



Uncovering organisational pride and psychological safety from glassdoor reviews

Ali Septiandri¹, Sanja Šćepanović^{1,2*}, Marios Constantinides^{1,5}, Licia Capra⁴ and Daniele Quercia^{1,3}

Handling Editor: Yelena Mejova

*Correspondence:

sanja.scepanovic@nokia-bell-labs.com

¹Nokia Bell Labs, Cambridge, UK

²University of Oxford, Oxford, UK

Full list of author information is available at the end of the article

Abstract

Understanding employee experiences and attitudes is crucial for promoting a positive work environment, and enhancing engagement, satisfaction, productivity, and innovation. Organisational culture, represented by constructs like organisational pride (OP) and psychological safety (PS), captures these experiences. OP reflects employees' emotional attachment and dedication to an organisation, while PS embodies the collective perception of safety for risk-taking and open communication. Together, these constructs offer a rich perspective, providing a top-to-bottom view of employee experiences and attitudes. To evaluate OP and PS, we developed a deep-learning framework utilising language embeddings and applied it on 430,000 employee reviews spanning 2008 to 2020, encompassing 318 major U.S. companies. Our analysis revealed significant sector-specific variations in these constructs, highlighting the unique strengths and challenges within each sector. We found that OP, which applies to the company as a whole, is high in utility and energy sectors but low in consumer and communications, while PS, which reflects the team level, is high in IT and low in communications and healthcare. Our automatic rationalisation of these organisational constructs paves the way towards the development of automated psychometric assessments at the workplace.

Keywords: Organisational pride; Psychological safety; Workplace; Deep learning; Automatic psychometric assessment; Companies

1 Introduction

“Psychological safety isn’t about being nice. It’s about giving candid feedback, openly admitting mistakes, and learning from each other” [1] — Amy Edmondson

Understanding employees' work experiences offers valuable insights into an organization's culture, contributing to employee engagement, job satisfaction, and overall productivity [2]. A positive workplace environment can enhance innovation and creativity by encouraging employees to share ideas and take initiative [3]. To maintain a supportive culture, many organizations proactively address workplace challenges through open communication channels, which helps improve employee well-being and reduce stress [4]. However, companies today face significant pressures from global disruptions such as the COVID-19 pandemic [5], geopolitical conflicts [6], and shifts in the talent market [7],

© The Author(s) 2025. **Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

which impact workforce dynamics and raise the need for strong organizational cultures. Consequently, a deep understanding of employees' perceptions—from their attitudes towards company strategy to team-level culture—allows organizations to strengthen their public reputation, align strategic goals, and foster resilience in a rapidly evolving business landscape.

Companies commonly rely on psychological constructs to measure various aspects of organisational culture. Organisational culture encompasses the shared values, beliefs, and practices of a company's employees [8], and it emerges from the interplay of top-down expectations and bottom-up norms [9]. Given its multidimensional nature [10], various aspects of organisational culture are captured, ranging from leadership, resolving conflicts and negotiating, independence, and freedom to make decisions, to training and teaching others [8]. Two widely discussed aspects of organisational culture are *organisational pride* (OP) and *psychological safety* (PS). Both OP and PS are significant aspects of organisational culture that shape employee values, beliefs, and practices [11–14]. Employees' pride in their workplace (OP) often indicates a positive and healthy organisational culture [11, 12, 15, 16]. OP reflects employees' emotional and attitudinal attachment to their organisation, influencing employee behaviour, creativity, turnover, and commitment. It serves as both an antecedent and a consequence of corporate success [11, 15, 17]. On the other hand, PS pertains to employees' shared perception that it is safe to take interpersonal risks, express ideas, and voice opinions without fear of negative consequences or judgement. It enables a supportive work environment where employees feel comfortable being themselves and contributing their perspectives [1, 13]. OP is acknowledged as a key component in shaping organisational culture, yet its impact has not been extensively quantified on a large scale. Our focus on OP stems from substantial qualitative evidence suggesting its significant role in company culture. However, OP alone does not fully capture the grassroots aspects of organisational culture. Indeed, organisational culture is not a mere aggregation of elements; it evolves from the intricate interaction of various factors [8], including both OP and PS. PS plays a crucial role in building organisational resilience and fostering a culture where employees can navigate uncertainty and challenges effectively [1, 14, 18–20]. Hence, we integrate PS with OP to comprehensively understand organisational culture's emergence from this dynamic and multifaceted interaction [12, 15, 21, 22]. Traditionally, these constructs have been assessed through surveys, which can be costly [23] and limited to self-selected participants, posing challenges to generalisability [24–26]. In this work, we developed automated methods for assessing aspects of organisational culture, and, in so doing, we made three contributions:

1. We implemented a deep-learning Natural Language Processing (NLP) framework that accurately captures OP and PS at company-level (§4).
2. We applied this framework to over 430,000 publicly available Glassdoor employee reviews, and validated it internally in terms of keywords expressed in high OP/PS companies, and externally in terms of associations of company OP/PS scores with the ratings. We also surfaced variations in OP/PS levels across different industry sectors from 318 of S&P 500 companies. Lastly, we uncovered that despite being correlated aspects of organizational culture, OP and PS each independently contribute to predicting employee satisfaction (§5).

Our automatic operationalisation of these two pivotal organisational constructs suggests avenues for the development of *automated psychometric assessment*, allowing for

analysing organisational culture across sectors and countries, in addition to exploring the interactions between various organisational constructs. Ultimately, the results generated from this process could contribute to the creation of more satisfying and productive work environments (§6).

2 Background and related work

First, we provide the reader with a background on psychological constructs commonly used to assess organisational culture, as discussed in the Organisational Behaviour Literature, with a focus on our two constructs of interest: *organisational pride* and *psychological safety*. Following that, we place our work within the broader context of existing literature on Natural Language Processing (NLP) research in Computational Social Science, which attempted at capturing aspects of organisational culture on a large scale.

2.1 Organisational behaviour literature

Organisational Culture (OC) refers to the shared values, beliefs, assumptions, norms, and behaviours that collectively shape the identity and character of an organisation [27]. It significantly impacts how employees align their actions with the organisation's goals and expectations, ultimately influencing work attitudes and overall performance [28].

2.1.1 *Organisational pride (OP)*

It refers to the emotional and attitudinal attachment that employees develop towards their organisation [15]. It encompasses both the emotional attachment triggered by perceiving a successful organisational event and the lasting cognitive attitude resulting from a general perception of the organisation [11, 29, 30]. As a psychological construct, OP serves as an indicator of both internal and external organisational reputation and can act as both a driver and an outcome of corporate success [11, 15, 17, 31].

Previous research has identified various factors that contribute to the development of employees' OP. Companies that provide job autonomy, job variety, rewards, recognition, and demonstrate success, social responsibility, and internal sustainability are more likely to evoke a sense of pride among their employees [17, 29]. In turn, OP influences employee behaviour in several positive ways. It leads to favourable job perceptions, customer-oriented behaviour, enhanced creativity, reduced turnover, increased commitment, and pro-environmental behaviours [11, 12, 31–36]. Furthermore, OP is positively correlated with organisational attractiveness and job pursuit intentions among prospective employees [12, 32–34]. It has also been found to mediate the relationship between corporate social responsibility (CSR) and organisational citizenship behaviour (OCB) [37, 38].

The measurement and monitoring of OP are crucial for organisations due to its impact on company reputation, employee outcomes, and turnover [33, 39, 40]. However, there is still a need for a comprehensive examination of the conceptual scope and structural properties of OP. Previous empirical studies have often relied on small survey samples from a limited number of organisations, raising concerns about unreliability [33, 34, 39–42]. While these studies have explored various antecedents and consequences of OP, such as length of affiliation, work-related emotions, and corporate social responsibility, the limited availability of data has hindered the examination of these consequences at higher-order levels of analysis [12, 33, 34, 40, 41, 43]. Consequently, potential macro-level influences and outcomes across industries and geographical regions have remained unexplored [33, 34, 44, 45].

2.1.2 Psychological safety (PS)

It encompasses the shared belief among employees that they can take interpersonal risks within the workplace [1, 13, 19, 20, 46]. This includes freely expressing ideas, voicing opinions, and being authentic, without the fear of facing negative consequences or judgement, encouraging employee engagement and the integration of diverse perspectives [13].

Traditionally, psychological safety has primarily been assessed at the team level using measurement items developed by Edmondson [1]. While a few studies have examined the measurement of psychological safety at the organisational level [14, 47, 48], they typically modified the team-level items by substituting “team” with “organization” [49]. However, reaching a consensus on psychological safety can be challenging in larger organisations, where employees may have diverse experiences within their respective teams [50, 51]—leadership style and team dynamics shape the perception of psychological safety [49]. For this reason, previous research indicates that psychological safety has a more substantial impact at the team level compared to the organisational level. Additionally, it has been suggested that smaller organisations with frequent collaboration and strong corporate cultures are more likely to foster a climate of psychological safety, although further empirical evidence is needed [49]. Furthermore, studies have explored potential adverse outcomes associated with excessively high levels of psychological safety within teams, such as unethical behaviours and reduced organisational performance [14, 46].

Researchers have traditionally relied on qualitative interviews to gain insights into psychological safety at both the team and individual levels. However, in order to evaluate levels of psychological safety within teams or organisations, the use of observational methods has been proposed as an alternative approach [49]. Despite its potential to provide valuable insights into the dynamics of psychological safety, including macro-level influences and outcomes across various industries and geographical regions, this complementary approach, similar to OP, has been scarce.

2.2 NLP research in computational social science

Das Swain et al. [8] introduced the use of natural language processing techniques on Glassdoor review data to automatically evaluate OC using psychological constructs similar to those in our study, such as the Organisation Cultural Inventory [52], the Organisation Culture Profile [53], Hofstede’s Organisation Culture Questionnaire [54], and the Organisation Culture Survey [55]. They compared word-level vector representations of descriptions of organisational constructs with embeddings of Glassdoor reviews and successfully demonstrated the face and construct validity of their language model-based assessment of organisational culture. Additionally, they provided empirical evidence, based on Fortune 500 companies, of the connection between OC and job performance. Similarly, Sen et al. [56] and Ceccarini et al. [57] utilised Glassdoor reviews to examine Internal Sustainability Efforts (ISEs) at the organisational level within S&P 500 companies. They confirmed the positive associations between ISEs and stock growth. Furthermore, Lee [58] quantified internal corporate social responsibility (CSR) and discovered a positive impact on Organisational Citizenship Behaviour (OCB), which encompasses voluntary actions by employees that go beyond their formal job requirements. In line with these findings, Šćepanović et al. [59] analysed thousands of Glassdoor reviews and uncovered that different types of stress contribute differently to companies’ success. Specifically, companies whose employees tend to perceive stress positively (i.e., describing it in terms of a high-growth and

collaborative work environment) enjoy higher stock growth compared to those whose employees perceive stress negatively.

At the time of writing, there is a noticeable gap in comprehensive empirical research that explores OP and PS, two constructs that hold significant importance in understanding employee experiences in the workplace. What sets them apart from other constructs like OC, CSR, and ISEs is their assessment using established psychological measures [1, 41]. Unlike these constructs, which have been measured on a large scale but with adapted definitions and descriptions, OP and PS metrics can rely on validated measures. Hence, they presents valuable opportunities to develop empirical methodologies that quantify organisational constructs based on validated psychological measures, while encompassing comprehensive aspects of the workplace, from the global organisation level to the team level, using online data.

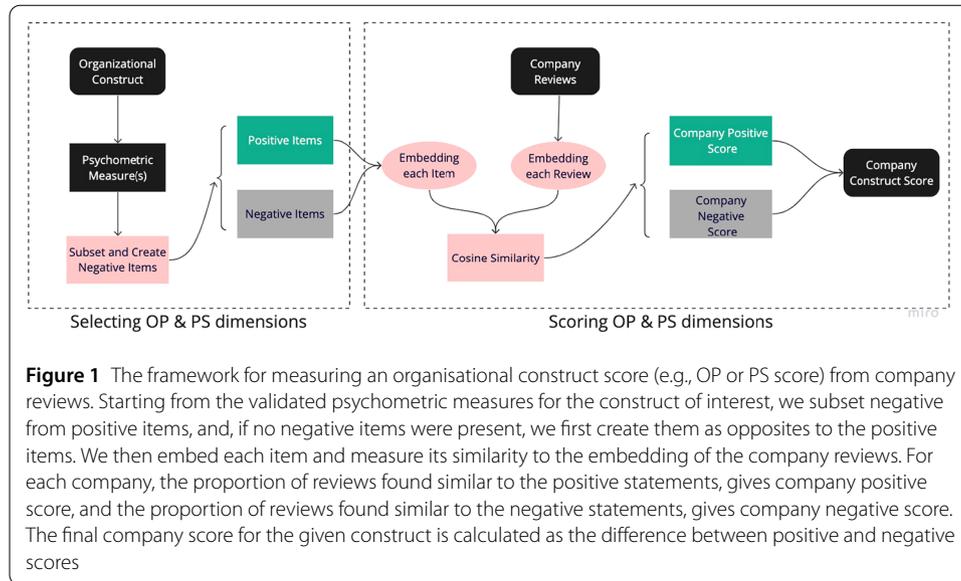
3 Datasets

3.1 Description and statistics

We collected data from Glassdoor, a well-known platform for company reviews, where individuals, including current and former employees, share their reviews and experiences regarding various aspects of their corporate life. The platform ensures high-quality reviews through three key measures. First, it employs both proprietary automated tools and manual moderation to oversee content, while also allowing users to flag inappropriate posts. This dual approach helps reduce the risk of fake reviews, such as those submitted under coercion from employers to post positive feedback. Second, users are required to contribute at least one review before gaining access to others' content. This policy fosters a more balanced range of opinions and helps mitigate *non-response bias*, where the views of respondents might significantly differ from those who choose not to participate. Lastly, the platform limits each employee to a single review per company per year. This cap prevents individuals from dominating the review landscape, addressing *sampling bias*, which occurs when certain perspectives are overrepresented in the data. As discussed in the related work section, this data collection has been previously employed in studies examining other organisational constructs on a large scale [8, 56, 59]. However, we acknowledge that Glassdoor reviews, like any source of self-reported data, are not free from bias and should be seen as a complementary source rather than a replacement for survey data.

In addition to being efficient and low-cost, a potentially unique advantage of using Glassdoor reviews is their relative independence from organizational-level psychological safety. In traditional surveys conducted within organizations, employees may feel pressured to inflate their assessments of organizational culture or job satisfaction due to fear of repercussions, especially in environments where psychological safety is low. This bias may be less prevalent in anonymous Glassdoor reviews, offering a distinct perspective on employee experiences.

The dataset used in this study consisted of Glassdoor reviews spanning a twelve-year period from 2008 to 2020. Each review in the dataset contained various components, including a title, a 'pro' section highlighting positive aspects of the company, a 'con' section discussing negative aspects, four ratings (ranging from 0 to 5) evaluating the company's *career opportunities*, *work-life balance*, *culture*, and *management*, as well as an *overall rating* of the company. Some reviewers also provided their location, enabling us to identify the states for a subset of the reviews. To ensure the reliability and robustness of our text



processing method, we applied specific criteria for dataset selection. We retained companies that had a minimum of 100 reviews, resulting in a final dataset of 430,788 reviews from 318 US-based companies, which represents 99.1% of the original dataset. Notably, 80% of these companies were listed in the S&P 500 index. To assess the representativeness of these 318 companies, we examined the distribution of industry sectors, which was found to be comparable to that of the S&P 500 companies. Additionally, we computed the correlation between the number of reviews per state and the total number of employment positions per state, and found that they are highly correlated ($r = 0.97, p < 0.001$) as shown in Fig. 7 in the [Appendix](#).

4 Framework for automatically assessing OP and PS

We employed a mixed-method approach to develop and validate an automated, language-based assessment of OP and PS, similar to Sen et al. [56]’s framework for measuring ISEs (Fig. 1). We identified key language features related to OP and PS by reviewing the organisational behaviour and psychometrics literature for existing assessments. We separated the positive and negative items and created opposing negative items when they were not initially available. We then employed the sentence-level BERT (SBERT) model [60] to generate embeddings for each item and subsequently calculated their similarity with the embeddings of company reviews. Specifically, we used `sentence-t5-xl` for its effectiveness in similarity comparisons while still being computationally efficient [61]. It has been trained on the Massive Text Embedding Benchmark, an evaluation benchmark across 58 datasets and 112 languages for embedding for tasks such as classification, clustering, and semantic textual similarity [62]. Adhering to the recommendations by Reimers and Gurevych [60], we used SBERT with its default settings and specifically set the maximum sequence length to 256 and employed token-based pooling.

SBERT is adept at capturing and processing the semantics of noisy reviews. This deep learning-based model is designed to interpret and process text that contains irregularities such as typos and grammatically inconsistent language by using subword embeddings [63, 64]. For example, the word “learning” could be divided into subwords like “learn”

and “ing”. The model then learns vector representations for each subword. This allows the model to infer meanings of words it has not seen before by understanding their component subwords. Hence, unlike traditional natural language processing techniques, SBERT allows for bypassing conventional preprocessing steps like stemming or lemmatization.

To assess the company’s positive score, we determined the proportion of reviews aligning with positive statements. Similarly, the negative score was determined by examining the proportion of reviews aligning with negative statements. Finally, we obtained the overall company score for the specific construct by taking the difference between the positive and negative scores.

4.1 Step 1: selecting the OP and PS dimensions

4.1.1 Curating positive OP and PS items

To begin with, we resorted to validated psychological scales to construct items for OP and PS. Next, we explain the curation process.

Organisational Pride. In Psychology, three widely used measures for organisational pride are developed by Jones [41], Turban et al. [12], and Cable and Turban [32]. These three measures have demonstrated high reliability (i.e., with reliability coefficients of at least .80). Considering that Glassdoor reviews were mostly expressed in the past tense, we deliberately omitted items that included future tense language. As a result, we chose four specific items from the measure developed by Jones [41] (OP1 to OP4 in Table 1).

Psychological Safety. To measure PS at both the company and team levels, we utilised the items from Edmondson [1], which is the most commonly used measure in the literature on organisational behaviour and psychometrics [49]. Note that we needed to paraphrase one item from “No one on this team would deliberately act in a way that undermines my efforts.” to “Everyone on this team would act in a supportive manner towards my efforts.” to avoid a double negative, which would otherwise introduce more noise when extracting embeddings from the deep learning model (PS1 to PS4 in Table 1).

4.1.2 Creating negative OP and PS items

Glassdoor requires users to separate their reviews into pros and cons sections. This helps to ensure a fair evaluation that covers both the positive and negative aspects of a company, allowing users to make informed decisions about job opportunities and company cultures. Therefore, this separation of pros and cons allowed us to align the items with their corresponding sentiments in the reviews, leading to a more precise evaluation of similarity. The original OP measure includes only positive items to assess OP. We manually developed additional items reflecting opposite valence to those in the original OP measure. These items assess ‘lack of organisational pride’ (LOP), showing absence of positive attachment and identification with the organisation (LOP1 to LOP4 in Table 1). Since the original measure includes both positive and negative items, for PS, we separated existing negative items from positive ones to gauge ‘lack of psychological safety’ (LPS), which involves shared perception of fear about potential negative consequences for speaking up (LPS1 to LPS3 in Table 1). Note that because the original PS scale has seven items total, we ended up with four positive and three negative items.

4.1.3 Final OP and PS items

Combining the previous two steps, we were left with a total of 15 items: 8 for measuring OP construct, and 7 for measuring PS construct (Table 1).

Table 1 Final selected or adapted items from psychometric measures for measuring OP and PS

Abbr.	Abbr.
Organisational Pride (OP) [41]	Psychological Safety (PS) [1]
OP1 I am proud to work for this firm.	PS1 Members of this team are able to bring up problems and tough issues.
OP2 People respect what this firm does.	PS2 It is safe to take a risk on this team.
OP3 I am proud to be associated with this firm.	PS3 Everyone on this team would act in a supportive manner towards my efforts.
OP4 I am proud of what this firm accomplishes.	PS4 Working with members of this team, my unique skills and talents are valued and utilized.
Lack of OP (LOP) adapted from [41]	Lack of PS (LPS) [1]
LOP1 I am displeased to work for this firm.	LPS1 If you make a mistake on this team, it is often held against you.
LOP2 People look down upon what this firm does.	LPS2 People on this team sometimes reject others for being different.
LOP3 I am displeased to be associated with this firm.	LPS3 It is difficult to ask other members of this team for help.
LOP4 I am dissatisfied with how this firm accomplishes what it does.	

4.2 Step 2: scoring the OP and PS dimensions

4.2.1 Calculating separate item scores

We examined the 15 items at the company-level (denoted as u) to quantify organisational behaviour across multiple companies. To achieve this, we defined the score $s(u, i)$ of the i th item for company u by calculating the fraction of reviews mentioning i within company u .

$$s(u, i) = \frac{\sum_{p \in R(u)} sim_t(v_p, v_i)}{|R(u)|} \tag{1}$$

Here, $R(u)$ represents the collection of reviews for company u , v_i denotes the SBERT (Sentence-BERT) vector corresponding to item i , and $sim_t(v_p, v_i)$ refers to the *thresholded* SBERT similarity score [60] between the SBERT vector of review p and the SBERT vector associated with item i . We applied this similarity thresholding technique to account for noise and potential variations in review expressions. For instances of missing data, specifically reviews without accompanying text or which are too short, this similarity thresholding effectively filtered out such instances since they yielded low similarity scores and were thus excluded from our analysis. Items with similarity scores below a certain threshold are considered dissimilar and assigned a score of 0, while those above the threshold retain their similarity scores. To be more specific, the definition of $sim_t(v_p, v_i)$ is as follows:

$$sim_t(v_p, v_i) = \begin{cases} sim(v_p, v_i), & \text{if } sim(v_p, v_i) > \max(\theta, p_{50}(i)) \\ 0, & \text{otherwise} \end{cases} \tag{2}$$

We chose the threshold as the maximum value between θ and $p_{50}(i)$, where θ is the average value of the median similarities (one for pros and one for cons), and $p_{50}(i)$ is the median similarities for item i .

Additionally, it is important to highlight that we empirically identified a connection between OP/PS and the pros, as well as between LOP/LPS and the cons, using similarity scores. Pairing the cons with OP and PS (or pros with LOP and LPS) was found to negatively affect the performance of our framework during the validation phase of our mixed-method approach. Specifically, when cons were paired with OP/PS, negative statements about the company or teams were sometimes misinterpreted as similar to our items,

thereby contributing to the construct scores despite their opposite meanings. This issue arises because the similarity measure struggles to handle negations effectively. To mitigate such inconsistencies, we naturally aligned positive construct items with pros and negative ones with cons.

4.2.2 Combining item scores

To derive composite scores for the constructs, namely OP and PS, we aggregated the scores of individual items. This aggregation enabled us to create more comprehensive representations of these underlying constructs. For OP, we calculated \bar{OP}_u for company u by averaging the scores of four items, as shown in Equation (3):

$$\bar{OP}_u = \frac{s(u, OP1) + s(u, OP2) + s(u, OP3) + s(u, OP4)}{4} \tag{3}$$

The same approach was then applied to calculate composite scores for PS, LOP, and LPS. Subsequently, we standardised the values for each category by subtracting the mean and dividing it by the standard deviation, yielding the z-score. This process is illustrated in the following equation.

$$z_{OP,u} = \frac{\bar{OP}_u - \mu(\bar{OP})}{\sigma(\bar{OP})} \tag{4}$$

$$z_{PS,u} = \frac{\bar{PS}_u - \mu(\bar{PS})}{\sigma(\bar{PS})} \tag{5}$$

$$z_{LOP,u} = \frac{L\bar{OP}_u - \mu(L\bar{OP})}{\sigma(L\bar{OP})} \tag{6}$$

$$z_{LPS,u} = \frac{L\bar{PS}_u - \mu(L\bar{PS})}{\sigma(L\bar{PS})} \tag{7}$$

Finally, the OP and PS scores for company u were obtained by subtracting the z-score of the pro construct category from the con construct category.

$$OP_{score,u} = z_{OP,u} - z_{LOP,u} \tag{8}$$

$$PS_{score,u} = z_{PS,u} - z_{LPS,u} \tag{9}$$

4.3 Step 3: validating the scoring of OP and PS

We focused on OP and PS since neither was previously studied and captured on a large scale. That means that there are no ground truth labels nor previous work to compare to. Hence, we assessed the effectiveness of our proposed method in two steps as follows. First, we implemented a baseline approach using extracted keywords from the OP and PS statements for the two construct measures. We automatically identified keywords and phrases from each construct item using the Rapid Automatic Keyword Extraction (RAKE) algorithm. This domain-independent method identifies key phrases—such as single keywords, bigrams and n-grams—by analyzing word frequency and co-occurrence patterns within the text as outlined in [65]. We then labelled a review as belonging to a certain construct, if we could find any of the phrases associated with that construct in the review, and scored the strength of the association based on the number of such phrases identified.

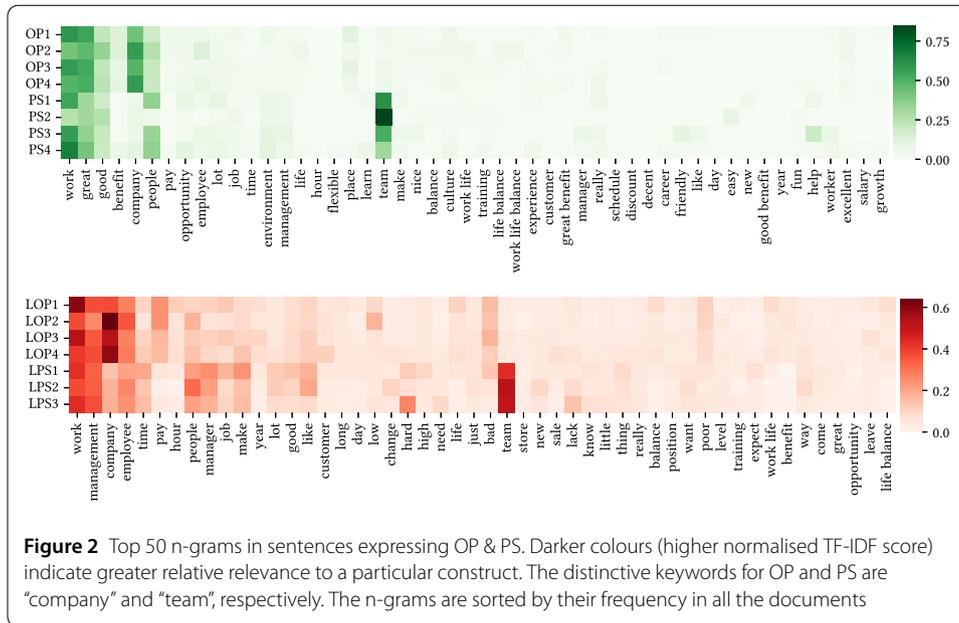
We randomly sampled 400 reviews: 200 from pros and 200 from cons. For pros, for each of the two positive constructs (OP and PS), we took 50 reviews with composite scores (Eq. (3)) above the 75th percentile (i.e., corresponding to a high construct value), and 100 with a composite score of 0 (i.e., not pertaining to any of the two constructs). Similarly, we sampled 200 reviews from cons. This stratified sampling ensured that we covered all the cases, i.e., reviews that were high or low in either of the constructs, and their combinations, or that did not express either of them. Three independent annotators then manually assessed the relevance of these reviews. Finally, we compared our proposed framework against the keyword-based method, as a lower bound baseline, on the set of manually labelled reviews. Additionally, we also benchmarked our method against predictions from GPT-4o-mini [66], a state-of-the-art Large Language Model [67], as a more competitive baseline.

4.4 Step 4: analysing OP and PS

To understand what characterises OP and PS, we adopted a linguistic approach, previously explored by Das Swain et al. [8], through which we identified the most frequent keywords, including 1, 2, 3, and 4-grams, from the reviews that our method deemed relevant for each construct. Since a review may pertain to multiple constructs, we divided the reviews into sentences. Next, we calculated the TF-IDF scores for the identified n-grams, with each document consisting of all shortlisted sentences for a specific construct. Finally, we ranked the keywords based on their TF-IDF scores, providing insights into the importance of certain terms for each construct as determined by our embedding-based method. This process was performed separately on the pro reviews paired with OP/PS and on the con reviews paired with LOP/LPS (as detailed in Sect. 4.1).

As previous literature suggests a positive correlation between both OP and PS and employee satisfaction [11, 13, 15, 17, 31, 49], we investigated the relationship between OP/PS scores and the company ratings through a linear regression analysis, while controlling for the number of reviews.

As we found that the distribution of OP and PS scores varies across different industry sectors (Fig. 4), we examined whether these two constructs relate differently to industry sectors. To determine whether the differences in the two scores across sectors are statistically significant, we conducted ANOVA followed by post-hoc analyses using Tukey's (honestly significant difference) HSD tests [68]. Note that we only analysed sectors with more than 10 companies to ensure the analysis possesses sufficient statistical power. A larger sample size within each sector allows for more reliable and generalisable findings, reduces the impact of outliers, and enables more nuanced subgroup analyses. We also conducted a qualitative analysis using the publicly available reviews. To identify reviews that encapsulate the predominant themes within each sector, we followed two steps that are commonly adopted in thematic analysis [69]. We first applied open coding to identify key concepts that emerged across multiple reviews; specifically, two of the authors read all the reviews that scored in the top 5% from OP/PS, and marked them with keywords that reflected the key concepts expressed in the text. They then used axial coding to identify relationships between the most frequent keywords, and select representative reviews. We summarise our qualitative analysis with respect to OP and PS. Any excerpts utilised to support the analysis were paraphrased to ensure anonymity and avoid the identification of specific reviews.



5 Results

We applied the framework (§4) on 430K employee reviews from 2008 to 2020 in 318 major U.S. companies, and showed that these two constructs are associated differently with organisational culture. Table 4 (Appendix) shows the summary statistics of all variables used in our analysis, and Table 7 (Appendix) shows the outcomes of our manual evaluation and the comparative analysis of our method versus the keyword- and LLM-based baselines. This table highlights the high accuracy of our approach and its clear advantage over the baseline method. Moreover, our method also outperformed or matched the LLM baseline, highlighting its state-of-the art performance for this task. In this section, we present our findings about the two constructs examined across companies and industries.

OP applies to a company as a whole, and it is high in utility and energy sectors and low in consumers and communications. When analysing the positive and negative construct items, which represent the pros and cons of review sections, we observed a consistent language pattern. Initially, we discovered that all items associated with both OP and PS shared common keywords such as *work*, *people*, and *employee* (Fig. 2). This outcome was expected since both aspects revolve around the employee experience in the workplace. However, OP-related items tended to focus more on the overall “company” aspect, indicated by words like *company*, *job*, and *benefit*. This finding aligns with the definitions of these concepts in organisational research, where it primarily encompasses employee experiences related to the organisation as a whole [15]. By exploring more specific keywords without the six most frequent keywords (*company*, *good*, *great*, *people*, *team*, and *work*) as shown in Fig. 3, we found that OP was now associated with (*good/great*) *benefits*, *culture*, *employees*, and (*good/great*) *place*. These findings are consistent with previous research indicating a positive relationship between OP and work-life balance [17], as well as the benefits of internal CSR practices for employees [70]. This suggests that our OP measure applies to a company as a whole. Examples of reviews with high and low OP/PS scores can be found in Table 2.

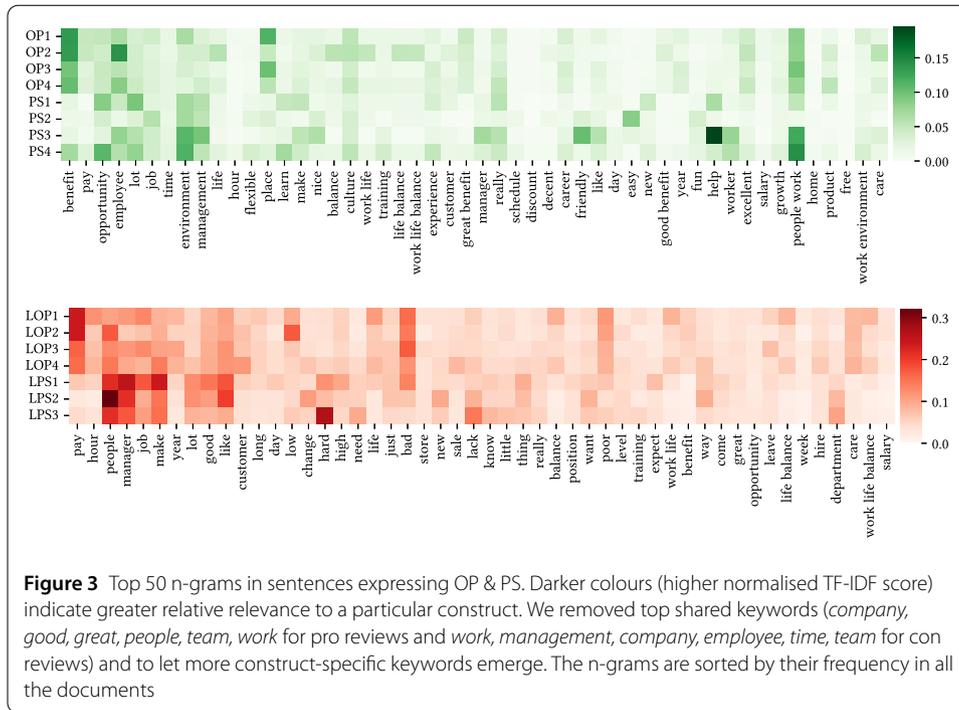
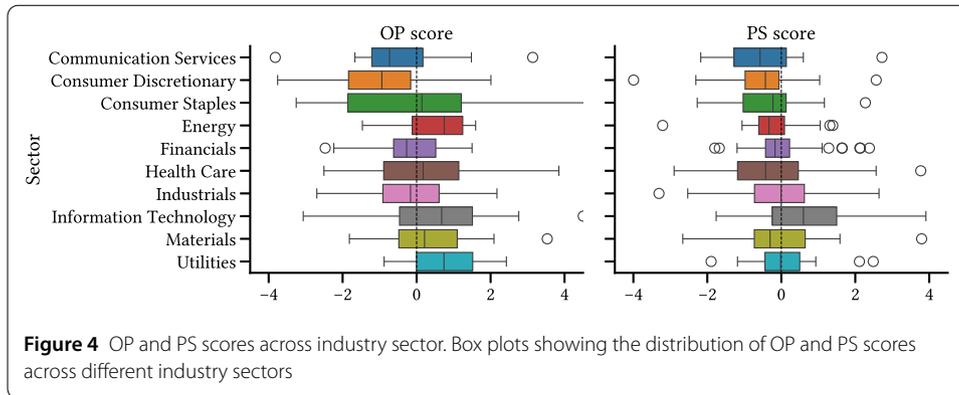


Table 2 Example reviews for high and low OP/PS scores

# High OP	# Low OP
1 Company that you are proud to work for.	1 One of the most unethical and corrupt companies I have ever worked for.
2 This company has some of the best and brightest talent and I am happy to work alongside of them.	2 Felt disjointed from the firm.
3 The company has an excellent reputation and I am proud to work here.	3 Constant lawsuits and bad press make one not want to admit to working there.
4 I love working for a company with such brand recognition.	4 One of the worst companies I have ever experienced.
# High PS	# Low PS
1 Excellent benefits, team members that are willing to assist with issues and questions.	1 If you make an error, they hold it against you.
2 Bunch of smart people working on difficult problems.	2 There's much confusion and miscommunication among teams.
3 Great team facing challenging environment. Strong commitment by all coworkers.	3 Lot of politics. If you end up being with a wrong team, your experience will be really bad.
4 Terrific people, strong team emphasis.	4 It is a team environment. Some team members are not as good as others.

OP across industry sectors. ANOVA indicated statistically significant differences in OP scores across sectors ($F = 4.11, p < 0.001$). Post-hoc analysis with Tukey's HSD identified significant mean differences ($p < 0.05$) between consumer discretionary and four other sectors: energy, health care, information technology, and utilities.

Our analysis showed above-average scores in the *utilities* and *energy* sectors (Fig. 4). The high OP in these sectors may be attributed to employees engaged in renewable energy sources, as they take pride in contributing to clean energy solutions and addressing pressing environmental issues. Example reviews: *Leading the way to a greener future. Leading the renewable and energy storage world transformation. Keeping the lights on means 24x7 dedication to customers* and *I can sense an exciting future that will reshape the company completely. The people here are great.* and *There is no slack anywhere in company X, which*



is a good thing for shareholders but means it can at times be quite exhausting working at company X.’

The *consumer staples* sector, which includes companies producing and selling essential consumer goods, exhibited substantial variability in OP scores. This can be attributed to the diverse nature of companies within this category. For example, companies offering food, beverages, or personal care products (e.g., Costco, Coca-Cola Company, and Procter & Gamble Co.) may experience different levels of pride (e.g., ‘*The CEO and CFO and in company X support their workers, and they believe in compensating and growing the people from within. They really care about their employees.*’) compared to those involved in the production and sale of alcohol or tobacco products (e.g., Philip Morris International Inc.). The latter group, particularly tobacco companies, faces growing external critique and societal pressures due to the prioritisation of healthier lifestyles, which can influence employee sentiments towards the organisations. Example reviews: ‘*New purpose to deliver smoke-free products jumpstarted a need to address the culture and explore new ways of working. [...] Cost pressure due to lower sales volumes results in budget cuts.*’ and ‘*[...] they did not mention that the role involved standing in a retail shop the entire day, approaching smokers and enrolling them to try smoke-free products. The management did not mention that the job involved “educating retailers” which in reality means bribing them to sell some-free in exchange for money vouchers, ipads and airpads.*’

Lastly, both the *consumer discretionary* and *communication services* sectors demonstrated significantly lower levels of OP. The *consumer discretionary* sector comprises companies that produce non-essential or discretionary goods and services, including industries with high turnover rates such as retail and e-commerce platforms. Example reviews: ‘*Very repetitive after a while!*’ and ‘*Like most retail businesses the pay is substandard. Many employees have second jobs to make ends meet even if they are working full-time here.*’ and ‘*Very stressful, lots of rude and antsy customers especially contractors.*’ and ‘*Staff keeps getting reduced, therefore not enough time to serve everyone well. Then “they” wonder why sales are down.*’ and ‘*Some poor policies that leave you dealing with some angry customers.*’ Moreover, ethical concerns surrounding sustainability, animal welfare, and labour practices, particularly within industries such as apparel and luxury goods, contribute to diminished feelings of pride among employees. Example reviews ‘*Minimal pay, rude customers, long holiday hours,*’ and ‘*Management seems clueless outside of their department.*’ and ‘*Retail is boring and monotonous*’ and ‘*long hours over Christmas - have to work weekends*’ and ‘*People no longer matter. Values and principles the company once was built on are no*

longer there.' and *'Hours are not as organised as they could be, targets sometimes unrealistic - management need to understand many customers will shop online as websites offer better discounts.'*

In the *communication services* sector, especially social media and telecommunication companies, negative public sentiment related to issues like data privacy, misinformation, or service problems can impact employees' pride in their organisation. This sector often grapples with ethical issues such as data privacy, content regulation, and the digital divide. If employees feel that their company is not effectively addressing these issues, it can reduce their sense of pride in the organisation. Example reviews: *'Low pay and questionable privacy practices.'* and *'You face angry customers, [...] The trainers were not helpful and you feel like you've been left to handle things on your own.'* Additionally, in large communication service companies, employees may perceive themselves as small contributors in a large system, leading to a diminished sense of pride in their work's impact on the company and its customers. Example reviews: *'Whilst trying to keep the best person in each position, the company creates too much turnover.'*, *'We faced frequent layoffs when projects would complete [...].'*

PS applies to the team as a unit, and it is high in IT and low in communications and healthcare. PS-related items were centred around the "team" aspect, evident from words like *team, people, help, and friendly*, as depicted in Fig. 2. This finding aligns with previous literature, suggesting that PS mainly pertains to employee experiences within their teams [49]. By exploring more specific words associated with PS (Fig. 3), we found that it exhibited stronger associations with keywords such as *environment, friendly, help, training, and management*. These findings align with previous studies that highlight the interaction between PS and leadership qualities and team characteristics [49, 50]. Overall, this shows that our PS measure applies to the team as a unit.

PS across industry sectors. ANOVA also confirmed statistically significant differences in PS scores across sectors ($F = 2.75, p < 0.01$). Post-hoc analysis with Tukey's HSD identified significant mean differences ($p < 0.05$) between information technology and three other sectors: consumer discretionary, consumer staples, and health care.

We found high levels of PS in the *information technology* sector (Fig. 4). This can be attributed to the career growth opportunities provided and the supportive work environment within the *information technology* sector companies. Example reviews: *'[...] it was all about what you could bring to the table and they have an environment that made you want to bring all that you had.'* and *'Great career growth opportunities, excellent horizontal industry mobility, and chance to truly own your career trajectory.'* and *'Amazing benefits, lots of time spent on your training and development.'* and *'Diversity among teams. [...].'*

On the other hand, the *utilities* sector exhibits lower levels of PS. As discussed above, the *utilities* sector is often undergoing transformations due to advances in renewable and green technologies, regulatory changes, and shifting energy policies. This environment of change and uncertainty may impact job security, thereby affecting psychological safety among employees. Example reviews: *'Difficult to change things for the better. [...] They rather conform everyone to the same standard rather than embracing new and better ideas.'* and *'Constant pressure. Micro managed in ways that have never been seen before.'* and *'Current efforts to downsize are disturbing.'* and *'Lack of communication between coworkers.'*

In the *communication services* sector, which includes telecommunications, broadcasting, media and entertainment, and internet service companies, lower levels of both PS

Table 3 Predicting company ratings using OP and PS scores. Standard errors are shown below the coefficients

	Career	Balance	Culture	Management	Overall
Intercept	2.650*** (0.058)	3.127*** (0.090)	2.731*** (0.065)	2.543*** (0.056)	2.914*** (0.047)
OP_{score}	0.201*** (0.010)	0.156*** (0.016)	0.228*** (0.011)	0.204*** (0.010)	0.227*** (0.008)
PS_{score}	0.025** (0.011)	0.052*** (0.018)	0.058*** (0.013)	0.056*** (0.011)	0.022** (0.009)
$OP_{score} \times PS_{score}$	0.010*** (0.004)	-0.010* (0.006)	-0.002 (0.004)	0.013*** (0.004)	0.001 (0.003)
$\log(n_{reviews})$	0.085*** (0.009)	0.028** (0.014)	0.089*** (0.010)	0.059*** (0.009)	0.069*** (0.007)
R^2	0.749	0.470	0.773	0.797	0.842
Adjusted R^2	0.745	0.463	0.770	0.794	0.840
AIC	-167.621	112.588	-95.732	-185.739	-301.657
BIC	-148.811	131.398	-76.922	-166.929	-282.847
Residual Std. Error	0.184	0.287	0.207	0.179	0.149
F Statistic	233.107***	69.263***	266.298***	306.434***	416.059***

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

and OP may result from the fast-paced and dynamic focus on innovation, adaptability, and customer demands. This constant change can create uncertainty and stress, which in turn impacts psychological safety. Mergers, acquisitions, and layoffs are common in this sector due to its dynamic nature, further contributing to a sense of job insecurity that negatively impacts psychological safety. Example reviews: ‘*So many redundancies, too often.*’ and ‘*You need to work hard.*’ and ‘*Terrifying with anxiety 24/7. [...] Your job is always on the line, you can be let go at any point.*’

OP and PS are related organizational culture aspects, but each independently predicts employee satisfaction. Although OP and PS scores are correlated (Appendix, Table 5), the interaction term in the joint regression is not significant for most satisfaction ratings, including overall satisfaction (Table 3), indicating that neither aspect amplifies nor diminishes the other’s effect. The exceptions are career and management ratings, where small but significant positive interaction terms show that higher OP or PS amplify each other’s effects. The joint regression results also show that OP effects are significantly stronger than PS effects in predicting ratings, sometimes by an order of magnitude (e.g., 0.227 for OP vs. 0.022 for PS when predicting overall satisfaction). Accordingly, we do not pursue a combined OP×PS analysis; instead, for an interested reader, we present an exploratory, construct-by-construct comparison of company standings across the two constructs in Appendix §C.

6 Discussion

Organisational pride and psychological safety provide a holistic understanding of employee experiences and attitudes throughout the organisational hierarchy. To assess these constructs automatically, we developed and validated an automated language-based framework, which we applied to hundreds of U.S. S&P 500 companies. By employing this framework, we identified significant variations in these constructs across industry sectors and categorised companies based on their respective levels of organisational pride and psychological safety. Next, we discuss the theoretical and practical implications of our work, while also acknowledging its limitations and outlining potential future directions.

6.1 Theoretical implications

From a theoretical perspective, our work contributes to developing automated ways of modelling workplace-related psychological constructs such as organisational culture. Our work builds upon previous research [8, 56] that aimed to automatically quantify organisational culture on a large scale, particularly by utilising Glassdoor data. We have made technical improvements to enhance the construct validity of this approach. In our study, we proposed a framework that captures organisational constructs from free-form employee reviews about the company by leveraging validated psychometric measures. While we demonstrated the framework's effectiveness in capturing OP and PS, it is important to note that the framework can be applied to any validated organisational construct as long as the corresponding psychometric measure exists. For constructs that have both positive and negative items in their psychometric measures (as was the case for PS in this study), our framework can directly incorporate them. However, for constructs with only affirmative items (such as OP in this study), the creation of negative items, as detailed in Sect. 4.1.2, can be beneficial.

Our study contributes to the theoretical understanding of organisational culture by examining two psychological constructs of OP and PS. To the best of our knowledge, this is the first study to report the varying degrees of OP and PS across different industry sectors, addressing the call for organisational behavioural research to disentangle context-specific theories rather than relying solely on context-general ones [71]. The consistent relationships between OP, PS, and ratings across the five Glassdoor categories suggest the general importance of these constructs in fostering employee satisfaction. As we demonstrated in Table 3, OP is a strong and robust predictor of all employee satisfaction ratings, with coefficients consistently larger than those for PS. Interestingly, for overall satisfaction, the effect of OP is an order of magnitude greater than that of PS (0.227 vs. 0.022). This likely reflects the fact that OP is measured at the company level and thus aligns closely with overall company satisfaction, while PS primarily captures team-level dynamics. This pattern holds across specific domains such as career, work-life balance, culture, and management. Although PS also significantly predicts satisfaction outcomes, its independent contribution is smaller. The interaction term between OP and PS is generally not significant, with the exception of career and management satisfaction, where the small but significant positive coefficients suggest a modest amplifying effect when both pride and safety are high. However, these interaction effects are limited in magnitude, and the primary drivers of employee satisfaction appear to be the main effects of OP and, to a lesser extent, PS. These findings indicate that initiatives aiming to increase organizational pride may have a particularly strong impact on employee satisfaction, though fostering psychological safety can provide additional benefits, especially in areas related to career advancement and management quality.

6.2 Practical implications

From a practical standpoint, our work offers immediate benefits to organisations by integrating automated psychological assessment into their workflows, conducting sector-specific analyses, and cultivating a positive and productive work environment.

Automated Psychological Assessment. Our automated framework for assessing organisational constructs, including organisational pride, overcomes several limitations associated with conventional self-report assessments. Traditional internal corporate surveys often

suffer from method bias and inconsistent implementation across companies. Our framework mitigates biases related to self-report data by using anonymous employee-provided information. This approach ensures more reliable and consistent measurements of organisational pride across diverse industry sectors, allowing organisations to obtain accurate insights into their employees' experiences and attitudes. A practical way to use our framework would be the integration with interactive dashboards where HR and managers can observe OP and PS scores over time, at both team and company levels, identifying trends and pinpointing areas for improvement. These dashboards, coupled with actionable insights and recommendations, can guide strategies to cultivate a positive company culture. For example, if a company scores low on psychological safety, HR could implement specific programs aimed at building trust, such as anonymous feedback tools or conflict resolution workshops. On the contrary, a company with low organisational pride might need to enhance internal communication regarding successes and milestones to boost pride, ensuring that the company's purpose is clearly communicated and that employees understand how their daily tasks fit into the bigger picture. The creation of educational tools and training materials focused on practices that foster organisational pride and psychological safety can educate both employees and leaders alike, and keep everyone engaged and informed about the best practices to enhance the company's cultural health.

Sector-Specific Analyses. By using our framework, any organisation, whether it be a small company or a large corporation, can identify the level of its workforce's organisational pride and psychological safety within its industry sector, and compare it across other sectors. This is important because organisations within specific sectors may face unique challenges and can leverage their strengths to address them effectively. By understanding the sector-specific dynamics related to organisational pride and psychological safety, organisations can tailor their strategies and interventions to create an environment that aligns with their industry's requirements and capitalise on their strengths [72].

Positive and Productive Work Environment. Organisations that prioritise creating a positive work environment can reap numerous benefits. Research has consistently shown that a positive work environment contributes to higher levels of employee well-being and performance [2, 3]. Strategies that foster emotional attachment, commitment, and psychological safety among employees not only can significantly enhance well-being and performance but also productivity. When employees feel valued, supported, and psychologically safe to take interpersonal risks and express their ideas, they are more likely to engage in proactive behaviours and contribute to the organisation's success [1].

Our framework's practical implications extend beyond organisational stakeholders to prospective employees and internal stakeholders. Job seekers can utilise our framework to compare organisational culture across companies and make more informed employment decisions (e.g., through visualisation tools; similar tools have been developed to allow prospective employees to compare their values with their prospective employers [57]). This can benefit both job seekers and current employees in reflecting on their assumptions and expectations regarding the work environment [73, 74].

6.3 Limitations and future work

Our work is subject to five main limitations, primarily related to the dataset used and how we measured the two psychological constructs studied.

The first three limitations pertain to the dataset itself. Firstly, our analysis focused on a subset of S&P 500 companies for which data were available, comprising 318 companies.

This focus may limit the applicability of our findings to smaller firms with fewer online reviews. For instance, the observed distinctions between OP and PS in relation to employee satisfaction might be less evident in smaller companies, which typically have fewer and less diverse teams compared to larger corporations. To broaden the applicability of our findings, future research could use stratified sampling to gather reviews from a more varied range of businesses, considering both economic status and cultural context. Secondly, it is important to acknowledge the ethno-cultural bias inherent in our dataset, as we exclusively considered U.S.-based companies and English-language reviews. This restricts the cross-cultural generalisability of our findings, and it is not clear if and which of our findings would be impacted. To enhance the cultural and linguistic breadth of future studies, researchers could use large language models to process and analyse organisational reviews in multiple languages from different platforms, thus offering a more nuanced understanding of organisational behaviour across diverse cultural backdrops. Lastly, although our dataset includes timestamps, we did not extensively analyse the temporal dynamics of OP or PS, partly because of the insufficient number of reviews over time for smaller companies, which would not allow for a comprehensive assessment. Future research, leveraging our methodology and larger temporal datasets, could incorporate longitudinal analysis to uncover time-sensitive effects, providing a more dynamic understanding of organisational culture over time. Examining these temporal aspects could reveal how perceptions and attitudes toward organisations change in response to macroeconomic fluctuations or significant internal changes.

The second two limitations pertain to the psychological constructs themselves. Firstly, our assessment of OP and PS focused on the organisation level, based on employee reviews about their overall company experience. While this approach provides a broad overview of psychological safety within an organisation, our findings suggest that PS may not be fully captured at this level because it often operates more meaningfully within smaller teams. PS is inherently relational and context-dependent, deeply influenced by immediate interpersonal dynamics, the support structures of team members, and leadership styles within teams [1]. This suggests that PS may be most accurately assessed within the context of these smaller, functionally cohesive units rather than solely at the organizational level. Our results reinforce this perspective by revealing that organizational-level PS may overlook significant variations across teams, which play a critical role in shaping employees' daily experiences. Treating PS solely as an organisation-wide construct may mask these intra-organisational nuances, as employees within certain departments or project teams could experience significantly different levels of psychological safety. For example, employees working in high-stress or high-risk areas, such as customer-facing roles, may feel varying degrees of PS compared to those in internal roles, depending on the unique culture and support systems within their teams. To address this limitation, future studies could adapt the framework we experimented with to assess PS at the team level by analysing reviews from employees grouped by team, department, or project if such data is available. Secondly, our measure of employee satisfaction was limited to the online company ratings data currently accessible to our research team. Future research could incorporate additional indices of success, such as those measured by industry standards organisations, annual revenue, or social media sentiment, to provide a more comprehensive understanding of employee satisfaction.

6.4 Ethical considerations

In this study, we used public and anonymised data obtained from a reputable company review website. The qualitative analysis was conducted using data from publicly accessible sources, and any excerpts utilised to support the analysis were paraphrased to ensure anonymity and avoid the identification of specific reviews. The statistics presented at the company level were derived from publicly available information.

7 Conclusion

This work highlights the importance of understanding employee experiences and attitudes through the lens of organisational pride and psychological safety to cultivate a positive work environment and drive organisational success. By utilising a deep-learning framework and analysing a vast dataset of employee reviews from 318 major U.S. companies, our findings emphasise the crucial role of organisational pride and psychological safety in shaping employee engagement, satisfaction, productivity, and innovation. By prioritising organisational pride and psychological safety, organisations can create a work environment that fosters employee well-being and drives long-term success.

Appendix A

A.1 Organisational pride (OP) & psychological safety (PS) psychometrics measures considered

OP items by Jones [41]:

1. "I am proud to work for [this company]."
2. "People respect what [this company] does."
3. "I am proud to be associated with [this company]."
4. "I am proud of what [this company] accomplishes."

OP items by Turban et al. [12]:

1. "I would be proud to be an employee at this firm"
2. "I would receive respect from others if I worked for this firm."
3. "My family would feel proud of me if I worked for this firm."

OP items by Cable and Turban [32]; adapted from [12]:

1. "I would feel proud to be an employee of this firm."
2. "I would be proud to tell others that I work for this firm."
3. "I would be proud to identify myself personally with this firm."

PS items by Edmondson [1]:

1. "If you make a mistake on this team, it is often held against you."
2. "Members of this team are able to bring up problems and tough issues."
3. "People on this team sometimes reject others for being different."
4. "It is safe to take a risk on this team."
5. "It is difficult to ask other members of this team for help."
6. "No one on this team would deliberately act in a way that undermines my efforts."
7. "Working with members of this team, my unique skills and talents are valued and utilised."

A.2 Summary statistics of the variables used in this study

Table 4 Summary statistics of analyzed variables

	Mean	Std. dev	Min	Median	Max	Histogram
<i>OP_{score}</i>	0.00	1.46	-4.78	-0.05	4.97	
<i>PS_{score}</i>	0.00	1.29	-3.99	-0.14	4.82	
# reviews	1354.68	2013.33	102.00	452.00	12,707.00	
Balance rating	3.30	0.39	2.16	3.33	4.51	
Career rating	3.21	0.37	2.16	3.19	4.47	
Culture rating	3.30	0.43	2.17	3.31	4.64	
Management rating	2.94	0.39	2.00	2.93	4.35	
Overall rating	3.36	0.37	2.25	3.36	4.60	

A.3 Correlations between OP, PS, and overall ratings

Table 5 Correlation matrix among OP and PS scores, and the overall company rating

	OP score	PS score	Overall rating
OP score	1.00		
PS score	0.69	1.00	
overall rating	0.89	0.67	1.00

A.4 Multicollinearity between OP, PS, and overall ratings

Table 6 Variance Inflation Factor (VIF) Results

Variable	VIF
OP score	1.904
PS score	1.896
overall rating	1.009

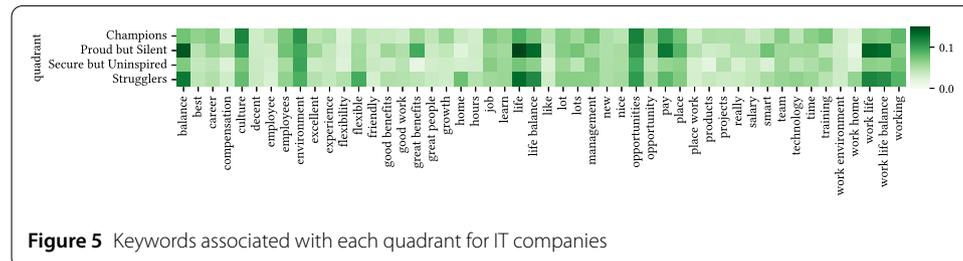
A.5 Results of manual assessment and baseline comparison

We validated our methods on a set of manually annotated reviews against two baselines: a phrase-based Rapid Automatic Keyword Extraction (RAKE), and a large language model (LLM)-based GPT-4o-mini.

First, we selected 400 reviews using stratified random sampling: 200 from pros, such that 50 were labeled with our method as OP, 50 as PS, and 100 that were labeled as not pertaining to either of the two constructs; and equivalently, 200 from cons such that 50 were labeled with our method as LOP, 50 as LPS, and 100 that were labeled as not pertaining to either of the two constructs. Then, three authors independently annotated the 400 reviews for the presence of each construct, as defined in the literature, keeping the construct items in mind. Specifically, they did so in the following way. For the reviews sampled from *cons*, they assigned one of four labels—LOP, LPS, LOP&LPS, or NA (inapplicable)—to match the automatic framework; likewise, reviews sampled from *pros* were labeled OP, PS, OP&PS, or NA. For example a cons review “HR department is protective of

Table 7 Comparison of our method and a keyword-based and an LLM (GPT-4o-mini) baselines

Method	F1 OP	F1 PS	F1 LOP	F1 LPS
GPT-4o-mini	39.75%	61.64%	28.32%	49.61%
Keyword-based	28.57%	40.91%	8.57%	43.75%
Our method	42.50%	57.80%	34.78%	52.46%



the company, not the employees. Hard to move up the food chain for some positions. Company is so big that individuals are sometimes unvalued.” would be labeled as LOP&LPS since it indicates both a lack of OP (individuals are unvalued) and PS (HR is protective of the company, not employees), whereas a pros review “Conservative and rarely have layoffs.” would be labeled with NA because it mentions job security but does not clearly express high OP or PS.

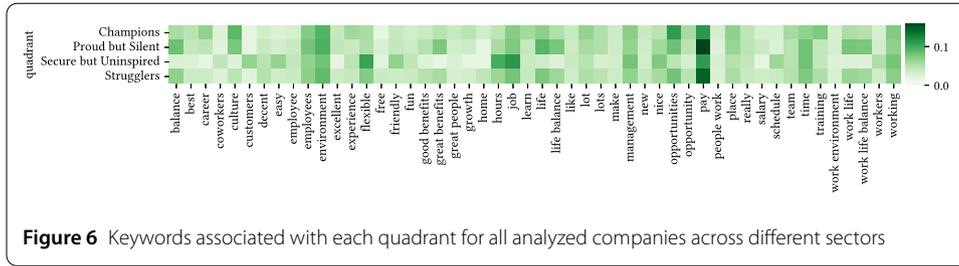
All three annotators initially agreed on 44.5% of the pros and 49.0% of the cons, while two out of the three annotators agreed on 80.3% of the pros and 80.8% of the cons. To ensure label quality, the annotators subsequently met to discuss and resolve disagreements. These primarily involved clarifying interpretations of the two constructs and assessing the level of implied information within brief reviews. For instance, some reviews mentioning low job safety were initially labeled as LOP by one annotator, but the group collectively agreed this did not equate to LPS. Similarly, reviews highlighting high wages or a good work-life balance were occasionally labeled as OP by an individual annotator, but upon discussion, it was concluded that these alone were insufficient to imply pride in the company. The final consensus labels were used for evaluating our method and the baselines.

Table 7 shows that our method outperforms the RAKE baseline implemented using rake-nltk Python library on the set of 400 manually annotated reviews. Moreover, our method also outperforms GPT-4o-mini in predicting three out of four constructs from the reviews (OP, LOP, LPS), and is at par for the last construct (PS).

A.6 IT sector analysis

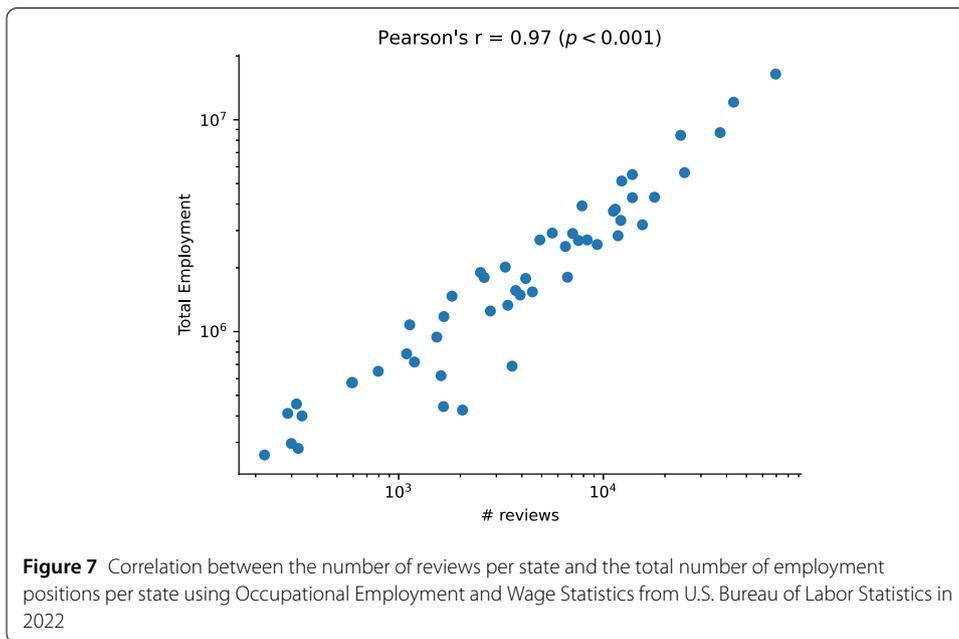
Upon excluding more generic keywords like ‘benefits’, ‘company’, ‘good’, ‘great’, ‘people’, and ‘work’, we found that IT companies with low public sentiment scores frequently had positive reviews mentioning “work-life balance” (Fig. 5). This contrasted with companies with Company and Teams Leadership (high OP, high PS), which did not emphasise work-life balance to the same degree, according to reviews. Moreover, “opportunities” and “environment” were frequently cited as advantages across reviews.

Although “work-life balance” was brought up in reviews for all analysed companies across different sectors, this theme did not emerge as strongly outside the IT industry (Fig. 6). Interestingly, when PS was low, employees tended to only mention “pay” as a positive factor rather than citing deeper advantages like work-life balance or opportunities



for growth. This suggests that a lack of psychological safety may limit employees’ ability to find meaning and fulfilment in their work beyond monetary compensation. In contrast, the prevalence of “work-life balance” comments for IT companies indicates that employees feel safe voicing appreciation for experiences beyond just pay when psychological safety is higher.

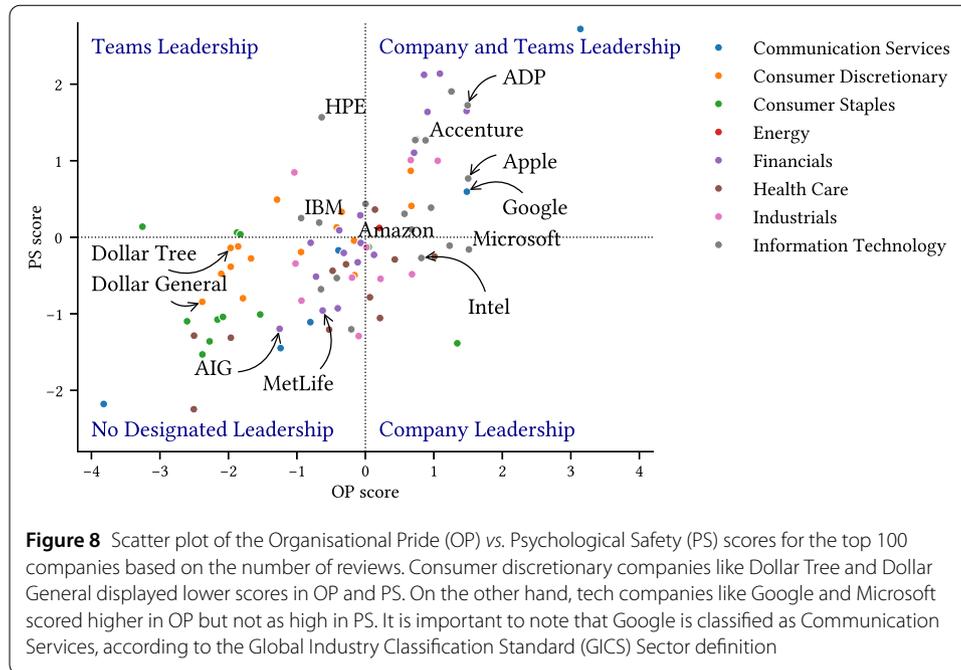
A.7 Sensitivity analysis



Appendix B: Analysing OP and PS in combination

To visualise the interplay between OP and PS, we projected the companies onto a Cartesian plane with two dimensions (Fig. 8). This method allowed us to succinctly and clearly represent each company’s positioning in relation to these two constructs.

To understand what characterises these four quadrants, we identified the most frequent keywords, including 1, 2, 3, and 4-grams, from the pro and con reviews. Next, we removed stopwords before calculating the TF-IDF scores for the identified n-grams. We only kept the top 100 keywords to identify the discriminative ones. Finally, we fit a multinomial logistic regression model using the top 100 keywords as the input to predict the quadrant of



the companies from which the reviews are coming. Finally, we identified the most prominent keywords by looking at the model’s coefficients for each quadrant.

Appendix C: Studying OP and PS in combination

As both OP and PS show significant correlations with overall company ratings (Appendix Table 5), with Pearson coefficients of $r = 0.89$ and $r = 0.67$, respectively, it is unsurprising that these constructs both emerged as significant *positive* predictors of the ratings individually (Table 3). However, given that OP and PS are also moderately correlated ($r = 0.69$), we sought to explore how the interplay between these two constructs relates to the company scores. Since the variance inflation factor (VIF) analysis found that the multicollinearity among our variables was low-to-moderate and, it did not require corrective measures (see Appendix Table 6), we began by examining a model that predicts employee ratings using OP and PS scores together, while controlling for the number of reviews. The results in Table 3 show that the model coefficients for OP are positive when predicting various ratings (ranging from 0.156 for balance to 0.228 for company culture), while the coefficients for PS are also positive but with a weaker effect (ranging from 0.022 for overall to 0.058 for culture). We also estimated a regression model incorporating an interaction term between OP and PS (see Appendix Table 3). The results remained consistent: OP and PS continued to show a positive association with the rating after controlling for the number of reviews, while the coefficient for their interaction term was not statistically significant.

We named the four quadrants shown in Fig. 8 as *No Designated Leadership (low OP & low PS)*, *Teams Leadership (low OP & high PS)*, *Company Leadership (high OP & low PS)*, and *Company and Teams Leadership (high OP & high PS)* for ease of reference.

We conducted an ANOVA on the mean company ratings across the four groups and found statistically significant differences ($F = 105.36, p < 0.001$). The post-hoc analysis presented in Table 8 further demonstrates that the differences are significant for most pairwise comparisons between the quadrants.

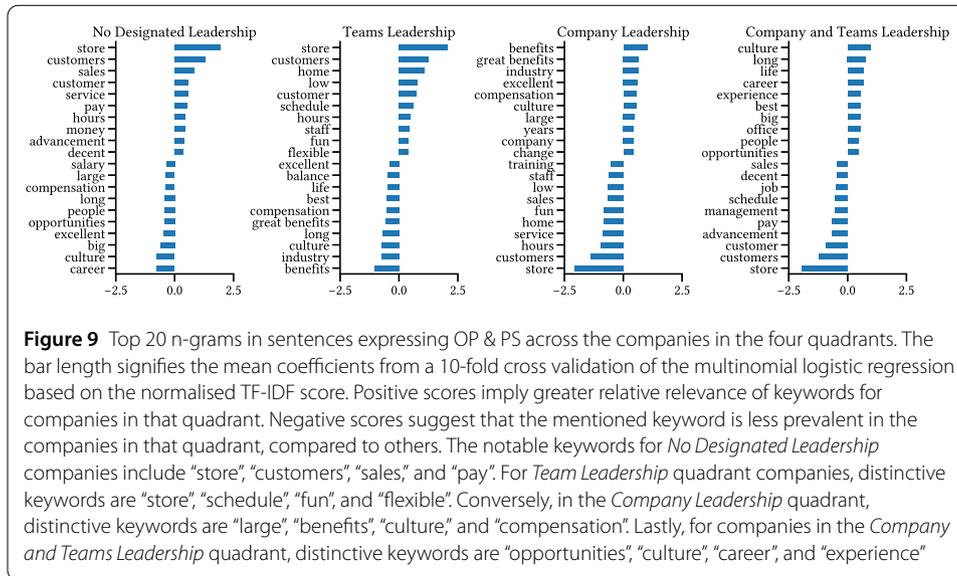


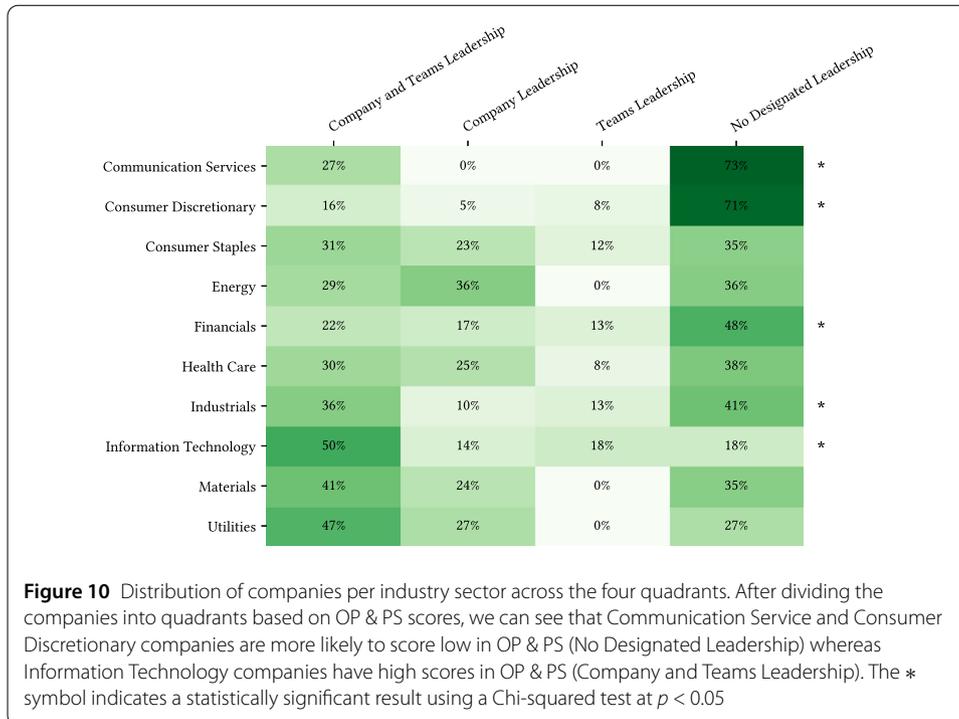
Table 8 Tukey HSD Test Results for Multiple Comparison of Means (FWER = 0.05)

Group 1	Group 2	Mean difference	p-adj	Reject
Company Leadership	Company & Teams Lead	0.1449	0.008	T
Company Leadership	No Designated Leadership	-0.4393	0.000	T
Company Leadership	Teams Leadership	-0.3889	0.000	T
Company & Teams Leadership	No Designated Leadership	-0.5843	0.000	T
Company & Teams Leadership	Teams Leadership	-0.5338	0.000	T
No Designated Leadership	Teams Leadership	0.0504	0.781	F

We next provide a definition for each quadrant, its characteristics derived from the most prominent keywords (Fig. 9), and the sectors that are related to it (Fig. 10).

No Designated Leadership. These companies exhibit a persistent challenge in cultivating a workplace culture that fosters pride and belonging (low OP) while also struggling to establish a psychologically supportive and safe environment (low PS). This is evident as employees predominantly focus on retirement (*401K* in Fig. 9) and significantly mention the following words less frequently compared to employees in the three other company types: *culture* (indicating low pride), and *friendly*, *work-life balance*, and *care* (indicating low safety). Notable companies in this category include MetLife, Nielsen, AIG, Dollar Tree, and Comcast, primarily representing sectors such as *financials*, *consumer discretionary*, and *communication services*.

Teams Leadership. These companies might face challenges in organisational leadership and fostering a strong sense of belonging (low OP), yet they excel in strong team leadership by creating a supportive and inclusive environment that prioritises employee well-being and open collaboration (high PS). This is apparent as employees frequently discuss *managers*, *coworkers*, and *flexible* work environments (signifying high safety). Conversely, there is less emphasis on discussions about *great benefits*, *career*, *projects*, and *culture* (indicating low pride). Companies such as Kmart and Staples in consumer staples ensure employee safety through a less demanding environment or more tightly knit teams, but this approach might inadvertently impact the pride associated with working in these organisations.



Company Leadership. These companies are distinguished by their robust organisational leadership, fostering strong pride and loyalty among employees towards the organisation (high OP). However, they encounter difficulties in establishing a psychologically safe workplace that promotes openness, vulnerability, and collaborative risk-taking (low PS). This is reflected in employee discussions about *culture*, *great benefits*, and *compensation*, (signifying high pride), while conversations about *coworkers*, *people work*, and *flexibility* are less frequent (indicating low safety). Companies such as BlackRock and JPMorgan in financials, Intel and Microsoft in information technology, and Pfizer and Merck in health care are examples, known for their challenging and dynamic environments, potentially affecting their employees’ psychological safety. As we have seen from the results on predicting ratings from both OP and PS, in companies with high OP, it is lower PS that contributes to a higher rating. Hence, many of the companies in this quadrant are highly rated by their employees.

Company and Teams Leadership. These companies are lauded for their outstanding leadership, both organizationally (high OP) and in team dynamics, fostering a supportive, inclusive, and empowering work culture (high PS). Employees in such environments are generally more motivated and experience greater job satisfaction, contributing to higher productivity. This is reflected in employee discussions about *projects*, *career*, *learning*, and *work-life balance*, contrasting with lesser focus on *hours* and *balance*, and seldom mentioning attributes like *easy* and *free*. As Šćepanović et al. [59] suggest, fulfilling workplaces often involve challenges and positive stress, rather than ease. These companies span diverse sectors, including notable ones like Apple and Nvidia in information technology, Halliburton in energy, Eastman Chemical Company in materials, and General Electric Company and FedEx in industrials.

Abbreviations

OP, Organisational Pride; PS, Psychological Safety; OC, Organisational Culture; NLP, Natural Language Processing; CSR, Corporate Social Responsibility; ISE, Internal Sustainability Efforts; OCB, Organisational Citizenship Behavior; S&P 500, Standard and Poor's 500; SBERT, Sentence-level BERT; SNLI, Stanford Natural Language Inference; LOP, Lack of Organisational Pride; LPS, Lack of Psychological Safety; TF-IDF, Frequency-inverse Document Frequency; CEO, Chief Executive Officer; CFO, Chief Financial Officer; IT, Information Technology; GICS, Global Industry Classification Standard; U.S., United States; HR, Human Resources; LLM, Large Language Model.

Acknowledgements

We thank Ke Zhou for his help with the data curation and Indira Sen for her advice on the methodology development.

Author contributions

AS collected the data and conducted the analysis. AS, SS, MC, and DQ conceived the experiments and wrote the manuscript.

Funding information

Nokia Bell Labs provided support in the form of salaries for authors [AS, SS, MC, DQ], but did not have any additional role in the study design, data collection and analysis, decision to publish, or preparation of the manuscript. MC also received support by the European Union's Horizon 2020 Research and Innovation Programme (grant agreement No. 739578). The specific roles of these authors are articulated in the 'author contributions' section.

Data availability

Data and code is available at: <https://social-dynamics.net/organizational-culture>.

Declarations

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Competing interests

The authors declare no competing interests.

Author details

¹Nokia Bell Labs, Cambridge, UK. ²University of Oxford, Oxford, UK. ³Politecnico di Torino, Turin, UK. ⁴University College London, London, UK. ⁵Cyens Centre of Excellence, Nicosia, Cyprus.

Received: 1 June 2024 Accepted: 17 July 2025 Published online: 05 August 2025

References

1. Edmondson A (1999) Psychological safety and learning behavior in work teams. *Adm Sci Q* 44(2):350–383
2. Harter JK, Schmidt FL, Hayes TL (2002) Business-unit-level relationship between employee satisfaction, employee engagement, and business outcomes: a meta-analysis. *J Appl Psychol* 87(2):268. <https://doi.org/10.1037/0021-9010.87.2.268>
3. Shin SJ, Zhou J (2003) Transformational leadership, conservation, and creativity: evidence from Korea. *Acad Manag J* 46(6):703–714. <https://doi.org/10.2307/30040662>
4. Spector PE, Jex SM (1998) Development of four self-report measures of job stressors and strain: interpersonal conflict at work scale, organizational constraints scale, quantitative workload inventory, and physical symptoms inventory. *J Occup Health Psychol* 3(4):356. <https://doi.org/10.1037/1076-8998.3.4.356>
5. Spicer A (2020) Organizational culture and covid-19. *J Manag Stud* 57(8):1737–1740. <https://doi.org/10.1111/joms.12625>
6. Vargas-Hernández JG, Noruzi MR, et al (2009) An exploration of the organizational culture in the international business relationships and conflicts era. *Am J Econ Bus Admin* 1(2):182–193. <https://doi.org/10.3844/ajebasp.2009.182.193>
7. Eversole BA, Venneberg DL, Crowder CL (2012) Creating a flexible organizational culture to attract and retain talented workers across generations. *Adv Dev Hum Resour* 14(4):607–625. <https://doi.org/10.1177/1523422312455612>
8. Das Swain V, Saha K, Reddy MD, Rajvanshy H, Abowd GD, De Choudhury M (2020) Modeling organizational culture with workplace experiences shared on glassdoor. In: Proceedings of the 2020 CHI conference on human factors in computing systems, pp 1–15. <https://doi.org/10.1145/3313831.3376793>
9. Cooke RA, Rousseau DM (1988) Behavioral norms and expectations: a quantitative approach to the assessment of organizational culture. *Group Organ Stud* 13(3):245–273
10. Prajogo DI, McDermott CM (2011) The relationship between multidimensional organizational culture and performance. *Int J Oper Prod Manag* 31(7):712–735
11. Gouthier MH, Rhein M (2011) Organizational pride and its positive effects on employee behavior. *J Serv Ind Manag* 22(5):633–649. <https://doi.org/10.1108/09564231111174988>
12. Turban DB, Lau C-M, Ngo H-Y, Chow IHS, Si SX (2001) Organizational attractiveness of firms in the People's Republic of China: a person-organization fit perspective. *J Appl Psychol* 86(2):194–206. <https://doi.org/10.1037/0021-9010.86.2.194>
13. Edmondson AC (2018) *The fearless organization: creating psychological safety in the workplace for learning, innovation, and growth*. Wiley, Hoboken

14. Higgins MC, Dobrow SR, Weiner JM, Liu H (2022) When is psychological safety helpful in organizations? A longitudinal study. *Acad Manag Discov* 8(1):77–102. <https://doi.org/10.5465/amd.2018.0242>
15. Sturm RE, Jolly PM, Williams SD (2022) It's a matter of organizational pride: how perceptions of organizational virtuousness and competence affect employee behaviors. *J Bus Psychol* 37(5):1079–1097. <https://doi.org/10.1007/s10869-021-09786-9>
16. Cheney G, Tompkins PK (1987) Coming to terms with organizational identification and commitment. *Cent States Speech J* 38(1):1–15. <https://doi.org/10.1080/10510978709368225>
17. Mas-Machuca M, Berbegal-Mirabent J, Alegre I (2016) Work-life balance and its relationship with organizational pride and job satisfaction. *J Manag Psychol* 31(2):586–602. <https://doi.org/10.1108/JMP-09-2014-0272>
18. Taylor C, Dollard MF, Clark A, Dormann C, Bakker AB (2019) Psychosocial safety climate as a factor in organisational resilience. In: Implications for worker psychological health, resilience, and engagement. *Psychosocial safety climate: a new work stress theory*, pp 199–228
19. Edmondson AC, Dillon JR, Roloff KS (2007) 6 three perspectives on team learning: outcome improvement, task mastery, and group process. *Acad Manag Ann* 1(1):269–314. <https://doi.org/10.1080/078559811>
20. Edmondson AC, Lei Z (2014) Psychological safety: the history, renaissance, and future of an interpersonal construct. *Annu Rev Organ Psychol Organ Behav* 1(1):23–43. <https://doi.org/10.1146/annurev-orgpsych-031413-091305>
21. Strueby S (2019) Are employees more likely to accept organizational change if
22. Abror A (2017) The relationship between psychological safety. Self-efficacy
23. Vaske JJ (2011) Advantages and disadvantages of Internet surveys: introduction to the special issue. *Hum Dimens Wildl* 16(3):149–153. <https://doi.org/10.1080/10871209.2011.572143>
24. Duda MD, Nobile JL (2010) The fallacy of online surveys: no data are better than bad data. *Human Dimens Wildl* 15(1):55–64. <https://doi.org/10.1080/10871200903244250>
25. Gigliotti LM (2011) Comparison of an Internet versus mail survey: a case study. *Human Dimens Wildl* 16(1):55–62. <https://doi.org/10.1080/10871209.2011.535241>
26. Fricker RD, Schonlau M (2002) Advantages and disadvantages of Internet research surveys: evidence from the literature. *Field Methods* 14(4):347–367. <https://doi.org/10.1177/152582202237725>
27. Ouchi WG, Wilkins AL (1985) Organizational culture. *Annu Rev Sociol* 11(1):457–483. <https://doi.org/10.1146/annurev.so.11.080185.002325>
28. Sokro E, et al (2012) Analysis of the relationship that exists between organisational culture, motivation and performance. *Probl Manag 21st Century* 3(1):106–119. <https://doi.org/10.33225/pmc/12.03.106>
29. Gunter B, Furnham A (1996) Biographical and climate predictors of job satisfaction and pride in organization. *J Psychol* 130(2):193–208. <https://doi.org/10.1080/00223980.1996.9915001>
30. Arnett DB, Laverie DA, McLane C (2002) Using job satisfaction and pride as internal-marketing tools. *Cornell Hotel Restaur Adm Q* 43(2):87–96. [https://doi.org/10.1016/S0010-8804\(02\)80035-0](https://doi.org/10.1016/S0010-8804(02)80035-0)
31. Kraemer T, Weiger WH, Gauthier MH, Hammerschmidt M (2020) Toward a theory of spirals: the dynamic relationship between organizational pride and customer-oriented behavior. *J Acad Mark Sci* 48(6):1095–1115. <https://doi.org/10.1007/s11747-019-00715-0>
32. Cable DM, Turban DB (2003) The value of organizational reputation in the recruitment context: a brand-equity perspective. *J Appl Soc Psychol* 33(11):2244–2266. <https://doi.org/10.1111/j.1559-1816.2003.tb01883.x>
33. Gauthier MHJ, Rhein M (2011) Organizational pride and its positive effects on employee behavior. *J Serv Manag* 22:633–649. <https://doi.org/10.1108/09564231111174988>
34. Raza A, Farrukh M, Iqbal MK, Farhan M, Wu Y (2021) Corporate social responsibility and employees' voluntary pro-environmental behavior: the role of organizational pride and employee engagement. *Corp Soc-Responsib Environ Manag* 28:1104–1116. <https://doi.org/10.1002/csr.2109>
35. Durrah O, Allil K, Gharib M, Hannawi S (2021) Organizational pride as an antecedent of employee creativity in the petrochemical industry. *Eur J Innov Manag* 24(2):572–588. <https://doi.org/10.1108/EJIM-08-2019-0235>
36. Appleberg KA (2005) The construction of a nomological network for organizational pride. *Benedictine University*
37. Hameed Z, Khan IU, Islam T, Sheikh Z, Khan SU (2019) Corporate social responsibility and employee pro-environmental behaviors: the role of perceived organizational support and organizational pride. *South Asian J Bus Stud*. <https://doi.org/10.1108/SAJBS-10-2018-0117>
38. Oo EY, Jung H, Park I-J (2018) Psychological factors linking perceived csr to ocb: the role of organizational pride, collectivism, and person–organization fit. *Sustainability* 10(7):2481. <https://doi.org/10.3390/su10072481>
39. Helm S (2013) A matter of reputation and pride: associations between perceived external reputation, pride in membership, job satisfaction and turnover intentions. *Br J Manag* 24(4):542–556. <https://doi.org/10.1111/j.1467-8551.2012.00827.x>
40. Ng TWH, Yam KC, Aguinis H (2019) Employee perceptions of corporate social responsibility: effects on pride, embeddedness, and turnover. *Pers Psychol* 72:107–137. <https://doi.org/10.1111/peps.12294>
41. Jones DA (2010) Does serving the community also serve the company? Using organizational identification and social exchange theories to understand employee responses to a volunteerism programme. *J Occup Organ Psychol* 83(4):857–878. <https://doi.org/10.1348/096317909X477495>
42. Masterson C, Liden RC, Kluemper D (2017) Peeling back the layers: a multi-method investigation of organizational pride, vol 2017. <https://doi.org/10.5465/ambpp.2017.189>
43. Jones DA, Willness CR, Madey S (2014) Why are job seekers attracted by corporate social performance? Experimental and field tests of three signal-based mechanisms. *Acad Manag J* 57(2). <https://doi.org/10.5465/amj.2011.0848>
44. O'Reilly C, Chatman J (1986) Organizational commitment and psychological attachment. The effects of compliance, identification, and internalization on prosocial behavior. *J Appl Psychol* 71(3):492–499. <https://doi.org/10.1037/0021-9010.71.3.492>
45. Mael F, Ashforth BE (1992) Alumni and their alma mater: a partial test of the reformulated model of organizational identification. *J Organ Behav* 13(2):103–123. <https://doi.org/10.1002/job.4030130202>
46. Pearsall MJ, Ellis AP (2011) Thick as thieves: the effects of ethical orientation and psychological safety on unethical team behavior. *J Appl Psychol* 96(2):401. <https://doi.org/10.1037/a0021503>
47. Baer M, Frese M (2003) Innovation is not enough: climates for initiative and psychological safety, process innovations, and firm performance. *J Organ Behav* 24(1):45–68. <https://doi.org/10.1002/job.179>

48. Carmeli A (2007) Social capital, psychological safety and learning behaviours from failure in organisations. *Long Range Plan* 40(1):30–44. <https://doi.org/10.1016/j.lrp.2006.12.002>
49. Newman A, Donohue R, Eva N (2017) Psychological safety: a systematic review of the literature. *Hum Resour Manage Rev* 27(3):521–535. <https://doi.org/10.1016/j.hrmr.2017.01.001>
50. Carmeli A, Reiter-Palmon R, Ziv E (2010) Inclusive leadership and employee involvement in creative tasks in the workplace: the mediating role of psychological safety. *Creat Res J* 22(3):250–260. <https://doi.org/10.1080/10400419.2010.504654>
51. Chen G, Tjosvold D (2012) Shared rewards and goal interdependence for psychological safety among departments in China. *Asia Pac J Manage* 29(2):433–452. <https://doi.org/10.1007/s10490-010-9201-0>
52. Cooke RA, Szumal JL (1993) Measuring normative beliefs and shared behavioral expectations in organizations: the reliability and validity of the organizational culture inventory. *Psychol Rep* 72(3):1299–1330. <https://doi.org/10.2466/pr0.1993.72.3c.1299>
53. O'Reilly CA III, Chatman J, Caldwell DF (1991) People and organizational culture: a profile comparison approach to assessing person-organization fit. *Acad Manag J* 34(3):487–516. <https://doi.org/10.2307/256404>
54. Hofstede G, Neuijen B, Ohayv DD, Sanders G (1990) Measuring organizational cultures: a qualitative and quantitative study across twenty cases. *Adm Sci Q*, 286–316. <https://doi.org/10.2307/2393392>
55. Glaser SR, Zamanou S, Hacker K (1987) Measuring and interpreting organizational culture. *Manag Commun Q* 1(2):173–198. <https://doi.org/10.1177/0893318987001002003>
56. Sen I, Quercia D, Capra L, Montecchi M, Šćepanović S (2023) Insider stories: analyzing internal sustainability efforts of major US companies from online reviews. *Humanit Soc Sci Commun* 10(1). <https://doi.org/10.1057/s41599-023-01672-4>
57. Ceccarini C, Bogucka EP, Sen I, Constantinides M, Prandi C, Quercia D (2022) Visualizing internal sustainability efforts in big companies. *IEEE Comput Graph Appl* 42(3):87–98. <https://doi.org/10.1109/MCG.2022.3163063>
58. Lee Y (2021) Bridging employee advocacy in anonymous social media and internal corporate social responsibility (csr). *Manag Decis*. <https://doi.org/10.1108/MD-01-2020-0101>
59. Šćepanović S, Constantinides M, Quercia D, Kim S (2023) Quantifying the impact of positive stress on companies from online employee reviews. *Sci Rep* 13(1). <https://doi.org/10.1038/s41598-022-26796-6>
60. Reimers N, Gurevych I (2019) Sentence-bert: sentence embeddings using Siamese bert-networks. In: Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP), pp 3982–3992. <https://doi.org/10.18653/v1/D19-1410>
61. Ni J, Hernandez Abrego G, Constant N, Ma J, Hall K, Cer D, Yang Y (2022) Sentence-T5: scalable sentence encoders from pre-trained text-to-text models. In: Muresan S, Nakov P, Villavicencio A (eds) Findings of the association for computational linguistics: ACL 2022. Association for Computational Linguistics, Dublin, pp 1864–1874. <https://doi.org/10.18653/v1/2022.findings-acl.146>
62. Muennighoff N, Tazi N, Magne L, Reimers N (2023) MTEB: massive text embedding benchmark. In: Vlachos A, Augenstein I (eds) Proceedings of the 17th conference of the European chapter of the association for computational linguistics. Association for Computational Linguistics, Dubrovnik, pp 2014–2037. <https://doi.org/10.18653/v1/2023.eacl-main.148>
63. Devlin J, Chang M-W, Lee K, Toutanova K (2018) Bert: pre-training of deep bidirectional transformers for language understanding. arXiv preprint. [arXiv:1810.04805](https://arxiv.org/abs/1810.04805)
64. Song X, Salcianu A, Song Y, Dopson D, Zhou D (2020) Fast wordpiece tokenization. arXiv preprint. [arXiv:2012.15524](https://arxiv.org/abs/2012.15524)
65. Rose S, Engel D, Cramer N, Cowley W (2010) Automatic keyword extraction from individual documents. Wiley, New York. <https://doi.org/10.1002/9780470689646.ch1>
66. Hurst A, Lerer A, Goucher AP, Perelman A, Ramesh A, Clark A, Ostrow A, Welihinda A, Hayes A, Radford A, et al (2024) Gpt-4o system card. arXiv preprint. [arXiv:2410.21276](https://arxiv.org/abs/2410.21276)
67. Vellum AI (2025) LLM Leaderboard. <https://www.vellum.ai/llm-leaderboard>. Accessed: 2025-05-02
68. Abdi H, Williams LJ (2010) Tukey's honestly significant difference (hsd) test. *Encycl Res Des* 3(1):1–5
69. Braun V, Clarke V (2006) Using thematic analysis in psychology. *Qual Res Psychol* 3(2):77–101. <https://doi.org/10.1191/1478088706qp063oa>
70. Youn H, Kim J-H (2022) Corporate social responsibility and hotel employees' organizational citizenship behavior: the roles of organizational pride and meaningfulness of work. *Sustainability* 14(4):2428. <https://doi.org/10.3390/su14042428>
71. Judge TA, Kammeyer-Mueller JD (2012) General and specific measures in organizational behavior research: considerations, examples, and recommendations for researchers. *J Organ Behav* 33(2):161–174. <https://doi.org/10.1002/job.764>
72. Barney JB, Hesterly WS (2015) Strategic management and competitive advantage: concepts and cases. Pearson, London
73. De Choudhury M, Counts S (2013) Understanding affect in the workplace via social media. In: Proceedings of the 2013 conference on computer supported cooperative work, pp 303–316. <https://doi.org/10.1145/2441776.2441812>
74. Shami NS, Yang J, Panc L, Dugan C, Ratchford T, Rasmussen JC, Assogba YM, Steier T, Soule T, Lupushor S, et al (2014) Understanding employee social media chatter with enterprise social pulse. In: Proceedings of the 17th ACM conference on computer supported cooperative work & social computing, pp 379–392. <https://doi.org/10.1145/2531602.2531650>

Publisher's note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.