Click Fraud Detection using Practical Memetics

Defcon 15, Aug 3-5, 2007

Broward Horne

http://www.realmeme.com

Overview

- Advertising Click Fraud
- Botnets (Agents Of Click Fraud)
- S-Curve
- Original Meme Theory (Dawkins)
- Empirical Meme Theory (Me! Ho!)
- Expanded Meme Mining Model
- MySpace Example
- IAFF.com Example
- Meme Seepage Theory
- Botnet Proof
- Gaming Conjectures
- Extrapolations

Click Fraud - Definition

From Wikipedia -

> “When a person, script or computer program imitates a legitimate user to generate a “charge per click” of an advertised product”

Is Click Fraud relatively unknown? Poll

Google Adsense Model

- www.google.com/adsense
- Click fraud risk

Botnets

- From Wikipedia-

"a collection of compromised computers ("zombies") running… under a common command-and-control infrastructure"

S-Curve (Physical Model)

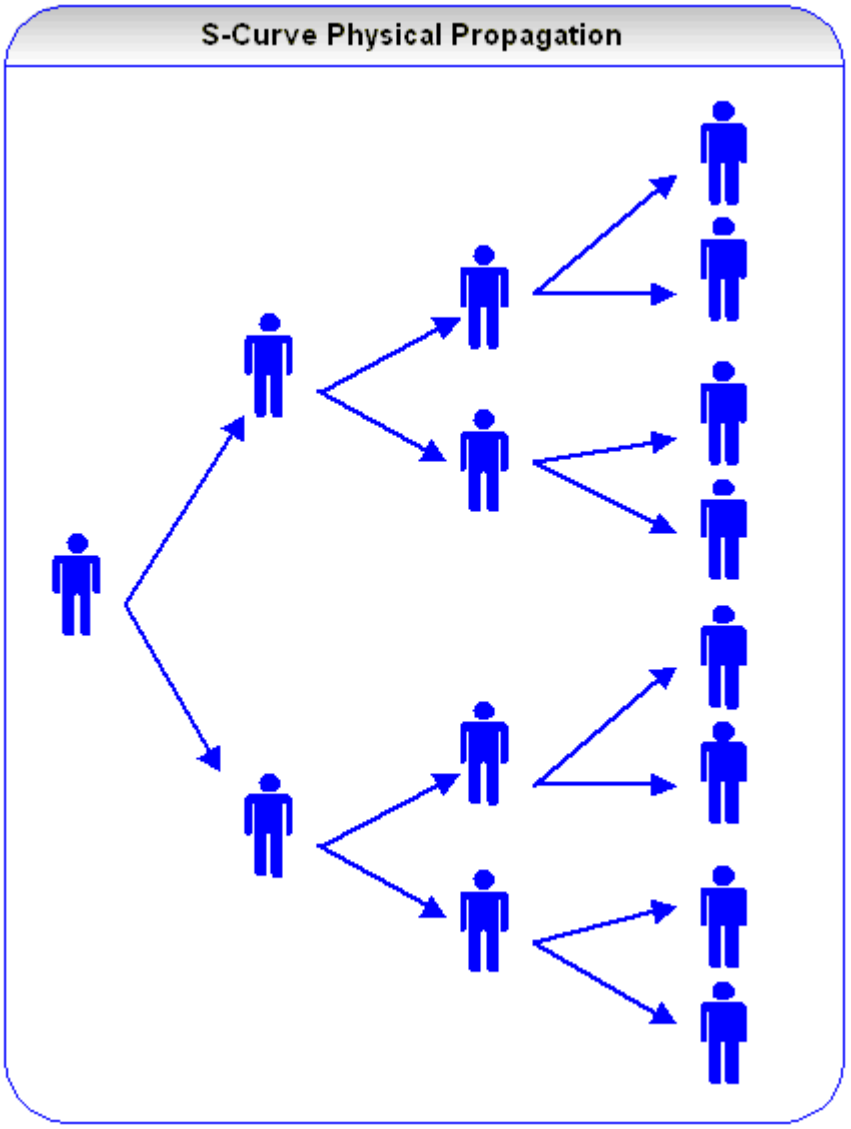

S-Curve (Math Model)

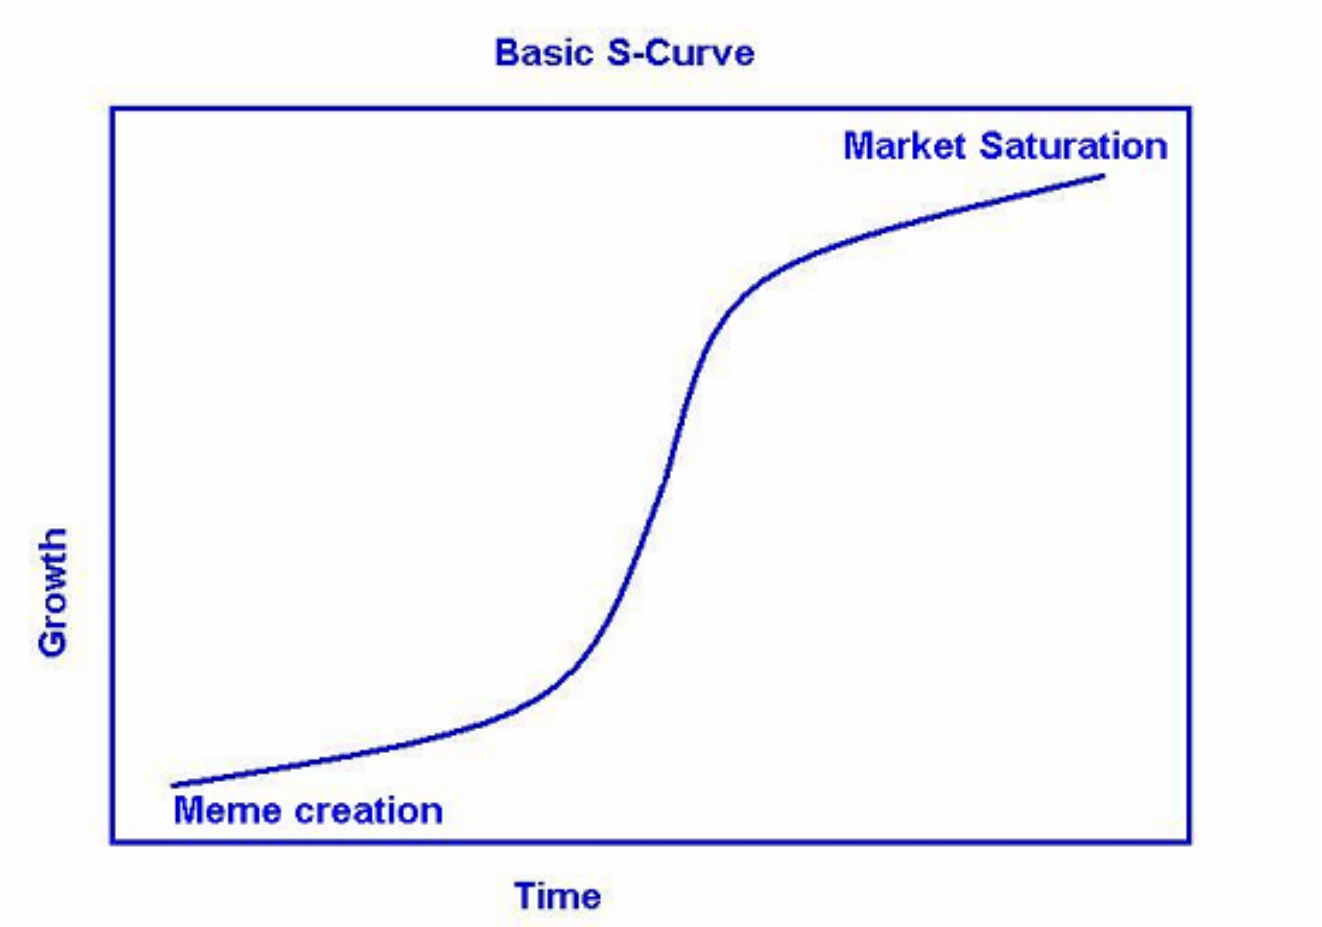

S-Curve Rate-Of-Change

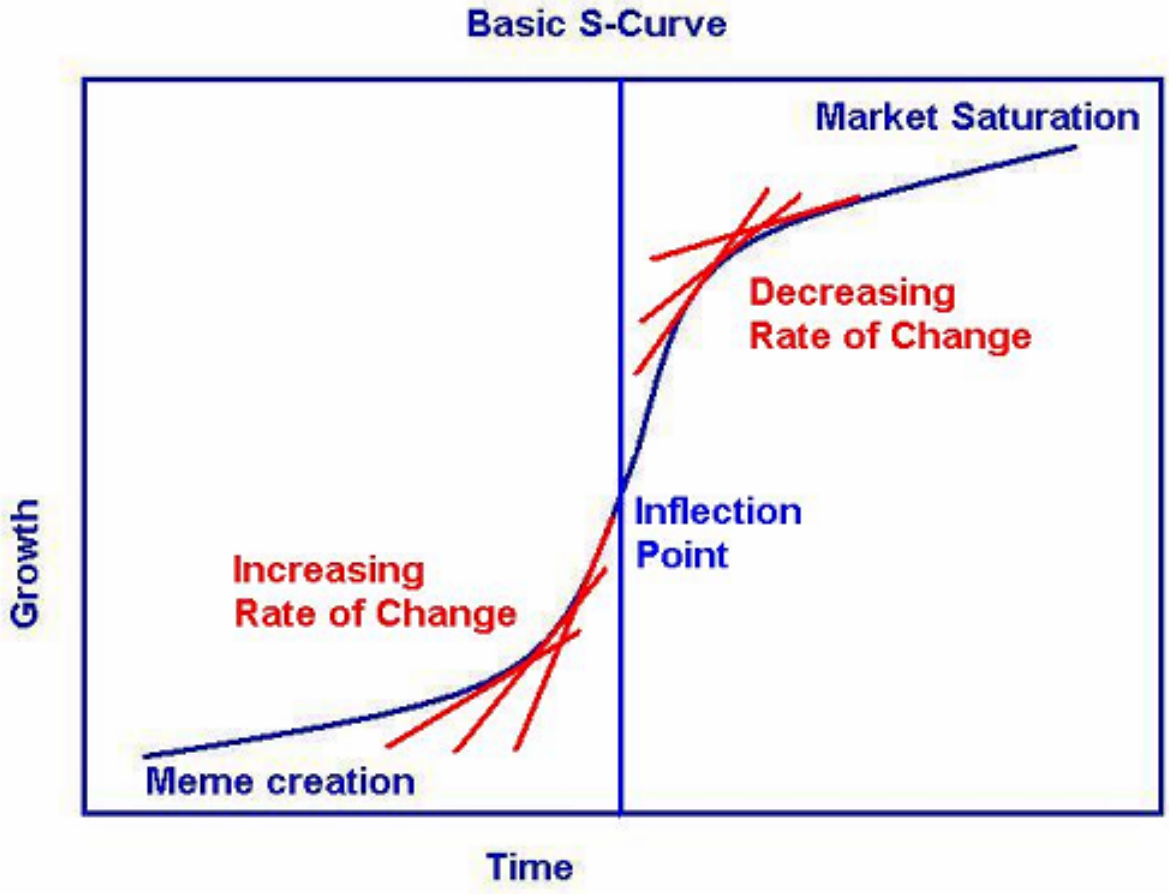

S-Curve Strategy

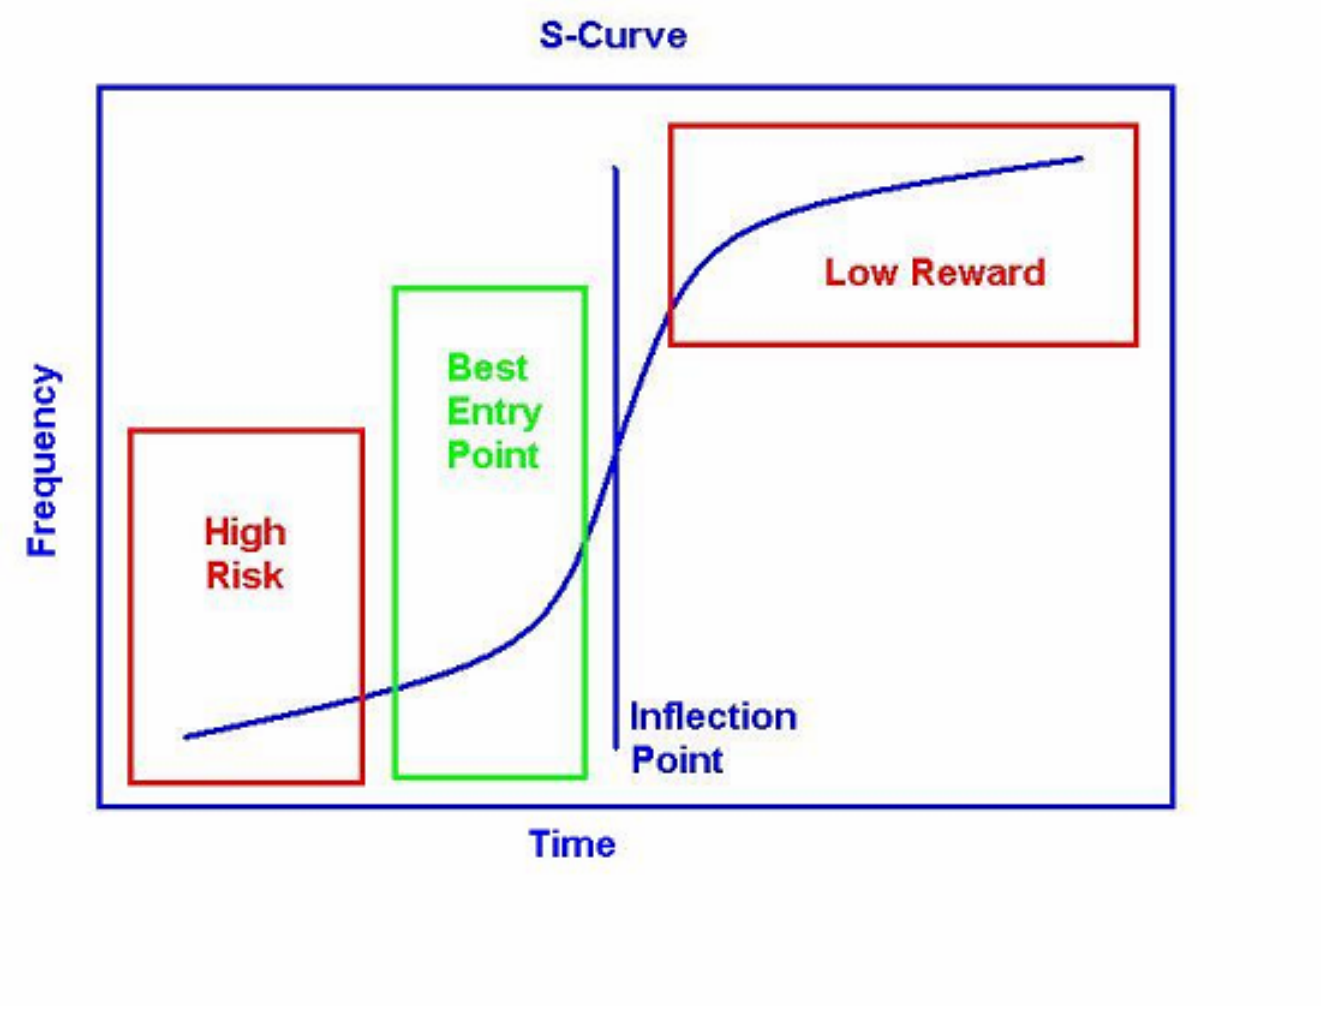

Dawkins Meme Theory

- Dawkins coined the term “meme” in 1976
- An idea like “I want a tattoo”
- Wikipedia definition -

“a unit of cultural information that propagates from one mind to another as a theoretical unit of cultural evolution”

Ideosphere

- The sum of all memes in circulation
- The “Global Human Consciousness”
- The Internet has a subset of the Ideosphere

Empirical Meme Theory (Me!)

- Original Meme Theory
- Keywords As Proxy For A Meme
- Electron Flow (E = I x R) (Networks)
- S-Curve
- Meme Miner (Dejanews.com)
- Google Trends tool
- Blogpulse.com

Example: Delphi

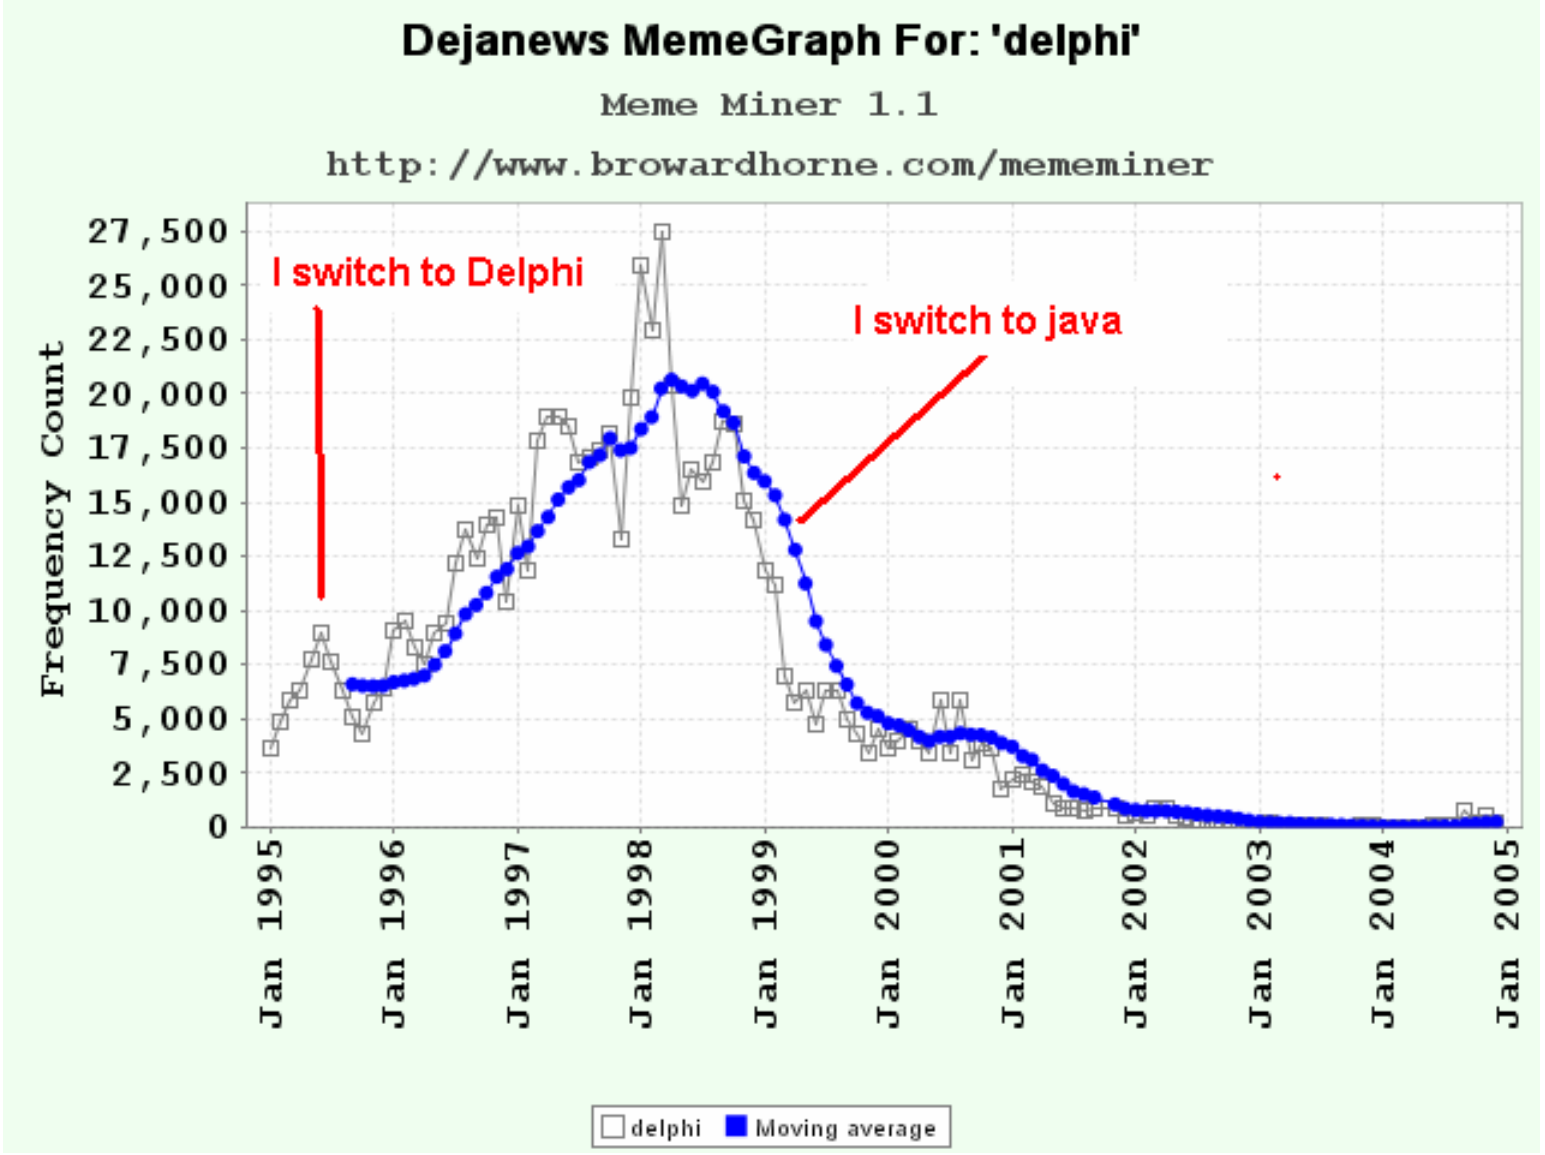

Example: Easter Bunny

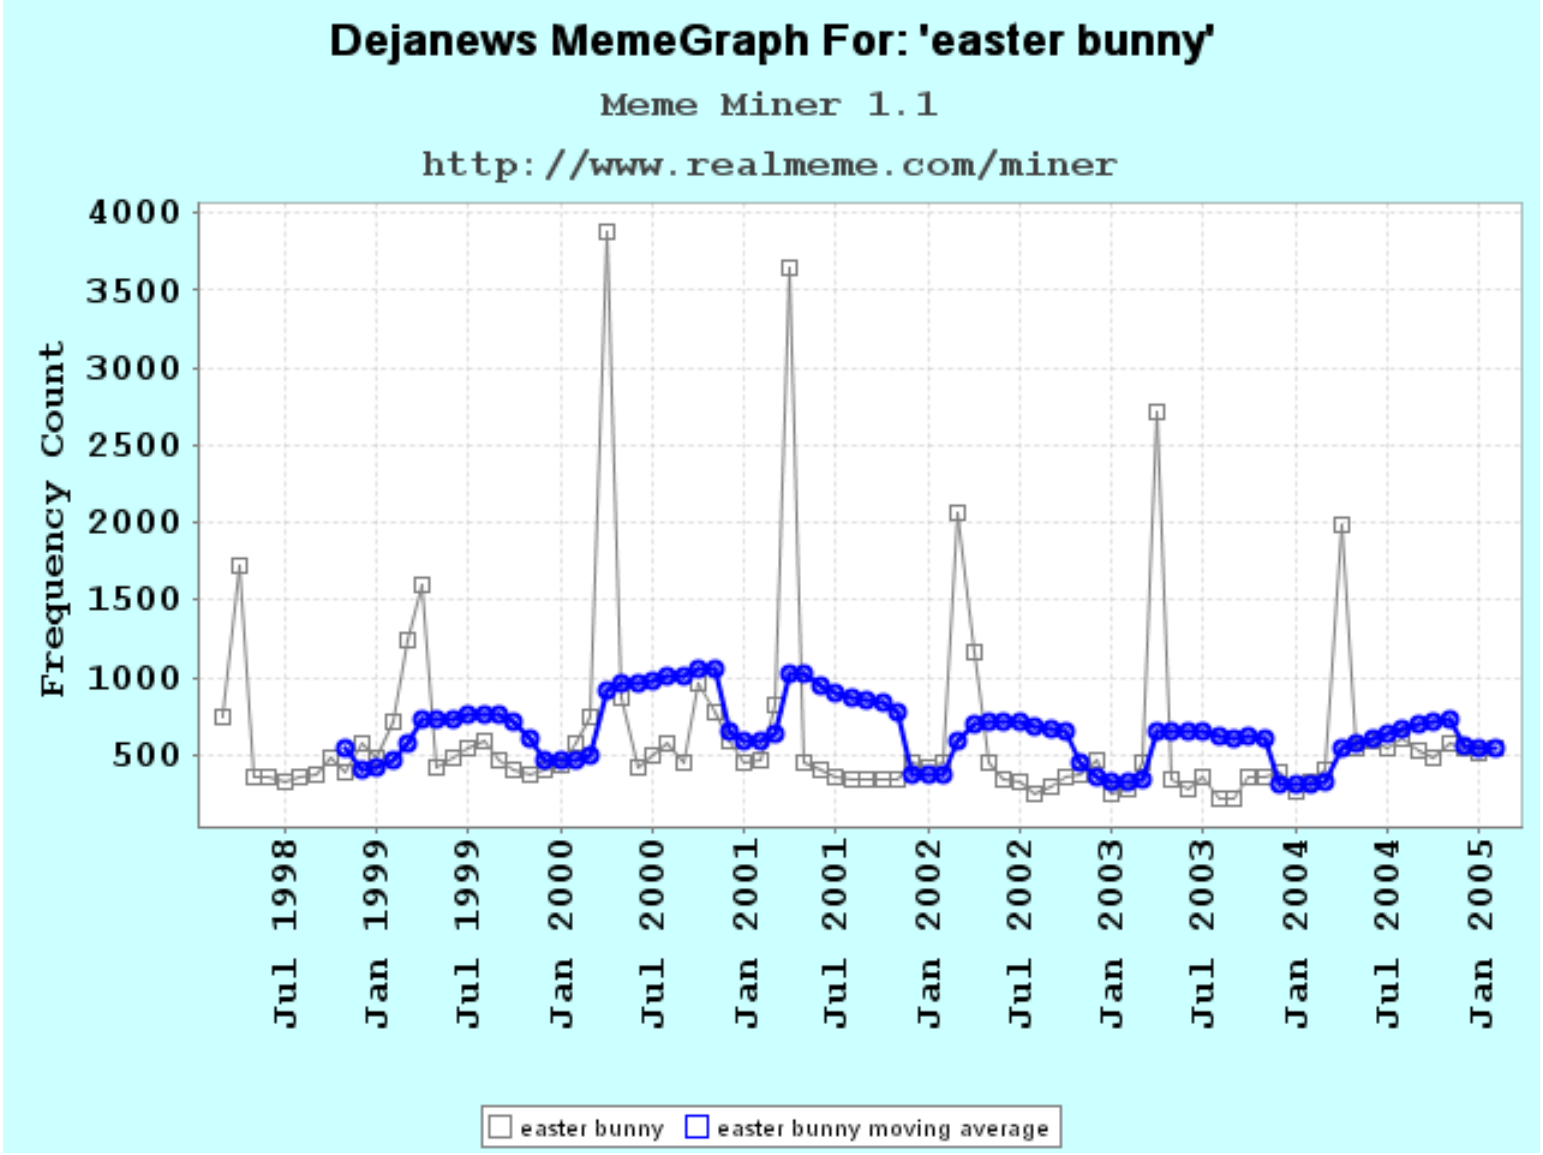

Example: Sex & Terrorism

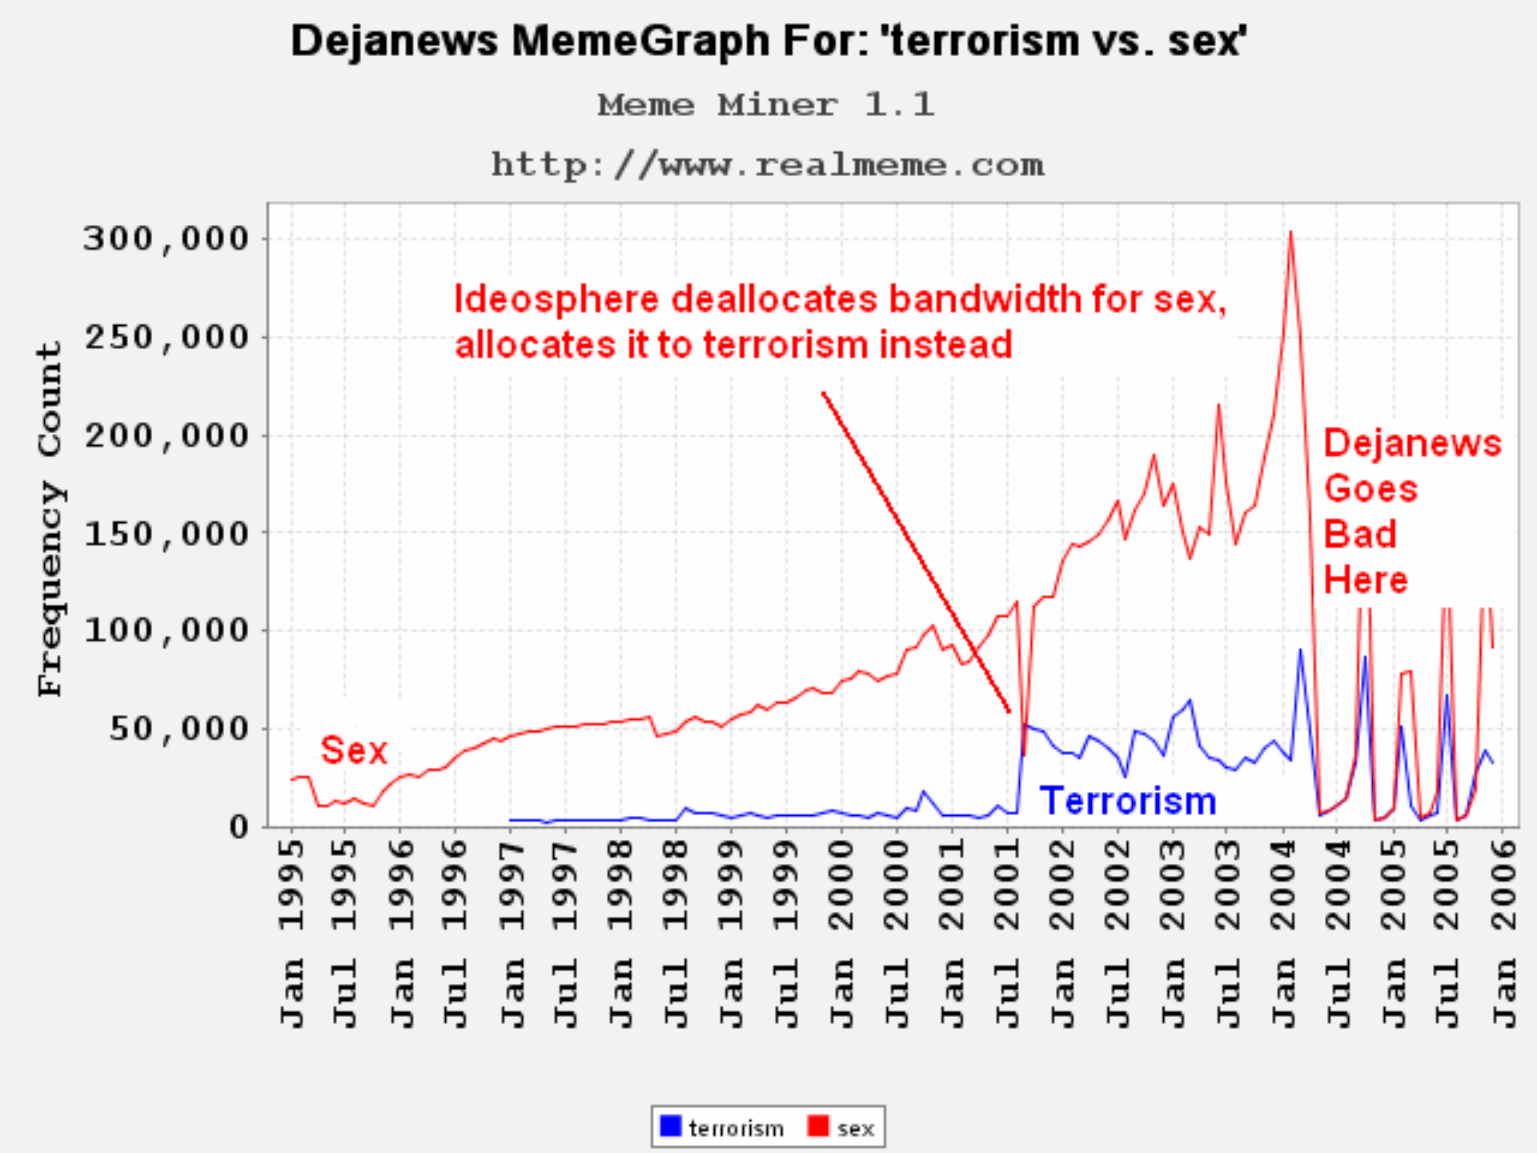

Meme Assumptions

- **Memes propagate as an S-Curve**
- **Memes propagate to most sites but at different amplitudes and latencies**

Meme Miner Inadequacies

- Single source
- Dejanews.com indexing revised by Google
- Older technology losing favor

Expanded Mining Model

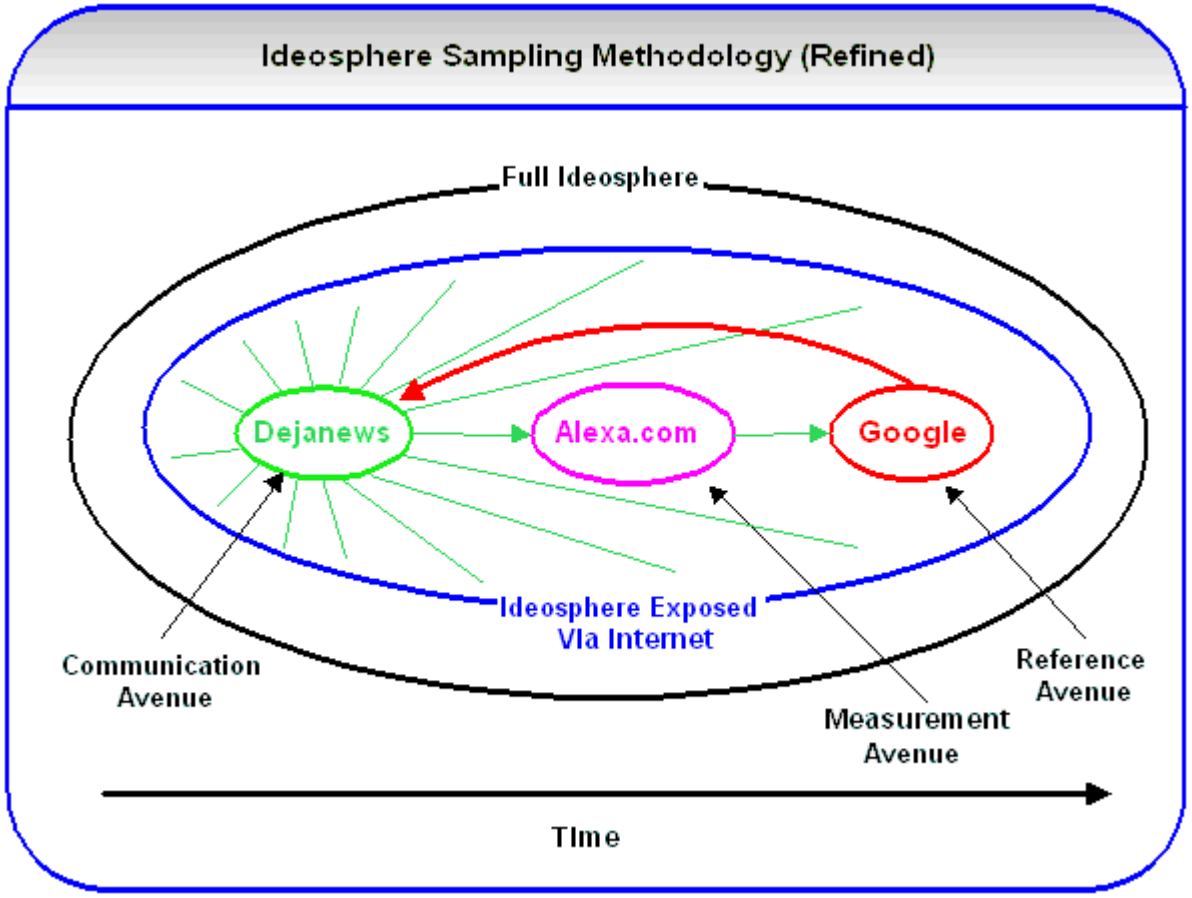

First tested on “MySpace” meme…

MySpace Meme (Dejanews)

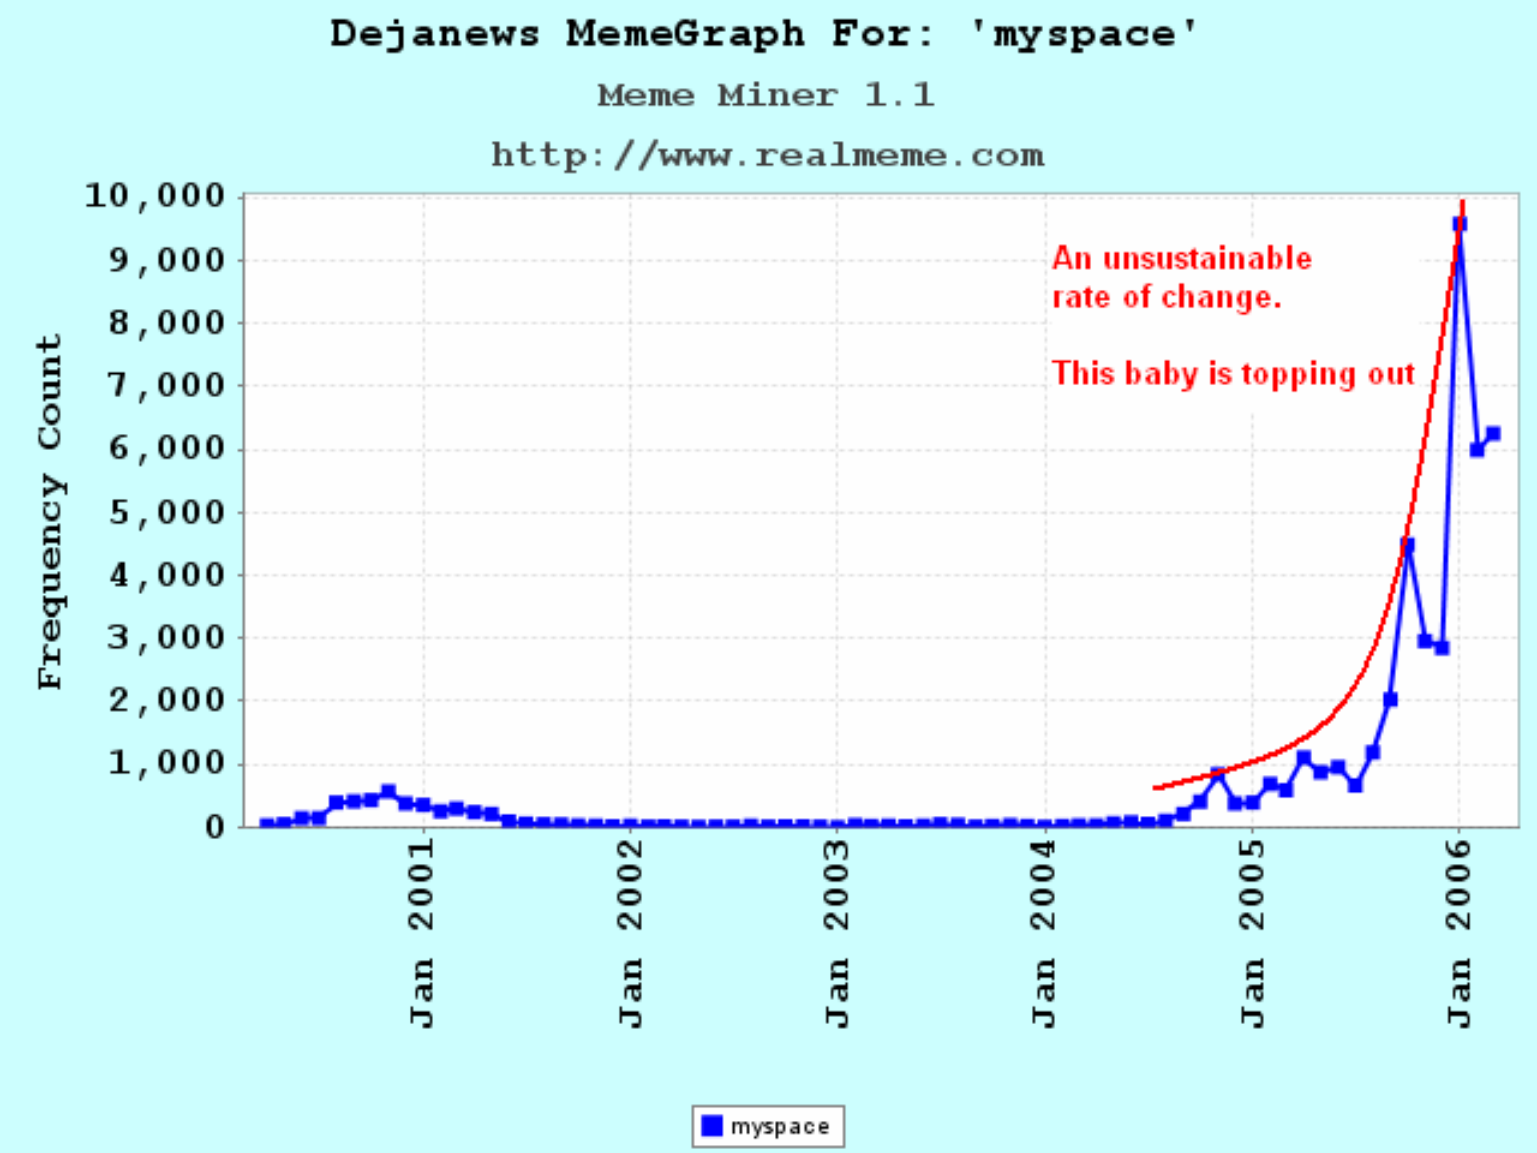

MySpace Meme (Alexa)

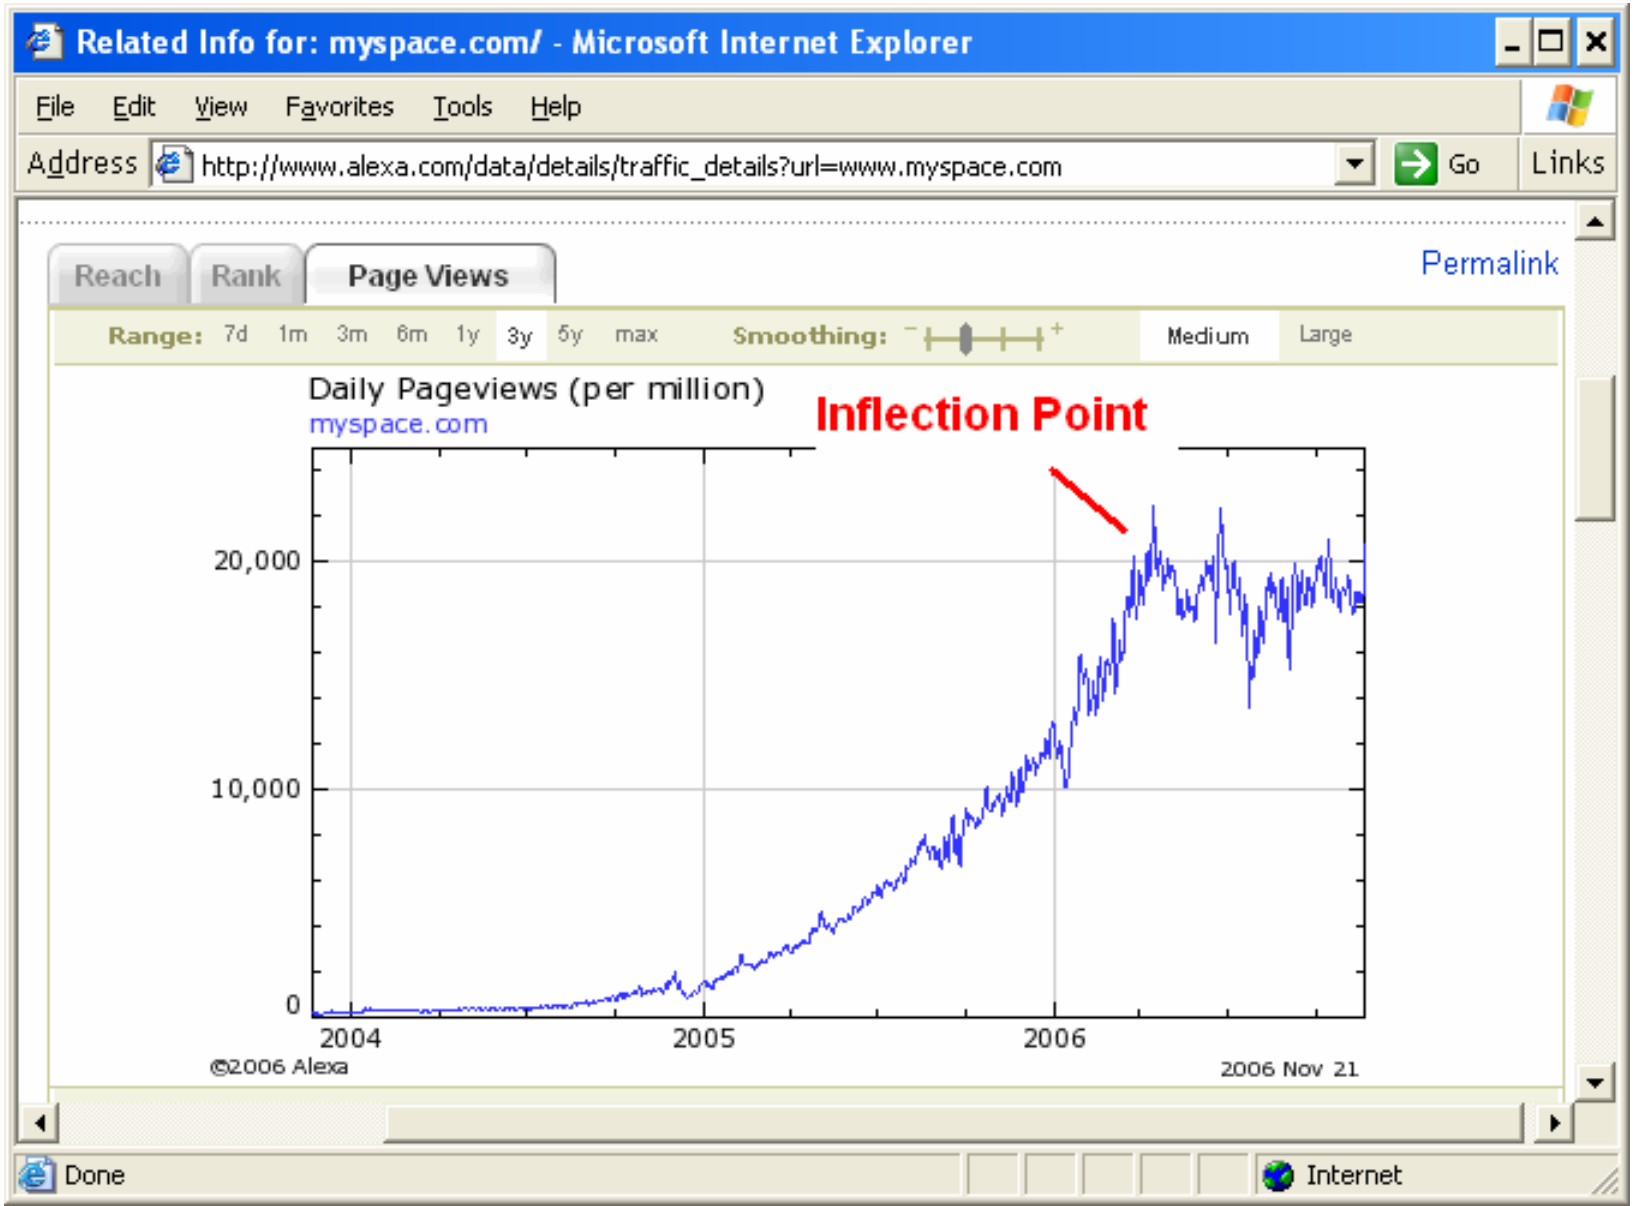

MySpace Meme (Google)

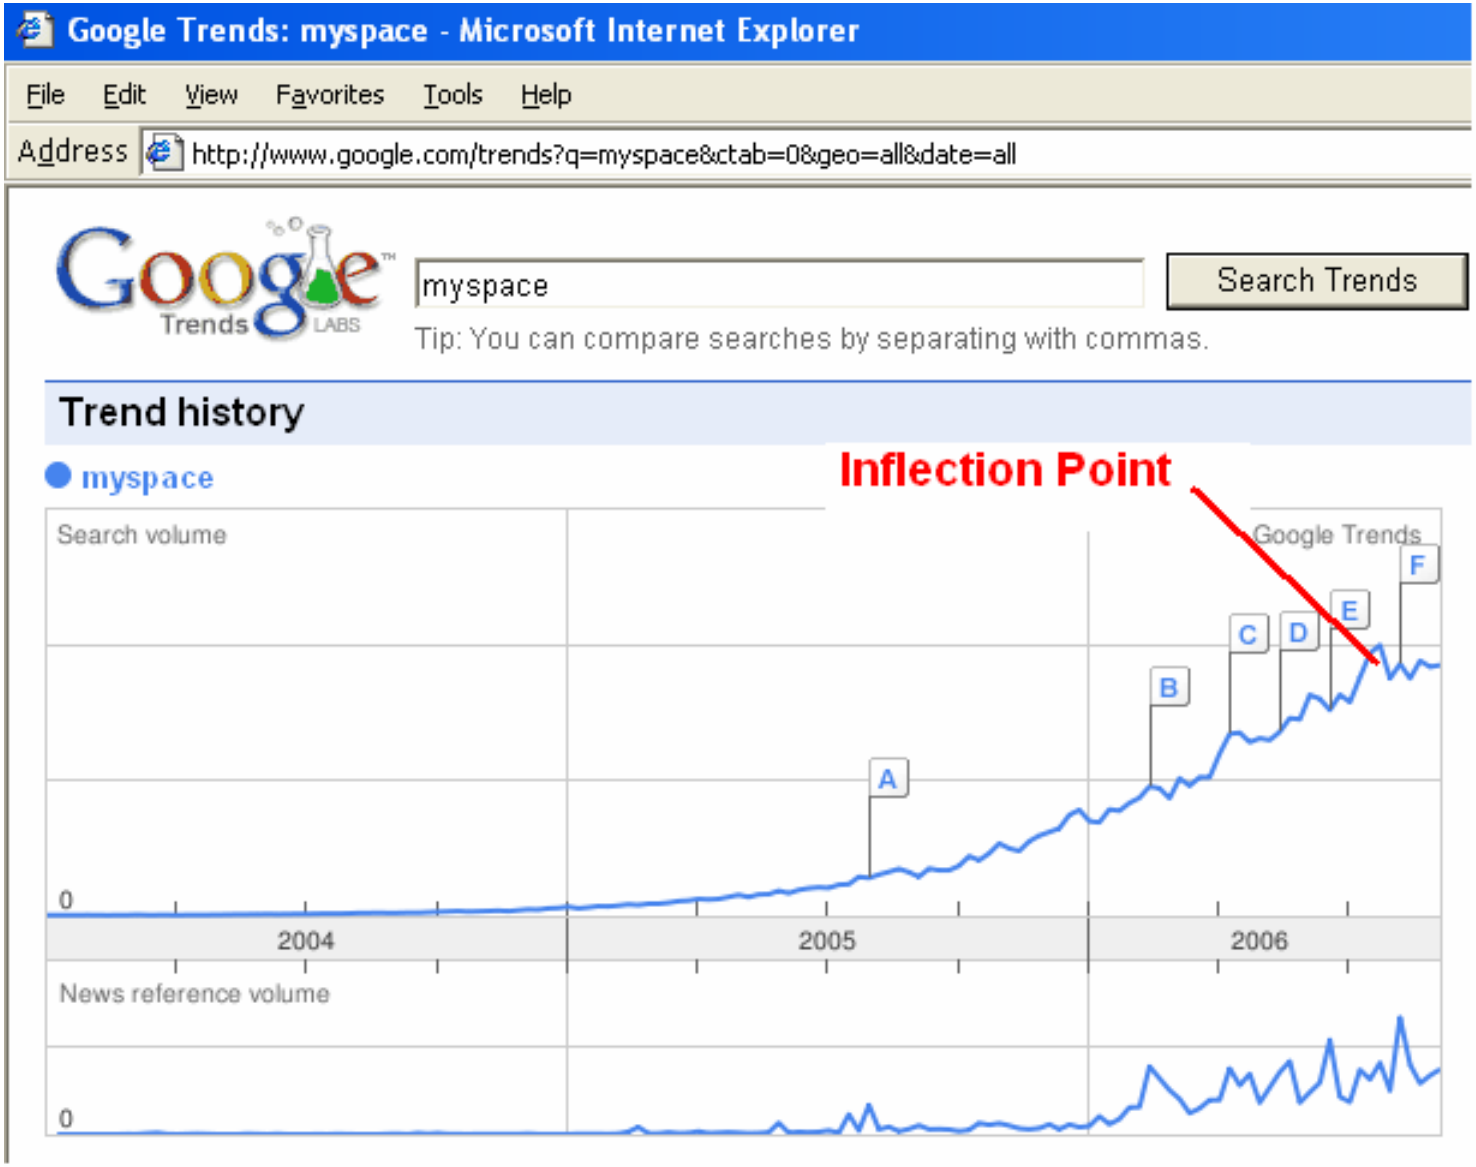

The New Model Worked…

For the “MySpace” meme, so I tested it against another case, a new site with a high growth rate…

“IAmFacingForeclosure.com” (IAFF.com)

But the results were different…

IAFF Meme (Alexa)

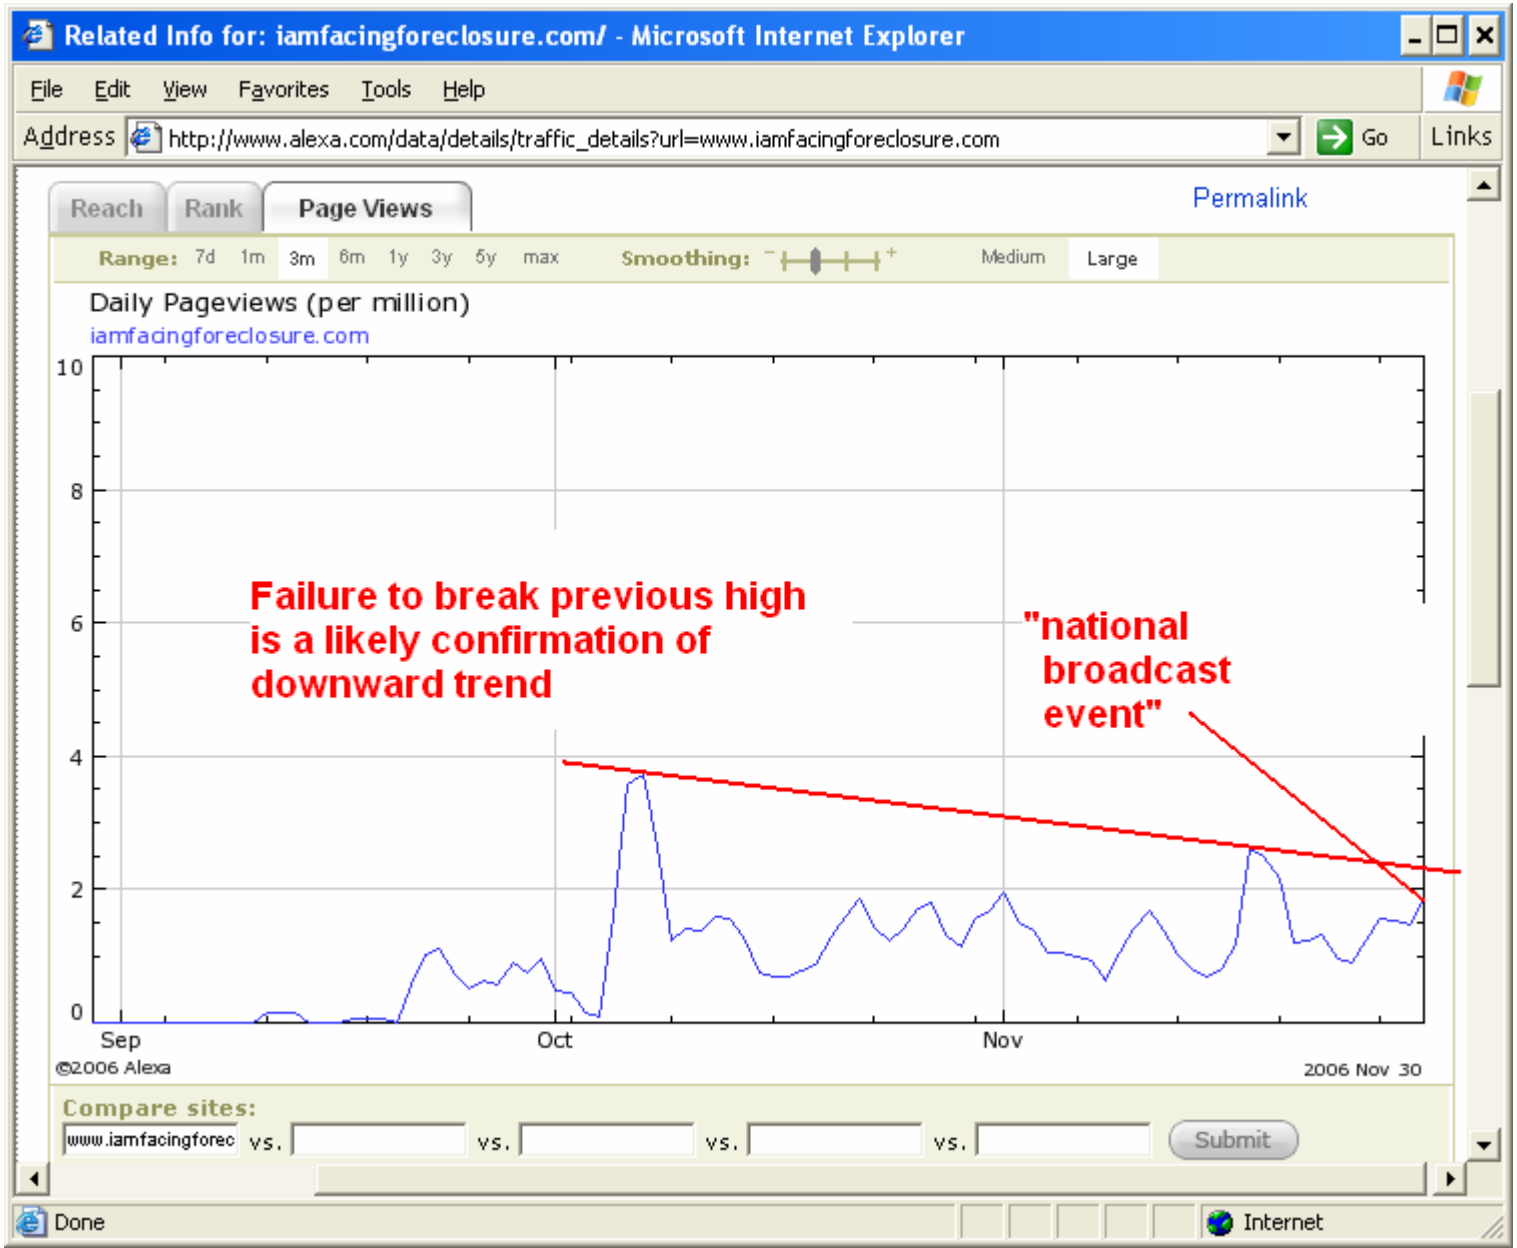

IAFF Meme (Dejanews)

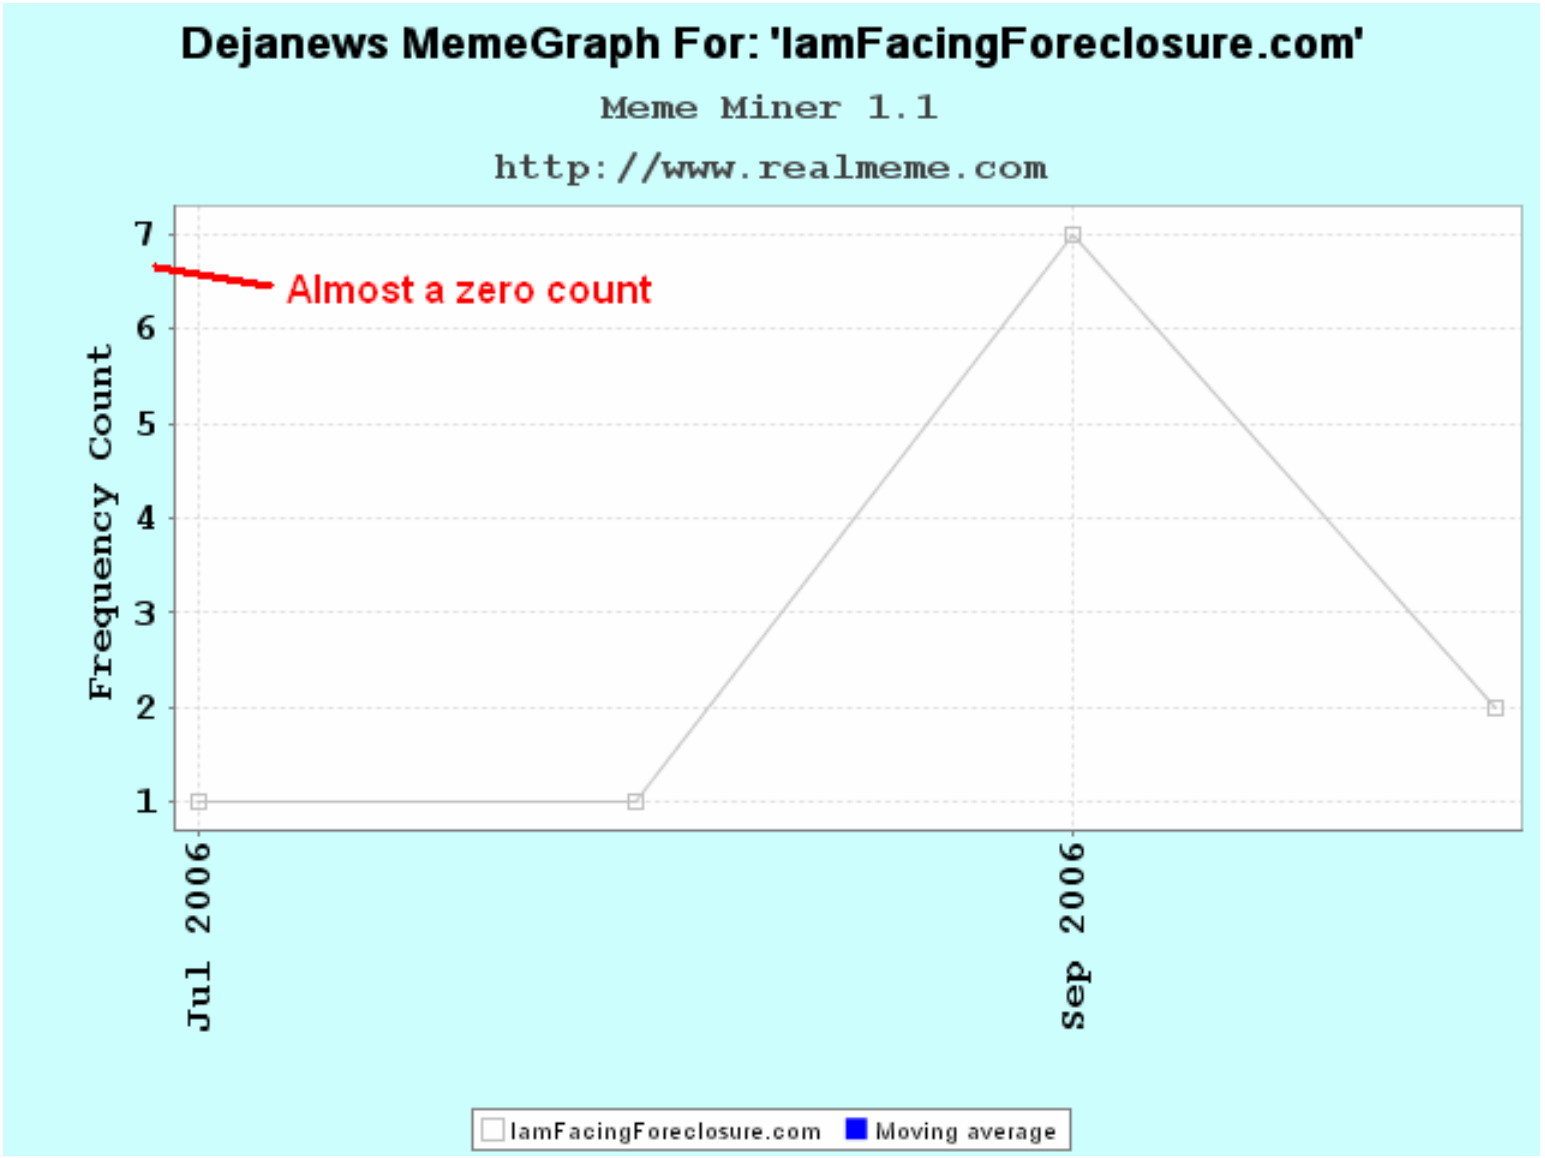

IAFF Meme (Google)

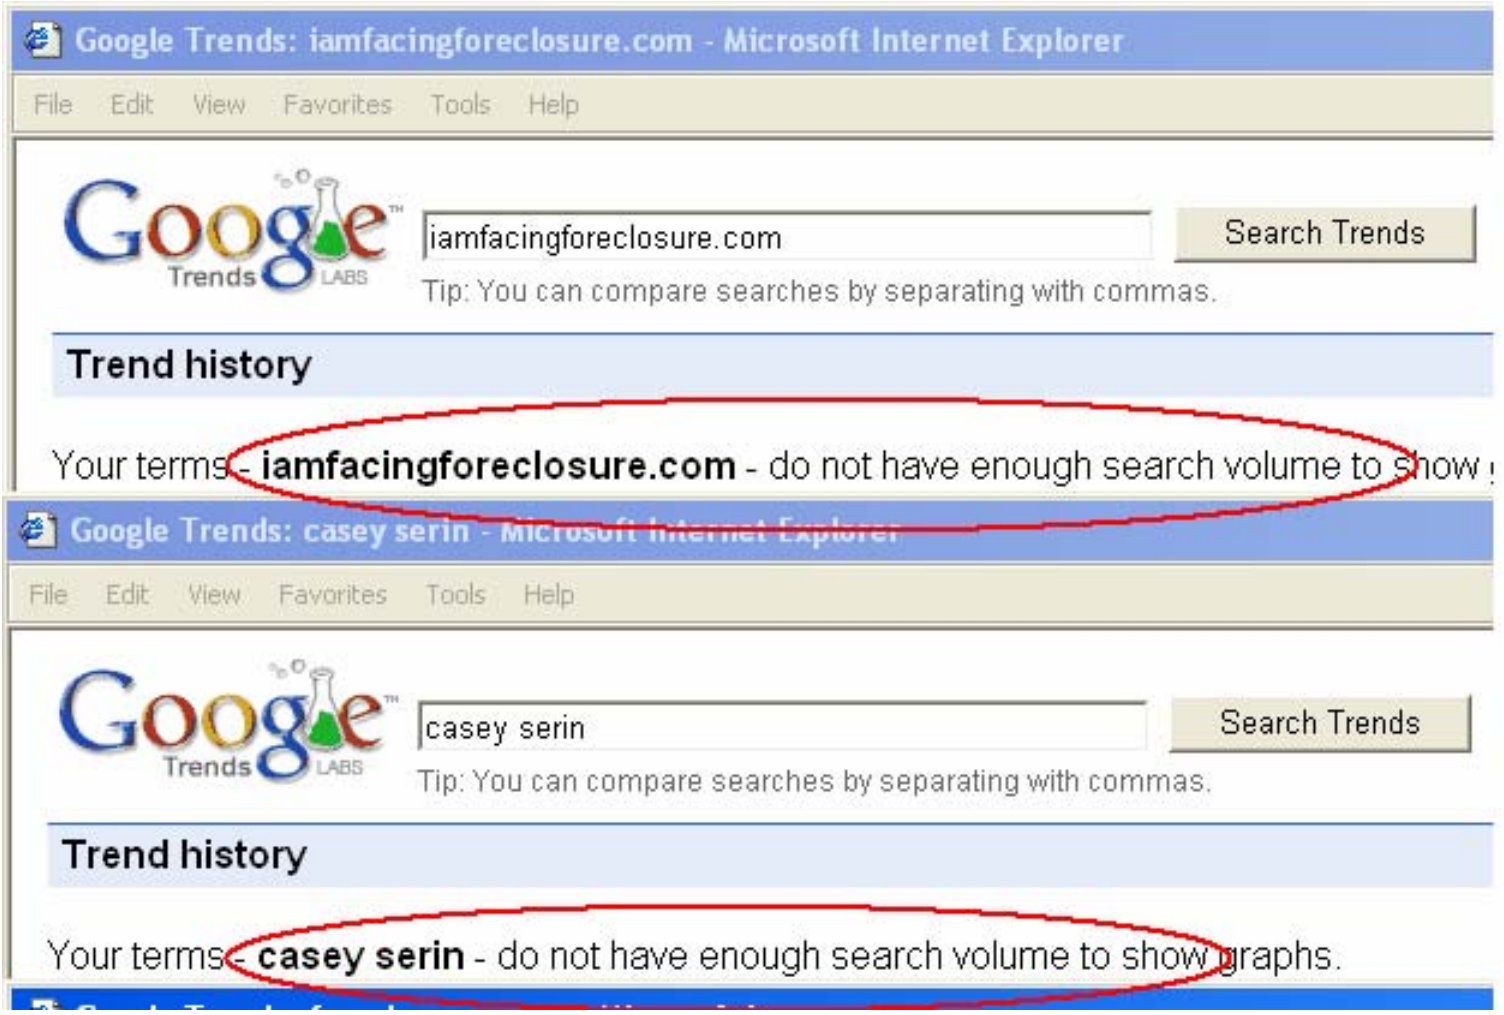

Meme Seepage Theory

Remember our Meme assumptions? Memes propagate as an S-curve and across most sites but with varying amplitudes and latencies.

If traffic increases to a primary site, then traffic to linked sites should increase proportionally (more or less)

If traffic increases to a primary site, then traffic to reference sites like Google should increase proportionally… and other sites in the top 10 result list should experience a lesser but measurable increase in traffic…

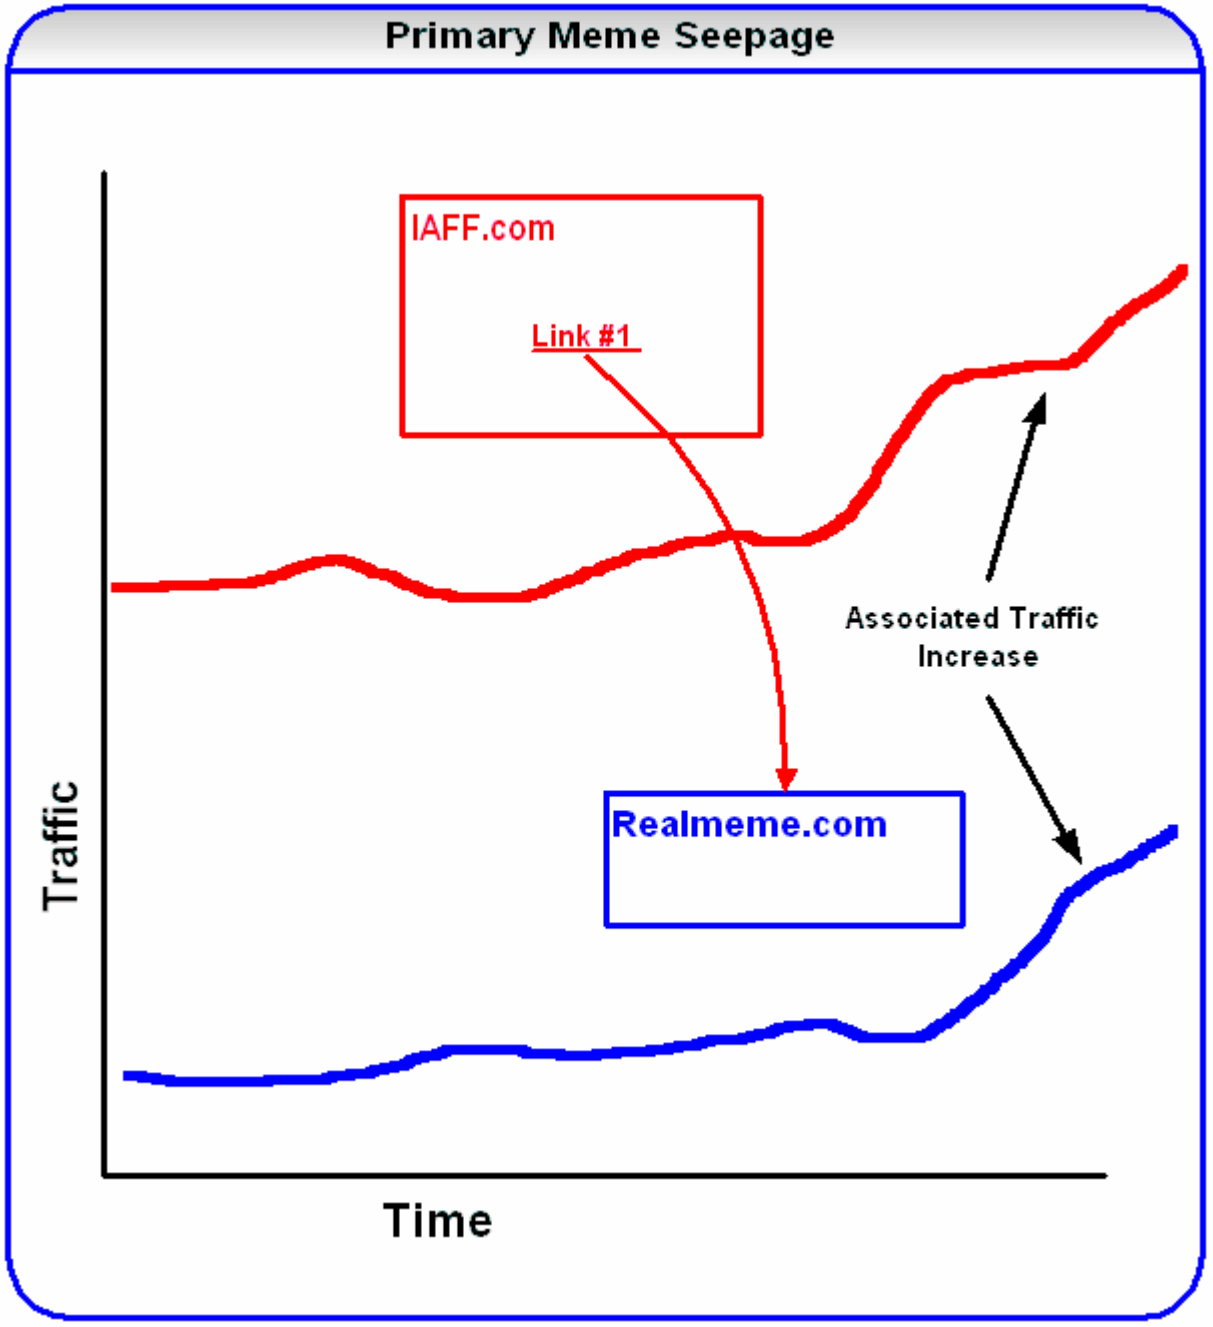
Primary Meme Seepage
IAFF.com
Link #1
Associated Traffic Increase
Realmeme.com
Traffic
Time

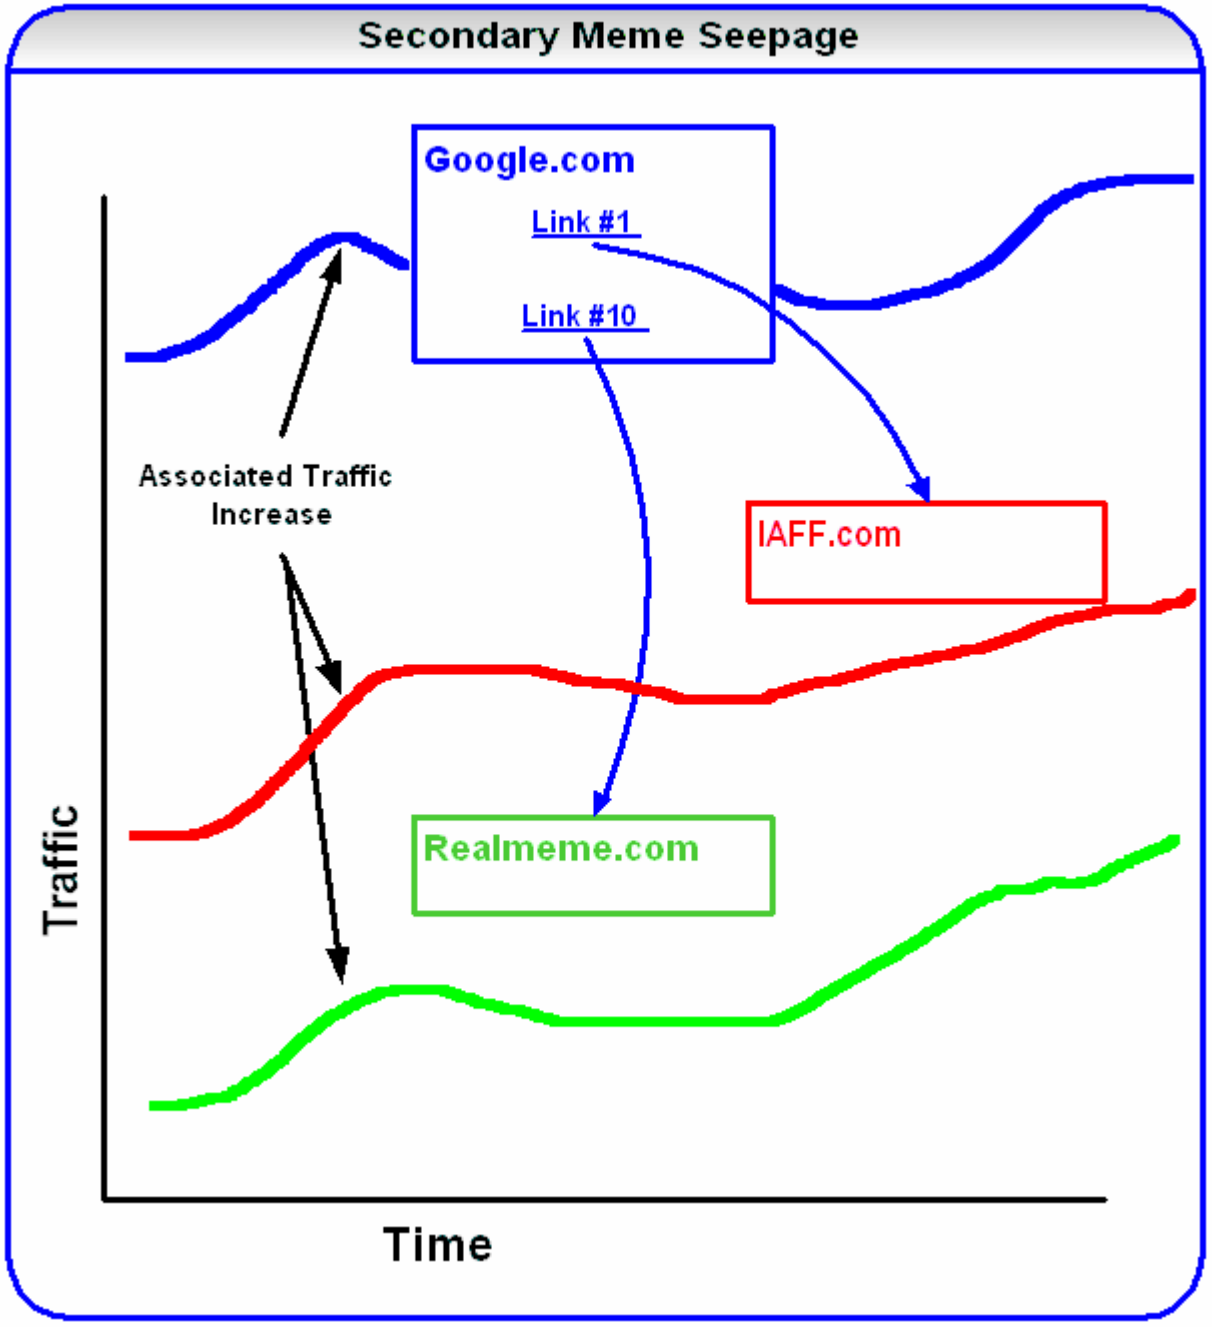
Secondary Meme Seepage
Google.com
Link #1
Link #10
Associated Traffic Increase
IAFF.com
Realmeme.com
Traffic
Time

So I Experimented..

I posted a link directly to IAFF.com, to tap off a slice of IAFF.com's traffic via meme seepage. Theoretically, a doubling of IAFF traffic should produce an equal ratio of redirected traffic to my own site, RealMeme.com

But the results were wrong again…

And I Experimented Again…

Alexa showed IAFF.com with a 25% increase in traffic but my site experienced no concurrent increase.

I was surprised and I posted the results to IAFF.com.

And a few days later, I did experience an anomalous increase in traffic but it didn't match IAFF.com's Alexa traffic delta. Here's what hit my site…

Botnet Proof

The following page hits are logs from my website.
It’s clear that they were artifically produced…

- Too many simultaneous operating systems per IP.
- The traffic is too dense and changeover too abrupt.
- Too many 2-page hits (my traffic is 90% 1-page)
- The page hits don't follow a human click flow.
- All hits are bookmarks, no blog entry points
- Most bookmarks are older

Botnet Logs (Same IP)

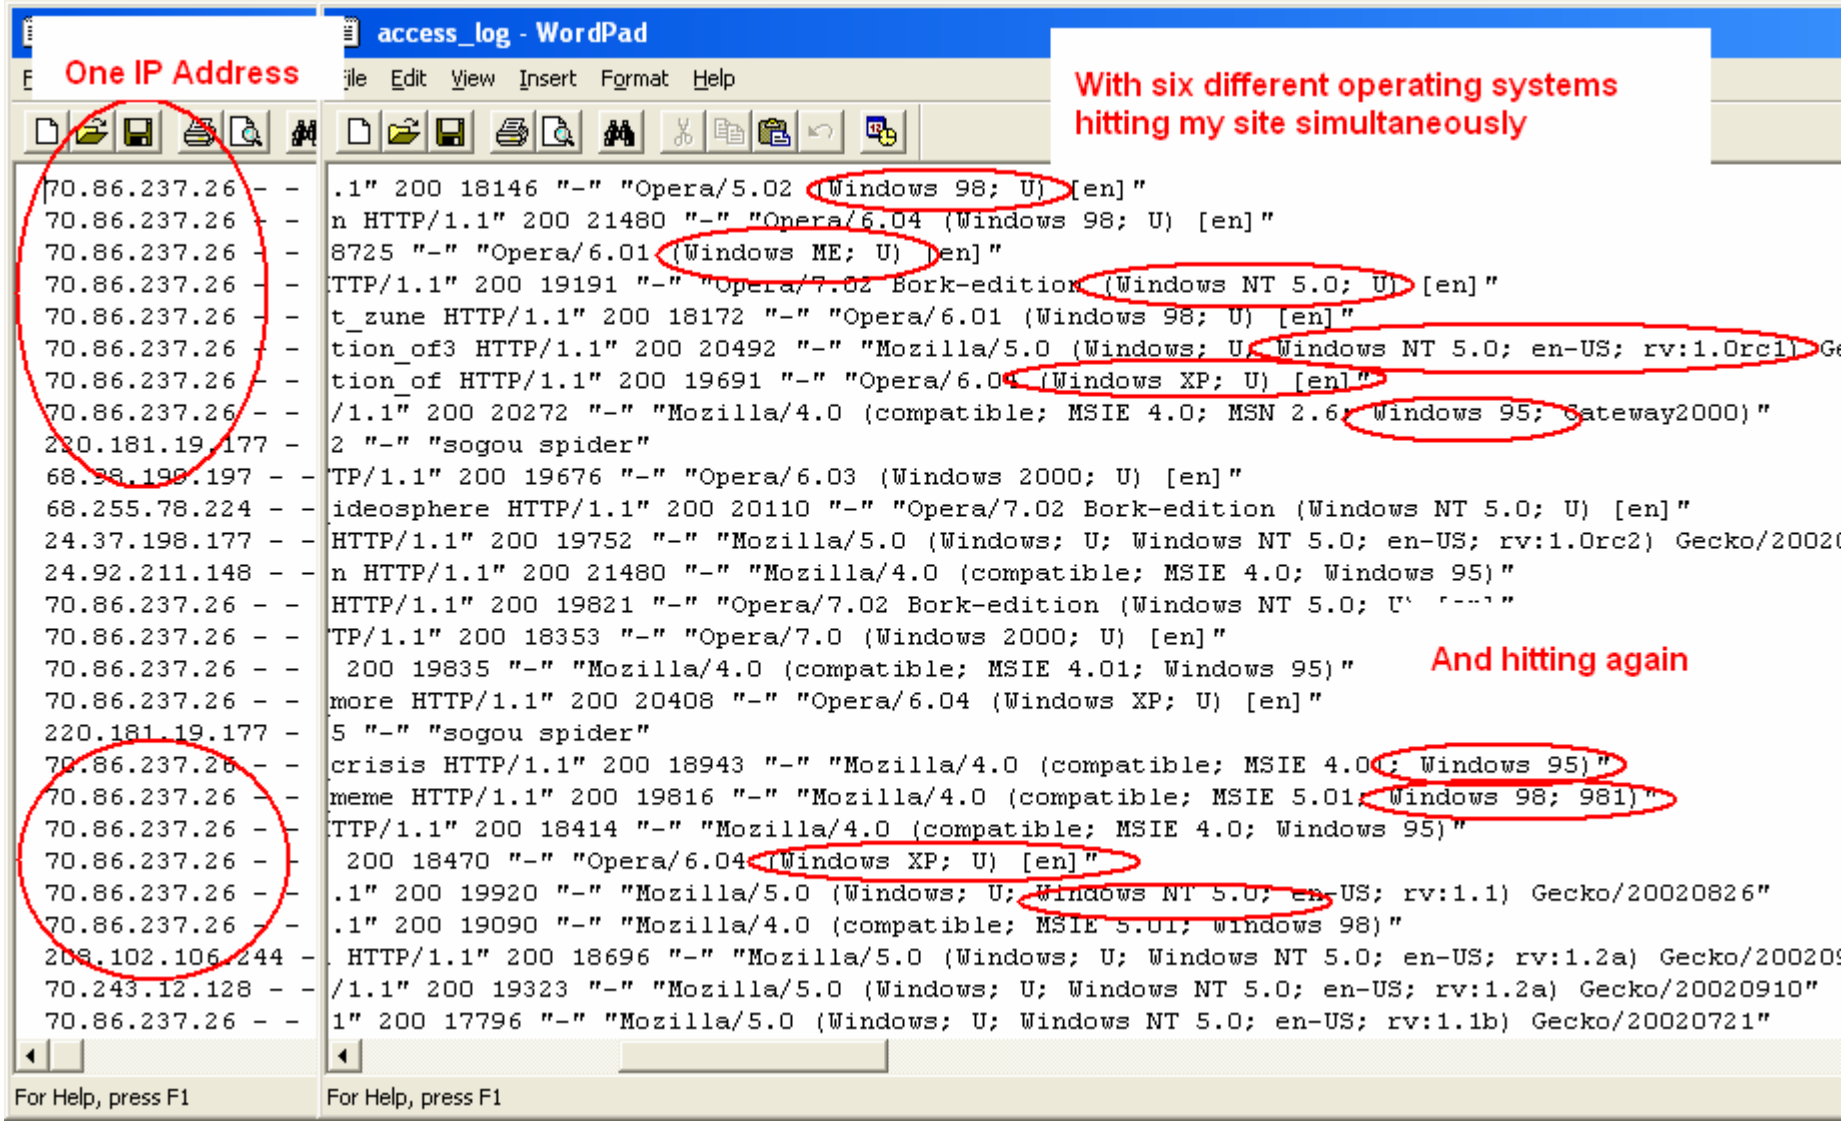

Botnet Logs (Density)

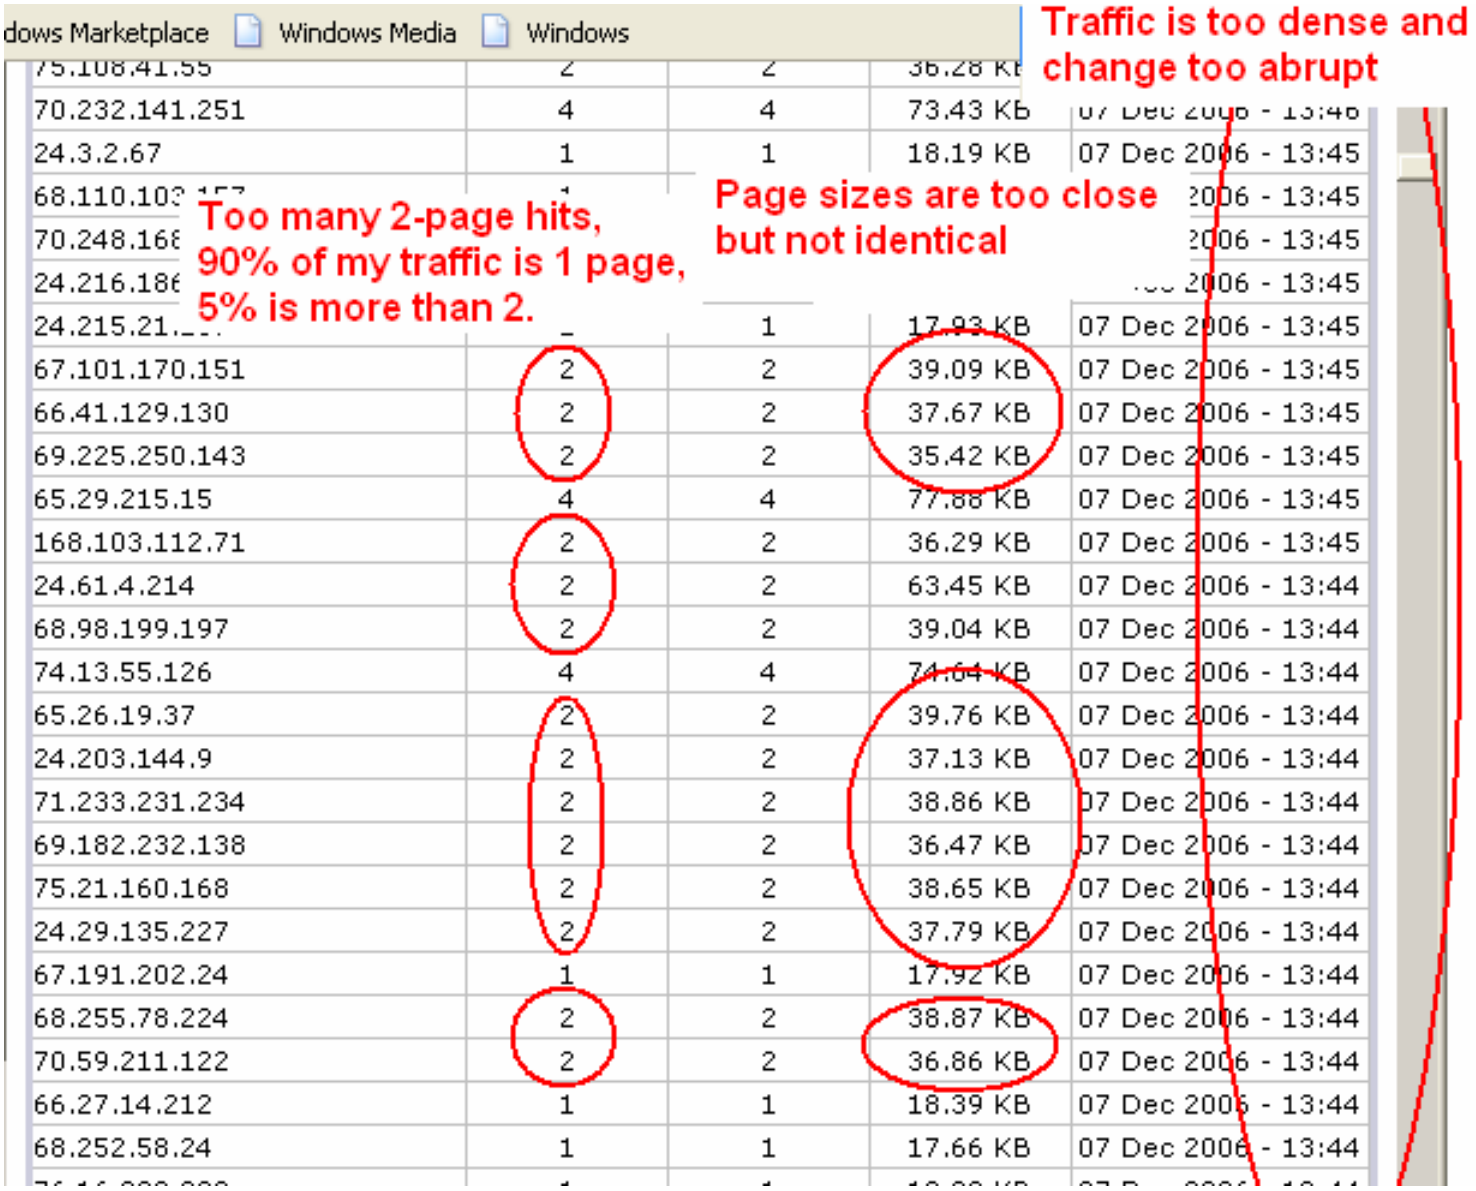

70.232.141.251	4	4	73.43 KB	07 Dec 2006 - 13:46
24.3.2.67	1	1	18.19 KB	07 Dec 2006 - 13:45
67.101.170.151	2	2	39.09 KB	07 Dec 2006 - 13:45
66.41.129.130	2	2	37.67 KB	07 Dec 2006 - 13:45
69.225.250.143	2	2	35.42 KB	07 Dec 2006 - 13:45
65.29.215.15	4	4	77.88 KB	07 Dec 2006 - 13:45
168.103.112.71	2	2	36.29 KB	07 Dec 2006 - 13:45
24.61.4.214	2	2	63.45 KB	07 Dec 2006 - 13:44
68.98.199.197	2	2	39.04 KB	07 Dec 2006 - 13:44
74.13.55.126	4	4	74.64 KB	07 Dec 2006 - 13:44
65.26.19.37	2	2	39.76 KB	07 Dec 2006 - 13:44
24.203.144.9	2	2	37.13 KB	07 Dec 2006 - 13:44
71.233.231.234	2	2	38.86 KB	07 Dec 2006 - 13:44
69.182.232.138	2	2	36.47 KB	07 Dec 2006 - 13:44
75.21.160.168	2	2	38.65 KB	07 Dec 2006 - 13:44
24.29.135.227	2	2	37.79 KB	07 Dec 2006 - 13:44
67.191.202.24	1	1	17.92 KB	07 Dec 2006 - 13:44
68.255.78.224	2	2	38.87 KB	07 Dec 2006 - 13:44
70.59.211.122	2	2	36.86 KB	07 Dec 2006 - 13:44
66.27.14.212	1	1	18.39 KB	07 Dec 2006 - 13:44
68.252.58.24	1	1	17.66 KB	07 Dec 2006 - 13:44

Botnet Logs (Pages)

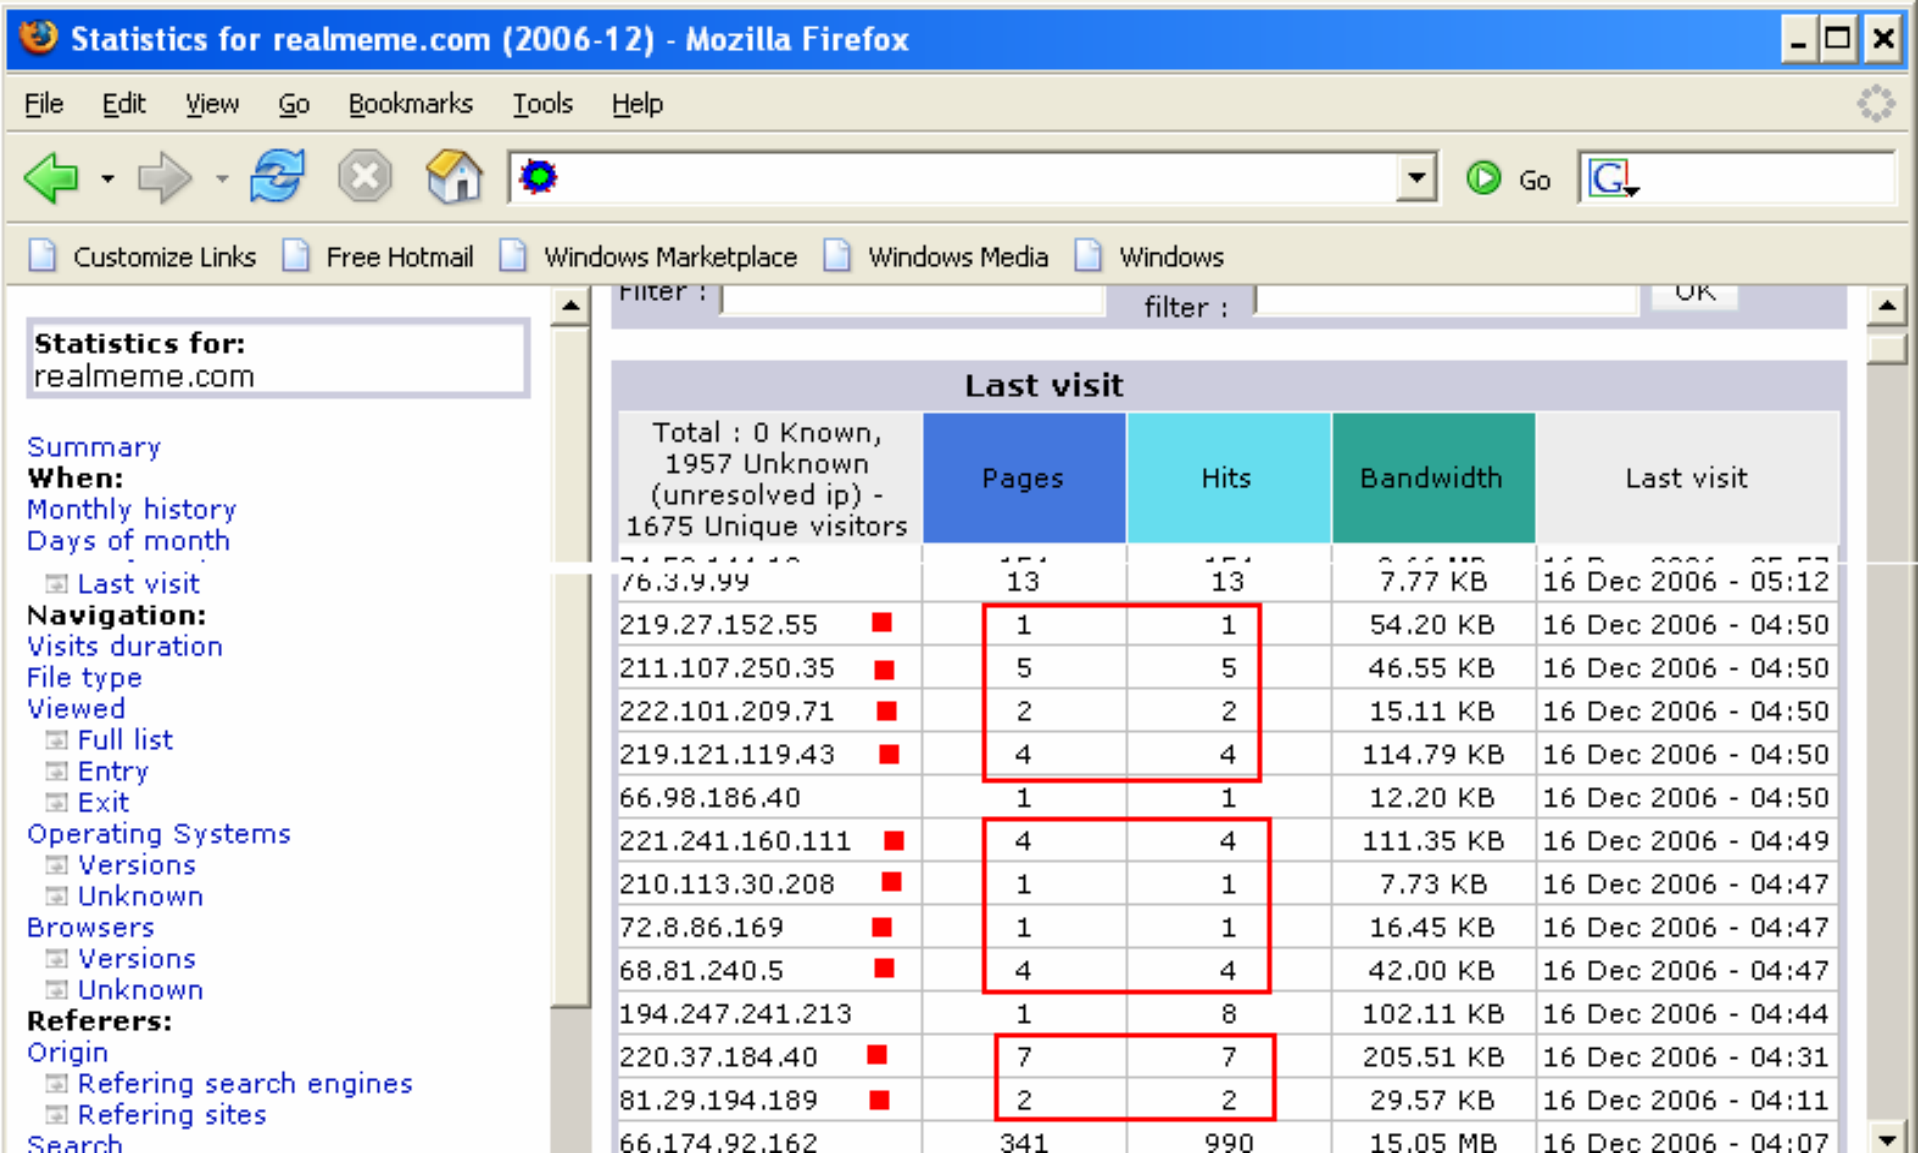

Botnet Logs (Entry Points)

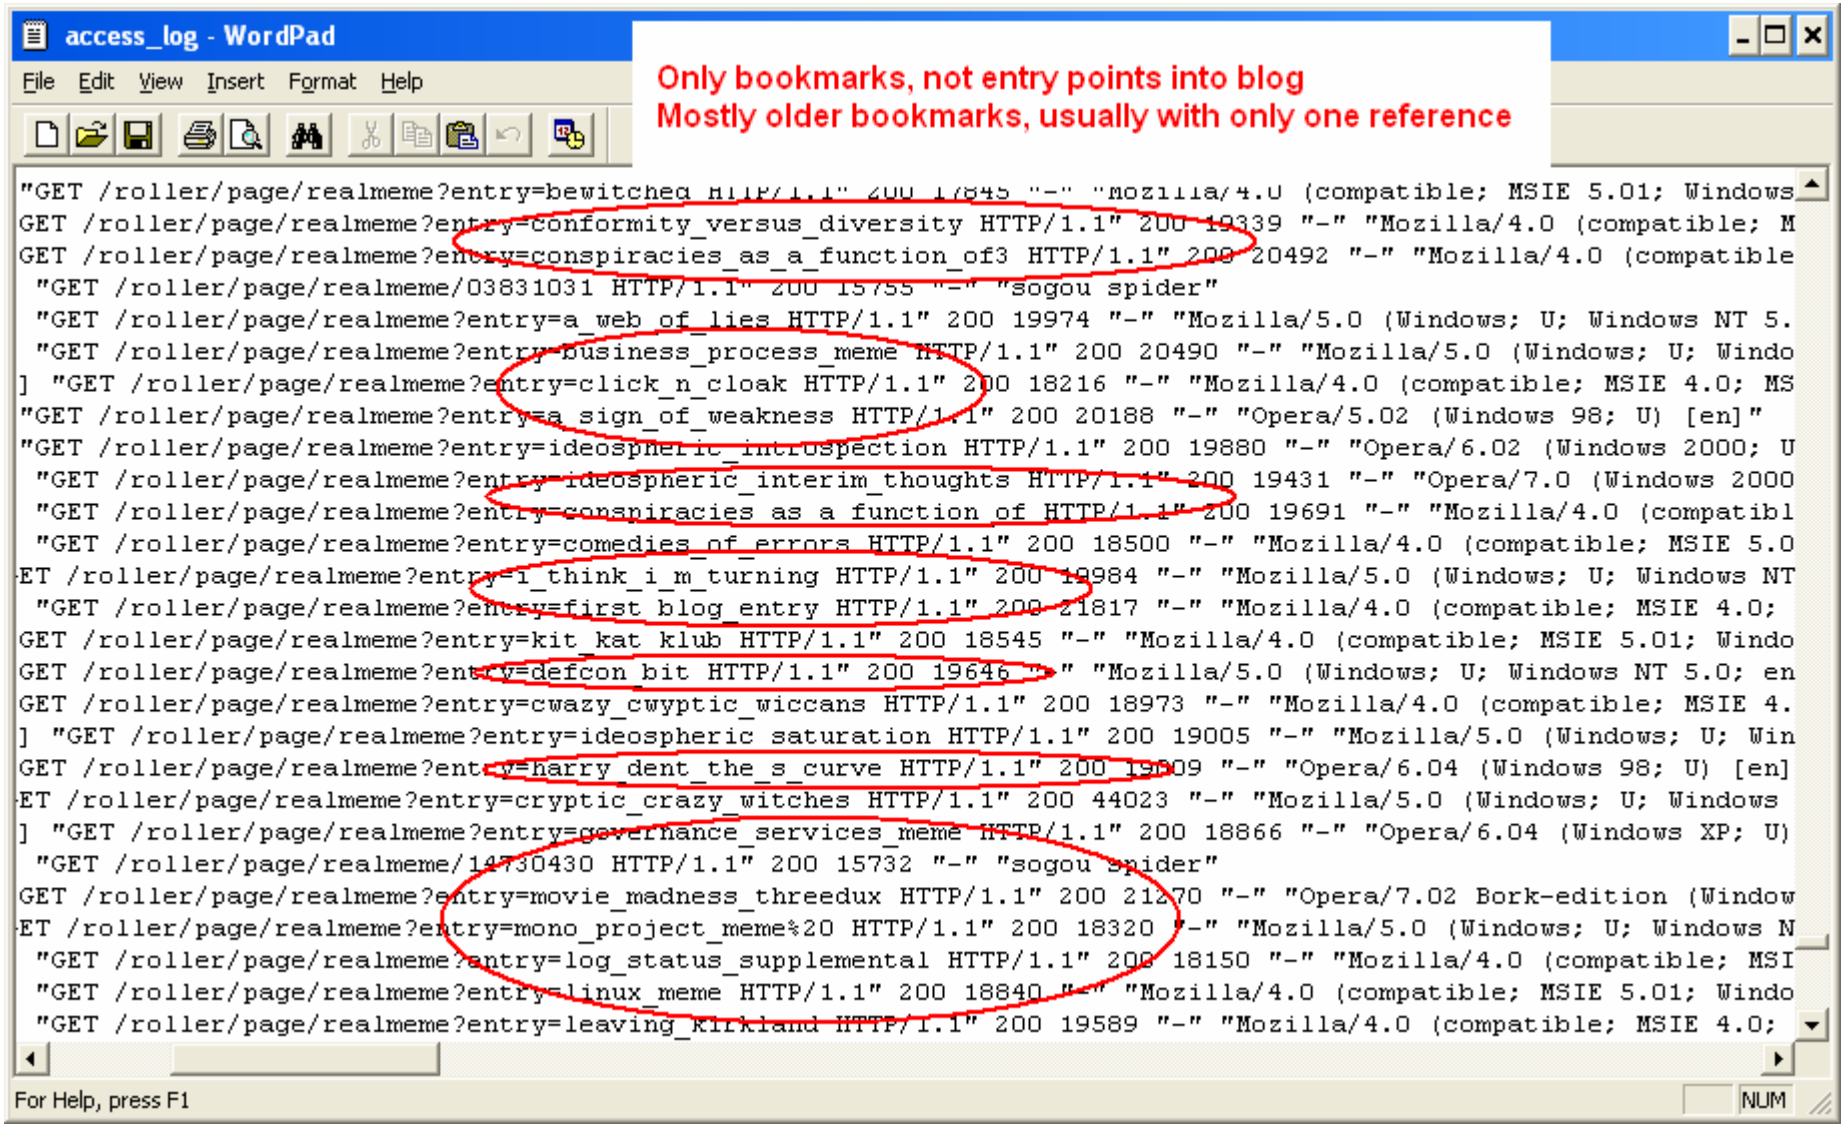

Botnet Epiphany

Okay, I’m not the smartest guy in the world but I eventually figured out that this new traffic was generated by bots.

But why?

So I tried another experiment…

Secondary Seepage Failure

I have a confession. My website was designed specifically for Google rankings and it's been surprisingly successful (my Defcon 16 presentation! Ho!) So I decided…

to induce a secondary seepage from IAFF.com to my site via Google…

Binding

I bound my site to IAFF.com by posting an IAFF.com analysis which was indexed by Google. At one point, I was the #7 Google result for

“IAmFacingForeclosure.com”

So now I’m getting a slice of traffic directly from IAFF.com AND from Google’s search results for IAFF.com

Binding Results

Once again, I saw major anomalies between IAFF.com's claimed traffic and the induced seepage to my site during "a major television event".

Can I prove fraud? No.

But I don't need to. I'm not Google or a Google advertiser.

My Theory

The botnets are mimicking meme seepage by generating traffic to linked secondary sites.

At first, I couldn’t figure out why. But as I worked out a methodology to expose botnet manipulation, I realized that LACK OF SEEPAGE is a major red flag. After all, that’s how I found these anomalies to begin with.

Gaming Google

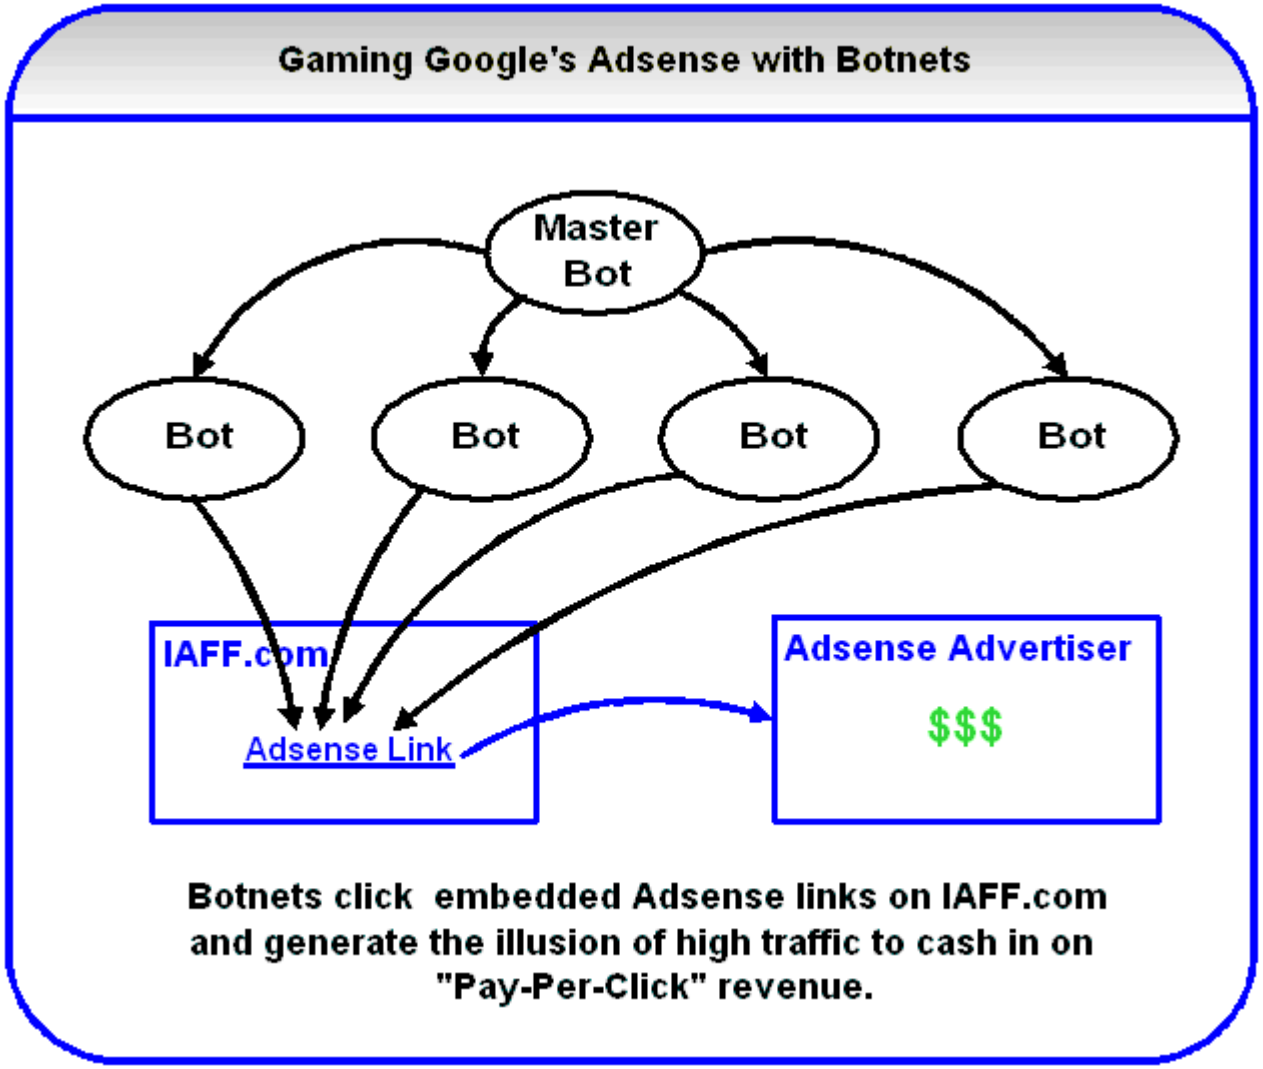

Gaming Alexa

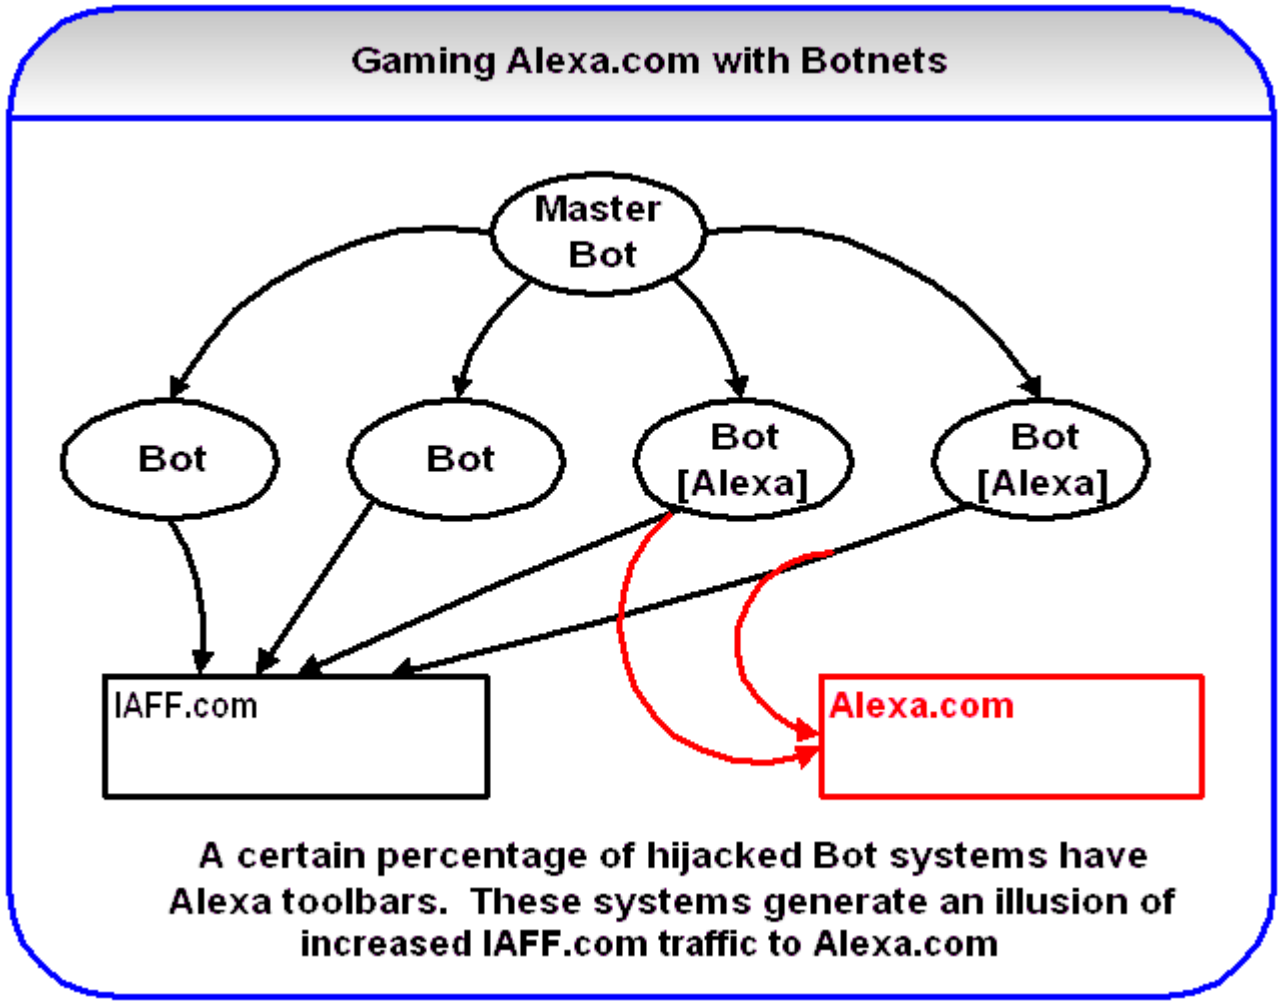

Gaming My Miner Model

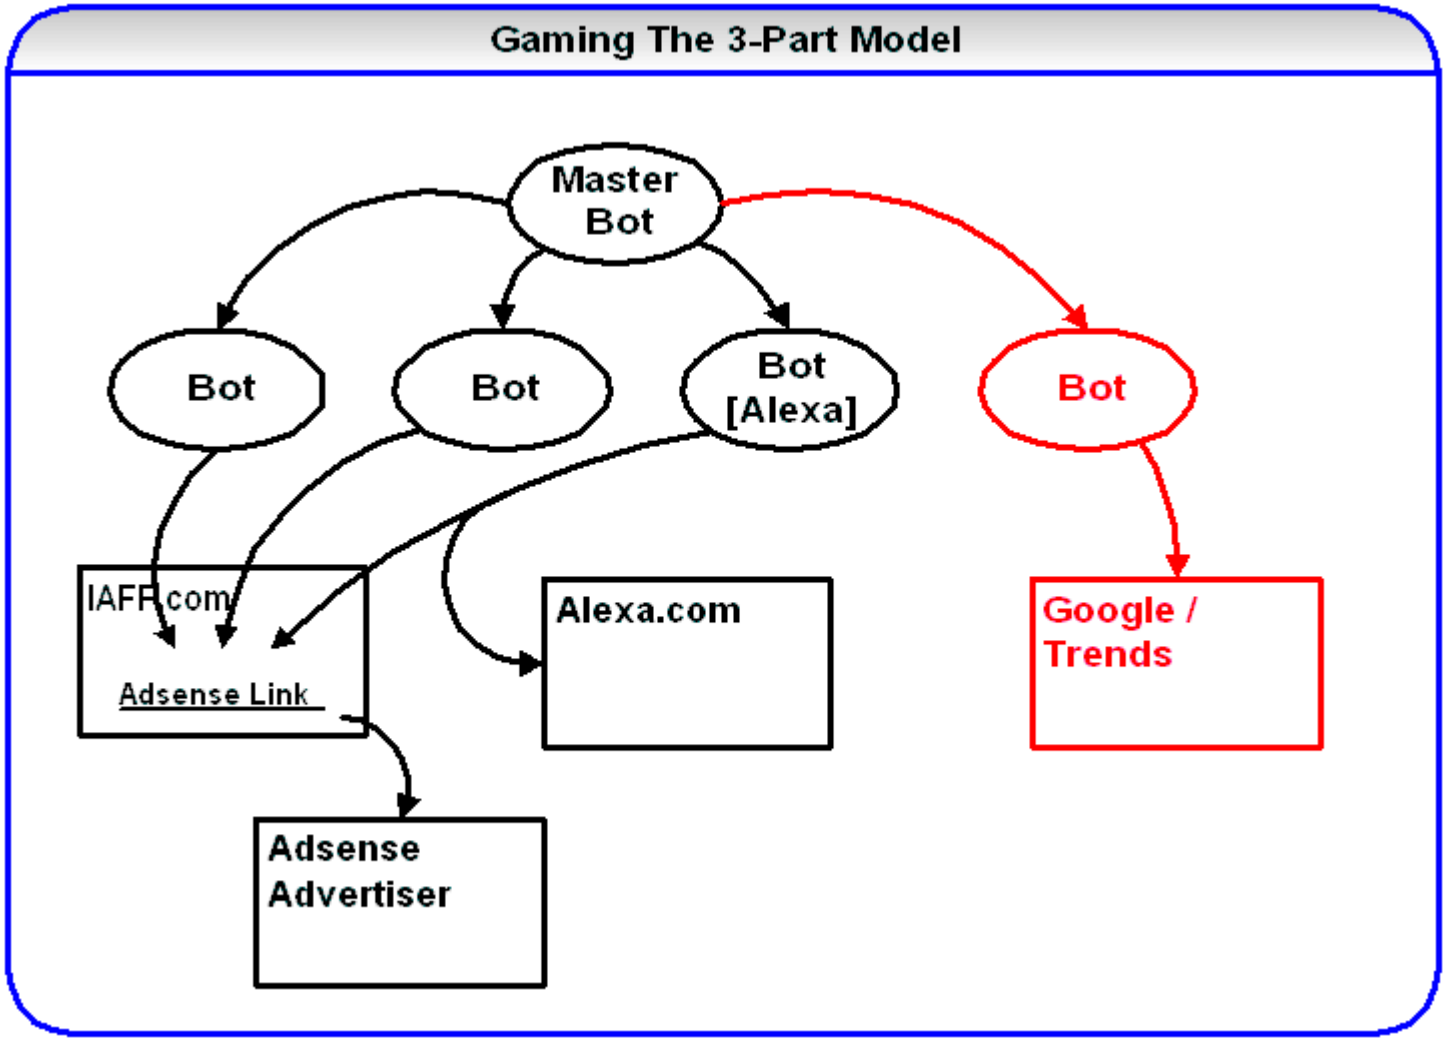

Meme Troubleshooting Table

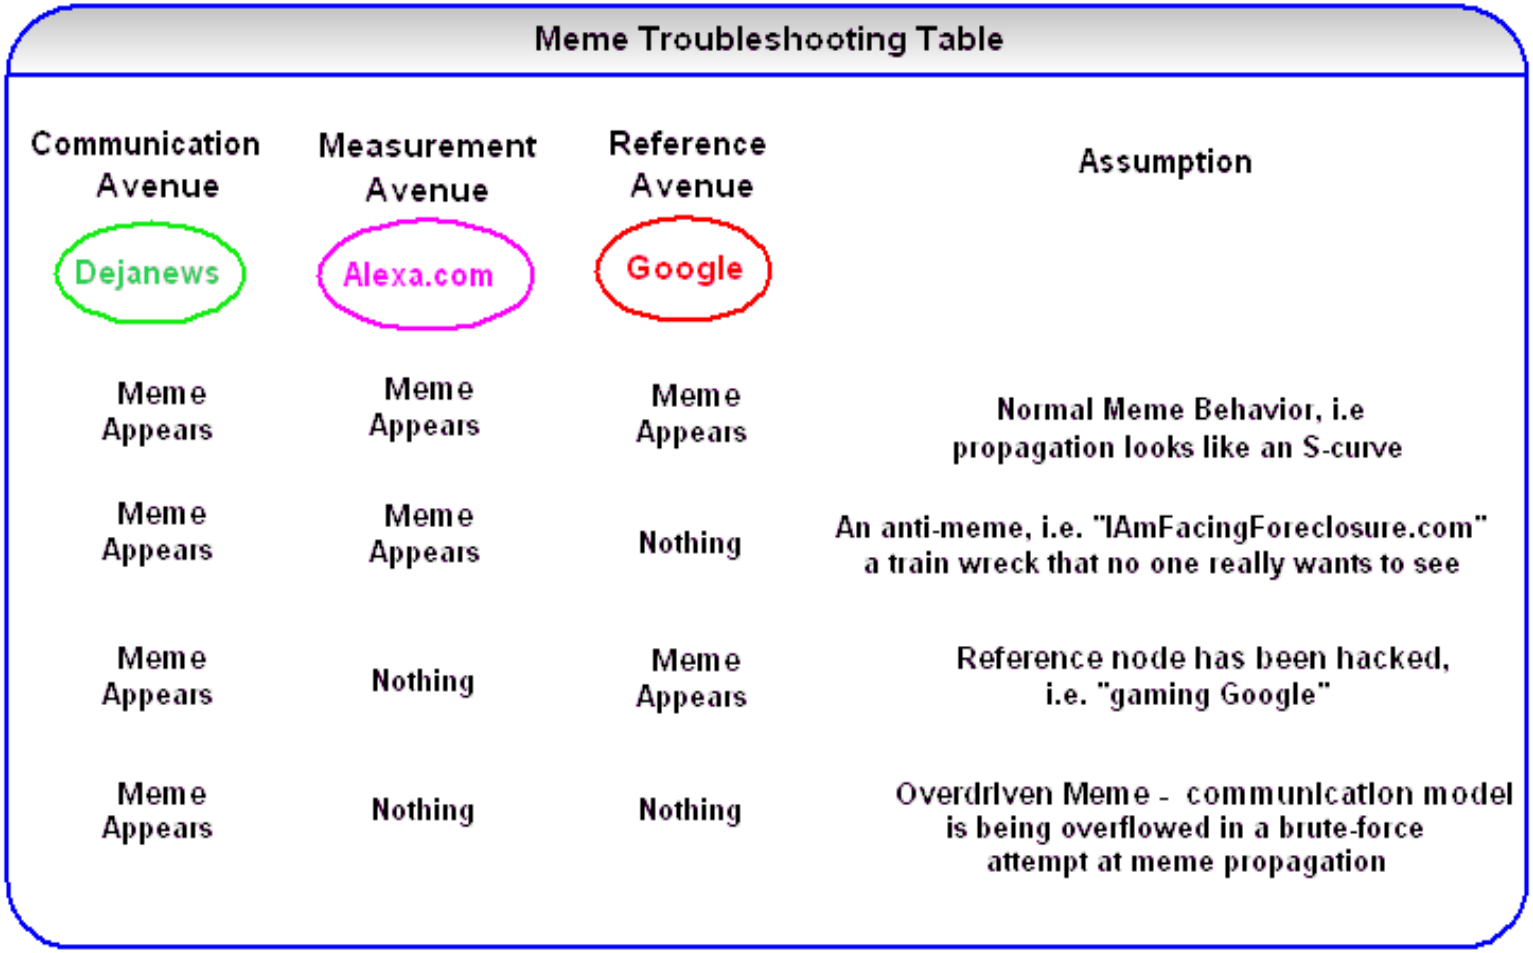

Humans versus Bots

Needed: a pervasive, immutable quality which is detectable in humans but which bots can never duplicate.

What is it?

Humans actually buy advertised products.

Conclusion

- Click fraud is more pervasive than reported
- It can be detected with memetic analysis…
- Botnets are a serious problem
- It will eventually become almost impossible to detect sophisticated bots.
- What will Google do?