

DO SEARCH ENGINES INFLUENCE MEDIA PIRACY? EVIDENCE FROM A RANDOMIZED FIELD STUDY

Liron Sivan, Michael D. Smith, and Rahul Telang

Heinz College, Carnegie Mellon University, Pittsburgh, PA 15213 U.S.A.
{sivan@andrew.cmu.edu} {mds@cmu.edu} {rtelang@andrew.cmu.edu}

Appendix A

Experiment 1: Empirical Approach

Endowment Effect

The experimental design we present in the paper is similar to that used previously by Tsai et al. (2011), which conducted a randomized experiment in which participants were solicited to “test a new search engine interface” (p. 259) and were paid \$45 for their participation. As in our study, Tsai et al.’s study asked its participants to use a search engine to find several products and to purchase them, allowing the participants to keep both the products and any money left over after the purchases were made. This design, in the words of the authors, “created a price incentive, encouraging participants to purchase from merchants with lower prices” (p. 262). Providing compensation to experimental participants is a common and accepted practice in the experimental economics, marketing, and information systems literatures (see, for example, Dhar et al. 2007; Haisley et al. 2008, Isaac and Davis 2006; Haws and Winterich 2013; Löschela et al. 2013; Rucker et al. 2011; Suk et al. 2012). As it pertains to our experimental design, it is important to note that although providing subjects with an unrestricted endowment of money may change the overall likelihood of any user making a purchase, it will do so equally for both the treated and control group users such that any difference in response between these two groups can reliably be attributed to the experimental intervention. Further, the fact that participants can retain the money after the experiment has ended, and use the money for any purpose they desire, will strengthen their perceived ownership of it and will do so equally for both the treated and control group users. Again, this means that any difference in response between these two groups can reliably be attributed to the experimental intervention.

We also note that our results in experiment 1, where our users are given a \$20 gift card, are consistent with results in experiment 3, where Amazon Mechanical Turk users are only compensated by \$1.50 for their time, further suggesting that the differences in responses between the control and treatment groups are not driven by an endowment effect.

Recall Bias

It is important to note that we can only directly observe user behavior while they are on our search engine: we do not observe their behavior outside of the search engine. Thus, our observation of where the user obtained their content is based on their survey answers. While this may introduce some recall bias, the degree of recall bias should not vary across the control and treatment conditions and the survey occurs immediately after the participant obtains the content, significantly reducing the overall possibility of recall bias. It is also possible, based on prior studies that found survey respondents are likely to underreport socially undesirable activities, that participants who choose a pirated option could be less likely to reveal that in their survey answers (Cannell et al. 1965; Means et al. 1992; Warner 1978; Wyner 1980). However, again, the propensity to misreport should not vary across the control and treatment conditions, and any underreporting bias would lead to an underestimate of the degree that infringing links induce more piracy. We also note that we can use observed search and clicking behavior to validate a user’s survey answers (including in some cases verifying price). In the results section we show that the users’ survey responses were consistent with their observed behavior.

Dropout Rates

In order to ensure that participants understood the task correctly, we included in our analysis only those users who claimed the \$20 virtual Visa card, used our custom-built search engine, and who reported the source from which they acquired the movie and its price. Approximately 1,000 participants were invited to participate in the study. Of these participants, 770 met the initial qualifications for participating in the study by expressing an interest in watching one of the movies on our list. These participants were invited to participate in our study and 632 of them logged into the system and were presented with the task details. Out of these participants, 235 completed the task as instructed, and 196 qualified for our study by also completing the post-experiment questionnaire. The dropout rates across each of the different stages described above are not statistically different across experimental conditions.

Comparison of User Characteristics and Attitudes toward Piracy across Treatment Conditions

Prior to analyzing the results of the experiment, we confirmed that the distribution of user characteristics (demographic characteristics, attitudes toward piracy) is similar across the control and treatment groups. The average values for each experimental condition are presented in Table A1. Chi squared tests show that there are no statistically significant differences in demographic characteristics or attitudes toward piracy between the experimental conditions (p-value > 0.05). This confirms that our randomization worked as intended.

	% Women	Average Age Group	Average Household Size	Average Household Income Group	Average Attitude Against Piracy (Likert Scale)	% Downloads Infringing
Condition 1: No manipulation	70.0%	4.20	2.85	2.82	4.60	33.3%
Condition 2: Legal content manipulation	59.2%	4.30	2.85	2.77	4.75	39.4%
Condition 3: Infringing content manipulation	64.6%	3.97	2.68	3.02	4.58	44.6%

Note: There are 7 possible age group values in the questionnaire: 1 (18–21), 2 (22–25), ..., 7 (61 and over); 10 possible household size values: 1, 2, ..., 9, 10 or more; and 6 possible household income group values: 1 (less than \$30,000), 2 (\$30,000– \$50,000), ..., 6 (over \$150,000).

To further verify similarity in preexisting attitudes toward piracy, we tested whether the participants’ initial intent for pirate or legal content is similar across the three experimental conditions. We did this by classifying whether their initial search terms reflected neutral, legal, or infringing intent according to the degree to which pirate or legal links were present in the (unmodified) search results for similar queries commonly issued by our users.

We discovered that search terms using only the movie’s name contained almost exclusively “neutral” results (i.e., results that neither promote legal or pirate sources), and thus we classify these searches as neutral. However, when search terms included the words “buy,” “rent,” or “purchase,” the search results contained 38% more legal links than pirate links, and when the search term contained a legal domain name (e.g., Amazon), the search results contained 78% more legal links than pirate links. Thus, we classify these search terms as representing “legal” intent. Conversely, when search terms included the words “download,” “stream,” or “full movie,” there were 33% more pirate links in the search results than legal links, and including the domain name of an infringing site (e.g., piratebay) resulted in search results that included 65% more pirate links than legal links. Because of this, we classify these search terms as representing “infringing” intent.

We then classify intent based on the initial intent reflected in each user’s search terms (or neutral if the user did not express intent in their searches). As above, reassuringly the distributions are not statistically different across the control and treatment conditions (see Table A2 for frequencies across conditions).

Table A2. Between-Conditions Comparison of the Initial Intent (based on the first keyword each user entered)

	N	Neutral Searches	Legal Intent	Infringing Intent
Condition 1: No manipulation	60	17	31	12
Condition 2: Legal content manipulation	68	28	23	17
Condition 3: Infringing content manipulation	65	21	33	11

Finally, we note that although the characteristics of users who participated in the experiment were similar across the control and treatment conditions, it is possible that our participant pool skews toward being more media or tech savvy than the general population. As such, in interpreting our results, one should focus on the difference between sales/piracy in the control and treatment conditions as opposed to the absolute levels of sales/piracy within any particular condition.

Appendix B

Logistic Regression

Although the tests reported in the body of the paper are sufficient to determine if there are differences between the control and treatment groups based on our experimental manipulations, we can also use a logistic regression model to control for and analyze differences between groups based on observed characteristics.

Specifically, we use the following logistic regression model to control for observable participant characteristics:

$$\log \frac{PR(Legal_i)}{1 - PR(Legal_i)} = \alpha + \beta_1 \cdot NI_i + \beta_2 \cdot I_i + \sum_{j=3}^5 \beta_j \cdot DC_i^j + \sum_{j=6}^8 \beta_j \cdot MCP_i^j + \sum_{j=9}^{10} \beta_j \cdot ATP_i^j + \beta_{11} \cdot intent + \varepsilon_i \quad (3)$$

where $Legal_i$ denotes whether the movie was acquired from a non-infringing source; NI_i is an indicator variable denoting whether participant i was treated with the non-infringing search condition; I_i is an indicator variable denoting whether participant i was treated with the infringing search condition; $\sum_{j=3}^5 \beta_j \cdot DC_i^j$ includes the following demographic characteristics: Gender (an indicator variable for whether the participant was a woman), Age ≤ 40 , Household size, and Income; $\sum_{j=6}^8 \beta_j \cdot MCP_i^j$ includes the following media consumption preferences: Time online (the average hours spent online per day), Acquires movies online (whether the participant ever downloaded or streamed a movie, including pirated movies), and Movies per year (the number of movies the participant watched in the last 12 months); $\sum_{j=9}^{10} \beta_j \cdot ATP_i^j$ includes the following attitude toward piracy variables: Against piracy (on a six-point Likert Scale ranging from 1, "There is nothing wrong with it," to 6, "It is the same as stealing"), Downloads infringing (whether the participant indicated that s/he uses torrents or other free online downloads/streaming/file-sharing); and *Non-infringing intent* _{i} (whether the first search term that the user entered indicates that his/her intent is to acquire a legal copy). We present the results of this model in Table B1.

These results are consistent with our means comparison results in that the treatment variables are statistically different from the control condition and that they have the expected sign (the legal treatment condition increases the likelihood of purchasing legally and the infringing treatment condition decreases the likelihood of purchasing legally). These results also confirm, as one would expect, that participants who consumed infringing content in the past are less likely to purchase the movie legally and that those who use a search term that implies legal intent are more likely to purchase legally.

Table B1. Logistic Regression Results for Equation (3)					
Dependent Variable: Acquired Legally	(1)	(2)	(3)	(4)	(5)
	Basic Model	Including Demographic Characteristics	Including Media Consumption Preferences	Including Attitude Towards Piracy	Including Intent to Acquire Legally
Constant	1.386*** (0.323)	1.637** (0.663)	-0.333 (1.288)	0.148 (1.511)	0.0575 (1.533)
Non-infringing mode	1.432** (0.608)	1.458** (0.613)	1.606** (0.634)	1.855*** (0.673)	1.905*** (0.685)
Infringing mode	-1.108*** (0.409)	-1.152*** (0.422)	-1.232*** (0.438)	-1.235** (0.507)	-1.261** (0.516)
Woman		-0.00896 (0.406)	0.0163 (0.416)	0.0167 (0.482)	0.0421 (0.491)
Age ≤ 40		-0.603 (0.398)	-0.744* (0.436)	-0.669 (0.491)	-0.492 (0.500)
Household size		-0.134 (0.146)	-0.216 (0.151)	-0.0772 (0.186)	-0.106 (0.193)
Income		0.126 (0.142)	0.132 (0.147)	-0.0719 (0.166)	-0.0333 (0.173)
Time online			0.0287 (0.213)	0.235 (0.249)	0.169 (0.250)
Acquires movies online			0.191 (0.231)	0.453* (0.275)	0.412 (0.282)
Movies per year			0.409** (0.178)	0.229 (0.210)	0.255 (0.213)
Against piracy				0.554 (0.486)	0.404 (0.504)
Downloads infringing				-2.424*** (0.516)	-2.410*** (0.523)
Non-infringing intent					1.266** (0.621)
Number of obs.	196	196	196	196	196
Pseudo R ²	0.1392	0.1611	0.1933	0.3529	0.3758

Notes: Standard errors in parentheses; ***p < 0.01, **p < 0.05, *p < 0.1

To see the effect of intent on user behavior, we use the method described in Appendix C to split the data based on user intent (following Question 15f in our survey) and rerun the regression. While the number of observations in the infringing intent are small, the results shown in Table B2 using the survey question provide the same insight as those presented in the body of the paper: Even users with infringing intent are affected by legal mode and vice versa.

Table B2. Logistic Regression Results for Equation (3) by Intent		
	(1)	(2)
	15F Legal Intent - % Legal b/se	15F Infringing Intent - % Legal b/se
Constant	-1.591 (1.809)	11.120 (8.246)
Non-infringing mode	2.228* (0.922)	1.795 (1.293)
Infringing mode	-1.334* (0.599)	-1.172 (1.711)
Woman	-0.324 (0.584)	1.329 (1.357)
Age ≤ 40	-0.999 (0.619)	-0.119 (1.213)
Household size	0.069 (0.224)	-0.485 (0.438)
Income	-0.051 (0.198)	0.067 (0.407)
Time online	0.503+ (0.305)	-0.507 (0.671)
Acquires movies online	0.531+ (0.306)	0.000 (.)
Movies per year	0.398+ (0.236)	-1.841 (1.577)
Against piracy	0.370 (0.572)	1.545 (1.700)
Downloads infringing	-2.357** (0.583)	0.000 (.)
Number of obs.	165	28
Pseudo R ²	0.3957	0.3152

Note: Standard errors in parentheses; ***p < 0.01, **p < 0.05, *p < 0.1

Appendix C

Classifying Legal and Infringing Intent Using the Post-Experiment Questionnaire

In the body of the paper, we analyzed the responses of users who, based on their keyword choices, expressed the intent to consume legal or infringing content. Another way to classify intent is based on a user’s prior experience with consuming pirated or legal content. To do this, we use question 15f in our post-experiment survey of users. This question asks users, “Of all the movies you watched, can you provide some details on what channel did you use to watch them?” One of the options was “Torrents (or other free online downloads/stream, file-sharing [channels]).” Users were able to choose “0–1 movies,” “2–5 movies,” “5–10 movies,” “10–15 movies,” or “15+ movies” consumed through this method.

In Table C1 we replicate the results from Table 2 in the body of the paper, but in this case instead of using search queries to classify intent, we classify users who stated that they had consumed 0–1 movies through torrents or other free download/streaming channels as those that have legal intent, and users who had consumed two or more movies through torrents or other free download/streaming channels as those having infringing intent.

Post-Experiment Questionnaire	Control Group	Legal Content Manipulation	Infringing Content Manipulation
Legal Intent	43/51 (84%)	57/57 (100%)	34/57 (60%)
Infringing Intent	5/8 (62%)	10/12 (83%)	3/8 (38%)

This table shows that, consistent with the results presented in the body of the paper, users with legal intent are significantly less likely to consume legally when placed in the infringing content manipulation (60%) than in either the legal (100%) or control (84%) conditions (t-tests: $t(56) = -5.34, p < 0.01$; $t(50) = -2.75, p < 0.001$, respectively). Likewise, users with infringing intent are significantly more likely to purchase legally in the legal content manipulation (83%) than in the infringing (38%) or control (62%) conditions (t-tests: $t(7) = 2.07, p < 0.05$; $t(7) = 1.057, p = 0.16$, respectively). These results confirm the results from the body of the paper that the ranking of pirated and legal search results matters even among users with an initial preference for pirate or legal channels.

As with the results in the body of the paper, we can use our clickstream data to test whether, when using the post-experiment questionnaire to classify user intent, users with low infringing intent search more when they are placed in the infringing treatment condition than they do in the control condition and likewise for users with high infringing intent who are placed in the legal treatment condition.

Consistent with the results presented in the body of the paper, our clickstream data suggest that users with high piracy intent search more when placed in the legal treatment condition than in the control condition and users with low piracy intent search more when placed in the piracy treatment condition than in the control condition. Specifically, the post-experiment survey data show that consumers with low piracy intent initiate on average 2.39 searches in the control condition ($n = 51$) and 3.42 searches in the infringing treatment condition ($n = 57$) (statistically significant at the 1% level). Likewise consumers with high piracy intent initiate, on average, 2.00 searches (9) in the control condition and 3.33 searches (12) in the legal treatment condition (insignificant). When we consider the average position of search instead of the number of searches, consumers with low piracy intent click on results at an average position of 4.21 (50) in the control condition and 8.45 (47) in the infringing treatment condition (1%), and consumers with high piracy intent click on results at an average position of 2.74 (9) in the control condition and 3.54 (12) in the legal treatment condition (insignificant). These results also hold if we consider the maximum position that each consumer clicks on during their searches. Averaging across consumers, consumers with low piracy intent search to a maximum position of 6.02 (50) in the control condition, and to a maximum position of 13.49 (47) in the infringing treatment condition (1%). Consumers with high piracy intent search to a maximum position of 4.22 (9) in the control condition and a maximum position of 6.58 (12) in the legal treatment condition (statistically insignificant).

Appendix D

Experiment 2: Empirical Approach and Selected Results

Data and Results

Following our approach in experiment 1, in our second experiment we included in our analysis only those users who claimed the \$20 virtual Visa card, used our custom-built search engine for at least one search, and reported the source from which they acquired the movie and how much they paid for it. A total of 234 participants qualified for our study.¹ In Table D1 we report average statistics for these participants across the treatment conditions. These statistics show that the distribution of demographic characteristics and attitudes toward piracy are similar across the control and treatment groups, as one would expect given the experimental assignment.

	% Women	Average Age Group	Average Household Size	Average Household Income Group	Average Attitude Against Piracy (Likert Scale)	% Downloads Infringing
Condition 1: No manipulation	57.14%	1.69	3.02	2.60	3.40	80.95%
Condition 2: Mild legal content manipulation	73.47%	1.71	3.18	2.35	3.73	69.39%
Condition 2a: Legal content manipulation	62.50%	1.75	2.79	2.13	3.85	62.50%
Condition 3: Mild infringing content manipulation	54.55%	1.68	3.41	2.18	3.66	75.00%
Condition 3a: Infringing content manipulation	66.67%	1.73	3.29	2.31	3.76	64.71%

Note: There are 7 possible age group values in the questionnaire: 1 (1–21), 2 (22–25), ..., 7 (61 and over); 10 possible household size values: 1, 2, ..., 9, 10 or more; and 6 possible household income group values: 1 (less than \$30,000), 2 (\$30,000–\$50,000), ..., 6 (over \$150,000).

As in experiment 1 (and using the method described in Appendix A), we compare the initial search intent expressed by users across the different treatment conditions (Table D2), finding no significant differences in expressed intent across conditions.

¹Approximately 650 participants were invited to participate in the study and expressed an interest in watching one of the movies in the study. A total of 550 participants logged into the system and were presented with the task details. Out of these participants, 270 completed the experimental task as instructed, and of these participants 234 (86 men, 148 women) qualified for our study by also completing the post-experiment questionnaire. The dropout rates in each of the different stages described above are not statistically different across experimental conditions.

Table D2. Between-Conditions Comparison of the Initial Intent (based on the first keyword each user entered)

	N	Neutral Searches	Legal Intent	Infringing Intent
Condition 1: No manipulation	42	6	23	13
Condition 2: Mild legal content manipulation	49	16	19	14
Condition 2a: Legal content manipulation	48	14	20	14
Condition 3: Mild infringing content manipulation	44	11	18	15
Condition 3a: Infringing content manipulation	51	15	27	9

Logistic Regression

We use a logistic regression model to control for and analyze differences between groups based on observed characteristics. This model is similar to the model from experiment 1, except that it includes dummy variables for the additional treatment conditions.

Specifically, we use the following logistic regression model to control for observable participant characteristics:

$$\log \frac{PR(legal_i)}{1 - PR(legal_i)} = \alpha + \beta_1 \cdot MNI_i + \beta_2 \cdot NI_i + \beta_3 \cdot MI_i + \beta_4 \cdot I_i + \sum_{j=5}^7 \beta_j \cdot DC_i^j + \sum_{j=8}^{10} \beta_j \cdot MCP_i^j + \sum_{j=11}^{12} \beta_j \cdot ATP_i^j + \beta_{13} \cdot intent + \varepsilon_i \tag{5}$$

where MNI_i is an indicator variable denoting whether participant i was treated with the mild non-infringing search condition; MI_i is an indicator variable denoting whether participant i was treated with the mild infringing search condition; and the other variables are the same as before. We present the results of this model in Table D3.

The results in Table D3 are consistent with our means comparison results in that the intense treatment variables are statistically different from the control condition and that all treatment variables have the expected signs (the non-infringing treatment condition increases the likelihood of purchasing legally and the infringing treatment condition decreases the likelihood of purchasing legally). As with experiment 1, the results show that participants who consumed infringing content in the past are less likely to purchase the movie legally, and that those who use a search term that reveals non-infringing intent are more likely to purchase legally.

Table D3. Logistic Regression Results					
	(1)	(2)	(3)	(4)	(5)
Dependent Variable: Acquired Legally	Basic Model	Including Demographic Characteristics	Including Media Consumption Preferences	Including Attitude Toward Piracy	Including Intent to Acquire Legally
Constant	0.486 (0.318)	0.772 (0.529)	3.077** (1.215)	3.265** (1.295)	2.932** (1.430)
Mild non-infringing mode	0.641 (0.460)	0.741 (0.470)	0.733 (0.479)	0.694 (0.484)	1.046** (0.529)
Intense non-infringing mode	1.912*** (0.611)	1.975*** (0.620)	1.977*** (0.628)	1.942*** (0.641)	2.321*** (0.671)
Mild infringing mode	-0.576 (0.438)	-0.644 (0.452)	-0.700 (0.462)	-0.732 (0.469)	-0.727 (0.528)
Intense infringing mode	-0.924** (0.428)	-0.932** (0.438)	-1.152** (0.469)	-1.228** (0.477)	-1.280** (0.530)
Woman		-0.626* (0.321)	-0.824** (0.341)	-0.940*** (0.350)	-1.178*** (0.390)
Younger (Age < 22)		-0.462 (0.338)	-0.593* (0.349)	-0.518 (0.356)	-0.383 (0.394)
Household size		0.0843 (0.0977)	0.0831 (0.1000)	0.0409 (0.102)	-0.00624 (0.109)
Income		-0.0100 (0.118)	-0.0267 (0.122)	-0.0174 (0.123)	-0.119 (0.136)
Time online			-0.317* (0.187)	-0.319* (0.190)	-0.229 (0.208)
Acquires movies online			-0.442* (0.264)	-0.209 (0.284)	-0.324 (0.302)
Movies per year			0.0780 (0.146)	0.0928 (0.150)	0.132 (0.165)
Against piracy				0.148 (0.410)	0.425 (0.448)
Downloads infringing				-0.968** (0.416)	-0.889* (0.458)
Non-infringing intent					3.127*** (0.693)
Number of obs.	234	234	234	234	234
Pseudo R ²	0.1321	0.1519	0.1755	0.1987	0.3107

Notes: Standard errors in parentheses; ***p < 0.01, **p < 0.05, *p < 0.1

Last, we use our post-experiment questionnaire to classify infringing intent as a robustness test for the classification based on initial query words used in the body of the paper. As with experiment 1, to conduct this classification we use question 15f in our post-experiment survey of users and classify users who stated that they had consumed 0–1 movies through torrents or other free download/streaming channels as those that have legal intent and users who had consumed two or more movies through torrents or other free download/streaming channels as those having infringing intent.

Table D4. Legal Purchase Rates across Treatment Conditions and Initial Intent			
Post-Experiment Questionnaire	Control Group	Legal Content Manipulation	Infringing Content Manipulation
Legal Intent	16/19 (84%)	32/32 (100%)	15/35 (43%)
Infringing Intent	10/23 (43%)	12/16 (75%)	5/16 (31%)

Table D4 shows that, consistent with the results presented in the body of the paper, users with legal intent are significantly less likely to consume legally when placed in the infringing content manipulation (43%) than in either the legal (100%) or control (84%) conditions (t-tests: $t(31) = -5.01, p < 0.001$; $t(18) = -2.91, p < 0.005$, respectively). Likewise, users with infringing intent are significantly more likely to purchase legally in the legal content manipulation (75%) than in the infringing (31%) or control (43%) conditions (t-tests: $t(15) = 2.49, p < 0.05$; $t(15) = 1.98, p < 0.05$, respectively). These results confirm the results from the body of the paper that the ranking of pirated and legal search results matters even among users with an initial preference for pirate or legal channels.

Appendix E

Experiment 3: Flagging Infringing Links, Empirical Approach and Results ██████████

In the first two experiments, there was no explicit differentiation between the legal and infringing links presented to users. In this experiment we make an explicit distinction between legal and infringing links to draw a stronger differentiation between these two types of content. We, do this in part to validate our results in experiments 1 and 2, and as a partial test of whether our results are driven by factors other than user indifference between legal and infringing content.

For this experiment, we recruited 666 participants from Amazon Mechanical Turk (AMT). We included the three treatment conditions that were used in Experiment 1 and added two additional treatment conditions in order to understand what drives our previous results. Similar to our previous experiments before the experiment started, each participant was randomly assigned to one of the following five search treatment conditions: no manipulation (control), non-infringing (legal) content manipulation, infringing (piracy) content manipulation, flagging infringing results, and both flagging and promoting infringing results. We flagged infringing results (Figure E1) in a similar manner to the way Google flags results that may harm one’s computer (Figure E2), but instead of saying “This site may be hacked” as Google does, our text indicated “This webpage may contain infringing materials.”

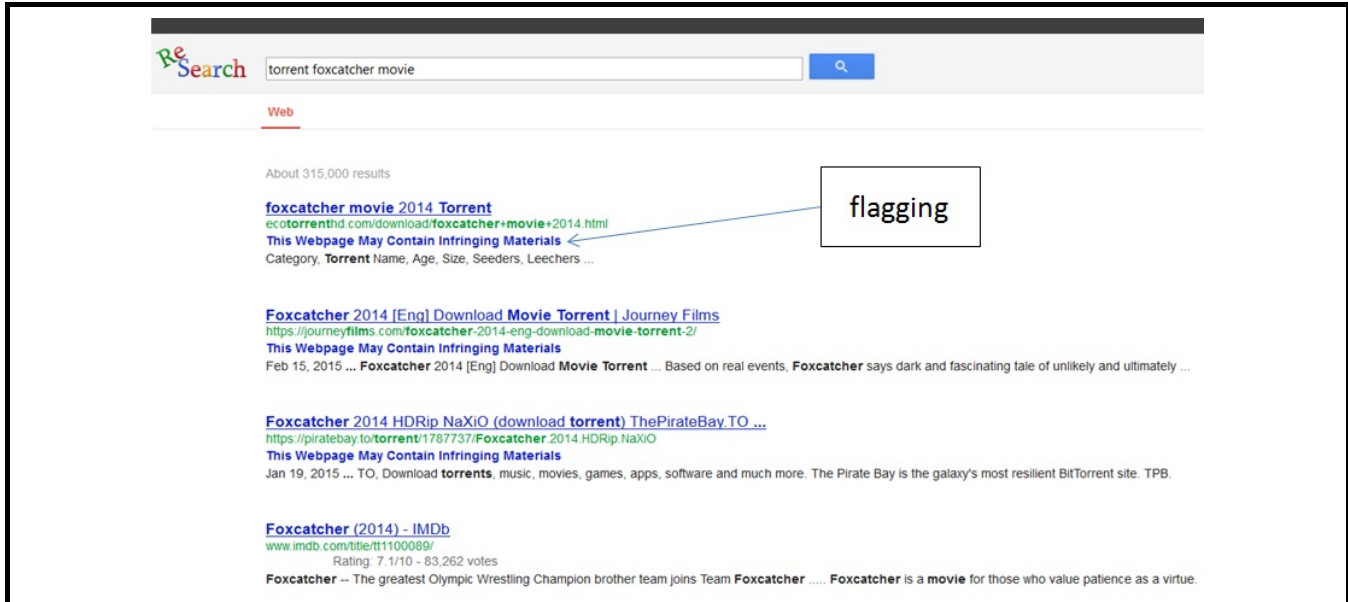


Figure E1. Flagging Infringing Results

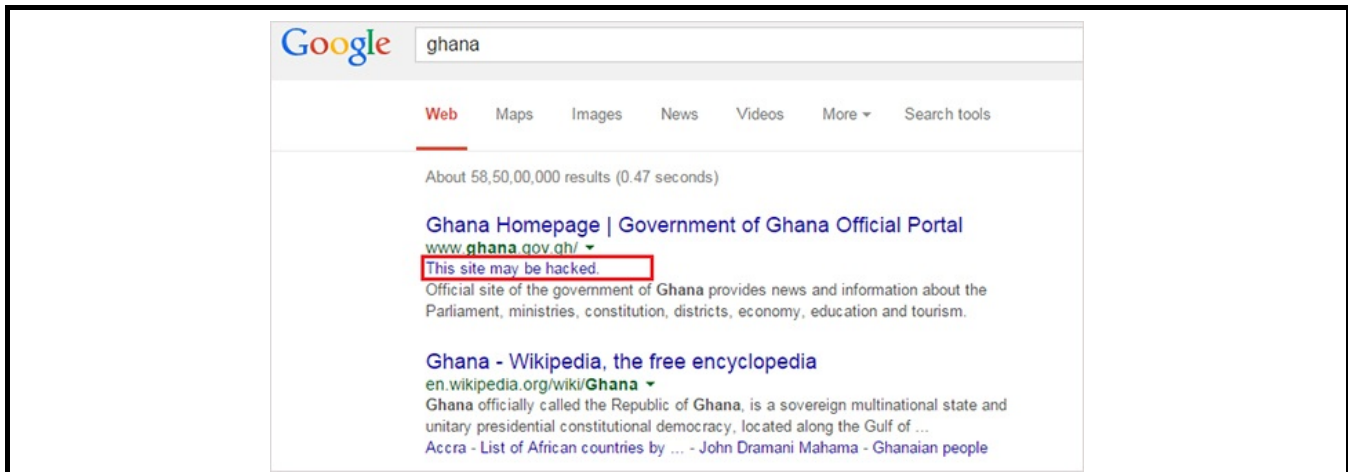


Figure E2. Google’s “This Site May Be Hacked” Warning

Experimental Design

There are three additional differences between this experiment and the two experiments in the body of the paper. First, the participants in this experiment did not receive the \$20 Visa virtual card used in the first two experiments. Instead they received \$1.50 through AMT for their participation. Second, the participants were not asked to actually acquire the movie of their choice, but to search as they normally would if they wished to watch the movie, and to find a desired source from which they would have downloaded/streamed/purchased/rented it. Third, in the post experiment questionnaire the participants were only asked to name the source they would have selected and the amount they would have spent.

Results

In Table E1, we compare the proportion of legal purchases across the different experimental conditions and search types. Consistent with experiments 1 and 2, Table E1 shows that users whose search terms express legal intent² are significantly *less* likely to purchase legally when placed in the infringing content manipulation (19%) than in either the legal (100%) or control (98.73%) conditions ($p < 0.001\%$ in both cases). Likewise, users whose search terms express infringing intent are significantly *more* likely to purchase legally in the legal content manipulation (81%) than in the infringing (15%) or control (53%) conditions ($p < 0.01\%$ in both cases).

However, we can also use these results to better understand whether user behavior changes when infringing links are highlighted in search results through the use of the flag. In this regard, Table E1 shows that the difference between the rate of legal purchases of users who have infringing intent in the control condition (53%) and in the treatment condition in which infringing results are being flagged (48%) is not statistically significant ($p > 0.3$). Meaning, although promoting legal results increased the purchase rate from 53% (control condition) to 81% (legal content manipulation), labeling infringing results without promoting them had no effect on behavior. This is also true when comparing the legal purchase rate of users who have infringing intent in the two infringing content manipulation conditions: the difference between the condition that includes flagging (17%) and the condition that does not include flagging (15%) is not statistically significant at any reasonable level of significance. Also, when comparing the legal purchase rate of users who have legal intent in the two infringing content manipulation conditions, the difference between the condition that includes flagging (27%) and the condition that does not flagging (19%), is not statistically significant ($p > 0.1$). Taken together, the results of this experiment suggest that making the infringing nature of some links more noticeable to users through the use of a flag has no impact on their propensity to consume infringing versus legal content. This, in turn, sheds light on the driver of our results. If our results were driven by a lack of perceived differentiation between legal and pirate links, then one would expect that user behavior would change if a flag were used to explicitly distinguish infringing links from other links. The fact that this does not occur, while not conclusive, is suggestive that users do perceive a difference between legal and infringing links.³ This inference is strengthened when combined with the clickstream data above suggesting that users with a stated preference for legal or infringing content search more intensely when placed in their non-preferred treatment condition. We discuss the implications of this result in more detail in the body of the paper.

Table E1. Purchase Rates across Treatment Conditions and Search Types

	N	% Legal	Legal Intent	Infringing Intent	Legal Intent - % Legal	Infringing Intent - % Legal
Condition 1: No manipulation	138	78.99%	79	59	98.73%	52.54%
Condition 1a: Flagging infringing	143	71.33%	66	77	98.48%	48.05%
Condition 2: legal content manipulation	110	94.55%	78	32	100.00%	81.25%
Condition 3: infringing content manipulation	131	17.56%	78	53	19.23%	15.09%
Condition 3a: Infringing content manipulation & Flagging infringing	144	24.31%	102	42	27.45%	16.67%

²Intent here is defined following the definition used in the body of the paper: When a user's search results produce no infringing links, we classified it as legal intent and likewise for infringing intent.

³We thank the Associate Editor for making this observation.

References

- Cannell, C. F., Fisher, G., and Bakker, T. 1965. "Reporting of Hospitalization in the Health Interview Survey," *Vital and Health Statistics*, PHS Pub. No. 1000, Ser. 2, No. 6, Washington, DC: U.S. Government Printing Office.
- Dhar, R., Huber, J., and Khan, U. 2007. "The Shopping Momentum Effect," *Journal of Marketing Research* (44:3), pp. 370-378.
- Haisley, E., Mostafa, R., and Loewenstein, G. 2008. "Myopic Risk-Seeking: The Impact of Narrow Decision Bracketing on Lottery Play," *Journal of Risk and Uncertainty* (37:1), pp. 57-75.
- Haws, K. L., and Winterich K. P. 2013. "When Value Trumps Health in a Supersized World," *Journal of Marketing* (77:3), pp. 48-64.
- Isaac, M. R., and Davis, D. D. (eds.). 2006. *Research in Experimental Economics, Volume 11: Experiments Investigating Fundraising and Charitable Contributions*, Bingley, UK: Emerald Group Publishing Limited.
- Löschela, A., Sturma, B., and Vogtc, C. 2013. "The Demand for Climate Protection—Empirical Evidence from Germany," *Economics Letters* (118:3), pp. 415-418.
- Means, B., Habina, K., Swan, G., and Jack, L. 1992. "Cognitive Research on Response Error in Survey Questions on Smoking," *Vital Health & Statistics*, Hyattsville, MD: National Center for Health Statistics, Series 6, Number 5.
- Rucker, D. D., Dubois, D., and Galinsky, A. D. 2011. "Generous Paupers and Stingy Princes: Power Drives Consumer Spending on Self Versus Others," *Journal of Consumer Research* (37:6), pp. 1015-1029.
- Suk, K., Lee, J., and Lichtenstein, D. R. 2012. "The Influence of Price Presentation Order on Consumer Choice," *Journal of Marketing Research* (49:5), pp. 708-717.
- Tsai J. Y., Egelman, S., Cranor, L., and Acquisti, A. A. 2011. "The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study," *Information Systems Research* (22:2), pp. 254-268.
- Warner, K. E. 1978. "Possible Increases in the Under-Reporting of Cigarette Consumption," *Journal of the American Statistical Association* (73), pp. 314-318.
- Wyner, G. 1980. "Response Errors in Self-Reported Number of Arrests," *Sociological Methods and Research* (9), pp. 161-177.