Online Appendix for "Amplifying Consumers' Voice: The FTC's ReportFraud Website Redesign"*

Michel Grosz
Federal Trade Commission
mgrosz@ftc.gov

Devesh Raval
Federal Trade Commission
draval@ftc.gov

July 9, 2024

*The views expressed in this article are those of the authors. They do not necessarily represent those of the Federal Trade Commission or any of its Commissioners.

C Appendix Tables and Figures

Figure OA1: Complaints by Week to the BBB and CFPB

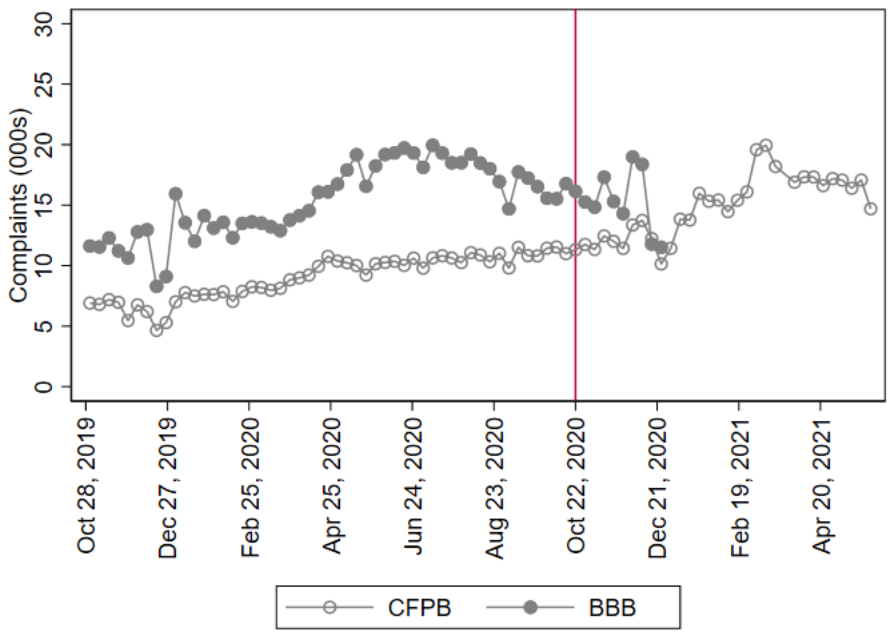

Notes: The figure shows the number of complaints, in thousands, logged each week between October 26, 2019 and June 19, 2021, across the BBB and CFPB sources. The weeks are defined as starting on Thursdays, since the website redesign was a Thursday. BBB complaints are limited to before January 1, 2021. The vertical line shows the date of the website redesign.

Figure OA2: Complaints by Week, Raw Counts

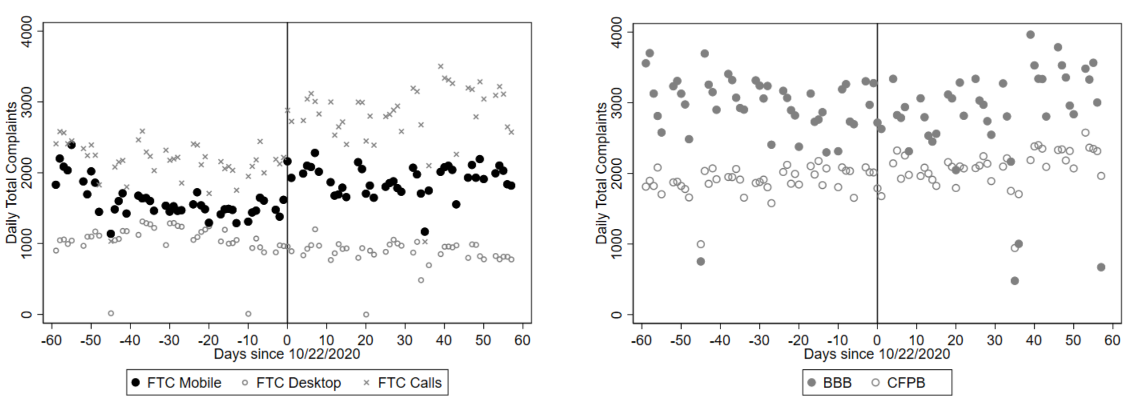

a) FTC

b) BBB/CFPB

Notes: The figure shows the number of complaints each day between from 60 days before and after October 22, 2020.

Figure OA3: RD Estimate of Website Redesign on Number of Complaints, BBB and CFPB

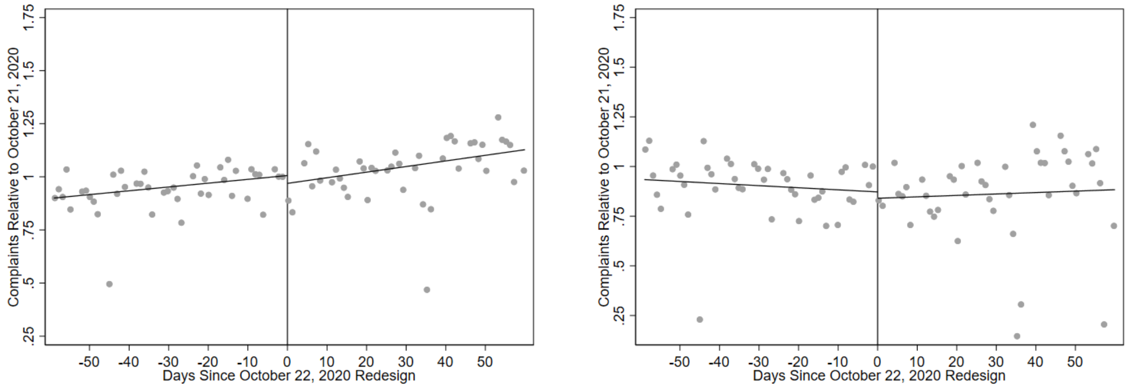

a) CFPB b) BBB

Notes: The figure shows the daily number of complaints report to the CFPB and BBB, from 60 days before and after the FTC's website redesign on October 22, 2020. For each panel, the number of complaints are expressed relative to the number of complaints on the day prior to the redesign, October 21, 2020, which are set to one. The vertical bar shows the date of the redesign. The fitted lines are an RD estimate the includes a first-degree polynomial and controls for the day of the week.

Figure OA4: Total and New Users to FTC Complaints Website

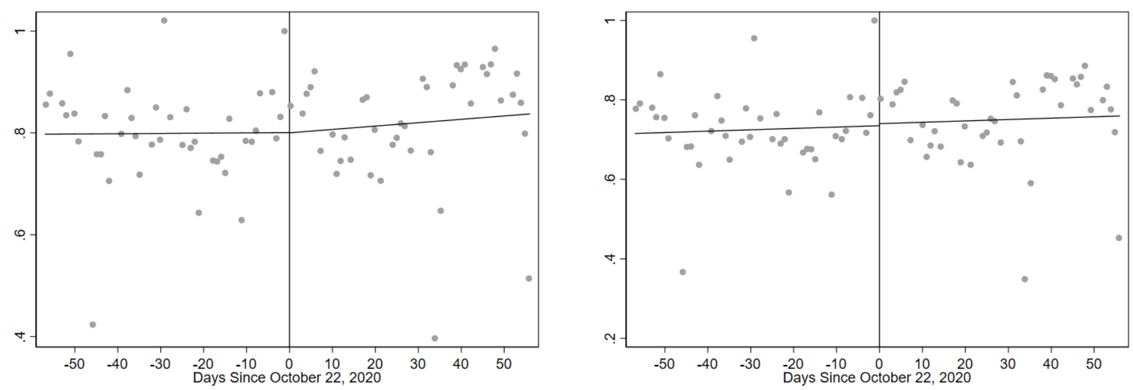

a) Total Users b) New Users

Notes: The figure shows the daily number of total users and new users to the FTC's complaint website, from 60 days before and after the FTC's website redesign on October 22, 2020. For each panel, the number of users are expressed relative to the number of users on the day prior to the redesign, October 21, 2020, which are set to one. The vertical bar shows the date of the redesign. The fitted lines are an RD estimate the includes a first-degree polynomial and controls for the day of the week.

Table OA1: Website metrics before and after redesign

	(1) October 2019 - October 20, 2020	(2) October 22, 2020 - July 31, 2021
Average session duration	00:05:33	00:05:30
Average time on page	00:01:10	00:02:28
Average page load time (sec)	1.72	3:29

Notes: The table shows Google Analytics metrics of the FTC website before and after the redesign.

Table OA2: Estimates of AMSE with Varying Polynomial Degree

Order	(1) FTC Online	(2) FTC Mobile	(3) FTC Desktop	(4) FTC Calls	(5) BBB/CFPB
1	0.017	0.001	0.005	0.004	0.030
2	0.030	0.006	0.009	0.006	0.054
3	0.057	0.011	0.017	0.011	0.136
4	0.118	0.031	0.032	0.021	0.289

Notes: The table shows the resulting estimates of the asymptotic mean squared error (AMSE) from estimates of equation (1), where the dependent variable is the log number of daily complaints. The specification includes polynomials up to the specified order and controls for day of the week. FTC Online refers to the sum of FTC mobile and desktop complaints. The data include complaints from 60 days before and after the FTC's website redesign on October 22, 2020. AMSE is calculated as in Pei et al. (2022).

Table OA3: RD Estimate of Website Redesign on Number of Complaints, Varying Polynomial Order

	(1) FTC Online	(2) FTC Mobile	(3) FTC Desktop	(4) FTC Calls	(5) BBB/CFPB
B. P=1					
Post	0.282***	0.256***	0.307***	-0.0856	0.0255
	(0.0379)	(0.0357)	(0.0395)	(0.0456)	(0.0677)
C. P=2					
Post	0.264***	0.266***	0.262***	0.0699	-0.0691
	(0.0577)	(0.0613)	(0.0518)	(0.0520)	(0.109)
D. P=3					
Post	0.395***	0.416***	0.374***	0.0808	0.185
	(0.0721)	(0.0657)	(0.0634)	(0.0531)	(0.104)
E. P=4					
Post	0.279**	0.239**	0.320***	-0.0388	-0.217*
	(0.0846)	(0.0707)	(0.0745)	(0.0567)	(0.109)

Notes: The table shows estimates of equation (1), where the dependent variable is the log number of daily complaints. The specification includes polynomials up to the specified order and controls for day of the week. FTC Online refers to the sum of FTC mobile and desktop complaints. The data include complaints from 60 days before and after the FTC's website redesign on October 22, 2020. Robust standard errors clustered at the daily level.
$*p < 0.05, **p < 0.01, ***p < 0.001$

Table OA4: RD Estimate of Website Redesign on Number of Complaints, 30-Day Window

	(1) FTC Online	(2) FTC Mobile	(3) FTC Desktop	(4) FTC Calls	(5) CFPB	(6) BBB
RD Estimate	0.298***	0.312***	0.285***	0.0315	-0.0262	-0.0291
	(0.0485)	(0.0395)	(0.0400)	(0.0395)	(0.0646)	(0.0477)

Notes: The table shows estimates of equation (1), where the dependent variable is the log number of daily complaints. The specification includes a first degree polynomial and controls for day of the week. FTC Online refers to the sum of FTC mobile and desktop complaints. The data include complaints from 30 days before and after the FTC's website redesign on October 22, 2020. Robust standard errors clustered at the daily level.
$*p < 0.05, **p < 0.01, ***p < 0.001$

Table OA5: RD Estimate of Website Redesign on Google Keyword Searches

Keyword	(1) Estimate	(2) SE	Keyword	(3) Estimate	(4) SE
online scams report	-47.90	19.63	online scam report	0.40	18.85
report online scams	-47.90	19.63	report a website for fraud	1.15	7.29
report scam business	-32.51	19.67	report business	1.40	5.85
ftc complaint assistant	-30.13	37.62	report ftc	1.93	5.32
report online fraud	-28.96	17.34	ftc report	1.93	5.32
f t c complai	-28.59	11.34	report website fraud	2.13	25.08
report internet scam	-26.20	13.71	fraud company	2.98	10.54
ftc complaint	-22.90	10.56	report a scam	6.03	6.35
report scam site	-22.65	12.34	fraud report	6.26	9.90
report fraud company	-17.88	16.49	report fraud	6.26	9.90
ftc report number	-15.18	11.93	scam reporting	10.19	15.12
report website scam	-15.04	24.21	report website	10.77	4.88
report scam website	-15.04	24.21	report scammer	10.84	10.92
ftc fraud	-14.41	16.52	fraud complaint	11.25	9.64
report online scammer	-13.68	10.74	report fraudulent website	12.38	13.04
ftc report scam	-13.33	13.27	how to report a fake website	12.72	9.50
ftc scam report	-13.33	13.27	report a scammer	12.95	10.46
report a scam website	-12.63	19.42	reporting a scammer	14.57	24.67
consprot	-11.47	6.85	how to report scammers	17.13	13.27
ftc fraud reporting	-6.78	13.85	federal fraud	18.30	15.59
business fraud	-5.72	5.96	how to report a fraud website	18.58	27.22
report company	-5.19	5.90	fraud reporting	18.96	17.36
fraud companies	-3.25	7.67	fraud claim	22.58	19.30
f t c fraud r	-2.65	6.51	reporting a scam	24.68	12.27
f t c report	-2.65	6.51	how to report fraud website	25.69	11.69
ftccomplaintassistant	-1.74	2.48	report a fake website	27.29	14.05
website complaint	-1.69	13.20	where to report scams	29.21	12.03
ftc report fraud	-1.36	15.48	fraud websites	31.45	19.06
ftc fraud report	-1.36	15.48	report fake website	37.40	10.04
f t c reporti	-1.06	6.96	website fraud	42.27	24.02
scam report	0.16	8.16	fraud website	42.27	24.02
report scam	0.16	8.16			

Notes: The table shows estimates of equation (1), where the dependent variable is the Google Trends popularity for 65 different keywords changed in the 10 weeks before and after the website redesign. The specification includes a first degree polynomial and controls for day of the week. Coefficients are sorted by their magnitude.

Table OA6: RD Estimate of March 3 2021 Press Release on Number of Complaints

	(1) FTC Online	(2) FTC Mobile	(3) FTC Desktop
Post	0.0420	0.0707*	0.0133
	(0.0386)	(0.0297)	(0.0353)

Notes: The table shows estimates of equation (1), where the dependent variable is the log number of daily complaints. The specification includes a first degree polynomial and controls for day of the week. FTC Online refers to the sum of FTC mobile and desktop complaints. The data include complaints from 60 days before and after the FTC's press release on March 3, 2021 about an initiative to encourage low-income communities to report fraud. Robust standard errors clustered at the daily level. $*p < 0.05, **p < 0.01, ***p < 0.001$

Table OA7: RD Estimates, Quality of Data

	(1) Has Zip	(2) Has Age	(3) Company Name	(4) Company Zip
A. FTC Online				
RD Estimate	0.0850***	0.127***	0.0620***	0.0137**
	(0.00372)	(0.00377)	(0.00872)	(0.00468)
B. FTC Mobile				
RD Estimate	0.0786***	0.128***	0.0553***	0.0243***
	(0.00516)	(0.00565)	(0.0133)	(0.00556)
C. FTC Desktop				
RD Estimate	0.0914***	0.126***	0.0687***	0.00315
	(0.00519)	(0.00518)	(0.0107)	(0.00574)

Notes: The table shows estimates of equation (1), where the dependent variable is the fraction of daily complaints that included a zipcode, included a consumer's age, included a defrauding company's name, or included a defrauding company's zipcode. Robust standard errors clustered at the daily level. $*p < 0.05, **p < 0.01, ***p < 0.001$

Table OA8: RD Estimates, Imputed Categories

	(1) Telemarket	(2) Text/Email	(3) Imposter	(4) Online Shopping	(5) Other/Misc
A. FTC Online					
RD Estimate	0.656***	0.0680	0.529***	0.351***	0.133***
	(0.0391)	(0.0541)	(0.0395)	(0.0381)	(0.0371)
A. FTC Mobile					
RD Estimate	0.720***	0.174*	0.560***	0.285***	0.160***
	(0.0523)	(0.0705)	(0.0453)	(0.0408)	(0.0377)
B. FTC Desktop					
RD Estimate	0.591***	-0.0377	0.497***	0.417***	0.106**
	(0.0488)	(0.0571)	(0.0420)	(0.0374)	(0.0369)

Notes: The outcomes are the log sum of predicted probabilities for each imputed category based on the text of complaints. The specification includes a first degree polynomial and controls for day of the week. Data includes complaints at the daily level across FTC desktop and mobile platforms. The data include complaints from 60 days before and after the FTC's website redesign on October 22, 2020. Robust standard errors clustered at the daily level. $*p < 0.05, **p < 0.01, ***p < 0.001$

Table OA9: Terms with Statistically Significant Changes

	(1) Positive	(2) Negative
FTC Mobile Only	requir, america, possibl, support	phone, address, messag, right, insur, numer
FTC Desktop Only	unemploy, subject	phone, address, account, receiv, number, anoth, charg, includ, report, amount, cancel, transfer, chang, follow, first, remov, found, clear, immedi, comput, verifi, appear, attempt, question, instruct, problem, addit, action, block, sever, associ, investig, bitcoin, anyon, search, longer, attorney, covid, submit, might, resolv, receipt involv practic unabl posit enter, approv, theft

Notes: The table shows the terms that saw a statistically significant decline or increase in use after the FTC's website redesign for only mobile or desktop complaints; terms with a significant decline for both types of complaints are included in Table 8. The set of all possible terms do not common stop words or words that were included in fewer than 1 percent or over 40 percent of FTC complaints. All words were destemmed to create the most popular terms. Each resulting term was then used as a dependent variable in estimates of equation (1), with a Bonferroni correction to account for multiple hypothesis testing. The specification includes a first degree polynomial and controls for day of the week. Data includes complaints at the daily level across FTC desktop and mobile platforms. The data include complaints from 60 days before and after the FTC's website redesign on October 22, 2020.

Table OA10: RD Estimate of Website Redesign on Selected Topics, Positive and Significant

Representation	Est.	(S.E.)	Total Complaints	Assigned Category
9_ifidonotreceivethebitcoin idefinitelywillsendoutyourvideorecordingto	1.96	(0.13)	3373	Hack Claim
49_unemployment claim_kansas dept_ks dept_unemployment benefit	1.89	(0.17)	1446	Identity Theft
21_unemployment claim_illinois dept_filed claim_unemployment benefit	1.73	(0.24)	2624	Identity Theft
23_unemployment fraud_unemployment claim_filed claim_fraudulent claim	1.67	(0.20)	2552	Identity Theft
187_data video_monitoring internet_uploaded trojan_letter video	1.67	(0.13)	485	Hack Claim
57_gift card_itunes gift_ebay gift_play gift	1.57	(0.17)	1308	Imposter
47_esto_cuando_que el_como	1.47	(0.20)	1508	Spanish Language Cluster
80_calls marriott_called marriott_marriott hotel_vacation marriott	1.43	(0.12)	840	Prize/Imposter
39_online interview_hiring manag_google hangout_interview	1.35	(0.18)	1729	Job
48_unemployment claim_filed claim_fraudulent claim_filed fraudul	1.27	(0.14)	1451	Identity Theft
91_number fraudul_ss number_robocall claim_scam robocal	1.24	(0.14)	766	Imposter
60_phone robo_number robo_robo claim_robo tel	1.03	(0.07)	1189	Imposter
181_update help_update just_update ne_update problem	1.03	(0.10)	2730	"Update" Word
112_regards outlook_outlook io_urgent task_email soon	1.02	(0.08)	600	Imposter
188_transaction ebay_deal ebay_buyer contact_ebay select	1.02	(0.08)	322	Online Shopping
1_icloud breach_saying icloud_icloud account_apple icloud	1.01	(0.11)	9218	Imposter

Notes: The table shows estimates of equation (1), where the dependent variable is the log number of complaints within each imputed topic area. For each topic, we include the representation of the topic, RD coefficient and standard error, the number of complaints in the two months before and after the redesign, and the category assigned by the authors to the topic. The representation of the topic includes a number (in order of size from 0 to 367) and representative words for the topic. The specification includes a first degree polynomial and controls for day of the week. Data includes complaints at the daily level across FTC desktop and mobile platforms. The data include complaints from 60 days before and after the FTC's website redesign on October 22, 2020. Robust standard errors clustered at the daily level.

Table OA11: RD Estimate of Website Redesign on Selected Topics, Positive and Significant (Continued)

Representation	Est.	(S.E.)	Total Complaints	Assigned Category
118_cashier check_receive check_payment_payment address	0.97	(0.09)	550	Advance Fee
95_sba fraud_sba loan_loan sba_contacted sba	0.96	(0.14)	713	Identity Theft
121_stop robo_number robo_robo_calls robocal	0.96	(0.13)	547	Unwanted Calls
144_paste text_unable copypast_copy text_paste don	0.94	(0.08)	448	Want to Copy Paste
152_foods market_food market_quality survey_survey	0.91	(0.11)	422	Job
170_robocall amazon_amazon robocal_scam robocal_robocall claim	0.91	(0.15)	371	Imposter
173_amazon sign_amazon password_respond amazon_https amazon	0.89	(0.13)	363	Imposter
37_order confirm_shipment_order ship_shipping spe	0.88	(0.15)	1780	Imposter
28_contacted ebay_fake ebay_ebay contact_ebay email	0.87	(0.08)	2366	Online Shopping
68_amazon fraud_fraud amazon_amazon fraudul_amazon account	0.76	(0.15)	1010	Imposter
125_foods market_shoppers work_shopper_store evalu	0.72	(0.13)	638	Job
92_amazon fraud_amazon robo_amazon account_robo amazon	0.69	(0.14)	725	Imposter
3_calls com_calls differ_caller_caller id	0.68	(0.07)	6612	Unwanted Calls
123_hacked taken_send video_video btc_traced hack	0.67	(0.19)	526	Hack Claim
17_federal bureau_payment_transfer_bureau investig	0.66	(0.18)	2826	Money Transfer
133_refund_fraud_asked pay_told pay	0.66	(0.18)	499	Job
0_ss number_number suspend_ssa_ssn suspend	0.65	(0.13)	28618	Imposter
62_windows defend_defender protect_protection plan_threat protect	0.64	(0.12)	1117	Tech Support
36_received voicemail_issue arrest_warrant arrest_left voicemail	0.60	(0.19)	1843	Imposter
27_iphone amazon_calling amazon_amazon fraud_contacted amazon	0.55	(0.20)	2402	Imposter
30_withdraw profit_withdraw money_withdraw fund_tried withdraw	0.55	(0.17)	2094	Investments
6_won million_scammer_prize money_told won	0.52	(0.15)	4068	Lottery/Prize
45_theft_attorney_court_fraud	0.52	(0.14)	1519	Unrelated Misc
19_charged amazon_charge amazon_amazon account_calling amazon	0.50	(0.17)	2771	Imposter
13_landlord_craigslist ad_tenant_craigslist org	0.46	(0.16)	3086	Rental
26_medicare inform_medicare report_medicare told_calling medicar	0.39	(0.11)	2422	Imposter
8_grant program_grant money_000 grant_free grant	0.29	(0.16)	3926	Imposter/Govt Grant
10_scammer_asked money_send money_asking money	0.27	(0.16)	3374	Imposter/Romance

Notes: The table shows estimates of equation (1), where the dependent variable is the log number of complaints within each imputed topic area. For each topic, we include the representation of the topic, RD coefficient and standard error, the number of complaints in the two months before and after the redesign, and the category assigned by the authors to the topic. The representation of the topic includes a number (in order of size from 0 to 367) and representative words for the topic. The specification includes a first degree polynomial and controls for day of the week. Data includes complaints at the daily level across FTC desktop and mobile platforms. The data include complaints from 60 days before and after the FTC's website redesign on October 22, 2020. Robust standard errors clustered at the daily level.

Table OA12: RD Estimate of Website Redesign on Selected Topics, Negative and Significant

Representation	Est.	(S.E.)	Total Complaints	Assigned Category
20_received check_applied job_offered job_send check	-0.42	(0.13)	2747	Job
2_spam text_unwanted text_received text_receiving text	-0.46	(0.16)	9213	Spam Text
59_order norton_contacted norton_subscription norton_norton antivirus	-0.52	(0.18)	1227	Imposter/Tech Support
114_survey mcguireresearch_selected survey_surveys random_survey text	-0.54	(0.17)	589	Spam Text
106_weight loss_losing weight_weightloss_loose weight	-0.67	(0.13)	635	Spam Text
24_unsubscribed email_unsubscribe email_emails unsubscrib_email unsubs	-0.69	(0.17)	2520	Spam Email
159_skin cream_skin product_skin trial_remove skin	-0.71	(0.17)	397	Spam Text
140_tumors ovarian_cancer femal_cancer_tumor	-0.80	(0.12)	486	Spam Text
46_giftcard_cards email_scammer_gift card	-0.90	(0.36)	1493	Imposter
119_money 3500_money 4000_money 1500_money 4500	-0.99	(0.17)	560	Spam Text
115_covid 19_topic covid_ss number_fraud	-1.03	(0.18)	573	Imposter
16_political text_text polit_political spam_messages polit	-1.14	(0.24)	2983	Spam Text
15_scam email_spam email_email threaten_hacked account	-1.31	(0.54)	2850	Hack Claim

Notes: The table shows estimates of equation (1), where the dependent variable is the log number of complaints within each imputed topic area. For each topic, we include the representation of the topic, RD coefficient and standard error, the number of complaints in the two months before and after the redesign, and the category assigned by the authors to the topic. The representation of the topic includes a number (in order of size from 0 to 367) and representative words for the topic. The specification includes a first degree polynomial and controls for day of the week. Data includes complaints at the daily level across FTC desktop and mobile platforms. The data include complaints from 60 days before and after the FTC's website redesign on October 22, 2020. Robust standard errors clustered at the daily level.

Table OA13: Differences between Takers and Compliers, FTC Mobile and Desktop Complaints

	(1)	(2)	(3)	(4)
	Mobile		Desktop	
	Takers	Compliers	Takers	Compliers
Demographic Characteristics				
Age under 40	0.447	0.509	0.279	0.157
Age 40-59	0.358	0.374	0.332	0.403
Age Over 60	0.195	0.118	0.389	0.439
Age 60-69	0.135	0.082	0.223	0.237
Age 70-79	0.051	0.026	0.129	0.162
Age Over 80	0.010	0.009	0.037	0.040
White	0.704	0.735	0.749	0.778
Black	0.107	0.102	0.1000	0.0811
Latino	0.121	0.0739	0.0792	0.0685
Asian/PI	0.0451	0.0621	0.0512	0.0503
Female	0.533	0.526	0.469	0.638
Demographic Weight	0.952	1.285	0.937	1.408
Text				
Reported Loss	0.251	0.185	0.208	0.212
Grade Level	9.644	2.745	8.678	7.758
Grade Level $<$ 8th	0.623	0.0173	0.575	0.382
Imputed Product Category				
Telemarketing	0.0616	0.198	0.0805	0.249
Unsolicited Text or Email	0.248	0.150	0.252	-0.0249
Imposter Scams	0.180	0.399	0.186	0.453
Online Shopping and Reviews	0.108	0.0848	0.0902	0.175
All Other and Misc.	0.402	0.169	0.391	0.149

Notes: The first column of the table shows the mean characteristics for FTC complaints in the 30 days prior to the website redesign on October 22, 2020. The second column shows the imputed means for complier complainants. These means are calculated using the pre-redesign mean, the coefficient estimate from equation (1), and the coefficient estimate on the number of complaints.

D Additional Details on Empirical Approaches

D.1 Flesch Reading Ease Score

The Flesch Reading Ease score assigns a text's readability a number between 1 (hardest) and 100 (easiest). The Flesch Reading Ease measure is defined as

$$206.835 - 1.015\left(\frac{words}{sentences}\right) - 84.6\left(\frac{syllables}{words}\right). \tag{1}$$

The scores can also be grouped into grade level difficulty, with lower than 70 being apt for 8th grade and above, and below 50 being college-level. We report results for the RD analysis on all three measures based on the Flesch Reading Ease Score in Table OA14.

Overall, median reading ease increased by 2.5 points, with the share of texts with at least an 8th grade or college reading ease declined by 5 to 6 percentage points. On the mobile site, median reading ease increased by 5 points, and the share of texts with at least an 8th grade or college reading ease declined by approximately 11 percentage points. We find small changes for the desktop site.

Table OA14: RD Estimates, Text Analysis Flesch Reading Ease

	(1) Median	(2) 8th gr	(3) college
A. FTC Online			
RD Estimate	2.401***	-0.0574***	-0.0488***
	(0.502)	(0.0111)	(0.0116)
A. FTC Mobile			
RD Estimate	4.969***	-0.117***	-0.113***
	(0.437)	(0.00871)	(0.00927)
B. FTC Desktop			
RD Estimate	-0.168	0.00185	0.0158*
	(0.373)	(0.00870)	(0.00752)

Notes: In the first column the dependent variable is the median Flesch-Kincaid Reading Ease Score, and the final two columns are the fraction of complaints above 8th grade or college according to the Flesch Reading Ease score. The specification includes a first degree polynomial and controls for day of the week. Data includes complaints at the daily level across FTC desktop and mobile platforms. The data include complaints from 60 days before and after the FTC's website redesign on October 22, 2020. Robust standard errors clustered at the daily level.
$*p < 0.05, **p < 0.01, ***p < 0.001$

D.2 Imputed Product Categories

The Consumer Sentinel database classifies FTC complaints into 30 product categories, which are themselves divided into about 100 more detailed product codes. However, this categorization, as well as how consumers provide the category of their complaint, changes with the website redesign. Thus, we use the text fields to predict complaint categories and examine how the predicted probabilities of complaints change with the redesign.

We develop our baseline categorization of complaints by taking the categories with at least a 10% share of desktop and mobile complaints in the two months after the redesign. Only three categories satisfy this criterion: Imposter Scams, Online Shopping / Reviews, and Unspecified Reports. Since "Unspecified Reports" is the largest category, we break it up into its product codes to create two additional categories with a share above 10%: Unwanted Telemarketing and Unsolicited Text or Email (which combines the Unsolicited Text and Unsolicited Email product codes). Finally, complaints with the "Other Misc." product code and complaints from categories below the 10% share above are categorized into a catch all "Other" category.

In order to further examine the performance of our baseline LLM predictions, we report the "confusion matrix" of these predictions in Table OA15 using data from the 10% test set not used for estimation. We assign each complaint to the category with the maximum probability, and then compare predicted categories (rows in the table) to actual categories (column in the table). In general, the most common actual category is the same as the predicted category. For example, 66% of complaints predicted to be about "Online Shopping" are actually categorized as "Online Shopping" in Sentinel. The main exception is Telemarketing, as a lot of the complaints that we categorize as Telemarketing based on the LLM predictions are actually characterized as Imposter Scams. These incorrect predictions may reflect that many imposter scams happen via telephone calls, and so share similarities with Unwanted Telemarketing Calls.

Table OA15: Confusion Matrix for Baseline Categorization

Predicted Category	Actual Category				
	Other	Telemarketing	Unsolicited Text/Email	Imposter Scams	Online Shopping
Other	57%	6%	26%	8%	2%
Telemarketing	8%	36%	0%	55%	1%
Unsolicited Text/Email	11%	0%	79%	8%	1%
Imposter Scams	6%	12%	15%	65%	2%
Online Shopping	7%	13%	6%	9%	66%

Notes: We use data from the 10% test set and predict categories based on the category with the maximum probability for each complaint. Each cell is the share of complaints assigned to a predicted category (the row) whose actual category in Consumer Sentinel is the given column.

We also develop a broader alternative categorization of complaints into 13 categories. We develop these categories by using the detailed product codes and including all product codes with at least a 1% share of desktop and mobile complaints in the two months after the redesign. We then combine all categories that are not included, as well as the "Other Misc." product code, into an "All Other" category. This process results in the following 13 categories: Unwanted Telemarketing; Unsolicited Text; Business Imposter; Online Shopping; Govt Imposter; Unsolicited Email; Tech Support; Job Scams; Prizes/Sweepstakes; Romance Scams; Misc Investments; Diet Plans / Centers; and All Other. We then fine tune the Large Language Model to predict these categories using data from the two months after the redesign, and hold out 10% of the sample as a test set.

Figure OA5 and Table OA16 display the RD estimates from this broader categorization. The predictive model has a predictive accuracy of 58% on the test set. Given the number of categories, we do not display the full confusion matrix. Instead, in Table OA17, we take each predicted category (assigned based on maximum probability), and then display the share of complaints whose actual category in Consumer Sentinel is the same as the predicted category, and the actual category with the highest share of complaints. The predicted category has the highest share of complaints from the actual category for 10 of the 13 categories; for Tech Support scams the highest share is Unwanted Telemarketing, and for Prize/Sweepstakes and Diet Plans/Centers the highest share is Unsolicited Text.

D.3 Topic Modeling

In this section, we provide more details on the topic modeling approach pursued in the paper. We estimate the topic model using the *BERTopic* package in Python (Grootendorst, 2022). This approach takes several steps:

1. The first step is to convert the documents into sentence embeddings. Here, we use the *All-MiniLM-L6-v2* model to do so. We also set a maximum length of 512 tokens, as the model cannot handle more than that number of tokens, and longer complaints.
2. The second step is to reduce the dimensionality of the resulting embeddings. Here, we apply the default in *BERTopic* of using the *UMAP* package.
3. The third step is to cluster the documents. Here we use the *HDBScan* package (the default), but a set a minimum cluster size of 120 documents.
4. The fourth step is to turn each cluster into one large document by combining all of the complaints in that cluster, and then developing a "bag of words" representation of each cluster. Here, we vectorize the words by using the SnowBall Stemmer from the *nltk* package, removing stopwords, and including both individual words and bigrams (combinations of two words).
5. The fifth step is to weight these words based on the relative frequency in a given cluster compared to other clusters. Here, we use a "C-TF-IDF" representation, which multiplies the frequency of a term in a cluster by the inverse of its overall frequency across all clusters. We use a class-based BM-25 weighting measure.
6. Finally, in the last step, we fine tune the topic representations using the *KeyBERTInspired* model.

Figure OA5: RD Estimate of Website Redesign on Imputed Product Category, More Detailed Categorization

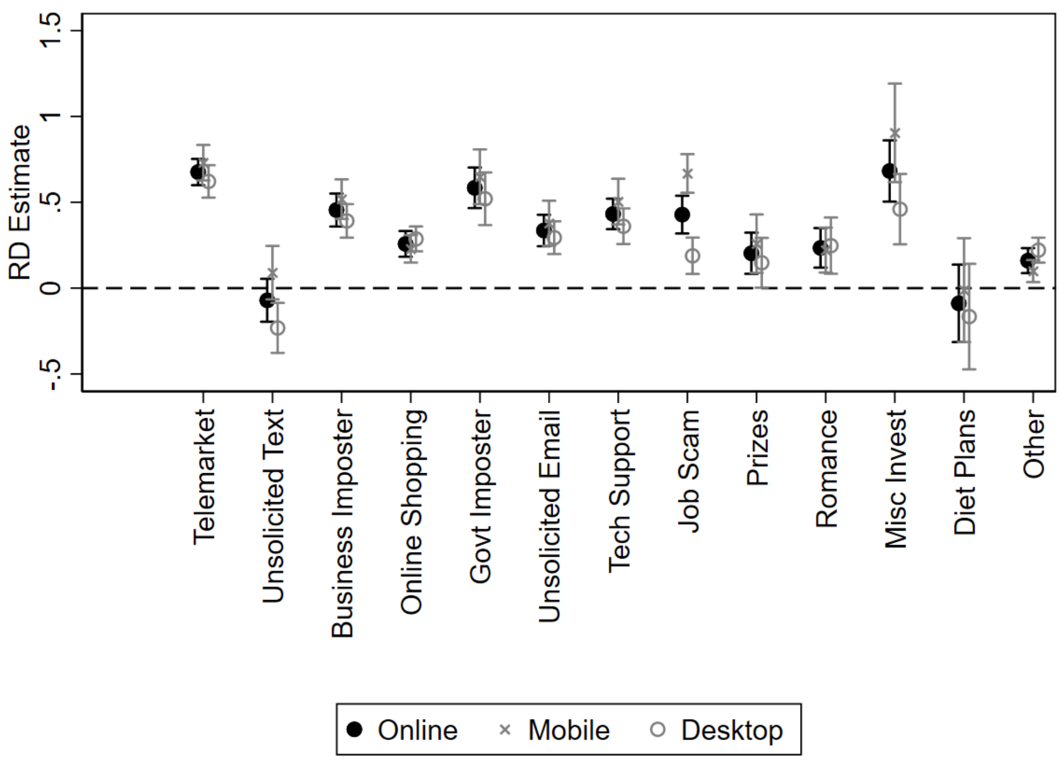

Notes: The figure shows point estimates and 95% confidence intervals for estimates of equation (1), where the dependent variable is the log sum of the predicted probability of each category using the text of consumer complaints. Robust standard errors clustered at the daily level. Table OA16 shows the point estimates and standard errors that correspond with this figure.

E Demographics of Complaining Consumers

E.1 Age

We first study the age of consumers who complain, the only demographic characteristic that the FTC explicitly asks for in its complaint submission process. Some scammers explicitly target older adults (Mclean, 2020); one reason may be older adults are at or past their lifecycle peak for wealth accumulation (DeLiema et al., 2020). In response, Congress passed the "Stop Senior Scams Act" in 2022 in order to prevent scams targeting seniors, and the FTC regularly publishes reports to Congress about efforts dedicated to protecting older adults (Federal Trade Commission, 2022).

The complaint data also include self-reported information about consumers' age. Figure OA6 shows the distribution of self-reported consumer age for consumers who filed complaints with FTC desktop and mobile systems. First, complaints that do not report age drop suddenly after the redesign. Second, the distribution of age conditional on reporting also seems to have changed. In particular, the share of consumers aged 60 or older, who might be the ones with the most difficulty in completing online complaint forms, go from being the group with the lowest share of complaints to being the highest.

Table OA18 shows how the number of complaints in each reported age band changed after the redesign. We have expressed each estimate relative to the overall estimates reported in Table 1. All of the age bands grew at or faster than the headline number because the share of consumers not reporting their age fell. For example, the under 40 band grew at 40%, which is only slightly higher, and not statistically significantly different from, the overall increase in complaints. We find similar increases in complaints across age bands; complaints from consumers below

Table OA16: RD Estimates, Imputed Categories, More Detailed Categorization

	(1) Online	(2) Mobile	(3) Desktop
All Other	0.160***	0.0993**	0.221***
	(0.0371)	(0.0327)	(0.0370)
Unwanted Telemarketing	0.676***	0.731***	0.622***
	(0.0390)	(0.0528)	(0.0483)
Unsolicited Text	-0.0707	0.0903	-0.232**
	(0.0638)	(0.0794)	(0.0744)
Business Imposter	0.455***	0.518***	0.392***
	(0.0490)	(0.0588)	(0.0499)
Online Shopping	0.258***	0.229***	0.287***
	(0.0381)	(0.0407)	(0.0370)
Govt. Imposter	0.584***	0.648***	0.520***
	(0.0603)	(0.0814)	(0.0782)
Unsolicited Email	0.336***	0.378***	0.294***
	(0.0468)	(0.0671)	(0.0487)
Tech Support	0.432***	0.504***	0.360***
	(0.0452)	(0.0679)	(0.0529)
Job Scams	0.428***	0.668***	0.188***
	(0.0560)	(0.0575)	(0.0540)
Prizes Sweepstakes	0.203**	0.259**	0.148*
	(0.0612)	(0.0868)	(0.0740)
Romance	0.235***	0.221**	0.248**
	(0.0588)	(0.0667)	(0.0836)
Misc. Investments	0.682***	0.905***	0.460***
	(0.0909)	(0.146)	(0.105)
Diet Plans/Centers	-0.0885	-0.0114	-0.166
	(0.115)	(0.154)	(0.157)

Notes: The outcomes are the log sum of predicted probabilities for each imputed category based on the text of complaints. The specification includes a first degree polynomial and controls for day of the week. Data includes complaints at the daily level across FTC desktop and mobile platforms. The data include complaints from 60 days before and after the FTC's website redesign on October 22, 2020. Robust standard errors clustered at the daily level. $*p < 0.05, **p < 0.01, ***p < 0.001$

40 increased by 40%, as did consumers 60+, compared to 46% for consumers aged 40-59.

These results may seem surprising given that we found a large increase in the share of complaints from consumers 60+ in the time series plot reported in Figure OA6. We thus plot the RD graphs by age group in Figure OA7; the jump at the discontinuity appears to be quite similar across all age groups. However, consumers that are 60+ or 70+ have a pronounced upward trend in complaints after the website redesign. While this increase might reflect longer run effects of the redesign, it could also be due to other factors affecting complaint rates such as changes in fraud victimization over time.

As shown previously, though, the redesign of the FTC's interface increased the proportion of consumers who recorded their age. We thus supplement the analysis above by using the consumer's reported first name to impute age. We use the number of births each year since 1900 with a particular first name from the Social Security Administration (SSA), as well as the likelihood of being alive in 2020 from the SSA actuarial tables. We combine these elements to calculate the median age for each name. This process makes the strong assumption that all names have the same expected life expectancy for a given birth year, and also does not account for immigration.

We first compile the number of births in the US with each first name since 1900, available from the Social Security Administration (SSA). We then calculate the fraction of people born in each year with each name who would be alive in 2020, using the SSA actuarial life tables. These tables only calculate survival rates for birth years on the decade (e.g. 2010, 2000), so for births in other years we interpolate. Because we do not know the sex at birth of each consumer, we calculate the survival rates as the average of male and female survival rates. This exercise gives us the

Table OA17: Predictive Accuracy, More Detailed Categorization

Predicted Category	Share Correct	Highest Share Category
All Other	56%	All Other
Unwanted Telemarketing	46%	Unwanted Telemarketing
Unsolicited Text	89%	Unsolicited Text
Business Imposter	34%	Business Imposter
Online Shopping	66%	Online Shopping
Govt. Imposter	73%	Govt. Imposter
Unsolicited Email	69%	Unsolicited Email
Tech Support	19%	Unwanted Telemarketing (32%)
Job Scams	52%	Job Scams
Prizes Sweepstakes	39%	Unsolicited Text (41%)
Romance	68%	Romance
Misc. Investments	66%	Misc. Investments
Diet Plans/Centers	2%	Unsolicited Text (93%)

Notes: We use data from the 10% test set and predict categories based on the category with the maximum probability for each complaint. Each row is a predicted category; the two columns are the share of complaints whose actual category in Consumer Sentinel is the same as the predicted category, and the actual category with the highest share of complaints (with that share in parentheses if it is different from the predicted category).

Figure OA6: FTC Online Complaints by Week and Age

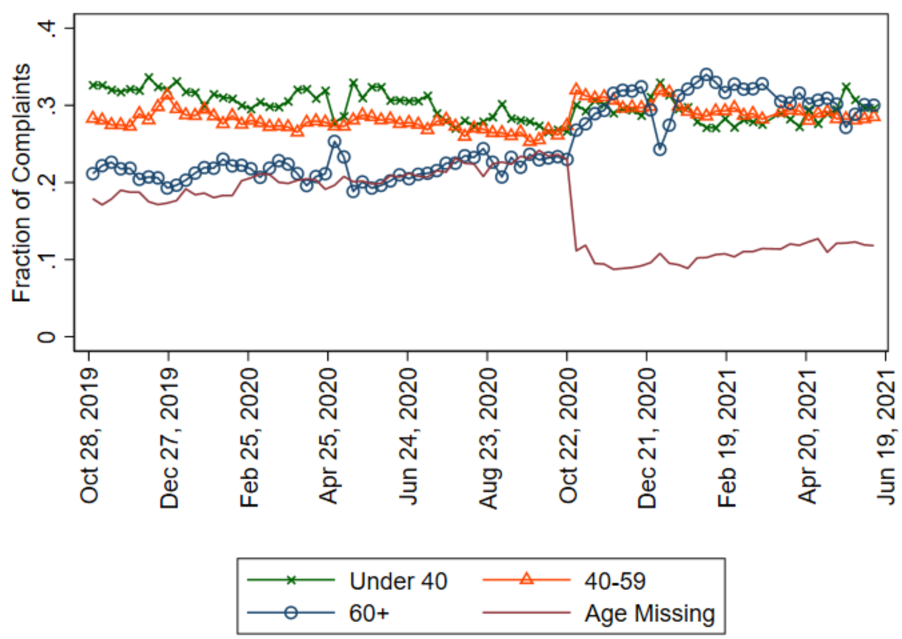

Notes: The figure shows the fraction of total complaints logged each week between October 26, 2019 and June 19, 2021, across FTC online sources and age bands. The weeks are defined as starting on Thursdays, since the website redesign was a Thursday.

median birth year for each name, from which we can calculate the median age in 2020 for each name. We use the age bands from this imputed median age in the table below. The key limitations of this approach are that we must assume that individuals with different names born in the same year have the same survival rates, and that there is no immigration.

These estimates, shown in Table OA19, are similar to our main results that use the actual age bands from the

Table OA18: RD Estimates, Log Complaints in Each Age Range

	(1) <40	(2) 40-59	(3) 60+	(4) 60-69	(5) 70-79	(6) 80+
A. FTC Online						
RD Estimate	0.404***	0.464***	0.404***	0.397***	0.416***	0.470***
	(0.0424)	(0.0496)	(0.0318)	(0.0353)	(0.0419)	(0.0736)
B. FTC Mobile						
RD Estimate	0.494***	0.474***	0.372***	0.370***	0.371***	0.490***
	(0.0431)	(0.0442)	(0.0437)	(0.0488)	(0.0664)	(0.134)
C. FTC Desktop						
RD Estimate	0.315***	0.455***	0.437***	0.424***	0.462***	0.450***
	(0.0453)	(0.0397)	(0.0380)	(0.0407)	(0.0481)	(0.0653)

Notes: The table shows estimates of equation (1), where the dependent variable is the log number of daily complaints by consumers in each age band. Consumer can also choose to not report their age. The age 60 and above is the sum of the other older age bands. Robust standard errors clustered at the daily level. $*p < 0.05, **p < 0.01, ***p < 0.001$

FTC website. The main difference is we find a much larger, albeit noisy, differences for older adults, with smaller increases for the 60-69 age band and large, albeit noisy, increases for the 80+ band. Thus, estimates using the median age of their name are broadly consistent with our findings from Table OA18.

Table OA19: RD Estimates, Log Complaints in Each Imputed Age Range, Using Vital Statistics

	(1) <40	(2) 40-59	(3) 60+	(4) 60-69	(5) 70-79	(6) 80+
A. FTC Online						
RD Estimate	0.412***	0.416***	0.403***	0.298***	0.356***	0.728***
	(0.0706)	(0.0899)	(0.0693)	(0.0379)	(0.0713)	(0.149)
B. FTC Mobile						
RD Estimate	0.364***	0.354***	0.383***	0.296***	0.277*	0.707***
	(0.0648)	(0.0779)	(0.0879)	(0.0472)	(0.112)	(0.179)
C. FTC Desktop						
RD Estimate	0.460***	0.477***	0.422***	0.300***	0.435***	0.749**
	(0.0673)	(0.0766)	(0.0736)	(0.0460)	(0.0797)	(0.231)

Notes: The table shows estimates of equation (1), where the dependent variable is the log number of daily complaints by consumers in each age band. The age bands are calculated combining the consumer name, SSA vital statistics on number of births each year with each name, and actuarial tables. The age 60 and above is the sum of the other older age bands. Robust standard errors clustered at the daily level. $*p < 0.05, **p < 0.01, ***p < 0.001$

Figure OA7: RD Estimate of Website Redesign on Number of Complaints by Age Group

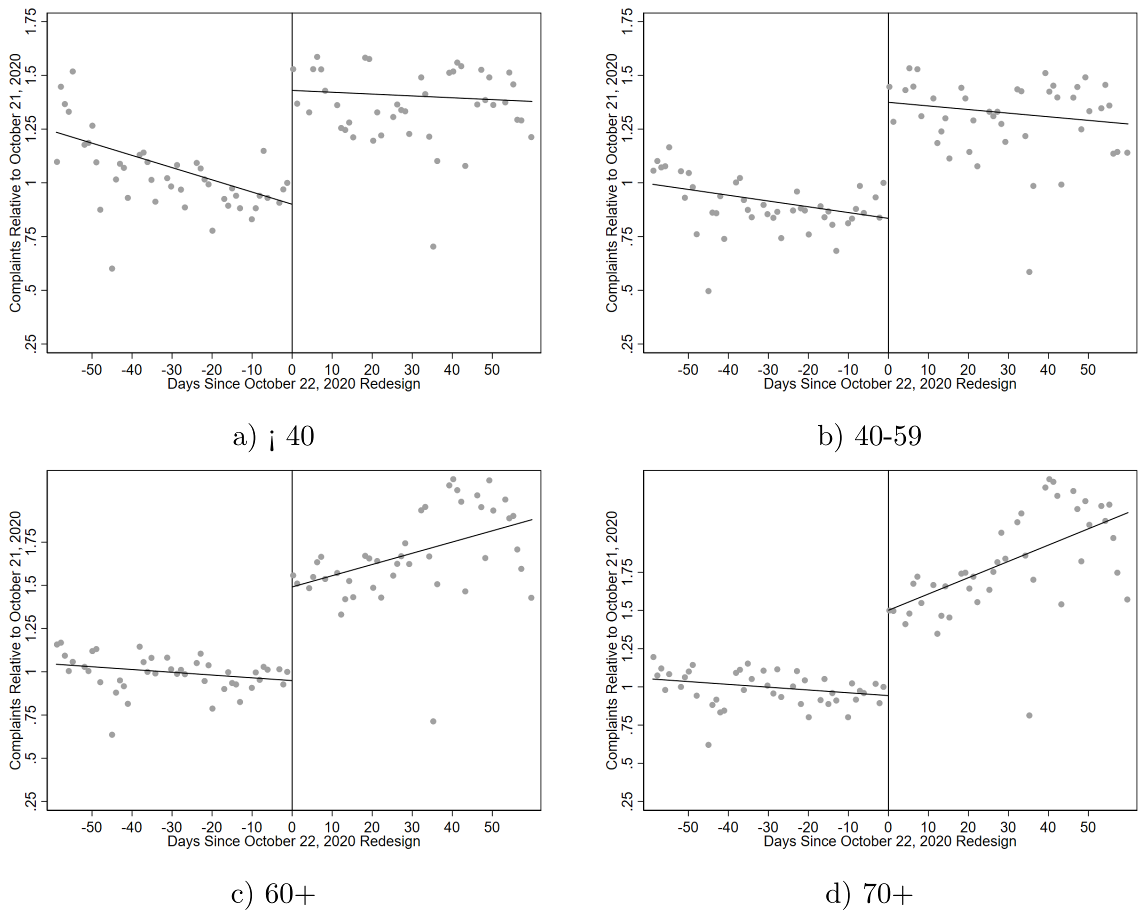

Notes: The figure shows the daily number of complaints report to the FTC by age group, from 60 days before and after the FTC's website redesign on October 22, 2020. For each panel, the number of complaints are expressed relative to the number of complaints on the day prior to the redesign, October 21, 2020, which are set to one. The vertical bar shows the date of the redesign. The fitted lines are an RD estimate that includes a first-degree polynomial and controls for the day of the week.

E.2 Race/Ethnicity

We next look at the racial and ethnic composition of complaining consumers. Raval (2020) found that consumers in both heavily Black and heavily Latino areas affected by fraud were less likely to complain than consumers in white areas, and posited that this difference in complaining was due to social alienation of non-white consumers.[1]

Because the Consumer Sentinel data do not explicitly ask for self-identified race or ethnicity, we impute race and ethnicity using consumer first and last names in a method analogous to the Bayesian Improved Surname Geocoding (BISG) (Consumer Financial Protection Bureau, 2014; Zhang, 2018). The availability of name data did not change before and after the redesign, because a non-missing name entry is required to complete the complaint. We match surnames to data from the Census on the distribution of race and ethnicity for more than 150,000 surnames. We also match first names to data from the Home Mortgage Disclosure Act (HMDA) on the distribution for more than 4,200 first names (Voicu, 2018).

For consumers with only first or last name matched to the Census or HMDA data, we calculated the probability that a consumer was a given race based upon the probability that their first or last name occurred in the population. For consumers with first and last name available, we used Bayes' Rule, with the probability an individual is of a particular race or ethnicity r given their first name f and last name s as:

[1] Sweeting et al. (2020) provides a summary of this work.

$$Pr(r|f,s) = \frac{p(r|f) * q(s|r)}{\sum_{r \in R}(p * q)}, \tag{2}$$

where $p(r|f)$ is the share of individuals in the HMDA data with that first name who are of that race, and $q(s|r)$ is the share of that race who has the surname.[2] We convert the resulting probabilities to proxy race and ethnicity by assigning the consumer a race according to their highest probability (Zhang, 2018).

Table OA20 shows how the number of complaints in each imputed race and ethnicity changed with the redesign, expressed relative to the overall increase in complaints. We find substantial increases in complaints among all groups. We do not, however, find disproportionate increases amongst Black and Latino consumers; that is, the website redesign did not affect pre-existing *disparities* in complaining. For all online complaints, we find a 31% increase from white consumers, compared to a 27% increase for Black consumers, 24% for Latino consumers, and 34% for Asian consumers. For mobile complaints, we find a much smaller increase for Latino consumers than other groups; mobile complaints from Latino consumers increase by only 21% after the redesign, compared to 31% for white consumers, 29% for Black consumers, and 39% for Asian consumers.

Table OA20: RD Estimates, Log Complaints in Each Imputed Race

	(1) White	(2) Black	(3) Latino	(4) Asian
A. FTC Online				
RD Estimate	0.311***	0.274***	0.241***	0.342***
	(0.0413)	(0.0360)	(0.0416)	(0.0404)
B. FTC Mobile				
RD Estimate	0.312***	0.290***	0.205***	0.390***
	(0.0439)	(0.0376)	(0.0469)	(0.0500)
C. FTC Desktop				
RD Estimate	0.309***	0.259***	0.278***	0.295***
	(0.0387)	(0.0402)	(0.0416)	(0.0517)

Notes: The table shows of equation (1), where the dependent variable is the fraction of daily complaints by consumers in each imputed race or ethnicity category. Race and ethnicity are imputed using a Maximum A Posteriori proxy based on the consumer's first and last name. Robust standard errors clustered at the daily level. $*p < 0.05, **p < 0.01, ***p < 0.001$

Our main race imputation uses consumer first and last names. However, race imputation using Bayesian techniques usually rely on geographic information as well (Voicu, 2018). We use the fraction of people with each first and last name who are of each race or ethnicity, as in the main results, but also include geographic data by zipcode from the 2010 Census. We calculate

$$p(r|s,f,g) = \frac{p(r|s)p(f|r)p(g|r)}{\sum_{r=1}^{6} p(r|s)p(f|r)p(g|r)}, \tag{3}$$

where $p(r|s,f,g)$ is the imputed probability of being of race or ethnicity r, given surname s, first name f, and geographic area g. For consumers who do not provide a zipcode, we use just the first and last name probabilities.

In Table OA21, we examine an analogous estimate where we also include the share of each race in the consumer's zip code in the 2010 Census to estimate the race imputation; complaints from white consumers rise by 32%, Black consumers by 20%, Latino consumers by 22%, and Asian consumers by 33%. These estimates are broadly similar to those in Table OA20, although we find larger gaps in the effect of the redesign between Black and Latino consumers compared to white and Asian consumers after using zip code for race imputation.

[2] Overall, 20% of complaints did not match to the first name data, and 14% did not match to the surname data. Only 6% did not match to either.

Table OA21: RD Estimates, Log Complaints in Each Imputed Race, Using Zipcode

	(1) White	(2) Black	(3) Latino	(4) Asian
A. FTC Online				
RD Estimate	0.317***	0.200***	0.216***	0.331***
	(0.0403)	(0.0432)	(0.0432)	(0.0426)
B. FTC Mobile				
RD Estimate	0.315***	0.279***	0.191***	0.381***
	(0.0425)	(0.0465)	(0.0478)	(0.0551)
C. FTC Desktop				
RD Estimate	0.319***	0.122*	0.242***	0.280***
	(0.0379)	(0.0559)	(0.0432)	(0.0519)

Notes: The table shows of equation (1), where the dependent variable is the fraction of daily complaints by consumers in each imputed race or ethnicity category. Race and ethnicity are imputed as the highest probability race or ethnicity group using posterior probabilities based on the consumer's first name, last name, and zip code. Robust standard errors clustered at the daily level. $*p < 0.05, **p < 0.01, ***p < 0.001$

E.3 Sex

We also use an imputation technique to assess whether there were differential effects by sex. The SSA publishes counts of sex assigned at birth by name going back over 100 years for names with at least five occurrences in a state. We aggregate these counts to calculate the fraction of occurrences for each name assigned to male and female individuals. We then assign a name a sex if at least 80 percent of occurences are of that sex, which we can do for 97 percent of all names in the SSA data.

Table OA22 shows estimates of the effect of the website redesign on the complaints from consumers who were imputed to be male and female.[3] Complaints increase more after the redesign for female consumers (32%) than male consumers (28%). However, we cannot reject that the increase for both female and male consumers is the same as the overall effect across all consumers.

[3]We report complaints for both sexes because sex is not possible to impute for a small minority of complaints, from consumers with rare names not listed in the SSA data, or names that we have not categorized as either male or female.

Table OA22: RD Estimates, Log Complaints in Each Imputed Sex

	(1) Female	(2) Male
A. FTC Online		
RD Estimate	0.316***	0.277***
	(0.0409)	(0.0387)
B. FTC Mobile		
RD Estimate	0.305***	0.271***
	(0.0433)	(0.0421)
C. FTC Desktop		
RD Estimate	0.328***	0.284***
	(0.0395)	(0.0362)

Notes: The table shows of equation (1), where the dependent variable is the fraction of daily complaints by consumers in each imputed sex. Sex is imputed using name counts from the Social Security Administration. Robust standard errors clustered at the daily level. $*p < 0.05, **p < 0.01, ***p < 0.001$

F Derivation of Complier Mean

Let Y be the variable of interest, $E(Y_{post})$ be the mean for Y after the redesign, and $E(Y_{pre})$ be the mean for Y before the redesign. N_{pre} and N_{post} are the number of complaints before and after the redesign.

Our object of interest is $E(Y_{complier})$, which is the mean of Y for compliers, consumers who are induced to complain because of the redesign. We assume that there are no defiers, so before the redesign all consumers are always takers and after the redesign consumers are either takers or compliers:

$$\begin{aligned} N_{post} &= N_{complier} + N_{taker} \\ N_{pre} &= N_{taker}. \end{aligned}$$

Our RD estimate (in percentages) γ for the total number of complaints identifies the percentage change in complaints:

$$1 + \gamma = \frac{N_{post}}{N_{pre}} = \frac{N_{complier} + N_{taker}}{N_{taker}}.$$

Our RD estimate (in levels) Δ for the change in the mean of Y identifies the mean change in Y from the redesign:

$$\Delta = E(Y_{post}) - E(Y_{pre}).$$

Finally, we can identify the mean for takers from the period before the redesign: $E(Y_{taker}) = E(Y_{pre})$.
We can then identify $E(Y_{complier})$ by rearranging the expression for $E(Y_{post})$ in terms of γ, Δ, and $E(Y_{complier})$. By definition,

$$E(Y_{post}) = \frac{E(Y_{complier})N_{complier} + E(Y_{taker})N_{taker}}{N_{complier} + N_{taker}}.$$

Using the definition of γ:

$$E(Y_{post}) = \frac{\gamma}{1+\gamma}E(Y_{complier}) + \frac{1}{1+\gamma}E(Y_{taker}).$$

Subtracting $E(Y_{taker})$ from both sides:

$$\begin{aligned} E(Y_{post}) - E(Y_{pre}) &= \frac{\gamma}{1+\gamma}E(Y_{complier}) - \frac{\gamma}{1+\gamma}E(Y_{taker}) \\ E(Y_{post}) - E(Y_{pre}) + \frac{\gamma}{1+\gamma}E(Y_{taker}) &= \frac{\gamma}{1+\gamma}E(Y_{complier}) \\ (E(Y_{post}) - E(Y_{pre}))\frac{1+\gamma}{\gamma} + E(Y_{taker}) &= E(Y_{complier}). \end{aligned}$$

The last expression is just:

$$E(Y_{complier}) = \Delta\frac{1+\gamma}{\gamma} + E(Y_{taker}).$$

In some cases, we estimate the RD effect in percentage terms rather than in level terms. That is, we estimate RD effect $1+\delta$, where:

$$1 + \delta = \frac{E(Y_{post}) - E(Y_{pre})}{E(Y_{pre})}.$$

In that case, the expression for the complier mean is:

$$E(Y_{complier}) = (1+\delta)E(Y_{taker})\frac{1+\gamma}{\gamma} + E(Y_{taker}).$$

References

Consumer Financial Protection Bureau. 2014. “Using publicly available information to proxy for unidentified race and ethnicity.”

DeLiema, Marguerite, Martha Deevy, Annamaria Lusardi, and Olivia S Mitchell. 2020. “Financial fraud among older Americans: Evidence and implications.” *The Journals of Gerontology: Series B*, 75(4): 861–868.

Federal Trade Commission. 2022. “Who experiences scams? A story for all ages.” Consumer Protection Data Spotlight.

Grootendorst, Maarten. 2022. “BERTopic: Neural topic modeling with a class-based TF-IDF procedure.” *arXiv preprint arXiv:2203.05794*.

Mclean, Bethany. 2020. “How International Fraud Rings Operate and Target Older Americans.” *AARP*.

Pei, Zhuan, David S Lee, David Card, and Andrea Weber. 2022. “Local polynomial order in regression discontinuity designs.” *Journal of Business & Economic Statistics*, 40(3): 1259–1267.

Raval, Devesh. 2020. “Whose voice do we hear in the marketplace? Evidence from consumer complaining behavior.” *Marketing Science*, 39(1): 168–187.

Sweeting, Andrew, David J Balan, Nicholas Kreisle, Matthew T Panhans, and Devesh Raval. 2020. “Economics at the FTC: fertilizer, consumer complaints, and private label cereal.” *Review of Industrial Organization*, 57: 751–781.

Voicu, Ioan. 2018. “Using first name information to improve race and ethnicity classification.” *Statistics and Public Policy*, 5(1): 1–13.

Zhang, Yan. 2018. “Assessing fair lending risks using race/ethnicity proxies.” *Management Science*, 64(1): 178–197.