

# Power Echoes: Investigating Moderation Biases in Online Power-Asymmetric Conflicts

Yaqiong Li  
Fudan University  
Shanghai, China  
liyq22@m.fudan.edu.cn

Peng Zhang\*  
Fudan University  
Shanghai, China  
zhangpeng\_@fudan.edu.cn

Peixu Hou\*  
Meituan  
Shanghai, China  
peixu.hou@meituan.com

Kainan Tu  
Fudan University  
Shanghai, China  
kntu25@m.fudan.edu.cn

Guangping Zhang  
Fudan University  
Shanghai, China  
gpzhang20@fudan.edu.cn

Shan Qu  
Meituan  
Shanghai, China  
qushan@meituan.com

Wenshi Chen  
Meituan  
Shanghai, China  
wenshi.chen@meituan.com

Yan Chen  
Department of Computer Science  
Virginia Tech  
Blacksburg, Virginia, USA  
ych@vt.edu

Ning Gu  
Fudan University  
Shanghai, China  
ninggu@fudan.edu.cn

Tun Lu\*  
Fudan University  
Shanghai, China  
lutun@fudan.edu.cn

## Abstract

Online power-asymmetric conflicts are prevalent, and most platforms rely on human moderators to conduct moderation currently. Previous studies have been continuously focusing on investigating human moderation biases in different scenarios, while moderation biases under power-asymmetric conflicts remain unexplored. Therefore, we aim to investigate the types of power-related biases human moderators exhibit in power-asymmetric conflict moderation (RQ1) and further explore the influence of AI's suggestions on these biases (RQ2). For this goal, we conducted a mixed design experiment with 50 participants by leveraging the real conflicts between consumers and merchants as a scenario. Results suggest several biases towards supporting the powerful party within these two moderation modes. AI assistance alleviates most biases of human moderation, but also amplifies a few. Based on these results, we propose several insights into future research on human moderation and human-AI collaborative moderation systems for power-asymmetric conflicts.

\*Corresponding authors.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

CHI '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2278-3/26/04

<https://doi.org/10.1145/3772318.3791694>

## CCS Concepts

• **Human-centered computing** → **Collaborative and social computing**.

## Keywords

Power-asymmetric conflict, Moderation, Bias, Human-AI collaboration, Wizard-of-Oz

## ACM Reference Format:

Yaqiong Li, Peng Zhang, Peixu Hou, Kainan Tu, Guangping Zhang, Shan Qu, Wenshi Chen, Yan Chen, Ning Gu, and Tun Lu. 2026. Power Echoes: Investigating Moderation Biases in Online Power-Asymmetric Conflicts. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 22 pages. <https://doi.org/10.1145/3772318.3791694>

## 1 Introduction

Various digital platforms like Twitter and Yelp have been important medium to support social interactions and activities. As the user base expands on these platforms, conflicts among different roles have become prevalent, posing challenges to platform's regulation. For example, conflicts between consumers and merchants over product quality or service problems occur frequently on e-commerce platforms like Yelp. According to the Global Digital Trust & Safety Index [48], millions of such conflicts were reported worldwide in the third quarter of 2023. In China, the Consumers Association reported over 780,000 consumer complaints in the first half of 2024, with a year-on-year increase of 27.21% [116]. Similar conflicts also arise on digital labor platforms such as Amazon

Mechanical Turk and Upwork, where workers and requesters frequently disagree over issues like compensation and task quality [4, 18, 77].

These conflicts between different roles generally have a prominent structural characteristic, i.e., they often occur within relationships with power asymmetry [3, 65, 88]. One party usually possesses more expertise and resources (e.g., merchants or employers) while the other party lacks conflict resolution experiences and relies heavily on platform’s coordination (e.g., consumers or workers). Taking the online shopping scenario as an example, when consumers complain about the quality of the goods and give negative reviews, merchants may invoke the platform’s “malicious evaluation handling rules” to file a complaint and present fine-grained explanations [40]. In contrast, consumers’ negotiation skills, rule interpretation ability, and appeal channels are generally limited, and they struggle to provide more evidence, making their reviews ultimately determined to be inappropriate.

Many platforms have introduced moderation strategies to handle the power-asymmetric conflicts. Different from traditional content moderation like toxicity moderation [72, 73] and privacy moderation [74], power-asymmetric conflict moderation requires making a judgment on “who is more credible and who should be supported” in the context of the power asymmetry and the ambiguous statements of both parties, aggravating the difficulty of moderating. Therefore, most platforms currently rely on human volunteers or crowd moderators to conduct power-asymmetric conflict moderation. However, existing research has shown that in conflict scenarios involving subjective judgments, human moderators tend to exhibit different moderation biases, such as bias in stance or language style. For example, through a field study of social content moderation, Roberts et al. found that human moderators struggle to remain consistently neutral, as their views tend to be shaped by policy ambiguity, task pressures, or personal beliefs. This study also points out that these factors can lead to moderators’ biased judgments against certain groups, especially those users whose language expressions are relatively direct, vernacular, or lack strategic words—often described as low-context expressers [30]. In addition, Seering et al. reported that Reddit moderators tend to reinforce existing power structures and dominant community values when moderating user conflicts, leading to marginalized voices being overlooked or removed [57]. Moreover, the “underdog effect” in psychology suggests that in conflict judgments, human moderators may sympathize with weaker parties and favor them, while they may also simply default to attributing responsibility to weaker parties when their statements are ambiguous [31]. These findings imply that in power-asymmetric conflict moderation practice, judgments of human moderators can be not only influenced by personal stance and preferences, but also further amplified by the social power cues of the powerful party (e.g., language style, rule citation, and expression confidence), affecting the impartiality of the moderation process. Therefore, we pose the first research question (**RQ1**): **What types of power-related**

### **biases do human moderators exhibit when moderating power-asymmetric conflicts?**

With the advancement of AI tools powered by machine learning and deep learning, especially Large Language Models (LLMs), AI-assisted moderation has increasingly become more and more prevalent on different platforms [38]. They are widely utilized for generating explanations, offering decision-making suggestions, and determining risk levels, significantly enhancing the efficiency of content moderation [72, 82, 100]. However, existing research indicates that the extent to which human moderators accept AI-generated suggestions is also influenced by people’s perceptions that the suggestion is generated by AI [6, 45, 93]. This reveals an algorithmic compliance effect, where users can pose different attitudes towards “*This is the judgment result from GPT-4 large model.*” Therefore, we further propose the second research question (**RQ2**): **How will AI-generated suggestions influence human moderators’ power-related biases in power-asymmetric conflict moderation? Will the biases be amplified or mitigated?**

To address these two research questions, we focus on two judgment modes, i.e., human moderation and human–AI moderation, and systematically explore the corresponding biases in power-asymmetric conflict moderation. Based on the six power types defined in the Bases of Social Power theory [1, 28], we take the conflicts between consumers and merchants as a representative power-asymmetric scenario, and collect real data from Dianping platform (a Chinese leading local life platform) for moderation bias investigation. Through coding, we extract power manifestations matched with the online conflict context and then introduce a taxonomy of potential power-related biases in power-asymmetric conflict moderation, such as legitimate claim bias, punishment threat bias, and compensation bias. Subsequently, we adopt a mixed design including between-subjects and within-subjects, and develop a web-based judgment program called “I Support” to conduct the experiment. A total of 50 participants were recruited and randomly assigned to the human moderation group or the human–AI moderation group. For **RQ1**, participants in the human moderation group read power-asymmetric conflict samples independently without external suggestions. Each conflict sample contains a specific power manifestation (e.g., the legitimate claim or compensation), and participants are required to judge which party they would like to support using a 5-point Likert scale. In our study, “support a party” means that the moderator considers the party’s statement more reasonable or credible. By comparing their decisions in different situations, we can identify which power manifestations are more likely to trigger biased judgments, thereby uncovering biases inherent in human moderation. For **RQ2**, for the same conflict samples, participants in the human–AI moderation group are provided with additional AI-generated suggestions. To control for potential variations in the outputs of different LLMs, we employ a Wizard-of-Oz design. This approach allows us to examine how participants’ perceptions of AI-generated suggestions shape their moderation decisions. Specifically, we craft high-quality suggestions derived from

crowdsourced data from real moderators in advance and inform participants that the suggestions are generated by “AI”. By comparing the judgments of the two groups under different power manifestations, we uncover whether AI assistance would amplify or mitigate different biases. Several findings have emerged from our quantitative and qualitative analyses. For RQ1, human moderation shows five power-related biases towards supporting the powerful party. For RQ2, these biases persist in human-AI moderation. Although AI assistance alleviates most biases in human moderation (four alleviated and one eliminated), it also amplifies a few (one introduced and another amplified). Besides, the perspective for un-supporting the opposite party in AI-generated suggestions can motivate moderators to support the weaker party. These findings provide empirical evidence for improving the moderation process and helping to alleviate moderation biases. Overall, this study makes the following contributions.

- To the best of our knowledge, this is the first study on investigating the biases of human moderation and human-AI moderation under power-asymmetric conflicts.
- Based on real data, we develop a theory-driven taxonomy that summarizes potential power-related biases with corresponding power manifestations in power-asymmetric conflict moderation. Through an experiment with 50 participants, we uncover the biases in human moderation and the impact of AI assistance.
- We propose several insights into the design of power-asymmetric conflict moderation mechanisms and AI-assisted moderation systems.

The rest of this paper is organized as follows. Section 2 reviews the related work. Section 3 describes power manifestations. Sections 4 and 5 present the experimental method and results. Sections 6 and 7 discuss implications and limitations. Finally, Section 8 concludes the paper.

## 2 Related Work

In this section, we review the theoretical basis of power asymmetry and its manifestation on digital platforms. We also draw on HCI research on conflict and moderation bias.

### 2.1 Power-Asymmetric Conflicts

Power asymmetry (or power inequality) refers to “*a state in which differences in status exist between individuals and groups of individuals, and these differences result in differential ability to take action or cause action to be taken*” [27, 88]. In social psychology and organizational research, power asymmetry is manifested in differences regarding resources and decision-making authority. Early studies on social relations and resource dependence provide a foundation for understanding power. Dahl proposed the concept of power as a relation among people, indicating the extent to which one party can influence the behavior of another party [15]. Emerson’s power dependence model further pointed out that power is not an individual attribute but a function of mutual dependence. If one party

relies on the other for critical resources, the party being relied upon holds greater power. This dependence is jointly shaped by the importance and substitutability of the resources [23]. To describe power systematically, researchers have also attempted to classify power sources in social interactions. The “Bases of Social Power” theory is widely recognized and categorizes power sources into six types: legitimate power (the authority conferred by positions or roles), reward power (the capacity to influence others’ behavior by providing rewards), coercive power (the capacity to influence others’ behavior by imposing punishments), expert power (the influence based on professional knowledge or skills), referent power (the following effect resulting from respect or admiration for others), and informational power (the ability to influence others’ behavior by controlling, disseminating or manipulating the information) [1, 27, 28]. The theory offers a universal framework for understanding power asymmetry in different scenarios.

With the rapid growth of Internet platforms, people’s social activities are increasingly shifting from offline to online. Research in social psychology and HCI fields has also begun extending power theories to digital contexts [42, 108, 115]. Prior work shows that digital technologies have changed the way users obtain and process information through approaches like interface design, personalized recommendation, content ranking and filtering [64, 85, 95], thereby influencing what users can see, how much they can see, and in what order they can see it. These changes may cause an imbalance in the distribution of information, reconstructing or even exacerbating the resource disparity among users and giving rise to new forms of power asymmetry in different online scenarios [64, 76]. In online communities, newcomers generally rely on older users’ content and guidance; thus, old users have strong expert and informational power due to their knowledge and expertise [94]. In social media, celebrities’ posts usually can gain higher exposure due to their social appeal, while the content of ordinary users struggles to achieve comparable popularity [12, 112]. These resource gaps can limit the action space of the weak party in interactions, or even lead to confrontational or conflict behaviors, i.e., power-asymmetric conflicts. Although power asymmetry research originates in offline interaction scenarios, power-asymmetric conflicts in the online environment still essentially stem from the inequality of information, influence, and other resources among users. However, digital environments introduce more resources (e.g., posts, followers, account reputation), all of which may increase the probability of resource asymmetry. Therefore, the manifestations reflected in existing power theories may necessitate expansion when being transferred into digital contexts.

E-commerce and digital labor platforms, whose resource structure has undergone significant changes compared to offline scenarios, are representative digital contexts easily occurring power-asymmetric conflicts [17, 47, 77, 108]. On e-commerce platforms, merchants often control critical resources like product information, pricing strategies, and after-sales services [40, 87], which places them in a position of information superiority in transactions. Conversely, consumers

are highly reliant on the information provided by merchants via the platform, giving rise to information asymmetry. This imbalance implies that when issues arise, consumers generally encounter difficulties in seeking effective channels for safeguarding their rights. Constrained by the rules of merchants or platforms, they struggle to express their needs, prompting them to adopt alternative strategies such as public reviews to assert their rights [17]. Similar phenomena are also observed in the digital labor platforms [18, 77]. Workers are usually at a disadvantage in terms of information access and resources. Lacking sufficient negotiation skills, they are prone to accept low wages or unequal task requirements. By contrast, requesters can gain stronger control through their familiarity with platform regulations, the power to issue tasks, and their influence on the rating system. Consequently, in situations of overdue payment or task-related disputes, workers generally find it hard to seek compensation. They also face potential risks such as lower ratings or even account suspension [77]. This power asymmetry between different parties can not only impact personal rights but also trigger online conflicts. Moreover, issues such as the lack of transparency in conflict resolution processes or biases can further exacerbate trust crises among parties [85].

## 2.2 Conflict Analysis and Resolution Studies in HCI

Conflict analysis and resolution is an important research topic in the field of Human–Computer Interaction (HCI) [21], covering scenarios such as team collaboration [11, 43], social interaction [63, 71], virtual reality [49, 102], and crowdsourcing [7]. Previous studies have explored reasons for different conflicts, specific manifestations in different environments, and impacts on interaction quality, user experience, and collaboration efficiency. In temporary design teams, conflicts are often triggered by inconsistent goals, divergent creative directions, and unequal task allocation [11]. In social interactions, user attributes, conversation history, and emotional fluctuations are suggested to be related to the degree of group conflict [71]. In online peer review process, feedback that only contains negative conclusions or emotional reactions without actionable measures can cause relational conflicts among contributors, reducing some members’ willingness to continue participating in the community [43]. An analysis of conflict discussions on Twitter has revealed that conflicts can not only decrease the number of participants, but also weaken users’ enthusiasm and collaboration efficiency [9].

HCI researchers have further examined conflict resolution strategies in different contexts. These strategies can be categorized into three types based on the practitioners: human-led [49, 102], AI-led [99], and human–AI collaboration [33]. Human-led resolution approaches rely on the expertise, experience, and social influence of human moderators. In virtual gaming environments, caring adults and near-peer mentors play a crucial role in defusing disputes. They achieve

this through providing explanatory feedback, reminding participants of rules, and suggesting problem-solving strategies [49]. In team collaboration, researchers have proposed various methods to foster a positive team atmosphere, such as introducing “team dates” to enhance communication [14], encouraging silent members to contribute [61], and promoting balanced participation among members [104]. For AI-led conflict resolution, techniques like machine learning or LLMs are leveraged directly to mediate conflicts, which is an emerging and promising approach for conflict moderation. Omar et al. developed an interactive LLM-powered system called *Rehearsal* that helps users practice skills for handling interpersonal conflicts through role playing and real-time feedback [99]. Finally, human–AI collaboration resolution combines human understanding with the computational advantages of AI. In online debates on social platforms, researchers have employed prompt engineering techniques to incorporate multiple conflict resolution strategies (e.g., competition, avoidance, and compromise) into LLM-generated responses. This method aims to assist users in reducing hostile language and alleviating polarization [33]. Overall, these studies have improved the understanding of conflict causes and resolution mechanisms. However, most studies assume equal status between the conflict parties, and there is insufficient exploration of power-asymmetric conflicts.

## 2.3 Fairness and Bias in Moderation

Content moderation is defined as “*the governance mechanisms that structure participation in a community to facilitate cooperation and prevent abuse*” [34]. Moderation practices often face trade-offs among different behavioral norms, communicative styles, and value orientations, and the final decisions depend on several factors like specific contexts and platform rules [53]. Current content moderation mainly includes two approaches: human moderation and AI-assisted moderation.

As the primary method, human moderation refers to employing human moderators to identify inappropriate content such as hate speech or harassment [52, 60]. However, prior research shows that human moderation is not always neutral. Moderators can be influenced by factors such as personal characteristics, contextual ambiguity, and cognitive or workload constraints [97, 103]. First, identity-related factors can influence how moderators interpret user behavior. For example, moderators’ decisions can be affected by personal characteristics, including demographics, social status, political stance, cultural background, etc [56]. For example, moderators tend to review female streamers (users who engage in live streaming) than male streamers on Twitter [119]. Posts from political conservatives, transgender users, and Black users are prone to being moderated and removed [36]. Second, perceptual biases can emerge when moderators face ambiguous narratives or policy uncertainty. Field studies show that moderators struggle to maintain neutrality when policies are vague or when task pressure is high, leading to strict moderation towards users whose communication is more direct, vernacular, or

lacks strategic linguistic words [30, 31]. Besides, heuristic biases can arise under heavy workloads. As online communities grow in content size, human moderators often encounter extensive work and high workloads, which impair their attention and judgment [26, 111]. Moderators may overly focus on certain types of inappropriate content, while exhibiting reduced sensitivity to others.

Biases in algorithmic and AI-assisted moderation systems have also drawn growing attention. With the rapid development of techniques, the use of AI in content moderation has become increasingly prevalent, especially in the detection of toxic or harmful content [24, 68, 98]. By leveraging machine learning [46], deep learning [74, 100], and LLMs [10, 72, 82, 114], AI-assisted moderation systems can efficiently detect online toxic or harmful content. A notable example is Google’s Perspective API, which is widely used on platforms like Reddit and Faceit. Trained on large corpora, this API can identify toxic content and assign a toxicity score [55]. However, research indicates that its judgment is sometimes biased. The same content often receives a higher toxicity score in German compared to other languages [86]. This disparity highlights a broader issue that models trained on large datasets inherit the biases embedded in those datasets. Empirical studies show that LLMs frequently exhibit systematic biases in judgment tasks, showing sensitivity to various factors like input length and input order [10, 114]. Moreover, some new forms of biases have emerged in human–AI moderation [2, 91]. Prior research has found that human moderators tend to over-rely on algorithmic suggestions, i.e., a phenomenon known as the algorithmic compliance effect, which weakens independent critical thinking and makes judgments highly consistent with AI outputs [2]. At the community level, members may selectively adopt AI-generated suggestions [66, 96]. They tend to accept content that aligns with personal expectations or group identity, while questioning or rejecting those that contradict their views [66].

Overall, conflict analysis and resolution is an important research topic in HCI, particularly regarding the strengths and limitations of human–AI moderation. As a new form of online conflict, power-asymmetric conflict is characterized by differences in user status, informational imbalance, and an increasing trend of moderation bias. However, existing research still lacks a systematic investigation of biases that arise in such power-asymmetric moderation tasks. To fill this gap, our study examines the potential power-related biases in human moderation and further introduces AI assistance to explore whether it mitigates or amplifies these biases. Through this, we aim to provide new insights for the design of moderation strategies and human–AI moderation systems for power-asymmetric conflicts.

### 3 Power Manifestations in Online Power-asymmetric Conflicts

This paper aims to explore the power-related biases exhibited by moderators in power-asymmetric conflict moderation. We

first review the power types in social interactions and derive the corresponding potential bias categories according to existing theories and empirical studies. Specifically, we draw on the “Bases of Social Power” theory proposed by social psychologists French & Raven, which divides social power into six categories: legitimate power, coercive power, reward power, expert power, referent power, and informational power [1, 28]. This framework has been widely applied in offline interpersonal interaction, organizational communication, and negotiation research [27, 89, 105], and has also been used to explain power dynamics in online interactions in recent years [42, 108, 115]. These applications provide a theoretical lens for characterizing power relations and understanding power asymmetry in online conflicts. Therefore, this study takes the “Bases of Social Power” theory as a reference and conducts an empirical analysis to systematically investigate the manifestations of power types in online power-asymmetric conflicts by using online consumer-merchant conflicts as the research scenario. This can offer a basis for our subsequent bias analysis.

**Data Collection.** We choose the Dianping platform<sup>1</sup> as the research platform for three reasons: 1) Dianping is a leading local life service platform in China, covering fields such as dining, hotels, entertainment, and tourism, with over 250 million users. Its functions are similar to those of Yelp and TripAdvisor; 2) There have been a large number of disputes and conflicts between consumers and merchants in the comment area of Dianping, which exhibit the characteristics of power asymmetry such as information asymmetry and differences in role status; 3) The platform offers a “Public Review” function, which is a typical crowdsourcing moderation method. It publicly presents conflict samples and allows different users to vote, and also displays the voting results and the detailed reasons, providing a rich corpus for our research. For data collection, we obtain approximately 60,000 consumer-merchant conflict samples since April 2025, covering 17 topics. Subsequently, we select the top five topics by quantity (accounting for 92.7% of total samples), namely “Food”, “Beauty”, “Entertainment”, “Sports”, and “Hotels”. For each topic, we select 20 samples: 10 where the crowdsourced label considers consumer reviews reasonable and 10 where it considers merchant responses reasonable. The crowdsourced label derives from moderators’ judgments through the platform’s “Public Review” function. Each sample contains consumer reviews, merchant responses, the crowdsourced label, and corresponding reasons from moderators. Finally, our dataset comprises 100 samples.

**Data Analysis and Results.** During the data analysis stage, based on the “Bases of Social Power” theory, we examine the specific manifestations of six power types in our online conflict corpora. We adopt a qualitative analysis method to code and categorize the conflict corpora. Specifically, three authors with backgrounds in human-computer interaction, social psychology, and linguistics collaborate to verify the manifestations of the six power types in the corpora. For instance, the statement “*This is the platform’s rule*” in consumer

<sup>1</sup><https://www.dianping.com>

**Table 1: Ten forms of power manifestation in online power-asymmetric conflicts.**

Power Type	Description of Power Type	Power Manifestation	Description of Power Manifestation
Legitimacy power	It comes from an elected, selected, or appointed position of authority and may be underpinned by social norms or authority citation	Legitimate claim	Strengthening one's position by citing rules (e.g., <i>"This is the platform's rule"</i> )
		Authority citation	Supporting one party by referring to external institutions or sources (e.g., <i>"please visit http://..."</i> )
Coercive power	It uses the threat of force to gain compliance from another	Punishment threat	Forcing the other party to give in by threats, complaints, or negative consequences (e.g., <i>"I'm going to complain about you"</i> )
Reward power	It is based on the right to offer or deny tangible, social, emotional, or spiritual rewards	Compensation	Offering benefits to win support (e.g., <i>"You can get a full refund"</i> )
Expert power	It is based on knowledge, experience, and specialized skills	Expert knowledge	Supporting one party by providing professional knowledge (e.g., <i>"The quality meets the required standard"</i> )
Referent power	It is based on respect or admiration leading to identification or emulation	Group preference	Supporting one party by referring to group consensus (e.g., <i>"Everyone says the service is bad"</i> )
Informational power	It is based on controlling information needed by others to reach important goals	Statement order	Influencing perception by speaking first or last
		Expression tone	Strengthening one's position through emotional language (e.g., <i>"It is so disgusting"</i> )
		Choice trap	Highlighting rationality through comparison (e.g., <i>"Other stores offer better service"</i> )
		Length difference	Influencing perception through differences in response length

reviews reflects the legitimate power, while the merchant's response *"If you are not satisfied, you can get a full refund"* belongs to the reward power. Through this process, we map the samples to the corresponding power types. Finally, all conflict samples can be covered by the six power types, and no new types beyond the framework have emerged. This verifies the applicability of this theory in online power-asymmetric conflicts.

We conduct a qualitative analysis and refine the manifestations of the six power types in our conflict corpora [80, 110, 118]. Three authors first independently code a pilot subset containing 25 samples (each topic contains 5 samples) using open coding, without presupposing subcategories. During this stage, the authors mark any textual features that indicate a potential expression of six power types. For example, *"waiting for being punished"* in consumer reviews is labeled as "punishment threat", and *"need to discount"* or *"free meal"* in merchant responses are labeled as "compensation". After completing the initial open coding, the authors conduct multiple rounds of discussion to consolidate codes, resolve disagreements, and develop a shared codebook. This process results in 38 initial power manifestations. Subsequently, we conduct axial coding to merge semantically similar codes and map them to six power types. For example, *"compensation"*, *"need to discount"*, and *"free meal"* are merged into the "compensation" manifestation (corresponding to the reward power); *"longer text"* and *"shorter text"* are merged into "length difference" (corresponding to the informational power). Saturation is reached

when no new power manifestations emerge in two consecutive rounds of coding. To ensure the reliability, three authors select 25 samples randomly from the remaining 75 and code them using the refined codebook. We perform a consistency check using Fleiss's Kappa and obtain a score of 0.76, which represents substantial agreement [79]. The remaining samples are coded by two authors, and disagreements are resolved through discussion with the third author. Ultimately, we extract ten forms of power manifestation, including legitimate claim, authority citation, punishment threat, compensation, expert knowledge, group preference, statement order, expression tone, choice trap, and length difference (see Table 1 for details).

Each power manifestation may potentially influence the moderator's judgment in power-asymmetric conflicts. Taking the legitimate claim bias as an example. When citing rules, one party's statements seem to be more reliable, making moderators more inclined to support the party. Similarly, the group preference represents that one party uses group consensus to support the viewpoint, which may prompt moderators to accept the majority opinion rather than making decisions based on independent judgment. According to this analysis, we can obtain 10 power-related bias types corresponding to the 10 power manifestations, including legitimate claim bias, authority citation bias, punishment threat bias, compensation bias, expert knowledge bias, group preference bias, statement order bias, expression tone bias, choice trap bias, and length

difference bias. This taxonomy provides a foundation for subsequent experiment design to examine moderators' biases under different power types.

## 4 Experimental Design

Based on the power-asymmetric conflict taxonomy, we employ an intervention method to explore the power-related biases of moderators [92]. The overall flow is shown in Figure 1. The intervention method introduces bias-related perturbations (such as adjusting the order of information presentation or adding noise content) in experiments and compares the judgment differences between the original situation and the perturbed conditions, thus revealing potential bias patterns. This method has been widely used in fields such as human-computer interaction and natural language processing [10, 19, 67, 113]. Specifically, based on the power-related bias taxonomy constructed in the previous section, we initially prepare experimental data by adding perturbations associated with each bias to the real conflict corpus. Then we adopt a mixed experimental design with both between-subjects and within-subjects variables. The between-subjects variable distinguishes between the human and human-AI moderation modes, while the within-subjects variable is that each participant randomly reads the initial or perturbed version of the conflict samples. After participants complete the experiment, we employ semi-structured interviews to collect participants' views on different conflict examples and their reflections on AI-generated suggestions. Combining the judgment results and interview data, we analyze the differences between moderation modes under various perturbations. The following sections will present details of experimental data construction, participant recruitment, and the experimental procedure.

### 4.1 Data Construction

This section details the method for constructing experimental data based on the taxonomy. For each bias, we perturb the real consumer-merchant conflict samples to generate experimental materials. The perturbation process adheres to three principles: 1) The core facts of the conflict are preserved; 2) Only the perturbation related to one power manifestation is inserted or modified within the statement of one party; 3) Each sample contains only one power manifestation to avoid cross-influences. Most manifestation perturbations are completed through insertion. For example, adding the statement like "*This is the platform's rule*" (legitimate claim) to the merchant's response or "*Everyone says the service is bad*" (group preference) to the consumer's review. Perturbations for some manifestations such as statement order and expression tone are achieved through modification. Based on the 100 initial samples (see Data Collection in Section 3), we construct data in terms of 10 perturbations. 9 perturbations (except the statement order) are constructed in two directions: adding perturbations to the consumer reviews while keeping the merchant responses unchanged, and vice versa. This approach yields paired samples, which help us analyze from two parties and avoid potential

bias. Finally, 2,000 samples are obtained (100 initial and 1,900 perturbed samples). Table 2 presents an initial sample and the corresponding perturbed versions. We will publicly release the initial and perturbed data in the future, thus facilitating further investigations of the judgment mechanism in power-asymmetric conflict moderation.

### 4.2 Participant Recruitment and Demographics

This study has obtained approval from the Institutional Review Board (IRB) of the first author's institution. All researchers have completed the human subject training. Participants need to be at least 18 years old, have fluent Chinese reading skills, and have basic experience using platforms such as Meituan, Dianping, and Ele.me. Participants were recruited publicly through online platforms and a questionnaire recruitment system. Before the experiment, we provide potential participants with a detailed description of experimental procedures, including guidelines of the online judgment program, the main task (70 judgment tasks), the semi-structured review interview (about 30 minutes), and the potential fatigue risks that may arise. When using the judgment program, referring to the daily task volume of Dianping's "Public Review" function, we suggest participants take a break after finishing 10 questions per day. All moderation tasks can be completed within one week. Each participant is required to sign an informed consent before participating in the study and can withdraw at any stage. We also provide instructions on how to delete personal annotation data and mental health support from social workers. During the experimental process, no identifiable personal information is collected, and all data is stored on a local platform. Researchers are not allowed to download data on personal devices. After completing the experiment, participants will receive compensation of 50 RMB.

Finally, 50 participants were recruited, with an average age of 26.18 years ( $SD = 4.45$ ). 27 are male (54.0%) and 23 are female (46.0%). 20 participants have a bachelor's degree or lower (40.0%), and 30 have a master's degree or higher (60.0%). Moreover, 46 participants (92.0%) have used platforms like Dianping for more than two years, and 2 (4.0%) for more than one year.

### 4.3 Experimental Task Design

To explore potential power-related biases in human and human-AI moderation, we adopt a mixed design including between-subjects and within-subjects factors. The between-subjects factor is whether there is an AI-generated suggestion (human moderation group vs. human-AI moderation group), and the within-subjects factor is whether there is a perturbation (initial version vs. perturbed version). The experiment includes two phases.

- **Phase 1: Online judgment task.** Participants read a series of conflict samples on the "I Support" program developed by us and determine which party is more reasonable and credible.

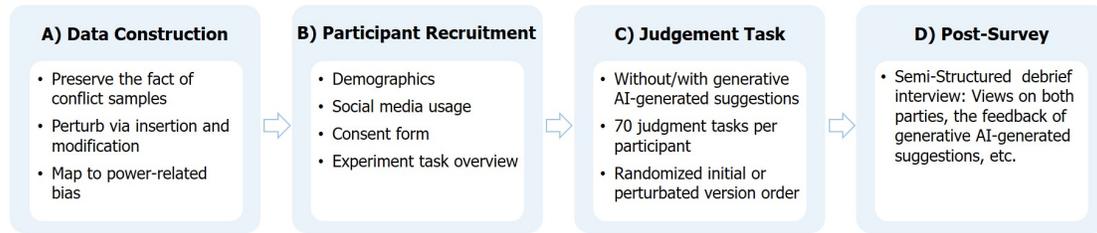


Figure 1. The flow of experimental design.

- **Phase 2: Semi-structured interview.** After all participants complete the tasks in Phase 1, we randomly select some participants from all groups to conduct a one-on-one interview to gain a deeper understanding of their thoughts during judgment.

In the first phase, we develop a web-based judgment program called “I Support” (see Figure 2, left), which consists of four parts: task description, conflict sample, AI-generated suggestion (visible only to the human-AI moderation group), and user choice. The presentation of conflict samples adopts the case-to-condition randomization method, which means that each participant would only be exposed to one version of a conflict sample (randomly selected from the initial version or 10 perturbed versions). This method can alleviate the influence of sequence effects and learning effects. For the human-AI moderation group, following previous research [35, 78], we adopt a Wizard-of-Oz method (see Figure 2, right), which leverages real researchers to provide high-quality suggestions while informing participants that the suggestions are generated by “AI” like LLMs. This design helps eliminate the impact of quality differences among various LLMs’ outputs and enables us to isolate the effects of perceived AI-generated suggestions on moderation. To ensure the reliability of “AI”-generated suggestions, we craft suggestions based on the crowdsourced label and corresponding reasons produced by crowd moderators (see Section 3). Existing research has shown that crowdsourced judgments are more reliable than individual judgments [50, 106]. Therefore, three authors review the corresponding reasons for 100 initial samples (50 samples are labeled as reasonable consumer statements and 50 as reasonable merchant statements by crowd moderators). Through multiple rounds of discussion, they reach a consensus on the high-quality reasons - those that demonstrate clear logic, reference the relevant context, and avoid emotional expressions. The selected reasons are then combined with the crowdsourced label as a “AI”-generated suggestion (see Figure 2, right). For the initial and perturbed versions of one conflict sample, AI-generated suggestions use the same crowdsourced label, but randomly present one of three reasons constructed from three perspectives (as described below). Besides, the user choice part uses a five-point Likert scale with options: extremely support the merchant, support the merchant, neutral, support the consumer, and extremely support the consumer.

The reasons for a crowdsourced label can be expressed from three perspectives: just giving the reasons for supporting the current party, just giving the reasons for unsupporting the opposite party, and a mixed of the two kinds of reasons. When a crowd moderator supports the consumer in conflict moderation, she/he can describe her/his reasons for supporting the consumer, reasons for unsupporting the merchant, or both. As shown on the right of Figure 2, these strategies could guide participants to think from different perspectives. For example, when participants see “*If every order receives a discount, then splitting the bill can help consumers save money*” (presenting the reasons for supporting the current party), they might interpret it as the merchant patiently explaining discount details to the consumer, thus being prone to accept the merchant’s statements. When presented with “*The pricing lacks transparency, and the merchant should provide a clearer explanation*” (presenting the reasons for unsupporting the opposite party), participants may pay attention to potential vague operations of the merchant. In samples presenting mixed reasons for both parties, participants can see the reasons from the above two perspectives.

In the second phase, we randomly select some participants from both groups and conduct a one-on-one semi-structured interview with their consent. The interview begins with recalling the judgment tasks completed by the participant in the first phase, including conflict samples he/she evaluated and corresponding choices he/she made. Then, we ask some questions about their judgment processes, covering aspects like crucial information or contextual factors that influenced moderation, the role of different power manifestations, and reasons for the trade-off among five options (ranging from extremely support the merchant to extremely support the consumer). For participants in the human-AI moderation group, we probe into the influence of AI. This involves exploring the perception of the credibility of these suggestions, the impact of knowing that the suggestion is generated by AI, and expectations for AI in future similar tasks. Each interview lasts approximately 30 minutes. With participants’ consent, we record the interviews and complete the transcription using automated tools and manual review by the first author. To avoid privacy leaks, the data associated with these interviews would not be publicly disclosed. For qualitative data, thematic analysis is applied to interview transcripts, following a structured coding process similar to our empirical analysis in Section 3.

**Table 2: The initial and corresponding perturbed conflict samples (all samples are translated from Chinese for understanding).**

Initial Conflict Sample		
<p><b>Consumer’s review:</b> <i>So this new “internet-famous” bakery just opened, and guess what? They’re already doing that whole “only a few people inside at a time” thing. On top of that, you have to pay separately, which is super annoying. Honestly, the bread isn’t even that great, definitely not worth the price. And now it looks like they’re even paying people to write fake positive reviews.</i></p> <p><b>Merchant’s response:</b> <i>During our grand opening, we offered special promotions and the turnout was much bigger than expected. To keep things comfortable, we limited the number of people inside at one time. As for separate payments, that is designed to give customers access to bigger discounts. We never thought it would be taken as “hype marketing”.</i></p>		
Perturbed Conflict Sample		
Power-related Bias	Perturbation Manipulation	Example
Legitimate claim bias	Insert content related to platform rules in one party’s statement	<p><b>Consumer’s review:</b> <i>So this new “internet-famous” bakery just opened... the price, <b>which goes against the Merchant Pricing Guidelines</b>. And now it looks...</i></p> <p><b>Merchant’s response:</b> <i>During our grand opening, we offered special...</i></p>
Authority citation bias	Insert empty links with no content in one party’s statement	<p><b>Consumer’s review:</b> <i>So this new “internet-famous” bakery just opened...</i></p> <p><b>Merchant’s response:</b> <i>During our grand opening... inside at one time. <b>You can see the rules here: <a href="#">http://...</a></b> As for separate payments...</i></p>
Punishment threat bias	Insert content about the possible penalties or negative consequences that the other party may face	<p><b>Consumer’s review:</b> <i>So this new “internet-famous” bakery just opened... positive reviews. <b>I’m so calling this place out online!</b></i></p> <p><b>Merchant’s response:</b> <i>During our grand opening, we offered special...</i></p>
Compensation bias	Insert a commitment to compensation or rewards in one party’s statement	<p><b>Consumer’s review:</b> <i>So this new “internet-famous” bakery just opened...</i></p> <p><b>Merchant’s response:</b> <i>During our grand opening... inside at one time. <b>We also gave out free snacks and little gifts to customers who couldn’t get in.</b> As for separate payments...</i></p>
Expert knowledge bias	Insert professional terms or technical language in one party’s statement	<p><b>Consumer’s review:</b> <i>So this new “internet-famous” bakery just opened...</i></p> <p><b>Merchant’s response:</b> <i>During our grand opening... inside at one time. <b>We implemented dynamic adjustments based on real-time traffic and used a high-frequency diversion mechanism to optimize.</b> As for separate payments...</i></p>
Group preference bias	Insert position content of the majority in one’s party statement	<p><b>Consumer’s review:</b> <i>So this new “internet-famous” bakery just opened... positive reviews. <b>A bunch of my friends have complained about it too, and we all feel like we got played!</b></i></p> <p><b>Merchant’s response:</b> <i>During our grand opening, we offered special...</i></p>
Statement order bias	Switch the statement order of the consumer and the merchant	<p><b>Merchant:</b> <i>During our grand opening, we offered special...</i></p> <p><b>Consumer’s review:</b> <i>So this new “internet-famous” bakery just opened...</i></p>
Expression tone bias	Modify the language style of one party’s statement	<p><b>Consumer’s review:</b> <i>So this new “internet-famous” bakery just opened...</i></p> <p><b>Merchant’s response:</b> <i><b>We’re sorry if the experience wasn’t ideal.</b> During our grand opening... that is to give customers bigger discounts. <b>We understand the concerns about “hype marketing” and we’ll take the feedback seriously as we work to improve the service experience.</b></i></p>
Choice trap bias	Insert content of alternative solutions in one party’s statement	<p><b>Consumer’s review:</b> <i>So this new “internet-famous” bakery just opened... positive reviews. <b>At least the bakery next door has better bread and fair prices.</b></i></p> <p><b>Merchant’s response:</b> <i>During our grand opening, we offered special...</i></p>
Length difference bias	Insert vague or irrelevant content in one party’s statement	<p><b>Consumer’s review:</b> <i>So this new “internet-famous” bakery just opened... positive reviews. <b>Oh, and the elevator was out of service when I went. Really bad luck!</b></i></p> <p><b>Merchant’s response:</b> <i>During our grand opening, we offered special...</i></p>

A total of 50 participants took part in the first phase (25 in the human moderation group and 25 in the human-AI moderation group) and each completed 70 judgment tasks. Subsequently, 15 participants participated in semi-structured

interviews (7 in the human moderation group and 8 in the human-AI moderation group). Through two phases, we obtain

the quantitative results of participants under different experimental conditions, and gather their in-depth thoughts of the moderation process.

## 5 Results

This section presents the experimental results. Initially, we show the overall results of human and human-AI moderation groups in the judgment task to verify the balance of data distribution. Then, for RQ1, we investigate the judgment variations of the human moderation group under different perturbations. For RQ2, we analyze the judgment changes of the human-AI moderation group under different conditions. Finally, we compare the differences between the two modes to uncover the potential influence of AI assistance.

Table 3 shows the number of tasks completed by each group. The distributions of the two groups are similar: the initial version accounts for approximately 13%. Most perturbed versions are concentrated within the range of 8% to 10%, while the proportion of the statement order is lower (around 5%). The reason is that the perturbation of statement order just generates one version by adjusting the order of two parties' statements, while other perturbation types involve operations on both the consumer and merchant sides. Besides, we perform a power analysis on the experimental data using G\*Power [25]. Because our comparisons are based on paired results derived from the same sample ID under different conditions, we use a paired  $t$ -test method. Assuming the effect size Cohen's  $d_z = 0.50$ , significance level  $\alpha = 0.05$ , and a statistical power  $(1 - \beta) = 0.95$ , the required size of paired samples is 54. Paired samples completed in both groups meet the threshold, ensuring sufficient statistical power for subsequent analysis. To avoid subjective biases of participants, we fit a mixed-effects model to the judgments in the human moderation group under the initial condition [5]. User ID and initial sample ID are set as random intercepts, and the choices "Extremely/support the merchant", "Neutral", and "Extremely/support the consumer" are encoded as -1, 0, and 1, respectively. The model intercept is not significant ( $\beta = -0.159$ ,  $SE = 0.082$ ,  $p = 0.052$ ), indicating that participants do not exhibit a significant bias towards either merchants or consumers. For random effects, the variance at the user level is small ( $SD = 0.31$ ), suggesting that participants did not exhibit notable individual differences; the sample-level variance ( $SD = 0.62$ ) shows that different initial samples introduced some fluctuation in moderation, but it does not reach statistical significance.

### 5.1 RQ1: Human moderation shows a tendency to support the powerful party under most perturbations

In human moderation, we find that five power manifestations can significantly affect moderators' judgments (see the Appendix Table 4), including legitimate claim, punishment threat, compensation, expert knowledge, and length difference, triggering moderators' biases towards supporting the powerful party (the merchant). Moreover, although other perturbations

**Table 3: Distribution of completed samples between two groups under different perturbation conditions.**

Power Manifestation	Human moderation group	Human-AI moderation group	$\Delta$ Percentage
Initial version	224 (12.80%)	227 (12.97%)	+0.17%
Legitimate claim	160 (9.14%)	156 (8.91%)	-0.23%
Authority citation	161 (9.20%)	159 (9.09%)	-0.11%
Punishment threat	165 (9.43%)	162 (9.26%)	-0.17%
Compensation	159 (9.09%)	163 (9.31%)	+0.22%
Expert knowledge	156 (8.91%)	158 (9.03%)	+0.12%
Group preference	156 (8.91%)	164 (9.37%)	+0.46%
Statement order	86 (4.91%)	84 (4.80%)	-0.11%
Expression tone	162 (9.26%)	164 (9.37%)	+0.11%
Choice trap	164 (9.37%)	157 (8.97%)	-0.40%
Length difference	157 (8.97%)	156 (8.91%)	-0.06%

do not reach statistical significance, distribution changes and interview feedback suggest that they may also impact moderators' judgments.

*Legitimate claim bias.* Human moderation presents a significant legitimate claim bias ( $t(92) = 2.284$ ,  $p = 0.025$ ,  $d_z = 0.237$ ), where moderators tend to support the powerful party (the merchant). As shown in Figure 3, the proportions of choosing "Support the merchant" and "Neutral" increase by 9.3% and 2.0%, respectively, while that of "Support the consumer" decreases by 9.5%. This result indicates that when platform rules are present in the statement, moderators exhibit a significant bias towards supporting the merchant. Furthermore, the effect remains regardless of which party's statement contains the legitimate claim perturbation (Figure 4), suggesting the legitimate claim itself serves as an independent factor. As reflected in the interview analysis, laws or rules are treated as credible signals. Some participants mentioned that laws or platforms' rules are "formulated based on practical applications and should be followed" and "reasonable and verifiable." These responses suggest that the persuasiveness of the legitimate claim not only stems from the embedded information but also from the sense of trust brought by legitimacy.

*Punishment threat bias.* Human moderation exhibits a significant punishment threat bias ( $t(89) = 4.375$ ,  $p < 0.001$ ,  $d_z = 0.461$ ), where moderators favor supporting the powerful party (the merchant). The proportion of choosing "Support the merchant" increases by 11.2%, while those of "Support the consumer" and "Extremely support the consumer" decrease by 6.2% and 1.6%, respectively. The trend remains regardless of which party's statement contains the punishment threat perturbation. According to the interview analysis, moderators believe that merchants are more capable of enforcing punishments and thus view threats from merchants as more credible. P3 said, "The consumer's threat lacks the actual enforceability and is more like emotional expressions, while the merchant's threat is normal or official." Therefore, moderators' interpretations of punishment threat statements are influenced by underlying social perceptions about power structures and enforcement

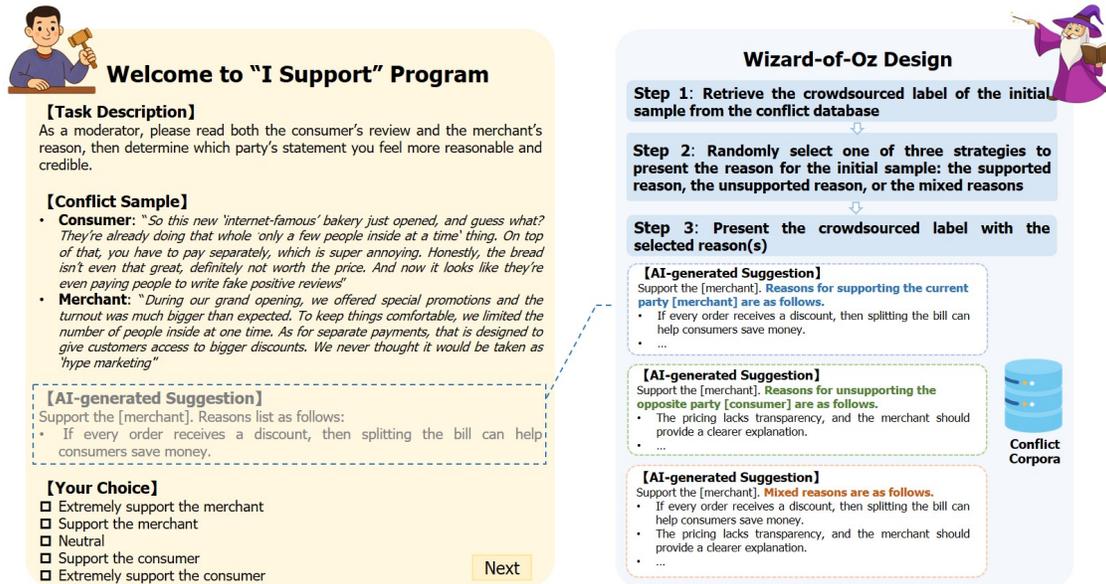


Figure 2. The design of "I Support" program (for presentation convenience, all descriptions are translated into English).

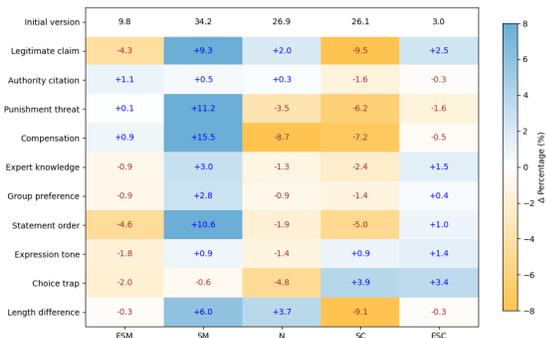


Figure 3. Choice variations without AI (initial version vs. perturbed version). For presentation convenience, "Extremely support the merchant", "Support the merchant", "Neutral", "Support the consumer", and "Extremely support the consumer" are abbreviated as "ESM", "SM", "N", "SC", and "ESC" respectively.

capabilities of different parties. This makes threats from the powerful party more likely to be recognized as credible signals.

**Compensation bias.** Human moderation shows a significant compensation bias ( $t(97) = 3.759, p < 0.001, d_z = 0.380$ ), where moderators tend to support the powerful party (the merchant). The proportion of choosing "Support the merchant" increases by 15.5%, while those of "Neutral" and "Support the consumer" decrease by 8.7% and 7.2%, respectively. Similarly, the phenomenon also remains regardless of which party's statement contains the compensation perturbation. The interview indicates the reasons for this phenomenon. Participants

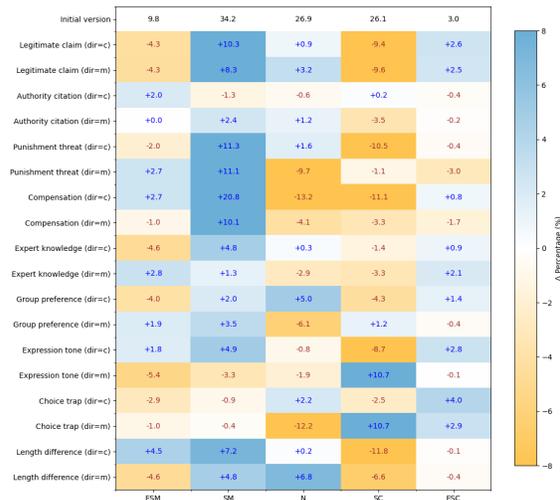


Figure 4. Choice variations without AI (initial version vs. perturbed version under two parties, using the same representations as Figure 3).

stated, "The merchant's compensation demonstrated his/her sincerity and a sense of responsibility" (P2), while "The discount request from the consumer is an unreasonable demand unrelated to the conflict fact." (P7) When coming from the powerful party, it represents a positive attitude to resolve the conflict, while when coming from the weaker party, it is more prone to being questioned as opportunism.

*Expert knowledge bias.* Human moderation presents a significant expert knowledge bias ( $t(93) = 2.357, p = 0.021, d_z = 0.243$ ), where moderators tend to support the powerful party (the merchant). Further analysis indicates a difference in terms of different parties. When the consumer's review inserts the expert knowledge, the proportion of choosing "Extremely support the merchant" decreases by 4.6%, while that of "Support the merchant" increases by 4.8%. When added in the merchant's response, the proportions of "Extremely support the merchant" and "Support the merchant" increased by 2.8% and 1.3%, respectively. This phenomenon suggests that moderators tend to trust professional expressions of merchants. Interviews also explain that the persuasiveness of expert knowledge is influenced by the party's identity or position. The professional terms used by merchants are often regarded as "detail supplementation" or "credible explanations", whereas similar expressions from consumers are easier to be interpreted as "intentional ornamentation" and "inconsistent with their level of knowledge." P5 remarked directly, "The professional terms cited by merchants are more persuasive, while consumers merely emphasize them deliberately to seek support." Professional explanations from the powerful party tend to be regarded as credible, while those from the weaker party may be seen as imitated or motivated by improper intentions.

*Length difference bias.* Human moderation shows a significant length difference bias ( $t(93) = 2.601, p = 0.011, d_z = 0.268$ ), where moderators tend to support the powerful party (the merchant). The proportion of choosing "Support the merchant" increases by 6.0%, while that of "Support the consumer" decreases by 9.1%. Furthermore, when consumers supply more content in their statements, the proportions of choosing "Extremely support the merchant" and "Support the merchant" increase by 4.5% and 7.2%, respectively. This suggests that moderators tend to support merchants even when consumers contain more content. The analysis also explains the reason behind it. Participants do not moderate based on the length, but make judgments by considering factors like content relevance and expression tone. Participants noted that longer content from merchants is perceived as "detailed process" or "more explanatory", whereas longer content from consumers contains "irrelevant", "emotional", or "exaggerated" content. Therefore, the impact of length difference is also influenced by content relevance and sources.

Prior work has shown that differences in demographics can lead to moderators' variations. So we conduct a more in-depth analysis by considering moderators' two basic demographic characteristics—age and gender. The samples are divided into younger and older groups based on a median age split (26 years old), and into male and female groups based on gender, respectively. Paired  $t$ -tests are then conducted between different groups, i.e., younger group VS older group, and male group VS female group. The results show that age and gender have different effects across different perturbations (see the Appendix). For age, both groups show significant effects under punishment threat and length difference perturbations. The younger group also exhibits significance under legitimate claim and

expert knowledge, while the older group shows significance under authority citation and compensation perturbations. For gender, both groups are affected by the punishment threat perturbation, with a stronger effect size in the female group. The female group also has significance under legitimate claim and punishment threat perturbations, while the male group presents significance under authority citation, compensation, expert knowledge, and length difference perturbations. These results suggest that different groups vary in their sensitivity to different perturbations. Younger users tend to focus on the logical consistency of content, and older users are more sensitive to authority or interest-related information. Females are more likely to pay attention to risk or social norms, while males focus more on authority or expertise cues.

**Overall, human moderation shows five power-related biases towards supporting the powerful party.**

## 5.2 RQ2: Power-related biases persist in human-AI moderation

*5.2.1 Human-AI moderation shows a tendency to support the powerful party under most perturbations.* In human-AI moderation, we find that five power manifestations can affect moderators' judgments (see the Appendix Table 4), including legitimate claim, authority citation, punishment threat, compensation, and length difference, triggering biases towards supporting the powerful party (the merchant). Since the interview analysis results under different perturbations are similar to those of the human moderation group, the following will not be presented to avoid repetition. Although other perturbations do not reach statistical significance, the proportion of choosing "Neutral" shows a declining trend, suggesting that these manifestations promote moderators to shift judgments from neutrality to clear stances.

*Legitimate claim bias.* Similar to human moderation, human-AI moderation presents a significant legitimate claim bias ( $t(94) = 4.375, p < 0.001, d_z = 0.521$ ), where moderators tend to support the powerful party (the merchant). As shown in Figure 5, the proportions of choosing "Extremely support the merchant" and "Support the merchant" increase by 6.9% and 3.4%, respectively, while those of "Neutral" and "Support the consumer" decrease by 1.9% and 8.5%, respectively. Furthermore, the effect remains regardless of which party's statement contains the legitimate claim perturbation (Figure 6).

*Authority citation bias.* Compared to human moderation, human-AI moderation presents a new bias - authority citation bias, with a significant level ( $t(96) = 4.375, p = 0.003, d_z = 0.309$ ), which indicates moderators tend to support the powerful party (the merchant). The proportion of choosing "Support the merchant" increases by 3.0%, while that of "Neutral" decreases by 3.1%. Furthermore, when the perturbation is inserted into the consumer's review, the proportion of "Support the consumer" increases by 3.3%, while that of "Neutral" decreases by 4.1%. Conversely, when similar content is added to the merchant's response, the proportions of "Extremely support the merchant" and "Support the merchant" increased by



Figure 5. Choice variations with AI (initial version vs. perturbed version, using the same representations as Figure 3).

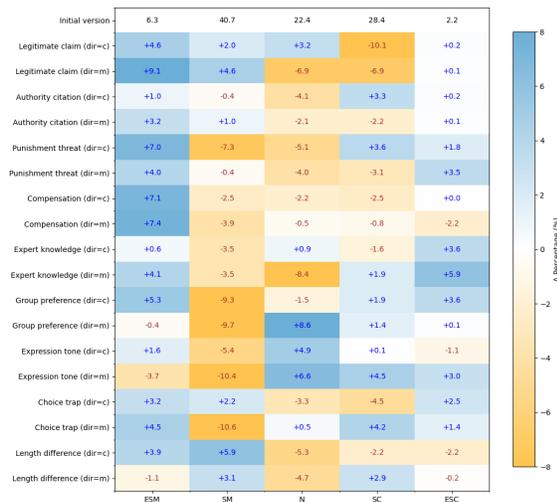


Figure 6. Choice variations with AI under two parties (initial version vs. perturbed version, using the same representations as Figure 4).

3.2% and 1.0%, respectively (Figure 6). This suggests that moderators tend to trust the party that contains authority citations. The interview analysis further indicates that the authority citation is regarded as a reliable cue. Even if the cited website is not verified, its existence can enhance moderators’ perceived reliability. P8 noted, “When I see citations, given the verification by AI, I perceive them as more trustworthy. Even if the link cannot be accessed, it may be visible on a particular platform.” Therefore, moderators’ judgments rely on only logical expressions but are also heuristically influenced by authoritative citations.

**Punishment threat bias.** Human-AI moderation exhibits a significant punishment threat bias ( $t(96) = 2.015, p = 0.047, d_z = 0.205$ ), where moderators tend to support the powerful party (the merchant). The proportions of choosing “Extremely

support the merchant” and “Extremely support the consumer” increase by 5.4% and 2.7%, respectively, while those of “Support the merchant” and “Neutral” decrease by 3.6% and 4.5%, respectively. When the perturbation is added to the consumer’s reviews, moderators tend to support the merchant.

**Compensation bias.** Human-AI moderation shows a significant compensation bias ( $t(99) = 3.718, p = 0.001, d_z = 0.358$ ), where moderators tend to support the powerful party (the merchant). The proportion of choosing “Extremely support the merchant” increases by 7.3%, while other options show a declining trend. Similarly, the effect remains consistent regardless of which party’s statement contains the compensation perturbation.

**Length difference bias.** Human-AI moderation presents a significant length difference bias ( $t(98) = 2.473, p = 0.012, d_z = 0.257$ ), where moderators tend to support the powerful party (the merchant). The proportions of choosing “Extremely support the merchant” and “Support the merchant” increase by 1.3% and 4.4%, respectively, while that of “Neutral” decreases by 5.0%. When there is more content in the consumer’s reviews, moderators tend to support the merchant. When merchants have more responses, moderators are prone to support consumers.

The perspective of what reasons are presented in AI’s suggestions can influence moderators’ judgments. Specifically, the reasons for unsupporting the opposite party have a significant effect ( $t(93) = -2.736, p = 0.008, d_z = -0.514$ ). The proportions of “Support the consumer” and “Extremely support the consumer” increase by 7.2% and 2.3%, respectively, while those of “Support the merchant” and “Neutral” decrease by 9.1% and 4.3%, respectively (Figure 7). This suggests that the reasons of this perspective can significantly shift moderators’ support to the consumer. Conversely, the reasons for supporting the current party do not show a significant difference, with the proportion of “Support the consumer” increasing by 2.4% and that of “Neutral” decreasing by 4.6%. Similarly, mixed reasons for both parties do not reach significance. In this case, the proportions of “Extremely support the merchant” and “Neutral” increase by 4.1% and 4.4%, respectively, while that of “Support the consumer” decreases by 8.3%. Interviews explain this. Most participants reported that the reasons of unsupporting the opposite party can “expose the logical issue directly” and encourage them to “think from the opposite perspective”. However, participants said the reasons for supporting the current party are “similar to my thoughts,” while mixed reasons are regarded as “too much information and lacking focus.”

**Therefore, human-AI moderation also exhibits power-related biases towards supporting the powerful party. Besides, the perspective for unsupporting the opposite party in AI-generated suggestions can motivate moderators to support the weaker party.**

5.2.2 AI alleviates most power-related biases while amplifying a few. The experimental results demonstrate that the introduction of AI has different impacts on moderators (see the Appendix Table 7) under different perturbations. For most



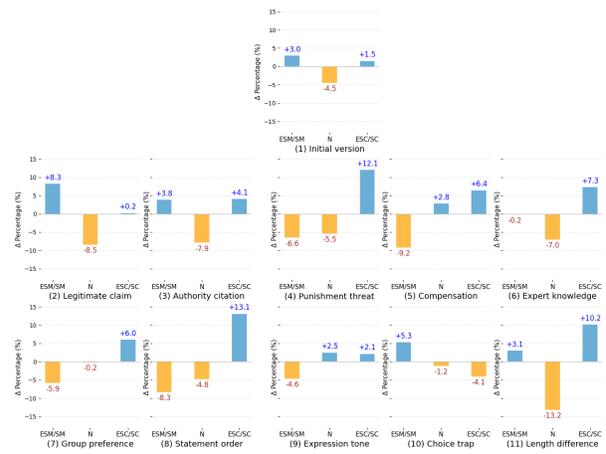
**Figure 7. Choice variations across three perspectives of reasons in AI-generated suggestions (initial version vs. perturbed versions, using the same representations as Figure 3).**

power-related biases existing in human moderation, the introduction of AI-generated suggestions alleviates moderators' biases in supporting the powerful party. However, under the legitimate claim condition, AI amplifies the existing bias.

*Alleviated biases.* The introduction of AI-generated suggestions alleviates several power-related biases in supporting the powerful party (the merchant), including punishment threat bias, compensation bias, expert knowledge bias, and length difference bias (Figure 8). For the punishment threat perturbation, AI-generated suggestions show statistical significance ( $t(86) = 2.520, p = 0.014, d_z = 0.270$ ). The proportion of “*Extremely/Support the consumer*” increases by 12.1%, while that of “*Neutral*” and “*Extremely/Support the merchant*” decreases by 5.5% and 6.6%, respectively. This suggests that the introduction of AI effectively alleviates the effect of merchants and shifts moderators' support towards consumers. The interviews further explain this. P12 noted, “*The reasons point to threatening expressions and remind me that the judgment should be unbiased rather than emotional influence.*” For the compensation perturbation, compared to human moderation, the proportion of “*Extremely/Support the consumer*” in human-AI moderation increases by 6.4%, while that of “*Extremely/Support the merchant*” decreases by 9.2%. Similar results are observed in expert knowledge and length difference perturbations. The interview analysis explains that AI-generated suggestions tend to be treated as reminders. When the expert knowledge from one party makes moderators confused, moderators can shift their focus to the conflict points identified by the AI, reducing the reliance on expert knowledge and identity. P10 said, “*The conflict points listed in LLM-generated suggestions enable me to directly compare the logic flaws in both parties' statements, rather than being swayed by professional terms.*”

*Introduced and Amplified biases.* As mentioned in the previous section, human-AI moderation introduces a new authority citation bias. Besides, the introduction of AI-generated suggestions amplifies the legitimate claim bias in supporting the

powerful party (the merchant). Compared with the results of human moderation, the proportion of “*Extremely/Support the merchant*” in human-AI moderation increases by 8.3%, while that of “*Neutral*” drops by 8.5%. This indicates that AI reduces neutral stances and strengthens the support for merchants, thus amplifying the tendency to favor the powerful party. The feedback analysis indicates that moderators believe that the rules cited by AI are credible. P9 supplemented, “*Since AI has cited rules, it indicates that they are reliable. So the claim by the merchant is more trustworthy.*”



**Figure 8. Choice variations with AI compared to without AI under different perturbed versions (using the same representations as Figure 3).**

*Perceptions of AI roles.* The interview analysis indicates, participants generally regard AI-generated suggestions as a reference. When their judgments are ambiguous or neutral, AI could help them form explicit stances quickly, saving time for thinking and reading. As P10 stated, “*The reason list supplements my thoughts from different perspectives. When facing ambiguous or lengthy statements from both parties, AI can quickly identify core conflict points.*” However, when AI-generated results are inconsistent with personal decisions, participants also express a cautious attitude. Some participants question AI's results and insist on their judgments. P12 noted, “*Sometimes AI also makes mistakes and lacks human empathy in ambiguous situations without standard answers.*”

*Expectations of AI assistance.* Our analysis also identifies users' expectations for future improvements: participants hope that AI can offer more targeted, analytical functions to better support their moderation, including providing legal evidence, highlighting conflict points, and simulating conflict processes, etc. For instance, P8 pointed out, “*AI should directly present details of relevant laws or rules, reducing burden of manual checks.*” P11 emphasized the need for a clearer comparative framework, “*AI should highlight logical contradictions in both parties' statements by changing colors or bolding texts.*” P13

further suggested dynamic scenario support, “*AI can simulate conflict prothe cess between both parties’ stances to help understand the scenario more intuitively.*”

**Overall, although AI assistance alleviates most biases of human moderation (four alleviated and one eliminated), it also amplifies a few (one introduced and another amplified).**

## 6 Discussion

Based on the above results, we identify several biases in power-asymmetric conflict moderation. In this section, we will discuss the potential mechanisms and effects of these biases. Based on that, multiple implications are proposed to inspire future research and design.

### 6.1 Findings

*6.1.1 Moderation biases stem from the perception of power asymmetry and heuristic thinking.* Our experimental results show that in both human and human-AI moderation, when conflict content contains perturbations including legitimate claim, punishment threat, compensation, and length difference, moderators tend to support the powerful party (the merchant). This indicates that bias in power-asymmetric conflicts is not incidental but rooted in long-standing social opinions. In other words, the bias does not arise from the labels “consumer” or “merchant”, but from the potential perception that merchants hold power and consumers are weaker in social interactions. However, the power relation between merchants and consumers is not fixed but can shift with factors such as evidence granularity, platform rules, or public discourse [39]. For instance, when public opinion favors supporting consumers, moderators’ perceptions that merchants are more powerful may be reduced. Such a perception difference can lead to the same language cues being interpreted differently on roles or contexts.

Power-related biases may arise from two main factors. First, default social perception of power asymmetry. The party with greater resources and interpretive authority is often viewed as more reliable. On e-commerce platforms, this perception privileges merchants, who are usually associated with symbols like brand reputation, service quality, and contract terms. When they explain or resolve conflicts with a service-oriented tone, their behaviors tend to be deemed legitimate. This reflects the System Justification Theory in social psychology, which suggests that individuals tend to hold existing social orders to reduce cognitive dissonance, even if it may harm personal interests [58]. This perception is deeply ingrained and rarely questioned. Second, moderators’ judgments tend to depend on heuristic thinking under constraints of limited information, heavy workload, and unverifiable evidence. It is typically a peripheral route processing according to the Elaboration Likelihood Model [90]. In moderating consumer-merchant conflicts, moderators are volunteers and lack sufficient time to conduct a comprehensive examination of the evidence presented by both parties, so expressions like rules or punishment threats are

perceived as surface cues by moderators for judgment. When default perceptions and heuristic thinking work together, the same language can be interpreted differently depending on the role: in merchants’ statements, they are perceived as conveying the meaning of rules compliance, formality, and sincerity; while in consumers’ statements, they are likely to be regarded as unreasonable, emotional, or manipulative.

*6.1.2 The perception of AI-generated suggestions: a double-edged sword.* Our research also reveals a double-edged sword role of the moderator’s perception of AI: it alleviates most biases that exist in human moderation, but also amplifies a few. Most power-related biases are mitigated to some extent under human-AI moderation, decreasing the supporting tendency towards the powerful party. The reasons might include two aspects. First, the suggestions help moderators quickly grasp core conflict points, reducing their reliance on heuristic cues like “*longer texts are more credible*”. Second, when moderators are neutral or lean towards merchants, AI-generated suggestions provide an alternative chance to reintroduce consumers’ perspectives into consideration, helping them break away from the thinking of default social perceptions and pay more attention to consumers’ statements. The situation is different in conflicts where the power relations are comparable, such as conflicts in online communities due to differences in contributions or conflicts in gaming communities caused by different understandings of rules [62]. There may be minor differences in action ability among parties, so moderators tend to focus more on the content itself rather than the party’s identity or power. In these scenarios, the main role of AI lies in helping moderators improve the moderation efficiency [117].

The effect of perceptions of AI is not always positive. Our study indicates that human-AI moderation amplifies the legitimate claim bias towards the powerful party and introduces a new authority citation bias. This may stem from moderators’ “legitimizing” effect on AI-generated content. When they encounter statements like rules in AI-generated suggestions, they potentially believe that the content is reliable. This effect creates a cumulative impact when moderators see the powerful party’s statements containing the legitimate claim or authority citation, and is further strengthened through heuristic thinking, enhancing trust in the powerful party. Our interview feedback echoes this conclusion, “*When I encounter citations from specific websites, given the verification by AI, I perceive them as more trustworthy.*” In essence, AI-generated suggestions are not only perceived as references but as a “completed verification,” weakening moderators’ motivation to re-examine. Besides, the perspective of what reasons are presented in AI-generated suggestions also shapes judgments. When suggestions adopt the perspective of unsupporting the opposite party, moderators tend to support the weaker party. The possible reason is that such unsupporting reasons highlight logic flaws in the opposing party’s statements, with a sense of accusation and negation, thus more likely to change moderators’ judgments. Conversely, other presentations (perspectives for supporting the current party and mixed perspectives) containing positive

arguments are less persuasive and insufficient to significantly change the original judgment.

These findings echo discussions in HCI regarding human-AI collaboration. People are often influenced by “illusion of validity” from LLMs [16, 59]. In scenarios like autonomous driving, seemingly plausible AI reminders are often mistakenly regarded as reliable [93]. In content understanding, deceptive explanations by LLMs can even enhance people’s trust in fake news compared to authentic explanations [16]. Our findings further extend current research. In power-asymmetric conflict moderation, AI does not merely serve the function of information assistance. Instead, its existence and presentation forms also shape moderators’ bias patterns. This double-edged sword effect highlights that future design of AI-assisted systems should extend beyond efficiency and incorporate mechanisms that alleviate moderation biases, thereby mitigating the impact of moderators’ “legitimizing” effect on AI.

**6.1.3 Commonalities and individual differences in power-asymmetric conflict moderation.** This study shows that moderators’ judgment of power-asymmetric conflicts follows a certain common pattern. They generally begin with looking at the consumer’s queries and check whether the merchant has addressed each conflict point. Then formulate an initial judgment by integrating the richness of details in aspects like product quality, platform rules, and service demands. For example, when a consumer claims, “*The staff had a bad attitude, and the food tasted ordinary,*” and the merchant responds with, “*All staff on duty that day have received training, and attached screenshots to prove the staff have been smiling throughout service. Besides, our food is prepared with standardized procedures...*” then the consumer’s claim “*bad attitude*” appears insufficient due to insufficient detail. Conversely, the merchant provides targeted responses to conflict points of service attitude and food quality. Tone also impacts moderators’ judgments. Objective and structured statements are deemed more persuasive, while emotional statements tend to trigger doubts. Overall, moderators tend to follow a specific pattern in moderation. They first check whether the responses match, then evaluate details and tone, and finally form a judgment.

However, moderators’ judgments exhibit some individual differences, which are related to factors such as their experience, demographics, and values. First, identity and personal experience can affect moderation. Moderators not only play the role of third-party judges but are also consumers in daily lives. This two-fold identity can subtly influence judgments. P2 remarked, “*I have personal experience of staying in hotels. They usually provide pick-up and drop-off services. It’s unreasonable for this hotel to claim otherwise.*” When P5 judged the same case, he stated, “*The primary function of a hotel is to provide accommodation. Pick-up and drop-off services are supplementary. If it is clearly communicated before booking that the service is not available, there’s no issue.*” This aligns with Instance-Based Learning Theory, which suggests that individuals adjust strategies based on the similarity between the current situation and past experience [32]. Second, judgment can be influenced by

human values. Moderators who prioritize the protection of the weaker party tend to support consumers, while those who emphasize order and institutional stability tend to support merchants. As P7 noted, “*These consumers’ reviews should be posted because others can gain insights into the service details of the merchant,*” while P9 said, “*Reviews must be grounded in facts and should not be emotional complaints.*” According to Schwartz’s value model, values serve as guiding principles in individual behavior. When statements align with moderators’ value system, they tend to view them as “*reasonable*” or “*acceptable*”. When the value claims in statements are inconsistent with their values, they tend to question or deny those statements. Third, differences in moderators’ imaginations on AI can influence moderation. These imaginations stem not from understandings of how AI actually works, but from the mental models of AI formed by moderators through daily technology use. Such mental models of AI refer to people’s internal beliefs about what AI can do, how it makes judgments, and why it produces certain outputs [29, 83]. Prior work shows that AI literacy shapes these mental models of AI. A higher level of AI literacy enables appropriate reliance on automated systems, while lower literacy exhibits over-reliance or excessive suspicion [22, 41]. This point is reflected in our interview. For example, P9 stated, “*As an important information assistant, LLM-driven AI now solves many problems in my daily work.*” Conversely, P12 commented, “*AI just grabs content from the internet and sometimes even fabricates it. It’s not reliable.*” Thus, varying levels of technical literacy influence how moderators imagine and trust AI, which in turn affects their moderation judgments.

## 6.2 Design implications

Based on the above findings, we propose the following design implications.

- Alleviate moderation biases stemming from default power asymmetry.** In power-asymmetric conflicts, moderation should pursue not only efficiency but also alleviate moderation biases. Two future directions are suggested based on our research. First, anonymization and symmetrical information presentation can be considered to reduce biases associated with roles. Conflicts between consumers and merchants can be reframed as “*Subject A*” and “*Subject B*,” with statements presented in a standardized format such as “*Assertion/Evidence/Claim/Other Language Cues*.” Automated sentence processing can further distill key information to minimize potential biases related to text length or rhetorical styles. This strategy echoes prior HCI research regarding the use of interface design to alleviate heuristic biases [44, 70, 101]. Second, bias reminders and counterfactual comparison mechanisms can be explored to strengthen moderators’ meta-awareness. When language cues like compensation or punishment threat appear in the powerful party’s statements, the system can generate two summaries with and without

these cues to highlight the potential influence of judging from surface signals. The approach aligns with principles of explainability and transparency emphasized in the content moderation field [51]. These features can help mitigate moderation biases stemming from default power asymmetry.

- **Control the “legitimizing” effect of AI-generated content on moderators.** Moderators often perceive AI-generated suggestions containing authority citation or rule details as verified evidence, reflecting a “legitimizing” effect. Future moderation systems are suggested to regulate this influence. First, uncertainty annotations like confidence intervals and multiple interpretations can be leveraged to replace deterministic expressions in systems, prompting moderators to think critically [75, 81]. Second, transparency can be enhanced through evidence chain tracing to prevent moderators’ misinterpretation of AI outputs as “*completed verification*.” For instance, add “*to be verified*” indicators and visually depict evidence sources or reasoning paths. Similar designs have been proven effective in reducing overreliance on AI in medical diagnosis [8], aligning with the call for moderate reliance in HCI [69]. These features do not undermine AI’s value but support balanced judgments that combine efficiency and alleviate moderation biases.
- **Improve moderation by incorporating both commonality and individuality.** Moderators under power-asymmetric conflicts show similar moderation patterns but also exhibit individual variances, indicating that systems should balance consistency with diversity. First, structured support can standardize core procedures. Our research shows moderators typically follow a sequence of “matching responses to conflict points, verifying details, assessing tone.” Introducing interactive checklists in moderation systems can guide this process by ensuring each conflict point is addressed. Second, systems should consider personal experiences and value orientations. Prior work also confirms that individual values shape moderation [54]. Therefore, functions like experience annotation (document past experiences) and configurable value-orientation settings (e.g., “*emphasis on norms*”) can capture these differences. Besides visualizations of distribution proportion, the group intelligence can also aggregate orientations using weighting algorithms, while measures like delayed or anonymous display can mitigate conformity effects [20, 106]. Third, it suggests considering variations in AI literacy. Hierarchical explanation interfaces can provide both concise and detailed justifications. Moreover, educational prompts upon initial use (e.g., “*the suggestions are only for reference!*”) and simulation practice on historical conflicts can further help moderators with lower AI literacy in shaping critical thinking.

### 6.3 Generalizability

This section discusses the generalizability of our work in terms of research findings and experimental design. Focusing on moderation bias in online power-asymmetric conflicts, we leveraged the consumer-merchant conflicts as a representative scenario, but the findings can be generalized into other contexts. First, the tendency of third-party moderators to support the powerful party is not limited to this scenario. In medical disputes, arbitrators might tend to accept doctors’ interpretations due to their professional authority [8]. In education, evaluators may support teachers because of their knowledge and institutional standing [37]. These phenomena, similar to consumer-merchant conflicts, show that power asymmetry has a universal impact on third-party moderators. Thus, our findings can contribute understanding of bias patterns in these contexts. Second, the “legitimizing” effect of AI-generated content is also reflected in other decision-making scenarios. For example, doctors may rely on LLMs’ outputs to reduce workload but overlook patient differences, and HR personnel amplify gender and racial biases when using LLMs for resume screening [84]. These show that AI-generated content containing plausible language may exacerbate misjudgments, offering implications for moderation and broader human–AI collaboration.

Moreover, our experimental design demonstrates generalizability and can be applied to other studies on human/AI bias and conflict judgments. First, the intervention approach used to examine bias is not limited to online consumer-merchant conflicts, but can be transferred to scenarios like workplace negotiations, community governance, and platform mediation. Second, the Wizard-of-Oz design (researchers acting as the “AI”) eliminates the differences in models’ outputs, allowing us to focus on moderators’ bias patterns rather than technical constraints. This offers a typical framework for studying how people perceive AI-generated content and can inform the design of AI-assisted systems independent of specific models.

## 7 Limitations and Future Work

This work has the following limitations. First, our study leverages consumer-merchant conflicts as a representative scenario of power asymmetry. While this scenario offers accessible data and a standardized resolution process, just the one scenario may limit the generalizability and reliability of our findings. Other representative contexts like workplace negotiations, community governance, or medical disputes involve multiple rounds of interaction, implicit power dynamics, and moral considerations, which may trigger different biases. Future research should validate the findings in more contexts. Furthermore, our data was obtained from Dianping, a Chinese local life service platform, with all participants from China. While this ensures contextual consistency, it limits cross-cultural generalizability. Prior studies have shown that content moderation can be influenced by cultural differences, including identity characteristics, ideology, and language, etc. [36, 86, 119]. These findings imply that the moderation biases identified in our

study could not be generalized to all cultures. So future studies should be carried out in diverse cultures to examine the universality and boundaries of the bias mechanisms uncovered in this work.

Besides, our study does not measure the participants' technical literacy level, which limits the ability to assess how it affects their imaginations of AI and moderation judgments. Future work could consider different levels of technical literacy among moderators to further explore how these differences shape their trust in AI during moderation. Moreover, to avoid technical differences, we adopt a Wizard-of-Oz method wherein researchers act as "AI" to provide suggestions, inevitably departing from real usage scenarios of AI models. Particularly, different LLMs' outputs vary significantly in tone, accuracy, hallucination frequency, and even minor changes in prompts can lead to different responses, affecting users' trust [13, 107, 109]. Suggestions from real LLMs may alter the manifestation of moderators' biases due to their expression style and output stability. More authoritative expressions may amplify biases, while inconsistent outputs among different LLMs could undermine moderators' trust. Future research could integrate real AI models for comparison and the corresponding influence on moderation.

## 8 Conclusion

In this paper, we investigated the power-related biases of human and human-AI moderation in power-asymmetric conflicts. Through a mixed design experiment involving 50 participants, we found that both moderation modes exhibit several power-related biases towards supporting the powerful party. Although AI-generated suggestion alleviates most biases of human moderation, it also amplifies a few. Moreover, presenting the reasons of unsupporting the opposite party in these suggestions also prompts moderators to support the weaker party to a certain extent. These findings provide several insights into the design of power-asymmetric conflict moderation mechanisms and AI-assisted moderation systems. In future work, we will further explore the differences in moderation biases across different cultural contexts and explore moderation designs as well as human-AI collaboration strategies to mitigate these biases.

## Acknowledgments

This research was supported by Meituan and National Natural Science Foundation of China (NSFC) under the Grant No. 62372113. Peng Zhang is a faculty of College of Computer Science and Artificial Intelligence, Fudan University. Tun Lu is a faculty of College of Computer Science and Artificial Intelligence, Fudan University, Shanghai Key Laboratory of Data Science, and Silver-X MOE Philosophy & Social Sciences Laboratory, Fudan Institute on Aging.

## References

- [1] Deema Alkilani. 2012. Power, Six Bases of Defining The Bases of Social Power That Leads to Socially Independent Change. *Encyclopedia of Leadership* (2012), 1242–1249.
- [2] Saar Alon-Barkat and Madalina Busuioac. 2023. Human-AI Interactions in Public Sector Decision Making: "Automation Bias" and "Selective Adherence" to Algorithmic Advic. *Journal of Public Administration Research and Theory* 33, 1 (2023), 153–169.
- [3] Vallet Améline, Locatelli Bruno, Barnaud Cécile, Makowski David, Quispe Conde Yésica, and Haroldv Levrel. 2020. Power Asymmetries in Social Networks of Ecosystem Services Governance. *Environmental Science & Policy* 114 (2020), 329–340.
- [4] Shestakofsky Benjamin. 2019. Book Review: Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass. By Mary L. Gray and Siddharth Suri. *ILR Review* 72, 5 (2019), 1283–1285.
- [5] Jules Brooks Blanchet, Megan E. Hillis, Yeongji Lee, Qijia Shao, Xia Zhou, Devin Balkcom, and David J. M. Kraemer. 2025. Enhancing the Educational Potential of Online Movement Videos: System Development and Empirical Studies with TikTok Dance Challenges. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Article 587, 19 pages.
- [6] Nattapat Boonprakong, Benjamin Tag, Jorge Goncalves, and Tilman Dingler. 2025. How Do HCI Researchers Study Cognitive Biases? A Scoping Review. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Article 473, 20 pages.
- [7] Michelle Brachman, Zahra Ashktorab, Michael Desmond, Evelyn Duesterwald, Casey Dugan, Narendra Nath Joshi, Qian Pan, and Aabhas Sharma. 2022. Reliance and Automation for Human-AI Collaborative Data Labeling Conflict Resolution. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2, Article 321 (2022), 27 pages.
- [8] Adrian Bussone, Simone Stumpf, and Dymna O'Sullivan. 2015. The Role of Explanations on Trust and Reliance in Clinical Decision Support Systems. In *2015 International Conference on Healthcare Informatics*. IEEE, 160–169.
- [9] Matt Canute, Mali Jin, Hannah Holtzclaw, Alberto Lusoli, Philippa Adams, Mugdha Pandya, Maite Taboada, Diana Maynard, and Wendy Hui Kyong Chun. 2023. Dimensions of Online Conflict: Towards Modeling Agonism. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. 12194–12209.
- [10] Guiming Hardy Chen, Shunian Chen, Ziche Liu, Feng Jiang, and Benyou Wang. 2024. Humans or LLMs as the Judge? A Study on Judgement Bias. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. Miami, Florida, USA, 8301–8327.
- [11] XinHui Chen, Xiang Yuan, Hui Zhang, Ruixiao Zheng, and Wanyi Wei. 2025. Maintaining "Balanced" Conflict: Proactive Intervention Strategies of AI Voice Agents in Online Collaboration of Temporary Design Teams. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Article 907, 19 pages.
- [12] Yu-Ling Chou, Hsuan-Jen Lee, Jie Tsai, En-Chi Liang, and Yung-Ju Chang. 2023. I Like Their Autonomy and Closeness to Me: Uncovering the Perceived Appeal of Social-Media Influencers. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Article 674, 19 pages.
- [13] Michelle Cohn, Mahima Pushkarna, Gbolahan O. Olanubi, Joseph M. Moran, Daniel Padgett, Zion Mengesha, and Courtney Heldreth. 2024. Believing Anthropomorphism: Examining the Role of Anthropomorphic Cues on Trust in Large Language Models. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. Article 54, 15 pages.
- [14] Petru L Curseu, Patrick Kenis, Jorg Raab, and Ulrik Brandes. 2010. Composing Effective Teams through Team Dating. *Organization Studies* 31, 7 (2010), 873–894.
- [15] Robert A Dahl. 1957. The Concept of Power. *Behavioral Science* 2, 3 (1957), 201–215.
- [16] Valdemar Danry, Pat Pataranutaporn, Matthew Groh, and Ziv Epstein. 2025. Deceptive Explanations by Large Language Models Lead People to Change their Beliefs About Misinformation More Often than Honest Explanations. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Article 933, 31 pages.
- [17] Janice Denegri-Knott. 2006. Consumers Behaving Badly: Deviation or Innovation? Power Struggles on The Web. *Journal of Consumer Behaviour: An International Research Review* 5, 1 (2006), 82–94.
- [18] Djellel Difallah, Elena Filatova, and Panos Ipeirotis. 2018. Demographics and Dynamics of Mechanical Turk Workers. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM '18)*. 135–143.
- [19] Kaushik Divyansh, Hovy Eduard, and Lipton Zachary. 2020. Learning The Difference That Makes A Difference With Counterfactually-Augmented Data. In *International Conference on Learning Representations*.

- [20] Claire Dolin, Ben Weinschel, Shawn Shan, Chang Min Hahn, Eurim Choi, Michelle L. Mazurek, and Blase Ur. 2018. Unpacking Perceptions of Data-Driven Inferences Underlying Online Targeting and Personalization. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. 1–12.
- [21] S. M. Easterbrook, E. E. Beck, J. S. Goodlet, L. Plowman, M. Sharples, and C. C. Wood. 1993. *A Survey of Empirical Studies of Conflict*. Springer London, 1–68.
- [22] Upol Ehsan, Samir Passi, Q. Vera Liao, Larry Chan, I-Hsiang Lee, Michael Muller, and Mark O Riedl. 2024. The Who in XAI: How AI Background Shapes Perceptions of AI Explanations. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. Article 316, 32 pages.
- [23] Richard M Emerson. 2019. Power-dependence Relations. In *Power in Modern Societies*. Routledge, 48–58.
- [24] Xianzhe Fan, Qing Xiao, Xuhui Zhou, Jiaxin Pei, Maarten Sap, Zhicong Lu, and Hong Shen. 2025. User-Driven Value Alignment: Understanding Users' Perceptions and Strategies for Addressing Biased and Discriminatory Statements in AI Companions. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Article 910, 19 pages.
- [25] Franz Faul, Edgar Erdfelder, Axel Buchner, and Albert-Georg Lang. 2009. Statistical Power Analyses Using G\* Power 3.1: Tests for Correlation and Regression Analyses. *Behavior Research Methods* 41, 4 (2009), 1149–1160.
- [26] Colin Ford, Dan Gardner, Leah Elaine Horgan, Calvin Liu, a. m. tsaasan, Bonnie Nardi, and Jordan Rickman. 2017. Chat Speed OP PogChamp: Practices of Coherence in Massive Twitch Chat. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17)*. 858–871.
- [27] Kyriaki Fousiani. 2022. Power Asymmetry, Negotiations and Conflict Management in Organizations. *Organizational Conflict* (2022).
- [28] John RP French, Bertram Raven, and Dorwin Cartwright. 1959. The Bases of Social Power. *Studies in Social Power* 150 (1959), 167.
- [29] Katy Ilonka Gero, Zahra Ashktorab, Casey Dugan, Qian Pan, James Johnson, Werner Geyer, Maria Ruiz, Sarah Miller, David R. Millen, Murray Campbell, Sadhana Kumaravel, and Wei Zhang. 2020. Mental Models of AI Agents in a Cooperative Game Setting. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [30] Ysabel Gerrard. 2020. Behind The Screen: Content Moderation in The Shadows of Social Media.
- [31] Peter Glick and Susan T Fiske. 1996. The Ambivalent Sexism Inventory: Differentiating hostile and benevolent sexism. *Journal of Personality and Social Psychology* 70, 3 (1996), 491–512.
- [32] Cleotilde Gonzalez, Javier F Lerch, and Christian Lebiere. 2003. Instance-based Learning in Dynamic Decision Making. *Cognitive Science* 27, 4 (2003), 591–635.
- [33] Jarod Govers, Eduardo Velloso, Vassilis Kostakos, and Jorge Goncalves. 2024. AI-Driven Mediation Strategies for Audience Depolarisation in Online Debates. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Article 803, 18 pages.
- [34] James Grimmelmann. 2015. The Virtues of Moderation. *Yale J L & Tech*. 17 (2015), 42.
- [35] Ken Gu, Madeleine Grunde-McLaughlin, Andrew McNutt, Jeffrey Heer, and Tim Althoff. 2024. How Do Data Analysts Respond to AI Assistance? A Wizard-of-Oz Study. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Article 1015, 22 pages.
- [36] Oliver L. Haimson, Daniel Delmonaco, Peipei Nie, and Andrea Wegner. 2021. Disproportionate Removals and Differing Content Moderation Experiences for Conservative, Transgender, and Black Social Media Users: Marginalization and Moderation Gray Areas. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2, Article 466 (2021), 35 pages.
- [37] Bingyi Han, Simon Coghlan, George Buchanan, and Dana McKay. 2025. Who is Helping Whom? Student Concerns about AI-Teacher Collaboration in Higher Education Classrooms. *Proceedings of the ACM on Human-Computer Interaction* 9, 2, Article CSCW206 (2025), 32 pages.
- [38] Gaole He, Gianluca Demartini, and Ujwal Gadiraju. 2025. Plan-Then-Execute: An Empirical Study of User Trust and Team Performance When Using LLM Agents As A Daily Assistant. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Article 414, 22 pages.
- [39] Natali Helberger, Orla Lynskey, Hans-W. Micklitz, Peter Rott, Marijn Sax, and Joanna Strycharz. 2021. *EU Consumer Protection 2.0: Structural Asymmetries in Digital Consumer Markets*. Technical Report. Institute for Information Law & Amsterdam School of Communication Research.
- [40] Vali Hessamedin, Jingjun Xu David, and Yildirim Mehmet. 2015. The Effect of Conflicting Consumer Reviews on the Accuracy of a Purchase Decisions. In *The 21st Americas Conference on Information Systems, AMCIS 2015*.
- [41] Michael C Horowitz and Lauren Kahn. 2024. Bending the Automation Bias Curve: A Study of Human and AI-based Decision Making in National Security Contexts. *International Studies Quarterly* 68, 2 (2024), sqae020.
- [42] Yoyo Tsung-Yu Hou, EunJeong Cheon, and Malte F. Jung. 2024. Power in Human-Robot Interaction. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*. 269–282.
- [43] Wenjian Huang, Tun Lu, Haiyi Zhu, Guo Li, and Ning Gu. 2016. Effectiveness of Conflict Management Strategies in Peer Review Process of Online Collaboration Projects. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. 717–728.
- [44] Jevan A Hutson, Jessie G Taft, Solon Barocas, and Karen Levy. 2018. Debiasing Desire: Addressing Bias & Discrimination on Intimate Platforms. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–18.
- [45] Angel Hsing-Chi Hwang and Andrea Stevenson Won. 2021. IdeaBot: Investigating Social Facilitation in Human-Machine Team Creativity. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Article 712, 16 pages.
- [46] Muhammad Okky Ibrohim and Indra Budi. 2019. Multi-label Hate Speech and Abusive Language Detection in Indonesian Twitter. In *Proceedings of the 3rd Workshop on Abusive Language Online*. 46–57.
- [47] Ridho Bramulya Ikhsan, Yudi Fernando, Anderes Gui, and Erick Fernando. 2024. The Power of Online Reviews: Exploring Information Asymmetry and Its Impact on Green Product Purchasing Behavior. *International Journal of Consumer Studies* 48, 3 (2024), e13050.
- [48] Q4 2023 Digital Trust & Safety Index. 2024. Dispute Data, Consumer Insights, and Emerging Trends. [https://sift.com/index-reports-disputes-q4-2023/?utm\\_source=chatgpt.com](https://sift.com/index-reports-disputes-q4-2023/?utm_source=chatgpt.com).
- [49] Krithika Jagannath, Katie Salen, and Petr Slovák. 2020. "(We) Can Talk It Out...": Designing for Promoting Conflict-Resolution Skills in Youth on a Moderated Minecraft Server. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW1, Article 49 (2020), 26 pages.
- [50] Surowiecki James. 2004. *The wisdom of crowds: Why the Many Are Smarter than the Few and How Collective Wisdom Shapes Business, Economies, Societies, and Nations*. Doubleday & Co.
- [51] Shagun Jhaver, Amy Bruckman, and Eric Gilbert. 2019. Does Transparency in Moderation Really Matter? User Behavior after Content Removal Explanations on Reddit. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–27.
- [52] Shagun Jhaver, Sucheta Ghoshal, Amy Bruckman, and Eric Gilbert. 2018. Online Harassment and Content Moderation: The Case of Blocklists. *ACM Transactions on Computer-Human Interaction* 25, 2, Article 12 (2018), 33 pages.
- [53] Jialun Aaron Jiang, Peipei Nie, Jed R. Brubaker, and Casey Fiesler. 2023. A Trade-off-centered Framework of Content Moderation. *ACM Transactions on Computer-Human Interaction* 30, 1, Article 3 (2023), 34 pages.
- [54] Jialun Aaron Jiang, Peipei Nie, Jed R. Brubaker, and Casey Fiesler. 2023. A Trade-off-centered Framework of Content Moderation. *ACM Transactions on Computer-Human Interaction* 30, 1, Article 3 (2023), 34 pages.
- [55] Jigsaw. 2023. Perspective API.
- [56] Seering Joseph, Kaufman Geoff, and Chancellor Stevie. 2022. Metaphors in Moderation. *New Media & Society* 24, 3 (2022), 621–640.
- [57] Seering Joseph, Wang Tony, Yoon Jina, and Kaufman Geoff. 2019. Moderator Engagement and Community Development in The Age of Algorithms. *New Media & Society* 21, 7 (2019), 1417–1443.
- [58] John T Jost and Jojanneke Van der Toorn. 2012. System Justification Theory. *Handbook of Theories of Social Psychology* 2 (2012), 313–343.
- [59] Harmanpreet Kaur, Harsha Nori, Samuel Jenkins, Rich Caruana, Hanna Wallach, and Jennifer Wortman Vaughan. 2020. Interpreting Interpretability: Understanding Data Scientists' Use of Interpretability Tools for Machine Learning. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. 1–14.
- [60] Sara Kiesler, Robert Kraut, Paul Resnick, and Aniket Kittur. 2012. Regulating behavior in online communities. Building Successful Online Communities: Evidence-Based Social Design.
- [61] Soomin Kim, Jinsu Eun, Joseph Seering, and Joohwan Lee. 2021. Moderator Chatbot for Deliberative Discussion: Effects of Discussion Structure and Discussant Facilitation. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1, Article 87 (2021), 26 pages.
- [62] Aniket Kittur and Robert E. Kraut. 2010. Beyond Wikipedia: Coordination and Conflict in Online Production Groups. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work*. 215–224.
- [63] Aniket Kittur, Bongwon Suh, Bryan A. Pendleton, and Ed H. Chi. 2007. He Says, She Says: Conflict and Coordination in Wikipedia. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI*

- '07). 453–462.
- [64] Rob Kling. 1996. *Computerization and Controversy: Value Conflicts and Social Choices*. Elsevier.
- [65] Samantha Kolovson, Calvin Liang, Sean A. Munson, and Kate Starbird. 2020. Personal Data and Power Asymmetries in U.S. Collegiate Sports Teams. *Proceedings of the ACM on Human-Computer Interaction* 4, GROUP, Article 05 (2020), 27 pages.
- [66] Yubo Kou and Xinning Gui. 2020. Mediating Community-AI Interaction through Situated Explanation: The Case of AI-Led Moderation. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2, Article 102 (2020), 27 pages.
- [67] Isadora Krsek, Kimi Wenzel, Sauvik Das, Jason I. Hong, and Laura Dabbish. 2022. To Self-Persuade or be Persuaded: Examining Interventions for Users' Privacy Setting Selection. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Article 623, 17 pages.
- [68] Vivian Lai, Samuel Carton, Rajat Bhatnagar, Q. Vera Liao, Yunfeng Zhang, and Chenhao Tan. 2022. Human-AI Collaboration via Conditional Delegation: A Case Study of Content Moderation. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Article 54, 18 pages.
- [69] John D Lee and Katrina A See. 2004. Trust in automation: Designing for appropriate reliance. *Human Factors* 46, 1 (2004), 50–80.
- [70] Kyungho Lee. 2025. Towards a Working Definition of Designing Generative User Interfaces. In *Companion Publication of the 2025 ACM Designing Interactive Systems Conference*. 489–495.
- [71] Sharon Levy, Robert E. Kraut, Jane A. Yu, Kristen M. Altenburger, and Yi-Chia Wang. 2022. Understanding Conflicts in Online Conversations. In *Proceedings of the ACM Web Conference 2022 (WWW '22)*. 2592–2602.
- [72] Yaqiong Li, Peng Zhang, Hansu Gu, Tun Lu, Siyuan Qiao, Yubo Shu, Yiyang Shao, and Ning Gu. 2025. DeMod: A Holistic Tool with Explainable Detection and Personalized Modification for Toxicity Censorship. *Proceedings of the ACM on Human-Computer Interaction* 9, 2, Article CSCW061 (May 2025), 24 pages.
- [73] Yaqiong Li, Peng Zhang, Lin Wang, Hansu Gu, Siyuan Qiao, Ning Gu, and Tun Lu. 2025. "Harmless to You, Hurtful to Me!": Investigating the Detection of Toxic Languages Grounded in the Perspective of Youth. *ArXiv abs/2508.02094* (2025).
- [74] Baoxi Liu, Peng Zhang, Yubo Shu, Zhengqing Guan, Tun Lu, Hansu Gu, and Ning Gu. 2022. Building a Personalized Model for Social Media Textual Content Censorship. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2, Article 499 (Nov. 2022), 31 pages.
- [75] Zhuoran Lu, Patrick Li, Weilong Wang, and Ming Yin. 2025. Understanding the Effects of AI-based Credibility Indicators When People Are Influenced By Both Peers and Experts. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Article 817, 19 pages.
- [76] M Lynne Markus. 1994. Finding A Happy Medium: Explaining the Negative Effects of Electronic Communication on Social Life at Work. *ACM Transactions on Information Systems (TOIS)* 12, 2 (1994), 119–149.
- [77] David Martin, Benjamin V. Hanrahan, Jacki O'Neill, and Neha Gupta. 2014. Being A Turker. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '14)*. 224–235.
- [78] David Maulsby, Saul Greenberg, and Richard Mander. 1993. Prototyping An Intelligent Agent Through Wizard of Oz. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems (CHI '93)*. 277–284.
- [79] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW, Article 72 (2019), 23 pages.
- [80] Matthew B Miles. 1994. *Qualitative Data Analysis: An Expanded Sourcebook*. Thousand Oaks (1994).
- [81] Tim Miller. 2023. Explainable AI is Dead, Long Live Explainable AI! Hypothesis-driven Decision Support using Evaluative AI. In *Proceedings of the 2023 ACM conference on fairness, accountability, and transparency*. 333–342.
- [82] Shyamal Mishra and Preetha Chatterjee. 2024. Exploring ChatGPT for Toxicity Detection in GitHub. In *Proceedings of the 2024 ACM/IEEE 44th International Conference on Software Engineering: New Ideas and Emerging Results (ICSE-NIER'24)*. 6–10.
- [83] Vikram Mohanty, Jude Lim, and Kurt Luther. 2025. What Lies Beneath? Exploring the Impact of Underlying AI Model Updates in AI-Infused Systems. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Article 539, 21 pages.
- [84] Dena F. Mujtaba and Nihar Ranjan Mahapatra. 2024. Fairness in AI-Driven Recruitment: Challenges, Metrics, Methods, and Future Directions. *ArXiv abs/2405.19699* (2024).
- [85] Isabel Munoz, Michael Dunn, Steve Sawyer, and Emily Michaels. 2022. Platform-mediated Markets, Online Freelance Workers and Deconstructed Identities. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–24.
- [86] Gianluca Nogara, Francesco Pierri, Stefano Cresci, Luca Luceri, Petter Törnberg, and Silvia Giordano. 2025. Toxic Bias: Perspective API Misreads German as More Toxic. *Proceedings of the International AAAI Conference on Web and Social Media* 19, 1 (2025), 1346–1357.
- [87] Gilbert N. Nyaga, Daniel F. Lynch, Donna Marshall, and Eamonn Ambrrose. 2013. Power Asymmetry, Adaptation and Collaboration in Dyadic Relationships Involving a Powerful Partner. *Journal of Supply Chain Management* 49, 3 (2013), 42–65.
- [88] Ben Hill Passmore. 2015. Inductive Reasoning, Information Symmetry, and Power Asymmetry in Organizations. In *Encyclopedia of Information Science and Technology, Third Edition*. IGI Global Scientific Publishing, 808–818.
- [89] Ken Petress. 2003. Power: Definition, Typology, Description, Examples, and Implications. <http://utshcsa.edu/gme/documents/powerdefinitionstypologyexamples.pdf> (2003).
- [90] Richard E Petty and John T Cacioppo. 1986. The Elaboration Likelihood Model of Persuasion. In *Advances in Experimental Social Psychology*. Vol. 19. Elsevier, 123–205.
- [91] Snehal Prabhudesai, Ananya Prashant Kasi, Anmol Manshing, Anindya Das Antar, Hua Shen, and Nikola Banovic. 2025. "Here the GPT made a choice, and every choice can be biased": How Students Critically Engage with LLMs through End-User Auditing Activity. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Article 1015, 23 pages.
- [92] Gerald Peter Quinn and Michael J Keough. 2002. *Experimental Design and Data Analysis for Biologists*. Cambridge University Press.
- [93] Emely Rosbach, Jonas Ammeling, Sebastian Krügel, Angelika Kießig, Alexis Fritz, Jonathan Ganz, Chloé Puget, Taryn Donovan, Andrea Klang, Maximilian C. Köller, Pompei Bolfa, Marco Tecilla, Daniela Denk, Matti Kiupel, Georgios Paraschou, Mun Keong Kok, Alexander F. H. Haake, Ronald R. de Krijger, Andreas F.-P. Sonnen, Tanit Kasantikul, Gerry M. Dorresteijn, Rebecca C. Smedley, Nikolas Stathonikos, Matthias Uhl, Christof A. Bertram, Andreas Riener, and Marc Aubreville. 2025. "When Two Wrongs Don't Make a Right" - Examining Confirmation Bias and the Role of Time Pressure During Human-AI Collaboration in Computational Pathology. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Article 528, 18 pages.
- [94] Akрати Saxena and Harita Reddy. 2022. Users Roles Identification on Online Crowdsourced Q&A Platforms and Encyclopedias: A Survey. *Journal of Computational Social Science* 5, 1 (2022), 285–317.
- [95] Tobias Schnabel, Saleema Amershi, Paul N Bennett, Peter Bailey, and Thorsten Joachims. 2020. The Impact of More Transparent Interfaces on Behavior in Personalized Recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 991–1000.
- [96] Jakob Schoeffer, Maria De-Arteaga, and Niklas Kühl. 2024. Explanations, Fairness, and Appropriate Reliance in Human-AI Decision-Making. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Article 836, 18 pages.
- [97] Joseph Seering. 2020. Reconsidering Self-Moderation: the Role of Research in Supporting Community-Based Models for Online Content Moderation. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2, Article 107 (2020), 28 pages.
- [98] Farhana Shahid. 2024. Human-AI Collaboration to Facilitate Culturally-Aware Content Moderation. In *Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing (CSCW Companion '24)*. 9–12.
- [99] Omar Shaikh, Valentino Emil Chai, Michele Gelfand, Diyi Yang, and Michael S. Bernstein. 2024. Rehearsal: Simulating Conflict to Teach Conflict Resolution. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Article 920, 20 pages.
- [100] Amit Sheth, Valerie L. Shalin, and Ugur Kursuncu. 2022. Defining and Detecting Toxicity on Social Media: Context and Knowledge are Key. *Neurocomputing* 490 (2022), 312–318.
- [101] Yubo Shu, Tian Tian, Peng Zhang, Hansu Gu, Yaqiong Li, Yiyang Shao, Tun Lu, and Ning Gu. 2025. SoulSearch: Applying Heuristic Optimization to Enhance Text-to-image Generation with Personalized Human-LMM Collaboration. *Science China Information Sciences* 68, 10 (2025).

- [102] George Spivack and Myrna B Shure. 2014. Interpersonal Cognitive Problem Solving (ICPS): A competence-building primary prevention program. In *The National Mental Health Association*. Routledge, 151–178.
- [103] Tim Squirrel. 2019. Platform Dialectics: the Relationships between Volunteer Moderators and End Users on Reddit. *New Media & Society* 21, 9 (2019), 1910–1927.
- [104] Yla R. Tausczik and James W. Pennebaker. 2013. Improving Teamwork Using Real-time Language Feedback. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. 459–468.
- [105] Leanne Ten Brinke and Dacher Keltner. 2022. Theories of Power: Perceived Strategies for Gaining and Maintaining Power. *Journal of Personality and Social Psychology* 122, 1 (2022), 53.
- [106] Ryosuke Ueda, Koh Takeuchi, and Hisashi Kashima. 2022. Mitigating Observation Biases in Crowdsourced Label Aggregation. In *2022 26th International Conference on Pattern Recognition (ICPR)*. 1171–1177.
- [107] Cunxiang Wang, Xiaozhe Liu, Yuanhao Yue, Qipeng Guo, Xiangkun Hu, Xiangru Tang, Tianhang Zhang, Cheng Jiayang, Yunzhi Yao, Xuming Hu, Zehan Qi, Wenyang Gao, Yidong Wang, Linyi Yang, Jindong Wang, Xing Xie, Zheng Zhang, and Yue Zhang. 2025. Survey on Factuality in Large Language Models. *Comput. Surveys* 58, 1, Article 13 (2025), 37 pages.
- [108] Panpan Wang, Qian Huang, and Robert M. Davison. 2020. How Do Digital Influencers Affect Social Commerce Intention? The Roles of Social Power and Satisfaction. *Information Technology & People* 34, 3 (2020), 1065–1086.
- [109] Joel Wester, Tim Schriels, Henning Pohl, and Niels van Berkel. 2024. "As an AI language model, I cannot": Investigating LLM Denials of User Requests. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. Article 979, 14 pages.
- [110] David Wicks. 2017. The Coding Manual for Qualitative Researchers. *Qualitative Research in Organizations and Management: An International Journal* 12, 2 (2017), 169–170.
- [111] Donghee Yvette Wohn. 2019. Volunteer Moderators in Twitch Micro Communities: How They Get Involved, the Roles They Play, and the Emotional Labor They Experience. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. 1–13.
- [112] Qing Xiao, Yuhang Zheng, Xianzhe Fan, Bingbing Zhang, and Zhicong Lu. 2025. Let's Influence Algorithms Together: How Millions of Fans Build Collective Understanding of Algorithms and Organize Coordinated Algorithmic Actions. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Article 1115, 19 pages.
- [113] Mingzhe Yang, Hiromi Arai, Naomi Yamashita, and Yukino Baba. 2024. Fair Machine Guidance to Enhance Fair Decision Making in Biased People. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Article 285, 18 pages.
- [114] Jiayi Ye, Yanbo Wang, Yue Huang, Dongping Chen, Qihui Zhang, Nuno Moniz, Tian Gao, Werner Geyer, Chao Huang, Pin-Yu Chen, Nitesh Chawla, and Xiangliang Zhang. 2025. Justice or Prejudice? Quantifying Biases in LLM-as-a-Judge. In *International Conference on Representation Learning*, Vol. 2025. 102351–102390.
- [115] Xin Ye and Lionel P Robert. 2025. Security Robot Power and Acceptance: Exploring French and Raven's Five Forms of Power. In *2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 1731–1736.
- [116] Lu Yingchao, Chang Hailing, and Zhou Yadan. 2025. The Impact of Consumer Complaints on Corporate Digital Transformation. *International Review of Economics & Finance* 101 (2025), 104203.
- [117] Zehui Yu, Lukas Otto, Dennis Assenmacher, and Claudia Wagner. 2024. A Systematic Review of The Effects of AI-Assisted Moderation on Individuals and Groups. *Human-Machine Communication* 9 (2024), 167–188.
- [118] Xuechen Zhang, Changyang He, Peng Zhang, Hansu Gu, Ning Gu, Qi Shen, Zhan Hu, and Tun Lu. 2025. RemiHaven: Integrating "In-Town" and "Out-of-Town" Peers to Provide Personalized Reminiscence Support for Older Drifters. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Article 1034, 20 pages.
- [119] Andrew Zolides. 2021. Gender Moderation and Moderating Gender: Sexual Content Policies in Twitch's Community Guidelines. *New Media & Society* 23, 10 (2021), 2999–3015.

## A Appendix

**Table 4: Between-group paired  $t$ -test results for the human and human-AI moderation groups. \*\*\* indicates  $p < 0.001$ , \*\* indicates  $p < 0.01$ , and \* indicates  $p < 0.05$ . Cohen's  $d_z$  reports the effect size on the dependent variables.**

Power Manifestation	Human moderation group				Human-AI moderation group			
	sd	t	p	$d_z$	sd	t	p	$d_z$
Legitimate claim	1.181	2.284	<b>0.025(*)</b>	0.237	0.884	5.077	<b>0 (***)</b>	0.521
Authority citation	1.346	1.897	0.061	0.194	0.972	3.08	<b>0.003(**)</b>	0.309
Punishment threat	1.12	4.375	<b>0(***)</b>	0.461	1.058	2.015	<b>0.047(*)</b>	0.205
Compensation	1.303	3.759	<b>0(***)</b>	0.380	1.008	3.718	<b>0.001(**)</b>	0.358
Expert knowledge	1.094	2.357	<b>0.021(*)</b>	0.243	1.024	1.233	0.22	0.125
Group preference	1.261	1.893	0.061	0.192	0.974	1.032	0.305	0.104
Statement order	1.284	1.966	0.053	0.225	1.144	0.486	0.629	0.054
Expression tone	1.216	1.215	0.228	0.130	1.021	1.187	0.238	0.119
Choice trap	1.369	0.377	0.707	0.039	1.03	1.577	0.118	0.16
Length difference	1.289	2.601	<b>0.011(*)</b>	0.268	1.016	2.473	<b>0.012(*)</b>	0.248

**Table 5: Paired  $t$ -test results between the initial and perturbed versions across age groups (using the same representations as Table 4).**

Power Manifestation	Younger group				Older group			
	sd	t	p	$d_z$	sd	t	p	$d_z$
Legitimate claim	1.181	3.224	<b>0.002(**)</b>	0.423	1.284	1.382	0.176	0.234
Authority citation	1.443	0.809	0.422	0.111	1.366	2.088	<b>0.043(*)</b>	0.322
Punishment threat	1.153	4.045	<b>0(***)</b>	0.572	1.136	2.532	<b>0.016(*)</b>	0.416
Compensation	1.303	1.982	0.052	0.263	1.284	3.169	<b>0.003(**)</b>	0.478
Expert knowledge	1.421	2.226	<b>0.030 (*)</b>	0.297	1.225	1.847	0.072	0.288
Group preference	1.363	0.791	0.432	0.107	1.253	1.598	0.120	0.278
Statement order	1.242	1.675	0.102	0.268	1.306	1.756	0.096	0.403
Expression tone	1.421	1.027	0.310	0.150	1.093	1.391	0.172	0.223
Choice trap	1.38	1.558	0.126	0.216	1.428	0	1.000	0
Length difference	1.243	2.729	<b>0.008(**)</b>	0.349	1.238	2.640	<b>0.013(*)</b>	0.467

**Table 6: Paired  $t$ -test results between the initial and perturbed versions across gender groups (using the same representations as Table 4).**

Power Manifestation	Male group				Female group			
	sd	t	p	$d_z$	sd	t	p	$d_z$
Legitimate claim	1.27	1.147	0.256	0.152	1.212	3.236	<b>0.003(**)</b>	0.518
Authority citation	1.327	2.234	<b>0.030(*)</b>	0.301	1.533	0.866	0.392	0.135
Punishment threat	1.176	2.418	<b>0.020(*)</b>	0.353	1.209	3.125	<b>0.004(**)</b>	0.514
Compensation	1.406	3.289	<b>0.002(**)</b>	0.411	1.212	1.071	0.291	0.174
Expert knowledge	1.319	2.785	<b>0.007(**)</b>	0.372	0.901	1.387	0.173	0.217
Group preference	1.387	0.550	0.585	0.076	1.223	0.970	0.338	0.153
Statement order	1.414	1.414	0.166	0.236	1.265	1.318	0.201	0.275
Expression tone	1.136	1.257	0.215	0.180	1.372	0.928	0.362	0.182
Choice trap	1.59	0.424	0.673	0.057	1.271	0.905	0.371	0.149
Length difference	1.252	2.446	<b>0.018(*)</b>	0.357	1.265	0.577	0.567	0.094

**Table 7: Paired  $t$ -test results comparing the human and human-AI moderation groups (using the same representations as Table 4).**

Power Manifestation	sd	t	p	$d_z$
Initial version	0.76	-0.449	0.654	-0.045
Legitimate claim	0.815	-0.523	0.602	-0.056
Authority citation	0.979	0.106	0.916	0.011
Punishment threat	0.886	2.520	<b>0.014 (*)</b>	0.270
Compensation	0.932	1.626	0.107	0.164
Expert knowledge	0.839	0.663	0.509	0.069
Group preference	1.022	1.165	0.247	0.119
Statement order	1.096	1.026	0.309	0.128
Expression tone	1.006	0.998	0.321	0.108
Choice trap	1.093	-0.716	0.476	-0.075
Length difference	0.896	1.022	0.309	0.106