

Optimising transparency, reliability and replicability: annotation principles and inter-coder agreement in the quantification of evaluative expressions

Matteo Fuoli¹ and Charlotte Hommerberg²

Abstract

Manual corpus annotation facilitates exhaustive and detailed corpus-based analyses of evaluation that would not be possible with purely automatic techniques. However, manual annotation is a complex and subjective process. Most studies adopting this approach have paid insufficient attention to the methodological challenges involved in manually annotating evaluation—especially concerning transparency, reliability and replicability. This article illustrates a procedure for annotating evaluative expressions in text that facilitates more transparent, reliable and replicable analyses. The method is demonstrated through a case study analysis of APPRAISAL (Martin and White, 2005) in a small-size specialised corpus of CEO letters published by the British energy company, BP, and four competitors before and after the Deepwater Horizon oil spill of 2010. Drawing on Fuoli and Paradis's (2014) model of trust–repair discourse, we examine how ATTITUDE and ENGAGEMENT resources are strategically deployed by BP's CEO in the attempt to repair stakeholders' trust after the accident.

Keywords: APPRAISAL theory, BP Deepwater Horizon oil spill, evaluation, inter-coder agreement, manual corpus annotation, reliability, replicability, transparency, trust repair.

1. Introduction

Quantifying evaluation using traditional corpus techniques presents several methodological challenges. Some of the main complexities arise from the fact that evaluation may be realised by an open-ended set of forms, may

¹ Centre for Languages and Literature, Lund University, Box 201, 221 00 Lund, Sweden.

² Department of Languages, Linnaeus University, 35195 Växjö, Sweden.

Correspondence to: Matteo Fuoli, *e-mail:* matteo.fuoli@englund.lu.se

span several words and may be expressed both explicitly and implicitly. In addition, it is a highly context-dependent phenomenon. Certain expressions may carry an evaluative meaning in some contexts, but not in others (for an overview, see Hunston, 2011: 12–19).

Traditionally, corpus-based studies have dealt with these challenges by focussing on a restricted range of language forms that tend to have a stable and predictable evaluative meaning (e.g., Biber, 2006; Biber and Finegan, 1989; Camiciottoli, 2013; Conrad and Biber, 2000; Hyland, 1998, 2005; and Kaltenbacher, 2006). Within this approach, pre-set lists of attitudinal markers are used to quantify evaluation using automatic techniques. While this method can be useful when dealing with large-scale corpora and when the goal is to identify broad quantitative patterns, it is limited both in terms of coverage (i.e., the range of phenomena that can be accounted for and accurately quantified) and in terms of the level of detail it allows.

An alternative approach that has been used with increasing frequency is manual corpus annotation (e.g., Bednarek, 2006, 2008; Carretero and Taboada, 2014; Fuoli, 2012; Hommerberg and Don, 2015; Lipovsky, 2008, 2013; Mackay and Parkinson, 2009; O'Donnell, 2014; Pounds, 2010, 2011; Ryshina-Pankova, 2014; Santamaría-García, 2014; and Taboada and Carretero, 2012 [2010]). Rather than using pre-defined lists of evaluative expressions, these are manually identified and coded in text, and quantitative data are derived from the annotations. Manual annotation allows exhaustive and detailed corpus-based analyses of evaluative language that would not be possible with purely automatic techniques. However, manually annotating evaluation is a complex and subjective task. This may hinder the transparency, reliability and replicability of analyses. With few exceptions (Fuoli, 2012; Hommerberg and Don, 2015; Read and Carroll, 2012; Ryshina-Pankova, 2014; and Taboada and Carretero, 2012 [2010]), these issues have been largely neglected in the literature. Explicit annotation principles are rarely formulated and, more often than not, no reliability test is performed, resulting in opaque analyses and procedures that are not replicable. Unless these issues are addressed, manual annotation cannot match the robustness of traditional corpus techniques and, as a result, the advantages that it offers are at least partially negated.

This paper illustrates a procedure for the manual annotation of evaluative language expressions that is designed to optimise transparency, reliability and replicability. Two strategies are implemented to achieve this goal. First, the annotation guidelines are explicitly formulated and formalised into an annotation manual that is made publicly available. Second, the reliability and replicability of the annotation procedure are validated by means of an inter-coder agreement test. We demonstrate this method with a case study analysis of APPRAISAL (Martin and White, 2005) in a small-size specialised corpus of CEO letters published by the British energy company, BP, and four competitors before and after the Deepwater Horizon oil spill of 2010. Drawing on Fuoli and Paradis's (2014) model of trust–repair discourse, the analysis examines how ATTITUDE and ENGAGEMENT resources are

strategically deployed by BP's CEO in the attempt to repair stakeholders' trust after the accident.

The primary goal of this paper is methodological. It aims to (i) emphasise the importance of transparency, reliability and replicability when subjective interpretations occupy a central role in the analysis, and (ii) offer a concrete illustration of how these aspects can be optimised. The study also contributes to the growing body of literature that focusses on BP's communicative response to the Deepwater Horizon oil spill (e.g., Breeze, 2012; Fuoli and Paradis, 2014; Harlow *et al.*, 2011; Muralidharan *et al.*, 2011; O'Connor, 2011; Schultz *et al.*, 2012; and Wickman, 2013), by investigating the company's post-crisis discourse from the point of view of trust repair.

The paper is organised as follows. Sections 2 and 3 provide the background for the study, outlining Fuoli and Paradis's (2014) model of trust–repair discourse and APPRAISAL theory. Section 4 describes the corpus. Section 5 presents the method, including a description and the results of the inter-coder agreement test. Section 6 presents the analysis. Section 7 concludes by assessing the method and offering suggestions for future work.

2. Trust–repair discourse

Trust is a valuable asset for a company. A high level of trust may, among other things, facilitate business transactions, increase employee commitment and customer satisfaction, and help to improve relations between the organisation and its stakeholders, thereby lowering the risk of unco-operative behaviour and conflict (Barney and Hansen, 1994; García-Marzá, 2005; Ingenhoff and Sommer, 2010; and Pirson and Malhotra, 2011). More generally, trust is key to ensuring a company's social legitimacy and, ultimately, its long-term viability (Poppo and Schepker, 2010: 124).

Trust is a dynamic construct that changes and develops over time (Cook, 2001; Linell and Marková, 2013; and Marková and Gillespie, 2008). Episodes of wrongdoing or negligence, controversies and scandals negatively affect stakeholders' trust in a company. A recent example of this is BP's Deepwater Horizon oil spill. The Deepwater Horizon oil spill is regarded as the largest accidental marine oil spill in history (Robertson and Krauss, 2010). It was triggered by the explosion and sinking of an offshore oil rig controlled by the British energy company, BP, in the Gulf of Mexico on 20 April 2010. The explosion resulted in eleven fatalities and the oil spill that followed has inflicted major damage on both the environment and the economy of the Gulf. The accident has severely impacted BP's image and reputation. In particular, the controversies surrounding BP's response to the spill undermined its trustworthiness, plunging the company into a crisis of public trust.

During a corporate crisis, discourse acquires a central role in controlling and minimising the damage to a company's image (Benoit, 1997).

Fuoli and Paradis (2014) present a theoretical framework for investigating and explaining trust–repair discourse. The authors propose that, when trust is at stake, the trust-breaker will shape their discourse to try to influence the trust-giver’s impressions of their ‘trustworthiness’. Trustworthiness has three main components: ability, integrity and benevolence (Mayer *et al.*, 1995). Ability concerns how competent and skilful a person (or company) is, integrity relates to how honest and sincere they are, and benevolence refers to the extent to which a person is believed to care about the trust-giver, beyond self-interested concerns. When a trust-breaking event occurs, the trust-breaker’s trustworthiness will be at stake along one or more of these dimensions. Their trust–repair discourse will be geared towards reshaping the trust-giver’s impressions accordingly.

Fuoli and Paradis (2014) identify two fundamental strategies that the trust-breaker may pursue to achieve this communicative goal. On the one hand, the trust-breaker may foreground their goodwill, sympathy and positive qualities—a strategy which the authors name ‘emphasise the positive’. On the other hand, the trust-breaker may seek to dialogically engage with and act upon the discourses that generate distrust. The authors name this strategy ‘neutralise the negative’. Both strategies may be used simultaneously and may interact in a single instance of trust–repair discourse.

Fuoli and Paradis (2014) propose that the emphasise-the-positive and neutralise-the-negative strategies are realised in text through attitudinal and dialogic engagement resources, respectively. Affective language may be used, for example, to communicate empathy and solidarity, promoting a positive impression of the trust-breaker’s benevolence, as in Example 1.

- (1) We are deeply sorry for the grief felt by their families and friends.
We know nothing can restore the loss of those men.

Evaluative expressions may be strategically deployed, for instance, to emphasise the trust-breaker’s ability, as in Example 2.

- (2) We recognize there is a great deal more to do, but I can report that BP finished its year of consolidation in robust shape.

Dialogic engagement concerns the linguistic devices by which speakers mark their stance towards other opinions and includes, for example, negation/denial, adversative discourse markers and epistemic modals (Martin and White, 2005; and White, 2003). These resources may be used to respond to, and to seek to neutralise, the discourses that represent a source of distrust, as the following example from the corpus shows:

- (3) Our fundamental purpose is to create value for shareholders, but we also see ourselves as part of society, not apart from it.

In Example 3, through the use of the contrastive and negation markers *but* and *not*, BP's CEO engages with and strongly rejects the view that the company is solely focussed on protecting shareholders' interests and is insensitive to society's legitimate concerns after the spill. By doing this, the CEO seeks to repair the company's integrity.

In the analysis presented in Section 6, we use Fuoli and Paradis's (2014) model to investigate the trust–repair discourse strategies employed by BP's CEO in his letters to shareholders after the Deepwater Horizon oil spill. The first step in the analysis consists of identifying and analysing the use of attitudinal and dialogic engagement expressions, which, as discussed above, are predicted by the model to play a crucial communicative role in trust–repair discourse. To accomplish this task, we employ APPRAISAL theory (Martin and White, 2005), which provides an integrated and coherent framework for the analysis of these discursive phenomena. The next section outlines the theory.

3. APPRAISAL theory

APPRAISAL theory offers a framework for analysing how speakers negotiate and manage inter-personal roles and relations in discourse. The type of linguistic phenomena accounted for by the theory have also been investigated under different labels from alternative and partly overlapping perspectives—for example, 'evaluation' (Hunston, 2011; Hunston and Thompson, 2000; and Thompson and Alba-Juez, 2014) and 'stance' (Biber, 2006; Biber and Finegan, 1988, 1989; and Conrad and Biber, 2000).

Grounded in Systemic Functional Linguistics (Eggins, 2004; and Halliday, 1994), the APPRAISAL framework is presented as a set of choices that are available to the speaker for expressing and negotiating inter-personal stances and value positions in discourse. The model includes three interactive components: ATTITUDE, ENGAGEMENT and GRADUATION. ATTITUDE concerns feelings, such as emotional reactions, judgments of behaviour and evaluations of things. ENGAGEMENT comprises a set of resources by which means speakers adopt a position with respect to alternative opinions and voices. GRADUATION is used for scaling the intensity of an ATTITUDE or the degree of speaker investment in a proposition (Martin and White, 2005: 35–9). These three categories are further developed into an extensive range of sub-categories organised in networks.

The analysis presented below is limited in scope to the ATTITUDE subcategories of AFFECT and JUDGEMENT, and ENGAGEMENT. The analysis of JUDGEMENT is further restricted to instances of '*self*-JUDGEMENT' (i.e., expressions that are used by the CEOs to evaluate their company members' qualities and behaviour; see Section 5.2.1). Based on Fuoli and Paradis's (2014) model, these resources can be expected to play a key role in trust–repair discourse, and their frequency can, thus, be anticipated

to significantly change in BP's CEO letters after the spill. In Section 6 we explore this hypothesis. While other APPRAISAL resources may also play a role in trust-repair discourse, we decided to focus on those that, based on Fuoli and Paradis's (2014) model, can be expected to be most directly relevant and, thus, provide the most meaningful insights into BP's trust-repair strategies. In light of these restrictions, the following sections provide an overview of AFFECT, JUDGEMENT and ENGAGEMENT only. A complete description of the APPRAISAL model can be found in Martin and White (2005).

3.1 ATTITUDE

ATTITUDE refers to the expression of feelings and evaluations in discourse. The APPRAISAL model suggests a division of this area of inter-personal meaning into three areas: AFFECT, JUDGEMENT and APPRECIATION. As mentioned above, APPRECIATION was not included in the analysis, and will not, thus, be considered in detail here.

AFFECT concerns emotions and states of mind. It can be realised by a variety of linguistic expressions belonging to any word class – for example, adjectives (e.g., *happy/sad*), nouns (e.g., *joy/sorrow*), verbs (e.g., *love/hate*) or adverbials (e.g., *happily/sadly*). Expressions of AFFECT may have positive or negative 'valence' (i.e., they may be used to communicate positive or negative emotions). The APPRAISAL model classifies AFFECT into four main types according to semantic/functional criteria: DIS/INCLINATION, UN/HAPPINESS, IN/SECURITY and DIS/SATISFACTION (Martin and White, 2005: 71). Bednarek (2008: 172) modifies the AFFECT sub-system by creating an independent category for SURPRISE (previously included in IN/SECURITY). The following is an example of AFFECT from our corpus.

- (4) The explosion and fire on the Deepwater Horizon rig shocked everyone within BP and we feel great sadness that 11 people died.

JUDGEMENT refers to positive and negative evaluations of people according to different parameters. The system consists of two main sub-categories: SOCIAL SANCTION, which refers to the moral evaluation of people's behaviour (veracity and propriety), and SOCIAL ESTEEM, which concerns their normality, capacity or tenacity. As with AFFECT, JUDGEMENT can be instantiated by an open-ended set of linguistic expressions, including adjectives (e.g., *skilled*, *clever* and *courageous*), manner adverbials (e.g., *cleverly* and *in a diligent and responsible way*) and nouns (e.g., *success* and *leader*). Example 5 shows an example of JUDGEMENT from the corpus.

- (5) With the talent and commitment of the people of ExxonMobil, we are strong, resilient and well-positioned for the future.

All the examples given so far account for what Martin and White (2005) term ‘inscribed’ ATTITUDE – that is to say, feelings and evaluations that are explicitly conveyed by manifestly positive or negative wordings. However, as the authors note, in certain contexts ‘the selection of ideational meanings [may be] enough to invoke evaluation, even in the absence of attitudinal lexis that tells us directly how to feel’ (Martin and White, 2005: 62). Thus, in a persuasive text such as a CEO letter, and in a situation where the trustworthiness of the company has been damaged, a factual statement like that in Example 6 may be interpreted as implying positive evaluation, even though no explicit assessment is expressed.

- (6) To encourage excellence in risk management throughout the organization, we are reviewing how we incentivize and reward people.

Martin and White (2005) refer to this type of implicit evaluation as ‘invoked’ ATTITUDE. While invoked realisations may be relevant to understanding the discursive dynamics of trust, it is very hard, if it is at all possible, to identify and quantify instances of this type reliably. This study is, therefore, restricted to inscribed ATTITUDE (see Section 5.2.1).

3.2 ENGAGEMENT

The APPRAISAL system of ENGAGEMENT comprises linguistic resources used by speakers to indicate their stance towards alternative opinions and viewpoints, and to anticipate and manage potential responses from interlocutors (White, 2012). ENGAGEMENT incorporates a wide range of diverse phenomena that have traditionally been referred to in the linguistics literature with labels such as ‘modality’, ‘polarity’, ‘evidentiality’ and ‘attribution’ (Martin and White, 2005: 94). The function of these resources in discourse is examined from a dialogic perspective (Bakhtin, 1981); they are used by speakers to signal whether they anticipate their proposition to be controversial or likely to be questioned by the audience, and manoeuvre potential reactions and responses.

ENGAGEMENT resources are organised into different categories based on their communicative/dialogic function. A major dividing line is drawn between ‘monoglossic’ and ‘heteroglossic’ utterances. Monoglossic propositions are those in which other viewpoints are not recognised (i.e., bare or categorical assertions). Heteroglossic resources are broadly subdivided into those that entail dialogic ‘expansion’ and those that involve dialogic ‘contraction’ (Martin and White, 2005: 102). Dialogic expansion incorporates resources by which the dialogic space is opened up for alternative viewpoints and voices, while dialogic contraction subsumes options that serve the communicative purpose of challenging or restricting the scope for alternative positions and voices.

The ENGAGEMENT category of dialogic expansion involves two subcategories. ENTERTAIN refers to options that signal that the position advanced is to be seen as just one voice among others on a particular issue (i.e., to indicate that the speaker takes into consideration the possible existence of alternative viewpoints in addition to the one they are advancing). The resources that are incorporated into this category include modal auxiliaries (*may*, *might* and *could*), modal adverbs (*perhaps* and *probably*), epistemic mental predicates (*think*, *suspect* and *doubt*) and certain evidentials (*it seems* and *it appears*) (Martin and White, 2005: 105). The following is an example of ENTERTAIN from the corpus.

- (7) Despite these actions, ConocoPhillips considers it possible that the recession could restrain energy demand and prices for several years.

The other dialogic expansion subcategory, ATTRIBUTION, refers to linguistic resources by which the proposition is attributed to an external source, and speakers present themselves as having no stake in it (i.e., as simply conveying information). This category may be realised, for example, by reporting structures ('x [claims, believes, suggests]'), nominalisations of such structures ('assertion that' and 'claim that') or adverbial adjuncts ('according to'). An example of ATTRIBUTION from the corpus is as follows.

- (8) The report stated that decisions made by multiple companies and work teams contributed to the accident.

In contrast to communicative strategies of dialogic expansion, which serve the purpose of opening up the dialogic space for alternative opinions, dialogic contraction operates to exclude these alternatives. The resources of dialogic contraction are subdivided into two broad categories. The category of DISCLAIM incorporates formulations that invoke alternative viewpoints only to directly reject or replace them (Martin and White, 2005: 118). DISCLAIM encompasses negation/denial, adversative discourse markers (e.g., *however*, *but* and *yet*) and other wordings by which speakers indicate that the natural expectations arising from a proposition are not fulfilled (e.g., *surprisingly* and *even though*) (Martin and White, 2005: 120–1). Example 9 provides an instance of this category:

- (9) Our investigation report was published on 8 September 2010, and found that no single factor caused the accident.

The other main subgroup of dialogic contraction is PROCLAIM. This category includes expressions which, rather than directly rejecting other viewpoints, acts to limit the scope for alternative positions (Martin and White, 2005: 121). Wordings that are classified under this category include those that signal agreement between the speaker and the addressee and thereby exclude

alternative voices (e.g., *of course*, *admittedly* and *obviously*), wordings whereby the presence of the authorial voice is emphasised so as to suppress any resistance that might exist (e.g., *I contend*, *indeed* and *we firmly believe that*) and expressions of endorsement of attributed propositions (e.g., ‘the report found that’ and ‘the studies demonstrate that’). Example 10 contains an instance of PROCLAIM from the corpus.

- (10) I want to make it absolutely clear that we are not seeking a return to business as usual.

The categories outlined in this section were included in the corpus annotation scheme, which is described in Section 5.2.1. The next section describes the corpus.

4. Data

For this study we compiled a specialised corpus of CEO letters published by BP and four other major oil companies during a period of two years before and after the Deepwater Horizon oil spill. In Section 6.1, we compare the BP letters published after the spill with those published before the accident and by the other companies. Our aim is to identify significant patterns of change in BP’s discourse and gain insights into the company’s trust–repair strategies. The corpus is presented in Section 4.2. The next section briefly describes the genre of CEO letters.

4.1 The genre of CEO letters

CEO letters are an important genre of corporate public discourse: they are ‘the most prominent and widely read part of an annual report’ (Hyland, 1998: 224). From a communicative point of view, these letters play a fundamental role in framing the information included in the reports, steering the reader’s response to the facts proffered, and promoting a positive corporate image (Garzone, 2005; and Hyland, 1998). CEO letters can be regarded as a promotional genre and their rhetorical purpose is primarily persuasive: they are written to convince stockholders of the validity of their investment in a company and to seek continued financial support at the same time (Breeze, 2012, 2013; and Hyland, 1998).

Unlike the more technical parts of the annual report, which include audited information about a company’s financial performance, CEO letters are not subject to any legal obligations to truthfulness and transparency (Breeze, 2012: 6), but, rather, represent the CEO’s own perspective on the company’s performance. CEO letters thus typically exhibit a more personalised and overtly evaluative writing style, which is manifested in structural features that are characteristic of the letter genre, such as a

Year	Companies				
	<i>BP (UK)</i>	<i>Chevron (US)</i>	<i>ConocoPhillips (US)</i>	<i>ExxonMobil (US)</i>	<i>Royal Dutch Shell (UK)</i>
2008	1,113	839	1910	964	886
2009	1,085	897	850	876	1,014
2010	2,118	993	1325	935	1,326
2011	1,840	942	1758	877	1,300

Table 1: Corpus details (number of words per text)

salutation or direct address and the ‘real’ CEO’s signature, the frequent use of personal pronouns, and evaluative expressions to express the speaker’s stance (Bhatia, 2004; Garzone, 2005; Gillaerts and Van de Velde, 2011; and Hyland, 1998). Compared to the rest of the annual report, whose chief function is to provide detailed information about a company’s financial standing, CEO letters can be seen, primarily, to perform an inter-personal function, putting a face to the company and establishing a dialogue with the readership. As a result, these letters have ‘enormous rhetorical importance in building credibility and imparting confidence, convincing investors that the company is pursuing an effective strategy’ (Hyland, 1998: 224).

4.2 The corpus

Table 1 provides a detailed breakdown of the corpus contents. The corpus includes twenty CEO letters,³ of an average length of 1,192 words. The total corpus size is 23,484 tokens. The letters are included in the companies’ annual reports, which are public documents that can be freely downloaded from their websites.

5. Methods

The analysis presented in Section 6 combines quantitative and qualitative methods. Quantitative analysis is used to (i) identify patterns of change in BP’s use of AFFECT, JUDGEMENT and ENGAGEMENT resources and (ii) guide the qualitative analysis of the BP letters published after the spill. The qualitative analysis consists in an interpretive close reading of BP’s letters after the spill, which aims to elucidate the rhetorical function of the annotated features and investigate their interplay with other discursive resources and

³ BP’s 2008 and 2009 annual reports did not include a letter to shareholders signed by the CEO, but instead featured an interview with him. Given that the interview genre is radically different from that of CEO letters, we included in the corpus the letters to shareholders signed by the company’s Chairman.

motives. The results of both analyses are interpreted in light of Fuoli and Paradis's (2014) model of trust–repair discourse, in order to seek to explain the patterns we observe and to determine the strategies deployed by the CEO to repair stakeholders' trust after the spill. The quantitative data are derived from the manual annotation of the corpus. This section describes the corpus annotation process and the inter-coder agreement test that was conducted to validate the reliability and replicability of the annotations. This is preceded by a brief discussion of the limitations of automatic corpus methods for the analysis of evaluation.

5.1 Limitations of automatic corpus methods for the analysis of evaluation

Quantifying evaluative expressions using automatic corpus methods involves several challenges. Some of the main complexities are due to well-known properties of evaluation (for a review, see Hunston, 2011: 12–19); in particular:

- Evaluation may be realised by an open-ended range of expressions. Thus, analyses based on pre-set lists of evaluative forms will not, in most cases, be exhaustive (i.e., they will not cover *all* evaluation in a text or corpus). This also applies to cases where pre-defined sets of lexico-grammatical patterns are used as a starting point (e.g., Bednarek, 2009; and Hunston and Sinclair, 2000).
- Context is crucial for decoding evaluative meaning. Context and co-text play a critical role in determining the evaluative or non-evaluative nature of an expression, and, thus, need to be taken into account to identify relevant forms correctly. This applies to both explicit and implicit evaluation. Compare, for example:

(11a) ExxonMobil is dedicated to minimizing adverse risks and impacts associated with our products.

[Explicitly evaluative: 'devoted to a cause, ideal, or purpose']

(11b) This may seem strange in a column dedicated to that very subject, but I think it is excellent advice.

[Non-evaluative: 'used for one particular purpose']

This also holds true for very frequent polysemous ENGAGEMENT markers such as *know*, *show* or *understand*, which perform an epistemic/dialogic function in certain co-texts only (Fuoli, 2012: 66). Classifying evaluative expressions is also often impossible without taking co-text into account. For instance, the adjective *great* in Example 12 can only be classified as an instance of JUDGEMENT based on the co-textual environment in which it appears, rather than on inherent semantic aspects of the word.

- (12) It has truly been an honour for me to serve 12 years as chairman of what is one of the world's great enterprises.

Therefore, corpus techniques that treat words and expressions in isolation from their original co-texts (e.g., word-list or *n*-gram analysis) will make accurate analysis difficult to achieve.⁴

- Evaluative items may span multiple words (e.g., 'long-term commitment to excellence'). Therefore, automatic methods based on single words or *n*-grams may fail to recognise their 'real' boundaries. This might affect accuracy. If wordlists were used as a starting point to quantify evaluation, then, for instance, *commitment* and *excellence* in the example above would be counted as two evaluative items; we believe they are better analysed as part of one single evaluative expression instead. Similarly, the words *considers* and *possible* in Example 7 might erroneously be counted as two independent instances.

Automatic semantic tagging systems are, to date, insufficiently accurate, and often produce false positives and coding errors (see, for example, Murphy, 2013). Therefore, they do not offer a reliable alternative to traditional corpus methods. Conversely, manual corpus annotation allows us to overcome the challenges described above. By manually annotating text, all evaluative expressions can be identified and counted.⁵ Identification and classification can be more accurate, as context and co-text are properly accounted for. However, manual corpus annotation is a complex and subjective process, and this poses challenges to achieving transparent, reliable and replicable analyses. The next sections describe the strategies adopted in this study to address these issues.

5.2 Annotation procedure

The annotation of our corpus comprised the following main steps:

(1) *Annotation scheme design*

As the first step in the process, we needed to establish principles for adapting the general-purpose APPRAISAL framework to the goals of our analysis and the specificities of the texts that

⁴ Clearly, one could repeatedly shift between word or *n*-gram lists and the original texts to study the context-specific meaning of evaluative expressions, but that seems unpractical, and defeats the purpose of using automatic techniques.

⁵ In the analysis presented here, all evaluative expressions belonging to selected APPRAISAL categories are identified and classified. The analysis is, therefore, limited in scope, but still exhaustive in the sense that it covers all expressions instantiating the categories considered, and is not limited to a pre-determined set of forms.

we investigated (Hommerberg and Don, 2015), and also to create context-specific guidelines to implement the framework to the annotation of our corpus. This phase involved decisions concerning the level of granularity to adopt in the analysis, what context-specific guidelines to follow when identifying and classifying instances, and what rules to apply when encountering ambiguous items. The final product of this phase was a context-specific annotation manual, which appears in Appendix A. The annotation scheme is presented in Section 5.2.1.

(2) *Annotator training and inter-coder agreement test*

Manually annotating APPRAISAL is a complex and subjective process. This may affect the reliability of the corpus annotation and of the quantitative data derived from it. In addition, it may hinder the replicability of the annotation procedure. Different people may produce different analyses, unless transparent and shared annotation principles are established. The first measure we took to tackle these matters was to create a manual that contains explicit annotation guidelines for all annotators to follow. In addition, the annotation of the whole corpus was preceded by an inter-coder agreement test. A detailed account of the test is provided in Section 5.3. One of the outcomes of the test was a jointly annotated ‘gold standard’ sample that was used as a guide for the annotation of the whole corpus.

(3) *Corpus annotation*

Matteo Fuoli manually annotated and classified all instances of the categories included in the annotation scheme (Figure 3). The annotation was based on the guidelines outlined in the annotation manual and the gold-standard sample created as a result of the inter-coder agreement test. Examples of APPRAISAL expressions included in Martin and White (2005) and Bednarek (2008) were also consulted during this task. Every text in the corpus was examined at least twice and at different points in time. The annotation required around twenty-five hours of work (i.e., approximately one hour per one-thousand words). The software program that we used for this task is the Content Annotation Tool or CAT (Bartalesi Lenzi *et al.*, 2012).⁶ CAT is a web-based annotation tool that provides a user-friendly interface for annotating words and expressions in text (see Figure 1). Coding

⁶ CAT can be accessed, free of charge, online at: <https://dh.fbk.eu/resources/catcontent-annotation-tool>. An alternative and widely used annotation program is the UAM Corpus Tool (O'Donnell, 2008). Some scholars (e.g., Bednarek, 2008) have used the Altova XMLSpy editor (see: <http://www.altova.com/xmlspy.html>). This tool, however, is not freely available and is not specifically designed for annotating corpora, which makes it less intuitive and easy to use compared to UAM or CAT. For a review of different manual corpus annotation tools, see O'Donnell (2014).

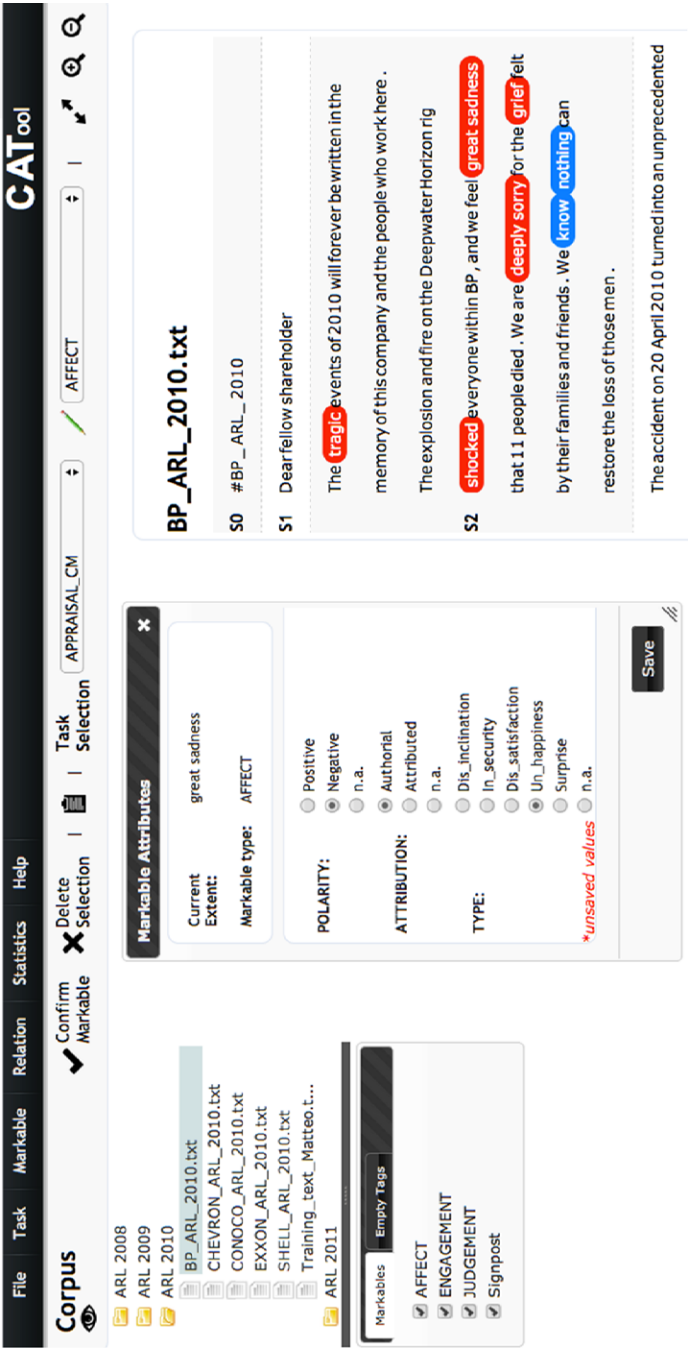


Figure 1: The CAT user interface

Tokens	Type	Polarity	Attribution	Comments
<i>proud</i>	Dis_satisfaction	Positive	Authorial	
<i>want</i>	Dis_inclination	Positive	Attributed	
<i>deeply appreciate</i>	Dis_satisfaction	Positive	Authorial	
<i>inspire</i>	Dis_inclination	Positive	Authorial	
<i>admired</i>	Dis_satisfaction	Positive	Attributed	
<i>confident</i>	In_security	Positive	Authorial	
<i>tragic</i>	Un_happiness	Negative	N/A	covert affect
<i>shocked</i>	Surprise	Negative	Authorial	negative surprise?
<i>great sadness</i>	Un_happiness	Negative	Authorial	
<i>deeply sorry</i>	Un_happiness	Negative	Authorial	
<i>grief</i>	Un_happiness	Negative	Attributed	

Figure 2: The converted CAT output

schemes can be fully customised to fit the annotation task at hand. CAT stores the annotations as stand-off XML. This format can easily be converted into a tabular ‘case-by-variable’ format (see Figure 2), which can be used with spreadsheet and statistical software for further processing and analysis.

5.2.1 Annotation scheme

As mentioned above, the analysis was restricted to the ATTITUDE subsystems of AFFECT and JUDGEMENT, and ENGAGEMENT. The annotation of JUDGEMENT was restricted to (i) positive instances and (ii) evaluative expressions that are used to praise the company and its members, excluding other evaluative ‘targets’. In this way, we could isolate those evaluative expressions that may directly realise the emphasise-the-positive strategy, in accordance with Fuoli and Paradis’s (2014) model. In addition, only explicit (inscribed) instances of APPRAISAL were annotated. Invoked instances were found to be extremely difficult to identify reliably and were, therefore, ignored in the annotation process. Further, a distinction was made between evaluative expressions that are included in ‘realis’ contexts, and those that are included in ‘irrealis’, non-factual contexts (i.e., expressions that are not used to afford an actual evaluation of the company but rather to refer to future or hypothetical desirable states of affair) (Taboada *et al.*, 2011: 278–9). Compare, for example:

- (13a) I would also like to thank our people for the tremendous effort, dedication and passion they have shown.

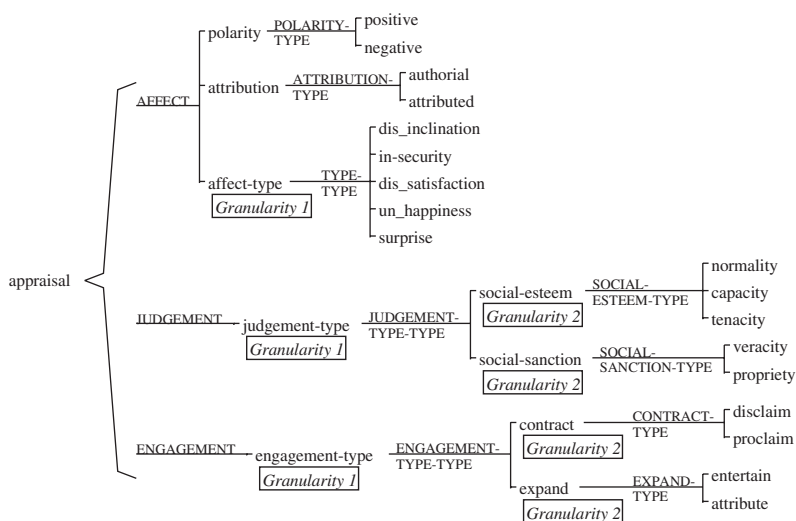


Figure 3: Annotation scheme⁷

- (13b) We can never eliminate every hazard, but we can become an industry leader in limiting and understanding risk.

While both types of expressions may play a role in trust management, unrealistic evaluations were eventually excluded from the analysis on the grounds that robust agreement on their identification could not be reached. More detailed information about coding decisions can be found in the annotation manual (see Appendix A).

5.3 Inter-coder agreement test

In the context of a corpus annotation task, ‘inter-coder agreement’ is a measure of the extent to which independent annotators make the same decisions when assigning pre-defined categories to units of text. Inter-coder agreement can be used as a criterion for evaluating the reliability of the corpus annotation, based on the assumption that ‘data are reliable if coders can be shown to agree on the categories assigned to units to an extent determined by the purposes of the study’ (Artstein and Poesio, 2008: 557). Demonstrating reliability is necessary to ensure the reproducibility⁸ of an analysis (Krippendorff, 2004: 215). If independent annotators using the same guidelines consistently make equivalent coding choices, we can expect other

⁷ The annotation scheme diagram was generated using the UAM Corpus Tool (O’Donnell, 2008).

⁸ This is the term used by Krippendorff (2004), though ‘replicability’ is an equivalent term.

annotators to produce similar interpretations, given similar circumstances. Conversely, a low level of agreement may indicate either that the coding scheme is defective or not sufficiently explicit, or that the annotators need more training. Reliability is a pre-requisite for the validity of a coding scheme (i.e., the extent to which it captures the ‘truth’ about the phenomenon under investigation) (Artstein and Poesio, 2008; and Krippendorff, 2004). As Krippendorff (2004: 213) remarks, however, ‘reliability is a necessary, but not a sufficient, condition for validity’. Independent analysts may, in fact, share similar interests, biases and prejudices that may affect the way they interpret the data in similar ways.

Inter-coder agreement measures have been widely used in Computational Linguistics, where a growing need for manually annotated corpora in the development of computational models of semantic and pragmatic phenomena has been accompanied by concerns about the subjectivity involved in the analysis of these aspects of language (Artstein and Poesio, 2008). Despite the importance of manual annotations in the creation of linguistic corpora, measures of inter-coder agreement are still rarely reported in the literature on corpus linguistics (Spooren and Degand, 2010; and Voormann and Gut, 2008). This also applies to research using the APPRAISAL framework where, in spite of criticisms concerning the arbitrariness of its taxonomy and the subjectivity of the analyses (Taboada and Carretero, 2012 [2010]: 278), explicit accounts of inter-coder agreement are not commonly provided. There is, however, a growing awareness of the importance of accounting for reliability in this area, and several recent studies based on APPRAISAL theory include a discussion of inter-coder agreement (Fuoli, 2012; Hommerberg and Don, 2015; Read and Carroll, 2012; Ryshina-Pankova, 2014; and Taboada and Carretero, 2012 [2010]).

Annotating APPRAISAL can be a very complex and subjective task. This is mainly due to the fact that the framework itself is conceived as a flexible interpretive tool, a ‘basic draft of categories’ (Hommerberg and Don, 2015), rather than a definitive model of evaluation that can be applied to any kind of text in a mechanical way. These complexities are heightened in the context of this analysis by the fact that the very communicative goal of the texts considered is that of promoting a positive image of the companies. This implies that all utterances included in the texts may be safely interpreted as conveying evaluation, which makes it very hard to confidently identify and annotate units of ‘explicit’ APPRAISAL. As a result, reaching high inter-coder agreement scores may, in this context, be a particularly challenging task, and the ideal of coders who ‘work completely independently and agree substantially’ (Spooren and Degand, 2010: 253) may be difficult to reach, too. According to Spooren and Degand (2010), there are three main strategies that can be applied in situations where reaching high inter-coder agreement between independent coders is challenging:

- (1) Double coding. Two annotators code the entire data independently and then discuss all the disagreements until full consensus is

reached. This strategy improves the quality of the annotation as it promotes cooperative rather than idiosyncratic coding strategies, and because the coders are forced to make their reasoning explicit and convince each other in case of disagreement. This is, however, a very time consuming procedure.

- (2) Partial overlap between two or more coders. A portion of the data is double coded, while the rest of the corpus is annotated by only one person. Inter-coder agreement is calculated on the double-coded sample. Disagreements are discussed and reconciled between the annotators to enhance the quality of the single-coded data.
- (3) One coder does all. The entire corpus is annotated by only one coder. While the annotator can be expected to adopt subjective annotation strategies, we can assume these strategies to be systematic. This means that if the annotator has a tendency to over-annotate one specific category, that category will be globally over-represented in the analysis, which should not present an obstacle to answering the research question.

Due to time and resource constraints, double coding the entire corpus was not a viable option. The one-coder-does-all alternative was also rejected, on the grounds that it is the weakest form of reliability check. Given the nature of the task at hand, it would leave open the possibility for individual biases to affect substantially the outcome of the analysis. Therefore, in this study, we have adopted the second of the solutions above.

5.3.1 Inter-coder agreement test design

The design of our inter-coder agreement test closely resembles that of Read and Carroll (2012). The test involved two expert annotators—that is to say, the authors of this paper. Both annotators have previous experience with corpus annotation using the APPRAISAL framework. The test consisted of two related annotation tasks:

- (1) Identification of expressions of APPRAISAL, a process which can be referred to as ‘markable identification’ (Artstein and Poesio, 2008) or ‘unitization’ (Krippendorff, 2004);
- (2) Classification of the expressions identified according to the scheme outlined in Section 5.2.1.

These tasks were performed based on the annotation manual (see Appendix A), using CAT. Before the test, a pilot study was carried out with the dual purpose of (i) calibrating the annotation guidelines, and (ii) training on the guidelines and the use of CAT. During the pilot test, the annotators independently tagged a small sample of texts with the same characteristics

of those included in the corpus. This was followed by a discussion session, during which the coders compared their annotations and reconciled disagreements, to the extent possible. Following Wiebe *et al.* (2005), the reconciled annotations from the pilot test were used as a gold standard reference for the subsequent annotations. After the training phase, the test took place. The coders annotated random paragraphs independently, amounting to approximately 25 percent of the corpus. The paragraphs were randomly selected using an automatic procedure.⁹ As per Read and Carroll (2012), the annotation was performed over three rounds, punctuated by an intermediate phase of analysis and a reconciliation of disagreements. These intermediate sessions also served to address unanticipated annotation problems and further refine the annotation guidelines. Inter-coder agreement, both prior to and after reconciliation, was measured and recorded. In total, the test required approximately sixty person-hours to complete.¹⁰ The next section reports the results of the test.

5.3.2 Inter-coder agreement results

For both the markable identification and classification tasks, we report the agreement scores obtained prior to, as well as after, reconciliation. However, it is important to note that only the scores obtained before reconciliation count as reliability data proper (Krippendorff, 2004: 219). Nevertheless, the agreement scores achieved through discussion and reconciliation provide a useful indication of how much agreement can be reached when collaboratively annotating our corpus, and when inaccuracies and inconsistencies due to annotator fatigue and distraction are filtered out from the data.

The results of the markable identification task are reported in Table 2. Following Read and Carroll (2012) and Taboada and Carretero (2012 [2010]), to measure the amount of agreement between annotators on this type of task we use ‘precision’ (PRE), ‘recall’ (REC) and ‘F measure’ (F-score) scores. Kappa statistics (Cohen *et al.*, 1960) are not suitable for this task, as it focusses on the identification of markables rather than the labelling of units of fixed length. If we take the annotations produced by Matteo Fuoli as the gold standard (i.e., the ‘correct’ annotations):

- PRE indicates the fraction of units identified by Charlotte Hommerberg that are actually relevant (or ‘correct’) – in other

⁹ The random paragraphs were selected using the freely available online tool called Random Line Picker. The tool is available online at: <http://textmechanic.com/Random-Line-Picker.html>

¹⁰ This figure accounts for the time needed to complete the annotation tasks, discussing and reconciling the disagreements, and analysing the agreement data. It sums the time spent by both annotators, thus on average each annotator devoted about thirty hours of their time to the test. Matteo Fuoli, however, was in charge of data preparation and analysis, so he devoted proportionally more time to the test.

	PRE		REC		F-score	
	Before	After	Before	After	Before	After
AFFECT	0.82	1.00	0.90	1.00	0.86	1.00
JUDGEMENT	0.71	0.92	0.71	0.98	0.71	0.95
ENGAGEMENT	0.77	0.97	0.85	1.00	0.81	0.98
<i>Mean</i>	0.77	0.96	0.82	0.99	0.79	0.98

Table 2: Inter-coder agreement results: markable identification task

words, those that have also been annotated by Matteo Fuoli. It is calculated by dividing the number of units identified by both annotators (i.e., the relevant units retrieved) by the number of units annotated by Charlotte alone (i.e., all the retrieved units).

- REC represents the fraction of relevant units that have been successfully identified by Charlotte. It is calculated by dividing the number of relevant units identified by Charlotte (i.e., those that have also been annotated by Matteo) by the total number of relevant units (i.e., the number of units annotated by Matteo, or the gold standard).
- F-score provides a synthetic measure of PRE and REC. It is calculated as the weighted harmonic mean of precision and recall.

For the sake of simplicity and following Read and Carroll (2012) and Wiebe *et al.* (2005), we consider overlapping annotated text units as matches.¹¹ As Table 2 shows, inter-coder agreement for the markable identification task before reconciliation is substantial, with an overall mean F-score of 0.79. The scores obtained from this test are higher than those reported in Read and Carroll (2012) and similar to those obtained by Taboada and Carretero (2012 [2010]). The table also shows that, while agreement for the categories of AFFECT and ENGAGEMENT is robust, JUDGEMENT was a more problematic category. This is not surprising. As mentioned earlier, one of the primary communicative goals of the texts analysed is precisely that of conveying a positive evaluation of the company. In a highly evaluative and positive context, discerning explicit from invoked positive JUDGEMENT and identifying bounded evaluative units was found to be a very complex task. Notwithstanding this, the scores obtained are higher than those reported for this category in Read and Carroll (2012).

Table 2 also shows the level of agreement reached after thoroughly discussing discrepancies between the annotations. After reconciliation, agreement becomes almost perfect, with an F-score of 0.98. The discussion and reconciliation sessions allowed annotators not only to identify the

¹¹ Initial efforts were spent to try to reach agreement on boundary placement, but this turned out to be an impracticable task.

	Observed agreement			
	Granularity 1		Granularity 2	
	Before	After	Before	After
AFFECT	0.93	0.96	n/a	n/a
JUDGEMENT	0.94	0.98	0.75	0.92
ENGAGEMENT	0.91	1.00	0.88	1.00
<i>Mean</i>	0.93	0.98	0.81	0.96

Table 3: Inter-coder agreement results: classification task

best solution for the disagreements collaboratively, but it also highlighted various instances of inconsistencies and of unproblematic annotations that were missed due to fatigue or distraction. These issues, which are an inevitable drawback of manual annotation, have clearly affected the inter-coder agreement scores obtained before reconciliation.

Table 3 reports the results for the classification task. The table shows the levels of agreement reached in the classification of the three main APPRAISAL categories included in the coding scheme and at two levels of granularity. The scores in the table refer to the ‘observed agreement’ between the annotators, both before and after reconciliation. Observed agreement is obtained by dividing the total number of judgments on which the annotators agree by the total number of markable items (i.e., the number of instances that were identified by both annotators independently). Observed agreement is not a very robust measure of agreement, as it does not account for chance. While Cohen’s kappa would be preferable, the highly skewed distribution of the categories in our data, where a disproportionate number of instances fall under one category, leads to the well-known ‘paradox’ that high levels of observed agreement correspond to very low kappa scores (Artstein and Poesio, 2008; and Di Eugenio and Glass, 2004). To avoid this problem, we report only observed agreement scores, with the warning that these figures are not corrected for chance and are, therefore, relatively poor indicators of reliability.

Table 3 shows very high levels of observed agreement on the classification task at the lower level of granularity, with all the coefficients above 0.90. Predictably, the scores decrease when the analysis becomes more fine-grained and more options are available (note that only one level of granularity was considered for AFFECT; see the coding scheme in Figure 3). Agreement for the category of JUDGEMENT at this level of granularity is particularly low, which is due to the ambiguity and under-specification of the classification criteria for these subcategories in the APPRAISAL literature, and to the fact that many instances of JUDGEMENT were found to be compatible with more than one reading/label. As shown in the table, upon reconciliation most disagreements could be solved and observed agreement neared 100 percent.

In sum, the inter-coder agreement test we carried out on a sample from the corpus indicates substantial agreement between independent annotators in the tasks of identifying and classifying instances of APPRAISAL in our corpus, based on the guidelines designed for this study (see Appendix A).

6. Analysis

As discussed above, Fuoli and Paradis (2014) propose that, when a trust-breaking event occurs, the trust-breaker's ability, integrity and benevolence may be at stake. As a result, the trust-breaker's discourse will be geared towards reshaping the trust-giver's impressions of their trustworthiness along these three dimensions, with the ultimate goal of restoring trust.

Based on the events that followed the Deepwater Horizon oil spill and the public reactions to BP's handling of the crisis, we may infer that, after the accident, the company's trustworthiness was at stake on all levels. BP's ability was being questioned due to its repeated failed attempts to control the spill.¹² Its integrity was undermined by its initial attempts to downplay the effects of the spill (Webb, 2010) and its biased reports of the amount of oil gushing into the sea (O'Connor, 2011; and Shogren, 2011). BP's benevolence came under intense scrutiny after some controversial public statements by the company's CEO ('I want my life back') and President ('We care about the small people'), which cast doubts on BP's real priorities and created the impression that it did not truly care about the people who were affected by the accident (O'Connor, 2011).

In this section, we combine quantitative and qualitative analysis of the use of APPRAISAL resources in BP's CEO letters to investigate how ability, integrity and benevolence are discursively (re-)constructed. The analysis addresses the following questions:

- (1) Does the frequency of AFFECT, JUDGEMENT and ENGAGEMENT expressions change in BP's discourse after the spill? If so, how?
- (2) How are these resources deployed in BP's letters to repair trust in the company's ability, integrity and benevolence? What aspects of the company's trustworthiness are emphasised? And,
- (3) Is there any difference between the trust-repair strategies adopted in the 2010 and 2011 letters?

In Section 6.1, we explore the distribution of the APPRAISAL categories that we considered in order to identify patterns of change in BP's discourse after the spill. The patterns we identified are used as a guide for the qualitative

¹² For a detailed account of all the attempts at stopping the spill, see http://www.nytimes.com/interactive/2010/05/25/us/20100525-topkill-diagram.html?_r=0

analysis of BP's texts, which is presented in Section 6.2. In Section 6.3, the trust–repair strategies that emerge from the analysis are discussed.

6.1 Quantitative analysis

In order to identify significant changes in the use of the APPRAISAL resources considered in BP's post-crisis discourse, the BP letters published after the spill were compared with those published before it and by the other oil companies in the same time period. As the other companies were only indirectly affected by the spill, their texts were treated as a baseline.

The frequency of expressions of AFFECT, ENGAGEMENT and JUDGEMENT in each text included in the corpus was compared using log-linear analysis. The frequency data were classified in a three-way contingency table according to three factors: APPRAISAL category, company and year. A log-linear model, including the main effects of these three factors, was fitted to the data to calculate the expected frequency counts. These values were compared to the observed frequencies.¹³

The mosaic plot in Figure 4 shows the results of the analysis. It represents schematically the observed frequency of the annotated features (cell size) and marks the statistically significant deviations from the expected frequencies through a shade/outline scheme. Shaded cells with a solid edge denote significantly over-represented values, while shaded cells with a dashed edge correspond to significantly under-represented values. The raw frequencies of the different APPRAISAL types, which were used for the log-linear analysis, are reported in the boxes.¹⁴

From the results presented in Figure 4, it becomes apparent that the use of APPRAISAL resources in BP's CEO letters changed after the spill. The plot shows that AFFECT and ENGAGEMENT are both significantly over-represented in BP's letters after the accident. The frequency of ENGAGEMENT markers in the 2010 letter shows a greater deviation from the expected values, with standardised residuals larger than four. Conversely, JUDGEMENT is significantly under-represented in the letter released the year following the accident. The letters published by BP after the spill thus appear to place greater emphasis on emotions and to shift towards a more overtly dialogic style. The pattern observed for JUDGEMENT suggests that, in the letter published the year after the spill, the CEO adopts a more neutral, factual style, with comparatively fewer explicit positive assessments of the company's performance and qualities. Notably, similar patterns cannot be observed for any of the other companies that we considered.

The results of the quantitative analysis provide useful insights into the CEO's discursive response to the spill, highlighting patterns of change and salient phenomena in the company's discourse after the accident. These results can be used as a guide for a more fine-grained, qualitative analysis of

¹³ All statistical analysis was performed using R, version 3.0.1 (<http://www.R-project.org/>).

¹⁴ Differences in text length are not problematic for the analysis.

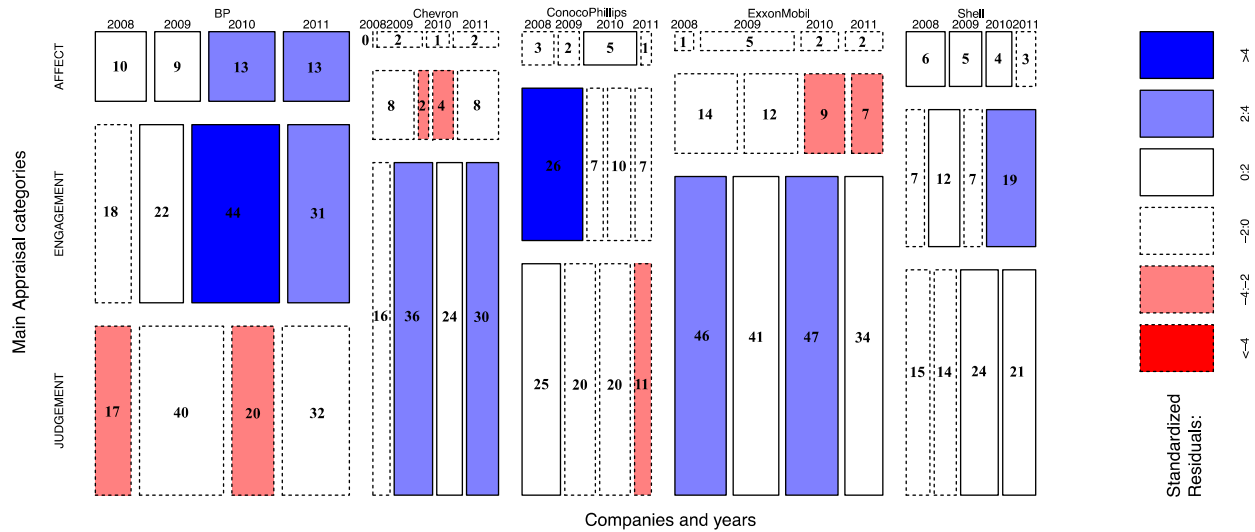


Figure 4: Mosaic plot: Analysis results. Standardised residuals (z-scores) are a statistical measure of the difference between observed and expected frequencies. Values higher than +2 or smaller than -2 indicate that this difference is significant.

BP's texts, which should provide a more complete picture of the function of the annotated features in context, and of the discourse strategies adopted by the CEO in the attempt to restore trust in BP. It is to this type of analysis that we now turn.

6.2 Qualitative analysis

As shown in Figure 4, AFFECT is over-represented in both BP's 2010 and 2011 letters. Unsurprisingly, most of the instances of this category found in the 2010 letter are used to convey negative emotions. They are placed in a very prominent position, at the beginning of the text, as shown in Example 14.

- (14) Dear fellow shareholder,
The tragic¹⁵ events of 2010 will forever be written in the memory of this company and the people who work here. The explosion and fire on the Deepwater Horizon rig shocked everyone within BP, and we feel great sadness¹⁶ that 11 people died. We are deeply sorry for the grief felt by their families and friends. We know nothing can restore the loss of those men.

The expressions of AFFECT included in this excerpt (underlined) refer to both the company's emotional reaction to the accident as well as to the emotions felt by the families and friends of the victims. These expressions serve the dual purpose of conveying BP's affective stance towards the accident, and to acknowledge the suffering caused by it. In this sense, the ultimate communicative function of these expressions in this particular context can be seen as that of showing empathy and care for the people affected by the spill. These feelings are further stressed in the following paragraph, where the CEO foregrounds his personal attachment to the areas hit by the spill by referring to his roots in one of the affected states, and recalling fond childhood memories on the Gulf coast.

- (15) And it all started in a part of the world that's very close to my heart. I grew up in Mississippi, and spent summers with my family swimming and fishing in the Gulf. I know those beaches and waters well. When I heard about the accident I could immediately picture how it might affect the people who live and work along that coast.

¹⁵ The word *tragic* in Example 14 was annotated as an instance of COVERT AFFECT (Bednarek, 2009).

¹⁶ Even though GRADUATION was not included in the analysis, GRADUATION markers such as *deeply* or *great* were by convention annotated as belonging to an AFFECT or JUDGEMENT unit, when used as modifiers. This choice has no impact on the quantitative figures. This choice is accounted for in the annotation manual.

If we interpret this passage in the light of the widespread resentment caused by some controversial declarations by the company's former *foreign* CEO and President in the aftermath of the spill (see Section 2), it is possible to see the CEO's emphasis on his American identity and affective attachment to the Gulf ('very close to my heart') as a strategic move aimed at re-negotiating solidarity and alignment with the victims of the spill and, at the same time, counteracting the widely held impression that the company did not truly care about them. In other words, the CEO's personal bond to the Gulf Coast can be taken as a 'guarantee' that the company of which he is the leader truly cares about the people affected by the spill, and will strive to protect them and the Gulf Coast's environment from its negative effects.

One conclusion that emerges from the discussion above is that, in the 2010 letter, AFFECT plays a key role in the CEO's attempt to communicate empathy, solidarity and care towards the victims of the spill. From a quantitative point of view, AFFECT is also a prominent category in the 2011 letter. In this text, however, expressions of AFFECT perform a different communicative function. Notably, while the vast majority of the instances of AFFECT found in the 2010 letter have a negative valence, the 2011 text includes a comparatively high number of instances of positive AFFECT. AFFECT is used here to foster confidence in the company's recovery and future prospects. The CEO declares, for instance, to be 'proud' of how BP responded to the accident, 'pleased' about several legal settlements reached during the preceding year, and 'confident' about BP's 'ability to design, engineer and operate large installations safely'.

As shown in Figure 4, ENGAGEMENT is significantly over-represented in both BP's letters after the spill and is particularly prominent in the 2010 letter. This finding indicates that the letters published after the accident are more overtly dialogic (i.e., the CEO explicitly engages with alternative propositions, opinions and voices more often). This tendency is particularly evident in the 2010 letter, where the CEO responds to the concerns and criticisms of several stakeholders. Towards the end of the text, for instance, he entertains an imaginary dialogue with the stakeholders and seeks to reassure them that the company 'has learned its lesson from the spill'.

- (16) I have heard people ask "Does BP 'get it'?" Residents of the Gulf, our employees and investors, governments, industry partners and people around the world all want to know whether we understand that a return to business-as-usual is not an option. We may not have communicated it enough at times, but yes, we get it. Our fundamental purpose is to create value for shareholders, but we also see ourselves as part of society, not apart from it.

A striking observation that can be made from Example 16 is that all the markers of ENGAGEMENT identified (underlined) are dialogically contracting (i.e., they act to reject or disallow alternative viewpoints and

opinions). The leading question and answer pair ‘Does BP “get it”? [...] yes, we get it’,¹⁷ the three instances of negation/denial and the counter-expectation marker, *but*, all contribute to ‘fending off’ and pre-empting adverse opinions and discourses. This pattern characterises the whole 2010 letter, where dialogically contracting ENGAGEMENT markers are used repeatedly to reject or pre-empt potential criticisms of the company’s decisions and actions, and to protect the company’s integrity from further damage.

ENGAGEMENT markers in the 2011 letter are also, for the vast majority, dialogically contracting. However, in this letter they perform a different communicative function. ENGAGEMENT resources are used by the CEO to boost the credibility of his promises and predictions about the future of the company, rather than to respond to criticisms and concerns. A clear example of this is reported in Example 17:

- (17) ● First and foremost, you will see a continuing, relentless focus on safety and risk management.
- You will see the company play to its strengths [...].
 - You will see a company that is simpler and more focused [...].
 - You will see a company that is organized effectively and applies its standards consistently.
 - You will see more visibility from us on our individual businesses.
 - You will be able to measure the effects of active portfolio management [...].
 - You will be able to measure the contribution of new upstream projects with higher margins, as they come onstream over the next three years.

The repeated, emphatic use of the evidential verbs *see* and *measure* in this example has a strong dialogically contracting effect, since these verbs act to rule out any doubt or concern that any scenario other than that envisaged by the CEO will occur. Put differently, the reader should trust the CEO’s statements, as they will be able to check them against tangible evidence. The reassuring, trust-building force of these utterances is bolstered by the repeated use of the second-person pronoun, *you*, which construes a direct, ‘face-to-face’ relationship between the CEO and the addressees, and indexes honesty and personal commitment. Interestingly, the second-person pronoun is never found in the other BP letters and is extremely rare in the other companies’ texts. In sum, ENGAGEMENT resources are more frequently used in BP’s letters after the accident. In the 2010 letter they mainly serve to protect the company’s integrity, whereas in the 2011 text they are primarily used to boost shareholders’ confidence in BP’s recovery.

¹⁷ The question-answer pair was annotated as a single unit of ENGAGEMENT:CONTRACT:DISCLAIM.

As the quantitative analysis has shown, JUDGEMENT is significantly under-represented in BP's 2010 letter. This result suggests that the evaluative force of this text is, comparatively, toned down. Explicit positive assessments of the company's behaviour and performance are rare, and positive evaluations tend to be invoked, rather than inscribed. Where positive evaluative expressions are used, they serve to stress BP's resilience to the crisis – as, for instance, in Example 18:

- (18) The sound underlying performance across our business continues to give us a solid foundation, and speaks volumes for the inner strengths of BP and our people.

The relative backgrounding of positive evaluation in BP's 2010 letter becomes strikingly evident if we compare this text with the letter published by one of the company's main competitors, ExxonMobil, in the same year.

- (19) To Our Shareholders
ExxonMobil continues to deliver superior long term shareholder value. We succeed by upholding the values that set us apart: a commitment to safety, operational excellence, and risk management; a disciplined, long term approach to investing; and the development and application of advanced technology and innovation.

As Figure 4 shows, and Example 19 confirms, positive JUDGEMENT in ExxonMobil's 2010 letter is strongly foregrounded. ExxonMobil's CEO adopts a highly evaluative writing style, characterised by a repeated use of positive evaluative expressions that emphasise the company's qualities and competitiveness. Interestingly, ExxonMobil was responsible in 1989 for one of the most catastrophic oil spills in history, Alaska's Exxon Valdez oil spill. In relation to this, it is noteworthy that the quantitative pattern for the use of APPRAISAL resources in ExxonMobil's 2010 letter is opposite to that observed for BP. ENGAGEMENT is significantly under-represented in ExxonMobil's text, leading to a very assertive, monologic writing style, whereas JUDGEMENT is over-represented. Considering that ExxonMobil was frequently mentioned in the news in the aftermath of the Deepwater Horizon accident, when comparisons were often drawn between the two spills, we might read these patterns as indicating ExxonMobil's strategic attempt to differentiate itself discursively from the competitor, in order to protect its image from the collateral damage that the BP spill may cause.

6.3 BP's trust–repair strategies

The analysis presented above has highlighted clear patterns of change in the use of APPRAISAL resources in BP's CEO letters following the Deepwater Horizon oil spill. In this section, we interpret the results of the analysis in

light of Fuoli and Paradis's (2014) model and seek to provide a concise answer to our research questions.

As shown above, BP's 2010 letter places substantial emphasis on emotions. Expressions of AFFECT are primarily used to show empathy and care towards the victims of the spill. In terms of Fuoli and Paradis's (2014) model, we may thus conclude that the use of AFFECT resources in this text is geared towards constructing benevolence. The highly dialogic nature of the 2010 text, evidenced by the higher-than-expected use of ENGAGEMENT expressions, can be explained in light of the numerous controversies and concerns generated by the accident. The CEO expends considerable effort in trying to neutralise the negative discourses about the company that circulated in the aftermath of the accident, and ENGAGEMENT resources play a central role in this endeavour. The use of ENGAGEMENT resources in the 2010 letter may thus be seen as configuring a defensive discourse strategy which aims to protect the company's integrity. Finally, the under-use of explicit positive evaluative language may also be seen as contributing to constructing integrity. Indeed, given the magnitude of the effects of the spill, frequent explicit positive self-assessments may have come across as inappropriate and insincere, further undermining BP's trustworthiness. The adoption of a relatively more factual and objective tone may thus be read as an attempt to communicate humbleness and credibility. The backgrounding of explicit positive evaluation may also be seen to imply that the ability facet of BP's trustworthiness is de-emphasised. Therefore, we may conclude that, in the 2010 letter, BP's CEO adopts a trust-repair strategy that emphasises the company's benevolence and integrity, while toning down ability.

The picture that emerges from the analysis of the 2011 letter is rather different. Ability is more strongly emphasised in this text, and optimism in the company's recovery and future performance is promoted by the CEO. As far as the integrity facet of BP's trustworthiness is concerned, credibility appears to be the main focus. The CEO aims to reassure the shareholders that the company will be able to overcome the crisis and return to profitability. Through the use of dialogically contracting expressions of ENGAGEMENT, he seeks to boost the credibility of his predictions, and remove uncertainties about the company's future. Compared to the 2010 letter, the benevolence facet of BP's trustworthiness appears to be less explicitly foregrounded in this text. Therefore, the 2011 letter shows a change in the CEO's trust-repair discourse strategy, with a stronger focus on the company's ability and a significant communicative effort to restore and promote shareholders' confidence.

7. Conclusion

Manual corpus annotation allows for exhaustive and detailed corpus-based analyses of evaluative language that would not be possible with purely automatic techniques, given the complex and context-dependent nature of

evaluation in discourse. As discussed above, transparency, reliability and replicability are crucial issues that need to be addressed when taking this kind of approach. We propose that these aspects can be optimised by adopting explicit annotation guidelines, and by assessing and reporting inter-coder agreement. The results of the inter-coder agreement test reported above are encouraging, as they show that, in spite of the complexity of the task at hand, robust agreement between independent coders can be reached, given explicit guidelines and appropriate training.

Clearly, the approach discussed here has several limitations, the most important being scalability (i.e., the extent to which it can be applied to larger corpora). Compared to purely automatic corpus methods, this type of approach is relatively time-consuming and resource-intensive. As a consequence, the amount of data that can be processed is, in normal circumstances, more limited than with traditional automatic techniques. We suggest that scalability may be improved by (i) using random sampling techniques as opposed to annotating full texts, and (ii) simplifying the coding scheme so as to make the manual annotation process faster. Another approach to address these limitations is to develop new, advanced software tools to assist, simplify and quicken the process of manual corpus annotation. We hope that this article will inspire more research and development in this area.

Acknowledgments

We would like to thank Joost van de Weijer (Lund University, Sweden) for his support with statistics and helpful comments on an earlier version of this article. We are also grateful to Carita Paradis (Lund University, Sweden) for her invaluable guidance throughout the preparation of the article. Thanks also go to Stina Ericsson (Linnaeus University, Sweden) and two anonymous reviewers for their constructive comments on the manuscript.

References

- Artstein, R. and M. Poesio. 2008. 'Inter-coder agreement for computational linguistics', *Computational Linguistics* 34 (4), pp. 555–96.
- Bakhtin, M. 1981. *The Dialogical Imagination*. Austin, Texas: University of Texas Press.
- Barney, J.B. and M.H. Hansen. 1994. 'Trustworthiness as a source of competitive advantage', *Strategic Management Journal* 15 (1), pp. 175–90.
- Bartalesi Lenzi, V., G. Moretti and R. Sprugnoli. 2012. 'CAT: the CELCT annotation tool', *Proceedings of the Eighth International Conference on Language Resources and Evaluation LREC 2012*, pp. 333–8. 23–25 May 2012, Istanbul. Accessed 12 August 2015 at: <http://www.lrec-conf.org/proceedings/lrec2012/index.html>

- Bednarek, M. 2006. *Evaluation in Media Discourse: Analysis of a Newspaper Corpus*. London and New York: Continuum International Publishing Group Ltd.
- Bednarek, M. 2008. *Emotion Talk across Corpora*. Houndmills, Basingstoke: Palgrave Macmillan.
- Bednarek, M. 2009. 'Language patterns and attitude', *Functions of Language* 16 (2), pp. 165–92.
- Benoit, W.L. 1997. 'Image repair discourse and crisis communication', *Public Relations Review* 23 (2), pp. 177–86.
- Bhatia, V. 2004. *Worlds of Written Discourse: A Genre-based View*. New York: Continuum.
- Biber, D. 2006. 'Stance in spoken and written university registers', *Journal of English for Academic Purposes* 5 (2), pp. 97–116.
- Biber, D. and E. Finegan. 1988. 'Adverbial stance types in English', *Discourse Processes* 11 (1), pp. 1–34.
- Biber, D. and E. Finegan. 1989. 'Styles of stance in English: lexical and grammatical marking of evidentiality and affect', *Text-Interdisciplinary Journal for the Study of Discourse* 9 (1), pp. 93–124.
- Breeze, R. 2012. 'Legitimation in corporate discourse: oil corporations after deepwater horizon', *Discourse and Society* 23 (1), pp. 3–18.
- Breeze, R. 2013. *Corporate Discourse*. London and New York: Bloomsbury Academic.
- Camiciottoli, B.C. 2013. *Rhetoric in Financial Discourse: A Linguistic Analysis of ICT-mediated Disclosure Genres*. Amsterdam and New York: Rodopi.
- Carretero, M. and M. Taboada. 2014. 'Graduation within the scope of attitude in English and Spanish consumer reviews of books and movies' in G. Thompson and L. Alba-Juez (eds) *Evaluation in Context*, pp. 221–39. Amsterdam and Philadelphia: John Benjamins.
- Cohen, J. 1960. 'A coefficient of agreement for nominal scales', *Educational and Psychological Measurement* 20 (1), pp. 37–46.
- Conrad, S. and D. Biber. 2000. 'Adverbial marking of stance in speech and writing' in S. Hunston and G. Thompson (eds) *Evaluation in Text: Authorial Stance and the Construction of Discourse*, pp. 56–73. Oxford: Oxford University Press.
- Cook, K.S. 2001. *Trust in Society*. New York: Russell Sage Foundation.
- Di Eugenio, B. and M. Glass. 2004. 'The kappa statistic: a second look', *Computational Linguistics* 30 (1), pp. 95–101.
- Eggins, S. 2004. *An Introduction to Systemic Functional Linguistics*. London and New York: Continuum International Publishing Group.

- Fuoli, M. 2012. 'Assessing social responsibility: a quantitative analysis of APPRAISAL in BP's and Ikea's social reports', *Discourse and Communication* 6 (1), pp. 55–81.
- Fuoli, M. and C. Paradis. 2014. 'A model of trust–repair discourse', *Journal of Pragmatics* 74, pp. 52–69.
- García-Marzá, D. 2005. 'Trust and dialogue: theoretical approaches to ethics auditing', *Journal of Business Ethics* 57 (3), pp. 209–19.
- Garzone, G. 2005. 'Letters to shareholders and chairman's statements: textual variability and generic integrity' in P. Gillaerts and M. Gotti (eds) *Genre Variation in Business Letters*, pp. 179–204. Bern: Peter Lang.
- Gillaerts, P. and F. Van de Velde. 2011. 'Metadiscourse on the move: the CEO's letter revisited' in G. Grazone and M. Gotti (eds) *Discourse, Communication and the Enterprise: Genres and Trends*, pp. 151–68. Bern: Peter Lang.
- Halliday, M. 1994. *An Introduction to Functional Grammar*. London: Edward Arnold.
- Harlow, W.F., B.C. Brantley and R.M. Harlow. 2011. 'BP initial image repair strategies after the Deepwater Horizon spill', *Public Relations Review* 37 (1), pp. 80–83.
- Hommerberg, C. and A. Don. 2015. 'Appraisal and the language of wine appreciation: a critical discussion of the potential of the Appraisal framework as a tool to analyse specialised genres', *Functions of Language* 22 (2), pp. 161–91.
- Hunston, S. 2011. *Corpus Approaches to Evaluation: Phraseology and Evaluative Language*. New York and London: Routledge.
- Hunston, S. and J. Sinclair. 2000. 'A local grammar of evaluation' in S. Hunston and G. Thompson (eds) *Evaluation in Text: Authorial Stance and the Construction of Discourse*, pp. 74–101. Oxford: Oxford University Press.
- Hunston, S. and G. Thompson. 2000. *Evaluation in Text: Authorial Stance and the Construction of Discourse*. Oxford: Oxford University Press.
- Hyland, K. 1998. 'Exploring corporate rhetoric: metadiscourse in the CEO's letter', *Journal of Business Communication* 35 (2), pp. 224–44.
- Hyland, K. 2005. *Metadiscourse: Exploring Interaction in Writing*. London and New York: Continuum.
- Ingenhoff, D. and K. Sommer. 2010. 'Trust in companies and in CEOs: a comparative study of the main influences', *Journal of Business Ethics* 95 (3), pp. 339–55.
- Kaltenbacher, M. 2006. 'Culture related linguistic differences in tourist websites: the emotive and the factual—a corpus analysis within the

- framework of APPRAISAL’, in G. Thompson and S. Hunston (eds) *System and Corpus: Exploring Connections*, pp. 269–92. London: Equinox.
- Krippendorff, K. 2004. *Content Analysis: An Introduction to its Methodology*. Thousand Oaks, London and New Delhi: Sage.
- Linell, P. and I. Markova. 2013. *Dialogical Approaches to Trust in Communication*. Charlotte, North Carolina: Information Age Publishing.
- Lipovsky, C. 2008. ‘Constructing affiliation and solidarity in job interviews’, *Discourse and Communication* 2 (4), pp. 411–32.
- Lipovsky, C. 2013. ‘Negotiating one’s expertise through appraisal in CVs’, *Linguistics and the Human Sciences* 8 (3), pp. 307–33.
- Mackay, J. and J. Parkinson. 2009. “‘My very own mission impossible’: an APPRAISAL analysis of student teacher reflections on a design and technology project”, *Text and Talk* 29 (6), pp. 729–53.
- Marková, I. and A. Gillespie. 2008. *Trust and Distrust: Sociocultural Perspectives*. Charlotte, North Carolina: Information Age Publishing.
- Martin, J. and P. White. 2005. *The Language of Evaluation: Appraisal in English*. London and New York: Palgrave Macmillan.
- Mayer, R., J. Davis and F. Schoorman. 1995. ‘An integrative model of organizational trust’, *Academy of Management Review* 20 (3), pp. 709–34.
- Muralidharan, S., K. Dillistone and J.-H. Shin. 2011. ‘The gulf coast oil spill: extending the theory of image restoration discourse to the realm of social media and beyond petroleum’, *Public Relations Review* 37 (3), pp. 226–32.
- Murphy, A.C. 2013. ‘On “true” portraits of letters to shareholders – and the importance of phraseological analysis’, *International Journal of Corpus Linguistics* 18 (1), pp. 57–82.
- O’Connor, E.O. 2011. ‘Organizational apologies: BP as a case study’, *Vanderbilt Law Review* 64, pp. 1959–91.
- O’Donnell, M. 2008. ‘Demonstration of the UAM corpus tool for text and image annotation’, *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics on Human Language Technologies: Demo Session*, pp. 13–16. Association for Computational Linguistics.
- O’Donnell, M. 2014. ‘Exploring identity through appraisal analysis: a corpus annotation methodology’, *Linguistics and the Human Sciences* 9 (1), pp. 95–116.
- Pirson, M. and D. Malhotra. 2011. ‘Foundations of organizational trust: what matters to different stakeholders?’, *Organization Science* 22 (4), pp. 1087–104.

- Poppo, L. and D.J. Schepker. 2010. 'Repairing public trust in organizations', *Corporate Reputation Review* 13 (2), pp. 124–41.
- Pounds, G. 2010. 'Attitude and subjectivity in Italian and British hard-news reporting: the construction of a culture-specific reporter voice', *Discourse Studies* 12 (1), pp. 106–37.
- Pounds, G. 2011. "'This property offers much character and charm": evaluation in the discourse of online property advertising', *Text and Talk* 31 (2), pp. 195–220.
- Read, J. and J. Carroll. 2012. 'Annotating expressions of appraisal in English', *Language Resources and Evaluation* 46 (3), pp. 421–47. (First published online in December 2010.)
- Robertson, C. and C. Krauss. 2010. 'Gulf spill is the largest of its kind, scientists say', *The New York Times* online. Accessed 14 February 2014 at: <http://www.nytimes.com/>
- Ryshina-Pankova, M. 2014. 'Exploring argumentation in course-related blogs through engagement' in G. Thompson and L. Alba-Juez (eds) *Evaluation in Context*, pp. 281–302. Amsterdam and Philadelphia: John Benjamins.
- Santamaría-García, C. 2014. 'Evaluative discourse and politeness in university students' communication through social networking sites' in G. Thompson and L. Alba-Juez (eds) *Evaluation in Context*, pp. 387–411. Amsterdam and Philadelphia: John Benjamins.
- Schultz, F., J. Kleinnijenhuis, D. Oegema, S. Utz and W. van Atteveldt. 2012. 'Strategic framing in the BP crisis: a semantic network analysis of associative frames', *Public Relations Review* 38 (1), pp. 97–107.
- Shogren, E. 2011. 'BP: a textbook example of how not to handle PR', Available online at: <http://www.npr.org>
- Spooren, W. and L. Degand. 2010. 'Coding coherence relations: reliability and validity', *Corpus Linguistics and Linguistic Theory* 6 (2), pp. 241–66.
- Taboada, M. and M. Carretero. 2012 [2010]. 'Contrastive analyses of evaluation in text: key issues in the design of an annotation system for attitude applicable to consumer reviews in English and Spanish', *Linguistics and the Human Sciences* 6 (1–3), pp. 275–95. DOI: 10.1558/lhs.v6i1-3.275.
- Taboada, M., J. Brooke, M. Tofiloski, K. Voll and M. Stede. 2011. 'Lexicon-based methods for sentiment analysis', *Computational Linguistics* 37 (2), pp. 267–307.
- Thompson, G. and L. Alba-Juez. 2014. *Evaluation in Context*. Amsterdam and Philadelphia: John Benjamins.
- Voormann, H. and U. Gut. 2008. 'Agile corpus creation', *Corpus Linguistics and Linguistic Theory* 4 (2), pp. 235–51.

- Webb, T. 2010. 'BP boss admits job on the line over Gulf oil spill', The Guardian online. Accessed 14 February 2014 at: <http://www.theguardian.com/business/>
- White, P. 2003. 'Beyond modality and hedging: a dialogic view of the language of intersubjective stance', *Text-Interdisciplinary Journal for the Study of Discourse* 23 (2), pp. 259–84.
- White, P.R. 2012. 'Exploring the axiological workings of "reporter voice" news stories—attribution and attitudinal positioning', *Discourse, Context and Media* 1 (2–3), pp. 57–67.
- Wickman, C. 2013. 'Rhetorical framing in corporate press releases: the case of British petroleum and the gulf oil spill', *Environmental Communication* 8 (1), pp. 3–20.
- Wiebe, J., T. Wilson and C. Cardie. 2005. 'Annotating expressions of opinions and emotions in language', *Language Resources and Evaluation* 39 (2), pp. 165–210.

Appendix A

The annotation manual can be found in the form of an appendix as supplementary material to the web-based version of this paper.

Your short guide to the EUP Journals
Blog <http://euppublishingblog.com/>

***A forum for discussions relating to
[Edinburgh University Press Journals](#)***



EDINBURGH
University Press

1. The primary goal of the EUP Journals Blog

To aid discovery of authors, articles, research, multimedia and reviews published in Journals, and as a consequence contribute to increasing traffic, usage and citations of journal content.

2. Audience

Blog posts are written for an educated, popular and academic audience within EUP Journals' publishing fields.

3. Content criteria - your ideas for posts

We prioritize posts that will feature highly in search rankings, that are shareable and that will drive readers to your article on the EUP site.

4. Word count, style, and formatting

- Flexible length, however typical posts range 70-600 words.
- Related images and media files are encouraged.
- No heavy restrictions to the style or format of the post, but it should best reflect the content and topic discussed.

5. Linking policy

- Links to external blogs and websites that are related to the author, subject matter and to EUP publishing fields are encouraged, e.g. to related blog posts

6. Submit your post

Submit to ruth.allison@eup.ed.ac.uk

If you'd like to be a regular contributor, then we can set you up as an author so you can create, edit, publish, and delete your *own* posts, as well as upload files and images.

7. Republishing/repurposing

Posts may be re-used and re-purposed on other websites and blogs, but a minimum 2 week waiting period is suggested, and an acknowledgement and link to the original post on the EUP blog is requested.

8. Items to accompany post

- A short biography (ideally 25 words or less, but up to 40 words)
- A photo/headshot image of the author(s) if possible.
- Any relevant, thematic images or accompanying media (podcasts, video, graphics and photographs), provided copyright and permission to republish has been obtained.
- Files should be high resolution and a maximum of 1GB
- Permitted file types: *jpg, jpeg, png, gif, pdf, doc, ppt, odt, pptx, docx, pps, ppsx, xls, xlsx, key, mp3, m4a, wav, ogg, zip, ogv, mp4, m4v, mov, wmv, avi, mpg, 3gp, 3g2*.