OWASP
AppSec Europe
London 2nd-6th July 2018

Detecting and Preventing Malicious Domain Registrations in the .eu TLD

Lieven Desmet

About me

- Senior Research Manager at KU Leuven
  - (Web) Application & DNS Security
  - Security Analytics

- Organizing committee of SecAppDev
- Board member of intigriti
- Board member of OWASP BE chapter

- *Joint research with EURid, registry of .eu*

Malicious use of domain names

- Domain names are often abused by cyber criminals
  - Spam, botnet C&C infrastructure, phishing, malware, …

- To avoid blacklisting, malicious actors often deploy a hit-and-run strategy
  - 60% are only active for 1 day after registration [Hao et al]

[Hao et al] “Understanding the Domain Registration Behavior of Spammers” IMC 2013

Research hypothesis:

“Malicious actors register domains in bulk, and do so for longer periods of time.”

Goal of this research

- “Can we identify such bulk behavior based on commonalities between individual registrations?”

- Understand the malicious domain registration ecosystem
- To detect and prevent malicious registrations

Outline of the talk

- Longitudinal campaign analysis
- Insights in malicious domain registrations
- Pro-active detection and prevention

Longitudinal campaign analysis

Domain name registrations in the .eu TLD

- **.eu** – 8th largest ccTLD (European Economic Area)
  - 3.8 million domain names

- Dataset used in this research:
  - 824,121 new registrations over 14 months (Apr 2015 – May 2016)
  - 20,870 registrations end up on blacklists (2.5%)

Available registration data

- Basic registration information
  - domain name, datetime of registration, and registrar
- Contact information of the registrant
  - company name, name, language, email address, phone, fax, as well as postal address
- Name server information
  - Name servers and/or glue records

Dataset enrichments

- Maliciousness of a domain name
  - Spamhaus DBL
  - SURBL multi list
  - Google Safe Browsing
- Geolocation information of name servers
  - MaxMind GeoLite2 Free database

Campaign identification process

- Start from maliciously flagged registrations
- Group registrations based on similarities between registration details
- Start heuristics:
  - Peaks in malicious registrations
  - Strong discrepancies between malicious and benign domains

Example campaign (c_11)

- Multiple fake registrant details
  - Combinations of

    2 email accounts,

    3 phone numbers,

    4 street addresses

- **8 months active**
  (Jun 3, 2015 – Feb 3, 2016)

- **1,275 blacklisted registrations**
  (= 53.96%)

Registration details used by c_11

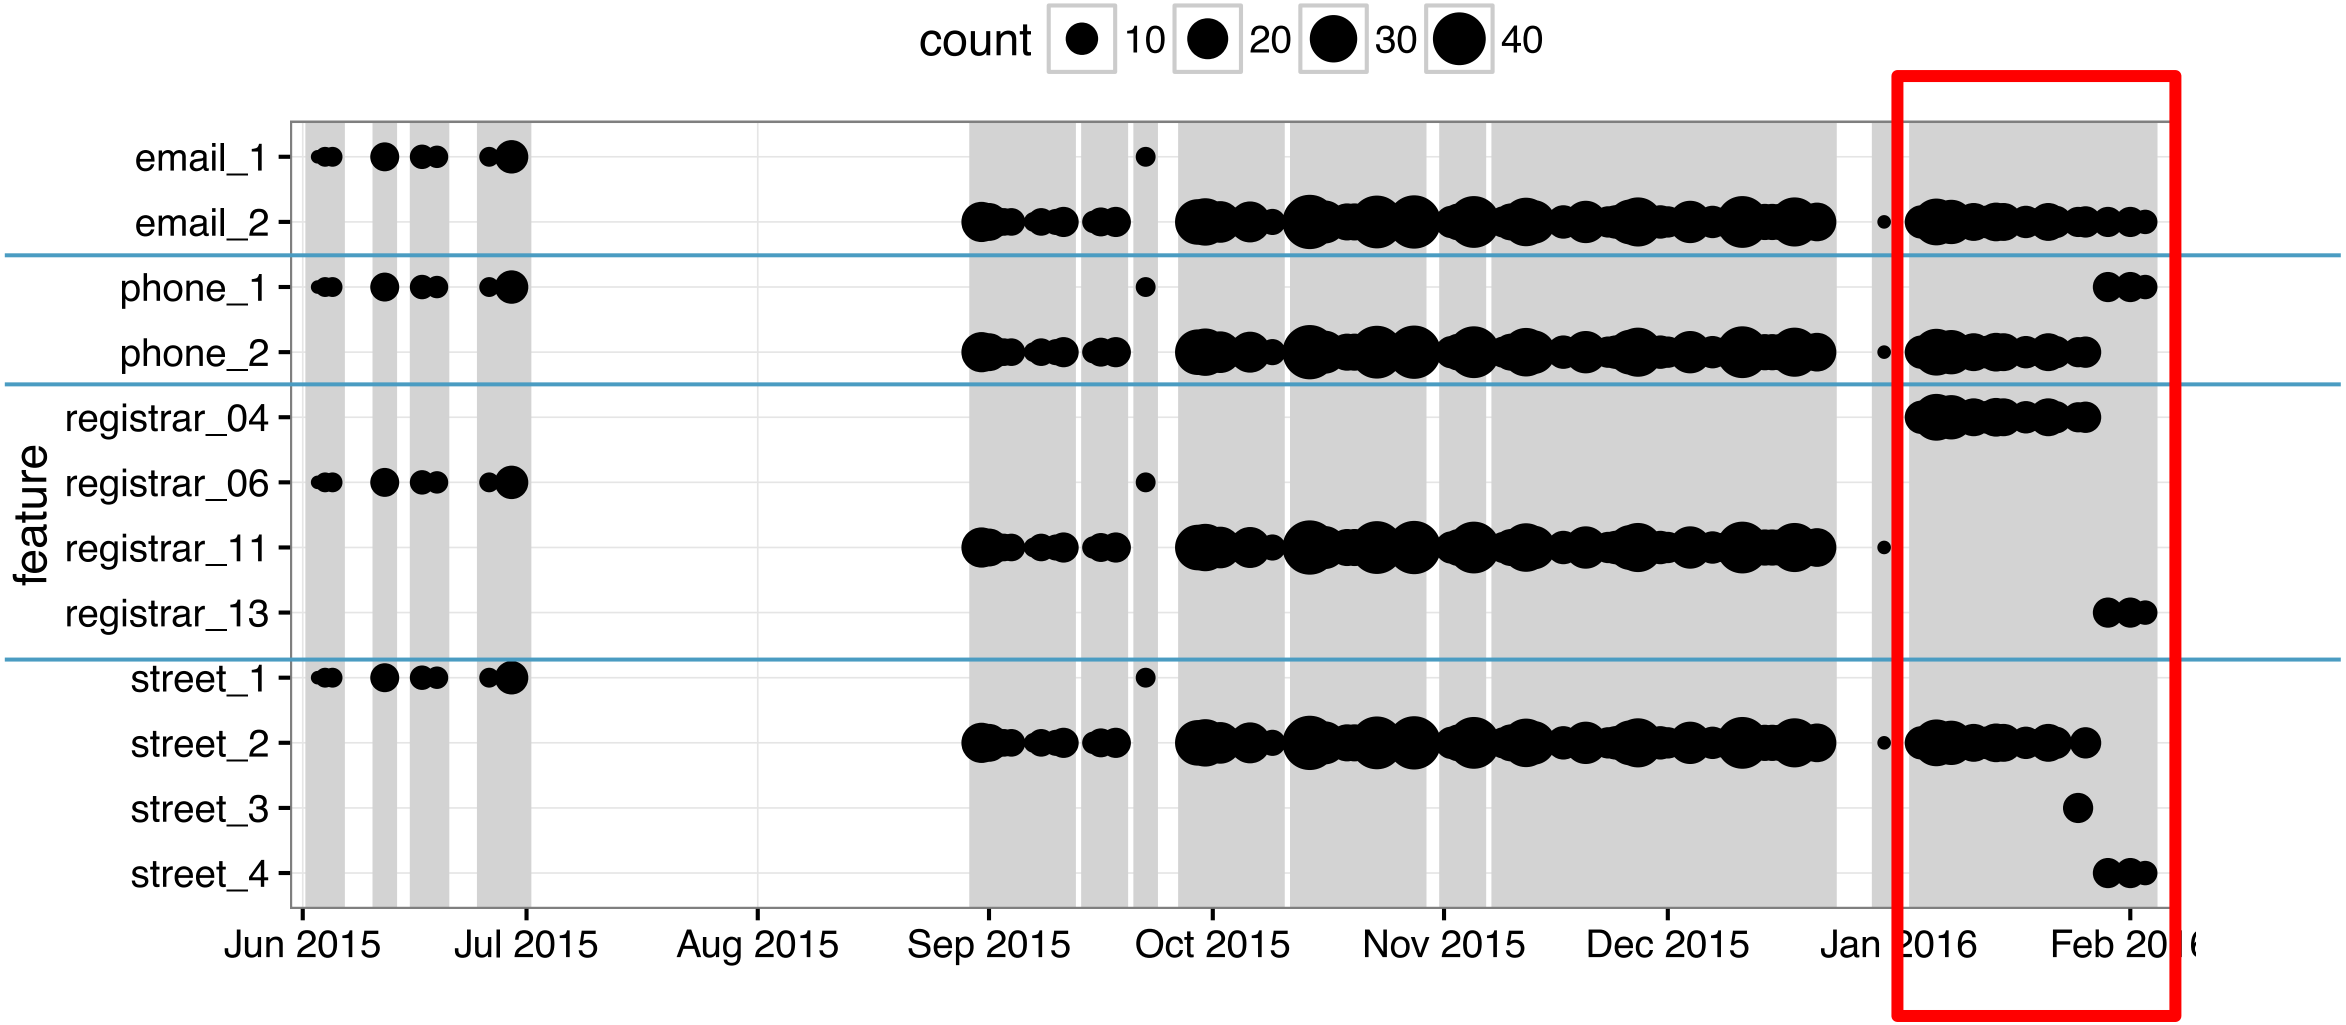

Example of an advanced campaign (c_15)

- Registrant details:
  - 98 fake registrants
  - Generated by Laravel Faker tool
- Domain names:
  - Consist out of 2-3 Dutch words
  - Dutch words are reused across registrants
- Batches of 8, 16, 24 or 32 registrations

- **8+ months active (Sep 16, 2015 – May 31, 2016)**
- **514 blacklisted registrations** ***(= 26.95%)***

Activity of identified campaigns

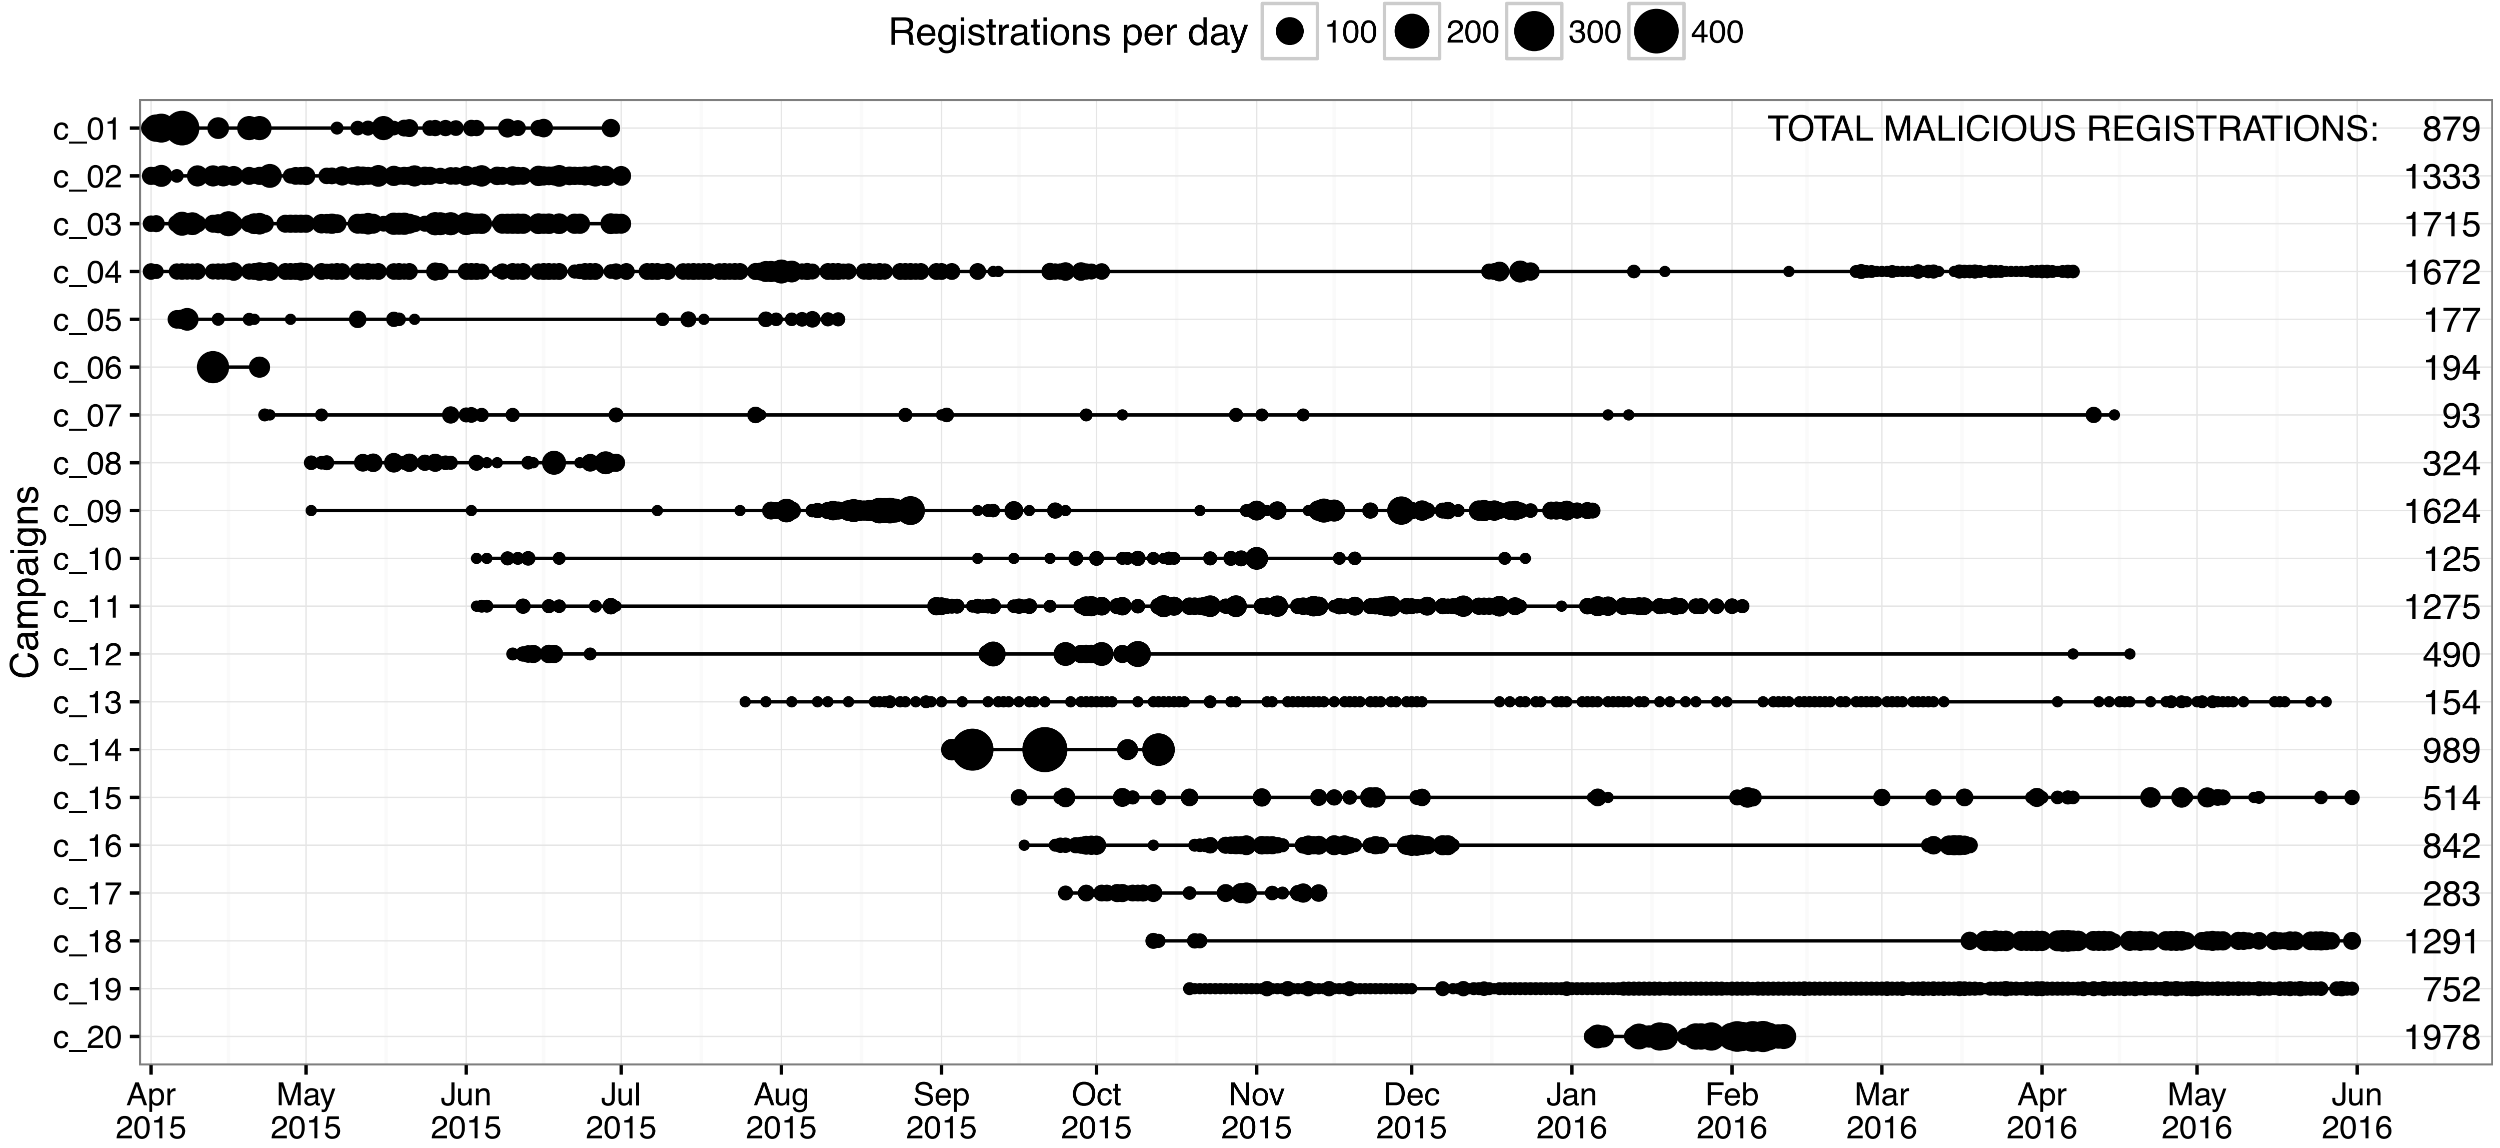

Campaign selection criteria

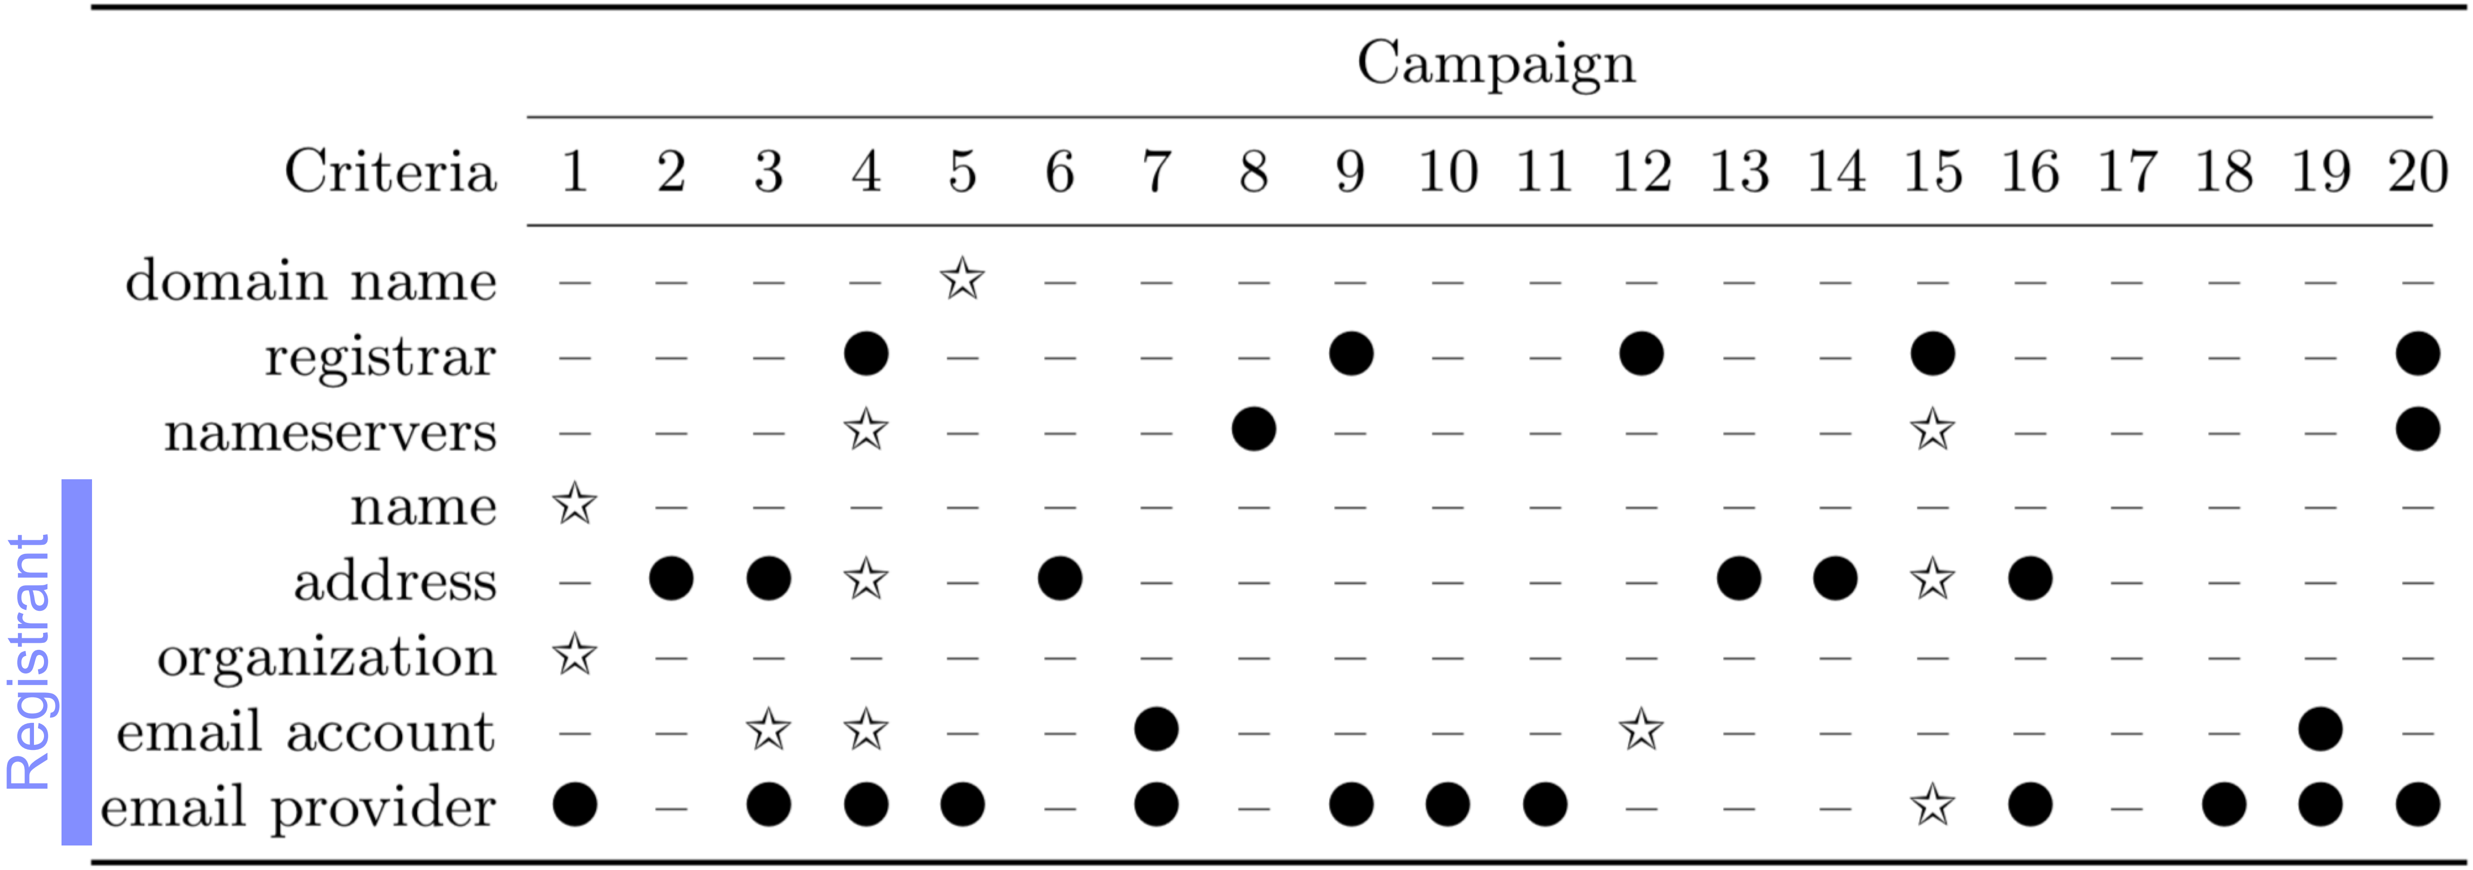

		Campaign																			
	Criteria	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
	domain name	–	–	–	–	☆	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	registrar	–	–	–	●	–	–	–	–	●	–	–	●	–	–	●	–	–	–	–	●
	nameservers	–	–	–	☆	–	–	–	●	–	–	–	–	–	–	☆	–	–	–	–	●
Registrant	name	☆	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	address	–	●	●	☆	–	●	–	–	–	–	–	–	●	●	☆	●	–	–	–	–
	organization	☆	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	email account	–	–	☆	☆	–	–	●	–	–	–	–	☆	–	–	–	–	–	–	●	–
	email provider	●	–	●	●	●	–	●	–	●	●	●	–	–	–	☆	●	–	●	●	●

● represents a string match, and ☆ a regular expression pattern

Insights in malicious domain registrations

Insight 1: Hit-and-run strategies

- Small window of opportunity:
  - Domain rendered useless once blacklisted
  - 73% is blacklisted 5 days after registration

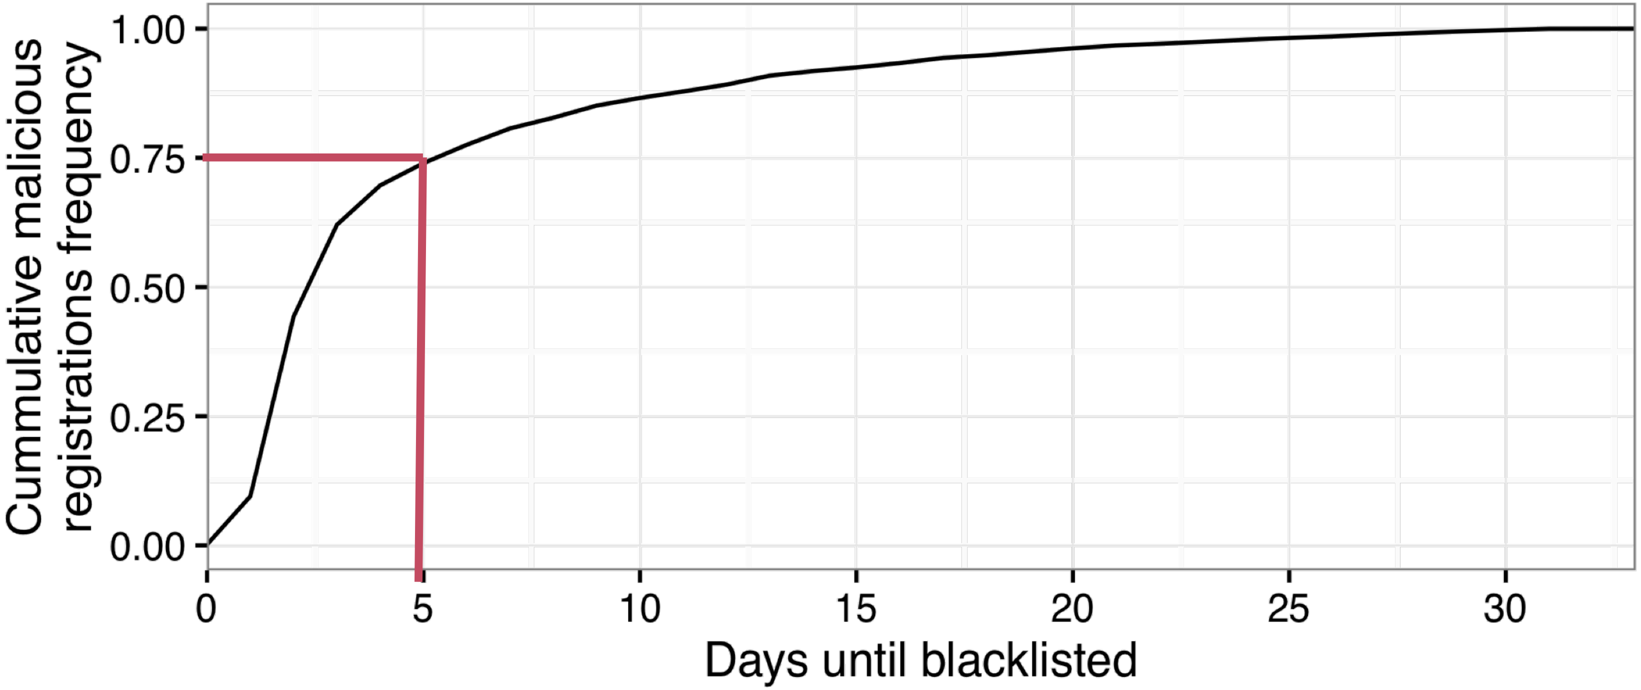

Insight 2: Campaigns are primarily linked to spam

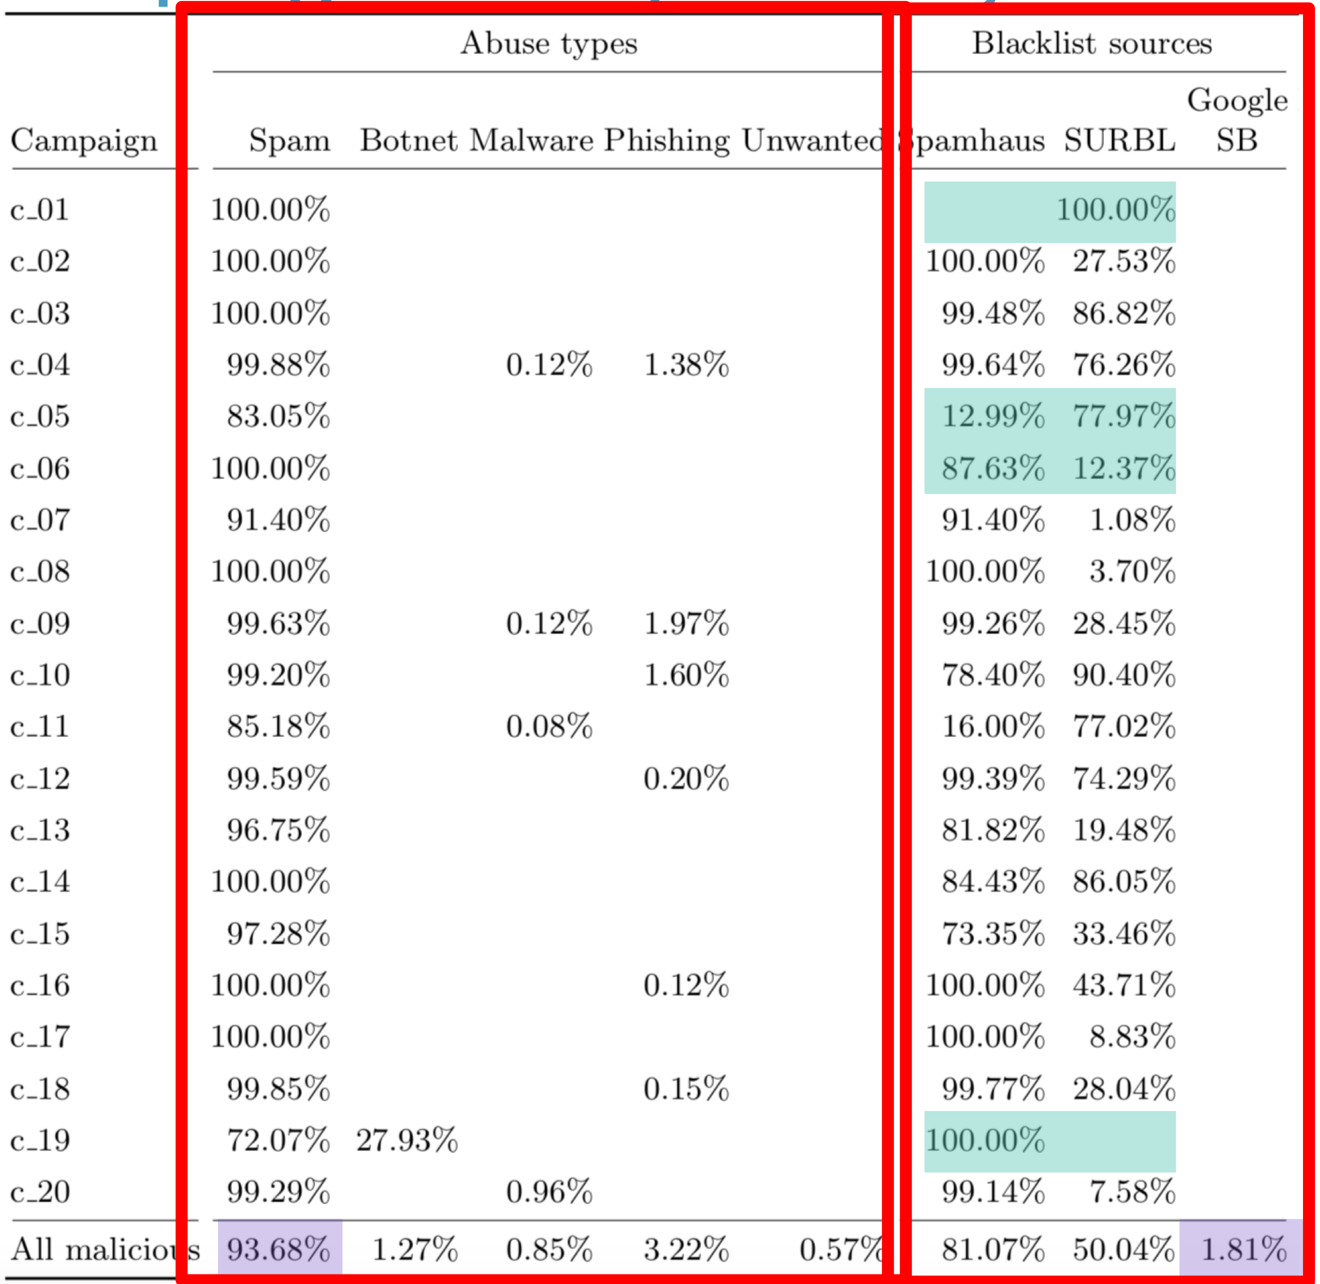

Campaign	Abuse types					Blacklist sources		
	Spam	Botnet	Malware	Phishing	Unwanted	Spamhaus	SURBL	Google SB
c_01	100.00%						100.00%	
c_02	100.00%					100.00%	27.53%	
c_03	100.00%					99.48%	86.82%	
c_04	99.88%		0.12%	1.38%		99.64%	76.26%	
c_05	83.05%					12.99%	77.97%	
c_06	100.00%					87.63%	12.37%	
c_07	91.40%					91.40%	1.08%	
c_08	100.00%					100.00%	3.70%	
c_09	99.63%		0.12%	1.97%		99.26%	28.45%	
c_10	99.20%			1.60%		78.40%	90.40%	
c_11	85.18%		0.08%			16.00%	77.02%	
c_12	99.59%			0.20%		99.39%	74.29%	
c_13	96.75%					81.82%	19.48%	
c_14	100.00%					84.43%	86.05%	
c_15	97.28%					73.35%	33.46%	
c_16	100.00%			0.12%		100.00%	43.71%	
c_17	100.00%					100.00%	8.83%	
c_18	99.85%			0.15%		99.77%	28.04%	
c_19	72.07%	27.93%				100.00%		
c_20	99.29%		0.96%			99.14%	7.58%	
All malicious	93.68%	1.27%	0.85%	3.22%	0.57%	81.07%	50.04%	1.81%

Insight 3: Variety in intensity and duration

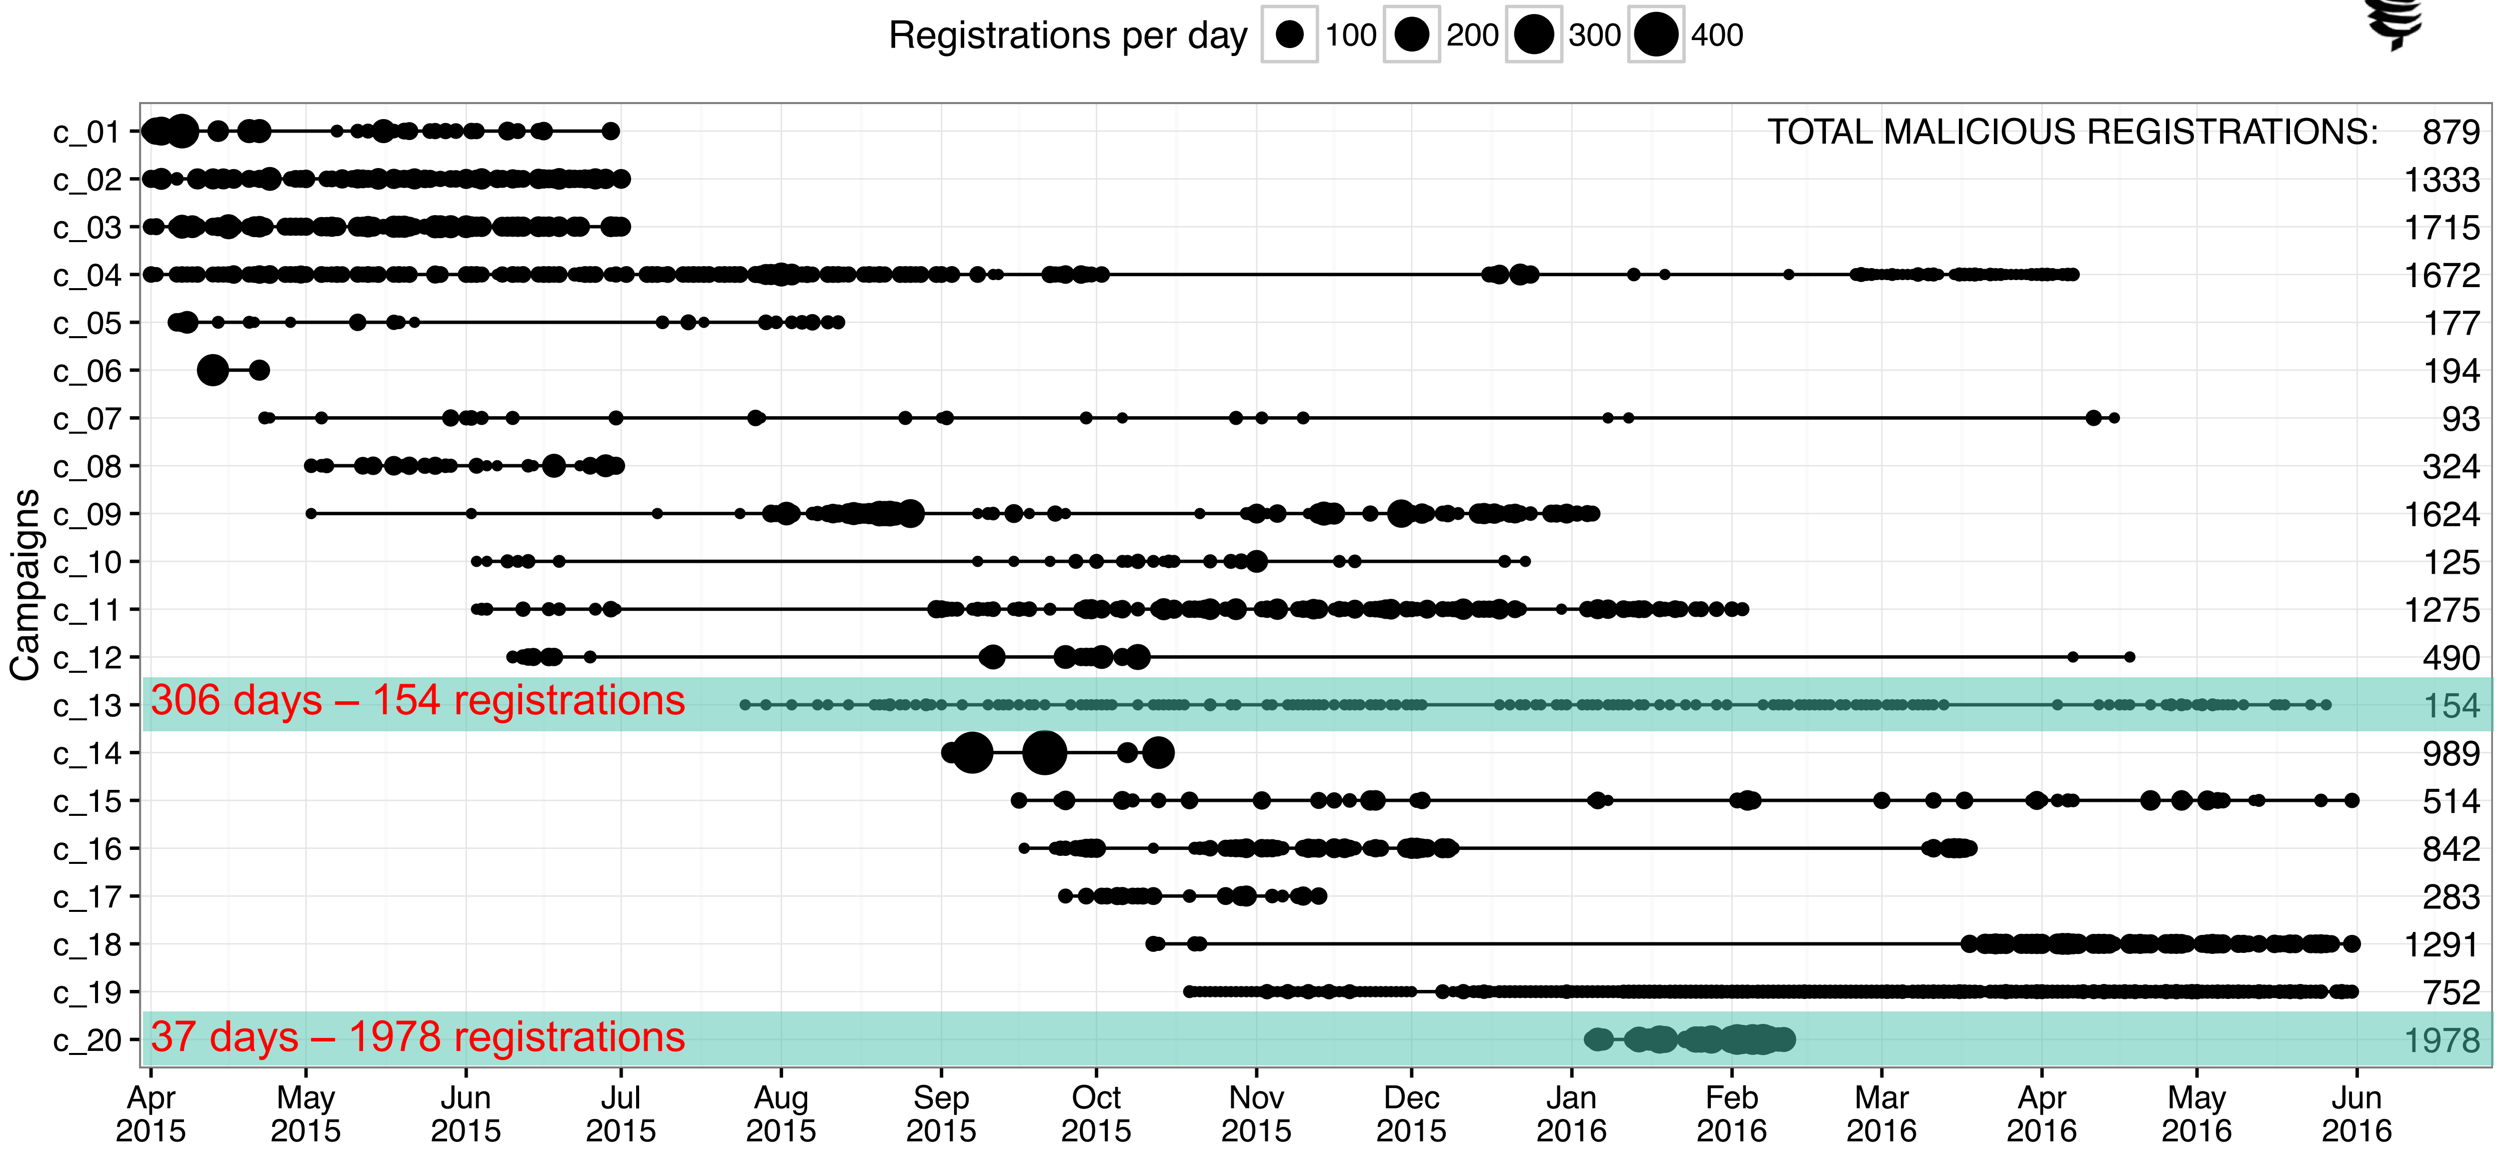

306 days – 154 registrations

37 days – 1978 registrations

Insight 4: Some campaigns are linked to each other

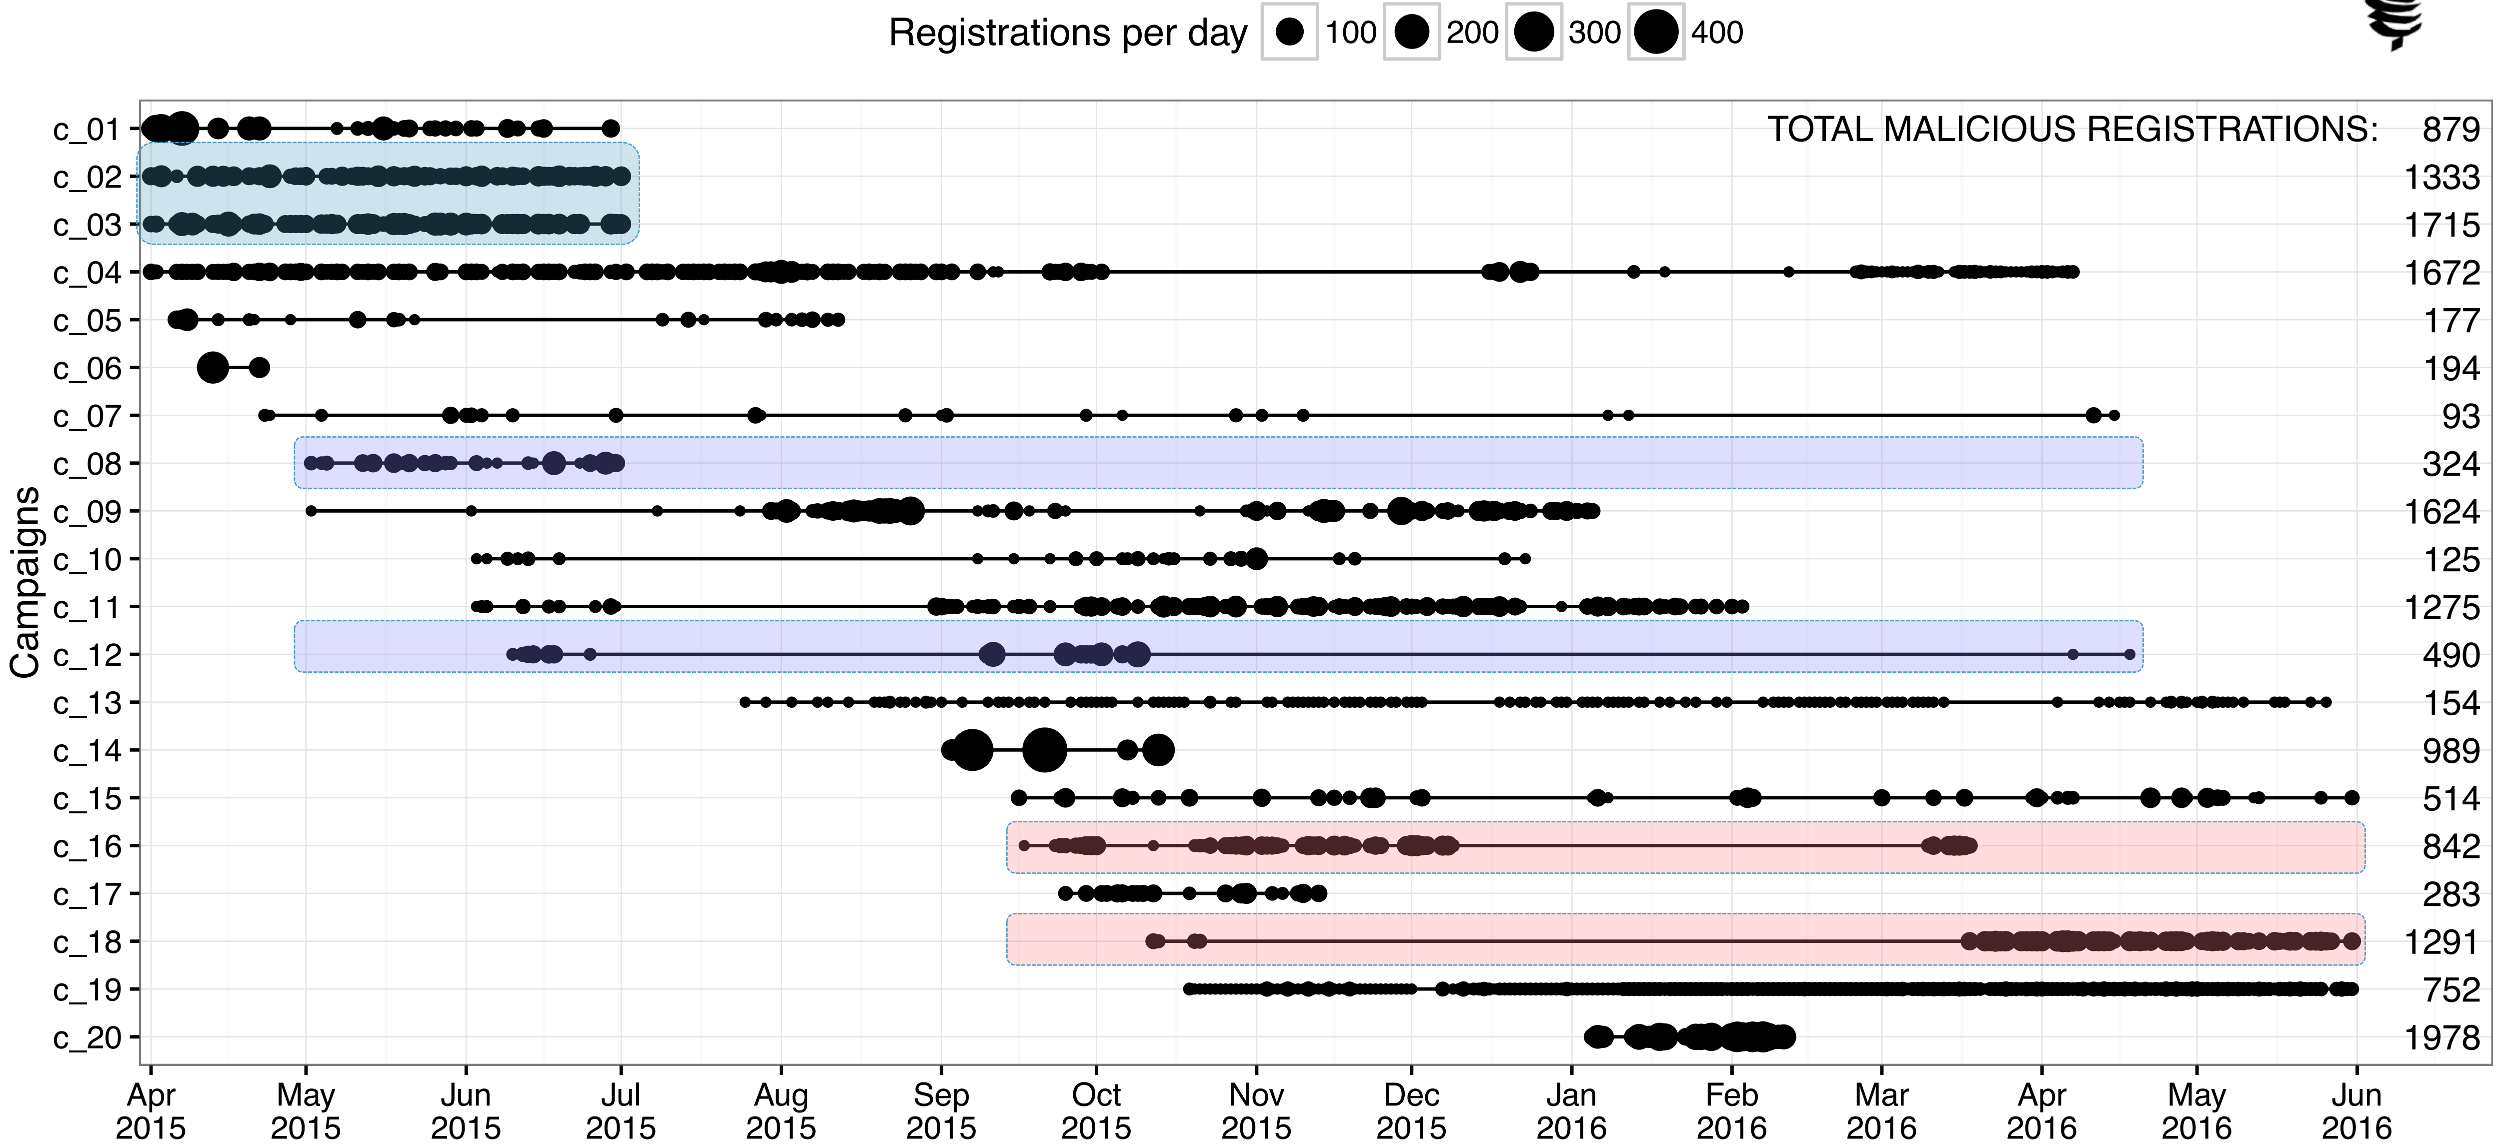

Insight 5: Some campaigns are fully automated

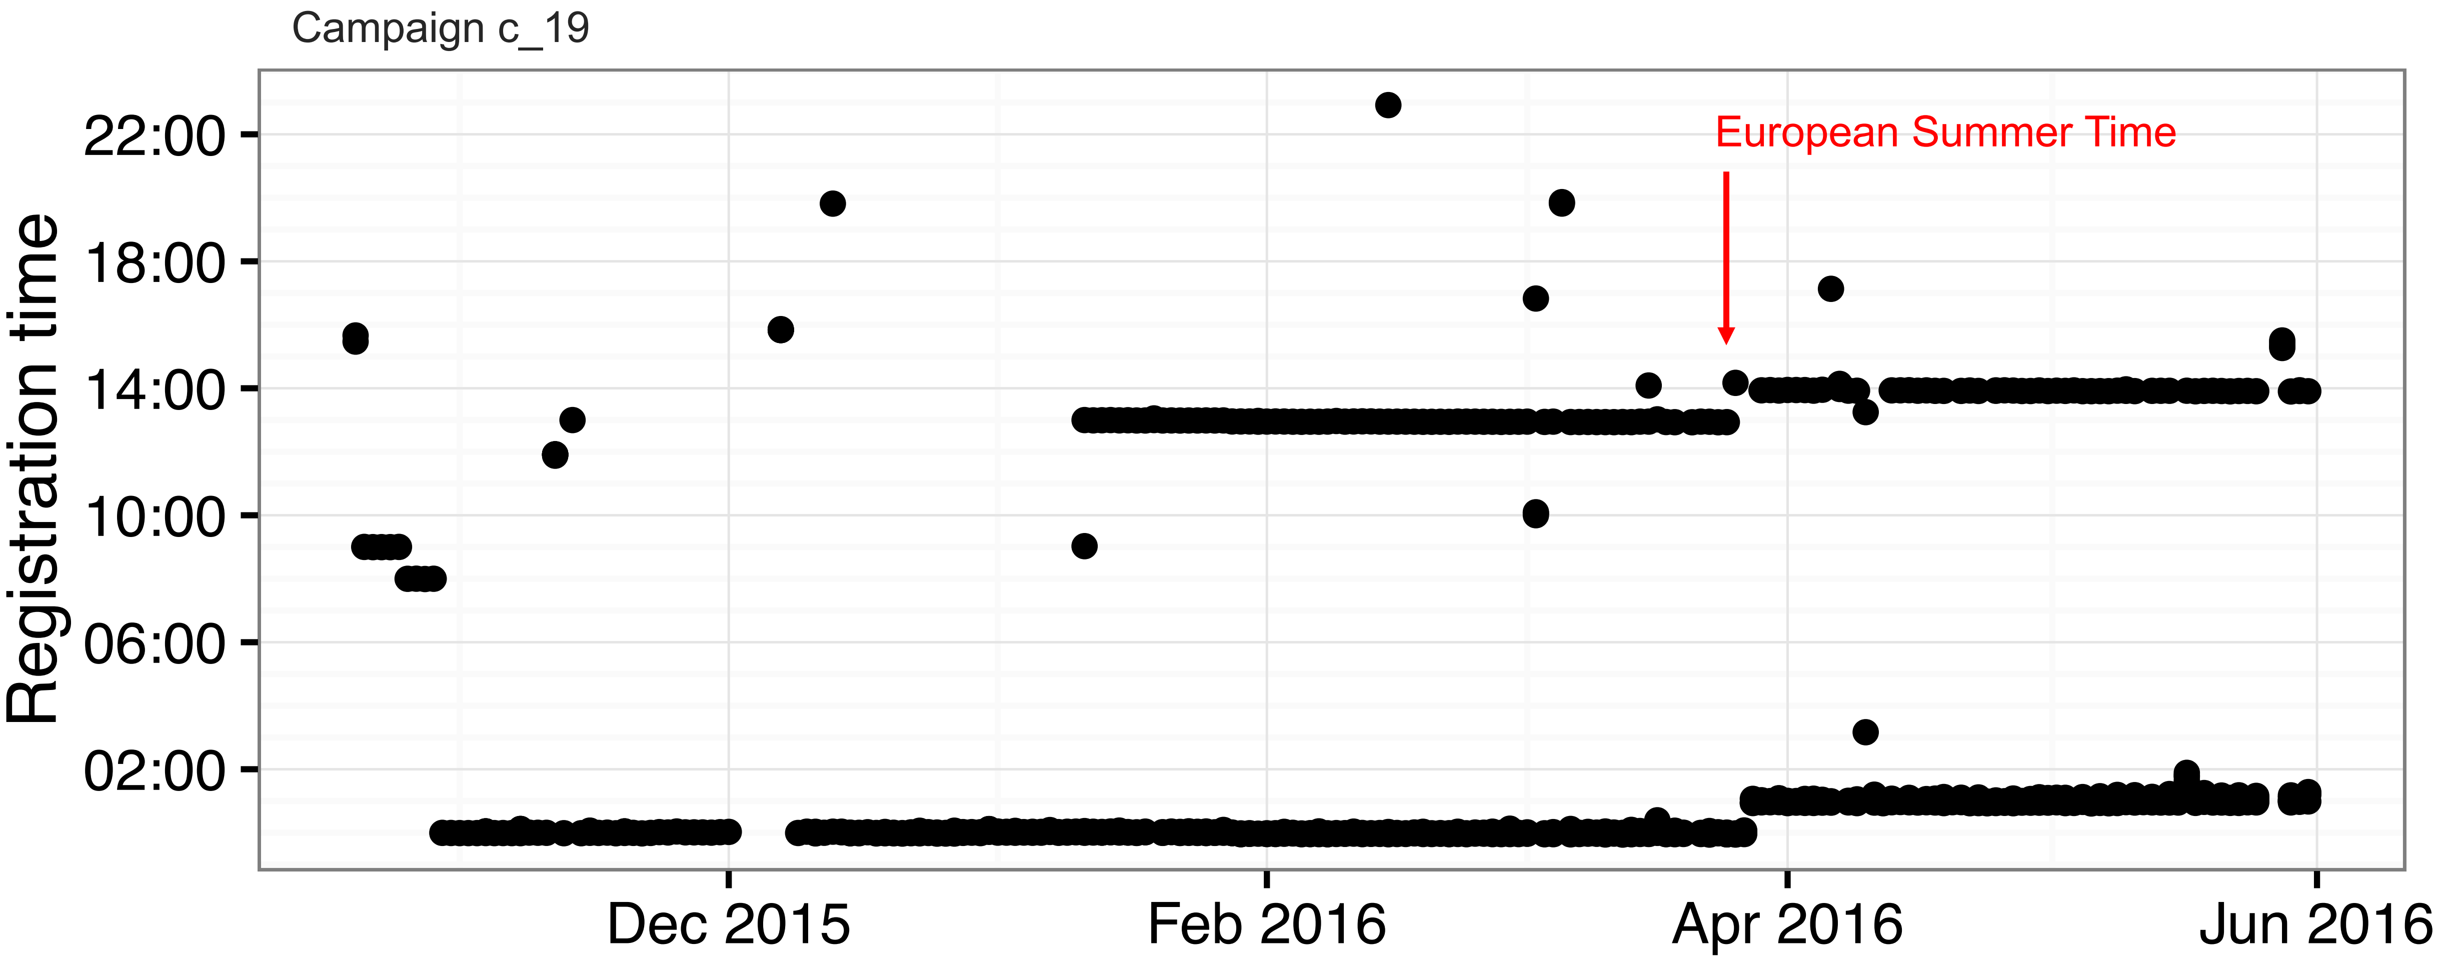

Insight 6: Some campaigns align with regular business activity patterns (1)

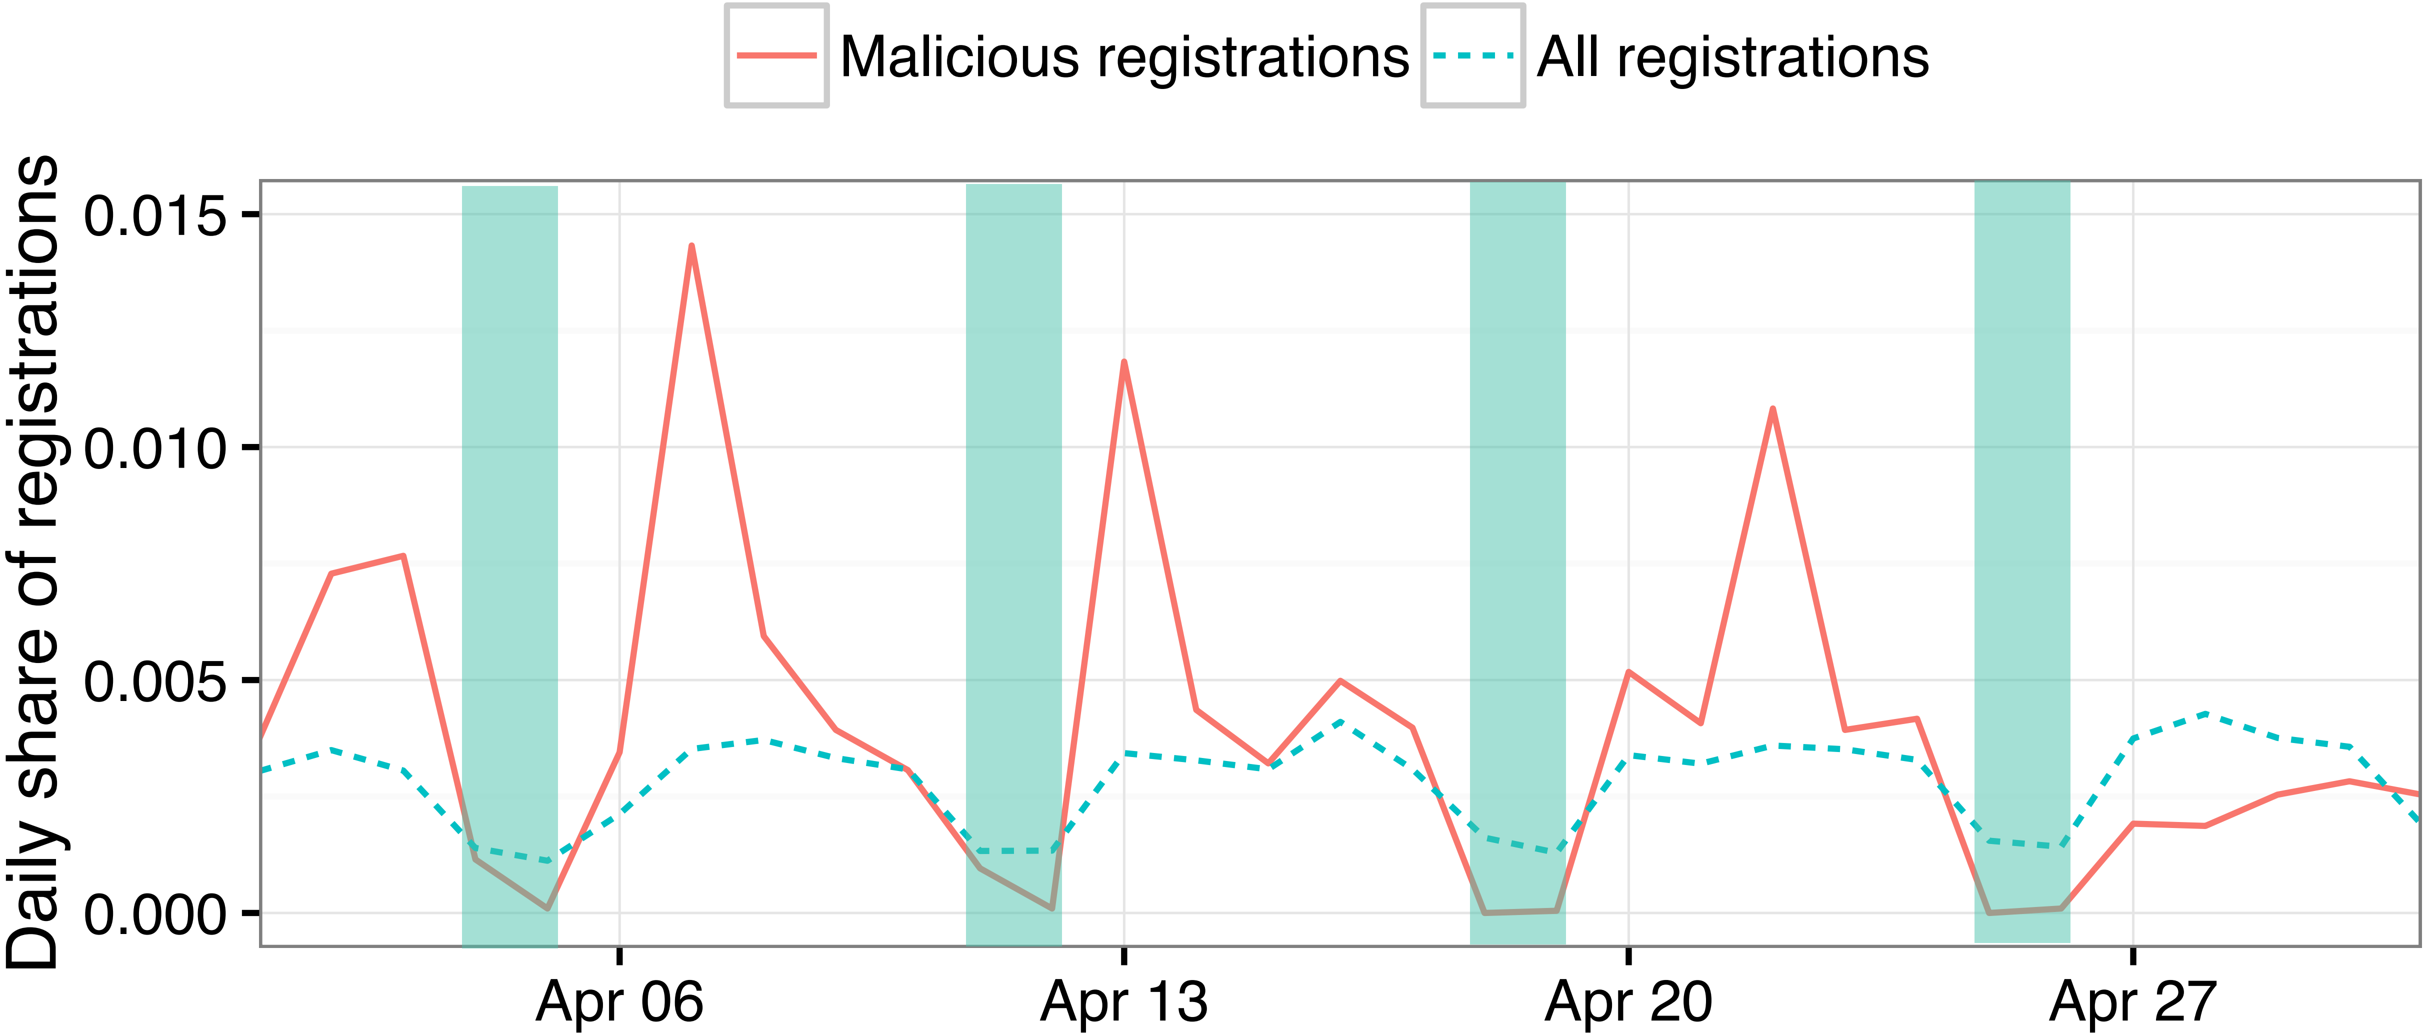

Insight 6: Some campaigns align with regular business activity patterns (2)

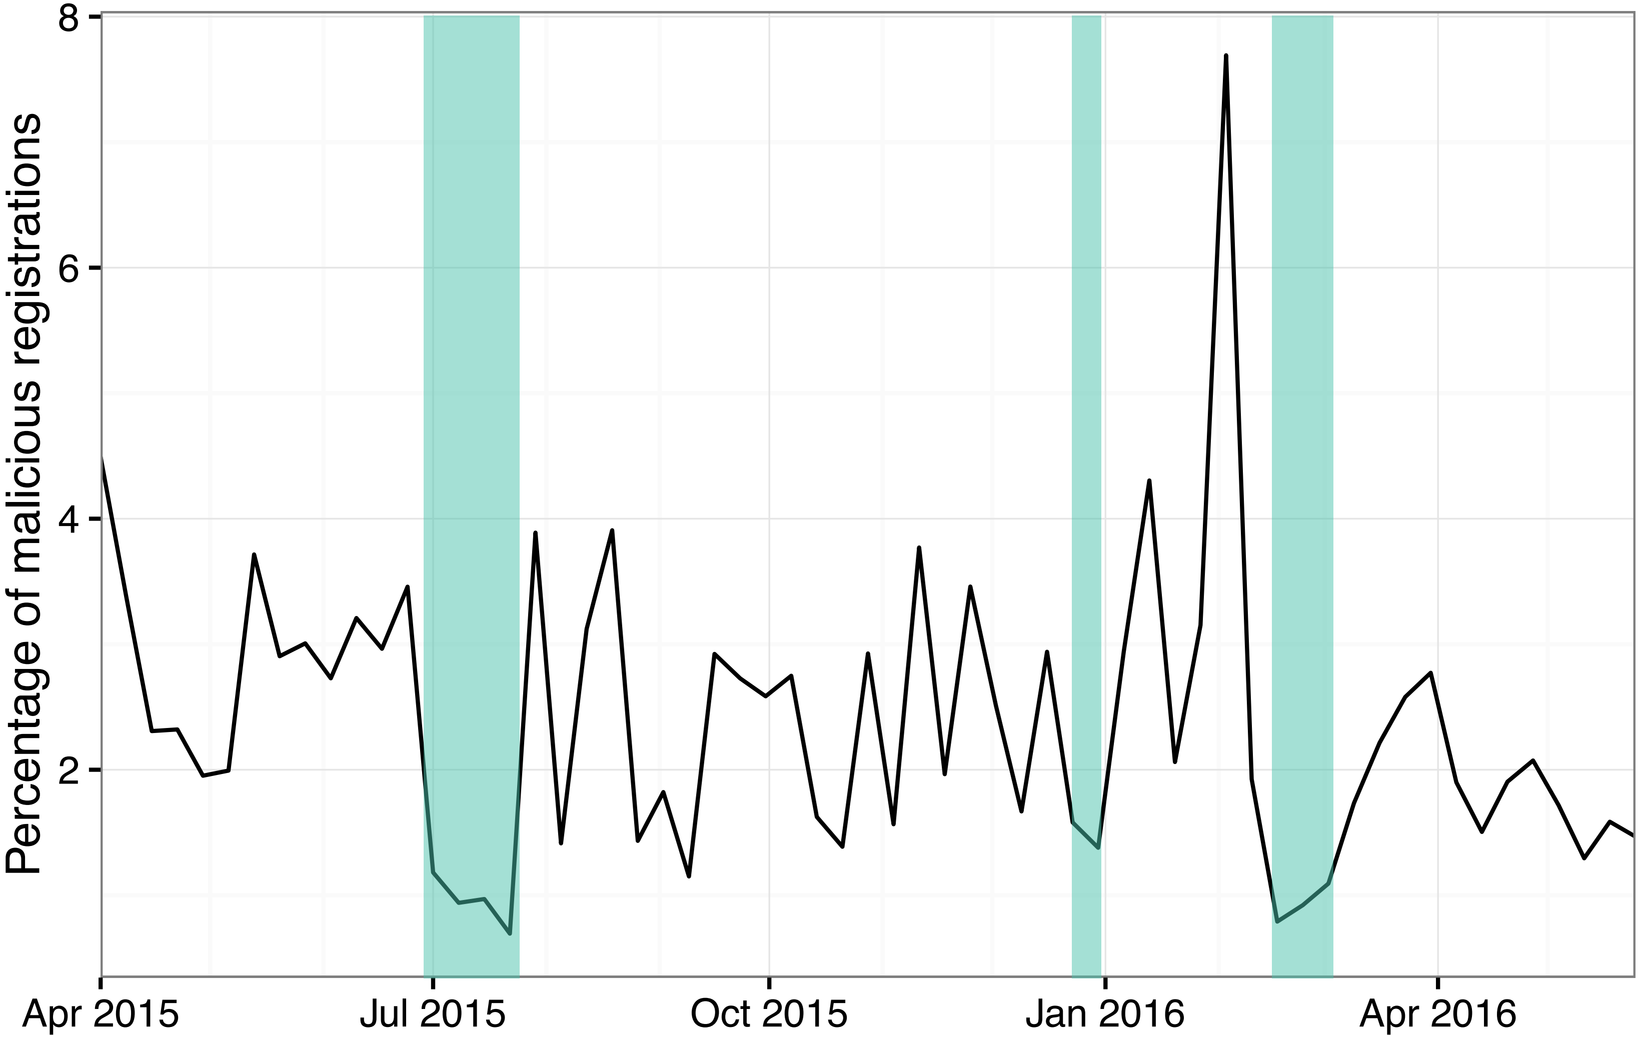

Insight 6: Some campaigns align with regular business activity patterns (3)

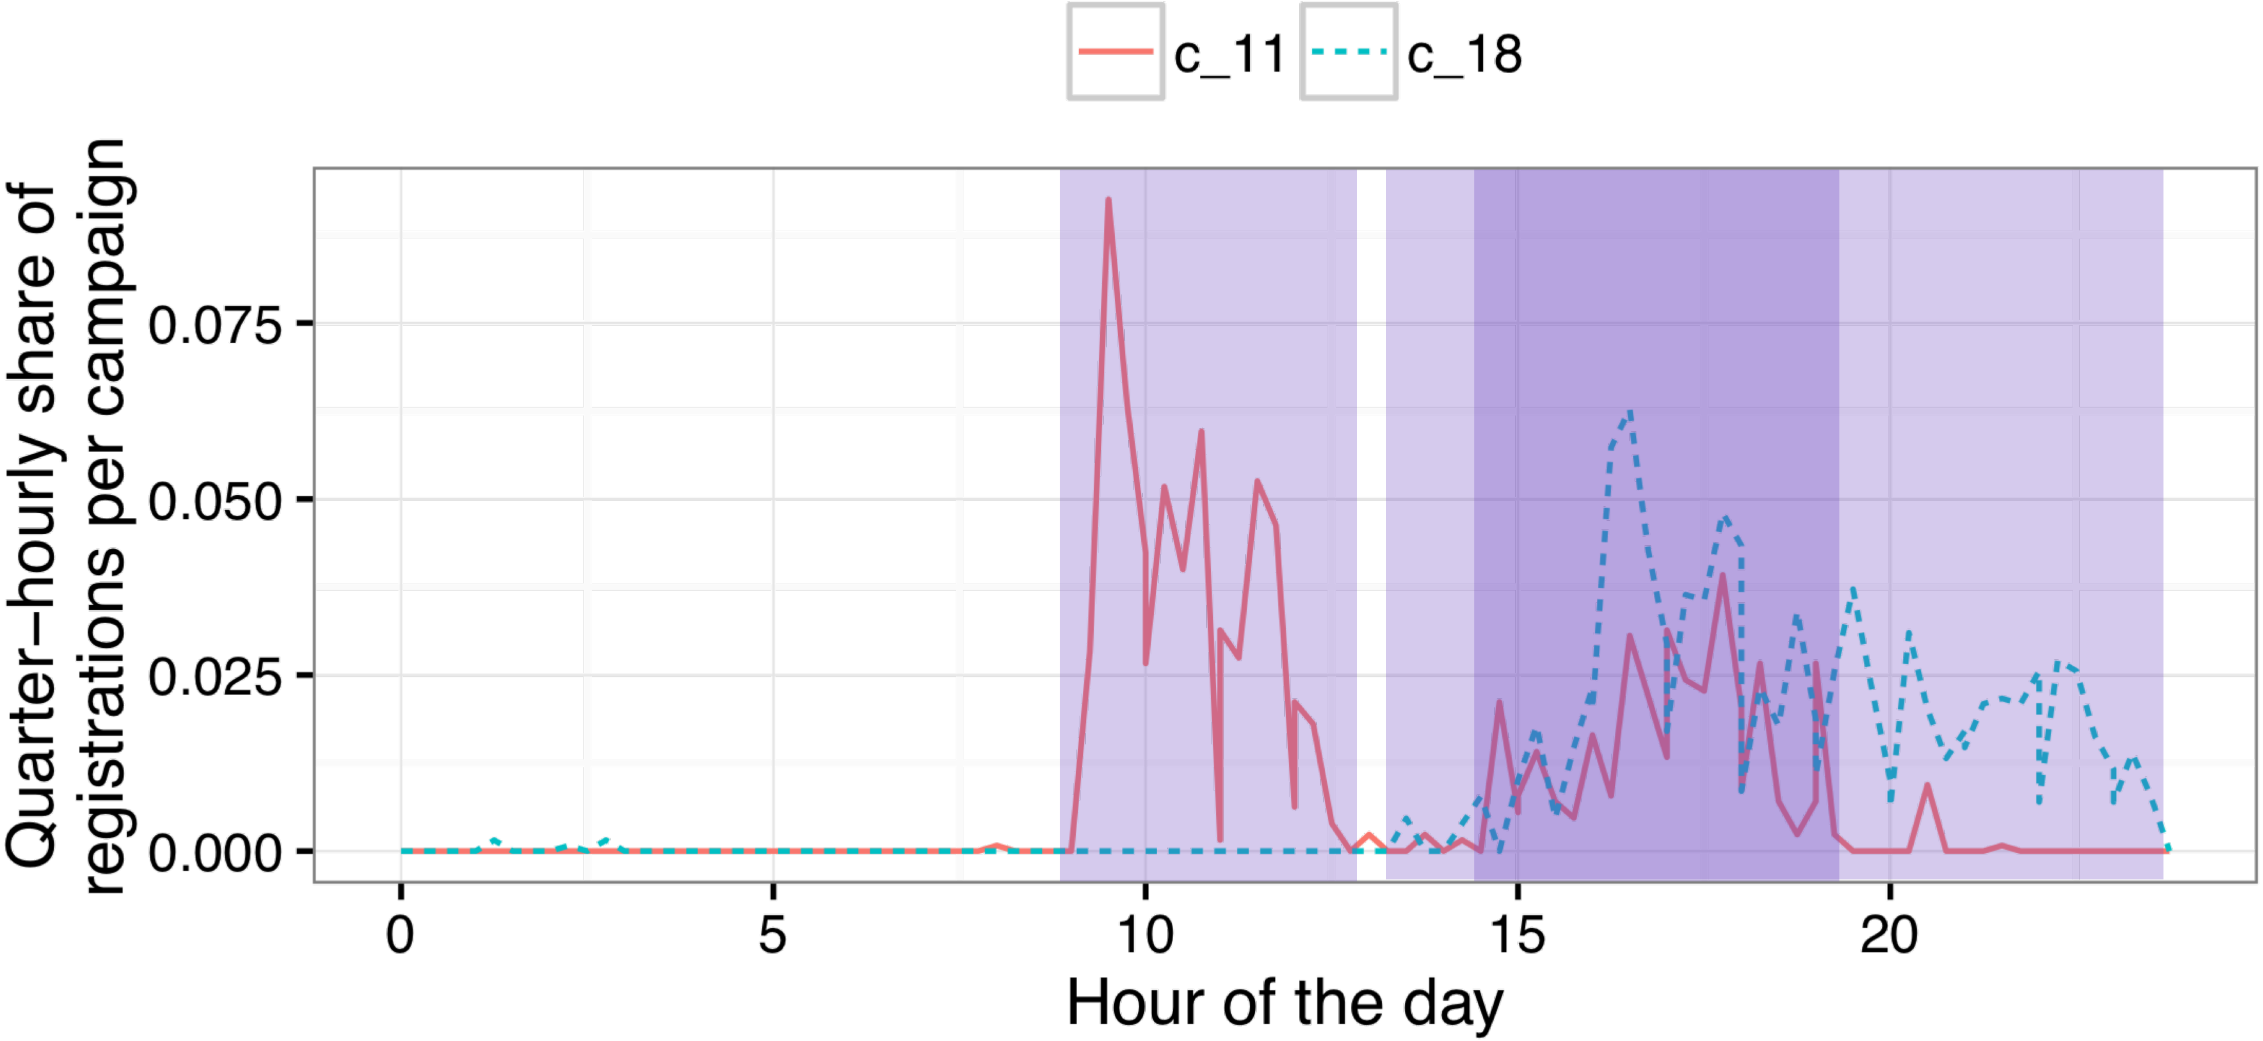

Insight 7: Top facilitators for malicious registrations

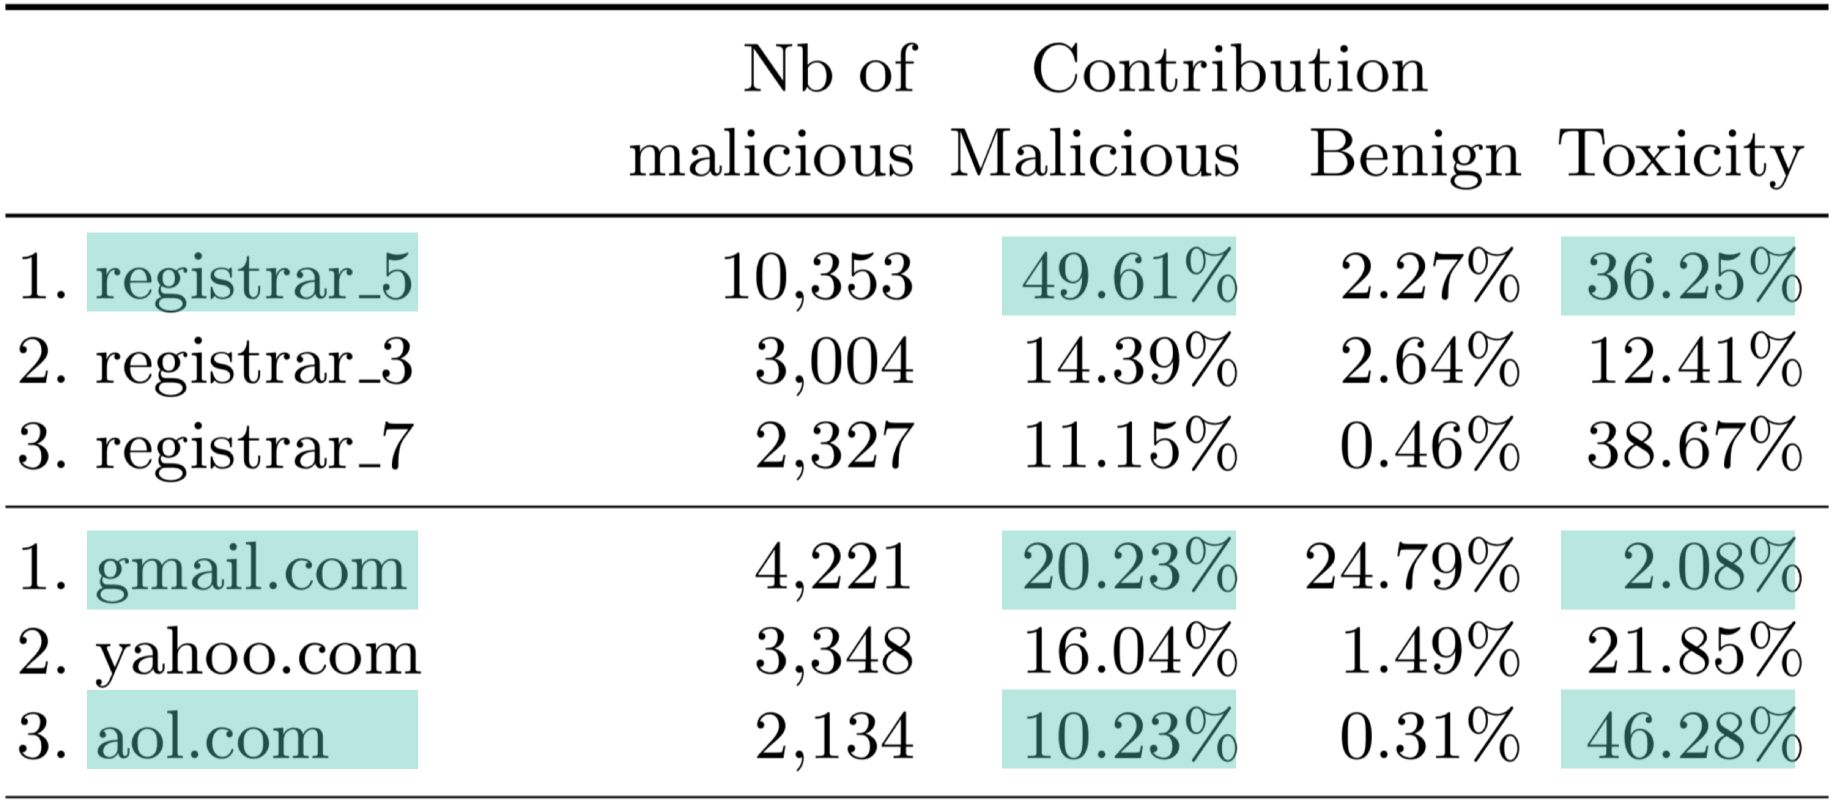

	Nb of malicious	Contribution Malicious	 Benign	 Toxicity
1. registrar_5	10,353	49.61%	2.27%	36.25%
2. registrar_3	3,004	14.39%	2.64%	12.41%
3. registrar_7	2,327	11.15%	0.46%	38.67%
1. gmail.com	4,221	20.23%	24.79%	2.08%
2. yahoo.com	3,348	16.04%	1.49%	21.85%
3. aol.com	2,134	10.23%	0.31%	46.28%

Insight 8: Adaptive campaign strategies

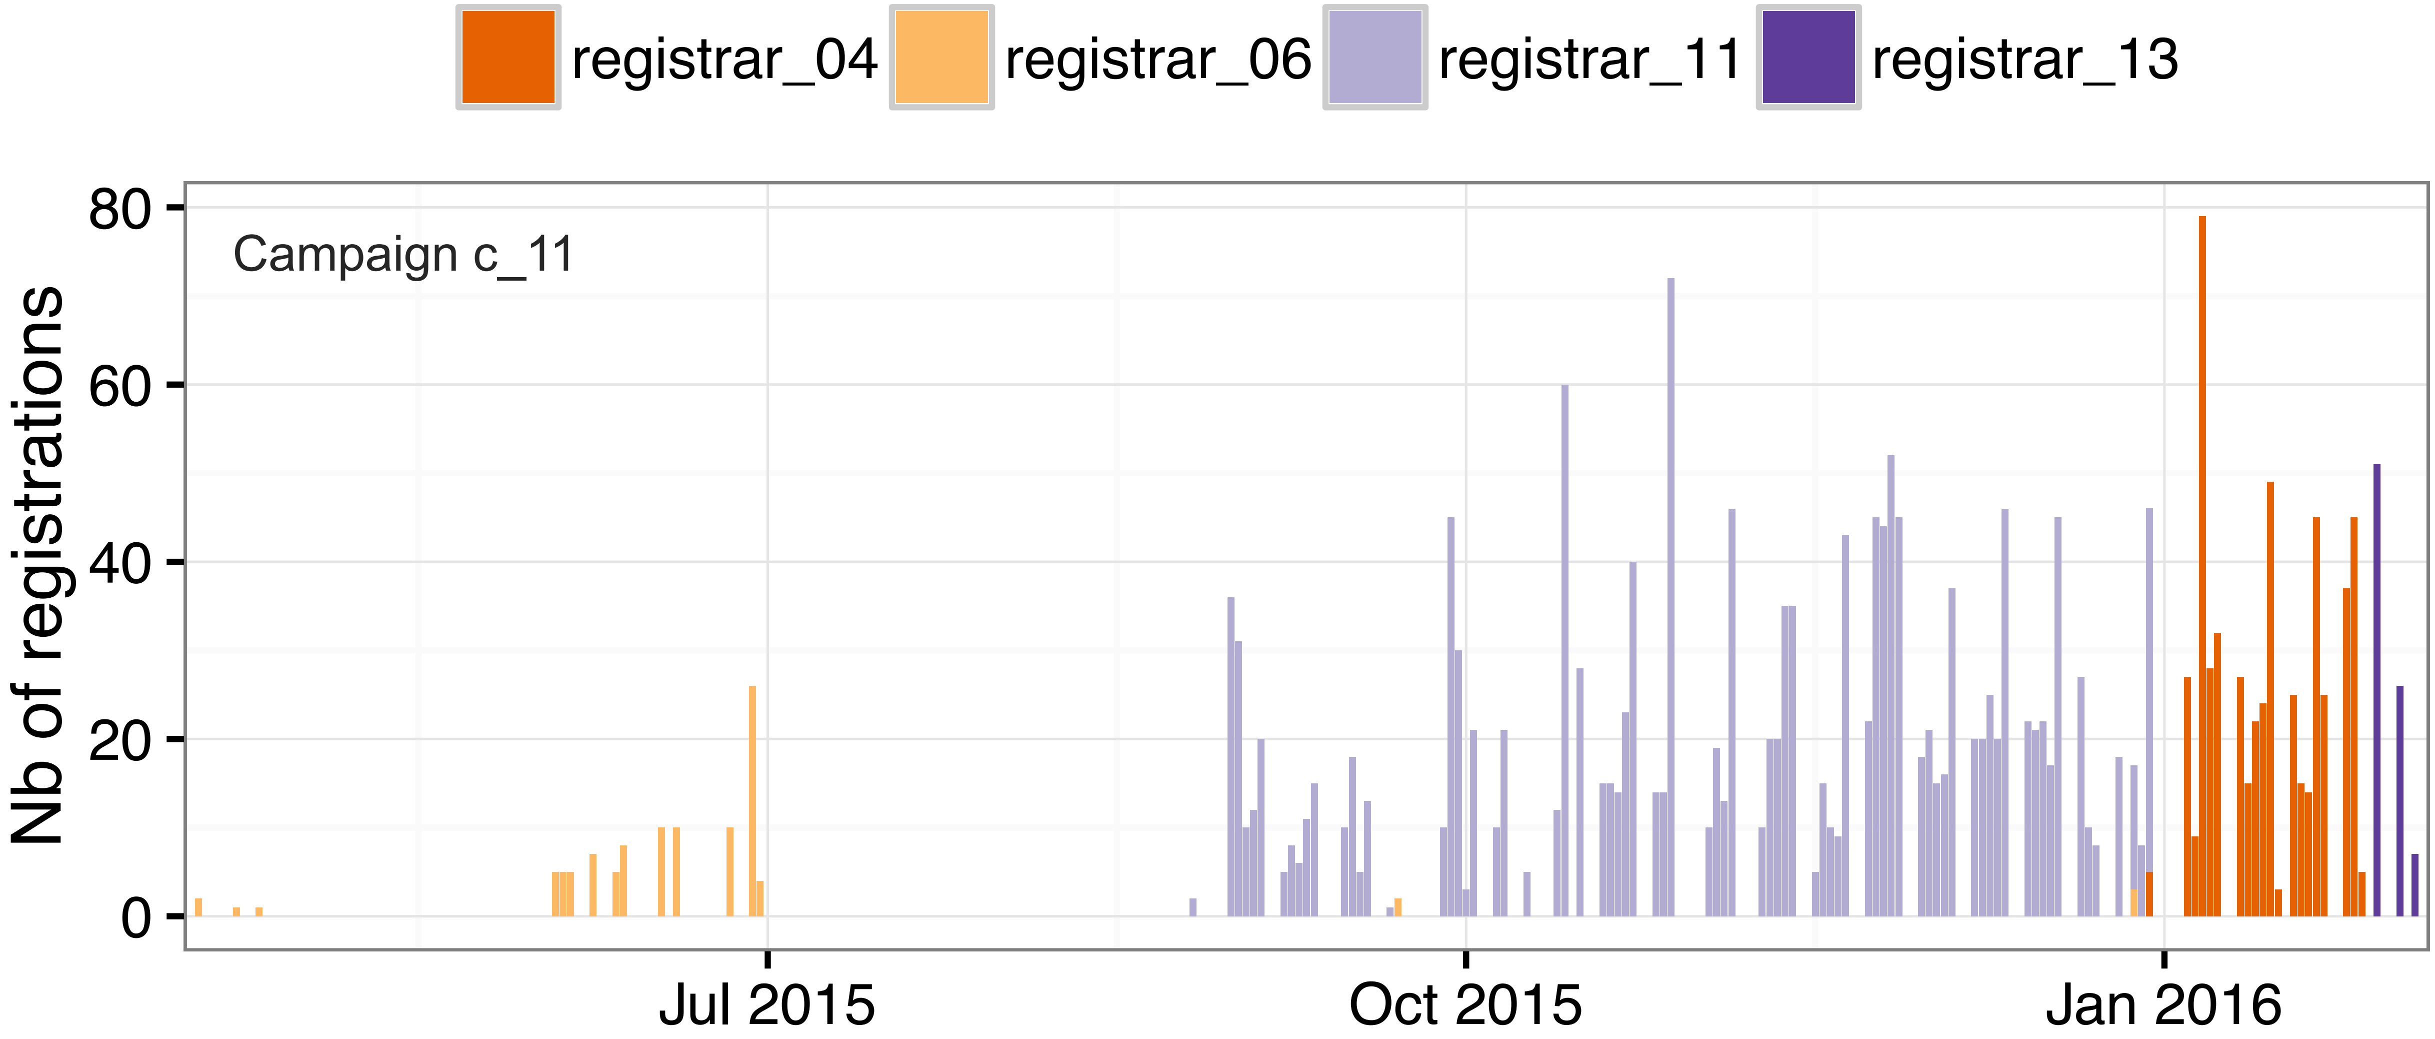

Insight 9: Campaigns vs blacklists

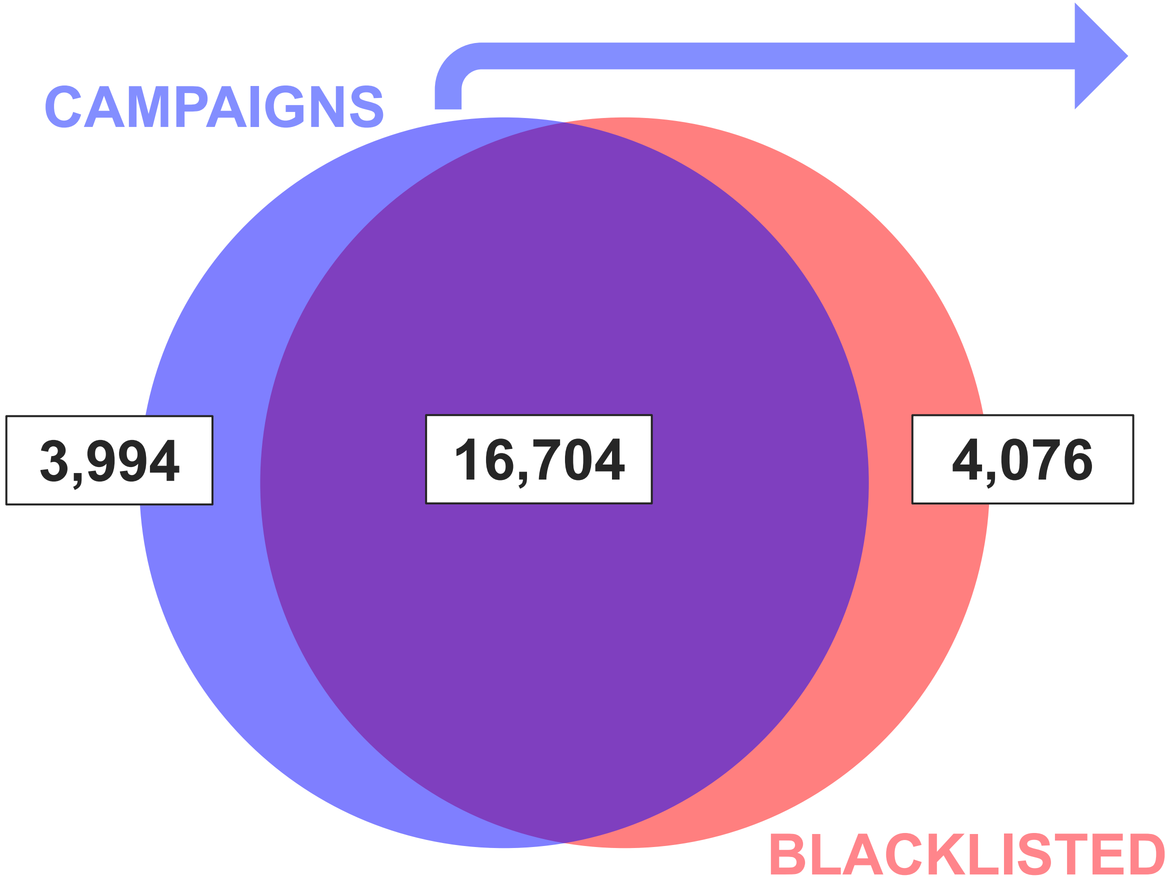

- Manual analysis of non-blacklisted domains
- Result: < 1% false positives
- About 20% extra on top of existing blacklists

Pro-active detection and prevention

“Could newly registered domains with malicious intent be detected or prevented at registration time?”

Pro-active detection and prevention*

- Based on previous domain registrations, prediction models are trained:
  - Similarity-based agglomerative clustering
  - Reputation-based classification
- For each new registration, the system predicts if the domain will be used for malicious activity
- Domains with malicious intent can be
  - Early detected
  - Prevented from being registered

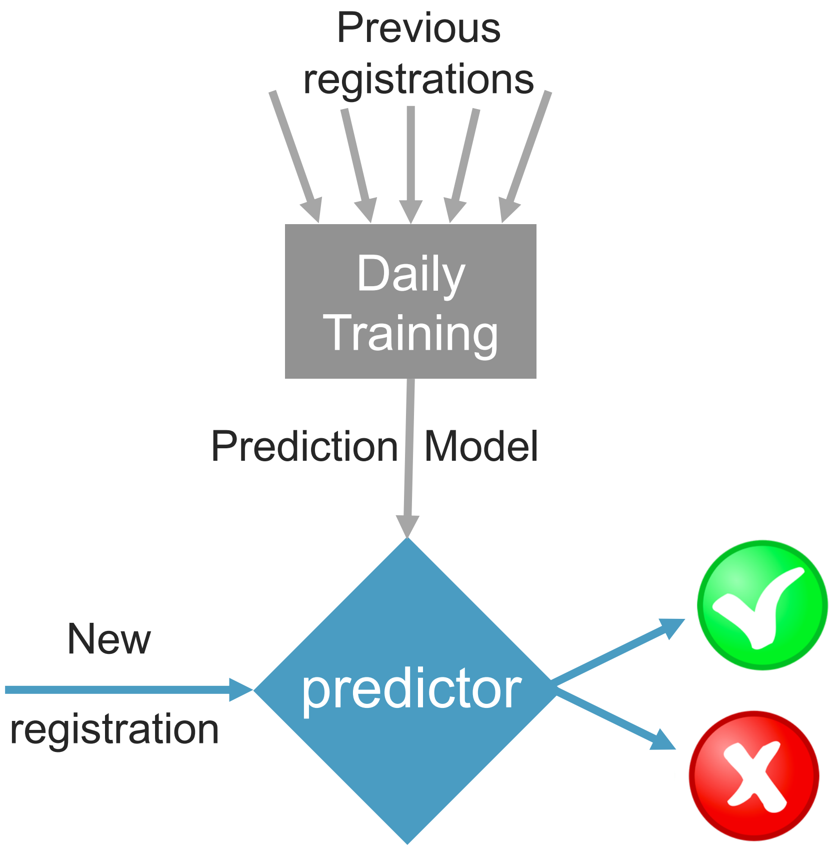

* Patent pending

Underlying assumptions/rationales

- Similarity-based agglomerative clustering
  - Domains belonging to the same campaign have very similar registration details
- Reputation-based classification
  - Domains belonging to registrants with a bad reputation, are likely to be malicious as well

In operation at EURid …

- Deployed as part of EURid’s Trust & Security program

- *Preliminary* results of first 110 days in production[1] :
  - 80% of malicious domain registrations have been predicted
  - 98% precision: ~1 false positive per day

[1] Results of the best performing predictor

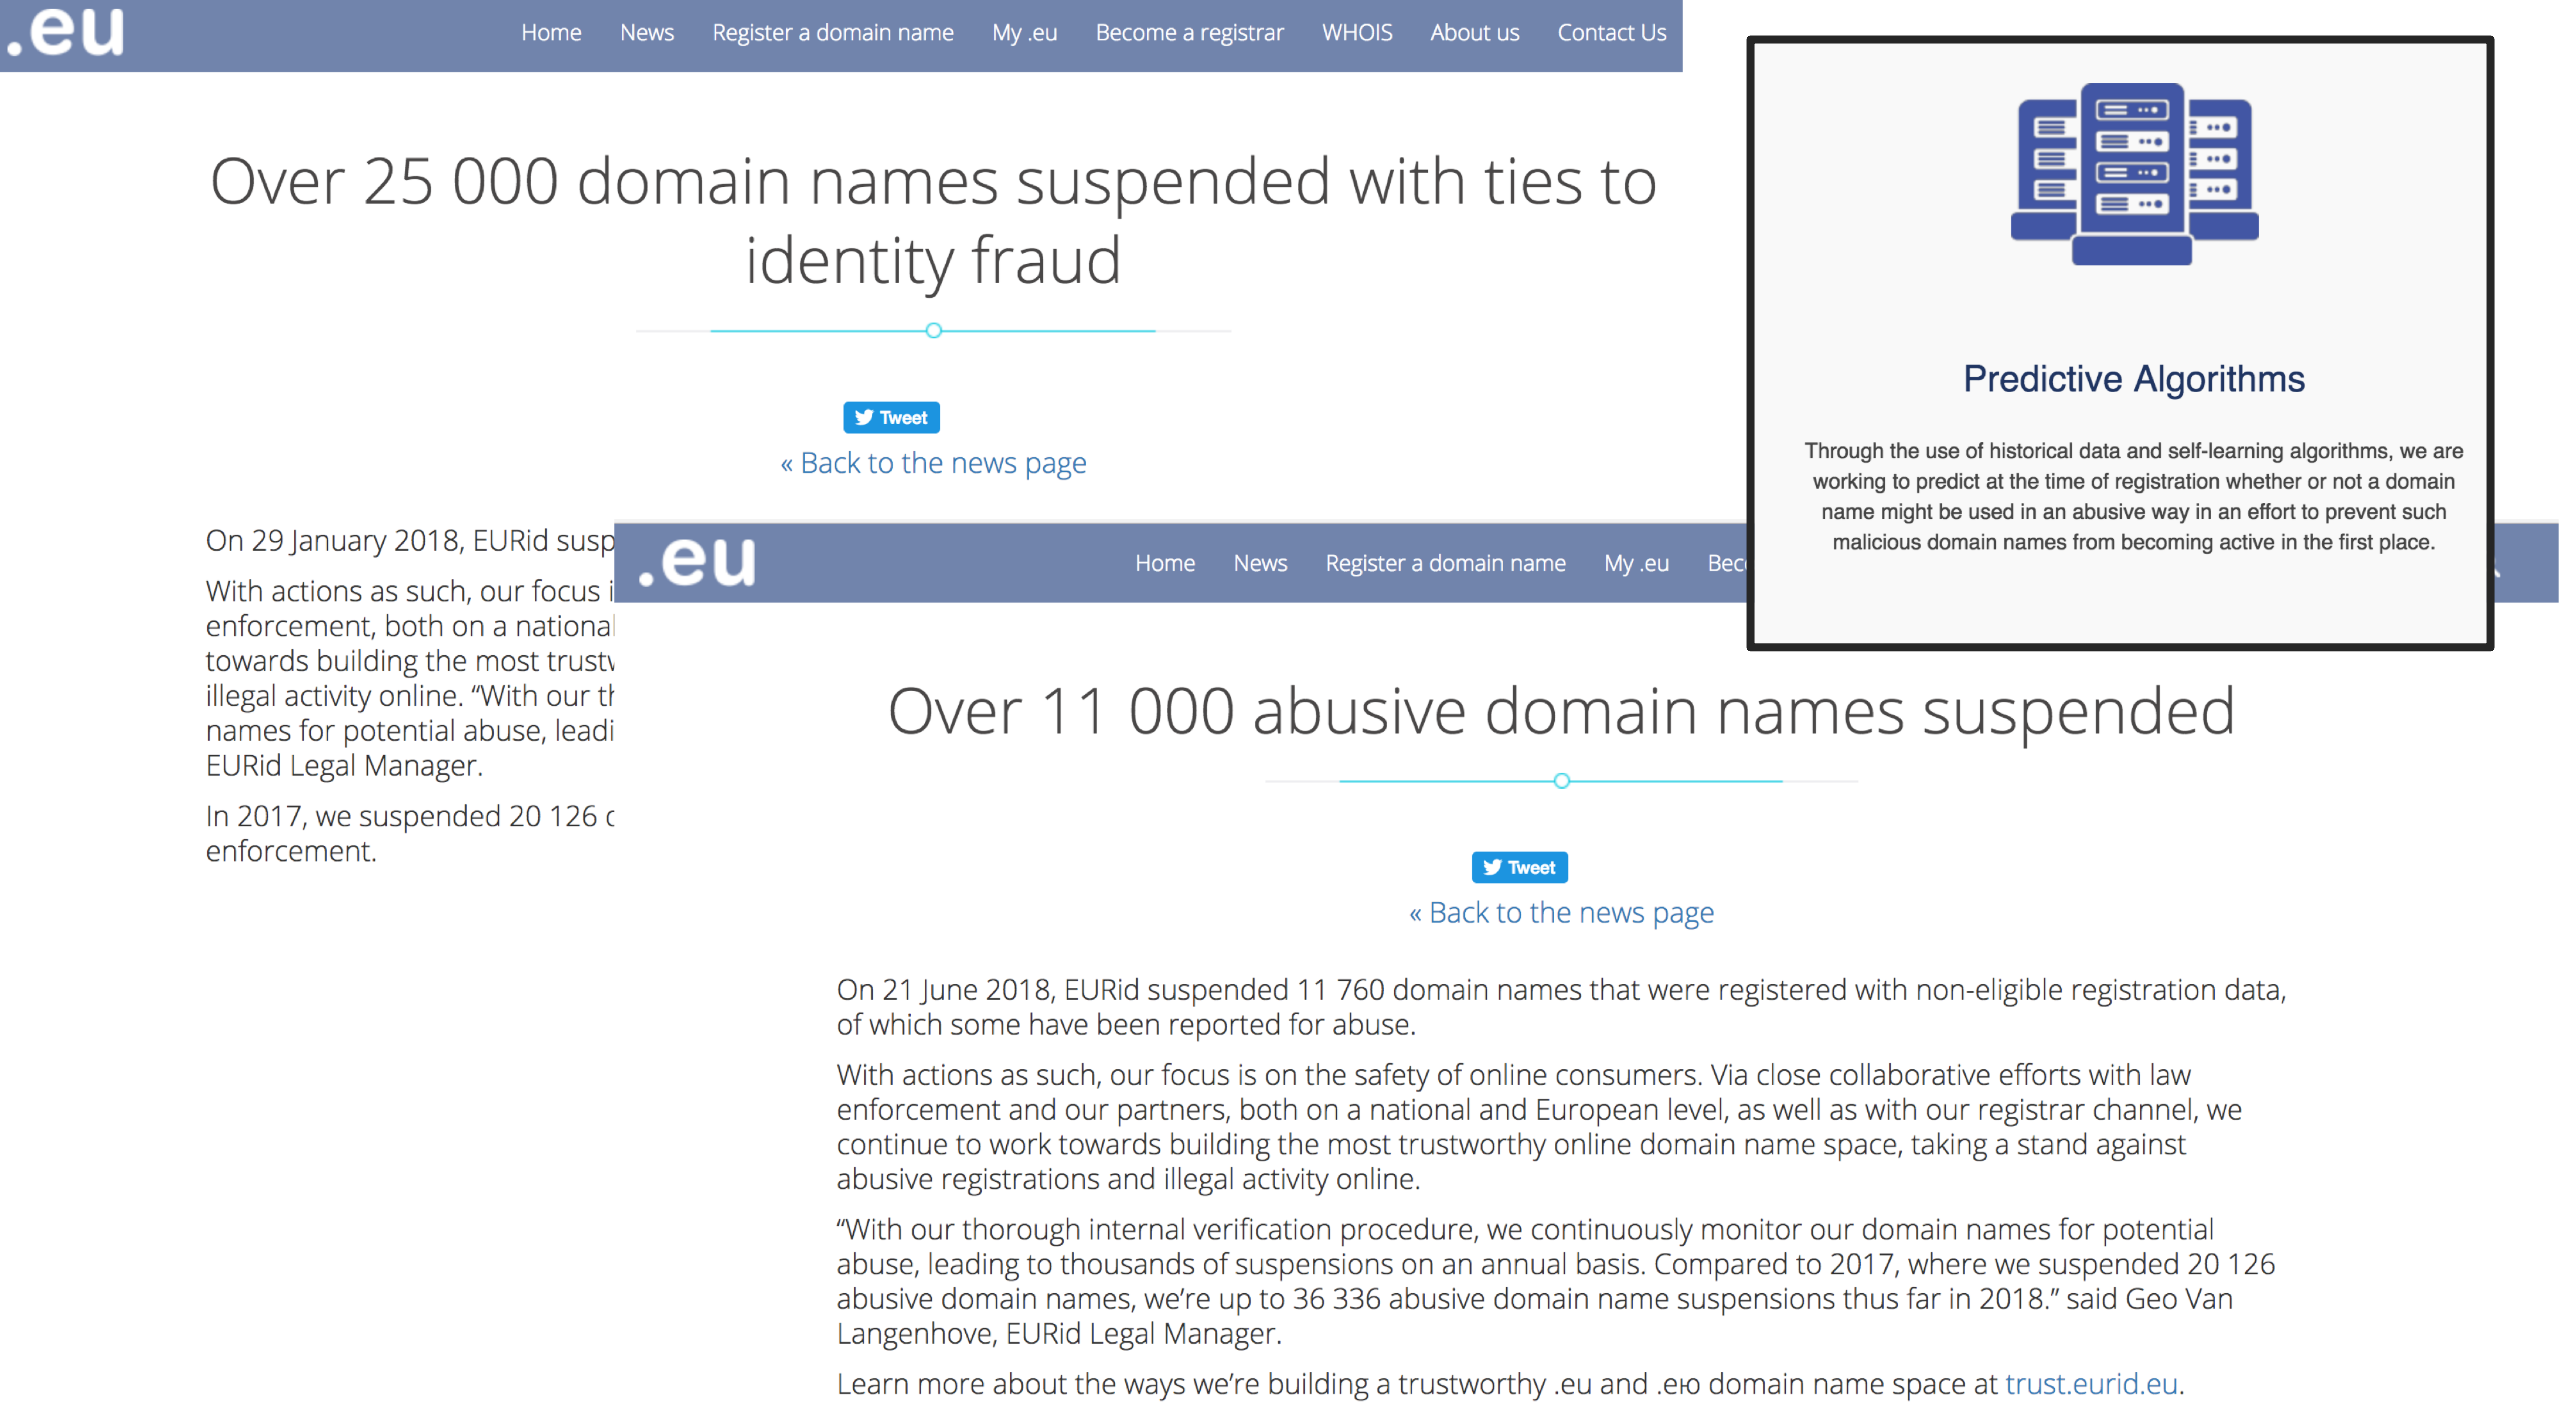

.eu Home News Register a domain name My .eu Become a registrar WHOIS About us Contact Us

Over 25 000 domain names suspended with ties to identity fraud

Tweet

« Back to the news page

On 29 January 2018, EURid susp

With actions as such, our focus i
enforcement, both on a national
towards building the most trustv
illegal activity online. "With our th
names for potential abuse, leadi
EURid Legal Manager.

In 2017, we suspended 20 126 c
enforcement.

Predictive Algorithms

Through the use of historical data and self-learning algorithms, we are working to predict at the time of registration whether or not a domain name might be used in an abusive way in an effort to prevent such malicious domain names from becoming active in the first place.

.eu Home News Register a domain name My .eu Bec

Over 11 000 abusive domain names suspended

Tweet

« Back to the news page

On 21 June 2018, EURid suspended 11 760 domain names that were registered with non-eligible registration data, of which some have been reported for abuse.

With actions as such, our focus is on the safety of online consumers. Via close collaborative efforts with law enforcement and our partners, both on a national and European level, as well as with our registrar channel, we continue to work towards building the most trustworthy online domain name space, taking a stand against abusive registrations and illegal activity online.

"With our thorough internal verification procedure, we continuously monitor our domain names for potential abuse, leading to thousands of suspensions on an annual basis. Compared to 2017, where we suspended 20 126 abusive domain names, we're up to 36 336 abusive domain name suspensions thus far in 2018." said Geo Van Langenhove, EURid Legal Manager.

Learn more about the ways we're building a trustworthy .eu and .ею domain name space at trust.eurid.eu.

Key takeaways

Rather small set of bad actors

- Up to 20 campaigns are responsible for 80% of malicious registrations

- Top facilitators:
  - About half of the malicious registrations via 1 registrar
  - 1 public email provider are malicious with a high toxicity

Cyber criminals are “human”….

- Lazy
  - Reuse same fake registrants
  - Use generators for registrant details
- Work force
  - Work 9 to 5
  - Take week-ends, holidays
  - Make mistakes (e.g. typos)
- Adapt over time

Pro-active detection and prevention

- Early results look promising
  - Captures the majority of malicious domain registrations
  - Operates at a low false-positive rate

- Interesting to see how this will impact the security landscape

Interested in more?

- Thomas Vissers, Jan Spooren, Pieter Agten, Dirk Jumpertz, Peter Janssen, Marc Van Wesemael, Frank Piessens, Wouter Joosen, Lieven Desmet, **Exploring the ecosystem of malicious domain registrations in the .eu TLD**, Research in Attacks, Intrusions, and Defenses, (RAID 2017), Atlanta, USA, September 18-20, 2017

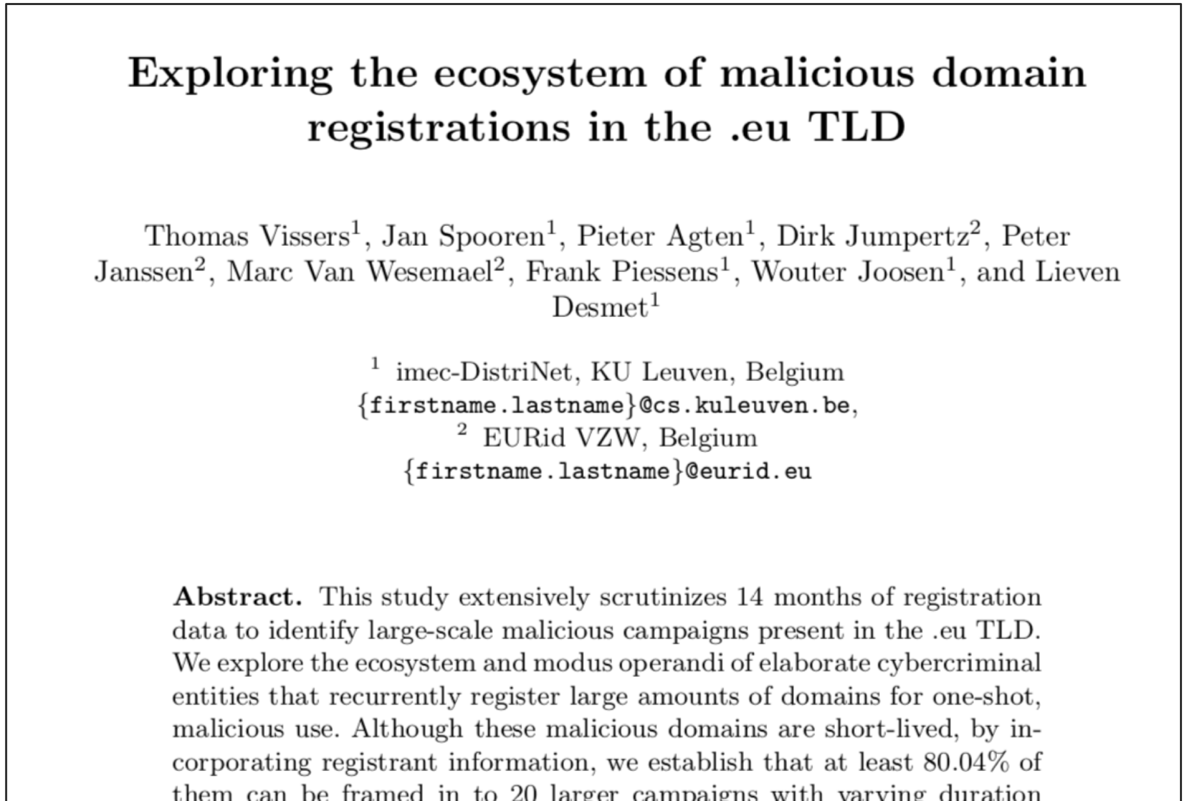

Exploring the ecosystem of malicious domain registrations in the .eu TLD

Thomas Vissers[1], Jan Spooren[1], Pieter Agten[1], Dirk Jumpertz[2], Peter Janssen[2], Marc Van Wesemael[2], Frank Piessens[1], Wouter Joosen[1], and Lieven Desmet[1]

[1] imec-DistriNet, KU Leuven, Belgium
{firstname.lastname}@cs.kuleuven.be,
[2] EURid VZW, Belgium
{firstname.lastname}@eurid.eu

Abstract. This study extensively scrutinizes 14 months of registration data to identify large-scale malicious campaigns present in the .eu TLD. We explore the ecosystem and modus operandi of elaborate cybercriminal entities that recurrently register large amounts of domains for one-shot, malicious use. Although these malicious domains are short-lived, by incorporating registrant information, we establish that at least 80.04% of them can be framed in to 20 larger campaigns with varying duration

Final version:
https://doi.org/10.1007/978-3-319-66332-6_21

OWASP
AppSec Europe
London 2nd-6th July 2018

Detecting and Preventing Malicious Domain Registrations in the .eu TLD

Lieven Desmet