Fashion Crimes: Trending-Term Exploitation on the Web

Tyler Moore[1], Nektarios Leontiadis[2], Nicolas Christin[2]

Computer Science Department, Wellesley College[1]
CyLab, Carnegie Mellon University[2]

ACM Conference on Computer & Communications Security
Chicago, Illinois
October 18, 2011

Outline

Outline

1. How trending search terms are abused
   - Monetizing traffic: malware or ads?
   - Research objectives
2. Measuring trending-term abuse
   - Data collection methodology
   - Incidence of abuse
   - How search-term characteristics affect abuse prevalence
3. Economics of trending-term exploitation
   - Estimating the exposed population
   - Revenue analysis: ad abuse
   - Revenue analysis: malware
4. What happens when search engines intervene?
   - Measuring the effect of Google's intervention
   - Cautionary tale on crackdowns

Search terms can be highly dynamic

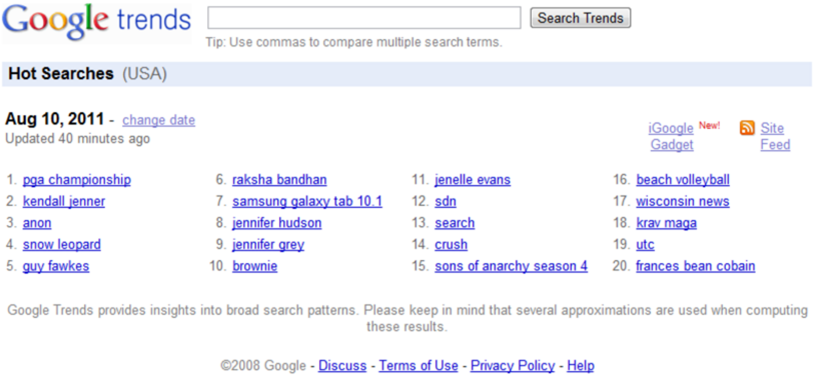

However, not all of the search results are relevant!

At best you may encounter ad-filled sites

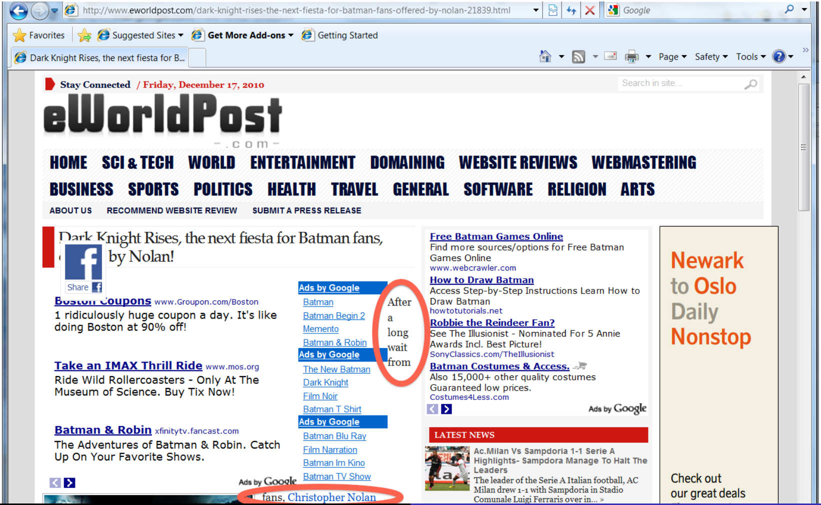

At worst you may encounter malware

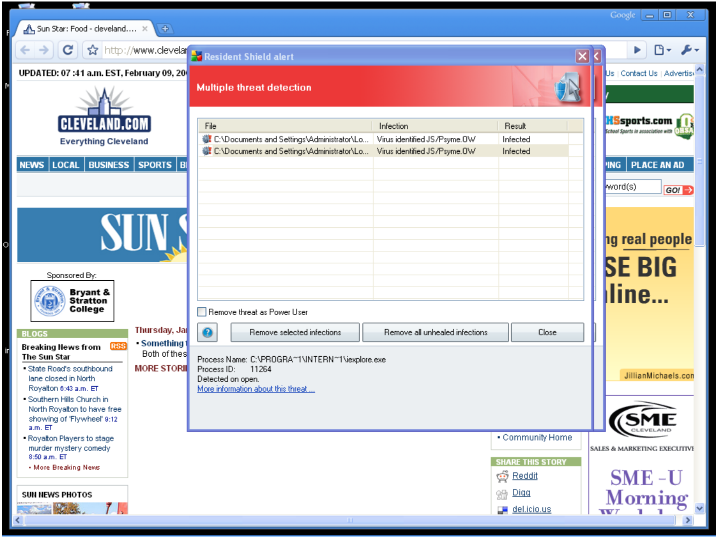

Research goals

1. Measure the prevalence of abuse in trending terms' search results relative to other terms
2. Identify whether certain types of search terms are more susceptible to abuse and why
3. Construct an economic model of revenue from malware and ads to understand the behavior of profit-minded adversaries
4. Measure the impact of a search-engine crackdown on low-quality, "made for Adsense" (MFA) sites

Why worry about dodgy advertising?

- Legal crackdowns on the underground economy might tempt criminals to shift to more reliable income sources
- Online advertising is a logical target
  - Ad platforms lack the incentive to detect fraud, since detection directly reduces profit
  - Advertisers struggle to monitor for abuse due to lack of transparency
  - Criminals already profit from online advertising: botnets carry out click-fraud, spyware games affiliate-marketing programs

Related work

- Empirical investigations of the underground economy
  - Underground fora (Franklin et al. CCS 2006, Caballero et al. USENIX Security 2011)
  - Email spam (Kanich et al. CCS 2008, Levchenko et al. S&P 2011)
  - Phishing (Moore and Clayton eCrime 2007)
  - Online social networks (Grier et al. CCS 2010)
- Empirical investigations of web-based scams
  - Social engineering (Christin et al. CCS 2010)
  - Drive-by downloads (Provos et al. USENIX Security 2007)
  - Web spam to promote fake antivirus (Rajab et al. LEET 2010, Cova et al. RAID 2010, Stone-Gross et al. WEIS 2011)
  - Web spam to promote ads (Wang et al. WWW 2007, Moore and Edelman FC 2010)
- Empirical investigations of trending abuse
  - Uncovering trending abuse tactics (John et al. USENIX Security 2011, Lu et al. CCS 2011)
  - Cloaking measurement (Wang et al. CCS 2011)

Outline

Data collection methodology

1. Construct a set of trending and control queries
   - *Trending set*: collect 20 Google Hot Trends hourly, and consider a term hot if it has appeared in last 72 hours
   - *Control set*: 495 persistently popular terms (most popular terms in 2010 for 27 categories according to Google)
2. Issue queries across multiple search engines
   - Gather top results from Google, Yahoo, Twitter every 4 hours
   - Over 60 million search results and tweets collected
3. Classify the search results as malicious or benign
   - *Malware*: Check each URL against Google's Safe Browsing API
   - *MFA*: Supervised machine-learning algorithm classifies websites appearing in results of more than 20 different trending terms

Total incidence of malware and MFA

	Terms			Results		
	Total	Infected	%	Total	Infected	%
Malware						
Web Search						
Trending set	6 946	1 232	18	9.8M	7 889	.08
Control set	495	123	25	16.8M	7 332	.04
Twitter						
Trending set	1 950	46	2.4	466K	137	.03
Control set	495	53	11	1M	139	.01
MFA sites						
Web Search						
Trending set	19 792	15 181	76.7	32.3M	954K	3.0
Twitter						
Trending set	1 950	1 833	94	466K	32 152	6.9

More control terms include at least 1 infected result

	Terms			Results		
	Total	Infected	%	Total	Infected	%
Malware						
Web Search						
Trending set	6 946	1 232	18	9.8M	7 889	.08
Control set	495	123	25	16.8M	7 332	.04
Twitter						
Trending set	1 950	46	2.4	466K	137	.03
Control set	495	53	11	1M	139	.01
MFA sites						
Web Search						
Trending set	19 792	15 181	76.7	32.3M	954K	3.0
Twitter						
Trending set	1 950	1 833	94	466K	32 152	6.9

More trending results are infected

	Terms			Results		
	Total	Infected	%	Total	Infected	%
Malware						
Web Search						
Trending set	6 946	1 232	18	9.8M	7 889	.08
Control set	495	123	25	16.8M	7 332	.04
Twitter						
Trending set	1 950	46	2.4	466K	137	.03
Control set	495	53	11	1M	139	.01
MFA sites						
Web Search						
Trending set	19 792	15 181	76.7	32.3M	954K	3.0
Twitter						
Trending set	1 950	1 833	94	466K	32 152	6.9

MFA sites much more pervasive than malware

	Terms			Results		
	Total	Infected	%	Total	Infected	%
Malware						
Web Search						
Trending set	6 946	1 232	18	9.8M	7 889	.08
Control set	495	123	25	16.8M	7 332	.04
Twitter						
Trending set	1 950	46	2.4	466K	137	.03
Control set	495	53	11	1M	139	.01
MFA sites						
Web Search						
Trending set	19 792	15 181	76.7	32.3M	954K	3.0
Twitter						
Trending set	1 950	1 833	94	466K	32 152	6.9

Malware incidence at any given point in time

- By viewing the exposure to malware through search at a single point in time, we find that the risk can be substantial
- Trending terms are more likely to be infected
- Trending terms are more likely to be undetected

	Terms		Results
	#	%	#
Trending terms – web search			
detected	12.8	4.4	14.8
top 10	2.9	1.0	3.2
undetected	6.2	2.1	7.6
top 10	1.2	0.4	1.5
Control terms – web search			
detected	9.5	1.9	14.1
top 10	3.1	0.6	3.9
undetected	1.0	0.2	1.0
top 10	0.1	0.0	0.1

Does abuse incidence vary by term category?

- News, Sports and Local account for over half of all terms
- We found no statistically significant difference in malware incidence across categories
- Some categories were more disposed to MFA (Food&Drink, Reference, Science, Shopping)
- Some categories were less disposed to MFA (e.g., Entertainment, Local, Health)

Category name	%	CPC	Malware % terms	MFA % terms	MFA % top 10	coef.
Arts & Humanities	2.7	$0.44	20.1	40.6	6.8	
Automotive	1.3	$0.67	16.0	29.2	5.2	−0.0062
Beauty & Personal Care	0.8	$0.76	19.6	32.5	6.9	
Business	0.4	$0.87	7.4	32.9	6.9	
Computers & Electronics	2.4	$0.61	14.5	31.7	5.9	
Entertainment	30.6	$0.34	18.6	41.0	6.4	−0.0043
Finance & Insurance	1.4	$1.26	20.2	30.4	5.6	
Food & Drink	2.9	$0.43	17.1	49.5	7.9	+0.0105
Games	2.3	$0.32	13.4	30.0	5.6	−0.0073
Health	2.5	$0.85	14.1	27.6	5.9	−0.0046
Home & Garden	0.5	$0.76	7.1	29.7	7.2	
Industries	1.6	$0.50	26.1	38.6	6.6	−0.0072
Internet	0.7	$0.49	7.7	43.7	6.0	
Lifestyles	4.5	$0.33	25.4	45.8	6.5	
Local	11.0	$0.51	21.8	39.2	6.9	−0.0027
News & Current Events	3.6	$0.39	19.7	45.0	7.0	
Photo & Video	0.2	$0.59	0.0	21.9	6.4	
Real Estate	0.2	$1.02	6.2	34.2	6.5	
Recreation	1.0	$0.43	13.7	43.5	6.5	
Reference	1.4	$0.43	14.5	55.4	8.7	+0.0203
Science	1.4	$0.40	16.0	44.9	9.1	+0.0095
Shopping	3.2	$0.56	11.6	43.7	8.8	+0.0106
Social Networks	0.5	$0.19	27.8	59.1	6.4	
Society	5.1	$0.62	15.2	33.7	5.6	−0.0085
Sports	15.4	$0.38	20.7	44.9	6.9	−0.0044
Telecommunications	0.8	$0.91	10.9	36.4	4.6	
Travel	1.7	$0.88	10.1	29.3	6.4	
Average (category)	**3.7**	**$0.59**	**18.4**	**38.3**	**6.6**	

How a term's popularity and ad price affect malware

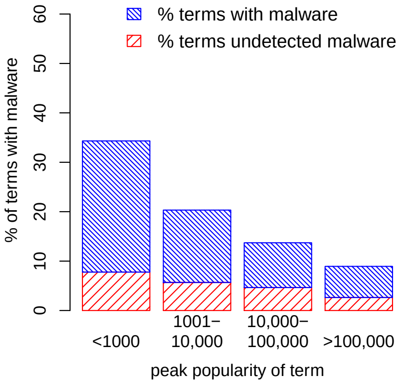

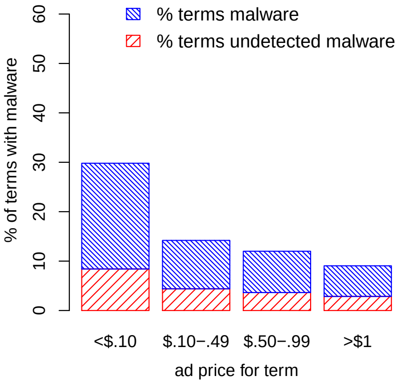

How a term's popularity and ad price affect MFA incidence

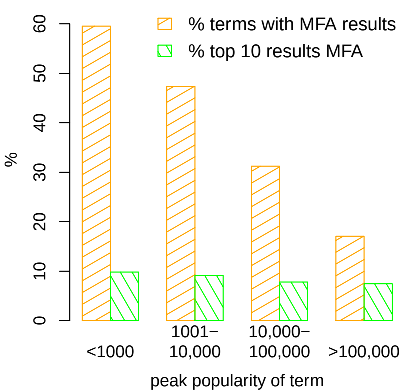

Outline

Estimated visits to MFA and malware sites

- V: # Visits to a website w from searching for s for time period t

$$V(w,s,t) = C\left(\mathrm{Rank}(w,s)\right) \cdot \mathsf{Pop}(s) \cdot \frac{4}{30 \times 24} \times t$$

- Click probability
- Website w position for term s
- Monthly peak popularity of term s

Estimated visits to MFA and malware sites

On 24 Sep 2010 5:00, a search for "dream act 2010 status" (72 600 searches per month), the following URL appears as the third result in Google:

`http://www.eworldpost.com/dream-act-2010-status-17168.html`

$$\boxed{V(w,s,t)} = \boxed{C}\left(\boxed{\text{Rank}(w,s)}\right) \cdot \boxed{\text{Pop}(s)} \cdot \frac{4}{30 \times 24} \times t$$

$$\boxed{V(\texttt{eworldpost.com}, \text{“dream act 2010 status”}, 1)} = \boxed{C}\left(\boxed{3}\right) \cdot \boxed{72\,600} \cdot \frac{4}{30 \times 24} \times 1$$

$$\boxed{V(\texttt{eworldpost.com}, \text{“dream act 2010 status”}, 1)} = 44 \text{ visits}$$

Estimated visits to MFA and malware sites

	# Visitors		
	Total	Period	Monthly Rate
MFA	39 274 200	275 days	4 284 458
Malware (trending set)			
detected	454 198	88 days	154 840
undetected	143 662	88 days	48 975
Malware (control set)			
detected	12 825 332	88 days	4 372 272
undetected	83615	88 days	28 505

Estimated daily victims to malware in search results

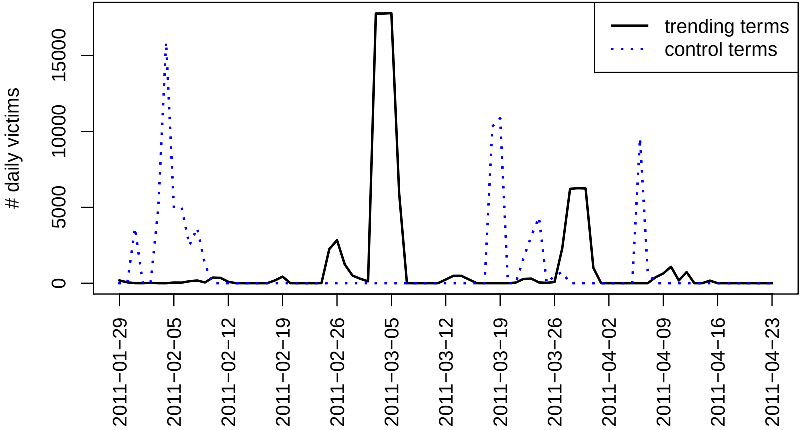

MFA Revenue Analysis

- Add up all MFA results
- Pay-per-click ads
- Visits per result

$$\sum_{w \in \mathrm{MFA}(s)} \sum_{s} \Big(V(w,s,t) \times \big(p_{\mathrm{PPC}} \cdot p_{\mathrm{clk}} \cdot r_{\mathrm{PPC}} + p_{\mathrm{banner}} \cdot r_{\mathrm{banner}} + p_{\mathrm{aff}} \cdot p'_{\mathrm{clk}} \cdot r_{\mathrm{aff}} \big) \Big)$$

- Banner ads
- Pay-per-conversion ads

MFA revenue for one month

- Pay-per-click ads
- Visits

$$\underbrace{4\,300\,000}_{\text{Visits}} \times \Big(\underbrace{0.94 \times 0.02 \times \$0.97}_{\text{Pay-per-click ads}} + \underbrace{0.53 \times \$0.00529}_{\text{Banner ads}} + \underbrace{0.33 \times 0.02 \times \$0.265}_{\text{Pay-per-conversion ads}} \Big)\Big)$$

- Banner ads
- Pay-per-conversion ads
- Monthly MFA revenue: $98 000

Malware revenue analysis

$$\sum_{w \in \mathrm{mal}(s)} \sum_s \left(V(w,s,t) \cdot p_{\mathrm{exp}} \cdot p_{\mathrm{pay}} \cdot r_{\mathrm{AV}} \right)$$

- Add up all malware results
- Expected # visits per result
- Expected revenue from fake antivirus

Malware revenue for one month

$$\boxed{49\,000} \cdot \boxed{1 \cdot 0.0216 \cdot \$58}$$

- Visits
- Expected revenue from fake antivirus (Source: Stone-Gross et al. WEIS 2011)
- Monthly malware revenue: $61 000

Outline

The New York Times

Business Day

The Dirty Little Secrets of Search

Photo illustration by The New York Times

By DAVID SEGAL
Published: February 12, 2011

PRETEND for a moment that you are Google's search engine.

Someone types the word "dresses" and hits enter. What will be the very first result?

There are, of course, a lot of possibilities. Macy's comes to mind. Maybe a specialty chain, like J. Crew or the Gap. Perhaps a Wikipedia entry on the history of hemlines.

O.K., how about the word "bedding"? Bed Bath & Beyond

The Official Google Blog

Insights from Googlers into our products, technology, and the Google culture.

Finding more high-quality sites in search

2/24/2011 06:50:00 PM

Our goal is simple: to give people the most relevant answers to their queries as quickly as possible. This requires constant tuning of our algorithms, as new content—both good and bad—comes online all the time.

Many of the changes we make are so subtle that very few people notice them. But in the last day or so we launched a pretty big algorithmic improvement to our ranking—a change that noticeably impacts 11.8% of our queries—and we wanted to let people know what's going on. This update is designed to reduce rankings for low-quality sites—sites which are low-value add for users, copy content from other websites or sites that are just not very useful. At the same time, it will provide better rankings for high-quality sites—sites with original content and information such as research, in-depth reports, thoughtful analysis and so on.

We can't make a major improvement without affecting rankings for many sites. It has to be that some sites will go up and some will go down. Google depends on the high-quality content created by wonderful websites around the world, and we do have a responsibility to encourage a healthy web ecosystem. Therefore, it is important for high-quality sites to be rewarded, and that's exactly what this change does.

It's worth noting that this update does not rely on the feedback we've received from the Personal Blocklist Chrome extension, which we launched last week. However, we did compare the Blocklist data we gathered with the sites identified by our algorithm, and we were very pleased that the preferences our users expressed by using the extension are well represented. If you take the top several dozen or so most-blocked domains from the Chrome extension, then this algorithmic change addresses 84% of them, which is strong independent confirmation of the user benefits.

Search This Blog

Search

powered by Google™

Follow us on Twitter

Site Feed

Google

578K readers

BY FEEDBURNER

Make Google your homepage

Blog Archive

Blog Archive

Labels

accessibility (30)
acquisition (19)
ads (101)
Africa (17)
Android (27)
apps (370)
April 1 (4)
Asia (31)
books + book search (44)
crisis response (24)
developers (111)

Measuring the effect of Google's intervention

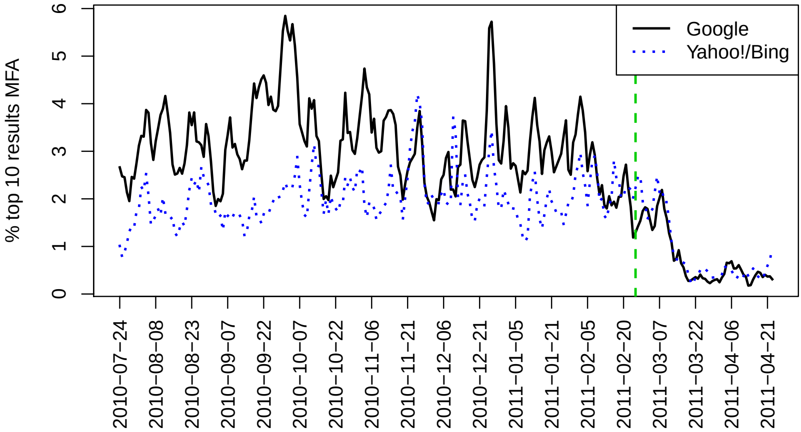

Implications of Google's intervention for search engines

	Monthly MFA visits		
	Pre-intervention	Post-intervention	% change
Google search	3 364 402	1 788 480	-47%
Yahoo!/Bing search	1 302 314	1 448 058	+11%
Total	**4 666 716**	**3 236 538**	**-31%**

Cautionary tale on crackdowns: typosquatting

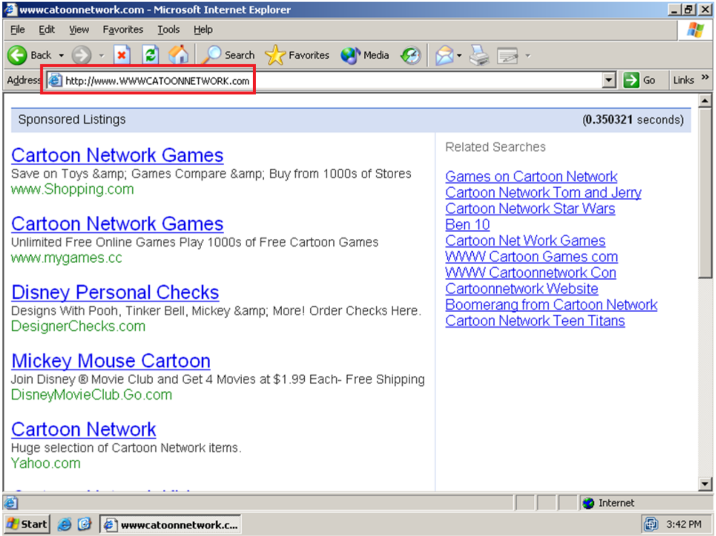

1 278 cartoonnetwork.com typos, including. . .

fartoonnetwork.com	cagtoonnetwork.com	cartlonnetwork.com	cartoonnestwork.com	cartoonnewotk.com
cartoonnetsork.com	cartoinnetwork.com	cartolnnetwork.com	cartoonnftwork.com	cartoonneywork.com
cartoonntewrk.com	cartoonnetlork.com	cartoonnetowok.com	crtonnetwork.com	cartoonnegwork.com
cargoonnetwork.com	carttoonnnetwork.com	cartoonnwetwork.com	cartoonetrwork.com	cartoonnetwodrk.com
cartoonnetwkor.com	catoonnnetwork.com	cartoooonnetwork.com	caoonnetwork.com	cartonbetwork.com
cartoonetgork.com	cartoonnetqork.com	cartoonneetwort.com	cartoonneetwork.com	catoonnetwrok.com
cartoomnetwoork.com	caryoonetwork.com	cartooonetwork.com	cantoonnetwork.com	cargoonnetworm.com
caretoonetwork.com	cartoonetwoork.com	cartoonnetwoer.com	carttoonnetwerk.com	chartoonnetwork.com
cartoonnetwokr.com	cartoonnetwokl.com	cartoonnetwoke.com	cartoownnetwork.com	cartoobetwork.com
cartoonnetworkcom.com	nartoonnetwork.com	cartoonnstwork.com	cartoounnetwork.com	cartoonework.com
carfoonnetwork.com	cartoonnotwork.com	cartoonnnetwok.com	cartoonnnetwor.com	cartonnetwortk.com
cartoopnetwork.com	cartoonnetwogk.com	cartoonetwaork.com	cartoonntewrok.com	cartoonetwoek.com
caqrtoonetwork.com	cartoonneework.com	cartppnnetwork.com	cartoonnetmwork.com	cartoooonework.com
cartoonntwoork.com	catoonneetwork.com	crattoonnetwork.com	cartoonnetweark.com	carttooonetwork.com
cartoonetwoirk.com	cartoonznetwork.com	cartoobnetwork.com	catoonnework.com	cartiinnetwork.com
cartoonnnnetwrk.com	cartoommetwork.com	cartoonnetwart.com	wwwcartonetwork.com	cartoonnttwork.com
cartoonhetwork.com	fcartoonnetwork.com	catoonnetwerk.com	artoonnetwor.com	cartoonnetwock.com
cartoonnetook.com	cartoonnetkwork.com	cartonnetwokr.com	carltonnetwork.com	cartoonetowrk.com
catoonnettwork.com	cartoo0nnetwork.com	cacrtoonnetwork.com	cartoonnetwoorkl.com	cartoonedtwork.com
cartoonnetwcrk.com	cartoonetwrk.com	cartoonnewark.com	cartoonnetwoirk.com	cartoknnetwork.com
cartooonnetwrk.com	cartoonnetbwork.com	caetooonetwork.com	cartoonknetwork.com	catoomnetwork.com
cartoonnexwork.com	carooonetwork.com	dartoonnetwork.com	certoonnetwork.com	cartoonetword.com
cartoonetworg.com	cartoonetworl.com	cartoonetworj.com	cartoonetwork.com	cartoonetwort.com
crattonnetwork.com	cartoonnewtokr.com	carntoonnetwork.com	caretoonnetwork.com	cartooonnetwoork.com
cartoonnerwort.com	cartoonnerwork.com	cartoonnerworl.com	cartoonnetfork.com	cartoonnetttwork.com
cartoonnetwar.com	cartoonnetwak.com	cartoonnekwork.com	cartooknetwork.com	cartoonegwork.com
cattoonnetwok.com	cartoonnetwwork.com	cartoonnetgor.com	cartoonnetwowk.com	wwwcatoonetwork.com
cartoolnnetwork.com	cartoonetworkcom.com	casrtoonetwork.com	cartoonnetswork.com	cartoonnedwort.com
cartoonnedword.com	cartoonnedwork.com	wwwcarttoonnetwork.com	cartoonerwork.com	cattoonnetwark.com
carttoonnetwook.com	cartoonnetwowrk.com	cartoonetwqork.com	crartoonnetwork.com	czrtoonnetwork.com

Conclusions

- Trending terms are successfully exploited by miscreants more often than consistently popular terms
- Pointing users to malware and ad-filled sites yields similar revenue levels
- Either way, miscreants are only modestly successful
- Google's quality crackdown worked, but beware unintended consequences for the attraction of malware to the bad guys
- For more, see `http://cs.wellesley.edu/~tmoore/`

Calibration tests balance comprehensiveness and efficiency

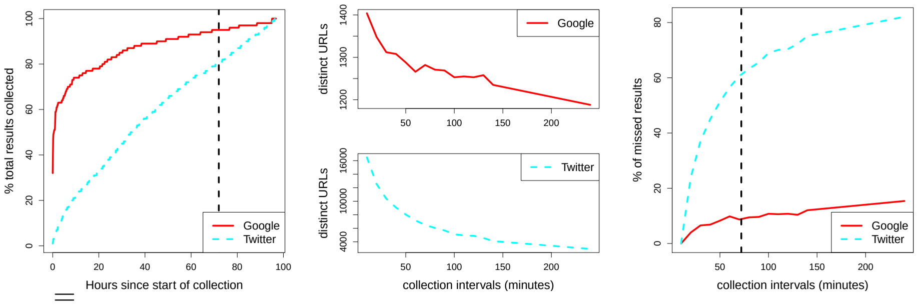

Regression model for malware prevalence

$$\text{logit}(p_{\text{HasMalware}}) = \beta + \text{AdPrice}x_1 + \log_2(\text{Popularity})x_2$$

	coef.	odds	Std. Err.	Significance
AdPrice	-0.509	$.601$	0.091	$p < 0.001$
$\log_2$(Popularity)	-0.117	0.889	0.012	$p < 0.001$

Regression model for MFA prevalence

$$\text{FracTop10MFA} = \beta + \text{AdPrice}x_1 + \log_2(\text{Popularity})x_2 + \text{Category}x_3 \ .$$

	coef.	Std. Err.	Significance
AdPrice	-0.0091	0.091	$p < 0.001$
$\log_2$(Popularity)	-0.004	0.012	$p < 0.001$
Coefficients for category variables in earlier Table, R^2: 0.1373			