Frontiers in Adaptive Markets Dynamics III: The Directed Network of News Attention and the Estimable Market for News

Avishek Bhandari*[1] and Hitesh Kumar Sahu[1]

[1]Indian Institute of Technology Bhubaneswar

June 2026

Abstract

News is produced, shared, and acted upon over social and informational networks, and a growing body of theory holds that the veracity of news is endogenous to the structure of those networks. Yet the network through which attention to news actually flows is seldom measured directly. We first estimate a directed network of news attention across fifteen topics spanning trade, financial, geopolitical, behavioural, and monetary-policy themes, using transfer entropy on an eight-year daily series drawn from a global open news-monitoring corpus. We then connect the measured network to the equilibrium market for news of Kranton and McAdams (2020), whose central objects, the visibility of true and false stories and a producer's incentive to invest in veracity, become estimable once the attention network is observed.

The estimated network is sparse and almost purely directional, with macro and monetary-policy topics tending to broadcast and geopolitical and trade topics tending to receive. A topic's connectedness is positively and significantly associated with the attention it commands, the empirical counterpart of the theory's prediction that visibility rises with social connectedness, and the network has grown denser over the sample. We report these results as suggestive rather than robust to multiple testing, and we are explicit about what transfer entropy can and cannot identify. The structural veracity predictions the theory delivers, including the sign of the density-veracity relationship, require a credibility-resolved and longitudinal extension. We specify that extension and describe NEURICX, a real-time news-intelligence engine now accumulating a forward panel designed to support it.

1 Introduction

News is not consumed in isolation. It is produced by outlets, shared across social ties, and acted upon by readers whose subsequent choices feed back into what is covered next. A defining feature of this process is that the veracity of news is not a fixed property of a story but an equilibrium

*Corresponding author. avishekb@iitbbs.ac.in. Computations were performed on the NEURICX research engine. We thank seminar audiences for helpful comments. All errors are our own.

outcome of the network over which the story travels. Kranton and McAdams (2020) make this point precise: in their market for news, producers choose story quality in anticipation of how widely a true versus a false story will be seen once consumers decide whether to share it, and the resulting equilibrium veracity depends on the connectedness of the social network and on the precision of the private signals readers hold. Misinformation, persuasion bias, and informational cascades are by now well documented as network phenomena (Vosoughi et al., 2018; Allcott and Gentzkow, 2017; DeMarzo et al., 2003; Bikhchandani et al., 1992; Banerjee, 1992), and the economics of news production has long been modelled as a market in which demand-side and reputational forces shape what gets reported (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2006). What remains scarce is direct measurement of the network through which attention to news actually flows.

This paper does two things. First, we measure a directed network of news attention across macro-financial topics. We treat the daily intensity of worldwide coverage of each topic as an observable series and ask, for every ordered pair of topics, whether past attention to one topic improves the prediction of attention to another beyond that topic's own past. The natural object for this question is directed transfer entropy (Schreiber, 2000), which we estimate with the nearest-neighbour mutual-information estimator of Kraskov et al. (2004) extended to the conditional case by Frenzel and Pompe (2007). The attention series are drawn from a global open news-monitoring corpus that indexes worldwide online coverage and reports, by day, the share of monitored coverage matching a query, a volume-intensity series (Leetaru and Schrodt, 2013). Second, we connect this measured network to the equilibrium market for news of Kranton and McAdams (2020). Their central comparative statics, that the visibility of a story rises in network connectedness and that the sign of the density-veracity relationship depends on how revenue is earned, become estimable objects once the attention network and a live cross-section of classified coverage are in hand.

We situate the exercise within two literatures. The first is the adaptive markets hypothesis (Lo, 2004, 2017), which holds that the informational efficiency of a market is itself a slow-moving, adapting quantity rather than a fixed benchmark in the sense of Fama (1970); the network of news attention is one channel through which that adaptation is mediated, and we expect its structure to shift across market regimes. The second is network economics (Jackson, 2014; Galeotti et al., 2010), which supplies both the centrality measures we apply to the attention network (Katz, 1953; Bonacich, 1987) and the strategic foundations for treating veracity as a network equilibrium (Ballester et al., 2006; Bramoullé et al., 2014; Jackson and Zenou, 2015).

The measured network has three features worth previewing, all reported here only as they appear in the empirical sections. The attention network is sparse and strongly directed: over the full sample its density is 0.0857 (18 of 210 ordered pairs are significant), and its mean edge asymmetry is 1.0, meaning every significant link runs in a single direction. Monetary-policy, macro and behavioural topics broadcast, with inflation, recession and supply-chain attention the leading out-strength sources, while geopolitical and trade topics receive, with election the dominant attention sink. The communities recovered from the attention network only partly track the five transmission channels (full-sample normalised mutual information 0.340), and this agreement is higher in the sparser early eras and falls as the network densifies. We also

report an honest null: in a live cross-section of classified coverage we detect no cross-sectional visibility premium for source credibility, with mean salience essentially flat across credibility tiers (the established-over-low gap is +0.0004). This null is informative. It says that within a single snapshot, salience tracks topic rather than source, and it motivates a structural test of the theory's visibility premium on a propagation-based measure rather than a contemporaneous salience measure.

That structural test requires a longitudinal instrument. We therefore describe a forward research programme built on a real-time news-intelligence pipeline that classifies articles into the five transmission channels with a salience score and a risk-sentiment reading, and persists each run to a cloud data warehouse, accumulating a longitudinal panel. As of the present writing the panel is live and accumulating but not yet deep enough for longitudinal transfer-entropy estimation; the multi-year open corpus supplies the estimation base today, and the panel is the instrument through which the density-veracity and visibility-premium predictions of the theory will be tested as it matures.

The remainder of the paper proceeds as follows. Section 2 sets out the market-for-news model and its primitives. Section 3 positions our contribution within the literatures on news, social learning, and networks. Section 4 derives the visibility premium that bridges the model to measurement, and Section 5 states the equilibrium and the density-veracity predictions we take to the data. Section 6 describes the data sources and the estimation of directed transfer entropy with its significance assessment. Section 7 reports the measured attention network, its centralities, its community structure against the channel taxonomy, and the connectedness-attention test. Section 8 describes the live cross-section, the visibility-premium null, and the forward longitudinal panel. Section 9 discusses what the estimates do and do not establish, and Section 10 concludes.

2 The market for news and the news-attention network

This paper joins two objects that have lived apart. The first is a theory of how a social network shapes the veracity of news, due to Kranton and McAdams (2020): a market in which producers choose story quality, consumers choose what to share, and the structure of who observes whom determines whether truth or falsehood becomes visible. The second is an estimable directed network, recovered from the daily intensity of news attention across a set of channel-tagged topics. Our claim is that the second is an empirical shadow of the first. The theory speaks of a directed observation network over which stories propagate; the data give us a directed flow of attention across the issues those stories are about. We set out both objects here so that the comparative statics of the theory can be confronted with the structure of the estimated network in the sections that follow.

2.1 The market for news (Kranton-McAdams primitives)

We state the primitives of Kranton and McAdams (2020) faithfully, retaining their notation. There are N consumers, of whom M generate revenue for producers (through views or through

actions, depending on the revenue model). Producers form a continuum of unit mass. A producer chooses story quality. A low-quality story is costless and is false with probability one; a high-quality story costs $c_R > 0$ and is true with probability one, where the reporting cost is drawn from a distribution $c_R \sim H$ on $(0, \infty)$. The share of producers who invest in quality fixes the ex-ante veracity of news, p_0, the probability that a circulating story is true.

Consumers are embedded in a directed social network. A link $i \to j$ means that consumer i observes what consumer j shares, that is, i follows j. The network is d-regular: each consumer follows exactly d others. The parameter d is the model's measure of social connectedness, and it is the object whose comparative static we take to the data.

The market operates over three phases $t \in \{0, 1, 2\}$. At $t = 0$ producers choose quality, fixing p_0; the story is broadcast and seen by each consumer independently with probability $b \in (0, 1]$. At $t = 1$ each consumer who has seen the story receives a private signal $s_i \in \{T, F\}$ of its veracity, with precision

$$\mathbb{P}(s_i = T \mid \text{true}) \;=\; \mathbb{P}(s_i = F \mid \text{false}) \;=\; \rho, \qquad \rho \in (\tfrac{1}{2}, 1), \tag{1}$$

and decides whether to share, comparing the posterior that the story is true against a sharing threshold normalised to $1/2$. At $t = 2$ consumers act on the stories they have received through the network, comparing their posterior against an action threshold also equal to $1/2$. A consumer's sharing behaviour is summarised by a sharing rule $\mathbf{z} = (z_T, z_F)$, the probabilities of sharing after a good signal $s_i = T$ and after a bad signal $s_i = F$ respectively. The Bayesian posteriors after each signal, $p_1(T; p_0)$ and $p_1(F; p_0)$, govern the best-response sharing rule: consumers never share when veracity is low, filter (share only after a good signal) at intermediate veracity, and always share when veracity is high, with the switch points located at $p_0 = 1 - \rho$ and $p_0 = \rho$.

The bridge between network structure and producer incentives is the relative visibility of true and false stories. A producer is rewarded only when a story reaches the revenue-generating consumers, and that reach depends on how the sharing rule interacts with signal precision and connectedness. We write $V_T(\mathbf{z})$ for the probability that a true story becomes visible and $V_F(\mathbf{z})$ for the probability that a false story becomes visible, and we define the visibility premium informally as

$$\Delta V(\mathbf{z}) \;=\; V_T(\mathbf{z}) - V_F(\mathbf{z}), \tag{2}$$

the differential reach of truth over falsehood. The premium is what a producer internalises when deciding whether to bear the reporting cost, and both V_T and V_F are increasing in connectedness d: a more connected audience makes any shared story more visible. The full expressions, the producer's best response, and the resulting equilibrium veracity are developed in Section 4 and Section 5. For the present purpose, two qualitative features matter. First, news veracity can never exceed the signal precision ρ, so some false news circulates in any equilibrium. Second, the sign of the relationship between connectedness and equilibrium veracity is not fixed: under a views-based revenue model it turns negative once the network is very dense (even false stories then spread widely), whereas under an actions-based model with high connectedness producers face a strong incentive to invest, because consumers' inferences sharpen as they observe more

neighbours. The connectedness-veracity sign is therefore an empirical question, and it is the question that organises Frontiers III.

2.2 The estimable news-attention network

The theory's directed observation network is latent: we do not see who follows whom in the population, nor which stories any individual shared. What we can observe, at scale and over years, is the intensity of public attention to the issues that stories are about. We take this from a global open news-monitoring corpus that indexes worldwide online coverage and reports, by day, the share of monitored coverage matching a query, a volume-intensity series, following the data conventions of Leetaru and Schrodt (2013). Each query is a topic, and we read the topic as a "story" or "issue" of the market for news in the sense of Kranton and McAdams (2020): an object of coverage that competes for attention and propagates through the audience.

Let $\mathcal{N}$ be the set of channel-tagged topics. We planned eighteen topics and analyse fifteen; three were not retrievable from the open corpus within the collection window and are reported as such rather than silently dropped (Section 6). Each topic is assigned to one of five transmission channels, the contagion-channel taxonomy that partitions $\mathcal{N}$ into $\mathcal{C} = \{$Trade, Financial, Geopolitical, Behavioural, Monetary Policy$\}$. The assignment of the analysed topics is

- **Trade**: tariff, trade war, supply chain, export controls;
- **Financial**: stock market, financial stability;
- **Geopolitical**: sanctions, military conflict, election, geopolitical risk;
- **Behavioural**: consumer confidence, recession, market volatility;
- **Monetary Policy**: inflation, central bank.

Every channel retains at least two analysed topics, so the partition remains informative for the community comparisons of Section 7.

The empirical object is a directed weighted adjacency matrix $\mathbf{W}$ over $\mathcal{N}$. We define W_{ij} as the directed transfer entropy from topic i to topic j, the reduction in uncertainty about tomorrow's attention innovation in topic j once the history of topic i is known, beyond what topic j's own past already explains (Schreiber, 2000). Entry W_{ij} is set to the estimated transfer entropy when the directed link survives a significance screen, and to zero otherwise; the estimator and the screen are specified in Section 6. The interpretation is direct: a positive W_{ij} records that attention to topic i today carries predictive content for attention to topic j tomorrow, the estimable counterpart of one story's propagation conditioning what the audience attends to next.

Two summaries of $\mathbf{W}$ carry the economic content. The out-strength of a topic,

$$s_i^{\text{out}} = \sum_{j \neq i} W_{ij}, \tag{3}$$

measures how much a topic broadcasts: the total directed attention flow it sends to the rest of

the network. The in-strength,

$$s_j^{\text{in}} = \sum_{i \neq j} W_{ij}, \tag{4}$$

measures how much a topic receives: the total directed flow it draws in, identifying attention sinks. Broadcasting and receiving are the empirical analogues of the asymmetric position a topic occupies in the latent observation network of the theory, and the contrast between them, which topics originate attention and which absorb it, is the structural fact we report in Section 7.

This mapping makes the bridge to the theory explicit. The topics of $\mathcal{N}$ are the stories and issues of the market for news; the directed attention flow encoded in $\mathbf{W}$ is the empirical shadow of the directed observation network over which those stories propagate; out-strength and in-strength stand in for the broadcasting and receiving roles that the theory assigns to positions in that network; and the connectedness of a topic in $\mathbf{W}$ is the estimable counterpart of the social connectedness d whose visibility comparative static the theory predicts. The remainder of the paper develops the theoretical side of this bridge (Sections 4 and 5) and the empirical side (Sections 6 and 7), and reads the two against one another.

3 Related literature

Our contribution sits at the intersection of four bodies of work. We review each in turn and then state where this paper adds.

Network games and centrality. A first strand studies how the architecture of interaction shapes equilibrium behaviour. Galeotti et al. (2010) characterise games in which a player's payoff depends on the actions of network neighbours, and Jackson and Zenou (2015) survey the resulting equilibrium and welfare properties. Ballester et al. (2006) show that, in games of strategic complements with local payoff interactions, equilibrium effort is proportional to Bonacich centrality, identifying the player whose removal most reduces aggregate activity. Bramoullé and Kranton (2007) establish the contrasting case of local public goods, where strategic substitutes generate specialisation, and Bramoullé et al. (2014) unify these forces through the spectral properties of the interaction matrix, relating stability and the sign of equilibrium response to its lowest eigenvalue. The centrality measures we use to read the estimated network, Katz status (Katz, 1953) and the family of measures generalised by Bonacich (1987), originate in this literature. We borrow its language: a directed weighted adjacency and the rankings it induces.

The economics of news, media, and misinformation. A second strand treats news as an economic good produced under incentives. Mullainathan and Shleifer (2005) model a market in which outlets slant reporting towards reader priors, and Gentzkow and Shapiro (2006) show how reputational concerns over accuracy discipline, or fail to discipline, that slant. The recent literature on online misinformation documents the empirical stakes: Vosoughi et al. (2018) find that false stories propagate faster and more widely than true ones, and Allcott and Gentzkow (2017) measure the circulation and plausible influence of fabricated news during an electoral cycle. Our theoretical anchor, Kranton and McAdams (2020), places this market on

a social network: consumers observe what their neighbours share, producers choose veracity in anticipation of how far true and false stories will travel, and the visibility premium that true stories enjoy over false ones governs the incentive to invest in quality. This strand supplies the object we wish to measure, but its primitives, sharing decisions and producer investment, are not directly observed in coverage data.

Social learning over networks. A third strand asks what agents infer when information arrives through others. Banerjee (1992) and Bikhchandani et al. (1992) show that sequential observation can produce informational cascades in which private signals are discarded, so that aggregate behaviour need not track fundamentals. Acemoglu et al. (2011) characterise when Bayesian agents observing a network of predecessors do, and do not, converge on the truth. The non-Bayesian tradition is equally instructive: DeMarzo et al. (2003) model persuasion bias under repeated averaging of neighbours' opinions, and Golub and Jackson (2010) establish conditions under which such naive updating delivers, or forfeits, the wisdom of crowds. This work clarifies why the structure of who attends to whom, the directed pattern we estimate, carries information about how beliefs form, while cautioning that propagation and accuracy can move apart.

Information-theoretic network inference. A fourth strand provides the estimator. Schreiber (2000) defines transfer entropy as a model-free, directional measure of predictive information flow, and Marschinski and Kantz (2002) adapt it to the heavy-tailed, weakly dependent series typical of financial data. Reliable estimation in continuous state spaces rests on the nearest-neighbour mutual-information estimator of Kraskov et al. (2004), extended to the conditional form transfer entropy requires by Frenzel and Pompe (2007). To read community structure off the directed weighted network we use the algorithm of Traag et al. (2019), which guarantees well-connected partitions, alongside the directed modularity of Leicht and Newman (2008). This strand is largely methodological and has, to our knowledge, rarely been brought to bear on the news market as an economic object.

Our contribution. We position this paper as a bridge between the second strand and the fourth. The economics of news (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2006; Kranton and McAdams, 2020) delivers a theory in which visibility on a directed social network governs the incentive to report truthfully, but its central quantity, the structure of attention flow, is treated as a primitive. Information-theoretic inference (Schreiber, 2000; Kraskov et al., 2004; Frenzel and Pompe, 2007) delivers a directional, model-free estimator of that flow, but has not been used to render the market for news estimable. We join the two: we estimate a directed network of news attention from an open news-monitoring corpus, read it with the centrality and community tools of the first strand, and treat the comparative statics of Kranton and McAdams (2020), in particular that visibility rises with social connectedness, as testable predictions. The social-learning literature (Golub and Jackson, 2010; Acemoglu et al., 2011) frames our interpretive caution: a more connected, more visible topic need not be a more accurate one, which is precisely the structural test we leave for the forward longitudinal instrument.

4 Visibility, the visibility premium, and connectedness

The bridge from the directed social network to the market for news runs through *visibility*: the probability that a story, once produced and broadcast, reaches a consumer through the sharing chain. Following Kranton and McAdams (2020), fix a sharing rule $\mathbf{z} = (z_T, z_F)$, the probabilities that a consumer who saw a story shares it after a good private signal (z_T) or a bad one (z_F). With broadcast probability $b \in (0, 1]$, private-signal precision $\rho \in (1/2, 1)$, and a d-regular following network in which each consumer follows d others, the visibility of a true and of a false story are, respectively,

$$V_T(\mathbf{z}) = 1 - (1 - b)\big(1 - b(\rho\, z_T + (1 - \rho)\, z_F)\big)^d, \tag{5}$$

$$V_F(\mathbf{z}) = 1 - (1 - b)\big(1 - b((1 - \rho)\, z_T + \rho\, z_F)\big)^d. \tag{6}$$

Each expression is one minus the probability that a consumer neither receives the direct broadcast (probability $1 - b$) nor is reached by any of d followed neighbours. The two differ only in the per-neighbour sharing probability: a true story is shared after a good signal, which arrives with probability ρ when the story is true; a false story triggers the same good signal only with probability $1 - \rho$. The bracketed terms therefore carry $\rho\, z_T + (1 - \rho)\, z_F$ for a true story and $(1 - \rho)\, z_T + \rho\, z_F$ for a false one.

The object that disciplines production is the *visibility premium*

$$\Delta V(\mathbf{z}) = V_T(\mathbf{z}) - V_F(\mathbf{z}), \tag{7}$$

the differential reach a producer earns by reporting truthfully rather than falsely. Revenue accrues from the M consumers who generate value, so a producer with reporting cost c_R produces a true story exactly when $c_R < M\, \Delta V(\mathbf{z})$. Integrating against the reporting-cost distribution H delivers the producer best response, the ex-ante share of true stories

$$p_0(\mathbf{z}) = H\big(M\, \Delta V(\mathbf{z})\big). \tag{8}$$

Equation (8) is the supply side of the market for news: veracity is demand-driven through the visibility premium, which is in turn shaped by how consumers share and by how connected they are.

We now record the two structural properties that the empirical network is built to probe. The first concerns connectedness d; the second concerns the sign of the premium in the region where consumers exercise judgement.

Proposition 1 (Visibility rises in connectedness; the premium is positive under filtering)**.** *Fix $b \in (0, 1]$, $\rho \in (1/2, 1)$, and a sharing rule $\mathbf{z} = (z_T, z_F) \in [0, 1]^2$ for which at least one component is positive.*

(i) Both $V_T(\mathbf{z})$ and $V_F(\mathbf{z})$ are strictly increasing in connectedness d.

(ii) If $z_T > z_F$ then $\Delta V(\mathbf{z}) > 0$. In particular, in the filtering region, where $\mathbf{z} = (1, 0)$, the premium is strictly positive.

Proof. Write $q_T = b(\rho\, z_T + (1-\rho)\, z_F)$ and $q_F = b((1-\rho)\, z_T + \rho\, z_F)$, so that $V_T = 1-(1-b)(1-q_T)^d$ and $V_F = 1-(1-b)(1-q_F)^d$. Because $b \in (0,1]$, $\rho \in (1/2,1)$, and $(z_T, z_F) \in [0,1]^2$ with at least one component positive, both q_T and q_F lie in $(0,1)$, whence $1-q_T$, $1-q_F \in [0,1)$.

(i) For any base $x \in [0,1)$ the map $d \mapsto x^d$ is strictly decreasing, so $(1-q_T)^d$ and $(1-q_F)^d$ are strictly decreasing in d. Since $1-b \geq 0$ and the bracketed power enters with a negative sign, both V_T and V_F are strictly increasing in d. (When $b=1$ visibility is one regardless of d; for $b<1$ the increase is strict.)

(ii) The arguments differ by

$$q_T - q_F = b\big[(\rho-(1-\rho))\, z_T + ((1-\rho)-\rho)\, z_F\big] = b\,(2\rho-1)\,(z_T - z_F).$$

With $\rho > 1/2$ we have $2\rho - 1 > 0$, so $z_T > z_F$ implies $q_T > q_F$, hence $(1-q_T)^d < (1-q_F)^d$ and therefore $\Delta V = V_T - V_F = (1-b)\big[(1-q_F)^d - (1-q_T)^d\big] > 0$. Setting $\mathbf{z} = (1,0)$ gives $z_T = 1 > 0 = z_F$, the filtering best response of Lemma 1 in Kranton and McAdams (2020), so the premium is strictly positive there. □

Part (ii) clarifies why the filtering region is the economically interesting one. When consumers never share ($\mathbf{z} = (0,0)$) or always share ($\mathbf{z} = (1,1)$), the two arguments coincide and $\Delta V = 0$: the network confers no truth-telling reward and producers minimise cost by reporting falsely. Only when consumers condition on their signal, sharing on a good signal and withholding on a bad one, does truthful reporting earn a strictly larger audience, and only then does equation (8) reward quality.

The magnitudes involved are easy to misread, because the maximiser of veracity does not call for unconditional sharing. Kranton and McAdams (2020) show (their Lemma 2) that $\arg\max_{\mathbf{z}} p_0(\mathbf{z}) = (\bar{z}_T, 0)$ for an interior $\bar{z}_T < 1$: dampening good-signal sharing widens the premium by holding down the incidental visibility of false stories. Their footnote-21 calibration, with $b = 1/2$ and $\rho = 2/3$, gives $\bar{z}_T \approx 0.823$ at $d = 5$ and $\bar{z}_T \approx 0.53$ at $d = 9$. The optimal sharing intensity thus falls as the network grows denser: with more neighbours observed, a smaller per-signal sharing probability already secures high true-story visibility, and restraint becomes the better instrument for protecting the premium.

Proposition 1(i) is the comparative static this paper takes to data. It predicts a positive association between a topic's connectedness in the news network and the attention it commands: more-connected topics should be more visible and so attract greater coverage. We do not observe consumer-level following, so we read connectedness off the estimated directed network, using each topic's network strength as the empirical analogue of d and its mean attention level as the analogue of visibility. The prediction, and the honest caveats attaching to a cross-sectional, correlational reading of it, are tested in Section 7.

5 Equilibrium veracity and the density-veracity sign

Section 4 fixed the producer's best response: invest in a true story when the reporting cost falls below the revenue-weighted visibility premium, $c_R < M\,\Delta V(\mathbf{z})$, so that best-response veracity

is $p_0(\mathbf{z}) = H(M\,\Delta V(\mathbf{z}))$. Consumers' sharing rule $\mathbf{z}$ in turn responds to the veracity they anticipate. A market equilibrium is a fixed point of this two-sided map: a sharing rule that is a best response to the veracity it induces, and a veracity that is a best response to that sharing. We state the equilibrium characterisation and its central comparative static, then translate each into a prediction for the directed attention network. All results in this section are due to Kranton and McAdams (2020); we reproduce them to make the bridge to the network estimates explicit.

5.1 Existence, uniqueness, and the veracity ceiling

Theorem 2 (Equilibrium veracity)**.** *A dysfunctional equilibrium, in which consumers never share and producers never invest, always exists. A functional equilibrium, in which true stories are shared and some producers invest, exists if and only if maximal veracity exceeds the never-share threshold, $\bar{p}_0 > 1-\rho$. When it exists, the functional equilibrium is essentially unique, with veracity*

$$p_0^* \;=\; \max\{\,1-\rho,\;\min\{\tilde{p}_0,\;\rho\}\,\}. \tag{9}$$

Equilibrium veracity can never exceed the signal precision, $p_0^* \leq \rho$, so some false news circulates in every equilibrium.

Three features of Theorem 2 carry directly to the empirics. First, coordination is two-sided: the dysfunctional outcome survives because, absent sharing, no producer earns a return on quality, and absent quality, no consumer gains from sharing. A news market can therefore sit in a low-veracity trap that is not a failure of any single participant. Second, the functional equilibrium exists only when the most favourable sharing rule would already push veracity past the never-share threshold $1-\rho$, so that better-informed consumers (higher ρ) widen the region in which a functional market can form. Third, equation (9) caps veracity at ρ: since private signals are imperfect, producers can never make truth-telling so attractive that falsehood disappears. The ceiling ρ, not unity, is the right benchmark against which to read any observed measure of news quality.

5.2 The density-veracity sign

The comparative static of interest is the sign of equilibrium veracity in social connectedness d, the empirical analogue of network density. Both true-story and false-story visibility rise in d (Section 4), so the net effect on quality turns on how the revenue model rewards reach. The source establishes that this sign is not universal but depends on what producers are paid for.

Remark 3 (Density-veracity sign, Kranton and McAdams, 2020)*.* Under views or advertising revenue, producers are paid for reach. In a sparse network more links raise equilibrium quality, because the filtering rule shares true stories preferentially and additional neighbours mainly propagate those. As the network becomes very dense, however, even false stories reach almost everyone, the visibility premium ΔV collapses, and equilibrium veracity is low: the density-veracity sign turns negative at high density. Under actions or voting revenue, producers are paid when consumers act correctly on what they read. Here high connectedness gives produc-

ers a strong incentive to invest, because consumers' inferences sharpen as they observe more neighbours, and the density-veracity sign is positive. The positive case is bounded above by the absence of any wisdom-of-the-crowd gain to the consumer.

The sign of the density-veracity relation is thus a discriminating prediction: the same densification of the attention network should depress veracity in advertising-funded segments and lift it in action-funded segments. This is the structural object the forward panel is built to estimate (Section 8).

5.3 Misinformation and the quantity of true news

The source also asks what a mass of committed misinformation agents does to the market. Let m denote the mass of agents who share only false stories.

Proposition 4 (Misinformation, Kranton and McAdams, 2020)**.** *There is a threshold $\bar{m} > 0$ such that a functional equilibrium exists if and only if $m < \bar{m}$. On that range equilibrium veracity $p_0^*(m)$ is non-increasing in m. The quantity of true news, $(1+m)\,p_0^*(m)$, is non-decreasing in m, and is strictly increasing on two subintervals $(0, m_1)$ and $(m_2, \bar{m})$: a small or a large mass of misinformation can raise the production of true news, because rational consumers respond by filtering more judiciously, lowering the false-signal sharing rate z_F and so raising the relative visibility of true stories.*

Proposition 11 separates two quantities that are easily conflated. The truth share p_0^* falls with misinformation, yet the volume of true news in circulation can rise, because tighter consumer filtering shifts the visibility premium back towards true stories. Counting true items and inferring veracity are therefore not interchangeable, a distinction the panel measurement design must respect.

5.4 Signal precision and gaslighting

A final comparative static concerns the precision ρ of consumers' private signals, the channel through which a campaign to degrade public discernment ("gaslighting") operates. Lowering ρ reduces equilibrium veracity when the market is already informative, $p_0^* > 1/2$, but can raise it when the market is poorly informative, $p_0^* < 1/2$: when consumers share almost indiscriminately, slightly noisier signals make them filter more, which can restore some producer incentive to invest. The welfare effect of degrading signal quality thus depends on the regime the market starts in.

5.5 Mapping the theory to network predictions

Table-free, we group the predictions by what the present data can and cannot reach.

Tested in this paper (density dynamics). Theorem 2 and Remark 3 both turn on network density, and the directed attention network lets us measure how density evolves. The estimated network densifies monotonically across eras, from a density of 0.0429 before 2020 to 0.0524

during the pandemic window and 0.0857 thereafter (Section 7). The accompanying alignment between attention communities and the channel taxonomy is higher in the sparser early eras and falls as the network densifies, with the normalised mutual information moving from about 0.60 to about 0.34. The structural reading is consistent with Remark 3: as coverage intensifies, attention spillovers cross channel boundaries more, the way denser networks let stories of either type reach everyone. The connectedness-attention association we report (Section 7) is the empirical counterpart of the prediction that visibility rises in d, and is consistent with it, though it is cross-sectional and correlational rather than a causal test.

Requiring the forward panel (the veracity sign). The veracity content of the theory, the level p_0^* in equation (9), the density-veracity *sign* of Remark 3, the misinformation comparative statics of Proposition 11, and the precision result above, cannot be read from an attention-volume series, which records how much a topic is covered, not whether the coverage is true. Each requires a measure of veracity paired with a propagation-based visibility measure and a revenue tag (views versus actions) at the level of the news segment. That is precisely the longitudinal instrument assembled in Section 8: a classified, source-attributed panel that accumulates over time and is designed to recover the density-veracity sign once it is deep enough to support directed estimation. Until then we report the density dynamics as confirmed and treat the veracity sign as the central forward hypothesis the programme is built to test.

6 Data and method

6.1 News-attention data

Our estimation base is a panel of daily news-attention series drawn from a global open news-monitoring corpus that indexes worldwide online coverage and reports, by day, the share of monitored coverage matching a query, a quantity we refer to as a volume-intensity series (Leetaru and Schrodt, 2013). Each series measures how much of the world's monitored news on a given day concerns a stated topic, so the panel is a direct, machine-readable proxy for the allocation of collective attention across topics. We treat each topic as a node and ask whether attention to one topic helps predict future attention to another, after accounting for the topic's own past.

We map topics to the five transmission channels of the contagion-channel taxonomy: Trade, Financial, Geopolitical, Behavioural, and Monetary Policy. We planned eighteen channel-tagged topics. Fifteen were retrievable and analysed; three (bond yields, banking crisis, and interest rate) could not be retrieved from the open corpus within the collection window and are reported as such rather than silently dropped. The fifteen analysed topics, grouped by channel, are: Trade (tariff, trade war, supply chain, export controls); Financial (stock market, financial stability); Geopolitical (sanctions, military conflict, election, geopolitical risk); Behavioural (consumer confidence, recession, market volatility); and Monetary Policy (inflation, central bank). Every channel retains at least two nodes, so the channel taxonomy remains a meaningful partition of the analysed node set $\mathcal{N}$ into channels $\mathcal{C}$. The sample spans 2018-01-01 to 2026-06-01 at daily frequency.

6.2 Stationarity

A volume-intensity level series is trending and serially persistent, so we work with the daily log-change of each topic's series, which we read as the day's attention innovation. Differencing in logs renders the series comparable across topics of very different baseline coverage and removes the slow secular drift in monitored volume. Across topics the median augmented Dickey-Fuller p-value is approximately 0.000: every daily log-change series rejects a unit root and is treated as $I(0)$. The directed dependence we estimate below is therefore dependence between stationary innovation series, not a spurious by-product of common trends.

6.3 Directed transfer entropy

We measure directed predictability by transfer entropy (Schreiber, 2000). For an ordered pair of topics $X \to Y$ with attention innovations entering at lag one,

$$\mathrm{TE}_{X \to Y} \;=\; \mathrm{I}(Y_t; X_{t-1} \,|\, Y_{t-1})\,, \tag{10}$$

the conditional mutual information between the future of Y and the past of X given the past of Y. Transfer entropy is a model-free measure of directed information flow: $\mathrm{TE}_{X \to Y} > 0$ exactly when the past of X carries information about the future of Y beyond what Y's own past already supplies, and it is asymmetric in general, so $\mathrm{TE}_{X \to Y}$ and $\mathrm{TE}_{Y \to X}$ can differ. Unlike a regression coefficient it captures nonlinear as well as linear dependence and requires no functional-form assumption.

We estimate the conditional mutual information with the k-nearest-neighbour mutual-information estimator of Kraskov et al. (2004), extended to conditional mutual information by Frenzel and Pompe (2007), under the Chebyshev (max-norm) metric with $k = 4$ neighbours. Series enter in their natural (raw) log-change units, matching the nearest-neighbour core used across this research programme; we note that the max-norm estimator is not scale-invariant, so the choice of raw units is deliberate and is held fixed across all topics and windows. We verified the estimator against three cases with known answers. On a bivariate Gaussian with a closed-form transfer entropy the estimate is 0.1434 against the analytic value 0.1438. On a coupled autoregressive system with a known driving direction the estimator recovers that direction, $\mathrm{TE}_{x \to y} = 0.1791$ far exceeding $\mathrm{TE}_{y \to x} = 0.0083$. On two independent series the estimate is 0.0071, close to the zero implied by independence. These checks confirm that the estimator is accurate where the answer is known and that it does not manufacture flow between unrelated series.

6.4 Significance

A positive transfer-entropy estimate is not by itself evidence of directed structure, because the k-nearest-neighbour estimator has a finite-sample bias and because two series can share linear autocorrelation that inflates apparent flow. We therefore assess each ordered pair against a null calibrated to that pair. For each source series we generate iterative amplitude-adjusted Fourier-transform (IAAFT) surrogates (Schreiber and Schmitz, 1996), which preserve the source

series' amplitude distribution and its power spectrum, and hence its linear autocorrelation, while destroying any nonlinear or cross-series temporal structure. Re-estimating transfer entropy from each surrogate source gives a null distribution against which the observed value is compared. A significant edge therefore reflects directional structure beyond what linear predictability alone would produce. We draw $B = 199$ surrogates per source and define the upper-tail p-value

$$p \;=\; \frac{1 + \#\{\,\text{surrogate TE} \geq \text{observed TE}\,\}}{B + 1}, \tag{11}$$

retaining the ordered pair as a directed edge when $p < 0.05$, that is, when the observed value exceeds the 95th surrogate percentile. We do not use a block bootstrap: resampling blocks reuses data points and so is invalid for a k-nearest-neighbour transfer-entropy statistic, whose value depends on the exact configuration of neighbours. As a multiplicity robustness check we additionally report, in Section 7, the number of edges that survive a Benjamini-Hochberg false discovery rate control (Benjamini and Hochberg, 1995) at $q = 0.05$.

6.5 Network objects

The significant edges define a directed weighted adjacency matrix $\mathbf{W}$, with $\mathbf{W}_{ij}$ equal to the estimated transfer entropy $i \to j$ when that ordered pair is significant and zero otherwise. From $\mathbf{W}$ we read the following objects. The out-strength of a topic is its total outgoing significant transfer entropy, which we interpret as broadcasting, and its in-strength is the total incoming significant transfer entropy, which we interpret as receiving. We rank topics by PageRank and by Katz centrality (Katz, 1953; Bonacich, 1987) to capture standing that accrues through indirect as well as direct links. We summarise the network with its density, the share of the $N(N-1)$ ordered pairs that are significant; its mean edge asymmetry, the degree to which significant links are one-directional rather than reciprocated; and its spectral radius $\lambda_{\max}(\mathbf{W})$, the largest-modulus eigenvalue, which is zero for an acyclic significant-edge network.

To ask whether attention clusters along the same lines as the transmission channels, we partition $\mathbf{W}$ into communities by Leiden modularity maximisation (Traag et al., 2019), which guarantees well-connected communities, and compare the recovered partition to the five-channel taxonomy by the normalised mutual information (NMI) and the adjusted Rand index (ARI). NMI measures shared information between the two partitions, and ARI measures their agreement adjusted for chance, so a value near zero indicates no more agreement than random labelling would give. We report these objects for the full sample and for three eras: pre-2020, the period from March 2020 to January 2022, and post-2022.

6.6 Live cross-section and forward panel

Alongside the multi-year attention series we draw on the NEURICX research engine, a real-time news-intelligence pipeline that classifies individual articles into the same five transmission channels and assigns each a salience score and a risk-sentiment label, persisting every run to a cloud data warehouse as a longitudinal panel. This panel supports a contemporaneous cross-sectional reading of channel salience and source credibility (Section 7) and is the forward

instrument for the structural veracity tests implied by the theory (Section 8). It is live and accumulating but not yet deep enough for longitudinal transfer-entropy estimation, so the multi-year open-corpus series carries the directed-network estimation in this paper while the panel matures.

7 The directed network of news attention

We now estimate the empirical object the model asks for: a directed network among topics, recovered from how attention to one topic anticipates attention to another. The node set is the 15 channel-tagged topics actually retrieved from the global open news-monitoring corpus (of 18 planned; the three series that could not be retrieved within the collection window are reported as such in Section 6, not silently dropped). Each topic's daily volume-intensity series is stationarised by its log-change, and across topics the median augmented Dickey-Fuller p-value is approximately 0.000, so every attention-innovation series rejects a unit root and enters the estimator as an $I(0)$ process. The directed adjacency $\mathbf{W}$ collects the significant transfer entropies, $\mathbf{W}_{ij}$ being the estimated $\mathrm{TE}_{i\to j}$ that survives the surrogate test of Section 6; an edge $i \to j$ means that attention to topic i carries predictive information about attention to topic j beyond j's own past. This is the estimable counterpart of the directed observation network of Kranton and McAdams (2020), in which a link $i \to j$ records that consumer i observes what j shares. We do not claim that topics "follow" one another in their sense; we claim only that the recovered edges are directed, and we let that directedness carry the structural reading.

7.1 Full-sample structure

Over the full sample of 3054 daily log-changes, the network is sparse and strikingly directional. Of the 210 ordered topic pairs, 18 carry transfer entropy that exceeds its surrogate null at the per-pair level (upper-tail $p < 0.05$), a density of 0.0857. The mean edge asymmetry is 1.0: every significant link is one-directional, with zero reciprocal pairs. The significant-edge network is therefore acyclic, and its spectral radius (the largest-modulus eigenvalue $\lambda_{\max}$ of $\mathbf{W}$) is 0.0. Acyclicity is an empirical, not an imposed, property; it says that the recovered predictive structure runs forward without feedback loops at this significance threshold.

We treat the 18 edges as suggestive directed structure rather than as multiplicity-robust claims. Applying the Benjamini-Hochberg procedure (Benjamini and Hochberg, 1995) at $q = 0.05$ leaves 0 edges surviving. This emptiness is in part a power limitation at this scale: with $B = 199$ surrogates per source the smallest attainable per-pair p-value is $1/200 = 0.005$, which exceeds the smallest Benjamini-Hochberg threshold, $0.05/210 \approx 0.0002$, so no edge can clear the most stringent rung of the step-up ladder regardless of its true strength. We report both numbers and read the per-pair edges as informative about direction and ordering, not as evidence of a particular edge surviving family-wise correction.

The strength rankings tell a coherent story about who broadcasts and who receives. Table 1 reports the leading topics by out-strength (the total weight a topic sends, our broadcasting measure), in-strength (the total weight it receives), PageRank, and Katz centrality (Katz,

Table 1: Leading topics by centrality, full sample. Out-strength is the broadcasting measure (total significant transfer entropy sent); in-strength is the receiving measure. PageRank and Katz centrality are reported as in the network literature. All values are from the directed transfer-entropy adjacency **W**.

	Out-strength		In-strength		PageRank		Katz	
Rank	Topic	Value	Topic	Value	Topic	Value	Topic	Value
1	inflation	0.0996	election	0.2339	central bank	0.1775	election	1.0234
2	recession	0.0797	export controls	0.0936	election	0.1707	export controls	1.0094
3	supply chain	0.0778	tariff	0.0922	tariff	0.1548	tariff	1.0092
4	central bank	0.0715	central bank	0.0480	trade war	0.0844	central bank	1.0049
5	stock market	0.0588	trade war	0.0455	export controls	0.0622	trade war	1.0046

1953; Bonacich, 1987). Macroeconomic and monetary-policy topics broadcast: inflation is the strongest source (0.0996), followed by recession (0.0797), supply chain (0.0778), central bank (0.0715), and the stock market (0.0588). Trade and geopolitical topics receive: election is by a wide margin the dominant attention sink (in-strength 0.2339), followed by export controls (0.0936), tariff (0.0922), central bank (0.0480), and trade war (0.0455). The two eigenvector-type measures concur on which topics sit where information accumulates. PageRank is led by central bank (0.1775), election (0.1707), tariff (0.1548), trade war (0.0844), and export controls (0.0622); Katz centrality orders election (1.0234), export controls (1.0094), tariff (1.0092), central bank (1.0049), and trade war (1.0046) first. Read together, the picture is one in which prices and policy (inflation, the central bank, the macro cycle) supply attention that flows into politically salient, slower-moving topics (elections, tariffs, export controls, sanctions-adjacent trade conflict) where it settles.

7.2 Community structure versus the channel taxonomy

A natural question is whether attention spillovers respect the five transmission channels (Trade, Financial, Geopolitical, Behavioural, Monetary Policy) into which the topics are tagged. We recover communities with the Leiden algorithm (Traag et al., 2019), the directed-network reading following Leicht and Newman (2008), and compare the partition with the channel taxonomy using normalised mutual information (NMI) and the adjusted Rand index (ARI). On the full sample the Leiden partition has 4 communities, with NMI 0.340 and ARI 0.007. The reading is deliberately modest. An ARI of 0.007 is essentially chance-level agreement, so the recovered communities do not reproduce the channel labels; the non-trivial NMI is carried by partial within-channel grouping rather than by a clean channel-by-channel correspondence. Attention spillovers, in short, only partly track the channel boundaries an economist would draw a priori. This is itself informative: the directed attention network is not a relabelling of the transmission-channel taxonomy but a distinct object, with cross-channel edges (for example, macro and monetary topics broadcasting into trade and geopolitical receivers) that the taxonomy does not anticipate.

Table 2: Per-era network statistics. Density is the share of significant ordered pairs; asymmetry is the mean edge asymmetry (1.0 denotes wholly one-directional); $\lambda_{\max}$ is the spectral radius. Communities are from the Leiden partition, with NMI and ARI computed against the five-channel taxonomy.

Era	Obs.	Density	Asymmetry	$\lambda_{\max}$	Communities	NMI	ARI
Pre-2020	789	0.0429	1.0	0.0	8	0.600	0.131
COVID	670	0.0524	1.0	0.0476	8	0.610	0.100
Post-2022	1595	0.0857	0.9413	0.0443	5	0.342	−0.044
Full sample	3054	0.0857	1.0	0.0	4	0.340	0.007

7.3 The network densifies, and channel alignment falls

Splitting the sample into the pre-2020, COVID (2020-03 to 2022-01), and post-2022 eras sharpens the picture. Table 2 reports the per-era statistics. The network densifies monotonically across the three eras, from a density of 0.0429 (9 edges) before 2020, to 0.0524 (11 edges) during the pandemic, to 0.0857 (18 edges) after 2022. Directionality is near-total throughout: mean edge asymmetry is 1.0 in the two earlier eras and 0.9413 post-2022, and the spectral radius stays small (0.0, 0.0476, 0.0443 across the eras), so even as edges accumulate the recovered structure remains close to acyclic.

Two readings follow, each tying the empirics back to the theory. First, the monotone densification is the empirical setting for the density comparative static of Section 5: in the model of Kranton and McAdams (2020) the sign of the relationship between connectedness and equilibrium veracity is not fixed but depends on the revenue margin, turning negative at high density on the views margin and positive on the actions margin. The data give us a network that has moved from sparse to dense over the sample, which is precisely the regime in which that sign question becomes live; our cross-section does not pin down the sign, and we return to it as the structural object the forward instrument of Section 8 is built to estimate. Second, channel alignment is higher in the sparser early eras and falls as the network densifies: NMI is 0.600 pre-2020 (ARI 0.131, 8 communities) and 0.610 during COVID (ARI 0.100, 8 communities), then falls to 0.342 post-2022 (ARI −0.044, 5 communities). As coverage intensifies, attention spillovers cross channel boundaries more, and the community partition drifts away from the taxonomy. The early-era agreement should not be over-read either: with few edges the partitions are unstable, and the negative post-2022 ARI is consistent with near-chance alignment once the network is dense.

7.4 Connectedness and attention

The model's central comparative static is that visibility, both of true and of false stories, rises in social connectedness d (Section 4), since a more connected consumer is more likely to encounter any given shared story. The empirical analogue we can test today is whether more network-connected topics command more attention. Across the 15 topics we compute Spearman rank correlations between each topic's network strength and its mean attention level. Total strength correlates positively with attention, $\rho = +0.668$ ($p = 0.0065$), and the relationship is driven

by broadcasting: out-strength gives $\rho = +0.638$ ($p = 0.0105$), while PageRank is weaker and marginal, $\rho = +0.495$ ($p = 0.061$). The receiving-side measures are not significant: in-strength and Katz centrality both give $\rho = +0.258$ ($p = 0.353$). The reading, consistent with the theory's connectedness-visibility prediction, is that more-connected topics, and specifically the topics that broadcast, command significantly more attention. We are careful about what this is and is not. It is a cross-sectional, correlational association, not a causal test of the comparative static; and heavier-covered topics may simply carry more estimable signal, a mechanical channel that would push the same correlations in the same direction and that we acknowledge rather than dismiss. The structural sign of the connectedness-veracity relationship, as opposed to this connectedness-attention association, is left to the forward longitudinal instrument of Section 8, which is designed to measure visibility by propagation rather than by coverage.

8 A live instrument and the visibility-premium analogue

The multi-year attention network estimated above rests on a single, retrospective open corpus, the global news-monitoring series of Section 6. The theory of Section 4, however, is ultimately a theory about article-level production: producers choose veracity in anticipation of how widely a true story will be seen relative to a false one. To move towards that object we built a forward-looking instrument, the NEURICX research engine: a real-time news-intelligence pipeline that reads live articles, assigns each to one of the five transmission channels, and records two article-level quantities. The first is a salience score in the unit interval, the pipeline's estimate of how prominently a story is positioned within current coverage. The second is a risk-sentiment flag ("risk-off" when the story carries an adverse market tone). Each run is persisted to a cloud data warehouse, so that the snapshots accumulate into a forward longitudinal panel rather than being discarded after inspection. The retrospective corpus supplies the estimation base today; the panel is the instrument that, as it deepens, will permit article-level and longitudinal tests of the propagation mechanism the model describes.

8.1 The live cross-section

We report the cross-section the instrument had accumulated at an early stage of its operation. It comprises 92 already-classified articles, drawn over 17 pipeline runs from 82 distinct sources, spanning the fortnight 2026-06-01 to 2026-06-14. This is a proxy cross-section, not a structural test; we describe what it shows and are explicit about what it does not.

Table 3 reports the channel composition. Counts are uneven, as one would expect from live coverage rather than a balanced design. Geopolitical stories carry both the highest mean salience (0.638) and the highest risk-off share (13 of 21), consistent with the prominence the multi-year network assigns to geopolitical and electoral attention. Behavioural topics are the most frequent (42 of 92) but the least salient on average (0.143): such stories are common but rarely positioned as headline coverage. Monetary-policy coverage is the most salient on average (0.833) but thin (three articles), so its mean is reported only for completeness and we draw no inference from it.

Table 3: Live cross-section by transmission channel (92 articles, 17 runs, 82 sources, 2026-06-01 to 2026-06-14). Salience is a mean over articles in the unit interval; risk-off counts articles flagged with an adverse market tone.

Channel	Articles	Mean salience	Risk-off
Trade	9	0.522	2
Financial	17	0.582	5
Geopolitical	21	0.638	13
Behavioural	42	0.143	9
Monetary Policy	3	0.833	1

Each article also carries a source-credibility tier, a coarse outlet-quality proxy assigned over source domains. Of the 92 articles, 58 originate from mixed-credibility outlets, 11 from established outlets, and 6 from low-credibility outlets (the remainder are unclassified at the domain level). The tiering is deliberately coarse and is used only to ask one question: does prominence in current coverage co-move with source credibility?

8.2 The visibility-premium analogue is a null

The model's central comparative static is the visibility premium $\Delta V(\mathbf{z}) = V_T(\mathbf{z}) - V_F(\mathbf{z})$ of Section 4: producers invest in veracity precisely when true stories are more visible than false ones. A genuine structural test of that object requires a propagation-based visibility measure, the chance that a story reaches a consumer through the directed sharing network, observed over time. The live cross-section offers neither: salience is a within-snapshot prominence proxy, not a propagation outcome, and a fortnight is not a longitudinal sample. What it can support is a much weaker, observational analogue: whether prominence in coverage tracks the credibility of the source.

It does not. Mean salience by credibility tier is 0.393 for established outlets, 0.398 for mixed outlets, and 0.393 for low-credibility outlets. The established-minus- low gap is +0.0004, which is negligible. The cross-section therefore detects no visibility premium: within a snapshot, how prominently a story is positioned tracks its topic, as Table 3 shows across channels, and is essentially flat in the credibility of its source. We report this honestly as a null result.

We caution against over-reading it. A flat salience-by-credibility profile is fully consistent with the theory rather than evidence against it, for three reasons. First, the object measured here (editorial prominence within a single snapshot) is not the model's visibility premium (the differential reach of true and false stories through a sharing network); a null on the former carries no implication for the latter. Second, salience reflects topic, and topic is plausibly orthogonal to outlet credibility in a short window, so a near-zero gap is the natural baseline. Third, the tiering is coarse and the established and low tiers are small (11 and 6 articles), so the design has little power to resolve a credibility gradient even if one existed. The null is informative chiefly in what it rules out as a shortcut: a salience-based cross-section is not a substitute for the propagation-based, longitudinal measure the structural test requires.

8.3 The forward panel

NEURICX supplies the forward panel. The pipeline persists each classified run to the cloud data warehouse, so the snapshots accumulate into a panel indexed by run and channel. As of the time of writing the panel holds 92 article rows and 85 channel-snapshot rows across 17 distinct runs, spanning 2026-06-01 to 2026-06-14, a window of 12.6 days. It is live and accumulating.

We are explicit about its present limits. A fortnight of runs is not deep enough for longitudinal transfer-entropy estimation, and a salience snapshot is not a propagation measure; the panel cannot yet bear the structural veracity tests the theory specifies. The multi-year open-corpus network estimated on the series of Section 6 remains the estimation base today. What the panel adds is direction: a forward, article-level record that, as it lengthens, can carry a propagation-based visibility measure and the longitudinal depth that a structural reading of ΔV demands. The contribution of this section is therefore the instrument and the honest baseline it already returns, not a claim resolved on a fortnight of data.

9 Discussion

It is worth being precise about what this paper establishes and, equally, about what it does not. The empirical layer recovers an object that the theory of Kranton and McAdams (2020) treats as primitive but that has not, to our knowledge, been measured directly for the macro-financial news environment: a directed, asymmetric network of attention across the topics that markets monitor, together with how that network moves over time.

What is identified. Three features of this network are estimable and are reported here as findings rather than as conjectures. First, the network is genuinely directed and asymmetric. Over the full sample its density is 0.0857 (eighteen of 210 ordered pairs significant), its mean edge asymmetry is 1.0 (every significant link runs in one direction, with no reciprocal pairs), and its spectral radius is 0.0 (the significant-edge graph is acyclic). Attention propagates, it does so directionally, and a small set of topics act as broadcasters while others act as sinks: across the full sample $\mathbf{W}$ has monetary-policy and macro topics broadcasting (out-strength led by inflation at 0.0996) and an electoral topic absorbing attention (in-strength 0.2339, the dominant sink). Second, the network has dynamics. Its density rises monotonically across the eras we study, from 0.0429 before 2020, to 0.0524 through the pandemic window, to 0.0857 thereafter, and its community structure tracks the channel taxonomy more closely when it is sparse (normalised mutual information near 0.60 in the early eras) than when it is dense (near 0.34 in the most recent era). Attention spillovers cross channel boundaries more as coverage intensifies. Third, and most directly tied to the theory, the connectedness comparative static is borne out. In the model both true-story and false-story visibility, V_T and V_F, are increasing in social connectedness d; the estimable analogue is that more-connected topics command more attention. Across the fifteen analysed topics the rank correlation between total network strength and mean attention is $\rho = +0.668$ ($p = 0.0065$), and it is the broadcasting margin that drives it (out-strength $\rho = +0.638$, $p = 0.0105$). The direction of this association is the one the visibility

mechanism predicts.

What is not identified. The mechanism in Kranton and McAdams (2020) runs from connectedness, through the visibility premium $\Delta V = V_T - V_F$, to producers' incentive to invest in veracity and hence to equilibrium news veracity p_0^*. Our data identify the first link in that chain but not the rest. We do not observe news veracity. We therefore cannot sign the structural relationship between density and veracity, which the theory shows is not even of constant sign: it is positive in sparse networks and where revenue accrues to informed action, and negative as networks become very dense and revenue accrues to views, because even false stories then spread widely. Nor can we estimate the visibility premium ΔV that drives producer investment, since a propagation-based measure of how far true relative to false stories travel requires a veracity label we do not yet have. The one cross-sectional proxy we can construct is null: mean article salience does not vary with a source-credibility tier (established 0.393, mixed 0.398, low 0.393; an established-minus-low gap of +0.0004). We read this as the absence of a detectable cross-sectional visibility premium in a thin snapshot, not as evidence against the mechanism, because salience is the wrong measure for ΔV and a single cross-section has no longitudinal depth. Identifying ΔV and the density-veracity sign requires a veracity or credibility layer attached to the network and a panel deep enough to track how visibility responds to it.

Limitations. Several caveats bound the interpretation. Transfer entropy, estimated here by the k-NN mutual-information approach of Kraskov et al. (2004) extended to the conditional case by Frenzel and Pompe (2007) in the sense of Schreiber (2000), measures directed predictability: that one topic's past sharpens the forecast of another's present beyond that topic's own past. It is not identified causation, and we do not claim it as such. The connectedness-attention association of the previous paragraph is correlational and cross-sectional, and heavier-covered topics may simply carry more estimable signal, a mechanical channel we acknowledge rather than dismiss. The volume-intensity series measure attention, the share of monitored coverage matching a query, not consumption or readership; a topic can be heavily covered without being heavily read. The credibility classification is a language-model outlet-tier judgement over source domains, useful as a first proxy but not a vetted veracity standard. Finally, the assignment of topics to the five transmission channels, and the choice of the topic set itself, are modelling decisions; the modest and era-dependent alignment between attention communities and channels (with adjusted Rand index near zero on the full sample) is consistent with attention being organised partly, but only partly, along the channel lines we impose. The multiple-testing position is correspondingly conservative: at the surrogate scale used here no edges survive a Benjamini-Hochberg correction (Benjamini and Hochberg, 1995), so the eighteen per-pair-significant edges are reported as suggestive directed structure rather than as multiplicity-robust claims.

The forward programme. These limitations are not permanent. The binding constraint on the structural tests is longitudinal depth in a panel that carries both network position and a credibility or veracity signal, and that panel is now accumulating. The live cross-section already

links classified articles to transmission channels with a salience and a risk-sentiment score, and each run is persisted, so the panel grows daily; at the cut-off it holds ninety-two article rows and eighty-five channel-snapshot rows across seventeen runs spanning 12.6 days. It is not yet deep enough for longitudinal transfer-entropy estimation, which is why the multi-year open-corpus series supplies the estimation base today, but it is the instrument designed to deliver what the cross-section cannot. As it matures it will support a propagation-based visibility premium, a direct test of whether connectedness raises veracity or, past some density, lowers it, and an examination of the misinformation comparative statics of Kranton and McAdams (2020, Proposition 1), under which a small or a large mass of misinformation agents can paradoxically raise the quantity of true news as consumers filter more judiciously. The theory specifies these tests; the data to run them are being collected.

Relation to the adaptive-markets view. We close by placing the result within the adaptive-markets perspective of Lo (2017). On that view market behaviour is not a fixed optimum but the outcome of agents adapting to an environment that itself changes. The market for news is a clear instance: the topics that markets attend to, and the directions in which attention flows among them, shift with the macro-financial environment, as the densification and the changing community structure across our eras illustrate. The contribution here is to suggest that the directed attention network is not merely a description of that adaptation but its substrate. The structure through which news travels is what producers respond to when they choose veracity, and is what determines, in the language of Kranton and McAdams (2020), whether more connection improves or degrades the information that markets ultimately price. Measuring that structure, and watching it move, is a prerequisite for understanding how the market for news adapts.

10 Conclusion

This paper set out to measure the directed network of news attention and to read it through an economic model of how news is produced and circulated. The measurement side rests on a global open news-monitoring corpus that reports, by day, the share of monitored coverage matching a query (Leetaru and Schrodt, 2013). Treating the daily log-change of each topic's volume-intensity series as an attention innovation, we estimated directed transfer entropy between topics with the Kraskov-Stoegbauer-Grassberger mutual-information estimator (Kraskov et al., 2004) extended to the conditional case by Frenzel and Pompe (2007), in the sense of Schreiber (2000), and assessed each ordered pair against amplitude-and-spectrum-preserving surrogates (Schreiber and Schmitz, 1996). The interpretive side borrows the estimable bridge of Kranton and McAdams (2020): the visibility of a true and of a false story, and the visibility premium that disciplines a producer's incentive to invest in veracity, are increasing in social connectedness. That comparative static is the object we carried from theory to data.

We are deliberate about what the evidence does and does not say. The directed attention network is sparse, almost entirely one-directional, and acyclic on its significant edges; its broadcasting topics are macroeconomic and behavioural, and its dominant attention sinks are

geopolitical and trade-related. Channel communities, recovered with the method of Traag et al. (2019), align with the contagion-channel taxonomy only partly, and that alignment weakens as coverage intensifies and spillovers cross channel boundaries. More connected topics command more attention, an association driven by broadcasting that is consistent with the connectedness-visibility prediction, but it is cross-sectional and correlational rather than causal, and we say so. The per-pair-significant edges are reported as suggestive directed structure, not as multiplicity-robust claims, because at the present surrogate scale the false-discovery control of Benjamini and Hochberg (1995) is power-limited. In the live cross-section we find no visibility premium for source credibility: within a snapshot, salience tracks topic, not outlet tier. We report that as an honest null.

The forward programme follows from these limits. The structural prediction of Kranton and McAdams (2020) concerns veracity, and veracity is identified by propagation over time, not by a single cross-section. NEURICX, the real-time news-intelligence engine that persists each classified run to a cloud data warehouse, is therefore the instrument we are building: it is live and accumulating, and as it deepens it will support a longitudinal, propagation-based test of the density-veracity sign and of the misinformation comparative statics. The multi-year open-corpus series furnishes the estimation base today; the panel is the bridge to tomorrow's structural test. In the adaptive-markets spirit of Lo (2004, 2017), we read the network of news attention not as a fixed object but as one that reorganises with the market it describes, and we leave its causal mechanics, and the veracity question proper, to the maturing instrument.

A Mathematical appendix: estimation and network statistics

This appendix collects, in self-contained form, the information-theoretic definitions behind our directed transfer-entropy estimator, the surrogate significance procedure, the network statistics reported in the main text, and a proof of the comparative static (Proposition 6) that motivates the connectedness-attention test.

A.1 Entropy, mutual information, and transfer entropy

For a discrete random variable X with probability mass function $p(x)$ the Shannon entropy is

$$\mathrm{H}(X) = -\sum_{x} p(x)\,\log p(x), \tag{12}$$

and for a continuous random vector X with density $f(x)$ the differential entropy is $\mathrm{H}(X) = -\int f(x)\,\log f(x)\,dx$, with the logarithm taken in nats throughout. The mutual information between two random vectors X and Y,

$$\mathrm{I}(X;Y) = \mathrm{H}(X) + \mathrm{H}(Y) - \mathrm{H}(X,Y) = \mathbb{E}\left[\log \frac{f(X,Y)}{f(X)\,f(Y)}\right], \tag{13}$$

is non-negative and equals zero if and only if X and Y are independent. The conditional mutual information given a third vector Z is

$$\mathrm{I}(X;Y \mid Z) = \mathrm{H}(X,Z) + \mathrm{H}(Y,Z) - \mathrm{H}(Z) - \mathrm{H}(X,Y,Z), \tag{14}$$

which measures the dependence between X and Y that remains after conditioning on Z, and is again non-negative with equality if and only if X and Y are conditionally independent given Z.

Transfer entropy quantifies directed, time-asymmetric information flow (Schreiber, 2000). For two stationary series $\{X_t\}$ and $\{Y_t\}$, with unit source and target lags as used throughout (lag 1), the transfer entropy from X to Y is the conditional mutual information between the target's present and the source's past, given the target's own past:

$$\mathrm{TE}_{X\to Y} = \mathrm{I}(Y_t\,;\, X_{t-1} \mid Y_{t-1})\,. \tag{15}$$

Equation (15) isolates the predictive information that the source's recent history carries about the target's present beyond what the target's own recent history already carries; it is directed ($\mathrm{TE}_{X\to Y} \neq \mathrm{TE}_{Y\to X}$ in general) and non-negative, vanishing when X_{t-1} adds nothing to the prediction of Y_t once Y_{t-1} is known. Higher-order conditioning on longer histories is a direct extension; we condition on a single lag in keeping with the daily attention-innovation series of the main text (Marschinski and Kantz, 2002).

A.2 The KSG estimator and its conditional extension

We estimate Equation (15) without binning or a parametric density, using the k-nearest-neighbour mutual-information estimator of Kraskov et al. (2004) extended to conditional mutual information by Frenzel and Pompe (2007). Both rest on the digamma function $\psi(\cdot)$ and on neighbour counts under the Chebyshev (max-norm) metric $\|u\|_\infty = \max_i |u_i|$.

Mutual information (Kraskov et al., 2004). Given n realisations of the joint vector (X,Y), fix a neighbour order k (we use $k=4$). For each point i, let ϵ_i be twice the Chebyshev distance from i to its k-th nearest neighbour in the joint (X,Y) space. Count the points strictly within $\epsilon_i/2$ of i in each marginal,

$$n_x(i) = \#\{\, j \neq i : \|X_j - X_i\|_\infty < \epsilon_i/2 \,\}, \qquad n_y(i) = \#\{\, j \neq i : \|Y_j - Y_i\|_\infty < \epsilon_i/2 \,\}, \tag{16}$$

where the inequality is strict (the strict-neighbour counting convention: a point exactly on the marginal margin is excluded). The estimator is

$$\widehat{\mathrm{I}}(X;Y) = \psi(k) + \psi(n) - \frac{1}{n}\sum_{i=1}^{n}\big[\psi(n_x(i)+1) + \psi(n_y(i)+1)\big]. \tag{17}$$

Conditional mutual information (Frenzel and Pompe, 2007). With the same fixed k in the joint (X,Y,Z) space and ϵ_i again twice the Chebyshev distance to the k-th neighbour there, count strict neighbours within $\epsilon_i/2$ in the (X,Z), (Y,Z), and Z subspaces, denoting these

$n_{xz}(i)$, $n_{yz}(i)$, $n_z(i)$. Then

$$\widehat{\mathrm{I}}(X;Y \mid Z) = \psi(k) - \frac{1}{n}\sum_{i=1}^{n}\big[\psi(n_{xz}(i)+1) + \psi(n_{yz}(i)+1) - \psi(n_z(i)+1)\big]. \tag{18}$$

Setting $X \mapsto X_{t-1}$, $Y \mapsto Y_t$, $Z \mapsto Y_{t-1}$ in Equation (18) yields our transfer-entropy estimate $\widehat{\mathrm{TE}}_{X\to Y}$. The series enter in their natural (raw) log-change units: the max-norm is not scale-invariant, so we do not standardise, matching the research programme's estimator core.

The estimator is consistent as $n \to \infty$ with k fixed (Kraskov et al., 2004; Frenzel and Pompe, 2007). We confirm it on three controls with known answers. On a bivariate Gaussian with an analytic mutual information of 0.1438 nats the estimate is 0.1434. On a unidirectionally coupled autoregressive pair the estimator recovers the direction, $\widehat{\mathrm{TE}}_{x\to y} = 0.1791$ against $\widehat{\mathrm{TE}}_{y\to x} = 0.0083$. On an independent pair the estimate is 0.0071, near the zero floor expected under independence.

A.3 Surrogate significance, the upper-tail p-value, and FDR control

Because the KSG estimator is biased at finite n and that bias depends on the unknown dependence structure, we assess significance against a null calibrated to each series rather than against a fixed threshold. For each ordered pair we generate $B = 199$ surrogate source series by the iterative amplitude-adjusted Fourier-transform (IAAFT) procedure (Schreiber and Schmitz, 1996). Each IAAFT surrogate is constrained to share, to numerical tolerance, the source series' empirical amplitude distribution (its sorted values) and its periodogram (its power spectrum), and therefore its linear autocorrelation function; only the nonlinear and directional structure linking the source to the target is destroyed. The null is thus "no directed coupling beyond what linear predictability and the marginal distribution already explain". Re-estimating Equation (18) on each surrogate source against the unaltered target yields a null distribution $\{\widehat{\mathrm{TE}}^{(b)}\}_{b=1}^{B}$, and the conservative upper-tail p-value is

$$p = \frac{1 + \#\{\, b : \widehat{\mathrm{TE}}^{(b)} \geq \widehat{\mathrm{TE}}^{\mathrm{obs}} \,\}}{B+1}. \tag{19}$$

An edge is retained when $p < 0.05$, that is when the observed estimate exceeds the 95th percentile of its surrogate null. With $B = 199$ the smallest attainable p is $1/200 = 0.005$.

Remark 5 (Block bootstrap is invalid here). A moving-block bootstrap, which resamples blocks of the series *with replacement*, is not a valid null for a k-nearest-neighbour transfer-entropy statistic. Resampling with replacement duplicates observations; under the max-norm metric duplicated points sit at Chebyshev distance zero from one another, collapsing the neighbour radii ϵ_i and inflating the digamma-based counts in Equations (17)–(18). The resulting statistic is biased far above its true value, so a bootstrap null understates significance spuriously. The IAAFT surrogate avoids this by permuting phases while preserving the amplitude distribution exactly, never duplicating a sample.

For multiplicity we report, as a robustness count, the number of edges that survive Benjamini-Hochberg control of the false discovery rate at $q = 0.05$ (Benjamini and Hochberg,

1995): ordering the per-pair p-values $p_{(1)} \le \cdots \le p_{(m)}$ over the m ordered pairs, reject the hypotheses indexed up to the largest r with $p_{(r)} \le (r/m)\, q$.

A.4 Network statistics

Let $\mathbf{W}$ be the directed weighted adjacency on the node (topic) set $\mathcal{N}$, with $\mathbf{W}_{ij}$ equal to the retained transfer entropy $i \to j$ and zero otherwise. The *out-strength* (broadcasting) and *in-strength* (receiving) of node i are the row and column sums

$$s_i^{\text{out}} = \sum_j \mathbf{W}_{ij}, \qquad s_i^{\text{in}} = \sum_j \mathbf{W}_{ij}^{\top} = \sum_j \mathbf{W}_{ji}. \tag{20}$$

Weighted PageRank is the stationary distribution $\boldsymbol{\pi}$ of the row-normalised weight matrix with damping α,

$$\boldsymbol{\pi} = \alpha\, \boldsymbol{\pi}\, \widetilde{\mathbf{W}} + (1-\alpha)\, \frac{1}{|\mathcal{N}|} \mathbf{1}^{\top}, \qquad \widetilde{\mathbf{W}}_{ij} = \mathbf{W}_{ij} / \textstyle\sum_k \mathbf{W}_{ik}, \tag{21}$$

ranking nodes by the attention they accumulate along directed paths. *Katz centrality* (Katz, 1953; Bonacich, 1987) weights walks of length ℓ by an attenuation factor β^{ℓ},

$$\mathbf{c} = \left(\mathbf{I} - \beta\, \mathbf{W}^{\top}\right)^{-1} \mathbf{1}, \tag{22}$$

well defined for $0 < \beta < 1/\lambda_{\max}(\mathbf{W})$, where $\lambda_{\max}(\mathbf{W})$ is the spectral radius; it credits a node both for direct incoming flow and for flow arriving through longer chains.

For community structure we partition $\mathcal{N}$ into groups and score a partition by its (directed) modularity, the share of edge weight falling within groups in excess of the share expected from out- and in-strengths alone (Leicht and Newman, 2008),

$$Q = \frac{1}{w} \sum_{i,j} \left(\mathbf{W}_{ij} - \frac{s_i^{\text{out}}\, s_j^{\text{in}}}{w} \right) \delta(c_i, c_j), \qquad w = \sum_{i,j} \mathbf{W}_{ij}, \tag{23}$$

where c_i is the group of node i and δ is the Kronecker delta. We maximise Q with the Leiden algorithm (Traag et al., 2019), which refines the Louvain move-and-aggregate scheme with a guarantee that recovered communities are internally connected. To compare a recovered partition with the five-channel taxonomy $\mathcal{C}$ we report two label-invariant agreement indices. The normalised mutual information between partitions A and B is

$$\mathrm{NMI}(A, B) = \frac{2\, \mathrm{I}(A; B)}{\mathrm{H}(A) + \mathrm{H}(B)} \in [0, 1], \tag{24}$$

treating group membership as a categorical variable, and the adjusted Rand index $\mathrm{ARI}(A, B)$ counts agreeing pairs of nodes corrected for chance, taking the value zero under random labelling and one for identical partitions.

A.5 Visibility is increasing in connectedness

We now prove the comparative static stated in Section 4: in the directed d-regular social network of Kranton and McAdams (2020), both the true-story visibility V_T and the false-story visibility V_F rise with social connectedness d. Recall the visibility identities, for broadcast probability $b \in (0,1]$, signal precision $\rho \in (1/2,1)$, and sharing rule $\mathbf{z} = (z_T, z_F) \in [0,1]^2$,

$$V_T(\mathbf{z}) = 1 - (1-b)\big(1 - b(\rho\, z_T + (1-\rho)\, z_F)\big)^d, \tag{25}$$

$$V_F(\mathbf{z}) = 1 - (1-b)\big(1 - b((1-\rho)\, z_T + \rho\, z_F)\big)^d. \tag{26}$$

Proposition 6 (Visibility rises in connectedness)**.** *Fix any sharing rule $\mathbf{z} \in [0,1]^2$, any $b \in (0,1]$, and any $\rho \in (1/2,1)$. Then $V_T(\mathbf{z})$ and $V_F(\mathbf{z})$ in Equations* (25)*–*(26) *are non-decreasing in the connectedness d, and strictly increasing whenever $b < 1$ and the relevant sharing probability is positive. Consequently the visibility premium upper bound is governed by the same monotone factor, and the comparative static carries to the producer's best response through Lemma 2.*

Proof. Both identities have the common form $V(d) = 1 - (1-b)(1-q)^d$, where

$$q_T = b(\rho\, z_T + (1-\rho)\, z_F) \quad \text{for } V_T, \qquad q_F = b((1-\rho)\, z_T + \rho\, z_F) \quad \text{for } V_F. \tag{27}$$

The bracket $\rho\, z_T + (1-\rho)\, z_F$ is a convex combination of $z_T, z_F \in [0,1]$ (weights ρ and $1-\rho$ sum to one and are non-negative since $\rho \in (1/2,1)$), so it lies in $[0,1]$; the same holds for the V_F bracket. Hence $q \in [0,b] \subseteq [0,1]$ and $1-q \in [0,1]$. Write $V(d) = 1 - (1-b)(1-q)^d$. Treating d as an integer, as the d-regular network requires, the forward difference is

$$V(d+1) - V(d) = (1-b)\big[(1-q)^d - (1-q)^{d+1}\big] = (1-b)\,(1-q)^d\, q \;\geq\; 0, \tag{28}$$

since $1-b \geq 0$ (because $b \leq 1$), $(1-q)^d \geq 0$, and $q \geq 0$. The difference is strictly positive when $b < 1$ (so $1-b > 0$), $q > 0$ (some sharing, so $1-q < 1$), which gives the strict statement. The same conclusion follows from the continuous extension: with $1-q \in (0,1]$,

$$\frac{\partial V}{\partial d} = -(1-b)\,(1-q)^d \log(1-q) \;\geq\; 0, \tag{29}$$

because $\log(1-q) \leq 0$ on $(0,1]$ while the prefactor $(1-b)(1-q)^d$ is non-negative, with strict inequality under the same conditions $b < 1$, $q > 0$. Applying the argument with $q = q_T$ gives the result for V_T and with $q = q_F$ for V_F. Finally, the producer's best-response veracity is $p_0(\mathbf{z}) = H(M\,\Delta V(\mathbf{z}))$ with $\Delta V = V_T - V_F$, and H non-decreasing; the maximiser $\arg\max_{\mathbf{z}} p_0$ of Lemma 2 inherits the d-dependence of the visibility map, so the comparative static passes through to equilibrium veracity. This monotone prediction, that more-connected attention commands higher visibility, is the relation the connectedness-attention test of Section 4 takes to the data. □

Remark 7 (Reading of the empirical analogue)*.* Proposition 6 predicts a positive association between a topic's network connectedness and its visibility. In the estimated attention network the empirical counterpart of connectedness is a node's total transfer-entropy strength, and

the counterpart of visibility is its realised attention level; the cross-sectional Spearman test of Section 4 examines exactly this monotone association. The proposition makes the sign prediction; it does not assert causality, and the empirical reading retains the mechanical-signal caveat discussed in the main text.

B Economic-theory appendix: the market for news

This appendix gives a self-contained statement of the market-for-news model of Kranton and McAdams (2020) on which the structural reading of our directed attention network rests, and then sets out the formal mapping from the model's primitives to the empirical analogues we estimate. We restate primitives, timing, posteriors, the sharing best response, the visibility equations, the producer best response, and the equilibrium and misinformation results, and we attribute each result to Kranton and McAdams (2020). Where we reproduce a derivation we mark it as our exposition of their result; the results themselves are theirs.

B.1 Primitives

There are N news consumers, of whom $M \leq N$ generate revenue for producers (for instance, the subset whose attention or actions a producer can monetise). Producers form a unit-mass continuum. A producer chooses story quality: a low-quality story is costless and is false with probability one, while a high-quality story costs $c_R > 0$ to report and is true with probability one. The reporting cost c_R is drawn from a distribution H on $(0, \infty)$, so producers with low realised costs are the ones willing to invest in truth.

Consumers are arranged on a directed social network. A directed link $i \to j$ means that consumer i observes what consumer j shares; equivalently, i follows j. The network is d-regular: every consumer follows exactly d others, and d measures social connectedness. The parameter d is the object whose comparative static our network test targets: Kranton and McAdams (2020) show that visibility rises in d, and we ask empirically whether more connected topics command more attention.

Each consumer who reads a story receives a private binary signal $s_i \in \{T, F\}$ of its veracity, with precision $\mathbb{P}(s_i = T \mid \text{true}) = \mathbb{P}(s_i = F \mid \text{false}) = \rho$, where $\rho \in (1/2, 1)$. The signal is informative but imperfect, and ρ governs how sharply a single reader can tell truth from falsehood before observing what others share.

B.2 Timing and posteriors

The economy runs over three phases $t \in \{0, 1, 2\}$.

$t = 0$ **(production).** Producers choose quality. The resulting share of true stories in circulation is the ex-ante news veracity p_0. Each consumer sees a producer's broadcast directly with probability $b \in (0, 1]$.

$t = 1$ **(sharing).** Each consumer who saw a story receives the private signal s_i and decides whether to share it with followers.

$t = 2$ **(action).** Consumers act on the stories they have seen, either directly broadcast or shared by those they follow.

A consumer who sees a story forms a posterior on its veracity. Writing $p_1(s; p_0)$ for the posterior truth probability after signal s given ex-ante veracity p_0, Bayes' rule gives

$$p_1(T; p_0) = \frac{p_0\,\rho}{p_0\,\rho + (1-p_0)(1-\rho)}, \tag{30}$$

$$p_1(F; p_0) = \frac{p_0\,(1-\rho)}{p_0\,(1-\rho) + (1-p_0)\,\rho}. \tag{31}$$

Equations (30)–(31) are the posteriors of Kranton and McAdams (2020); a good signal raises the posterior above p_0 and a bad signal lowers it, with the size of the shift governed by ρ.

The sharing threshold p^S and the action threshold p^A are each normalised to $1/2$: a consumer shares (respectively acts on) a story when the relevant posterior exceeds one half. A pure sharing rule is the pair $\mathbf{z} = (z_T, z_F)$, where z_T is the probability of sharing after a good signal and z_F the probability of sharing after a bad signal.

B.3 The sharing best response (Lemma 1)

Comparing the posteriors (30)–(31) to the threshold $1/2$ partitions the veracity line into three regions. We state the resulting best response, due to Kranton and McAdams (2020).

Lemma 8 (Sharing best response, Kranton and McAdams, 2020)**.** *Let $Z(p_0)$ denote a consumer's best-response sharing rule given ex-ante veracity p_0. Then*

$$Z(p_0) = \begin{cases} (0,0) & \text{if } p_0 < 1-\rho \quad \textit{(never share)}, \\ (1,0) & \text{if } p_0 \in (1-\rho,\ \rho) \quad \textit{(filter)}, \\ (1,1) & \text{if } p_0 > \rho \quad \textit{(always share)}. \end{cases}$$

That is, when prior veracity is low a consumer shares nothing; in the intermediate range a consumer filters, sharing only after a good signal; and when prior veracity is high a consumer shares regardless of signal.

Exposition of the cut-points. The intermediate, filtering region is the economically interesting one, and it is where the comparative statics of misinformation operate below. The lower cut $p_0 = 1-\rho$ is the value at which a good-signal reader is just indifferent: substituting $p_0 = 1-\rho$ into (30) gives $p_1(T; p_0) = 1/2$, so below it even a good signal does not clear the sharing threshold. The upper cut $p_0 = \rho$ is where a bad-signal reader is just indifferent: substituting $p_0 = \rho$ into (31) gives $p_1(F; p_0) = 1/2$, so above it even a bad signal clears the threshold and the reader shares unconditionally. This indifference reasoning is our exposition; the result is Lemma 1 of Kranton and McAdams (2020).

B.4 Visibility (the estimable bridge)

A story's visibility is the probability that a representative revenue consumer ends up seeing it, either by direct broadcast (probability b) or through at least one of the d neighbours that consumer follows. Aggregating over the d-regular network for a fixed sharing rule $\mathbf{z} = (z_T, z_F)$ gives the visibility of true and of false stories,

$$V_T(\mathbf{z}) = 1 - (1-b)\big(1 - b(\rho\, z_T + (1-\rho)\, z_F)\big)^d, \tag{32}$$

$$V_F(\mathbf{z}) = 1 - (1-b)\big(1 - b((1-\rho)\, z_T + \rho\, z_F)\big)^d. \tag{33}$$

Equations (32)–(33) are equations (1)-(2) of Kranton and McAdams (2020). The two expressions differ only in how a true and a false story trigger sharing: a true story is shared after a good signal with probability ρ and after a bad signal with probability $1-\rho$, and the roles reverse for a false story, which is why ρ and $1-\rho$ swap places between (32) and (33). Both V_T and V_F are increasing in connectedness d: a more connected consumer has more chances to receive a story through the network, so any story that some neighbour shares becomes more visible. This monotonicity in d is the comparative static our network test addresses.

The visibility premium is the gap that a producer captures by investing in a true rather than a false story,

$$\Delta V(\mathbf{z}) = V_T(\mathbf{z}) - V_F(\mathbf{z}), \tag{34}$$

and it is positive precisely when readers filter, that is when $z_T > z_F$, so that good signals propagate truth more than bad signals propagate falsehood.

B.5 The producer best response (eq. 3) and the optimal sharing rule (Lemma 2)

A producer earns the visibility premium on the M revenue consumers if and only if a true story is worth its cost. Hence a producer reports truthfully when $c_R < M\,\Delta V(\mathbf{z})$, and the best-response veracity is the mass of producers whose realised cost clears that bar,

$$p_0(\mathbf{z}) = H\big(M\,\Delta V(\mathbf{z})\big). \tag{35}$$

Equation (35) is equation (3) of Kranton and McAdams (2020). Because H is increasing, producer veracity rises one-for-one with the visibility premium ΔV: anything that widens the premium pulls more producers across their cost threshold into truthful reporting.

It follows that the sharing rule which maximises producer veracity is the one that maximises the premium ΔV. Kranton and McAdams (2020) characterise it.

Lemma 9 (Veracity-maximising sharing rule, Kranton and McAdams, 2020)**.** *The sharing rule that maximises best-response veracity sets $z_F = 0$ and an interior good-signal sharing probability, namely* $\arg\max_{\mathbf{z}}\, p_0(\mathbf{z}) = (\bar{z}_T,\, 0)$ *for some $\bar{z}_T \in (0, 1]$. For $b = 1/2$ and $\rho = 2/3$, the maximiser is $\bar{z}_T \approx 0.823$ at connectedness $d = 5$ and $\bar{z}_T \approx 0.53$ at $d = 9$.*

The two numerical values are footnote 21 of Kranton and McAdams (2020). They make the

central tension concrete: the veracity-maximising amount of good-signal sharing falls as the network grows denser (from about 0.823 at $d = 5$ to about 0.53 at $d = 9$), because in a denser network a story that is shared at all becomes widely visible, so unrestrained sharing erodes the premium that rewards truth.

B.6 Equilibrium veracity (Theorem 1) and misinformation (Proposition 1)

An equilibrium is a sharing rule and a producer quality choice that are mutually best responses: consumers play $Z(p_0)$ from Lemma 8 given the veracity they expect, and producers play (35) given the sharing rule. Kranton and McAdams (2020) characterise the equilibrium set.

Theorem 10 (Equilibrium veracity, Kranton and McAdams, 2020)**.** *A dysfunctional equilibrium, in which no consumer shares and no producer invests in truth, always exists. A functional equilibrium, in which sharing and investment are positive, exists if and only if the maximal attainable veracity $\bar{p}_0$ exceeds $1 - \rho$. When a functional equilibrium exists it is essentially unique, with equilibrium veracity*

$$p_0^* = \max\{\, 1 - \rho,\ \min\{\tilde{p}_0,\ \rho\}\,\},$$

where $\tilde{p}_0$ is the veracity sustained under filtering. In particular $p_0^* \leq \rho$: equilibrium news veracity can never exceed the signal precision ρ, so some false news circulates in every equilibrium.

The ceiling $p_0^* \leq \rho$ is the result's sharpest economic content: even at its best, the market for news cannot be more truthful than a single reader's private signal, because sharing aggregates signals of precision ρ rather than ground truth. This is the structural reason that a forward veracity layer, rather than a target of zero misinformation, is the right empirical object to track.

The sign of the relationship between connectedness and equilibrium veracity is not monotone, and Kranton and McAdams (2020) show it depends on what producers monetise. When revenue comes from views or advertising, adding links raises equilibrium quality in sparse networks, because true stories that are filtered in get viewed more often; but as the network becomes very dense, even false stories spread widely, the premium collapses, and the density-veracity sign turns negative. When revenue instead comes from consumer actions such as voting, and connectedness is high, producers have a strong incentive to invest, because a consumer who observes many neighbours draws sharper inferences; here the density-veracity sign is positive, bounded above by the absence of any wisdom of the crowd. The empirical counterpart we can speak to today is the first link in this chain, that visibility rises in connectedness; the revenue-contingent sign of the density-veracity relationship is a forward object for the longitudinal instrument.

Finally, Kranton and McAdams (2020) study an inflow of agents who only ever spread falsehood.

Proposition 11 (Misinformation, Kranton and McAdams, 2020)**.** *Introduce a mass m of misinformation agents who share only false stories. A functional equilibrium exists if and only if $m < \bar{m}$. On that range, equilibrium veracity $p_0^*(m)$ is non-increasing in m, yet the quantity of true news $(1 + m)\, p_0^*(m)$ is non-decreasing in m and can strictly increase on two sub-intervals $(0, m_1)$ and $(m_2, \bar{m})$.*

The mechanism is filtering. A larger mass of false-only sharers raises the chance that a bad-signal share is misinformation, so rational consumers cut back bad-signal sharing (the optimal z_F falls). By (34) and (35) this widens the visibility premium ΔV and pulls more producers into truthful reporting, which is why a small or a large amount of misinformation can, perversely, raise the absolute production of true news even as the share $p_0^*(m)$ falls. A closely related comparative static (Kranton and McAdams, 2020) concerns signal precision: lowering ρ lowers veracity when $p_0^* > 1/2$ but may raise it when $p_0^* < 1/2$, so degrading readers' private discernment is not unambiguously harmful in an already low-veracity market.

B.7 From model objects to estimable analogues

The model is a theory of how a directed sharing network governs which stories become visible and, through visibility, how truthful production is. Our empirical layer estimates a directed network over news topics and a live cross-section of coverage, so the bridge is direct in places and deliberately forward in others. Table 4 states the correspondence and marks each row as tested now or forward.

Three features of the mapping deserve emphasis. First, the connectedness comparative static is the link we can speak to with confidence today: more connected topics command more attention, and the association is carried by out-strength (broadcasting) rather than in-strength, for which the cross-sectional Spearman correlation is $+0.258$ ($p = 0.353$) and is not significant. We read this as consistent with the prediction that visibility rises in connectedness, while noting that it is correlational and cross-sectional, and that heavier-covered topics may carry more estimable signal, a mechanical channel we acknowledge rather than rule out.

Second, the visibility premium maps to a quantity we can already measure but for which we presently find a null: in the live cross-section, salience tracks topic rather than source credibility, and the credibility-to-salience gap across tiers is negligible. We report this as a null that motivates a structural test on the maturing panel using a propagation-based rather than a salience-based measure of visibility, which is closer to V_T and V_F as the model defines them.

Third, the veracity objects, ex-ante p_0, equilibrium p_0^*, the misinformation mass m, and the signal precision ρ, are forward by construction. The ceiling $p_0^* \leq \rho$ of Theorem 10 and the non-monotone quantity-of-truth result of Proposition 11 are statements about how truthful production responds to network structure over time, and they require a longitudinal record of coverage veracity rather than a single attention cross-section. That record is the role of the forward instrument: a live, accumulating, classified panel that supplies, as it deepens, the veracity series the structural tests demand. The estimation base today is the multi-year attention network $\mathbf{W}$, which identifies the directed spillover structure the model takes as its primitive; the panel is the instrument through which the model's veracity predictions become estimable.

Table 4: Mapping from the Kranton and McAdams (2020) model to estimable empirical analogues. "Now" rows are addressed by the present data; "forward" rows require the maturing longitudinal instrument.

Model object	Empirical analogue	Status
Directed follower network ($i \to j$, d-regular)	Directed transfer-entropy attention network $\mathbf{W}$ over 15 topics (W_{ij} the significant directional information flow $i \to j$ in daily attention innovations)	Now
Social connectedness d	Topic network strength (out-strength, in-strength, $\mathbf{W}$-based centralities)	Now
Visibility V_T, V_F (eqs. 32–33)	Attention propagation: how strongly a topic's coverage transmits to and from the rest of the network	Now
Comparative static: V_T, V_F increasing in d	Cross-sectional association of network strength with mean attention across the 15 topics (positive for total strength, $\rho = +0.668$, $p = 0.0065$, and for out-strength, $\rho = +0.638$, $p = 0.0105$)	Now
Visibility premium $\Delta V = V_T - V_F$ (eq. 34)	Salience-credibility cross-section: a credibility-to-salience gap across outlet tiers (the live snapshot shows a flat, null premium, established minus low $= +0.0004$)	Now (null)
Producer best response $p_0(\mathbf{z}) = H(M\,\Delta V)$ (eq. 35)	Outlet investment in truthful coverage as a function of the premium	Forward
Channel taxonomy ($\mathcal{C}$, five transmission channels)	Community structure of $\mathbf{W}$ versus the five-channel partition (modularity, normalised mutual information, adjusted Rand index)	Now
Ex-ante veracity p_0; equilibrium veracity $p_0^* \le \rho$	Forward veracity layer: a classified, time-stamped record of coverage truth status on the longitudinal panel	Forward
Misinformation mass m; quantity of true news $(1+m)\,p_0^*(m)$	Tracked inflow of low-credibility, false-leaning coverage versus the volume of high-credibility coverage on the panel	Forward
Signal precision ρ	Average discriminating power of per-article veracity signals on the panel	Forward

References

Acemoglu, D., Dahleh, M. A., Lobel, I., & Ozdaglar, A. (2011). Bayesian learning in social networks. *Review of Economic Studies*, 78(4), 1201–1236. https://doi.org/10.1093/restud/rdr004.

Allcott, H., & Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, 31(2), 211–236. https://doi.org/10.1257/jep.31.2.211.

Ballester, C., Calvó-Armengol, A., & Zenou, Y. (2006). Who's who in networks. Wanted: the key player. *Econometrica*, 74(5), 1403–1417. https://doi.org/10.1111/j.1468-0262.2006.00709.x.

Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3), 797–817. https://doi.org/10.2307/2118364.

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: a practical and

powerful approach to multiple testing. *Journal of the Royal Statistical Society B*, 57(1), 289–300. https://doi.org/10.1111/j.2517-6161.1995.tb02031.x.

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5), 992–1026. https://doi.org/10.1086/261849.

Bonacich, P. (1987). Power and centrality: a family of measures. *American Journal of Sociology*, 92(5), 1170–1182. https://doi.org/10.1086/228631.

Bramoullé, Y., & Kranton, R. (2007). Public goods in networks. *Journal of Economic Theory*, 135(1), 478–494. https://doi.org/10.1016/j.jet.2006.06.006.

Bramoullé, Y., Kranton, R., & D'Amours, M. (2014). Strategic interaction and networks. *American Economic Review*, 104(3), 898–930. https://doi.org/10.1257/aer.104.3.898.

DeMarzo, P. M., Vayanos, D., & Zwiebel, J. (2003). Persuasion bias, social influence, and unidimensional opinions. *Quarterly Journal of Economics*, 118(3), 909–968. https://doi.org/10.1162/00335530360698469.

Fama, E. F. (1970). Efficient capital markets: a review of theory and empirical work. *Journal of Finance*, 25(2), 383–417. https://doi.org/10.2307/2325486.

Frenzel, S., & Pompe, B. (2007). Partial mutual information for coupling analysis of multivariate time series. *Physical Review Letters*, 99(20), 204101. https://doi.org/10.1103/PhysRevLett.99.204101.

Galeotti, A., Goyal, S., Jackson, M. O., Vega-Redondo, F., & Yariv, L. (2010). Network games. *Review of Economic Studies*, 77(1), 218–244. https://doi.org/10.1111/j.1467-937X.2009.00570.x.

Gentzkow, M., & Shapiro, J. M. (2006). Media bias and reputation. *Journal of Political Economy*, 114(2), 280–316. https://doi.org/10.1086/499414.

Golub, B., & Jackson, M. O. (2010). Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1), 112–149. https://doi.org/10.1257/mic.2.1.112.

Jackson, M. O. (2014). Networks in the understanding of economic behaviors. *Journal of Economic Perspectives*, 28(4), 3–22. https://doi.org/10.1257/jep.28.4.3.

Jackson, M. O., & Zenou, Y. (2015). Games on networks. In H. P. Young & S. Zamir (Eds.), *Handbook of Game Theory with Economic Applications* (Vol. 4, pp. 95–163). Amsterdam: Elsevier. https://doi.org/10.1016/B978-0-444-53766-9.00003-3.

Katz, L. (1953). A new status index derived from sociometric analysis. *Psychometrika*, 18(1), 39–43. https://doi.org/10.1007/BF02289026.

Kranton, R., & McAdams, D. (2020). Social networks and the market for news. Working paper.

Kraskov, A., Stögbauer, H., & Grassberger, P. (2004). Estimating mutual information. *Physical Review E*, 69(6), 066138. https://doi.org/10.1103/PhysRevE.69.066138.

Leetaru, K., & Schrodt, P. A. (2013). GDELT: Global Data on Events, Location, and Tone, 1979–2012. Paper presented at the International Studies Association Annual Convention, San Francisco.

Leicht, E. A., & Newman, M. E. J. (2008). Community structure in directed networks. *Physical Review Letters*, 100(11), 118703. https://doi.org/10.1103/PhysRevLett.100.118703.

Lo, A. W. (2004). The adaptive markets hypothesis: market efficiency from an evolutionary perspective. *Journal of Portfolio Management*, 30(5), 15–29. https://doi.org/10.3905/jpm.2004.442611.

Lo, A. W. (2017). *Adaptive Markets: Financial Evolution at the Speed of Thought*. Princeton, NJ: Princeton University Press.

Marschinski, R., & Kantz, H. (2002). Analysing the information flow between financial time series. *European Physical Journal B*, 30(2), 275–281. https://doi.org/10.1140/epjb/e2002-00379-2.

Mullainathan, S., & Shleifer, A. (2005). The market for news. *American Economic Review*, 95(4), 1031–1053. https://doi.org/10.1257/0002828054825619.

Schreiber, T. (2000). Measuring information transfer. *Physical Review Letters*, 85(2), 461–464. https://doi.org/10.1103/PhysRevLett.85.461.

Schreiber, T., & Schmitz, A. (1996). Improved surrogate data for nonlinearity tests. *Physical Review Letters*, 77(4), 635–638. https://doi.org/10.1103/PhysRevLett.77.635.

Traag, V. A., Waltman, L., & van Eck, N. J. (2019). From Louvain to Leiden: guaranteeing well-connected communities. *Scientific Reports*, 9, 5233. https://doi.org/10.1038/s41598-019-41695-z.

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151. https://doi.org/10.1126/science.aap9559.

Acknowledgements

The author thanks the Google Cloud Education Programs for a research credits award under the Google Cloud Research Credits Programme (credit reference wilsonjessica-485604931), which funds the cloud infrastructure, AI services, and data-warehouse layer of the deployed platform. The computational infrastructure was procured under Institute Seed Grant SP128 (Innovative Approaches for Analyzing Financial Market Dynamics: Computational Financial Analysis with Wavelets, Networks, and Statistical Learning Powered by High-Performance Computing), funded by the Indian Institute of Technology Bhubaneswar (Sponsored Research and Industrial Consultancy, sanction dated 18 November 2023).