

A Peek into the Political Biases in Email Spam Filtering Algorithms During US Election 2020

Hassan Iqbal, Usman Mahmood Khan, Hassan Ali Khan, Muhammad Shahzad

Department of Computer Science, North Carolina State University

Raleigh, North Carolina, USA

{hiqbal,ukhan3,hakhan,mshahza}@ncsu.edu

ABSTRACT

Email services use spam filtering algorithms (SFAs) to filter emails that are unwanted by the user. However, at times, the emails perceived by an SFA as unwanted may be important to the user. Such incorrect decisions can have significant implications if SFAs treat emails of user interest as spam on a large scale. This is particularly important during national elections. To study whether the SFAs of popular email services have any biases in treating the campaign emails, we conducted a large-scale study of the campaign emails of the US elections 2020 by subscribing to a large number of Presidential, Senate, and House candidates using over a hundred email accounts on Gmail, Outlook, and Yahoo. We analyzed the biases in the SFAs towards the left and the right candidates and further studied the impact of the interactions (such as reading or marking emails as spam) of email recipients on these biases. We observed that the SFAs of different email services indeed exhibit biases towards different political affiliations. We present this and several other important observations in this paper.

CCS CONCEPTS

• Information systems → Email; • Networks → Network measurement.

KEYWORDS

US Elections, Emails, Spam, Bias, Political Bias, Algorithm Bias

1 INTRODUCTION

The spam filtering algorithms (SFAs) in the widely-used email services of today such as Gmail, Outlook, and Yahoo do not provide any transparency on their internal workings. Given the lack of this transparency, an important question to study is whether these SFAs hold any biases towards certain political affiliations. This question is motivated by the growing body of evidence suggesting that the biases in online algorithms can influence undecided voters. For example, Epstein *et al.* showed that the bias in search rankings can shift the voting preferences of the undecided voters by as much as 20% without those voters being aware of the manipulation [23]. Furthermore, several US political candidates in the 2020 US election raised concerns that the email clients were filtering out the campaign emails they were sending to their constituents [1].

Research Questions: In this paper, we attempt to assess the fairness of the SFAs of three dominant email services, Gmail, Outlook, and Yahoo, in the context of the 2020 US election. Specifically, we study the following four research questions:

- **Q1:** Do SFAs of email services exhibit aggregate political biases? How do these biases compare across email services?
- **Q2:** Do SFAs treat similar emails from senders with different political affiliations in the same way?
- **Q3:** Do the interactions of the users with their email accounts, such as reading emails, impact the political biases of SFAs?
- **Q4:** Do SFAs exhibit different political biases for recipients belonging to different demographic?

To the best of our knowledge, there is no prior published work that has examined biases in SFAs towards political campaigns.

Proposed Methodology: To answer these questions, we conducted an extensive study during the 2020 US election over a period of 5 months from July 1, 2020 to November 30, 2020 on Gmail, Outlook, and Yahoo. We created 102 email accounts and subscribed to 2 Presidential, 78 Senate, and 156 House candidates. To accurately estimate the political biases and mitigate any potential effects of demographics (ethnicity, age, and gender), we created multiple email accounts with different combinations of demographic factors and designed two experiments. The first experiment studies the general trends of biases in SFAs across the email services for the Presidential, Senate and House candidates. The second experiment studies the impact of different email interactions such as reading the emails, marking them as spam, or vice versa on the biases in SFAs. We designed an automated process to perform all the subscriptions, and took periodic backups to keep all the email accounts active as well as to keep track of the correct number of spam emails received over the course of data collection for each of the three services.

We made several important observations in our study. For example, as an *aggregate trend*, Gmail leaned towards the left while Outlook and Yahoo leaned towards the right. Yahoo retained about half of all the political emails in inbox (up to 55.2% marked as spam) while outlook filtered out the vast majority of emails (over 71.8%) from all political candidates and marked them as spam. Gmail, however, retained the majority of left-wing candidate emails in inbox (< 10.12% marked as spam) while sent the majority of right-wing candidate emails to the spam folder (up to 77.2% marked as spam). We further observed that the percentage of emails marked by Gmail as spam from the right-wing candidates grew steadily as the election date approached while the percentage of emails marked as spam from the left-wing candidates remained about the same. We present these and several other important observations in this paper.

Key Contributions:

This is a pre-print. The final version is published in The ACM Web Conference 2022 (WWW '22). For the published version and citation, please refer to DOI: <https://doi.org/10.1145/3485447.3512121>

- To the best of our knowledge, this paper presents the first study that extensively explores the political biases in SFAs.
- We used the propensity score matching approach [32] to determine whether the SFA of any given email service provided same treatment to similar emails from candidates of different political affiliations.
- We have aggregated and analyzed a large data set of over 318K political emails across the three email services. This data set is available at [11].

Paper Organization: Next, we discuss the related work in § 2 and describe our methodology in § 3. We then present the extensive analysis of our data set in § 4 to study our four research questions. Finally, in § 5, we summarize our observations, discuss their implications and provide some suggestions, and conclude the paper.

2 RELATED WORK

Researchers have proposed several different definitions of spam emails. Butler defines spam as any unsolicited email that comes from an entity that the recipient is not already aware of or has no interest in knowing about [18]. Cormack *et al* similarly defines spam as any unsolicited or unwanted email that is sent indiscriminately, and has no current relation to the recipient [20]. Similar definitions have been proposed in other related articles [14] [34]. In contrast, Google defines spam as any content that is unwanted by the user [35]. This is significantly different from the criteria proposed by the previous research in that the spam email does not have to meet any of the explicitly defined conditions so long as there is a reason to believe that the email may be unwanted by the recipient [35]. Other email clients such as Outlook and Yahoo have not made their definitions of spam public. In this study, our objective is to study how these email services treat emails from the political websites that the recipient has subscribed to, and if that treatment has any biases, irrespective of how these services define spam.

One prior work examines the manipulative tactics in political emails [2]. The focus of this work is on finding different manipulative tactics that the campaigns use to encourage readers to open emails and donate. It categorizes manipulative tactics as click baits (forward referencing, sensationalism, urgency), and user interface manipulation (obscured names, ongoing thread, subject manipulation). To the best of our knowledge, there is no past published work on examining the political biases in the SFAs of different email services. However, prior work has examined the biases of several other web-based algorithms. For example, the past studies have demonstrated that the online advertising results may be personalized [16] [22]. Hannak *et al* found significant personalization in Google web search based on the account login status and the geolocation of the IP address [25]. Another study reported that 15.8% of the Bing search results were personalized for different users [26]. Puschmann analyzed Google search results for German political parties and candidates and found that the results were personalized based on several factors including location, language, and time [31]. Huyen *et al* showed that these results are further personalized based on the user’s browsing history [29]. However, to the best of our knowledge, no such study exists for email SFAs in the context of political campaign emails.

3 METHODOLOGY

3.1 Emails and Demographics

We used three email services, Gmail, Outlook, and Yahoo, and created 102 email accounts, 34 on each of the three services. To accurately estimate the political biases and mitigate the potential effects of demographic factors such as ethnicity and age, we created our email accounts with different combinations of these factors. As email services do not explicitly collect ethnicity information, we assigned a different name to each email account that we randomly picked from a database of common names associated with White, African-American, Hispanic, Asian, and South Asian ethnicities. For age, we assigned each email account to one of the three age groups of 18-40, 40-65, and 65+. Finally, we randomly assigned male and female genders to the email accounts. To conform to the ethical standards, none of the email accounts that we created belonged to any real users and all the accounts are new with no prior history. We manually created all the accounts by following the account creation procedure of the three email services.

3.2 Subscribed Candidates

We subscribed our email accounts to presidential, Senate, and House candidates, described next.

3.2.1 Presidential Candidates: This category includes the two Presidential candidates, one from the left, *i.e.*, Joe Biden (Democrat), and one from the right, *i.e.*, Donald Trump (Republican).

3.2.2 Senate and House Candidates: This category combines candidates from both the United States Senate and the House of Representatives, as shown in Fig. 1. The blue and red circles represent the left and right candidates, respectively, in the US House of Representatives that we picked for this study. Similarly, the blue and red squares represent the left and right Senate candidates. In some states, we subscribed to a different number of left and right candidates due to four reasons. First, different states have different number of seats in the House depending on factors like the population of the state. Second, each Senate and House election is contested by different number of candidates from both Democratic (left) and Republican (right) parties. Third, some candidates only had government affiliated websites (with .gov domain names), which are prohibited to send out campaign emails [6, 12]. Fourth, due to our automated subscription methodology (will be described shortly),




Figure 1: Distribution of Senate & House Candidates in our subscriptions.

we could not subscribe to the campaign websites that required filling out a CAPTCHA form. We initially chose the campaign websites of all the left and the right Senate and House candidates in the 50 states. After filtering those websites based on the third and the fourth reasons stated above, we ended up with an unequal number of campaign websites of the left and the right candidates, in various states, for both the Senate and the House.

To reduce the gap between the number of left and right candidates, our subscription methodology was as follows. If any state had more than one but unequal number of left and right candidates, we subscribed to the maximum number of candidates such that the counts of the left and right candidates were the same. To keep as many states in our analysis as possible, we did not use this approach in states like Alaska (with only 1 Republican senate candidate), where we found candidates with non-government affiliated websites from only one of the parties. There were 11 such states. As a result, there is a small difference between the counts of the left and the right House and Senate candidates in our subscriptions. In total, we subscribed to Senate candidates from 36 states with 78 subscriptions (44 L and 34 R), and House candidates from 42 states with 156 subscriptions (81 L and 75 R). With both Senate and House subscriptions combined, we were able to cover all 50 states.

3.3 Experiment Design

To answer our research questions, we designed two experiments, described next.

3.3.1 Baseline Experiment (E1). This experiment brings forth the true trends, which have not been subjected to any personalizations, of the biases in the SFAs of the three email services. The observations from this experiment also serve as the baseline for the comparison of observations from the next experiment. This experiment involved 66 email accounts, 22 accounts per service. For each email service, we assigned 6 accounts to White Americans, 6 to African Americans, 4 to Hispanic Americans, 3 to Asian Americans, and 3 to South Asian Americans. Among each set of 22 email accounts, the three age groups of 18-40, 40-65, and 65+ had 8, 8, and 6 accounts, respectively. Each email account subscribed to all of the Presidential, Senate, as well as House candidates included in our study. We kept the email accounts in the baseline experiment untouched and did not subject them to any interactions.

3.3.2 Interaction Experiment (E2). This experiment studies the impact of different interactions with the email accounts on the biases in SFAs. It contains 12×3 email accounts that subscribe to all the Presidential, Senate, and House candidates in our study. In this experiment, we randomly assigned ethnicity, age, and gender to the 12 email accounts of each service. We split this experiment into three groups containing 4×3 accounts each, and performed three different interactions, one per group. These three interactions included reading all emails, moving all emails from inbox to spam folder, and moving all emails from spam folder to inbox. We chose these three interactions for four reasons. First, reading is the most common action that one performs on an email of interest. Second, moving emails from inbox to spam folder reflects the user's preference that the user is no longer interested in such emails. Third,

moving emails from spam folder to inbox reflects the user's preference that the user is interested in such emails, and the SFA wrongly marked them as spam. Fourth, we choose to move all emails instead of randomly selecting them to avoid introducing any bias based on the content or sender of the email. We performed the reading interaction on 4×3 email accounts every 24 hours, and the other two interactions of moving from inbox to spam and spam to inbox every 5 days on their respective 4×3 email accounts.

3.4 Subscription Process

As each email account subscribes to 236 websites (2 Presidential, 78 Senate, and 156 House) on average, we had to complete 24,072 subscriptions. To automate the subscription process, we wrote a Python script that scrapes websites using Selenium library [10], and automatically fills out the subscription form to subscribe email accounts to the campaign websites. However, since there are 236 unique websites in our experiments, it was not feasible to write separate scraping codes for each of these websites. To address this challenge, we developed a general algorithm that extracts all forms from a website, determines if one or more of these forms are related to subscription, and then fills them out. To ensure that we did not miss any subscriptions, we generated the logs for the failed subscription attempts that we later completed manually by visiting the corresponding websites.

3.5 Data Set

We started our data collection on July 1st, 2020 and ended it on February 28th, 2021. However, we observed that the volume of emails from the campaign websites significantly dropped after November 20th, 2020. Therefore, we truncated the data set on November 30th, 2020 and conducted analysis on emails that we collected over these 5 months (153 days). We collected 318,108 emails across the three services. The content that we have collected for each email contains mail header fields such as Subject, From, To, Date, Message-ID, Delivered-To, Received-SPF, Received-by, Content-Type, MIME-Version, Content-Type, and message body.

4 ANALYSIS

In this section, we study the biases in the SFAs of Gmail, Outlook, and Yahoo. In the following four subsections, we present our results and observations to answer the four questions mentioned in § 1.

4.1 Political Biases in Spam Filtering Algos.

In this section, we study whether the SFAs exhibit political biases and how these biases compare across different email services. Specifically, we examine whether the SFAs in each of the three email services: 1) lean towards the right by sending more campaign emails of the left to the spam folder, 2) lean towards the left by sending more campaign emails of the right to spam, or 3) remain neutral by giving similar treatment to the campaign emails from both left and right. We answer this question by analyzing the campaign emails of the Presidential, Senate, and House candidates that we received in the 22 accounts of each email service in the baseline experiment (E1). We first analyze the aggregate political bias in the SFAs in terms of the percentage of the left and the right campaign emails

that are marked as spam. After that, we conduct a temporal evaluation of the political bias at weekly intervals. Last, we analyze the biases in the SFAs for the campaign emails from individual Senate and House candidates.

4.1.1 Aggregate Political Bias. We observed that the SFAs of the email services indeed exhibit political biases: they treat the left and the right campaign emails differently. Gmail leans towards the left as it marks a higher percentage of the right emails as spam. Outlook and Yahoo, on the other hand, lean towards the right. Each blue line in Fig. 2 shows the cumulative distribution (CDF) of the percentage of left emails marked as spam in each of the 22 email accounts of the corresponding email service. The red lines show the same for the right emails. We observe that each CDF line rises rapidly, which demonstrates that the SFA of each email service is fairly consistent across the 22 email accounts in its treatment of emails as spam. We further observe that Gmail marks a significantly higher percentage (67.6%) of emails from the right as spam compared to the emails from left (just 8.2%). Outlook is unfriendly to all campaign emails, more unfriendly to the left than to the right. It marks a higher percentage of left (95.8%) emails as spam than those of right (75.4%). Yahoo marks 14.2% more left emails as spam than the right emails. Each of these numbers above represents the average across the 22 accounts of the corresponding services. Onward, we will refer to these observations about Gmail leaning towards the left and Outlook and Yahoo towards the right as the *aggregate trend*.




Figure 2: Cumulative distribution of the %age of left (blue) and right (red) emails marked as spam in each of the 22 email accounts of each service.

4.1.2 Temporal Evolution of Political Bias. Next, we discuss 1) whether the spam percentage varies over time for the left and right campaign emails, and 2) whether there is any correlation between the spam percentage and the number of received emails. We present our observations using Fig. 3. The blue and the red solid lines in this figure show the average percentage (averaged across the 22 email accounts) of the left and the right emails, respectively, marked as spam each week. The shaded bands around the solid lines show the standard deviation for each week. The blue and the red vertical bars show the total number of emails (inbox + spam) received each week from the left and the right candidates, respectively.




Figure 3: Left y-axis: average (line plot) and standard deviation (shaded bands) of the %age of emails marked as spam each week. Right y-axis: # of left and right emails (vertical bars) received each week.

We make three important observations from Fig. 3. First, in Gmail, we observe an increasing trend in the right spam percentage over

time, whereas the left spam percentage did not vary much and remained under 15%. For the right spam, we also observe that the spam percentage increased with an increase in the number of right campaign emails. The left spam emails also show this trend, albeit, it is less apparent in Fig. 3. Thus, we plot Fig. 4, where we show the Pearson correlation (r) of the number of emails from left and right received in a week to the percentage of the left and right emails marked as spam in that week, respectively, in the 22 email accounts of each service. We observe from this figure that Gmail has a positive correlation for both the left ($r = 0.45$) and the right ($r = 0.44$). This shows that Gmail marks a larger fraction of emails as spam as the volume of emails increases. Second, Outlook is almost indifferent to the left spam percentage as it stayed above 85% for all the weeks. The variation in the number of left emails did not cause noticeable variation in the spam percentage ($r = 0.15$). The right spam percentage shows more volatility over the 22-week period, but still, there is no correlation ($r = -0.15$) between the right spam percentage and the number of emails from the right. Third, in Yahoo, the spam emails for both the left and the right initially increased and then decreased over time for most of the weeks. With the exception of the last two weeks, when the elections were over and the candidates significantly reduced the number of emails, the spam percentage for the left emails remained higher than that of the right emails. For the right candidates, the spam percentage decreased with the increase in the number of right emails ($r = -0.33$). Contrarily, the left spam percentage increased with the increase in the number of left emails ($r = 0.34$).




Figure 4: Pearson correlation between the weekly average spam %age and the weekly # of emails in the 22 accounts.

To summarize, the *aggregate trend* that we observed in §4.1.1 holds over weekly interval as well, *i.e.* Gmail leans towards the left whereas, Outlook and Yahoo towards the right. In Gmail and Yahoo, the number of emails from the left and from the right have a noticeable influence on the percentages of their emails marked as spam. However, such influence is not seen in the case of Outlook.

4.2 Impact of Political Affiliation

In §4.1, we saw that there are indeed aggregate biases in the SFAs of different email services. However, an important question still remains: *do these biases exist even when we consider only those emails from the left and the right candidates that have very similar attributes?* In other words, does the political affiliation of the sender alone play a significant role in getting an email marked as spam? An answer in affirmative would be worrisome because a sizable chunk of voting population heavily relies on these email services, and these biases could sway their decisions about who to vote for and whether to even cast their votes. There is growing evidence that the online interactions of people shape their political opinions. For example, Hu *et al.* showed that custom Google search snippets consistently amplify political bias [27].

An ideal way to obtain the answer to this question would be to have both the left and the right candidates send the same set of emails to our email accounts. Then, by comparing the percentage of those emails marked as spam when sent by the left candidates with the percentage of those same emails marked as spam when sent by the right candidates, we could answer this question. While ideal, unfortunately, this approach is impractical because it was beyond our control to decide what emails different candidates sent. However, due to the large volume of emails in our data set, it is still possible to obtain approximately the same effect as the ideal method described above using the well-known statistical method of propensity score matching (PSM) [30].

PSM is a popular statistical method that is used to preprocess data from observational studies in which it is not feasible to conduct a randomized controlled trial, *i.e.*, the studies where it is not possible to control who gets to be the member of the treatment group and who gets to be the member of the control group. PSM essentially takes various attributes, commonly known as *covariates*, of the members of the treatment and the control groups and selects appropriate members from the two groups to create a new treatment and a new control group, also known as *matched* groups, such that the distribution of any given covariate of the members of the matched treatment group is similar to the distribution of that covariate of the members of the matched control group. As a result of this similarity of distributions, the observations that one makes from the two matched groups about the effects of the treatment can be approximated to be the observations from a randomized controlled trial [30]. Note that the original treatment and control groups are also known as *unmatched* groups. We have provided a quick primer on PSM in the supplementary material, §A.1. For more details on PSM, we refer the interested readers to [19, 30, 32].

Next, we first map our emails problem to PSM and describe the covariates that we have selected. After that, we present our observations from the matched groups that PSM creates and study whether significant biases exist in SFAs even in the matched groups.

4.2.1 Mapping Emails Problem to PSM. PSM takes all the emails, selects a subset of emails from them based on the values of appropriate covariates (which we will discuss shortly), and creates the two *matched* groups. All the emails in one group are from the left and the other are from the right. An important property of these groups is that each email in one group has a corresponding email in the other group such that the values of the covariates of the two emails are very similar. Thus, the emails in the two matched groups are very similar in terms of the selected covariates. Once PSM creates the matched groups, we can then study whether or not the SFA of any given email service marked a comparable percentage of the emails in the two matched groups as spam.

We applied PSM on our email data set collected during the baseline experiment (E1). In our application of PSM on emails in any given email service, the unmatched treatment group is comprised of the emails from that political affiliation whose emails were marked more as spam. For example, for Gmail, we considered emails from the right candidates as the treatment group and the emails from the left candidates as the control group because more emails from right were classified by Gmail as spam compared to the emails from left. The reason behind considering emails from the disadvantaged

political affiliation (*i.e.*, political affiliation for which a larger percentage of emails was marked as spam) as the treatment group is that our goal is to determine whether that group of emails has been *treated* unfairly by the SFA compared to the emails from the other affiliation. Table 1 summarizes which political affiliation’s emails we considered as the treatment group and which affiliation’s as the control group in applying PSM on each email service.

Table 1: Assignment of emails to treatment group ($Z = 1$) and control group ($Z = 0$) for the three services.

Z	Gmail	Outlook	Yahoo
1	Right	Left	Left
0	Left	Right	Right

4.2.2 Selection of Covariates. The covariates whose values PSM uses should be the features that the SFAs use in determining whether any given email is spam or not. Unfortunately, none of the three email services providers in our study have publicly disclosed what these feature are. However, researchers have studied SFAs in the past and have identified five types of features that appear to influence the decisions of SFAs [17, 21]. These include 1) the meta data about email content, 2) the actual content of the email, 3) the attributes of the sender, 4) the reaction of the recipient (such as reading an email, replying to it, marking an email as spam or not-spam), and 5) the demographics of the recipient. Among these five types of features, the values for the last two types of features are not determined by the senders rather by the receivers (*i.e.*, us), and thus do not need PSM. We will analyze the impacts of the reactions of the recipients in §4.3 and of their demographics in §4.4. The values of the first three types of features, however, are beyond our control as they depend on who is sending the emails and what content are they including in the emails. Thus, these three types of features require the use of PSM.

For the third type of features, *i.e.*, the attributes of the sender, the only information we have is the IP address of the SMTP server used by the sender. Our analysis revealed that over 80% of the emails from the right were sent using just four digital marketing organizations, namely BlueHornet [8], Acoustic [3], Amazon-SES [4], and MailGun [7]. Similarly, over 80% of the emails from the left were sent using Blue State Digital [5], NGP VAN [9], Amazon-SES [4], and MailGun [7]. As these digital marketing organizations are among the largest in the world, it is highly unlikely that the SFAs would mark emails as spam just because they were sent using one of their SMTP servers. Thus, we do not perform PSM using the IP address of the sender’s SMTP server as a covariate.

This leaves us with the first two feature types. Next, we describe the covariates that we selected for these two types of features.

T1: Meta Data based Covariates. The meta data based covariates capture the properties of the contents of the email instead of the actual contents of the email. We calculated values for ten meta data covariates, listed and defined in Table 2, from each email. In selecting these covariates, one of the properties that we considered was that the distribution of any given covariate in the unmatched treatment group should have some overlap with the distribution of that covariate in the unmatched control group. A covariate whose distributions do not have any overlap in the two groups would cause

the propensity scores of the otherwise similar emails across the two groups to become very different. This could keep PSM from creating good email pairs (see §A.1 to note that each pair has one email from unmatched treatment group and the other email from the unmatched control group such that the difference between the values of their covariates is under a threshold).

Table 2: The ten T1 Covariates

Covariate	Description
Content Lexicon	The number of words in the text of the email.
# Sentences	The number of sentences in the text of the email.
Readability Score	Calculated using the widely used Gunning Fog index [24, 28, 33].
Social Media	The number of times social media platforms are mentioned in the text body.
Thread	Whether or not the subject starts with "Re:" or "Fwd:"
Upper Case	The number of upper case words in the subject and body.
Special Characters	The number of special characters such as !, @, * in the subject and the content.
# HREFs	The number of HREFs present in the raw body of the email.
# HTTP	The number of HTTP(s) links in the raw email body.
# Images	The number of images referred or attached in the raw body.

T2: Content based Covariates. To apply PSM on emails based on their contents, we need to create pairs of matched groups such that the topics of the emails in any given pair of matched groups are closely related and the text of the emails in that pair of matched groups has similar terms. To obtain covariates that can result in such pairs of matched groups, we use results from a recent study on election emails [2]. In this study, the authors applied a structural topic model on the content of over 105K emails and determined that there are 65 unique topics that political campaign emails are about. They partitioned these 65 topics into six categories: C1) political campaigns and events (topics such as Trump MAGA, primary debate, etc.), C2) political issues (LGBTQ, guns, etc.), C3) voter mobilization (winning races, voting, etc.), C4) explicit fundraising (donations, fundraising deadlines, etc.), C5) ceremonial niceties (social media, holiday wishes, etc.), and C6) miscellaneous (signing petitions, surveys, etc.). For each topic, the authors presented 15 FREX terms (terms that are frequent as well as exclusive for a given topic) and 15 probability terms (terms with the highest likelihood of appearing in an email on a given topic). For the complete list of topics as well as the FREX and the probability terms for each topic, we refer the interested readers to Table 2 in [2].

In our PSM analysis where we consider the contents of the emails, we create six pairs of matched groups, one pair per category. To create a pair of matched groups corresponding to any given category, we use the sum of the frequency of the FREX terms and the sum of the frequency of the probability terms of each topic in any given email as covariates. To clarify with an example, consider a hypothetical category that has four topics. For this category, we will have eight covariates, two per topic (one from FREX terms and the other from probability terms). The value of the covariate from FREX terms for any given topic for any given email is calculated

by counting the number of times the FREX terms for that topic appear in the content (subject and body) of that email. The value of the covariate from probability terms is calculated in the same way by counting the number of appearances of the probability terms.

4.2.3 Applying the PSM. Next, we apply PSM to create the matched treatment and control groups of emails. We emphasize that we do not use all the covariates of the two types (i.e., T1 and T2) together. We separately apply PSM on T1 covariates and on the covariates of each of the six categories in T2. To make the paper self-contained, we have presented the technical details of how our implementation of PSM worked in the supplementary material, §A.2.

4.2.4 Observations from the Matched Groups. Now that we have created the matched groups, we study whether the SFAs demonstrate similar biases in the matched groups as we saw in §4.1.1. Fig. 5 plots the difference between the percentage of emails marked as spam in the matched treatment group and the percentage of emails marked as spam in the matched control group (i.e., $Treatment\ Spam\ \% - Control\ Spam\ \%$) for each of the three email services and for each of the 7 matched groups (one matched group generated using T1 covariates and six matched groups generated using the covariates of the six categories in T2.) Recall from Table 1 that the treatment group for Gmail is comprised of the emails from right candidates, while for Outlook and Yahoo, it is comprised of emails from the left candidates. For comparison, Fig. 5 also plots the absolute difference (green line) between the percentage of emails marked as spam in the unmatched treatment group and the percentage of emails marked as spam in the unmatched control group for each of the three email services.




Figure 5: $Treatment\ Spam\ \% - Control\ Spam\ \%$. The color of any given bar represents the political affiliation of the emails in the corresponding treatment group.

From this figure, we observe that the aggregate trend that we observed in §4.1.1 holds even in the matched groups and the values of $Treatment\ Spam\ \% - Control\ Spam\ \%$ in the matched groups are fairly close to those observed in the unmatched groups (shown with green lines). In Gmail, for the matched groups obtained using C1 and C2, we observed a 17.1% and 16.2% decrease, respectively. This happened because in the unmatched groups, Gmail was marking almost all the emails on the topics of *Lindsey Graham* and *South Carolina* in C1 and on *radical left* in C2 sent by the candidates from the right as spam. Since there were very few emails on these topics from the left candidates, the matched groups did not contain many emails on these topics, which helped reduce the value of $Treatment\ Spam\ \%$. Nonetheless, while there was a decrease, the values of $Treatment\ Spam\ \% - Control\ Spam\ \%$ were still $> 40\%$.

4.3 Impact of Interactions

In this section, we study whether the interactions of the users with their email accounts cause the biases to decrease or increase. For

this, we analyze the emails collected during the Interaction Experiment E2 (§3.3.2) for the three different types of interactions that we performed on the campaign emails: 1) reading all emails, 2) moving all emails in inbox to spam folder (I→S), and 3) moving all emails in spam folder to inbox (S→I). We started the reading interaction with the campaign emails on August 3, 2020 and repeated daily, and the I→S and S→I interactions on September 13, 2020 and repeated every 5 days. Recall from §3.3.2 that we performed these three interactions on three different sets of email accounts.

To study the impact of these interactions on the political biases in the SFAs, we present the observations from this interaction experiment (E2) relative to the observations from the baseline experiment (E1). We measure this impact in two ways. First, to observe any changes in the percentages of emails marked as spam as a result of the interactions, we compute the difference between the left (right) spam percentage in E2 and the left (right) spam percentage in E1. Second, to analyze whether the interaction increased or decreased the political bias in the SFAs, we compute the Bias-Index (*BI*) defined as:

$$BI = \frac{|E2 \text{ Left Spam } \% - E2 \text{ Right Spam } \%|}{|E1 \text{ Left Spam } \% - E1 \text{ Right Spam } \%|}$$

The values of *BI* can lie in three ranges, interpreted as:

- $0 < BI < 1$: the bias in E2 dropped lower than the bias in E1 in response to the given interaction.
- $BI = 1$: the bias in E2 stayed the same as E1.
- $BI > 1$: the bias in E2 increased compared to E1.

4.3.1 Reading All Emails (Inbox + Spam). When a user reads political emails, the spam percentage should decrease because, arguably, the user is showing interest in the received content. However, we observed that the SFAs in the three email services reacted differently to the reading interaction. Fig. 6 presents the impact, at weekly intervals, of reading all emails. The negative values of spam percentage difference in this figure demonstrate a decrease in spam percentage relative to the baseline experiment (E1) and vice versa. We observe from this figure that in Gmail, the spam percentage marginally decreased for both the left and the right emails while still maintaining Gmail’s leaning towards the left. Due to only a marginal impact, the *BI* stayed approximately at 1. In Outlook, the percentage of right spam kept decreasing over time while that of left spam stayed unchanged. This increased the right-leaning of Outlook further, which resulted in a slightly increasing trend in *BI* overtime. In Yahoo, we observed a counter-intuitive trend: the spam percentage of both the left and the right emails slightly increased due to the reading interaction. The increase was more for right emails compared to the left emails, which resulted in slight decrease in the *BI*. Nonetheless, the changes in the spam percentages across all three services were minimal.




Figure 6: Impact of reading interaction on the left and right spam percentages (left y-axis) and the political bias index (right y-axis).

To see the net impact of the interactions on the bias, for each email service, Fig. 7 shows the average percentage of the left emails

and of the right emails marked as spam in the baseline experiment, after all reading interactions, after all I→S interactions, and after all S→I interactions. This figure makes it clear that after the reading interactions, for all three email services, while there are minor changes in the percentages of the left and the right emails marked as spam, when compared to the baseline experiment, the magnitudes of the changes are negligible. Thus, the reading interaction did not have any significant impact on the political bias of any of the three email services.




Figure 7: Percentage of left and right emails marked as spam in baseline experiment and after the reading, I→S, and S→I interactions.

4.3.2 Moving All Inbox Emails to Spam Folder. Fig. 8 presents the impact of the I→S interaction on spam percentage and political bias index for the three services. We observe from this figure that Gmail starts marking a significantly higher percentage of left emails as spam in response to the I→S interaction, reducing its left-leaning. The percentage of right emails marked as spam also increased. These increases are intuitive because when a user moves certain emails to the spam folder, the user is expressing that, in future, such emails should not appear in the inbox. The increase in left spam percentage was significantly higher (by 45%) compared to the right spam percentage (by 11%), which is also intuitive because, in the baseline experiment, the percentage of right spam emails in Gmail was already a lot higher than the percentage of left spam emails. Consequently, the *BI* of Gmail reduced significantly, showing that the I→S interaction significantly increased the fairness of Gmail towards the left and right emails when compared to the observations from the baseline experiment (E1). Fig. 7 shows that after all five I→S interactions, Gmail marked 54.2% and 83.9% emails from left and right, respectively, as spam across the 4 email accounts assigned to the I→S interaction for Gmail in E2. While Gmail’s biasness reduced, it still stayed slightly left-leaning.




Figure 8: Impact of I→S interaction on the left and right spam percentages and the bias index. X-axis shows I→S interaction number.

Yahoo demonstrated a similar trend where the spam percentage of both left and right emails increased. As the increase in the right spam percentage was more than the left spam percentage compared to the baseline experiment, Yahoo’s *BI* improved (*i.e.*, decreased) in response to the I→S interaction. Fig. 7 shows that after the I→S interactions Yahoo became almost unbiased. Outlook, however, did not show a significant impact on the percentage of emails marked as spam. Consequently, Outlook’s *BI* improved only marginally in response to the I→S interaction. This was expected as Outlook was already marking most left (96%) as well as right (81%) emails as spam (Fig. 2), and thus the room for marking more emails as spam was relatively small.

To summarize, due to the I→S interaction, the political bias in all services improved, significantly in Gmail, moderately in Yahoo, while only marginally in Outlook.

4.3.3 Moving All Spam Emails to Inbox. When a user moves emails from spam to inbox, the spam percentage should decrease because the user is showing interest that such emails should appear in the inbox. The response of Gmail to the S→I interaction follows this intuition while that of Outlook and Yahoo does not. Fig. 9 shows how much the spam percentage and bias index for the three services changed after each S→I interaction. The negative values for both the left and right emails for Gmail show that Gmail starts putting a higher percentage of emails from both sides in inbox after just the first S→I interaction. Fig. 7 shows that after the five S→I interactions, on average, Gmail marks just 5.34% of the right emails as spam (compared to 67.6% in the baseline experiment) and 0% of the left emails as spam (compared to the 8.2% in the baseline experiment). Thus, Gmail still maintains its left leaning, but not very strongly anymore (*BI* also dropped significantly).




Figure 9: Impact of S→I interaction on spam percentages and *BI*. X-axis shows S→I interaction number.

The response of Outlook to the S→I interaction was counter-intuitive. Although Outlook marked the largest number of both left and right emails as spam in the baseline experiment, its reduction in spam percentages in response to the S→I interaction was only marginal for both the left (by 5%) and the right (by 8.3%) emails. This resulted in an increase in Outlook’s right-leaning further (*BI* increased). Yahoo demonstrated similar behavior as Outlook: while it marginally decreased its spam percentage, the decrease was slightly more for the right emails compared to the left emails, increasing its right-leaning (*BI* increased here as well).

To summarize, due to the S→I interaction, the political bias in Gmail reduced significantly. However, unexpectedly, it increased in both Outlook and Yahoo because neither of the two services reacted noticeably to user’s desire to not mark the emails as spam that the two services were marking as spam.

4.4 Impact of Demographics

Recall from §3.1 that our email accounts are comprised of multiple combinations of three age groups, five ethnicities, and two genders. We observed from our data that neither the age group, nor the ethnicity, nor the gender of the account holder had any impact on how SFAs treated the emails. Due to this and due to the space constraints, we have not shown any corresponding figures.

5 CONCLUDING DISCUSSION

In this paper, we conducted a large-scale measurement study to examine biases in SFAs of Gmail, Outlook, and Yahoo during the 2020 US elections. We subscribed to a large number of left and right Presidential, Senate, and House candidates using several email accounts on the three email services. Next, we first summarize the

answers that our analysis has revealed to the four questions we posed in §1, and after that provide a concluding discussion to wrap up the study.

Summary. Our observations in §4.1 to answer **Q1** revealed that all SFAs exhibited political biases in the months leading up to the 2020 US elections. Gmail leaned towards the left (Democrats) whereas Outlook and Yahoo leaned towards the right (Republicans). Gmail marked 59.3% more emails from the right candidates as spam compared to the left candidates, whereas Outlook and Yahoo marked 20.4% and 14.2% more emails from left candidates as spam compared to the right candidates, respectively.

Our analysis in §4.2 to answer **Q2** showed that the aggregate biases that we observed in the complete email data set persisted even when we considered only those emails from the left and right candidates that had very similar covariates. Thus, it appears that the political affiliation of the sender plays role in getting an email marked as spam.

Our observations in §4.3 to answer **Q3** showed that Gmail responded significantly more rapidly to user interactions compared to Outlook and Yahoo. While the political bias in Gmail stayed unchanged after the reading interaction, it decreased significantly due to the I→S and S→I interactions. Contrary to this, the political bias in Outlook, increased due to the reading and S→I interactions while it remained almost unchanged after the I→S interaction. In Yahoo, the bias decreased due to the reading and I→S interactions while it increased due to the S→I interaction. While the political biases changed in response to various interactions, Gmail maintained its left leaning while Outlook and Yahoo maintained their right leaning in all scenarios.

Finally, our observations in §4.4 to answer **Q4** showed that the demographic factors, including age, ethnicity, and gender, did not influence the political bias of SFAs.

Discussion. We conclude the paper with three thoughts. The observation that the aggregate trend that we observed in §4.1.1 for the unmatched groups showed up in the matched groups as well is rather worrying because this implies that the SFAs of email services do have quite a noticeable bias. They mark emails with similar features from the candidates of one political affiliation as spam while do not mark similar emails from the candidates of the other political affiliation as spam. Thus the political affiliation of the sender appears to play some role towards the decision of the SFAs. Arguably, there is also this possibility that the SFAs of email services learnt from the choices of some voters marking certain campaign emails as spam and started marking those/similar campaign emails as spam for other voters. While we have no reason to believe that there were deliberate attempts from these email services to create these biases to influence the voters, the fact remains there that their SFAs have learnt to mark more emails from one political affiliation as spam compared to the other. As these prominent email services are actively used by a sizable chunk of voting population and as many of the voters today rely on the information they see (or don’t see) online, such biases may have an unignorable impact on the outcomes of an election. It is imperative for the email services to audit their SFAs to ensure that any properties of the sender that they consider in determining whether any given email is spam or not are not, unknowingly, putting one side at an advantage compared to the other.

Second, the general perception is that when a user reads emails, marks them as spam, or moves them from the spam folder to inbox, the SFA adapts to user’s preferences. While our observations agree with this perception, very strongly for Gmail and to a smaller extent for Outlook and Yahoo, this adaptation does not necessarily eliminate the political bias. Some interactions do reduce the bias, but that effect is not consistent across all the email services. In other words, we did not find any consistent actions that one could recommend to users to help them reduce the bias in the way the SFA treats political emails that are sent to them.

Third, if an undecided voter receives too many emails from one political party, there is a likelihood that they may get swayed towards that party. As users open their spam folders very infrequently, it is unlikely that most undecided voters will undertake the effort to open the spam folder and mark some campaign emails as not-spam to make the SFA unbiased. Therefore, it is important for the SFAs to be unbiased at the outset without relying on explicit user feedback.

We conclude by noting that fairness of spam filtering algorithm is an important problem that needs dedicated attention from email service providers, particularly due to the enormous influence that electronic communication has in our lives today and will have going forward. However, it is not an easy problem to solve. Attempts to reduce the biases of SFAs may inadvertently affect their efficacy. Therefore, there is an imminent need to develop techniques that reduce the biases of SFAs while simultaneously ensuring that the users do not receive unwanted emails.

REFERENCES

- [1] 2020. GOP Congressman Asks Google CEO Sundar Pichai Why His Campaign Emails Always Go to Spam. <https://www.thewrap.com/greg-steube-gmail-spam-sundar-pichai/>
- [2] 2020. Manipulative tactics are the norm in political emails: Evidence from 100K emails from the 2020 U.S. election cycle. <https://electionemails2020.org>
- [3] 2022. Acoustic Digital Marketing. <https://www.acoustic.com/>
- [4] 2022. Amazon Simple Email Service. <https://aws.amazon.com/ses/>
- [5] 2022. Blue State. <https://www.bluestate.co/>
- [6] 2022. General Prohibition Against Using Official Resources for Campaign or Political Purposes. <https://ethics.house.gov/campaign>
- [7] 2022. MailGun. <https://www.mailgun.com/>
- [8] 2022. Mapp’s Bluehornet. <https://mapp.com/bluehornet/>
- [9] 2022. NGP VAN. <https://www.ngpvan.com/>
- [10] 2022. Selenium driver. <https://www.selenium.dev/>
- [11] 2022. US Presidential Elections Campaign Emails Dataset. https://drive.google.com/drive/u/1/folders/1DgL_K1Vmfj7v6Rvz3BtL6PG7sPlwZ9xl
- [12] 2022. US Senate Ethics. <https://www.ethics.senate.gov/public/>
- [13] Alan Agresti. 2003. *Categorical data analysis*. Vol. 482. John Wiley & Sons.
- [14] I Androutsopoulos and et al. 2000. An experimental comparison of naive bayesian and keyword-based anti-spam filtering with personal e-mail messages. *SIGIR* (2000).
- [15] Peter C Austin. 2011. An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivariate behavioral research* 46, 3 (2011), 399–424.
- [16] R Balebako and et al. 2012. Measuring the effectiveness of privacy tools for limiting behavioral advertising. *In Web 2.0 Workshop on Security and Privacy* (2012).
- [17] Enrico Blanzieri and Anton Bryl. 2008. A survey of learning-based techniques of email spam filtering. *Artificial Intelligence Review* 29, 1 (2008), 63–92.
- [18] M Butler. 2003. Spam – the meat of the problem. *Computer Law and Security Report* 19 (2003), Issue 5.
- [19] Marco Caliendo and Sabine Kopeinig. 2008. Some practical guidance for the implementation of propensity score matching. *Journal of economic surveys* 22, 1 (2008), 31–72.
- [20] G Cormac and T Lynam. 2005. Spam corpus creation for TREC. *Proceedings of second conference on email and anti-spam* (2005).
- [21] Emmanuel Gbenga Dada, Joseph Stephen Bassi, Haruna Chiroma, Adebayo Olu-sola Adetunmbi, Opeyemi Emmanuel Ajibuwa, et al. 2019. Machine learning for

- email spam filtering: review, approaches and open research problems. *Heliyon* 5 (2019).
- [22] A Datta and et al. 2015. Automated experiments on ad privacy settings: A tale of opacity, choice, and discrimination. *PETS* (2015).
- [23] R Epstein and R Robertson. 2015. The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. *PNAS* 112 (2015), Issue 33.
- [24] Robert Gunning et al. 1952. Technique of clear writing. (1952).
- [25] A Hannak and et al. 2013. Measuring personalization of web search. *WWW* (2013).
- [26] A Hannak and et al. 2017. Measuring personalization of web search. (2017).
- [27] D Hu and et al. 2019. Auditing the Partisanship of Google Search Snippets. *WWW* (2019).
- [28] Yaser Keneshloo, Shuguang Wang, Eui-Hong Han, and Naren Ramakrishnan. 2016. Predicting the popularity of news articles. *In Proceedings of the 2016 SIAM International Conference on Data Mining*. SIAM, 441–449.
- [29] H Le and et al. 2019. Measuring Political Personalization of Google News Search. *WWW* (2019).
- [30] Wei Pan and Haiyan Bai. 2015. *Propensity score analysis: Fundamentals and developments*. Guilford Publications.
- [31] C Puschmann. 2017. How significant is algorithmic personalization in searches for political parties and candidates? (2017).
- [32] Paul R Rosenbaum and Donald B Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* 70, 1 (1983), 41–55.
- [33] Philipp Singer, Emilio Ferrara, Farshad Kooti, Markus Strohmaier, and Kristina Lerman. 2016. Evidence of online performance deterioration in user sessions on Reddit. *PLoS one* 11, 8 (2016), e0161636.
- [34] Spamhaus. 2005. The definition of spam. (2005).
- [35] B Taylor, D Fingal, and D Aberdeen. 2007. The War Against Spam: A report from the front line. (2007).
- [36] Robert Tibshirani. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* 58, 1 (1996), 267–288.

A SUPPLEMENTARY MATERIAL

A.1 PSM – A Primer

To create the matched treatment and control groups, PSM performs two steps. In the first step, it calculates a propensity score for each member of the unmatched treatment and control groups using the values of the covariates of that member. Propensity score for any given member is defined as the probability that that member belongs to the treatment group conditioned on the member’s covariates. Formally, for any member i , let Z_i represent the indicator random variable that is 1 if the member i belongs to the treatment group and 0 if the member i belongs to the control group. Let the number of covariates that can be observed for any given member be K and let X_{ij} represent the j^{th} covariate of the i^{th} member. The propensity score for the i^{th} member, represented by e_i , is then given by the following equation.

$$e_i = Pr(Z_i = 1 | X_{i1}, X_{i2}, \dots, X_{iK})$$

In the second step, PSM uses the scores that it calculated for all the members of the unmatched treatment and control groups, and creates matched treatment and control groups. More specifically, for each member i of the unmatched treatment group, PSM finds a member j in the unmatched control group such that the absolute difference between the scores of the two members (*i.e.*, $|e_i - e_j|$) is below a certain threshold. If it is able to find such a member j in the control group, it puts the member i of the unmatched treatment group into the matched treatment group and the member j of the unmatched control group into the matched control group. If it is not able to find a member in the control group for which $|e_i - e_j|$ is below the threshold, it discards the member i of the unmatched treatment group and moves on to member $i + 1$ and repeats the step above. After iterating through all the members of the unmatched treatment group, PSM returns the new treatment




Figure 10: Histograms of the propensity scores of the emails in the unmatched treatment (T) and the unmatched control (C) groups using T1 covariates.

and control groups, where the distribution of any given covariate is very similar in the two matched groups.

A.2 Applying PSM on Emails Data Set

Next, we describe how we apply the two steps of PSM (mentioned in §A.1) to create the matched treatment and matched control groups of emails from the unmatched treatment and unmatched control groups. We emphasize that we do not use all the covariates of the two types described above together. Instead, we separately apply the two steps of PSM on T1 covariates and on the covariates of each of the six categories in T2.

Step 1: Propensity Score Estimation. To estimate the propensity score for each email, we first create a logistic regression model using covariates as the independent variables and Z as the dependent variable. Recall from §A.1 and Table 1 that Z is equal to 1 for any email that belongs to the treatment group and is equal to 0 for emails that belong to the control group. As different covariates can have different amounts of correlations with Z , we incorporate lasso regularization when creating the logistic regression [13, 36]. Lasso regularization discounts the effect of covariates that do not have a noticeable correlation with Z . After creating the logistic regression model, for each email, we feed its covariates to this model and the model outputs a value between 0 and 1, which is the propensity score for the email. The method that we just described is one of the most common methods in literature to estimate the propensity scores [15, 30].

Step 2: Matching. Fig. 10 shows the distributions of the propensity scores in the unmatched treatment (marked as T) and control (marked as C) groups for the three services when using T1 covariates. We observe from this figure that for each email service, there is a decent overlap between the distributions of the left and the right emails. This shows that we should be able to create matched treatment and control groups with large enough number of emails in each group such that any observations that we make from them are statistically significant. We made very similar observations for the T2 covariates.

As with the propensity score estimation, several methods have been proposed in literature to create matched groups. We again employed one of the commonly used methods, known as caliper matching [15]. Caliper matching not only creates well-matched pairs of emails but also excludes any emails from the treatment group for which a good match in the control group is not available. Caliper matching matches any given email i in the unmatched treatment group with that email j in the unmatched control group




Figure 11: Histograms of the propensity scores of the emails in the matched treatment (T) and the matched control (C) groups using T1 covariates.

for which the absolute difference between propensity scores of the two emails is minimum and at the same time less than a threshold B . The threshold B , also known as the caliper width, is defined as $B = \kappa \times \sigma$ where σ is the standard deviation of all the propensity scores in the unmatched treatment and control groups, and κ is a multiplicative constant. In [15], the authors suggested to set $\kappa = 0.2$.

The quality of the matched treatment and control groups is quantified in terms of standardized mean difference (SMD), given by the following equation:

$$\text{SMD} = (\bar{e}_T - \bar{e}_C) / \sqrt{(s_T^2 + s_C^2) / 2}$$

where \bar{e}_T and s_T^2 represent the mean and variance of the propensity scores in the matched treatment group and \bar{e}_C and s_C^2 represent the same for the matched control group. The quality of the matched groups is considered good when the absolute value of the SMD is ≤ 0.1 [15].

To create the matched treatment and control groups, we start with $\kappa = 0.2$ and find an email in the unmatched control group corresponding to each email in the unmatched treatment group as per the caliper matching criteria mentioned above. Next, we check the quality of the resulting matched treatment and control groups by calculating the value of SMD. If the absolute value of SMD turns out to be greater than 0.1, we tighten the caliper width by decrementing κ in the steps of 0.02, and repeat the steps again until $|\text{SMD}|$ drops below 0.1. At this point, we obtain the final matched treatment and control groups that have the same number of emails, and the distribution of the propensity scores as well as that of any covariate is very similar across the two groups. For example, Fig. 11 shows the distribution of the propensity scores in the matched treatment and in the matched control groups for the three services using the T1 covariates. We indeed observe that for each email service, the distribution of the propensity scores in the matched treatment group very closely overlaps with the distribution of the propensity scores in the matched control group. This shows that the two matched groups for each email service that PSM has created for us are indeed comprised of two very similar sets of emails. We made very similar observations for the T2 covariates.