

# Work in Progress: A Glance at Social Media Self-Censorship in North America

Wei Hu  
University of Waterloo  
Waterloo, ON, Canada  
wei.hu1@uwaterloo.ca

Diogo Barradas  
University of Waterloo  
Waterloo, ON, Canada  
dbarrada@uwaterloo.ca

**Abstract**—There is a growing trend of social backlash and ostracism for thoughts and opinions shared online. Coupled with the rise of strict content moderation or digital nudges discouraging unpopular opinions, this trend calls into question whether users feel comfortable expressing their views freely on social media. Self-censorship can be defined as the “act of intentionally and voluntarily withholding information from others in the absence of formal obstacles” [4].

In this work, we sought to understand the self-censorship behavior of Canadian and United States social media users through an online survey. Our analysis suggests that users’ exhibit different degrees of concern when posting controversial content, and that these differences can be explained by demographic, psychometric and political orientation factors. Our results also suggest that there seems to be a consensus on the type of content that is more prone to be self-censored.

## 1. Introduction

Self-censorship is the act of limiting or controlling one’s own expression or behavior to avoid offending or upsetting others, or to conform to social or cultural norms [4]. As such, it can be perceived as a form of repression that imposes challenges on the proper functioning of a democratic society. Indeed, self-censorship undermines freedom of speech,<sup>1</sup> which is of paramount importance to the flow of information and fair democratic elections. For instance, Ong et al. [23] modelled how fear of state surveillance, harassment, and legal prosecution in Southeast Asian countries can reduce the expected utility of online expression, quieting dissenters and discouraging collective action. Likewise, following the 2016 Turkish coup attempt, Turkish citizens engaged in self-censorship, expressing less of their opinions on social media and removing old posts unfavourable to the government, due to fear of persecution [32]. In short, self-censorship is a major risk for authoritarianism and for autocratization.

Much of the existing self-censorship literature analyses this phenomenon within repressive countries with an history of conducting Internet censorship to quash dissent [23], [32], [5], [2], [10]. Less effort has been made to understand this behavior in democratic North American countries [24], [11]. Despite the perceivable absence of open political persecution in these countries, North American social media users are still affected by social norms

1. We also note that the distinction between self-censorship and self-moderation is nuanced, e.g., research on “chilling effects” [33] describes how experts and academics often self-moderate to avoid frictions.

that impose unique incentives for individuals to engage in self-censorship. For instance, social media platforms are increasingly taking a stance on what to censor (e.g., banning a sitting President’s accounts [3]), along with soft moderation (e.g., attaching warning labels to users’ posts that question elections’ integrity [38] or vaccines’ side-effects [28]). These mechanisms discourage free speech and call into question whether users feel comfortable expressing their views truthfully on social media [22]. Thus, it comes as no surprise that Reddit users have been resorting to throwaway accounts when discussing divisive political events in the United States [21].

In this work, we conducted a user survey and examined the types of controversial content that North American social media users, specifically those located within Canada and the United States (CAN-US), are most hesitant to share. We presented users with statements resembling controversial topics that led to cancellation in the past, and thus likely to be self-censored in the future. Following an approach based on Elo rating models, we asked users to rank these statements to establish a hierarchy of content prone to be self-censored.

From the insights produced by our study, we highlight that there is an apparent consensus on the types of content that are more prone to be self-censored, and that other factors like demographics and psychometrics can explain a variance in respondents’ concerns about posting online.

## 2. Methodology

We recruited 50 individuals to participate in an online survey. The survey involved gathering a number of psychometric and demographic indicators, examining participants’ political orientation, and asking participants about the kinds of content they are more willing to post on social media, so as to infer which opinions different clusters of individuals are more hesitant to share. We now introduce our research questions and our participants’ recruitment procedure. Then, we detail the design of our survey and discuss ethical considerations tied to its implementation.

*Research questions.* We aim to shed light on the prevalence of self-censoring behaviour on social media, exercised by Canadian and United States citizens. Toward this goal, we seek answers to three research questions:

*RQ1:* Is there a consensus in what sort of statements CAN-US social media users tend to self-censor?

*RQ2:* Are statements perceived as sympathetic to some controversial content equally self-censored as statements that are blatantly controversial?

*RQ3*: Does self-censorship vary meaningfully when grouped according to demographics and psychometrics?

*Recruitment of participants.* Our survey was deployed on Amazon Mechanical Turk (MTurk) and it was set up in such a way that users are directed to the survey's questionnaire through a link to a Google Form. MTurk was configured to filter respondents by age group and by location (so as to target participants based in CAN-US). We deployed the survey five times, specifying a different age group (18-25; 25-30; 30-35; 35-45; and 45-55 years old) each time. The questionnaire for each age group is identical. 10 participants were recruited from each age group, amounting to a total of 50 participants. The average time for participants to complete the survey was 20 minutes and 43 seconds (close to our estimate of 20 minutes). Each participant was paid \$2.50 (USD).

*Survey design.* The overarching goal of our survey is to study what different factors, including online presence, psychometrics, and demographics, lead CAN-US users to conduct self-censorship in social media platforms. The survey refrained from collecting any personally identifiable information about the respondents, and was organised in five major sections, which we describe below. Our survey's questionnaire can be found in Appendix E, and the collected data and analysis code is publicly available [12].

1) *Online presence.* This section is composed of 7 questions related to what degree the respondent uses social media and worries about potential consequences of her posts. In each question, respondents are asked to state whether they agree with a given statement using a 7-point Likert scale ranging from *strongly disagree* to *strongly agree*. An example of such a statement is "I am worried my social media posts can damage my friendships."

2) *Psychometrics.* This section is composed of 19 questions that aim to quantify each respondent's personality and IQ. Inspired by Tennant [34], we used a 10-item Big Five test [25] for quantifying respondents' personality across five dimensions: openness; conscientiousness, extraversion; agreeableness; and neuroticism (which we refer to as OCEAN). In each of the items of the Big Five test, respondents are asked to state whether they agree with a given statement using a 7-point Likert scale ranging from *strongly disagree* to *strongly agree*. An example statement is "I see myself as extroverted and enthusiastic".

To quantify respondents' IQ scores, we used a 12-item Raven's progressive matrices (RPM) test [26] which measures non-verbal fluid intelligence. We chose to administer only 9 out of 12 questions in the RPM booklet to reduce our survey's estimated completion time.

3) *Statement comparison.* This section is composed of 30 pairs of statements whose individual statements will be ranked according to three options – a) more willing to post, b) more hesitant to post, or c) equally willing to post either statement. This section aims at gathering data to enable us to compare participants' hesitance to post various categories of statements on social media. Each statement is representative of one of five considered categories, whose examples are given in Table 1.

Since different participants are expected to be more emphatic towards different statements, we ask respondents to imagine themselves as someone who believes both statements, and then to decide which of the two they

would be more willing to post on social media. Respondents are also given the option to select "I am equally willing to post either statement.", which effectively skips over the question. This section's statements were designed to capture a broad range of controversial content, and topics that have led to cancellation in the past.

4) *Political Orientation.* This section is composed of 5 questions asking respondents to use a 7-point Likert scale to state how much they agree with a political statement. These questions were chosen to be intentionally divisive and target core differences between conservatives and liberals [19]. An example of such a statement is "I should not contribute more than I expect to receive in return."

5) *Demographics.* The last section includes 3 questions to gather additional information from the survey's respondents, namely their profession, education level, and age.

*Ethical considerations.* Statements that individuals self-censor are often controversial and may be perceived as offensive. The study conducted in this paper has been reviewed by our University's Research Ethics Board (REB) and received an ethics clearance. During the ethics review, we worked with our REB to reduce the level of risk to respondents (e.g., by modifying some statements to be milder) while maintaining a valid methodology for our study. Indeed, if we simply omitted statements found to be potentially offensive, our survey would miss controversial statements that often overlap with these that social media users may refrain to post due to self-censorship (or that do post but are later ostracized for). Despite having implemented the aforementioned risk mitigation process, we are aware that some of the statements included in our survey may still be perceived as offensive. However, the REB has determined that the risk imposed by our study is justified by the tangible benefits our study would yield, i.e., to help gain a better understanding of the extent to which social media users within CAN-US feel comfortable in freely expressing their views in such platforms.

### 3. Analysis of Survey Data

Our analysis is geared at a) assessing whether there is consensus about self-censored statements using an Elo rating-based approach (*RQ1*); b) comparing the controversy level of different statement groups using statistical methods (*RQ2*), and; c) comparing the level of concern exhibited by different user groups about social media usage, when users are grouped together based on their psychometrics and demographics (*RQ3*).

*Data pre-processing.* We cleaned the data gathered by the survey and filtered out the likely low-quality, and thus uninformative, responses from 7 participants. Our criteria to accept responses as valid was based on the assumption that the respondents' success on the RPM test should be reasonably higher than what would be obtained through random guessing. We note there were no discrepancies between participants' self-reported age on the survey and the age ranges requested from Mechanical Turk.

We were also able to verify that the remaining respondents' responses to the statement comparison section were transitive, i.e., when a respondent chose statement *A* over statement *B* and statement *B* over statement *C*, then the respondent has also chosen statement *A* over statement

TABLE 1: Categories of statements (and examples) used in the Statement Comparison section of the survey. Participants are asked to fill in the prompts in square brackets in a manner that makes the statement most true for them.

Statement Category	Statement
Blatantly Political	[NAME] was the most influential and effective Presidents in modern history.
Sympathetic to political stance	Immigrants should be better welcomed.
Blatantly discriminatory	[RACIAL GROUP] people are insufferable.
Sympathetic to discrimination	Critical Race Theory does not have any academic rigor.
Attacks on closely-held beliefs	[COLLEGE DEGREE] has no academic rigor. There are better paths to unemployment.

C. The absence of transitivity discrepancies offers further evidence that these respondents did not answer randomly.

### 3.1. Consensus on Posts Self-Censorship

**3.1.1. Analysis Methodology.** To assess whether there is consensus in which posts respondents are most inclined to self-censor, we used the Elo rating system. Through this approach, we generated a ranking that shows which statements were more (and less) likely to be censored. Unlike qualitative approaches such as thematic analysis [6], [7], Elo models allows us to objectively measure whether there is a transitive hierarchy of censored content, thus avoiding other potential pitfalls tied to personal biases [14].

*Elo ratings.* Elo systems see many applications in skill-based games, such as two-player games where there is a transitive hierarchy of competence. These systems assign a score (i.e., an Elo rating) to each player, so as to measure their relative skill-level. Even in scenarios where skill fluctuates and there is a “luck” component, like Scrabble, certain physical sports or video games, Elo is nonetheless an effective predictor of game winners [13], [16].

When determining an Elo for a player, the outcome of the player’s games as well as the Elo of her opponents are taken into account as follows. Given two ratings,  $R_A$ ,  $R_B$  for players  $A, B$  respectively, the expected win-rate  $E_A$  for player  $A$  can be given by:

$$E_A = \frac{1}{1 + 10^{(R_B - R_A)/400}} \quad (1)$$

The equation’s constants are FIDE Chess’ default [31], and their purpose is to set an appropriate scale for the ratings [13]; a player with 400 more points than her opponent is 10 times more likely to win a match.

After playing a series of games, players’ Elos can be updated. Let  $S_A$  be the true win-rate attained by player  $A$  over a series of games (if only one game was played, then  $S_A = 1$  if  $A$  won, otherwise it is 0). The new rating  $R'_A$  is adjusted proportionally to the difference between the expected win-rate and the true win-rate.

$$R'_A = R_A + K \cdot (S_A - E_A) \quad (2)$$

Here,  $K$  is a hyper-parameter. The lower  $K$  is set, the more stable a player rating is. When building our model, we used an initial rating of 1500. Section 3.1.2 details how different values of  $K$  impact the fitting of our model.

*Elo-based consensus over self-censored posts.* We used the notion of Elo to examine whether there is a consensus between respondents on what content is most worth self-censoring. We begin with a prior that respondents are equally inclined to self-censor all statements, which is reflected in us assigning each statement an initial Elo rating of 1500. From there, suppose a respondent is shown two statements,  $A$  and  $B$ , and then selected  $A$  as the statement

TABLE 2: Elo scores for each statement. Statements classified as *sympathetic to discrimination* are shown in light-blue while *blatantly discriminatory* statements are shown in dark-blue.

Elo	Statement
1636	Think for yourself instead of constantly trusting in government to do what’s best for you.
1636	Immigrants should be better welcomed.
1631	I support [POLITICAL PARTY NAME].
1606	[NAME] was the most influential and effective Presidents in modern history.
1606	Hoping [POLITICAL CANDIDATE NAME] wins the next elections!
1597	When it comes to climate change, it is challenging to separate science from the politics.
1597	[SPORTS TEAM NAME] are an embarrassment to our nation!
1550	It is unclear to what degree vaccination protects against Omicron.
1547	[COLLEGE DEGREE] has no academic rigor. There are better paths to unemployment.
1520	“Until 1973 the American Psychiatric Association defined being gay as having a mental illness.”
1512	Critical Race Theory does not have any academic rigor.
1476	Colonialism has done more good than harm.
1450	The influx of foreigners has done irreparable damage to our community.
1435	People who buy [PRODUCT] name, shame on you!
1421	It is wrong that there exists a word that only members of one race can use.
1418	[GENDER] should not work as [JOB TITLE].
1397	Systemic [Racism/Sexism] is not sufficiently substantiated.
1388	Blackface should not be considered cultural appropriation.
1377	BlackLivesMatter was a movement that did more harm than good.
1371	[RACIAL GROUP] people are insufferable.
1328	If companies stopped interviewing [RACIAL GROUP] people, they could probably streamline their interview process.

they are more willing to post. We would then treat  $A$  as having won the game over  $B$ . This approach yields a list of 1290 games (43 respondents  $\times$  30 statements pairs). We iterate through this list in a random order, and update Elo scores at every step, according to Equation 2, ensuring the 30 statement pairs joined all the statements into one connected component. Since the order through which this iteration is performed makes a difference on the final Elos, we shuffled our pairings and ran the model 1000 times.

For each statement, we assign its final Elo as its mean rating across all runs. We then calculate the Mean Absolute Error (MAE) of the model by taking the difference between the true win-rate and  $E_A$  from Equation 1. A low MAE suggests there is consensus. Indeed, if respondents had mostly conflicting opinions, then the Elo model, which assumes transitivity, would not have been sufficiently flexible to represent the relationship between statements. In such a case, we would not expect to see a hierarchy emerge, but to observe very similar Elos instead.

*Consensus over self-censored posts categories.* Instead of treating each individual statement as a player and assign them a rating, we may also consider the five statement categories as players and assign those a rating instead. Under the aforementioned experimental setup, if the respondent chose statement  $A$  over  $B$ , then the group containing  $A$  would be seen as having beaten the group containing  $B$ . An Elo score would be assigned to each of the five groups.

**3.1.2. Results.** We start by exploring the results of our *Elo analysis between individual statements*. In Table 2, we present the hierarchy of statements according to their calculated Elos when the model’s  $K$  hyperparameter is set to 20. (This corresponds to the  $K$  value used by FIDE Chess for most players whose performance have stabilized [20].) To get some intuition for the scores, we can use Equation 1 and calculate that the most self-censored statement (with an Elo of 1328) is expected to be chosen over the least self-censored statement (with an Elo 1636) only  $\approx 14.5\%$  of the times. As we can see, the rating difference is somewhat significant. Indeed, if the Elo model has resulted in a good fit, these results suggest that the respondents have mostly agreed on one statement being more worth self-censoring than the other.

*Elo model validation.* To determine whether the model was a good fit, we compared the true win-rate to the win-rate predicted by the model and observed a Mean Absolute Error (MAE) of 4.5% across the pairings. This low error rate suggests that the model has successfully captured the respondents’ post withholding preferences. Along with the low MAE, the presence of a clear hierarchy and meaningful differences in Elo between statements provides further evidence that there is some degree of consensus between respondents about which statements are most worth self-censoring.

To assess whether our model’s hyper-parameter choice ( $K = 20$ ) was a suitable prior, we calculated the Mean Absolute Error (MAE) of the model for various values of  $K$ . Our results revealed that our choice was fairly close to the optimal  $K = 22$  (Appendix A). The results of further validation experiments also suggested that our Elo model was able to successfully avoid overfitting (Appendix B).

*Elo analysis between statement categories.* To understand whether there is consensus on which of the five groups of statements in Table 1 are perceived as most worthy to self-censor, we computed an Elo rating for each group by viewing the selection of a pair of statements to be a game between the two groups to which the statements belong. In this experiment, our fitted Elo model achieved a MAE of 16.5% across the five groups. The higher error at the statement group level is expected: there are statements within the same group that are  $>150$  Elo points apart implying a smaller consensus (e.g., the statements with an Elo of 1520 and 1377 in Table 2).

When calculating the win-rate between statement groups, we observed a meaningful difference between age groups. For example, some groups viewed statements that are “attacks on closely-held beliefs” to be more worth self-censoring than statements “sympathetic to a political stance” while others did not. The difference in win-rate between two groups that disagreed the most (25-30 year olds vs. 35-45 year olds) was 39.3%. The average difference between the two most disagreeing age-groups across all pairings was 25.6%. These observations appear to be consistent with the within-group consensus observed in the previous section.

As a result, we calculated the statement groups’ Elos separately by age group. Table 3 shows the resulting hierarchy: each statement group is assigned a value, ranging from 1 to 5, in ascending order of its Elo score within the age group, i.e., the statement group with score 1 is

TABLE 3: Statement categories ranked by age groups.

Statement Categories	Age Groups				
	18-25	25-30	30-35	35-45	45-55
Blatantly Discriminatory	3	1	1	1	1
Sympathetic to discrimination	1	2	2	3	2
Attacks on closely-held beliefs	2	5	3	2	3
Blatantly political	4	3	5	4	5
Sympathetic to political stance	5	4	4	5	4

most self-censored by that age group. The overall order of the statement groups in this table is sorted based on each group’s average rank across the 5 age groups in ascending order. In this setting, the MAE averaged 13.4%. The actual Elo scores used to rank the statement groups were omitted, because the Elo magnitudes are meaningless when compared across different groups.

Despite the above, we were able to observe some consensus about what statement groups are most worth self-censoring between specific age groups. Indeed, most groups agreed that statements which are *Blatantly discriminatory*, *Sympathetic to discrimination*, *Attacks on closely-held beliefs* are more worth self-censoring than political statements. However, the rankings of the 18-25 group had meaningful differences from the others. This difference may be explained due to added variance, since the 18-25 group turned out to consist of the smallest sample size after our initial pre-processing step to rule out participants who obtained very low RPM scores.

**Answer to RQ1.** There is an overall consensus on the content perceived as most worthy to self-censor. Using the Elo model, we presented a hierarchy of statements which reflects users’ willingness to post such statements. Our results suggest a larger consensus amongst individuals from the same age group.

## 3.2. Comparing Levels of Controversy

**3.2.1. Analysis Methodology.** We wish to understand whether there is a consensus in that “blatant” statements are more self-censored than “sympathetic to” statements. Considering the hierarchy of statements introduced previously in Table 2, we leverage the Mann-Whitney U (MWU) test [17] to compare the *ranks* of the statements included in the “blatant” category to those in the “sympathetic to” category. Here, the rank of a statement is simply the order it appears in Table 2. The test assigns a probability that the median of the ranks obtained by the statements belonging to each category is the same. The p-value of the test states how likely we are to see the observed difference between the groups or an even greater difference, if the groups indeed came from the same distribution. A smaller p-value means it is less likely the groups came from the same distribution.

If the MWU test shows that it is unlikely that the two groups have the same median rank, then we can conclude that respondents deem one group of statements to be more worth self-censoring than the other.

**3.2.2. Results.** Table 2 depicts statements that are *Blatantly discriminatory* in dark-blue, and those *Sympathetic to discrimination* in light-blue. The first group of statements had a mean rating of 1391.8 (with stdv 53.5), while

the second had a mean of 1429.8 (with stdv 56.5). This difference suggests blatantly discriminatory statements are more self-censored. However, we were not able to confirm a statistically significant difference between the two using the MWU test. It yielded a p-value = 0.23, which means there is a 23% chance of observing a greater or equivalent difference even if the two groups were sampled from the same distribution. The p-value is higher than the  $\alpha = 0.05$  often accepted for this test [36], [1], [30]. We note that our sample size is modest (only 10 statements). In practice, the MWU test does not have much power with small sample sizes, and would never return a p-value less than 0.05 whenever the sample size is  $\leq 7$ .

We similarly examined the self-censorship of statements that are blatantly political compared to those that seem to sympathise with a political stance. The mean Elo for the groups were 1614.3 and 1604.8 respectively, which does not suggest a meaningful difference. We observed a p-value = 1 in the MWU test, and conclude that respondents are equally likely to self-censor either group.

**Answer to RQ2.** Respondents seem equally likely to self-censor blatantly political statements and those that perceived as taking on a political position. Meanwhile, *Blatantly discriminatory* statements seem more censored than those *Sympathetic to discrimination*. However, we did not observe a statistically significant difference through the Mann-Whitney U test.

### 3.3. Influence of Demographics and Psychometrics on Self-Censorship Behavior

**3.3.1. Analysis Methodology.** To understand how self-censorship behavior varies between different groups of respondents, we grouped participants based on their *Demographics* and *Psychometrics* sections' answers, and compared them on the *Online Behaviour* section.

*Splitting participants into groups.* We considered two approaches for dividing participants into groups based on their demographics and psychometrics information. The first was to directly split the respondents according to a specific metric, such as their age group, their IQ measurements (based on the RPM section results), their score on a particular Big Five personality dimension, or the score given on the political orientation section. The second approach is to form groups by effectively clustering participants together based on the combination of a variety of these metrics.

To convert the respondents' responses into a comparable set of scores for each metric, we converted the qualitative appreciation scores used in the survey (which ranged from *strongly disagree* to *strongly agree*) to a 7-point numeric scale. Scores for questions measuring the same metric are combined and min-max normalized. Normalization is important before clustering, otherwise features with larger magnitudes would take on a disproportional weight. At this point, each respondent was given a score between 0 and 1 on each of the following metrics: extroversion, agreeableness, conscientiousness, neuroticism, openness, IQ, and political orientation.

*Clustering participants with k-means.* We used *k*-means [18] to cluster participants into  $k=4$  groups. Our

choice of *k* was informed by the elbow method, which refers to the highest *k* beyond which the observed reduction of the MAE slows down significantly.

*Comparing groups of participants.* Once the groups are formed, we aggregated the scores for each question in the *Online Behaviour* section for each group. We then looked for statistical differences between the groups' scores using the Kruskal-Wallis test, a generalized version of the MWU test (employed in Section 3.2) that accommodates more than 2 groups. The p-value of the Kruskal-Wallis test measures the probability of seeing a difference more extreme than (or equivalent to) what we observed if the groups had, in fact, all come from the same distribution.

**3.3.2. Results.** We start by exploring our results on the *comparison of groups across a single dimension*. We observed statistically significant differences in groups' answers to the *Online Behavior* portion of the survey. In Appendix C–Table 6, we show the set of statements and sources of concern for which we observed the largest differences between groups. As for demographics, we observed that respondents aged between 18-25 were more concerned about the consequences of their social media posts (Appendix D–Table 7). We did not find a similar difference across groups when comparing education level or employment status. Then, we grouped respondents based on their scores on the Big Five, IQ, and political orientation sections. For each dimension, we formed three groups: (1) respondents in the top 20-percentile, (2) respondents in the bottom 20-percentile, (3) the remaining respondents. Our results show that the group with high RPM scores was less concerned about consequences (Appendix D–Table 8), and that the groups which ranked high for conservatism and disagreeableness revealed more concerns (Appendix D–Tables 9 and 10, respectively).

*Comparing clusters formed across multiple dimensions.* An analysis of the groups formed through *k*-means revealed that one of the clusters deviated meaningfully from all others (Appendix D – Table 11). This deviant cluster had mostly neutral scores (around 4 on the Likert scale) on the metrics in the *Online Behaviour* section considered in our study, whereas the other three clusters had low scores indicating having little concern. Thus, the identified cluster consists of the respondents most concerned about the consequences of their social media use.

*Validating k-means results.* To verify that the deviating cluster does not simply consist of a set of people who did not bother to answer the survey properly, we further examined the centroid of that cluster. Indeed, it could be that there are survey respondents that chose “Neither Agree nor Disagree” for the entire survey, in which case they would all show up in the same cluster.

Table 4 shows the centroid of the four clusters, where cluster ID=1 (highlighted in blue) is the deviating cluster. We clear our suspicion by observing that it is not the case that cluster ID=1 has only neutral scores, indeed showing an IQ and agreeableness scores that somewhat stand out. This observation is consistent with Tables 9 and 10 in Appendix D. It seems plausible that participants who have a less agreeable nature may have a tendency to post more controversial content and thus end up being more concerned about potential consequences.

**Answer to RQ3.** We found participants from ages 18-25 to be more concerned about the effects of their social media opinions. We also found conservative and disagreeable respondents to be more concerned, and those who scored highly on the RPM test to be less concerned. Thus, we conclude demographics and psychometrics is predictive of self-censorship behavior.

#### 4. Limitations and Future Work

*Limited sample size.* Our modest sample of 43 participants (after filtering) and 21 statements made it difficult to derive more significant conclusions from our analysis. For instance, it is likely that a larger number of statement comparisons would allow us to see a larger spread of Elo ratings, thus making the statement hierarchy more evident. We will also aim for greater diversity in our sampling, by finding participants with varying degrees of digital literacy and age of first exposure to social media. Our future work will extend our study to accommodate a wider range of statements and a larger number of participants, further mitigating selection-bias. Currently, we ensure respondent age groups are balanced, and we control for psychometrics and political leanings. We also plan to extend our survey to other North American countries.

*Draws in the Elo system.* We ignored all pairings where respondents’ answer with “I am equally hesitant/willing to post either statement.” This lack of consideration for draws can contribute to an inaccurate assessment of the true differences in self-censorship between statements. In a future version of our study, we plan to apply Elo models which are specially geared at handling ties [31].

*Abstraction efforts.* The questionnaire requires participants to engage in hypotheticals that require multiple layers of abstraction. It could be difficult for a participant to first imagine herself as someone who agrees with both statements, and then to further imagine how she would feel posting each. In the future, we can control for this factor by asking participants to what degree they agree with the various statements shown.

#### 5. Related Work

*Causes of self-censorship on social media.* One important cause of self-censorship on social media is to avoid interpersonal conflicts. Powers et al. [24] examined American college students’ view of social media discourse and showed that students preferred to discuss their political views offline, mostly due to a rather politically homogeneous nature of social networks and the desire to avoid frictions. Gibson and Sutherland [11] further revealed that 40% of Americans engage in self-censorship behaviour because they worry that expressing unpopular views will alienate people from their close circles.

Another major cause of engaging in self-censorship in social media is that of avoiding professional repercussions. For instance, Aktas et al. [2] describe Turkish academics’ self-restraint in posting on social media. Larsen et al. [15] describe how journalists in Central America abstain from using social media to express their views due to job se-

TABLE 4: Centroid of each of the four clusters. Cluster 1 (shaded) deviated substantially from the other clusters.

Cluster ID	Big Five Personality						IQ	Political
	O	C	E	A	N			
1	0.69	0.67	0.54	0.28	0.56	0.45	0.59	
2	0.87	0.94	0.91	0.93	0.91	0.73	0.37	
3	0.79	0.91	0.23	0.81	0.82	0.71	0.53	
4	0.49	0.53	0.08	0.53	0.38	0.82	0.31	

curity concerns. Rudnik [27] shows that similar concerns lead Russian and Belarusian bloggers to self-censor.

Other concerns include safety and privacy. In effect, there is a vast body of literature on the analysis of self-censorship in countries ruled by repressive regimes [15], [23], [8], [5], [10] where, for safety reasons, journalists avoid publishing or exchanging information about certain topics. In addition, Warner and Wang [35] revealed that the self-censoring behavior of individuals living in the United Kingdom has increased as new online surveillance methods were introduced by intelligence agencies. We expand on this conversation by linking one’s willingness to self-censor to psychometrics and demographic factors.

*Measuring self-censorship on social media.* It is challenging to keep track of self-censorship events on social media as such an effort involves analysing social media posts that never actually materialized. Towards this goal, Das and Kramer [9] used Facebook internal data to capture content that users started writing but ultimately refrained from posting. Yet, it is possible that many self-censored statements do not make it that far, as users might immediately dismiss the idea of writing a social media post due to fearing social repercussions. Instead, Sleeper et al. [29] proposed that users keep a log of statements they wanted to post but ultimately did not post. However, the inconvenience of keeping such a log may have prevented users from sharing self-censored posts with researchers.

A different approach to study self-censorship could be to examine posts that people wrote, regretted sharing, and then deleted. For example, Xia et al. [37] used `polititweet.org` – a service that tracks messages that were posted on Twitter, but later deleted – to understand how deleted tweets helped spread disinformation.

In our study, we applied the Elo model to measure the degree of censorship across different topics. We had participants imagine scenarios where self-censorship might occur, instead of waiting for actual retractions, thus learning how the same participant perceives different topics.

#### 6. Conclusion

In this study, we sought to understand the phenomenon of social media self-censorship behaviour exercised by users within Canada and United States, both qualitatively and quantitatively. We designed a survey which enabled us to compare the likelihood of participants to discuss different topics with the potential to generate some sort of controversy. The results of our analysis over participants’ responses suggest that there is a consensus about what kinds of content are more prone to be self-censored by CAN-US social media users. We also identified a relationship between users’ demographics, psychometrics, and political orientation and their concerns about discussing certain contents online.

## Data Availability

The raw data provided by the respondents of the survey as well as the code used to process said data has been made publicly available [12]. The survey participants' personally identifiable information has been anonymized.

## Acknowledgements

We thank the anonymous reviewers for their insightful feedback and suggestions that helped improve our paper. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

## References

- [1] William Aiken, Hyoungshick Kim, Jungwoo Ryoo, and Mary Beth Rosson. An Implementation and Evaluation of Progressive Authentication Using Multiple Level Pattern Locks. In *Proceedings of the 2018 16th Annual Conference on Privacy, Security and Trust (PST)*, pages 1–6, 2018.
- [2] Vezir Aktas, Marco Nilsson, and Klas Borell. Social scientists under threat: Resistance and self-censorship in Turkish academia. *British Journal of Educational Studies*, 67(2):169–186, 2019.
- [3] Meysam Alizadeh, Fabrizio Gilardi, Emma Hoes, K Jonathan Klüser, Mael Kubli, and Nahema Marchal. Content Moderation As a Political Issue: The Twitter Discourse Around Trump's Ban. *Journal of Quantitative Description: Digital Media*, 2, 2022.
- [4] Daniel Bar-Tal. Self-censorship as a socio-political-psychological phenomenon: Conception and research. *Political Psychology*, 38:37–65, 2017.
- [5] Svetlana S Bodrunova, Anna Litvinenko, and Kamilla Nigmatullina. Who is the censor? Self-censorship of Russian journalists in professional routines and social networking. *Journalism*, 22(12):2919–2937, 2021.
- [6] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101, 2006.
- [7] Virginia Braun and Victoria Clarke. *Thematic analysis*. American Psychological Association, 2012.
- [8] Anne Shann Yue Cheung. *Self-censorship and the struggle for press freedom in Hong Kong*. Brill Nijhoff, 2021.
- [9] Sauvik Das and Adam Kramer. Self-censorship on Facebook. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 7(1), pages 120–127, 2013.
- [10] Abdelmalek El Kadoussi. The perception of self-censorship among Moroccan journalists. *The Journal of North African Studies*, 27(2):231–263, 2022.
- [11] James L Gibson and Joseph L Sutherland. Keeping Your Mouth Shut: Spiraling Self-Censorship in the United States. Available at SSRN 3647099, 2020.
- [12] Wei Hu and Diogo Barradas. Our study's code repository. <https://github.com/TheWeiHu/self-censorship-user-study>, 2022.
- [13] Lars Magnus Hvattum and Halvard Arntzen. Using ELO ratings for match result prediction in association football. *International Journal of forecasting*, 26(3):460–470, 2010.
- [14] Mostafa Javadi, Koroush Zarea, et al. Understanding thematic analysis and its pitfall. *Journal of client care*, 1(1):33–39, 2016.
- [15] Anna Grøndahl Larsen, Ingrid Fadnes, and Roy Krøvel. *Journalist Safety and Self-Censorship*. Routledge, 2021.
- [16] Robert Lehmann and Klaus Wohlrahe. Who is the 'Journal Grand Master'? A new ranking based on the Elo rating system. *Journal of Informetrics*, 11(3):800–809, 2017.
- [17] Thomas W MacFarland and Jan M Yates. Mann-Whitney U Test. In *Introduction to Nonparametric Statistics for the Biological Sciences Using R*, pages 103–132. Springer, 2016.
- [18] James MacQueen et al. Some methods for classification and analysis of multivariate observations. In *Proceedings of the 5th Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 281–297, 1967.
- [19] Dan Meegan. *America the Fair: Using Brain Science to Create a More Just Nation*. Cornell University Press, 2019.
- [20] Gabriel Alves Vasiljevic Mendes, Leonardo Cunha de Miranda, and Erica Esteves Cunha de Miranda. MasterMind Chess: Design and Implementation of Classic, Capablanca and Fischer Modes with Real Time Match Observation. In *Proceedings of the 2014 Brazilian Symposium on Computer Games and Digital Entertainment (SBGAMES)*, pages 220–228, 2014.
- [21] Rishab Nithyanand, Brian Schaffner, and Phillipa Gill. Measuring Offensive Speech in Online Political Discourse. In *Proceedings of the 7th USENIX Workshop on Free and Open Communications on the Internet (FOCI 17)*, 2017.
- [22] Pippa Norris. Closed Minds? Is a 'Cancel Culture' Stifling Academic Freedom and Intellectual Debate in Political Science? Technical report, HKS Working Paper No. RWP20-025, Available at SSRN: <https://ssrn.com/abstract=3671026>, 2020.
- [23] Elvin Ong. Online repression and self-censorship: Evidence from southeast Asia. *Government and Opposition*, 56(1):141–162, 2021.
- [24] Elia Powers, Michael Koliska, and Pallavi Guha. "Shouting Matches and Echo Chambers": Perceived Identity Threats and Political Self-Censorship on Social Media. *International Journal of Communication*, 13:20, 2019.
- [25] Beatrice Rammstedt and Oliver P John. Measuring personality in one minute or less: A 10-item short version of the Big Five Inventory in English and German. *Journal of Research in Personality*, 41(1):203–212, 2007.
- [26] John C Raven and JH Court. *Raven's progressive matrices*. Western Psychological Services Los Angeles, CA, 1938.
- [27] Alesia Rudnik. Why do bloggers keep silent? Self-censorship in social media: cases of Belarus and Russia. Master's thesis, Södertörn University (Sweden), 2020.
- [28] Filipo Sharevski, Alice Huff, Peter Jachim, and Emma Pieroni. (Mis) perceptions and engagement on Twitter: COVID-19 vaccine rumors on efficacy and mass immunization effort. *International Journal of Information Management Data Insights*, 2(1):100059, 2022.
- [29] Manya Sleeper, Rebecca Balebako, Sauvik Das, Amber Lynn McConahy, Jason Wiese, and Lorrie Faith Cranor. The Post that Wasn't: Exploring Self-Censorship on Facebook. In *Proceedings of the 2013 ACM Conference on Computer Supported Cooperative Work*, pages 793–802, 2013.
- [30] Kresimir Solic, Tena Velki, and Tomislav Galba. Empirical study on ICT system's users' risky behavior and security awareness. In *Proceedings of the 2015 38th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*, pages 1356–1359, 2015.
- [31] Leszek Szczecinski and Aymen Djebbi. Understanding draws in Elo rating algorithm. *Journal of Quantitative Analysis in Sports*, 16(3):211–220, 2020.
- [32] Rima Tanash, Zhouhan Chen, Dan Wallach, and Melissa Marschall. The Decline of Social Media Censorship and the Rise of Self-Censorship after the 2016 Failed Turkish Coup. In *Proceedings of the 7th USENIX Workshop on Free and Open Communications on the Internet*, 2017.
- [33] Jaime A. Teixeira da Silva. How to shape academic freedom in the digital age? Are the retractions of opinionated papers a prelude to "cancel culture" in academia? *Current Research in Behavioral Sciences*, 2, 2021.
- [34] John Matthew Tennant. It's Not You, It's Me: Sex, Context, and Individual Differences in Experimental Partner Choice. Master's thesis, University of Toronto (Canada), 2017. pp. 37.
- [35] Mark Warner and Victoria Wang. Self-censorship in social networking sites (SNSs)—privacy concerns, privacy awareness, perceived vulnerability and information management. *Journal of Information, Communication and Ethics in Society*, 2019.

- [36] Dominik Wermke, Nicolas Huaman, Christian Stransky, Niklas Busch, Yasemin Acar, and Sascha Fahl. Cloudy with a chance of misconceptions: exploring users’ perceptions and expectations of security and privacy in cloud office suites. In *Proceedings of the 16th Symposium on Usable Privacy and Security (SOUPS 2020)*, pages 359–377, 2020.
- [37] Yiping Xia, Josephine Lukito, Yini Zhang, Chris Wells, Sang Jung Kim, and Chau Tong. Disinformation, performed: Self-presentation of a Russian IRA account on Twitter. *Information, Communication & Society*, 22(11):1646–1664, 2019.
- [38] Savvas Zannettou. ”I Won the Election!”: An Empirical Analysis of Soft Moderation Interventions on Twitter. In *Proceedings of the 15th International AAAI Conference on Web and Social Media (ICWSM 2021)*, pages 865–876, 2021.

## A. Elo Model’s Hyper-parameters

Figure 1 depicts the variation of the mean absolute error (MAE) obtained by the Elo model while increasing the value of  $K$ . We can observe from the figure that the default value of  $K$  we chose sits close to the optimal value of  $K=22$ .




Figure 1: MAE’s variation according to the increase of  $K$ . A higher  $K$  translates into a stabler player’s rating.

## B. Verifying Overfitting on the Elo Model

Due to the flexibility to learn a rating for each statement, there was the possibility that the Elo model was overfitting. To verify that this was not the case, we had the model predict on out-of-sample data, using leave-one-out cross-validation (LOOCV). This process involved calculating the statement Elos while leaving one pairing of statements for which we had the true win-rate. Then, using the predicted Elos, we calculated the expected win-rate and compared it to the true win-rate. After applying LOOCV, we observed an increase in MAE to  $\approx 11\%$ , as shown at the bottom of Table 5. We see that the Elo model performed better than a simple baseline predictor that always predicts 0.5. Indeed, if a) users had selected answers randomly or b) there is no transitive hierarchy and the model can not learn from the other pairings, then the best a predictor could do is to predict 0.5. Because the model was able to perform better than just predicting 0.5 across the statements, it seems likely that neither a) nor b) is true.

Furthermore, to examine whether there is higher consensus within the same age group, we perform the analysis separately for each age group. When doing so, we notice a further improvement in performance against the baseline, which suggests a stronger sense of consensus within the same group.

Overall, the low error-rate predicting on out-of-sample pairings and the better performance against the baseline is suggestive that our Elo model is not overfitting.

TABLE 5: Comparing the performance of the LOOCV Elo model against the baseline model.

Age Groups	Elo MAE - Mean	Elo MAE - Stdv	Baseline MAE
18-25	10.0%	0.2%	12.1%
25-30	17.1%	0.2%	20.5%
30-35	14.6%	0.3%	17.6%
35-45	15.1%	0.2%	19.3%
45-55	15.2%	0.2%	19.2%
<b>Combined</b>	11.3%	0.6%	14.1%

## C. Statements

Table 6 depicts the statements and sources of concern for which we observed the largest differences between groups’s answers to the *Online Behavior* surveys’ portion.

## D. Group Comparison Tables

Some demographic and psychometric groups’ degree of concern about the consequences of posting on social media stood out. Tables 7 to 11 compare the differing groups to the average across respondents. The *Avg. Score* column is the average across participants.

## E. Survey Questions

This section of the appendix contains a summary of the questionnaire that was filled by respondents.

### E.1. Survey Introduction

The purpose of this academic survey is to get a better understanding of self-censorship in North America. The results from the study will help researchers better understand whether users still feel comfortable expressing their views freely on social media. In this survey, we will collect some demographics data, administer a few short tests about how you think and some information regarding political orientation. From there, participants will be given pairs of intentionally controversial statements and asked to indicate which they are more hesitant to post on social media.

### E.2. Information Consent

Please read the information and consent form and confirm:

*I’ve read the information consent form and give my consent to take part in the survey.*

### E.3. Online Behavior

Let’s start with some basic questions about your online social media behavior!

- I consider myself a frequent social media user.
  - Yes
  - No
- I worry about the employment-related repercussions that can come from my social media posts.
 

7-point Likert scale ranging from *strongly disagree* to *strongly agree*.
- I am concerned my social media posts can damage my friendships.



TABLE 6: Subset of statements exemplifying different categories of users’ concerns about posting on social media.

Statement	Source of concern
I’ve faced consequences / been cancelled for content I’ve posted on social media.	Previous consequences and cancellations
I fear legal consequences for my social media posts.	Legal consequences
I worry that my post on social media can jeopardize my romantic relationships.	Personal relationship consequences
I worry about being banned or cancelled if I don’t restrain myself in terms of what I post.	Fear of being deplatformed

TABLE 7: Comparing 18-25 year old age group to the average.

Source of Concern	Score	Avg. Score	P-Value
Previous Consequences and Cancellations	4	2.3	0.097
Legal Consequences	5	3.1	0.050
Romantic Consequences	3	2.7	0.30
Fear of Being Deplatformed	3.2	3.2	0.49

TABLE 8: Comparing high scorers on the RPM to the average.

Source of Concern	Score	Avg. Score	P-Value
Previous Consequences and Cancellations	1.5	2.3	0.0030
Legal Consequences	1.9	3.1	0.0088
Romantic Consequences	1.7	2.7	0.031
Fear of Being Deplatformed	3.3	3.2	0.21

TABLE 9: Comparing disagreeable respondents to the average.

Source of Concern	Score	Avg. Score	P-Value
Previous Consequences and Cancellations	4.3	2.3	9.54E-05
Legal Consequences	4.8	3.1	0.0049
Romantic Consequences	3.7	2.7	0.11
Fear of Being Deplatformed	4.9	3.2	0.0049

TABLE 10: Comparing conservative respondents to the average.

Source of Concern	Score	Avg. Score	P-Value
Previous Consequences and Cancellations	3.5	2.3	0.041
Legal Consequences	4.3	3.1	0.047
Romantic Consequences	3.4	2.7	0.40
Fear of Being Deplatformed	4.3	3.2	0.032

TABLE 11: Deviating  $k$ -means cluster compared to the average.

Source of Concern	Score	Avg. Score	P-Value
Previous Consequences and Cancellations	4.2	2.3	7.70E-05
Legal Consequences	4.6	3.1	0.005
Romantic Consequences	3.9	2.7	0.016
Fear of Being Deplatformed	4.3	3.2	0.058

7-point Likert scale ranging from *strongly disagree* to *strongly agree*.


- 4) I fear legal consequences for my social media posts.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.
- 5) I’ve faced consequences / been “cancelled” for content I’ve posted on social media.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.
- 6) I worry that my post on social media can jeopardize my romantic relationships.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.

- 7) I worry about being banned or cancelled if I don’t restrain myself in terms of what I post.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.

#### E.4. Pattern Matching

Here are 9 short puzzles for a quick test of your pattern matching skills!

- 1) Select the option the best completes the image:



[Authors’ note: There are an additional 8 similar questions taken from [26] in the questionnaire that was distributed.]

#### E.5. Personality Test

Let’s get a sense of your personality! Please indicate to what degree you agree with each of the following statements.

- 1) I see myself as extraverted, enthusiastic.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.
- 2) I see myself as open to new experiences, complex.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.
- 3) I see myself as critical, quarrelsome.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.
- 4) I see myself as dependable, self-disciplined.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.
- 5) I see myself as anxious, easily upset.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.
- 6) I see myself as reserved, quiet.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.
- 7) I see myself as sympathetic, warm.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.
- 8) I see myself as disorganized, careless.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.

- 9) I see myself as calm, emotionally stable.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.
- 10) I see myself as conventional, uncreative.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.

### E.6. Comparing Willingness to Post

The following statements are intentionally controversial and are not necessarily the views of the researchers but are being used to answer the research question. Please indicate which of the two statements provided you'd be more willing to post.

Try to fill in the [PROMPT]s in a manner that makes the statements most true for you. N.B.: We are not looking to measure whether you believe any of the statements to be true. Instead, we hope when you are filling the questionnaire that you would emphasize and imagine yourself as someone who believes in each of the statements: if you believed both statements, which would you be more willing to post?

- 1) I would be more willing to post:
  - A. Immigrants should be better welcomed.
  - B. Critical Race Theory does not have any academic rigor.
  - C. I am equally willing to post either statement.
- 2) I would be more willing to post:
  - A. Systemic [Racism/Sexism] is not sufficiently substantiated.
  - B. Hoping [POLITICAL CANDIDATE NAME] wins the next elections!
  - C. I am equally willing to post either statement.
- 3) I would be more willing to post:
  - A. It is wrong that there exists a word that only members of one race can use.
  - B. Blackface should not be considered cultural appropriation.
  - C. I am equally willing to post either statement.

[Authors' note: There are 12 additional comparisons in this section asking respondents which statement they are more willing to post.]

### E.7. Comparing Hesitance to Post

The following statements are intentionally controversial and are not necessarily the views of the researchers but are being used to answer the research question. Please indicate which of the two statements provided you would be more hesitant to post.

Try to fill in the [PROMPT]s in a manner that makes the statements most true for you. N.B.: We are not looking to measure whether you believe any of the statements to be true. Instead, we hope when you are filling the questionnaire that you would emphasize and imagine yourself as someone who believes in each of the statements: if you believed both statements, which would you be more hesitant to post?

- 1) I am more hesitant to post:
  - A. Hoping [POLITICAL CANDIDATE NAME] wins the next elections!

- B. People who buy [PRODUCT] name, shame on you!
  - C. I am equally hesitant to post either statement.
- 2) I am more hesitant to post:
    - A. Think for yourself instead of constantly trusting in government to do what's best for you.
    - B. Hoping [POLITICAL CANDIDATE NAME] wins the next elections!
    - C. I am equally hesitant to post either statement.
  - 3) I am more hesitant to post:
    - A. BlackLivesMatter was a movement that did more harm than good.
    - B. Systemic [Racism/Sexism] is not sufficiently substantiated.
    - C. I am equally hesitant to post either statement.

[Authors' note: There are 12 more comparisons in this section asking respondents which statement they are more hesitant to post. Flipping the question (i.e., asking about hesitance vs. willingness) had an effect on how respondents treated the prompt when we initially tested out the survey pre-deployment.]

### E.8. Political Orientation

We're hoping to get some insights on your thoughts regarding some issues.

- 1) It feels wrong when an employee who needs their job, is fired.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.
- 2) It's desirable when employees who contribute more to the success of the company receive a larger share.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.
- 3) How would your friends describe your political position?  
7-point Likert scale ranging from *very left* to *very right*.
- 4) I often have tender, concerned feelings for people less fortunate than me.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.
- 5) I value social status and prestige, control or dominance over people and resources.  
7-point Likert scale ranging from *strongly disagree* to *strongly agree*.

### E.9. Demographics Information

Please answer a few final questions for statistical analysis purposes only. You will remain entirely anonymous.

- 1) Select your highest degree of completed education:
  - A. Did not graduate high school
  - B. High school or equivalent
  - C. Associate / Bachelor's Degree
  - D. Graduate Degree
- 2) Select your employment status:
  - A. Salaried Employee
  - B. Part-Time Employee
  - C. Student
  - D. Unemployed
  - E. Self-Employed
- 3) Select your age group:
 

A. 18-25	B. 25-30
C. 30-35	D. 35-45
E. 45-55	F. 55+