Indonesian Journal of Educational Science and Technology (Nurture)
Vol. 5, No. 2 2026: 71-92

Grok, State Capital and Conspiracy Narratives: A Game-Theoretic Analysis of Indonesia's Free Meal Discourse

Agus Widiyarta[1], Catur Suratnoaji[2], Indri Anjar Kartika Sari[3*]
[1,2]Faculty of Social and Political Sciences, Universitas Pembangunan Nasional "Veteran" Jawa Timur, Surabaya, Indonesia
[3]Master's Programme in Communication Science, Universitas Pembangunan Nasional "Veteran" Jawa Timur, Surabaya, Indonesia
Corresponding Author: Indri Anjar Kartika Sari
25067020011@student.upnjatim.ac.id

A R T I C L E I N F O

Keywords: Grok, Platform Capitalism, Conspiracy Narratives, Game Theory, Free Nutritious Meal

Received : 15 March
Revised : 20 April
Accepted: 30 May

©2026 Widiyarta, Suratnoaji, Sari: This is an open-access article distributed under the terms of the Creative Commons Atribusi 4.0 Internasional.

A B S T R A C T

Indonesia's Free Nutritious Meal programme (MBG), now the largest item in the national budget, is debated on X, where Grok answers citizens. This study examines how state capital, platform capital and a sceptical public met in MBG discourse. We analysed 5,272 posts from March to May 2026, a mention network, 93 Grok posts and 94 coded suspicion posts, using a term co-occurrence map, a retrospective corroboration check and a game-theoretic model. Grok mostly repeated official rationales and endorsed the agency head weeks before he became a corruption suspect. Documented accountability claims outnumbered conspiracy narratives, focusing on budgets, not foreign powers. Opacity, generalised suspicion and oracle endorsement form an equilibrium without conspiracy that only disclosure of procurement records can break

DOI: https://doi.org/10.55927/nurture.v5i2.7
E-ISSN: 2985-7287
https://journalnurture.my.id/index.php/nurture/index

INTRODUCTION

Few public programmes anywhere have grown as fast as Indonesia's Free Nutritious Meal programme (MBG). Launched in January 2025 with IDR 71 trillion, it received a 2026 ceiling of IDR 335 trillion, later revised to IDR 268 trillion, and its implementing agency, the National Nutrition Agency (Badan Gizi Nasional, BGN), has been allocated IDR 240.2 trillion for 2027, making it the largest central-government spending unit, ahead of the Ministry of Defence (IDR 189 trillion) and the National Police (IDR 139.2 trillion) (Investortrust, 2026a; CNBC Indonesia, 2026; Gebrak, 2026). It feeds tens of millions of children and pregnant women through more than 27,000 local kitchens, the Satuan Pelayanan Pemenuhan Gizi (SPPG).

Scale of this kind turns nutrition policy into political economy: every kitchen needs a partner, a foundation and suppliers, each a market created by the state. By mid-2026 the Corruption Eradication Commission (KPK), the Ombudsman, Transparency International Indonesia and Indonesia Corruption Watch had warned about conflicts of interest in appointing kitchen partners and about the involvement of the armed forces and police (Bisnis.com, 2026a; Ombudsman RI, n.d.; Transparency International Indonesia, 2025; JawaPos, 2026a). On 2 June 2026 the President dismissed the founding head of BGN; the next day the Attorney General's Office named him and two former deputies as suspects in an alleged corruption case involving affiliated foundations and the procurement of electric motorcycles, shoes, tablets and televisions (Hukumonline, 2026; Investortrust, 2026b). The case was still under investigation at the time of writing, and all three are presumed innocent until a court rules otherwise.

The programme is also geopolitical: the Agreement on Reciprocal Trade signed with the United States on 19 February 2026 committed Indonesia to remove tariffs on nearly all US goods and to buy US wheat, soybeans and beef in large volumes (Setkab, 2026; Agrolatam, 2026), while the President has warned against dependence on imports from countries caught up in geopolitical tension (Kontan, 2026a).

Earlier work on the same corpus found MBG criticism on X often sarcastic, centred on budget and politics, and circulating in a sparse, one-directional network addressed upwards to the President, with the implementing agency absent (Sari et al., 2026a); a companion study found that a deployed lexical early warning rule saturated on busy days and fired on denials as readily as on reports (Sari et al., 2026b). Two observations remained unexplained: the most active mentioning account was Grok, the AI chatbot X embeds in public conversations, and a vocabulary of theft, rent, hidden agendas and foreign capital ran through the corpus, sometimes as documented criticism and sometimes closer to conspiracy narrative.

This article reads MBG discourse as the meeting point of two capitalisms and a sceptical public: the state-capitalist kitchen, a fiscal machine that creates markets, distributes rents and claims legitimacy through nutrition and sovereignty (Bremmer, 2010; Kurlantzick, 2016; Alami & Dixon, 2020), and the platform-capitalist oracle, a privately owned, foreign-controlled answer engine that citizens summon to adjudicate truth in public (Srnicek, 2017; van Dijck et al.,

2018). It asks what Grok did in MBG discourse and with what reach; how the vocabulary of suspicion, rent and geopolitics appeared and how much of it was documented accountability rather than conspiracy narrative; which suspicions voiced between March and May 2026 were later taken up by legal or official action; and what an early warning system sees of all this.

The article contributes a triangular framework of kitchen, oracle and crowd, with the early warning system as arbitrating alarm, in which the line between warranted suspicion and conspiracy narrative is a matter of evidence over time rather than vocabulary; to our knowledge, the first analysis of a generative-AI chatbot as a network actor in a Global South policy controversy and the first retrospective corroboration analysis of suspicion in Indonesian policy discourse; and three additions to early warning design: a governance-discontent tier, oracle-reply monitoring and a corroboration ledger. "Conspiracy" in the title is analytical: we allege no conspiracy but study how talk of it, and the conditions that make it plausible, appear in public discourse.

LITERATURE REVIEW

State Capitalism and the Food State

"State capitalism" names arrangements in which the state owns, directs or creates markets rather than merely regulating them (Bremmer, 2010; Kurlantzick, 2016); Alami and Dixon (2020) recommend studying its concrete relations between state power and capital accumulation rather than a single model. MBG is such a relation: a state-created market in which private partners, foundations and suppliers earn returns from public money, coordinated by an agency answering to the President. Such markets in Indonesia have long been shaped by official–business alliances (Robison & Hadiz, 2004), within a "restrained" state capitalism of state enterprises and directed credit under fiscal limits (Kim, 2021). Danantara and the Agrinas food companies (InfoPublik, 2025; Bisnis.com, 2025a) extend state market-making into food, the 2026 and 2027 ceiling cuts show the fiscal constraint, and the programme sits in a food regime binding national policy to global commodity chains (McMichael, 2009).

Platform Capitalisma and the AI Oracle

Platform capitalism describes firms that provide the infrastructure on which others interact and extract value from that interaction (Srnicek, 2017), shaping what counts as public conversation (Gillespie, 2018; van Dijck et al., 2018). Embedded chatbots add a layer: when an X user writes "@grok is this true?" beneath a post, the platform's own model answers publicly in the thread. We call this role the *oracle*: an answer engine whose authority rests on speed, apparent neutrality and the promise of synthesis, but whose outputs are probabilistic text without guaranteed grounding (Crawford, 2021). Grok's record is sobering: the Digital Forensic Research Lab found inaccurate and contradictory fact-checks during the Israel–Iran war (Arab News, 2025), and a 2026 study of 766 viral World Cup claims found that Grok answered fewer than 40% and that a third or more of its fact-checks were wrong or only partly right (Maldita.es, 2026).

The oracle is also owned: Grok is developed by xAI and integrated into X, both controlled by Elon Musk, and its system prompt has repeatedly been altered on political topics (see Results). It is thus foreign infrastructure inside the domestic public sphere, to which debates about digital sovereignty apply directly (Pohle & Thiel, 2020).

Conspiracy Narratives, Warranted Suspicion and Opacity

Conspiracy theories explain events as the secret work of powerful actors against the public interest; they appeal when people need certainty or control and flourish under uncertainty and distrust (Douglas et al., 2017; Uscinski & Parent, 2014). Philosophers warn, however, against treating every suspicion of collusion as irrational. Keeley (1999) notes that real conspiracies exist and that the problem with "unwarranted" theories lies in their structure, such as resistance to disconfirmation; Sunstein and Vermeule (2009) define conspiracy theories partly by lack of evidence, so the boundary moves as evidence arrives; and cultural theorists read them as a response to real concentrations of power and secrecy (Fenster, 2008), "the poor person's cognitive mapping" of a system too large and opaque to see whole (Jameson, 1988).

We therefore distinguish operationally. A *documented accountability claim* names an actor, an action and usually a source such as an audit body, a court or a civil-society monitor. A *generalised rent accusation* asserts, often through wordplay, that the programme exists to enrich insiders without naming a specific act. A *hidden-agenda narrative* attributes the programme to a concealed plan by a powerful group, foreign or domestic, that public evidence does not support. Whether a claim is warranted is settled not by its wording but, if at all, by what happens afterwards (an investigation, a denial, a verdict or silence), which is why the corroboration analysis looks forward from the study window.

Early Warning and the Framework

Early warning from social media treats posts as signals of events (Sakaki et al., 2010), but such signals are distorted by query design, platform behaviour and media attention (Lazer et al., 2014). For a contested programme an early warning system therefore does more than detect; it arbitrates, deciding which posts deserve a human reader now. In the framework (Figure 1), the state-capitalist *kitchen* claims legitimacy through nutrition and sovereignty while creating rents; the platform-capitalist *oracle* claims it through speed, neutrality and "truth-seeking" while serving its owner's commercial and political interests; the *crowd* suspects both; and the *alarm* must decide what to escalate.

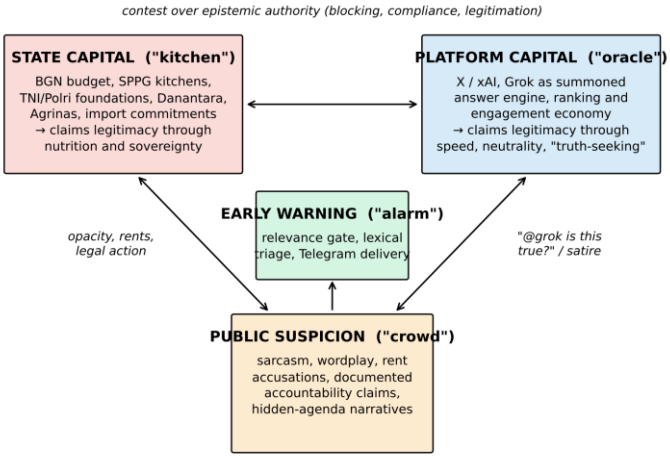

Figure 1. Analytical framework: the state-capitalist kitchen, the platform-capitalist oracle and the suspicious crowd, with the early warning system as the alarm that arbitrates between them.

Three propositions guided the analysis. P1: an AI oracle embedded in a foreign platform will act less as a verifier than as a general-purpose interlocutor, and when it addresses the policy it will tend to reproduce official framing. P2: suspicion vocabulary will contain more documented accountability claims and generalised rent accusations than hidden-agenda narratives, and a substantial share of specific suspicions will later be taken up by official action. P3: a lexical early warning rule built around medical and hygiene incidents will treat most suspicion posts as Safe and so will not register governance discontent as a risk.

The MBG Case in 2026

Each SPPG is run by a partner under the umbrella of a foundation and receives public funds per portion served. In May 2026 the Attorney General's Office noted that the police ran about 1,376 kitchens and the armed forces 452 of roughly 27,000 (Republika, 2026a), and a civil-society review of 102 partner foundations found 28 (27.45%) formally linked to political parties (JawaPos, 2026a). The KPK warned in April 2026 of a high potential for conflicts of interest because partner selection was centralised and procedures unclear (Bisnis.com, 2026a), and in October BGN's new leadership reported that some foundations levied kitchen partners for every tray served (CNN Indonesia, 2026). Because the posts analysed predate the June investigation and these disclosures, we can ask, without hindsight in the data, which suspicions anticipated later official action.

On 10 January 2026 the Ministry of Communication and Digital Affairs (Komdigi) temporarily blocked Grok, the first country in the world to do so, because it was being used to manipulate photographs of real people into sexual images (detikInet, 2026a; MetroTV, 2026). Access was restored at the beginning of February after X Corp gave a written commitment to comply with its obligations as a registered electronic system operator (detikInet, 2026b; Gigazine, 2026), three weeks before the study window opened. The episode established that the state regarded Grok as a foreign service subject to national sovereignty and that X accepted that framing in order to keep operating.

METHODOLOGY

Corpus, Sub-Corpora and Lexicons

The analysis uses the corpus and network of the companion studies (Sari et al., 2026a, 2026b). Public posts were collected with NodeXL Pro using the query "Makan Bergizi Gratis" OR "MBG" OR "Makan Siang Gratis" OR #MakanBergiziGratis OR #MBG, restricted to Indonesian and excluding pure reposts, for 2 March to 24 May 2026: 5,272 posts by 2,421 author handles, each with engagement counts and a model-assigned emotion label. The mention network contains 971 accounts linked by 666 unique directed ties. Period 1 (2 March–30 April) and Period 2 (1–24 May) are distinguished, the latter by a sharp rise in volume. Two companion findings condition what follows: a blind human coding of 150 random posts found that only 49 (32.7%) concerned the programme, because the acronym MBG is shared with unrelated products, fandoms and advertising, and the IndoBERT emotion labels agreed with human coding at chance level (Cohen's $\kappa = 0.05$) (Sari et al., 2026a). Emotion labels are therefore reported only to show how an automated pipeline reads the material, and suspicion analyses are restricted to the 1,899 posts containing at least one programme keyword (MBG, "makan bergizi", BGN, SPPG, ompreng or the name of the former agency head).

All 93 posts by @grok were extracted; 37 were exact duplicates, leaving 56 unique texts posted between 2 March and 22 May 2026. No other post in the corpus contains the handle @grok, so we analyse what the oracle said and how far it travelled, not the prompts that summoned it. The third author (I.A.K.S.) coded each unique post into five inductive categories: legitimation or policy explanation, verification, fiscal arithmetic, satire and image play, and off-topic. We built five transparent Indonesian lexicons from a first reading and public commentary (corruption and rent; hidden agenda; foreign and geopolitical; security apparatus; platform and AI), matched as whole lower-cased words; they capture explicit vocabulary, not implicit insinuation. The 94 programme-keyword posts matching at least one of the first three lexicons formed the suspicion sub-corpus (59 unique posts), which the same coder assigned to seven categories derived from the review: generalised rent accusation, documented accountability claim, geopolitical-economic critique, hidden-agenda narrative, pro-government counter-narrative, business-opportunity talk and other. Supplementary analyses (a term co-occurrence map in the manner of VOSviewer (van Eck & Waltman, 2010), keyword counts of actors and fiscal frames, and an engagement comparison) use the unique programme-keyword posts other than Grok's.

Corroboration, Early Warning Replay and Ethics

For each specific suspicion theme we searched news and official statements from 25 May to 15 November 2026 for subsequent state or oversight action, classifying it as *taken up by legal action*, *confirmed by official admission*, *officially denied* or *unresolved*; taken up by legal action does not mean proven, since suspects are presumed innocent. To ask what an early warning system would see, we applied the deployed triage rule of the project's Telegram alert job (Sari et al., 2026b): keracunan, mual, muntah or diare yield Critical; otherwise basi,

ulat, lalat or susu yield Danger; otherwise korupsi, anggaran, SPPG or ompreng yield Watch; otherwise Safe. Our re-implementation reproduces the published distribution exactly (4,863 Safe, 263 Watch, 46 Danger and 100 Critical posts).

The third author designed the early warning bot evaluated here and collected the corpus for her thesis, an interest we counter by reporting its blind spots prominently and reproducing its published replay exactly before extending it. Only public posts were analysed; quotations are paraphrased or shortened so they cannot be traced, only public institutions, officials in their official capacity and @grok are named, and claims about named persons are limited to what official bodies and reputable news organisations have reported.

RESEARCH RESULT

The Oracle: What Grok Did

Grok was the most active mentioning account (out-degree 42) and hub of the second-largest community (43 accounts) after the President's (Sari et al., 2026a), yet of its 56 unique posts 40 (71.4%) were off-topic answers that matched a keyword, such as recipes, translations and jokes, and 16 (28.6%) concerned the programme (Table 1). Reach was negligible: 8 likes, 4 replies, no reposts and 8,389 views across 93 posts, with 42 accounts reached directly but only 2 at the second step.

Table 1. Grok's posts in the corpus by category.

Category	Unique posts	% of unique	All posts (incl. duplicates)	Views
Off-topic	40	71.4	70	4,793
Legitimation / policy explanation	7	12.5	9	765
Satire and image play	4	7.1	7	2,381
Verification	3	5.4	3	342
Fiscal arithmetic	2	3.6	4	108
Total	56	100.0	93	8,389

Note: across all 93 posts Grok received 8 likes, 4 replies and no reposts. Model emotion labels assigned "disgust" to 73 of the 93 posts.

Legitimation and policy explanation was the largest programme group (7). Asked whether money should go to parents instead of meals, Grok said cash might be spent on cigarettes or sugar, and elsewhere backed a hybrid with an unattributed claim that "World Bank data" showed cash 20–30% more efficient. On 22 April 2026 it called the then agency head's leadership "a good step for long-term human-capital investment", adding that "blind praise is not my style", and on 13 May described him as an apt example of a figure who could support small businesses; three weeks later he was dismissed and named a suspect. Grok could not have known of an unannounced investigation; the point is structural: an oracle synthesising public text reproduces officials' standing at the moment of asking and lends it machine neutrality. *Verification* posts confirmed a ministry housing plan for kitchen workers and dismissed a viral claim as "just a joke". *Fiscal arithmetic* inherited its assumptions: in March Grok cited "15 million children", far below the agency's reported 62 million beneficiaries by end-April

(ANTARA, 2026a). Grok also *co-produced satire*, calling the programme a "weapon of mass destruction" because eaters "fainted"; these four posts drew 2,381 views (28% of Grok's views), more than the legitimation posts. P1 is supported in part: Grok was mainly a general-purpose interlocutor that, on the programme, more often reproduced official rationales than checked claims, but also amplified mockery.

Who Writes the Oracle's Prompt, and Why

An answer from Grok is assembled from four layers, each controlled by a different party with a different motive (Figure 2). The first is the model, controlled by xAI, with a stated aim of "maximal truth-seeking" and incentives of engagement, subscriptions and training data. The second is the *system prompt* of the @grok reply bot, public since May 2025 (xAI, 2025), whose history shows how an oracle can be steered: in February 2025 it briefly told the model to ignore sources accusing Musk and Donald Trump of misinformation, blamed on an employee acting without approval (Fortune, 2025a); in May 2025 an "unauthorised modification" inserted claims about "white genocide" in South Africa (TechCrunch, 2025a); in July 2025 instructions to treat media as biased and not shy from politically incorrect claims preceded antisemitic replies and a suspension (Fortune, 2025b); and Grok 4 was then found to search its owner's posts before answering contested questions, until xAI barred answers stemming from Musk's or the company's beliefs (TechCrunch, 2025b; The Decoder, 2025). Each change was reversed under public pressure.

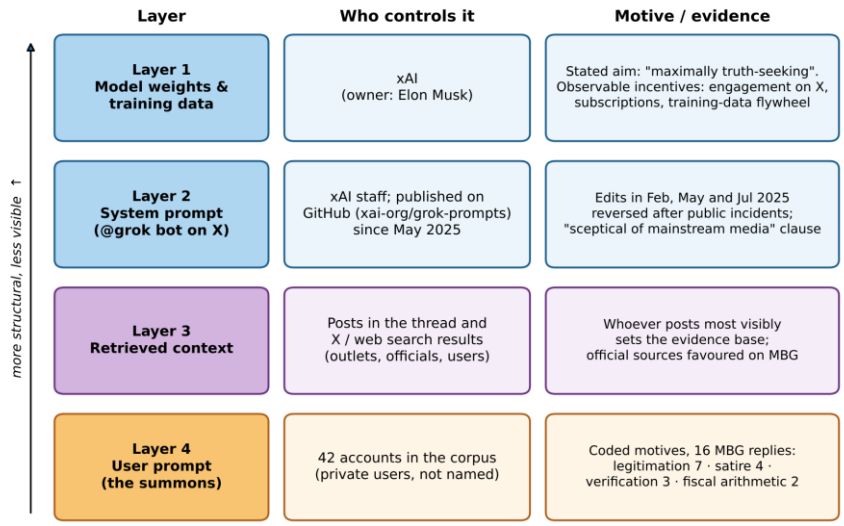

Figure 2. Four layers that shape an answer from @grok, the party that controls each layer, and the motive stated or observed. Layer 4 is coded from the 16 programme-related replies in the corpus.

The third layer is *retrieved context*: on the programme questions in our corpus the visible sources were official statements and national news, the most plausible reason why the legitimation answers reproduced official rationales, with no hidden instruction needed. The fourth is the *user prompt*: of the 16 programme-related replies, seven were summoned to explain or justify the policy, four for

jokes or images, three to verify a claim or video and two to compute sums of public money; the 42 accounts answered were all ordinary users, only two of them connected to anyone else in the network. Placing the whole answer in one hand overstates the company's control over any reply about school meals; treating the oracle as neutral understates its demonstrated capacity to steer what millions read.

The Crowd: Suspicion and Geopolitics Vocabulary

Among the 1,899 programme-keyword posts, 94 (4.9%) used the vocabulary of corruption and rent, hidden agendas or foreign relations. Corruption and rent vocabulary dominated (67 posts), led by "korupsi" (38 occurrences), "maling" (thief, 15) and "cuan" (profit, 11); hidden-agenda vocabulary was rare (8 posts) and foreign vocabulary appeared in 26. Suspicion rose with attention, not faster: 74 of the 94 posts appeared in the last three weeks, when overall volume also surged (Figure 3), and their share stayed near one in twenty (5.5% of 348 posts in Period 1; 4.9% of 1,546 in Period 2) (3.5–6.5% under stricter definitions).

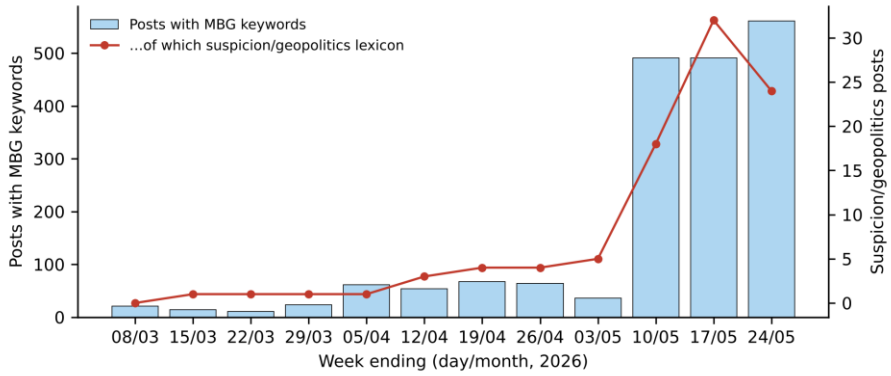

Figure 3. Weekly posts containing programme keywords (bars, left axis) and those using suspicion or geopolitical vocabulary (line, right axis), 2 March–24 May 2026.

Table 2 gives the coded composition of the 59 unique suspicion posts. *Generalised rent accusations* were the largest category (15) and relied on wordplay recoding the programme's acronyms: MBG as "Maling Berkedok Gizi" (thieves disguised as nutrition) or "Makanan Beracun Gratis" (free poisonous food), and SPPG as "Satuan Penjilat Prabowo Gibran" (unit of Prabowo–Gibran sycophants). *Documented accountability claims* (10) cited institutions or data: the KPK's list of eight corruption risks (20 April), the procurement of more than 25,000 electric motorcycles (16 April), an ICW report to the KPK on halal certification procurement (8 May), the President's instruction to investigate corruption gaps (11 May) and a prediction that the agency head could become a suspect (29 April). *Geopolitical-economic critiques* (8) linked the programme to the rupiah, US soybean dependence and the dairy market. *Hidden-agenda narratives* were rare (3): one presented a US philanthropic foundation's support for school meals as a possible "new threat", one claimed a deliberate plot to destroy the country, and one asked what was planned for 2030; all show Keeley's (1999) structure of a hidden intentional agent and an explanation absorbing any counter-evidence. Nine *pro-government counter-narratives* accused critics of

serving "the mafia" or dismissed a claim of IDR 4 trillion in corruption as a hoax, so the accusation of conspiracy runs both ways. P2 is supported: documented claims and generalised rent accusations together (25) outnumbered hidden-agenda narratives (3) by more than eight to one.

Table 2. Hand-coded categories of the 59 unique suspicion posts and their early warning levels.

Category	Posts	Safe	Watch	Danger	Critical
Generalised rent accusation	15	9	6	0	0
Documented accountability claim	10	1	9	0	0
Geopolitical-economic critique	8	5	2	1	0
Hidden-agenda narrative	3	1	2	0	0
Pro-government counter-narrative	9	5	4	0	0
Business-opportunity talk	3	2	1	0	0
Other	11	7	2	0	2
Total	59	30	26	1	2

Suspicion posts were seen more than other programme posts (median 348 views against 93; Mann–Whitney $p < 0.001$) and more often reposted (38% against 24%), but among the 15 accounts that wrote both kinds the gap narrows (296 against 218), so much of the difference lies in who writes suspicion. The suspicion is also fiscal and domestic: the state budget appeared in 83 unique programme posts, the local economy in 57, prices in 56, cuts elsewhere in 46 and the rupiah in 31, against 21 foreign references, mostly Japan as a benchmark for hygienic kitchens (10); China and the platform's owner did not appear. In the term co-occurrence map (64 terms, four clusters), "korupsi" and "maling" attach to "anggaran" (budget) and "uang" (money), not food or children, while a May 2026 meme reading MBG as "Mas Bahlil Ganteng" made up 18% of unique keyword posts in the last week, counted by a keyword monitor as programme talk.

What the Alarm Saw

Applied to the 94 suspicion posts, the deployed rule returned 46 Safe, 42 Watch, 3 Danger and 3 Critical; half the 59 unique posts (30) were Safe (Table 2). Watch was triggered mainly by "korupsi", so documented claims, which tended to use it, were mostly Watch (9 of 10), whereas generalised rent accusations, which preferred "maling" and "cuan", were mostly Safe (9 of 15); none of the acronym recodings contains a Watch-tier word. The three Critical posts were triggered by poisoning vocabulary, not suspicion. The alarm's view of governance discontent thus hangs on a single word, supporting P3. The IndoBERT model labelled 66 of the 94 suspicion posts as disgust, 20 as interest and none as anger: a public angry at no one, the opposite of what the posts say.

Retrospective corroboration

Table 3 follows the eight specific suspicion themes forward. Three were taken up by legal action within six months (one also confirmed by the agency's own admission), one was confirmed as policy rather than wrongdoing, two remain

unresolved and two have no public support. Those taken up concerned procedures leaving documentary traces (procurement, partner verification, foundation affiliation) already flagged by oversight bodies. After the window, the agency called a claim that programme profits were shared with the President a hoax (June; Republika, 2026b), and Komdigi did the same for a claim that Elon Musk would fund the programme (July; detikInet, 2026c), a sign that in public imagination the platform's owner and the programme had merged.

Table 3. Suspicion themes voiced in the study window and their subsequent status (to 15 November 2026).

Theme in window (first seen)	**Subsequent development**	**Status**
Mark-ups in procurement, including electric motorcycles (16–20 Apr)	Former head and two deputies named suspects; procurement of 21,801 e-motorcycles, 32,000 pairs of shoes, 31,000 tablets and 5,400 televisions under investigation (3 Jun)	Taken up by legal action
The agency head "could become a suspect" (29 Apr)	Dismissed 2 Jun; named suspect 3 Jun; detention extended (25 Jun)	Taken up by legal action
Kitchens run by insiders and their circle; profit-sharing (11 May)	Prosecutors allege affiliated foundations passed verification improperly; agency reports foundations levying kitchens per tray (2 Oct)	Taken up by legal action; confirmed by official admission
Halal certification procurement reported to KPK by ICW (8 May)	Listed among governance concerns by civil-society coalition; no public charge identified	Unresolved
Dependence on US soybeans and imported milk (16–21 May)	Trade agreement purchase commitments; dairy deficit acknowledged by government	Confirmed as policy fact (not wrongdoing)
IDR 4 trillion corruption claim (disputed 14 Apr)	No official figure on losses published; a figure of IDR 10 trillion attributed to a state audit was reported in October	Unresolved
Foreign philanthropy as a "new threat" (23 Mar)	No evidence of any such plan	Unsupported
Programme as a plot to destroy the country (10 May)	No evidence	Unsupported

Sources: Hukumonline (2026); Investortrust (2026b); Okezone (2026a); Kontan (2026c); CNN Indonesia (2026); Okezone (2026b); Agrolatam (2026); Katadata (2026a). Named suspects are presumed innocent.

DISCUSSION

Two Capitalisms and a Third Party

In the MBG case two forms of capital compete to define what is true: the kitchen through budgets, press releases, hoax denials and prosecutions, the oracle through instant answers carrying an aura of machine neutrality. In January the state blocked the oracle to protect citizens; by April the oracle was explaining the programme and praising its head. Sovereignty meant conditional inclusion, within which the oracle reproduced official framing, not through co-optation (it also produced satire and an unsupported statistic) but because it draws on what is most available: official communication. When that account changed overnight in June, its earlier pronouncements stayed in public threads uncorrected (Human Cooperation Lab, 2026; Maldita.es, 2026).

The oracle is thus a third party in state–society relations: *summoned*, so its agenda is its summoners'; *registral* rather than positional, legitimising the programme in one thread and ridiculing it in another; and *unaccountable in time*, its statements persisting unrevised under a brand rather than a person. No coordination was needed for it to become, in some threads, an unpaid spokesperson for official policy: two forms of capital that both benefit from a calm public sphere converge structurally, which conspiracy narratives misread as a plot. We saw no sign that Grok was adjusted for the programme, but a foreign company holds that capacity and has used it politically elsewhere; controlling access to a foreign AI service is not controlling its role in domestic deliberation (Pohle & Thiel, 2020).

Suspicion as a Distorted Map of Real Opacity

The crowd's suspicion was mostly not conspiratorial, and several specific suspicions anticipated official action by weeks. In Fenster's (2008) and Jameson's (1988) terms, a programme that creates tens of thousands of kitchen contracts, routes them through politically tied foundations and gives hundreds of kitchens to the security forces is too opaque to see whole, and "thieves disguised as nutrition" is a compressed map of that opacity. A government that labels all suspicion hoax discredits itself when its own prosecutors take some up; answering documented claims with data and treating only hidden-agenda narratives as disinformation better preserves trust (Kontan, 2026e; Republika, 2026b).

Implications for Early Warning, the State and Platforms

Three additions to the relevance-gated pipeline (Sari et al., 2026b) follow. A *governance-discontent tier* built from the suspicion lexicons and a sarcasm signal would show the rise of governance suspicion that preceded the June crisis, triggering a transparency response rather than a health-facility check. *Oracle-reply monitoring* should flag AI-account replies that verify, deny or praise, recording date and sources, and route them to analysts without raising alerts alone. A *corroboration ledger* should record each theme's first appearance and later status (Table 3), so that a theme "unsupported" after six months is not escalated as new. All tiers should pass a validation gate that requires independent confirmation, collapses duplicates, checks for coordinated posting and sends

every Critical alert to a human. Evaluation should include governance episodes such as 2–3 June 2026, when suspicion posts quadrupled (19 to 75) at a near-constant 5% share, and score oracle replies for accuracy and timeliness.

For the agency, the findings point to a searchable register of kitchens, foundations, partners and suppliers, letting documented claims be checked and removing the premise of generalised accusations that nothing can be known; to reserving hoax labels for fabricated claims; and to machine-readable information for the chatbots citizens consult. Platforms' embedded answer engines should date and source their answers, decline to endorse named officials and be correctable. Given the use of the Electronic Information and Transactions Law against critics, analysts should report only what official bodies and reputable outlets state, describe persons under investigation as suspects and treat conspiracy content as data.

A Game-Theoretic Synthesis: Equilibrium without Conspiracy

To ask why the observed behaviour might be stable, we formalise the triangle as a sequential game solved by backward induction (Fudenberg & Tirole, 1991), drawing on models of strategic communication (Crawford & Sobel, 1982). The model is stylised: it identifies which parameters policy can change, without estimating them. The *state* S moves first and chooses whether to *disclose* procurement and ownership records, at cost k, or remain *opaque*; $\Delta = \Pi_O - \Pi_D$ is the rent that opacity protects. The *crowd* C chooses among *trust*, *generalised suspicion*, which yields an attention return e and carries a small legal risk r_G, and a *documented claim*, which yields reward b with corroboration probability q at effort c_d and larger legal risk r_D; corroboration is more likely when records are public ($q_D > q_O$). The *oracle* O either *endorses* the most prominent, usually official, account or *verifies* at cost v; a later-contradicted endorsement carries penalty β with probability p_c, and endorsement gives the state a legitimacy gain ρ under opacity and $\rho' \leq \rho$ under disclosure. Generalised suspicion costs the state λ in legitimacy and a corroborated documented claim costs it μ (Figure 4). Solving backwards gives three conditions:

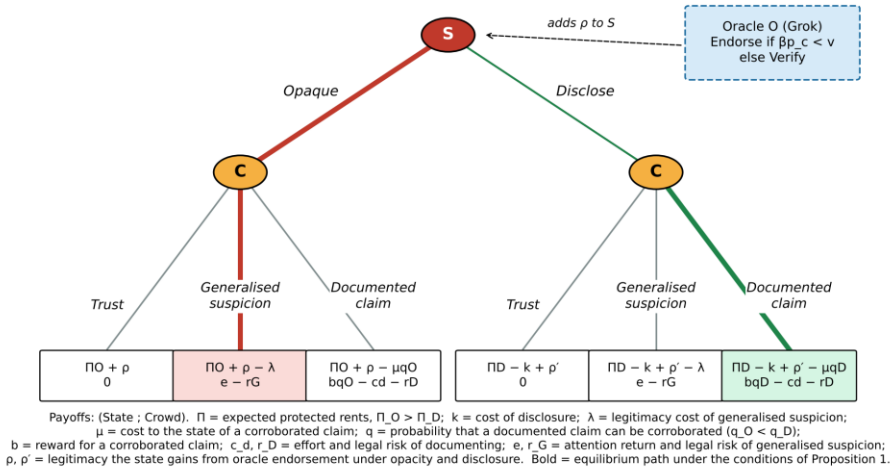

Figure 4. The disclosure–suspicion game. The state (S) chooses disclosure or opacity; the crowd (C) chooses trust, generalised suspicion or a documented claim; the oracle (O) endorses or verifies. Terminal boxes show payoffs to the state and the crowd. Bold branches mark the equilibrium path under the conditions of Proposition 1 (left) and the alternative path under disclosure (right).

$$\textit{Oracle endorses:}\quad \beta p_c < v \tag{1}$$

$$\textit{Crowd generalises under opacity:}\quad e - r_G > bq_O - c_d - r_D \tag{2}$$

$$\textit{State stays opaque:}\quad \lambda < \Delta + k + \mu q_D + (\rho - \rho') \tag{3}$$

Proposition 1 (suspicion trap). The profile (Opaque; generalised suspicion under opacity and documented claims under disclosure; Endorse) is a subgame-perfect equilibrium if (i) condition (1) holds, (ii) condition (2) holds while $bq_D - c_d - r_D > \max(0, e - r_G)$, and (iii) condition (3) holds. Generalised suspicion is cheap for the state because it names no act and can be dismissed as hoax; the crowd prefers it because, without records, a documented claim is costly, risky and unlikely to be confirmed. The corpus is consistent with each element: generalised accusations outnumbered documented claims (15 against 10), hidden-agenda narratives were rarest (3), suspicion reached more viewers ($e > r_G$), Grok endorsed official rationales in seven of its sixteen programme posts and never corrected its endorsement of the agency head (low β), and, as far as we could establish, the state had not published a register of kitchen ownership by the end of the window (opacity).

Proposition 2 (complementarity). An increase in the oracle's endorsement value ρ widens the set of disclosure costs and legitimacy costs for which opacity is an equilibrium; raising its liability β above v/p_c removes the endorsement and narrows that set by $\rho - \rho'$. State and platform capitalism are thus allied without contact: under opacity the most visible sources an answer engine retrieves are official statements, so its endorsement is worth more to an opaque state, with no transfer or pressure needed (Figure 5a).

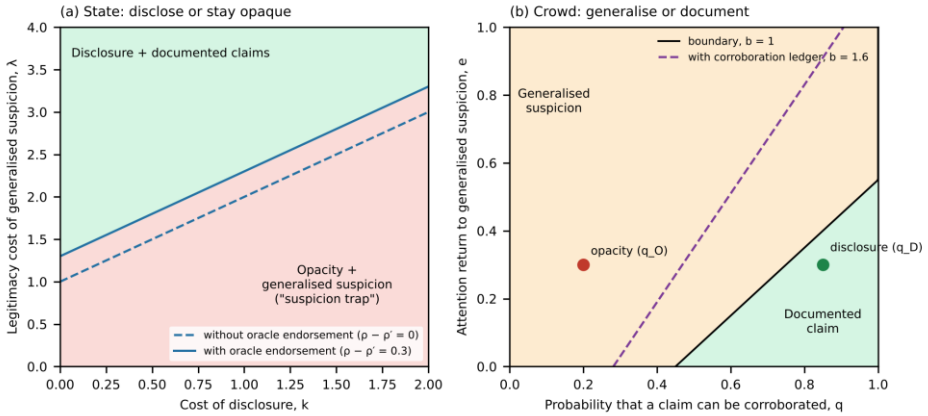

Figure 5. Comparative statics of the disclosure–suspicion game with illustrative parameter values ($\Delta + \mu q_D = 1.0$; $c_d = 0.3$; $r_D = 0.2$; $r_G = 0.05$). (a) The state's choice as a function of disclosure cost k and the legitimacy cost of generalised suspicion λ, with and without oracle endorsement. (b) The crowd's choice as a function of the probability of corroboration q and the attention return to generalised suspicion e; the dashed line shows the effect of a corroboration ledger that raises the reward to corroborated claims.

Proposition 3 (limits of monitoring). Measures that change only the crowd's payoffs (b, c_d, r_D) move the crowd to documented claims only if q_O exceeds the threshold in equation (4); when q_O is near zero no feasible reward does so (Figure 5b).

$$q_O > (e - r_G + c_d + r_D) / b \quad (4)$$

A corroboration ledger therefore helps once records exist but cannot substitute for them, consistent with Brazilian evidence that public release of municipal audit reports let voters punish corruption (Ferraz & Finan, 2008). Leaving the trap requires raising q, which only the state can do by disclosing, or raising λ through aggregation of suspicion by theme and region and through the ledger's track record: three of eight themes taken up by legal action make dismissal costlier.

Proposition 4 (credible threat). If the state's threat to block a foreign platform is credible and the platform's compliance cost m is below the market's value M, the unique subgame-perfect outcome is compliance after, at most, a temporary block (Schelling, 1960), as in January 2026. The state has not used this lever on the oracle's answers about public money, which by Proposition 2 benefit it, so demand for oracle accountability must come from outside the state. Whether the June naming of the former agency head, which raises λ, moved the game toward disclosure or repression is for later data. Conspiracy narratives about the programme thus need neither a conspiracy to explain them nor dismissal as irrational: they are what a rational crowd says when it cannot document what it suspects, facing a state that gains little from disclosure and an oracle that loses nothing by repeating power.

CONCLUSIONS AND RECOMMENDATIONS

Public talk about Indonesia's state-capitalist meal programme in 2026 ran on a foreign platform whose AI chatbot was its most active voice. The oracle acted mainly as a general-purpose interlocutor but, on the programme, more often reproduced official rationales than checked claims, endorsed the agency head weeks before he became a suspect and amplified satire. Public suspicion was mostly not conspiratorial: documented claims and generalised rent accusations far outnumbered hidden-agenda narratives, three of eight specific themes were taken up by legal action within six months, and the hidden-agenda themes found no support. A food-safety early warning rule saw half of this suspicion as Safe, and an unvalidated emotion model saw none of it as anger. Opacity, generalised suspicion and oracle endorsement sustain one another without coordination, and only disclosure of procurement records, which makes corroboration possible, breaks the equilibrium. Early warning for contested programmes must arbitrate between state, oracle and citizens and remember how past suspicions resolved.

ADVANCED RESEARCH

The keyword corpus inherits the acronym's noise; the Grok sub-corpus lacks the summoning prompts; a single coder precluded inter-coder reliability; the lexicons miss implicit insinuation and images; mention counts are not stances; and engagement without follower numbers cannot separate reach from audience size. Corroboration rests on sources to 15 November 2026; investigations may change any theme's status and suspects may be acquitted. The study covers one platform and period, X users are not representative of Indonesians (Mellon & Prosser, 2017), the model may since have changed, and the account of Grok's instruction layers rests on company statements and press reports. Future work should capture oracle prompts and reactions, extend the ledger across programmes and platforms, and test whether a governance-discontent tier warns of crises earlier than incident-based rules.

ACKNOWLEDGMENT

The authors received no specific funding for this work and declare no competing interests. The corpus is shared with two companion studies (Sari et al., 2026a, 2026b); the Grok, suspicion, corroboration and replay analyses are new.

Use of Generative AI

Claude (Anthropic), a large language model, was used to help draft and structure this manuscript from the authors' data, earlier manuscripts and notes, and to produce Figures 1–5. The authors checked all text, figures, numbers and sources and take full responsibility for the content.

Author Contributions

Conceptualisation and game-theoretic framing: A.W., C.S. and I.A.K.S. Data curation, formal analysis and writing of the original draft: I.A.K.S. Supervision, methodology and writing review and editing: A.W. and C.S. All authors approved the final manuscript.

Data Availability

The corpus, coded Grok and suspicion sub-corpora, lexicons and replay scripts are available in the project repository (https://github.com/digimetateam-hash/tesis) and the archived release (https://doi.org/10.5281/zenodo.11029482). The co-occurrence map is provided as VOSviewer files with the figure scripts. The full coding tables (Grok's programme-related posts, coding rules, the 59 coded suspicion posts and the chronology of events) are in the supplementary file.

REFERENCES

Agrolatam. (2026). US–Indonesia trade deal: Agricultural exports and tariffs. https://www.agrolatam.com/news/us-indonesia-trade-deal-agricultural-exports-tariffs/

Alami, I., & Dixon, A. D. (2020). State capitalism(s) redux? Theories, tensions, controversies. *Competition & Change, 24*(1), 70–94. https://doi.org/10.1177/1024529419881949

ANTARA. (2026a, May 19). Pemerintah pangkas anggaran MBG jadi Rp268 triliun pada 2026 [Government cuts MBG budget to IDR 268 trillion in 2026]. ANTARA News.

Arab News. (2025, June 25). Grok shows "flaws" in fact-checking Israel–Iran war: Study. https://www.arabnews.com/node/2605704

Bisnis.com. (2025a, November 10). Ekonom kritik rencana Danantara danai proyek peternakan ayam Rp20 triliun [Economists criticise Danantara's IDR 20 trillion poultry plan]. https://ekonomi.bisnis.com/read/20251110/12/1927479/

Bisnis.com. (2026a, April 20). Catatan dan usulan KPK untuk program MBG [KPK's notes and recommendations on the MBG programme]. https://kabar24.bisnis.com/read/20260420/16/1967750/

Bremmer, I. (2010). *The end of the free market: Who wins the war between states and corporations?* Portfolio.

CNBC Indonesia. (2026, September 3). BGN siapkan anggaran MBG tahun depan Rp232,5 triliun [BGN prepares IDR 232.5 trillion for MBG next year]. https://cnbcindonesia.com/news/20260903161816-4-764731/

CNN Indonesia. (2026, October 2). BGN temukan dapur MBG "dipalak" yayasan "gabut" [BGN finds MBG kitchens "extorted" by idle foundations]. https://cnnindonesia.com/ekonomi/20261002211641-92-1411125/

Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. *Econometrica, 50*(6), 1431–1451. https://doi.org/10.2307/1913390

detikInet. (2026a, January 10). Komdigi resmi blokir akses Grok AI imbas marak deepfake asusila [Komdigi blocks Grok AI over sexual deepfakes]. https://inet.detik.com/law-and-policy/d-8300492/

detikInet. (2026b). Komdigi buka blokir Grok AI, ini janji platform milik Elon Musk [Komdigi lifts Grok block after platform's pledge]. https://inet.detik.com/law-and-policy/d-8335512/

detikInet. (2026c, July 9). Elon Musk disebut akan tanggung anggaran MBG, Komdigi bilang hoax [Claim that Elon Musk will fund MBG is a hoax, says Komdigi]. https://inet.detik.com/cyberlife/d-8567350/

Douglas, K. M., Sutton, R. M., & Cichocka, A. (2017). The psychology of conspiracy theories. *Current Directions in Psychological Science, 26*(6), 538–542. https://doi.org/10.1177/0963721417718261

Fenster, M. (2008). *Conspiracy theories: Secrecy and power in American culture* (Rev. ed.). University of Minnesota Press.

Ferraz, C., & Finan, F. (2008). Exposing corrupt politicians: The effects of Brazil's publicly released audits on electoral outcomes. *Quarterly Journal of Economics, 123*(2), 703–745. https://doi.org/10.1162/qjec.2008.123.2.703

Fortune. (2025a, February 24). xAI chief engineer blames former OpenAI employee for Grok blocking Musk and Trump misinformation sources. *Fortune*. https://fortune.com/2025/02/24/xai-chief-engineer-blames-former-openai-employee-grok-blocks-musk-trump-misinformation/

Fortune. (2025b, July 8). Elon Musk's Grok AI shows conservative tilt after system prompt change. *Fortune*. https://fortune.com/2025/07/08/elon-musk-grok-ai-conservative-bias-system-prompt/

Fudenberg, D., & Tirole, J. (1991). *Game theory*. MIT Press.

Gebrak. (2026, August). Anggaran BGN tembus Rp240,2 triliun di RAPBN 2027, paling besar di atas Kemhan dan Kesehatan [BGN budget reaches IDR 240.2 trillion, above Defence and Health]. https://www.gebrak.id/2026/08/

Gigazine. (2026, February 2). Indonesia lifts block on access to Grok. https://gigazine.net/gsc_news/en/20260202-indonesia-lift-ban-grok

Gillespie, T. (2018). *Custodians of the internet: Platforms, content moderation, and the hidden decisions that shape social media*. Yale University Press.

Hukumonline. (2026, June 3). Diduga korupsi program MBG, Kejagung tahan eks Kepala BGN [Attorney General's Office detains former BGN head over alleged MBG corruption].

Human Cooperation Lab. (2026, February 3). @Grok is this true? LLM-powered fact-checking on social media [Working paper announcement]. https://groups.google.com/g/human-cooperation-lab/c/6ADjTadDT44

InfoPublik. (2025). Peran strategis Agrinas Pangan Nusantara perkuat ketahanan pangan nasional [The strategic role of Agrinas Pangan Nusantara in national food security]. https://infopublik.id/kategori/sorot-ekonomi-bisnis/932775/

Investortrust. (2026a, September 17). BGN dapat anggaran Rp240 triliun, alokasi terbesar disalurkan ke MBG [BGN receives IDR 240 trillion, mostly for MBG]. https://investortrust.id/national/116704/

Investortrust. (2026b, June 3). Daftar proyek program MBG yang dikorupsi Dadan cs [List of MBG procurements in the case against Dadan and others]. https://investortrust.id/national/105217/

Investortrust. (2026c, February 20). Indonesia secures zero tariffs on 1,819 products in U.S. trade pact. https://investortrust.id/market/95035/

Jameson, F. (1988). Cognitive mapping. In C. Nelson & L. Grossberg (Eds.), *Marxism and the interpretation of culture* (pp. 347–360). University of Illinois Press.

JawaPos. (2026a, June 14). Carut marut program MBG libatkan orang dekat Prabowo hingga yayasan terafiliasi parpol [MBG turmoil: insiders and party-affiliated foundations]. https://jawapos.com/kasuistika/2606140194/

Katadata. (2026a). Indonesia masih defisit susu 6,72 juta ton di tengah kebutuhan MBG [Indonesia's milk deficit of 6.72 million tonnes amid MBG demand].

Keeley, B. L. (1999). Of conspiracy theories. *The Journal of Philosophy, 96*(3), 109–126.

Kim, K. (2021). Indonesia's restrained state capitalism: Development and policy challenges. *Journal of Contemporary Asia, 51*(3), 419–446. https://doi.org/10.1080/00472336.2019.1675084

Kontan. (2026a). Prabowo peringatkan bahaya ketergantungan impor di tengah konflik global [Prabowo warns of import dependence amid global conflict].

Kontan. (2026b, June 7). Kementan kejar produksi susu lokal, andalkan impor sapi perah untuk dukung MBG [Ministry of Agriculture relies on dairy cattle imports for MBG].

Kontan. (2026c, June 3). Yayasan terafiliasi eks pimpinan BGN diduga terima insentif miliaran dari MBG [Foundations affiliated with former BGN leaders allegedly received billions].

Kontan. (2026e, April 24). BGN sebut banyak hoaks soal program MBG di medsos [BGN says many hoaxes about MBG on social media].

Kurlantzick, J. (2016). *State capitalism: How the return of statism is transforming the world*. Oxford University Press.

Lazer, D., Kennedy, R., King, G., & Vespignani, A. (2014). The parable of Google Flu: Traps in big data analysis. *Science, 343*(6176), 1203–1205. https://doi.org/10.1126/science.1248506

Maldita.es. (2026, October 27). Grok is not reliable for fact-checking: Study of 766 queries about viral World Cup claims. https://maldita.es/investigaciones/20261027/grok-artificial-intelligence-chatbts-world-cup/

McMichael, P. (2009). A food regime genealogy. *The Journal of Peasant Studies, 36*(1), 139–169. https://doi.org/10.1080/03066150902820354

Mellon, J., & Prosser, C. (2017). Twitter and Facebook are not representative of the general population. *Research & Politics, 4*(3). https://doi.org/10.1177/2053168017720008

MetroTV. (2026, January 14). Indonesia jadi negara pertama blokir AI Grok di aplikasi X [Indonesia becomes first country to block Grok].

Okezone. (2026a, June 25). Korupsi MBG, penahanan Dadan Hindayana cs diperpanjang 40 hari [Detention of Dadan Hindayana and others extended].

Okezone. (2026b, October 15). Kejagung didorong usut tuntas dugaan korupsi Dadan cs [Prosecutors urged to complete the investigation].

Ombudsman RI. (n.d.). Ombudsman RI temukan empat potensi maladministrasi dalam program Makan Bergizi Gratis [Ombudsman finds four potential maladministration issues]. https://ombudsman.go.id/news/download/

Pohle, J., & Thiel, T. (2020). Digital sovereignty. *Internet Policy Review, 9*(4). https://doi.org/10.14763/2020.4.1532

Republika. (2026a, June 5). Kejagung tak pilih kasih, SPPG TNI-Polri jika tersangkut korupsi MBG akan diusut [Prosecutors will investigate military and police kitchens if implicated].

Republika. (2026b, June 13). BGN: Tuduhan pembagian keuntungan program MBG ke Presiden adalah hoaks [BGN: claim of profit-sharing with the President is a hoax].

Robison, R., & Hadiz, V. R. (2004). *Reorganising power in Indonesia: The politics of oligarchy in an age of markets*. RoutledgeCurzon.

Sakaki, T., Okazaki, M., & Matsuo, Y. (2010). Earthquake shakes Twitter users: Real-time event detection by social sensors. In *Proceedings of the 19th International Conference on World Wide Web* (pp. 851–860). https://doi.org/10.1145/1772690.1772777

Sari, I. A. K., Suratnoaji, C., & Widiyarta, A. (2026a). *Public discourse on Indonesia's Free Nutritious Meal (MBG) programme: Social network analysis and an IndoBERT approach* [Manuscript submitted for publication].

Sari, I. A. K., Suratnoaji, C., & Widiyarta, A. (2026b). Relevance-gated early warning with Telegram alerts for free nutritious meal discourse on X. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.7526942

Schelling, T. C. (1960). *The strategy of conflict*. Harvard University Press.

Setkab. (2026, February 20). Indonesia, US reach agreement on reciprocal trade, impose zero percent tariffs to 1,819 Indonesian products. Cabinet Secretariat of the Republic of Indonesia.

Srnicek, N. (2017). *Platform capitalism*. Polity.

Sunstein, C. R., & Vermeule, A. (2009). Conspiracy theories: Causes and cures. *Journal of Political Philosophy, 17*(2), 202–227. https://doi.org/10.1111/j.1467-9760.2008.00325.x

TechCrunch. (2025a, May 16). xAI blames Grok's obsession with white genocide on an "unauthorized modification". *TechCrunch*. https://techcrunch.com/?p=3007866

TechCrunch. (2025b, July 10). Grok 4 seems to consult Elon Musk to answer controversial questions. *TechCrunch*. https://techcrunch.com/2025/07/10/grok-4-seems-to-consult-elon-musk-to-answer-controversial-questions/

The Decoder. (2025, July 15). xAI says Grok 4 is no longer searching for Musk's views before it answers. *The Decoder*. https://the-decoder.com/xai-says-grok-4-is-no-longer-searching-for-musks-views-before-it-answers/

Transparency International Indonesia. (2025). *Risiko korupsi di balik hidangan Makan Bergizi Gratis* [Corruption risks behind the free nutritious meal]. https://ti.or.id/wp-content/uploads/2025/06/

Uscinski, J. E., & Parent, J. M. (2014). *American conspiracy theories*. Oxford University Press.

van Dijck, J., Poell, T., & de Waal, M. (2018). *The platform society: Public values in a connective world*. Oxford University Press.

van Eck, N. J., & Waltman, L. (2010). Software survey: VOSviewer, a computer program for bibliometric mapping. *Scientometrics, 84*(2), 523–538. https://doi.org/10.1007/s11192-009-0146-3

xAI. (2025). *grok-prompts: Prompts for the Grok chat assistant and the @grok bot on X* [Computer software]. GitHub. https://github.com/xai-org/grok-prompts