Exposure to online self-harm content and self-harm among secondary school students: associations and quasi-experimental evidence

Lucas Sempé[1,2,*] Mina Fazel[2] Galit Geulayov[2]
Emma Soneson[2]

[1] International Initiative for Impact Evaluation (3ie), United Kingdom.
[2] Department of Psychiatry, Oxford University, Oxford, United Kingdom.

[*] Correspondence: Lucas Sempé <lsempe@3ieimpact.org>

Abstract

Background. Exposure to self-harm content on social media has been linked to normalisation and promotion of self-harm among adolescents, but whether this relationship is causal remains unknown. We examined (1) associations between different dimensions of online self-harm content exposure and self-harm, and (2) the causal effect of content exposure on self-harm using an instrumental variable approach.

Methods. We used data from 24,909 students (Years 7–13, ages 11–19) in the 2023 OxWell Student Survey, a school-based cross-sectional survey across four English regions. Self-harm was classified using structured self-report questions and expert panel review. We examined four dimensions of self-harm content exposure: frequency of encountering content (5-level ordinal), type of online content consumed (11 categories), mode of encountering content (active search, feed, accidental, shared), and motivations for seeking content. We fitted multilevel logistic regression models adjusting for gender, year group, ethnicity, care status, deprivation, and neurodivergence. To address reverse causality, we employed two-stage least squares (2SLS) instrumental variable analysis using daily time spent on the internet as the instrument, restricting the outcome to recent self-harm (past month).

Findings. Frequency of self-harm content exposure showed a strong dose–response association with self-harm, with daily exposure associated with an adjusted OR of 17.9 (95% CI 14.1–22.8). Active searching for self-harm content showed the strongest association among exposure modes (aOR 11.4, 95% CI 8.8–14.7), with a predicted probability of self-harm of approximately NaN%.

The instrumental variable analysis yielded estimates consistent with a causal effect: the 2SLS linear probability model indicated a 5.8 percentage-point increase in the probability of self-harm per unit increase in content exposure ($p < 0 \cdot 001$), and the control-function approach for categorical exposure yielded an OR of 54.9 ($p < 0 \cdot 001$). First-stage F-statistics of 438 and 498 confirmed instrument strength.

Interpretation. Exposure to online self-harm content is strongly associated with self-harm among secondary school students, and quasi-experimental analysis supports a causal component to this relationship. Active engagement with such content shows substantially stronger associations than passive or accidental exposure, suggesting that platform-level content moderation alone is insufficient: young people who deliberately seek self-harm content also need clinical support addressing the underlying motivations. These findings support both regulatory interventions targeting algorithmic amplification and clinical approaches for active seekers.

Funding. *[To be completed]*

Research in Context

Evidence before this study. Existing reviews and meta-analyses have found associations between exposure to online self-harm content and self-harm behaviour among adolescents, but virtually all evidence comes from cross-sectional or qualitative studies that cannot separate cause from consequence. Whether individuals who already self-harm are more likely to seek or be algorithmically served this content is a plausible and largely unexamined confound. No study has applied instrumental variable methods or other quasi-experimental approaches to estimate the causal effect of online self-harm content exposure on self-harm in a large adolescent sample.

Added value of this study. To our knowledge, this is the first study to combine a large school-based sample (n > 24,000) with an instrumental variable design to provide quasi-experimental evidence on the effect of online self-harm content exposure on self-harm among adolescents. We document a strong dose–response association between exposure frequency and self-harm, show that active searching carries markedly higher odds than passive or feed-based exposure, and demonstrate through two IV specifications that the association is not attributable to reverse causation alone. First-stage F-statistics confirm instrument validity, and sensitivity analyses address potential violations of the exclusion restriction.

Implications of all the available evidence. The quasi-experimental evidence strengthens the case for regulatory and platform-level interventions to reduce the availability and algorithmic amplification of self-harm content online. The finding that active searching shows the strongest associations points to the limits of content moderation alone: many at-risk young people are deliberately seeking this material and need therapeutic support, not just restricted access. These results are directly relevant to ongoing policy debates around the UK Online Safety Act and equivalent legislation internationally.

Introduction

Self-harm among adolescents is a growing public health problem. Incidence has risen across the United Kingdom and other Western countries over the past two decades (Morgan et al. 2017; Griffin et al. 2018), with hospital presentations for self-harm among young people reaching historically high levels (Geulayov et al. 2022). The consequences extend beyond the immediate episode: self-harm in adolescence is a strong predictor of repeated self-harm, psychiatric morbidity, and suicide in adulthood, making it both a clinical concern and a population-level priority.

A growing body of research has linked social media use to adverse mental health outcomes among young people (Odgers and Jensen 2020; Valkenburg, Meier, and Beyens 2022), including depressive symptoms (Cunningham, Hudson, and Harkness 2021), social isolation (O'Day and Heimberg 2021), and cyberbullying (Giumetti and Kowalski 2022; Hamm et al. 2015). One prominent candidate mechanism is exposure to self-harm content online: encountering such material may normalise self-harm, provide instruction on methods, and — in online communities organised around self-harm — actively encourage it (Biernesser et al. 2020; Dyson et al. 2016). A recent meta-analysis documented significant associations between online self-harm content exposure and self-harm behaviour (Nesi et al. 2021), but this literature relies almost entirely on cross-sectional or qualitative designs.

Two important gaps therefore remain. First, the direction of causality is unresolved. Reverse causation — whereby individuals who already self-harm are more likely to encounter or seek out this content — is a plausible and largely unaddressed confound. No study to date has applied instrumental variable methods or other quasi-experimental approaches to estimate the causal effect of online self-harm content exposure on self-harm in a large adolescent sample. Second, exposure is rarely disaggregated by pathway. Whether content is encountered through active searching, algorithmic feeds, accidental discovery, or interpersonal sharing may carry different risk implications, but these pathways are seldom examined separately.

Understanding the mode of exposure is important because it has different policy and clinical implications. Algorithmic amplification of self-harm content can be addressed through platform regulation and content moderation. Accidental exposure may be reduced through better filtering. But if active searching is the dominant pathway, then interventions need to address the underlying motivations — such as help-seeking and desire for peer connection — rather than relying solely on access restriction. The existing literature does not distinguish among these pathways.

We address both gaps using data from the 2023 OxWell Student Survey. We first examine associations between self-harm and four dimensions of online self-harm content exposure:

frequency of encountering such content, type of content consumed, mode of exposure, and motivations for seeking it. We then estimate the causal effect of exposure on self-harm using an instrumental variable (IV) approach, with daily time spent on the internet as the instrument for content exposure. This study has been pre-registered (Sempé et al. 2024).

Methods

Study design and participants

We used data from the 2023 wave of the OxWell Student Survey, a school-based cross-sectional survey conducted annually across four regions in England (Mansfield et al. 2021). Schools are recruited through local authority and multi-academy trust partnerships; within participating schools, all students in school Years 7–13 (ages approximately 11–19) are invited to take part. Participation is voluntary, with no monetary incentives. The 2023 wave was completed by 37,952 students. After excluding respondents whose self-harm status could not be determined — those answering "Prefer not to say", "Not sure what this means", or with missing gateway responses — the primary analytic sample comprised 24,909 students from 177 schools. For the instrumental variable analysis, we further restricted to respondents who reported their most recent self-harm as "in the last week" or "in the last month", aligning the temporal frame of the outcome with the content-exposure question (past month), and excluded those with missing data on the instrument or covariates; the resulting IV sample size is reported in the model output. The study was approved by the University of Oxford Medical Sciences Interdivisional Research Ethics Committee (Ref: R67020/RE001). Headteachers gave institutional consent; students provided individual assent; parents were notified in advance and could withdraw their child. The data are not publicly available owing to the sensitive content and the age of participants; access to anonymised data may be requested from the OxWell Study Team under a data-sharing agreement.

Measures

Self-harm was identified through a two-step process designed to capture cases that structured items alone may miss (Geulayov et al. 2022). Respondents who endorsed a gateway item on deliberate self-harm and confirmed at least one method (self-injury or self-poisoning) were classified as self-harmers. In addition, an expert panel independently reviewed open-text responses and classified additional cases; disagreements were resolved by consensus.

We examined four dimensions of online self-harm content exposure. Frequency of exposure was assessed by asking respondents whether they had come across content about self-harm on online platforms in the last month, with response options ranging from "No" to "Yes, daily". Type of online content was captured by asking respondents which types of content they typically viewed online from a list of 11 categories (e.g., entertainment and humour, health and fitness, forums, family and friends); a separate model was fitted for each content type. Mode of encountering content was assessed among those who had seen self-harm content, covering active searching, appearing in their feed, accidental discovery, or sharing by others, using questions drawn from the instrument developed by Skripkauskaite and Fazel (2022) and the OxWell Study Team. Finally, motivations for seeking content were assessed among those who reported actively seeking self-harm content, including finding support, connecting with others who had similar experiences, or looking up facts and figures.

All models adjusted for gender (four categories: Boy, Girl, Gender Diverse, Gender Non-Disclosing), school year group, self-reported ethnicity (ONS categories ("Developing Admin-Based Ethnicity Statistics for England and Wales" 2023)), care status, socioeconomic deprivation (a standardised composite derived from principal component analysis of household poverty indicators; see Table 8), and self-reported neurodivergence. We did not adjust for concurrent anxiety or depressive disorder, which was assessed in the survey using the short Development and Well-Being Assessment (DAWBA). Because anxiety and depression are both consequences of the exposures and precursors of self-harm, conditioning on them could introduce collider or mediator bias, attenuating the very associations of interest; we therefore treat them as potential mediators rather than confounders. The unadjusted distribution is reported in Table 8 for transparency.

Because skip-logic in the survey directed different subsets of respondents to different exposure questions, the analytic sample varies across models. The dose–response model includes all respondents with valid self-harm classification and content-exposure data; the content-type models are restricted to those who answered the content-type battery; the mode-of-exposure models are restricted to those who reported having seen self-harm content; and the motivation model is further restricted to those who actively searched. Sample sizes are reported in each table.

Statistical analysis

To account for the nesting of students within schools, all associational models used multilevel logistic regression with school-level random intercepts, estimated via `glmmTMB` with Normal(0, 3) priors on fixed effects to aid convergence. We report odds ratios and marginal predicted probabilities. For families of related tests (11 content-type models, 5 mode-of-exposure models), p-values were adjusted using the Benjamini–Hochberg false-discovery-rate procedure; all individually reported associations remain significant after correction.

To move beyond association and provide quasi-experimental evidence on the direction of the relationship, we employed an instrumental variable (IV) approach. The key challenge is endogeneity: individuals who self-harm may be more likely to seek out or be algorithmically served self-harm content, giving rise to reverse causality that conventional regression cannot address. The instrument used is daily time spent on the internet (hours, derived from the survey item X2320), which plausibly increases exposure to online self-harm content but is assumed to have no direct effect on self-harm behaviour except through such exposure. To further minimise reverse causality, we restricted the self-harm outcome to respondents who reported their most recent self-harm as occurring "in the last week" or "in the last month", aligning temporally with the content exposure question (past month).

We fitted two IV specifications. In the first, content exposure is treated as a continuous variable (scored 0–30 to reflect monthly frequency). We estimated a two-stage least squares (2SLS) model using `AER::ivreg`, which implements a linear probability model (LPM) with correct standard errors and provides built-in diagnostics for weak instruments and endogeneity (Wu–Hausman test). For comparison, we also report naïve logit estimates from the same sample. In the second specification, content exposure is treated as a categorical variable (five levels). Because standard 2SLS is not directly applicable to a multi-category endogenous variable, we used a control-function approach: the first stage is a logistic regression of exposure on the instrument and covariates; the generalised residual from this stage is then included in the second-stage logit alongside the original exposure categories. A significant coefficient on the residual provides a direct test of endogeneity (Wooldridge 2015). Standard errors in the control-function approach are not automatically corrected for the two-step estimation; we report them as approximate and note that bootstrapped standard errors would be a useful extension. Neither IV specification includes school-level random effects; the associational models do.

The strength of the instrument is assessed using the first-stage F-statistic (or chi-squared statistic for the logistic first stage), with values above 10 generally taken to indicate a sufficiently strong instrument.

The causal structure motivating the IV analysis is represented in Figure 1. The instrument Z (daily time spent on the internet) affects the endogenous exposure X (frequency of self-harm content encountered) but is assumed to affect the outcome Y (self-harm) only through X. The vector $\mathbf{C}$ represents measured confounders (gender, year group, ethnicity, care status, deprivation, neurodivergence), while U represents unmeasured confounders that may jointly influence X and Y (e.g., underlying psychological distress).

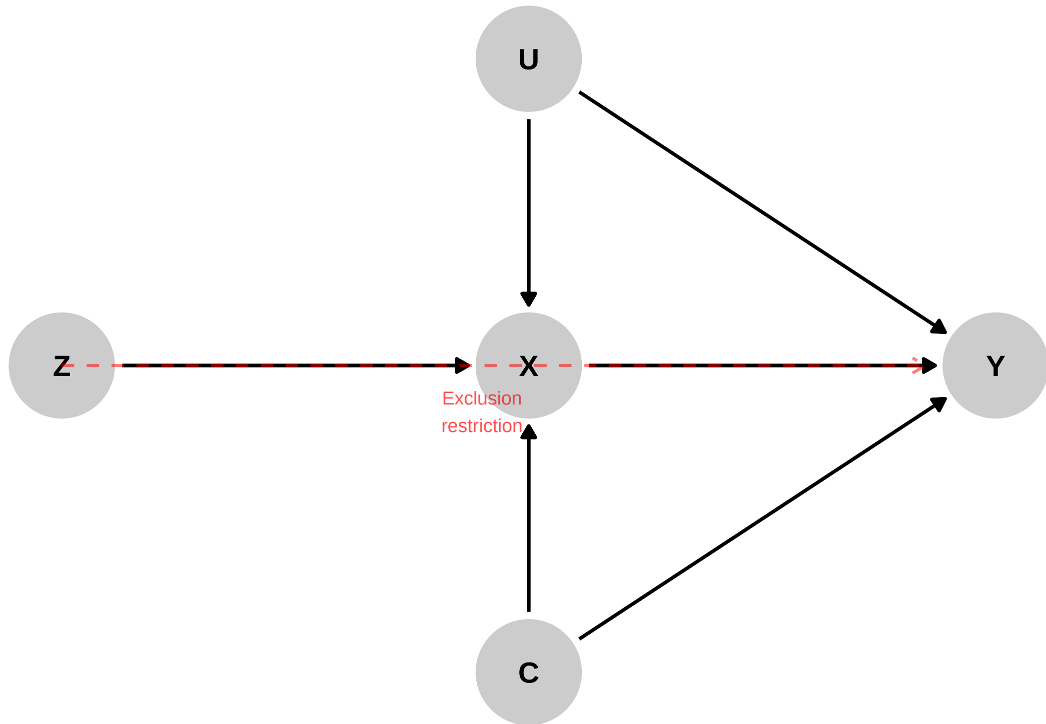

Figure 1: Directed acyclic graph (DAG) representing the instrumental variable assumptions. Z = daily time spent on the internet (instrument); X = frequency of self-harm content exposure (endogenous); Y = self-harm (outcome); C = measured confounders; U = unmeasured confounders. The dashed arrow from Z to Y represents the pathway that must be absent for the exclusion restriction to hold.

The IV approach requires three assumptions: (1) *relevance* — Z is associated with X (testable via the first-stage F-statistic); (2) *independence* — Z is independent of U; and (3) *exclusion restriction* — Z affects Y only through X. We discuss potential threats to assumptions (2) and (3) in the limitations.

Results

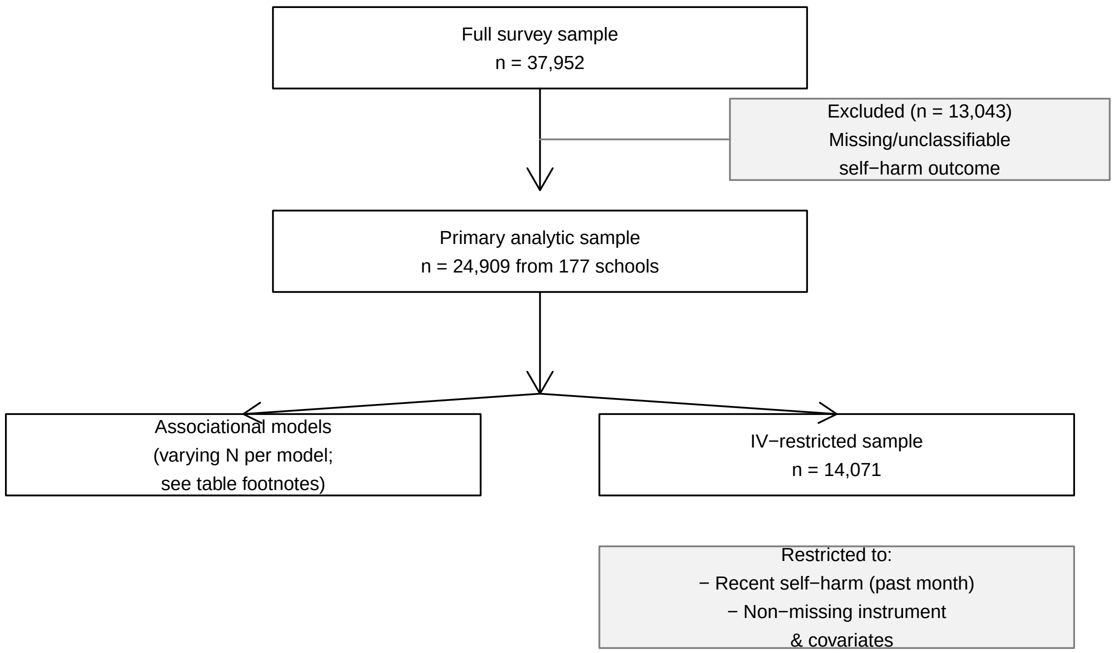

Figure 2: Participant flow diagram showing derivation of analytic samples.

Table 1 presents the characteristics of the sample stratified by self-harm classification. Among self-harmers, 100% had encountered self-harm content online at least once in the past month, with 10% reporting daily exposure. The most common mode of encountering this content was through social media feeds (0%), followed by accidental exposure (0%) and active searching (0%). Among those who reported seeking self-harm content, common motivations included finding support for themselves, connecting with others who had similar experiences, and looking up factual information.

Table 1: Descriptive characteristics by self-harm classification

Characteristic	N	Yes N = 4,479	No N = 20,430	(Missing) N = 13,043
Have you ever deliberately self-harmed?	33,737			
Yes		4,479 (100%)	0 (0%)	1,707 (19%)
No		0 (0%)	20,430 (100%)	0 (0%)
Not sure what this means		0 (0%)	0 (0%)	883 (10%)
Prefer not to say		0 (0%)	0 (0%)	4,164 (47%)
Skipped by respondent		0 (0%)	0 (0%)	2,074 (23%)
Missing		0	0	4,215
Year group	37,952			
Y05		214 (4.8%)	2,198 (11%)	1,977 (15%)
Y06		249 (5.6%)	2,317 (11%)	1,726 (13%)
Y07		579 (13%)	3,354 (16%)	2,347 (18%)
Y08		659 (15%)	3,219 (16%)	1,861 (14%)
Y09		728 (16%)	3,103 (15%)	1,630 (12%)
Y10		620 (14%)	2,227 (11%)	1,256 (9.6%)
Y11		686 (15%)	2,041 (10.0%)	1,085 (8.3%)
Y12		449 (10%)	1,192 (5.8%)	744 (5.7%)
Y13		295 (6.6%)	779 (3.8%)	417 (3.2%)
Gender	37,569			
Boy		1,225 (28%)	10,274 (51%)	5,829 (45%)
Girl		2,775 (63%)	9,574 (47%)	6,268 (49%)
Gender Diverse (GD)		197 (4.5%)	70 (0.3%)	143 (1.1%)
Gender Non-Disclosing (GND)		221 (5.0%)	370 (1.8%)	623 (4.8%)
Missing		61	142	180
Ethnic background	37,946			
Asian/Asian British (aggregated)		451 (10%)	3,156 (15%)	1,682 (13%)
Black/Black British/African/Caribbean (aggregated)		161 (3.6%)	991 (4.9%)	580 (4.4%)
Mixed/Multiple Ethnic Groups (aggregated)		298 (6.7%)	1,015 (5.0%)	753 (5.8%)
Other ethnic group		129 (2.9%)	849 (4.2%)	674 (5.2%)
White (aggregated)		2,711 (61%)	10,850 (53%)	6,715 (52%)
Skipped by respondent		729 (16%)	3,569 (17%)	2,633 (20%)
Missing		0	0	6
Child in care, looked after, or fostered?	37,938			
Yes		159 (3.5%)	646 (3.2%)	534 (4.1%)
No		3,482 (78%)	14,026 (69%)	7,650 (59%)
I don't know what this means		230 (5.1%)	985 (4.8%)	793 (6.1%)
Not now, but I used to be in care		80 (1.8%)	126 (0.6%)	128 (1.0%)
Prefer not to say		61 (1.4%)	106 (0.5%)	174 (1.3%)
Skipped by respondent		467 (10%)	4,541 (22%)	3,750 (29%)
Missing		0	0	14
Neurodivergent	37,938			
Yes		1,405 (31%)	2,225 (11%)	2,156 (17%)
No		1,469 (33%)	10,822 (53%)	4,498 (35%)
Not sure		1,061 (24%)	2,631 (13%)	2,313 (18%)
Prefer not to say		73 (1.6%)	197 (1.0%)	299 (2.3%)
Skipped by respondent		471 (11%)	4,555 (22%)	3,763 (29%)
Missing		0	0	14
Deprivation index	36,292	5.73 (3.44)	5.44 (3.39)	5.34 (3.45)

(continued)

Characteristic	N	Yes N = 4,479	No N = 20,430	(Missing) N = 13,043
Missing		351	763	546
Come across self-harm content online (last month)?	30,778			
Yes, daily		342 (8.4%)	137 (0.7%)	187 (2.3%)
Yes, weekly		301 (7.4%)	211 (1.1%)	150 (1.9%)
Yes, a few times		846 (21%)	1,107 (5.9%)	660 (8.2%)
Yes, once or twice		1,018 (25%)	2,753 (15%)	972 (12%)
No, never		982 (24%)	9,479 (51%)	1,751 (22%)
Skipped by respondent		569 (14%)	4,950 (27%)	4,363 (54%)
Missing		421	1,793	4,960
How encountered: I searched for it	20,896			
Yes		511 (15%)	95 (0.7%)	138 (3.7%)
Skipped by respondent		2,978 (85%)	13,592 (99%)	3,582 (96%)
Missing		990	6,743	9,323
How encountered: It came up on my feed	20,896			
Yes		1,497 (43%)	2,147 (16%)	1,024 (28%)
Skipped by respondent		1,992 (57%)	11,540 (84%)	2,696 (72%)
Missing		990	6,743	9,323
How encountered: Accidental exposure	20,896			
Yes		862 (25%)	1,971 (14%)	772 (21%)
Skipped by respondent		2,627 (75%)	11,716 (86%)	2,948 (79%)
Missing		990	6,743	9,323
How encountered: Someone else shared with me	20,896			
Yes		205 (5.9%)	319 (2.3%)	157 (4.2%)
Skipped by respondent		3,284 (94%)	13,368 (98%)	3,563 (96%)
Missing		990	6,743	9,323
How encountered: Other	20,896			
Yes		123 (3.5%)	275 (2.0%)	202 (5.4%)
Skipped by respondent		3,366 (96%)	13,412 (98%)	3,518 (95%)
Missing		990	6,743	9,323
Why were you looking for self-harm related content?	30,012			
Find support for yourself		148 (6.0%)	16 (<0.1%)	36 (0.3%)
Find support for someone you know		16 (0.6%)	8 (<0.1%)	10 (<0.1%)
To connect with others who have had similar experiences		125 (5.0%)	6 (<0.1%)	20 (0.2%)
To look up facts and figures about self-harm		52 (2.1%)	19 (0.1%)	19 (0.2%)
Another reason		150 (6.0%)	37 (0.2%)	39 (0.3%)
Skipped by respondent		1,992 (80%)	16,231 (99%)	11,088 (99%)
Missing		1,996	4,113	1,831

[1] n (%); Mean (SD)

Figure 3 shows exposure frequency by self-harm status. Among non-self-harmers, 0% reported no exposure in the past month; among self-harmers, this fell to just 0%. At the other extreme, 10% of self-harmers reported daily exposure.

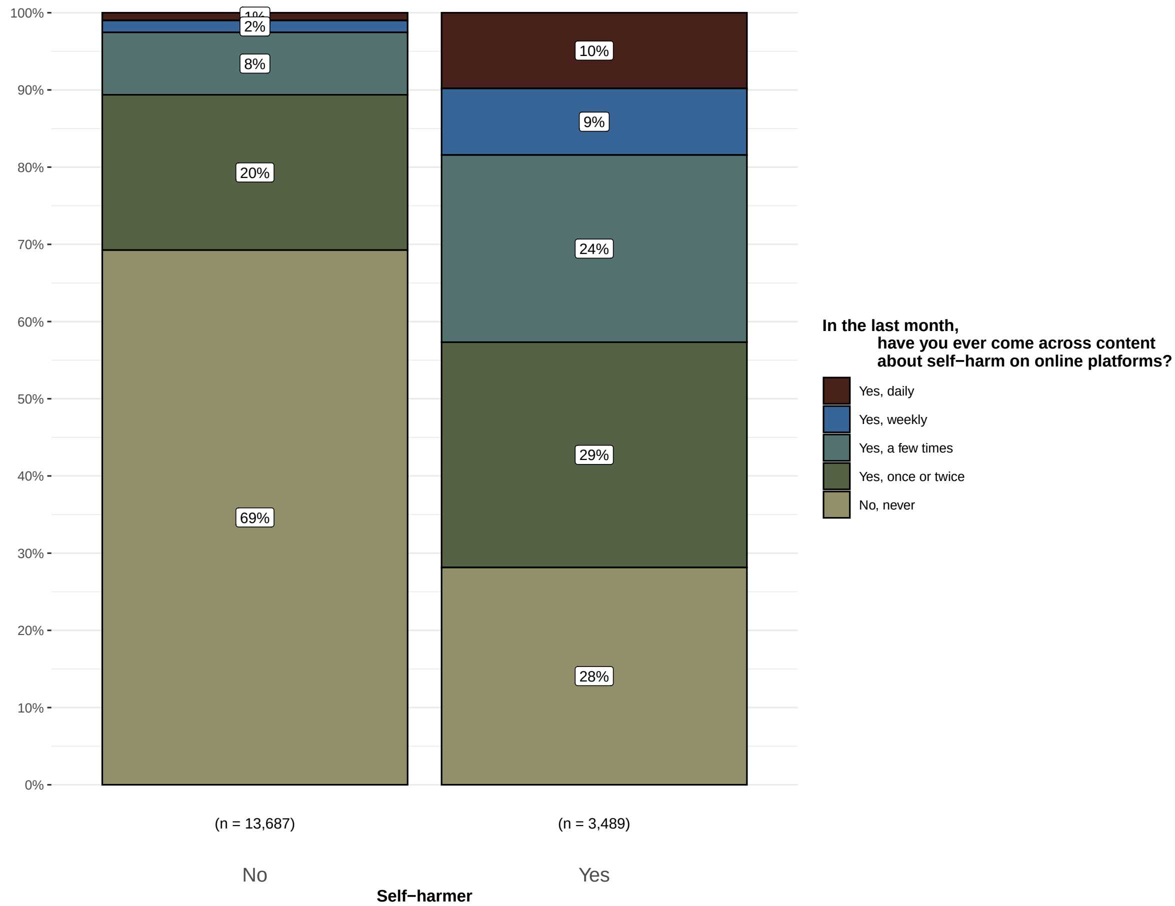

Figure 3: Frequency of encountering online self-harm content by self-harm classification

Figure 4 presents responses on how participants who had viewed self-harm content came across such material. Panel I shows that 0% of those who self-harmed actively searched for self-harm content, compared with only 0% of those who did not. Panel II reveals that approximately 0% of self-harmers had self-harm content appear in their feeds, compared with 0% of non-self-harmers.

Panel III indicates that 0% of self-harmers accidentally encountered self-harm content, versus 0% of non-self-harmers. Panel IV shows no statistically significant association between having self-harm content shared by someone else and self-harm classification.

```
          used  (Mb) gc trigger   (Mb)  max used   (Mb)
Ncells   3904929 208.6    6842008  365.5   5614890  299.9
Vcells 110897943 846.1  170941395 1304.2 170915200 1304.0
```

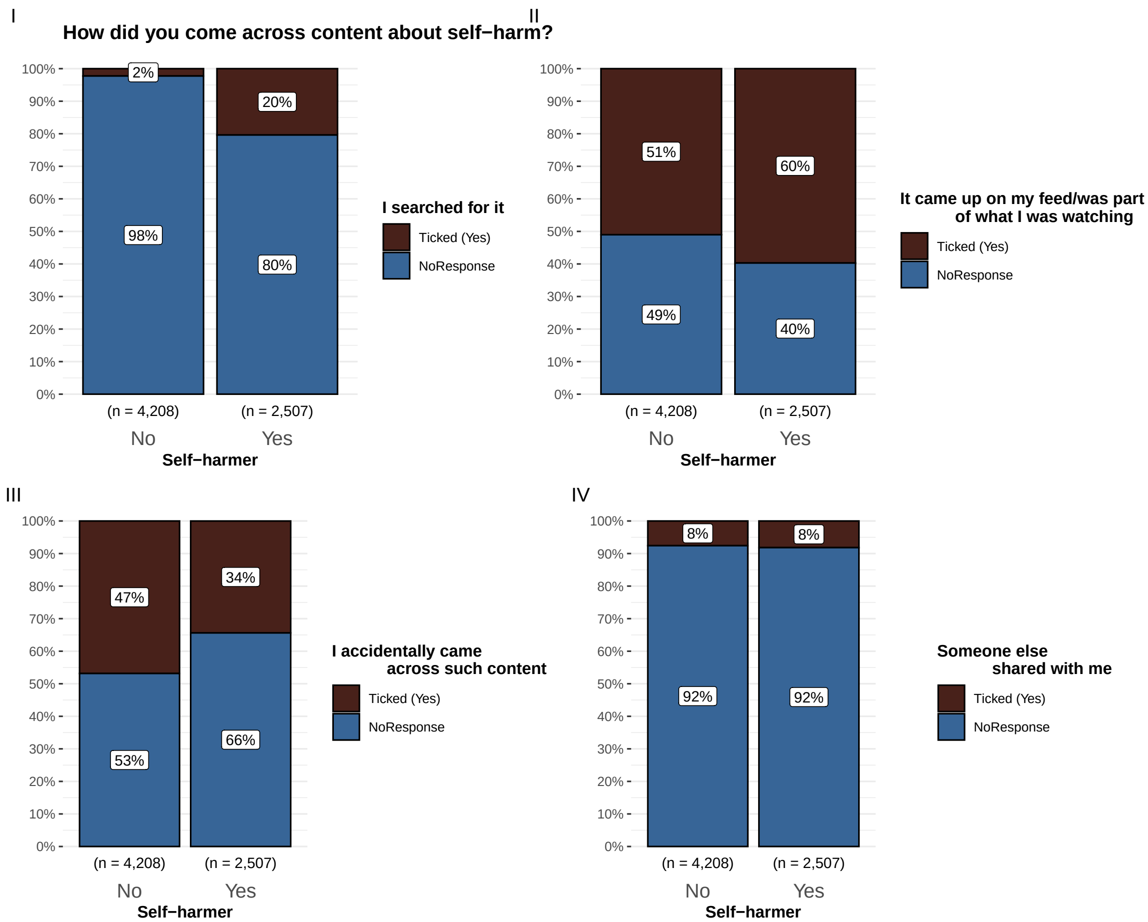

Figure 4: Mode of encountering self-harm content by self-harm classification

Odds of self-harm rose consistently with exposure frequency in the adjusted multilevel models (Table 2). Daily exposure carried the highest adjusted odds ratio (OR 17.9, $p < 0 \cdot 001$).

Figure 5 displays the predicted probabilities. The model yields a clear dose–response pattern: respondents who reported no exposure to self-harm content had an estimated probability of self-harm of approximately %, rising to around 67% for those reporting daily or weekly exposure.

Table 2: Adjusted odds ratios for self-harm by frequency of online self-harm content exposure

	(1)
Self-harm content exposure (Once or twice)	3.005***
	(0.169)
	(<0.001)
Self-harm content exposure (A few times)	5.866***
	(0.378)
	(<0.001)
Self-harm content exposure (Weekly)	10.466***
	(1.142)
	(<0.001)
Self-harm content exposure (Daily)	17.902***
	(2.204)
	(<0.001)
Num.Obs.	16 062

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Note: Model adjusted for gender, year group, ethnicity, care status, deprivation, and neurodivergence.

Reference category: No exposure in the past month.

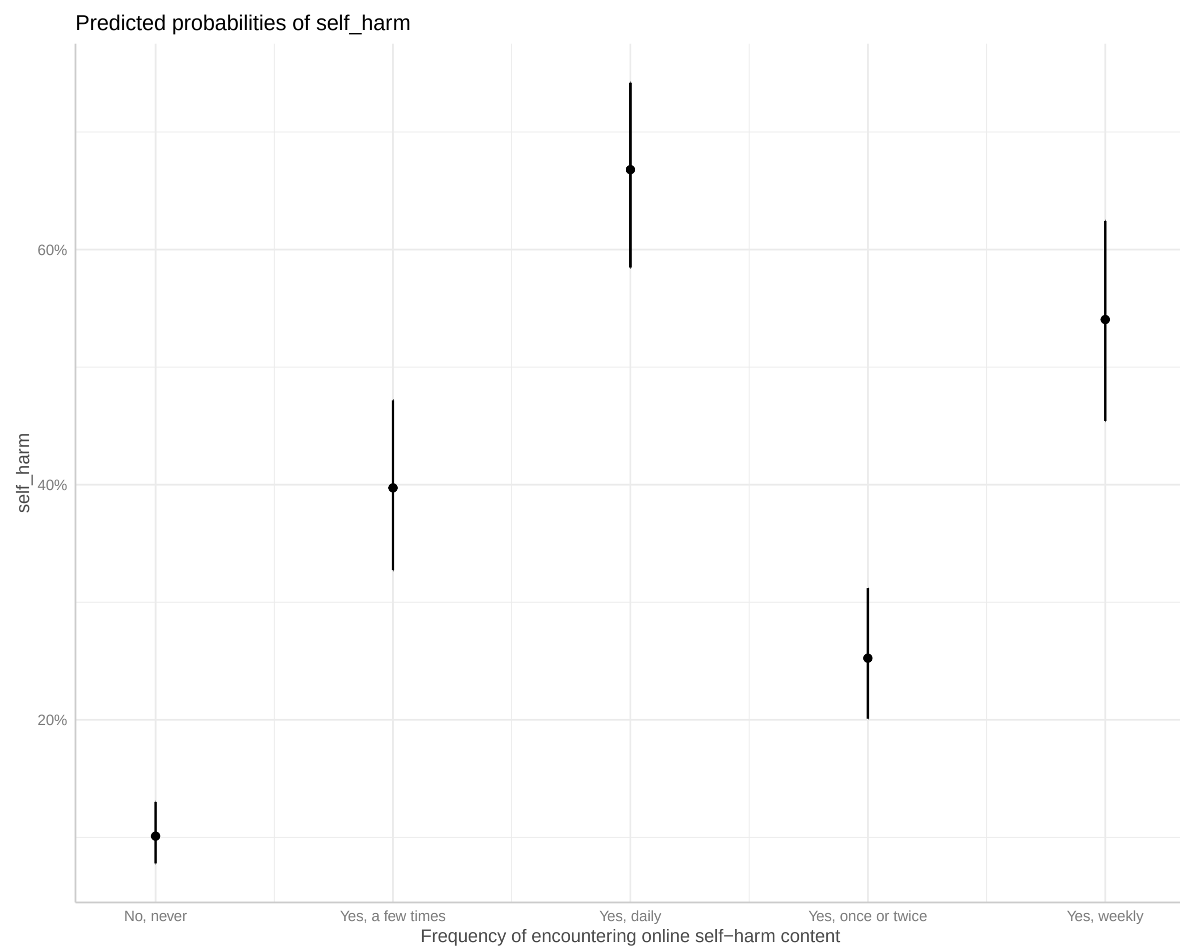

Figure 5: Predicted probabilities of self-harm by frequency of online self-harm content exposure. Error bars represent 95% confidence intervals.

Table 3: Adjusted odds ratios for self-harm by type of online content viewed

	Self-harm (Yes)			Self-harm		
	OR	std.error	pvalue	OR	std.error	pvalue
Family and friends	1.000	0.044	0.998			
Entertainment (e.g. film, TV, music) & Humour				0.772	0.044	<0.001
Education				0.917	0.044	0.069
Forums	1.759	0.157	<0.001			
Lifestyle	1.266	0.057	<0.001			
Health and fitness & Support and wellness	1.494	0.066	<0.001			
Beauty and fashion	1.530	0.074	<0.001			
Celebrities/Influencers	1.245	0.055	<0.001			
Other	0.896	0.047	0.036			

Note: Each row represents a separate model. All models adjusted for gender, year group, ethnicity, care status, deprivation, and neurodivergence.
Num. Obs. = 20487

Models for each of 11 content categories showed clear variation in the strength of association (Table 3). Viewing Forums.X2330dTicked (Yes) and Beauty and fashion.X2330hTicked (Yes) content was associated with the highest odds (OR 1.8 and 1.5, $p < 0 \cdot 001$), while Entertainment (e.g. film, TV, music) & Humour.X2330biNoResponse content was associated with lower odds (OR 0.77, $p < 0 \cdot 001$). Because the exposure question asks about "content about self-harm" without distinguishing format or intent — ranging from graphic imagery to recovery stories to clinical fact sheets — these associations should be interpreted as reflecting engagement with broad content categories rather than with specific harmful material.

Figure 6 presents the predicted probabilities, ranging from approximately 15% to 26% across content categories.

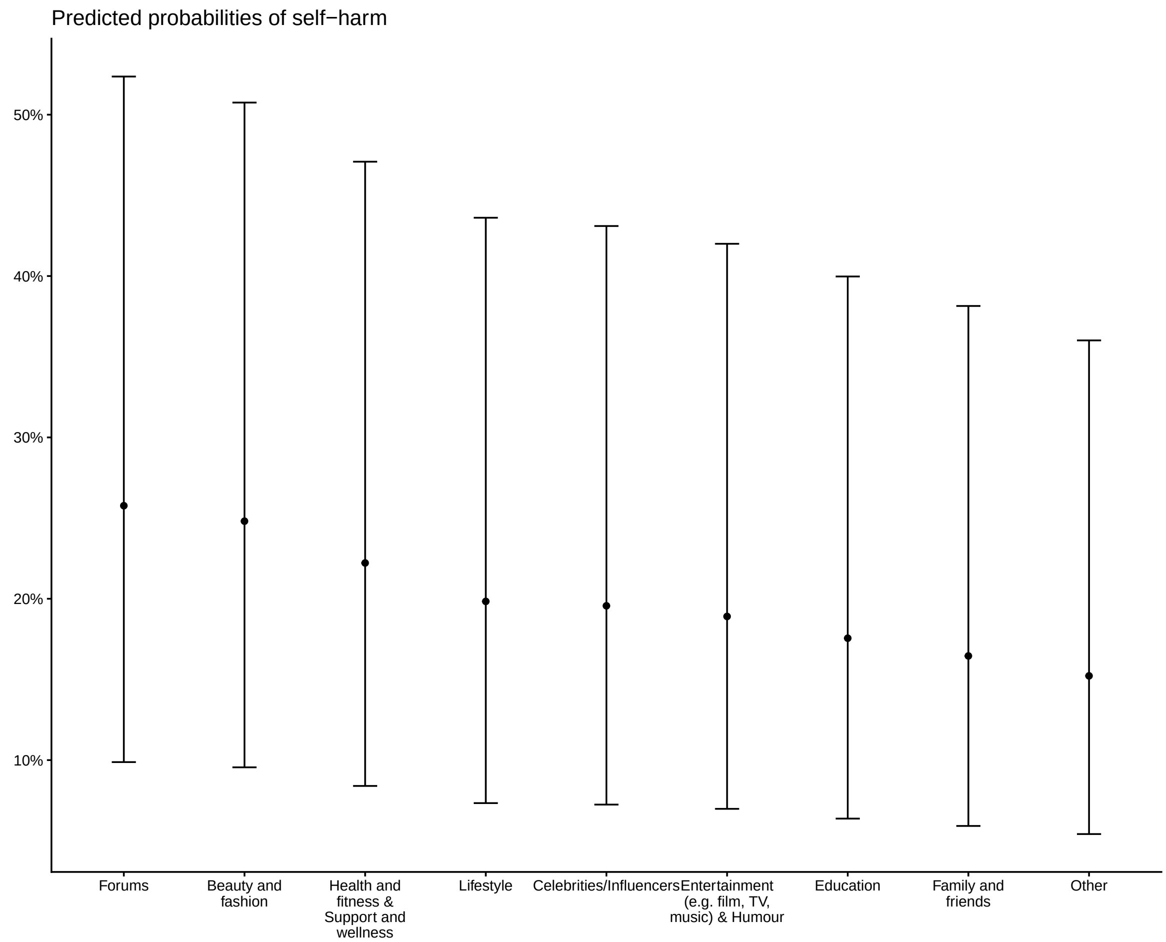

Figure 6: Predicted probabilities of self-harm by type of online content viewed. Error bars represent 95% confidence intervals.

The mode of exposure mattered considerably (Table 4). Active searching carried by far the highest odds (OR 11.4, $p < 0 \cdot 001$), followed by accidental exposure. Content appearing in the feed was associated with slightly lower odds (OR 0.77, $p < 0 \cdot 001$) — passive algorithmic delivery does not appear to carry the same strength of association as deliberate engagement.

Figure 7 displays the predicted probabilities. Those who actively searched for self-harm content had a predicted probability of self-harm of approximately NaN%.

Table 4: Adjusted odds ratios for self-harm by mode of encountering self-harm content

	Self-harm (Yes)			Self-harm		
	OR	std.error	pvalue	OR	std.error	pvalue
I searched for it	11.361	1.486	<0.001			
It came up on my feed/was part of what I was watching				0.765	0.046	<0.001
I accidentally came across such content				1.643	0.099	<0.001
Someone else shared it with me	1.057	0.117	0.615			
Other	0.774	0.099	0.045			

Note: Each row represents a separate model. All models adjusted for gender, year group, ethnicity, care status, deprivation, and neurodivergence. Num. Obs. = 8048

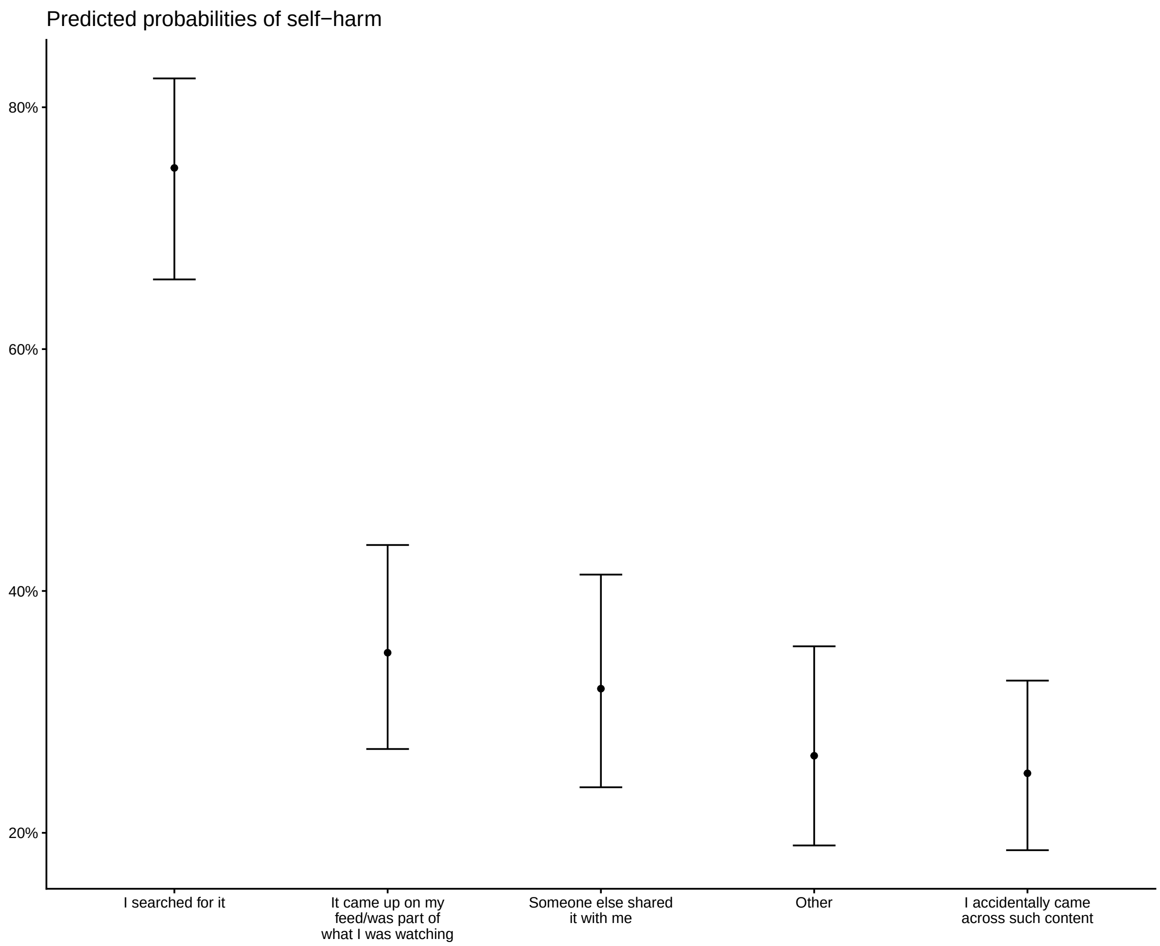

Figure 7: Predicted probabilities of self-harm by mode of encountering self-harm content. Error bars represent 95% confidence intervals.

Among those actively seeking self-harm content, the stated motivations were strongly predictive. Seeking content "to find support for yourself" and "to connect with others who have had

Table 5: Adjusted odds ratios for self-harm by motivation for seeking self-harm content

	(1)
Find support for someone you know	1.473
	(0.999)
	(0.568)
Find support for yourself	2.526*
	(1.110)
	(0.035)
To connect with others who have had similar experiences	4.951**
	(2.662)
	(0.003)
To look up facts and figures about self-harm	0.664
	(0.288)
	(0.345)
Num.Obs.	507

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

similar experiences" were each associated with markedly elevated odds (OR 2.5 and 5, respectively; Table 5). Predicted probabilities for both motivations approached 56%–72% (Figure 8), among the highest observed across all models.

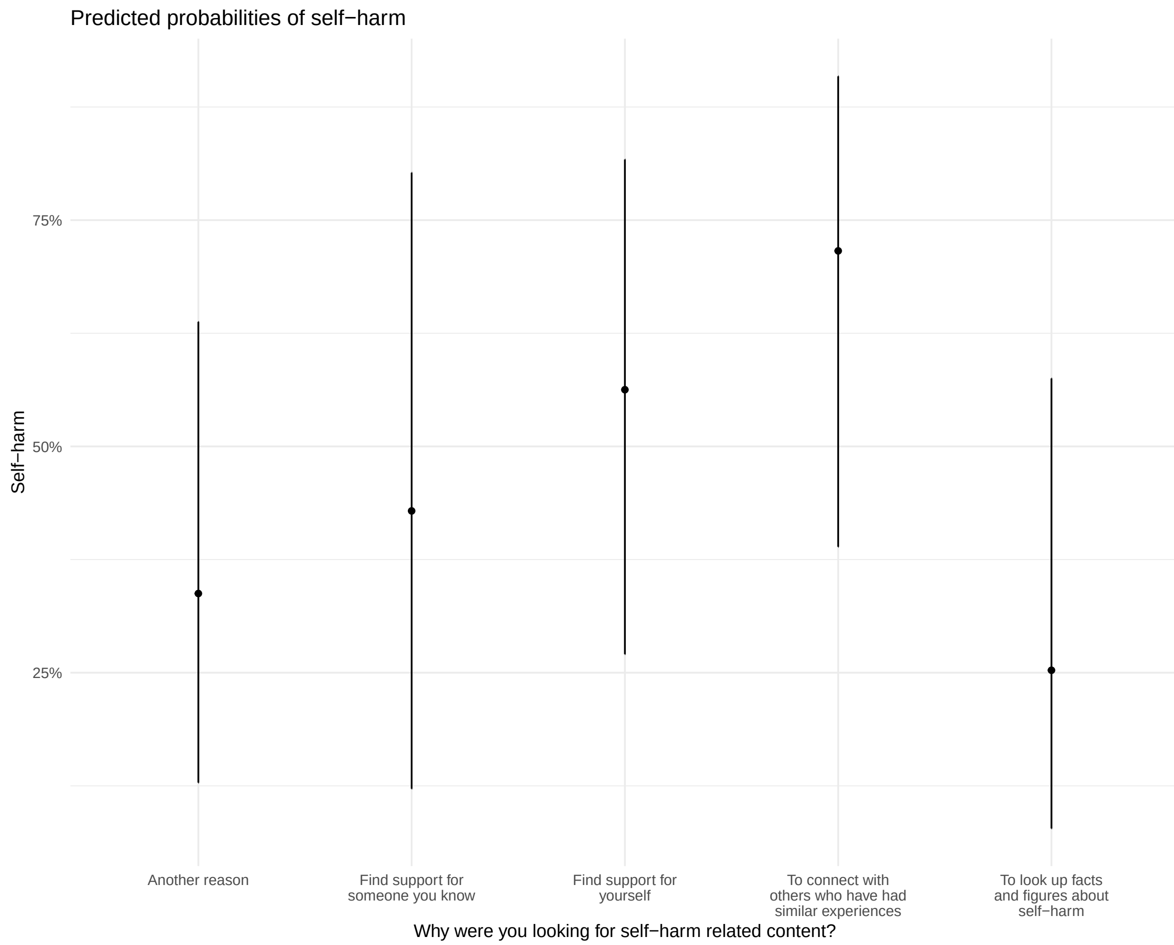

Figure 8: Predicted probabilities of self-harm by motivation for seeking self-harm content. Error bars represent 95% confidence intervals.

The associations above are consistent and strong, but cross-sectional data cannot rule out reverse causation: young people who self-harm may be more likely to seek or encounter this content. The IV approach addresses this directly.

Table 6 presents results from two sets of IV models. In Models 1–3, self-harm content exposure is treated as a continuous variable. The first-stage regression confirms that the instrument (daily time spent on the internet) is a strong predictor of self-harm content exposure (F-statistic = 438, well above the conventional threshold of 10). The 2SLS estimate from `ivreg` — a linear probability model (LPM) with correct standard errors — indicates that a one-unit increase in self-harm content exposure is associated with a 5.8 percentage-point increase in the probability of self-harm (95% CI 5.1–6.5 pp, $p < 0 \cdot 001$). This is a substantial effect: over the observed range of the exposure variable, it implies meaningfully elevated risk. The Wu–Hausman test for endogeneity was significant ($p < 0 \cdot 001$), confirming that the naïve logit estimates are inconsistent and that the IV approach is warranted. Because the LPM coefficients are on a different scale from the logistic models, direct comparison of coefficient magnitudes across Models 1–3 and 4–6 requires caution.

In Models 4–6, exposure is treated as a categorical variable using a control-function approach. The first-stage chi-squared statistic confirms instrument strength (^{2} = 498). The control-function residual is statistically significant (p $< 0 \cdot 001$), confirming endogeneity and justifying the control-function correction. The second-stage estimate yields an odds ratio of 54.9 (p $< 0 \cdot 001$). Standard errors for this specification are approximate; bootstrapped standard errors would be a useful extension. Both specifications — using different functional forms and different definitions of the endogenous variable — converge on the same substantive conclusion: the association between content exposure and self-harm is not attributable entirely to reverse causation.

The ICC for the school-level random intercept was 0.017, confirming meaningful between-school variation.

For the associational models, we computed E-values (VanderWeele and Ding 2017) to quantify the minimum strength of association that an unmeasured confounder would need to have with both the exposure and the outcome to fully explain away the observed association. For the dose–response model at the daily exposure level (aOR 17.9), the E-value is 35.3, indicating very high robustness. For active searching (aOR 11.4), the E-value is 22.3. These values suggest that unmeasured confounding alone is unlikely to account for the largest observed associations, though it may attenuate smaller effects.

The key untestable assumption of the IV analysis is the exclusion restriction: that daily time on the internet affects self-harm only through exposure to self-harm content. Plausible violations include pathways through sleep disruption, displacement of face-to-face social interaction, cyberbullying, and social comparison. If such pathways exist, the IV estimate would capture a composite effect rather than the isolated effect of content exposure.

To quantify the sensitivity of the IV estimate to violations of the exclusion restriction, we implemented the union-of-confidence-intervals (UCI) approach of Conley, Hansen, and Rossi (2012). This method allows the instrument to have a direct effect γ on the outcome (i.e., $Z \to Y$ with coefficient γ), and re-estimates the IV model for a range of plausible values of γ. Figure 9 displays how the IV estimate for content exposure changes as γ increases from 0 (perfect exclusion restriction) to $0 \cdot 05$ (a direct effect of 5 percentage points per hour of internet use on the probability of self-harm). The IV estimate remains positive up to $\gamma =$ 0.025, meaning that the instrument would need a direct effect of at least 0.025 probability units per hour on self-harm — bypassing content exposure entirely — to explain away the observed IV estimate. The UCI across all examined values of γ spans [-0.0807, 0.0654] on the coefficient scale.

Table 6: Instrumental variable regression models estimating the effect of self-harm content exposure on self-harm

	Naïve logit	First stage (OLS)	2SLS (ivreg)	Naïve logit	First stage (logit)	Control function
Online self-harm exposure (continuous)	1.124***		0.058***			
	(0.006)		(0.004)			
Time spent on the internet (Instrument)		0.388***			1.225***	
		(0.019)			(0.011)	
Self-harm content (Once or twice)				3.649***		54.896***
				(0.324)		(17.431)
Self-harm content (A few times)				7.925***		113.166***
				(0.763)		(35.726)
Self-harm content (Daily)				37.623***		466.734***
				(5.848)		(153.005)
Self-harm content (Weekly)				18.397***		253.470***
				(2.665)		(84.162)
Control-function residual						0.061***
						(0.019)
Num.Obs.	11 659	11 659	11 659	11 659	11 659	11 659
R2		0.071	−0.099			
R2 Adj.		0.069	−0.101			
AIC	6339.1	67 823.0	7226.6	5827.7		5750.9
BIC	6508.4	67 999.8	7403.3	6019.1		5949.8
Log.Lik.	−3146.525	−33 887.521		−2887.839		−2848.470
F	58.183	40.417		64.478	53.089	62.966
RMSE	0.28	4.43	0.33	0.27	0.45	0.27

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: All models adjusted for gender, year group, ethnicity, care status, deprivation, and neurodivergence.

Instrument: daily time spent on the internet (X2320).

Models 1–3: continuous exposure; 2SLS (ivreg) reports LPM coefficients (probability changes).

Models 4–6: categorical exposure, control-function approach; coefficients exponentiated (ORs).

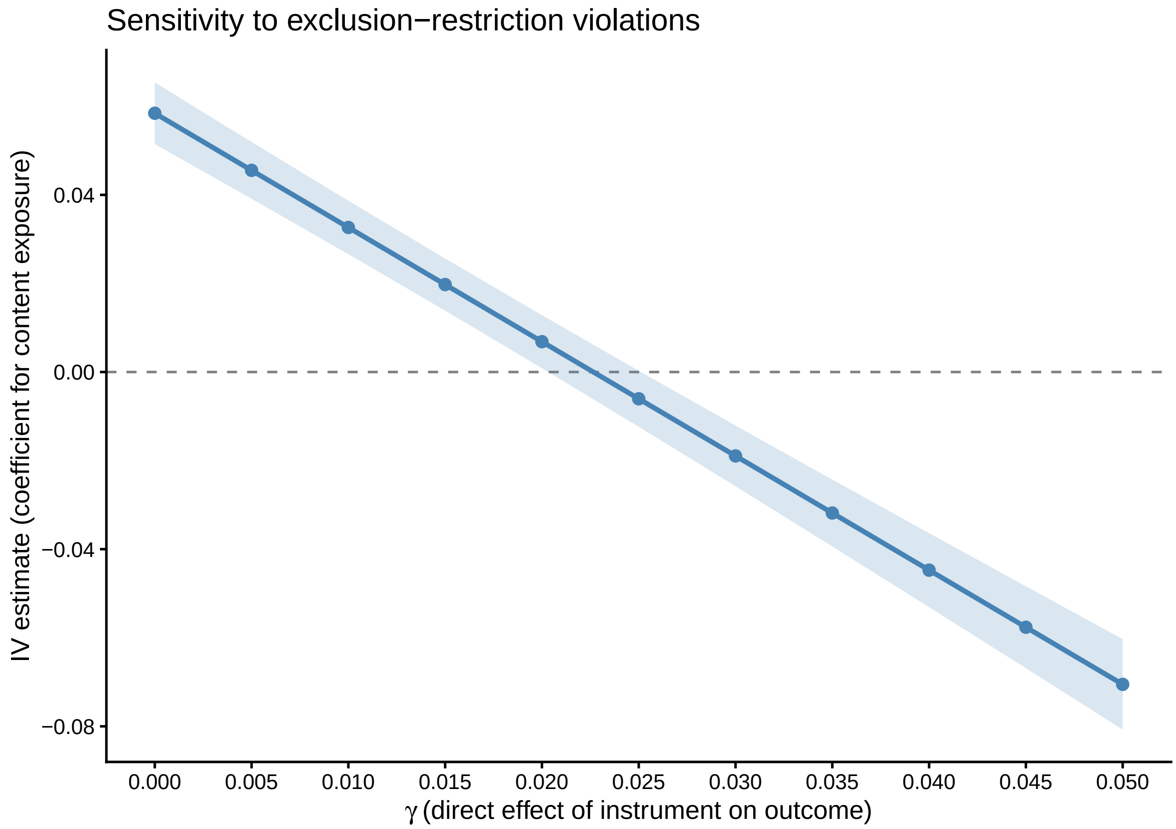

Figure 9: Conley et al. (2012) plausibly-exogenous sensitivity analysis. The IV estimate for the effect of self-harm content exposure on self-harm is plotted against gamma, the assumed direct effect of the instrument (daily internet time) on self-harm. Error bars represent 95% confidence intervals. The dashed line indicates zero effect.

Discussion

This study has two main findings. First, exposure to self-harm content online is strongly and consistently associated with self-harm, with the mode of exposure mattering considerably. Second, instrumental variable analysis provides quasi-experimental evidence that this relationship has a causal component and is not merely a reflection of at-risk young people being more likely to encounter such content.

Active searching for self-harm content showed by far the strongest association with self-harm, with an odds ratio of 11.4 and predicted probabilities approaching NaN%. Accidental exposure and content appearing through algorithmic feeds showed significant but smaller associations, while interpersonal sharing showed no meaningful relationship. These distinctions matter for intervention. Content moderation and algorithmic curation can reduce accidental and passive exposure, but the strong association with active searching indicates that many at-risk young people are deliberately seeking this material, probably as part of coping. Restricting access alone is unlikely to be sufficient for this group; therapeutic support and alternatives to online coping are more likely to help. Among those who actively searched for self-harm content, the most strongly associated motivations were seeking personal support and connecting with others who shared similar experiences, both associated with very high predicted probabilities of self-harm (approximately 56%). Online self-harm communities may offer a sense of belonging, but they also appear to reinforce and normalise the behaviour. Building platforms that meet the genuine need for peer connection while reducing exposure to harmful content is a realistic target for digital health intervention.

The IV analysis is the most novel contribution of this study. Both specifications yielded positive and statistically significant effect estimates consistent with a causal interpretation. The two IV models use different scales — the continuous specification (LPM via ivreg) reports probability changes while the categorical specification (control-function logit) reports log-odds — making direct comparison of magnitudes uninformative. However, both confirm the same direction of effect, supporting the conclusion that the observed associations are not driven entirely by reverse causation. Daily time spent on the internet is a plausible instrument: more time online increases the probability of encountering self-harm content. However, the exclusion restriction — that time online affects self-harm *only* through content exposure — cannot be tested empirically. If internet time affects self-harm through other pathways (sleep disruption, cyberbullying, social comparison), the IV estimates capture a composite effect rather than the isolated effect of content exposure. The Conley et al. sensitivity analysis reported in the Results quantifies the robustness of the IV estimate to such violations.

These findings have regulatory and clinical implications, though they should be interpreted with the caveats discussed below. Policy debates around the UK Online Safety Act, the EU Digital Services Act, and similar US legislation have proceeded largely on associational grounds. The present quasi-experimental findings suggest that reducing exposure to self-harm content online could contribute to reducing self-harm, not merely correlate with it. Content removal, algorithmic de-amplification, and interstitial warnings on self-harm search queries are all supported by these data. Active searching, however, presents a different clinical problem. Young people who deliberately seek self-harm content are likely doing so within a broader pattern of self-harm and distress-related coping. Content suppression alone will not reach them. Clinical assessment should routinely ask about online self-harm content engagement, and therapeutic work should address the motivations behind it — particularly the desire for peer connection and support. The help-seeking motivations identified here represent both a risk and a point of intervention. The need is real; current online environments meet it poorly. Digital services that redirect young people toward evidence-based resources — crisis lines, moderated peer support, NHS-endorsed tools — could address the underlying need while reducing exposure to harmful material.

This study has several strengths: a large school-based sample (n > 24,000), a validated combined self-harm classification, multilevel modelling with school-level random effects, simultaneous examination of multiple exposure dimensions, and — most notably — an instrumental variable analysis that provides quasi-experimental evidence in a field dominated by cross-sectional associations.

The main limitations are as follows. First, the exclusion restriction — the key untestable assumption of the IV analysis — cannot be empirically verified. If daily internet time affects self-harm through pathways other than content exposure (sleep disruption, cyberbullying, social comparison, displacement of face-to-face interaction), the IV estimates capture a composite effect of internet time on self-harm rather than the isolated causal effect of content exposure. The sensitivity analysis comparing IV and logit estimates across specifications provides partial reassurance about the direction of the effect, and the Conley et al. (2012) plausibly-exogenous analysis reported above quantifies the degree of exclusion-restriction violation required to nullify the IV estimate, but neither fully eliminates this concern.

Second, although the survey assessed concurrent anxiety and depressive disorder (via the DAWBA), we did not include this variable as a covariate. Anxiety and depression are plausibly both consequences of social media exposure and precursors of self-harm; conditioning on them could introduce collider or mediator bias that would attenuate the exposure–outcome associations. We therefore treat them as potential mediators rather than confounders. However, this means that the associational models may partly reflect indirect pathways through internalising symptoms, and the estimates should be interpreted as total rather than direct effects. Third, all measures are cross-sectional and self-reported, introducing potential recall bias and temporal ambiguity. Restricting the IV analysis to recent self-harm (past month) reduces the sample and may introduce selection bias toward more severe or persistent cases; a comparison of the full and IV-restricted samples would be informative. Fourth, the exposure question asks about "content about self-harm" without distinguishing format or intent — clinical fact sheets, recovery stories, and graphic imagery are conflated. The associations reported here therefore relate to broad engagement with self-harm-related content rather than to specific harmful material. The survey covers four English regions and may not generalise to other contexts. Finally, school-based sampling may miss young people whose mental health prevents attendance, potentially underestimating prevalence and associations in the most severely affected group.

Exposure to online self-harm content is strongly associated with self-harm among secondary school students, and the mode of exposure matters: active searching shows substantially stronger associations than passive or accidental exposure. Instrumental variable analysis provides quasi-experimental evidence that this relationship has a causal component and is not driven solely by the tendency of those who self-harm to seek out such content, though the strength of this inference depends on the validity of the exclusion restriction. These findings support platform-level and regulatory interventions targeting the availability and algorithmic distribution of self-harm content, while also pointing to the need for clinical support for young people who actively seek this material.

References

Biernesser, Candice, Craig J. R. Sewall, David Brent, Todd Bear, Christina Mair, and Jeanette Trauth. 2020. "Social Media Use and Deliberate Self-Harm Among Youth: A Systematized Narrative Review." *Children and Youth Services Review* 116 (September): 105054. https://doi.org/10.1016/j.childyouth.2020.105054.

Conley, Timothy G., Christian B. Hansen, and Peter E. Rossi. 2012. "Plausibly Exogenous." *Review of Economics and Statistics* 94 (1): 260–72. https://doi.org/10.1162/REST_a_00139.

Cunningham, Simone, Chloe C. Hudson, and Kate Harkness. 2021. "Social Media and Depression Symptoms: A Meta-Analysis." *Research on Child and Adolescent Psychopathology* 49 (2): 241–53. https://doi.org/10.1007/s10802-020-00715-7.

"Developing Admin-Based Ethnicity Statistics for England and Wales." 2023. https://www.ons.gov.uk/peoplepopulationandcommunity/culturalidentity/ethnicity/articles/developingadminbasedethnicitystatisticsforenglandandwales/2020#ethnic-breakdown.

Dyson, Michele P., Lisa Hartling, Jocelyn Shulhan, Annabritt Chisholm, Andrea Milne, Purnima Sundar, Shannon D. Scott, and Amanda S. Newton. 2016. "A Systematic Review of Social Media Use to Discuss and View Deliberate Self-Harm Acts." *PLOS ONE* 11 (5): e0155813. https://doi.org/10.1371/journal.pone.0155813.

Geulayov, Galit, Rohan Borschmann, Karen L. Mansfield, Keith Hawton, Paul Moran, and Mina Fazel. 2022. "Utilization and Acceptability of Formal and Informal Support for Adolescents Following Self-Harm Before and During the First COVID-19 Lockdown: Results from a Large-Scale English Schools Survey." *Frontiers in Psychiatry* 13 (June). https://doi.org/10.3389/fpsyt.2022.881248.

Giumetti, Gary W., and Robin M. Kowalski. 2022. "Cyberbullying via Social Media and Well-Being." *Current Opinion in Psychology* 45 (June): 101314. https://doi.org/10.1016/j.copsyc.2022.101314.

Griffin, Eve, Elaine McMahon, Fiona McNicholas, Paul Corcoran, Ivan J. Perry, and Ella Arensman. 2018. "Increasing Rates of Self-Harm Among Children, Adolescents and Young Adults: A 10-Year National Registry Study 2007–2016." *Social Psychiatry and Psychiatric Epidemiology* 53 (7): 663–71. https://doi.org/10.1007/s00127-018-1522-1.

Hamm, Michele P., Amanda S. Newton, Annabritt Chisholm, Jocelyn Shulhan, Andrea Milne, Purnima Sundar, Heather Ennis, Shannon D. Scott, and Lisa Hartling. 2015. "Prevalence and Effect of Cyberbullying on Children and Young People: A Scoping Review of Social Media Studies." *JAMA Pediatrics* 169 (8): 770–77. https://doi.org/10.1001/jamapediatrics.2015.0944.

Mansfield, Karen Laura, Stephen Puntis, Emma Soneson, Andrea Cipriani, Galit Geulayov, and Mina Fazel. 2021. "Study Protocol: The OxWell School Survey Investigating Social, Emotional and Behavioural Factors Associated with Mental Health and Well-Being." *BMJ Open* 11 (12): e052717. https://doi.org/10.1136/bmjopen-2021-052717.

Morgan, Catharine, Roger T. Webb, Matthew J. Carr, Evangelos Kontopantelis, Jonathan Green, Carolyn A. Chew-Graham, Nav Kapur, and Darren M. Ashcroft. 2017. "Incidence, Clinical Management, and Mortality Risk Following Self Harm Among Children and Adolescents: Cohort Study in Primary Care." *BMJ* 359 (October): j4351. https://doi.org/10.1136/bmj.j4351.

Nesi, Jacqueline, Taylor A. Burke, Alexandra H. Bettis, Anastacia Y. Kudinova, Elizabeth C. Thompson, Heather A. MacPherson, Kyle A. Fox, et al. 2021. "Social Media Use and Self-Injurious Thoughts and Behaviors: A Systematic Review and Meta-Analysis." *Clinical Psychology Review* 87 (October): 102038. https://doi.org/10.1016/j.cpr.2021.102038.

O'Day, Emily B., and Richard G. Heimberg. 2021. "Social Media Use, Social Anxiety, and Loneliness: A Systematic Review." *Computers in Human Behavior Reports* 3 (January): 100070. https://doi.org/10.1016/j.chbr.2021.100070.

Odgers, Candice L., and Michaeline R. Jensen. 2020. "Annual Research Review: Adolescent Mental Health in the Digital Age: Facts, Fears, and Future Directions." *Journal of Child Psychology and Psychiatry* 61 (3): 336–48. https://doi.org/10.1111/jcpp.13190.

Sempé, Lucas, Mina Fazel, Emma Soneson, and Simona Skripkauskaite. 2024. "Exposure to Self-Harm Content Online and Self-Harm Behaviour in Secondary School Pupils," April. https://doi.org/10.17605/OSF.IO/5UXMJ.

Skripkauskaite, Simona, and Mina Fazel. 2022. "Time Spent Gaming, Device Type, Addiction Scores, and Well-Being of Adolescent English Gamers in the 2021 OxWell Survey: Latent Profile Analysis." *JMIR Pediatrics and Parenting* 5 (4): e41480. https://doi.org/10.2196/41480.

Valkenburg, Patti M., Adrian Meier, and Ine Beyens. 2022. "Social Media Use and Its Impact on Adolescent Mental Health: An Umbrella Review of the Evidence." *Current Opinion in Psychology* 44 (April): 58–68. https://doi.org/10.1016/j.copsyc.2021.08.017.

VanderWeele, Tyler J., and Peng Ding. 2017. "Sensitivity Analysis in Observational Research: Introducing the E-Value." *Annals of Internal Medicine* 167 (4): 268–74. https://doi.org/10.7326/M16-2607.

Wooldridge, Jeffrey M. 2015. *Introductory Econometrics: A Modern Approach*. 6th ed. Boston, MA: Cengage Learning.

Appendix

Age-group interactions

Table 7 presents models examining the interaction between age group and the mode through which self-harm content was encountered. These models test whether the association between content exposure pathways and self-harm varies across different year groups. The results do not indicate statistically significant interaction effects, suggesting that the associations reported in the main analyses are broadly consistent across the age range covered by the survey (Years 7–13, ages approximately 11–19). This is noteworthy because it runs counter to the common assumption that younger adolescents are disproportionately vulnerable to the effects of online self-harm content; the present data suggest a more uniform pattern of association regardless of age.

Table 7: Age-group interaction models for mode of encountering self-harm content

	(1)	(2)	(3)	(4)
(Intercept)	0.452***	0.551***	0.809	0.637**
	(0.066)	(0.080)	(0.117)	(0.090)
	(<0.001)	(<0.001)	(0.141)	(0.001)
X2371aTicked (Yes)	12.874***			
	(2.190)			
	(<0.001)			
X2371aSkipped(GW_NoResponse)	0.520***			
	(0.085)			
	(<0.001)			
X2371aSkipped(GW_Choice)	0.246***			
	(0.016)			
	(<0.001)			
X2371aSkipped(GW_Stopped_vA)	0.592***			
	(0.053)			
	(<0.001)			
ygY10_11	1.375***	1.320**	1.212*	1.314***
	(0.093)	(0.128)	(0.096)	(0.086)
	(<0.001)	(0.004)	(0.015)	(<0.001)
ygY12_13	1.844***	1.979***	1.537**	1.749***
	(0.207)	(0.322)	(0.220)	(0.192)
	(<0.001)	(<0.001)	(0.003)	(<0.001)
GENDER4GD	8.904***	8.913***	8.944***	9.001***
	(1.529)	(1.509)	(1.515)	(1.522)
	(<0.001)	(<0.001)	(<0.001)	(<0.001)
GENDER4Girl	2.133***	2.070***	2.079***	2.099***
	(0.098)	(0.094)	(0.094)	(0.095)
	(<0.001)	(<0.001)	(<0.001)	(<0.001)
GENDER4GND	3.356***	3.431***	3.416***	3.468***
	(0.361)	(0.359)	(0.358)	(0.362)
	(<0.001)	(<0.001)	(<0.001)	(<0.001)
CareNot now, but I used to be in care	1.723*	1.551*	1.560*	1.558*
	(0.381)	(0.332)	(0.335)	(0.334)
	(0.014)	(0.040)	(0.038)	(0.039)
CareNo	1.006	0.902	0.916	0.911
	(0.116)	(0.099)	(0.101)	(0.100)
	(0.958)	(0.345)	(0.426)	(0.395)
CareI don't know what this means	0.976	0.866	0.874	0.865
	(0.138)	(0.118)	(0.119)	(0.117)
	(0.862)	(0.287)	(0.323)	(0.285)
CarePrefer not to say	2.165***	1.858**	1.850**	1.861**
	(0.501)	(0.421)	(0.422)	(0.422)
	(<0.001)	(0.006)	(0.007)	(0.006)
CareNoResponse	0.692	0.581	0.589	0.584
	(0.405)	(0.336)	(0.340)	(0.336)
	(0.529)	(0.347)	(0.359)	(0.350)
ETHNIC.ONSMixed/Multiple Ethnic Groups (aggregated)	1.112	1.108	1.121	1.119
	(0.092)	(0.091)	(0.092)	(0.091)
	(0.200)	(0.210)	(0.162)	(0.170)
ETHNIC.ONSAsian/Asian British (aggregated)	0.671***	0.687***	0.686***	0.691***
	(0.047)	(0.047)	(0.047)	(0.047)

Full sample descriptives

Table 8 provides descriptive statistics for the full analytic sample. The survey captured a wide range of sociodemographic and behavioural characteristics, including self-harm classification, frequency and recency of self-harm behaviours, online behaviour, and exposure to self-harm content. These descriptive data complement the stratified summary in Table 1 and provide a comprehensive overview of the study population.

Table 8: Descriptive statistics for the full analytic sample

Characteristic	**N = 24,909**
Self-harm classification	4,479 / 24,909 (18%)
Year group	
Y05	2,412 / 24,909 (9.7%)
Y06	2,566 / 24,909 (10%)
Y07	3,933 / 24,909 (16%)
Y08	3,878 / 24,909 (16%)
Y09	3,831 / 24,909 (15%)
Y10	2,847 / 24,909 (11%)
Y11	2,727 / 24,909 (11%)
Y12	1,641 / 24,909 (6.6%)
Y13	1,074 / 24,909 (4.3%)
Gender	
Boy	11,499 / 24,706 (47%)
Girl	12,349 / 24,706 (50%)
Gender Diverse (GD)	267 / 24,706 (1.1%)
Gender Non-Disclosing (GND)	591 / 24,706 (2.4%)
Missing	203
Ethnic background	
Asian/Asian British (aggregated)	3,607 / 24,909 (14%)
Black/Black British/African/Caribbean (aggregated)	1,152 / 24,909 (4.6%)
Mixed/Multiple Ethnic Groups (aggregated)	1,313 / 24,909 (5.3%)
Other ethnic group	978 / 24,909 (3.9%)
White (aggregated)	13,561 / 24,909 (54%)
Skipped by respondent	4,298 / 24,909 (17%)
Child in care, looked after, or fostered?	
Yes	805 / 24,909 (3.2%)
No	17,508 / 24,909 (70%)
I don't know what this means	1,215 / 24,909 (4.9%)
Not now, but I used to be in care	206 / 24,909 (0.8%)
Prefer not to say	167 / 24,909 (0.7%)
Skipped by respondent	5,008 / 24,909 (20%)
Neurodivergent	
Yes	3,630 / 24,909 (15%)
No	12,291 / 24,909 (49%)
Not sure	3,692 / 24,909 (15%)
Prefer not to say	270 / 24,909 (1.1%)

(continued)

Characteristic	N = 24,909
Skipped by respondent	5,026 / 24,909 (20%)
Deprivation index	5.49 (3.40)
Missing	1,114
Worry about not having enough money for family needs	
Never or hardly ever	17,054 / 24,909 (68%)
Often	1,654 / 24,909 (6.6%)
Some of the time	5,859 / 24,909 (24%)
Skipped by respondent	342 / 24,909 (1.4%)
Family uses food banks	
Never or hardly ever	22,878 / 24,909 (92%)
Often	310 / 24,909 (1.2%)
Some of the time	1,151 / 24,909 (4.6%)
Skipped by respondent	570 / 24,909 (2.3%)
House is cold and/or damp	
Never or hardly ever	22,162 / 24,909 (89%)
Often	462 / 24,909 (1.9%)
Some of the time	1,876 / 24,909 (7.5%)
Skipped by respondent	409 / 24,909 (1.6%)
Unable to afford uniform, equipment, or school trips	
Never or hardly ever	22,527 / 24,909 (90%)
Often	620 / 24,909 (2.5%)
Some of the time	1,412 / 24,909 (5.7%)
Skipped by respondent	350 / 24,909 (1.4%)
Unable to afford to eat at school	
Never or hardly ever	22,845 / 24,909 (92%)
Often	532 / 24,909 (2.1%)
Some of the time	1,153 / 24,909 (4.6%)
Skipped by respondent	379 / 24,909 (1.5%)
Insufficient space at home for homework or relaxation	
Never or hardly ever	21,231 / 24,909 (85%)
Often	965 / 24,909 (3.9%)
Some of the time	2,330 / 24,909 (9.4%)
Skipped by respondent	383 / 24,909 (1.5%)
No or poor internet access at home	
Never or hardly ever	22,563 / 24,909 (91%)
Often	357 / 24,909 (1.4%)
Some of the time	1,598 / 24,909 (6.4%)
Skipped by respondent	391 / 24,909 (1.6%)
Goes to bed hungry due to food insecurity	
Never or hardly ever	23,443 / 24,909 (94%)
Often	247 / 24,909 (1.0%)
Some of the time	804 / 24,909 (3.2%)
Skipped by respondent	415 / 24,909 (1.7%)
Free school meals (% in school)	20 (16)
Missing	1,114
Anxiety and depressive disorder	4,611 / 24,909 (19%)

(continued)

Characteristic	N = 24,909
Age at first self-harm	11.58 (2.36)
Missing	20,987
Have you ever deliberately self-harmed?	4,479 / 24,909 (18%)
Frequency of self-harm	
Daily	566 / 24,909 (2.3%)
Weekly	637 / 24,909 (2.6%)
A few times	1,579 / 24,909 (6.3%)
Once or twice	770 / 24,909 (3.1%)
Prefer not to say	453 / 24,909 (1.8%)
Skipped by respondent	20,904 / 24,909 (84%)
When did you last self-harm?	
In the last week	1,048 / 24,909 (4.2%)
In the last month	932 / 24,909 (3.7%)
In the past 3-6 months	839 / 24,909 (3.4%)
6 months to a year ago	503 / 24,909 (2.0%)
Over a year ago	643 / 24,909 (2.6%)
Skipped by respondent	20,944 / 24,909 (84%)
Self-harm method: Self-injury	
Yes	3,776 / 24,909 (15%)
Skipped by respondent	21,133 / 24,909 (85%)
Self-harm method: Overdose	
Yes	848 / 24,909 (3.4%)
Skipped by respondent	24,061 / 24,909 (97%)
Self-harm method: Other	
Yes	757 / 24,909 (3.0%)
Skipped by respondent	24,152 / 24,909 (97%)
Hours per day on social networking sites or forums	
0 hrs	1,681 / 23,237 (7.2%)
30 mins	1,643 / 23,237 (7.1%)
1 hr	1,778 / 23,237 (7.7%)
1 hr 30 mins	1,313 / 23,237 (5.7%)
2 hrs	2,346 / 23,237 (10%)
3 hrs	2,387 / 23,237 (10%)
4 hrs	1,884 / 23,237 (8.1%)
5 hrs	1,396 / 23,237 (6.0%)
6 hrs	957 / 23,237 (4.1%)
7 hrs	460 / 23,237 (2.0%)
8 hrs or more	1,548 / 23,237 (6.7%)
Skipped by respondent	5,844 / 23,237 (25%)
Missing	1,672
Publicly available social media account	
Yes	6,866 / 23,018 (30%)
No	10,231 / 23,018 (44%)
Skipped by respondent	5,921 / 23,018 (26%)
Missing	1,891
Come across self-harm content online (last month)?	
Yes, daily	479 / 22,695 (2.1%)

(continued)

Characteristic	N = 24,909
Yes, weekly	512 / 22,695 (2.3%)
Yes, a few times	1,953 / 22,695 (8.6%)
Yes, once or twice	3,771 / 22,695 (17%)
No, never	10,461 / 22,695 (46%)
Skipped by respondent	5,519 / 22,695 (24%)
Missing	2,214
How encountered: I searched for it	
Yes	606 / 17,176 (3.5%)
Skipped by respondent	16,570 / 17,176 (96%)
Missing	7,733
How encountered: It came up on my feed	
Yes	3,644 / 17,176 (21%)
Skipped by respondent	13,532 / 17,176 (79%)
Missing	7,733
How encountered: Accidental exposure	
Yes	2,833 / 17,176 (16%)
Skipped by respondent	14,343 / 17,176 (84%)
Missing	7,733
How encountered: Someone else shared with me	
Yes	524 / 17,176 (3.1%)
Skipped by respondent	16,652 / 17,176 (97%)
Missing	7,733
How encountered: Other	
Yes	398 / 17,176 (2.3%)
Skipped by respondent	16,778 / 17,176 (98%)
Missing	7,733
Why were you looking for self-harm related content?	
Find support for yourself	164 / 18,800 (0.9%)
Find support for someone you know	24 / 18,800 (0.1%)
To connect with others who have had similar experiences	131 / 18,800 (0.7%)
To look up facts and figures about self-harm	71 / 18,800 (0.4%)
Another reason	187 / 18,800 (1.0%)
Skipped by respondent	18,223 / 18,800 (97%)
Missing	6,109

[1] n / N (%); Mean (SD)

Role of the funding source. *[To be completed — The funder had no role in study design, data collection, data analysis, data interpretation, or writing of the report.]*