Measuring the Social Acceptability of Infrastructure Investments

A natural language processing approach

February 2023

EDHECinfra
Research publication

Contents

Authors

Jianyong Shen is a Senior Research Engineer focusing on quantitative analysis and Machine Learning solutions in infrastructure assets' pricing/risk estimation and ESG impact. He has 15 years experience in the quantitative analysis in financial industries. He holds a Master in theoretical physics and a PhD in cosmology from Fudan University.

Tim Whittaker is a Research Director at the EDHEC Infrastructure Institute and Head of Data Collection. He holds a Master of Business (Financial Management) and a PhD in Finance from Griffith University.

1. Introduction

Infrastructure development has the power to enrich communities, providing critical public services such as electricity, transport and water. However, it also possesses the potential to be very disruptive to the communities that it serves. Such disruptions include loss of amenity, increased pollution (both noise and air) as well as impacts on the local wildlife. These effects might be considered be minor for society at large but, for the local community, they can be very significant, creating negative sentiment and diminishing public support for the infrastructure development. As a result, a failure to identify and react to deteriorating public endorsement towards infrastructure projects has the potential to lead to delays or even project cancellation.

McCarthy et al. (2020) show that, in the US, community opposition is one of the major factors that can lead to the cancellation of road infrastructure projects.

In Australia, the East West Link road in Melbourne was cancelled, despite contracts already being agreed, following a deterioration in support for the project that triggered protests Alcorn (2014). Victorian Auditor General's (2015) found that this cancellation cost Victorian taxpayers AUD$1.1 billion.

In the UK, meanwhile, the HS2 rail project has reported to parliament that protests have already added £200m to its costs (Plummer and Pickard (2022)).

For wind farms, as with other infrastructure, failure to obtain public support can lead to large costs. In 2013, the La Compagnie du Vent wind farm was ordered to remove four turbines from an already operating wind farm, due to their impact on local residents (see (Dodd, 2013)). Both in Ireland and Australia, local residents have had significant legal wins against wind farm developers either stopping a wind farm development, or requiring compensation (see (independent, 2018) and (Costa, 2020)) and Duxbury (2021) report that in Sweden, despite the requirements for greener energy, the public is turning against further developments.

As a result, monitoring public attitudes, or "social sentiment [1]", towards infrastructure projects and sectors is important. The traditional approach, involving detailed field work and surveys to identify and respond to issues, is a costly and time-consuming process. In this paper, we propose a different method, namely sentiment analysis on newspaper articles, to create an index of social sentiment around infrastructure projects. This method enables an immediate and relatively cheap measure of sentiment to be calculated, a contrast with the traditional approach of surveys and fieldwork.

In this paper, we develop a measure of social sentiment for wind farms. This is a timely approach for analysis amid the current push towards cutting emissions from electricity production, especially given that wind power is a relatively mature renewable technology. We first identify articles that contain news regarding wind farms and then filter these by content. Specifically, we are interested in articles that contains topics related to the EDHEC*infra* Environmental, Social and Governance (ESG) Risk Exposure Profile[2] for wind farms. Once the

1 - In this study "social sentiment" is employed as a proxy for the social acceptance of the infrastructure. If society does not accept an infrastructure development, then it will firstly express sentiment negatively and then is likely to move and regulate to control the development, as the different "acceptability" level shown in Busse and Siebert (2018). However, by monitoring sentiment, it is possible to determine how accepting a society is of a development.

2 - The EDHEC*infra* ESG Exposure Profiles (EPs) aim to identify the main material factors explaining the direct and indirect exposures to ESG risks of different types of infrastructure assets. The EPs follow the definition of ESG risks and impacts developed

articles are identified, their sentiment is then modelled and this data is used to create an index of wind farms in three different countries: Australia, the US and UK.

This paper shows that the results correlate well with those of contemporaneously conducted opinion surveys, indicating that our wind farm sentiment index does indeed provide a measure of social sentiment.

by EDHEC*infra* Manocha et al. (2022), and the definition of infrastructure assets introduced by the TICCS classification standard EDHEC*infra* (2022). The EPs set a minimum standard for ESG risk exposure assessment and provide asset owners, asset managers, investors and other stakeholders with a parsimonious view of the ESG profile of infrastructure asset types.

2. Literature Review

Social acceptance of renewables can be broadly divided into three categories, according to Wüstenhagen et al. (2007). These are socio-political acceptance, community acceptance and market acceptance. Socio-political acceptance relates to the policies that drive the adoption of new technologies. Community acceptance relates to the siting and building of renewable assets. Finally, market acceptance relates to renewable assets as investments and reliable suppliers.

Wüstenhagen et al. (2007) states that socio-political acceptance can be gauged by conducting opinion polls of the public's relative support. A population-wide gauge makes sense because the socio- political aspect is concerned with how society at large broadly views the renewables industry. However, for community acceptance, hyper-localised opinion measurement needs to be conducted, as this element is more closely concerned with how specific projects are developed and who is directly affected. Wüstenhagen et al. (2007) also note that this level of acceptance can be time varying; the public can be less keen on projects early in their development cycle but become more supportive as time passes and the projects begin to deliver benefits. Finally, for market acceptance, one must gauge the willingness of businesses to invest in renewable energy developments or companies, or indeed become their customers.

When measuring the first two, socio-political and community acceptance, the typical approach has been to conduct surveys of public opinion (see Hall et al. (2013)). For example, Ribeiro et al. (2014) polled 3,646 people in Portugal whilst Sütterlin and Siegrist (2017) employ the market research data of 1,211 people. Designing and conducting these surveys is a costly and time consuming process. As a result, these surveys are only able to be conducted at a point in time, with longitudinal surveys rare. One alternative to conducing opinion polls is to mine opinions of people mentioned in written text. Increasingly, with the growth of social media, there now exists a large collection of textual data that is ripe for mining for opinions about different products and services. This has allowed more recent methodological approaches to be developed to measure public opinion and hence, social acceptance from text data. These approaches employ a Natural Language Processing (NLP) technique called sentiment analysis. Liu (2012) state these data sources can be mined to rapidly determine the sentiment expressed by a significant number of people, a process termed sentiment analysis. Algaba et al. (2020) defines sentiment in this context as:

> "Sentiment is the disposition of an entity toward an entity, expressed via a certain medium."

In practical terms, sentiment analysis involves developing statistical or lexicographic models to identify the subject of a passage of text and then measure its polarity in order to determine the sentiment expressed. The lexicographic approach involves linking dictionaries or phrases with their corresponding polarity valence whilst the statistical approach involves constructing statistical models that 'learn' the polarity of the text passage.[3] The use of sentiment modelling has broad applications across multiple fields. Godbole et al. (2007) demonstrated that the large sets of text data online can be used to generalise these techniques to many different areas of academic study. These include politics (see O'Connor et al.

3 - The polarity of the passage of text is determined by whether it expresses a positive, neutral or negative view.

(2010), public perceptions on crime Prichard et al. (2015), finance Loughran and McDonald (2011) Manela and Moreira (2017), economics Shapiro et al. (2020) or even the public reputation of firms Colleoni et al. (2011).

The increasing sophistication of sentiment analysis tools has fuled fuelled its popularity. It can be a powerful tool to measure social acceptance as it is a way of measuring public opinion at high frequency. For instance, Bollen et al. (2011) employ sentiment analysis on Tweets to determine the public mood around certain holidays and events. Whilst Oliveira et al. (2017) found that there was a strong correlation between sentiment mentioned in tweets and public political opinion polls. Etter et al. (2018) demonstrate that the use of online blogs and social media aid in this measurement process. Sentiment analysis is increasingly used to examine the social perceptions of specific industries - including renewable energy - to understand where they stand in terms of public legitimacy (see Nuortimo et al. (2018), Nuortimo and Härkönen (2018), Jiang et al. (2016) and Dehler-Holland et al. (2022)).

In this paper we aim to measure the "social acceptance" of infrastructure industries by examining the public sentiment around different aspects of the releveant infrastructure assets. We use sentiment analysis tools to monitor the sentiment around factors that are related to public acceptance of infrastructure assets. We use our findings to develop tools to examine the role of sentiment in the future development of the renewable energy industry.

3. Methods

The creation of our sentiment index is a three-step process. First, the relevant articles need to be identified. Secondly, the sentiment for these articles needs to be measured. Thirdly, an index showing the sentiment of the articles over time needs to be constructed using this data. Finally, the constructed index will be validated against the information reflecting the public sentiment. Figure 1 illustrates these stages.

This section discusses how we implement Stages 1 to 3, as in Figure 1. Chapter 5 presents the validations of the constructed indices.

3.1 Identifying relevant articles

To identify articles containing the topics of interest, we must first identify the topics closely related to the infrastructure sector. The relevant topics are identified from the EDHEC*infra* ESG exposure profile[2]. This is done to ensure articles that are labelled are closely related to the EDHEC*infra* ESG exposure profile for the wind farm infrastructure sector.

We use topic modelling to identify relevant topics within our news corpus. Specifically, we employ the Gallagher et al. (2017) CorEx topic modelling methodology because it does not assume a data generating model - in contrast to traditional topic model approaches such as Latent Dirichlet Allocation (LDA). Instead, CorEx offers the ability to provide a "seed" keyword to the model and, as a result, the topic search model is effectively anchored to find articles that contain the "seeded" keywords.

The CorEx approach is implemented in three steps.4 First, the corpus is examined for common topics. In this stage we set the number of topics to be examined at 50, similar to Gallagher et al. (2017). Once these initial topics are identified, they are examined to determine whether they are related to the ESG subject as identified from the EDHEC*infra* ESG exposure profile. Any topics judged as not relevant are then excluded and topics are then re-selected to ensure 50 groups of topics are retained. Finally, the selected articles are then validated employing both a rules based as well as human annotator input.

Following the identification of the articles of interest, these raw news articles are pre-processed to normalise formatting and remove any unnecessary noise (such as special characters, meaningless words and text ie. advertisements) to create the final corpus of news articles for further analysis. The final corpus contains 37,290 wind farm ESG news articles over the period October 1996 to August 2021. The analysis includes five countries: the UK, US, Australia, New Zealand and Canada. This is to provide more diversity of coverage of ESG issues.

3.2 Measuring sentiment

As Shapiro et al. (2020) state, there are two main approaches to determining sentiment from written text, a lexicon-based or a statistical (or machine learning) approach. Both require significant input from researchers in order to develop the datasets. The lexicon approach, as described by Shapiro et al. (2020), involves curating a set of pre-defined words or phrases (a dictionary) that are ranked by their valence 1,0 and -1 for positive, neutral and negative sentiment. This list and its sentiment scores are generally domain specific and compiled by subject matter experts. Employing dictionaries to estimate the sentiment expressed in an article essentially resembles a Bag of Words (BOWs) approach, according to Shapiro et al. (2020); it effectively matches and

4 - Further details of the full implementation can be found in Appendix A

Figure 1: Major steps of sentiment index

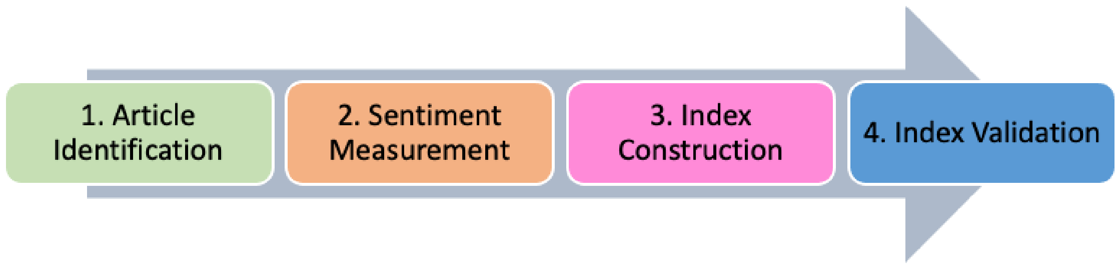

counts occurrences of words and summarises the expressed sentiment based on the valence of the words in the dictionary. For the lexicon approach, the construction of the dictionary is the most labour- intensive effort. Once it is created, it is possible to estimate the sentiment of articles by counting the occurrences and polarities of the words.

The simplicity of the lexicon approach is also its major weakness. By only including words in the pre-specified dictionary, it loses information around the context of how the words are used in a sentence. Specifically, it loses what Shapiro et al. (2020) refer to as 'degree modifiers' (words such as very, severe and slightly), which can have a significant impact on the inferred sentiment in the passage of text. To mitigate such problems, researchers have switched from employing words to employing phrases in dictionaries, but this increases the complexity of building a lexicon.

Another issue with building lexicons is that they are not stable over time as the meaning of words changes over time. Hamilton et al. (2016) and Lukes and Søgaard (2018) both demonstrate that the lexicographic features of text also change over time. This is an issue for both the lexicon and machine learning approaches (discussed more in detail later). Furthermore, lexicons cannot be generalised from one news domain to another. Loughran and McDonald (2011) show that those developed to analyse other domains are not then able measure sentiment in financial news. Instead, the authors develop their own lexicon to improve the sentiment analysis of 10-k financial reports.

The second approach to creating a dataset for sentiment analysis is to employ Natural Language Programming techniques from machine learning. These techniques develop statistical models to determine the polarity of passages of text, based on pre-annotated text provided by researchers. There are two approaches to develop the pre-annotated text datasets, by employing either naturally labelled or professionally curated text.

For naturally labelled text, the labels are automatically attached because of how the text is collected, from sources such as movie reviews, customer reviews and user-tagged messages on social media.

Naturally labelled datasets are typically very large; however, having a label attached to the text makes the collection of sentiment labelled data both easy and quick. Furthermore, the labels paired with the data are chosen by the actual author, and therefore can be assumed to accurately paired to the text data. That said, there are issues with employing such datasets for wider sentiment analysis. First, the domain coverage is generally narrow. Review datasets allow for the comparison of movies, products etc. but the specificity means that these datasets cannot be used to quantify the sentiment expressed in news articles. Second, the datasets, whilst large and easy to collect, are not as carefully curated as datasets that have researchers involved. For instance, the text and the rating may be inconsistent, with the text expressing a positive (or negative) sentiment with the label expressing a negative (positive) experience. If these inconsistencies remain in the data, then the resulting model may have little to no predictive performance for unseen text.

An alternative approach to naturally labelling is to manually annotate the collected text data. As most text exists without any sentiment labels, this is a popular way to build datasets. An individual reads each passage and labels it as positive, negative or neutral, resulting in a data set that is similar to naturally labelled data. However, the cost involved in employing a human annotator (in practice multiple annotators) means that this method of data curation is significantly more expensive. Another issue, as summarised by Paullada et al. (2021), is that annotation work is interpretive. As a result, failure to properly train and supervise the annotators can result in biases creeping into the dataset. To combat any bias, one method of creating a 'gold standard' of tagged articles is to have each piece of text annotated by multiple people. A set level of agreement between the annotators on the sentiment expressed must be met for the text to be included in the dataset.

In this paper, we use a lexicographic approach, as described in Shapiro et al. (2020), to determine the sentiment of an article. For this method we first employ the dictionary in the VADER sentiment analysis model (see (Hutto and Gilbert, 2014)). This dictionary is augmented with the sentiment dictionaries from Loughran and McDonald (2011), Hu and Liu (2004) and the Harvard General Inquirer.

The VADER model is then applied to score the sentiment of all articles in the corpus. As the VADER model outputs a numeric score from -1 to 1 they need to be converted into sentiment polarities (positive, neutral and negative). This is achieved by applying score thresholds, which are searched effectively on a labelled, or "ground truth" dataset to maximise the macro-F1 score of the VADER model. As a result, an article is viewed as positive (or negative) if its score is larger than 0.1 (or smaller than 0.07). If its score is between 0.07 and 0.1, the article is classified as neutral.

These two score thresholds (0.1 for positive versus neutral and 0.07 for neutral versus negative) are calculated on the 'ground truth' dataset. The following sequence is run to maximise the F1 scores.

- Compute every word's point-wise mutual information (PMI) against the polarities (positive, neutral and negative).
- Compute every word's sentiment score by the difference between the word's PMIs of positive and negative.
- Compute the article sentiment score by averaging the words' sentiment scores in the article.

To evaluate the VADER model, we compare its performance with that of a 'Ground Truth' dataset. We use two test statistics, the Spearman's rank-based collection and the Macro-F1 statistic. The Spearman's ranking-based correlation examines the correlation between the VADER model derived sentiment scores and 'ground truth' polarity scores. A high correlation implies a higher agreement between the two, whilst a lower sentiment indicates that the model poorly predicts the sentiment inferred in the article. The macro-F1 test statistic examines the ability of the model to correctly classify the polarity of the article. This test is a standard benchmark employed in classification problems. However, it does require the conversion of the continuous score output by the VADER model to the discrete sentiment categories (negative, neutral or positive) before identifying the score thresholds.

Once the model for sentiment measurement has been created and deemed sufficiently robust to determine the sentiment of the articles in the ground truth sample, it is then applied to all articles in the dataset. Once all articles have had their sentiment modelled, we can now create our index of sentiment for the renewable assets.

3.3 Creating the sentiment index

We follow Shapiro et al. (2020) to construct the sentiment index for the renewable energy technologies. It is assumed that the article's sentiment score at time T can be split into systematic part and idiosyncratic part as expressed by the following equation:

$$s_a(T) = f_s(T) + \sum_k (f_I(k,T,a)x_I(k,T,a) + \varepsilon_a(T)) \tag{3.1}$$

where:

- $s_a(T)$ is the a-th article sentiment score at the time T;
- $f_s(T)$ is the systematic effect at the time T;
- $f_I(k,T,a)$ and $x_I(k,T,a)$ is the k-th idiosyncratic effect and feature of the a-th article to reflect the article's own features; and,
- $\varepsilon_a(T)$ is the noise at time T.

Shapiro et al. (2020) view the systematic effect $f_s(T)$ as the sentiment index directly. The regression is done at each time step T. The idiosyncratic factor $x_I(k,T,a)$ in their work is a flag to indicate it the news article is editorial or not.

Shapiro et al. (2020) show the above method can build a meaningful sentiment index which is closely tracking the University of Michigan Consumer Sentiment Index. However, we notice the shortcomings of obtaining the systematic effects $f_s(T)$ through the independent regression. First, $f_s(T)$ is independent across time steps and shows very volatile behaviour. This is due to the how the methodology operates, not actually new information of how $f_s(T)$ should evolve. If their sentiment index reflects the reality of the consumer sentiment, it is hard to believe such sentiment is varying significantly in such a short time step. Secondly, when we have a relatively small news article corpus (e.g. renewable energies), it is not uncommon that there are no news articles published at some of time steps. In such case, we cannot obtain $f_s(T)$ by the independent regressions. Therefore, we view the articles' sentiment scores as a time series of signal and modify Shapiro et al. (2020)'s equation to the following:

$$\tilde{s}(T) = f_s(T) + \sum_k \beta_k x_k(T) + \varepsilon(T) \tag{3.2}$$

$$f_s(T) = f_s(T-1) + \eta(T) \tag{3.3}$$

where:

- $\tilde{s}(T)$ is the averaged articles' sentiment scores of at the time T;
- $f_s(T)$ is the systematic effect at the time T and has the random movement $\eta(T)$;
- β_k and $x_k(T)$ is the k-th idiosyncratic effect and averaged feature of the articles at the time T; and,
- $\varepsilon(T)$ is the observation noise at time T.
- both signal $\eta(T)$ and noise $\varepsilon(T)$ are following a random walk of normal distribution $N(0,\sigma_\eta)$ and $N(0,\sigma_\varepsilon)$ respectively.

In our case, we employ the article features $x_k(T)$ including the type of news agencies, sub-regions, keywords for sub-infrastructure types. We calibrate the systematic effect $f_s(T)$ over the country and infrastructure sector subset of the whole corpus for different months from January 2005. The Kalman Filter in the State-Space model is used in the calibration and the smoothed value of the systematic effect $f_s(T)$ are the indices for the sector in the countries as presented in Section 5.

4. Data

We use articles from Dow Jones' Factiva as the main source of textual data for this study. This data source database was chosen as it includes a broad range of news articles from around the world published by international and local news agencies in different languages. Figure 2 shows the steps to extract the relevant news articles from the raw data source.

4.1 Data Filters and Preprocessing

We extract the articles from the raw data source with the following criteria:

- Industrial sector. We use Factiva's industry tags to identify the infrastructure-related news that is relevant to our study. For wind farm news articles, we select those with the tags *iwind, i16, i16101, i163, i1, ieutil, iutil, icre, iconst, i502, i5020044* and *irenewf*. Only articles with one of these industry codes are selected.
- Countries. We use Factiva tags to select only the countries we are interested in studying, namely Australia, the US and UK. The Factiva tags for each country include its sub-regions5.
- News agencies. The same event can be reported by multiple news agencies (ie. newspapers and websites with a local, national or global perspective).
- Contents. We are interested in news articles specifically related to the ESG topics.

The extracted news articles are preprocessed to remove the noise due to the issues of formatting, special character encoding, trash URL links and advertisements. The preprocessed articles (i.e. Infrastructure ESG News Database in Figure 2) form the news dataset employed to calculate the sentiment index. However, in order to train and evaluate the sentiment models, a smaller data sample is employed. This dataset, otherwise known as the 'ground truth' data set, is compiled and annotated by human annotators.

4.2 Ground Truth Dataset

To develop a dataset to validate the performance of the sentiment analysis model, a 'ground-truth' dataset is compiled. This dataset is created employing human annotators who are asked to determine the sentiment expressed in a selection of news articles. Sentiment is expressed across five different categories, Strongly Positive, Positive, Neutral, Negative and Strongly Negative, to measure the overall polarity and strength of an article. The annotators are provided with clear guidelines Balahur and Steinberger (2009) to label the articles; these are:

- Make the judgement purely based on the text. Do not try to use your general knowledge.
- Focus on the sentiment towards the given infrastructure sector (e.g. windfarm), NOT other topics in the articles.
- Try to be objective and separate the good/bad news from your emotive reactions. In other words, the sentiment polarity should not be labelled according to how the news affects the annotator.
- Expressions (including verbs, nouns...) of attitude and intentions usually carry sentiment in the most scenarios, e.g. "try to avoid" versus "avoid".
- Ask yourselves "do you like or dislike the infrastructure more after reading the text?" to help you make the judgement.
- Label Neutral when you still cannot judge after trying all the above.

To prepare the dataset for "ground truth" labelling, duplicated articles are excluded (based on cosine similarity). Furthermore, a balanced sample is created. The articles are sampled

5 - New Zealand and Canada are included in the 'ground truth' data collection and sentiment modelling process. However, a lack of validation data means that they are excluded from the analysis of the indices.

Figure 2: The flow of data extraction

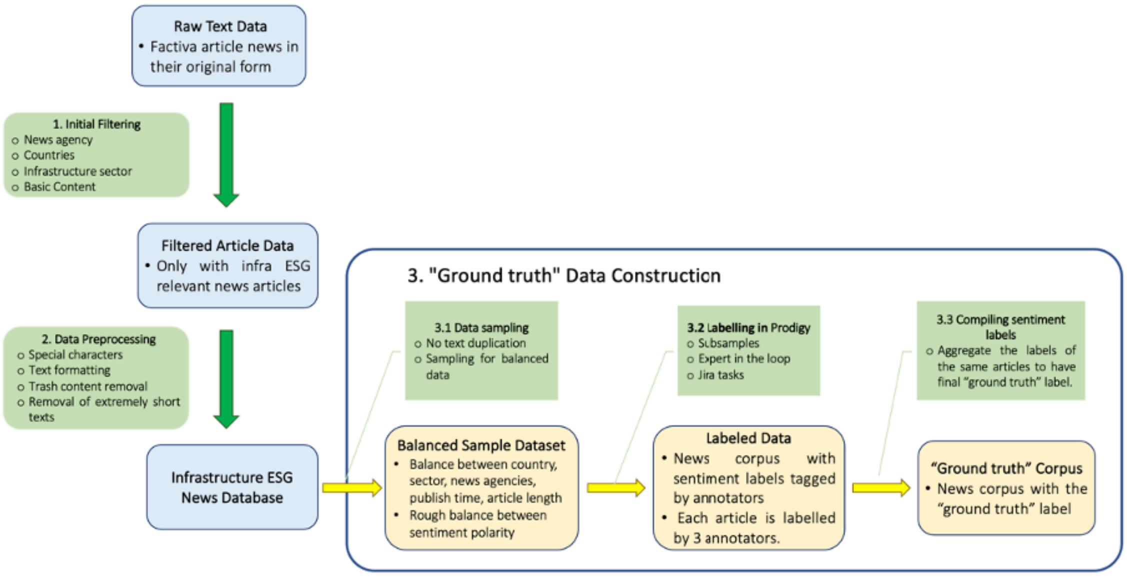

according to Sector, Country, News agency, Date of article, Length, expected sentiment and whether they contain ESG topics. The intention with this sampling approach is to obtain a sample of articles that is relatively balanced across the six elements to develop a ground truth sample of 138 articles with which to assess the performance of the sentiment model.6 Once the ground truth dataset is obtained, it is held out of the dataset to allow for proper consideration and calibration of the sentiment models developed. To manage disagreement amongst annotators articles are scored for sentiment by a minimum of three annotators, with the sentiment rankings then averaged to create the "ground-truth" measure of sentiment.

6 - In practice, having more than 100 articles to tune the VADER model proves sufficient in this study.

5. Results

5.1 Measuring Sentiment

We present the results examining the performance of the Pointwise Mutual Information model (PMI, Shapiro et al. (2020)) to discern sentiment from the 'ground truth' dataset. Table 1 presents the summary statistics of the goodness-of-fit measures for the sentiment models with their corresponding lexicographic dictionaries. The different sentiment models are compared against the 'ground truth' human labelled sentiment to determine which model most closely represents human performance. As the manually annotated dataset employs multiple annotators, and the annotators are free to label the articles differently, we observe that there is some difference (as evidenced by the correlation and F1 score not being 1). The worst-performing sentiment classifier is the dictionary from Hu and Liu (2004) with negation rules (HL + negation). As the manually annotated dataset is small, it is not possible to generalise this model across the dataset; for this reason the PMI model, the next best performing sentiment classifier, is employed in this study.

Figure 3 presents graphically the agreement between the manually annotated and the modelled article sentiment. Whilst there is a large amount of dispersion in the predicted sentiment, the average and median sentiment measures for both models correlate. With the greater polarity of the sentiment, (-2.0,+2.0) there is greater agreement, whilst in the neutral sentiment rating (0.0) there is much more significant dispersion in the recorded sentiment. This probably stems from the complexity of gauging sentiment in long articles that cover multiple topics.

From the evidence presented in Figure 3, as well as Table 1, we can conclude that the model is sufficiently capable of correctly labelling the sentiment of the news articles at an accuracy appropriate for the purposes of the research question. As a result, the sentiment index is constructed according to the methods described in Eq. 3.3 and is presented in the next section.

5.2 Sentiment Index

In this section we present the wind farm sentiment index for the three countries of interest (Australia, the US and the UK) as well as the results of validation tests conducted to confirm that the index is capable of measuring the sentiment of the general public towards wind farms. For the sentiment charts, the key point to note is that, as the sentiment index measures systematic sentiment, the trend is the most important factor, not the actual scores. As the trend tends upwards (downwards), it implies increasingly positive (negative) sentiment towards wind farms. The raw score of the index does not provide any information on the "point estimate" of sentiment.

5.2.1 Australia

Figure 4 presents the time-series of the Australian wind farm sentiment index. The sentiment changes over time and there are clear cycles in which the index implies increasingly negative sentiment. The first of these occurs at the beginning of the index in 2005, and then from 2007 the index implies increasingly positive sentiment, with the peak in 2010, before dropping again. Finally, in more recent years, the sentiment index becomes increasingly positive. At the end of the analysis period, news sentiment is positive or in the neutral zone of the model.

The changing sentiment can be associated with events that occurred in Australia. In the first fall in sentiment in 2006, a wind farm project was

Table 1: Goodness-of-fit of model sentiment scores

Model	Rank Correlation	Macro-F1
HL + negation	48%	n/a
VADER + HL	55%	56%
PMI (Based on VADER + HL)	59%	58%
Human (averaged)	73%	70%

Figure 3: Correlation between human and model rating

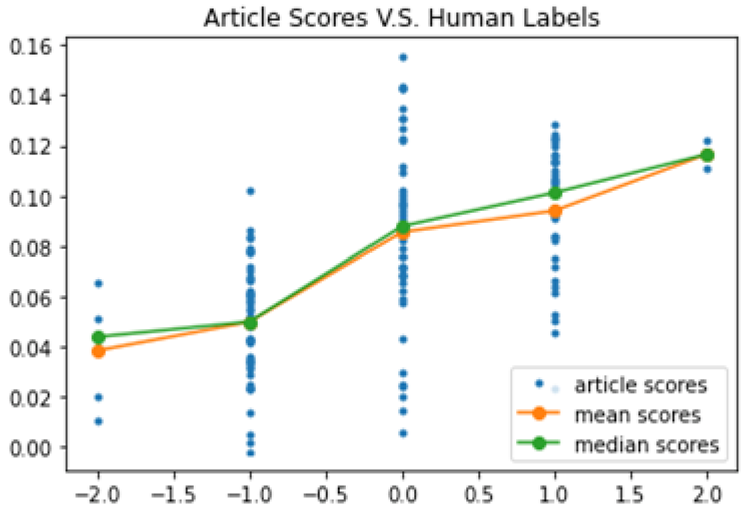

Figure 4: Australian Wind Farm Sentiment Index

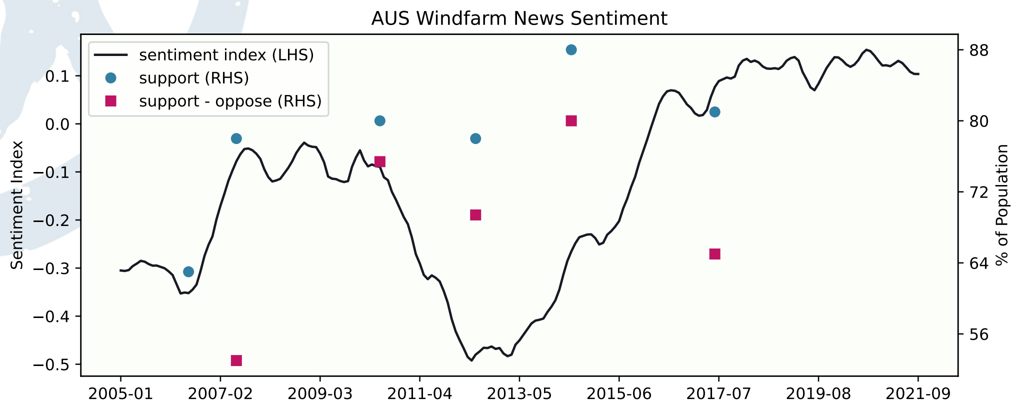

cancelled due to its threat to native bird populations. This is reflected in the news coverage and article sentiment scores in 2006; some 70% of articles were found to exhibit negative sentiment. The next cyclically shift downwards in sentiment (from 2010 onwards) corresponds with changes in the political discourse with a lack of support for measures to mitigate climate change. In 2012, 61% of articles were found to exhibit negative sentiment. Finally, the last long-term increase in public sentiment, from the period 2015 to 2021 corresponds with increasing roll out of wind capacity.

5.2.2 The US

We look next at the index for the US, Figure 5. Like the Australian index it shows long-term cycles in news sentiment, but it also exhibits a lot more volatility. Sentiment is neutral at the start of the index in 2005 but becomes more positive, reaching a local peak in 2009. Much like Australia, it then declines at this point. After 2011 there is significant volatility in the index, with sharp rises and falls in sentiment.

Figure 5: United States of America Wind Farm Sentiment Index

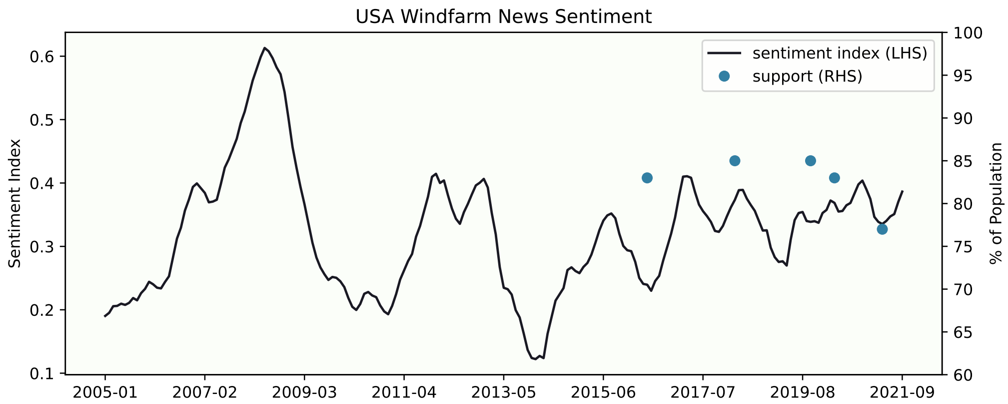

5.2.3 The UK

Finally, results for the construction of the Wind Farm Sentiment Index for the UK are presented in Figure 6. Like the Australian and US results, there is a general increase in positive news sentiment for wind farms over the course of the period of analysis. As with the Australian and US results, there is a period of decreasing support until 2011, after which there is sustained positive sentiment in the index.

5.3 Index Validation

The index of news sentiment is informative but what is reported in the news does not necessarily correlate with wider public opinion. Newspapers can have editorial biases which dictate the type and tone of the coverage of issues, especially with topics such as wind farms. As a result, it is necessary to perform analysis examining the ability of the sentiment indices to represent broader public opinion.

For the purposes of this paper, opinion polling was obtained examining the public's views of wind farms as well as associated technologies and industries from the countries of interest. For Australia, a time series of polling was obtained from Lothian (2020). For the US, several surveys were used to obtain a measure for the public acceptance of wind farms; these were Rosentiel (2010) and Tyson et al. (2021). Finally, for the UK we employed the BEIS Public Attitudes Tracker, which includes surveys of the general public's support for on-shore as well as off-shore wind generation. It should be noted that not all the polling cited in this study is conducted in a consistent manner, with differences of methodologies that may impact on the comparability of the results. Furthermore, the time series of polling only really exists for the UK survey, where it has been conducted consistently since 2012; the Australian and US data is not tracked consistently so can only be considered 'point-in-time' analysis, not full tracking of public sentiment.

5.3.1 Australia

The polls cited in Lothian (2020) were intended to discover the public reaction to the landscape with both and without wind farms. This survey covers only one of the dimensions of public acceptance as defined in the literature review (community acceptance) and even then is a very narrow focus. The results of the surveys are plotted in Figure 4 alongside the sentiment index, with the dots representing the proportion of support for the index, whilst the square market representing the proportion of the survey respondents against the impact of wind farms on the scenery. Whilst the sentiment index covers more topics than just scenery impact, it is possible to see that, at the points that the survey is conducted, there is an element of correlation between the respondents that support wind farms and the sentiment index.

Figure 6: Great Britain Wind Farm Sentiment Index

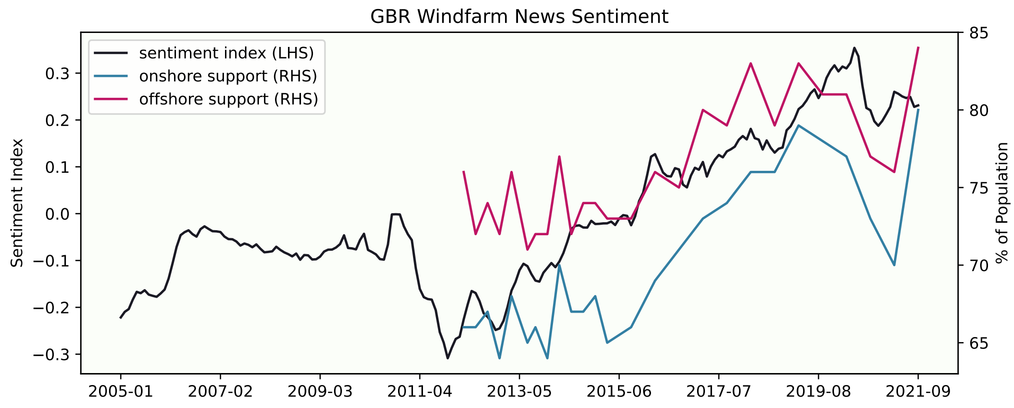

5.3.2 The US

The measure of broader public opinion is obtained from Pew Research Center surveys. These examine the level of support for various energy sources including oil & gas, coal, solar, wind power and nuclear energy. By allowing comparisons between energy sources as well as identifying political leanings of respondents, this approach differs significantly from the survey conducted in Australia. Furthermore, their surveys do not cover the full period of analysis, nor can be considered tracking surveys of public opinion for wind power. However, despite these differences in methodology, as with the Australian study, Figure 5 shows a strong correlation between support for wind farms and the news sentiment index. That said, the opinion polling appears to be significantly less volatile than the news sentiment index. This could be due to the frequency with which the opinion polls are conducted compared with the calculation frequency of the news index. It does however indicate that the sentiment index does provide a methodology to track public opinion in the US for wind farms.

5.3.3 The UK

Finally, the UK wind farm sentiment index is compared with polling produced by the BEIS Public Attitudes Tracker. This polling provides significant information around public perceptions of multiple topics, including initiatives to reduce carbon emissions and renewable energy sources including solar and wind. Unlike the other surveys, the BEIS Public Attitudes Tracker also asks respondents about their support for on-shore as well as off-shore wind farms. The granularity of the questions, as well as the regularity with which the survey is conducted, make it ideal for validating the UK news sentiment index. As with the Australian and US indices, the UK sentiment index exhibits a close relationship with the polled tracking of attitudes to wind farms, Figure 6. The performance of the news sentiment index closely tracks the results of the off-shore wind farm wind farm acceptance polling, with the on-shore polling exhibiting a similar trend to both but beginning at a lower level of acceptance. From the results presented in Figure 6 its possible to conclude that the UK news sentiment index closely tracks public sentiment towards wind farms. This provides important evidence supporting the methodology employed in this paper.

In summary, the estimated indices are highly time-varying, but increasingly wind farms are viewed positively in the three countries analysed. Furthermore, the news sentiment indices estimated correlate with actual public opinion as measured by opinion polls. This provides evidence that this methodology is appropriate for measuring the public opinion towards wind farms.

6. Conclusions

This paper examines the potential for measuring evolving public sentiment towards wind farms from news articles. Such articles are a valuable source of information of public opinion on multiple topics. By identifying and measuring the sentiment of articles containing Environmental, Social or Governance (ESG) topics, our hypothesis was that that an index measuring the sentiment of public opinion of wind farms could be created.

By applying a sentiment index construction methodology first introduced in economics by Shapiro et al. (2020), we were able measure public opinion sentiment at large in Australia, the US and UK. The index produced showed that sentiment during the analysis period (2005-2021) was time- varying, with periods of negative sentiment that turned more positive in recent years.

To validate the results, the index was compared with polling data, also from Australia, the US and UK. The polling results supported our intention that the index did measure the public sentiment towards wind farms. This conclusion however, should be taken with caution for two main issues. First, whilst the sentiment index is measured every time a new news article is identified, the polling occurs at discrete intervals, so the sentiment index is significantly more volatile than measured public opinion. Secondly, the polling identified in the three countries of analysis is not conducted in the same manner, nor asks the same questions. As a result, conclusions generalising findings across the three countries is not possible.

With those caveats on the validation, it is encouraging that the results of the BEIS Public Attitudes Tracker, the opinion poll with the longest and most consistent questions on public attitudes to wind farms, does correlate well with the sentiment index constructed for the UK. This does provide some comfort to our view that the index does indeed measure public sentiment in a reliable manner.

The results in this study provide areas for future research. Firstly, the sentiment index relies on the assumption that ESG topics do correlate with public sentiment. This assumption was not tested in this research and is a key area for future research to ensure the results are able to be interpreted in the manner anticipated. Secondly, whilst the results are encouraging for examining wind farms, it is possible to broaden this analysis to other infrastructure sectors. Third, the methodology employed in this paper does enable identification of which factors/events have had a large positive or negative impact of public opinion. As a result, it has the potential to provide a cost-effective methodology to identify public opinion and assist in its management. Finally, this methodology could be adapted to other sources of textual data such as social media (Twitter, Facebook etc.). This has the potential to improve the instantaneous measurement of opinion whilst also allowing for more localised measurement of public sentiment. This would be of great advantage to the owners and operators of infrastructure assets in identifying and reacting to changes in public sentiment in the future.

7. Appendix A

7.1 Extracting news articles about infrastructure

As the index is designed to measure sentiment regarding infrastructure, it is necessary to fierst extract those news articles that focus on this topic . Dow Jone's Factiva is used as the source in this study because it provides a global selection with a relatively wide time span. Factiva also provides the industrial tags for each article which identify the industries mentioned in it according to Factiva's own industry taxonomy. We manually searched the infrastructure related tags by mapping the Factiva taxonomy and EDHEC*infra* Taxonomy EDHEC*infra* (2022). Then the news articles mentioning the desired type of infrastructures were extracted based on the mapping. finally, we filtered the extracted articles by a set of keywords reflecting the infrastructure asset type to increase the relevance.

7.2 Identifying ESG relevant articles

We used the the CorEx topic modelling approach, introduced by Gallagher et al. (2017), to identify the articles that contain the topics of interest, as itr offers the ability to anchor keywords when conducting a topic search. This enables the identification of specific topics of interest from keywords raised from the EDHEC*infra* ESG exposure profile[2]. (The purpose of these Exposure Profiles is explained in Manocha and Blanc-Brude (2021)). This is done to ensure the articles that are analysed are closely related to the EDHEC*infra* ESG exposure profile for the different infrastructure sectors. As the CorEx approach requires a pre-defined dictionary, we pick the most frequently used words (top 1200[1]) plus those expressed in the Exposure Profile[2] about the ESG aspects.

The article selection process is conducted in three steps. First, the corpus is examined for hidden common topics, which are represented by the lists of words to show potential common topics mentioned. Next, having summarised the hidden topics, we pair them down to ensure they are closely related to ESG topics included in the EDHEC*infra* ESG exposure profile. Finally, the selected articles are then validated, employing both rules based as well as human annotator input. The process is described in further detail below.

7.2.1 Exploring the hidden common topics

It is necessary to browse the hidden topics in the corpus before any further action. A relatively large number of hidden topics was set (N=50) to understand what topics the corpus contains. From our analysis, we find that that the number hidden topics should be set to less than 50 as a number of topics appeared with empty keywords. During this step, it was observed that some topics identified were closely related to the topics described in the EDHEC*infra* ESG exposure profile for the different infrastructure sectors. These topics were then included in the next step.

7.2.2 Refining the hidden topics

The next step is to refine the hidden topics identified in the first stage. Here the keywords identified in the first stage were inspected to pick those ($\sim$ 4 keywords) mostly representing the ESG aspects based on the infrastructure sector's Exposure Profile[2]. These selected keywords are used as anchoring words to represent the ESG related topics in another topic modelling exercise to refine the results and identify relevant articles. In addition to the identified keywords,

1 - These words have about a 97.5% chance of being mentioned in the news articles in our corpus. Furthermore, the words are first stemmed and the dictionary is the list of these stems. This is to reduce the sparsity of the words due to the plurality, prefix or suffix

keywords representing the missing ESG aspects (according to the Exposure Profile) are added as the anchoring words to see if these aspects are mentioned in the corpus but with low frequencies. Finally, we keep the total number of the hidden topics at 50 to catch other non-ESG topics. This is for two reasons. Firstly, it may identify other ESG topics not represented by the anchoring words and, secondly, these non-ESG topics provide illustrative examples for reviewers of non-ESG articles.

The model is then run, with the refined anchoring words and the anchoring strength being slowly increased to examine the stability of the hidden topics. If the topic turns out to be non-ESG, then the corresponding anchoring words are removed. However, if the topic is ESG related even when it is not included in the current anchoring words list, the article's keywords are added as new anchoring words. This process is run for several iterations until the topics are stable and a refined set of anchoring words related to the ESG topics of interest are obtained.

7.2.3 Validating the output

The final step in the CorEx approach is to validate the articles selected in the second step, employing both rules and human annotators. It was found that articles selected in the second step did contain mentions of ESG topics. However, some contained only a few words about an ESG topic, which cannot therefore be considered their main topic of discussion. To obtain articles with more of a focus on the ESG topics, they are broken into sentences and the trained model from the second step is employed to determine the number of sentences in each one which contain ESG topics. An article is considered to be focussed on ESG only if more than 20% of its sentences contain ESG topics. Finally, human annotators are employed to read and annotate whether the article is concerned with the topics within the EDHEC*infra* ESG exposure profile[2].

Following the identification of the articles of interest, the raw articles are pre-processed to normalise their format and remove any unnecessary noise (such as special characters, meaningless words and text) in order to create a usable news article database for further analysis. The resulting database contained 37,290 wind farm ESG news articles from the period October 1996 to August 2021. We analysed five countries (the UK, US, Australia, New Zealand and Canada) to provide more diversity of coverage of ESG issues.

- Alcorn, G. (2014, Oct). East west link: The case for and against melbourne's $6.8bn road.
- Algaba, A., D. Ardia, K. Bluteau, S. Borms, and K. Boudt (2020). Econometrics meets sentiment: An overview of methodology and applications. *Journal of Economic Surveys 34*(3), 512–547.
- Balahur, A. and R. Steinberger (2009). Rethinking sentiment analysis in the news: from theory to practice and back. *Proceeding of WOMSA 9*, 1–12.
- Bollen, J., H. Mao, and A. Pepe (2011). Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. In *Proceedings of the international AAAI conference on web and social media*, Volume 5, pp. 450–453.
- Busse, M. and R. Siebert (2018). Acceptance studies in the field of land use—a critical and systematic review to advance the conceptualization of acceptance and acceptability. *Land Use Policy 76*, 235–245.
- Colleoni, E., A. Arvidsson, L. K. Hansen, and A. Marchesini (2011). Measuring corporate reputation using sentiment analysis. In *Proceedings of the 15th international conference on corporate reputation: navigating the reputation economy*.
- Costa, J. (2020, Aug). Residents win historic legal battle against contentious wind farm.
- Dehler-Holland, J., M. Okoh, and D. Keles (2022). Assessing technology legitimacy with topic models and sentiment analysis–the case of wind power in germany. *Technological Forecasting and Social Change 175*, 121354.
- Dodd, J. (2013, Jun). La compagnie du vent ordered to remove turbines by french court.
- Duxbury, C. (2021, Feb). Sweden's fading enthusiasm for onshore wind farms.
- EDHEC*infra* (2022). The infrastructure company classification standard (ticcs). *EDHECinfra Research Publication*.
- Etter, M., E. Colleoni, L. Illia, K. Meggiorin, and A. D'Eugenio (2018). Measuring organizational legitimacy in social media: Assessing citizens' judgments with sentiment analysis. *Business & Society 57*(1), 60–97.
- Gallagher, R. J., K. Reing, D. Kale, and G. Ver Steeg (2017). Anchored correlation explanation: Topic modeling with minimal domain knowledge. *Transactions of the Association for Computational Linguistics 5*, 529–542.
- Godbole, N., M. Srinivasaiah, and S. Skiena (2007). Large-scale sentiment analysis for news and blogs. *Icwsm 7*(21), 219–222.
- Hall, N., P. Ashworth, and P. Devine-Wright (2013). Societal acceptance of wind farms: Analysis of four common themes across australian case studies. *Energy Policy 58*, 200–208.

- Hamilton, W. L., K. Clark, J. Leskovec, and D. Jurafsky (2016). Inducing domain-specific sentiment lexicons from unlabeled corpora. In *Proceedings of the conference on empirical methods in natural language processing. conference on empirical methods in natural language processing*, Volume 2016, pp. 595. NIH Public Access.

- Hu, M. and B. Liu (2004). Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 168–177.

- Hutto, C. and E. Gilbert (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the international AAAI conference on web and social media*, Volume 8, pp. 216–225.

- independent (2018, Mar). Locals win supreme court battle against wind farm.

- Jiang, H., M. Qiang, and P. Lin (2016). Assessment of online public opinions on large infrastructure projects: A case study of the three gorges project in china. *Environmental Impact Assessment Review 61*, 38–51.

- Liu, B. (2012). Sentiment analysis and opinion mining. *Synthesis lectures on human language technologies 5*(1), 1–167.

- Lothian, A. (2020). A survey of the visual impact and community acceptance of wind farms in australia. *Australian Planner 56*(3), 217–227.

- Loughran, T. and B. McDonald (2011). When is a liability not a liability? textual analysis, dictionaries, and 10-ks. *The Journal of finance 66*(1), 35–65.

- Lukes, J. and A. Søgaard (2018). Sentiment analysis under temporal shift. In *Proceedings of the 9th workshop on computational approaches to subjectivity, sentiment and social media analysis*, pp. 65–71.

- Manela, A. and A. Moreira (2017). News implied volatility and disaster concerns. *Journal of Financial Economics 123*(1), 137–162.

- Manocha, N. and F. Blanc-Brude (2021). Towards a Scientific Approach to ESG for Infrastructure Investors. Technical report, EDHEC Infrastructure Insititute.

- Manocha, N., D. Marcelo, and F. Blanc-Brude (2022, March). Esg taxonomy of imapcts and risks for infrastructure companies, version 2.0. *EDHECinfra Research Publication*.

- McCarthy, L. N., L. Bolaños, J. Y. Kweun, and J. Gifford (2020). Understanding project cancellation risks in us p3 surface transportation infrastructure. *Transport Policy 98*, 197–207.

- Nuortimo, K. and J. Härkönen (2018). Opinion mining approach to study media-image of energy production. implications to public acceptance and market deployment. *Renewable and Sustainable Energy Reviews 96*, 210–217.

- Nuortimo, K., J. Härkönen, and E. Karvonen (2018). Exploring the global media image of solar power. *Renewable and Sustainable Energy Reviews 81*, 2806–2811.

References

- O'Connor, B., R. Balasubramanyan, B. R. Routledge, and N. A. Smith (2010). From tweets to polls: Linking text sentiment to public opinion time series. In *Fourth international AAAI conference on weblogs and social media*.
- Oliveira, D. J. S., P. H. d. S. Bermejo, and P. A. dos Santos (2017). Can social media reveal the preferences of voters? a comparison between sentiment analysis and traditional opinion polls. *Journal of Information Technology & Politics 14*(1), 34–45.
- Paullada, A., I. D. Raji, E. M. Bender, E. Denton, and A. Hanna (2021). Data and its (dis) contents: A survey of dataset development and use in machine learning research. *Patterns 2*(11), 100336.
- Plummer, G. and J. Pickard (2022, May). Hs2 protests set to cost high-speed uk rail line up to £200mn.
- Prichard, J., P. Watters, T. Krone, C. Spiranovic, and H. Cockburn (2015). Social media sentiment analysis: A new empirical tool for assessing public opinion on crime? *Current Issues in Criminal Justice 27*(2), 217–236.
- Ribeiro, F., P. Ferreira, M. Araújo, and A. C. Braga (2014). Public opinion on renewable energy technologies in portugal. *Energy 69*, 39–50.
- Rosentiel, T. (2010). Support for alternative energy and offshore drilling.
- Shapiro, A. H., M. Sudhof, and D. J. Wilson (2020). Measuring news sentiment. *Journal of econometrics*.
- Sütterlin, B. and M. Siegrist (2017). Public acceptance of renewable energy technologies from an abstract versus concrete perspective and the positive imagery of solar power. *Energy Policy 106*, 356–366.
- Tyson, A., B. Kennedy, and C. Funk (2021). Gen z, millennials stand out for climate change activism, social media engagement with issue.
- Victorian Auditor General's (2015, Dec). East west link project. Technical report.
- Wüstenhagen, R., M. Wolsink, and M. J. Bürer (2007). Social acceptance of renewable energy innovation: An introduction to the concept. *Energy policy 35*(5), 2683–2691.

Recent Publications (2019-2021)

- Blanc-Brude, F. & A. Gupta. "Robust Benchmarks for investors in private infrastructure funds" (November 2021)
- Blanc-Brude, F., A. Gupta, L. Lu & A. Wee "The Volatility of Unlisted Infrastructure Investments" (May 2021)
- Blanc-Brude, F., A. Gupta, "The Fair Value of Investments in Unlisted Infrastructure equity" (April 2021)
- N. Manocha, N. & F. Blanc-Brude "Towards a Scientific Approach to ESG for Infrastructure Investors" (March 2021)
- Amenc, N., F. Blanc-Brude & A. Gupta. "Strategic Asset Allocation with Unlisted Infrastructure" (February 2021)
- Blanc-Brude, F., L. Lum, R. Tan & T. Whittaker, "The choice of performance test benchmark of Superannuation funds' investments in infrastructure" (December 2020)
- Amenc, N. & F. Blanc-Brude. "The Cost of Capital of Motorway Concessions in France - A Modern Approach to Toll Regulation" (September 2020)
- F. Blanc-Brude & A. Gupta. "Unlisted Infrastructure Performance Contribution, Attribution & Benchmarking" (July 2020)
- Whittaker, T. & R. Tan. "Anatomy of a Cash Cow: An In-depth Look at the Financial Characteristics of Infrastructure Companies." (July 2020)
- Amenc, N., F. Blanc-Brude, A. Gupta, L. Lum. "Investors Should Abandon Absolute Returns Benchmarks - Lessons from the Covid-19 Lockdowns" (June 2020)
- Amenc, N., F. Blanc-Brude, A. Gupta, J-Y. Lim. "2019 Global Infrastructure Investor Survey - Benchmarking Trends and Best Practices" (April 2019)
- Whittaker, T., S. Garcia. "ESG Reporting and Financial Performance: The case of infrastructure." (March 2019)
- Blanc-Brude, F, J-L. Yim. "The Pricing of Private Infrastructure Debt - A dynamic Approach" (February 2019)
- Blanc-Brude, F., C. Tran. "Which Factors Explain Unlisted Infrastructure Asset Prices?" (January 2019)
- S. Garcia, F. Blanc-Brude, T. Whittaker. "Tome La Siguiente Salida (Take the Next Exit) - A Case Study of Road Investments Gone Wrong, Spain, 1998-2018" (March 2018)

These publications are available on our website: edhecinfra.com

For more information, please contact:
Nataliia Gaidarenco
e-mail: nataliia.gaidarenco@edhec.edu

EDHEC Infrastructure Institute
EDHEC Asia-Pacific
One George Street - #15-02
Singapore 049145
Tel.: +65 6653 8575

EDHEC Europe
10 Fleet Place
London EC4M 7RB
Tel.: +44 20 7332 5601

edhec.infrastructure.institute