British Journal of Educational Studies

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/rbje20

WHO RESPONDS TO PHISHING EMAILS? AN INTERNATIONAL INVESTIGATION OF 15-YEAR-OLDS USING PISA DATA

John Jerrim

To cite this article: John Jerrim (01 Sep 2023): WHO RESPONDS TO PHISHING EMAILS? AN INTERNATIONAL INVESTIGATION OF 15-YEAR-OLDS USING PISA DATA, British Journal of Educational Studies, DOI: 10.1080/00071005.2023.2234456

To link to this article: https://doi.org/10.1080/00071005.2023.2234456

British Journal of Educational Studies
Vol. 00, No. 00, XXXX 20232023, pp. 1–24

WHO RESPONDS TO PHISHING EMAILS? AN INTERNATIONAL INVESTIGATION OF 15-YEAR-OLDS USING PISA DATA

by JOHN JERRIM, *UCL Social Research Institute, UCL, London, UK*

ABSTRACT: Young people are facing an ever-increasing array of online dangers. One of the most common is receipt of a phishing email. This paper presents new evidence on the characteristics of young people most likely to respond to such emails. I find approximately one-in-seven 15-year-olds are at risk of responding to a phishing email, rising to one-in-five amongst those from disadvantaged socio-economic backgrounds. Such risks are particularly high amongst young people with low levels of cognitive skill. Unfortunately, students who are taught about the dangers posed by phishing emails at school are just as likely to take inappropriate actions following their receipt as their peers who have not. I thus conclude that greater emphasis and higher quality instruction needs to be provided to young people about the online risks they face, particularly to those from disadvantaged socio-economic backgrounds and low academic achievers.

Keywords: PISA, socio-economic inequality, phishing, cyber-fraud

1. INTRODUCTION

There has long been interest in inequalities across a range of social, educational and labour market outcomes. A now extensive literature has illustrated how individuals from less advantaged social backgrounds and with lower levels of educational achievement do not accomplish the same outcomes as their more advantaged peers. Although such inequalities have their roots in the early years (Cattan *et al*., 2022), adolescence is widely regarded as another critical period in young people's development (Viner *et al*., 2015). It is a time when young people are faced with many challenges, and when their propensity to engage in risky behaviours increase (Kipping *et al*., 2015). Ensuring those from disadvantaged backgrounds successfully navigate this part of early adulthood – equipping themselves with the knowledge and skills they need for the future – is vital to ensuring they flourish.

However, with the increasing digitisation of modern society, young people perhaps now encounter a greater array of challenges – and risks – than ever. They – like all of us – will be subject to attempted cyber-fraud. Although there are now sophisticated ways to con individuals through 'deep fakes', some attempts at cyber-fraud continue to be quite basic. Yet some groups may still

ISSN 0007-1005 (print)/ISSN 1467-8527 (online)

https://doi.org/10.1080/00071005.2023.2234456
http://www.tandfonline.com

be at risk from even quite rudimentary scams, thus potentially suffering the serious consequences cyber-crime can bring.

Surprisingly, there has been relatively few studies into inequalities in young people's susceptibility to cyber-fraud, with the literature tending to focus on older adults. Yet this is a critical issue facing young people as well. If those from less advantaged socio-economic backgrounds or with weaker cognitive skills are more likely to be fooled by online scams, then this may hold back their prospects for the future. Indeed, it would suggest that those individuals already facing a precarious situation may also be at greatest risk of the negative consequences associated with being defrauded. However, there may also be a clear route – providing lessons about cyber-fraud as part of the school curricula – that could help such inequalities to be reduced.

The goal of this paper is to present new cross-national evidence on this matter. It does so by focusing on 15-year-olds susceptibility to one specific type of cyber fraud – responding to a phishing email. Differences in young people's propensity to respond to such phishing solicitation is presented across countries, socio-economic groups and between young people with different cognitive skills. I also consider the extent to which 15-year-olds who have received instruction from their schools about the risks associated with phishing are less likely to respond to phishing solicitation, and thus whether such instruction might (in its current form) be helping to address inequalities in this area.

2. Literature and Research Questions

A host of studies have explored the susceptibility of individuals to various forms of online fraud, including exploring differences across demographic groups. These have, however, mostly focused on adults rather than school-aged children. I review a selection of these studies below, focusing on those related to the characteristics explored in this paper (cognitive skills, socio-economic background and differences across countries).

Education Level and Cognitive Skills

Zhang and Ye (2022) presented empirical evidence that highlighted how less educated, more impulsive and more trusting individuals were more likely to become victims of online frauds in China. In two field experiments, Wood *et al*. (2018) found that adults with lower levels of education were at increased risk of falling for mass marketing scams. Ebner *et al*. (2020) found that 75–89-year-olds with lower levels of cognitive skill were more susceptible to phishing emails than those with higher levels of cognitive skill. However, the same relationship (between cognitive skills and susceptibility to phishing) did not emerge for younger adults (18–37-year-olds). The work of Gavett *et al*. (2017, p. 10) supports this view, with their empirical study finding that '*people with*

more years of education tended to be more suspicious of the phishing attempts'. In contrast, a pan-European study found that adults with lower levels of education were less likely to report being the subject of fraud than more educated individuals (European Commission, 2020). In a survey of over 10,000 adults, Whitty (2019) found that individuals with higher levels of education were more likely to be victims of cyber-fraud. In their study of susceptibility of investment fraud, Mueller *et al*. (2020, p. 169) report that '*education was not a significant predictor of scam susceptibility*'. Lee and Geistfeld (1999) echo this finding, failing to find support for their hypothesis that less educated consumers would be more receptive to telemarketing scams. This is consistent with a study of older American's by DeLiema *et al*. (2020), who failed to find education, cognitive ability or financial literacy to be associated with the risk of financial fraud. Evidently, then, previous research into how the link between education and fraud – including cyber-fraud – is somewhat mixed.

Socioeconomic Status

Several studies have investigated differences in online fraud by socio-economic status. Based upon survey data conducted in England and Wales, the Office for National Statistics (2022) estimates that three percent of the population had responded to or clicked on a link within a phishing message, with 11% of those that did then providing information that could be used by the recipients. The percentage of respondents who clicked links in phishing messages or responded to them was markedly higher amongst those living in disadvantaged areas (5% in the most deprived areas versus 2% in the least deprived areas) and those living in social housing (7%). A pan-European study found that individuals who were struggling financially were less likely to report being the subject of fraud than individuals with higher levels of income (European Commission, 2020). DeLiema *et al*. (2020) failed to find an association between wealth and risk of suffering financial fraud amongst those aged 50 and above. Anderson (2019) reached a similar conclusion from a survey conducted across the United States, failing to find a relationship between income and being the victim of various forms of fraud. In a study of older adults (average age 81), Glover *et al*. (2023) found childhood socio-economic status to interact with cognitive function to predict scam susceptibility.

Variation Across Countries

While several studies have explored differences across countries in the extent of – or susceptibility to – cyber-crime, evidence based upon large, nationally representative samples remains limited. For instance, Alseadoon and Othman (2021) explore cultural differences in susceptibility to phishing emails, reaching the conclusion that '*users' vulnerability to phishing emails is different across cultures*' (Alseadoon and Othman, 2021, p. 240). Their analysis was however based upon

a convenience sample of just 213 participants. The European Commission (2020) conducted a survey of various types of fraud across Europe, with approximately 1,000 respondents per country. They found Western Europeans to be more likely to report having been the subject of a scam than Eastern Europeans, particularly those in Ireland, the United Kingdom and Denmark. The OECD (2021) argue that there are large socio-economic differences in susceptibility to digital fraud (in the form of phishing emails) across industrialised countries. Cook *et al*. (2023) explores fear of economic cybercrime Europe, with sample sizes of around 1,000 adults per country. They demonstrate how there are substantial cross-national differences in fear of cybercrime, being lowest in Sweden, Estonia and the Netherlands, and highest in Ireland, Lithuania, Latvia, Romania and the Czech Republic. Chen *et al*. (2023) use IP addresses from a blocklist to explore prevalence of cybercrime across the world. High-income regions were found to host most cybercrime IPs and lower-middle-income regions the least. Reep-van den Bergh and Junger (2018) explore victims of various forms of cybercrime across Europe, drawing upon a range of country specific surveys. They warn, however, that '*prevalence estimates between countries are incomparable due to, most of all, question wording*' (Reep-van den Bergh and Junger, 2018, p. 12). Together, the above illustrates how cross-national comparisons of cyber-fraud – and the risks posed by phishing emails in particular – remain somewhat limited, not least due to the challenging data requirements (large, random samples drawn across multiple countries with the questions worded in the same way).

Research Questions

The aforementioned studies have provided important insights into the characteristics of those most likely to be tricked into taking unwise actions by online scams. Yet some clear gaps in the evidence base remain. Many studies draw upon small convenience samples within a single national setting. There has been limited consideration of inequalities in responding to phishing emails amongst young people – including across socio-economic groups – and the extent that this may be related to differences in their cognitive skills. Likewise, little is currently known about the efficacy of school's attempts to teach young people about how to recognise and react appropriately to the online dangers they face (such as phishing emails).

I thus attempt to fill these gaps in the literature by addressing three research questions.

To begin, cross-national evidence is presented comparing the percent of 15-year-olds who are at risk of responding to phishing solicitation. This thus seeks to explore the generalisability of results across national settings, and whether young people in certain countries are more likely to respond to phishing solicitation than others. Results are presented with and without controlling for cross-country differences in demographic characteristics and cognitive skills to

explore the extent that this can 'explain' the results. This contributes to the existing literature by providing new evidence on the risks associated with phishing across countries amongst young people, and the extent that any such cross-national differences can be explained by differences in demographic characteristics and cognitive skills. My first research question is thus:

Research question 1. In which countries are teenagers most likely to respond to a phishing email? To what extent can differences across countries be explained by differences in cognitive skills?

I then turn to exploring inequalities in such risks by common background characteristics, including socio-economic status. As noted by Hanoch and Wood (2021) the evidence on the link between demographic variables and susceptibility to phishing emails (and online fraud more generally) is 'patchy' at best. This is consistent with Whitty (2019, p. 279), who states that '*with respect to socio-demographic characteristics . . . .in general, the literature is fairly sketchy, with much of the research focussing on the elderly*'. Hanoch and Wood (2021, p. 262) have consequently called for '*further examination of demographic variables*' with the suggestion that '[future] *studies should include country-specific and cross-national comparisons*'. I thus contribute this to the literature by establishing (a) whether socio-economic differences in the likelihood of responding to phishing solicitation are of similar magnitude amongst young people across the industrialised world and (b) the extent that such variation across groups are related to differences in cognitive skill.

Research question 2. How is the likelihood of responding to phishing emails related to the demographic characteristics of young people, including socio-economic status? To what extent are any such differences driven by differences in cognitive abilities?

Finally, the education system offers one potential route to reducing young people's susceptibility to online fraud, including amongst those with lower levels of cognitive skills and socio-economically disadvantaged groups. Indeed, many schools now recognise the risk posed to young people from online fraud. Lessons are hence sometimes provided – usually as part of personal and social education classes – about the dangers of phishing and other malicious communications, and how young people should respond. But do such lessons – as currently provided – work? Are those teenagers who are warned about online dangers and how to handle them at less risk of being duped by phishing emails or not? This is clearly an important issue as, if current provision is not working,

a strong case can be made for greater investment – and more curriculum time – to be devoted to this area. Yet little is currently known about the efficacy of the provision currently provided by schools. My final research question is therefore:

Research question 3. To what extent can teaching students in school make them less susceptible to responding to phishing emails?

3. Data and Methodology

The data used are drawn from the 2018 round of the OECD PISA study. This is an international assessment of 15-year-olds achievement in reading, science and mathematics conducted every three years. My analysis includes all members of the OECD. Within each country at least 150 schools are selected with probability proportional to size, with approximately 35–40 15-year-olds randomly selected from within. Response rates for both schools and pupils are generally quite high (OECD averages of 92% and 90%) though with some variation across countries (see OECD, 2020 for further details). The PISA dataset is supplied with a set of pupil and Balanced Repeated Replication (BRR) weights that – when applied – fully account for the complex sample design (including the nesting of pupils within schools). These are converted into so-called 'senate weights' (Jerrim *et al*., 2017) so that each country carries equal weight in the analysis.

After completing a two-hour cognitive test, pupils answer a background questionnaire. In 2018, this included the following question:

> Reading Task: You have received a message in your inbox from a well-known mobile phone operator telling you that you are one of the winners of a smartphone. The sender asks you to click on the link to fill out a form with your data so they can send you the smartphone.
>
> In your opinion, how appropriate are the following strategies in reaction to this email?

Participants were then presented with five actions they may take in response to such an email, being asked how appropriate they felt each would be using a six-point scale (1 = '*not appropriate at all*' to 6 = '*very appropriate*'):

- Answer the email and ask for more information about the smartphone.
- Check the sender's email address.
- *Click on the link to fill out the form as soon as possible*.
- Delete the email without clicking on the link.
- Check the website of the mobile phone operator to see whether the smartphone offer is mentioned.

My specific interest is in the third (italicised) statement, where the respondent indicates they would be likely to click on the link and fill in the form. This,

obviously, should never be done. Anyone who takes this course of action is in danger of being scammed. Throughout most of the analysis I focus on differences between pupils who deemed clicking the link to be 'appropriate' (values 4, 5 or 6 on the six-point scale) versus those who deemed it to be inappropriate (values of 1, 2 or 3 on the six-point scale). Below I discuss various sensitivity analyses that have been conducted – operationalising the outcome measure in alternative ways – to test the robustness of my results. Note that participants did not actually click on any email – and thus no actual behaviour was recorded – just their intentions.

One limitation with such a survey-based approach is that one is capturing how teenagers say they would act rather than observing what they would actually do. Relatedly, there is a risk that responses to this question are subject to undesirable survey behaviour (e.g., careless responses, inconsistent response patterns, low effort). Thus, to reduce the potential impact of such problems, a set of sample selection rules are applied. An overview of these rules is provided in Table 1.

To begin, any student who reported that they did not take the PISA study seriously is excluded.[1] I then remove from the sample pupils displaying extreme response styles – individuals who rated all five options as four and above or two and below across the six-point scale. As a computer-based study, information on pupils' response times is also available. This allows me to exclude from the sample 'rapid responders' – those who provided their answers so quickly it is unlikely that they have read the whole question properly.[2] Likewise, I also exclude a very small number of very slow respondents who spent more than two minutes answering the question. Finally, pupils whose responses were inconsistent across the response options are also excluded – i.e., those who indicated that clicking the link immediately and deleting the email without clicking the link were both appropriate responses (selecting four, five or six for both questions on the six-point response scale). After applying these sample selection rules, the final analytic sample size is 176,186[3] (60% of the full sample).

TABLE 1: Sample restrictions

Sample selection criteria	Observations	% of full sample
0 Full sample	294,527	100%
1 Low effort removed	244,079	83%
2 All high responses removed	228,414	78%
3 All low responses removed	214,990	73%
4 Rapid responders removed	182,816	62%
5 Slow responders removed	181,762	62%
6 Inconsistent respondents removed	176,186	60%

Source = PISA 2018 database. Analysis includes OECD countries only. Observations refers to the total unweighted observations pooling data across all OECD countries.

The background questionnaire also included a set of other questions used in the analysis. For instance, pupils were asked:

'At school, have you ever been taught the following things?'

- '*How to detect phishing or spam emails*'

I explore how these reports – and those of student's school peers – are related to whether they deem clicking the link to be an appropriate response to an unsolicited phishing email.

Methodology

I begin by presenting the percent of 15-year-olds who indicated that clicking the link and providing personal information would be an appropriate response by country (research question 1) or by group (research question 2). These unconditional estimates are then supplemented by the following regression model:

$$G_{ijk} = \alpha + \beta.SES_{ijk} + \gamma.Effort_{ijk} + \delta.Time_{ijk} + \theta.PISA_{ijk} + \rho.u_k + \varepsilon_{ij} \quad (1)$$

Where:

G_{ij} = A binary indicator, coded 0 if the respondent indicated they would not click on the link and coded 1 if they indicated that they would.

SES_i = Quartiles of the PISA Economic, Social and Cultural Index (ESCS) scale. This is a multidimensional measure of family background, combining information on parental education, occupation and home possessions.

$Effort_i$ = Students' reports of the amount of effort they put into the PISA study on a 0–10 scale.

$Time_i$ = The amount of time the student spent reading and responding to the phishing email survey question.

$PISA_i$ = A vector of PISA reading, science and mathematics scores.

u_k = Country fixed effects.

ε_{ij} = Random error term.

i = Student i.

j = School j.

With respect to research question 1, the ρ parameter related to the country fixed-effects (u_k) reflect differences across countries in the risk of responding to the phishing email conditional on the other factors included in the model. In other words, to what extent can cross-country differences be explained by cross-national variation in cognitive skills, socio-economic status and survey effort? Turning to my second research question, the parameters of interest are β and θ. These capture differences in responding to the phishing email between socio-economic status and PISA achievement groups.

Three specifications of the model presented in Equation (1) are estimated. The first does not include any controls. The second then adds controls for the (self-reported) effort the student put into the PISA test and the amount of time they spent reading and responding to the scenario about the phishing emails. This is to investigate whether any of the remaining socio-economic difference can be explained by apparent differences in survey response behaviour (over and above the sample exclusions made based upon such variables, as discussed above). Then, in the final model, PISA scores are added as additional controls. This will in turn reveal the extent that there continues to be socio-economic differences in response to the phishing email, over and above differences in cognitive skills.

Finally, to address the last research question, a regression model is specified of the form:

$$G_{ijk} = \alpha + \beta.\, Inst_{ijk} + \emptyset.SES_{ijk} + \gamma.Effort_{ijk} + \delta.Time_{ijk} + \theta.PISA_{ijk} + u_k + \varepsilon_{ij} \quad (2)$$

Where:

$Inst_{ijk}$ = An indicator coded 1 if the student has been taught about how to detect phishing emails during their time at school, and 0 otherwise.

From the model presented in Equation (2), β captures the extent that students who had received teaching about the dangers associated with unsolicited emails were more or less likely to believe clicking the link in the phishing email was an appropriate response. To test the robustness of results, estimates are presented for specifications with and without including controls. Likewise, the sensitivity of results is presented to using the average response of students within the school (i.e., where *most* students in the school said that they had received such lessons about phishing) rather than using just students' own reports.

Robustness Tests

I test the robustness of my results to using three alternative outcome measures. First, I create a binary outcome that is coded 1 if the respondent reported clicking the link to be the most (or joint most) appropriate response to the phishing email, and coded 0 otherwise. In other words, there is no option that the young person rated more highly than clicking on the link. These results are presented in Appendix A.

Second, there is another clearly inappropriate action in the question presented to students – '*answer the email and ask for more information about the smartphone*'. I hence create an alternative binary variable coded as 1 if the respondent indicated either of these responses (clicking the link or responding to the email) to be appropriate (selected 4, 5 or 6 for either of these options), and coded 0 otherwise. These results are presented in Appendix B.

Finally, the survey organisers include within the PISA dataset a quasi-continuous scale score which reflects the appropriateness of different responses to the phishing email. This is based upon the views of expert reviewers, who suggested the following ordering of the appropriate action to take: 1. Check the sender's email address (most appropriate); 2. Delete the email without clicking on the link; 3. Check the website of the mobile phone operator; 4. Answer the email and ask for more information about the smartphone; 5. Click on the link to fill out the form as soon as possible (least appropriate). The scale created ranges from 0 to 1, capturing the proportion of the student ordering of options that is consistent with the expert ordering. I have standardised this scale to mean 0 and standard deviation 1 and reversed its direction (so that higher scores on the scale indicate less appropriate responses). Estimates from this analysis can hence be interpreted in terms of effect sizes. These results are presented in Appendix C.

3. Results

Research question 1. In which countries are teenagers most likely to respond to the phishing email?

Figure 1 illustrates the percent of teenagers across OECD countries who believe clicking links and providing personal details is an appropriate response to an unsolicited phishing email. On average, around one-in-seven (14%) young people are at risk of responding to the phishing email. Variation across countries is relatively modest, with most countries falling between 10–20%. Notable outliers include the Scandinavian countries of Denmark, Sweden and Finland, where the percent of teenagers likely to respond to phishing solicitation is significantly lower than in other developed countries (6–7%).[4] At the other extreme sits Mexico (30%) and Chile (27%) – the two upper-middle income countries with OECD membership, along with Colombia (20%) – where around a quarter of young people are at risk of responding to a phishing email. Otherwise, there are few clear clusters of countries within the results. For instance, looking towards Asia, while Japanese teenagers are the least likely to be respond to unsolicited emails anywhere in the world (4%) the same cannot be said for their peers in South Korea (19%). Likewise, although Australia (10%) and the UK (9%) sit towards the bottom of Figure 1, in other Anglophone countries the percentage is notably higher, including New Zealand (13%), Canada (14%), Ireland (15%) and the United States (15%). Thus, while Scandinavian countries and upper-middle income countries are notable outliers (sitting at opposite extremes) cross-country variation in the risk of responding to phishing emails amongst teenagers is otherwise relatively limited.

Figure 2 illustrates the extent that the cross-country pattern illustrated in Figure 1 remains intact once gender, socio-economic status, cognitive skills and

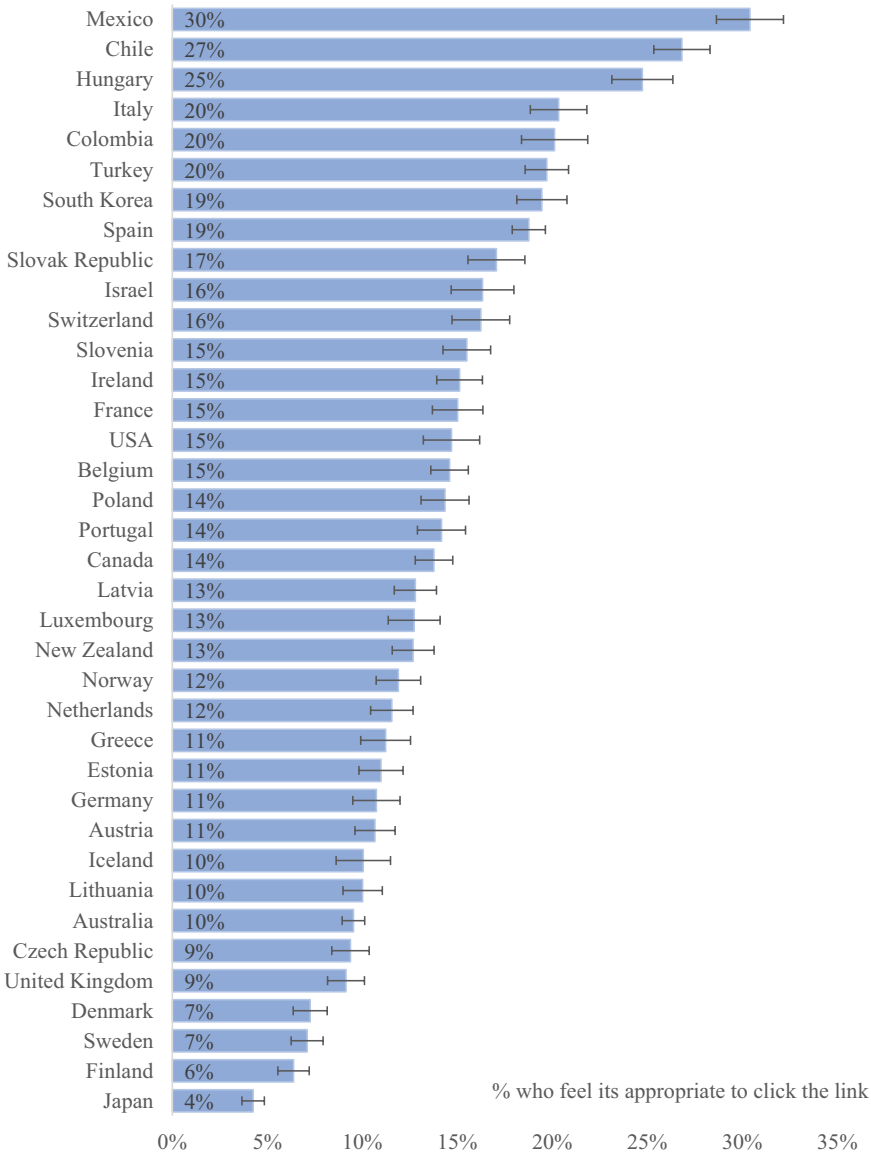

Figure 1: The percentage of 15-year-olds at risk of responding to phishing emails across countries

Figures refer to the percentage of 15-year-olds in each country who believe it is appropriate to click on a link within an unsolicited email. Thin line through the centre of each bar illustrates the estimated 95% confidence interval.

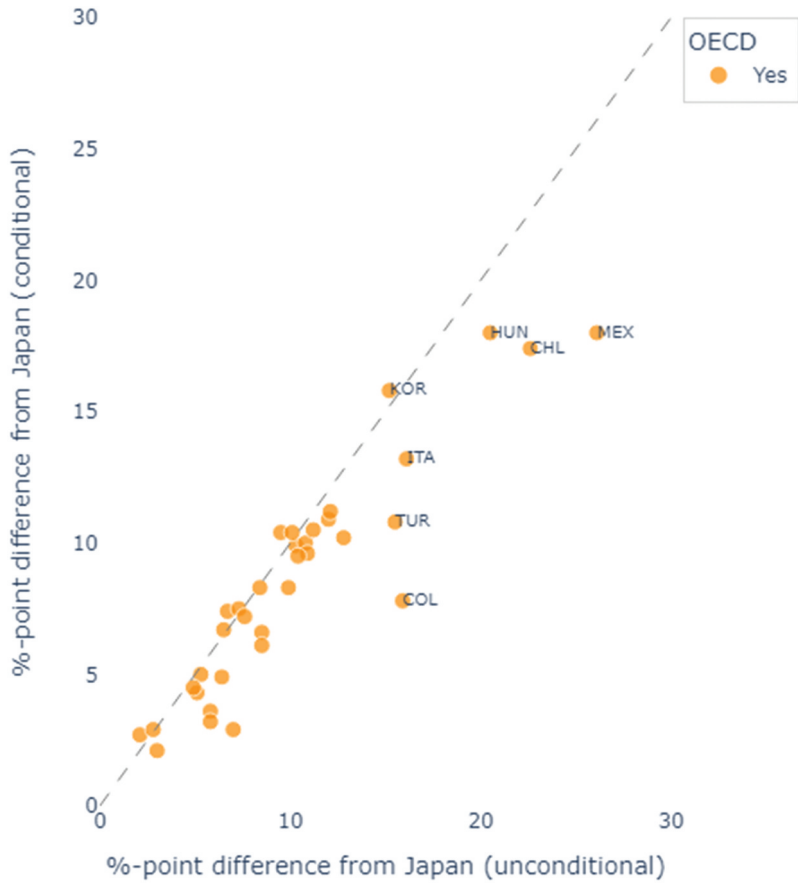

Figure 2: Cross-country differences in the likelihood teenagers respond to phishing solicitation. Conditional versus unconditional estimates

An interactive version of this graph is available within online Appendix D. Figures thus refer to percentage point differences in the probability of responding to the phishing email relative to Japan. Unconditional estimates presented on the horizontal axis, and conditional estimates presented on the vertical axis (controlling for gender, socio-economic status, survey effort and cognitive skills). 45-degree line illustrates where the conditional and unconditional estimates are equal. Pearson correlation = 0.92.

effort put into the survey have been controlled. The estimates presented capture how much more likely 15-year-olds are in each country to respond to the phishing email than their peers in Japan as the reference group (Japan had the smallest proportion of young people − 4% - who felt clicking the link was an appropriate response). Figures thus refer to percentage point differences in the probability of clicking the link relative to Japan. Unconditional estimates are

presented on the horizontal axis, and conditional estimates on the vertical axis. An interactive version of this graph – allowing readers to zoom and identify individual countries – is provided in Appendix D.

The key finding from Figure 2 is that the cross-country correlation across the two sets of results is very strong; the Pearson correlation is 0.92. This illustrates that, on the whole, the cross-country differences presented in Figure 1 cannot be explained by differences in cognitive skills or survey effort. Colombia (COL) is a notable exception; once cognitive skills and survey effort have been controlled, the percent of teenagers who would click on the phishing email moves much closer to the international average. It is interesting to contrast Colombia to South Korea in Figure 2; two countries where the unconditional estimates are very similar, but the conditional estimates are rather different. In other words, South Korea has a large percentage of teenagers who report that they would respond inappropriately to a phishing email, despite their high-level of cognitive skills. Otherwise, it remains clear that Italy, Hungary, Chile and Mexico continue to stand out as countries where young people are most likely to respond to phishing solicitation, even once cognitive skills, socio-economic composition and survey effort have been controlled.

Research question 2. How is the likelihood of responding to phishing emails related to young people's characteristics, including socio-economic background and PISA scores?

Table 2 begins by presenting a set of descriptive statistics illustrating the percentage of young people who report they would click on links within the phishing email across demographic groups. These estimates are based upon the data pooled across all OECD nations, with each country taking equal weight.

Table 2 illustrates how there is no gender difference; boys are just as likely to respond to phishing solicitation as girls. Likewise, differences between immigrants and country natives are also relatively muted. In contrast, there are clear differences by socio-economic status; teenagers from socio-economically disadvantaged backgrounds are markedly more likely to click links in unsolicited emails than their more advantaged peers. However, by far the biggest gap is between young people with different levels of cognitive skill. While a quarter (25%) of low-achieving students (bottom quintile of PISA reading scores) suggest they believe clicking the link to be an appropriate response, this fall to just one-in-twenty (5%) of those in the top quintile of PISA reading scores.

Table 3 digs further into the socio-economic gap by presenting results from the regression models. In particular, it illustrates the extent that socio-economic differences can be explained by residual differences in survey response behaviour (effort, response time) in model M1, and the extent that it can be explained by differences in their cognitive skills – as measured by PISA scores – in model M2.

TABLE 2: Responding to phishing solicitation by teenagers background characteristics

Variable	Group	Percent
Gender	Male	14%
	Female	14%
Socio-economic status quartile	Low SES	17%
	Q2 SES	15%
	Q3 SES	13%
	High SES	11%
Reading ability quintile	Low achievement	24%
	Q2 achievement	19%
	Q3 achievement	14%
	Q4 achievement	9%
	High achievement	5%
Immigrant status	Country native	14%
	First generation immigrant	18%
	Second generation immigrant	15%

Figures refer to the percentage of pupils who believe it is appropriate to click the link in an unsolicited email. Each country carries equal weight in the analysis (senate weights applied).

TABLE 3: Socio-economic inequality in responding to phishing solicitation. Regression model estimates

	M0		M1		M2	
	Beta	SE	Beta	SE	Beta	SE
Q1 Most disadvantaged (Ref)	-	-	-	-	-	-
Q2 SES	−2.0%*	0.4%	−1.6%*	0.3%	0.2%	0.3%
Q3 SES	−3.9%*	0.3%	−3.3%*	0.3%	−0.1%	0.3%
Q4 Most advantaged quartile	−6.5%*	0.3%	−5.5%*	0.3%	−0.4%	0.3%
N	**150,362**		**150,344**		**150,344**	
Controls						
Effort	-		Y		Y	
Response time	-		Y		Y	
Country fixed effects	-		Y		Y	
PISA test scores	-		-		Y	

Estimates based on linear probability models pooled across OECD countries with senate weights applied (each country carrying equal weight). Beta column refers to the percentage point difference in believing clicking on the link within an unsolicited email is appropriate compared to the reference group (pupils from the most disadvantaged backgrounds). SE refers to the estimated standard error. Spain excluded from the sample due to missing information on participants reading test scores. * indicates statistical significance at the 5% level.

In the unconditional model (M0) socio-economically advantaged students are found to be 6.5% points less likely to say they would click the link than the most disadvantaged pupils. This falls only slightly- to 5.5% points – when test effort, question response time and country fixed effects are added to the model.

It hence seems unlikely that the socio-economic gap in responding to phishing solicitation can be explained by differences in survey response behaviour. In contrast, once controls for PISA scores have been added to the model (M2) the socio-economic gradient to responses disappear and are no longer statistically significant at conventional levels. Thus, it seems that the primary reason why socio-economic groups are more likely to fall prey to phishing emails is their lower levels of cognitive skills.

Figure 3 presents the picture in terms of individual countries. Figures refer to the magnitude of the socio-economic gap, in terms of the percentage point difference in the likelihood of responding to the phishing email. Unconditional estimates are presented along the horizontal axis, with conditional estimates – controlling for test effort and cognitive skills – presented on the vertical axis. The dashed lines represent where the estimated socio-economic gap is zero.

In terms of individual countries, Colombia, Iceland, Estonia and Japan stand out as not having a sizeable socio-economic gap even in the unconditional model. Colombia and Portugal also stand out as nations where – after controlling for differences in cognitive skills – the most advantaged socio-economically group are significantly more likely to be at risk from phishing emails than the most disadvantaged group (with a conditional difference of around five percentage points or more). At the other extreme, Switzerland is the only country where socio-economically disadvantaged students are significantly more likely to be at risk of clicking on links in phishing emails than their more advantaged peers even after differences in cognitive skills have been controlled (a conditional difference of five percentage points).

Research question 3. To what extent can teaching students in school make them less likely to respond to phishing emails?

To conclude, I consider whether students who receive instruction from their school about the dangers associated with phishing emails are at less risk of being fooled by phishing solicitation. In particular, Table 4 illustrates the percentage point change in the likelihood that young people believe that clicking a link in an unsolicited email is an appropriate response if they have been taught about such online dangers at school. To test the robustness of the results, estimates are presented using both pupils' own reports and the reports of their peers about whether they have been taught about phishing emails at schools.

Overall, there is no clear evidence that students who are taught about the dangers of phishing at school are less likely to respond to unsolicited emails. The estimated differences reported in Table 5 are always small, with the point estimate changing direction depending on whether students' own reports or the school-average report is used. Indeed, there is no country where consistent evidence is found of a sizeable and statistically significant difference across

Figure 3: Socio-economic differences in the likelihood of responding to a phishing email by country. Conditional versus unconditional estimates

Notes: An interactive version of this plot is supplied in Appendix E. Figures refer to the percentage point difference in young people responding to the phishing email between the most advantaged and least advantaged socio-economic groups. The unconditional gap is presented on the horizontal axis, with the conditional gap (controlling for cognitive abilities and survey effort) on the vertical axis. Dashed vertical and horizontal line indicates where there is no difference in the likelihood of responding to the email between socio-economic groups. Blue markers indicate where neither the conditional or unconditional estimates are statistically significant. The circular (square) markers indicate where the conditional (unconditional) estimates are statistically significant.

the different approaches (see Table 5 for further details). Hence, as currently implemented, schools across industrialised nations appear to be achieving little in helping young people to respond appropriately to common online dangers – such as phishing emails.

TABLE 4: The link between receiving instruction about phishing in school and inappropriately clicking links in unsolicited emails

	Pupil report				School average report			
	Unconditional		Conditional		Unconditional		Conditional	
	Beta	SE	Beta	SE	Beta	SE	Beta	SE
% point change	1.1%*	0.2%	0.5%*	0.2%	−0.7%*	0.2%	−0.2%	0.2%
N	**171,773**		**149,687**		**172,573**		**150,344**	

Notes: Figures refer to the percentage point change in believing clicking the link is an appropriate response to an unsolicited email. Positive figures indicate that students who reported receiving lessons about phishing were more likely to say they would click the link. For reference, the cross-country average is 14%. Estimates based upon the PISA 2018 sample pooled across all OECD countries with data available. Senate and BRR weights applied so that each country in the analysis carries equal weight. Conditional estimates include controls for gender, socio-economic status, self-reported effort on the PISA test, country fixed effects and PISA reading, science and mathematics test scores. * indicates statistical significance at the 5% level.

Robustness Tests

In Appendix A, B and C I present a series of alternative estimates for each research question defining the outcome measure a different way. This includes investigating whether students selected clicking the link as the most (or joint most) appropriate option (Appendix A), whether they selected either clicking the link or responding to the email as appropriate responses (Appendix B) and using a continuous scale based upon a hierarchical ordering of the five options (Appendix C).

With regards research question 1, the cross-country pattern of results remains similar regardless of which approach is used. Countries such as Mexico, Hungary, South Korea and Colombia consistently stand out as having a comparatively high proportion of inappropriate responses. In contrast, Japan, the United Kingdom and several Scandinavian countries have comparatively low levels of inappropriate responses. The cross-national ordering of countries is again largely insensitive to whether measures of cognitive ability or survey effort are controlled or not, with a handful or partial exceptions (e.g., Colombia, Mexico).

Turning to the second research question, few differences continue to be observed between genders and between immigrants and natives. Bigger differences are found between socio-economic groups and children with different cognitive abilities. However, under all three alternative approaches, the socio-economic gap in young people's responses to the phishing email is almost entirely explained by differences in their cognitive skills. This is consistent with the key findings presented above.

Finally, I continue to find little association between instruction provided by schools about the dangers posed by phishing emails and the appropriateness of

TABLE 5: The link between receiving instruction about phishing in school and inappropriately clicking links in unsolicited emails. Country estimates

	Pupil report		School average report	
Country	Beta	SE	Beta	SE
Chile	7.2*	1.8	−7.9*	3.3
Mexico	4.3*	1.4	0.6	1.7
Norway	3.8*	1.8	5.1	7.2
Turkey	2.9*	1.5	1.0	3.4
Iceland	2.6	1.6	−0.9	1.8
Israel	2.5	1.7	7.3	4.1
Slovak Republic	2.3	1.4	−2.5	1.7
Italy	2.0	1.8	−4.4*	2.1
Estonia	1.5	1.2	−0.3	1.1
Sweden	1.4	1.1	−2.8	1.5
Greece	1.3	0.9	1.4	1.4
Finland	1.1	0.9	−0.6	0.8
Germany	1.0	1.3	−1.5	2.0
Ireland	1.0	1.3	1.3	3.5
Latvia	0.9	1.3	0.3	1.3
Czech Republic	0.8	0.9	0.7	1.1
New Zealand	0.7	1.3	0.0	0.9
Netherlands	0.3	1.4	−3.0	3.9
Portugal	0.2	1.1	0.2	1.2
Canada	0.2	0.9	−0.8	1.3
United States of America	0.1	1.3	0.3	1.4
France	0.0	1.3	−1.6	1.9
Australia	−0.2	0.7	−0.3	0.8
Austria	−0.2	1.0	1.0	0.9
Lithuania	−0.3	0.9	0.2	1.0
Denmark	−0.4	0.9	−1.1	1.1
United Kingdom	−0.4	0.8	0.6	0.9
Colombia	−0.6	2.0	2.3	2.4
Japan	−0.7	0.8	−0.4	0.6
Poland	−0.8	1.2	−0.8	1.3
Switzerland	−1.3	1.6	−3.2	2.0
Hungary	−1.3	1.6	0.9	1.7
Slovenia	−1.8	1.4	−1.2	1.4
Belgium	−1.9	1.2	−0.2	1.5
Luxembourg	−2.2	1.3	0.2	1.6
South Korea	−3.6*	1.4	−0.3	3.0

Figures refer to the percentage point difference in the likelihood of responding to the phishing email if the pupil was taught about phishing at school. Positive figures indicate that students who reported receiving lessons about phishing were more likely to say they would click the link. SE column refers to the estimated standard errors. * indicates statistically significant difference from zero at the 5% level.

student's responses to the PISA phishing question. Across the three alternative approaches, effect sizes are either very small, statistically insignificant, or inconsistent in their direction depending upon whether pupil's own reports or the school-average report is used. This supports the conclusion that there is little evidence that schools across the OECD currently provide effective instruction regarding the dangers of phishing emails.

4. Conclusions

Phishing remains one of the most common attempts at cyber fraud, with it estimated that around 3.4 billion spam emails are sent every day (AAG, 2023). Yet, despite the prominence of this problem, there are gaps in our knowledge about who is likely to respond to phishing emails. For instance, to what extent does it vary by country and across key demographic groups? To what extent can such variation be explained by differences in cognitive abilities? Does the instruction currently provided by schools help reduce the chances that young people get snared by such traps?

This paper has sought to present new international evidence on such issues. In doing so, it presents novel evidence on the likelihood that young people across the industrialised world are at risk of being fooled by a phishing attempt. I find that around one-in-seven young people are at risk of responding inappropriately to phishing emails. There is limited evidence of cross-country variation, though with the percentage at risk being particularly low across Scandinavia, Japan and the United Kingdom, while being notably higher in the upper-middle income members of the OECD (e.g., Mexico, Chile). While there are no clear differences across genders, socio-economically disadvantaged groups are – at least in some countries – at greater risk from phishing attacks than their more advantaged peers. This, however, is largely being driven by socio-economic differences in cognitive abilities. Unfortunately, current attempts by schools to address this issue do not seem to be particularly effective. Indeed, teenagers who have been taught at school about the risks posed by phishing emails appear just as likely to take inappropriate action in response to one as their peers who have not received any such instruction.

How do these findings compare to those within the existing literature? The cross-national pattern of countries observed does not bear much relation to the ordering of countries in studies such as Cook *et al*. (2023) or Chen *et al*. (2023). This, however, is not particularly surprising, given my focus on teenager's propensity to respond to phishing emails, while Cook *et al*. (2023) investigated fear of economic crime and Chen *et al*. (2023) on the prevalence of cyber-crime across countries. What this does go to show, however, is a greater need for research that links the behaviour of teenagers in response to attempted online fraud to how frequently they face such threats. For instance, are teenagers in countries that face more cyber-crime any more thoughtful in their responses to

phishing solicitation than their peers where cyber-crime is comparatively rare? Future work, collecting internationally comparable data on both the risks teenagers face and their responses to attempted online fraud, is an area ripe for future research.

My finding of a strong association between cognitive skills and response to phishing solicitation is consistent with some previous studies into this issue, but not others. Most previous work that has found a relationship between cognitive skills and the probability of being defrauded has done so in older adults; for instance, Ebner *et al*. (2020) reported such an association for those aged 75–89, but not for younger (age 18–37) individuals. Likewise, previous findings regarding educational achievement have been mixed, with some studies finding an association (although not always in the hypothesised direction – e.g., Lee and Geistfeld, 1999), while others have failed to do so. What many of these studies have lacked, however, is a very high-quality measure of participants cognitive ability. The PISA data provides a major advance in this area, given that it has measured teenagers' cognitive skills across three domains using a two-hour test. It hence provides perhaps the strongest evidence to date that low cognitive ability is associated with an increased risk of responding to phishing solicitation even amongst younger age groups. In terms of future research, this finding points towards the need for future studies investigating the link between cognitive abilities and the chances of being the victim of cyber-crime to ensure a high-quality measure of participants cognitive skills is collected.

Finally, turning to socio-economic status, the previous literature has proven to be somewhat inconclusive. An important distinction of course needs to be made between being targeted by cyber-crime and the actions individuals take when defrauding is attempted. For instance, the Office for National Statistics (2022) shows that higher socio-economic groups are more likely to be targeted by phishing emails than lower socio-economic groups, but are then also less likely to reply to the email or click on the links within it. My results speak only to the action that young people from different socio-economic backgrounds take and, even then, the picture differs across countries. For instance, socio-economically disadvantaged young people were much more likely to click the link within the hypothetical phishing email than their more advantaged peers in countries such as Chile, while there is next to no socio-economic difference in Estonia and Japan. Moreover, in almost every country, the socio-economic gap in responding to phishing solicitation disappears once young people's PISA test scores have been controlled. This in-turn suggests that the reason why young people from disadvantaged backgrounds are more likely to respond to phishing emails is their weaker cognitive skills. Collecting high-quality measures of cognitive skills is hence also vital in studies seeking to decipher the independent risk factors associated with cyber-crime.

Readers should consider these findings considering the limitations of this work. Four issues stand out. First, the analysis is based upon responses to survey questions. These capture how young people say they would act under a hypothetical scenario, rather than being based upon observations of how they would actually respond when such an unsolicited email is received. Although a study where many teenagers are sent a (harmless) phishing email to investigate their actual behaviour is likely to pose significant ethical and logistical challenges, it would also offer unparalleled opportunities to further research in this area. Future studies should hence seek to obtain ethical approval to deploy such a study design, providing the next generation of evidence needed to better understand teenagers' reactions to phishing attempts. Second, I have investigated students' responses to one specific online danger in the form of a phishing email. Yet, with young people facing a range of ever more sophisticated online threats, future studies should hence seek to establish how easily tricked young people may be to more subtle online scams. Third, the PISA data do not contain any information about the *quality* of instruction students have received about dangers online – just simply whether they had received any lessons at all. More clearly needs to be done to better understand what 'quality' provision looks like in this area, and whether the best current approaches are any more effective at reducing the risk of young people responding to phishing emails (and online fraud more generally). Finally, my analysis has used cross-sectional data only, and not been able to establish cause and effect. Future research should seek to conduct longitudinal studies that build our understanding of the ages, stages and drivers of young people's developing understanding of online risks and dangers.

With these caveats in mind, the findings nevertheless hold some important implications for education policy and practice. More needs to be done to help young people to navigate what is becoming an increasingly complex and dangerous online world. This is particularly true for some of the most vulnerable groups – those from disadvantaged backgrounds and with lower levels of cognitive abilities – who are most at risk of falling for attempts at digital fraud. There is a clear gap for greater emphasis in schools for providing effective instruction in this area, including greater quantity and quality of time devoted to it. However, before this can be realistically achieved, a much stronger evidence base first needs to be developed. What, exactly, does a 'quality' curriculum look like in helping young people to recognise online risks (including – but not limited to – phishing emails)? It is vital now that interventions and initiatives are tested in schools to establish what is effective in helping teenagers to better understand phishing attacks and – when they are faced with them – the most appropriate cause of action to take. Only then will we stand a realistic chance of truly building young people's competencies in this important area and, in-turn, minimise susceptibilities that certain unscrupulous groups seek to exploit.

Disclosure Statement

No potential conflict of interest was reported by the author(s).

Supplementary data

Supplemental data for this article can be accessed online at https://doi.org/10.1080/00071005.2023.2234456.

Notes

[1] Pupils were asked to say how much effort they put into the PISA test using a 0–10 scale. I exclude any pupils reporting their effort to be 4 or below.

[2] I exclude any pupil who responded to all five statements in 15 seconds or less. The total length of the question is approximately 120 words. Based upon the average reading rate of adults of approximately 240 words per minute (Brysbaert, 2019), the question would take an average person around 30 seconds to read (without any time allocated for thinking and responding). In this context, response times of 15 seconds would be exceedingly quick – with concerns about whether pupils have read the question properly and have taken the task seriously.

[3] Spain is dropped from parts of the analysis. This is due to issues with PISA reading scores for this country in the 2018 round, which is one of the controls included in the statistical models. See OECD (2020) for further details.

[4] Norway – the other Scandinavian country – also sits just below the OECD average at 12%.

ORCID

John Jerrim http://orcid.org/0000-0001-5705-7954

References

AAG (2023) *The Latest 2023 Phishing Statistics*. Available at: https://aag-it.com/the-latest-phishing-statistics/#:~:text=Phishing%20is%20the%20most%20common,sent%20in%202021%20were%20spam.

Alseadoon, I. M. and Othman, M. F. I. (2021) Cultural comparison towards users' susceptible to phishing emails, *International Journal of Computer Science & Network Security*, 21 (10). doi:10.22937/IJCSNS.2021.21.10.33.

Anderson, K. B. (2019). *Mass-market consumer fraud in the United States: A 2017 update [Staff report]*. Federal Trade Commission, Bureau of Economics. https://www.ftc.gov/system/files/documents/reports/mass-market-consumer-fraud-united-states-2017-update/p105502massmarketconsumerfraud2017report.pdf

Brysbaert, M. (2019) How many words do we read per minute? A review and meta-analysis of reading rate. *Journal of Memory and Language*, 109 (104047), 1–30. doi:10.1016/j.jml.2019.104047.

Cattan, S., Fitzsimons, E., Goodman, A., Phimister, A., Ploubidis, G. B. and Wertz, J. (2022) Early childhood and inequalities, *IFS Deaton Review of Inequalities*. Available at: https://ifs.org.uk/inequality/early-childhood-inequalities-chapter.

Chen, S., Hao, M., Ding, F., Li, Y. and Zhang, Y. (2023) Exploring the global geography of cybercrime and its driving forces. *Humanities and Social Sciences Communications*, 10 (1), 71. doi:10.1057/s41599-023-01560-x.

Cook, S., Giommoni, L., Trajtenberg Pareja, N., Levi, M. and Williams, M. L. (2023) Fear of economic cybercrime across Europe: A multilevel application of routine activity theory. *The British Journal of Criminology*, 63 (2), 384–406. doi:10.1093/bjc/azac021.

DeLiema, M., Deevy, M., Lusardi, A. and Mitchell, O. S. (2020) Financial fraud among older Americans: Evidence and implications. *The Journals of Gerontology: Series B*, 75 (4), 861–868. doi:10.1093/geronb/gby151.

Ebner, N. C., Ellis, D. M., Lin, T., Rocha, H. A., Yang, H., Dommaraju, S., Soliman, A., Woodard, D. L., Turner, G. R., Spreng, R. N. and Oliveira, D. S. (2020) Uncovering susceptibility risk to online deception in aging. *Journals of Gerontology, Series B: Psychological Sciences & Social Sciences*, 75 (3), 522–533. doi:10.1093/geronb/gby036.

European Commission (2020) *Survey on scams and fraud experienced by consumers: final report*. Available at: https://commission.europa.eu/system/files/2020-01/survey_on_scams_and_fraud_experienced_by_consumers_-_final_report.pdf.

Gavett, B. E., Zhao, R., John, S. E., Bussell, C. A., Roberts, J. R. and Yue, C. (2017) Phishing suspiciousness in older and younger adults: the role of executive functioning. *PLoS ONE*, 12 (2), e0171620. doi:10.1371/journal.pone.0171620.

Glover, C. M., Yu, L., Stewart, C. C., Wilson, R. S., Bennett, D. A., Lamar, M. and Boyle, P. A. (2023) Childhood socioeconomic status interacts with cognitive function to impact scam susceptibility among community-dwelling older adults. *Aging & Mental Health*, 27 (4), 765–770. doi:10.1080/13607863.2022.2087206.

Hanoch, Y. and Wood, S. (2021) The scams among us: Who falls prey and why. *Current Directions in Psychological Science*, 30 (3), 260–266. doi:10.1177/0963721421995489.

Jerrim, J., Lopez-Agudo, L. A., Marcenaro-Gutierrez, O. D. and Shure, N. (2017) What happens when econometrics and psychometrics collide? An example using the PISA data, *Economics of Education Review*, 61, 51–58.

Kipping, R. R., Smith, M., Heron, J., Hickman, M. and Campbell, R. (2015) Multiple risk behaviour in adolescence and socio-economic status: Findings from a UK birth cohort. *European Journal of Public Health*, 25 (1), 44–49. doi:10.1093/eurpub/cku078.

Lee, J. and Geistfeld, L. V. (1999) Elderly consumers' receptiveness to telemarketing fraud. *Journal of Public Policy & Marketing*, 18 (2), 208–217. doi:10.1177/074391569901800207.

Mueller, E. A., Wood, S. A., Hanoch, Y., Huang, Y. and Reed, C. L. (2020) Older and wiser: age differences in susceptibility to investment fraud: The protective role of emotional intelligence. *Journal of Elder Abuse & Neglect*, 32 (2), 152–172. doi:10.1080/08946566.2020.1736704.

OECD (2020) *PISA 2018 technical report* (Paris, OECD). Available at: https://www.oecd.org/pisa/data/pisa2018technicalreport/.

OECD (2021) *Are 15-year-olds prepared to deal with fake news and misinformation?* Available at: https://www.oecd-ilibrary.org/education/are-15-year-olds-prepared-to-deal-with-fake-news-and-misinformation_6ad5395e-en.

Office for National Statistics (2022) *Phishing attacks – who is most at risk?* Available at: https://www.ons.gov.uk/peoplepopulationandcommunity/crimeandjustice/articles/phishingattackswhoismostatrisk/2022-09-26.

Reep-van den Bergh, C. M. M. and Junger, M. (2018) Victims of cybercrime in Europe: A review of victim surveys. *Crime Science*, 7 (1), 5. doi:10.1186/s40163-018-0079-3.

Viner, R. M., Ross, D., Hardy, R., Kuh, D., Power, C., Johnson, A., Wellings, K., McCambridge, J., Cole, T. J., Kelly, Y. and Batty, G. D. (2015) Life course epidemiology: recognising the importance of adolescence. *Journal of Epidemiology & Community Health*, 69 (8), 719–720. doi:10.1136/jech-2014-205300.

Whitty, M. T. (2019) Predicting susceptibility to cyber-fraud victimhood. *Journal of Financial Crime*, 26 (1), 277–292. doi:10.1108/JFC-10-2017-0095.

Wood, S., Liu, P.-J., Hanoch, Y., Xi, P. M. and Klapatch, L. (2018) Call to claim your prize: perceived benefits and risk drive intention to comply in a mass marketing scam. *Journal of Experimental Psychology: Applied*, 24 (2), 196–206. doi:10.1037/xap0000167.

Zhang, Z. and Ye, Z. (2022) The role of social-psychological factors of victimity on victimization of online fraud in China, *Frontiers in Psychology*, 13, 1030670. doi:10.3389/fpsyg.2022.1030670.

Correspondence
John Jerrim
University College London, 55-59 Gordon Square, Bloomsbury, London
Email: J.Jerrim@ucl.ac.uk