3.563	0.002

Note: Granger causality F -tests from bivariate VAR models. Lag order selected by AIC. Bold indicates significance at the 5% level.

Mechanisms: Who Enters, What They Say, How Long They Stay

The volume surge was not an amplification of existing voices but an influx of new ones. In the 14 days before the tweet, 20,340 unique users participated. In the eight days after, that number nearly quadrupled to 79,885, of whom 70,112 (87.8 percent) had never previously tweeted about the protests. These new entrants accounted for 71.5 percent of all post-shock tweets. Figure 6 decomposes the daily tweet count into contributions from new and returning users. Before the event, returning users dominate. After the event, new users dwarf returning users at a ratio of roughly three to one. The existing community’s volume barely changes; what changes is the addition of a massive transient population.

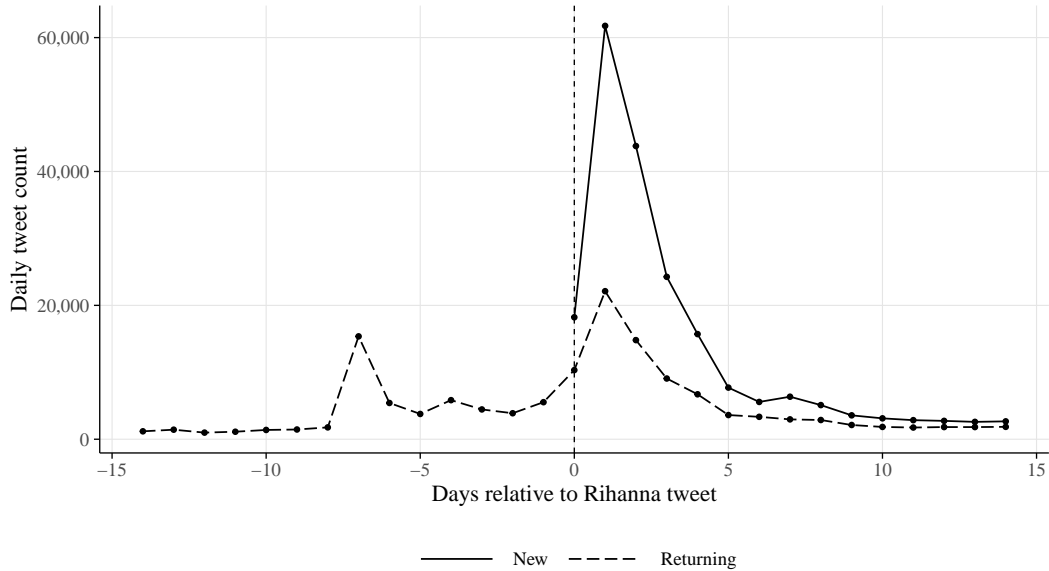


Figure 6: Daily tweet count around the Rihanna tweet, decomposed by user type. Solid line: users whose first appearance is on or after 2 February 2021. Dashed line: users who had tweeted previously. The surge is almost entirely attributable to new entrants.

The new entrants were not discussing the same subjects. Their tweets were 38.2 percent celebrity/meta-content and only 4.2 percent policy substance. They were having a different conversation while using the same hashtag. Table 4 shows that this pattern varies by entry event. Bharat Bandh entrants produced 8.2 percent policy content and only 5.7 percent celebrity content, and 33.9 percent were still active 30 days later. Rihanna entrants showed the inverse pattern, with only 12.5 percent remaining after 30 days and 65.6 percent tweeting on only a single day. Farm law repeal entrants were 50.4 percent policy content, confirming that when the precipitating event was a policy action, participants discussed policy. The content focus of new participants is determined by the nature of the event that drew them.

Table 4: Participant Quality by Entry Cohort

Entry event	Users	One-day	Active 30d	Policy %	Meta %
Rihanna tweet	54,161	65.6%	12.5%	4.2%	38.2%
Republic Day rally	6,349	75.0%	8.6%	3.8%	11.7%
Bharat Bandh (Dec)	4,760	55.0%	33.9%	8.2%	5.7%
Lakhimpur Kheri	1,291	78.2%	10.7%	6.0%	4.3%
Farm law repeal	5,225	98.7%	—	50.4%	4.2%

Note: Entry cohorts defined as users whose first tweet falls within a three-day window of each event. The farm law repeal cohort has no 30-day persistence measure because the dataset ends two days after the repeal.

User profile data reveals who these entrants were. Table 5 reports account-level characteristics by entry cohort. Three findings stand out. First, the Rihanna cohort is dramatically less likely to report an identifiably Indian location: only 16.0 percent, compared to 33.4 percent of core users, 31.2 percent of Bharat Bandh entrants, and 44.5 percent of Lakhimpur Kheri entrants. The celebrity shock drew participants from outside the Indian political sphere at a rate no other event approached. Second, only 4.4 percent of Rihanna entrants had political or movement-related keywords in their profile descriptions, compared to 12.2 percent of core users and 16.6 percent of Lakhimpur Kheri entrants. These were not politically engaged individuals participating in a new cause; they were apolitical users drawn by a celebrity signal. Third, the Rihanna entrants are not newly created accounts or bots: their median account age is 3,109 days (approximately 8.5 years), slightly older than the core community. They are established, moderate-follower users who had simply never engaged with Indian agricultural politics before and, in the vast majority of cases, never would again.

Table 5: User Profile Characteristics by Entry Cohort

	Core	Bharat Bandh	Rihanna	Lakhimpur	Repeal
Users	52,957	13,820	56,371	1,291	5,225
Median followers	164	70	131	79	87
Median account age (days)	2,113	449	3,109	954	1,387
India location (%)	33.4	31.2	16.0	44.5	49.7
International location (%)	6.8	7.7	5.8	7.7	5.9
Political/movement bio (%)	12.2	11.9	4.4	16.6	13.0
One-day participants (%)	21.4	29.1	63.1	78.2	98.7

Note: Core users entered the conversation before December 2020. Location classified by keyword matching on self-reported location field; users not matching Indian or international patterns are classified as unknown. Political/movement bio indicates presence of political, activist, or agricultural keywords in the user’s profile description.

Figure 7 presents the mention network before and after the Rihanna tweet. Panel (a) shows a relatively tight cluster organized around established protest actors: the most-mentioned accounts are narendramodi, DelhiPolice, UNHumanRights, and Kisanektamorcha. Panel (b) shows the network in the three days following the tweet. The core community remains visible but is surrounded by a dispersed cloud of new entrants. The most-mentioned accounts have shifted to Rihanna, GretaThunberg, akshaykumar, and KanganaTeam. The network did not merely expand; its centre of gravity shifted from policy institutions to celebrities.

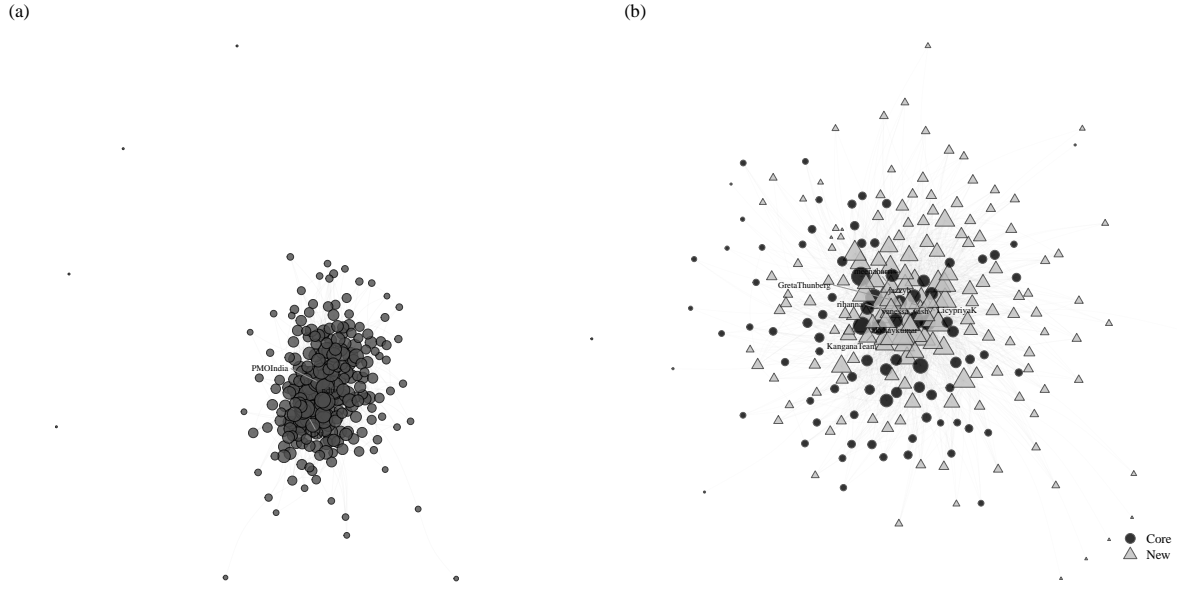


Figure 7: Mention network before and after the Rihanna tweet. Panel (a): seven days before, organized around policy actors. Panel (b): three days after, with the core community (filled circles) surrounded by new entrants (triangles) and focal accounts shifted to celebrities. Node size proportional to degree; Fruchterman-Reingold layout.

Discussion

The central finding of this paper is that international celebrity attention to a domestic protest movement involves a trade-off that existing scholarship has not measured. The visibility benefits are real: Rihanna’s tweet generated a 22-fold increase in conversation volume, drew global media coverage, and placed the farmer protests on the international agenda. The farmer protest movement ultimately succeeded, and international visibility may well have contributed to the political calculus that led to the farm law repeal in November 2021. But the attention came with a cost that is invisible in the volume metrics scholars and activists typically use: a 45-percentage-point increase in celebrity/meta-commentary, a 10-percentage-point decrease in policy substance, and the entry of 70,000 participants, disproportionately non-Indian and apolitical in profile, who talked about Rihanna rather than MSP and departed within days. The cost was not merely aesthetic. It was strategically

exploitable.

Several implications follow. The first concerns the strategic calculus of social movements. The evidence suggests that celebrity attention windfalls come with costs that are not immediately apparent. The 70,000 users who entered the conversation created the appearance of massive support, yet it was shallow, expressed through engagement with celebrity dimensions rather than substantive demands, and fleeting, with two-thirds departing within 24 hours. For a movement that needed sustained political pressure on specific policy questions, the net value of the celebrity attention surge is at best ambiguous.

The second implication concerns the relationship between the attention trap and authoritarian populist governance. India under the BJP has experienced what scholars have characterized as a steady erosion of democratic norms, including restrictions on press freedom, the criminalization of dissent through sedition and anti-terror statutes, and the strategic deployment of social media for propaganda purposes (Ganguly, 2019; Jaffrelot, 2021; Varshney, 2019). Chacko (2025) documents how the BJP systematically deployed conspiracy theories to delegitimize the farmer protests, framing them as the work of Khalistani separatists, foreign NGOs, and international provocateurs. The Rihanna episode provided the BJP with its most potent ammunition for this framing. It was the celebrity attention itself, and the discourse displacement it produced, that shifted the conversational terrain from the relatively unfavourable ground of agricultural policy (where the farmers had expertise, evidence, and broad public sympathy) to the much more favourable ground of nationalist sentiment and foreign interference (where the BJP could deploy its well-honed rhetorical machinery). The attention trap, in this sense, is not merely an incidental cost of celebrity involvement; it is a vulnerability that authoritarian populist regimes can exploit. Garg (2025) argues that what sustains protest movements over time is not visibility per se but ‘media field embeddedness,’ a concept that captures a movement’s position within the broader network of activist media and organizational infrastructure. Celebrity attention produces visibility without embeddedness, and the results of this paper suggest that the distinction matters.

The third implication concerns the direction of information flow between social media and mainstream media. The finding that Twitter Granger-causes Times of India coverage (but not the reverse) means that the discourse displacement that occurred on Twitter did not stay on Twitter. If mainstream media took cues from the social media conversation, then the shift from policy to celebrity content on Twitter may have propagated into newspaper framing of the protests. This is consistent with accounts from the period describing how the mainstream media’s coverage shifted from the farm laws themselves to the Rihanna controversy and the toolkit conspiracy in the days following the tweet.

Several limitations should be acknowledged. The dataset was collected using pro-farmer hashtags, which means I observe the protest conversation but not the full two-sided discourse including the pro-government counter-campaign. Future work with access to the complete Twitter firehose could examine how the attention trap interacts with the government’s coordinated counter-mobilization. The content classification relies on keyword matching, which captures explicit mentions of policy terms and celebrity names but may miss more subtle forms of substantive engagement or meta-commentary. More sophisticated approaches using large language models for content classification could improve measurement precision. Finally, the analysis is limited to a single platform (Twitter) and a single protest movement. Whether the attention trap generalizes to other platforms, other countries, and other types of movements is an empirical question that warrants further investigation.

Despite these limitations, the core finding is robust across multiple empirical approaches: local projections, Granger causality, user-level decomposition, and network analysis all converge on the same conclusion. International celebrity attention to domestic protest movements can be a trap. It generates visibility, but visibility is not the same as substance. The 70,000 users who flooded into the farmer protest conversation after Rihanna’s tweet were, in the most literal sense, not talking about farming. They were talking about Rihanna. And in doing so, they temporarily displaced the voices that were.

Statements and Declarations

Ethical considerations

This study uses publicly available social media data collected under a CC0 Public Domain license and publicly available newspaper articles. No human subjects were recruited or contacted. Ethics approval was not required.

Disclosure statement

The author reports there are no competing interests to declare.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Data availability

The Twitter dataset is publicly available on Kaggle as the Farmers Protest Tweets Dataset (CC0 Public Domain license). The Times of India dataset was collected by the author. Replication code is available upon reasonable request.

Use of generative AI

Claude Opus (Anthropic, version 4) was used during the preparation of this manuscript for coding assistance with R scripts for data analysis and visualization, and for grammatical checks on the manuscript text. The author reviewed and edited all outputs and takes full responsibility for the content of the publication.

References

- Barberá, P., Wang, N., Bonneau, R., Jost, J. T., Nagler, J., Tucker, J., and González-Bailón, S. (2015). The critical periphery in the growth of social protests. *PloS one*, 10(11):e0143611.
- Behl, N. (2022). India's farmers' protest: an inclusive vision of indian democracy. *American Political Science Review*, 116(3):1141–1146.
- Benford, R. D. and Snow, D. A. (2000). Framing processes and social movements: An overview and assessment. *Annual review of sociology*, 26(2000):611–639.
- Bennett, W. L. and Segerberg, A. (2023). The logic of connective action: Digital media and the personalization of contentious politics. In *Handbook of digital politics*, pages 287–338. Edward Elgar Publishing.
- Bourdieu, P. (2018). The forms of capital. In *The sociology of economic life*, pages 78–92. Routledge.
- Brockington, D. (2014). *Celebrity advocacy and international development*. Routledge.
- Chacko, P. (2025). Conspiracy theories and india's transnational authoritarian populism: Ngos, khalistanis and soros. *The British Journal of Politics and International Relations*, page 13691481251393570.
- Driessens, O. (2013). Celebrity capital: redefining celebrity using field theory. *Theory and society*, 42(5):543–560.
- Entman, R. M. et al. (1993). Framing: Towards clarification of a fractured paradigm. *McQuail's reader in mass communication theory*, 390:397.
- Ganguly, S. (2019). India under modi: Threats to pluralism. *Journal of democracy*, 30(1):83–90.

- Garg, A. (2025). *Seeds of Dissent: Organizational Capacity, Media Embeddedness, and the Divergent Trajectories of the Indian Farmers' Movement*. University of Chicago.
- González-Bailón, S., Borge-Holthoefer, J., Rivero, A., and Moreno, Y. (2011). The dynamics of protest recruitment through an online network. *Scientific reports*, 1(1):1–7.
- Jaffrelot, C. (2021). *Modi's India: Hindu nationalism and the rise of ethnic democracy*. Princeton University Press.
- Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review*, 95(1):161–182.
- Kristofferson, K., White, K., and Peloza, J. (2014). The nature of slacktivism: How the social observability of an initial act of token support affects subsequent prosocial action. *Journal of Consumer Research*, 40(6):1149–1166.
- Kulichkina, A., Righetti, N., and Waldherr, A. (2025). Protest and repression on social media: Pro-navalny and pro-government mobilization dynamics and coordination patterns on russian twitter. *new media & society*, 27(9):5433–5454.
- McCombs, M. E. and Shaw, D. L. (1972). The agenda-setting function of mass media. *Public opinion quarterly*, 36(2):176–187.
- Mishra, D., Akbar, S. Z., Arya, A., Dash, S., Grover, R., and Pal, J. (2021). Rihanna versus bollywood: twitter influencers and the indian farmers' protest. *arXiv preprint arXiv:2102.04031*.
- Morozov, E. (2012). *The net delusion: The dark side of Internet freedom*. PublicAffairs.
- Narayanan, S. (2021). Understanding farmer protests in india. *Journal of Academics Stand Against Poverty*, 1(1):133–140.
- Newey, W. K. and West, K. D. (1987). Hypothesis testing with efficient method of moments estimation. *International Economic Review*, pages 777–787.

- Simon, H. A. (1996). Designing organizations for an information-rich world. *International Library of Critical Writings in Economics*, 70:187–202.
- Steinert-Threlkeld, Z. C. (2017). Spontaneous collective action: Peripheral mobilization during the arab spring. *American Political Science Review*, 111(2):379–403.
- Thrall, A. T., Lollo-Fakhreddine, J., Berent, J., Donnelly, L., Herrin, W., Paquette, Z., Wenglinski, R., and Wyatt, A. (2008). Star power: Celebrity advocacy and the evolution of the public sphere. *The international journal of press/politics*, 13(4):362–385.
- Tufekci, Z. (2017). *Twitter and tear gas: The power and fragility of networked protest*. Yale university press.
- Varshney, A. (2019). Modi consolidates power: Electoral vibrancy, mounting liberal deficits. *Journal of Democracy*, 30(4):63–77.
- Wheeler, M. (2013). *Celebrity politics*. Polity.
- Wu, T. (2017). *The attention merchants: The epic scramble to get inside our heads*. Vintage.