

# Suspended Accounts in Retrospect: An Analysis of Twitter Spam

**Kurt Thomas**, Chris Grier,  
Vern Paxson, Dawn Song

University of California, Berkeley  
International Computer Science Institute

# Motivation

- Social networks are regular targets for abuse
  - 26% of URLs on Twitter lead to spam
  - 10% of URLs on Facebook
- Underground economy targeting social networks remains obscure
  - Account behavior
  - Challenges faced by spammers

# Our Study

- We examine 7 months of Twitter spam:
  - Accounts, techniques used to distribute spam
  - Emerging support infrastructure
  - Sophistication of campaigns
- Characterize performance of existing defenses
  - Hostile environment for spammers
  - Low barriers, persistent attacks

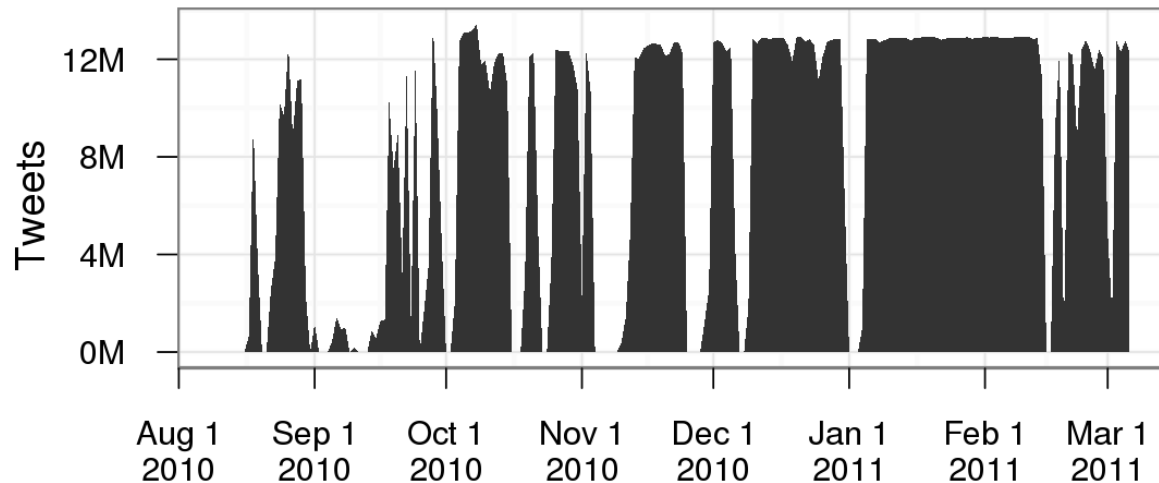
# Outline

- Dataset
- Accounts, URLs, Tools
- Spam as a service
- Campaigns

**DATASET**

# Sample Size

- Daily sample from Twitter, 12M tweets/day
  - Limit to tweets with URLs
  - Some days missing: network outages, crashes

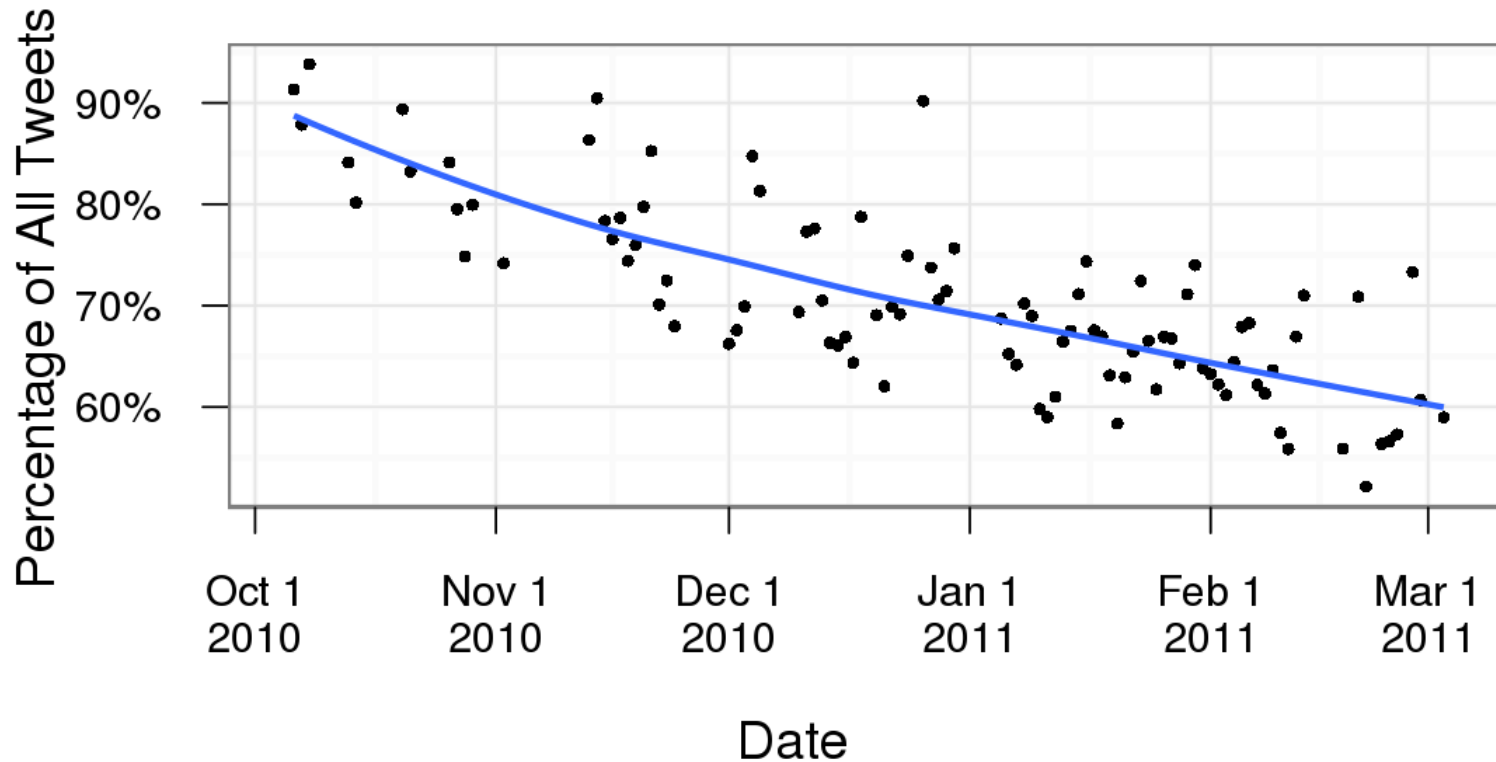


**Tweets:** 1.8 billion

**Accounts:** 32.8 million

# Sample Rate

- Fraction of all tweets with URLs:



# Identifying Spam

- Which of 1.8 billion tweets are spam?
- Previous approaches:
  - Blacklisting [*Grier et al, CCS 2010*]
  - Heuristics (bursty, distributed) [*Gao et al. IMC 2010*]
- Our approach:
  - Identify accounts suspended by Twitter
  - **Only 8% of their URLs appear in blacklists**



# Suspended Accounts

## Who goes there?

Sorry, the account you were headed to has been suspended due to strange activity. [Mosey along now](#), nothing to see here.



All done here?

Take me home!


...or if you're curious as to why an account might be suspended, [head over this way](#) for the juicy details.

# Suspended.equals(Spam)?

- **True positives:** 93% of suspended accounts were spammers
  - Remaining 7% aggressive marketing, following
  - None were false positives
- **False negatives:** 6% of active accounts were uncaught spammers

**Twitter catches 37% of spam accounts**

# Spam Dataset



**Spam Accounts:** 1.1 million

**Spam Tweets:** 80 million

**Spam URLs:** 37 million

# **TOOLS OF THE SPAM TRADE**

# Taxonomy of Social Spam

- Three components to social network spam:



**Cassie Aguada**

**@CassieAguada973**

[view full profile](#) →

**5**

Tweets

**0**

Following

**0**

Followers

**0**

Listed

## Recent Tweets



**CassieAguada973** Cassie Aguada

[@Tobiasgrundesjo](#) Omg this kind of good curves

[bit.ly/oSpCbt?=-xk5i](https://bit.ly/oSpCbt?=-xk5i)

27 minutes ago



**CassieAguada973** Cassie Aguada

This is how to earn money? [bit.ly/oSpCbt?=-ed8o](https://bit.ly/oSpCbt?=-ed8o)

[#WeCantWait](#)

1 hour ago

# Taxonomy of Social Spam

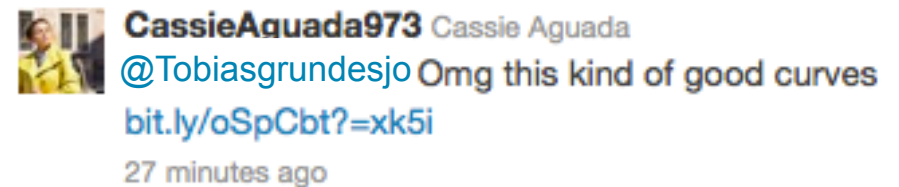
- Three components to social network spam:
  - Accounts



A screenshot of a Twitter profile card for Cassie Aguada. The card features a profile picture of a woman in a yellow jacket, her name 'Cassie Aguada', and her handle '@CassieAguada973'. Below the name is a link to 'view full profile'. At the bottom, there are four statistics: 5 Tweets, 0 Following, 0 Followers, and 0 Listed.

5	0	0	0
Tweets	Following	Followers	Listed

## Recent Tweets



A screenshot of a recent tweet from Cassie Aguada (@CassieAguada973). The tweet text is '@Tobiasgrundesjo Omg this kind of good curves bit.ly/oSpCbt?=xk5i' and it was posted 27 minutes ago.



A screenshot of another recent tweet from Cassie Aguada (@CassieAguada973). The tweet text is 'This is how to earn money? bit.ly/oSpCbt?=ed8o #WeCantWait' and it was posted 1 hour ago.

# Taxonomy of Social Spam

- Three components to social network spam:
  - Accounts
  - Engagement
    - Social graph



**Cassie Aguada**

**@CassieAguada973**

[view full profile](#) →

**5**  
Tweets

**0**  
Following

**0**  
Followers

**0**  
Listed

## Recent Tweets



**CassieAguada973** Cassie Aguada

[@Tobiasgrundesjo](#) Omg this kind of good curves

[bit.ly/oSpCbt?=-xk5i](https://bit.ly/oSpCbt?=-xk5i)

27 minutes ago



**CassieAguada973** Cassie Aguada

This is how to earn money? [bit.ly/oSpCbt?=-ed8o](https://bit.ly/oSpCbt?=-ed8o)

[#WeCantWait](#)

1 hour ago

# Taxonomy of Social Spam

- Three components to social network spam:
  - Accounts
  - Engagement
    - Social graph
    - Mention



**Cassie Aguada**

**@CassieAguada973**

[view full profile →](#)

**5**  
Tweets

**0**  
Following

**0**  
Followers

**0**  
Listed

## Recent Tweets



**CassieAguada973** Cassie Aguada

[@Tobiasgrundesjo](#) Omg this kind of good curves

[bit.ly/oSpCbt?=-xk5i](https://bit.ly/oSpCbt?=-xk5i)

27 minutes ago



**CassieAguada973** Cassie Aguada

This is how to earn money? [bit.ly/oSpCbt?=-ed8o](https://bit.ly/oSpCbt?=-ed8o)

[#WeCantWait](#)

1 hour ago



# Taxonomy of Social Spam

- Three components to social network spam:
  - Accounts
  - Engagement
    - Social graph
    - Mention
    - Hashtag



**Cassie Aguada**  
**@CassieAguada973**  
view full profile →

**5** Tweets | **0** Following | **0** Followers | **0** Listed

## Recent Tweets



**CassieAguada973** Cassie Aguada  
**@Tobiasgrundesjo** Omg this kind of good curves  
[bit.ly/oSpCbt?=-xk5i](https://bit.ly/oSpCbt?=-xk5i)  
27 minutes ago



**CassieAguada973** Cassie Aguada  
This is how to earn money? [bit.ly/oSpCbt?=-ed8o](https://bit.ly/oSpCbt?=-ed8o)  
[#WeCantWait](#)  
1 hour ago

# Taxonomy of Social Spam

- Three components to social network spam:
  - Accounts
  - Engagement
    - Social graph
    - Mention
    - Hashtag
  - Monetizable URLs



**Cassie Aguada**

**@CassieAguada973**

[view full profile →](#)

**5**

Tweets

**0**

Following

**0**

Followers

**0**

Listed

## Recent Tweets



**CassieAguada973** Cassie Aguada

**@Tobiasgrundesjo** Omg this kind of good curves

[bit.ly/oSpCbt?=-xk5i](https://bit.ly/oSpCbt?=-xk5i)

27 minutes ago



**CassieAguada973** Cassie Aguada

This is how to earn money? [bit.ly/oSpCbt?=-ed8o](https://bit.ly/oSpCbt?=-ed8o)

**#WeCantWait**

1 hour ago

# Accounts

- Of 1.1 million spam accounts, median stats:

**Lifetime:** **1** day

From first sampled tweet to last tweet

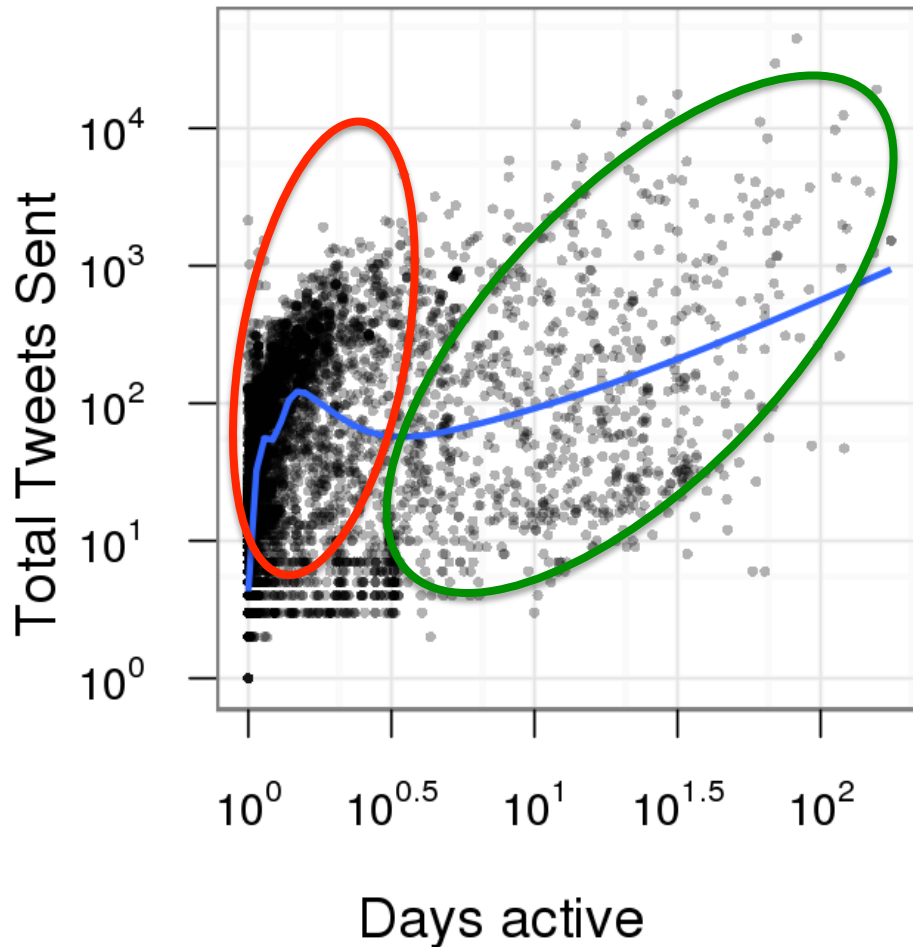
**Followers:** **2**

Users receiving a spammers content

**Tweets:** **13**

Total tweets sent; not just those sampled

# Tweets




- Two strategies for tweeting:
  - Short-lived, **bursty** posting
  - Long-lived, **low daily** posting


# Engagement

- Spam accounts forgo social graph for distribution
  - 89% of accounts < 10 followers
- Unsolicited @mentions: 52% of accounts
- #Hashtags, trends: 17% of accounts
- Alternatives:
  - Direct messages
  - Search

# Social Activity




**Regular Twitter User**



**Spam Account**

# Spam URLs

- Of 37 million URLs, two strategies:



# URLs – Shorteners

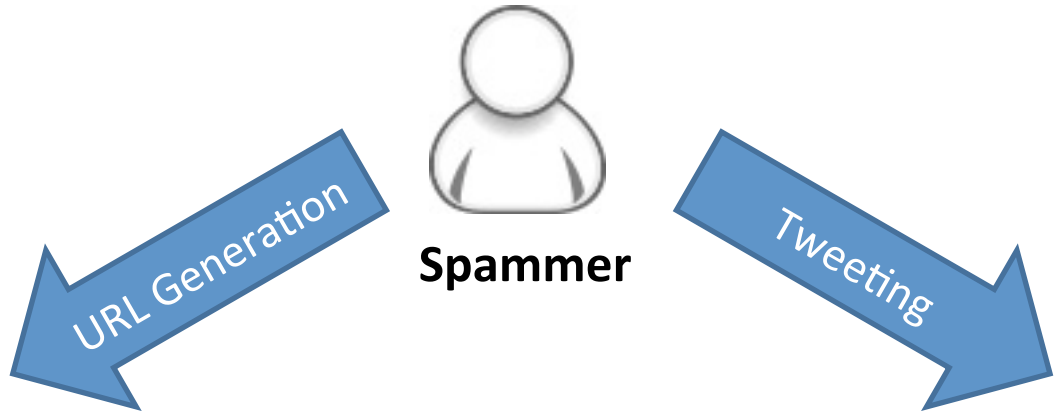
- 60% of spam tweets contain shortened URLs
  - Over 317 shorteners

Shortening Service	Fraction of Tweets
Bit.ly	35%
Tinyurl.com	7%
Is.gd	2%
Goo.gl	2%
Ow.ly	2%



**SPAM AS A SERVICE**

# Centralized Distribution



## Spam Content



## Twitter Account



**Nicole C.**  
**@CheapCialisNow**

# Spam-as-a-Service



**Content Controller**

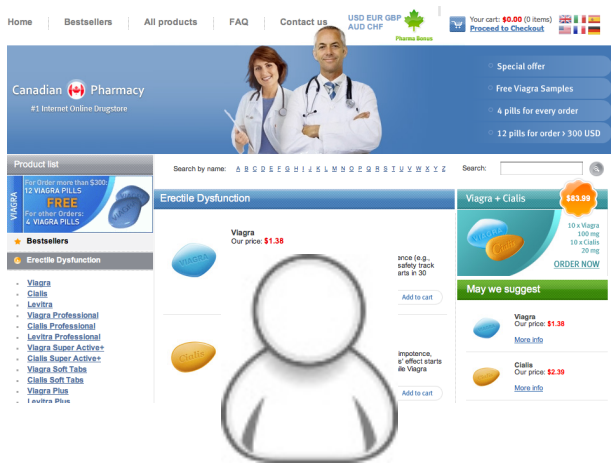


**Nicole C.**  
**@CheapCialisNow**




**Account Controller**

# Spam-as-a-Service



Content Controller



**Nicole C.**  
**@CheapCialisNow**



Account Controller



Affiliate Service



# Affiliate Programs

- Examples:
  - Clickbank.com – affiliate paid per click
  - Amazon.com – affiliate paid per purchase
- Identify mapping between affiliates, Twitter accounts

# Affiliate Programs

Program	Tweets	Twitter Accounts	Affiliates
Clickbank.com	3 million	16,309	203
Amazon.com	1 million	8,129	919

- Clickbank affiliate
  - Median: 565 tweets
  - Max: 217,686 tweets
- Amazon affiliate
  - Median: 2 tweets
  - Max: 324,613 tweets

# Alternative Services

- Account sellers, controllers
  - Ex: spn.tw, Assetize
  - Mediate access to accounts
  - 200,000 tweets, 1600 accounts
- Ad-based URL shorteners
  - Ex: vur.me, eca.sh
  - Embed shortened URL in iFrame, surround by ads
  - 360,000 tweets, 400 accounts

# **SPAM CAMPAIGNS**



# All Together Now – Campaigns

- Cluster accounts, tweets into campaigns
  - Automated heuristics
  - Manual validation
- Identify 5 of the largest Twitter spam campaigns
  - 145,000 accounts (13% of dataset)
  - 18 million tweets (22% of dataset)

# Campaign: *Amazon Abuse*

## Duration:

1 month

## Accounts:

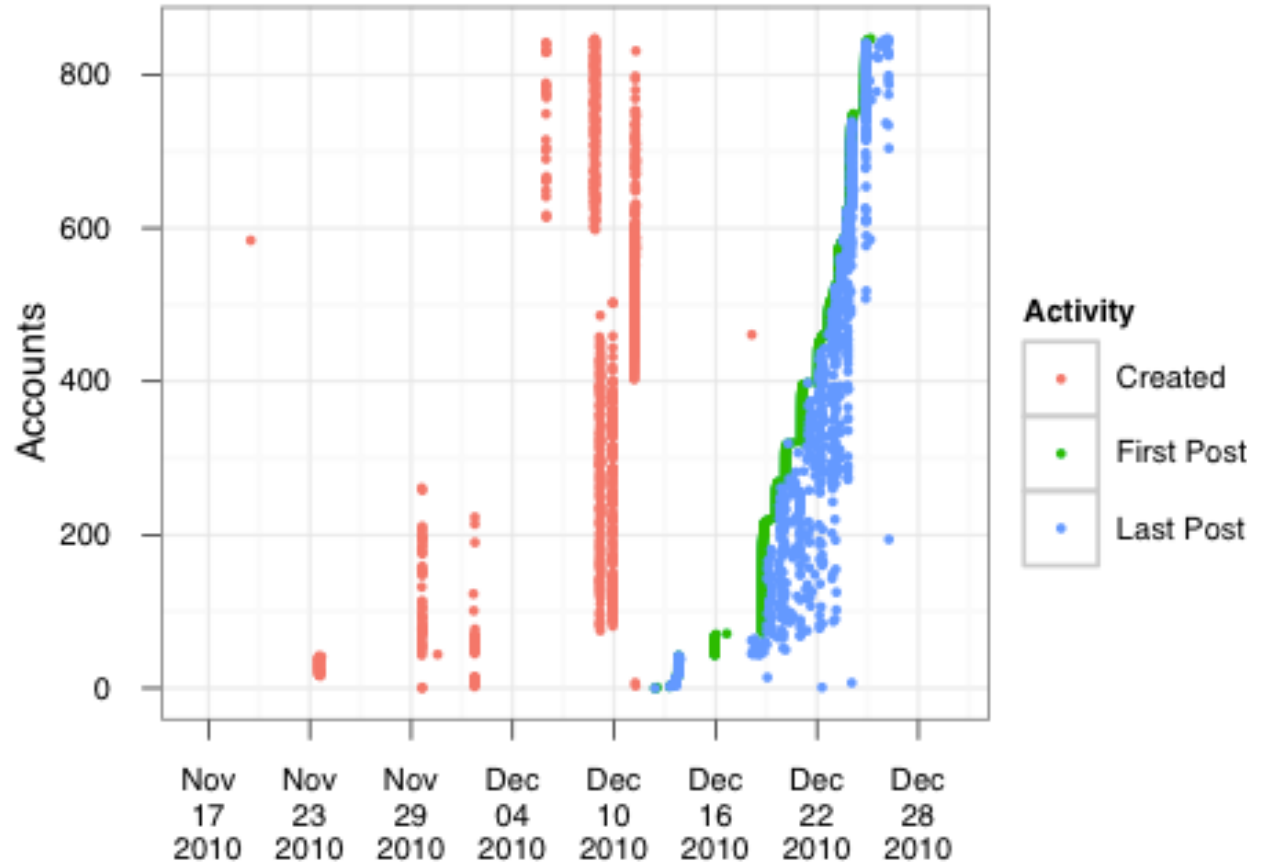
848

## Tweets:

129,600

## Engagement:

Mentions



# Campaign: *Clickbank Abuse*

## Duration:

7 months

## Accounts:

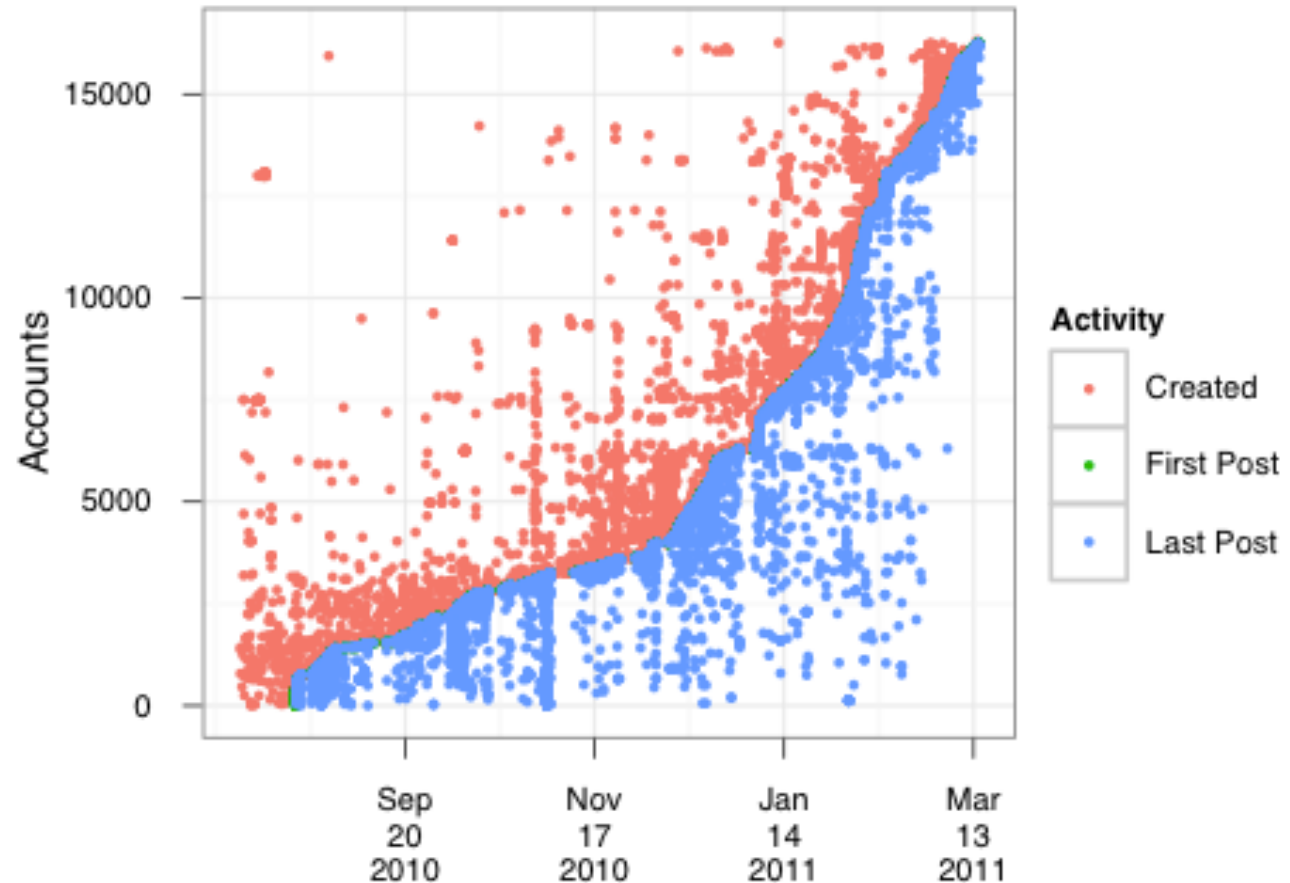
16,000

## Tweets:

3 million

## Engagement:

Mentions,  
hashtags



# Campaign: *Pharmacy*

## Duration:

1 day

## Accounts:

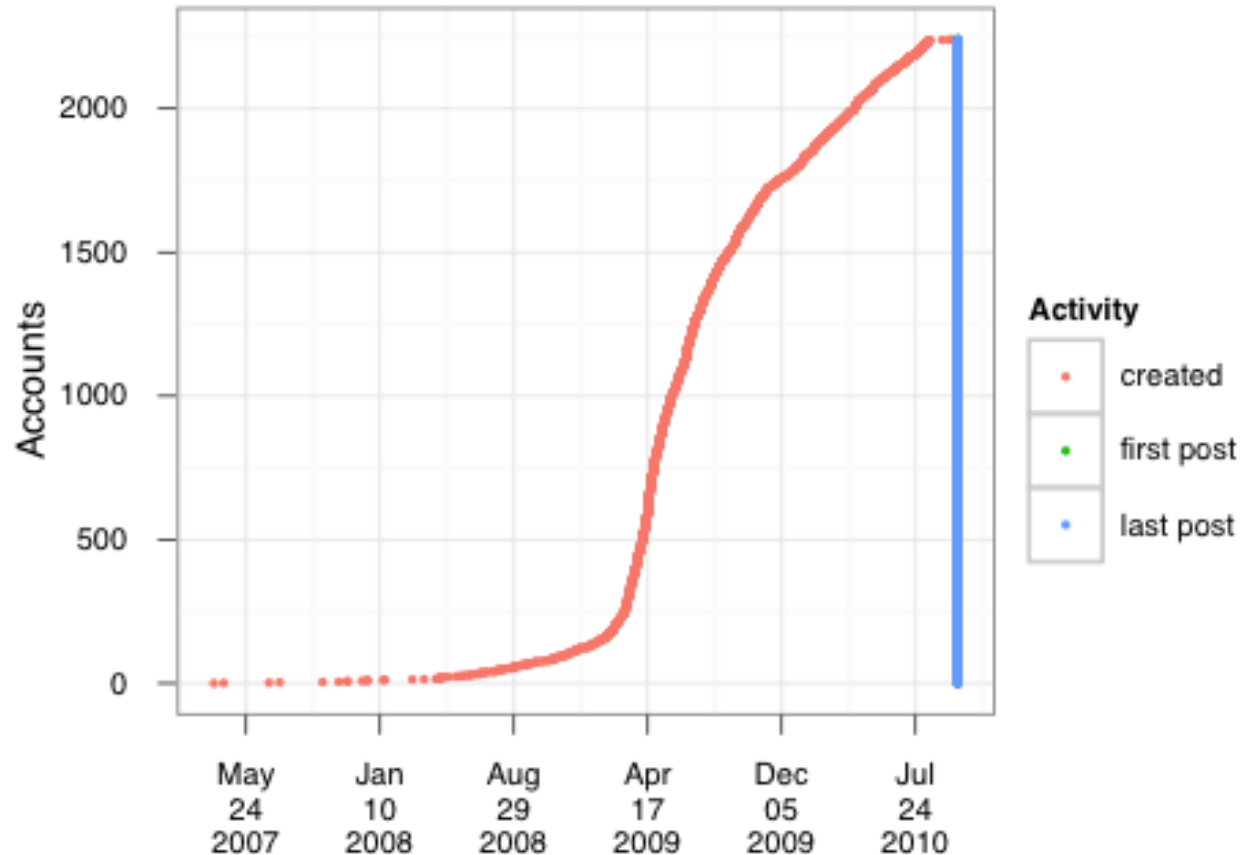
2,200

## Tweets:

130,000

## Engagement:

Trend  
hijacking



# Conclusion

- Variety of spam strategies
- Fledgling spam-as-a-service market
  - Affiliate programs
  - Account providers
- Complex spam campaigns
  - Low barrier to creating accounts
  - Weak defenses, slow response

**Solutions must target both social networks  
and the support infrastructure**