

“Link-Tax” Policy and News Media Market Concentration

Seth Kushniryk

Department of Economics, Queen’s University

October 19, 2026

Abstract

Australia’s 2021 News Media Bargaining Code (NMBC) induces Google and Meta to negotiate payments with news publishers, yet the distributional consequences of these agreements are unknown. Using novel monthly web traffic data for 425 Australian news outlets from Semrush, I find that the NMBC produced a transitory spike in outlet-level market concentration of approximately 170 HHI points, which reverted within roughly 16 months. Computed at the parent-company level, concentration instead rose persistently: the agreement-holding parents captured share from non-agreement parents and held it. At the outlet level, a propensity-score-matched difference-in-differences design shows that outlets whose parent companies secured platform agreements received roughly 82 per cent more organic traffic than non-agreement outlets over a one-year window, with the result confirmed by a triple-difference design using three control countries. Reconciling these results: the agreement-holding parents captured between-parent share, then redistributed the traffic gains within their portfolios in a way that lifted their smallest regional and digital sub-properties while leaving their flagship outlets untouched. The policy compressed the size distribution among agreement holders, but a substantial tail of below-threshold publishers falls outwith the bargaining framework and the present analysis alike.

1 Introduction

Google and Meta together account for the majority of referral traffic to news websites in the English-speaking world. The dependence of news publishers on these two platforms has produced what former ACCC chair Rod Sims characterises as “a clear market failure that needed to be addressed” (Sims, 2022). Governments have responded with a class of policies commonly termed “link taxes,” which seek to redirect revenue from digital platforms to the news organisations from whose content those platforms profit. The European Union’s Article 15, Australia’s News Media Bargaining Code (NMBC), and Canada’s *Online News Act* each instantiate different versions of this idea. Australia’s is the oldest functioning regime, enacted in March 2021, and the one with the most developed track record.

The NMBC does not impose a tax in the conventional sense. It creates a bargaining framework: the Australian Treasury may “designate” a platform, whereupon mandatory arbitration provisions apply. No platform has been designated to date. The threat alone has been sufficient to induce Google and Meta to enter roughly thirty voluntary agreements with Australian news businesses. The value of these agreements is highly unequal. Nine Entertainment reportedly receives AUD\$30 million per year, whereas Country Press Australia (a collective bargaining body representing 160 regional newspapers) pays its members only AUD\$31,000–\$62,000 each from its Google contract (Grueskin, 2022), a per-outlet gap of two to three orders of magnitude. Small outlets are ineligible altogether, and the terms of each agreement are secret.

The stated purpose of the NMBC is to sustain Australian journalism broadly. Yet the mechanism through which it operates is a series of private bilateral transfers to outlets that negotiate. Outlets that fail to secure an agreement receive nothing. If the platforms, having agreed to pay some publishers, reallocate traffic towards those publishers and away from others, then the policy could increase market concentration rather than reduce it. Inasmuch as media concentration has downstream consequences for democratic accountability and civic participation (Gao, Lee, & Murphy, 2020), the distributional effects of link-tax regimes are a first-order policy question.

This paper asks two things. First, did the NMBC increase news media market concentration in Australia? Secondly, did platform agreements redistribute web traffic from non-agreement to agreement outlets, and is this the mechanism behind any observed change in concentration?

I answer both questions using a novel dataset of monthly web traffic for 425 Australian news outlets, drawn from Semrush, over the period January 2017 to June 2024. Ownership

data, compiled from the EIN Presswire World Media Directory and archival research, allows me to compute outlet-level Herfindahl–Hirschman indices. Agreement status comes from an Australian Treasury disclosure.

A segmented regression on the outlet-level HHI time series shows a level shift of approximately 170 points preceding NMBC passage. The effect is transitory: a counter-vailing trend change of -10.61 points per month erases the initial spike within roughly 16 months. Computed at the ownership level, however, the picture reverses: an ownership-level HHI shows a smaller level shift of $+93$ points at passage and a positive trend of $+13.26$ per month thereafter, rising persistently. The two measures move in opposite directions after the NMBC because the reversion at the outlet level reflects compression of the size distribution within the agreement group: the smaller regional and digital sub-properties of agreement-holding parents grew rapidly while their flagship outlets did not, an apparent puzzle taken up in detail in Section 6.4. Turning to the outlet level, my preferred specification is a propensity-score-matched difference-in-differences regression that re-weights non-agreement outlets to balance pre-period traffic level, volatility, and trend. Agreement outlets received 46 per cent more organic web traffic than the matched non-agreement comparison group in the six months following the NMBC, rising to 82 per cent over a one-year window. The unmatched specification yields almost identical estimates of 48 and 90 per cent. The matched estimates pass Oster (2019) bounds for selection on unobservables ($|\delta| > 10$) and hold in a triple-difference framework that uses Canadian, New Zealand, and British outlets as controls.

The puzzle of a transitory aggregate effect alongside persistent outlet-level divergence is resolved by heterogeneity within the treated group. When I split agreement outlets into pre-period size quartiles, the smallest quartile shows a treatment effect of 2.73 log points (a roughly fourteenfold increase in traffic), whereas the largest quartile shows no significant effect at all. The NMBC did not so much advantage large incumbents as allow the smaller regional and digital sub-properties of the agreement-holding parents to catch up to their flagship outlets, compressing the size distribution and pushing HHI back down. The aggregate concentration increase reversed precisely because of that pattern.

The Australian experience also raises an external-validity concern that Sandrini and Somogyi (2024) term the “going-dark” constraint: platforms may withdraw from carrying news rather than pay for it. Meta did so temporarily in Australia (February 2021) and permanently in Canada (August 2023), producing an 85 per cent decline in Canadian news engagement on the platform (Media Ecosystem Observatory, 2024). With Meta declining to renew its Australian agreements in February 2024 (Meta, 2024) and the Treasury responding in December 2024 with a proposed News Bargaining Incentive scheme

(Jones & Rowland, 2024), the NMBC may be approaching its natural endpoint, but the distributional questions it raises apply equally to whatever succeeds it.

The existing empirical literature on link-tax policy addresses aggregate traffic effects following the Spanish Google News exit (Calzada & Gil, 2020; Majó-Vázquez, Cardenal, & González-Bailón, 2017) and content substitution within platform feeds (Freimane, 2025); theoretical work on news aggregation (Jeon & Nasr, 2016) considers market expansion and substitution but not the distributional consequences of negotiated transfers. The outlet-level distributional question, namely how a bargaining code reshapes the allocation of audience across outlets, is the one I take up here, exploiting agreement-level data that have not previously been available.

2 Institutional Background

2.1 Platform intermediaries and the news industry

The shift in advertising revenue from print media to digital platforms is by now a familiar story. What matters here is the narrower question of traffic dependence. News websites in Australia, as in most English-speaking countries, receive a large share of their audience through two channels: Google Search (and its subsidiary, Google News) and Meta’s Facebook news feed. The precise shares vary by outlet and over time. What is invariant is that these platforms function as gatekeepers. An outlet that loses visibility on Google or Facebook will, in most cases, lose readers.

The Australian Competition and Consumer Commission recognised this dependence in its 2019 *Digital Platforms Inquiry* (Australian Competition and Consumer Commission, 2019), which concluded that a “significant imbalance of bargaining power” existed between digital platforms and Australian news businesses. The inquiry recommended legislation to address the imbalance. The result was the Treasury Laws Amendment (News Media and Digital Platforms Mandatory Bargaining Code) Act 2021 (Parliament of Australia, 2021), commonly referred to as the News Media Bargaining Code or NMBC.

2.2 The News Media Bargaining Code

The NMBC entered into force in March 2021. Its design is unusual among link-tax policies, and the distinction from European-style copyright amendments is worth drawing clearly.

The European approach, exemplified by Article 15 of the 2019 Directive on Copyright in the Digital Single Market (European Parliament and Council, 2019), modifies copyright

law to grant publishers a neighbouring right over the use of news snippets. Spain's earlier implementation (repealed in 2021) went further, making the right non-waivable. Google responded by shutting down Google News in Spain rather than pay. The NMBC takes a different route. It does not alter copyright. Instead, it establishes a bargaining framework backed by the threat of compulsory arbitration.

The mechanism works as follows. The Australian Treasury may "designate" a digital platform if it determines that a "significant bargaining power imbalance" exists and that the platform has not made a "significant contribution to the sustainability of the Australian news industry through commercial agreements." Once designated, a platform becomes subject to mandatory negotiation, mediation, and final-offer arbitration with eligible news businesses. The arbitration provisions are deliberately punitive: final-offer arbitration, in which the arbitrator must choose one party's proposal in its entirety, creates strong incentives to settle.

No platform has been designated. The threat alone, sometimes described as a "Sword of Damocles" in policy commentary, has been sufficient to induce voluntary agreements. Google and Meta have each entered into agreements with numerous Australian publishers. The Treasury's 2022 disclosure identifies approximately thirty such agreements (Treasury, 2022). Eight large media businesses hold agreements with one or both platforms: the ABC, Australian Community Media, Guardian Australia, Network Ten, News Corp, Nine Entertainment, SBS, and Seven West Media. Four smaller entities (Out Publications, Region Group, Schwartz Media, and The Conversation) hold agreements with Google, in some cases alongside a Meta agreement. Country Press Australia, a collective bargaining body representing 160 regional newspapers, holds agreements with both Google and Meta.

The distribution of value across these agreements is sharply unequal. Nine Entertainment reportedly receives AUD\$30 million per year, whereas Country Press Australia's Google contract pays roughly AUD\$31,000–\$62,000 per member newspaper (Grueskin, 2022). The per-outlet gap is two to three orders of magnitude. Outlets with annual turnover below AUD\$150,000 are ineligible for the NMBC's protections altogether (Parliament of Australia, 2021), which excludes many small community and independent publishers from the regime a priori. The terms of each agreement are confidential, imposed at the platforms' insistence. One anonymous media executive described the process to the *Columbia Journalism Review* as follows: "It's like a brown paper bag gets stuffed with money, is shoved across the table, and then the platforms can say, 'Now just shut the fuck up.'" (Grueskin, 2022). Bossio and Barnett (2023) document the unequal distribution of outcomes and the frustration of excluded publishers in detail.

Two further features of the NMBC bear on the analysis. The Code permits collective bargaining; three collective groups formed: Country Press Australia, which negotiated as a bloc for its regional members; a second representing special-interest outlets; and a third, representing radio stations, which secured no deals. Separately, Meta blocked all Australian news content on its platform for approximately one week in February 2021 (Easton, 2021), just prior to the Code's passage, in an apparent attempt to demonstrate that it could withdraw from the Australian news market entirely. The blockade was lifted after negotiations with the government, and Meta subsequently entered into agreements. The episode foreshadowed a more consequential withdrawal in Canada two years later.

2.3 Recent developments

The Australian regime has entered a period of uncertainty. On 29 February 2024, Meta announced that it would deprecate its Facebook News tab in Australia and the United States, and would not enter into new commercial deals for traditional news content in either market (Meta, 2024); the existing three-year Australian agreements, signed in 2021, were on a path to expire without renewal. On 12 December 2024, the federal government responded by announcing the News Bargaining Incentive, a charge-and-offset scheme designed to incentivise platforms with Australian revenues above AUD\$250 million to negotiate with publishers, with the charge offset by the value of any commercial agreements struck (Jones & Rowland, 2024). Following consultation through 2025 and early 2026, the government announced the scheme's final design in August 2026: a 2.5 per cent charge on covered platforms' Australian advertising revenue, offset by the value of commercial agreements with qualifying news businesses, with legislation before Parliament at the time of writing.

Meta's withdrawal is not idiosyncratic; it instantiates the going-dark constraint discussed in Section 1. Freimane (2025) adds evidence of platform substitution short of withdrawal: within Australia, the NMBC was associated with an increase in the share of foreign content displayed on Google News, suggesting that Google may have partially substituted away from domestic news thereanent.

The NMBC may therefore be approaching the end of its operational life. The distributional questions it raises, however, are not specific to the Australian regime. Canada, the European Union, New Zealand (where a *Fair Digital News Bargaining* bill is under consideration), and other jurisdictions face the same structural problem: any policy that channels platform payments through bilateral negotiation will tend to favour publishers with the resources and bargaining position to secure deals, potentially at the expense of

those excluded from it.

3 Related Literature

Three strands of empirical work bear on the relationship between digital platforms, link-tax policies, and news media markets.

The theoretical starting point is Jeon and Nasr (2016), who model a news aggregator that expands the total market for news by reducing search costs for readers, but may also substitute for direct visits to publishers' own websites. Whether the net effect on any given publisher is positive or negative depends on the publisher's outside option: large, well-known outlets are more likely to lose direct visits, while smaller outlets benefit from the aggregator's reach. This prediction is relevant to the heterogeneity results I report below.

The Spanish experience provides the main empirical evidence on what happens when an aggregator disappears. In December 2014, a new intellectual property law required news aggregators to pay publishers for displaying snippets. The right was non-waivable: publishers could not opt out of the fee or grant Google free use. Google responded by closing Google News in Spain. Calzada and Gil (2020) estimate that the closure reduced traffic to Spanish news sites by 8 to 14 per cent, with larger effects for smaller and less internationally oriented outlets. Majó-Vázquez et al. (2017) find that news consumption became more fragmented after the exit, with readers concentrating on a smaller set of directly accessed sites. Calzada, Duch-Brown, and Gil (2023), in related work, study Google News algorithm updates and find that changes to the aggregator's ranking algorithm reduce market concentration by approximately 1 per cent, suggesting that the algorithm's design has a measurable bearing on competitive outcomes.

Athey, Mobius, and Pal (2021) study the introduction of Google News in various countries and find that it increased total news consumption while partially cannibalising direct navigation to publishers. The effect is heterogeneous across outlet types, consistent with the theoretical predictions of Jeon and Nasr (2016).

Direct evidence on the effects of link-tax regimes is thin. The Spanish case is sometimes cited as link-tax evidence, but the mechanism there (aggregator exit) is quite different from the Australian one (negotiated transfers). The Canadian *Online News Act* produced a different form of platform withdrawal: Meta blocked all Canadian news rather than negotiate, and Google agreed to pay into a collective fund. Media Ecosystem Observatory (2024) document the consequences of Meta's withdrawal for Canadian news engagement; Song and Manchanda (2025) provide complementary evidence on how carrying news

affects engagement with non-news content on social media platforms.

The closest study to the present paper is Freimane (2025), who examines the effect of the NMBC on the composition of content displayed in Google News Australia. Freimane finds that the share of foreign (non-Australian) content increased after the Code’s passage, consistent with Google partially substituting towards content for which it need not pay. That work addresses content supply on the platform side; mine addresses traffic outcomes on the publisher side. The two are complementary, inasmuch as platform-side content substitution is one channel through which outlet-level traffic might be redistributed.

No existing study examines the effect of a link-tax policy on media market concentration, nor the outlet-level distributional consequences of platform agreements.

The concern motivating this paper is that link-tax policies might increase concentration in news markets. The question of why media concentration matters has a substantial literature of its own, largely outwith the scope of this paper, but a few results bear directly on the welfare interpretation of my findings. Gao et al. (2020) show that newspaper closures (an extreme form of local concentration) are associated with higher municipal borrowing costs, consistent with a loss of government accountability. The connexion between media pluralism and democratic outcomes, while difficult to quantify precisely, provides the normative backdrop against which distributional effects of platform policy should be evaluated.

4 Data

4.1 News outlet universe

The sampling frame for Australian news outlets is the EIN Presswire World Media Directory (WMD), a global index of news organisations. The WMD lists 466 Australian news outlets as of 2024, spanning newspapers (312), broadcast media (69), internet-native outlets (34), magazines (9), and one press agency. These outlets are distributed across all Australian states and territories, with the largest concentrations in New South Wales (138) and Victoria (84). Forty-seven outlets are classified as national in scope.

For the triple-difference analysis described in Section 5, I also draw on WMD listings for three control countries: Canada (1,076 outlets), New Zealand (268), and the United Kingdom (1,174). These countries share a common language, broadly similar media market structures, and, at the time of the NMBC’s passage, no comparable link-tax policy in force. Canada’s *Online News Act* was not enacted until 2023, and New Zealand’s proposed legislation remains under consideration.

Table 1: Sample Composition

<i>Panel A: Media type</i>		<i>Panel B: Location</i>	
Newspaper	312	New South Wales	138
Broadcast	69	Victoria	84
Internet	34	Queensland	66
Magazine	9	AU National	47
Press agency	1	South Australia	38
Total	425	Western Australia	31
		Tasmania	10
		Northern Territory	7
		Capital Territory	4
		Total	425

Notes: Australian news outlets listed in the EIN Presswire World Media Directory with usable Semrush data.

Table 1 summarises the composition of the Australian sample by media type and location.

4.2 Web traffic data

I obtain monthly web traffic data from Semrush, a commercial analytics platform that estimates organic search traffic to websites based on clickstream data and search engine results page modelling. For each outlet listed in the WMD, I collect Semrush’s estimate of the monthly average daily organic traffic originating from within Australia, over the period January 2017 to June 2024. Of the 466 WMD-listed Australian outlets, 425 have usable Semrush data for at least part of the sample period. The remaining 41 are sites for which Semrush reports no data, typically because the site is too small to appear in Semrush’s estimation model or because the domain has changed.

Sample composition matters for interpreting the non-agreement comparison group. The outlets that Semrush cannot model are typically very small operations with negligible organic search presence, which is roughly the same population that the NMBC’s AUD\$150,000 annual-turnover eligibility threshold excludes from the bargaining framework. The 425 outlets analysed below are therefore close to the universe of outlets potentially eligible to hold an agreement, and the within-sample non-agreement comparison is best read as “eligible but unagreed” rather than “all outlets without agreements.” The substantial tail of below-threshold publishers (community newsletters, hyperlocal blogs, niche specialist titles) falls outwith the dataset entirely.

Semrush traffic estimates are not equivalent to server-side analytics. They are modelled

quantities, derived from a combination of clickstream panels and search engine ranking data. The estimates are most reliable for sites with substantial search-engine-derived traffic and less reliable for sites whose audience arrives primarily through direct navigation, social media referral, or mobile applications. For the mechanism this paper studies, organic search is the relevant channel: platform agreements with Google bear most directly on search rankings, so any agreement-driven reallocation of audience should show up first in search-derived traffic. What matters is that Semrush provides a consistent measure of organic search visibility across a large number of outlets over time. Measurement error in levels is less concerning than it would be in a cross-sectional comparison, because the identifying variation comes from within-outlet changes over time.

Tables 2 and 3 present summary statistics for traffic and market shares, broken out by media type and by pre- and post-NMBC periods. The heterogeneity across outlets is extreme. Among newspapers, the median outlet received roughly 5,700 daily visits in the pre-period, while the largest received over 8 million. The distribution is similarly skewed for broadcast and internet-native outlets. The broadcast class shows a striking pre-to-post-period increase in traffic at every percentile (the median jumps from 199 to 11,291 daily visits, the 75th from 1,034 to 185,088), in contrast to the near-stable newspaper distribution. Three forces plausibly contribute: the COVID-19 news-consumption surge (the pre-period predominantly precedes March 2020, the post-period is fully after); the rapid expansion of broadcaster digital sub-properties over 2020–2024 (the ABC’s regional rollout is the most publicly documented case, with NMBC revenue committed to expanded regional and rural journalism); and small broadcast outlets crossing Semrush’s detection threshold mid-sample, so that they appear in the post but not the pre. Market shares, computed at the outlet level as each outlet’s traffic divided by total traffic across all 425 outlets, are correspondingly concentrated: the median newspaper’s share is 0.01 per cent in the pre-period, while the largest newspaper commands over 9 per cent.

4.3 Ownership data

Ownership information enters the analysis in two places. Multiple outlets may belong to the same parent company, which matters for interpreting concentration at the ownership level even though the HHI is computed at the outlet level. More directly, platform agreements are struck at the parent-company level, so an outlet’s treatment status is determined by the agreements held by its owner.

I compile ownership histories from the WMD listings and from archival research using the Wayback Machine, covering the period 2016 to 2022. Where ownership changed

Table 2: Summary Statistics: Daily Organic Traffic by Media Type

		25th	50th	75th	95th	Max	Mean
Broadcast	Pre	64	199	1,034	3,157,920	18,600,000	546,739
	Post	1	11,291	185,088	4,360,525	24,600,000	1,449,002
Internet	Pre	515	14,490	171,257	9,724,129	9,724,129	792,799
	Post	701	11,249	181,686	11,000,000	11,000,000	893,943
Newspaper	Pre	1,281	5,716	21,350	610,187	8,313,923	161,880
	Post	1,244	4,721	22,654	662,230	7,612,015	171,352

Notes: Monthly averages of daily organic traffic from Semrush, for Australian news outlets listed in the EIN Presswire World Media Directory. Pre-period: January 2017 to February 2021. Post-period: March 2021 to June 2024. Magazines ($n = 4$) and the single press agency are omitted for brevity.

Table 3: Summary Statistics: Outlet-Level Market Shares (%) by Media Type

		50th	75th	90th	95th	99th	Max
Broadcast	Pre	0.00	0.00	0.26	2.64	23.79	23.79
	Post	0.00	0.00	0.24	2.28	26.56	26.56
Internet	Pre	0.00	0.23	0.69	13.10	13.10	13.10
	Post	0.00	0.20	0.74	13.29	13.29	13.29
Newspaper	Pre	0.01	0.03	0.17	0.88	4.41	9.01
	Post	0.00	0.02	0.13	0.82	4.61	8.31

Notes: Market shares computed as each outlet’s organic traffic divided by total traffic across all 425 Australian outlets in the sample. Percentiles computed over outlet-level period means.

during the sample period, I assign the outlet to its owner as of each date. The major ownership groups in the Australian news market are News Corp, Nine Entertainment, the ABC (publicly owned), Seven West Media, and Australian Community Media. A large tail of independently owned outlets, many of them regional newspapers, accounts for the bulk of the outlet count but a small share of total traffic.

4.4 Agreement data

Agreement status comes from a disclosure by the Australian Treasury in its 2022 review of the NMBC (Treasury, 2022). Table 4 reproduces the list of businesses known to hold agreements, along with their Google and Meta agreement status. Thirteen entities are listed: eight major media businesses with agreements with one or both platforms, four smaller entities, and Country Press Australia.

A limitation of these data is that the precise timing of each agreement is unknown. The Treasury disclosure identifies which businesses hold agreements, but not when those

Table 4: Platform Agreements by News Business

News Business	Google	Meta
<i>Major media businesses</i>		
ABC	Yes	Yes
Australian Community Media	Yes	Yes
Guardian Australia	Yes	Yes
Network Ten	No	Yes
News Corp	Yes	Yes
Nine	Yes	Yes
SBS	Yes	No
Seven West	Yes	Yes
<i>Smaller entities</i>		
Out Publications Pty Ltd	Yes	No
Region Group	Yes	No
Schwartz Media	Yes	Yes
The Conversation	Yes	No
<i>Collective bargaining</i>		
Country Press Australia (160 papers)	Yes	Yes

Source: Australian Treasury (2022).

agreements were executed. I therefore treat the NMBC’s passage date (March 2021) as the treatment date for all agreement outlets. Insofar as some agreements were struck before the formal passage (in the period of intense negotiation during January and February 2021) and others after, this introduces measurement error in the treatment timing. The direction of the resulting bias depends on the temporal pattern of agreement formation; I discuss this further in Section 7.

5 Empirical Strategy

5.1 Market concentration: segmented regression

To test whether the NMBC affected aggregate market concentration, I estimate a segmented regression on the monthly outlet-level Herfindahl–Hirschman Index. The HHI is computed each month as

$$HHI_t = \sum_{j=1}^N s_{jt}^2$$

where s_{jt} is the traffic share of outlet j in month t , calculated as outlet j 's organic traffic divided by total traffic across all 425 outlets. I compute the index at the outlet level rather than at the ownership-group level. An ownership-group index, which would aggregate co-owned outlets under a single parent, is the conventional measure for antitrust purposes. I prefer the outlet-level version because it is sensitive to within-group reallocation across individual sites, which turns out to be central to the results. The ownership-group index would mask this variation. This choice is reinforced by the structure of much of the Australian media industry. ABC's state and regional units and ACM's regional titles retain substantial editorial and operational independence, each maintaining its own newsroom, leadership, and audience strategy. An ownership-group HHI would conflate these as a single competitive unit when in practice they operate quite separately. I nonetheless report an ownership-level HHI alongside the outlet-level series in Section 6.2: the two measures move in opposite directions after the NMBC, and the contrast is itself diagnostic of the mechanism.

The segmented regression takes the form

$$HHI_t = \beta_0 + \beta_1 post_t + \beta_2 time_t + \beta_3 (time_t \times post_t) + \varepsilon_t \quad (1)$$

where $post_t$ is an indicator equal to one from March 2021 onwards, and $time_t$ is a linear time trend centred at the treatment date. The parameter β_1 captures the level shift in HHI at NMBC passage, though much of the shift in fact leads that date; β_2 absorbs the pre-treatment trend, and β_3 captures any change in the trend thereafter. I estimate the model using ordinary least squares with Newey–West standard errors (Newey & West, 1987), allowing six lags to account for serial correlation in the monthly HHI series. The sample comprises 91 monthly observations.

5.2 Outlet-level mechanism: difference-in-differences

The aggregate HHI analysis cannot identify the mechanism behind any change in concentration. To test whether platform agreements redistributed traffic across outlets, I estimate a difference-in-differences specification at the outlet level:

$$Y_{it} = \beta_0 + \beta_1 Treat_{it} + \alpha_i + \gamma_t + \varepsilon_{it} \quad (2)$$

where Y_{it} is the inverse hyperbolic sine of organic traffic for outlet i in month t , α_i are outlet fixed effects, γ_t are year-month fixed effects, and $Treat_{it} = Post_t \times Agreement_i$ is the interaction of a post-NMBC indicator with an indicator for whether outlet i belongs to

a parent company that holds an agreement with Google, Meta, or both. The coefficient β_1 captures the average change in log traffic for agreement outlets relative to non-agreement outlets, net of common time effects and time-invariant outlet characteristics.

I estimate equation (2) over two symmetric windows around the treatment date of March 2021: a six-month window (September 2020 to September 2021) and a one-year window (March 2020 to March 2022). The shorter window reduces the risk of confounding from other events; the longer window allows the treatment effect to develop more fully.

Inference is by wild cluster-robust score bootstrap with 99,999 replications, clustered at the outlet level (MacKinnon, Nielsen, & Webb, 2023). I cluster at the outlet level for two reasons. First, although the major Australian media groups (ABC, ACM, News Corp, Nine, Seven West, etc.) hold NMBC agreements at the parent-company level, the agreements do not automatically cover every outlet under a given parent, so treatment status is properly observed at the outlet rather than the parent level. Secondly, parent-level clustering would yield only thirteen treated clusters, a configuration in which the wild cluster bootstrap is known to be conservative (MacKinnon & Webb, 2018), with the Country Press Australia collective (sixteen outlets bargaining as one) further inflating cluster-size heterogeneity. The parent-clustered estimates are reported as a robustness check in Appendix A.3.

5.3 Propensity score matching

Agreement outlets are not a random subset of Australian news outlets. They tend to be larger, with more established brands and greater visibility. If large outlets were growing faster than small outlets for reasons unrelated to the NMBC, the DiD estimate would be biased upward. I therefore match treated and untreated outlets on pre-period characteristics using propensity score methods.

The propensity score is estimated from a probit model in which the dependent variable is agreement status and the covariates are three summary statistics of each outlet's traffic in the twelve months prior to treatment (January to December 2020): the mean of log traffic, the standard deviation of log traffic, and the linear trend in log traffic. These covariates capture the level, volatility, and trajectory of an outlet's audience in the pre-period.

Following Smith and Todd (2005), I implement nearest-neighbour matching with replacement and ties, using `psmatch2` in Stata. The resulting propensity score weights are applied to the DiD specification in equation (2). The matched estimates should be interpreted as the average treatment effect on the treated (ATT) under the assumption that, conditional on the propensity score, agreement status is independent of potential outcomes.

5.4 Robustness: Oster bounds

A residual concern with any observational DiD is that unobserved confounders, correlated with both agreement status and traffic growth, could account for the estimated treatment effect. I assess the sensitivity of the results to such confounding using the bounding procedure of Oster (2019).

The procedure computes a parameter δ , defined as the ratio of selection on unobservables to selection on observables that would be required to drive the estimated treatment effect to zero. If $|\delta| > 1$, unobserved factors would need to be more strongly correlated with the outcome than the included controls are, in order to explain away the result. I compute δ by comparing a “short” regression (treatment indicator plus time fixed effects only) with a “long” regression (treatment indicator plus both outlet and time fixed effects), setting $R_{\max} = 1$.

5.5 Robustness: triple difference

A second concern is that the DiD may confound the effect of NMBC agreements with a general tendency for larger outlets to grow faster. If this tendency is not specific to Australia, it can be identified and netted out using outlets in countries without a link-tax regime as an additional control group.

I construct a triple-difference estimator using outlet-level traffic data from Australia, Canada, New Zealand, and the United Kingdom. In each control country, I assign “pseudo-treatment” status to the largest outlets by pre-period traffic, matching the Australian treatment rate of approximately 65 per cent. The specification is

$$Y_{ict} = \beta_7 (AU_c \times Treat_i \times Post_t) + \beta_6 (Treat_i \times Post_t) + \alpha_i + \gamma_{ct} + \varepsilon_{ict} \quad (3)$$

where AU_c is an indicator for Australia, $Treat_i$ is agreement status in Australia (pseudo-treatment in control countries), $Post_t$ is the post-NMBC indicator, α_i are outlet fixed effects, and γ_{ct} are country-by-time fixed effects. The parameter of interest is β_7 , which isolates the Australia-specific treatment effect. The parameter β_6 captures any global tendency for large outlets (treated or pseudo-treated) to grow faster in the post-period, regardless of country.

The identifying assumption is that, absent the NMBC, the difference in traffic growth between large and small outlets would have evolved similarly in Australia and in the control countries. This is plausible insofar as the control countries share a common language, similar media market structures, and exposure to the same global trends in

digital news consumption, but did not have a link-tax regime in place during the treatment window. Canada’s *Online News Act* was not enacted until June 2023, well after the one-year post-treatment window closes. Canada was, however, actively drafting the Act through 2021 and 2022, and the United Kingdom was at earlier stages of similar regulatory discussion. If large outlets in either country were anticipating future payments and gaining traffic accordingly, $\hat{\beta}_6$ would absorb that effect and attenuate $\hat{\beta}_7$. The $\hat{\beta}_6$ estimates reported in Table 9 are not statistically distinguishable from zero in either window, which is consistent with control-country large outlets not having been on an anticipation-driven trajectory during the treatment window.

6 Results

6.1 Outlet-level HHI

Figure 1 plots the outlet-level HHI for Australia over the full sample period, with a vertical line at March 2021. The series is volatile but shows a visible upward spike in the months surrounding the NMBC’s passage, followed by a gradual decline. The month-to-month volatility reflects a combination of news-cycle effects (a single dominant story can briefly inflate one outlet’s share), Semrush measurement noise in monthly estimates, and occasional outlets entering or exiting the index near detection thresholds. Figure 2 overlays the segmented regression fit.

Table 5 reports the segmented regression estimates. The level-shift coefficient is 170 HHI points (95% CI: [105, 235]), about 17 per cent of the pre-treatment intercept of 973. The pre-treatment series was rising at +4.28 HHI points per month (CI: [2.97, 5.58]), capturing the late-pre-period ramp visible in Figure 1.

The trend-change coefficient $\hat{\beta}_3$ is -10.61 (CI: $[-12.06, -9.16]$), large enough to offset the level shift within roughly 16 months. The concentration increase was transitory: a sharp spike followed by reversion to the pre-treatment trajectory. The negative trend coefficient captures compression of the size distribution within the agreement group, as smaller agreement outlets gained share from their flagship stablemates; the mechanism is taken up in Section 6.4.

The HHI begins rising visibly in mid-2020, well before the March 2021 enactment date. The federal government had announced the mandatory bargaining code in April 2020 (Frydenberg & Fletcher, 2020) and the ACCC published a draft in July (Australian Competition and Consumer Commission, 2020), putting large eligible publishers on a clear path to forthcoming transfers and signalling to platforms that the regime was inevitable.




Figure 1: Outlet-Level HHI, January 2017 to June 2024
Vertical line at March 2021 (NMBC passage).

Both sides could plausibly have begun adjusting in this period: publishers investing in capacity ahead of expected payment, platforms favouring agreement-bound publishers in algorithmic surfaces. The pandemic-era flight of news consumption to established brands offers an alternative explanation, but operated globally rather than specifically in Australia, and is therefore largely netted out at the outlet level by the triple-difference design. Closer to the formal date, Meta’s week-long news blockade in late February 2021 temporarily concentrated traffic in search-driven channels that favour large brands. The segmented regression’s pre-trend coefficient absorbs this rising late-pre-period ramp explicitly, so the 170-point level shift estimated at March 2021 captures only the incremental jump at formal enactment. To the extent that any portion of the +4.28 HHI-points-per-month pre-trend was itself policy-attributable rather than pandemic-attributable, the level shift is a lower bound on the policy’s total peak displacement.

6.2 Ownership-level HHI

Platform agreements are signed by parent companies rather than by individual outlets. A natural counterpart to the outlet-level series aggregates traffic to the parent before computing shares. Independent outlets retain their own labels, and the Country Press Australia bloc is treated as one unit since it negotiated as such. Figure 3 plots the resulting




Figure 2: HHI with Segmented Regression Fit
Pre- and post-treatment linear fits from equation (1). Newey–West standard errors with 6 lags.

series.

The ownership-level HHI sits at approximately 2.7 times the outlet-level series, consistent with a market in which a handful of large parents account for most traffic. The top four parents alone command 92 per cent of total Australian news traffic in the pre-treatment period (against 53 per cent for the top four individual outlets). The second column of Table 5 reports the segmented regression. The pre-treatment trend is -7.70 HHI points per month (CI: $[-9.65, -5.75]$); between-parent concentration was declining in the lead-up to the NMBC. At passage, the level shift is $+92.80$ (CI: $[+13.02, +172.59]$), positive but smaller in magnitude than the outlet-level shift of $+170$. The post-treatment trend change is $+13.26$ per month (CI: $[+10.28, +16.24]$), reversing the pre-treatment direction: concentration between parents rose persistently after the NMBC, with no analogue to the outlet-level reversion.

The two HHI series therefore move in opposite directions after the NMBC. At the outlet level, concentration rises briefly and falls back. At the parent level, it rises and keeps rising. Section 6.4 traces the reallocation that produces this pattern.

Both HHI series trend before treatment, so the segmented-regression estimates could be sensitive to the choice of pre-period. Table A.3 estimates equation (1) with pre-periods beginning in each January from 2017 to 2020. As the start moves later, $\hat{\beta}_2$ steepens and

Table 5: Segmented Regression: Effect of NMBC on HHI

	Outlet-level HHI	Ownership-level HHI
Level shift ($\hat{\beta}_1$)	169.85*** (32.99) [105.19, 234.52]	92.80** (40.14) [13.02, 172.59]
Trend change ($\hat{\beta}_3$)	-10.61*** (0.74) [-12.06, -9.16]	13.26*** (1.50) [10.28, 16.24]
Pre-treatment trend ($\hat{\beta}_2$)	4.28*** (0.66) [2.97, 5.58]	-7.70*** (0.98) [-9.65, -5.75]
Constant	973.03*** (21.70)	2607.01*** (31.76)
N	91	91

Notes: Newey-West standard errors with 6 lags in parentheses. 95% confidence intervals in brackets. Treatment date: March 2021. The outlet-level HHI computes shares across all 425 individual outlets; the ownership-level HHI aggregates traffic to the parent company before computing shares. The qualitative pattern is robust to using the top-4 concentration ratio (C4) in place of HHI: outlet-level C4 spikes +3.18 percentage points at passage and reverts at -0.18 points per month; ownership-level C4 spikes +0.97 points and rises at +0.04 points per month, all significant at the 1 per cent level. *** $p < 0.01$, ** $p < 0.05$.

the level shift shrinks from 170 to 111 points at the outlet level. Excluding the anticipation window (April 2020 to February 2021) so that the pre-trend is fitted only on months before the policy was announced increases the level shift to roughly 203 outlet-level and 149 ownership-level points. The ownership-level level shift is not robust to pre-periods beginning in 2019 or later, but the ownership-level finding depends on the post-period trend change, not the level shift. $\hat{\beta}_3$ remains positive and significant in every specification with at least two pre-treatment years.

6.3 Difference-in-differences

Table 6 reports DiD estimates from equation (2) on the propensity-score-matched (preferred) and unmatched samples.

The matched specification is preferred because agreement outlets are systematically larger and more established than the universe of non-agreement outlets; without rebalancing on observables, the comparison conflates the effect of holding an agreement with the pre-existing size gap. The matching procedure (Section 5.3) weights non-agreement outlets by propensity score, balancing on pre-period traffic level, volatility, and trend. The



Figure 3: Ownership-Level HHI, January 2017 to June 2024

PSM-weighted treatment coefficient over the six-month window is 0.376 (bootstrap 95% CI: [0.124, 0.630]), or approximately 46 per cent more organic traffic for agreement outlets ($e^{0.376} - 1 \approx 0.46$). Over the one-year window, the coefficient rises to 0.601 (CI: [0.250, 0.959]), or about 82 per cent. Both regressions use outlet and month fixed effects, with R^2 values above 0.94.

The unmatched estimates are almost identical in magnitude: 0.394 (CI: [0.256, 0.532], or 48 per cent) at six months and 0.644 (CI: [0.421, 0.868], or 90 per cent) at one year. That the headline PSM result barely moves under the simpler specification provides a transparency check: the agreement-traffic premium is not an artefact of the matching procedure. Matching slightly narrows the six-month confidence interval and widens the one-year one, but the point estimates move very little.

The fact that the treatment effect grows from the six-month to the one-year window is informative about the nature of the effect. A transitory shock, such as Meta’s February 2021 news blockade, would produce a large short-run effect that attenuates over time. The opposite pattern obtains here: the divergence between agreement and non-agreement outlets widens with the horizon. This is consistent with a sustained reallocation of traffic, or with agreement outlets using their new resources to build audience over time, rather than with a one-off disruption.

Figures 4 and 5 display the mean log traffic trajectories for agreement and non-

Table 6: Difference-in-Differences: Effect of Platform Agreements on Log Traffic

	PSM-Weighted		Unmatched	
	6-month	1-year	6-month	1-year
Treatment	0.376***	0.601***	0.394***	0.644***
Bootstrap 95% CI	[0.124, 0.630]	[0.250, 0.959]	[0.256, 0.532]	[0.421, 0.868]
$\approx$ % effect	46%	82%	48%	90%
Outlet FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
N	4,940	9,500	5,161	9,925
R^2	0.972	0.942	0.983	0.960
Adj. R^2	0.969	0.939	0.982	0.958

Notes: Dependent variable is $\sinh^{-1}$ (organic traffic). Wild cluster-robust score bootstrap confidence intervals with 99,999 replications (Webb weights, jackknifed residuals), clustered by outlet. Percentage effects computed as $e^{\hat{\beta}} - 1$. *** $p < 0.01$.

agreement outlets over the six-month and one-year windows, weighting observations by the propensity score. The pre-treatment trends are closely parallel, lending credibility to the identifying assumption; the post-treatment divergence is visible and grows with the horizon. The corresponding unweighted plots appear in the appendix (Figures A.1 and A.2).

Figure 6 extends the static estimates by reporting event-study coefficients month by month over the one-year window. The PSM-weighted specification is identical to the headline DiD except that the post indicator is replaced with a full set of event-time interactions, normalised to zero at $t = -1$. Two features are informative. First, the treatment effect rises through month +5 and plateaus at roughly +0.55 log points (approximately 73 per cent more traffic) over the remaining horizon. Secondly, a positive bump appears at months -5 to -3 , peaking at +0.38 (marginally significant) in December 2020; this corresponds to the period after the ACCC draft was published and serious bilateral negotiations were under way, and the segmented regression's pre-trend coefficient absorbs it. Coefficients at months -2 and -1 return to near zero: the reference month is -1 by construction, and month -2 (January 2021) shows no significant deviation, consistent with the anticipation bump being a temporary artefact of the late-2020 negotiation period rather than a persistent pre-trend. The pre-trends visible in the single-country event study are largely common across the four countries in the sample and are netted out by the triple-difference event study reported in Section 6.7.




Figure 4: Mean Log Traffic by Agreement Status, PSM-Weighted (6-Month Window)

6.4 Reconciling the HHI and DiD results

The preceding results present a puzzle. The outlet-level HHI spike is transitory, fully unwound within about 16 months. The ownership-level HHI, by contrast, persists: it shifts upward at NMBC passage and keeps rising afterwards. At the same time, the outlet-level DiD coefficient shows agreement outlets persistently outperforming non-agreement outlets, with the gap widening over the horizon. How are these three findings consistent?

The ownership-level result already does part of the work. Concentration between parents rose persistently after the NMBC; the agreement-holding parents captured share from non-agreement parents and held it. The outlet-level reversion is therefore not a story of agreement holders giving back their gains. The traffic moved between parents and stayed moved; the question is what happened to it within the parents who received it.

Part of the answer is that the outlet-level HHI and the DiD weight outlets very differently. The DiD coefficient averages the treatment effect across outlets, while the HHI is a sum of squared market shares in which a handful of large outlets carry nearly all the weight. If the treatment effect varies systematically with outlet size, the two measures can point in opposite directions: the DiD can register a positive average even as the outlet-level HHI declines.

One possibility is compositional: if the HHI spike was driven by Meta’s February 2021 news blockade (which temporarily zeroed Meta-sourced traffic for all Australian




Figure 5: Mean Log Traffic by Agreement Status, PSM-Weighted (1-Year Window)

outlets), while the DiD captures a separate, sustained post-resolution effect, then the two results reflect different phenomena operating at different time scales. The more interesting possibility is heterogeneity within the treated group, and this is the one the data support.

6.5 Heterogeneity by outlet size

I split agreement outlets into quartiles based on their pre-period mean log traffic (computed over the twelve months from March 2020 to February 2021) and estimate a modified version of equation (2) in which the treatment indicator is interacted with quartile dummies:

$$Y_{it} = \sum_{q=1}^4 \beta_q (Treat_{it} \times Quartile_q) + \alpha_i + \gamma_t + \varepsilon_{it}$$

Table 7 reports the results over the one-year window. Treatment effects fall monotonically with pre-period size, from a roughly fourteenfold gain in Q1 to a statistically negligible effect in Q4 (F -test of $Q1 = Q4$: $F = 94.0$, $p < 0.0001$).

The Q1 growers are dominated by ABC's regional and topical digital properties: ABC Perth, ABC Adelaide, ABC Newsradio, ABC Canberra, ABC Hobart, ABC Darwin, Australia Plus, ABC Queensland, and Radio Australia, among others. News Corp's Queensland community papers (Caboolture News, Brisbane News, the Satellite) and a small




Figure 6: Event Study, PSM-Weighted, 1-Year Window. Coefficients on treated $\times$ event-time interactions; $t = -1$ omitted. 95% CIs use outlet-clustered SEs.

number of regional ACM titles round out the list. The pattern is one in which the major agreement-holding parents grew their previously low-traffic regional and digital sub-properties; ABC, in particular, publicly committed its NMBC revenue to expanding its regional and rural journalism operations (Treasury, 2022). The Q1 estimate is somewhat sensitive to the largest five growers in the quartile: trimming them reduces the coefficient from 2.73 to roughly 2.30 log points (from a fourteenfold to a ninefold gain), though the qualitative pattern survives well beyond standard sensitivity thresholds.

The NMBC's benefits accrued overwhelmingly to the smallest outlets within the agreement group. The largest agreement outlets, which already commanded substantial audiences, gained nothing measurable. This resolves the HHI puzzle. The index is a sum of squared market shares, so it is far more sensitive to movements at the top of the distribution than at the bottom. When small agreement outlets gained traffic, their shares increased from a very low base, contributing negligibly to the sum of squares. Meanwhile, the growth in total market traffic compressed the shares of the largest outlets, whose absolute traffic had not changed, through a pure denominator effect. The HHI fell because the treatment effect was concentrated where its weight in the index was smallest, while the denominator effect eroded the shares that dominate the index. The initial HHI spike is better attributed to the transitory disruption of Meta's February 2021 news blockade;

Table 7: Heterogeneity: Treatment Effect by Pre-Period Size Quartile (1-Year Window)

	Coefficient	Std. err.	$\approx$ % effect
Q1 (smallest treated)	2.734***	(0.279)	+1,440%
Q2	0.932***	(0.216)	+154%
Q3	-0.092	(0.098)	≈ 0
Q4 (largest treated)	0.060	(0.089)	≈ 0
Test $\beta_{Q1} = \beta_{Q4}$:	$F = 94.02, p < 0.0001$		
Outlet FE		Yes	
Month FE		Yes	
N		9,925	
R^2		0.968	

Notes: Quartiles based on pre-period mean $\sinh^{-1}(\text{traffic})$ over the 12 months preceding treatment. Robust standard errors clustered by outlet in parentheses. *** $p < 0.01$.

the sustained reallocation documented in the DiD was, in its concentration consequences, equalising.

The inverse size-gradient also bears on mechanism. Algorithmic boosting by platforms, if it were the primary channel, would favour large, high-traffic outlets whose content generates the most engagement. The opposite pattern is observed. The results are more consistent with a resource channel: agreement payments enabled small outlets to invest in capacity, and their traffic grew substantially from a low base as a result. Large outlets, already operating near their capacity frontier, had less scope for marginal improvement.

6.6 Oster bounds

Table 8 reports the Oster (2019) δ values for the main DiD estimates. In the propensity-score-matched sample, $\delta = 22.46$ for the six-month window and $\delta = 10.93$ for the one-year window. In the unmatched sample, $\delta = 2.52$ and $\delta = 1.65$, respectively. All values exceed unity. Unobserved confounders would need to explain more residual variation in traffic than the outlet fixed effects already do, in order to eliminate the estimated treatment effect. The matched estimates are particularly resistant: the large δ values reflect the fact that matching on pre-period characteristics already accounts for the most obvious sources of selection, leaving little room for further confounding.

Table 8: Oster (2019) Bounds for Omitted Variable Bias

	PSM-Weighted		Unmatched	
	6-month	1-year	6-month	1-year
Short $\hat{\beta}$	0.287	–	4.116	4.198
Short R^2	0.002	–	0.114	0.113
Long $\hat{\beta}$	0.376	0.601	0.394	0.644
Long R^2	0.972	0.942	0.983	0.960
δ	22.46	10.93	2.52	1.65
$R_{\max}$	1			

Notes: δ is the ratio of selection on unobservables to selection on observables required to drive $\hat{\beta} \rightarrow 0$. Short regression: treatment + time FE. Long regression: treatment + outlet FE + time FE. $|\delta| > 1$ indicates the result is unlikely to be explained by omitted variable bias.

6.7 Triple difference

Table 9 reports the triple-difference estimates from equation (3). The Australia-specific treatment effect $\hat{\beta}_7$ is 0.377 over the six-month window (CI: [0.191, 0.563]) and 0.743 over the one-year window (CI: [0.475, 1.010]). In percentage terms, agreement outlets in Australia received roughly 46 per cent more traffic over six months and 110 per cent more over one year, relative to what would be predicted by global size-based trends alone.

The coefficient $\hat{\beta}_6$, which captures the treatment-by-post interaction common to all four countries, is 0.022 ($p = 0.71$) in the six-month specification and -0.087 ($p = 0.21$) in the one-year specification. Neither is statistically distinguishable from zero. Large outlets in Canada, New Zealand, and the United Kingdom did not grow faster than small outlets in the post-NMBC period. The differential growth of agreement outlets in Australia is specific to Australia, and therefore attributable to the policy rather than to a global size-growth correlation.

The triple-difference sample comprises 2,397 outlets across the four countries, observed over the relevant windows (31,161 outlet-month observations in the six-month specification, 59,925 in the one-year specification). The pseudo-treatment assignment in control countries is based on pre-period traffic rank, with the top 65 per cent of outlets by pre-period mean traffic assigned to the pseudo-treated group, matching the approximate treatment rate in Australia.

The country-by-month fixed effects in equation (3) absorb any Australia-wide post-period shock. Were there any effect from platform retaliation against non-agreement outlets, the within-Australia comparison would overstate the treatment effect. Estimating

Table 9: Triple Difference: Australia vs. Canada, New Zealand, United Kingdom

	6-month	1-year
$\hat{\beta}_7$ (AU $\times$ Treat $\times$ Post)	0.377*** (0.095) [0.191, 0.563]	0.743*** (0.137) [0.475, 1.010]
$\approx$ % effect	46%	110%
$\hat{\beta}_6$ (Treat $\times$ Post)	0.022 (0.060) $p = 0.713$	-0.087 (0.069) $p = 0.205$
AU $\times$ Post	0.032 (0.073)	0.042 (0.092)
Outlet FE	Yes	Yes
Country $\times$ Month FE	Yes	Yes
N	31,161	59,925
Outlets	2,397	2,397
R^2	0.976	0.961
Adj. R^2	0.974	0.959

Notes: Pseudo-treatment in control countries assigned to the top 65% of outlets by pre-period mean traffic, matching the approximate Australian treatment rate. Robust standard errors clustered by outlet in parentheses. 95% confidence intervals in brackets. *** $p < 0.01$.

the triple difference with just month fixed effects makes the $AU \times Post$ term identified and provides a direct test. The coefficient is positive and significant (+0.180 over six months, +0.231 over one year), meaning Australian non-agreement outlets outperformed their control-country counterparts after the NMBC enactment, the opposite of what retaliation would produce, and consistent with the market-wide traffic growth documented in Section 6.5.

Triple-difference event study and pre-trend robustness

Figure 7 reports the triple-difference event-study coefficients, isolating the Australia-specific dynamics from common cross-country movements. The picture is sharper than the single-country event study. Pre-treatment coefficients are negative and statistically significant from month -12 through month -7, then rise towards zero by month -6. Australian agreement outlets were on a relatively weak trajectory compared with their cross-country peers in the lead-up to the NMBC; netting out the global pattern of large-outlet traffic growth during COVID (which drives the positive pre-trend in the single-country event study) reveals an Australia-specific pre-trend running the opposite direction. A




Figure 7: Triple-Difference Event Study, 1-Year Window. Coefficients on $au \times treated \times$ event-time interactions; $t = -1$ omitted. Sample stacks Australia with Canada, New Zealand, and the United Kingdom; pseudo-treatment in control countries follows the size-rank rule described above. 95% CIs use outlet-clustered SEs.

short-lived positive bump at month -3 ($+0.20$, CI $[0.05, 0.36]$) parallels the anticipation signal in the PSM event study, but subsides by month -2 . Post-treatment coefficients ramp sharply upward, reaching $+0.67$ by month $+12$.

A placebo exercise corroborates the reading. Re-estimating the static DiD with the treatment date set to March 2019 yields a coefficient of -0.371 (SE 0.076 , $p < 10^{-5}$); the same exercise at March 2020 yields -0.213 (SE 0.110 , $p = 0.053$). Both placebo coefficients are negative; whatever bias the pre-trend induces in the static DiD would attenuate the headline estimate, not inflate it. The real treatment coefficient ($+0.532$ on the unmatched sample over the same window) reverses the pre-NMBC direction.

7 Discussion

7.1 Channels

The DiD estimates establish that agreement outlets gained traffic relative to non-agreement outlets after the NMBC. Three channels could account for the gain, and the evidence weighs differently on each.

Resource transfer through parent-company intermediation is the channel with the most direct documentation. NMBC payments flow to parent companies, which choose how and where to spend them. The ABC’s submission to the Treasury’s 2022 review committed that its NMBC revenue would fund expanded regional and rural public interest journalism (Treasury, 2022), and the Q1 heterogeneity results bear this out: ABC’s regional and topical digital properties dominate the list of largest gainers. News Corp’s Queensland community papers and ACM’s regional titles show the same pattern of small-sub-property growth, though without the same level of public documentation. The triple-difference reinforces the channel at the cross-country level: the near-zero $\hat{\beta}_6$ indicates that large outlets did not grow faster than small outlets in the control countries, so the Australian divergence is specific to the policy rather than to a global size-growth correlation.

Capacity headroom at the sub-property level is a permissive condition rather than a competing channel. The small regional and digital outlets that grew most were ones with the most room to grow; the flagship metropolitan outlets of the same parents, already at or near their capacity frontier, gained nothing measurable. Capacity headroom does not by itself drive the result, since non-agreement outlets with similar headroom did not see the gains. It does explain why the resource channel produced compression within parents rather than uniform proportional gains across each parent’s portfolio.

Algorithmic boosting by the platforms (preferential ranking of agreement-outlet content in search results or news feeds) cannot be observed in the data available here, since platform-side ranking and feed data are not accessible. The heterogeneity results weigh against it as the primary channel: if platforms were systematically boosting agreement outlets, the largest outlets would be expected to benefit most, since their content generates the greatest engagement. The data show the opposite pattern. A residual role for algorithmic boosting cannot be excluded; if Google gave agreement outlets preferential placement in Google News panels, that would co-occur with the resource channel and produce the same observable pattern. Disentangling the two would require platform-side data not available here.

Meta’s February 2021 news blockade is a one-off disruption rather than a channel. It could have contributed to the initial HHI spike, but cannot account for effects measured over six-month and one-year windows, particularly given that the treatment effect *grows* with the horizon rather than attenuating.

7.2 Policy implications

The NMBC created a binary partition of the Australian news market: outlets with agreements, and those without them. The agreements delivered substantial traffic gains to the former group, concentrated among the smaller outlets within it. The policy did not, on net, increase market concentration in a lasting way, because the within-group compression of the size distribution offset any between-group divergence.

Large incumbents did not capture the benefits of the NMBC. The worry that they would, leaving independents worse off in both absolute and relative terms, finds no support in the data. Small agreement outlets were the primary beneficiaries. The concern, rather, is for outlets that were excluded from the agreement framework altogether. The NMBC's eligibility threshold and the platforms' preference for negotiating with established brands left many small, community, and independent outlets with no agreement and no recourse. My estimates do not show that these outlets lost traffic in absolute terms, but they fell behind in relative terms, and the secrecy of the negotiation process gave them no basis on which to bargain.

The going-dark constraint is the more fundamental limitation. Sandrini and Somogyi (2024) observe that bargaining codes are only effective to the extent that platforms choose to participate. Meta's permanent withdrawal from Canadian news, and its non-renewal of Australian agreements, demonstrate that participation is voluntary in a meaningful sense. A platform that concludes that news content is not worth the regulatory cost will simply stop carrying it, with consequences for publishers that may exceed any benefit the bargaining code provided. The Australian Treasury's apparent pivot towards a direct levy (Jones & Rowland, 2024) acknowledges this problem, though it introduces others: a levy that is not tied to the display of news content removes the platform's incentive to carry news at all.

7.3 Limitations

Treatment timing is the most material qualification. I do not observe when each individual agreement was struck and so treat March 2021 as the uniform treatment date for all agreement outlets, which introduces classical measurement error in the treatment indicator. If some agreements predate March 2021, my estimates understate the true effect in early post-treatment months; if others post-date it, the post-treatment window includes untreated months for those outlets. The net direction is ambiguous, though the consistency of results across windows and specifications offers some reassurance.

The data inputs themselves carry measurement error. Semrush captures organic search

traffic, not the full audience: readers arriving via direct navigation, social-media links, email newsletters, or mobile applications are measured imprecisely or not at all. If agreement outlets gained disproportionately on those other channels, the estimates here are conservative; if the NMBC primarily redirected search traffic while leaving other channels unchanged, the estimates may overstate the total audience effect. Ownership data, compiled from archival sources, carries its own errors; some ownership changes during the sample period were likely missed, especially for smaller outlets. Such errors would affect the HHI calculation but are unlikely to correlate with agreement status in a way that would bias the DiD.

The triple-difference design rests on pseudo-treatment assignment in control countries, based on pre-period traffic rank. This is a reasonable proxy for the characteristics that predict agreement status in Australia, but not a perfect one: agreement status also depended on ownership, bargaining power, and relationships with platform representatives. To the extent that pseudo-treatment misclassifies control-country outlets, the triple-difference estimates will be attenuated.

8 Conclusion

Australia's News Media Bargaining Code produced a transitory spike in news media market concentration, measured by the Herfindahl–Hirschman Index, of approximately 170 points. The spike was followed by a countervailing trend that erased it within roughly 16 months. At the outlet level, the mechanism is clear: outlets whose parent companies secured agreements with Google or Meta received 46 to 82 per cent more organic web traffic than the matched non-agreement comparison group, depending on the time horizon. The unmatched specification yields nearly identical estimates of 48 to 90 per cent, and the result holds in a triple-difference design using outlets from three control countries.

That the aggregate effect was transitory while the outlet-level effect persisted reflects a single underlying pattern: the smaller regional and digital sub-properties of the agreement-holding parents gained substantially while their flagship outlets did not. The NMBC functioned as a partial equaliser within the set of outlets it reached, compressing the size distribution among agreement holders and thereby unwinding the initial concentration increase. The aggregate HHI reverted not because the policy's effects dissipated, but because the effects were concentrated among the smallest agreement outlets, whose gains compressed the share distribution from below.

The distributional question, however, extends beyond the agreement group. Within the eligible universe analysed here, non-agreement outlets fell behind agreement holders in

relative terms. The substantial population of publishers below the AUD\$150,000 turnover eligibility threshold (community newsletters, hyperlocal blogs, niche specialist titles) is largely outwith the sample and is excluded from the bargaining framework by design. Whether they were made actively worse off, or merely failed to share in the gains, is a question the present data cannot answer directly.

As further jurisdictions adopt or contemplate similar regimes, the Australian experience offers a cautionary lesson. A policy that channels platform payments through bilateral negotiation will produce unequal outcomes, shaped by bargaining power and firm size no less than by the platforms' own strategic interests. The going-dark constraint compounds the problem: Meta's withdrawal from Canadian news and its non-renewal of Australian agreements demonstrate that platforms retain an effective exit option that no bargaining code can foreclose. The distributional consequences documented here are not peculiar to the Australian design. They are, rather, features of any regime that relies on voluntary platform participation and private negotiation to sustain public-interest journalism.

References

- Athey, S., Mobius, M., & Pal, J. (2021). *The impact of aggregators on internet news consumption* (NBER Working Paper No. 28746). Cambridge, MA: National Bureau of Economic Research.
- Australian Competition and Consumer Commission. (2019). *Digital platforms inquiry: Final report* (Tech. Rep.). Canberra: Australian Competition and Consumer Commission.
- Australian Competition and Consumer Commission. (2020, July). *News media and digital platforms mandatory bargaining code: Draft legislation and explanatory memorandum* (Tech. Rep.). Canberra: Australian Competition and Consumer Commission.
- Bossio, D., & Barnet, B. (2023). The News Media Bargaining Code: Impacts on Australian journalism one year on. *Policy & Internet*, 15(4), 611–626.
- Calzada, J., Duch-Brown, N., & Gil, R. (2023). *Do search engines increase concentration in media markets?* (CESifo Working Paper No. 10671). Munich: CESifo.
- Calzada, J., & Gil, R. (2020). What do news aggregators do? evidence from Google News in Spain and Germany. *Marketing Science*, 39(1), 134–167.
- Easton, W. (2021, February). *Changes to sharing and viewing news on Facebook in Australia*. Meta Newsroom.
- European Parliament and Council. (2019). *Directive (EU) 2019/790 of the European Parliament and of the Council on copyright and related rights in the Digital Single Market*. Official Journal of the European Union, L 130/92.
- Freimane, M. (2025). *Substituting away? The effect of platform bargaining regulation on content display* (MSI Discussion Paper No. 2506). Leuven: KU Leuven.
- Frydenberg, J., & Fletcher, P. (2020, April). *ACCC to develop mandatory code of conduct for digital platforms*. Joint media release, Treasurer and Minister for Communications, Canberra.
- Gao, P., Lee, C., & Murphy, D. (2020). Financing dies in darkness? The impact of newspaper closures on public finance. *Journal of Financial Economics*, 135(2), 445–467.
- Grueskin, B. (2022, March). Australia pressured Google and Facebook to pay for journalism. Is America next? *Columbia Journalism Review*.
- Jeon, D.-S., & Nasr, N. (2016). News aggregators and competition among newspapers on the internet. *American Economic Journal: Microeconomics*, 8(4), 91–114.
- Jones, S., & Rowland, M. (2024, December). *Albanese Government to establish News Bargaining Incentive*. Joint media release, Assistant Treasurer and Minister for Communications, Canberra.

- MacKinnon, J. G., Nielsen, M. Ø., & Webb, M. D. (2023). Cluster-robust inference: A guide to empirical practice. *Journal of Econometrics*, 232(2), 272–299.
- MacKinnon, J. G., & Webb, M. D. (2018). The wild bootstrap for few (treated) clusters. *Econometrics Journal*, 21(2), 114–135.
- Majó-Vázquez, S., Cardenal, A. S., & González-Bailón, S. (2017). Digital news consumption and copyright intervention: Evidence from Spain before and after the 2015 “link tax”. *Journal of Computer-Mediated Communication*, 22(5), 284–301.
- Media Ecosystem Observatory. (2024). *Old news, new reality: A year of Meta’s news ban in Canada* (Tech. Rep.). Montreal: McGill University and University of Toronto.
- Meta. (2024, February). *An update on Facebook News*. Meta Newsroom.
- Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708.
- Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, 37(2), 187–204.
- Parliament of Australia. (2021). *Treasury laws amendment (News media and digital platforms mandatory bargaining code) act 2021 (Cth)*.
- Sandrini, L., & Somogyi, R. (2024). *News media bargaining codes* (ZEW Discussion Paper No. 24-043). Mannheim: ZEW Leibniz Centre for European Economic Research.
- Sims, R. (2022). *Instruments and objectives: Explaining the News Media Bargaining Code*. Judith Neilson Institute for Journalism and Ideas, Sydney.
- Smith, J. A., & Todd, P. E. (2005). Does matching overcome LaLonde’s critique of nonexperimental estimators? *Journal of Econometrics*, 125(1–2), 305–353.
- Song, Y., & Manchanda, P. (2025). Frontiers: Does carrying news increase engagement with non-news content on social media platforms? *Marketing Science*, 44(4), 748–759.
- Treasury. (2022). *News media and digital platforms mandatory bargaining code: The code’s first year of operation* (Tech. Rep.). Canberra: Australian Government.

Appendix

A.1 Propensity Score Matching Diagnostics

Table A.1 reports the probit model used to estimate propensity scores for matching treated (agreement) and untreated outlets. The covariates are computed from the twelve months immediately preceding the treatment date (January to December 2020).

Pre-period mean log traffic is the strongest predictor of agreement status, with a coefficient of 0.154 ($p < 0.001$). This confirms the selection concern: agreement outlets are, on average, substantially larger than non-agreement outlets. Pre-period volatility (the standard deviation of log traffic) also predicts agreement status, with a coefficient of 0.828 ($p < 0.001$). The pre-period linear trend is positive but not statistically significant ($p = 0.425$), indicating that agreement and non-agreement outlets were not on systematically different growth trajectories in the year before treatment.

The pseudo R^2 of the probit is 0.200, suggesting that the three covariates capture a meaningful share of the variation in agreement status without overfitting. Nearest-neighbour matching with replacement and ties produces propensity score weights that are applied to the DiD specification. Of the 397 outlets in the pre-period sample, 17 lack a sufficiently close match and receive zero weight. The matched estimates reported in Tables 6 and 8 use these weights.

Table A.1: Propensity Score Estimation: Probit Model

	Agreement Status
Pre-period mean (log traffic)	0.154*** (0.018)
Pre-period std. dev. (log traffic)	0.828*** (0.189)
Pre-period trend (log traffic)	0.596 (0.747)
Constant	−0.947*** (0.166)
<i>N</i>	397
Log likelihood	−205.24
Pseudo R^2	0.200

Notes: Pre-period covariates computed over January–December 2020. Matching implemented via nearest-neighbour with replacement and ties (`psmatch2`). *** $p < 0.01$.

A.2 Unweighted DiD Figures

Figures A.1 and A.2 show the unweighted DiD event-study plots that correspond to the PSM-weighted versions in the main text (Figures 4 and 5). The pre-treatment trends are visibly less parallel in the unweighted sample, which motivates the propensity-score matching procedure described in Section 5.3. The post-treatment divergence is of similar magnitude across the two samples.




Figure A.1: Mean Log Traffic by Agreement Status, Unweighted (6-Month Window)




Figure A.2: Mean Log Traffic by Agreement Status, Unweighted (1-Year Window)

A.3 Parent-Clustered DiD Robustness

Table A.2 reports the DiD estimates from equation (2) under wild cluster bootstrap inference clustered at the parent-company level rather than the outlet level. Point estimates are unchanged from Table 6; only the inference differs.

Under parent clustering, only thirteen of the eighty-six total clusters are treated. The wild cluster bootstrap is known to be conservative in this configuration (MacKinnon & Webb, 2018), and the Country Press Australia collective (sixteen outlets bargaining as one parent cluster) further inflates cluster-size heterogeneity. The two-sided p -values reported below are consistent with this conservatism rather than with a null effect: the unmatched specifications remain significant at the 10% level for both windows, and the headline PSM 1-year estimate sits at $p = 0.093$. I therefore use outlet-level clustering for the main inference in Table 6.

Table A.2: Parent-Clustered DiD Robustness

	PSM-Weighted		Unmatched	
	6-month	1-year	6-month	1-year
Treatment	0.376	0.601	0.394	0.644
Bootstrap p -value	0.115	0.093	0.078	0.068
Outlet FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
N	4,940	9,500	5,161	9,925
Total clusters	86	75	86	75
Treated clusters	13	13	13	13

Notes: Wild cluster bootstrap with 99,999 replications (Webb weights, jackknifed residuals), clustered at the parent-company level. Point estimates match Table 6. Two-sided p -values reported.

A.4 Segmented-Regression Pre-Period Robustness

Table A.3 reports the segmented-regression estimates of equation (1) under alternative pre-periods, holding the post-period fixed. The final row of each panel excludes the anticipation window (April 2020 to February 2021) so that the pre-trend is estimated only on months preceding the policy's announcement. Section 6.2 discusses the pattern.

Table A.3: Segmented-Regression Estimates under Alternative Pre-Periods

Pre-period start	$\hat{\beta}_1$ level shift	$\hat{\beta}_2$ pre-trend	$\hat{\beta}_3$ trend change	N
<i>Panel A: Outlet-level HHI</i>				
2017m1 (baseline)	169.9 (33.0)	4.28 (0.66)	−10.61 (0.74)	91
2018m1	152.2 (33.0)	5.45 (0.94)	−11.78 (0.95)	79
2019m1	124.4 (31.1)	7.95 (0.99)	−14.28 (0.95)	67
2020m1	110.6 (34.7)	10.34 (2.92)	−16.67 (2.76)	55
2017m1, excl. 2020m4–2021m2	202.8 (31.0)	3.32 (0.67)	−9.65 (0.96)	81
<i>Panel B: Ownership-level HHI</i>				
2017m1 (baseline)	92.8 (40.1)	−7.70 (0.98)	13.26 (1.50)	91
2018m1	79.5 (42.5)	−6.71 (1.54)	12.26 (1.94)	79
2019m1	35.3 (37.0)	−2.72 (1.64)	8.27 (2.01)	67
2020m1	19.1 (39.7)	−0.95 (4.12)	6.50 (4.28)	55
2017m1, excl. 2020m4–2021m2	149.0 (43.1)	−9.31 (1.00)	14.86 (1.45)	81

Notes: Newey–West standard errors (6 lags) in parentheses. All specifications end in June 2024 with treatment at March 2021. HHI points; trends in HHI points per month.

A.5 Ownership Distribution

Figure A.3 displays the distribution of outlets across ownership groups as of March 2021. Australian Community Media (ACM) is the largest single group by outlet count, followed by the independently owned outlets, News Corp, and the ABC. The concentration of outlet counts does not map directly onto traffic concentration: News Corp and Nine command a disproportionate share of total traffic relative to their share of the outlet count.




Figure A.3: Outlet Count by Ownership Group (March 2021)