Do Social Disorders Facilitate Social Engineering? A Case Study of Autism and Phishing Attacks

Ajaya Neupane
Department of Computer Science
University of California Riverside
ajaya@ucr.edu

Kiavash Satvat
Department of Computer Science
University of Alabama at Birmingham
kiavash@uab.edu

Nitesh Saxena
Department of Computer Science
University of Alabama at Birmingham
saxena@uab.edu

Despina Stavrinos
Department of Psychology
University of Alabama at Birmingham
dstavrin@uab.edu

Haley Johnson Bishop
Department of Psychology
University of Alabama at Birmingham
haleyj89@uab.edu

ABSTRACT

Social engineering is a well-established and well-studied threat especially against healthy computer users. Little studied, however, is the level of vulnerability to social engineering attacks against people with medical conditions. Social disorders in particular may make people more susceptible to such attacks. In this paper, as an initial line of investigation into this understudied research line, we launch a study of *phishing*, a prominent social engineering attack, against people suffering from *autism spectrum disorder*, a unique developmental disorder characterized by hampered social skills and communication.

We present a study of phishing detection with two groups of participants each with 15 participants, one diagnosed with autism and other without autism, in which they were asked to distinguish real versions of certain websites from their fake counterparts. Given the known gullibility and social vulnerability of users with autism, our study is designed to test the hypothesis that individuals with autism will be more prone to such attacks in contrast to healthy participants of prior studies. Our results, however, do not support this hypothesis demonstrating that participants with autism are not more vulnerable to phishing attempts. We attribute this result to the unique characteristics of users with autism including attention to detail, strong memory of factual information and diverse way of thinking, which are skills that folklore assumes may actually make users with autism highly qualified for cybersecurity careers. Overall, our work serves to demonstrate that targeted (spear) phishing attacks against Internet users suffering from autism may not be more successful compared to untargeted attacks against the user population without autism. It also highlights that social disorders may not necessarily facilitate social engineering attacks.

ACM Reference Format:
Ajaya Neupane, Kiavash Satvat, Nitesh Saxena, Despina Stavrinos, and Haley Johnson Bishop. 2018. Do Social Disorders Facilitate Social Engineering? A Case Study of Autism and Phishing Attacks. In *2018 Annual Computer Security Applications Conference (ACSAC '18), December 3–7, 2018, San Juan, PR, USA.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3274694.3274730

1 INTRODUCTION

Phishing is a type of social engineering attack where the attacker directs users to fake websites having the look and feel similar to that of a real website, with the intention of obtaining their sensitive and private credentials. The attackers then exploit the information gathered from such an attack to misuse the victim's account.

Social engineering attacks are a well-established and well-studied threat in the literature especially against normal computer users. Researchers have conducted a number of human-centered phishing detection studies with such users (e.g., [13–15, 29, 32, 34]), to understand their task performance in identifying phishing scams and the effectiveness of anti-phishing toolbars, security indicators and phishing warnings in preventing them. However, little work has been conducted to investigate the level of vulnerability to social engineering attacks against people with different medical conditions. The Internet is used by a wide variety of people, and it is therefore important to understand how these people perform against such security attacks. For example, a recent study by Oliveria et al. involving aged and young people revealed that older men and women might be more susceptible to spear-phishing attacks [27]. Such studies will help us understand these specific populations and develop specialized security parameters, controls and warnings.

In particular, social disorders may make people more susceptible to social engineering attacks. In this paper, as an initial line of investigation into this understudied research line, we launch a study of phishing, a prominent social engineering attack, against individuals with Autism Spectrum Disorder (ASD), a unique developmental disorder characterized by hampered social skills and communication. ASD is a neurodevelopmental disorder characterized by deficits in social communication and social interaction, the presence of repetitive behaviors and restricted interests, and a complex combination of diminished and intact cognitive abilities [9]). ASD is one of the fastest-growing developmental disabilities in the United States with a prevalence that has increased drastically, from 1 in 88 children in 2008, to 1 in 68 children in 2014 [5]. The increasing prevalence rates of ASD mean that more and more people diagnosed with ASD in their early childhood are now transitioning into adulthood where

ACSAC '18, December 3–7, 2018, San Juan, PR, USA
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-6569-7/18/12…$15.00
https://doi.org/10.1145/3274694.3274730

they will soon be faced with the challenges of living independently. One of these challenges includes the ability to make decisions and discern the intentions of other people. As social interactions (i.e., conversational rules), non-verbal (i.e., body language) and non-literal (i.e., sarcasm, deception) are known areas of impairment for individuals with ASD, they are often taken advantage of [16, 35]. In this context, we conjecture that the users with autism might be more vulnerable to the social engineering attacks, especially phishing.

Our Contributions: In this paper, we contribute to the scientific literature the first study of autism and phishing. Specifically, we conduct a phishing detection study with 15 users with autism and 15 users without autismand report on their task accomplishment when asked to identify the real and fake versions of websites. In line with the prior studies of phishing detection [13, 25], we designed and conducted our phishing detection study to understand the performance of users with autism. We carefully enrolled the participants with similar demographics to our study. First, we perform a quantitative analysis on the task performance of the participants with and without autism in the phishing detection task. In contrast to our intuition, our results suggest that the users with autism performed no less than the users without autism. Our results therefore do not support our hypothesis based on the fact that the users with autism are susceptible to deception. Second, we perform the qualitative analysis of the responses of participants with and without autism regarding security scams and phishing detection. The results we obtained from our phishing detection task are explained by the answers the participants with autism reported to our post-test questionnaire. We found out that the participants with autism noticed the missing security locks/certificates, logos, and obfuscated URL in the fake websites presented to them.

We attribute the above result to the several unique traits of users with autism including paying attention to detail, being methodical, having strong memory of factual information and a diverse way of thinking. Folklore assumes that these are needed skills in the cybersecurity workforce for which users with autism are in high demand [8]. Broadly, our work suggests that targeted spear phishing attacks against Internet users with autism may not be more successful compared to generalized attacks against typical user population.

2 BACKGROUND & PRIOR WORK

In this section, we provide the background necessary to understand our work and review relevant prior work.

2.1 Autism Spectrum Disorder

Autism Spectrum Disorder (ASD or simply autism) is a neurodevelopmental disorder characterized by deficits in social communication and social interaction, the presence of repetitive behaviors and restricted interests, and a complex combination of diminished and intact cognitive abilities [9]). ASD is one of the fastest-growing developmental disabilities in the United States with a prevalence that has increased drastically from 1 in 88 children in 2008, to 1 in 68 children in 2014 [5]. Though we were unable to find the exact population statistics of the adults with autism, we argue that the population will be huge in future as these children with ASD grow up. And as the children with ASD become more independent in adolescence and adulthood, more subtle and potentially dangerous forms of deception become commonplace.

2.2 Related Work

Researchers have conducted a number of human-centered phishing detection studies (e.g., [13–15, 25, 26, 29, 32, 34]), which focus on users' task performance in identifying phishing scams and the effectiveness of anti-phishing toolbars, security indicators and phishing warnings. Dhamija et al. recruited 22 participants and asked them to identify real and fake websites and performed the quantitative and qualitative analysis to understand people fall for phishing attacks [13]. Egelman et al. examined the effectiveness of the phishing warnings and analyzed why they fail users [14]. Friedman et al. performed user study across regions to understand their conceptions of web security [15]. Schechter et al. performed a study to evaluate the website authentication and the effect of role playing on usability studies with mostly university students [29]. Sunshine et al. performed a study on the effectiveness of SSL warnings [32]. Wu et al. performed a study to understand the effectiveness of security toolbars in preventing phishing attacks [34]. Neupane et al. [25, 26] reported on the neural activities of the users when they are performing phishing detection and warnings tasks. However, while our study specifically focuses on the users with autism, all of these studies were conducted with general population. They have usually revealed that such users do not pay attention to the browser-based phishing cues and often make incorrect choices.

Recently, researchers have been performing user studies on specific sets of populations [19–24, 27]. Oliveria et al. conducted a phishing study on aged and young people revealed that older men and women might be more susceptible to spear-phishing attacks [27]. Such studies will help us understand these specific population and develop specialized security parameters, and warnings. Lerner et al. conduct a study of encrypted email client with the lawyers and journalists and report that they may have diverse operational constraints and threat models, and the one-size-fits-all solution may not be useful to them [19]. Mcgregor et al. report on the usable security study on the successful collaboration of journalists in keeping the investigation of Panama paper secret before their release [24]. Mcgregor et al. also conducted user studies with journalists and reported on the practical and cultural constraints that can limit the computer security and privacy practices of the journalism community [23] and how existing security tools may not be effective for journalists [22]. These studies suggest that the current security tools may not be suitable for all segments of the population and demands more investigation to understand on needs of special target populations. Marne et al. [21] studied the usability of learning system-assigned passwords on the people with learning disabilities. Mare et al. [20] analyzed the security and privacy challenges in designing services for low-literate users in developing regions.

Although relatively little research has been conducted to examine the topic of gullibility and deception in the autism population, findings have been consistent in demonstrating that individuals with ASD are more trusting [36] and more easily deceived than their typically developing counterparts [12, 35]. This susceptibility to lies and deception can be harmful to users with autism in the context of school-age relationships and bullying [31]. Unlike these studies, our study shows that the users with autism might not be as susceptible to phishing attacks as one may have expected. Their ability to notice small changes, and the detailed oriented nature might be helpful in identifying such attacks [2, 4].

Several of the impairments seen in individuals with ASD may contribute to this susceptibility to deception including impairments in executive function (e.g., theory of mind), social interactions (e.g., conversational rules, facial expressions), and non-literal (e.g., sarcasm) [10, 12, 16, 18, 35]. Deficits in theory of mind (ToM) have been well-studied in the ASD population and refers to an individual's ability to infer the mental states of others [10, 16]. This ability has also been studied in relation to a larger group of cognitive abilities known as executive functions, a group of cognitive abilities used to accomplish goals (e.g., planning, inhibition, cognitive flexibility, working memory) [33]. Impairments in these areas often result in isolation and inability to form and maintain peer relationships among children and adolescents with ASD [17]. Kloosterman et al. investigated the association between these executive functioning impairments and peer victimization among adolescents with ASD, and found that poorer executive functioning predicted higher amounts of physical, social and verbal peer victimization [18]. Although executive functioning impacts the ability to determine the mental states and motivations of others, social interactions and reading the non-verbal cues of others may also determine an individual's ability to detect deception [31].

Dennis et al. the examined how the social interaction of children with high-functioning ASD impacted the detection of real and deceptive emotion [12]. Findings suggest that children with ASD were less able than normal children to correctly identify the emotions people feel, deceptive emotional expression in the face, and the motivation behind the deception [12]. This inability to detect deception and motivations behind it is often manifested in the form of victimization and bullying in adolescents with ASD [17, 31]. School-age children and adolescents with ASD report significantly higher rates of victimization and bullying compared to their typically developing counterparts [17].

In this study, we focus on understanding how users with autism react to the phishing websites presented to them. Internet deception and fraudulent websites, more commonly referred to as phishing scams, are prime examples of this and have become extremely problematic as internet access and use are so prevalent. Due to the vulnerability of individuals with ASD to deception in the context of peer relationships, these individuals may also be more vulnerable to an increasingly more common deceptive scheme – Internet deception. To the best of our knowledge, the current study is among the first to examine the susceptibility of individuals with ASD to fraudulent websites and phishing scams compared to typically developing individuals.

3 STUDY DESIGN

Phishing attack involves stealing a user's private credentials by presenting a fake website as a trustworthy entity. The attackers exploit the users' trust in the appearance of a website by designing a visually similar fake website. The victims are lead to these fake websites through the emails impersonated to come from the real websites/companies. These websites may have logo and website address (URL) similar to the real website. So the users should check for the domain name and look at the indicators of a valid secure connection and the details of the certificate to distinguish between real and fake websites. However, the users, even though they may know what to look for, may not always notice these parameters. To test these skills and to test the strength of the users with autism in identifying real and impersonated websites, we designed our phishing detection task.

Fully in line with the design of prior phishing detection studies [13, 26], we assumed that the users have already reached the fake website and need to identify its validity. To this end, we explicitly asked the participants to detect the fake websites from real websites, and focused on determining their performance, and on understanding their real-vs-fake decision-making.

We selected websites from the top 100 rank of Alexa [1]. e.g., Amazon, eBay, PayPal, Facebook, and Citibank. We then created fake website addresses by using URL obfuscation techniques, URL shortening techniques, and insertion of IP address. For URL obfuscation, we replaced domain names with similar looking words. For URL shortening, we shortened the URL with Google and Bitly link shortener, and for the IP address technique, we replaced URL with IP addresses. We then modified or removed the logo from some websites to make them look distinctly fake. We also borrowed some real-world phishing website addresses from *www.phishtank.com*. We, however, could not directly use the websites from *phishtank.com* in our study as the phishing websites relating to the brands we had selected for our study were already taken down. The hypothesis for the design was that *the users may notice IP address and modified logo in a website and identify them as the fake one, while they may not notice the changes in website address and may fail to detect them.*

The participants' task in this study was to distinguish between real and fake websites. The design of our phishing detection task is in line with the ones previously employed in [13, 26]. We developed an in-house software to execute the phishing detection task in the Firefox browser (the study was limited to Firefox given its popularity). The participants interacted with websites displayed in the browser very much like a real-world environment. In order to protect the privacy of the participants, while being subjected to real-world phishing sites, we pre-downloaded these sites for offline use and hosted them on our local web-server.

The fake websites (denoted "Fake"), which differ from the real websites (denote "Real") only in the URL are called "difficult fake (DFake)", assuming they might be difficult to detect. The other fake websites, which differ from real websites in more than one factor, such as layout, logo, fonts and URL, were referred to as "easy fake (EFake)", assuming these might be easier to detect. The Table 1 presents the sample of the URLs used in our study.

There were 60 randomized trials in this experiment: 30 each corresponding to real and fake websites. For fake websites, we had 15 each of easy fake and difficult fake websites. The experiment started with the Firefox browser loading the instructions page (specifying the tasks participants were to perform), which lasted for 60 seconds. This was followed by the trials pages, each displayed for 16s. Each trial consisted of a webpage (corresponding to a fake/real website) shown for 6s, followed by a 10s long response page. The response

Table 1: Examples of Fake Website URLs used in the study.

URL	Type	Brands
http://178.24.25.189-secure-bestbuy.com	EFake	Bestbuy
http://account.login-facebook.com/confirmation-account/index.html	EFake	Facebook
http://accounts.fspace.cc/login/index.html	EFake	Facebook
http://bit.ly/1UPcDZS/index.html	EFake	Pinterest
http://www.drobbox.net/db/box/index.htm	EFake	Dropbox
http://secure.update.chaseonline.securityupdate.pulaskiymca.org/	EFake	Chase
http://www.tim2via.com/~mutu/acc/index.htm	EFake	Google
http://www.yourinsta-128.48.52.81-secure.com/	EFake	Instagram
http://127.0.0.1//fNIR_web/Phishing_Easy/CitiBank/th4739/	EFake	Citibank
http://www.hawkeyess.in/Ama-2015/index.html	EFake	Amazon
http://www.bk-dentalarts.com/bbs/Amazon/Amazon/index.html	DFake	Amazon
http://www.watch-movies-in-netflix.com	DFake	Netflix
http://apple.supportteam5.co.uk/gb/index.htm	DFake	Apple
http://www.enigma3productions.com/k/dropbox/index.htm	DFake	Dropbox
http://www.societyofboa.com/bulk/bankofamerica.com/index.htm	DFake	BankOfAmerica
http://www.rosseta.com.vn/wp-admin/WELLSFARGO/WellsFargo/index.htm	DFake	WellsFargo
http://p4yp41.moldex.org/Paypal/provide-update/paypal.htm	DFake	Paypal
http://vissahome.clave.com/Formulario/index.htm	DFake	VISA
http://www.gmailingsalert.com	DFake	Google
http://www.dev.rodmkt.com.br/sir/outlook/index.htm	DFake	Bing
http://www.styeaerts.in/mail/login.yahoo.com/index.htm	DFake	Yahoo
http://www.walmart.comyr.com/	DFake	Walmart
http://p4yp41.moldex.org/Paypal/provide-update/paypal.htm	DFake	Paypal
http://www.ebaysignalcompte980654160132.hostfull.com/	DFake	Ebay

page had a dialog box with question: “Do you think the shown website is real?” and the “Yes” and No” buttons. A break/rest page of 2s (+ sign shown at the center of a blank page), after each trial was added, during which participants were asked to relax. Also, we extended the time of break to 6 seconds every 20 trials. The experiment then ended after 60 trials with the goodbye note, displayed for 5s. Figure 1 provides the flow diagram of the task we conducted.

4 STUDY PROTOCOL

In this section, we discuss the experimental protocol we implemented for our study.

4.1 Ethical and Safety Considerations

The study was approved by our University's IRB after a full review given the study focused on people with disorders. The participation in the study was strictly voluntary. The participants were given the option to withdraw from the study at any point of time. The standard best practices were followed to protect the confidentiality and privacy of participants' data (survey responses, task responses) acquired during the study. No names and other identifiable information were collected during the study.

4.2 Recruitment and Preparation Phase

The participants were recruited through a participant registry from our laboratory and via advertisements across campus and on online media. Interested participants contacted the research lab, and were screened for eligibility by the IRB trained, research assistants. To be eligible to participate, participants had to be 18 years or older and use the Internet regularly. Also, the participants with autism had an ASD diagnosis from a medical doctor or a clinical psychologist. Once it was determined that an individual was eligible, they were scheduled for an in-person study visit.

Fifteen participants with and without autism each were recruited for the study. Each participant took about a total of 1 hour to complete the study. During the preparation phase of the study, informed consent was obtained from each participant. The participants then provided their demographic information (such as age, gender, and education level). Table 2 summarizes the demographic information of our participants. Our pool for participants with autism was comprised of 66.6% male and 33.3% female, 20% were above the age of 30 and belonged to fairly diverse educational levels. Similarly, the pool for the control participants, the participants without autism, was comprised of 66.6% male and 33.3% female, 26% were above the age of 30. Our sample, especially in terms of age, was closer to the group of users who use the Internet frequently [7], and who are supposedly more vulnerable to phishing attacks [30], and hence are a good target of our study. Our participant sample is also well-aligned with the samples used in related prior studies [25, 26].

Also, in line with other prior studies [13, 25, 26], the participants were not told anything regarding the security relevance of the experiments, in order to avoid explicit security priming of the individuals which may impact their task performance.

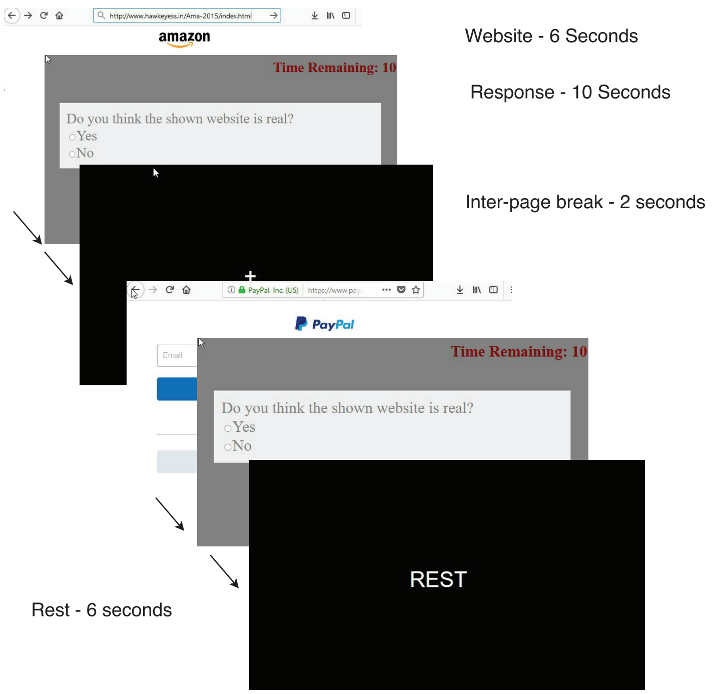

Do you think the shown website is real?
○Yes
○No

Do you think the shown website is real?
○Yes
○No

Figure 1: Sample Experimental Stimuli and Timing Flow Diagram. The stimuli (websites) were randomly presented to the participants. Each website was presented for 6 seconds, and 10 seconds were given to the participants to identify the legitimacy of the website.

Table 2: Participant Demographics

Details		Autism	Control
N		15	15
Gender	Male	66.6%	66.6%
	Female	33.3%	33.3%
Age	18-24	60.0%	33.3%
	25-30	20.0%	40.0%
	31+	20.0%	26.6%
Employment	Employed	33.3%	93.4%
	Unemployed	66.6%	6.6%
Background	Vocational	6.6%	0.0%
	Grammar School	6.6%	0.0%
	High School	53.3%	0.0%
	Bachelors	13.3%	46.6%
	Masters	0.0%	53.3%
	Others	20.0%	0.0%

4.3 Data Collection Phase

After filling out the questionnaire surveying demographic information, the participants completed the computer-usage-experience questionnaire to assess their familiarity with computers and their computer expertise, and the Autism-spectrum Quotient (AQ) to assess the ASD symptom severity [10]. AQ is a standard paper-and-pencil test comprising of 50 questions to measure the autistic characteristics in users. The higher the score in AQ, the higher the presence of symptoms consistent with autism.

The mean AQ score of participants with autism in our study was 29.2 with the standard deviation of 2.2, and the mean AQ score of the participants without autism was 17.06 with the standard deviation of 3.63. The AQ is not a diagnosis for Autism, but a score of 20 is considered as intermediate level of autism and a score above 32 is considered higher symptomatic level of autism [10]. It took about 15 minutes for each of the participants to answer ASQ. To recall, the participants we recruited for our study had ASD diagnosis from a medical doctor or a clinical psychologist. After completing these pen-and-paper questionnaires, participants were briefed on the instruction to perform the phishing detection task. The task required participants to identify a series of sample websites as real or fake based on appearance. This task was completed in 20 minutes.

4.4 Post-Test Phase

After performing the task, the participants were given a post-test questionnaire with questions designed to determine their familiarity with the websites used in the experiment. The participants were asked to answer if they were familiar with the websites presented in the study and if they had an account with any of them. The data collected from the post-test questionnaire was used to analyze the phishing detection accuracy for the websites on which the participants had accounts in. The hypothesis was that the participants will have higher accuracy of phishing detection for familiar websites. Also, in the post-test questionnaire, we had asked participants about their concerns of security and privacy on the web, web browsing behavior, knowledge of secure connection, computer skills, knowledge and experience with spam and phishing, and the strategies they used to identify the phishing websites. After the completion of this phase, each participant was paid a $20 gift card for their participation in the study.

5 QUANTITATIVE ANALYSIS

To recall, in our experimental task, participants were shown real and fake websites, and were asked to answer if the websites shown to them were real ("yes" response) or fake ("no" response). We had logged the participants responses and response times during the experiment. A participant's response was marked as correct if she had marked *yes for the real website*, and *no for the fake website*. Otherwise, the response was marked as incorrect. We then calculated the accuracy with which participants correctly identified real, easy fake and difficult fake websites. *Accuracy* is defined as the ratio of the total number of correctly identified websites to the total number of websites presented to each participant for each category of websites. We computed the average accuracy of correctly identified real or fake websites, and average response time, *first*, for all the websites, and *second* for only familiar websites.

We also performed statistical analysis to measure the statistical significance of the results we achieved. First, we used the Kolmogorov-Smirnov test to measure normalcy of the data. Our data set was non-normal so we used the Friedman's test and the Wilcoxon Singed-Rank Test (WSRT) for measuring statistical differences in the mean accuracy of real and fake websites within a group and Mann-Whitney U Test for measuring the statistical differences between the two groups (participants with autism and without autism) underlying our analysis. We conducted the power analysis of a Mann-Whitney U test to determine the power of our sample size (N=15) using an alpha of 0.05, a large effect size (d = .8), and two tails and observed the statistical power of 64%. This means that our study had 64% chance of detecting the difference in performance between participants with autism and without autism if a difference is there to be detected. Holm-bonferroni correction was used during post-hoc analysis for multiple comparisons. We also used Spearman's correlation coefficient (for non-normal distribution) and Pearson's correlation coefficient (for normal distribution) to measure the correlation between the autism scores and the performance accuracy.

5.1 Participants without Autism (Control)

In this section, we analyze the accuracy and response time we measured during our study for the participants without the autism spectrum disorder. This population serves as the control group for our study.

5.1.1 Performance Analysis: All Websites. In this analysis, we considered all the websites (familiar and unfamiliar websites) presented to the participants in the experiment. We then calculated the percentage accuracy of participants and the average response times for the different types (real, easy fake and difficult fake) of websites. Table 3 summarizes our results. The overall accuracy of correctly identifying a website is around 77%. Even though the average accuracy of easy fake websites looks higher than other the average accuracy of other websites in Table 3, we did not observe statistically significant difference in the mean accuracies and response time across real, fake, efake, and dfake trials on the Friedman's test.

Table 3: Control participants: Average (μ) and standard deviation (σ) of the accuracy and the response time across all websites

Trial		Accuracy μ (σ)	Response Time μ (σ)
Real		76.4 (17.3)	1.8 (0.4)
Fake	Overall	77.8 (18.5)	1.9 (0.5)
	EFake	81.3 (16.3)	2.1 (0.6)
	DFake	74.6 (25.4)	1.7 (0.6)
Overall		77.0 (13.9)	1.9 (0.3)

5.1.2 Performance Analysis: Familiar Websites. In this analysis, we calculated the accuracy and response times corresponding to the detection of the legitimacy of only the familiar websites. We considered both real and fake versions of the websites deemed familiar by the users. Table 4 summarizes our results. The overall accuracy of correctly identifying a familiar website is around 77%, which is better than random guessing (50%). We observe that the accuracy of identifying the easy fake website is higher than other websites from table 4, we did not find statistically significant difference in mean accuracies and response time across real, fake, easy fake and difficult fake websites on the Friedman's test.

Table 4: Control participants: Average (μ) and standard deviation (σ) of the accuracy and the response time across familiar websites

Trial		Accuracy μ (σ)	Response Time μ (σ)
Real		76.2 (20.4)	1.9 (0.5)
Fake	Overall	78.2 (19.2)	1.9 (0.5)
	EFake	83.4 (15.1)	2.1 (0.5)
	DFake	72.8 (29.1)	1.8 (0.7)
Overall		77.2 (14.1)	1.9 (0.3)

5.2 Participants with Autism

In this section, we analyze the accuracy and response time we measured during our study for the participants with the autism. The group of participants with autism serves as a target class for our study.

5.2.1 Performance Analysis: All Websites. For this analysis, similar to the analysis of the control participants, we consider all the websites (familiar and unfamiliar websites) presented in the experiment. We then calculated the percentage accuracy and the response times related to different types of websites. Table 5 summarizes our results.

Table 5: Participants with Autism: Average (μ) and standard deviation (σ) of the accuracy and the response time across all websites

Trial		Accuracy μ (σ)	Response Time μ (σ)
Real		80.6 (17.4)	2.0 (0.5)
Fake	Overall	68.9 (25.1)	1.7 (0.5)
	EFake	72.6 (21.2)	1.8 (0.5)
	DFake	64.9 (30.8)	1.6 (0.6)
Overall		74.8 (19.0)	1.8 (0.4)

From Table 5, we see that the overall accuracy of phishing detection is around 74%, which is better than random guessing (50%). Also, we observe that the accuracy of identifying real websites is more than the accuracy of identifying fake websites. However, we did not find statistically significant differences among the mean accuracies of real, fake, easy fake and difficult fake websites on using Friedman's test.

5.2.2 Performance Analysis: Familiar Websites. To enable website familiarity analysis, we had asked participants, in the post-experiment phase to answer if they had an account with the websites presented in the study. In this analysis, we calculated the accuracy and response times corresponding to the detection of the legitimacy of only familiar websites. We considered both real and fake versions of the websites deemed familiar by the users. Table 6 summarizes our results. As seen in Table 6, the overall accuracy of phishing detection for familiar websites was around 75%, which is only better than a random guess. We assumed that the users might be better at detecting the phishing attacks against the websites they were familiar with. However, our results do not show much difference.

We did not observe statistically significant difference in mean accuracies across Real, Fake, Easy Fake and Difficult Fake websites on the Friedman's test. However, the Friedman's test showed a statistically significant difference in mean response time across ($\chi^2(3) = 13.8, p < .005$). On further contrasting the response times across different types of websites with WSRT, we found that the participants spent statistically significantly higher time on real websites than the fake websites ($p<.005$).

Table 6: Participants with Autism: Average (μ) and standard deviation (σ) of the accuracy and the response time across familiar websites

Trial		Accuracy μ (σ)	Response Time μ (σ)
Real		86.6 (16.0)	2.3 (0.8)
Fake	Overall	69.1 (29.9)	1.6 (0.7)
	EFake	69.6 (29.0)	1.6 (0.7)
	DFake	63.4 (39.7)	1.6 (1.0)
Overall		75.2 (22.9)	1.9 (0.7)

5.3 Performance Comparison: Participants with Autism vs. Control Participants

To analyze the differences in the performance of the participants with and without autism in phishing detection task, we performed Mann-Whitney U Test. Mann-Whitney U Test is a non-parametric test used to compare if two sample means from two different groups of the same population are equal or not. In our analysis, on Mann-Whitney U Test, we did not observe statistically significant difference in the mean accuracy of identifying phishing websites between the participants with autism and without autism. This shows that the participants with autism may not be more susceptible to phishing attacks.

On further inspection, we found that the previous studies [4, 8] have reported that the users with autism are methodical. Also, Carper et al. [11], in their neuroscience study, reported that the symmetry in the brains of users with autism might make them more detail oriented. Also, cybersecurity practitioners have reported that the people with autism have strengths including cognitive pattern recognition, attention to detail, logical and methodical thinking, focus and integrity which might help corporate and/or intelligence sector [3, 6]. In light of these aspects, the participants with autism might have been careful, and might have noticed the differences in URL in making their decisions, and might have identified many of the websites correctly. Indeed, in the post-test questionnaire (see Section 6), the participants with autism reported that they noticed differences in URL, and differences in logo on the websites presented to them. Nevertheless, our study concludes that the users with autism might not be more vulnerable to phishing attacks compared to typical users without autism.

5.4 Correlation: Autism Quotient vs. Accuracy and Response Time

We investigated how the autism scores drawn from AQ score correlates to the accuracy and the response time of detecting the real, fake, easy fake, and difficult fake websites. We used the Spearman's correlation coefficient to calculate the correlation of these metrics in both categories of websites – all websites, and only the familiar websites.. And we did not observe any statistically significant correlation between the AQ score, and the accuracy and the response time of the websites presented to the participants. This means that the ability of detecting phishing website may be independent of the autism condition of the participants.

6 QUALITATIVE ANALYSIS

To recall, after completing the phishing detection experimental task, the colorblue both the participants with and without autism were asked to fill out a post-test questionnaire. This questionnaire was designed to determine participants' knowledge of computer security and privacy. We asked the participants about phishing, spam, and the browser indicators. We had also asked them regarding their strategy implemented to identify the phishing websites presented to them during the phishing detection. Our post-test survey was deliberately presented at the end of the actual experiment to prevent any explicit priming due to our asking security-pertinent questions.

We summarize the responses of the participants to the questions they were asked in the post-test questionnaires below:

- *Concern of Security and Privacy on the Web*: Since our study was a security study, we started our post-test questionnaire with basic security questions. One of these questions polled for participants' concern regarding the security and the privacy on the Internet. It might be correlated with the amount of attention the users pay while identifying malicious web activities and attack. 100% of our participants (both with and without autism) reported that they were concerned about their security and privacy on the web. Also, 100% of the participants (both with and without autism) mentioned that they were concerned about the leakage of the passwords they use on the Internet websites.

- *Web Browsing Behavior*: To ensure that participants relate themselves with the brands of the websites presented in our study, we asked them how often they browse the Internet and for what purposes. 100% of the participants (both with and without autism) reported that they browse Internet daily for making online purchases, online bill payments, reading news articles, online banking and social networking. Most of the phishing websites are targeted to the brands involved in these domains. We also asked the participant familiarity to these websites, and found that on average, the participants with autism were familiar to 90% of these websites, and had account in 28% of them. However, the participants without autism were familiar with 90% of the websites, and had account in 47% of them.

- *Knowledge of Secure connection and certificates:* To measure the participants' knowledge of secure connection, we had asked them if they had heard about secure shell and/or configured firewall. Only 6.6% of the users (2 out of 30 participants) knew how SSH and firewall were configured. Interestingly both these participants had autism, and these users also turned out to be the best performers of the phishing detection task with the accuracy above 95%. None of the participants without autism in our study had heard about secure shell and/or configured the firewall.

- *Computer Skills*: We wanted to compare the computer skills of the users with their phishing detection performance. Computer skills relate to programming and designing of websites, among other things. The average self-rating of the computer skill of the participants with autism was about 7 on a scale of 10, and the participants without autism was 6 out of 10. We did not find any statistically significant correlation between computer skill, and the users phishing detection performance, however. Also, to know the extent of their technical capacity, we had asked the participants if they had ever registered a domain or ever designed a website. Only 15% of the participants (all belonging to the group with autism) had registered a domain, and had designed the website. We did not have sufficient data to perform statistical analysis.

- *Knowledge and Experience with Spam and Phishing*: The participants had to answer if they had ever seen a spam email and explain why they are sent. 20% of the participants with autism and 6% of the participants without autism reported that they had never seen any spam email and do not know why they are sent. The rest of the participants (both with and without autism) mostly explained spam as emails sent to try to get users' information. Some of them mentioned: "spam emails are sent to trick other people in believing in them"; "shotgunning – if you send out enough, you'll eventually get a hit (a person will click on it)"; "people want to trick other people". One of the participants believed that spams represent the bugs on the Internet. Upon using pearson's correlation, we observed strong statistically significant correlation between their knowledge of spam and their phishing detection performance in our study ($r = .63, p < .05$).
  We also asked the participants if they knew what phishing was and why phishing emails are used. 60% of the participants with autism, and 94% of the participants without autism reported that they had heard of phishing and properly defined phishing attacks. Some reported phishing as: "an email that will send you a link to go to a fake website to try to steal your information", while some defined phishing as: "a fake website that looks like a real one to get a person's information and steal/sell that info". Some even defined phishing as an impostor: "where someone poses as an established company in order to gain access to credentials under false pretenses". We performed pearson's correlation test to understand relationship between phishing detection accuracy and the phishing knowledge among all of the participants, and observed strong statistically significant correlation between them ($r = .65, p < .05$).

- *Strategies of Phishing Detection:* We had asked the participants if they had noticed any differences in the websites as they were repeatedly presented to them in the experiment. This would help us know if the participants actually noticed the differences in the websites presented to them or were randomly marking them as real and fake. 80% of the participants with autism reported that they detected some differences on the websites repeatedly shown to them. They said: "some did not have the same logo',"URLs and logos were different","sometimes they have different certificates, URLs". 90% of the participants without autism also reported that they had detected some differences. Some of them said: 'the logo on the tab of the browser, and the logo included in the website were different','some websites were secure, and some were not secure'.
  We had also asked the participants about the strategy they adopted to identify the real website. Among the participants with autism, 90% of them reported that they had looked at the website logo, look and feel of the website, and the address of the website, and 60% of them reported that they looked at the security lock/certificate. The users with autism were careful enough to notice the obfuscated URLs, modified logs, and provide response accordingly. On the other hand, among the participants without autism, 60% of them said that they looked at the website logo, look and feel of the website, security certificate, and the address of the website, 13% of them reported they used other techniques (e.g., looked at the links by hovering in it, the logo on the tab and the browser) to

identify the fake websites, and 26% reported they only looked at the URLs, and/or security lock/certificate. Overall, the responses provided by both sets of participants reflect that they were carefully examining the websites to identify their legitimacy even though they were only partially successful in it as demonstrated by our quantitative analysis (see Section 5).

- *Other Interesting Comments*: On asking the participants with autism, if having autism was an obstacle to learn computer skills, many of them said it was not a problem at all. Also, they reported that they spend most of the time playing video games on the computer, and are interested in web security and privacy and also mentioned that some of the young individuals with autism are pretty good with breaking firewall and getting through a website.

In summary, our post-test questionnaires reflect that the participants of our study, were pretty careful in identifying the websites presented to them. The participants with autism, similar to the participants without autism, had used their knowledge properly and had identified the phishing websites correctly with the accuracy better than random guessing. The responses provided by the participants with autism demonstrates that they might not be more vulnerable than the participants without autism in phishing detection task.

7 DISCUSSION AND FUTURE WORK

In this section, we summarize and further discuss the main findings of our study. We also outline the strengths and limitations of our study.

Users with autism may be equally capable to healthy individuals in detecting phishing websites: In this study, we analyzed both the quantitative and qualitative data collected for the phishing detection task with the with and without autism spectrum disorder. The prior psychology studies have reported that the users with autism are susceptible to deception and find it difficult to make decisions and discern the intentions of other people [16, 35]. In this light, we had hypothesized that they may be more susceptible to phishing attacks compared to the participants without autism. However, our quantitative results demonstrates that the users with autism are no less than healthy participants in identifying the phishing websites. Even though the education level of our participants with autism on average was lower than the participants without autism, there was no statistically significant difference in the accuracy. Their good performance in the phishing detection task was explained by their answers to the qualitative questionnaires. Many participants reported that they had noticed missing security locks/certificates, obfuscated URLs, and distorted logos in some of the websites presented to them. The previous studies on the users with autism also report them as the detail oriented and methodological people with the capability to notice minute changes [2, 4]. Our study suggests that the users with autism are capable of noticing the minute differences, analyzing them, and making correct decisions, which may enable them to detect phishing attacks.

Users with autism may make for good cybersecurity professionals: The previous neuroscience studies have reported that the brain of users with autism is inherently capable to pick up tiny details [4]. A report prepared by The National Autistic Society [2] states the users with autism have a methodological approach, good memory for factual information, strong problem solving skills, specialist knowledge and skills, and pattern recognition skills. We believe these skills may have made them successful in identifying the real vs. fake websites presented to them during the study. These skills might also be helpful in other security tasks, e.g., anomaly detection, malware detection and unusual behavior of the computer systems. The detail–oriented nature of the users with autism may help identify attacks, loopholes in defense mechanisms, and bugs in the secure codes. Given this, it is not surprising that there are several openings in cybersecurity calling for users with autism to apply [8]. However, more focused studies are needed to further understand the capability of users with autism in cybersecurity. Nevertheless, our study highlights the positive aspects of users with autism and show that such users may make them skilled at cybersecurity, at least in the field of detection of phishing scams.

Potential future directions: Recently, researchers have been conducting studies to understand the usability of security tools and parameters to specific set of population e.g., journalists and lawyers [19, 24]. The universally acceptable usable privacy and security is challenging, and hence the studies to analyze the performance of users in different groups of population is important. The future study could be conducted to analyze how usable CAPTCHAs, phishing toolbars, and verification images are with respect to such specialized groups of users, including those suffering from social and mental disorders. Also, the password policies are being stronger lately with the users required to mix up their password with different symbols, alphabets and digits. Such passwords may not be usable for aged users as it might be difficult for them to remember the complex stuffs. There is need for password security research focused on special samples of users.

8 STUDY STRENGTHS AND LIMITATIONS

Our study has several strengths. *First*, we set-up our study with real working websites, unlike images used in previous studies [13, 26]. We also followed standard URL obfuscation methodology, URL shortening methodology, and use of IP address to design our fake URLs. To make them the URLs realistic, we borrowed real-world fake URLs from *www.phishtank.org*. We recruited participants with and without autism and conducted a 'between-group' study to analyze the performance of users with autism.

Similar to prior lab-based studies, we also had some limitations. We showed multiple trials (60 websites) in the span of 25 minutes to collect data. The users may not see that many websites within such a short time frame in real-world. Also we did not consider phishing websites with domain validation certificates. The domain validation certificates would have made the phishing websites harder to identify and might have reduced the accuracy of phishing website detection across both sets of participants. Our hypothesis was that the users with autism might be easier to deceive with respect to phishing attacks, so we had not introduced the difficultly in our study. Second, the participants were explicitly and directly asked to identify real and fake websites. In real-world, they are directed to these pages from emails or other methods. However, the users have to eventually make the decision of real and fake implicitly and they may perform worse in such implicit scenarios than they did in our explicit scenario. Although we tried to imitate the real-world scenario of web-browsing and computer usage in our study, the participants' performance might have been affected just due to the fact that the study was conducted in lab. Third, we had a small participant sample size of fifteen in each group. Our sample may not be representative of a wider population, although it is representative of active computer users [7] and phishing-prone populations [30]. Also, it is difficult and challenging to recruit such specific set of participants, and other studies on autism [28], learning disabilities [21] and journalists [19] and lawyers [19] also had similar sample sizes of participants. Fourth, the participants without autism were relatively better educated than the participants with autism. Even then the participants with autism performed equally to the participants without autism in identifying phishing websites. Hence the participants should at least perform in par with the participants without autism and with only high school education. The control participants without autism, above the age of 19 and with only high school education or lower was difficult to find.

9 CONCLUDING REMARKS

In this paper, we presented a phishing detection study to understand the performance of users with autism in identifying real and fake websites. The users with autism have been found to be gullible in previous studies. However, they are also reported to be methodological and detail-oriented. Our results show that the gullibility of users with autism may not make them more susceptible to phishing attacks, and, in fact, their detail oriented nature may make them better equipped to combat phishing attacks, when contrasted with the individuals without autism.

10 ACKNOWLEDGMENT

The authors thank: Cierra Hopkins and Karan Patel for their help with data collection, Mariah Stasiak Eubank for her help with recruitment and scheduling of the participants, and Rajesh Kana for his feedback in the early stages of this work. We also thank Gang Wang (our shepherd) and ACSAC'18 anonymous reviewers for their constructive comments and guidance.

REFERENCES

[1] Alexa. http://www.alexa.com. [5-15-2018].

[2] Autism and careers in cyber security: A short guide for employers. https://www.iaac.org.uk/wp-content/uploads/2017/07/Autism-and-careers-in-cyber-security-FINAL.pdf. [1-30-2018].

[3] Autistic recruiter. https://www.forbes.com/sites/kateoflahertyuk/2018/09/10/how-employing-autistic-people-can-help-stop-cyber-attacks/#e2e22b54b503. [9-11-2018].

[4] The brain difference that allows people with autism to pick up tiny detailse. https://www.huffingtonpost.com/entry/autism-brain-symmetry_us_5841d0b1e4b0c68e0480b469. [1-30-2018].

[5] Centers for disease control and prevention [cdc]. (2015, 24 february 2015), autism spectrum disroder (asd): Facts about asd. http://www.cdc.gov/ncbddd/autism/data.html. [1-30-2018].

[6] How can someone with autism specifically enhance the cyber security workforce? https://www.blackhat.com/us-18/briefings.html#how-can-someone-with-autism-specifically-enhance-the-cyber-security-workforce. [9-08-2018].

[7] Internet Users Demographics. http://www.pewinternet.org/fact-sheet/internet-broadband/. [9-11-2018].

[8] Nice webinar: Cybersecurity careers for autistic people. https://www.nist.gov/news-events/events/2017/10/nice-webinar-cybersecurity-careers-autistic-people. [1-30-2018].

[9] Association, A. P. *Diagnostic and Statistical Manual of Mental Disorders: Fifth Edition (DSM-V)*. American Psychiatric Association, 2014.

[10] Baron-Cohen, S., Wheelwright, S., Skinner, R., Martin, J., and Clubley, E. The autism-spectrum quotient (aq): Evidence from asperger syndrome/high-functioning autism, malesand females, scientists and mathematicians. *Journal of autism and developmental disorders 31*, 1 (2001).

[11] Carper, R. A., Treiber, J. M., DeJesus, S. Y., and Müller, R.-A. Reduced hemispheric asymmetry of white matter microstructure in autism spectrum disorder. *Journal of the American Academy of Child & Adolescent Psychiatry 55*, 12 (2016).

[12] Dennis, M., Lockyer, L., and Lazenby, A. L. How high-functioning children with autism understand real and deceptive emotion. *Autism 4*, 4 (2000).

[13] Dhamija, R., Tygar, J. D., and Hearst, M. Why phishing works. In *Proceedings of the SIGCHI conference on Human Factors in computing systems* (2006), ACM, pp. 581–590.

[14] Egelman, S., Cranor, L. F., and Hong, J. You've been warned: an empirical study of the effectiveness of web browser phishing warnings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2008), ACM, pp. 1065–1074.

[15] Friedman, B., Hurley, D., Howe, D. C., Felten, E., and Nissenbaum, H. Users' conceptions of web security: a comparative study. In *CHI'02 extended abstracts on Human factors in computing systems* (2002), ACM, pp. 746–747.

[16] Kleinman, J., Marciano, P. L., and Ault, R. L. Advanced theory of mind in high-functioning adults with autism. *Journal of autism and developmental disorders 31*, 1 (2001), 29–36.

[17] Kloosterman, P. H., Kelley, E. A., Craig, W. M., Parker, J. D., and Javier, C. Types and experiences of bullying in adolescents with an autism spectrum disorder. *Research in Autism Spectrum Disorders 7*, 7 (2013), 824–832.

[18] Kloosterman, P. H., Kelley, E. A., Parker, J. D., and Craig, W. M. Executive functioning as a predictor of peer victimization in adolescents with and without an autism spectrum disorder. *Research in autism spectrum disorders 8*, 3 (2014), 244–254.

[19] Lerner, A., Zeng, E., and Roesner, F. Confidante: Usable encrypted email: A case study with lawyers and journalists. In *Proceedings of the IEEE European Symposium on Security and Privacy* (2017), IEEE, pp. 385–400.

[20] Mare, S., Vashistha, A., and Anderson, R. Security and privacy design considerations for low-literate users in developing regions. In *Proceedings of the Thirteenth Symposium on Usable Privacy and Security* (2017), USENIX Association.

[21] Marne, S. T., Nasrullah, M., and Wright, M. Learning system-assigned passwords: A preliminary study on the people with learning disabilities. In *Proceedings of the Thirteenth Symposium on Usable Privacy and Security* (2017), USENIX Association.

[22] McGregor, S. E., Charters, P., Holliday, T., and Roesner, F. Investigating the computer security practices and needs of journalists. In *Proceedings of the USENIX Security Symposium* (2015), USENIX.

[23] McGregor, S. E., Roesner, F., and Caine, K. Individual versus organizational computer security and privacy concerns in journalism. *Proceedings on Privacy Enhancing Technologies 2016*, 4 (2016), 418–435.

[24] McGregor, S. E., and Watkins, E. A. When the weakest link is strong: Secure collaboration in the case of the panama papers. In *Proceedings of the USENIX Security Symposium* (2017), USENIX Association.

[25] Neupane, A., Rahman, M. L., Saxena, N., and Hirshfield, L. A Multimodal Neuro-Physiological Study of Phishing and Malware Warnings. In *ACM Conference on Computer and Communications Security* (2015), ACM.

[26] Neupane, A., Saxena, N., Kuruvilla, K., Georgescu, M., and Kana, R. Neural signatures of user-centered security: An fMRI study of phishing, and malware warnings. In *Proceedings of the Network and Distributed System Security Symposium* (2014), NDSS.

[27] Oliveira, D., Rocha, H., Yang, H., Ellis, D., Dommaraju, S., Muradoglu, M., Weir, D., Soliman, A., Lin, T., and Ebner, N. Dissecting spear phishing emails for older vs young adults: On the interplay of weapons of influence and life domains in predicting susceptibility to phishing. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (2017), ACM.

[28] Ranick, J., Persicke, A., Tarbox, J., and Kornack, J. A. Teaching children with autism to detect and respond to deceptive statements. *Research in Autism Spectrum Disorders 7*, 4 (2013), 503–508.

[29] Schechter, S. E., Dhamija, R., Ozment, A., and Fischer, I. The emperor's new security indicators. In *Proceedings of the IEEE Symposium on Security and Privacy* (2007), IEEE.

[30] Sheng, S., Holbrook, M., Kumaraguru, P., Cranor, L. F., and Downs, J. Who falls for phish?: a demographic analysis of phishing susceptibility and effectiveness of interventions. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (2010), ACM.

[31] Sofronoff, K., Dark, E., and Stone, V. Social vulnerability and bullying in children with asperger syndrome. *Autism 15*, 3 (2011), 355–372.

[32] Sunshine, J., Egelman, S., Almuhimedi, H., Atri, N., and Cranor, L. F. Crying wolf: An empirical study of ssl warning effectiveness. In *Proceedings of the USENIX Security Symposium* (2009), USENIX.

[33] Welsh, M. C., and Pennington, B. F. Assessing frontal lobe functioning in children: Views from developmental psychology. *Developmental neuropsychology 4*, 3 (1988), 199–230.

[34] Wu, M., Miller, R. C., and Garfinkel, S. L. Do security toolbars actually prevent phishing attacks? In *Proceedings of the Conference on Human Factors in computing systems* (2006), ACM.

[35] Yang, Y., Tian, Y., Fang, J., Lu, H., Wei, K., and Yi, L. Trust and deception in children with autism spectrum disorders: A social learning perspective. *Journal of autism and developmental disorders 47*, 3 (2017), 615–625.

[36] Yi, L., Pan, J., Fan, Y., Zou, X., Wang, X., and Lee, K. Children with autism spectrum disorder are more trusting than typically developing children. *Journal of experimental child psychology 116*, 3 (2013), 755–761.