



GDI

Global
Disinformation
Index

Risk Assessment: South African Media Market

www.disinformationindex.org

Author: Lucas Wright.

Researchers: Chris Roper (Code for Africa), Amanda Strydom (Code for Africa) and Scarlet Dawson.

Design: www.designbythink.co.za

The Global Disinformation Index is a UK-based not-for-profit that operates on the three principles of neutrality, independence and transparency. Our vision is a world in which we can trust what we see in the media. Our mission is to restore trust in the media by providing real-time automated risk ratings of the world's media sites through a Global Disinformation Index (GDI). The GDI is non-political. Our Advisory Panel consists of international experts in disinformation, indices and technology. For more information, visit www.disinformationindex.org

GDI Global
Disinformation
Index



January 2020. Published under a Creative Commons License (CC BY-NC-SA 4.0)

Table of contents

Preface	4
Introduction	5
The South African media market: Key features and scope	8
The GDI framework: Theory and methodology	9
Disinformation risk ratings	11
Conclusion	16
Annex: Methodology	17
Endnotes	18

Preface

Since the invention of the web, how we live our lives online – and off – has changed in countless ways. This includes how news is funded, produced, consumed and shared.

With these shifts in the news industry have come risks. Disinformation is one of them. Disinformation has been used as a tool to weaponise mass influence and disseminate propaganda. No country or media market is immune from this threat.

To combat disinformation, we need to find ways to defund and disrupt the system. This is where the Global Disinformation Index (GDI) has set its focus.

At the GDI, we believe that an independent, trusted and neutral risk rating of news sites' disinformation risks is needed. These risk ratings can be used by advertisers and ad tech companies to ensure that where they direct their online ad spends is aligned with their own brand safety and risk mitigation strategies for disinformation.

The GDI aspires to offer a trusted and neutral assessment about a news domain's risk of disinforming. By looking at structural, content, operational and context indicators, the GDI will provide a domain-level rating about a news site's risk of disinforming a user.

We have designed the organisation – and our risk ratings – based on the three pillars of neutrality, independence and transparency.

Neutrality: We are apolitical, global, and evidence-based. We are establishing a governance structure which aspires to the highest standards of global corporate governance.

Independence: The GDI is established as a not-for-profit entity. We receive no benefit from the risk ratings we give to a particular site. We exist solely to assess online news domains' risk of disinforming their readers. We are advised by a panel of international experts: Anne Applebaum (London School of Economics), Peter Pomerantsev (London School of Economics) and Miguel Martinez (Signal Media).

Transparency: The GDI's rating criteria, assessments and methodology will be community-driven and made publicly auditable. A dispute mechanism will be developed and made available for the owners of domains that disagree with their rating. In completing the report, all sites were contacted when the assessment began and were provided with their individual scores to discuss, review and adjust where relevant.

The following report presents the results of a piloting of the risk rating methodology in South Africa. South Africa has been chosen given its historical and respected media market; its high-level of readers consuming their news online; its robust and growing programmatic advertising market; and its past experience with countering disinformation campaigns targeting online readers and public debates.

We consider the findings from the pilot as the start of a discussion among news sites, advertisers and ad tech companies on how the GDI risk ratings can be used to strengthen the funding of an independent, diverse and trusted media. Please join us in this journey.

Introduction

The harms of disinformation¹ are proliferating around the globe – threatening our elections, our health, and our shared sense of accepted facts.

Websites masquerading as news outlets are driving and profiting financially from the situation. The goal of the Global Disinformation Index (GDI) is to cut off the revenue streams that incentivise and sustain the spread of disinformation. Using both artificial and human intelligence, the GDI has created an assessment framework to rate the disinformation risk of news domains. This work was done in collaboration with ten noted experts working on media, fact-checking and disinformation.² The GDI ratings will give greater control to advertisers and ad tech companies over where adverts appear.

Companies' programmatic adverts are a critical financial lifeline to disinforming sites. They create a perverse incentive for financially motivated actors to traffic in disinformation to get more clicks. The GDI estimates that at least US\$235 million in programmatic adverts annually is placed by ad tech companies on known disinformation sites.³ The GDI's risk ratings can cut off this funding by providing a neutral and independent assessment of disinformation risk that can inform the brand safety and risk mitigation strategies of advertisers and ad tech companies.

The GDI risk ratings are relevant for more established and new media sites that are part of a country's broader media landscape. This includes highly-acclaimed and traditional news outlets that have online presences. The shifting nature of news consumption to online interfaces has meant a shift in the editorial and advertising policies and practices. As a result, there are risks and challenges now posed by disinformation actors to these sites that merit assessment. This is the objective of the risk rating and rationale for the pilot.

For example, existing operational policies governing user-generated content may not have kept pace with the areas that they cover, such as how hate speech and privacy protections are handled on the comments section of a site. New issues like artificially and algorithmically-generated content (stories as well as videos and other visuals) are so fresh that sites may not even have thought of having policies in place.

Moreover, the current advertising model based on user clicks to generate site revenue has created challenges for media markets across countries and regions. This has meant the need for a news site to capture the clicks of an online reader whose attention is increasingly bombarded by different news headlines and sites of varying quality, reputation and credibility. In countries where readers rely on online news via their phones, this information overload is challenging at best. In South Africa, it is estimated that 90 percent of readers get their news online, and mostly via their mobile phones (76 percent).⁴

More online consumption of news does not necessarily equal more trusted content – or trust in media. Overall, the media are the least trusted institution in the world.⁵

Recent research suggests that 41 percent of people in South Africa distrust the media and that 43 percent don't trust the news they see on social media.⁶ While trust in the media has historically been strong in South Africa, it is now in decline.⁷ Disinformation has a role to play. Findings show that 70 percent of those surveyed in South Africa are concerned by what news is real and “fake” online.⁸

Worries over media quality and disinformation also have led some advertisers to bluntly block news sites entirely due to concerns about risks to their brand safety. This has exacerbated the declining advertising revenues that news sites have suffered in recent years due to the exponential increase in content and news-related content online.

There is a need for advertisers to have a more trusted, neutral and independent standard to understand the nuances and contours of a news site's disinformation risks. The GDI risk rating addresses this need and provides advertisers and ad tech companies with greater information about a range of disinformation flags related to a site's structure, content, operations and context. As a result, the GDI risk ratings will also cover the top-end of highly-used news sites across all types of media and advertising markets – as well as the “long tail” of smaller news sites. The GDI hopes that by providing a more rigorous assessment of disinformation risk for all news sites, formerly “news-averse” advertisers may be persuaded to support quality (“low-risk”) news to a greater extent.

The following report presents preliminary findings pertaining to disinformation risks for the media market in South Africa, based on a pilot study that covered 30 of some of the country's top news domains.⁹ The data provides an initial snapshot of the overall strengths and challenges that these sites have to mitigate disinformation risks. **The report is the first time such a scoping and scoring for the South African media market has been done.** The report and its findings of the pilot are intended to be discussed and debated – and should be taken in that spirit. We will use exchanges with the sites to improve our methodology and to help sites reduce their disinformation risks.¹⁰

We have taken care at every stage to minimise bias and provide objective data. We recognise the need to work together across the media and ad industry to provide an independent, transparent and trusted assessment of news sites.¹¹

Key Findings: South Africa

In looking at the media landscape for South Africa, GDI's assessment found that:

The current advertising model to drive traffic and clicks is creating potential disinformation risks for a respected and vibrant media market.

- Many mainstream and local news sites in South Africa present some degree of disinformation risk. These flags tend to concern their operational and editorial integrity as well as perceptions of trust in their site. In contrast, these same sites show relatively low disinformation risks in regards to how they present content in the stories which we reviewed.
- There appears to be a generalised trend when it comes to a site's prevalent use of clickbait headlines. The media experts surveyed in South Africa perceive that most South African sites in the sample peddle in clickbait titles.¹² This trend likely reflects the need to use eye-catching titles that help to drive traffic to news sites which then can be converted to revenue streams from a site's adverts.

Many sites do not have all of the operational checks and balances in place which are needed to create safeguards against disinformation risks. When policies do exist, they are often hidden away or hard to find.

- These policies, as outlined by the Journalism Trust Initiative, include having a statement of editorial independence and comprehensive policies for user-generated and automatically-generated content.¹³

Perceptions of brand trust in South African news sites reflect the overall crisis of confidence in media by its users.

- Expert perceptions of accurate news and overall trust in South African news sites is relatively modest, affecting overall risk levels for the market. According to our independent survey of media sector experts, many do not feel that the sites generally carry highly accurate content or regularly correct their errors.
- On a positive note, there appears to be strong perceptions that South African news sites generally do differentiate clearly between opinion and straight news articles. This is important for readers to avoid mistaking online opinion pieces as factual news.

South African sites seem to overall provide content that limits sensationalism and negative coverage.

- Looking across the sample, very few sites were found to use articles to negatively target specific individuals or groups. These sites also rarely publish articles that make use of overly charged or emotionally-driven content.

All of these findings come from the pilot research led by the GDI from September to November 2019. The market analysis is based on 13 disinformation flags that were assessed for South Africa by an analyst and by an independent survey of over 50 experts.¹⁴ While we present the average scores for the market sample, no specific sites are named or individually scored in this report.¹⁵ The scores should be seen as offering initial insights into the South African media market and its overall levels of potential disinformation risk. The pilot results are being debated and refined with stakeholders from news sites, advertisers and ad tech industry. (The annex of this report outlines the assessment framework).

The South African media market: Key features and scope

South Africa’s news consumption is increasingly dominated by the internet. According to the Reuters Institute,¹⁶ 90 percent of South Africans access news online, far more than TV (68 percent) or print (40 percent).

The market for online news is dominated by News24 and SABC. Based on the most recent survey, the number of South Africans who say they have accessed these two online sources in the past week is 70 percent and 45 percent, respectively.¹⁷ Recent research suggests that 41 percent of people in South Africa distrust the media and that 43 percent don’t trust the news they see on social media.¹⁸

The country also has a growing market for online advertising. In 2018, nearly R5 billion (US\$3.5 million) was spent on digital display adverts, more than a 24 percent increase over 2017.¹⁹ In South Africa, this combination of a robust demand for online news and a growing market for ad monies provides opportunities to direct more online ad revenues to trustworthy news sites – but it also offers increased incentives for actors trying to make money from the clicks generated by disinformation.

For this study, we defined the South African media market based on an initial list of nearly 70 news sites, which included well-known national outlets, tabloids, regional newspapers, and blogs. We then worked with local media experts to refine the list based on each site’s reach and relevance. We defined reach and relevance based on a site’s Alexa rankings, Facebook followers, and Twitter followers. We also consulted with local experts to identify domains with lower reach but high relevance among decision makers and niche audiences (see Figure 1).²⁰

Figure 1. South African media sites assessed

1. Black Opinion	11. Mail & Guardian	21. The Benoni City Times
2. Business Live	12. Money Web	22. The Citizen
3. BusinessTech	13. MyBroadband	23. The Daily Maverick
4. City Press	14. Netwerk24	24. The Daily Sun
5. ENCA	15. News24	25. The Kempton Express
6. Fin24	16. SABC	26. The New Soweto
7. Herald Live	17. Sowetan Live	27. The Roodepoort Record
8. IOL (includes Star, Voice and Isolezwe)	18. Sunday Times / Times Live	28. The South African
9. Live Report	19. Tech Central	29. The South Coast Herald
10. Low Velder	20. The Alberton Record	30. The Zululand Observer

The GDI framework: Theory and methodology

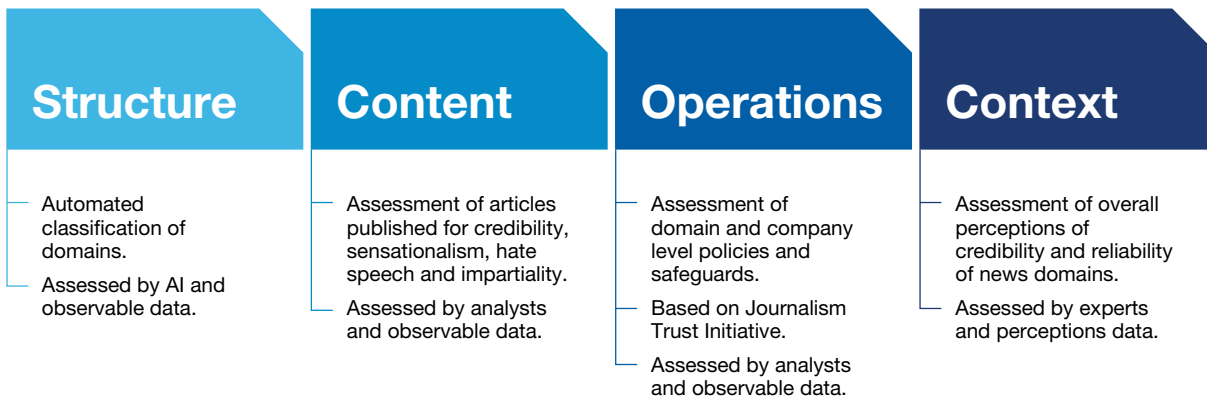
Our index is not an attempt to identify truth and falsehoods. It does not label any site as a disinformation site – or, inversely, as a trusted news site. Rather, our approach is based on the idea that a range of signals, taken together, can indicate a site’s risk of carrying disinformation.

Disinformation risk flags relate to the structure (i.e. technical features), content (i.e. reliability of content), operations (i.e. operational and editorial integrity) and context (i.e. expert perceptions of brand trust) of the site. The output of the index is therefore the site’s overall disinformation risk level, rather than the truthfulness or journalistic quality of the site. The aim is to provide advertisers and ad tech companies with a trusted, independent and neutral assessment of potential disinformation risks for a site. The assessment is applicable to all media markets and all types of sites, including mainstream media and more developed media markets.

A site’s disinformation risk level is based on that site’s aggregated score across all of the pillars and indicators (see Figure 2). A site’s overall score ranges from 0 (maximum risk level) to 100 (minimum risk level). This report presents the findings of disinformation risks from the pilot run for the South African media market. The findings are based on the three pillars that were manually reviewed: **‘Content’**, **‘Operations’**, and **‘Context’**.²¹

The ‘Content’ and ‘Operations’ pillars are designed to capture discrete, observable features of a domain by analysing a snapshot of a particular moment in time. This approach is effective at mitigating bias and standardising our analysis across domains and countries, but it is limited in scope. Historical information about a domain’s content and practices is not captured by these pillars – nor are less observable disinformation flags (such as regularly disinforming readers by saying nothing about a story or topic). Both of these limitations are addressed by the ‘Context’ pillar, which assesses long-term trends and indicators that are harder to measure. In this report, two-thirds of a domain’s score is based on a snapshot of observable features (through the ‘Content’ and ‘Operations’ pillars), while the final third comes via an independent expert survey that contextualises our findings. Over 50 media experts were asked a series of questions about domains which they knew operationally and editorially. The survey was conducted by the respected global opinion and data company YouGov.

Figure 2. Overview of the GDI disinformation risk assessment



The ‘Structure’ pillar is used as the first step to determine which sites require a deeper, human review. Sites are assessed by a machine-learning algorithm prototype that is trained on metadata from thousands of websites known for regularly disinforming readers. It identifies these domains according to technical features²² of the website itself, and currently produces a binary assessment: it either is or is not a high-risk disinformation site.

The GDI’s work builds on initial experiments performed in 2018 under a grant from the Knight Foundation Prototype Fund. In this demonstration, we trained a neural network classifier on hundreds of pre-labelled high-risk and low-risk news sites, focusing on over twenty different technical metadata signals. The prototype classifier correctly identified 98.8 percent of the domains that had been pre-labelled as “high-risk” for carrying disinformation. This prototype demonstrates the predictive power of metadata and other computational signals to rate the disinformation risk of news sites.

For this pilot of the index, the structural indicators were used only as a filter for selecting domains in need of human review. As such, their scores on this pillar were not used to calculate the final risk rating.²³ None of the domains included in our sample were identified as high-risk disinformation sites by the automated classifier. As the sample is composed of some of the most popular sites in the South African media market, they would not be expected to share structural features with high-risk sites.

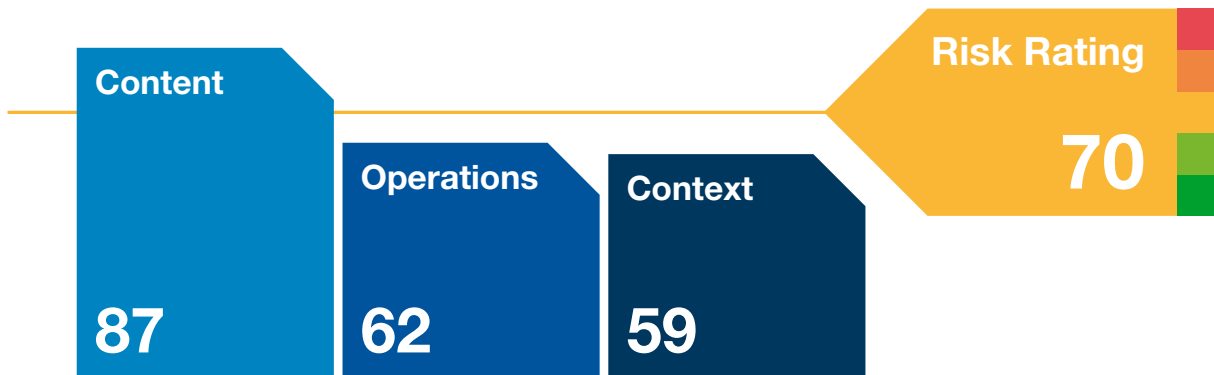
The highest-risk domains within the sample consist largely of sites that score well on their content, but entirely fail to meet standards for editorial and operational policies.

Overall, our review found limited disinformation risk across the 30 site sample in terms of the reliability of content, but there are latent risks when it comes to the disclosure of operational and editorial policies and independent expert perceptions of site trustworthiness

(see Figure 4 below). The current scores capture the risks posed by potential integrity breakdowns, which could eventually trigger higher content-related risks for the stories covered on these sites.

For example, if a new owner buys the publication, the lack of these policies could open the path to pressure being put on staff to inaccurately portray stories or negatively target groups.

Figure 4. Overall market scores, by pillar

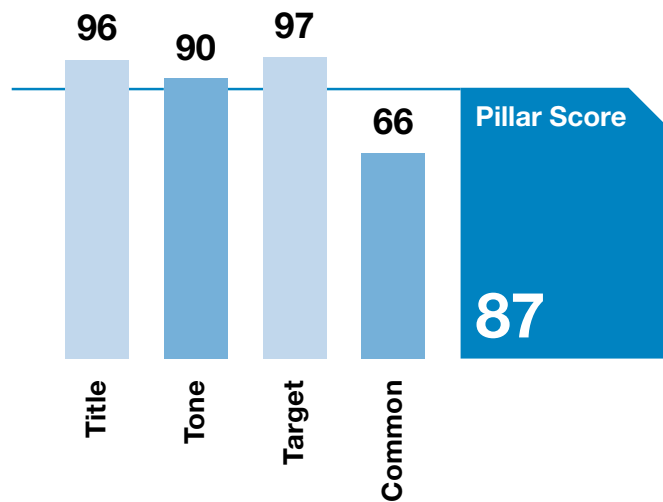


Pillar Overview

CONTENT PILLAR

This pillar focuses on the reliability of the content provided on the site. For the South African media market, we find that while news domains tend to excel on this pillar, many do show risks arising from limited common coverage of their sites' news articles.²⁴

Figure 5. 'Content' pillar indicators and scores



Our analysis for the 'Content' pillar is based on ten anonymised articles for each domain. All domains score over 75 when it comes to the use of headlines that accurately reflect the actual content of the article. Sites also scored relatively high for not negatively targeting groups or individuals, and for not using emotional or biased language in their reporting.


However, we found a wide distribution of scores when it came to the common coverage of news stories based on the articles reviewed, with several domains scoring greater than 90 and several others falling below 50. While there may be reasons why an article would cover a topic not being widely given press elsewhere, the lack of common coverage is an important disinformation flag for whether the portrayal of stories could be used as a vehicle for spreading disinformation.²⁵ This flag has a high potential for a site's disinformation risk when combined with weak operational indicators for that site.

OPERATIONS PILLAR

This pillar assesses the operational and editorial integrity of a news site. The operational indicators are the quickest wins to reduce a site's disinformation risk rating, as they represent policies that domains can directly establish and make public.²⁶ However, most sites in our market sample are found to be lacking many of these policies.

For example, very few of the domains we reviewed have published a statement of editorial independence. For the 'Operations' pillar, this specific disinformation flag has the poorest performance across the market sample and poses the greatest risk for South African domains. This problem is evident among new media (blogs, aggregators, etc.) as well as traditional media (national and regional newspapers). Our review found that only five sites out of our sample of 30 meet this standard.

Figure 6. 'Operations' pillar indicators and scores



Additionally, every domain in the sample lacks a policy on the publication of algorithmically-generated content, and almost one-third of the sites have no policies outlining rules for user-generated comments. A majority of domains also lack guidelines for when and how they issue corrections to their stories. With that said, the market does well on disclosures of ownership and funding sources, with a majority of domains scoring a full '100' on both indicators.


All 30 sites in our sample have the potential to score perfectly on all the indicators of the 'Operations' pillar if they adopt and disclose such operational policies and information. The indicators for the 'Operations' pillar are taken from the standards which have been set by journalists as part of the Journalism Trust Initiative (JTI).²⁷ As the JTI points out,²⁸ adopting these standards raises credibility in the eyes of the public, compels traditional media to reassess their practices in the digital age, and encourages new media outlets to be more transparent about their business models.

Sites performing poorly on this pillar include blogs and news aggregators as well as a number of more professional news outlets. This suggests that in order to minimise risk in the South African media market, all publishers should rethink their standards for public disclosure of the JTI's key policies.

CONTEXT PILLAR

A site's performance on this pillar is a good measure of expert perceptions of brand trust in a given media site. 'Context' pillar scores have significant room for improvement for many domains, although shifting expert perceptions can only occur over the medium- to long-term. This is partly due to the fact that perceptions can be "sticky" and take time to realign with a site's current realities. With that said, our statistical analysis indicates that expert perceptions do reflect several of the analyst findings for the content and operations indicators, so adopting the content and operational standards measured in those pillars may have the additional effect of increasing perceptions in the eyes of experts.

Figure 7. 'Context' pillar indicators and scores



The 'Context' pillar's perception-based indicators reveal significant risks of disinformation. South African domains are viewed by independent media experts as doing well at labeling opinion and news, but poorly when it comes to issuing corrections to their stories and using clickbait titles in their headlines.

This is not to say that sites are not issuing corrections. Rather the corrections may not be visible or clearly done, leading to perceptions that this practice is not prevalent.

Also, our analysis of articles found that headlines typically do reflect the body of the article – despite experts' perception that domains often publish clickbait headlines. This difference could be the result of expert perceptions having not yet caught up to a site's current practices.²⁹

Conclusion

Our assessment of the disinformation risk of news sites in South Africa finds that the country's vibrant and relatively trusted media market still presents various disinformation risks that should be addressed.

The analysis shows that, overall, many of the 30 sites tend towards a common level of disinformation risk, which places many of them within the middle of the scores. Domains typically demonstrate low risk in our framework when it comes to indicators that assess the reliability of content. Still, these domains' overall ratings are brought down by either operational shortcomings or low levels of brand trust in them.

News sites could address these shortcomings by taking actions that:

- focus on adopting journalistic and operational standards like those set by the Journalism Trust Initiative.
- ensure sites publish a statement of editorial independence, guidelines for issuing corrections, and policies for user- and algorithmically-generated content.
- improve and make more visible a site's corrections practices. It is important that such site corrections are clearly seen and understood, rather than being hidden on a web page below the fold.

- attempt to address the challenge of the “clickbait” culture and its race-to-the-bottom by presenting headlines that are clear and which accurately reflect the text of a story. This includes working with advertisers and ad tech companies over the long-term to shift the incentives of the overall online advertising business model.

Longer-term challenges remain for all news outlets that are primarily funded by advertising in a world of ever increasing amounts of content – and algorithmic amplification of that which is most engaging. This reality underpins much of the crisis of trust in news sites.

The GDI risk ratings are an attempt to address this shortfall. The need for a trustworthy, independent rating of disinformation risk is pressing. The launch of this risk-rating framework will provide crucial information to policymakers, news websites, and the ad tech industry. The end goal is to enable key decision-makers to stem the tide of money that incentivises and sustains disinformation.

In 2020, we will reassess these 30 domains and expand our index to a greater number of countries around the world. We look forward to engaging with news sites and the tech industry throughout this process.

Annex: Methodology

Pillar scoring

The ‘Content’ pillar produces a score based on a series of indicators reviewed by a dedicated country analyst across ten articles published by a domain. These ten were randomly selected from among that domain’s most shared articles within a two-week period, and then stripped of any information that could identify the publisher. The indicators included in the final risk rating for the pilot are: title representativeness, author attribution, article tone, and common coverage of the topic or story among other domains.

The ‘Operations’ pillar is scored at the domain level by the same country analyst. We selected five indicators from the Journalism Trust Initiative’s list of trustworthiness signals in order to capture the risk associated with a domain’s potential financial conflicts of interest, vulnerability to disinformation in its comments sections, and editorial standards. This is not meant to capture actual quality of journalism, as this pillar rates a domain based on its public disclosure of operations, which may differ from actual operations. The indicators included are: disclosure of true beneficial owners, transparency in funding sources, published policies for comments

sections and the publication of algorithmically-generated content, a clear process for error reporting, and a public statement affirming editorial independence.

The ‘Context’ pillar score is based on results from a survey of local media experts’ perceptions of a domain’s content and operations. Incorporating survey data in calculating the risk rating is essential because it captures a wider range of opinions, and because experts’ perceptions are based on long-term behaviour and performance – a good complement to our ‘Content’ pillar, which goes more in-depth but analyses only ten articles. The survey captures four indicators: accuracy, clear differentiation of news and opinion articles, use of clickbait headlines, and error reporting.

Domains are placed into one of five risk categories based on their final risk score. The cutoffs for the categories are determined by combining the risk ratings for domains in all countries in the index and calculating this global sample’s mean and standard deviation. Domains are placed into a category based on the number of standard deviations that separate their rating from the global mean score. The table below shows each category and its cut-offs.

Figure A.1: Risk Levels and Categories

TOTAL DOMAIN SCORE	DISINFORMATION RISK LEVEL	DISINFORMATION RISK CATEGORY
< -1 SD from mean	5	Maximum risk
≥ -1 and < -0.5 SD from mean	4	High risk
≥ -0.5 and ≤ 0.5 SD from mean	3	Medium risk
> 0.5 and ≤ 1SD from mean	2	Low risk
> 1 SD from mean	1	Minimum risk

Data collection

The South African domains were split between two analysts who were trained by GDI staff on our framework according to a codebook that provides detailed instructions for assessing each indicator. We also hired a third analyst to review articles from ten domains so that we could test for inter-rater reliability and ensure the quality of our data.

The survey was conducted by YouGov using a snowball sampling method and includes 50 media experts from academia, civil society, and industry. Each respondent was asked a series of questions about domains that they indicated they were familiar with but do not actively contribute to. Eleven domains which were reviewed by fewer than ten experts had their 'Context' pillar scores excluded from their final risk rating. In the full 2020 version of the index, we will increase the sample size to avoid small samples for lesser-known domains.

Endnotes

1 We define disinformation in terms of the verb 'to disinform': "to deliberately mislead; opposite of inform."

2 The human review elements of the framework were developed in collaboration with Alexandra Mousavizadeh (head of insights for Tortoise Media and co-founder of the GDI). The framework was advised by, vetted by, and finalised with the support of a technical advisory group (TAG), including Ben Nimmo (Graphika), Camille Francoise (Graphika), Miguel Martinez (Signal AI), Nick Newman (Reuters Institute of Journalism), Olaf Steinfadt, (Reporters without Borders), Cristina Tardaguila (the Poynter Institute's International Fact-Checking Network), Amy Mitchell (Pew Research Center), Scott Hale (Meedan and Credibility Coalition), Finn Heinrich (OSF) and Laura Zommer (Chequeado).

3 Figures are estimates by the GDI as of September 2019. See: https://disinformationindex.org/wp-content/uploads/2019/09/GDI_Ad-tech_Report_Screen_AW16.pdf.

4 See: <http://www.digitalnewsreport.org/survey/2019/south-africa-2019/>.

5 2019 Edelman Trust Barometer reveals 'My Employer' is the most trusted institution. (2019, January 20). Retrieved from www.edelman.com/news-awards/2019-edelman-trust-barometer-reveals-my-employer-most-trusted-institution.

6 See: www.edelman.com/sites/g/files/aatuss191/files/2019-03/2019_Edelman_Trust_Barometer_Global_Report.pdf. Also see: https://d25d2506sfb94s.cloudfront.net/cumulus_uploads/document/uabvypb2cf7/Globalism2019_Social_Media_General.pdf.

7 See "South Africa": https://reutersinstitute.politics.ox.ac.uk/sites/default/files/inline-files/DNR_2019_FINAL.pdf.

8 See: https://reutersinstitute.politics.ox.ac.uk/sites/default/files/inline-files/DNR_2019_FINAL.pdf.

9 See: https://reutersinstitute.politics.ox.ac.uk/sites/default/files/inline-files/DNR_2019_FINAL.pdf.

10 The findings for South Africa will be updated in 2020 based on inputs on this report, its findings and another round of country level assessment of the 'Content' and 'Operations' pillars. All media sites in the sample were originally contacted in September 2019 about their assessment. Five sites in South Africa responded with additional questions or queries. In November 2019, all the sites were given their individual performance findings from the pilot. As of 11 December, one site in South Africa has responded. As a result of this dialogue, some site's operational scores may be updated. However, this should not affect the overall market level findings.

11 The GDI looks forward to working with the "whole-of-industry" in this effort. There is a notable concern that there is a demand for such a risk assessment of sites which less trusted and independent actors may seek to fill.

12 We define clickbait as: “Sensational, outrageous or emotional headlines online that are aimed at stirring the reader’s interest to click on a story.”

13 For more on the Journalism Trust Initiative, please see: <https://jti-rsf.org/en/>.

14 A total of 50 experts were surveyed in South Africa. These were drawn from the media sector and included experienced journalists, senior-level media researchers, academics, executives from non-governmental organisations (NGOs) and think tanks, and media industry business leaders.

15 This research is based on a pilot of the GDI risk rating tool. The scores are preliminary and will be updated in 2020 when the assessment is re-run.

16 See: <http://www.digitalnewsreport.org/survey/2019/south-africa-2019/>.

17 Ibid.

18 See: www.edelman.com/sites/g/files/aatuss191/files/2019-03/2019_Edelman_Trust_Barometer_Global_Report.pdf. Also see: https://d25d2506sfb94s.cloudfront.net/cumulus_uploads/document/uabvvpb2cf7/Globalism2019_Social_Media_General.pdf.

19 <https://www.pwc.co.za/en/press-room/entertainment--media-report-2019.html>.

20 As part of the pilot, we intentionally included several regional outlets and domains that had been criticised for disinformation in the past in order to test how our framework would respond.

21 For more on our methodology, see the appendix and methodology at: <https://disinformationindex.org/research/>.

22 For example, use of ads.txt, security protocols, and site-specific email aliases.

23 In the scaled-up version of the index to be released in 2020, the ‘Structure’ pillar will produce a score that will be factored into a domain’s risk rating.

24 This is not to say that South African domains do not publish disinformation. A few high-risk domains do score poorly on this pillar. This pillar should be interpreted as a measure of risk based on a sample of articles from a domain, and understood within the broader context of the assessment (as per how a site scores on the other pillars).

25 For example, this may be due to hyper-local news coverage by a site, or to investigative reporting that covers seldom-profiled stories.

26 The ‘Operations’ pillar looks at whether relevant policies are in place. It does not assess the level of robustness of the policy based on good practice, and does not look at how the policies are being implemented. However, other indicators in the framework do capture some of the relevant practices, such as by measuring expert perceptions on how often sites correct errors or are viewed as carrying accurate content.

27 For more information on the JTI, which has adopted an ISO standard for the industry, please see: <https://jti-rsf.org/en/>.

28 <https://www.cen.eu/news/workshops/Pages/WS-2019-013.aspx>.

29 It is important to note that scores for eleven of the South African domains are not included in this pillar due to the low number of responses from the experts. Low response rates compromise the statistical robustness of the ‘Context’ pillar findings for these eleven sites. As a result, this means there are not context indicator scores for eight of the twelve lowest-scoring domains in our sample.



www.disinformationindex.org