Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Beyond Influence Centrality: A Cross-Context Evaluation of Structural Brokerage and Node Irreplaceability in Health-Information Networks

Mkululi Sikosana

mkululi.sikosana@stu.mmu.ac.uk

Manchester Metropolitan University

Research Article

Keywords: digital connectivity, social network analysis, cross-context brokerage, centrality measures, structural substitutability, misinformation propagation, network resilience, node-removal simulation

Posted Date: November 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11641986/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Title: Beyond Influence Centrality: A Cross-Context Evaluation of Structural Brokerage and Node Irreplaceability in Health-Information Networks

Mkululi Sikosana*[1]

[1] Manchester Metropolitan University, Manchester, United Kingdom

* Corresponding author: mkululi.sikosana@stu.mmu.ac.uk (MS)

ORCID: 0009-0008-3244-5764

Abstract

Conventional centrality measures identify structurally prominent nodes in online social networks but do not explicitly quantify their brokerage across distinct information contexts or the availability of alternative nodes performing equivalent bridging roles. This study introduces two complementary context-sensitive centrality measures: Cross-Context Two-Step Brokerage (CCTB), which quantifies the proportion of directed incoming–outgoing interaction opportunities spanning different contexts, and Irreplaceable Context Brokerage (ICB), which weights ordered context-pair connections according to the scarcity of alternative bridging nodes. The measures were evaluated using FibVID and Monant, with exploratory replication using MC-Fake, through chronological holdout evaluation, comparisons with established centrality measures and predictive modelling of subsequent cross-context brokerage. Incorporating CCTB and ICB alongside degree centrality and context breadth increased mean Average Precision from 0.2966 to 0.3542 in FibVID, from 0.0847 to 0.0925 in Monant, and from 0.0655 to 0.1190 in MC-Fake. Under a fixed 5% node-removal budget, simulations examined the potential of centrality-guided targeting to disrupt structural pathways relevant to misinformation diffusion. Degree-based targeting achieved the greatest reduction in future two-step cross-context brokerage opportunities in FibVID (98.73%), context-breadth targeting performed best in Monant (20.77%), while CCTB and ICB achieved the highest reductions in MC-Fake (96.81% each). These findings suggest that context-sensitive brokerage measures provide complementary information for identifying structurally important intermediaries and may inform targeted strategies aimed at limiting misinformation propagation across online social networks. However, the simulated reductions represent disruption of potential structural diffusion pathways rather than empirically observed reductions in misinformation transmission or established causal intervention effects.

Keywords: digital connectivity, social network analysis; cross-context brokerage; centrality measures; structural substitutability; misinformation propagation; network resilience; node-removal simulation

1. Introduction

Research on social diffusion commonly identifies important accounts by how many interactions they attract, whether they occupy short paths between others, or whether influence can flow outward through the network. These objectives are valuable but not identical. A node can be active within one article's discussion without facilitating contact between distinct medical topics. Another can participate in multiple otherwise weakly overlapping conversations while having relatively modest overall degree. If network analysis aims to understand the architecture of health-information exchange, it should distinguish among popularity, diffusion reach, and cross-context brokerage rather than combine them into a single ill-defined influence score [4–10].

Structural centrality measures identify nodes based on connectivity, proximity, intermediary positioning, and diffusion-related influence [1]. Yet they need not distinguish intermediary nodes linking distinct information contexts from nodes whose context-pair connections have few observed substitutes. This study evaluates Cross-Context Two-Step Brokerage (CCTB) and Irreplaceable Context Brokerage (ICB) as complementary, explicitly context-aware measures. The evaluation compares the proposed brokerage measures with established structural and diffusion-oriented centralities.

This question is motivated by the heterogeneity of the two primary datasets. FibVID links Twitter posts to fact-checked claims and provides parent relationships, user identifiers and timestamps [2]. Monant comprises medical articles, associated fact-checking resources and discussion records [3]. However, the available Monant comment records do not provide independently adjudicated article–claim stance annotations capable of establishing whether individual commenters endorse, reject or correct misinformation. Although reply relationships permit the reconstruction of observed discussion networks, they do not establish commenters' beliefs or the veracity of individual replies. Consequently, the proposed brokerage measures can be evaluated against subsequent interaction patterns, but not against unobserved psychological states, verified misinformation transmission or causal intervention effects.

The main contribution is a formal and reproducible operationalisation of cross-context two-step brokerage and the scarcity of other users performing an equivalent context-pair connection. CCTB quantifies the proportion of a node's directed incoming–outgoing interaction opportunities that connect distinct information contexts. ICB discounts a context pair when many other nodes can bridge the same ordered pair. These definitions go beyond static degree and simple context count while remaining distinct from DIC's temporal weighting, MVC's susceptibility-related scores and PC's diffusion-kernel interpretation. Their contribution lies in the explicit representation of directed context transitions and the availability of alternative nodes for each ordered context pair.

To avoid circular evaluation, both metrics are computed on earlier interactions and assessed on later interactions. The evaluation distinguishes generic future activity from the more specific outcome of future cross-context brokerage. A second experiment compares whether the two new features improve predictive Average Precision beyond node degree and the number of previously observed distinct contexts. This design treats a null result as informative: if degree and breadth already capture all useful information, introducing new formulas cannot by itself establish scientific utility.

2. Theory and related research

The interpretation of centrality depends on the structural property being measured. This section reviews connectivity, reachability, shortest-path mediation and recursive prominence before considering temporal influence, propagation-related centrality and vulnerability-based formulations. It then examines structural hole theory and the role of alternative brokerage pathways to establish the specific gap addressed by CCTB and ICB: quantifying the directed connections between labelled information contexts and the availability of other nodes that can form those connections.

The starting point is the distinction between the structural properties that conventional centralities quantify. Degree centrality measures the number of distinct immediate connections and is widely used as an indicator of local prominence [7–9]. Its central limitation in this study is that a highly connected node can interact exclusively within the context of a single claim or article. Closeness

centrality measures directed path accessibility, but its meaning depends on reachability, treatment of disconnected nodes and approximation in large networks [7–9]. Betweenness centrality counts participation in shortest paths and is conceptually related to brokerage; nevertheless, shortest-path intermediation does not specify whether the paths connect different information contexts or whether other nodes can bridge those same context pairs [7–9]. Eigenvector centrality gives greater weight to links from prominent neighbours, capturing recursive structural prestige rather than context-diverse mediation [7–9]. In directed, reducible networks, its scores may additionally depend on orientation, convergence and component structure. Accordingly, these four measures serve as complementary structural baselines rather than direct estimators of context-specific brokerage.

Although betweenness centrality quantifies the extent to which a node lies on shortest paths between other nodes [7–9], it does not explicitly distinguish paths that traverse different information contexts or evaluate the availability of alternative nodes that can bridge the same ordered pairs of contexts. Brokerage and structural-hole theory recognise the importance of intermediary network positions and alternative connections [5], but these concepts do not, by themselves, operationalise the context-specific properties examined here. The proposed CCTB and ICB metrics address these distinct structural dimensions. CCTB quantifies the proportion of directed incoming–outgoing interaction opportunities across different contexts, independently of global shortest-path membership, whereas ICB weights observed context-pair connections by the number of alternative bridging nodes. Consequently, the proposed metrics complement rather than replace conventional betweenness centrality by incorporating contextual differentiation and context-pair substitutability into the assessment of node importance.

Beyond static topology, the temporal ordering and context of interactions influence how network positions should be interpreted. Temporal network research shows that the order of interactions affects which processes can occur across a changing graph [4]. A parent-to-child interaction recorded after another event cannot retroactively cause the earlier event. However, time-respecting paths alone are insufficient when the aim is to study movement across distinct discussion contexts. A network may contain a large number of time-ordered paths entirely within one article or claim. Our context-labelled formulation partitions a node's observed interaction opportunities by claim identifier in FibVID and by article URL in Monant, thereby addressing a different dimension from temporal recency.

A further perspective comes from brokerage and structural-hole theory [5]. [5]'s structural-hole account links brokerage to positions bridging otherwise separated contacts [5]. A directed two-step mediator is one elementary operationalisation of such a position, but conventional brokerage need not distinguish whether the contacts belong to the same health topic, claim or article. CCTB introduces explicit context inequality for incoming and outgoing edge labels. The interpretation is structural: a node creates a potential connection between two observed interaction contexts. The score alone cannot establish that medical misinformation moved between those contexts.

The availability of alternative intermediaries introduces a related but distinct consideration. A node may connect two contexts that dozens of others already connect, or may be the only observed mediator of a particular ordered context pair. These situations have different implications for network fragility and access to diverse discussions. ICB assigns fractional credit to the set of alternative mediators of each ordered pair. This is conceptually related to coverage and structural holes, rather than to predicting the cumulative number of retweets. It does not require estimating user susceptibility or assigning emotions to posts.

Distinguishing descriptive network position from causal intervention effects is equally important. [6] formalised influence maximisation under explicit diffusion models and intervention assumptions. Removing a node in a data-derived reply graph does not establish that its downstream conversation would not occur under a real intervention: nodes can seek alternative channels, and the observational edge is not a randomised exposure event. The present analysis therefore deliberately avoids claims about percentage reductions in misinformation following node targeting. Any intervention application would require additional dynamic and counterfactual validation.

Against this theoretical background, propagation, vulnerability and dynamic-influence centralities provide relevant additional comparators. [1] examined Propagation Centrality (PC), Misinformation Vulnerability Centrality (MVC) and Dynamic Influence Centrality (DIC) alongside conventional centralities in health-information networks. PC is intended to characterise propagation-related prominence, but a diffusion-oriented ranking need not distinguish within-context exposure from connections between distinct claims or articles. MVC introduces vulnerability-related terms, but such proxies should not be treated as direct observations of individual beliefs or psychological susceptibility, nor do they count alternative nodes connecting the same context pair. DIC accounts for changing structural influence, yet temporal influence and cross-context intermediation are different constructs: a dynamically prominent node need not connect multiple information contexts or hold a scarce bridging position. These limitations concern the constructs measured, not defects in the original methods. The present study therefore tests two explicitly defined quantities: the proportion of directed incoming–outgoing interaction opportunities spanning different contexts (CCTB) and the scarcity of substitutable nodes for the same ordered context pairs (ICB). Importantly, the PC, MVC, and DIC columns in the current exploratory evaluation are documented operational proxies and are not claimed to reproduce every stage of the previously published algorithms [1].

3. Materials and Methods

The study uses two health-information interaction datasets to investigate whether context-sensitive brokerage measures capture structural information beyond that captured by established centralities. The methodological sequence begins with data reconstruction and directed network construction, followed by a chronological partitioning of interactions, a mathematical specification of the nine measures, and a set-based comparison of their ranked node selections. Finally, the analysis evaluates subsequent observed activity and brokerage using node-disjoint predictive assessments. Unless otherwise specified, all centrality scores are derived solely from the initial 75% of chronologically ordered interactions.

3.1 Datasets, Network Construction and Preprocessing

This study primarily employed two publicly available health-information datasets, FibVID [2] and Monant [3], to investigate cross-context brokerage in directed online interaction networks, with MC-Fake [18] included as an exploratory third dataset to assess the reproducibility of the findings. FibVID captures claim-associated social media propagation, whereas Monant provides discussion interactions associated with medical articles. The datasets were processed to reconstruct directed parent–child interactions, identify their corresponding information contexts and establish comparable network representations. Table 1 summarises the dataset characteristics, preprocessing procedures, network construction and chronological evaluation partitions.

Table 1. Characteristics, preprocessing and network construction of the FibVID and Monant datasets.

Characteristic	FIBVID	MONANT
Data source	COVID-19-related social media propagation dataset [2]	Medical article discussion dataset [3]
Primary content	Claim-associated Twitter posts	Medical article comments and replies
Original data records	221,253 propagation posts	710,993 reconstructed comment records
Additional source characteristics	1,353 news claims; 1,774 originating tweets	Article-associated discussion records
Network nodes	Pseudonymised user identifiers	Comment-author identifiers
Directed edges	Parent-post author → child-post author	Parent-comment author → replying author
Context definition	Claim category and claim identifier	Medical article URL
Preprocessing	Exclusion of root records without parents and self-directed interactions	Reconstruction of logical CSV records, resolution of parent-comment links and exclusion of self-directed interactions
Retained directed interactions	218,248	263,469
Observed information contexts	236 claim contexts	15,244 article URLs
Network-construction partition (75%)	163,686 interactions	197,601 interactions
Subsequent evaluation partition (25%)	54,562 interactions	65,868 interactions
Nodes in network-construction partition	110,627	61,083
Nodes with degree ≥ 3	12,602	17,147
Analytical scope	Claim-associated interaction structures	Article-associated discussion structures
Interpretative limitation	Recorded interactions do not establish endorsement or belief	Replies do not establish misinformation endorsement or transmission

MC-Fake was included as an exploratory third dataset to assess whether the structural and predictive findings from FibVID and Monant could be reproduced in an independent, multi-topic information network. Introduced in [18], the dataset contains news-related social media interactions. Following reconstruction and deduplication, the analytical network comprised 1,952,282 directed interactions across 26,052 news contexts. News identifiers defined the information contexts, while reply and repost relationships were used to construct directed user interactions. Unlike FibVID and Monant, MC-Fake was treated as a multi-topic replication dataset rather than an exclusively health-information network.

Following preprocessing, each interaction was represented as (u, v, c, t), where u and v denote source and target nodes, respectively, c identifies the associated information context, and t represents the interaction timestamp. Parallel interactions were retained for context-sensitive brokerage calculations, whereas conventional adjacency-based centrality measures were calculated using collapsed directed connections.

For each dataset, interactions were arranged chronologically, with the initial 75% used to construct the network and calculate centrality scores. The remaining 25% were reserved for evaluating subsequent node activity and cross-context brokerage. Nodes with at least three incident interactions in the network-construction partition were eligible for predictive evaluation. Because the three datasets represent different interaction processes and contextual structures, their network statistics were analysed separately rather than interpreted as directly equivalent observations of misinformation transmission. For MC-Fake, chronological ordering relied on timestamps inferred from tweet identifiers.

The analysis used secondary research datasets containing pseudonymised node identifiers. Although these identifiers do not directly disclose personal identities, potential linkability to publicly available content requires appropriate data governance. Accordingly, the study reports aggregate network characteristics and does not infer individual beliefs, psychological susceptibility, medical conditions or endorsement of misinformation from observed interactions. Reproducibility materials and derived node-level outputs should be shared in accordance with the original datasets' access and reuse conditions.

3.2 Mathematical development of brokerage measures

3.2.1 Incoming and outgoing context counts

Let $G = (V, E)$ be an observed directed multigraph whose edges have labels $c \in C$. For node v, let $I(v,c)$ denote the number of edges ending at v in context c and $O(v,c)$ the number of edges leaving v in context c. Write $I(v)=\Sigma_c I(v,c)$ and $O(v)=\Sigma_c O(v,c)$. A directed two-step opportunity $a \rightarrow v \rightarrow b$ is potentially cross-context when its first and second edges have different context labels. The number of such labelled opportunities is $N_{cross}(v)=I(v)O(v) - \Sigma_c I(v,c)O(v,c)$. The subtraction removes within-context combinations, so a high degree confined to one claim does not automatically yield a high cross-context score.

3.2.2 Cross-Context Two-Step Brokerage (CCTB)

The node-score scripts used for the nine-metric comparisons calculate $CCTB(v) = N_{cross}(v)/\max\{1,I(v)O(v)\}$. This is the fraction of all directed incoming–outgoing edge combinations at v that involve two different contexts, bounded between zero and one. This fraction-based normalisation distinguishes cross-context opportunity share from the absolute number of two-step interaction opportunities. Both scores are zero when either direction is absent or no cross-context combination exists. The observed two-step opportunities do not prove actual transmission along each two-edge sequence.

3.2.3 Irreplaceable Context Brokerage (ICB)

Let $P(v)$ be the set of ordered pairs (c,d), $c \neq d$, for which $I(v,c)>0$ and $O(v,d)>0$. Let $m(c,d) = |\{v \in V : (c,d) \in P(v)\}|$ be the number of observed brokers capable of connecting that same ordered pair. Define $ICB(v)=\Sigma_{(c,d)\in P(v)} 1/m(c,d)$. A context connection for which v is the only recorded broker contributes one unit; if ten nodes connect the same ordered pair, each contributes one tenth. Thus the total credit distributed for any represented context pair equals one. This conservation property

makes the score interpretable as a fractional allocation of observed context-pair coverage across alternative brokers.
The evaluation encompasses nine measures in three methodological families: degree, closeness, betweenness and eigenvector centrality as traditional structural measures; PC, MVC and DIC as propagation, vulnerability and dynamic-influence comparators; and the proposed CCTB and ICB measures.

The conceptual differences between CCTB and ICB are illustrated in Figure 1, which highlights how the proposed centrality measures distinguish between the extent of cross-context interaction opportunities and the availability of alternative nodes capable of connecting the same information contexts.

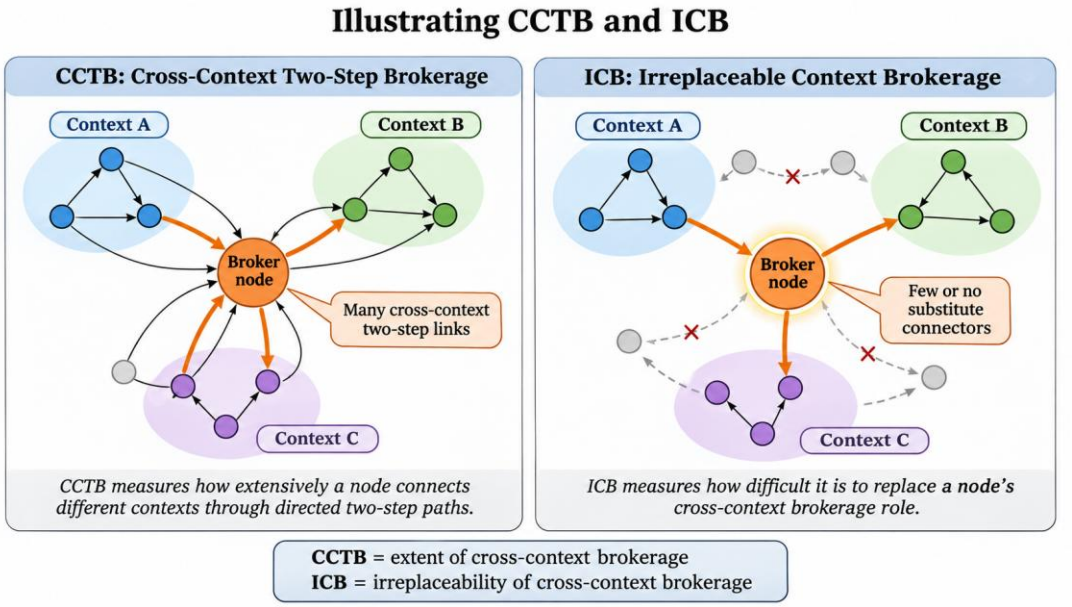


Figure 1. CCTB and ICB. CCTB quantifies the proportion of directed incoming–outgoing interaction opportunities across different information contexts, whereas ICB weights ordered context-pair connections based on the scarcity of alternative bridging nodes. The illustration is schematic and does not represent observed interactions from FibVID or Monant.

3.2.4 Formal distinction from DIC, MVC and PC

CCTB is built from cross-labelled two-step edge products and normalised activity. ICB allocates context-pair credit based on observed alternative brokers. Neither relies on a temporal influence-decay coefficient, a psychological vulnerability proxy, nor a path-diffusion kernel. The values are therefore not alternative parameterisations of DIC, MVC or PC. They are also not identical to PageRank, degree or ordinary betweenness. However, elements such as local brokerage and structural redundancy are established concepts [5,7–9]; the novel aspect claimed is the explicit paired-context formulation and its empirical assessment, not the first recognition of brokerage in network science.

3.2.5 Invariance and boundary properties

Relabeling node identifiers without changing the network structure or context labels leaves both CCTB and ICB unchanged. Uniform replication of all recorded interactions also preserves CCTB because the cross-context opportunity count and the total incoming–outgoing opportunity count are multiplied by the same squared replication factor. Similarly, ICB remains invariant under uniform interaction replication, provided the observed ordered context pairs and their corresponding sets of

bridging nodes remain unchanged. When all interactions are assigned the same context label, both measures equal zero. Furthermore, introducing an additional node capable of bridging an existing ordered context pair reduces the fractional ICB contribution assigned to each previously identified broker of that pair. These invariance and boundary properties establish predictable mathematical behaviour under changes in node labelling, interaction multiplicity and context-pair substitutability.

3.2.6 Computational feasibility and approximation

Computing CCTB requires the directed node/context incidence tables and a join on shared context labels; it is linear in the number of observed edges plus the number of nonzero node-context incidences. ICB requires enumerating pairs of incoming and outgoing contexts for each eligible intermediary. Most nodes have small context sets, but highly active accounts can generate many pairs. For this exploratory study, node pair enumeration was capped at 10,000 combinations by truncating each side to 100 sorted contexts for pathological cases. This affected no FibVID nodes and 37 Monant nodes. Because sorted truncation is not an unbiased estimator, ICB results for Monant should be regarded as approximation-sensitive; a fully scalable implementation should use sketches, exact sparse pair aggregation, or explicitly randomised subsampling with error bounds.

3.3 Mathematical definitions of all nine centrality measures

Let $d \in \{\text{FibVID, Monant}\}$ index a dataset and $G^d = (V^d, E^d)$ denote its training-period directed node graph. $A^d_{uv} = 1$ indicates at least one observed edge $u \rightarrow v$ in the collapsed binary graph, and $n^d = |V^d|$. For a node v, $i^d(v)$ and $o^d(v)$ are its incoming and outgoing degrees. Directed shortest-path distance is $dist^d(s,t)$, σ^d_{st} counts shortest directed $s \rightarrow t$ paths, and $\sigma^d_{st}(v)$ counts those traversing v. All scores use only the first 75% of chronological interactions.

1. ***Degree centrality (DC)***. DC quantifies the number of distinct incoming and outgoing node connections, indicating immediate structural connectivity. Parallel recorded interactions are collapsed to one directed adjacency relation.

$$DC^d(v) = i^d(v) + o^d(v) = \Sigma_u A^d_{uv} + \Sigma_w A^d_{vw} \quad (1)$$

2. ***Closeness centrality (CC).*** CC captures how readily a node can reach other nodes along directed shortest paths. The operational definition is outward harmonic closeness, which handles unreachable nodes by assigning them no contribution. Large-graph values are sampled-landmark approximations using the reversed graph and may differ in scale from the full-graph expression.

$$CC^d(v) = \Sigma_u \in V^d \setminus \{v\}, dist^d(v,u) < \infty \; 1 / dist^d(v,u) \quad (2)$$

3. ***Betweenness centrality (BC).*** BC measures how frequently a node lies on shortest directed paths between other node pairs, capturing shortest-path intermediation rather than specifically cross-context brokerage. Reported values were estimated using sampled-source Brandes calculations.

$$BC^d(v) = \Sigma_s \neq v \neq t, \sigma^d_{st} > 0 \; \sigma^d_{st}(v) / \sigma^d_{st} \quad (3)$$

4. ***Eigenvector centrality (EC).*** EC assigns higher scores to nodes receiving ties from other highly scored nodes, measuring recursively defined incoming-neighbour prominence. Power iteration is used, with convergence and reducibility limitations acknowledged.

$$\lambda^d x^d(v) = \Sigma_u \in V^d A^d_{uv} x^d(u), \quad EC^d(v) = x^d(v) \quad (4)$$

5. ***Propagation Centrality comparator (PC).*** The implemented PC column is a PageRank-style proxy measuring the stationary prominence of nodes under repeated link-following and teleportation. It uses damping α = 0.85 and redistributes dangling-node mass. This is not claimed to reproduce every operation of the previously published PC algorithm.

$$PC^{d}(v) = (1 - \alpha)/n^{d} + \alpha[\Sigma_{u}: o^{d}(u) > 0\ A^{d}{}_{uv}\, PC^{d}(u)/o^{d}(u) + D^{d}/n^{d}] \quad (5)$$

In Equation (5), $D^{d} = \Sigma_{u}:o^{d}(u)=0\ PC^{d}(u)$ denotes the total stationary probability of dangling nodes.

6. ***Misinformation Vulnerability Centrality comparator (MVC)***. The implemented exploratory MVC proxy combines incoming degree with a seeded, uniformly drawn random factor. It is not a measured psychological susceptibility or a clinical attribute. Some Monant scores are divided by the maximum value, without changing their within-dataset ranking.

$$MVC^{d}(v) = [i^{d}(v)]^{5}\, q^{d}{}_{v}, \quad q^{d}{}_{v} \sim Uniform(0,1) \quad (6)$$

7. ***Dynamic Influence Centrality comparator (DIC)***. The implemented proxy iteratively propagates incoming-neighbour influence across the historical adjacency structure, normalising at each step. This recursion does not reproduce the published time-decay mechanism and should not be interpreted as an exact implementation of that method.

$$z^{d(t)} = x^{d(t)} + (A^{d})^{T}x^{d(t)}; \quad x^{d(t+1)} = z^{d(t)} / max\{1, max_{w}\, z^{d(t)}(w)\} \quad (7)$$

$$x^{d(0)} = 1; \quad DIC^{d}(v) = x^{d(10)}(v) \quad (8)$$

8. ***Cross-Context Two-Step Brokerage (CCTB)***. CCTB measures the proportion of a node's incoming–outgoing event combinations that join different information contexts. Let $I^{d}(v,c)$ and $O^{d}(v,c)$ count incoming and outgoing historical events of context c and let $I^{d}(v) = \Sigma c\ I^{d}(v,c)$, $O^{d}(v) = \Sigma c\ O^{d}(v,c)$. Unlike the binary adjacency centralities, CCTB uses event multiplicity.

$$N^{d}cross(v) = I^{d}(v)O^{d}(v) - \Sigma c\, I^{d}(v,c)O^{d}(v,c) \quad (9)$$

$$CCTB^{d}(v) = N^{d}cross(v) / max\{1, I^{d}(v)O^{d}(v)\} \quad (10)$$

9. ***Irreplaceable Context Brokerage (ICB)***. ICB assesses how difficult it is to substitute a node as a broker of ordered context pairs. Each observed context pair contributes a total weight of one distributed equally across all nodes eligible to bridge it. The recorded 10,000-pair enumeration cap affected 37 Monant nodes, so their ICB estimates may change under full enumeration.

$$P^{d}(v) = \{(c_1, c_2) : c_1 \neq c_2, I^{d}(v,c_1) > 0, O^{d}(v,c_2) > 0\} \quad (11)$$

$$h^{d}(c_1, c_2) = |\{w \in V^{d} : (c_1, c_2) \in P^{d}(w)\}| \quad (12)$$

$$ICB^{d}(v) = \Sigma(c_1, c_2) \in P^{d}(v)\ 1 / h^{d}(c_1, c_2) \quad (13)$$

3.4 Set-theoretic selection and agreement

For each dataset d, let $\mathcal{M}$ = {DC, CC, BC, EC, PC, MVC, DIC, CCTB, ICB} and $T^{d}{}_{m}(k)$ be the k highest-scoring nodes for measure m, with descending score and node ID ascending used to resolve score ties. Set operations are performed on the selected node IDs, not on the numerical centrality scores.

$$U^{d}(k) = \bigcup_{m} \in \mathcal{M}\ T^{d}{}_{m}(k) \quad (14)$$

Equation (14) is the union of all selected nodes. The intersection in Equation (15) identifies nodes selected by both metrics a and b.

$$H^{d}(a, b; k) = T^{d}{}_{a}(k) \cap T^{d}{}_{\beta}(k) \quad (15)$$

Exclusive selection uses set difference to exclude every other top-k metric list:

$$E^{d}{}_{m}(k) \;=\; T^{d}{}_{m}(k) \;\setminus\; [\bigcup_{j} \in \mathcal{M} \setminus \{m\}\, T^{d}{}_{j}(k)] \quad (16)$$

To group nodes by their exact combination S of identifying measures, the intersection must exclude nodes also appearing in any metric outside S:

$$G^{d}(S,k) \;=\; [\bigcap_{m} \in S\, T^{d}{}_{m}(k)] \;\setminus\; [\bigcup_{j} \in \mathcal{M} \setminus S\, T^{d}{}_{j}(k)] \quad (17)$$

In Equation (17), $S \subseteq \mathcal{M}$ denotes a subset of the evaluated centrality measures. The equation identifies nodes selected by exactly the measures in S, excluding nodes also selected by measures outside this subset. Table 8 presents the nonempty exact-overlap classes for $|S| \geq 2$ at $k = 10$, whereas Table 9 identifies nodes selected exclusively by CCTB or ICB. Accordingly, the intersection notation in Table 8 represents exact membership classes defined by Equation (17), rather than unrestricted intersections. All top-10 selections follow the specified deterministic ranking and tie-breaking procedure.

3.5 Experimental Evaluation and Reproducibility

The evaluation protocol is specified in Algorithm 1. Network-construction and evaluation partitions are formed chronologically, while predictive training and testing are separated at the node level. The procedure distinguishes subsequent network activity from cross-context brokerage and evaluates the additional predictive contribution of CCTB and ICB beyond basic network characteristics.

Algorithm 1. Evaluation of Subsequent Node Activity and Cross-Context Brokerage

Input: Datasets d ∈ {FibVID, Monant}; timestamped interactions E^d; training-period scores DC, CB, CCTB, ICB and seven comparator measures; seeds R = {11, 23, 37, 41, 51}

Output: Future-activity and brokerage outcomes; raw-score ROC-AUC, Average Precision (AP), top-decile yields; five-split incremental model performance

```
for each dataset d ∈ {FibVID, Monant} do
    sort E^d by timestamp; define E^dtrain = first 75% and E^deval = remaining 25%
    compute all node scores from E^dtrain only; retain nodes satisfying the predefined eligibility rule
    for each eligible node v do
        set yactivity(v) = 1 if v occurs in at least one interaction in E^deval; otherwise 0
        set ybroker(v) = 1 if v has an incoming–outgoing event pair with distinct contexts in E^deval; otherwise 0
    for each raw ranking score m ∈ {DC, CB, CCTB, ICB} do
        compute ROC-AUC and AP for subsequent activity and brokerage
        compute observed event yield among the top 10% of ranked eligible nodes
    for each random seed r ∈ R do
        split eligible nodes into disjoint 70% training and 30% testing subsets
        fit log(1 + x) transformations and standardisation as specified, using training nodes only
        fit logistic model M0 with features {DC, CB}
        fit logistic model M1 with features {DC, CB, CCTB, ICB}
        fit additional comparator models {DC, CC, BC, EC, PC, MVC, DIC}, then + ICB, then + ICB + CCTB
        record held-out ROC-AUC and AP for all model specifications
    aggregate results across the five seeds as mean and standard deviation
    report simple-baseline augmentation in Table 3 and comparator augmentation in Table 5
return raw-ranking statistics and predictive evaluation summaries
```

Implementation notes. Future brokerage is an observed structural outcome, not evidence of causal information transfer. The five node splits share the same network-construction graph and are not independent network replications. Average Precision is particularly informative for infrequent positive outcomes. Alternative temporal cutoffs, context-label permutations and full context-pair enumeration remain prospective sensitivity analyses, not completed tests.

3.6 Experimental implementation of measures of centrality

The experimental procedure is summarised in Algorithms 2 and 3. Algorithm 2 describes the construction of the network-construction-period graphs and the computation and ranking of all nine centrality measures. Algorithm 3 specifies the set-theoretic comparison of centrality rankings, the identification of exclusive CCTB and ICB node selections, and the evaluation of subsequent cross-context brokerage among the highest-ranked nodes. The mathematical definitions of the centrality measures and exact-overlap sets are provided in Sections 3.3 and 3.4, respectively. Algorithms 2 and 3 describe the primary FibVID and Monant evaluation; MC-Fake was examined separately as an exploratory extension using the corresponding historical scoring and prospective evaluation principles.

Algorithm 2: Construction of Nine-Metric Centrality Rankings

Input: Timestamped interactions $E^d = \{(u, v, c, t)\}$ for $d \in \{\text{FibVID, Monant}\}$; metric set $M = \{\text{DC, CC, BC, EC, PC, MVC, DIC, CCTB, ICB}\}$

Output: Node-score matrix S^d, top-10 and top-200 sets $T^d_m(k)$, and approximation/tie diagnostics

```
for each dataset d ∈ {FibVID, Monant} do
    validate timestamps, parent references and node identifiers
    exclude roots without resolved parents and self-interactions
    sort eligible interactions by timestamp t
    partition E^d into first 75% E^d_train and final 25% E^d_eval
    construct binary directed adjacency A^d from E^d_train
    retain labelled event multiplicities I^d(v,c), O^d(v,c)
    for each node v ∈ V^d_train do
        compute DC, CC, BC and EC on A^d
        compute operational PC, MVC and DIC comparator scores
        compute CCTB from cross-context incoming–outgoing event pairs
        compute ICB from weighted ordered context-pair coverage
        record seeds, approximation settings and convergence flags
    for each metric m ∈ M do
        sort nodes by (−score_m(v), node ID)
        save T^d_m(10), T^d_m(200) and cutoff-tie counts
return S^d, T^d_m(k) and diagnostics
```

Algorithm 3. Set-Based Comparison and Prospective Brokerage Evaluation

Input: Historical node scores S^d, top-k sets $T^d_m(k)$, evaluation interactions E^d_eval, metric set $\mathcal{M}$, $k \in \{10, 200\}$

Output: Exact-overlap classes, CCTB/ICB-exclusive nodes, top-200 brokerage counts and tie diagnostics

```
for each dataset d ∈ {FibVID, Monant} do
    derive y^d(v), the subsequent cross-context brokerage indicator, from E^d_eval
    compute exact top-10 overlap classes G^d(S,10) using Equation (17)
```

identify exclusive sets E^d_CCTB(10) and E^d_ICB(10) using Equation (16)
report exact overlaps and exclusive selections in Tables 8 and 9
for each centrality measure $m \in \mathcal{M}$ do
compute $B^d_m(200) = \sum_{v \in T^d_m(200)} y^d(v)$
record the number of nodes exhibiting subsequent cross-context brokerage
report top-200 brokerage counts in Table 4
record ranking-cutoff ties and selection uncertainty
return overlap classes, exclusive sets, brokerage counts and diagnostics

Implementation note. All node rankings are computed from the initial 75% of chronological interactions. Subsequent brokerage is evaluated only on the remaining 25%. Deterministic node-ID tie-breaking supports reproducibility but does not resolve substantive uncertainty from tied scores.

3.7. Centrality-Guided Node-Removal Simulations

Node-removal simulations were conducted across FibVID, Monant and MC-Fake to evaluate the potential of centrality-guided interventions for disrupting misinformation propagation pathways. Interactions were chronologically divided into historical (75%) and future (25%) observations. Five targeting strategies were compared: degree centrality, context breadth, a PageRank-based PC proxy, CCTB and ICB. The top-ranked 5% of historical nodes were selected for simulated removal, and reductions in future two-step cross-context brokerage opportunities were evaluated. The intervention budget was rounded down to the nearest whole node, and intervention outcomes were calculated by excluding future brokerage opportunities involving the selected nodes. Directed interactions and context identifiers were preserved, while deterministic rankings and tie-breaking sensitivity checks supported reproducibility. The unmodified future network served as the counterfactual baseline.

Intervention effectiveness was calculated as:

$$\mathbf{R_m\ (\%) = [(B_0 - B_m) / B_0] \times 100}$$

where B_0 denotes the baseline number of future two-step cross-context brokerage opportunities and B_m denotes the number remaining after excluding the selected nodes. The resulting percentages represent simulated reductions in structural brokerage opportunities, rather than observed decreases in misinformation transmission, with the comparative intervention outcomes across FibVID, Monant and MC-Fake presented in Table 11.

4. Results

The empirical evaluation examines the structural characteristics of the reconstructed FibVID and Monant networks and assesses the extent to which context-sensitive brokerage measures provide complementary information beyond established centrality measures. The results are presented sequentially, beginning with network reconstruction and individual centrality performance, followed by incremental predictive evaluation, cross-metric node selection and the structural relationships between degree centrality and cross-context brokerage. This organisation distinguishes between descriptive network characteristics and prospective node-selection performance, providing a consistent basis for interpreting the proposed CCTB and ICB measures.

4.1 Reconstruction quality and observation units

As documented in Table 1, FibVID contained 218,248 eligible directed node interactions after excluding roots without parents and self-edges. These occurred in 236 claim contexts with at least one such interaction. The Monant parser recovered 710,993 logical comment records from 47 line-split files and retained 263,469 directed parent–child interactions between different authors. In each case, the observed node and edge counts depend on the operational definition of the graph. There is no directly comparable news-veracity outcome in the Monant comment-only sample.

4.2 Structural distribution

The network-construction populations reported in Table 1 were substantially larger than the subsets eligible for predictive evaluation. Applying the minimum requirement of three observed incident interactions retained 12,602 FibVID nodes and 17,147 Monant nodes for subsequent evaluation. Analysis of the complete network-construction graphs revealed substantial sparsity in context-sensitive brokerage, with 3,199 of 110,627 FibVID nodes (2.89%) and 4,622 of 61,083 Monant nodes (7.57%) exhibiting positive CCTB and ICB scores. Although both measures identified the same positive-score nodes, they quantify different structural characteristics and therefore produce distinct rankings. Nodes with zero brokerage scores should not be interpreted as lacking social influence, since the proposed measures specifically capture observed cross-context interaction opportunities.

4.3 Individual Centrality Performance for Subsequent Activity and Brokerage

The discriminative performance of degree centrality, context breadth, CCTB and ICB was examined against two subsequent network outcomes: observed node activity and cross-context brokerage. Table 2 presents the corresponding ROC-AUC values and top-decile event yields for both datasets. The dataset-specific performance patterns are further illustrated in Figures 2 and 3 for FibVID and Monant, respectively.

Table 2. Raw-score future ranking metrics, using a chronological 75%/25% edge split.

Dataset	Score	AUC activity	AUC brokerage	Top 10% activity	Top 10% brokerage
FIBVID	degree	0.635	0.808	0.545	0.182
FIBVID	context_breadth	0.683	0.735	0.723	0.167
FIBVID	CCTB	0.567	0.795	0.631	0.218
FIBVID	ICB	0.566	0.790	0.651	0.204
MONANT	degree	0.571	0.651	0.128	0.048
MONANT	context_breadth	0.667	0.849	0.162	0.073
MONANT	CCTB	0.631	0.788	0.137	0.057
MONANT	ICB	0.639	0.806	0.165	0.075

The results in Table 2 show that context breadth achieved the highest ROC-AUC for subsequent activity in both datasets, reaching 0.683 for FibVID and 0.667 for Monant. By comparison, CCTB achieved values of 0.567 and 0.631, while ICB achieved 0.566 and 0.639, respectively. The corresponding brokerage results indicate a different performance pattern. In FibVID, degree centrality achieved the highest brokerage ROC-AUC of 0.808, whereas context breadth achieved the highest value in Monant at 0.849. Although CCTB and ICB did not consistently outperform these simpler measures, their top-decile brokerage yields indicate differences in the types of nodes prioritised by each ranking criterion.

The differing discrimination patterns for subsequent activity and brokerage in FibVID are illustrated in Figure 2.

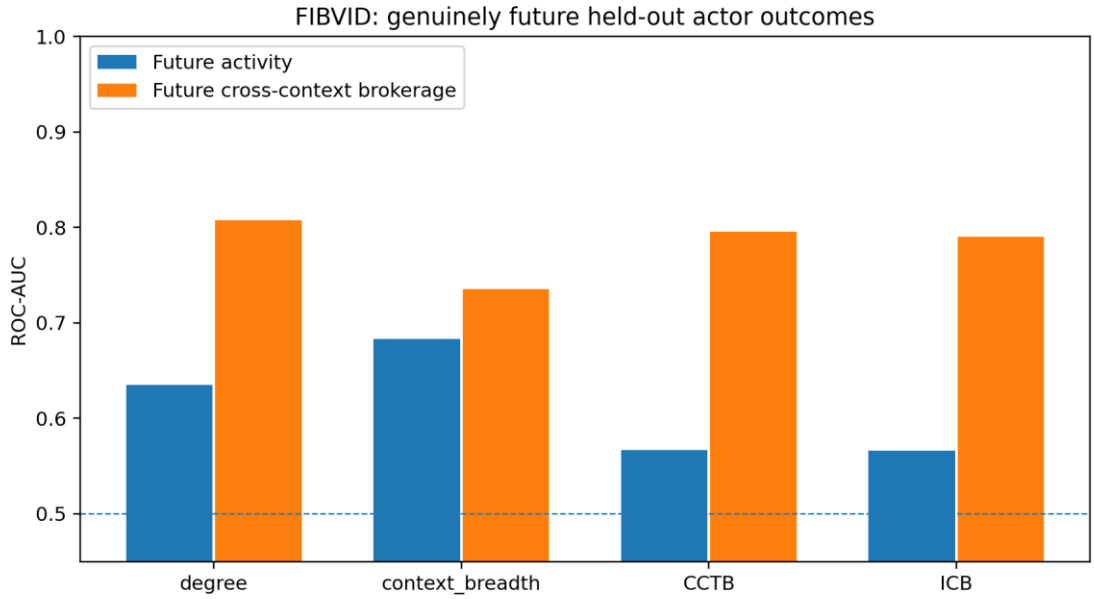


Figure 2. FibVID ranking against future observed activity and future cross-context brokerage.

Figure 3 presents the corresponding Monant results, allowing the relative performance of connectivity, context breadth and the two proposed brokerage measures to be examined within the discussion-reply network.

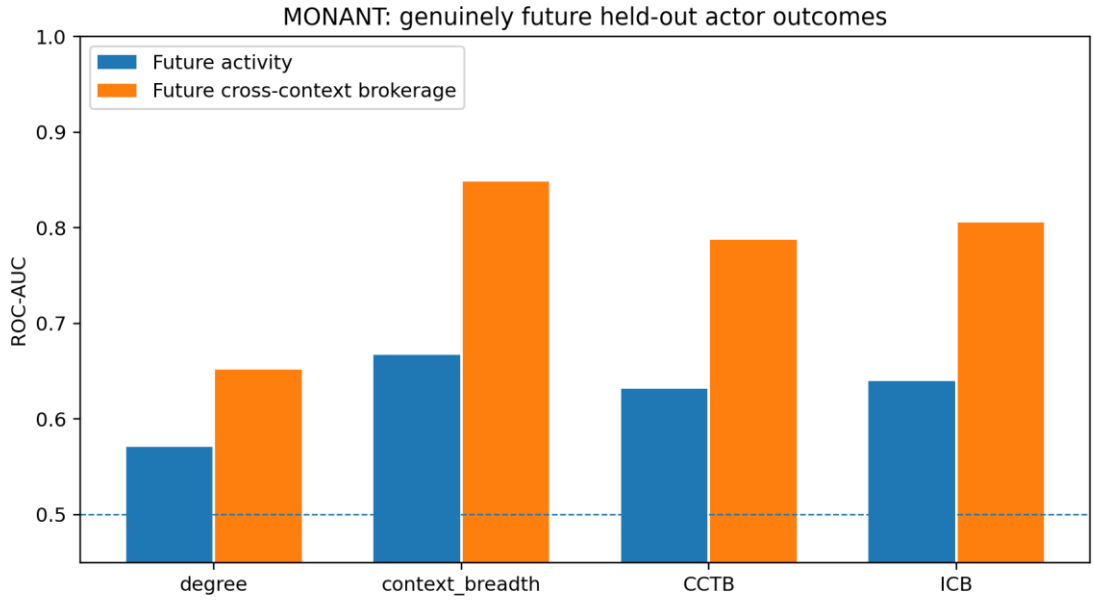


Figure 3. Monant ranking against future observed activity and future cross-context brokerage.

Considered together, Figures 2 and 3 demonstrate that the relative performance of the four ranking measures depends on both the dataset and the subsequent outcome being evaluated. Context breadth is particularly informative for general activity, whereas brokerage-oriented measures provide alternative patterns of prioritisation. These observations motivate the incremental predictive analysis that follows, which examines whether CCTB and ICB contribute information beyond degree and context breadth when used as additional model features.

4.4 Incremental Predictive Performance of CCTB and ICB

To assess whether CCTB and ICB contribute predictive information beyond basic network connectivity and contextual participation, incremental logistic regression models were evaluated using degree centrality and context breadth as baseline predictors. The augmented models additionally incorporated CCTB and ICB. Table 3 presents the mean ROC-AUC and Average Precision (AP), together with their standard deviations across five node-disjoint evaluation splits. Figure 4 illustrates the corresponding changes in AP for FibVID and Monant, providing a visual comparison of predictive performance between the baseline and augmented models.

Table 3. Forward-time node-held-out incremental prediction of future cross-context brokerage. Five node split seeds.

Dataset	Features	AUC mean	AUC SD	AP mean	AP SD
FIBVID	Degree + breadth	0.7954	0.0169	0.2966	0.0447
FIBVID	+ CCTB + ICB	0.8184	0.0186	0.3542	0.0395
MONANT	Degree + breadth	0.7989	0.0227	0.0847	0.0041
MONANT	+ CCTB + ICB	0.7965	0.0241	0.0925	0.0106

The predictive results in Table 3 show differences in the contribution of CCTB and ICB across the two datasets. For FibVID, incorporating CCTB and ICB increased mean ROC-AUC from 0.7954 to 0.8184 and mean Average Precision (AP) from 0.2966 to 0.3542, corresponding to absolute improvements of 0.0230 and 0.0576, respectively. For Monant, mean AP increased from 0.0847 to 0.0925, representing an absolute improvement of 0.0078, whereas mean ROC-AUC decreased slightly from 0.7989 to 0.7965. These results suggest that the proposed brokerage measures provide additional predictive information beyond degree centrality and context breadth, particularly for FibVID. However, the contrasting ROC-AUC results demonstrate that the improvements are not consistent across all performance criteria.

To examine these differences visually, Figure 4 compares the mean AP obtained from the baseline and augmented models across both datasets. The accompanying error bars represent standard deviations across the five node-disjoint evaluation splits, illustrating the variability in predictive performance.

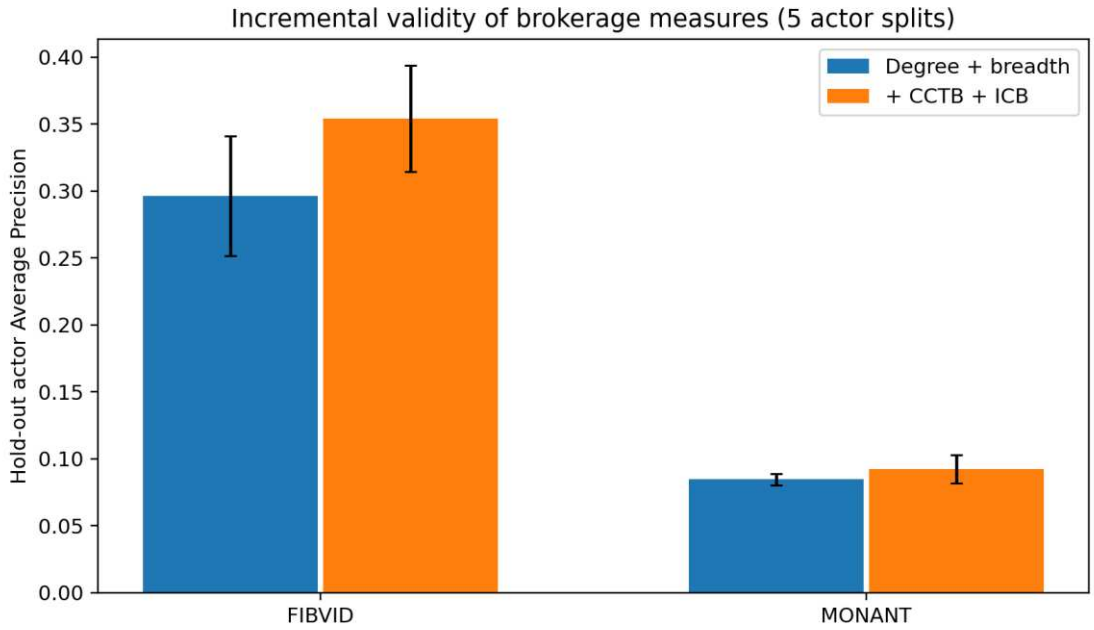


Figure 4. Incremental Average Precision relative to degree and context breadth. Error bars are standard deviations across five node splits, not confidence intervals over independently sampled networks.

The graphical comparison in Figure 4 reinforces the differences observed between the two datasets. The improvement in Average Precision is more pronounced for FibVID, while the smaller difference in Monant is accompanied by overlapping variability across the evaluation splits. These findings suggest that the contribution of context-sensitive brokerage features is dataset-dependent and should be interpreted alongside the ROC-AUC results reported in Table 3.

4.5 Comparative Evaluation of Nine Centrality Measures

To evaluate the information captured by distinct centrality constructs, degree, sampled closeness, sampled betweenness, eigenvector, PC, MVC, DIC, CCTB and ICB were compared within matched historical node graphs. The first 75% of interactions were used for scoring and the subsequent 25% for observed future brokerage. The analysis is exploratory: shortest-path metrics use approximations, MVC uses simulated vulnerability, and identical node identities are not assumed across datasets.

Table 4 shows how many nodes in each metric's top-200 ranking subsequently exhibit brokerage in the evaluation partition.

Table 4. Future brokers among the top 200 ranked nodes using historical-only scores.

Metric	FibVID (n/200)	Monant (n/200)
Degree	70	23
Closeness*	44	0
Betweenness*	59	15
Eigenvector	47	0
PC	4	36
MVC†	45	23
DIC	47	0
CCTB‡	50	11
ICB	105	30

Note. *Closeness and betweenness centralities were estimated using approximate shortest-path calculations. † MVC represents simulated vulnerability scores. ‡ Monant CCTB top-200 ranking is strongly tie-dependent (977 nodes at the cutoff). Differences in the graph and outcome definitions preclude interpreting these figures as causal influence or misinformation endorsement.

Beyond the future-brokerage counts summarised in Table 4, set-based comparison of the top-200 rankings identifies nodes selected by the proposed brokerage measures but absent from the combined comparator rankings. At top 200, ICB recovered 86 nodes outside the union of the combined rankings of the four traditional structural centralities and the three centralities introduced in 2025 in FibVID, of whom 35 subsequently brokered across contexts (40.7%). In Monant, 30 ICB-exclusive nodes were recovered, and 8 were subsequently brokered (26.7%). CCTB recovered 119 exclusive nodes in FibVID, of whom 17 subsequently brokered (14.3%). The Monant CCTB exclusivity result is not reliably interpretable because a large group of nodes has an identical cutoff score. These set-based findings describe complementary rankings; they do not show that each selected node caused the diffusion of information.

To determine whether the proposed measures provide predictive information beyond the seven comparator centralities, Table 5 compares the baseline logistic regression model with specifications incorporating ICB and the combined ICB and CCTB features.

Table 5. Incremental predictive performance for future brokerage, averaged over five node-disjoint splits.

Dataset	Predictors	Mean AP	Mean ROC-AUC
FIBVID	Baseline metrics	0.0952	0.8854
FIBVID	Baseline + ICB	0.1144	0.8929
FIBVID	Baseline + CCTB + ICB	0.1137	0.8929
MONANT	Baseline metrics	0.0576	0.7507
MONANT	Baseline + ICB	0.0859	0.7990
MONANT	Baseline + CCTB + ICB	0.0843	0.7996

Note: Baseline metrics comprise Degree Centrality (DC), Closeness Centrality (CC), Betweenness Centrality (BC), Eigenvector Centrality (EC), Propagation Centrality (PC), Misinformation Vulnerability Centrality (MVC), and Dynamic Influence Centrality (DIC). PC, MVC and DIC are represented by operational proxies rather than exact reproductions of their original formulations.

As shown in Table 5, incorporating ICB increased mean Average Precision from 0.0952 to 0.1144 for FibVID and from 0.0576 to 0.0859 for Monant. Mean ROC-AUC also increased in both datasets. However, subsequent addition of CCTB did not yield further improvements in Average Precision, which decreased slightly to 0.1137 for FibVID and 0.0843 for Monant. These results suggest that ICB contributes complementary predictive information beyond the seven comparator measures, while the additional contribution of CCTB is limited under the evaluated model specifications.
Figure 5 illustrates the corresponding differences in mean Average Precision across the comparator-only, ICB-augmented and fully augmented model specifications.

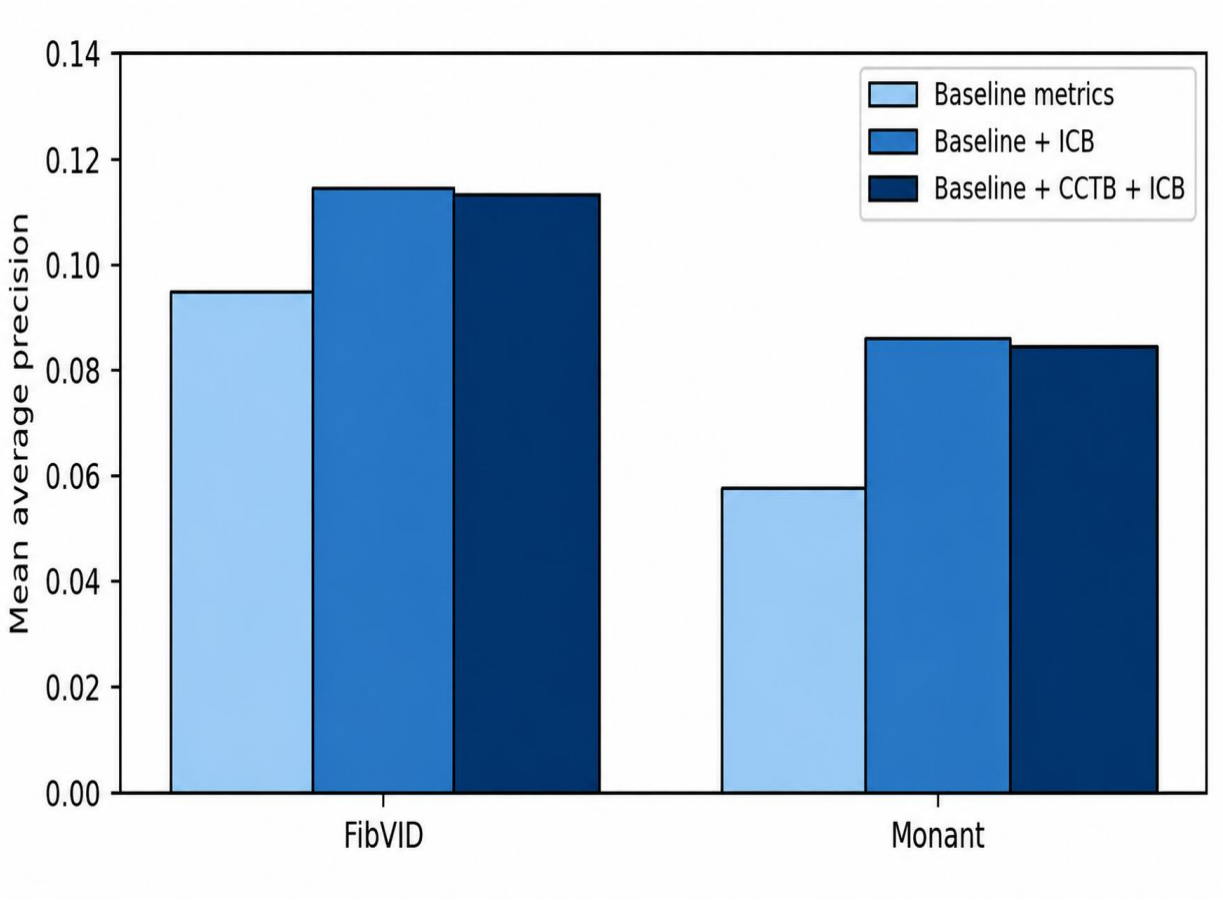


Figure 5. Incremental average precision for future cross-context brokerage on FibVID and Monant. The improvements illustrated in Figure 5 must be interpreted in relation to the computational approximations and ranking uncertainties associated with the evaluated measures. Monant CCTB exhibits 977 ties at the top-200 boundary; a 2,000-draw random tie audit produced a central 95% range of 6–18 future brokers among selected groups of 200. The structural approximations require additional random-seed stability checks, and the simulated MVC values do not indicate observed psychological susceptibility. Results therefore support exploratory complementary prediction, not definitive ranking superiority or a deployable intervention claim.

4.6 Top-10 Centrality Rankings, Metric Agreement and Exclusive Node Selection

For interpretability, Tables 6 and 7 report the top 10 node IDs for each of the four traditional structural measures (degree, closeness, betweenness and eigenvector), the three measures introduced in 2025 (PC, MVC and DIC), and the proposed CCTB and ICB. Scores are calculated from the historical interaction snapshots. Nodes are ordered by descending score; an ascending node-ID tie-break is applied solely for reproducible display and does not establish a unique scientific ranking when scores are tied.

The FibVID top-10 selections for all nine measures are presented in Table 6, including the extent of ties in scores at the selection boundary.

Table 6. FibVID top-10 nodes across the nine centrality measures.

Measure	Top 10 node IDs (descending scores)	Cutoff tie status
Degree	51591, 27011, 9225, 93248, 95758, 143737, 117831, 55209, 40617, 97748	Unique
Closeness	51591, 27011, 94466, 40617, 55209, 63018, 93248, 4535, 143876, 92676	Unique
Betweenness	51591, 27011, 45031, 93248, 55209, 70018, 36831, 95758, 9225, 40617	Unique
Eigenvector	51300, 1488, 51591, 120624, 82232, 115953, 27149, 69775, 119676, 23297	Unique
PC	33000, 42728, 13530, 137692, 29039, 62094, 80885, 92906, 88557, 61732	Unique
MVC	119676, 82232, 51591, 1488, 115953, 16307, 44102, 43963, 51300, 123321	Unique
DIC	51300, 1488, 51591, 120624, 82232, 115953, 27149, 69775, 119676, 23297	Unique
CCTB	10183, 104042, 105547, 106030, 106682, 10836, 109288, 110882, 111699, 112115	89 tied at rank 10
ICB	51591, 20991, 23297, 16307, 122641, 67140, 93248, 4163, 78487, 25915	Unique

Table 6 reveals considerable overlap among several established centrality measures, with nodes such as 51591, 93248 and 27011 appearing in multiple top-10 rankings. In particular, EC and DIC identify the same set of 10 nodes, suggesting substantial agreement in their node selections despite their different operational formulations. The proposed ICB measure also prioritises several nodes identified by conventional measures, including 51591 and 93248, while selecting additional nodes absent from the other displayed top-10 lists. By contrast, the displayed CCTB selections do not overlap with those of the other measures. However, 89 nodes share the CCTB score at the tenth-rank boundary, meaning that the displayed selections depend on deterministic tie-breaking and should not be interpreted as definitive evidence of uniquely identified nodes.

The corresponding Monant rankings are presented in Table 7, providing a second network context for examining agreement and divergence among the nine centrality measures.

Table 7. Monant top-10 nodes across the nine centrality measures.

Measure	Top 10 node IDs (descending scores)	Cutoff tie status

Degree	219535, 141370, 165439, 169499, 137586, 133750, 206955, 133746, 226828, 170726	Unique
Closeness	138581, 133746, 164008, 165999, 134643, 165562, 164341, 176408, 133908, 135573	15 tied at rank 10
Betweenness	219535, 226828, 236860, 220964, 169499, 141370, 224069, 248691, 222471, 133750	Unique
Eigenvector	165439, 170726, 137586, 164149, 141370, 133746, 164470, 168598, 184720, 170029	Unique
PC	219535, 169499, 206955, 226828, 141370, 214429, 209235, 165439, 137586, 133750	Unique
MVC	141370, 219535, 165439, 137586, 226828, 206955, 133750, 169499, 164149, 150630	Unique
DIC	165439, 170726, 137586, 164149, 141370, 133746, 164470, 168598, 184720, 170029	Unique
CCTB	125432, 125515, 125581, 125708, 125735, 125876, 125893, 125901, 125964, 125967	977 tied at rank 10
ICB	219535, 165439, 133750, 141370, 137586, 202198, 164149, 169499, 170726, 133746	Unique

The top-10 rankings in Tables 6 and 7 provide the basis for examining agreement among the nine centrality measures across the two datasets. To distinguish nodes jointly prioritised by multiple measures from those selected independently, the exact-overlap formulation in Equation (17) was applied to the displayed top-10 sets. Table 8 reports the resulting mutually exclusive intersection groups, identifying the specific combinations of measures selecting each shared node.

Table 8. Exact intersections of displayed top-10 centrality sets in FibVID and Monant.

Dataset	Centrality measures selecting nodes	Node IDs	No. of nodes
FIBVID	DC ∩ CC ∩ BC ∩ EC ∩ MVC ∩ DIC ∩ ICB	51591	1
	DC ∩ CC ∩ BC ∩ ICB	93248	1
	DC ∩ CC ∩ BC	27011, 40617, 55209	3
	EC ∩ DIC ∩ ICB	23297	1
	EC ∩ MVC ∩ DIC	1488, 51300, 82232, 115953, 119676	5
	DC ∩ BC	9225, 95758	2
	EC ∩ DIC	27149, 69775, 120624	3
	MVC ∩ ICB	16307	1
MONANT	DC ∩ BC ∩ EC ∩ PC ∩ MVC ∩ DIC ∩ ICB	141370	1
	DC ∩ EC ∩ PC ∩ MVC ∩ DIC ∩ ICB	137586, 165439	2
	DC ∩ BC ∩ PC ∩ MVC ∩ ICB	133750, 169499, 219535	3
	DC ∩ CC ∩ EC ∩ DIC ∩ ICB	133746	1
	DC ∩ BC ∩ PC ∩ MVC	226828	1
	DC ∩ EC ∩ DIC ∩ ICB	170726	1
	EC ∩ MVC ∩ DIC ∩ ICB	164149	1
	DC ∩ PC ∩ MVC	206955	1

Dataset	Centrality measures selecting nodes	Node IDs	No. of nodes
	EC ∩ DIC	164470, 168598, 170029, 184720	4

The exact-overlap analysis presented in Table 8 identifies 17 distinct FibVID nodes and 15 Monant nodes appearing in the top-10 rankings of at least two centrality measures. The greatest cross-metric agreement was observed for FibVID node 51591 and Monant node 141370, each identified by seven measures. In FibVID, the overlapping selections were distributed across eight distinct combinations of centrality measures, compared with nine combinations in Monant. Notably, EC and DIC selected identical sets of 10 nodes across both datasets, indicating complete agreement on their top-10 membership, although not necessarily on their ranking order. Several nodes were also jointly identified by degree, closeness and betweenness centralities, particularly in FibVID, whereas Monant exhibited more combinations involving degree, eigenvector, PageRank-based and other centrality measures. These findings reveal both convergence and variation in node prioritisation across the evaluated metrics, demonstrating that different centrality formulations can identify overlapping sets of structurally prominent nodes.

While Table 8 establishes the extent of agreement among the evaluated centrality measures, it does not identify nodes selected by only one measure. To examine this complementary dimension, Table 9 presents the exclusive top-10 selections for CCTB and ICB, identifying nodes absent from all eight alternative rankings in each dataset.

Table 9. Exclusive top-10 nodes under CCTB and ICB, relative to the other eight measures.

Dataset	Measure	Exclusive node IDs	Count
FIBVID	CCTB	10183, 104042, 105547, 106030, 106682, 10836, 109288, 110882, 111699, 112115	10
	ICB	20991, 122641, 67140, 4163, 78487, 25915	6
MONANT	CCTB	125432, 125515, 125581, 125708, 125735, 125876, 125893, 125901, 125964, 125967	10
	ICB	202198	1

Note. Exclusivity denotes $E^{d}(m) = T_{m}^{(d)} \setminus \bigcup_{j \in M \setminus \{m\}} T_{j}^{(d)}$, for the displayed ten-node selections in Tables 6 and 7. This does not establish uniqueness across the complete rankings. CCTB cutoff ties can change the selected nodes.

Table 9 identifies nodes appearing exclusively in the displayed top-10 rankings of CCTB or ICB, relative to the other eight centrality measures. In FibVID, CCTB selected 10 nodes absent from all other displayed top-10 lists, whereas ICB identified 6 such nodes. The corresponding Monant results were ten and one, respectively. This indicates that CCTB exhibited no overlap with the comparator top-10 selections in either dataset, whereas ICB showed greater agreement with the existing rankings, particularly in Monant. However, the apparent exclusivity of CCTB must be interpreted with caution, as 89 FibVID nodes and 977 Monant nodes were tied at their respective top-10 cutoffs. Consequently, alternative tie-breaking rules could produce different exclusive-node selections. These findings describe differences in top-ranked node membership rather than establishing that the proposed metrics identify previously undiscovered influential nodes or demonstrate superior predictive performance.

4.7 Structural Relationships Between Degree Centrality and Cross-Context Brokerage

The preceding analyses indicate that conventional degree centrality and the proposed context-sensitive brokerage measures capture different structural characteristics of network nodes. To examine these relationships, Figures 6 and 7 compare mean CCTB and log-transformed ICB scores across observed degree categories in FibVID and Monant, respectively. Presenting the measures separately enables their associations with node connectivity to be examined without combining centrality scores with different mathematical scales.

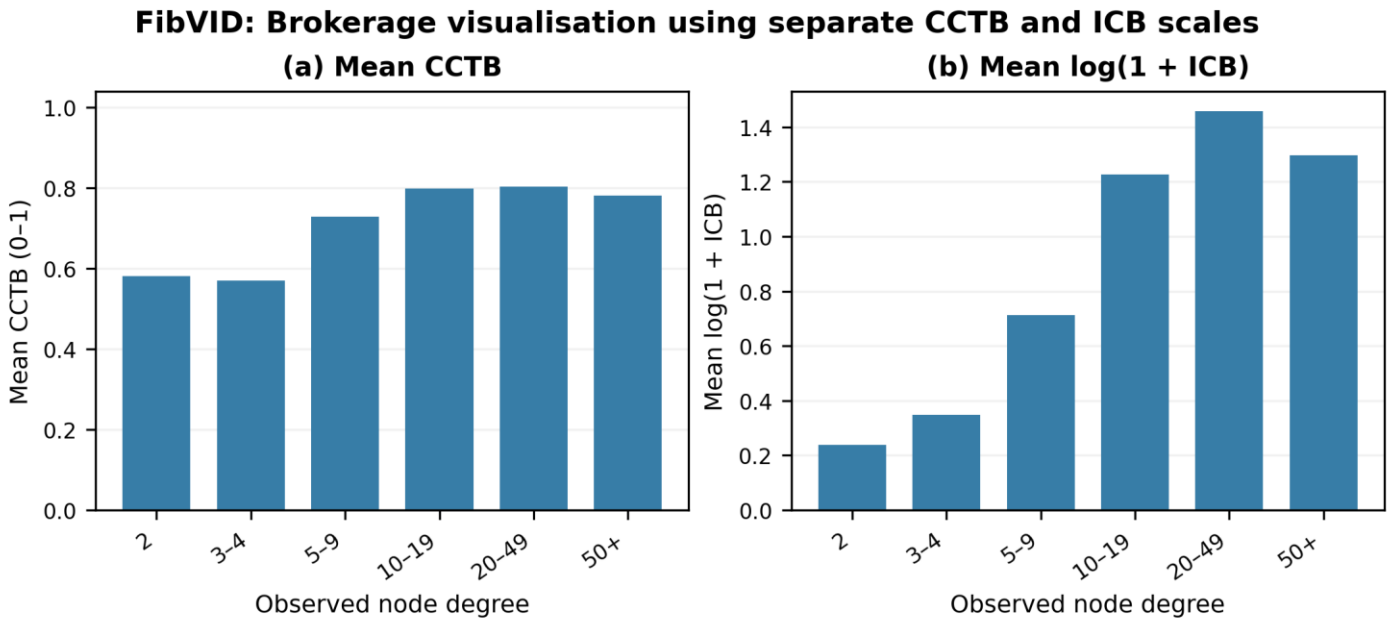


Figure 6. FibVID: Brokerage visualisation using separate CCTB and ICB scales. Panel (a) presents mean normalised CCTB across observed degree categories, while panel (b) presents mean log(1 + ICB). Both panels are restricted to nodes exhibiting positive brokerage scores. The comparison illustrates how the relative extent of cross-context opportunities and scarcity-weighted context-pair coverage vary with node connectivity.

Figure 6 illustrates distinct relationships between node connectivity and the proposed brokerage measures in FibVID. Mean CCTB generally increases across the lower and intermediate degree categories, from approximately 0.58 among degree-two nodes to 0.80 in the 20–49 category, before declining slightly to approximately 0.78 among nodes with degree 50 or greater. In contrast, mean log(1 + ICB) increases more consistently across the lower degree categories, reaching its highest value in the 20–49 category before declining modestly in the highest degree category. These patterns indicate that the relative proportions of cross-context interaction opportunities and scarcity-weighted context-pair coverage do not vary in the same way with connectivity. Nevertheless, the results are descriptive and conditional on positive brokerage scores; they do not establish that either metric provides independent predictive information beyond degree centrality.

Figure 7 extends the analysis to Monant by comparing the mean CCTB and log-transformed ICB scores across observed node-degree categories. The two-panel presentation allows the relationship between connectivity and each proposed brokerage measure to be examined separately, while accounting for their different mathematical scales.

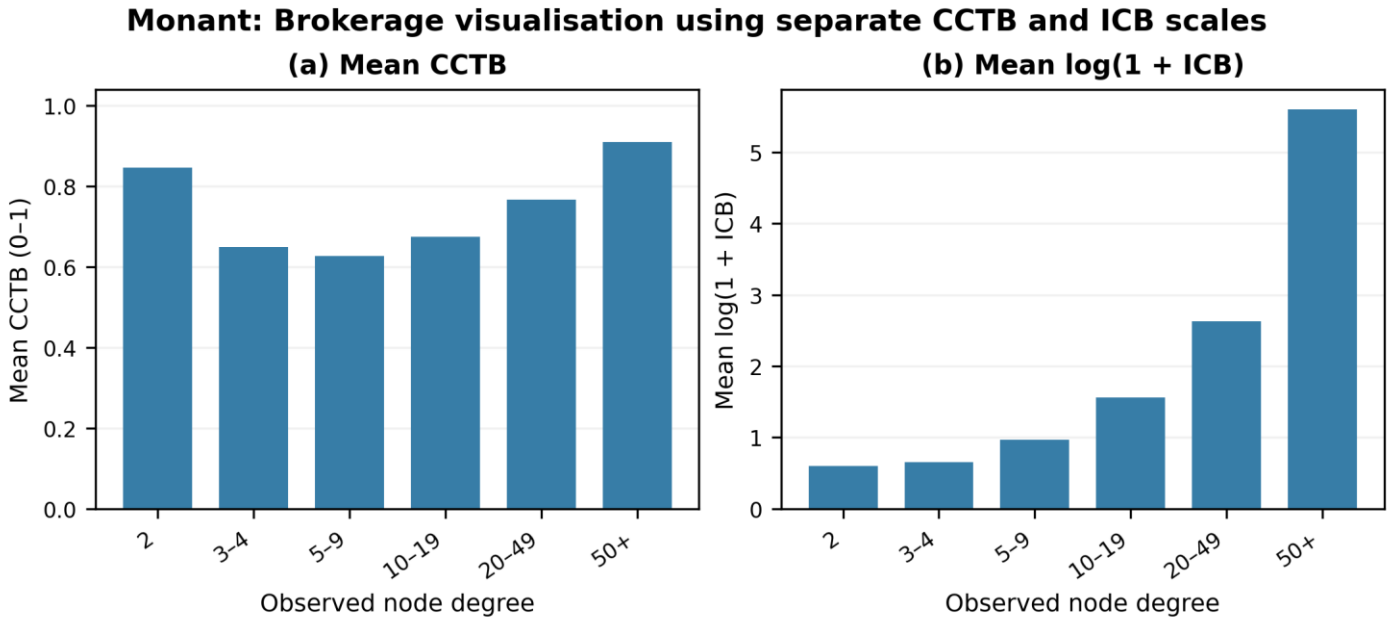


Figure 7. Monant: Brokerage visualisation using separate CCTB and ICB scales. Panel (a) presents mean normalised CCTB across observed degree categories, while panel (b) presents mean log(1 + ICB). Both panels are restricted to nodes exhibiting positive brokerage scores. The comparison illustrates the relationship between connectivity and two distinct dimensions of context-sensitive brokerage in the Monant network.

Figure 7 reveals different relationships between node connectivity and the two proposed brokerage measures in Monant. Mean CCTB does not increase monotonically across degree categories, declining from approximately 0.85 among degree-two nodes to 0.63 in the 5–9 category before increasing to approximately 0.91 among nodes with degree 50 or greater. By contrast, mean log(1 + ICB) generally increases with node degree, with the highest values observed in the 50+ category. These patterns suggest that the relative proportion of cross-context interaction opportunities and scarcity-weighted context-pair coverage respond differently to increasing connectivity. However, the comparisons are descriptive, restricted to brokerage-positive nodes and do not establish independent predictive value beyond degree centrality.

4.8. Cross-Dataset Replication and Generalisability Assessment Using MC-Fake

To assess the generalisability of the findings from FibVID and Monant, CCTB and ICB were additionally evaluated on the MC-Fake dataset [18], an existing misinformation-related dataset comprising news records and associated social media interactions. Following reconstruction and deduplication, the analytical network contained 1,952,282 directed interactions across 26,052 information contexts. Consistent with the primary evaluation protocol, centrality scores were derived from the first 75% of chronologically ordered interactions, while subsequent interactions were reserved for predictive assessment.

Building on the primary predictive comparisons reported in Table 2, Table 10 summarises the incremental contribution of CCTB and ICB across FibVID, Monant and MC-Fake. The inclusion of the proposed measures increased mean Average Precision by 0.0576, 0.0078 and 0.0535,

respectively, indicating a consistent direction of improvement across the three independently evaluated networks. Additional MC-Fake analyses indicated that ICB individually achieved higher Average Precision than CCTB or their combined specification under the corresponding baseline. The structural comparison also revealed identical top-10 selections between eigenvector centrality and the operational DIC proxy, consistent with the findings from FibVID and Monant.

Table 10. Cross-dataset comparison of the incremental predictive contribution of CCTB and ICB to future cross-context brokerage.

Dataset	Baseline AP	Baseline + CCTB + ICB AP	Absolute ΔAP	Relative AP increase
FIBVID	0.2966	0.3542	+0.0576	+19.4%
MONANT	0.0847	0.0925	+0.0078	+9.2%
MC-FAKE	0.0655	0.1190	+0.0535	+81.7%

Note. Baseline comprises degree centrality and context breadth. Values are reported mean Average Precision (AP) for within-network, node-disjoint evaluations. Absolute ΔAP is the augmented AP minus the baseline AP; relative increase is ΔAP divided by baseline AP. Because outcome prevalence and network characteristics differ, AP values should not be interpreted as directly comparable across datasets. The MC-Fake results are exploratory and subject to the methodological qualifications described in the text.

These findings provide preliminary evidence that the complementary predictive contribution of context-sensitive brokerage extends across heterogeneous information networks. Nevertheless, the results represent separately fitted within-network evaluations rather than direct cross-dataset model transfer. Furthermore, approximate closeness and betweenness computations, operational centrality proxies, extensive CCTB score ties and limited ICB truncation introduce additional uncertainty into the MC-Fake comparison. Consequently, the findings support qualified cross-dataset replication rather than definitive generalisability.

4.9. Centrality-Guided Node-Removal Simulation Results

The supplementary node-removal simulations revealed dataset-dependent differences in intervention effectiveness (Table 11). Degree centrality achieved the greatest reduction in future cross-context brokerage opportunities in FibVID (98.73%), while context breadth performed best in Monant (20.77%). In MC-Fake, CCTB and ICB achieved the highest reductions (96.81% each). These findings indicate that the effectiveness of centrality-guided targeting varies across network structures, with the reported percentages representing simulated structural reductions rather than observed decreases in misinformation propagation.

5. Discussion

The findings highlight the distinction between general network prominence and context-sensitive structural brokerage in health-information networks. CCTB quantifies the proportion of directed two-step interaction opportunities that connect different information contexts, whereas ICB measures the scarcity of alternative nodes that can bridge the same ordered context pairs. These measures therefore capture complementary aspects of contextual mediation and structural substitutability that are not explicitly represented by conventional centrality measures.

The predictive results demonstrate that context breadth is informative for subsequent network activity, while CCTB and ICB provide complementary information for future brokerage (Tables 2 and 3). The expanded nine-metric evaluation particularly supports the incremental contribution of ICB, although adding CCTB did not consistently improve performance (Table 5). Thus, the proposed measures offer context-sensitive predictive information rather than universal superiority over established centralities.

Differences in predictive performance may partly reflect variation in network organisation and interaction processes. FibVID captures posting relationships associated with fact-checked claims, whereas Monant records replies within discussions surrounding medical articles. These interactions do not necessarily represent equivalent information-sharing behaviours, and their differing contextual structures may influence the availability of observable brokerage opportunities and alternative bridging nodes. Accordingly, differences in predictive performance should be interpreted in relation to network characteristics rather than as evidence of misinformation endorsement, persuasion or susceptibility.

The MC-Fake extension provides preliminary evidence that the predictive contribution of context-sensitive brokerage is reproducible beyond the two primary networks. Incorporating CCTB and ICB increased mean Average Precision from 0.0655 to 0.1190, consistent with the direction of improvement observed in FibVID and Monant (Table 10). However, MC-Fake ablation results indicated that ICB, when evaluated individually, achieved stronger predictive performance than the combined specification under the corresponding baseline. The variation in the magnitude of improvements and feature contributions suggests that the usefulness of context-sensitive brokerage depends on the network and evaluation setting. These results support qualified cross-dataset replication rather than direct transfer of predictive models.

The structural comparisons further demonstrate how context-sensitive brokerage differs from conventional network prominence. Degree centrality measures immediate connectivity, closeness centrality characterises network accessibility, betweenness centrality identifies positions along shortest paths [15], and eigenvector centrality accounts for the prominence of connected neighbours [16]. Although these measures capture established dimensions of node importance, they do not explicitly differentiate the information contexts associated with interactions. This distinction is relevant to brokerage, where intermediary positions and alternative connections influence structural roles [17]. The ranking and overlap analyses (Tables 4 and 6–9) indicate that ICB can identify nodes not prioritised by conventional measures, while Figures 6 and 7 demonstrate that CCTB and ICB exhibit distinct patterns across degree categories. Nevertheless, ICB represents the scarcity of observed substitutes, not proof that a node is indispensable, since unrecorded interactions may provide alternative connections.

The proposed measures also complement the established PC, MVC and DIC metrics, which address propagation-related prominence, vulnerability-related characteristics and dynamic influence, respectively [1]. CCTB and ICB instead operationalise cross-context mediation and structural substitutability. These conceptual differences do not establish universal superiority, particularly because the operational PC, MVC and DIC comparators incorporate the computational approximations described in Section 3.

5.1. Centrality-Guided Interventions for Mitigating Misinformation Propagation

The supplementary node-removal simulations extend the predictive evaluation by examining whether targeting historically ranked nodes reduces future cross-context brokerage opportunities. Table 11 compares the structural intervention outcomes of degree centrality, context breadth, the operational PC comparator, CCTB and ICB under a fixed 5% historical-node removal budget across the three networks. The matrix also identifies the strongest-performing strategy in each dataset and summarises its potential public health application and interpretative limitations. These structural counterfactuals are distinct from the node-disjoint Average Precision results in Table 10 and should not be interpreted as observed public health intervention outcomes.

Table 11. Cross-dataset structural intervention outcomes and public health implications

FibVID | Highest reduction: Degree				
Degree	**Context breadth**	**PC**	**CCTB**	**ICB**
98.73%	**97.05%**	**86.14%**	**97.70%**	**97.70%**
Potential public health application: Prioritise highly connected interaction pathways for monitoring and appropriately targeted corrective information.				
Interpretative caution: Degree is the strongest strategy. Simulated structural reductions do not establish changes in misinformation exposure.				
Monant | Highest reduction: Context breadth				
Degree	**Context breadth**	**PC**	**CCTB**	**ICB**
20.13%	**20.77%**	**20.33%**	**20.63%**	**20.58%**
Potential public health application: Identify users participating across distinct discussion environments to prioritise verification and contextualised information responses.				
Interpretative caution: The differences between strategies are small. Participation in multiple contexts does not demonstrate cross-context transmission.				
MC-Fake | Highest reduction: CCTB / ICB				
Degree	**Context breadth**	**PC**	**CCTB**	**ICB**
93.08%	**92.16%**	**47.45%**	**96.81%**	**96.81%**
Potential public health application: Identify cross-context brokers whose pathways may warrant claim verification and targeted corrective messaging.				
Interpretative caution: Tie-breaking affects results, and MC-Fake is not exclusively health-related; generalisation to public health requires validation.				

Note. Values indicate percentage reductions in simulated future cross-context brokerage opportunities following removal of the top-ranked 5% of historical nodes. These findings suggest that targeting structurally influential nodes could potentially reduce pathways facilitating misinformation propagation by the reported percentages. However, equivalent reductions in the spread of actual misinformation were neither directly observed nor established.

Table 11 demonstrates that the effectiveness of centrality-guided intervention targeting varies across network contexts. In FibVID, degree centrality achieved the greatest reduction in simulated cross-context brokerage opportunities (98.73%), marginally exceeding CCTB and ICB (97.70% each). In Monant, context breadth performed best (20.77%), although differences between the evaluated measures were small. Conversely, CCTB and ICB achieved the highest reductions in MC-Fake (96.81% each), substantially outperforming the operational PC proxy (47.45%). These findings suggest that conventional network prominence and cross-context brokerage capture distinct structural properties, and that the relative effectiveness of each targeting strategy depends on

network characteristics. Importantly, removing the same proportion of nodes does not imply equivalent intervention costs or numbers of disrupted interactions. The comparisons therefore establish relative effectiveness under a fixed node budget, rather than under an equal-edge-removal or equal-cost constraint.

Although CCTB and ICB represent distinct concepts of brokerage magnitude and structural irreplaceability, their nearly identical intervention outcomes do not establish an additional practical advantage for ICB under the evaluated conditions. From a public health perspective, the findings support context-sensitive prioritisation of information monitoring, claim verification and corrective messaging rather than reliance on a single centrality measure. However, the reported reductions represent simulated structural counterfactuals, not observed decreases in misinformation transmission or exposure. The operational PC approximation, ranking sensitivity and the broader topical coverage of MC-Fake further limit generalisability. These limitations and outstanding robustness requirements are addressed in Section 6.

6. Limitations and robustness requirements

The interpretation of the findings is constrained by the information available in the source datasets. Monant lacks independently adjudicated article–claim stance mappings, preventing the classification of individual replies as endorsing, rejecting or correcting misinformation. Similarly, although FibVID provides claim-level categories, these do not establish individual users' agreement with the associated content. Consequently, the reconstructed networks capture observable health-information interactions rather than verified misinformation endorsement or transmission.

Temporal coverage and statistical dependence introduce further limitations. Although chronological partitioning restricted feature construction to the first 75% of observed interactions, the available timestamps, claim identifiers and article URLs may not fully represent the evolution of information contexts. Moreover, the five node-disjoint evaluation splits originated from the same reconstructed networks, meaning that reported standard deviations reflect variation in node allocation rather than uncertainty across independent network samples. Shared neighbours, interactions and contexts further limit the assumption of statistical independence. The node-removal simulations additionally assume that the observed future network remains unchanged apart from the excluded targets. They do not model behavioural adaptation, alternative communication pathways or replacement brokerage following intervention.

Centrality estimation also introduces computational uncertainty. Deterministic context-pair truncation affected 37 Monant nodes and 11 MC-Fake nodes, potentially altering ICB scores and the representation of alternative bridging pathways. In MC-Fake, closeness and betweenness centralities were approximated, while PC, MVC and DIC were implemented as operational proxies rather than exact reproductions of their original formulations. Extensive ties in CCTB scores further complicated the interpretation of top-ranked node selections, particularly in Monant and MC-Fake. These limitations restrict direct comparisons between the proposed and established centrality measures.

The MC-Fake replication presents additional data-provenance limitations. Interaction timestamps were inferred from tweet identifiers without independent validation against platform metadata, while conflicting parent-user mappings introduce uncertainty into the reconstructed network. Although chronological CCTB stability, feature ablation, activity threshold, and selected model-sensitivity analyses supported the persistence of certain predictive improvements, the findings do

not fully establish harmonised cross-network performance. Moreover, subsequent interactions and cross-context brokerage measure observable network activity rather than persuasion, changes in perceived veracity, health behaviour or causal misinformation diffusion.

Future research should prioritise exact or statistically validated centrality estimation, context-label permutation tests, claim- or article-block resampling, alternative context-identification procedures and independent network reconstruction. Direct cross-dataset predictive transfer and evaluation under alternative chronological partitions would further establish whether the observed brokerage patterns persist across network structures and information environments. Until such validation is undertaken, the findings support structural characterisation and qualified cross-dataset replication, rather than universal generalisability, causal transmission or mitigation efficacy.

7. Conclusions

Two novel, complementary context-sensitive brokerage centrality measures were mathematically specified and evaluated using reconstructed FibVID and Monant interaction networks, with additional replication undertaken using MC-Fake. CCTB quantifies the proportion of directed two-step interaction opportunities across different information contexts, whereas ICB weights ordered context-pair connections based on the scarcity of alternative bridging nodes. Unlike the established DIC, MVC and PC measures, the proposed metrics explicitly address contextual mediation and structural substitutability rather than conventional network prominence or propagation-related characteristics.

The experimental findings indicate that context breadth is a strong individual predictor, while incorporating CCTB and ICB improves Average Precision for future brokerage beyond the effects of degree centrality and context breadth. The MC-Fake extension provided additional evidence of this pattern, with positive Average Precision improvements observed across all three heterogeneous networks, although the magnitude varied considerably. Furthermore, ICB frequently demonstrated a stronger individual predictive contribution than CCTB, and their combination did not consistently outperform ICB alone.

Overall, the findings support the complementary value of context-sensitive brokerage measures for identifying future cross-context interactions. The node-removal simulations further demonstrate that effective structural targeting strategies vary across network contexts, with CCTB and ICB performing particularly well in MC-Fake. Nevertheless, the results establish neither causal misinformation transmission nor real-world mitigation effectiveness. Future research should prioritise independent network validation, improved centrality estimation and intervention modelling that accounts for alternative brokerage pathways.

Declarations

Declaration of generative artificial intelligence use

OpenAI ChatGPT was used to assist with manuscript organisation, language refinement, Python code development and debugging, and the preparation of reproducibility documentation. AI-generated suggestions and computational procedures were subject to author review against the available research data, analytical outputs and methodological specifications. AI assistance was not treated as independent empirical evidence or verification of the study's findings. The author retained responsibility for the design, implementation, interpretation and reporting of the research and takes full responsibility for the accuracy, integrity and final content of the manuscript.

Ethics approval and consent to participate.
This study is a computational secondary analysis of previously distributed research datasets. No participants were recruited, contacted, re-identified, or subjected to an experimental intervention. Institutional and dataset-specific reuse conditions must be confirmed before redistributing any source records.

Data availability.
The study analysed three previously published datasets. FibVID comprises claim-associated social-media propagation records, Monant contains medical article discussions and associated interaction records, and MC-Fake provides news-related social-media interactions across multiple topics. The original datasets were not generated by the author and should be obtained from their respective providers in accordance with applicable access and reuse conditions. Dataset reconstruction procedures, derived analytical specifications and supporting reproducibility documentation are provided in the accompanying Zenodo repository (https://doi.org/10.5281/zenodo.23868500) [19]. Public availability of original interaction records and node-level outputs remains subject to the relevant dataset permissions and privacy considerations.

Code availability
Python implementations of Cross-Context Two-Step Brokerage (CCTB) and Irreplaceable Context Brokerage (ICB), network preprocessing procedures, centrality calculations, chronological evaluation scripts, predictive modelling workflows, mathematical validation tests and reproducibility documentation are archived on Zenodo (https://doi.org/10.5281/zenodo.23868500) [19]. The repository provides instructions for reproducing the computational procedures and documents relevant implementation approximations and dataset-specific limitations. Reproduction of the reported analyses requires access to the original datasets under their respective reuse conditions and the dataset-specific configurations documented in the accompanying repository.

References

[1] Sikosana M, Maudsley-Barton S, Ajao O. Analysing health misinformation with advanced centrality metrics in online social networks. PLOS Digit Health. 2025;4:e0000888. doi:10.1371/journal.pdig.0000888.
[2] Kim J, Aum J, Lee S, Jang Y, Park E, Choi D. FibVID: comprehensive fake news diffusion dataset during the COVID-19 period. Telemat Inform. 2021;64:101688. doi:10.1016/j.tele.2021.101688.
[3] Srba I, Pecher B, Tomlein M, Moro R, Stefancova E, Simko J, Bielikova M. Monant medical misinformation dataset: mapping articles to fact-checked claims. SIGIR. 2022:2949–2959. doi:10.1145/3477495.3531726.
[4] Holme P, Saramäki J. Temporal networks. Phys Rep. 2012;519:97–125. doi:10.1016/j.physrep.2012.03.001.
[5] Burt RS. Structural Holes: The Social Structure of Competition. Harvard University Press; 1992.
[6] Kempe D, Kleinberg J, Tardos É. Maximizing the spread of influence through a social network. KDD. 2003:137–146. doi:10.1145/956750.956769.
[7] Freeman LC. A set of measures of centrality based on betweenness. Sociometry. 1977;40:35–41. doi:10.2307/3033543.
[8] Page L, Brin S, Motwani R, Winograd T. The PageRank citation ranking: bringing order to the web. Stanford InfoLab Technical Report; 1999.

[9] Newman MEJ. Networks: An Introduction. Oxford University Press; 2010.
[10] Valente TW. Social Networks and Health: Models, Methods, and Applications. Oxford University Press; 2010.
[11] Gneiting T, Raftery AE. Strictly proper scoring rules, prediction, and estimation. J Am Stat Assoc. 2007;102:359–378. doi:10.1198/016214506000001437.
[12] Sikosana M, Maudsley-Barton S, Ajao O. Linguistic patterns in pandemic-related content. Front Artif Intell. 2025;8:1627522. doi:10.3389/frai.2025.1627522.
[13] Kim J, Aum J, Lee S, Jang Y, Park E, Choi D. FibVID: comprehensive fake news diffusion dataset during the COVID-19 period. Zenodo dataset. 2021. https://doi.org/10.5281/zenodo.4441377.
[14] Srba I, Pecher B, Tomlein M, Moro R, Stefancova E, Simko J, Bielikova M. Monant Medical Misinformation Dataset. Zenodo. 2022. https://doi.org/10.5281/zenodo.5996864.
[15] Freeman, LC. Centrality in social networks: Conceptual clarification. Social Networks, 1979;1(3), 215–239. https://doi.org/10.1016/0378-8733(78)90021-7.
[16] Bonacich, P. Power and centrality: A family of measures. American Journal of Sociology, 1987;92(5), 1170–1182. https://doi.org/10.1086/228631.
[17] Burt, R. S. (1992). Structural Holes: The Social Structure of Competition. Harvard University Press.
[18] Min E, Rong Y, Bian Y, Xu T, Zhao P, Huang J, et al. Divide-and-conquer: Post-user interaction network for fake news detection on social media. In: Proceedings of the ACM Web Conference 2022. New York: Association for Computing Machinery; 2022. p. 1148–1158. doi: 10.1145/3485447.3512163
[19] Sikosana M. Context Brokerage Centrality: Reproducibility Code and Documentation for CCTB and ICB [software]. Zenodo; 2026. Available from: https://doi.org/10.5281/zenodo.23868500