Reflections on Measurement Research: Crooked Lines, Straight Lines, and Moneyshots

Vern Paxson

EECS Department, University of California

International Computer Science Institute

Lawrence Berkeley National Laboratory

Berkeley, California USA

August 16, 2011

First, some acknowledgments:

- LBL: Van Jacobson, Sally Floyd
- UC Berkeley: Domenico Ferrari
- ICSI: Scott Shenker, Mark Handley, Mark Allman, Christian Kreibich, Robin Sommer, Nicholas Weaver, Chris Grier
- UC San Diego: Stefan Savage, Geoff Voelker
- … and 130+ other coauthors

What I’m Going To Try To Convey

- Illuminate* Internet measurement as an empirical science
  - What’s its effective role in network research?
  - What makes its practitioners tick?
    - Hint: no one does it because they LIKE taking measurements
  - The role of nimble opportunism
  - Why measuring Badness is one of today’s most interesting, fun, and challenging areas

* A personal perspective

Part I:

The Crooked Path to Straight Lines

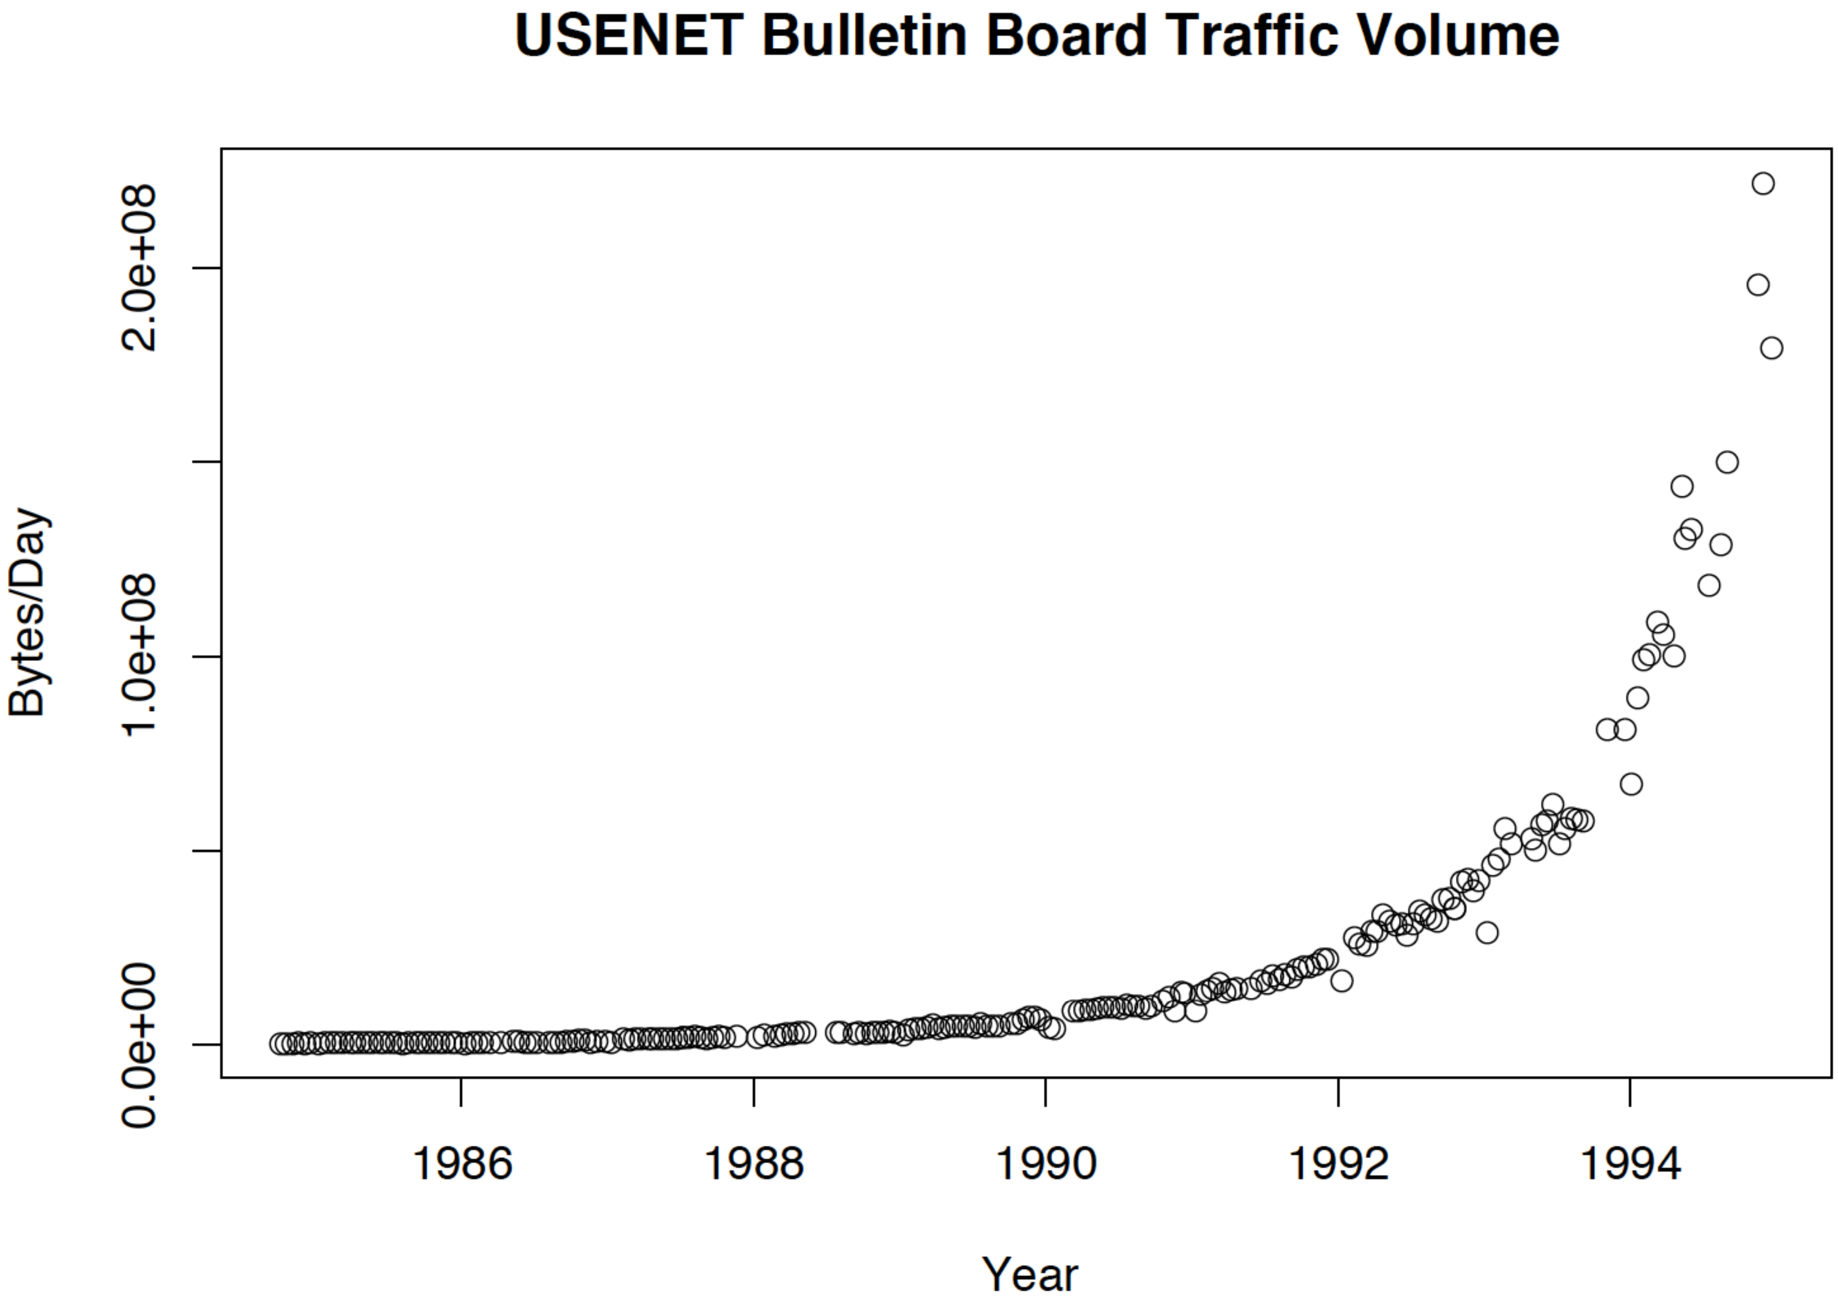
USENET Bulletin Board Traffic Volume
Bytes/Day
0.0e+00
1.0e+08
2.0e+08
1986
1988
1990
1992
1994
Year

USENET Bulletin Board Traffic Volume

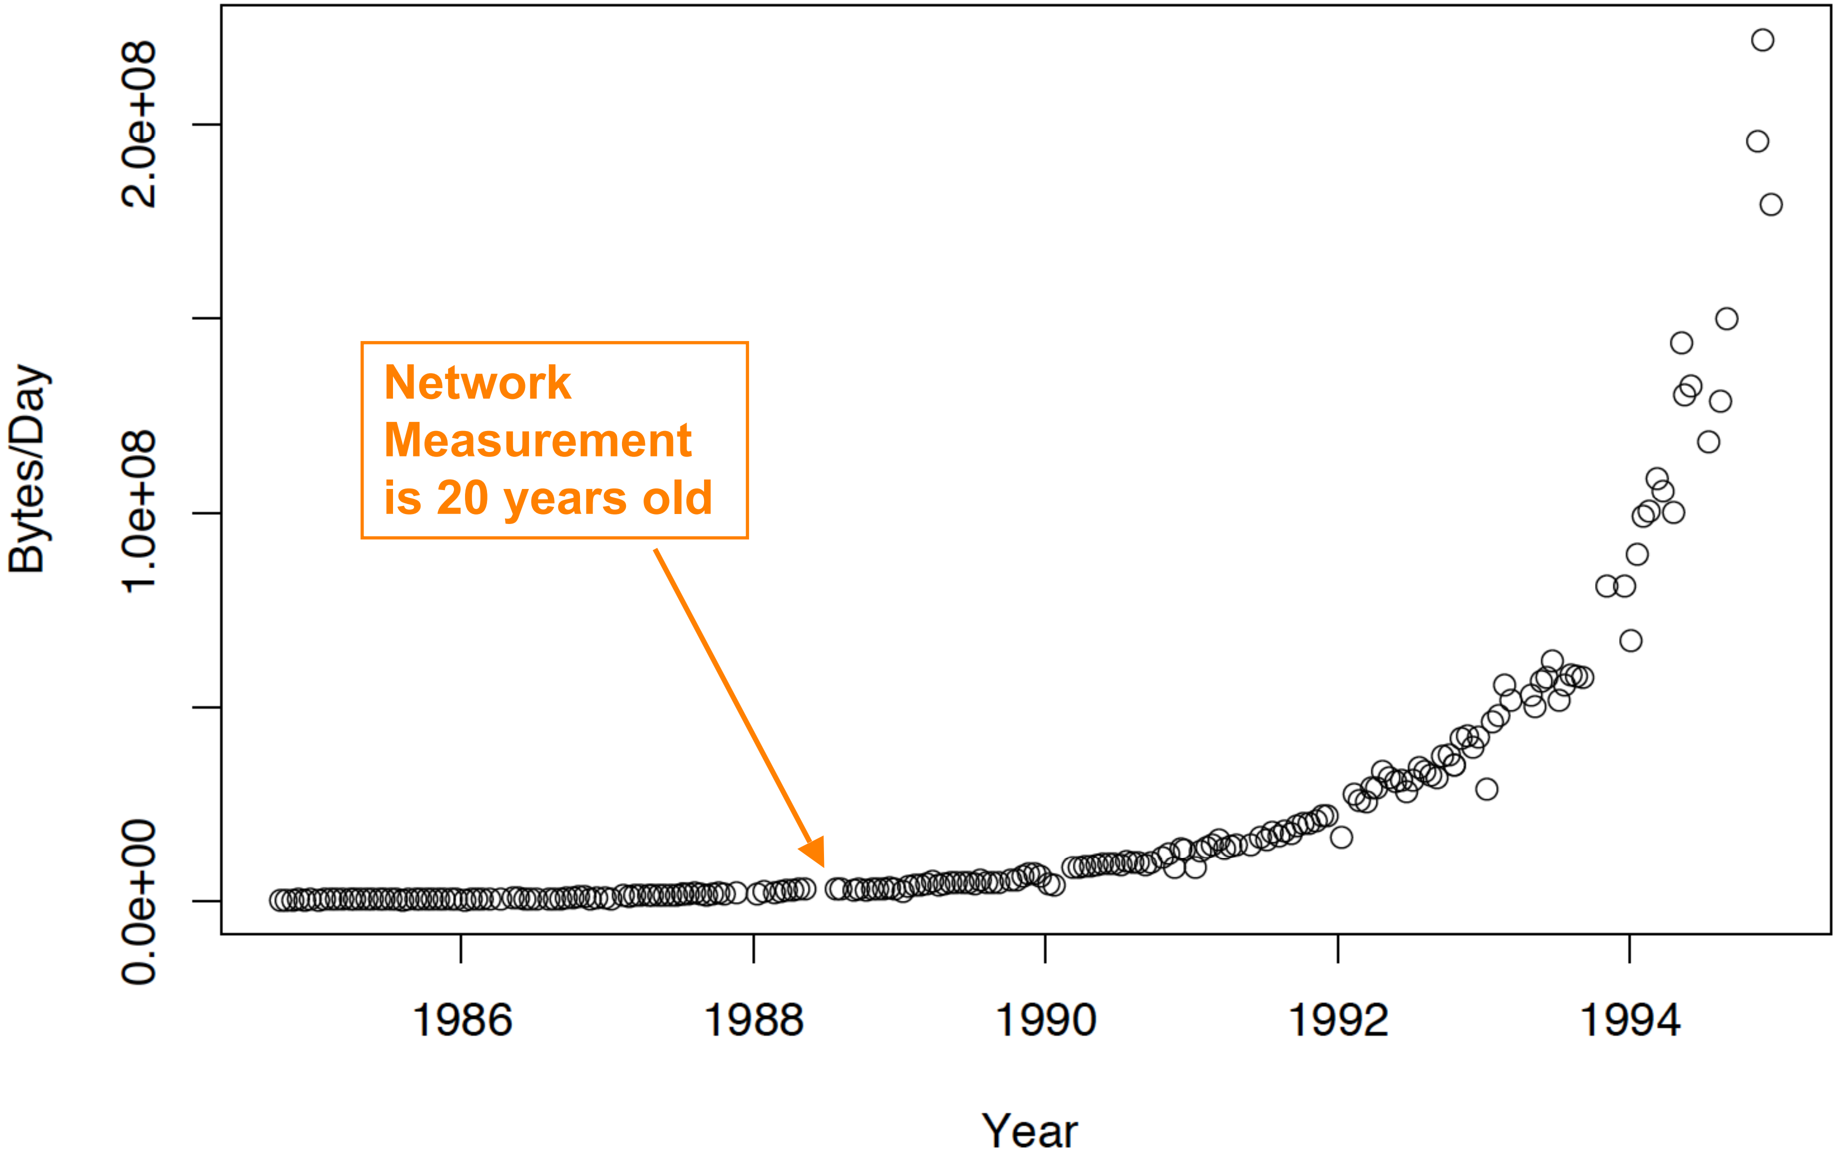

Network Measurement is 20 years old

ITC 12 Torino, June 1988

TWO DECADES OF DATA TRAFFIC MEASUREMENTS: A SURVEY OF PUBLISHED RESULTS, EXPERIENCES AND APPLICABILITY

Peter F. PAWLITA

Siemens AG
Data Systems Division
Munich, Federal Republic of Germany

This survey of published data traffic measurements and their results covers some two decades, from 1966 to 1987. The measurements are classified, reviewed and compared, concerning, e.g., traffic variables, user and system characteristics, statistical and modeling aspects. Emphasis is placed on identification of those variables validated by measurements, on general applicability of results, and a critical view of the present status of the subject.

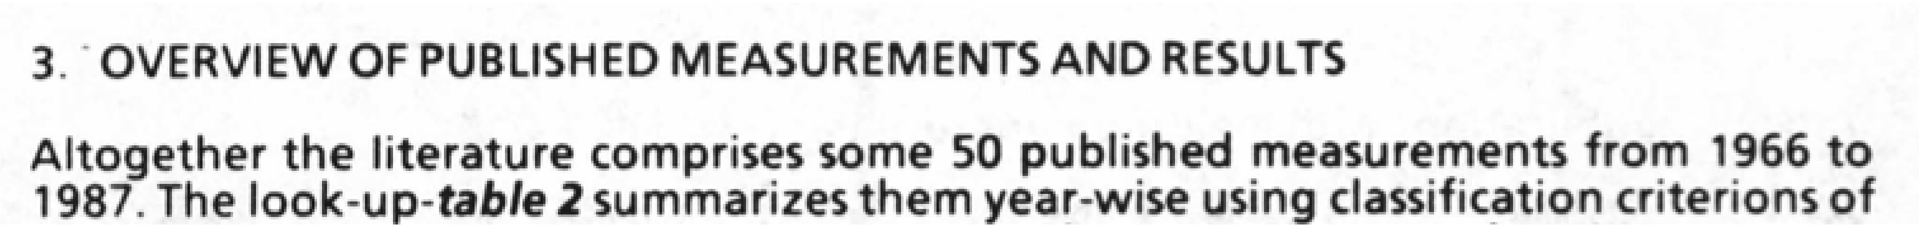

3. OVERVIEW OF PUBLISHED MEASUREMENTS AND RESULTS

Altogether the literature comprises some 50 published measurements from 1966 to 1987. The look-up-***table 2*** summarizes them year-wise using classification criterions of

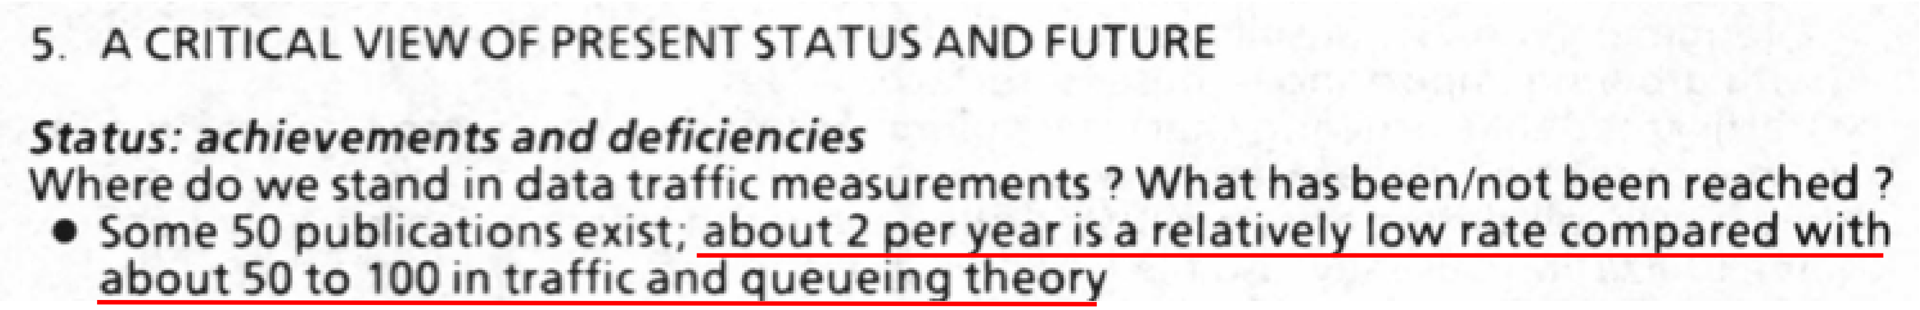

5. A CRITICAL VIEW OF PRESENT STATUS AND FUTURE

Status: achievements and deficiencies
Where do we stand in data traffic measurements ? What has been/not been reached ?

- Some 50 publications exist; about 2 per year is a relatively low rate compared with about 50 to 100 in traffic and queueing theory

USENET Bulletin Board Traffic Volume

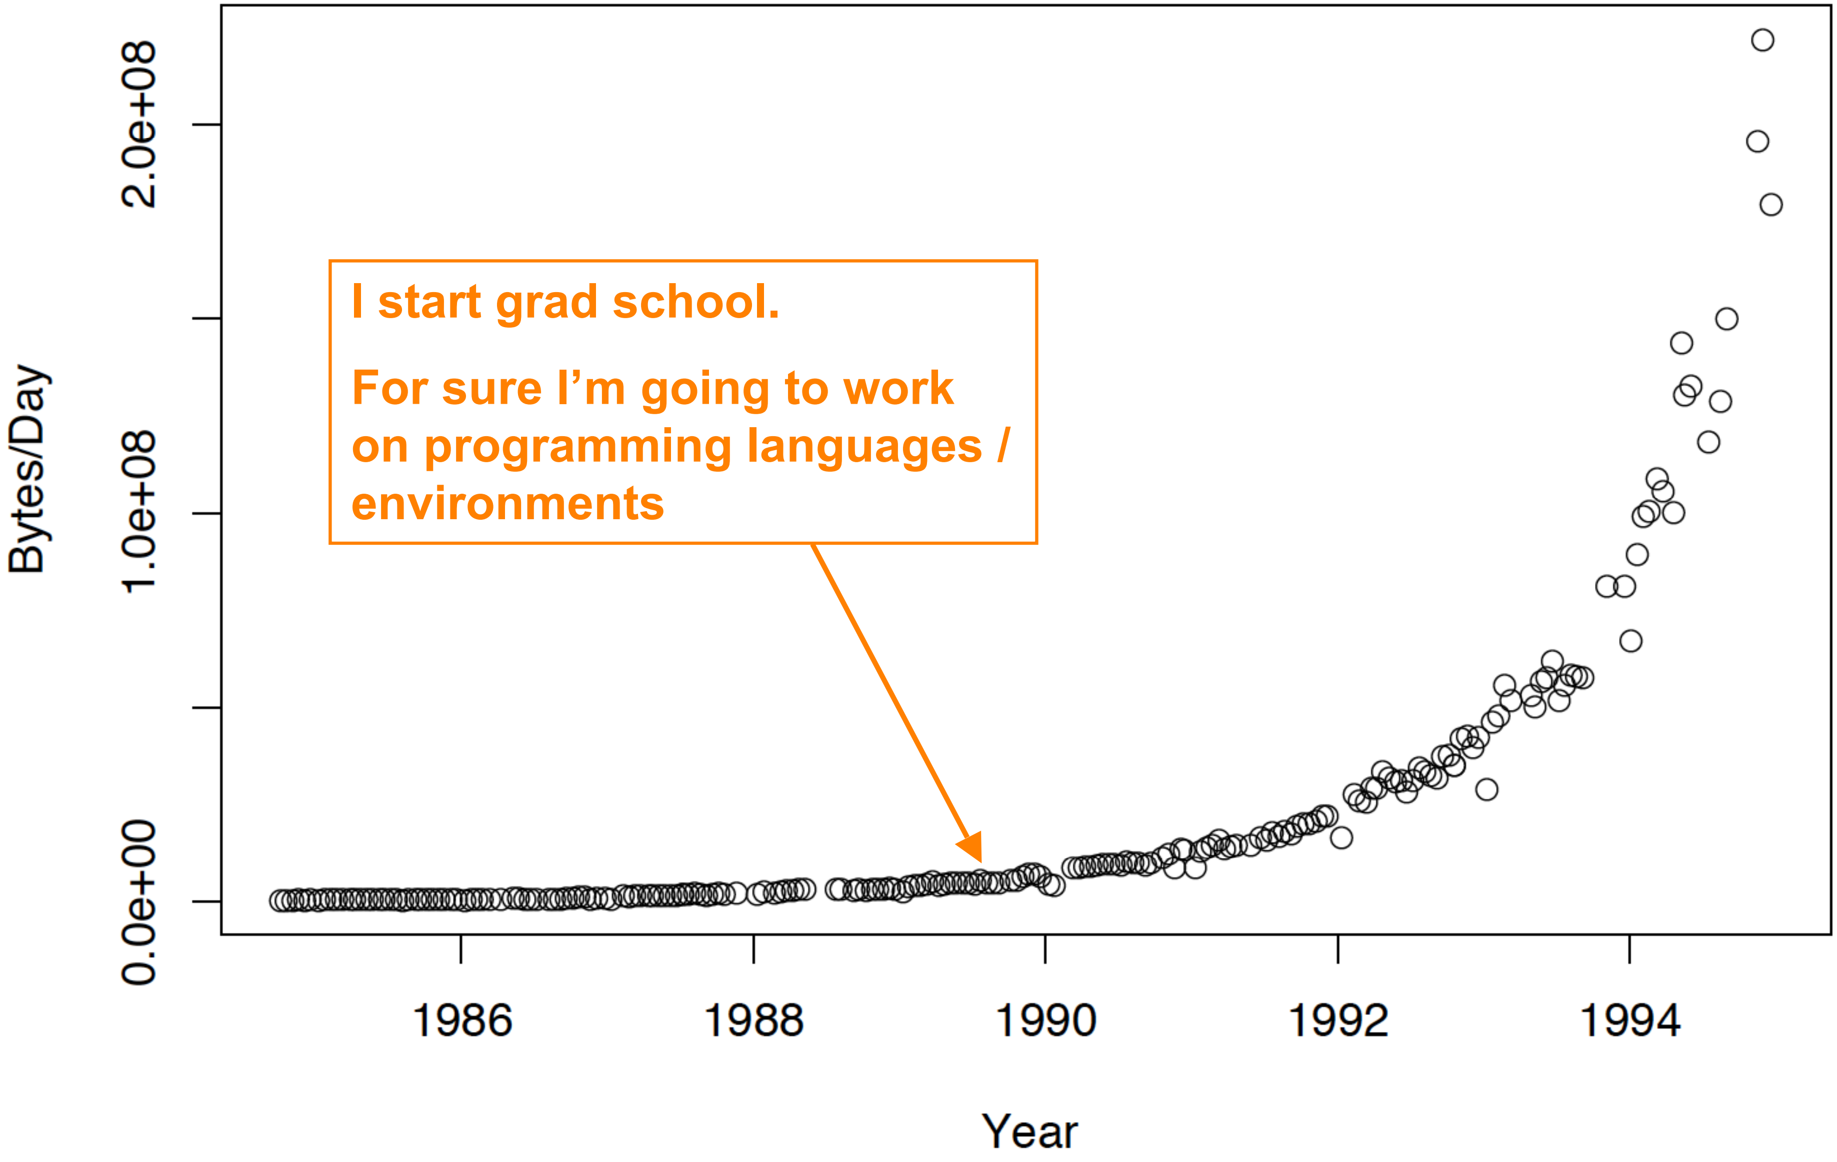

I start grad school.

For sure I’m going to work on programming languages / environments

USENET Bulletin Board Traffic Volume

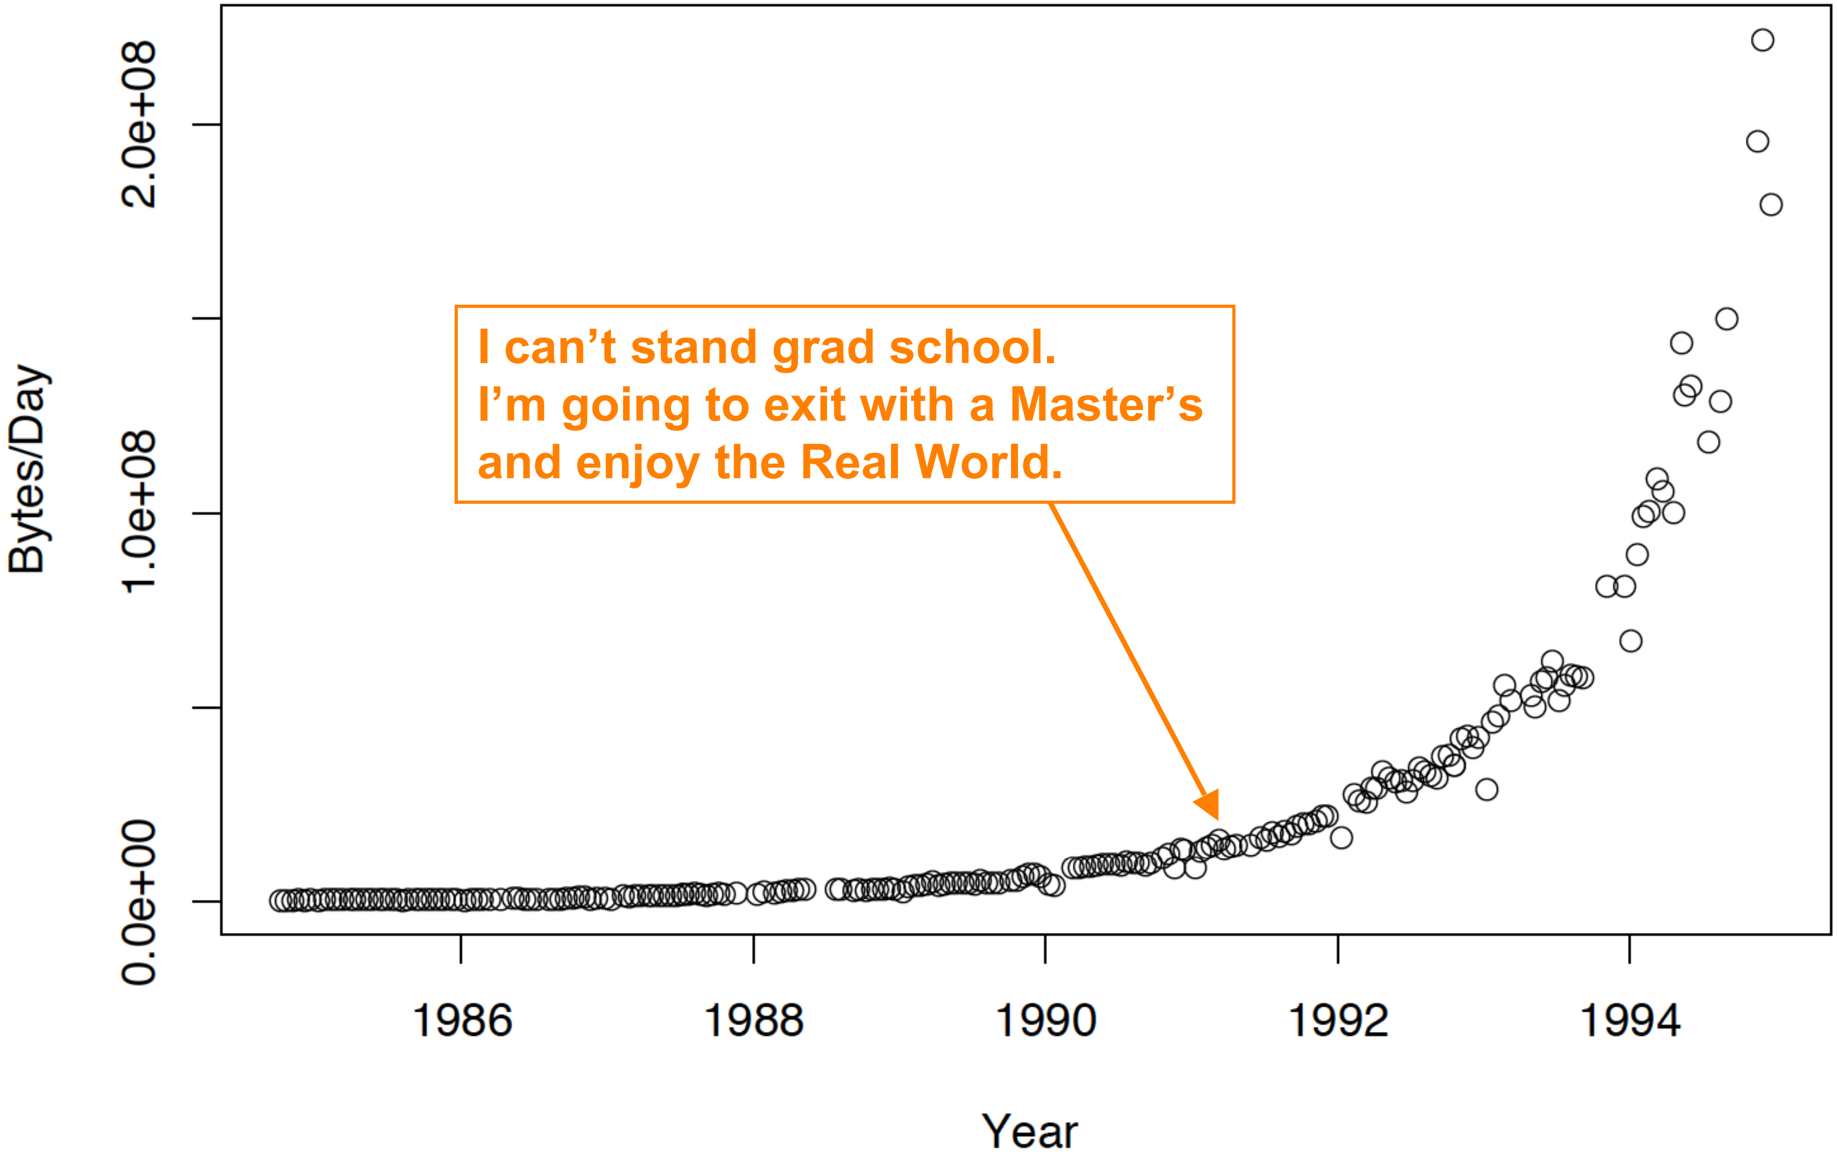

I can’t stand grad school.
I’m going to exit with a Master’s
and enjoy the Real World.

USENET Bulletin Board Traffic Volume

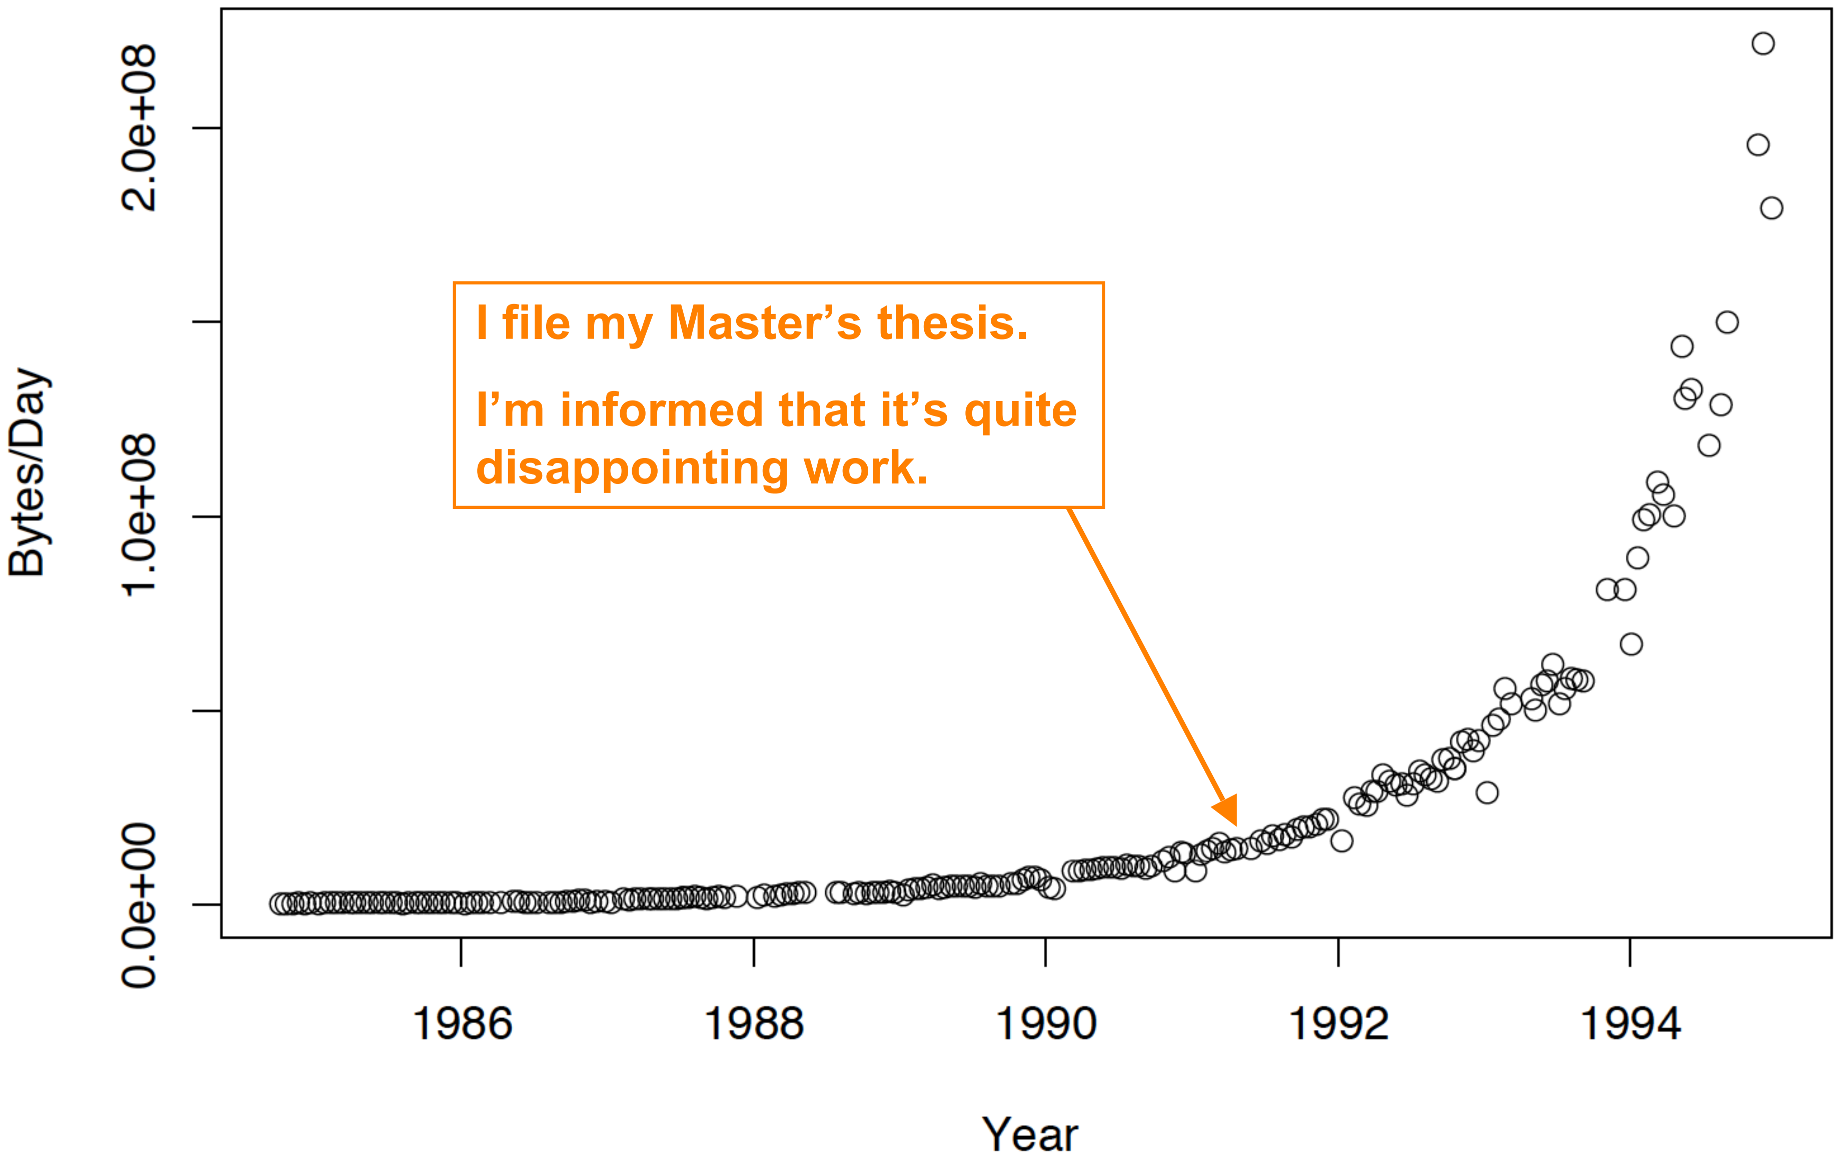

I file my Master’s thesis.

I’m informed that it’s quite disappointing work.

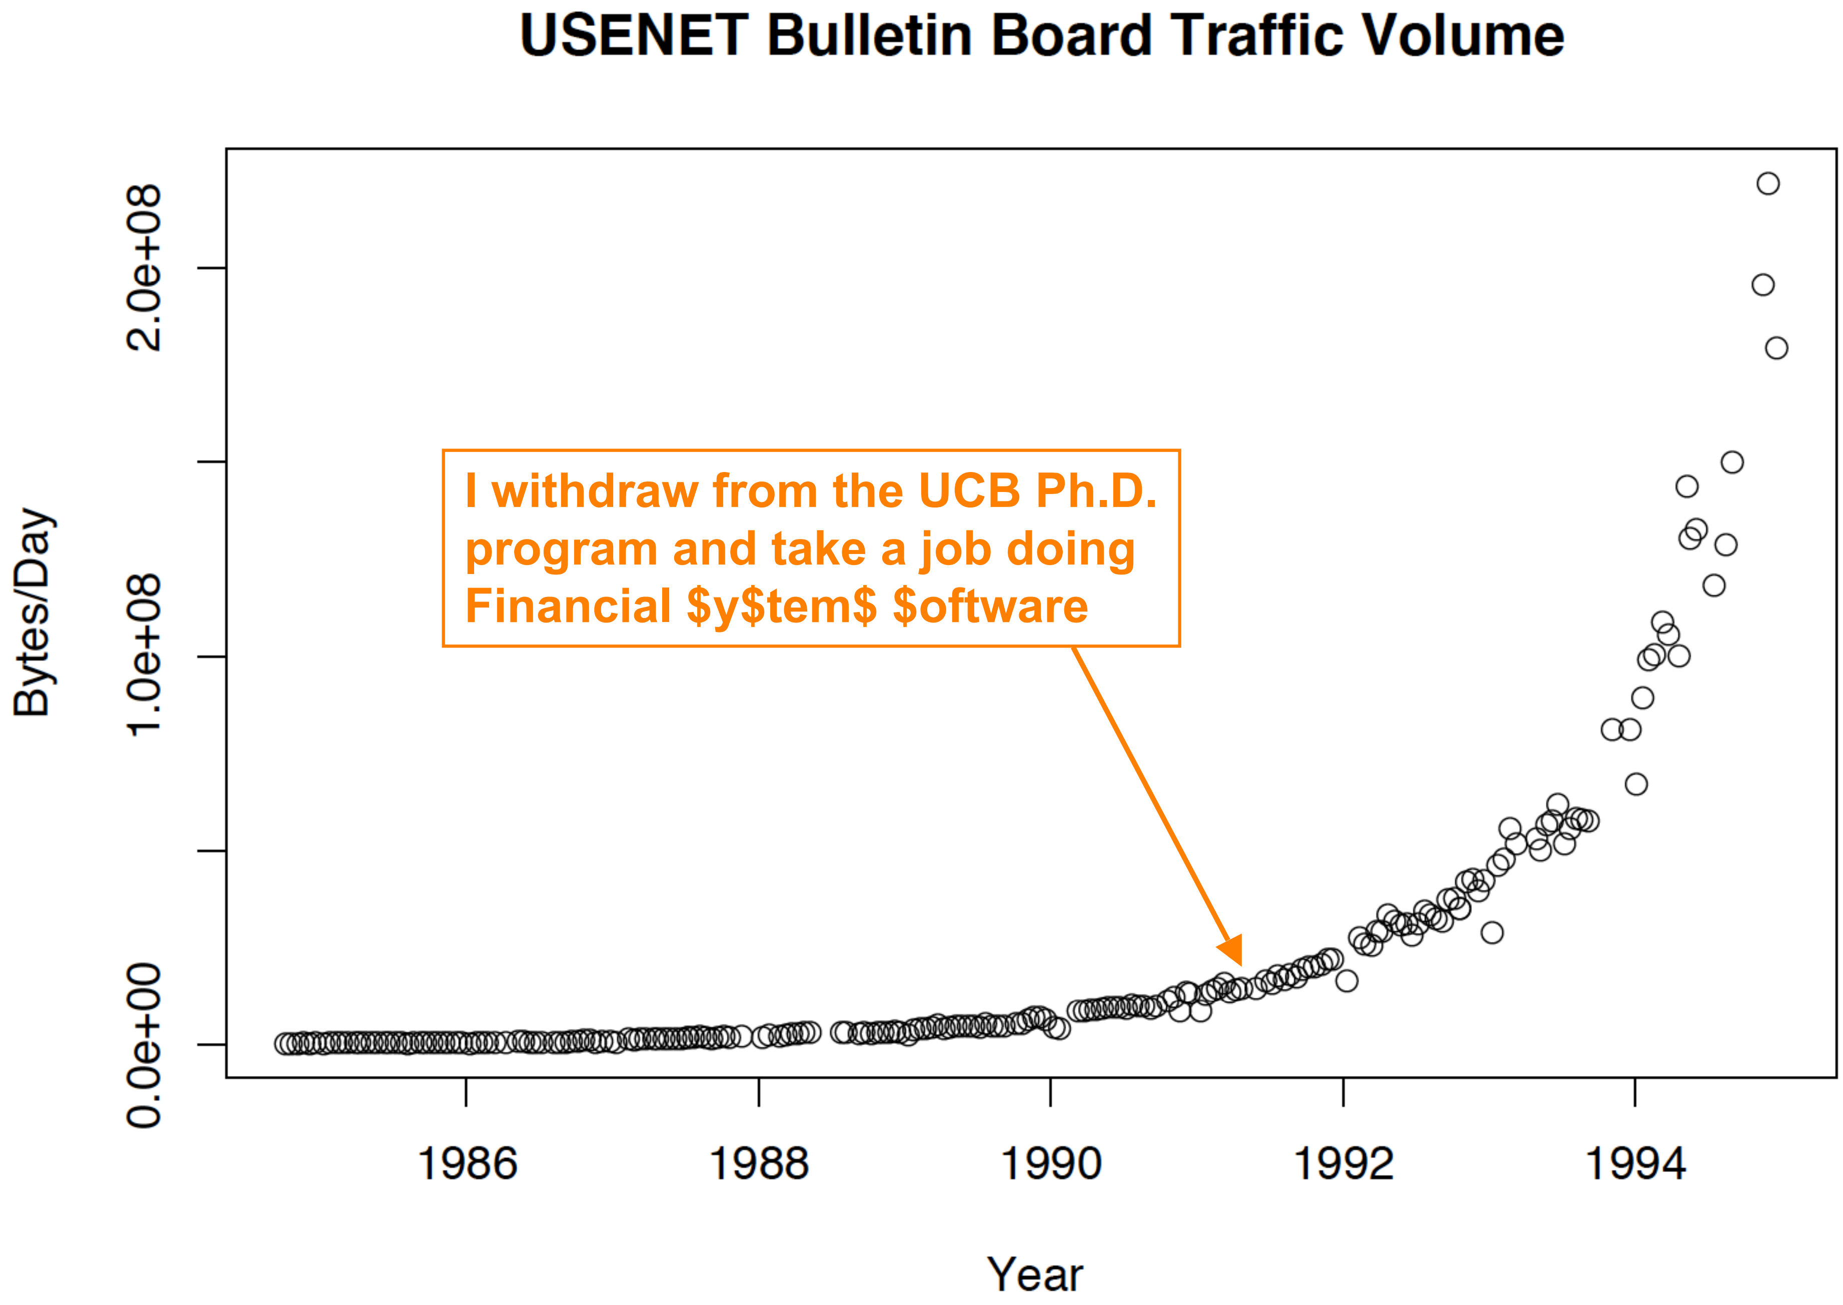

I withdraw from the UCB Ph.D. program and take a job doing Financial $y$tem$ $oftware

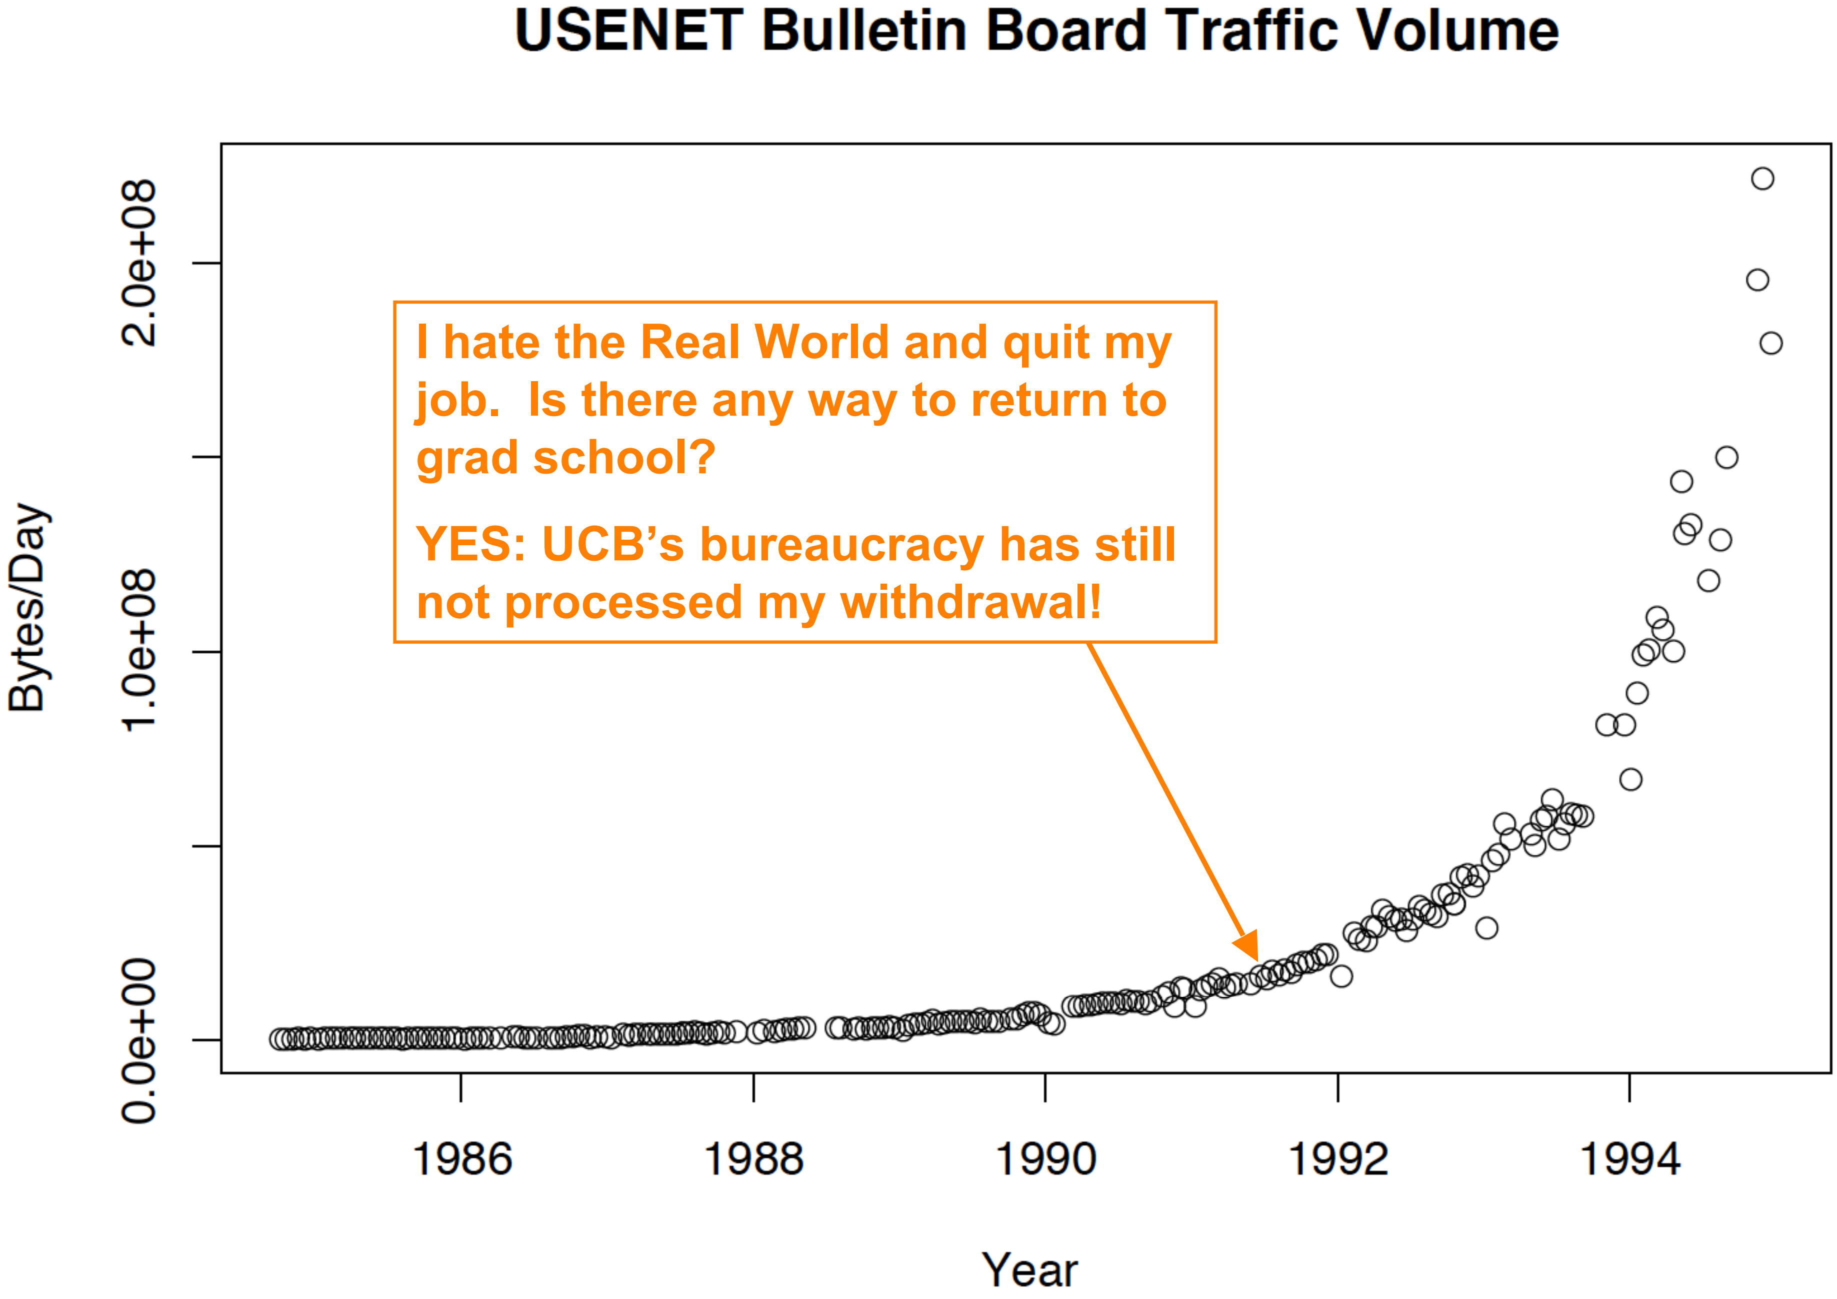

I hate the Real World and quit my job. Is there any way to return to grad school?

YES: UCB's bureaucracy has still not processed my withdrawal!

USENET Bulletin Board Traffic Volume

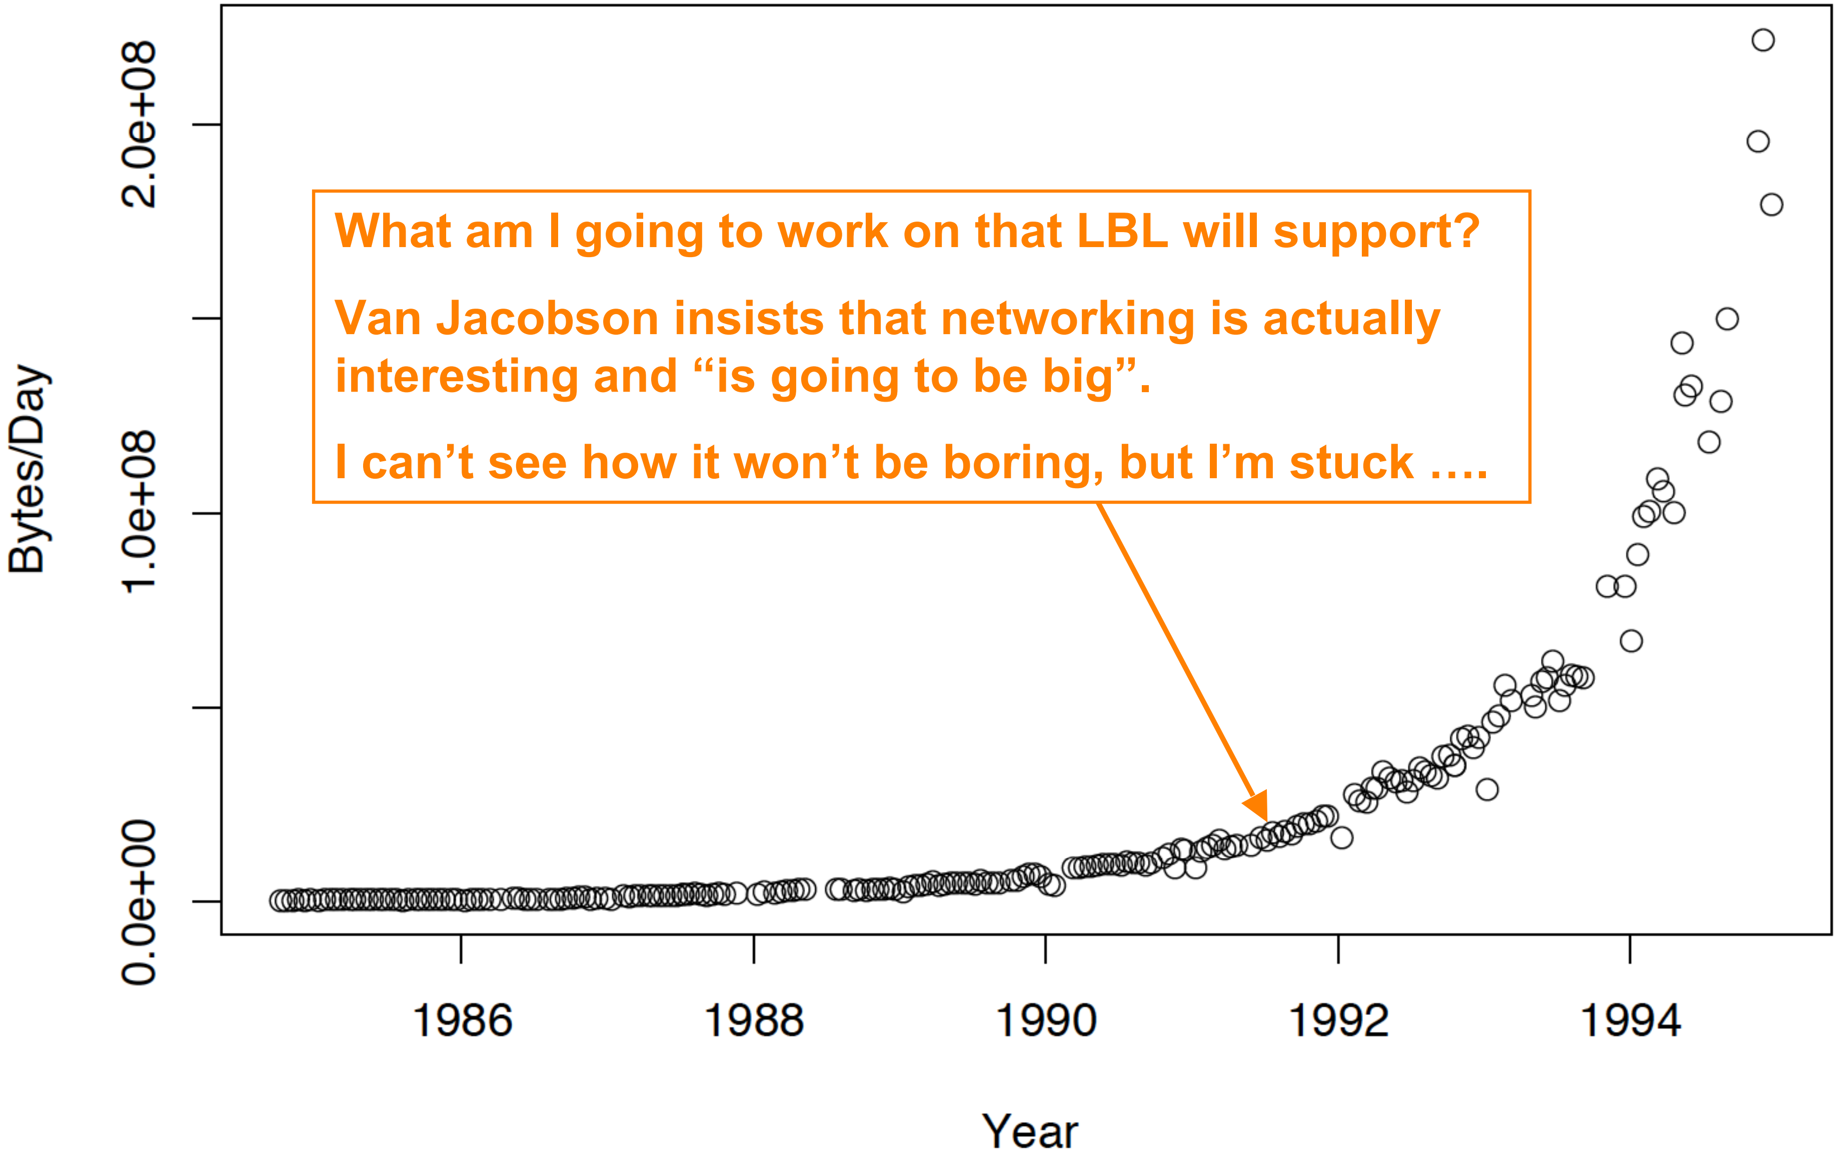

What am I going to work on that LBL will support?

Van Jacobson insists that networking is actually interesting and “is going to be big”.

I can’t see how it won’t be boring, but I’m stuck ….

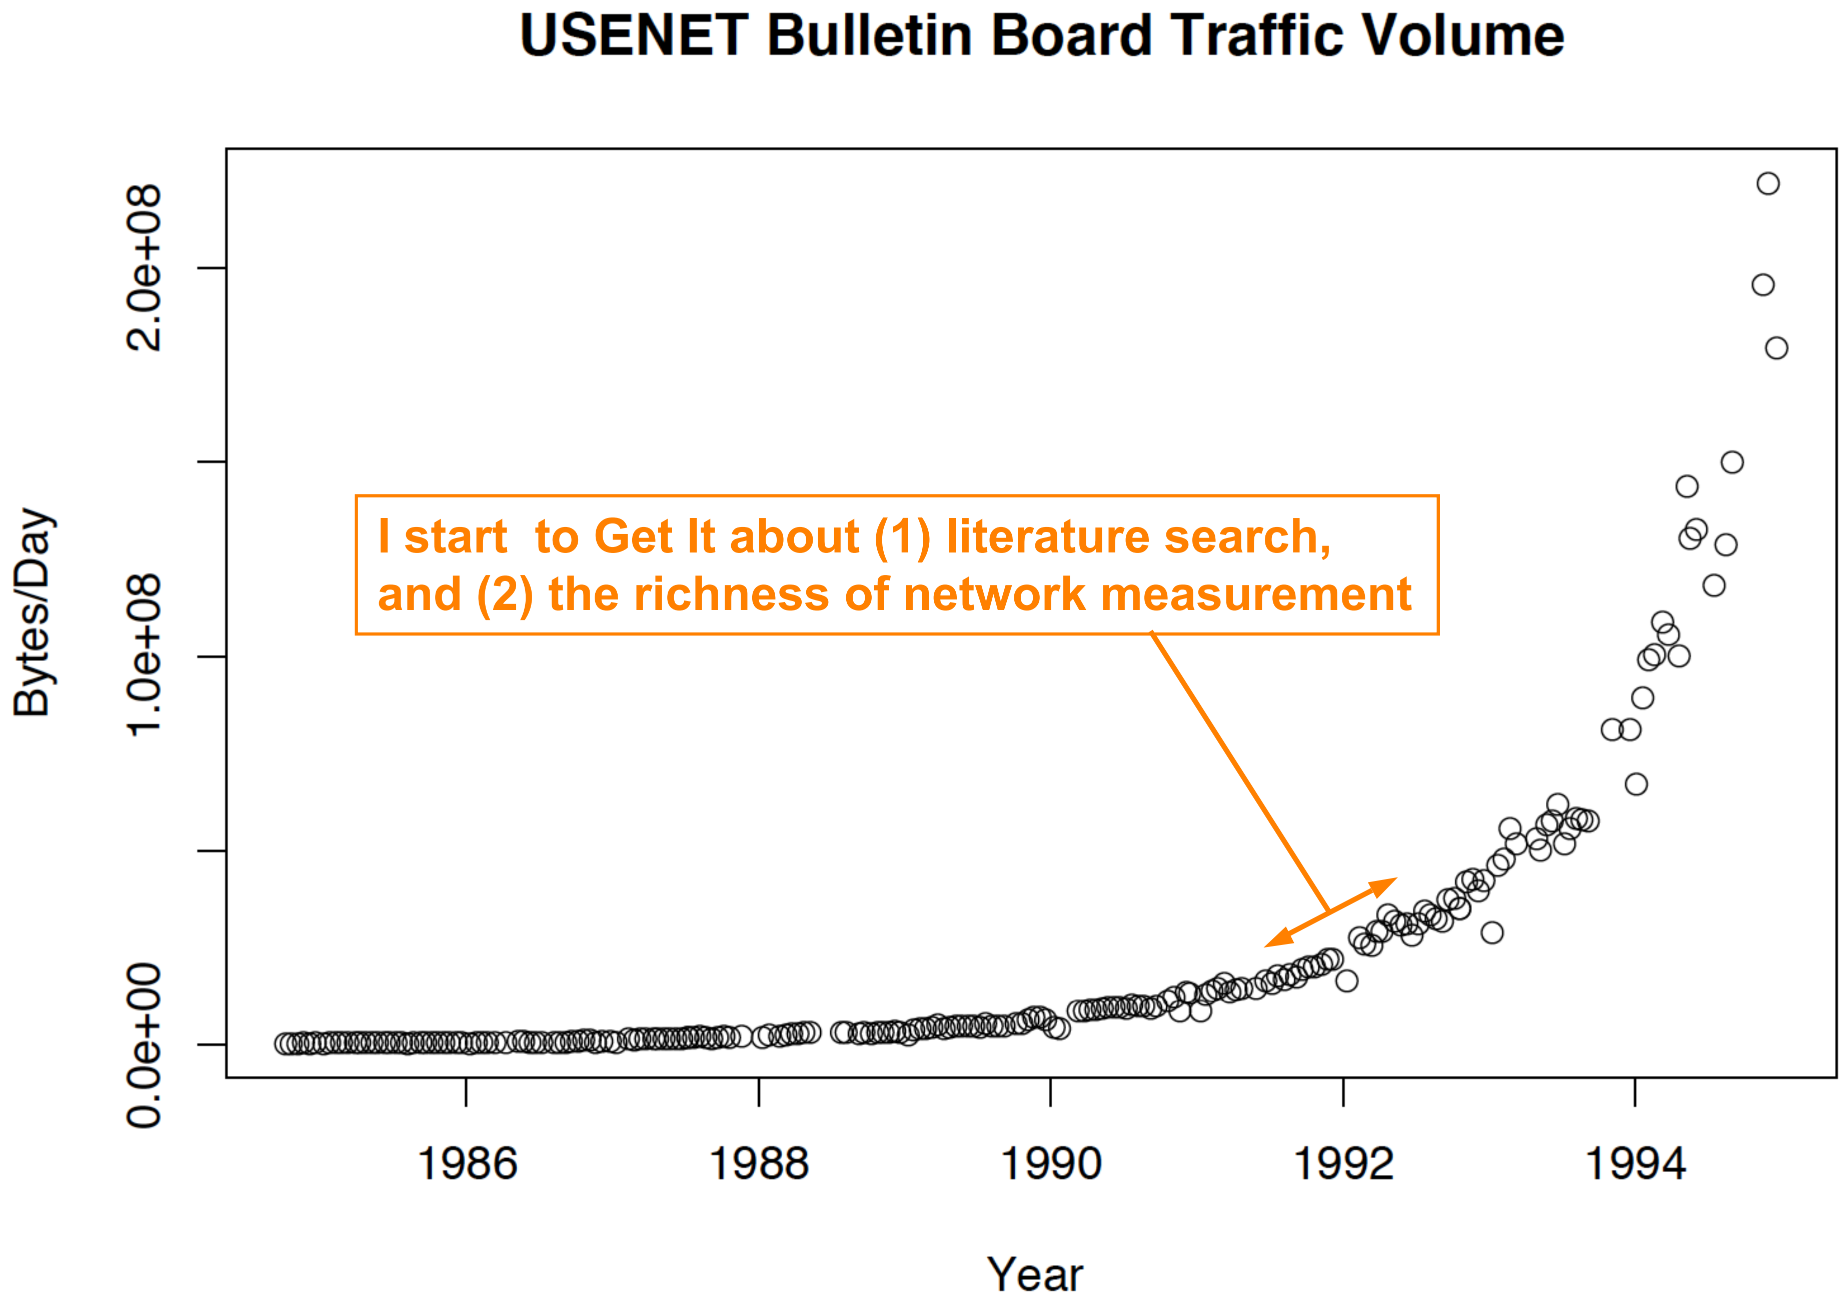

I start to Get It about (1) literature search, and (2) the richness of network measurement

USENET Bulletin Board Traffic Volume

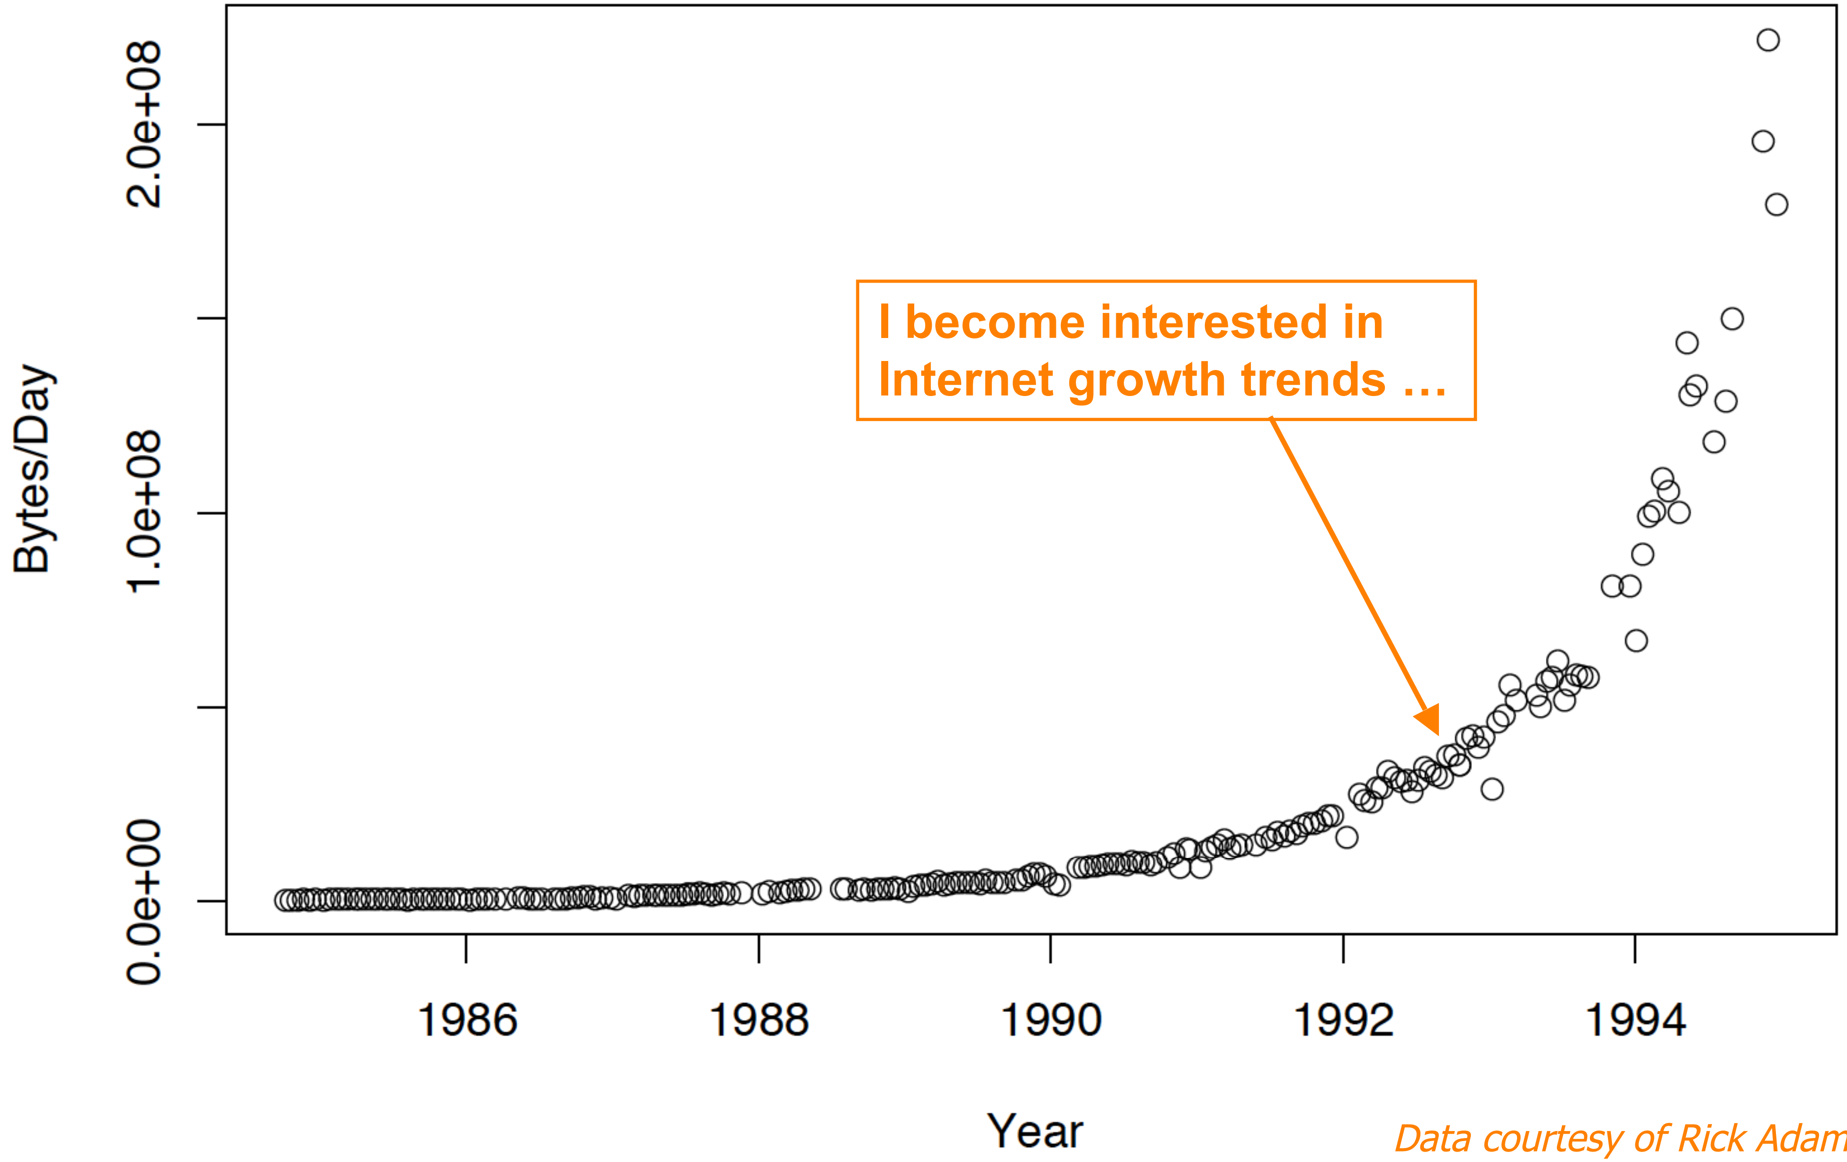

I become interested in Internet growth trends …

Data courtesy of Rick Adams

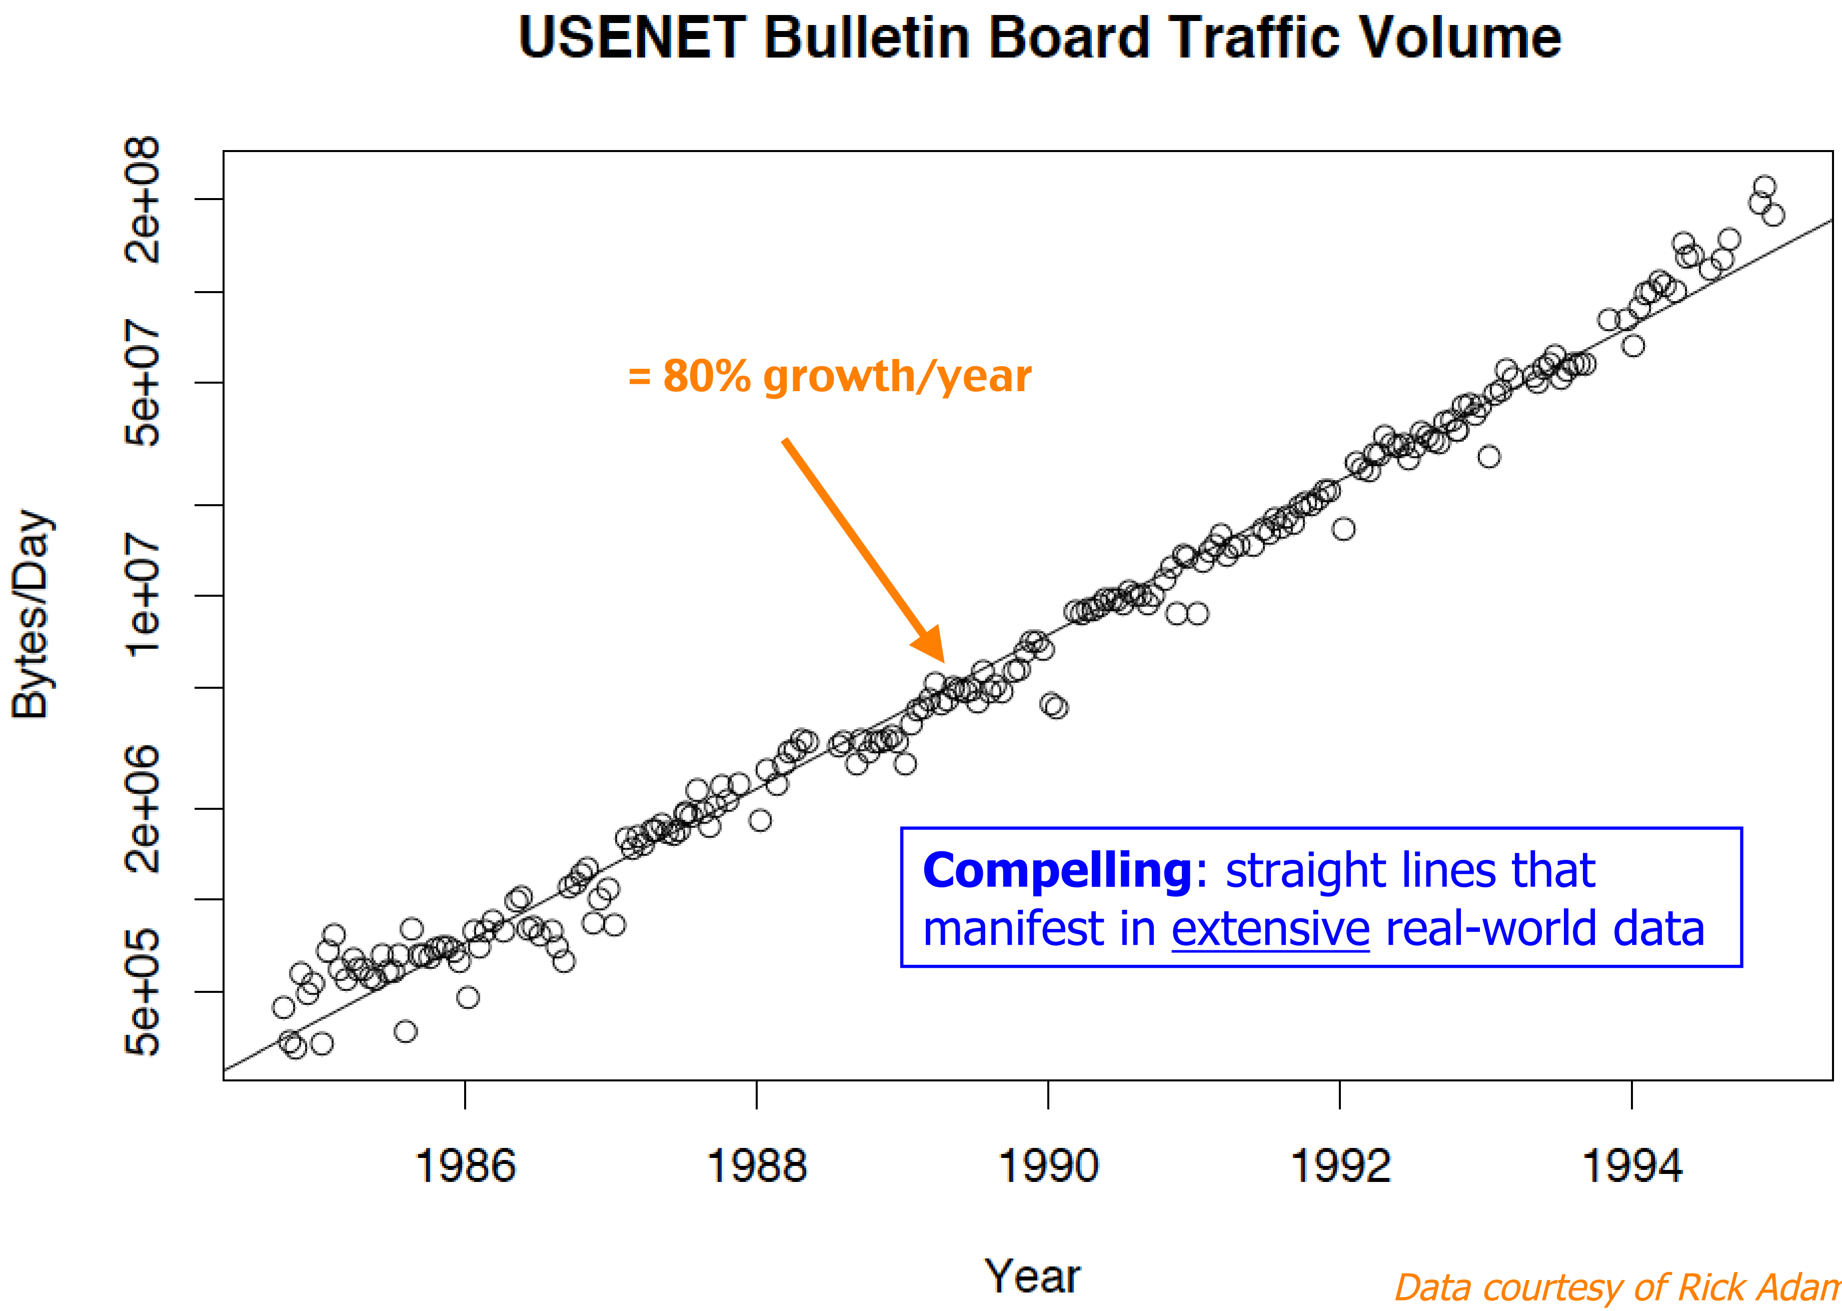

= 80% growth/year

Compelling: straight lines that manifest in extensive real-world data

Data courtesy of Rick Adams

USENET Bulletin Board Traffic Volume

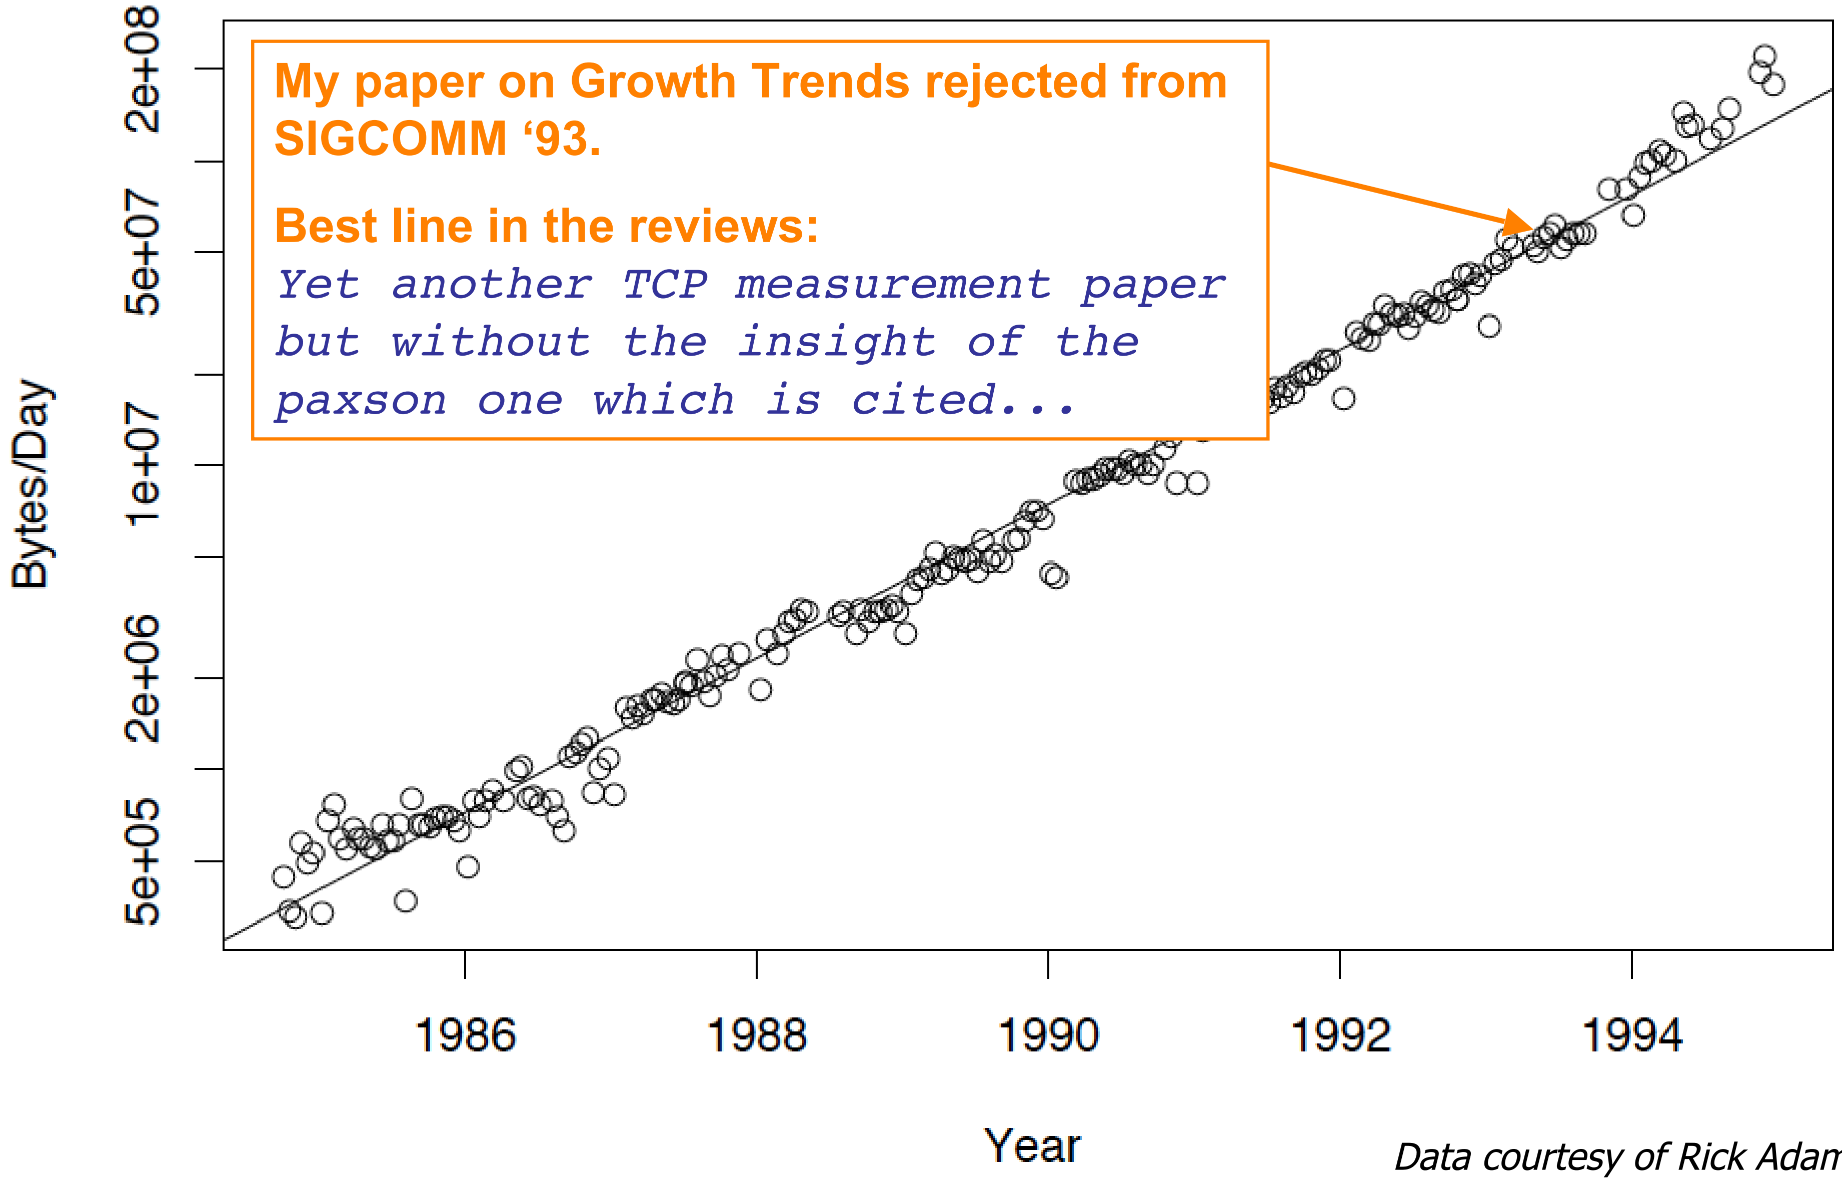

My paper on Growth Trends rejected from SIGCOMM '93.

Best line in the reviews:

Yet another TCP measurement paper but without the insight of the paxson one which is cited...

Data courtesy of Rick Adams

USENET Bulletin Board Traffic Volume

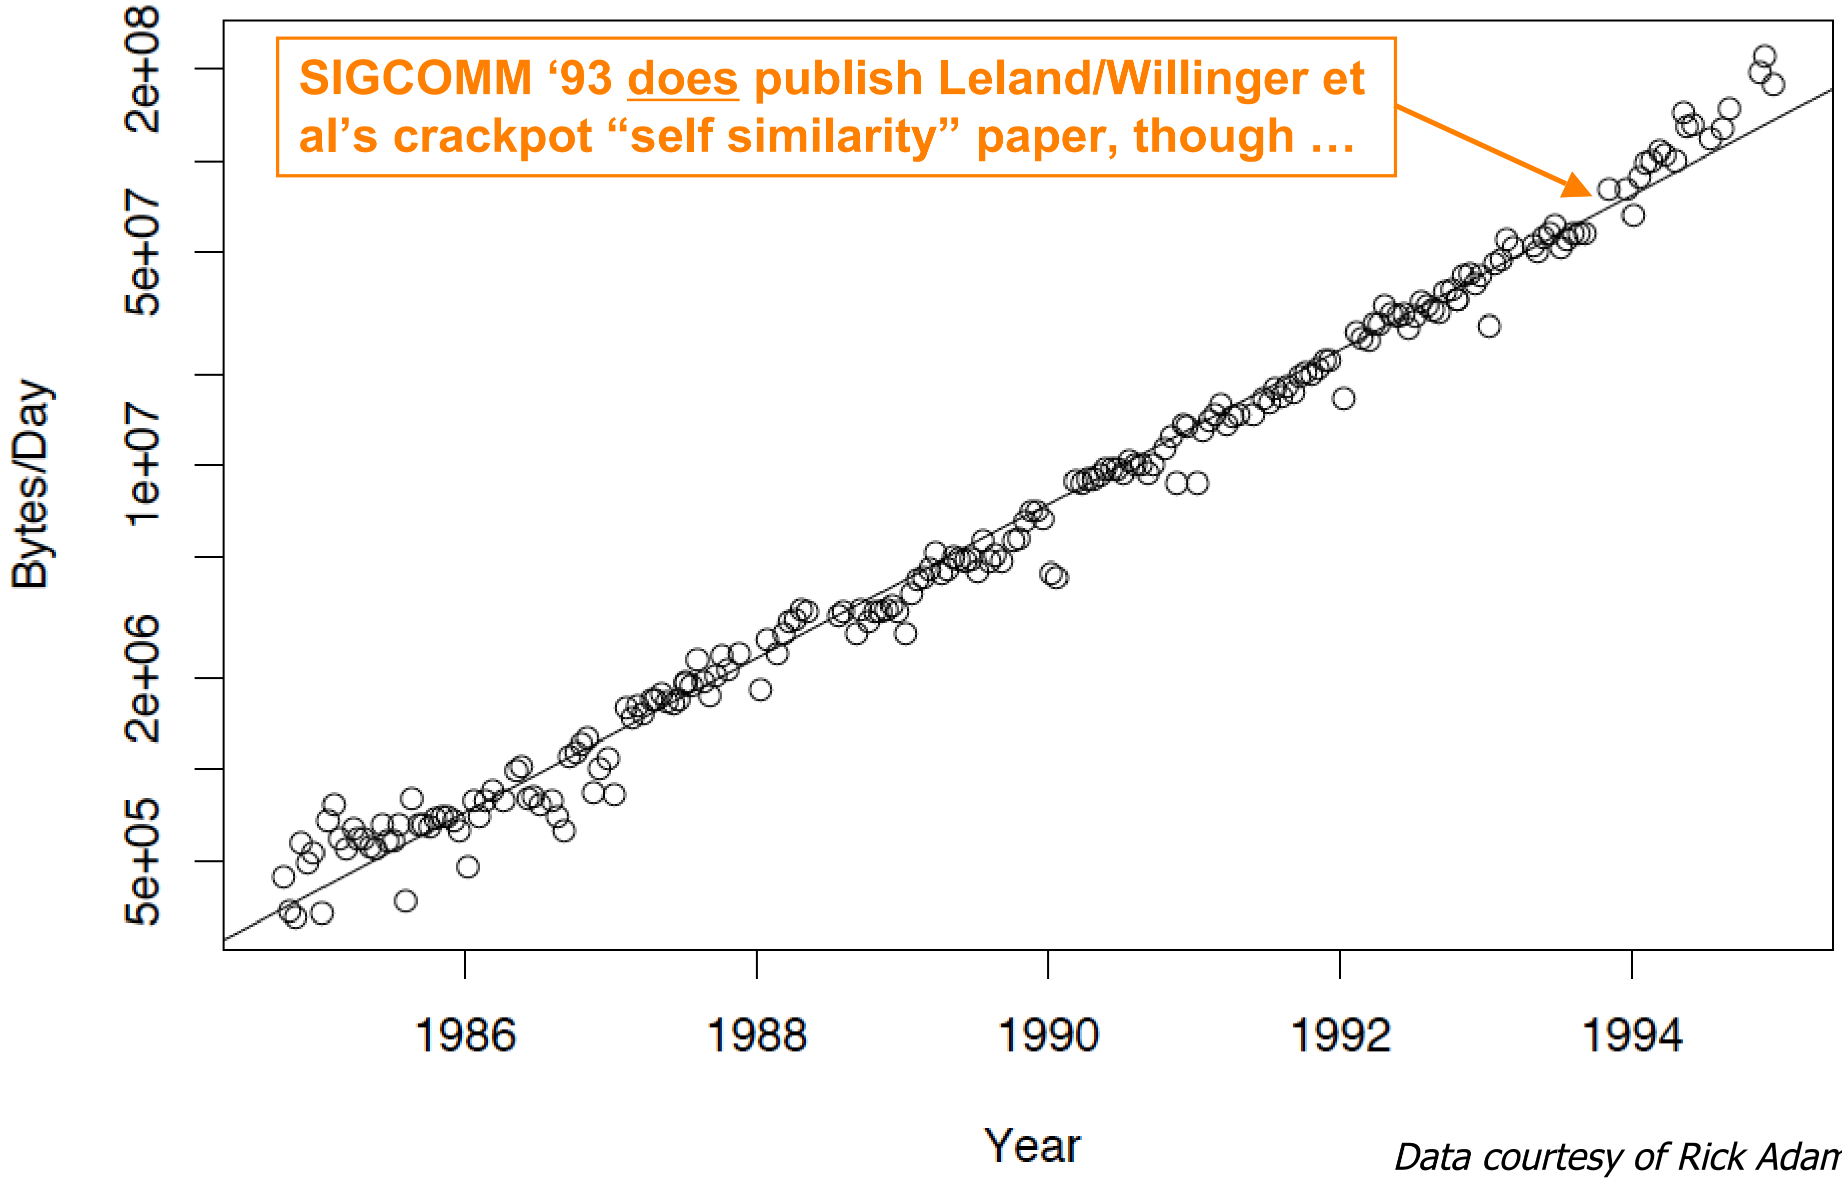

SIGCOMM '93 does publish Leland/Willinger et al's crackpot "self similarity" paper, though …

Data courtesy of Rick Adams

[On the self-similar nature of Ethernet traffic (extended version)](#)
WE Leland, MS Taqqu, W Willinger… - Networking, IEEE/ACM …, 1994 - ieeexplore.ieee.org
… The main objective of this paper is to establish in a statistically rigorous manner the **self-similarity** characteristic of the very **high** quality, **high** time-resolution Ethemet LAN Manuscript received July 1, 1993; revised January 15, 1994; approved by IEEE/ACM TRANSACTIONS ON …
Cited by 4759 - Related articles - All 69 versions

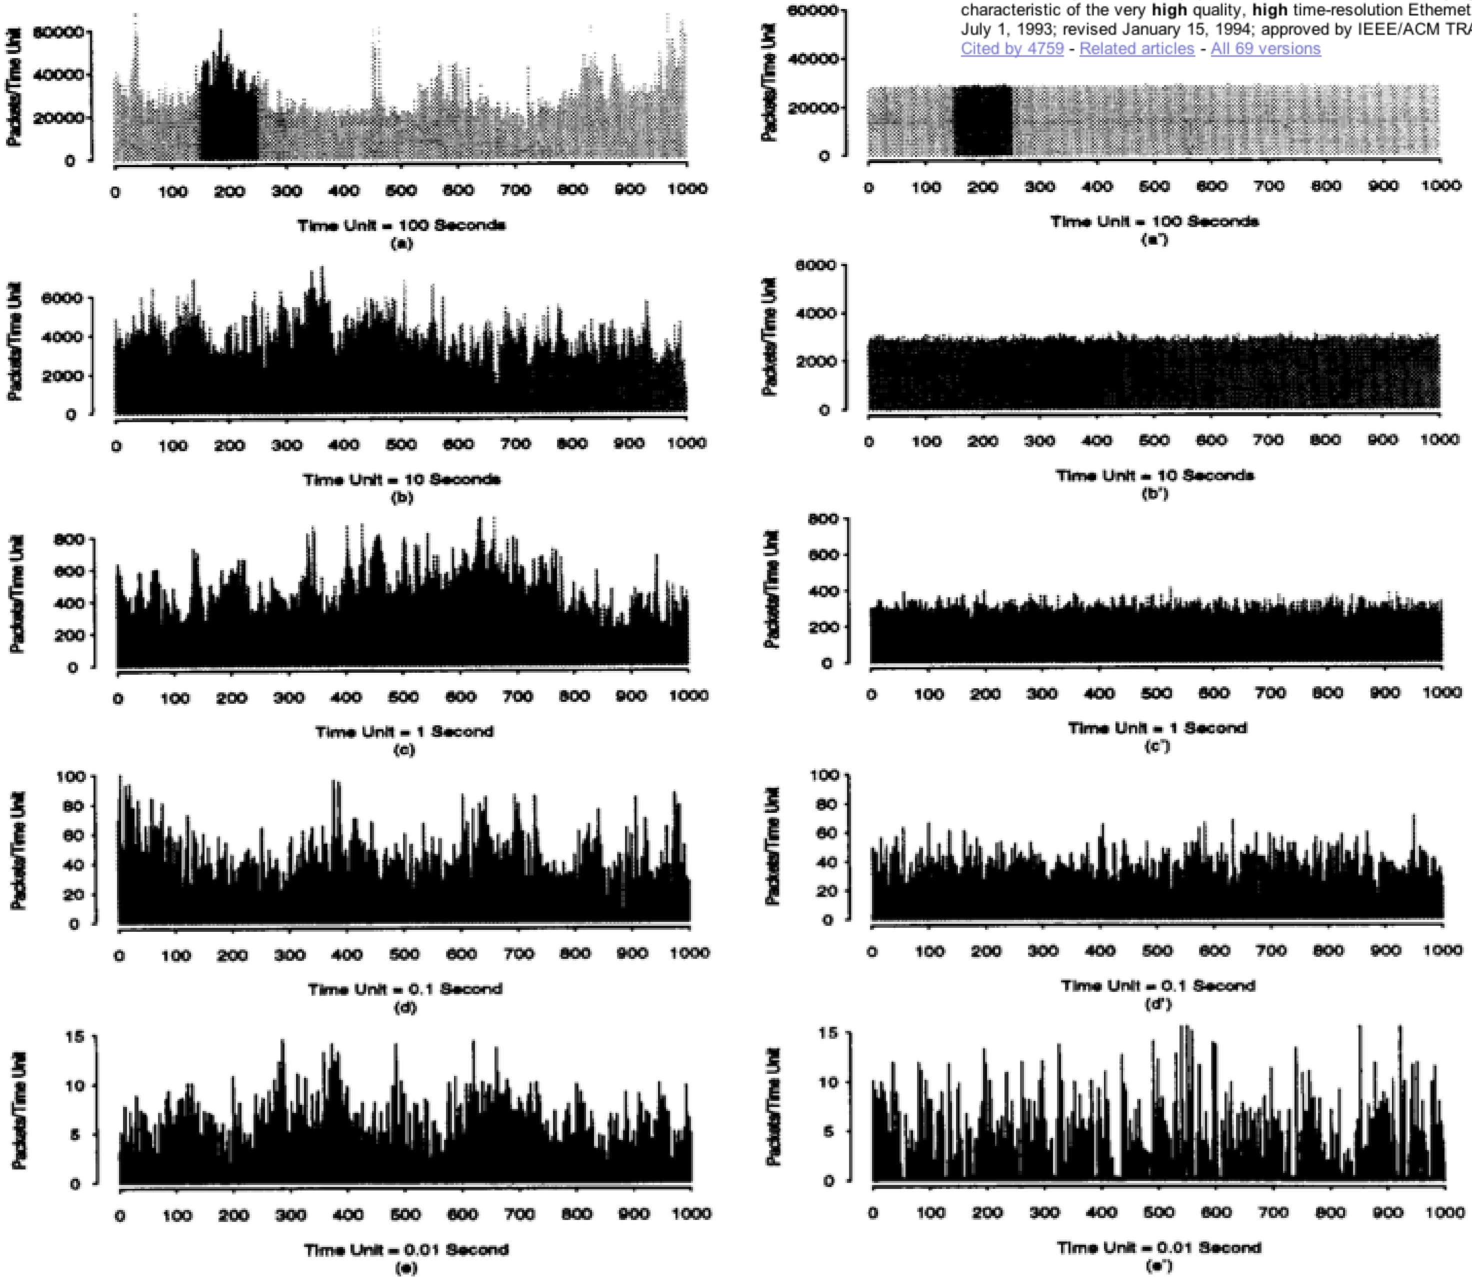

Fig. 4. Pictorial "proof" of self-similarity: Ethernet traffic (packets per time unit) on five different time scales (a)–(e). For comparison, synthetic traffic from an appropriately chosen compound Poisson model on the same five different time scales (a')–(e').

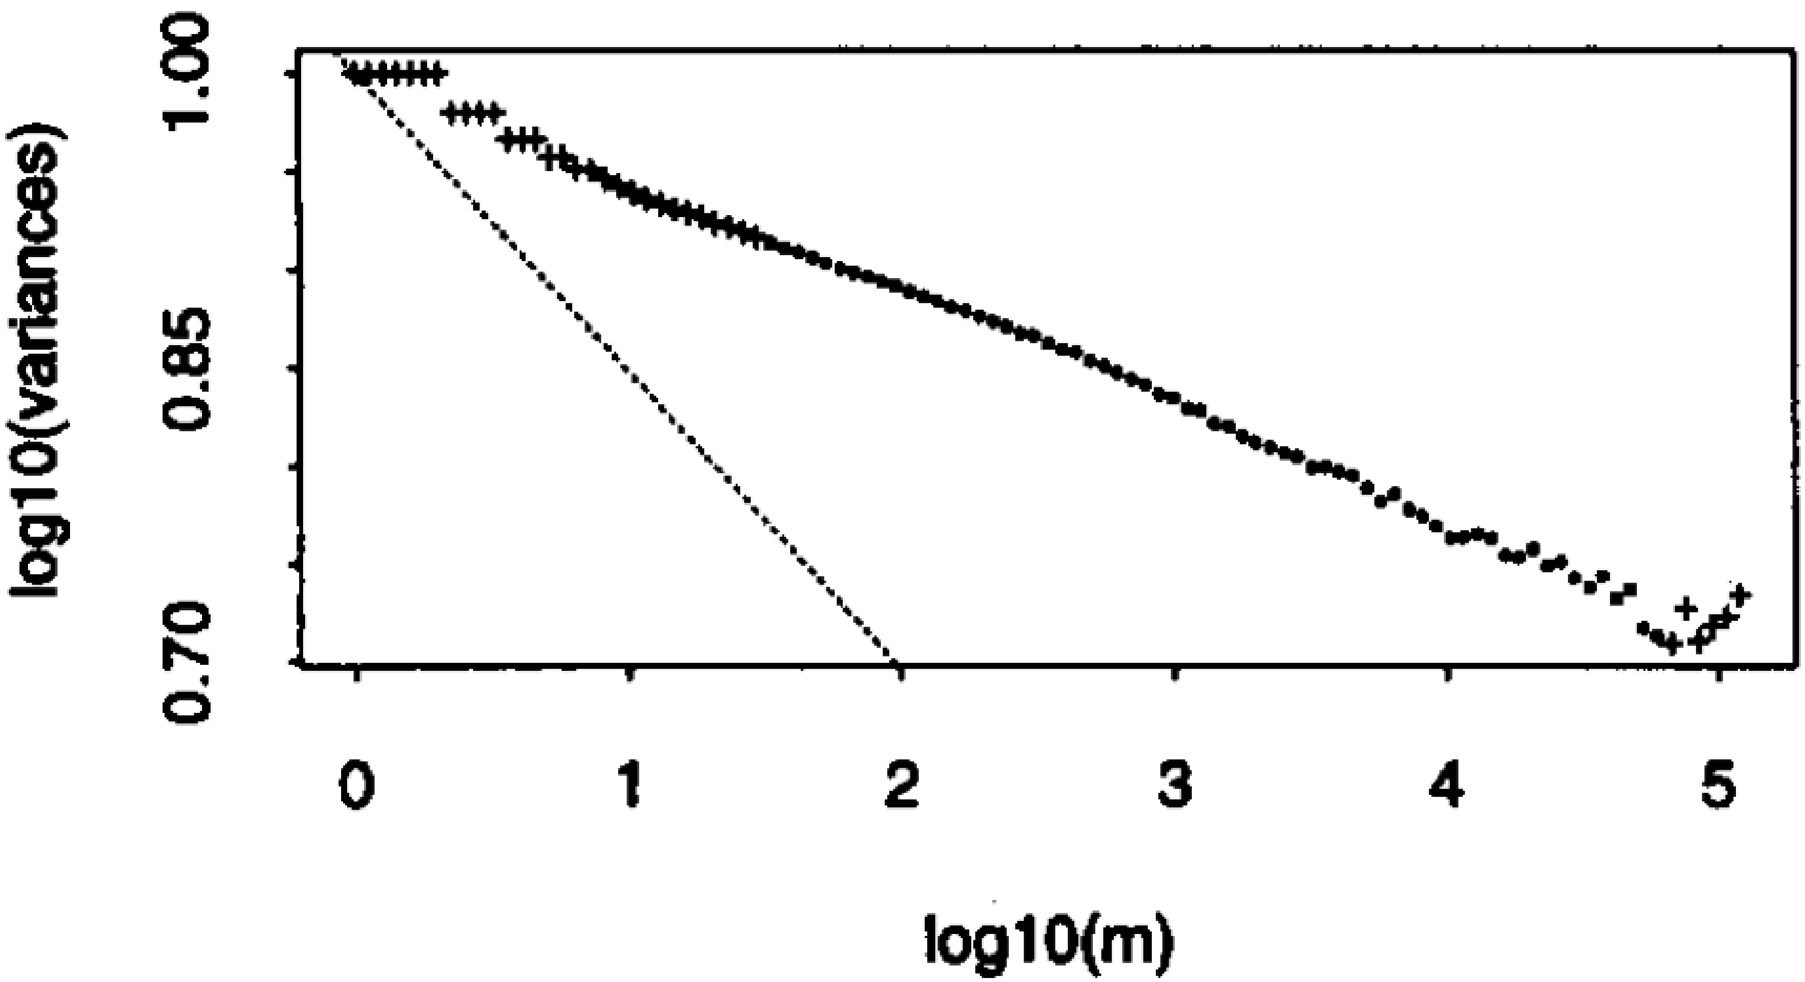

m = scale of aggregation, e.g., m=10^2 aggregates at 100 msec
variance = σ^2 for packet (or byte) arrival process at that aggregation

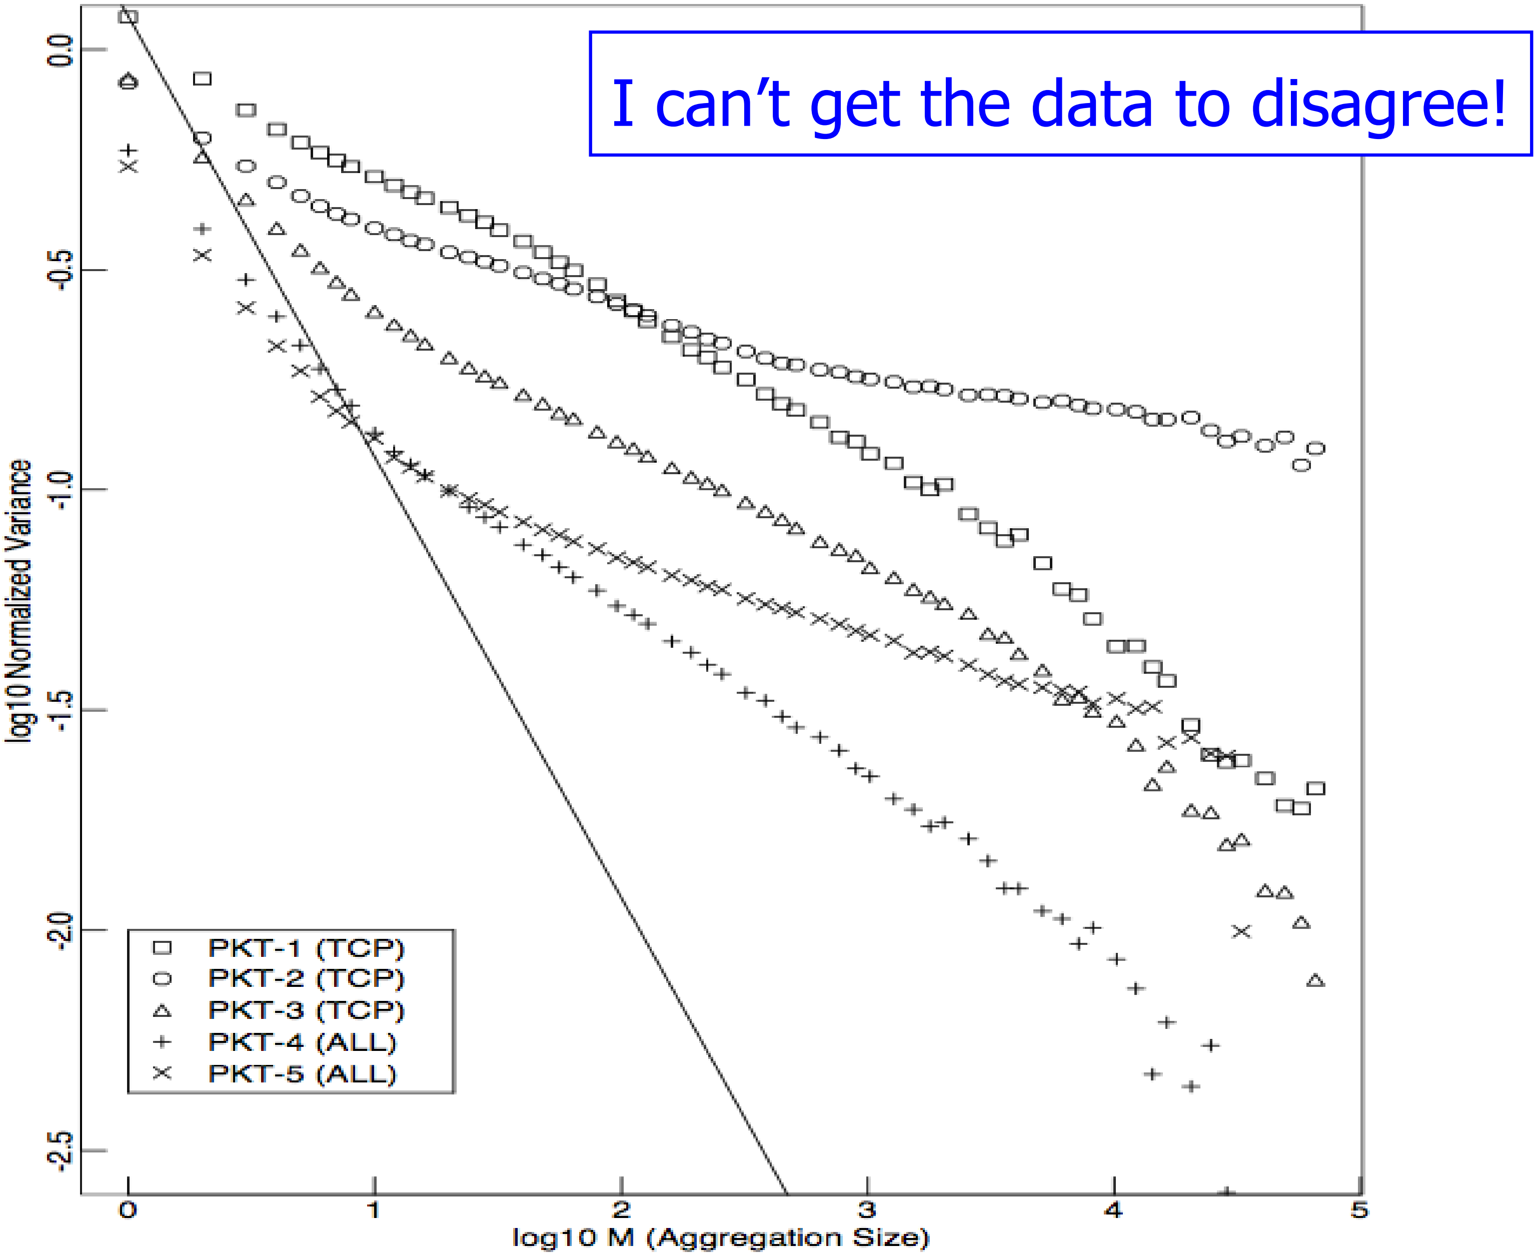

I can’t get the data to disagree!

[On the self-similar nature of Ethernet traffic (extended version)]
WE Leland, MS Taqqu, W Willinger… - Networking, IEEE/ACM …, 1994 - ieeexplore.ieee.org
lf-similarity
eived

Their work also proposed an explanation for self-similarity that is **predictive**: Network activity should be marked by sizes/durations that are *heavy-tailed*.

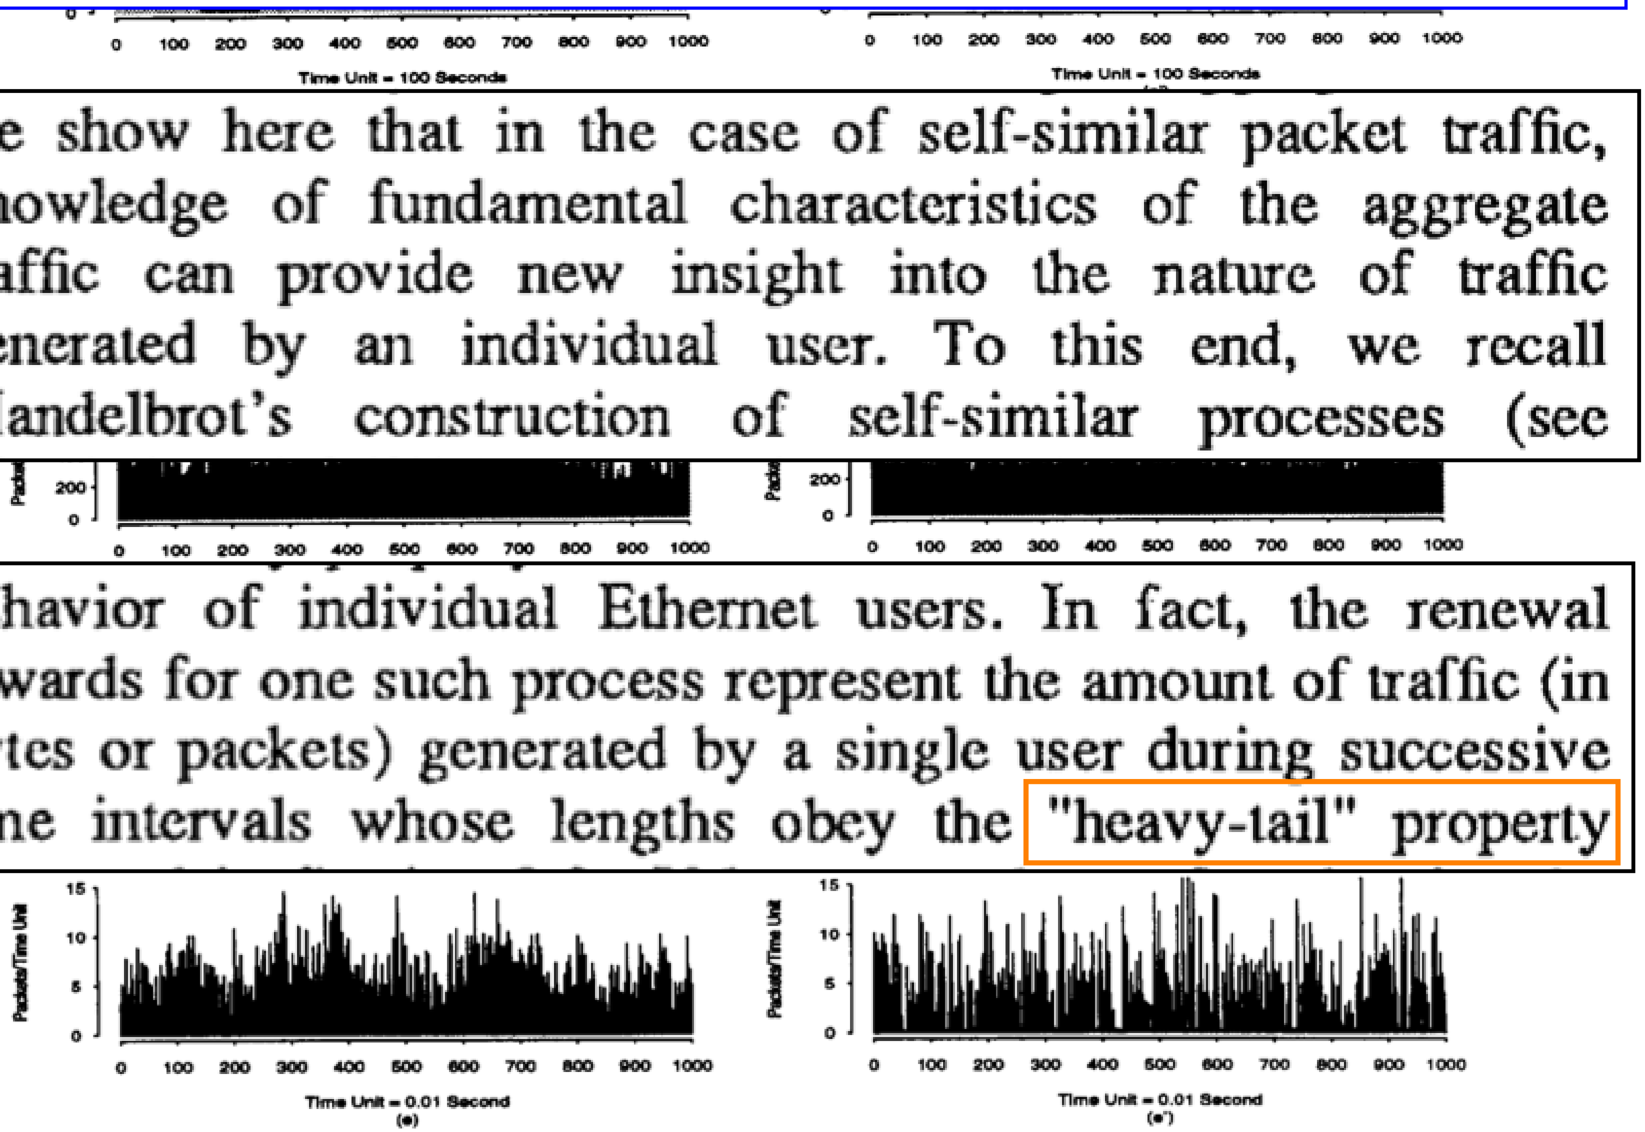

we show here that in the case of self-similar packet traffic, knowledge of fundamental characteristics of the aggregate traffic can provide new insight into the nature of traffic generated by an individual user. To this end, we recall Mandelbrot's construction of self-similar processes (see

behavior of individual Ethernet users. In fact, the renewal rewards for one such process represent the amount of traffic (in bytes or packets) generated by a single user during successive time intervals whose lengths obey the "heavy-tail" property

Fig. 4. Pictorial "proof" of self-similarity: Ethernet traffic (packets per time unit) on five different time scales (a)–(e). For comparison, synthetic traffic from an appropriately chosen compound Poisson model on the same five different time scales (a')–(e').

On the self-similar nature of Ethernet traffic (extended version)
WE Leland, MS Taqqu, W Willinger… - Networking, IEEE/ACM …, 1994 - ieeexplore.ieee.org
... The main objective of this paper is to establish in a statistically rigorous manner the **self-similarity** characteristic of the very **high** quality, **high** time-resolution Ethemet LAN Manuscript received July 1, 1993; revised January 15, 1994; approved by IEEE/ACM TRANSACTIONS ON ...
Cited by 4759 - Related articles - All 69 versions

Their work also proposed an explanation for self-similarity that is **predictive**: Network activity should be marked by sizes/durations that are *heavy-tailed*.

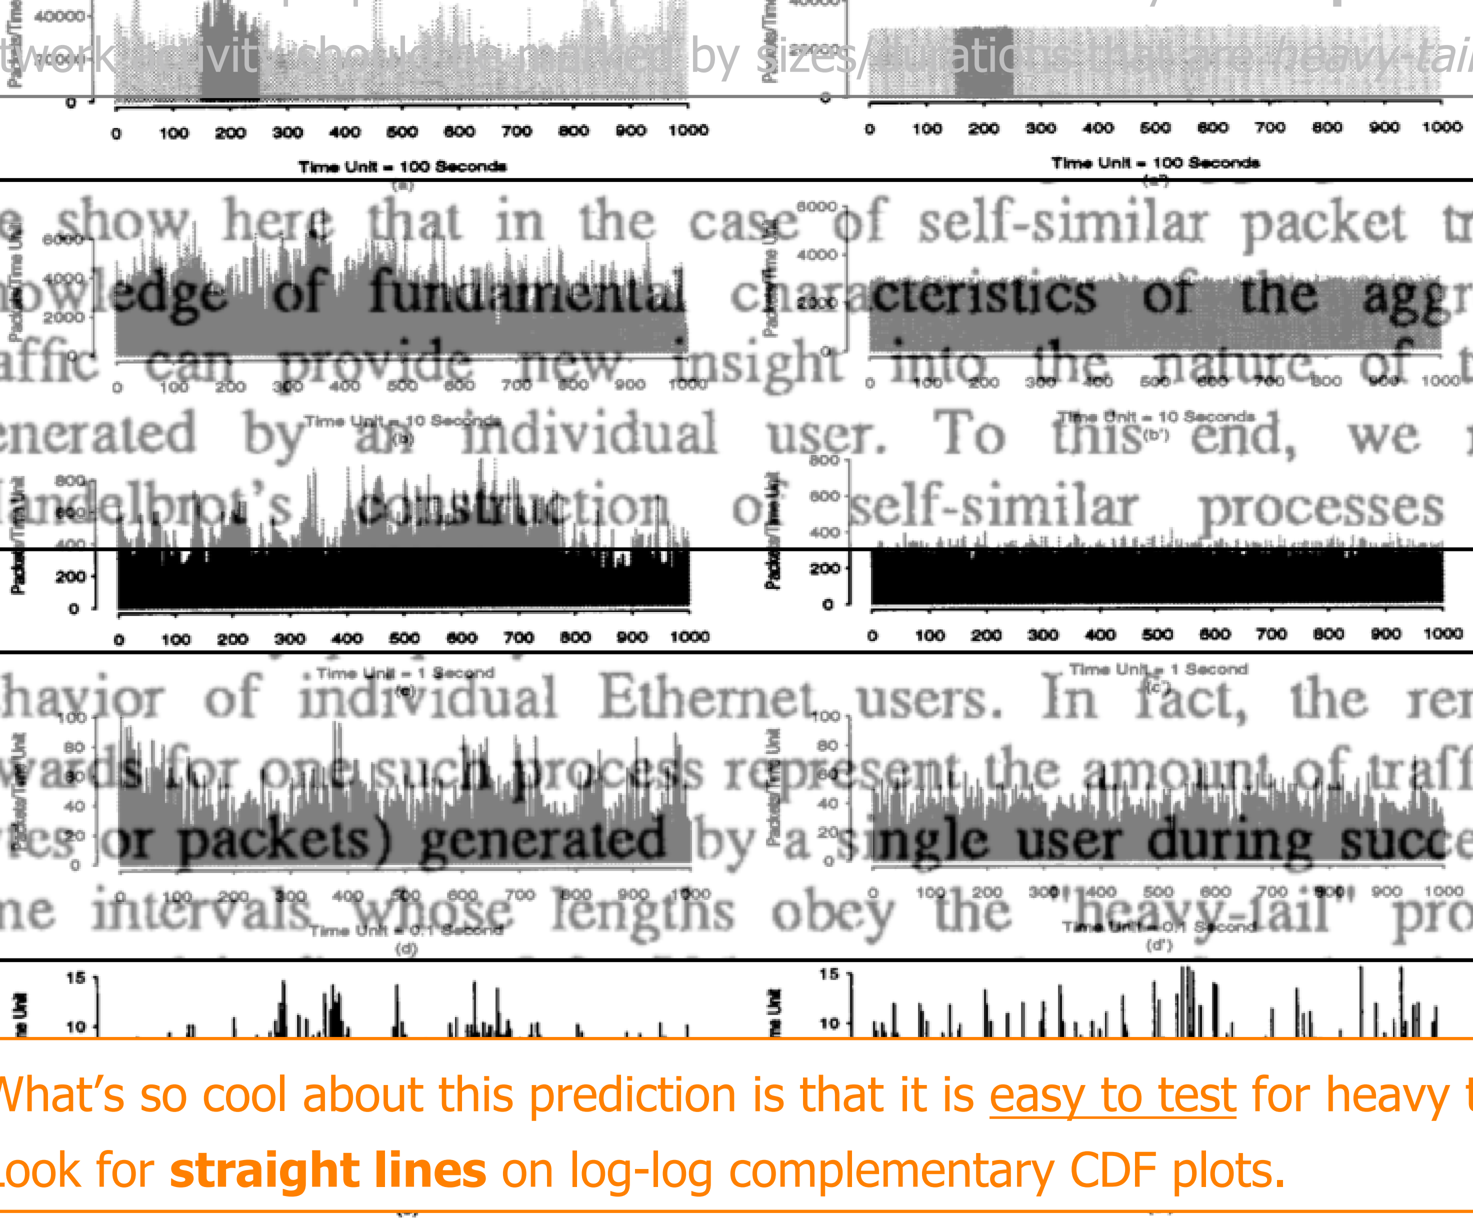

we show here that in the case of self-similar packet traffic, knowledge of fundamental characteristics of the aggregate traffic can provide new insight into the nature of traffic generated by an individual user. To this end, we recall Mandelbrot's construction of self-similar processes (see

behavior of individual Ethernet users. In fact, the renewal rewards for one such process represent the amount of traffic (in bytes or packets) generated by a single user during successive time intervals whose lengths obey the "heavy-tail" property

What's so cool about this prediction is that it is easy to test for heavy tails: Look for **straight lines** on log-log complementary CDF plots.

Fig. 4. Pictorial ''proof'' of self-similarity: Ethernet traffic (packets per time unit) on five different time scales (a)–(e). For comparison, synthetic traffic from an appropriately chosen compound Poisson model on the same five different time scales (a')–(e').

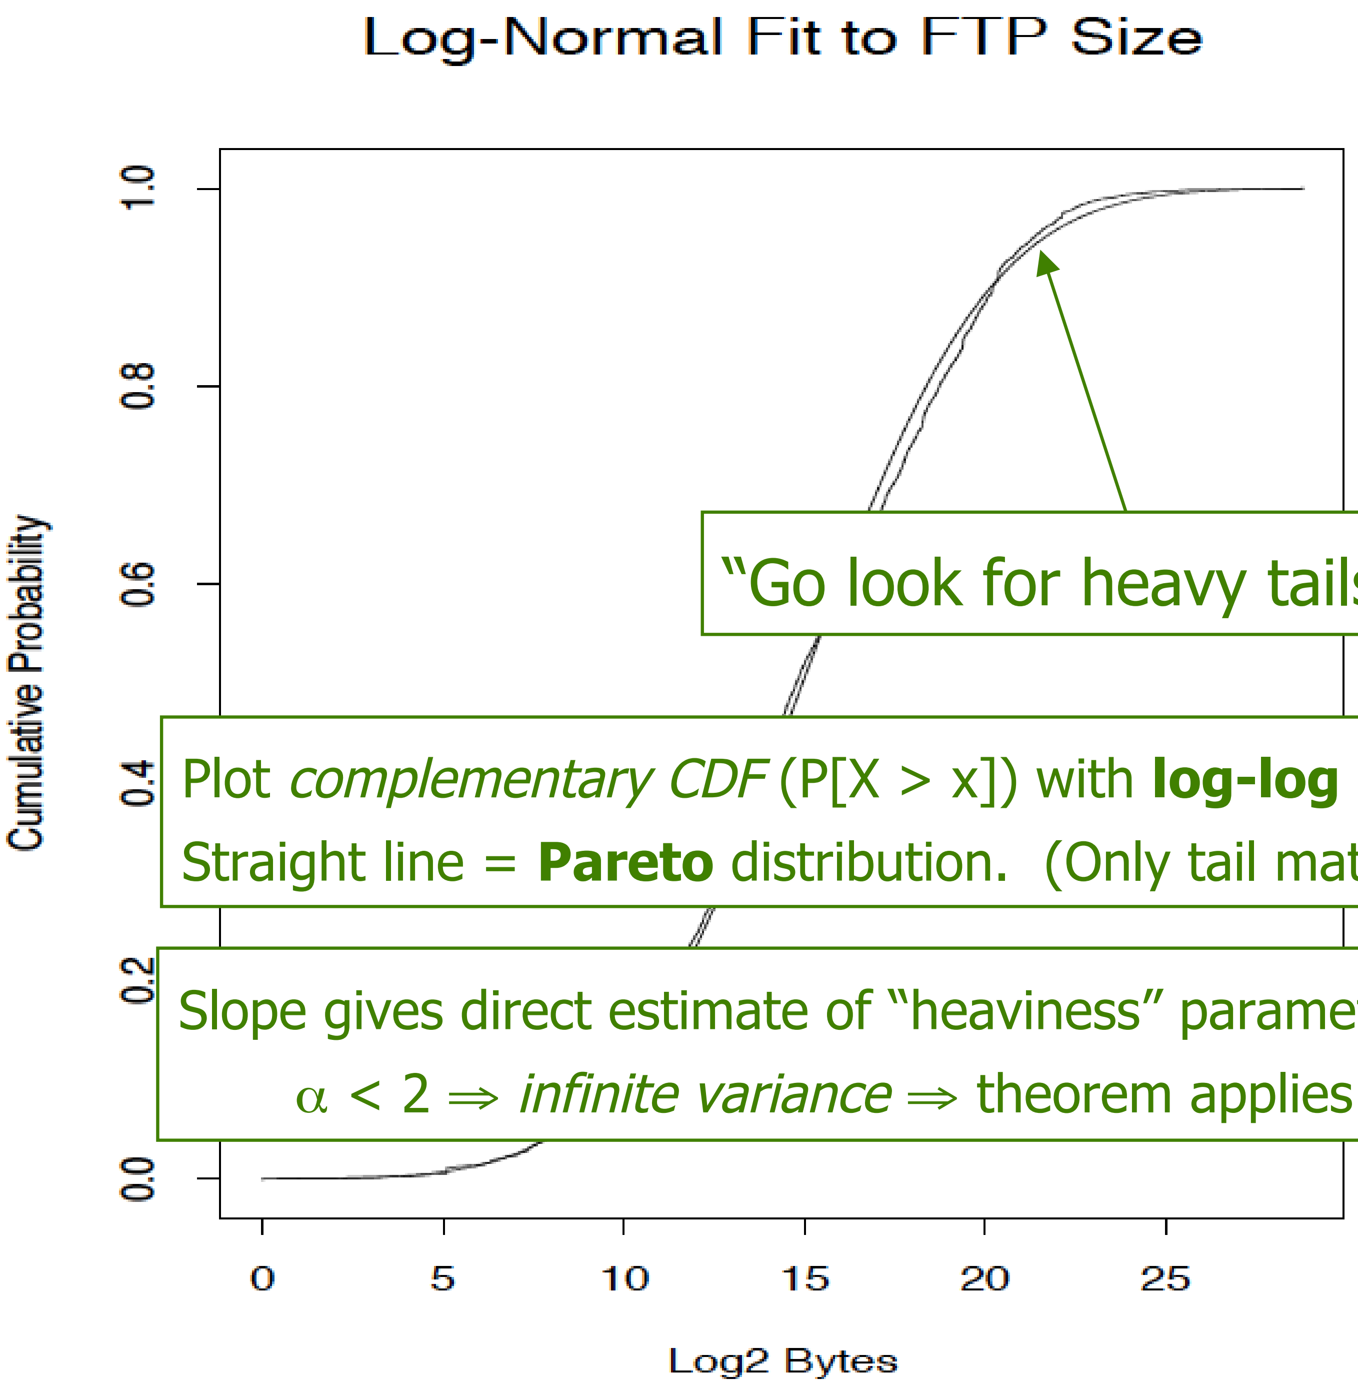

"Go look for heavy tails"

Plot *complementary CDF* (P[X > x]) with **log-log** scaling.
Straight line = **Pareto** distribution. (Only tail matters.)

Slope gives direct estimate of "heaviness" parameter **α.**
$\alpha < 2 \Rightarrow$ *infinite variance* $\Rightarrow$ theorem applies!

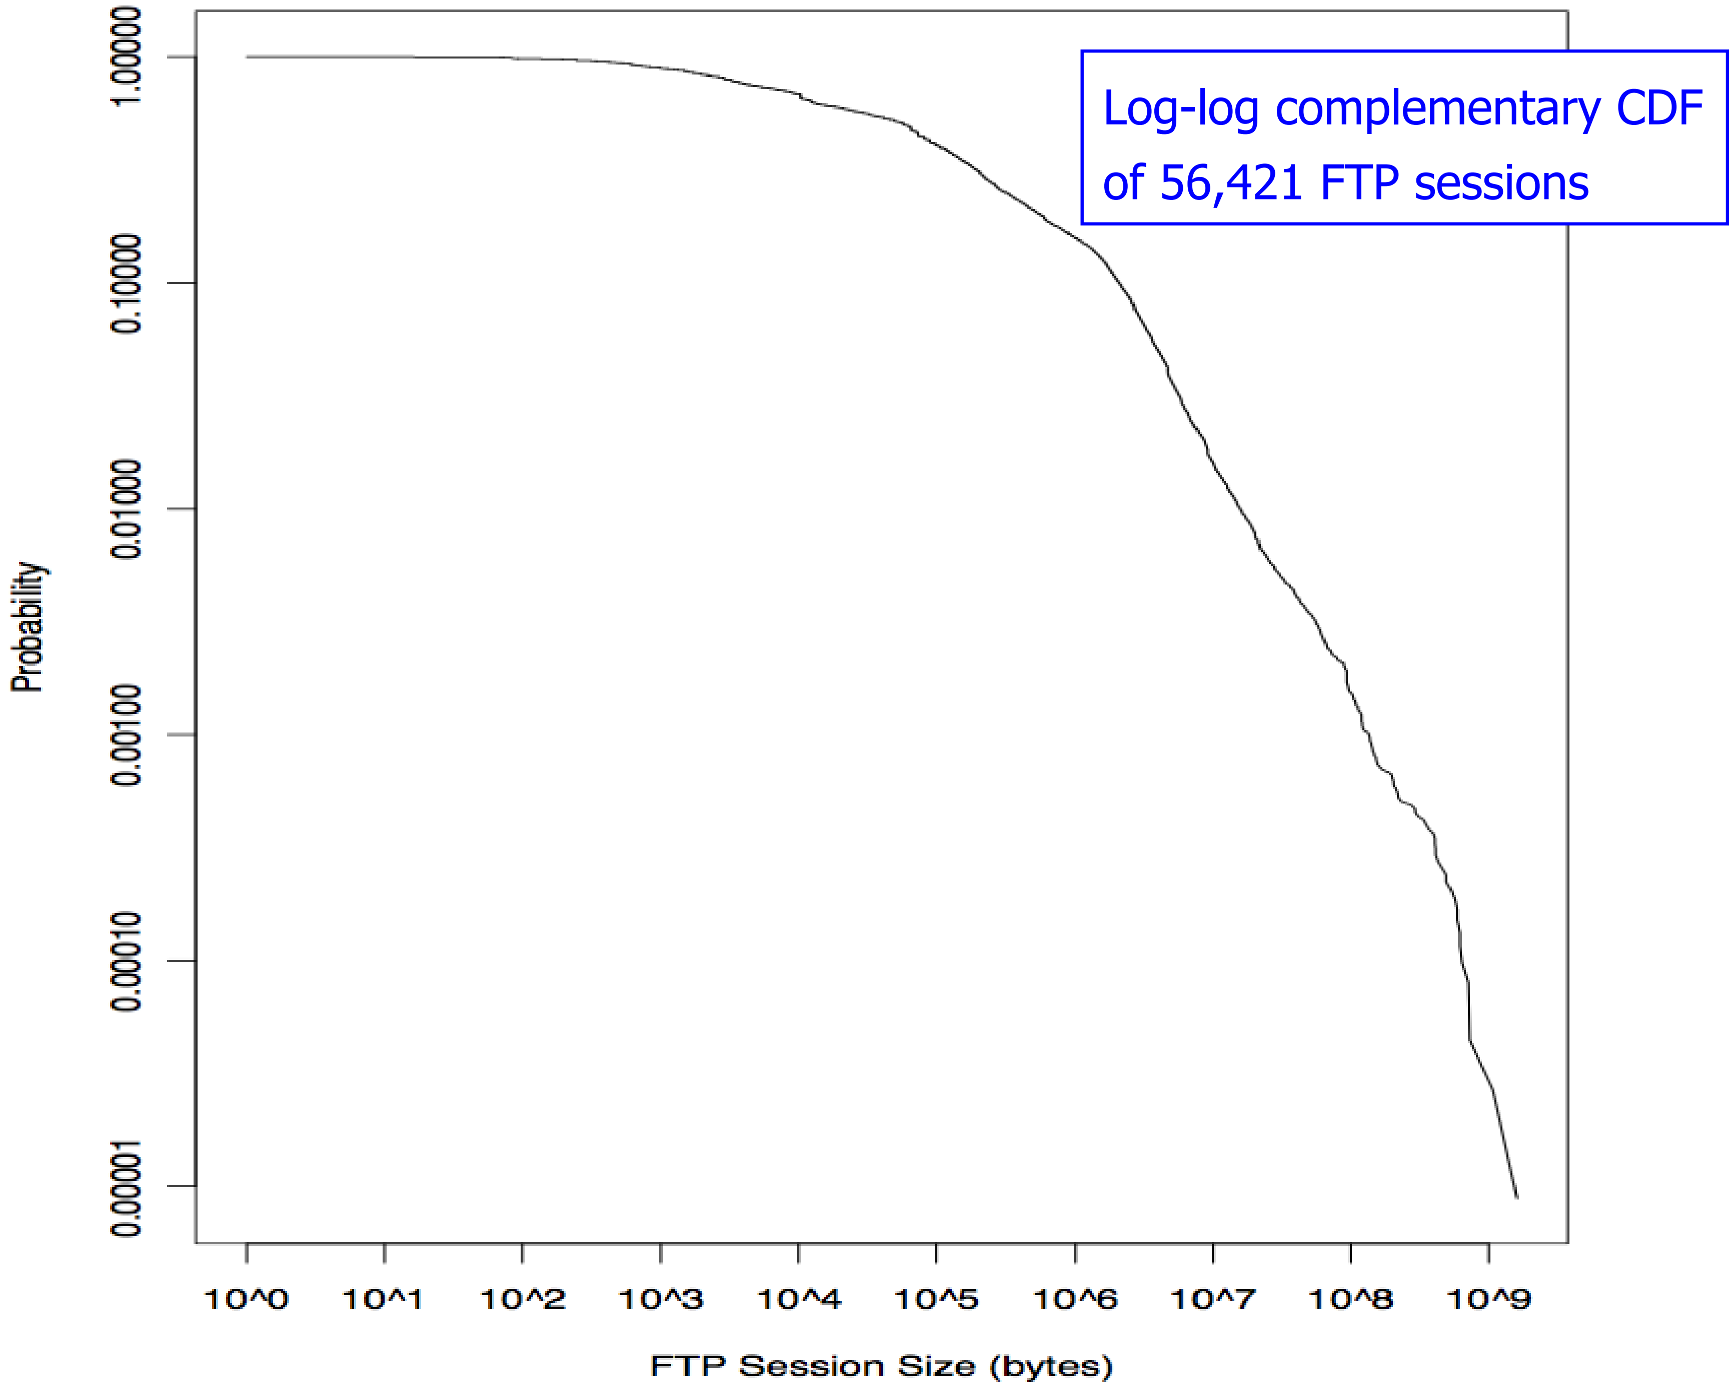

Log-log complementary CDF of 56,421 FTP sessions

Same, upper 10% tail

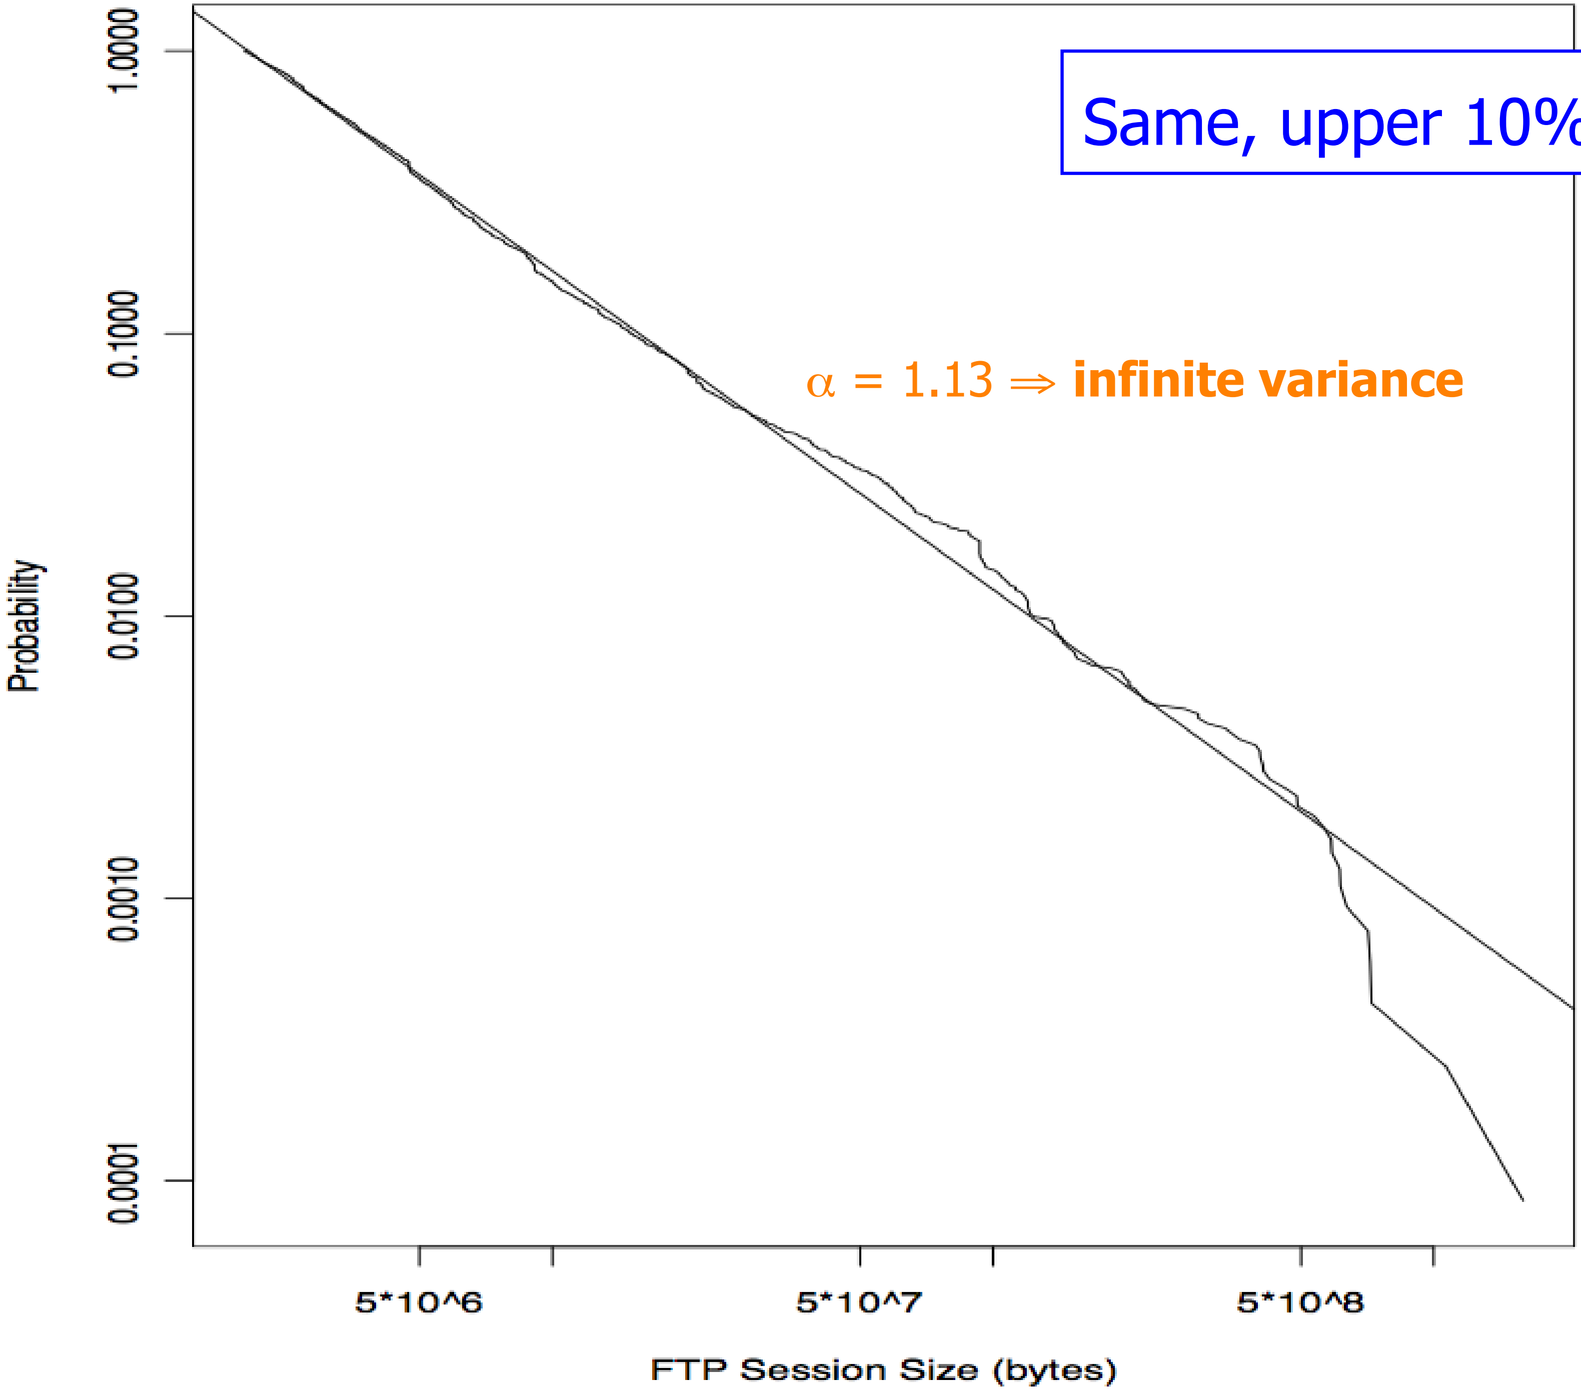

$\alpha = 1.13 \Rightarrow$ **infinite variance**

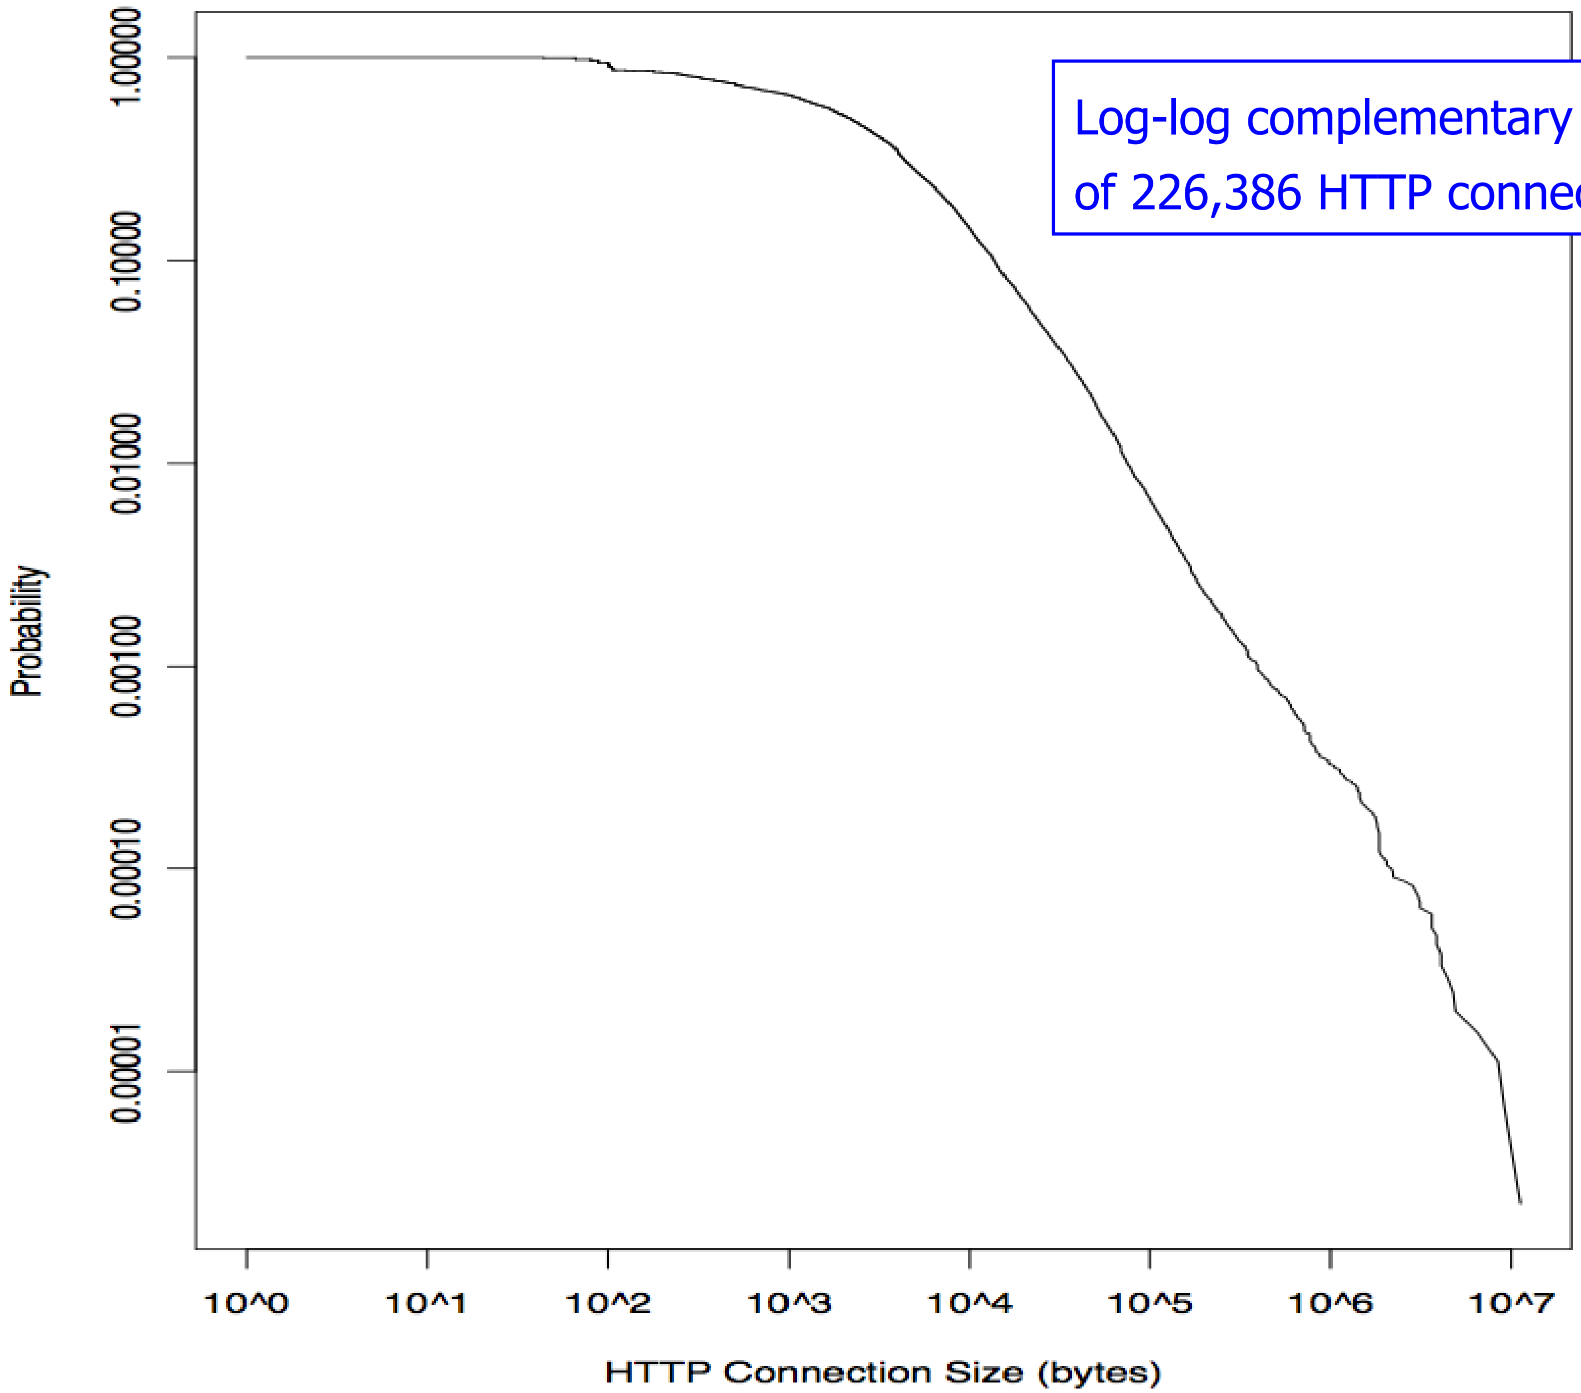

Log-log complementary CDF of 226,386 HTTP connections

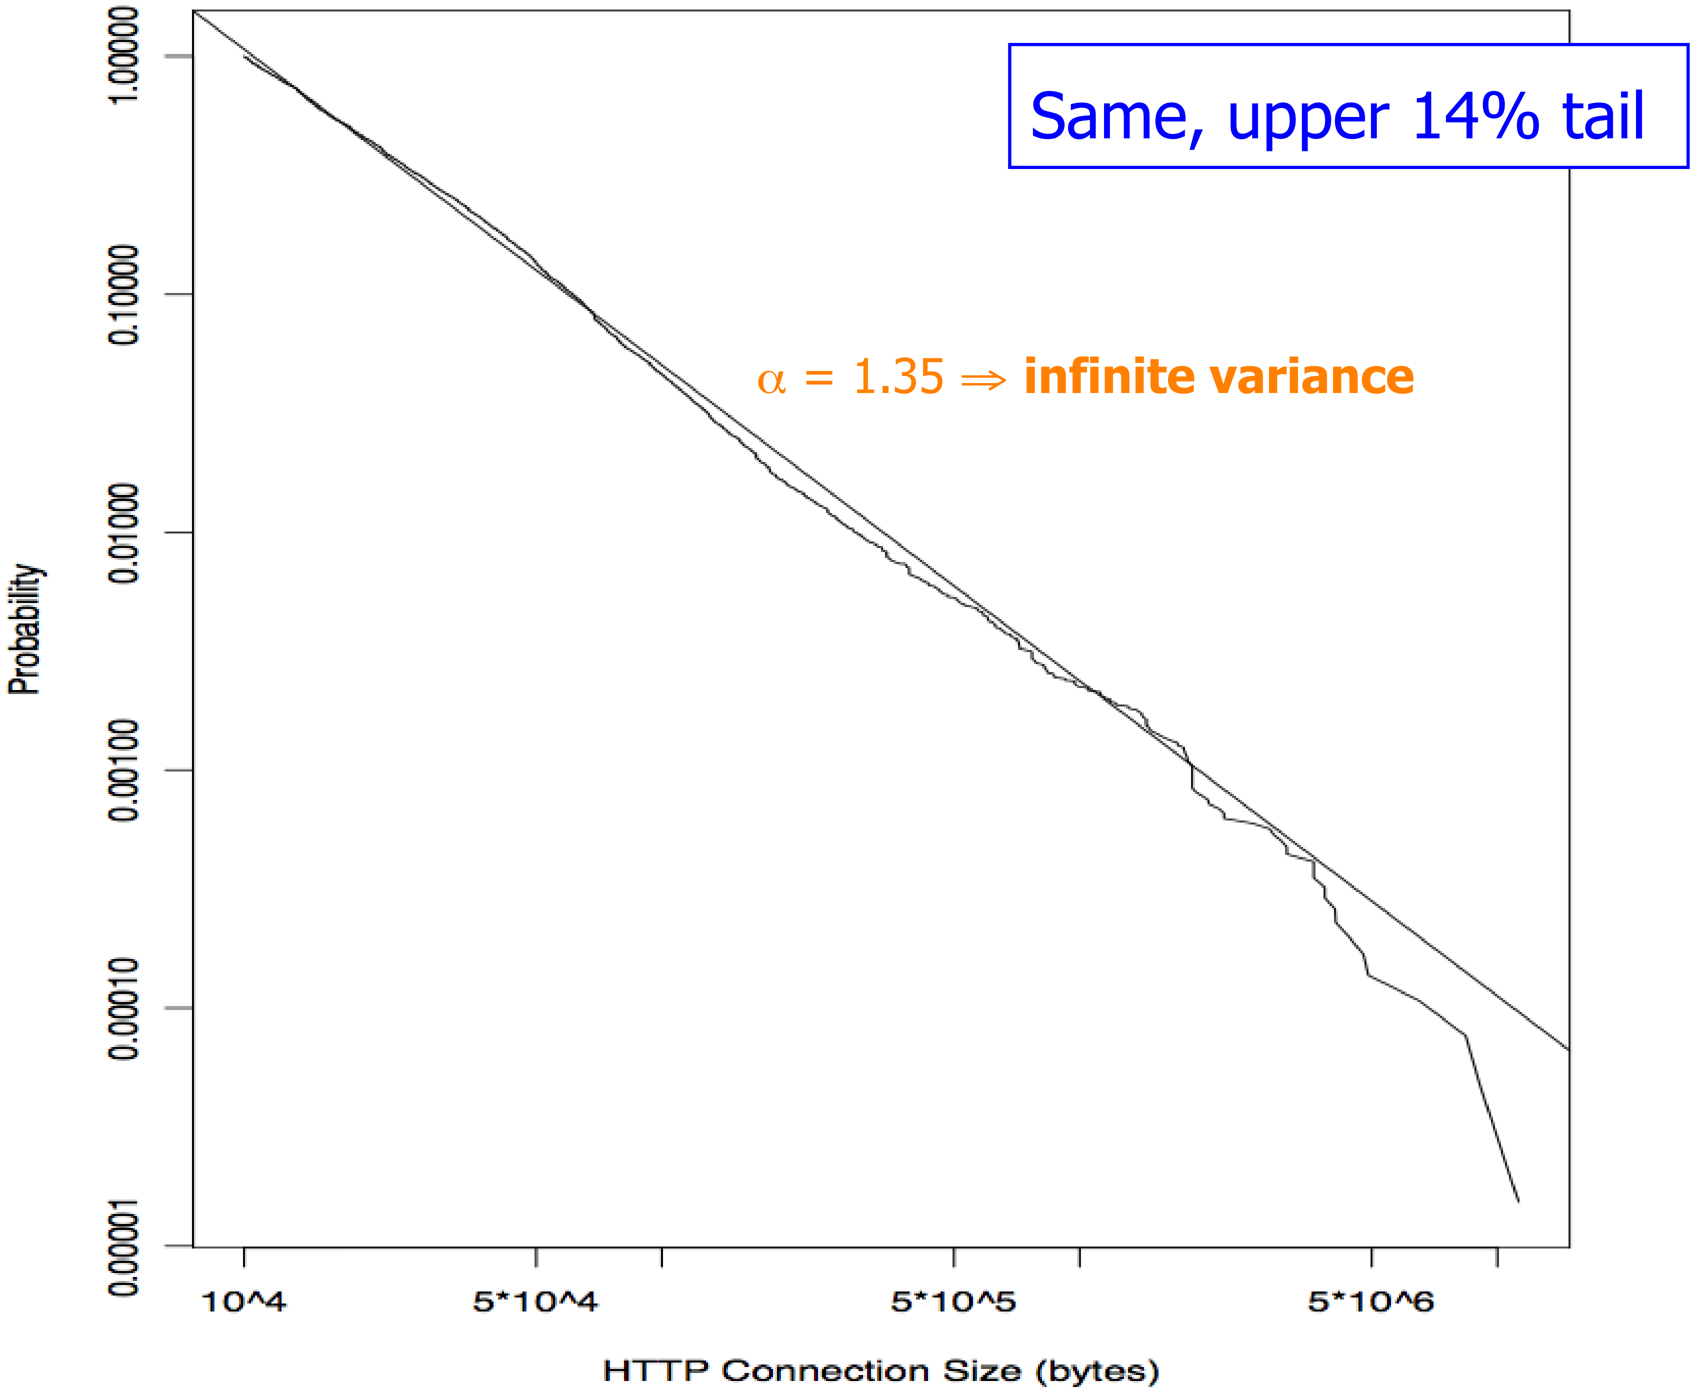

Same, upper 14% tail

$\alpha = 1.35 \Rightarrow$ **infinite variance**

The Danger of Mental Models

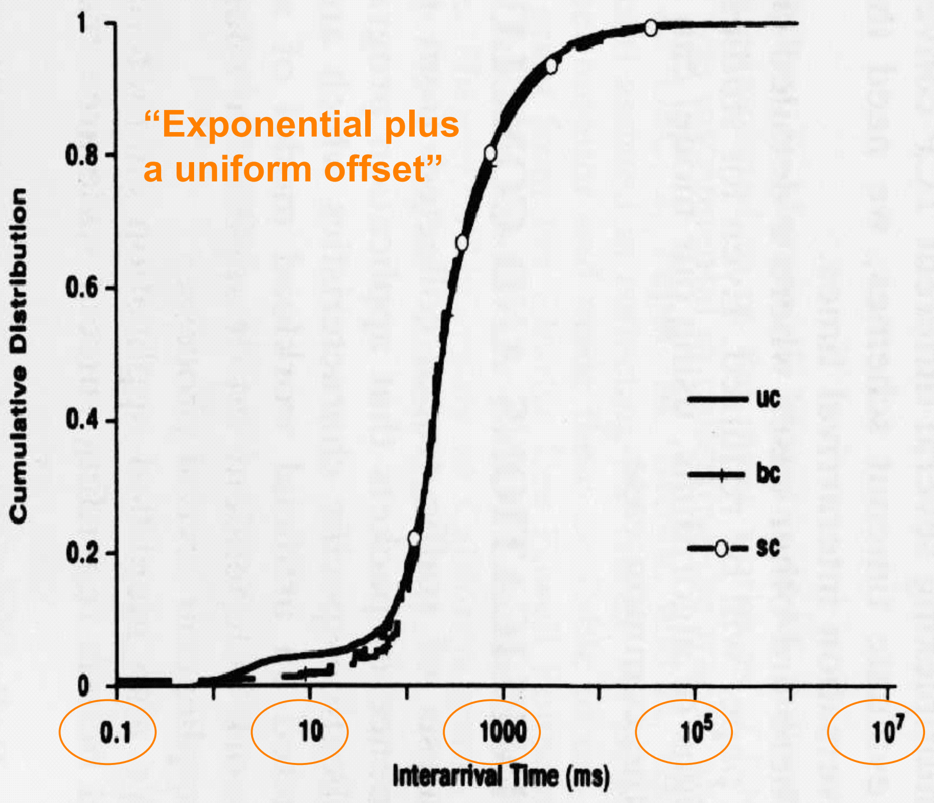

"Exponential plus a uniform offset"

Figure 19. Distribution of TELNET packet interarrivals

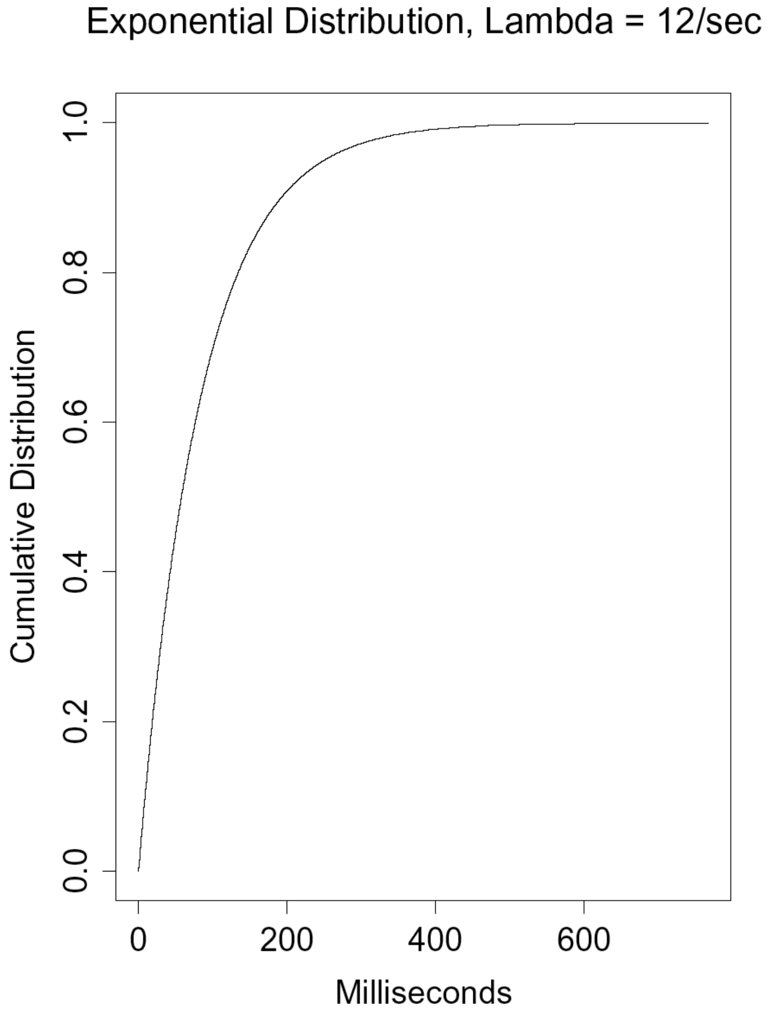
Exponential Distribution, Lambda = 12/sec

A uniform plus exponential distribution best models interarrival times of packets belonging to interactive applications.

Table 3: Selected Observations.

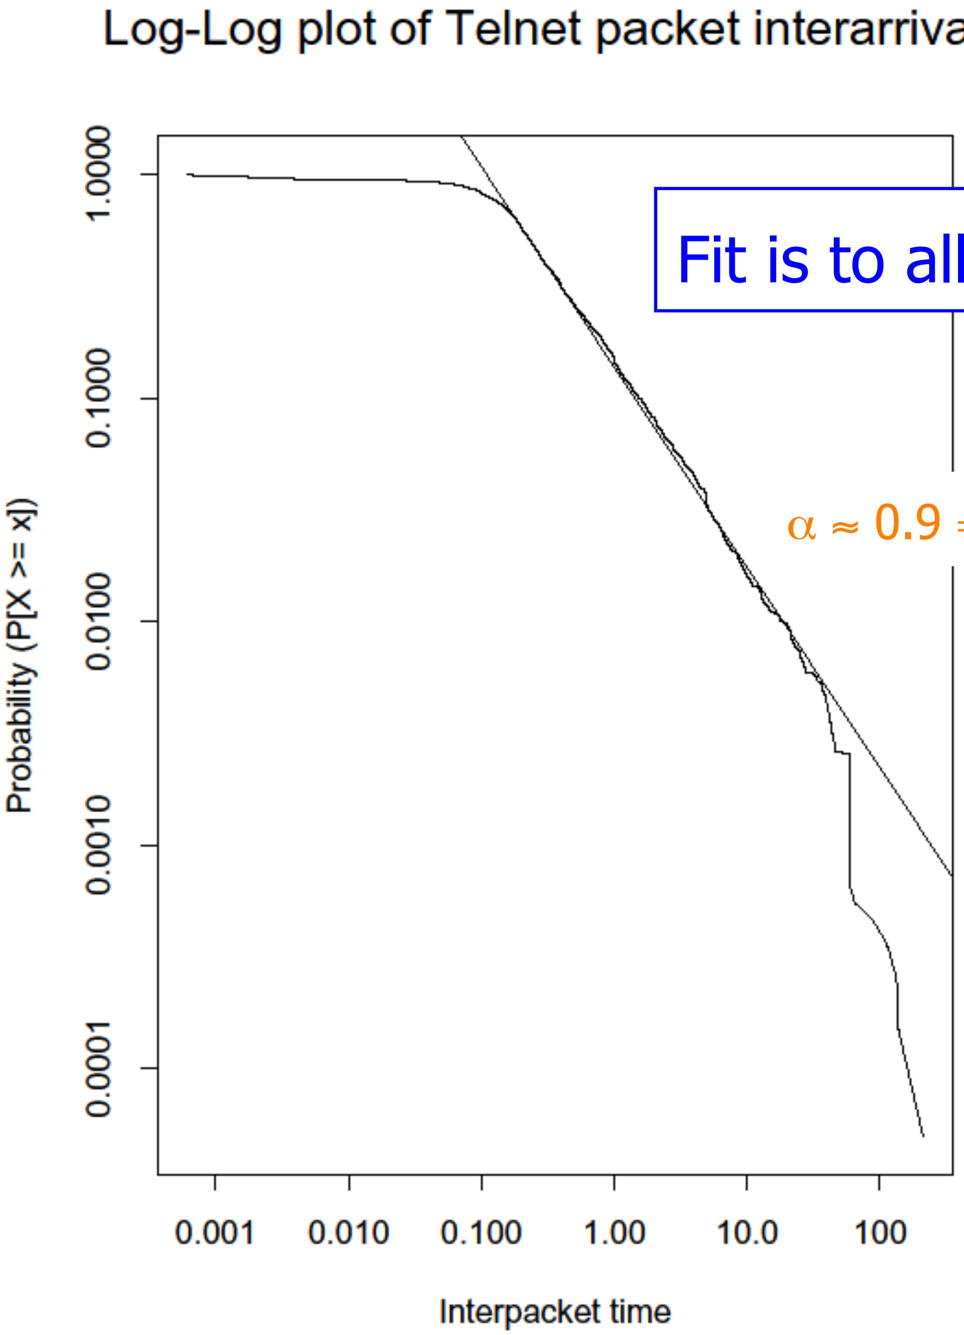

Fit is to all but lowest 25%!

$\alpha \approx 0.9 \Rightarrow$ **infinite mean**

USENET Bulletin Board Traffic Volume

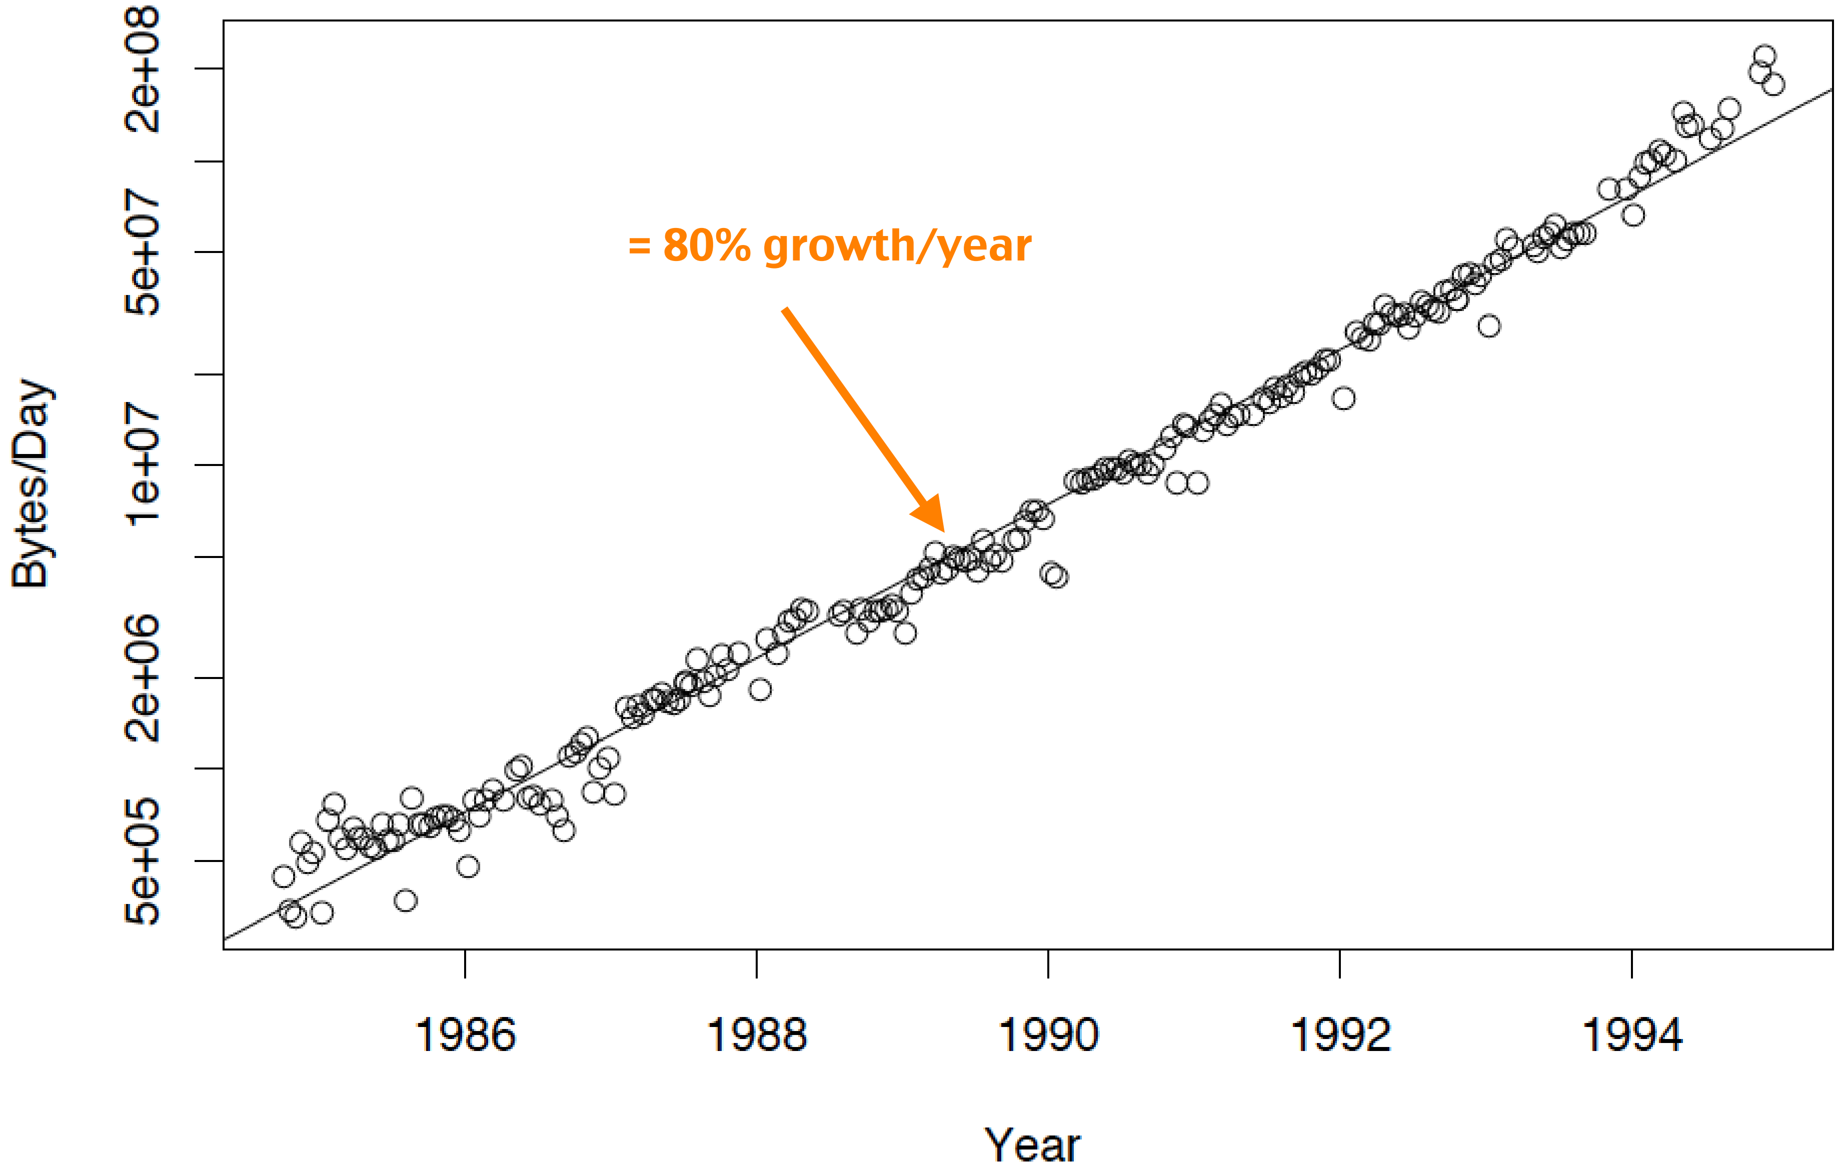

= 80% growth/year

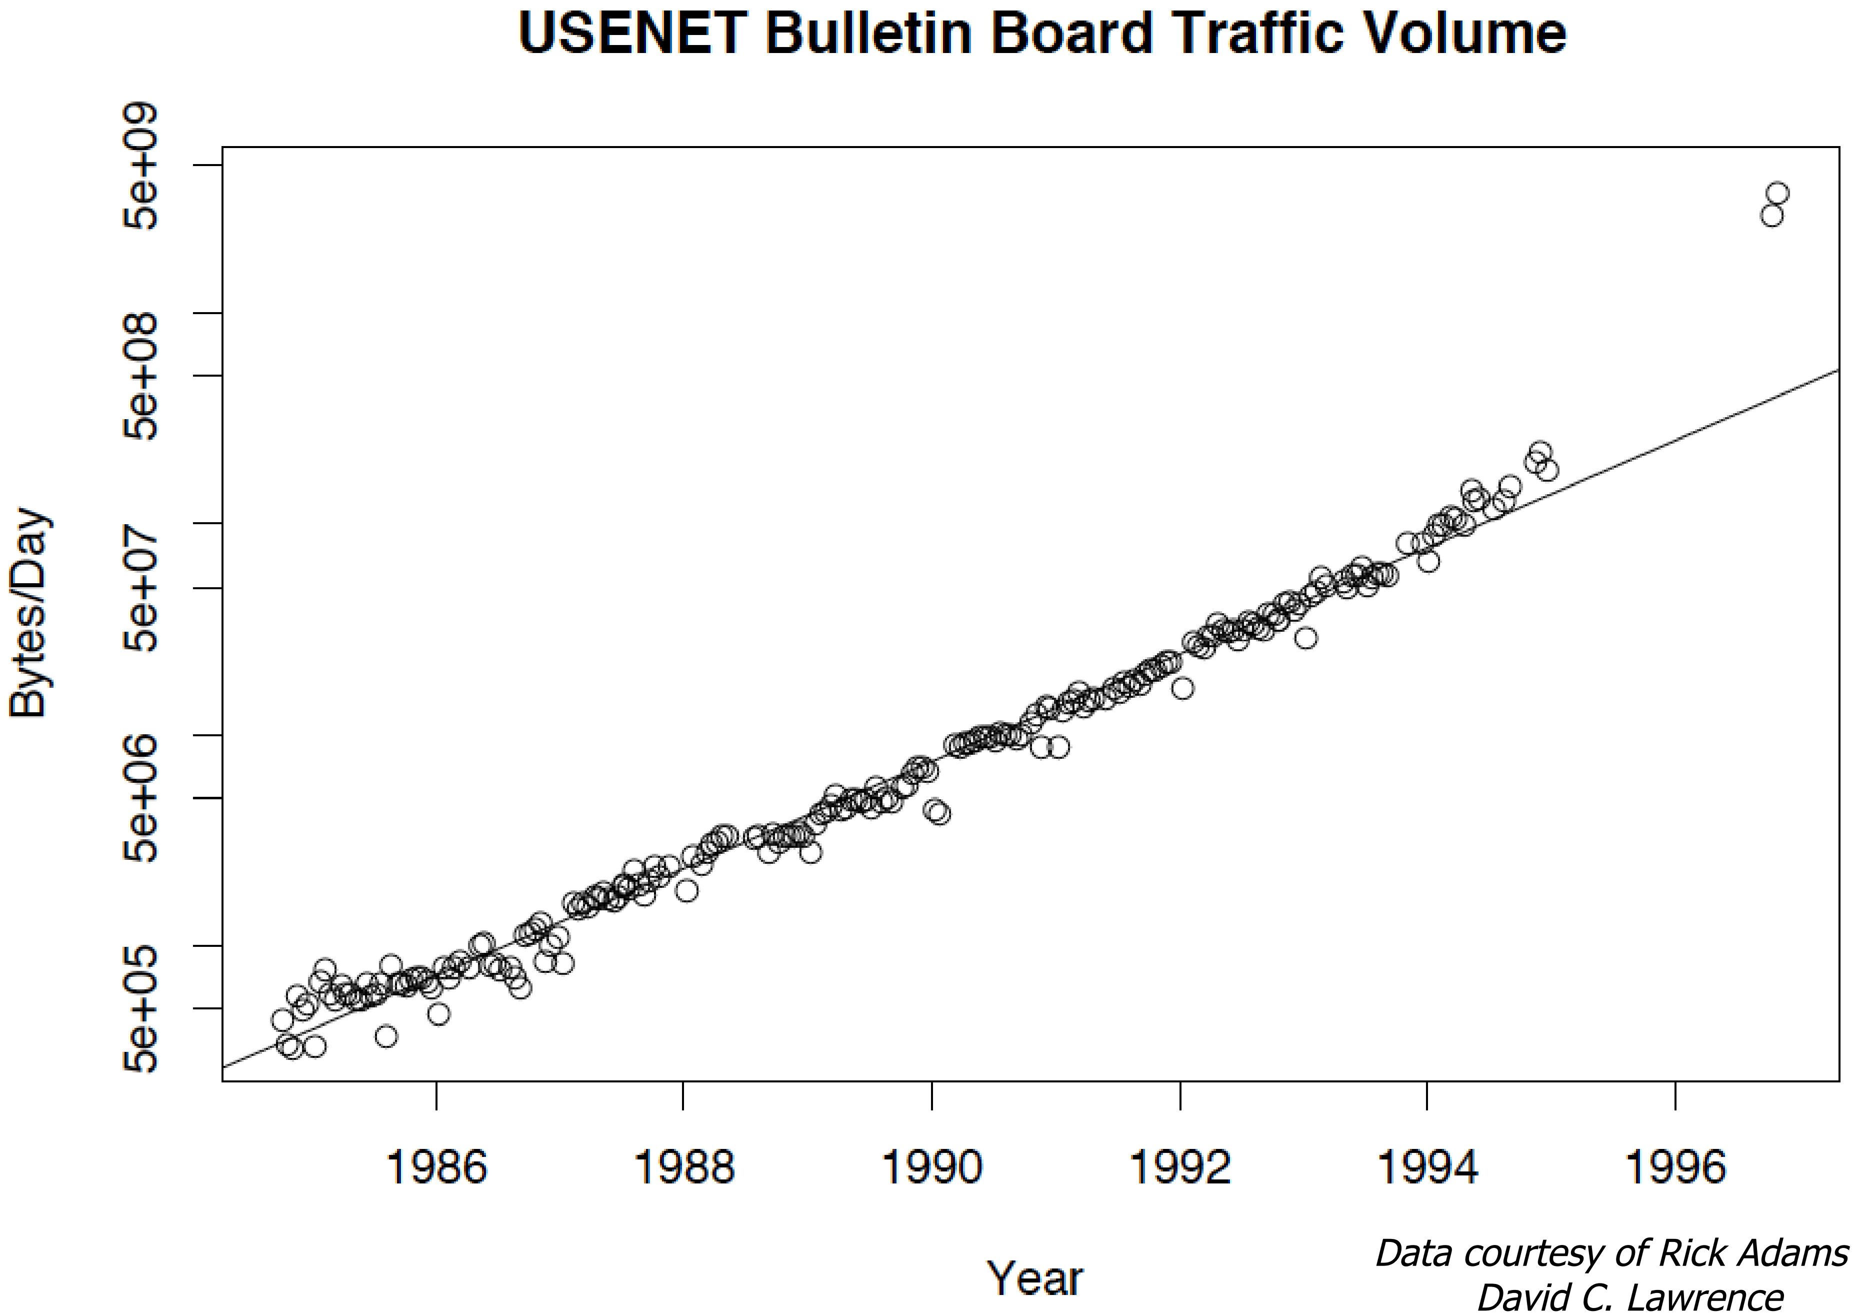

Data courtesy of Rick Adams &
David C. Lawrence

USENET Bulletin Board Traffic Volume

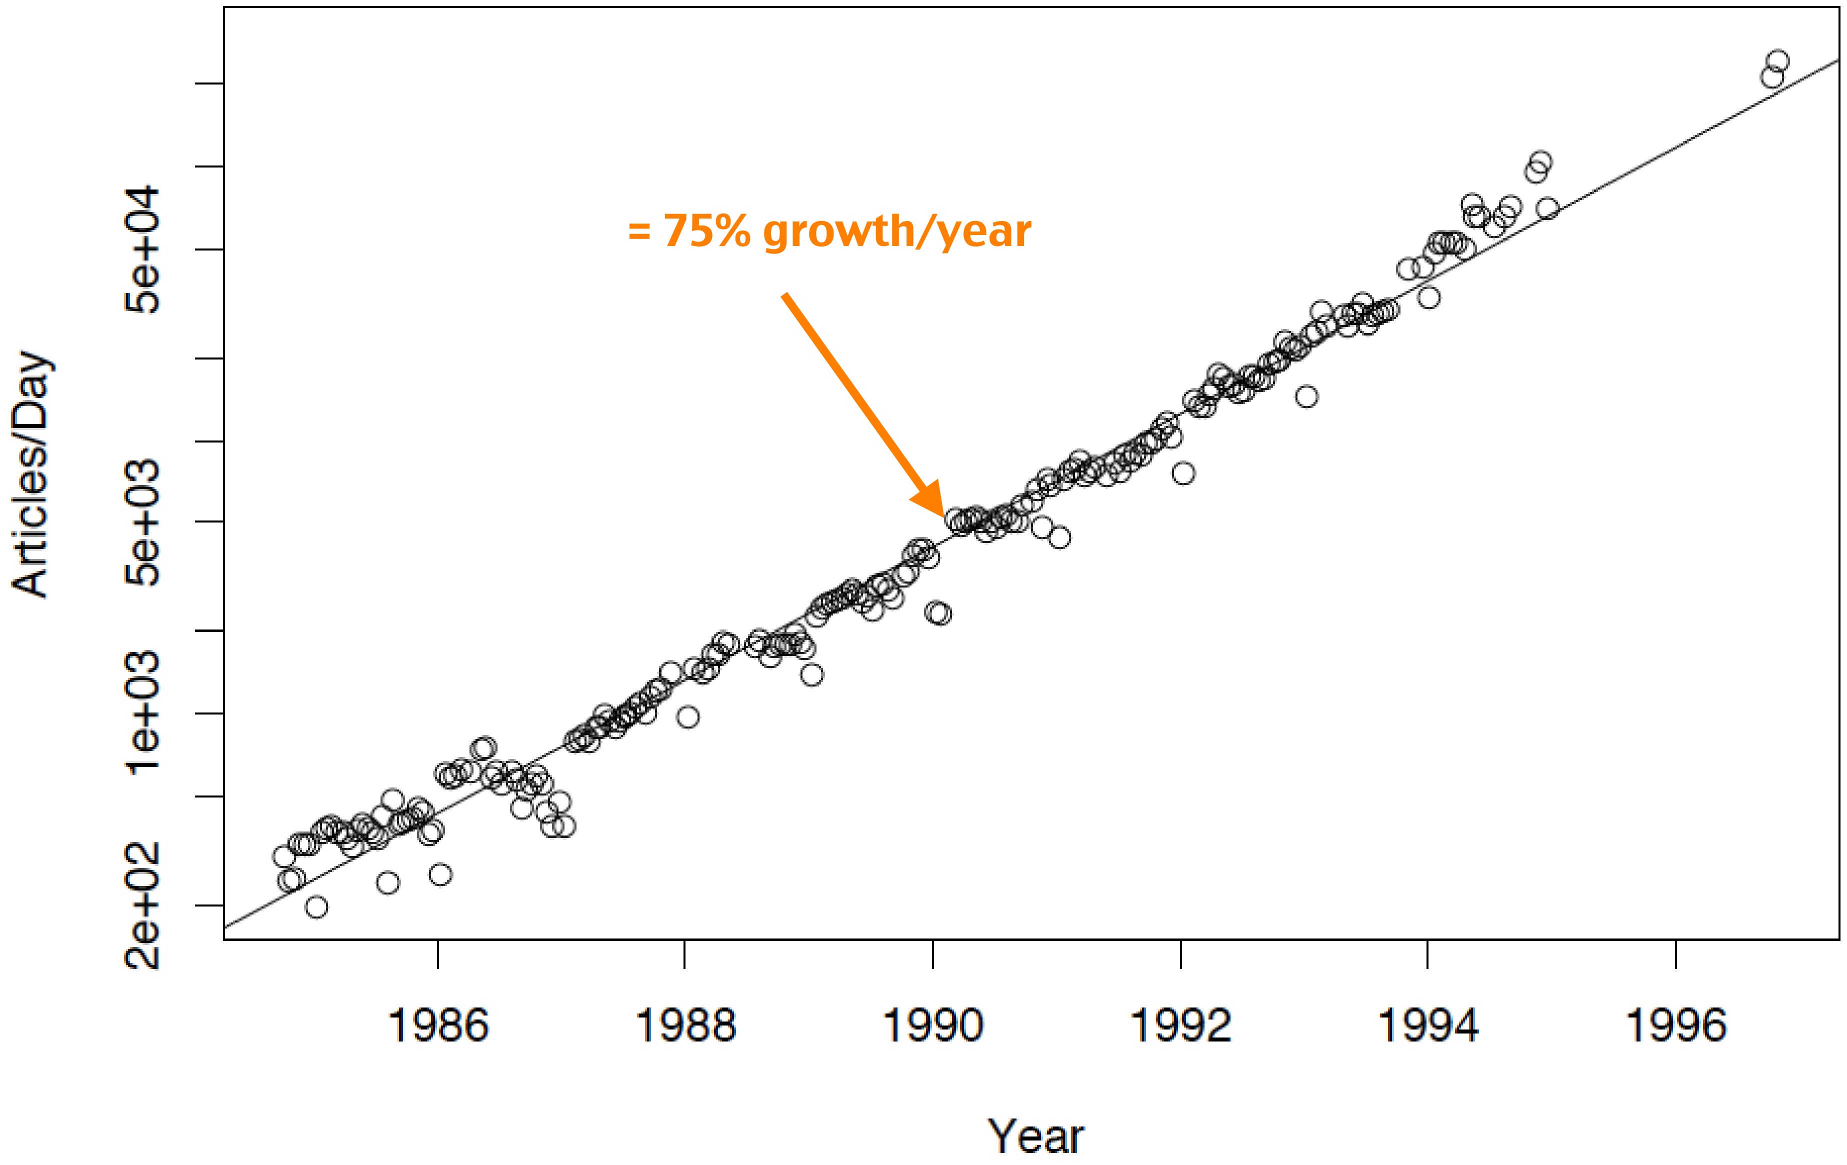

= 75% growth/year

USENET Bulletin Board Traffic Volume

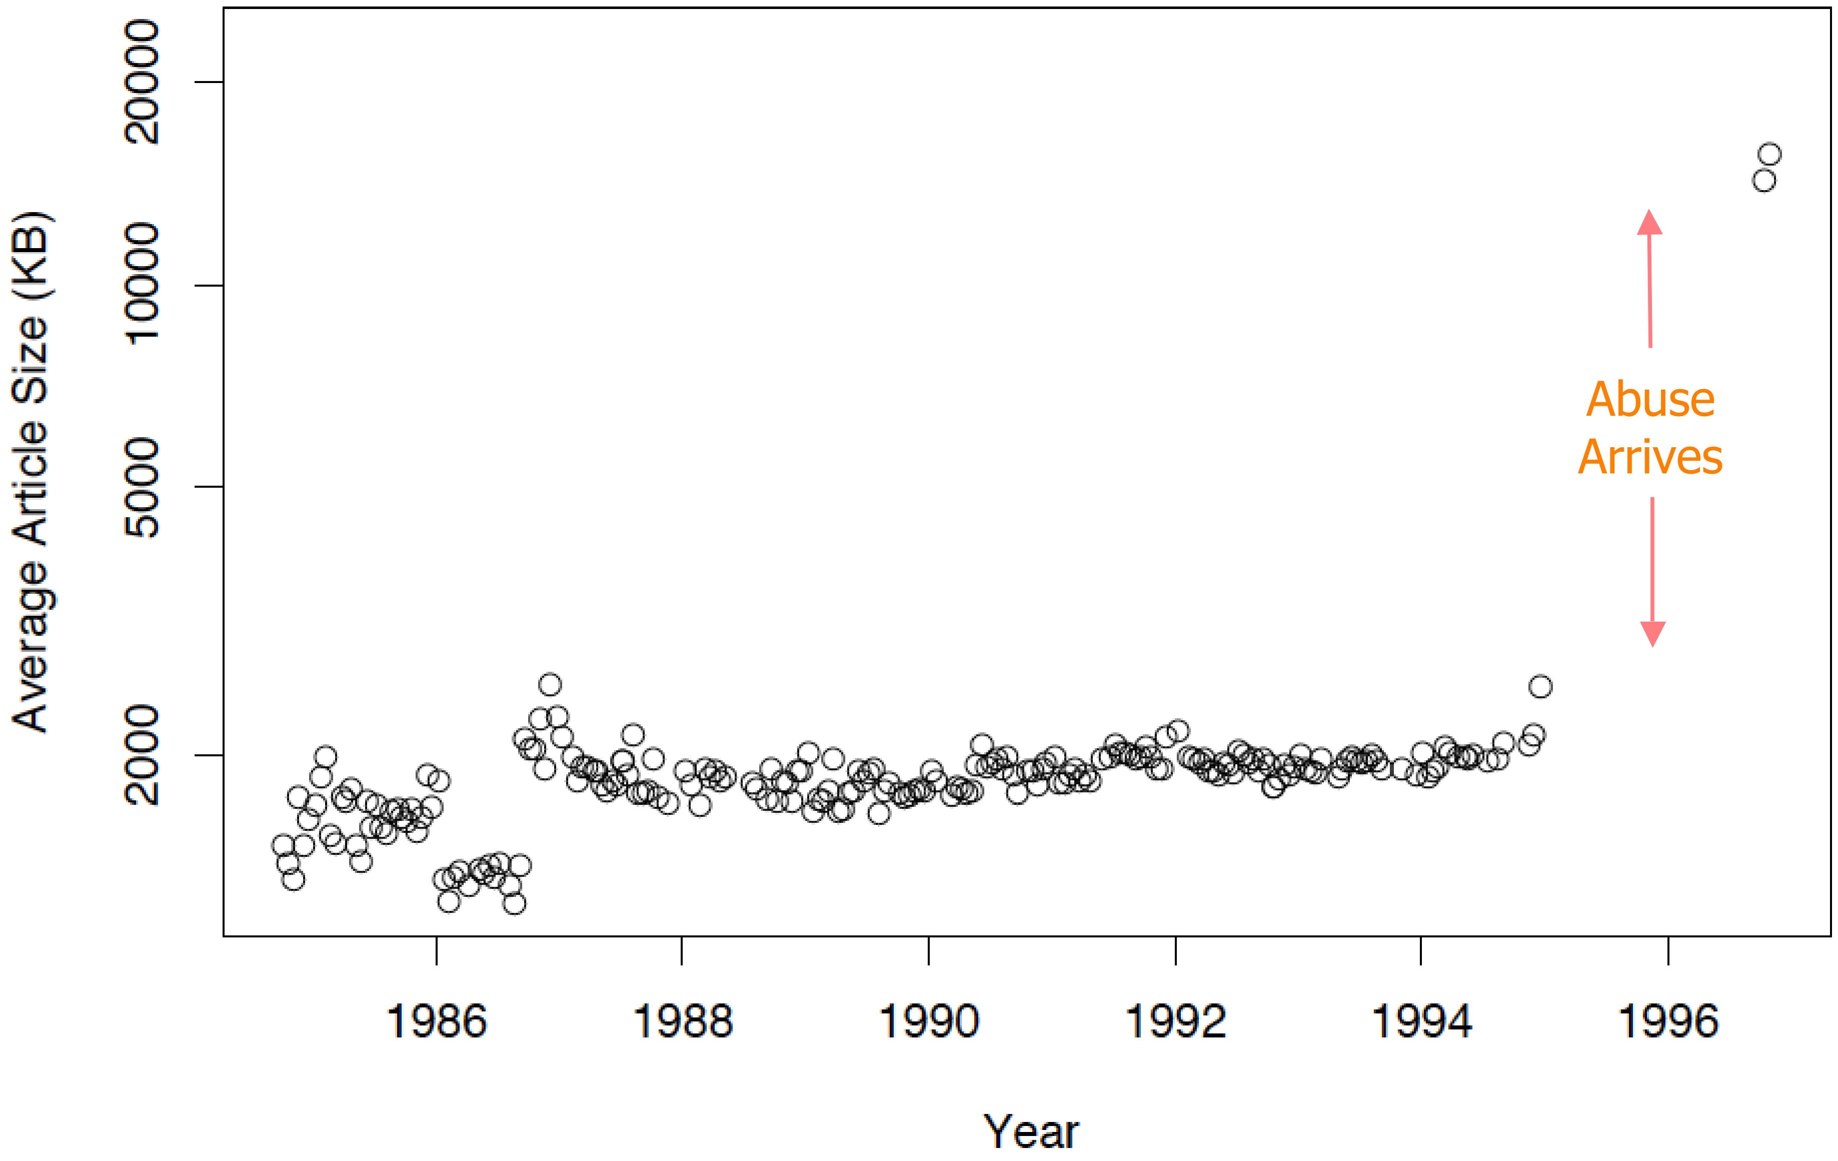

Part II:

Measuring Malice & Crooks

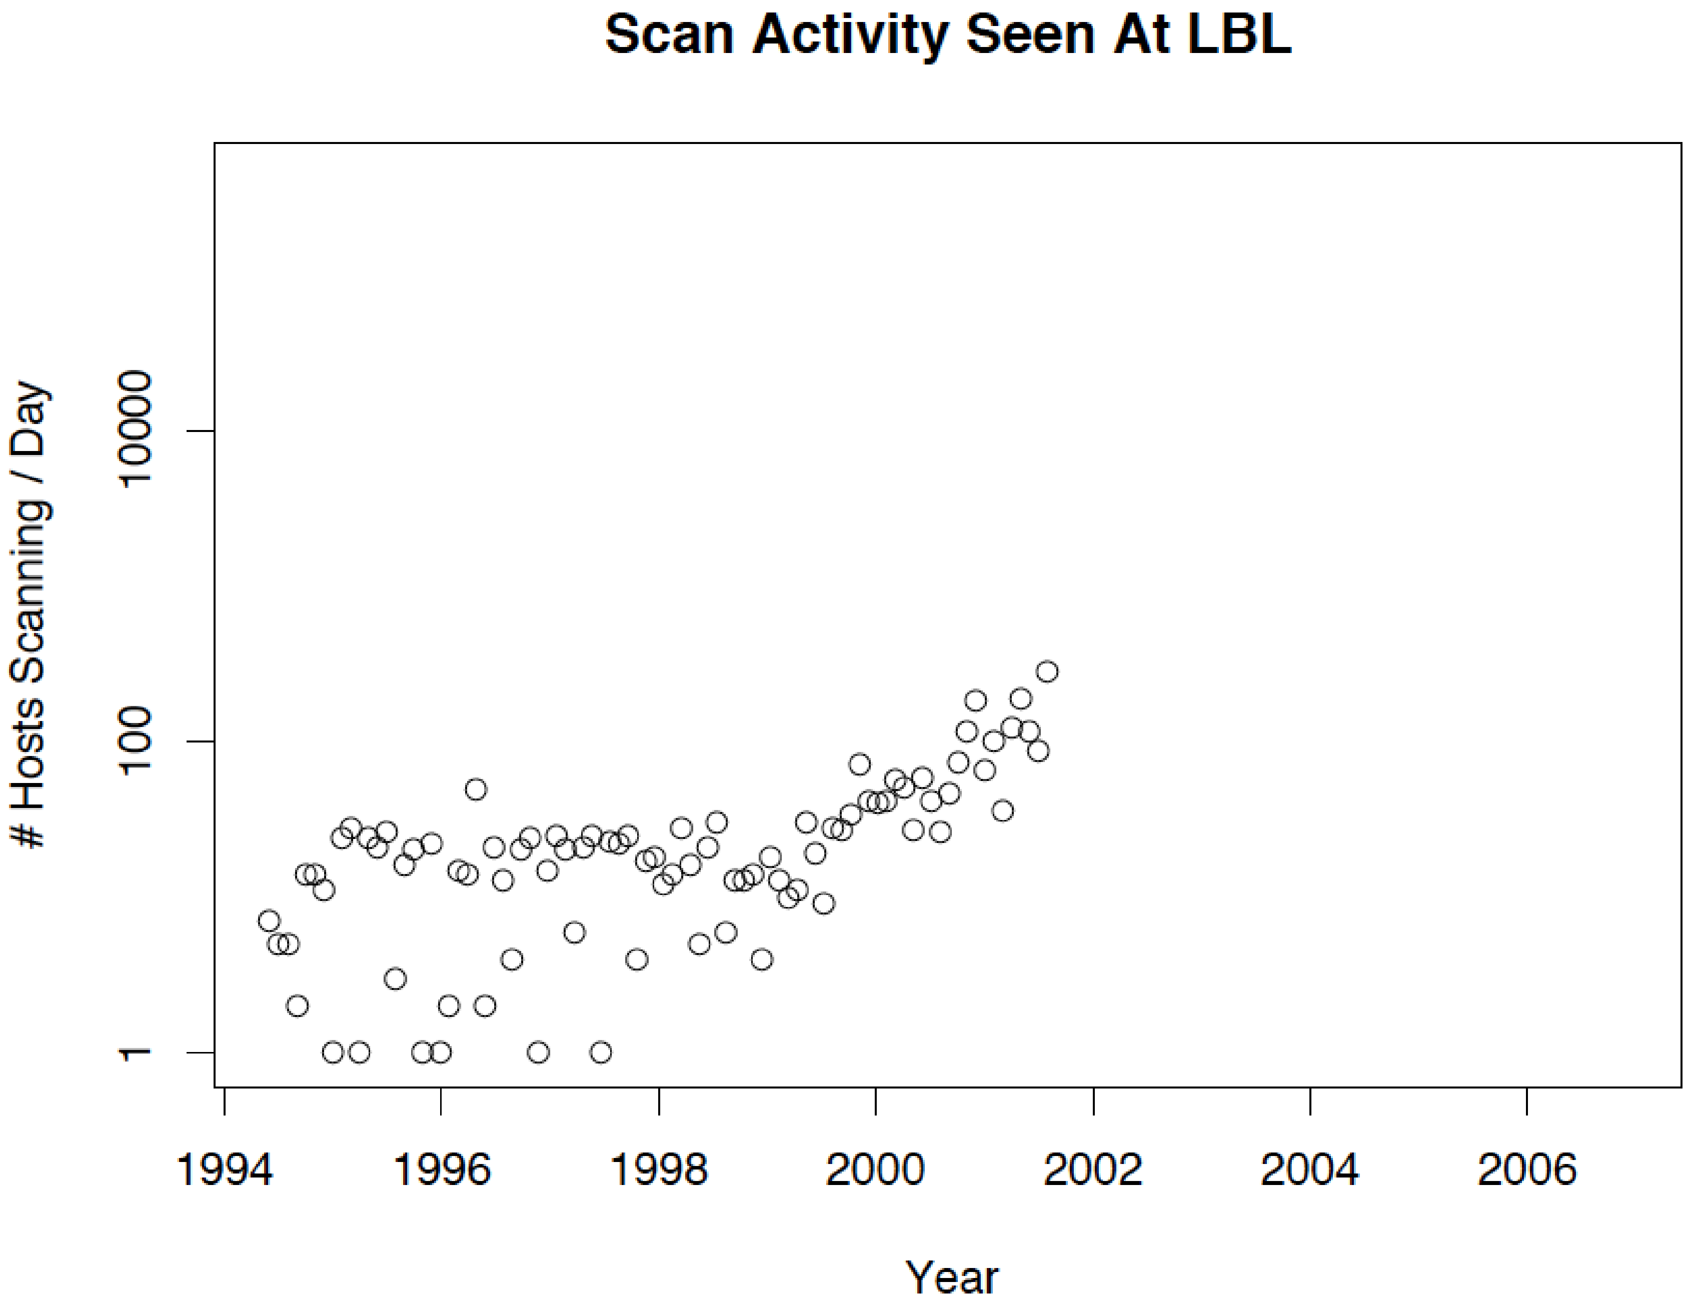
Scan Activity Seen At LBL
# Hosts Scanning / Day
1
100
10000
Year
1994
1996
1998
2000
2002
2004
2006

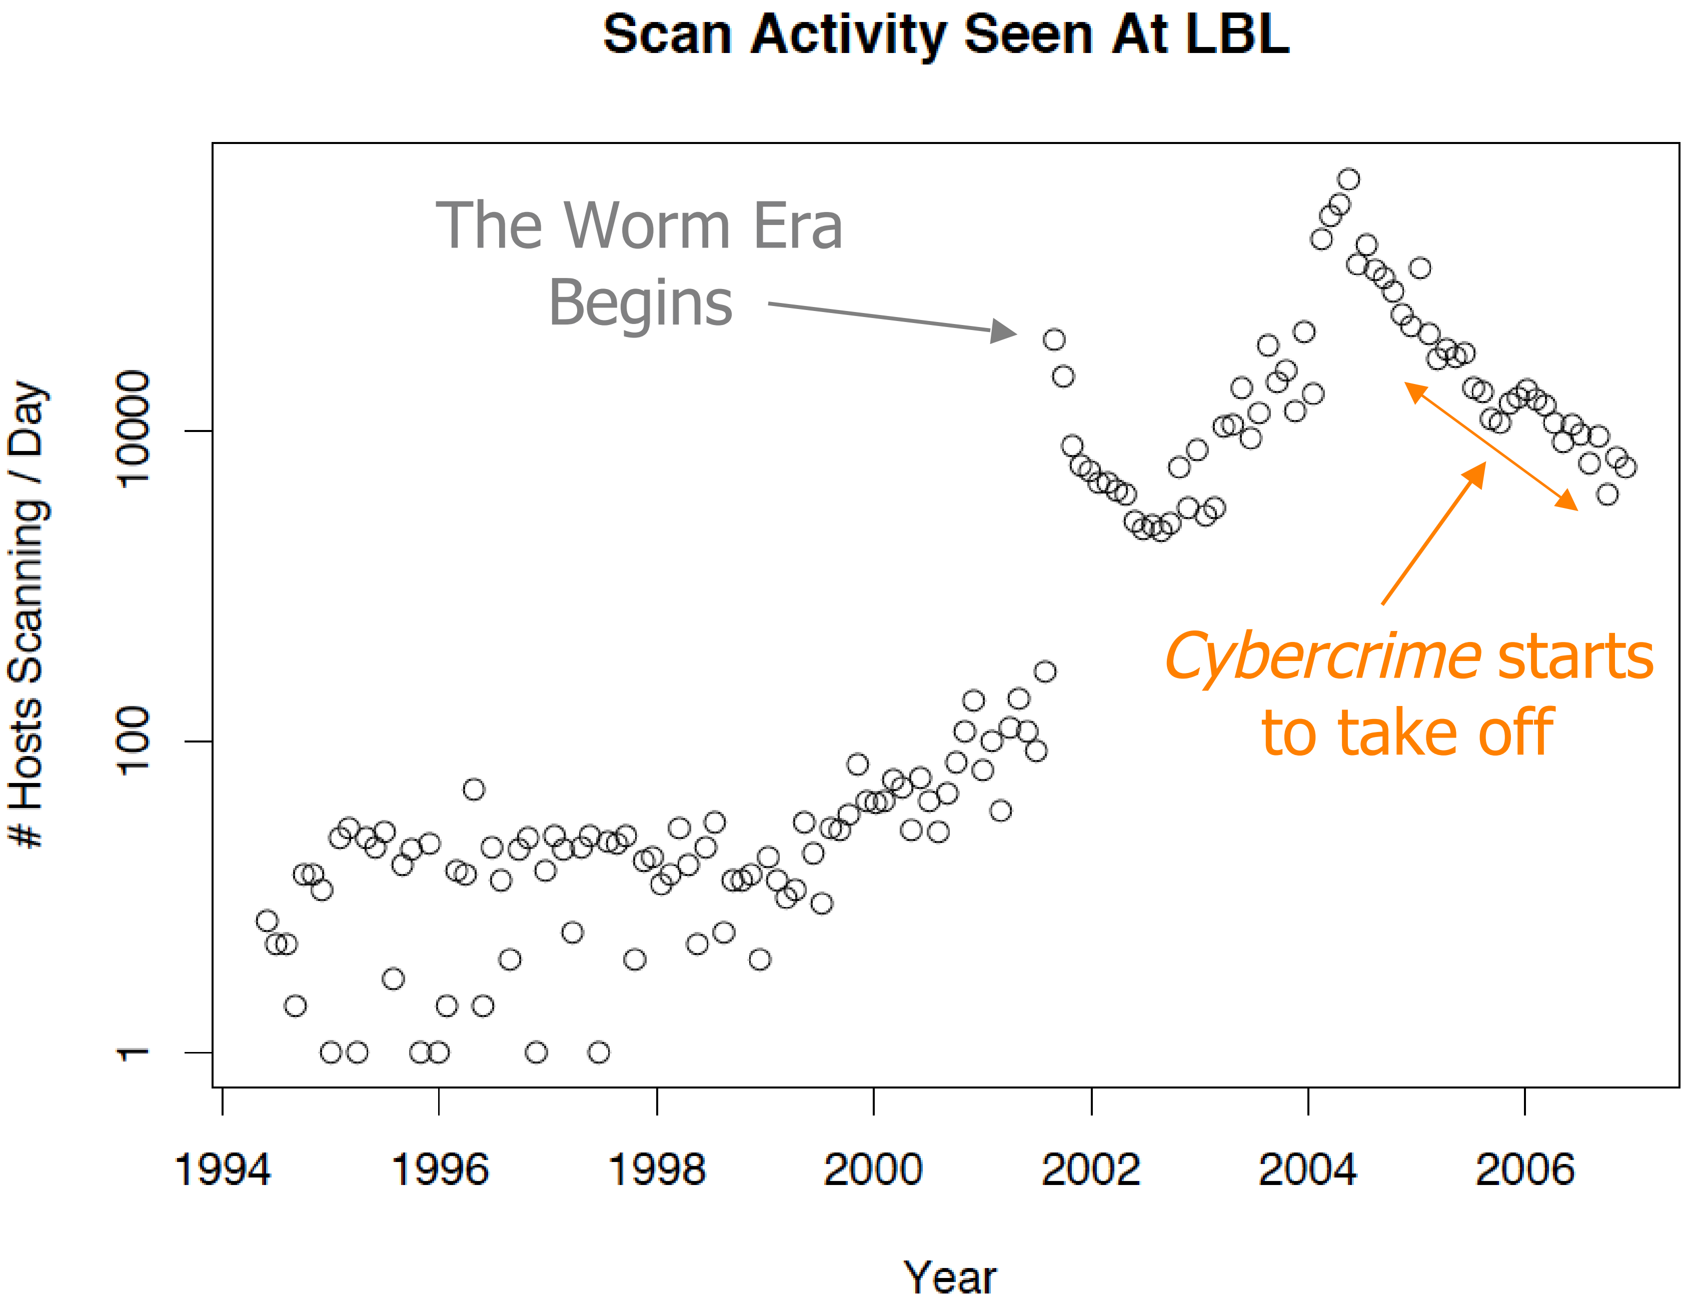

The Worm Era Begins

Cybercrime starts to take off

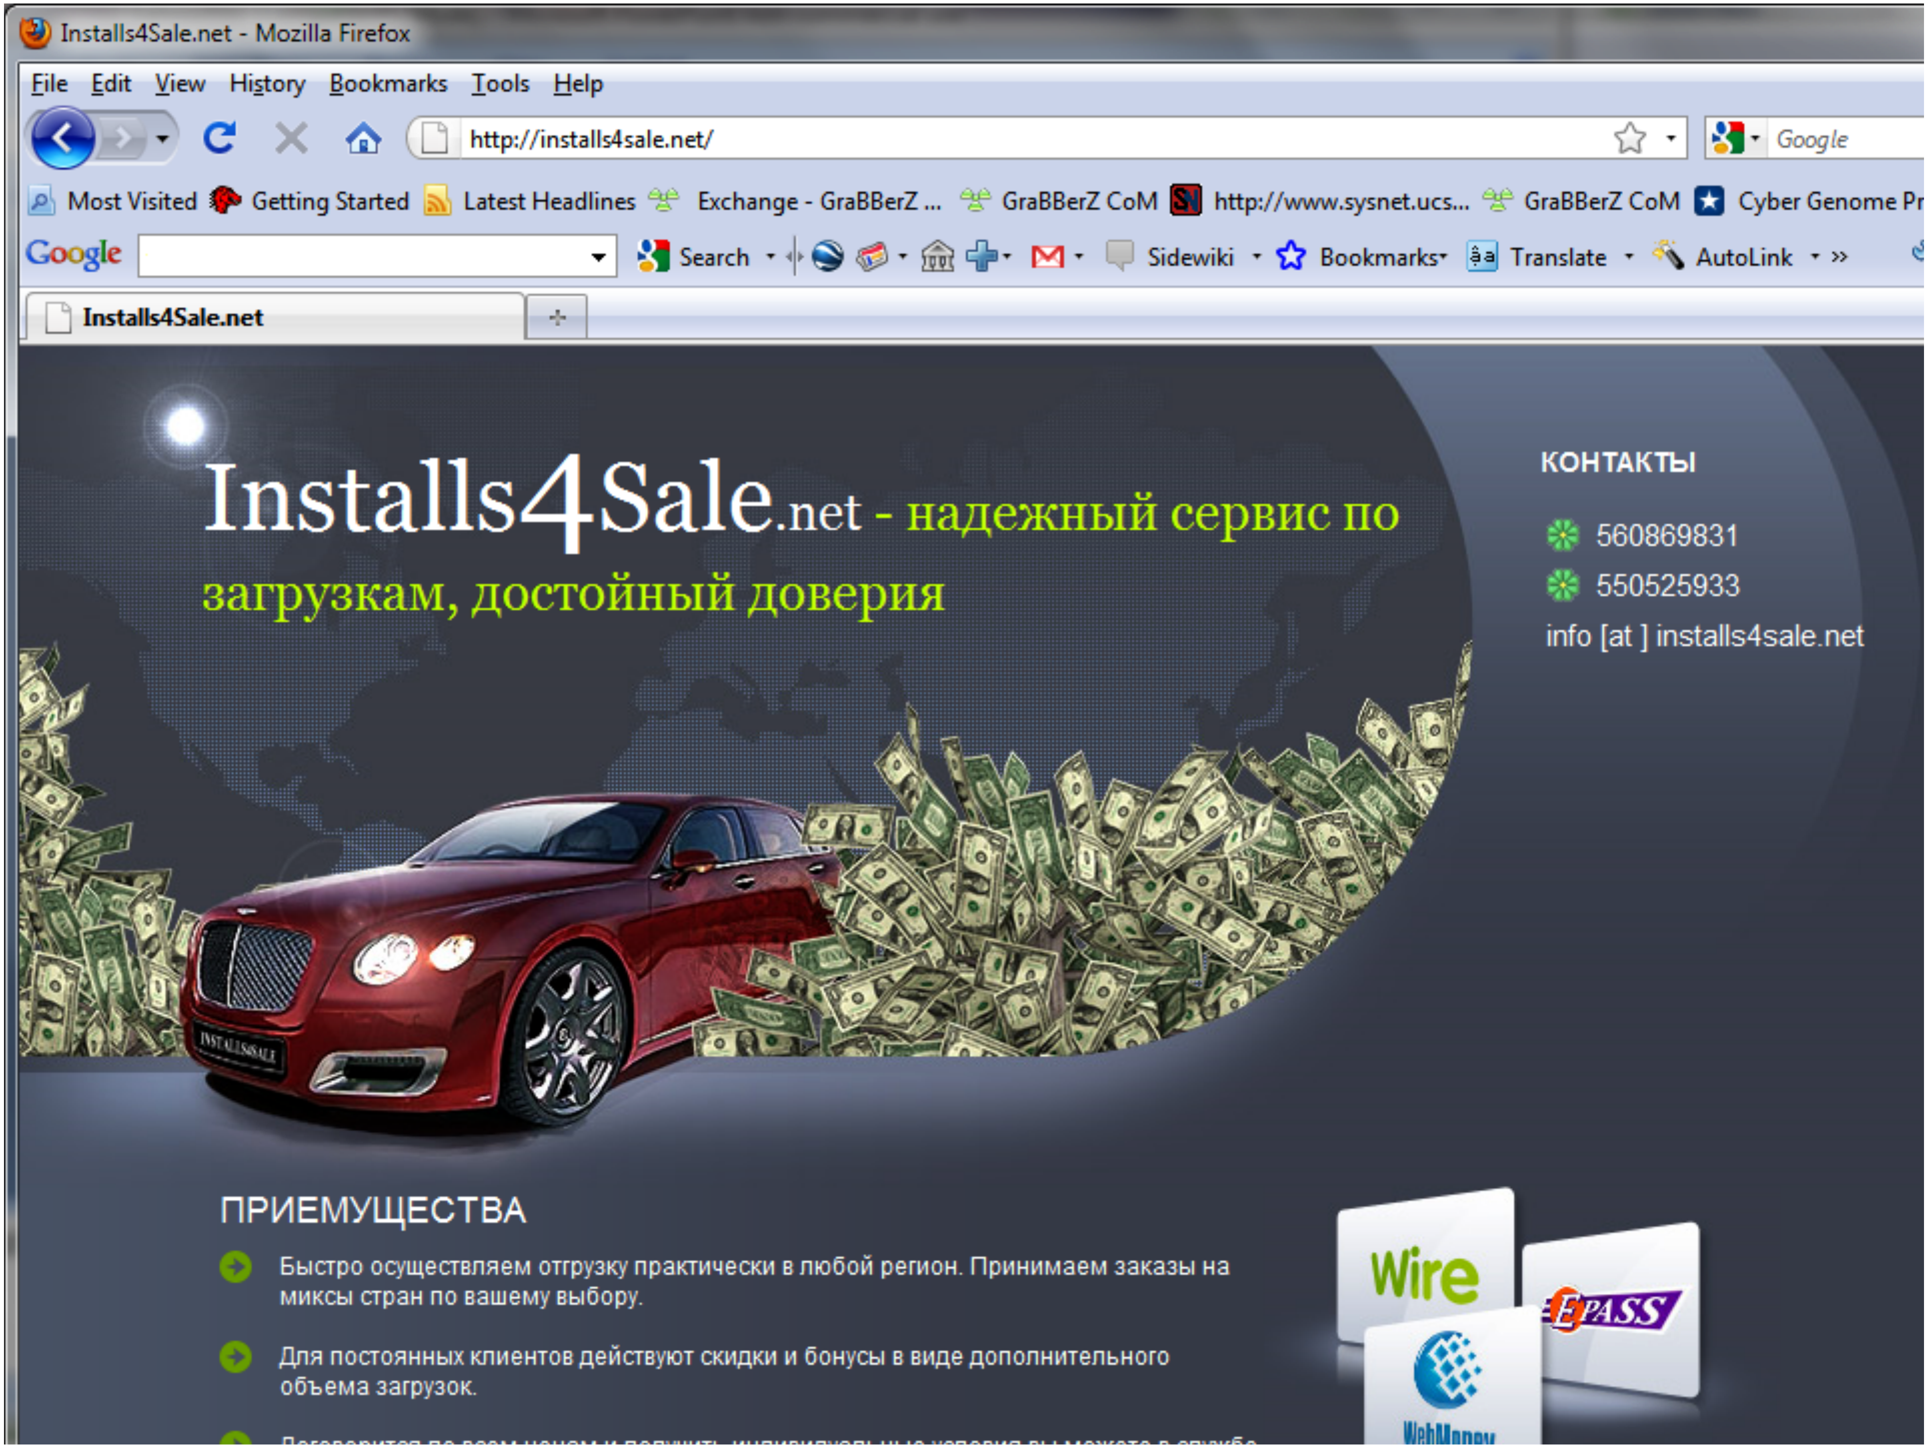

Installs4Sale.net - надежный сервис по загрузкам, достойный доверия

КОНТАКТЫ

560869831

550525933

info [at] installs4sale.net

ПРИЕМУЩЕСТВА

- Быстро осуществляем отгрузку практически в любой регион. Принимаем заказы на миксы стран по вашему выбору.
- Для постоянных клиентов действуют скидки и бонусы в виде дополнительного объема загрузок.
- Договорится по всем ценам и получить индивидуальные условия вы можете в службе

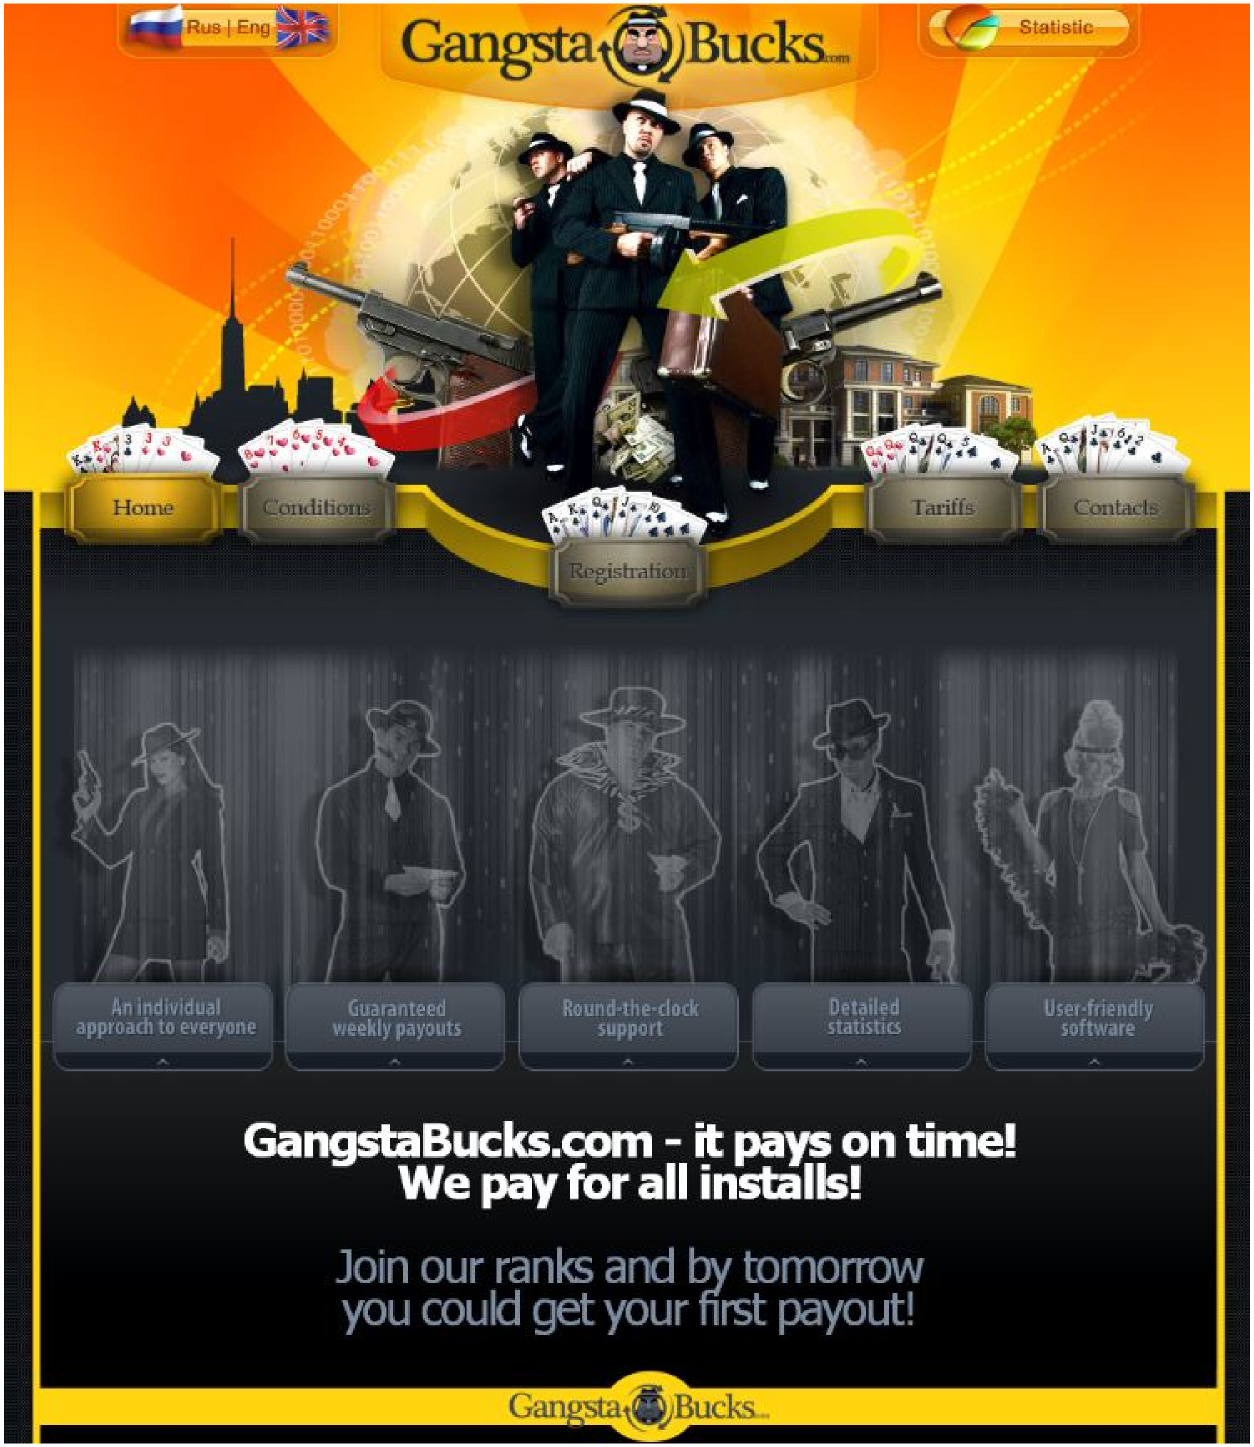

GangstaBucks.com - it pays on time!
We pay for all installs!

Join our ranks and by tomorrow
you could get your first payout!

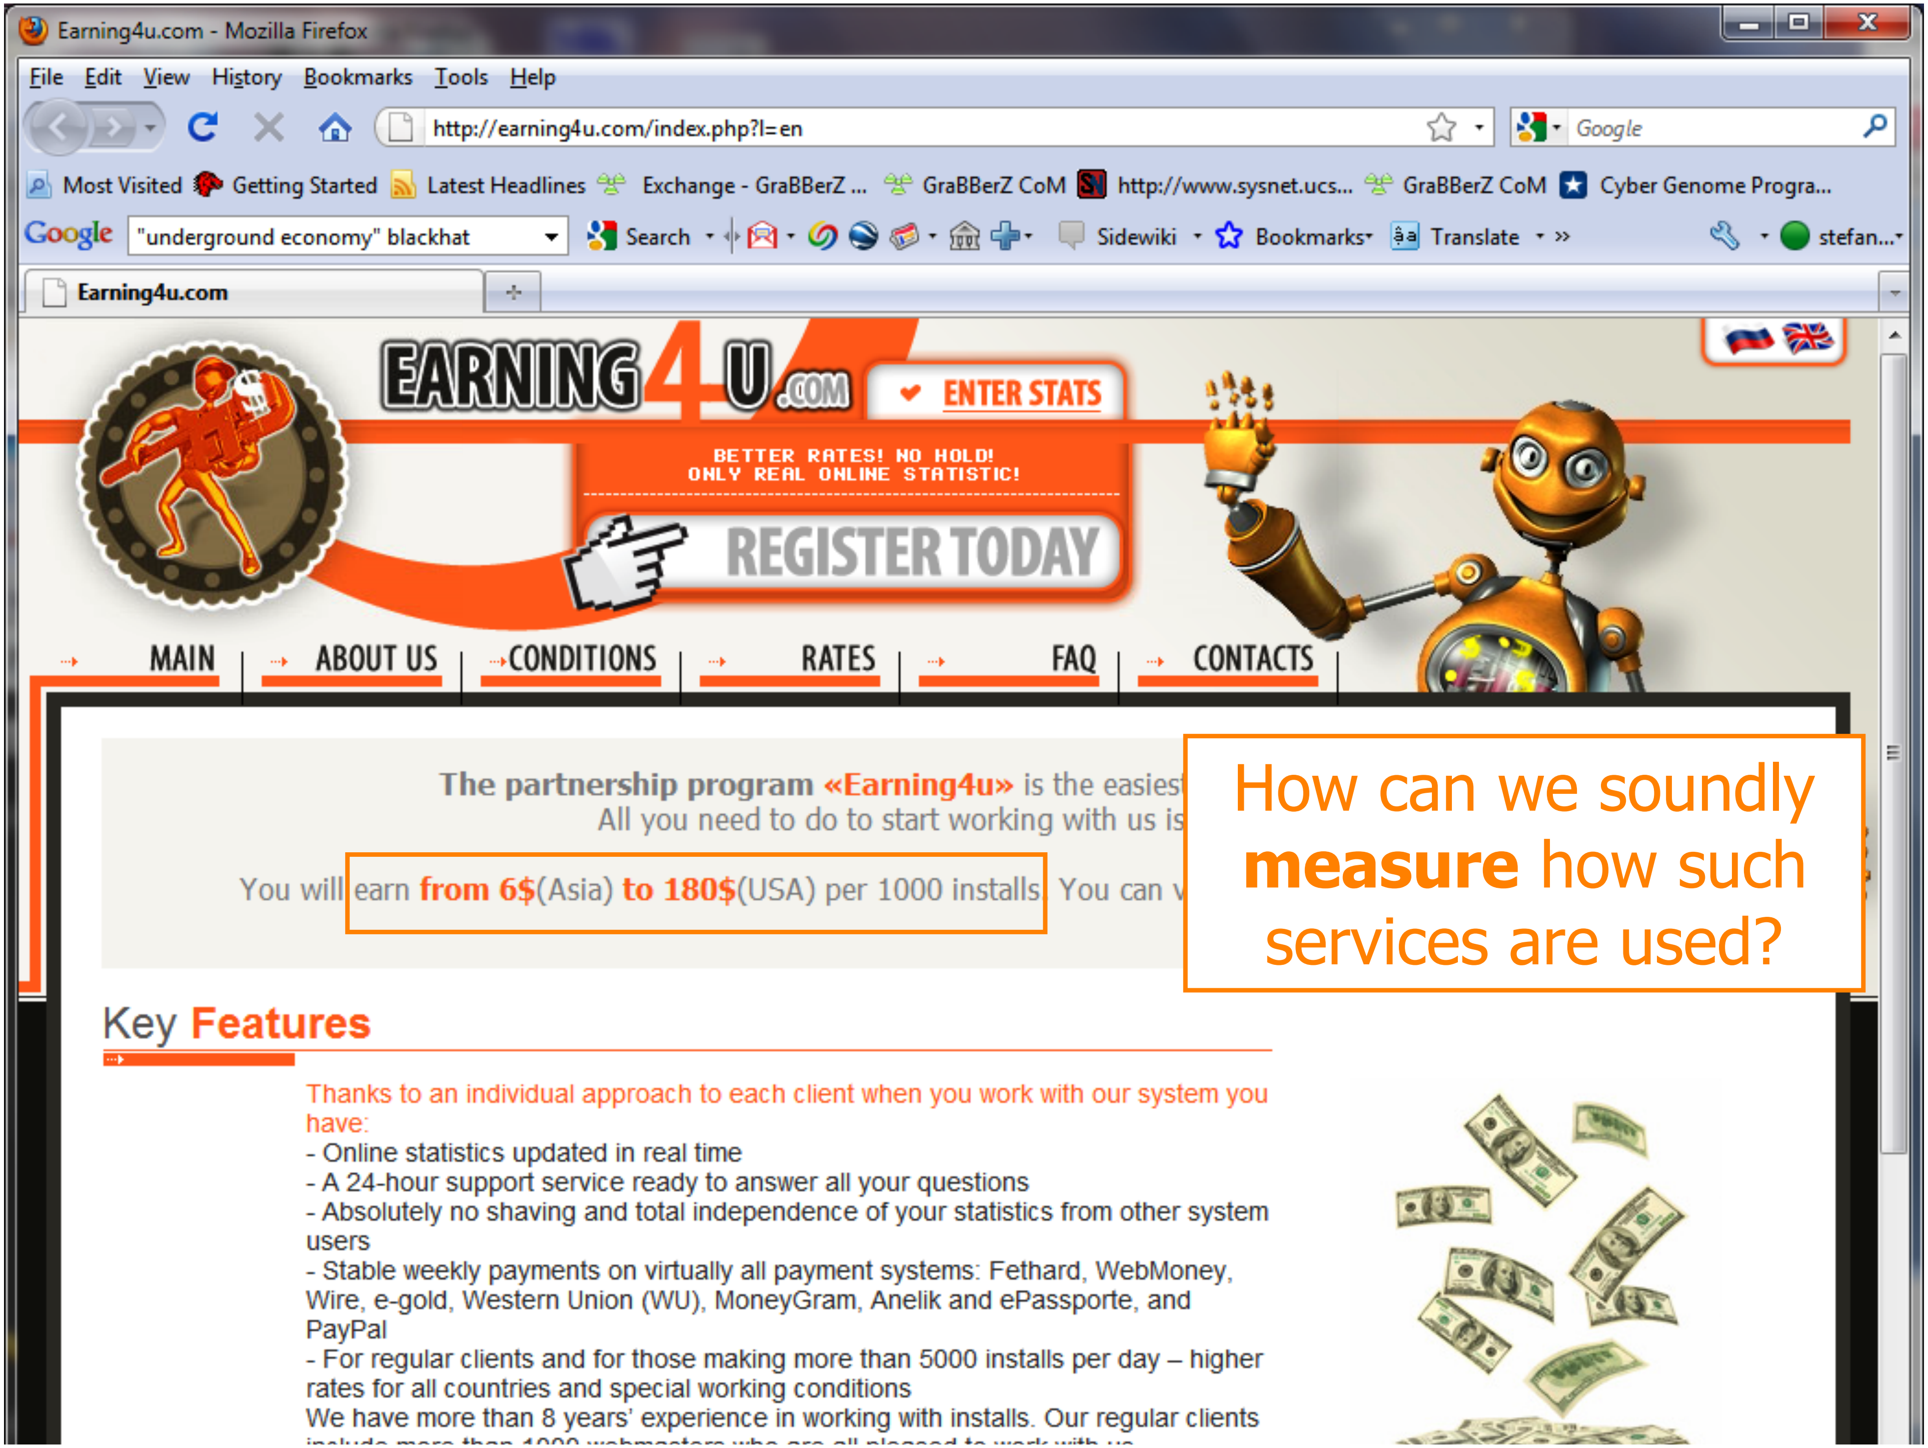

The partnership program **«Earning4u»** is the easies
All you need to do to start working with us is

You will earn **from 6$**(Asia) **to 180$**(USA) per 1000 installs. You can v

Key **Features**

Thanks to an individual approach to each client when you work with our system you have:

- Online statistics updated in real time
- A 24-hour support service ready to answer all your questions
- Absolutely no shaving and total independence of your statistics from other system users
- Stable weekly payments on virtually all payment systems: Fethard, WebMoney, Wire, e-gold, Western Union (WU), MoneyGram, Anelik and ePassporte, and PayPal
- For regular clients and for those making more than 5000 installs per day – higher rates for all countries and special working conditions

We have more than 8 years' experience in working with installs. Our regular clients include more than 1000 webmasters who are all pleased to work with us

How can we soundly **measure** how such services are used?

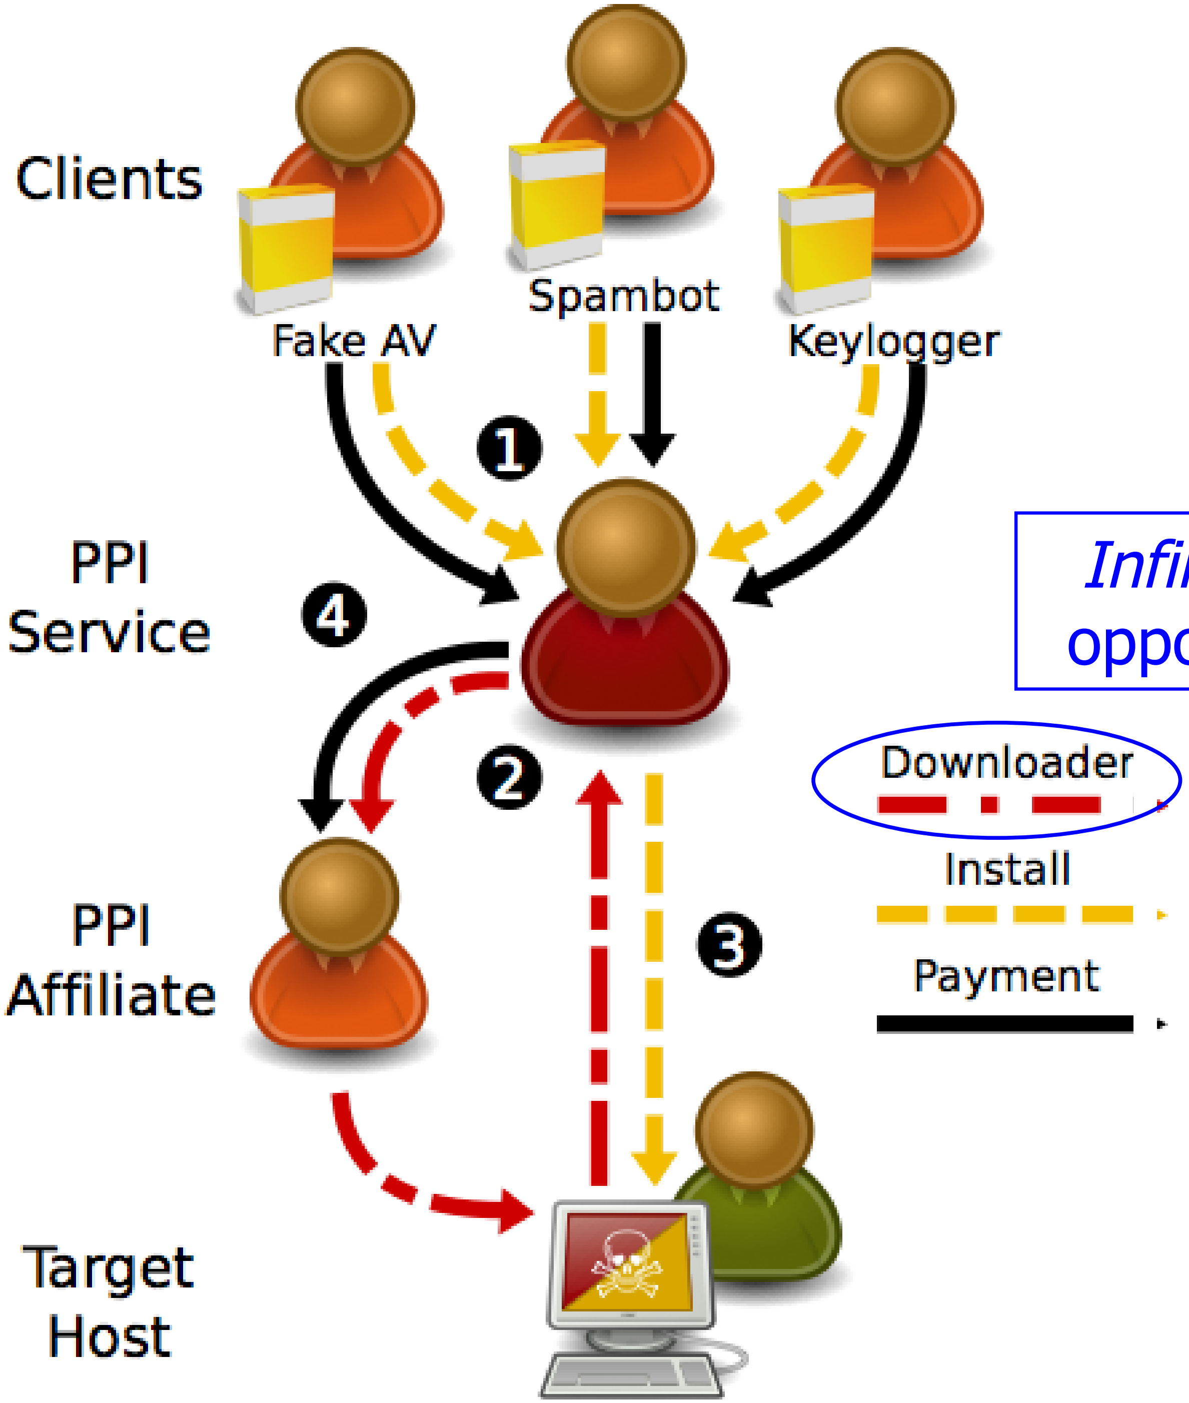

Infiltration opportunity

Advanced malware intelligence via PPI infiltration

Milking = mimic downloader, repeatedly ask PPI service for next program to install

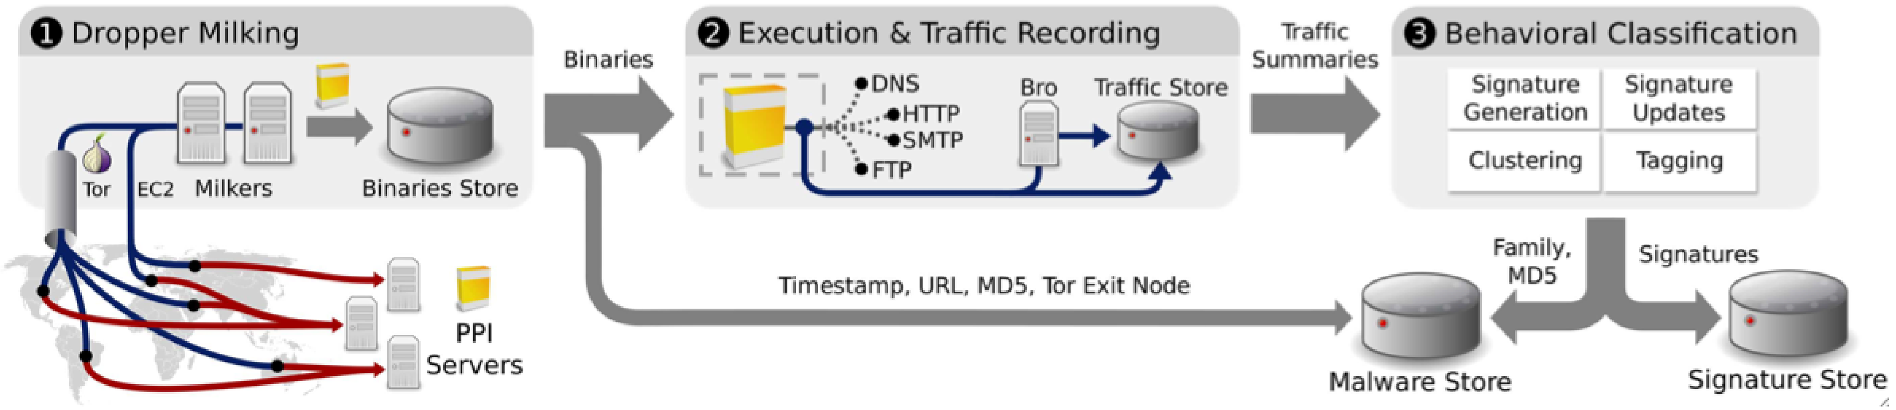

Running since August 2010, we downloaded > 1M binaries (9K distinct) from 4 different affiliate programs

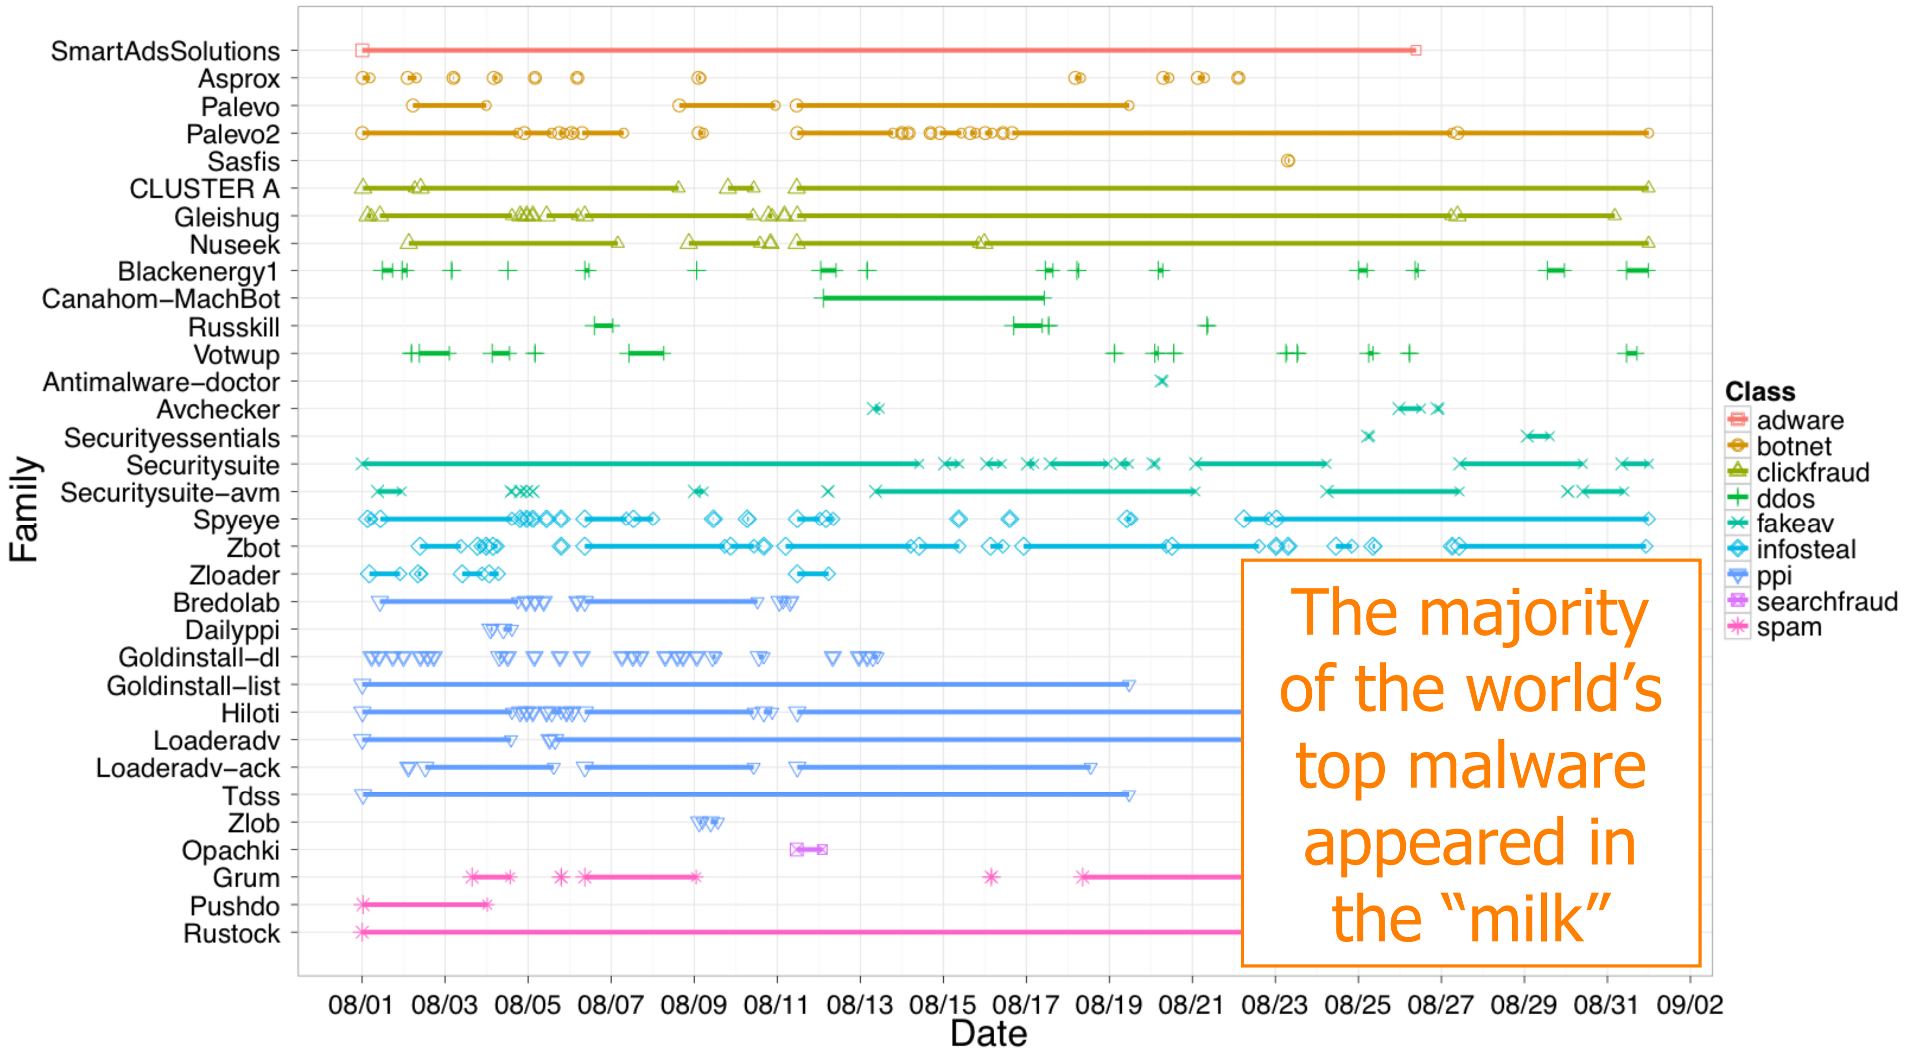

The majority of the world’s top malware appeared in the “milk”

PPI distribution of malware during August 2010

Phases of the Spam Value Chain

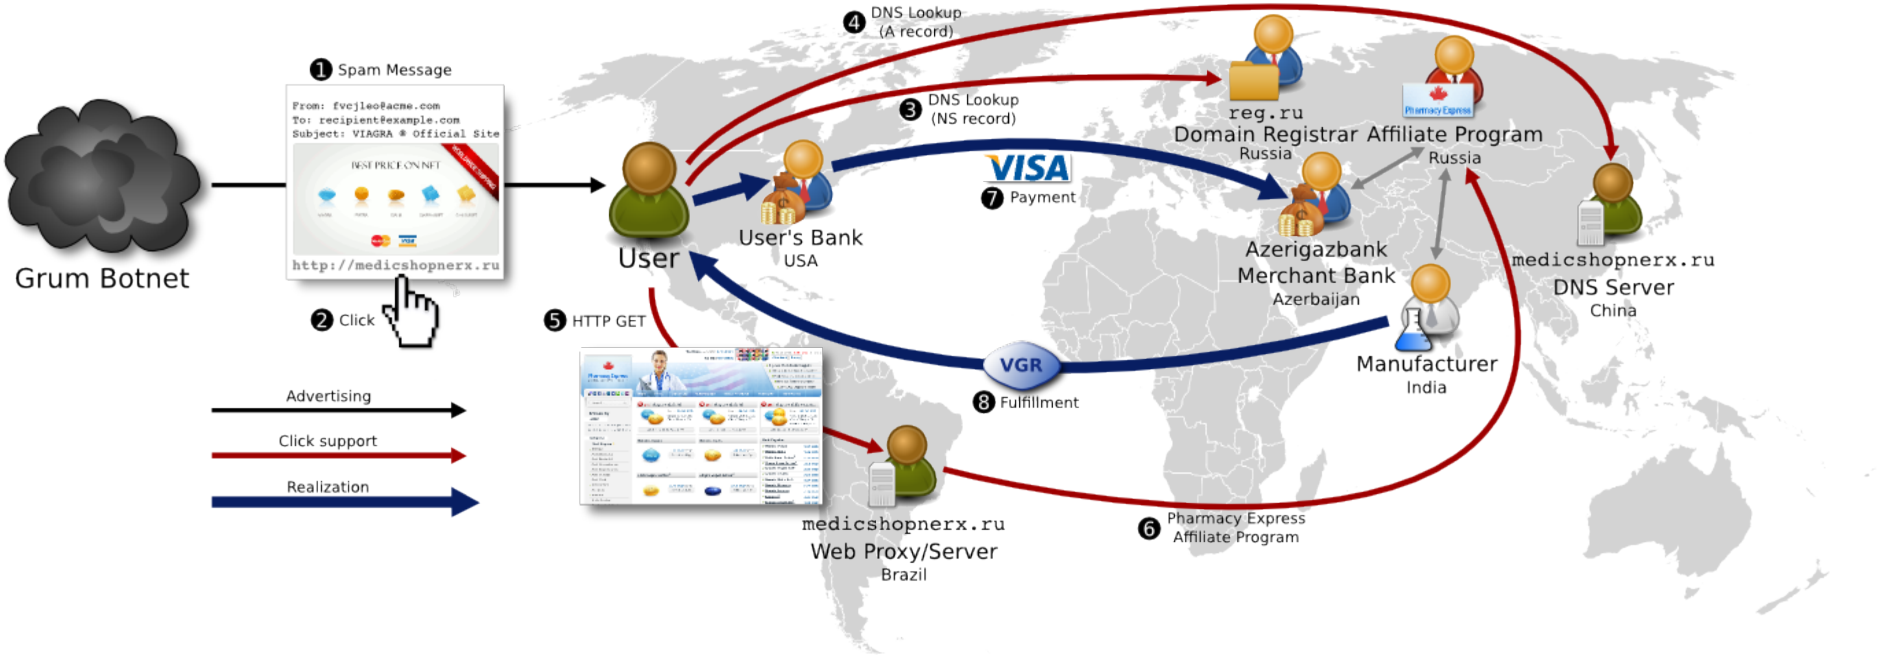

From: fvcjleo@acme.com
To: recipient@example.com
Subject: VIAGRA ® Official Site

http://medicshopnerx.ru

Measuring URLs, DNS servers, HTTP redirection, etc. all a matter of energetic crawling & recording.

But **purchases** / **banks** / "**fulfillment**" ??

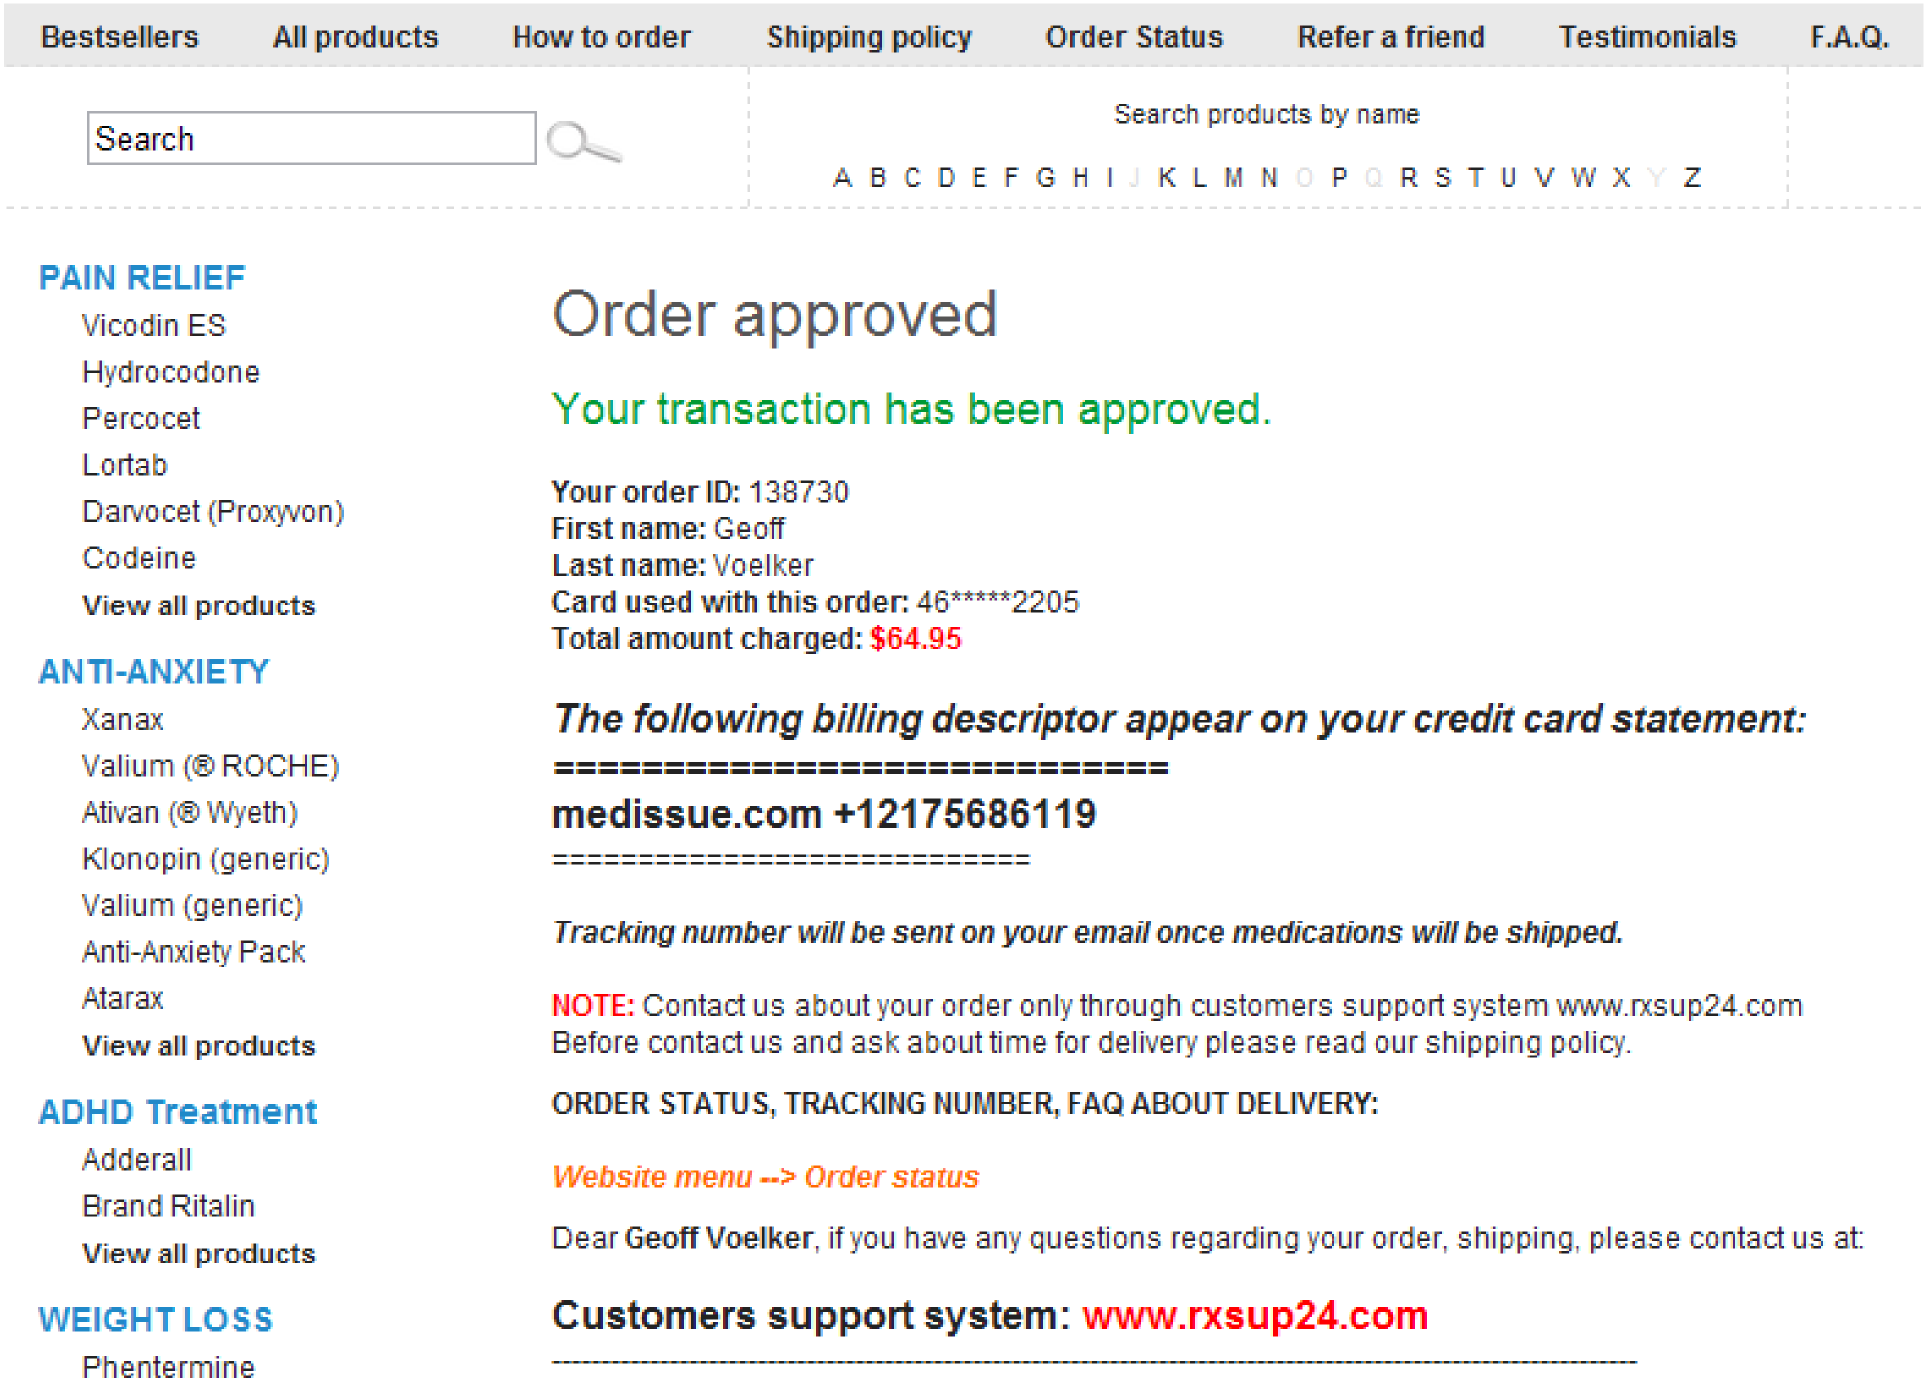

Order approved

Your transaction has been approved.

Your order ID: 138730
First name: Geoff
Last name: Voelker
Card used with this order: 46*****2205
Total amount charged: **$64.95**

The following billing descriptor appear on your credit card statement:

===============================

medissue.com +12175686119

===============================

Tracking number will be sent on your email once medications will be shipped.

NOTE: Contact us about your order only through customers support system www.rxsup24.com
Before contact us and ask about time for delivery please read our shipping policy.

ORDER STATUS, TRACKING NUMBER, FAQ ABOUT DELIVERY:

Website menu --> Order status

Dear **Geoff Voelker**, if you have any questions regarding your order, shipping, please contact us at:

Customers support system: www.rxsup24.com

DELL

40
50
30
20
10
18

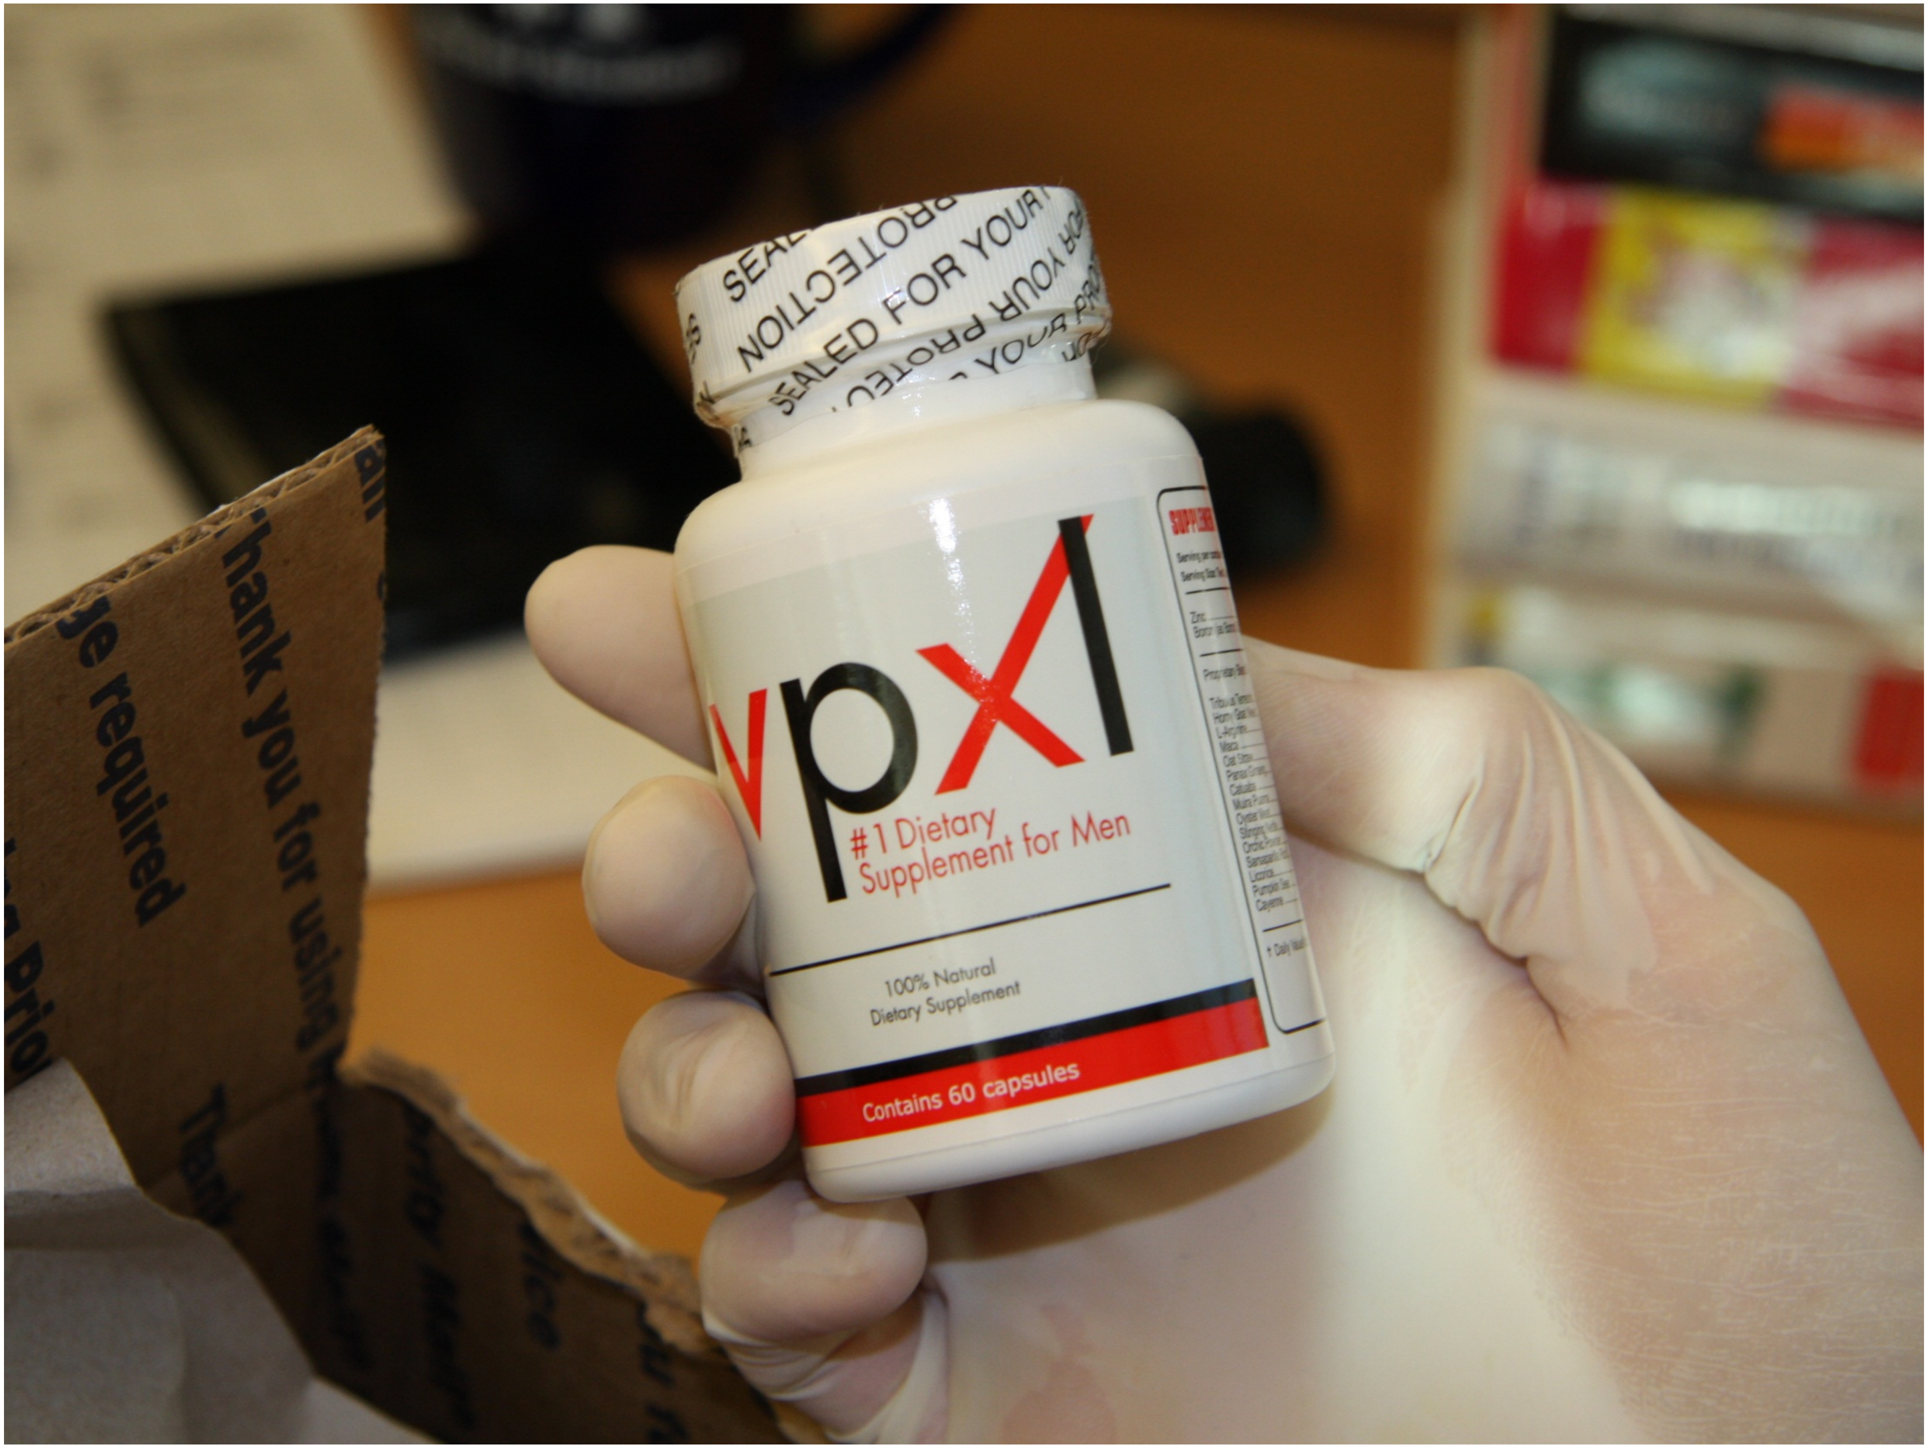
SEALED FOR YOUR PROTECTION
vpXl
#1 Dietary Supplement for Men
100% Natural Dietary Supplement
Contains 60 capsules
thank you for using
ge required

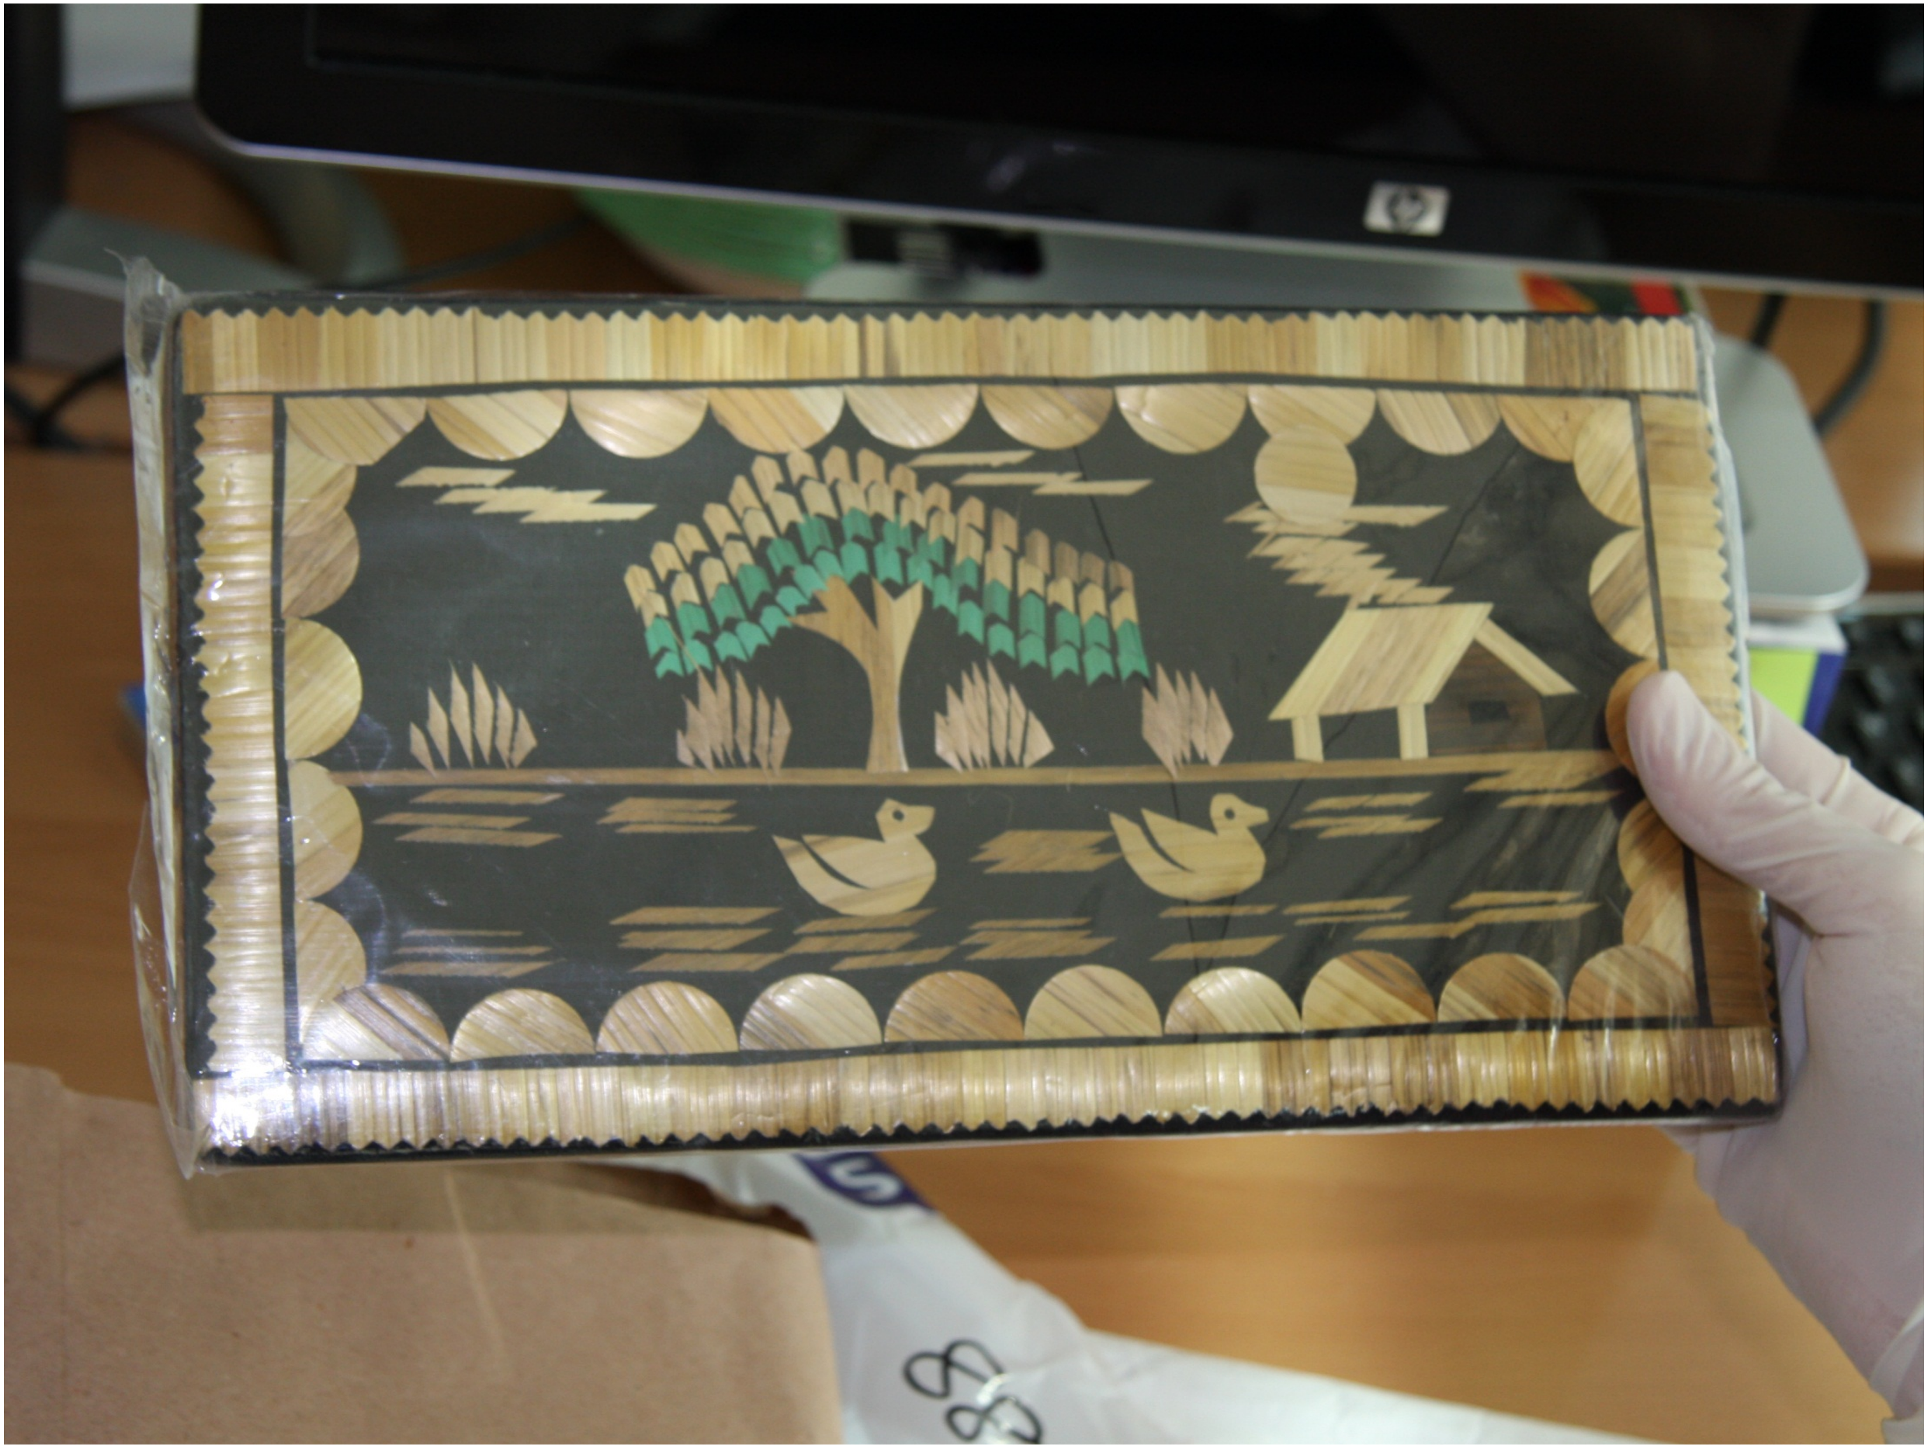
88

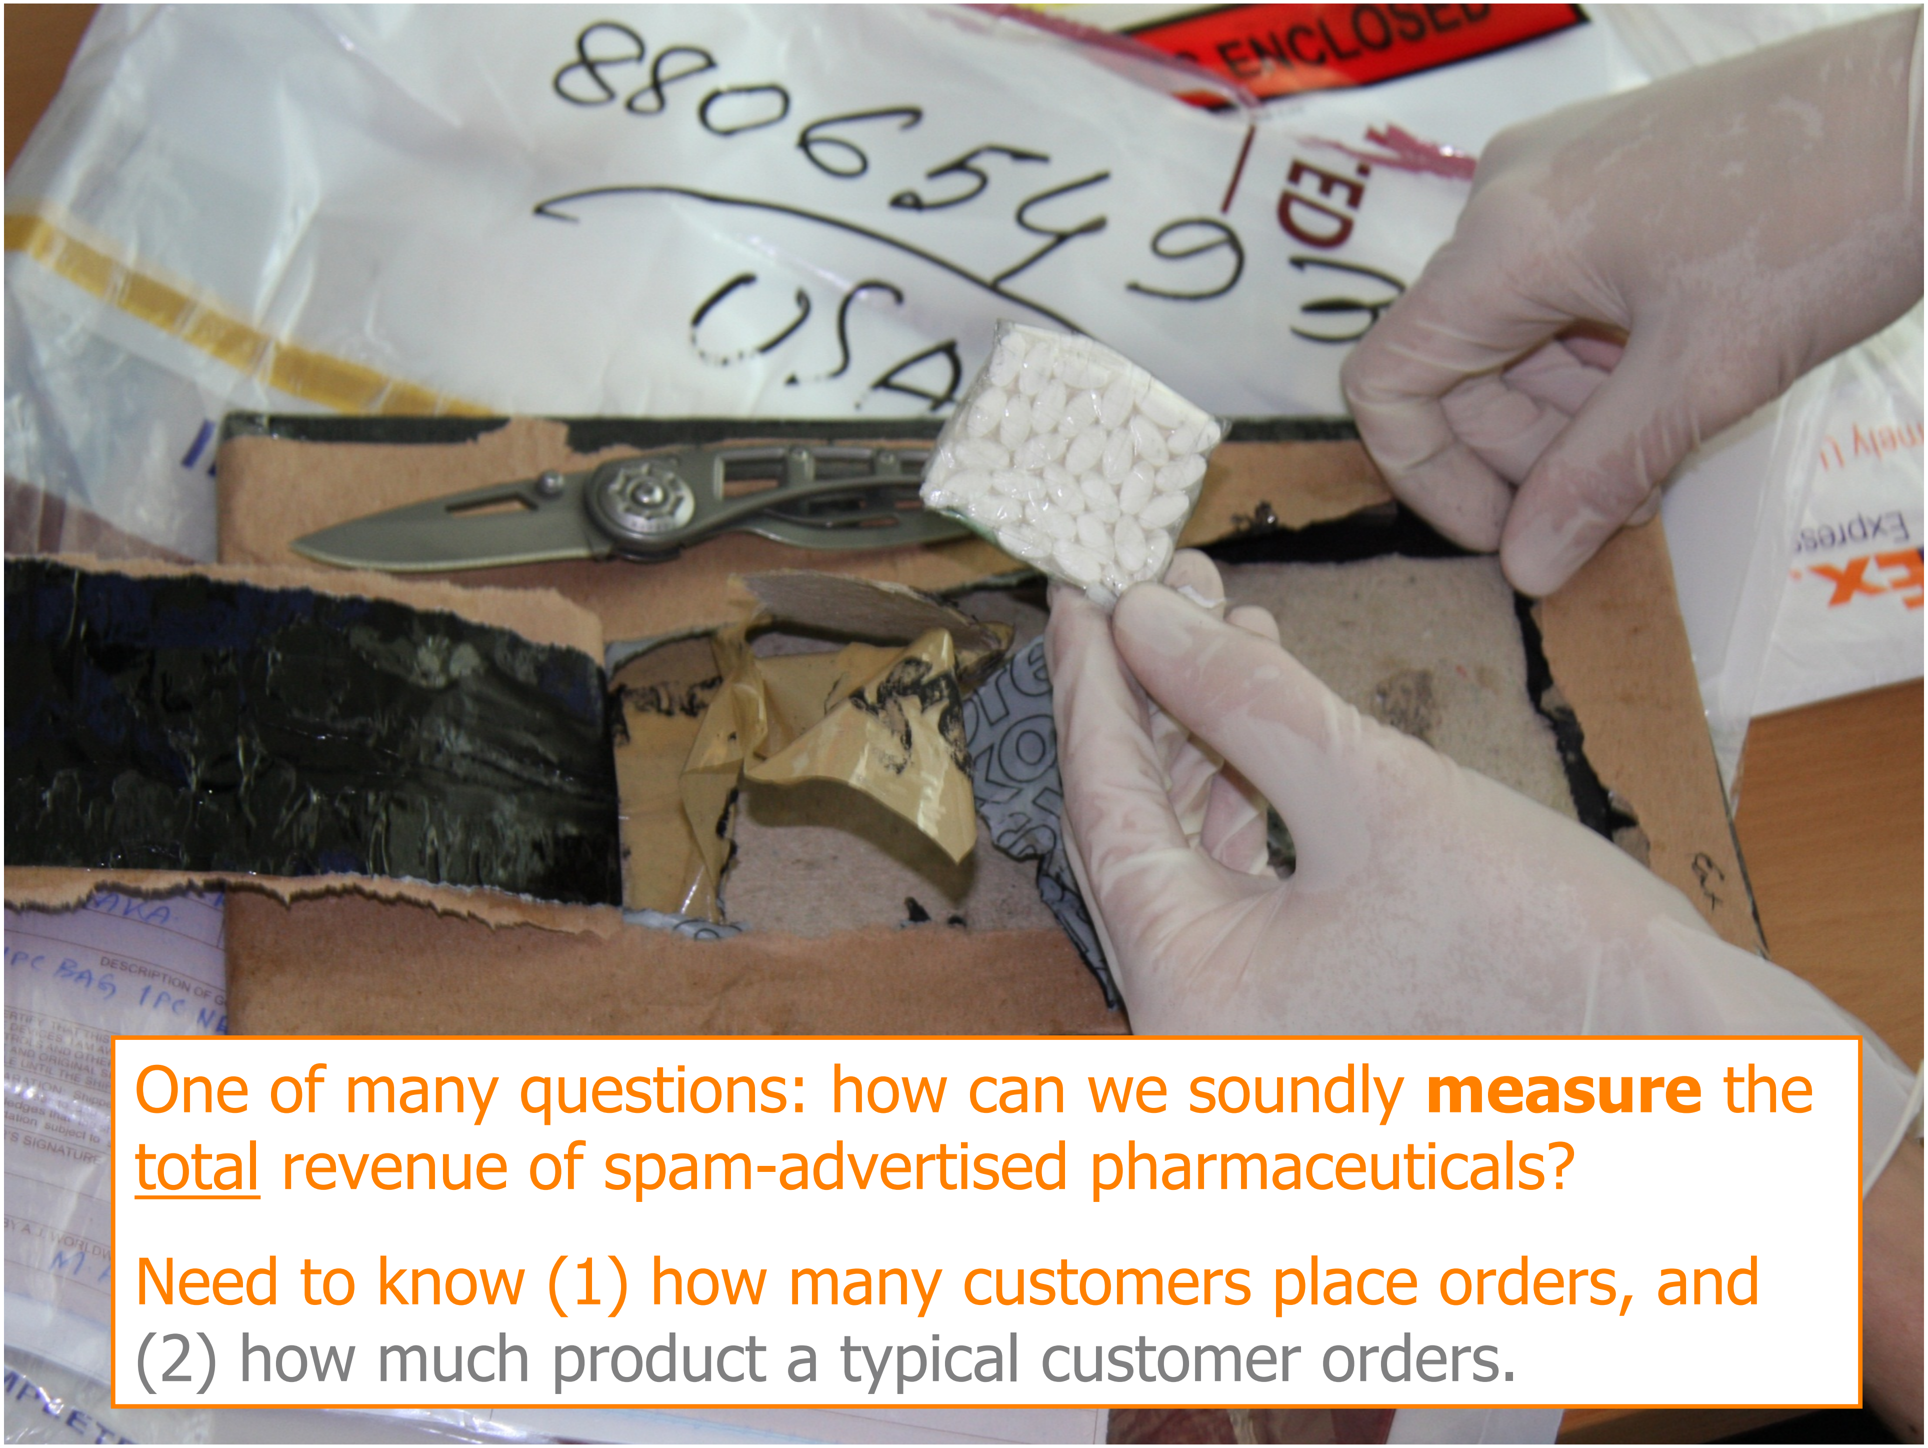

One of many questions: how can we soundly **measure** the total revenue of spam-advertised pharmaceuticals?

Need to know (1) how many customers place orders, and (2) how much product a typical customer orders.

Part III:

A Seeming Digression re the IP ID Field

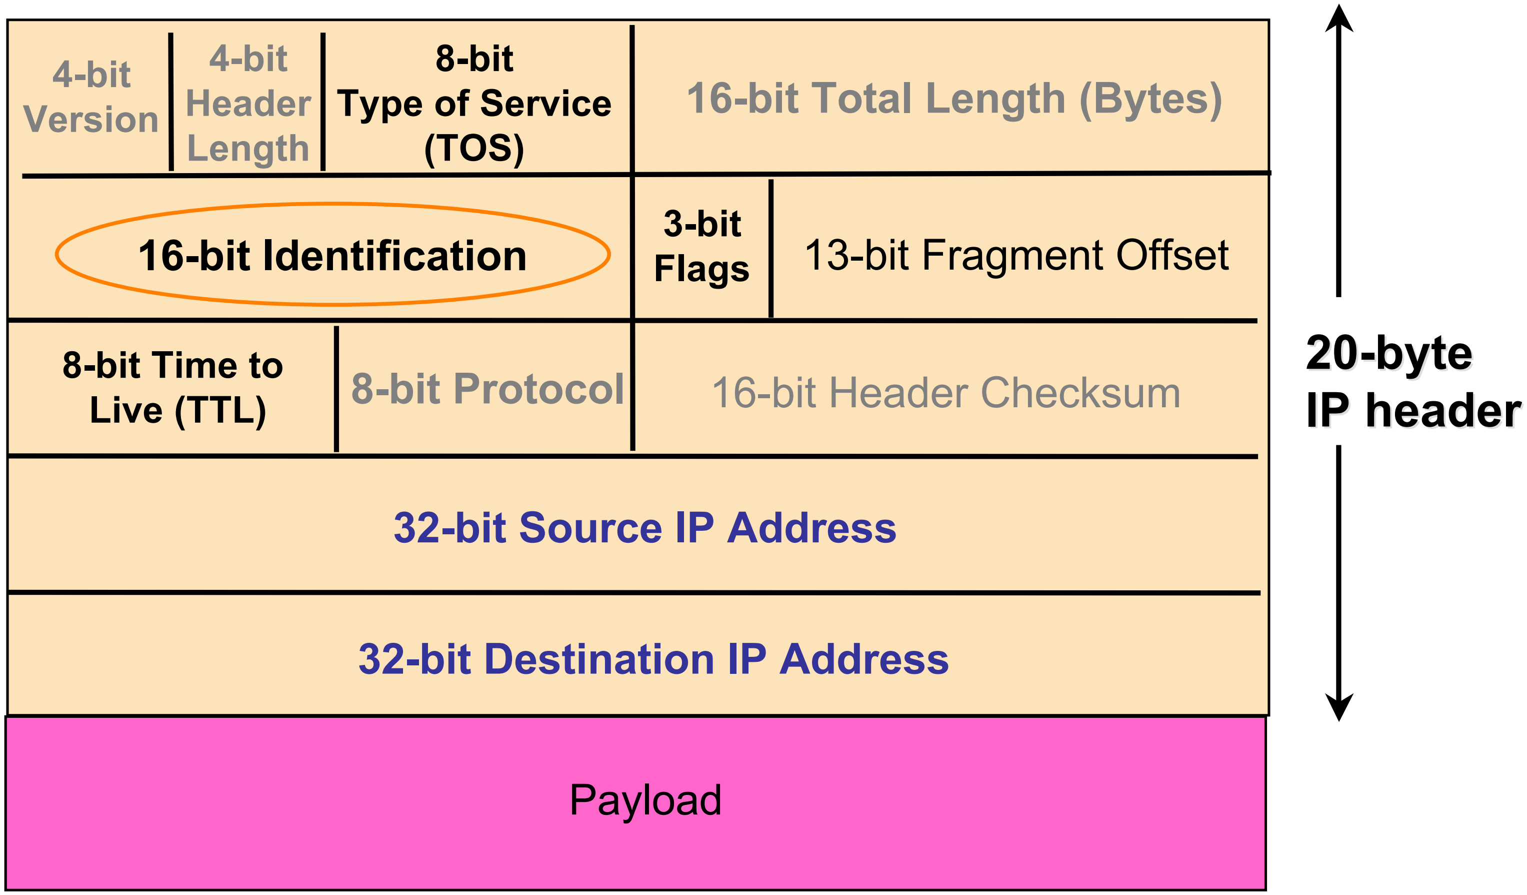

4-bit Version

4-bit Header Length

8-bit Type of Service (TOS)

16-bit Total Length (Bytes)

16-bit Identification

3-bit Flags

13-bit Fragment Offset

8-bit Time to Live (TTL)

8-bit Protocol

16-bit Header Checksum

32-bit Source IP Address

32-bit Destination IP Address

Payload

20-byte IP header

- Many systems increment it per packet globally sent ⇒ *side channel*
- Enables inference of quantity of traffic sent between two points in time, otherwise unobserved
- (Side channel even enables NAT detection and stealthy port scanning)

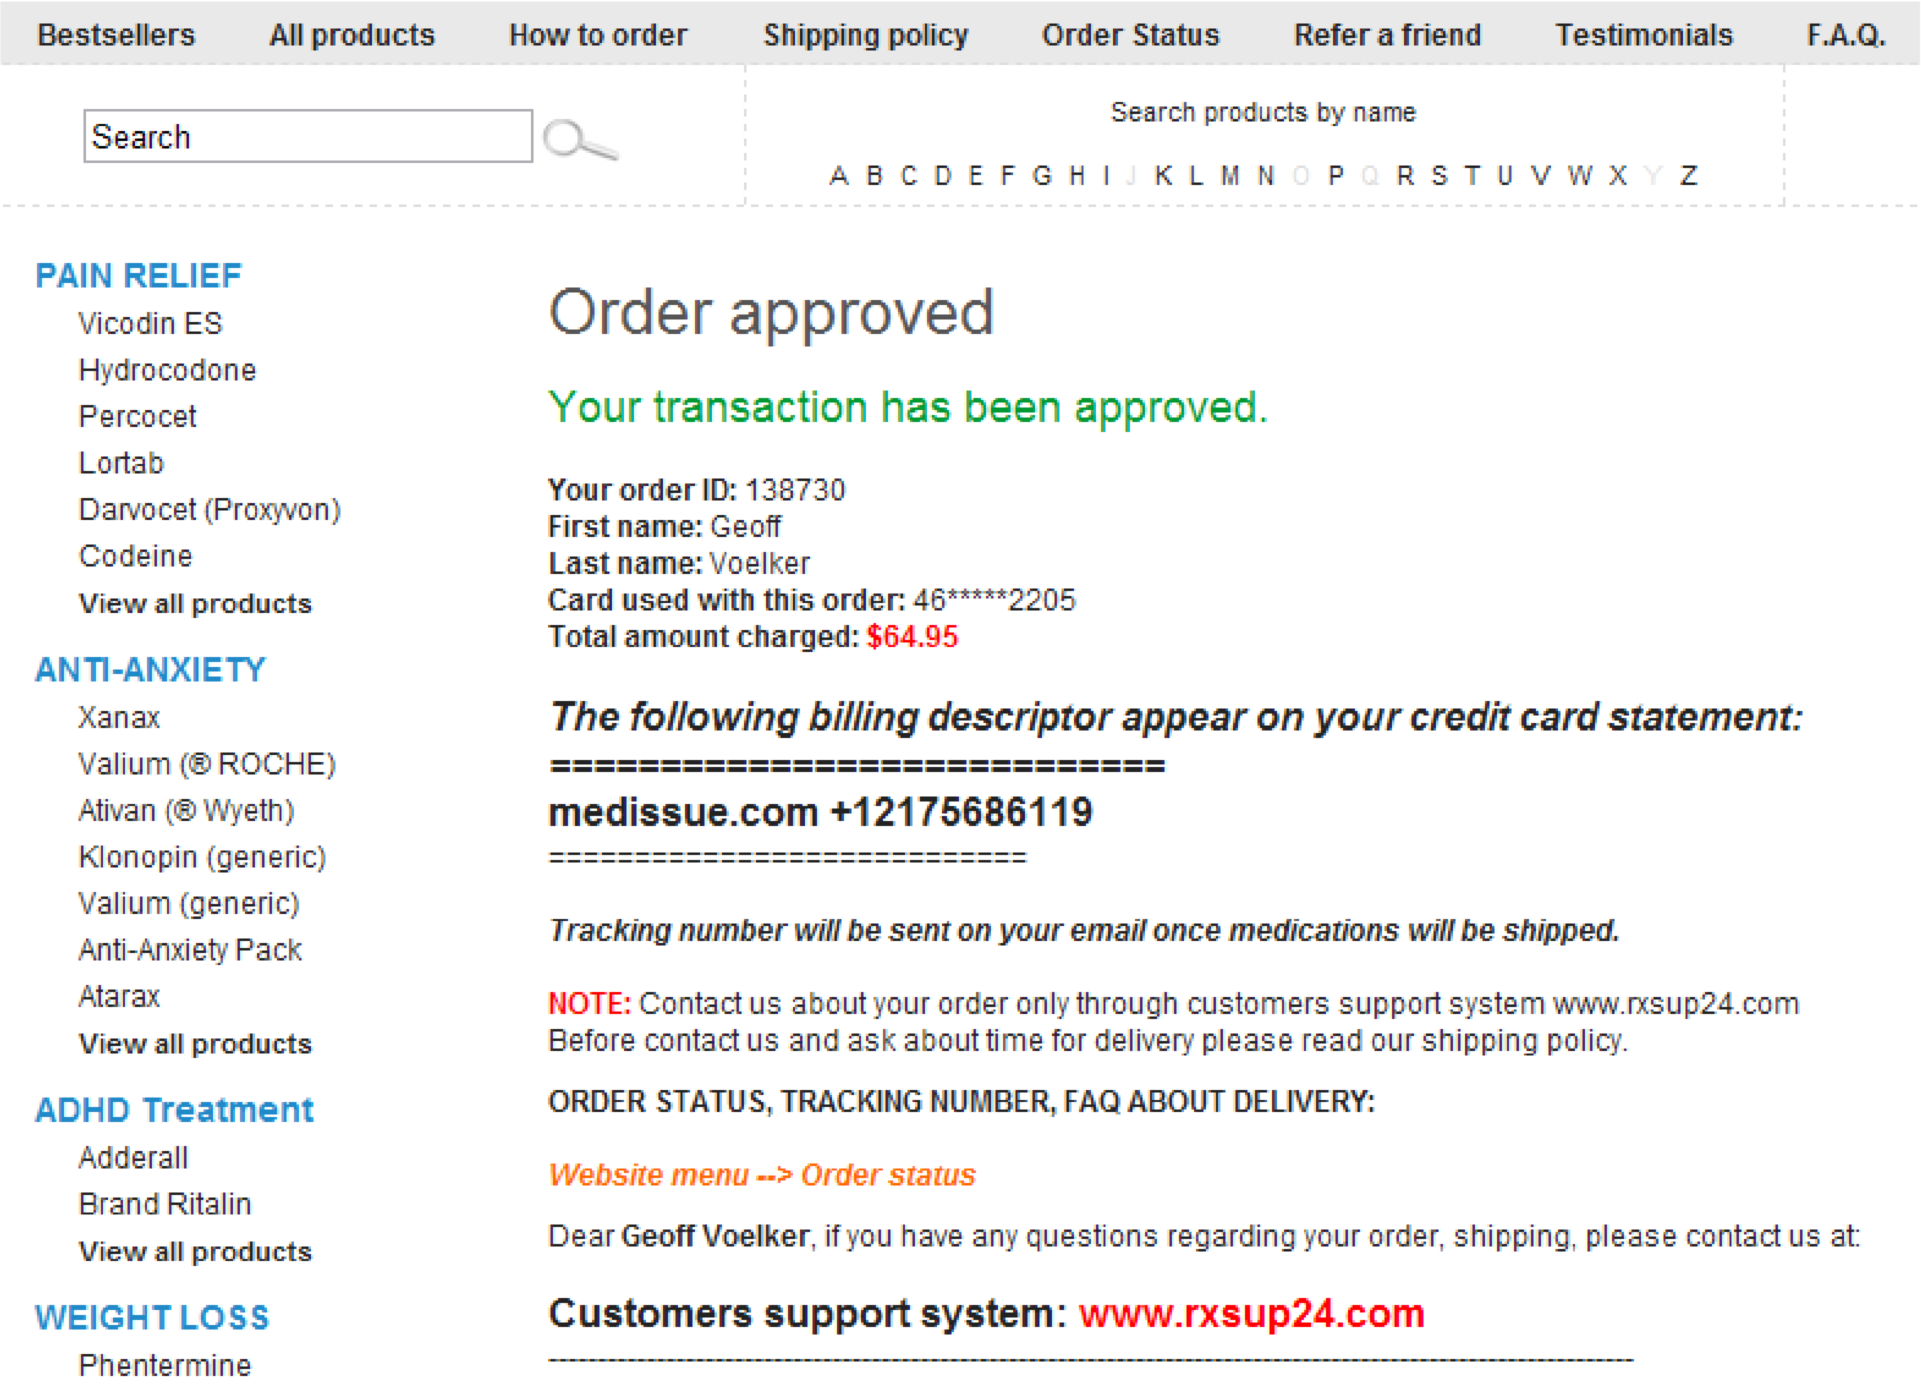

Order approved

Your transaction has been approved.

Your order ID: 138730
First name: Geoff
Last name: Voelker
Card used with this order: 46*****2205
Total amount charged: **$64.95**

The following billing descriptor appear on your credit card statement:

============================

medissue.com +12175686119

=============================

Tracking number will be sent on your email once medications will be shipped.

NOTE: Contact us about your order only through customers support system www.rxsup24.com
Before contact us and ask about time for delivery please read our shipping policy.

ORDER STATUS, TRACKING NUMBER, FAQ ABOUT DELIVERY:

Website menu --> Order status

Dear **Geoff Voelker**, if you have any questions regarding your order, shipping, please contact us at:

Customers support system: www.rxsup24.com

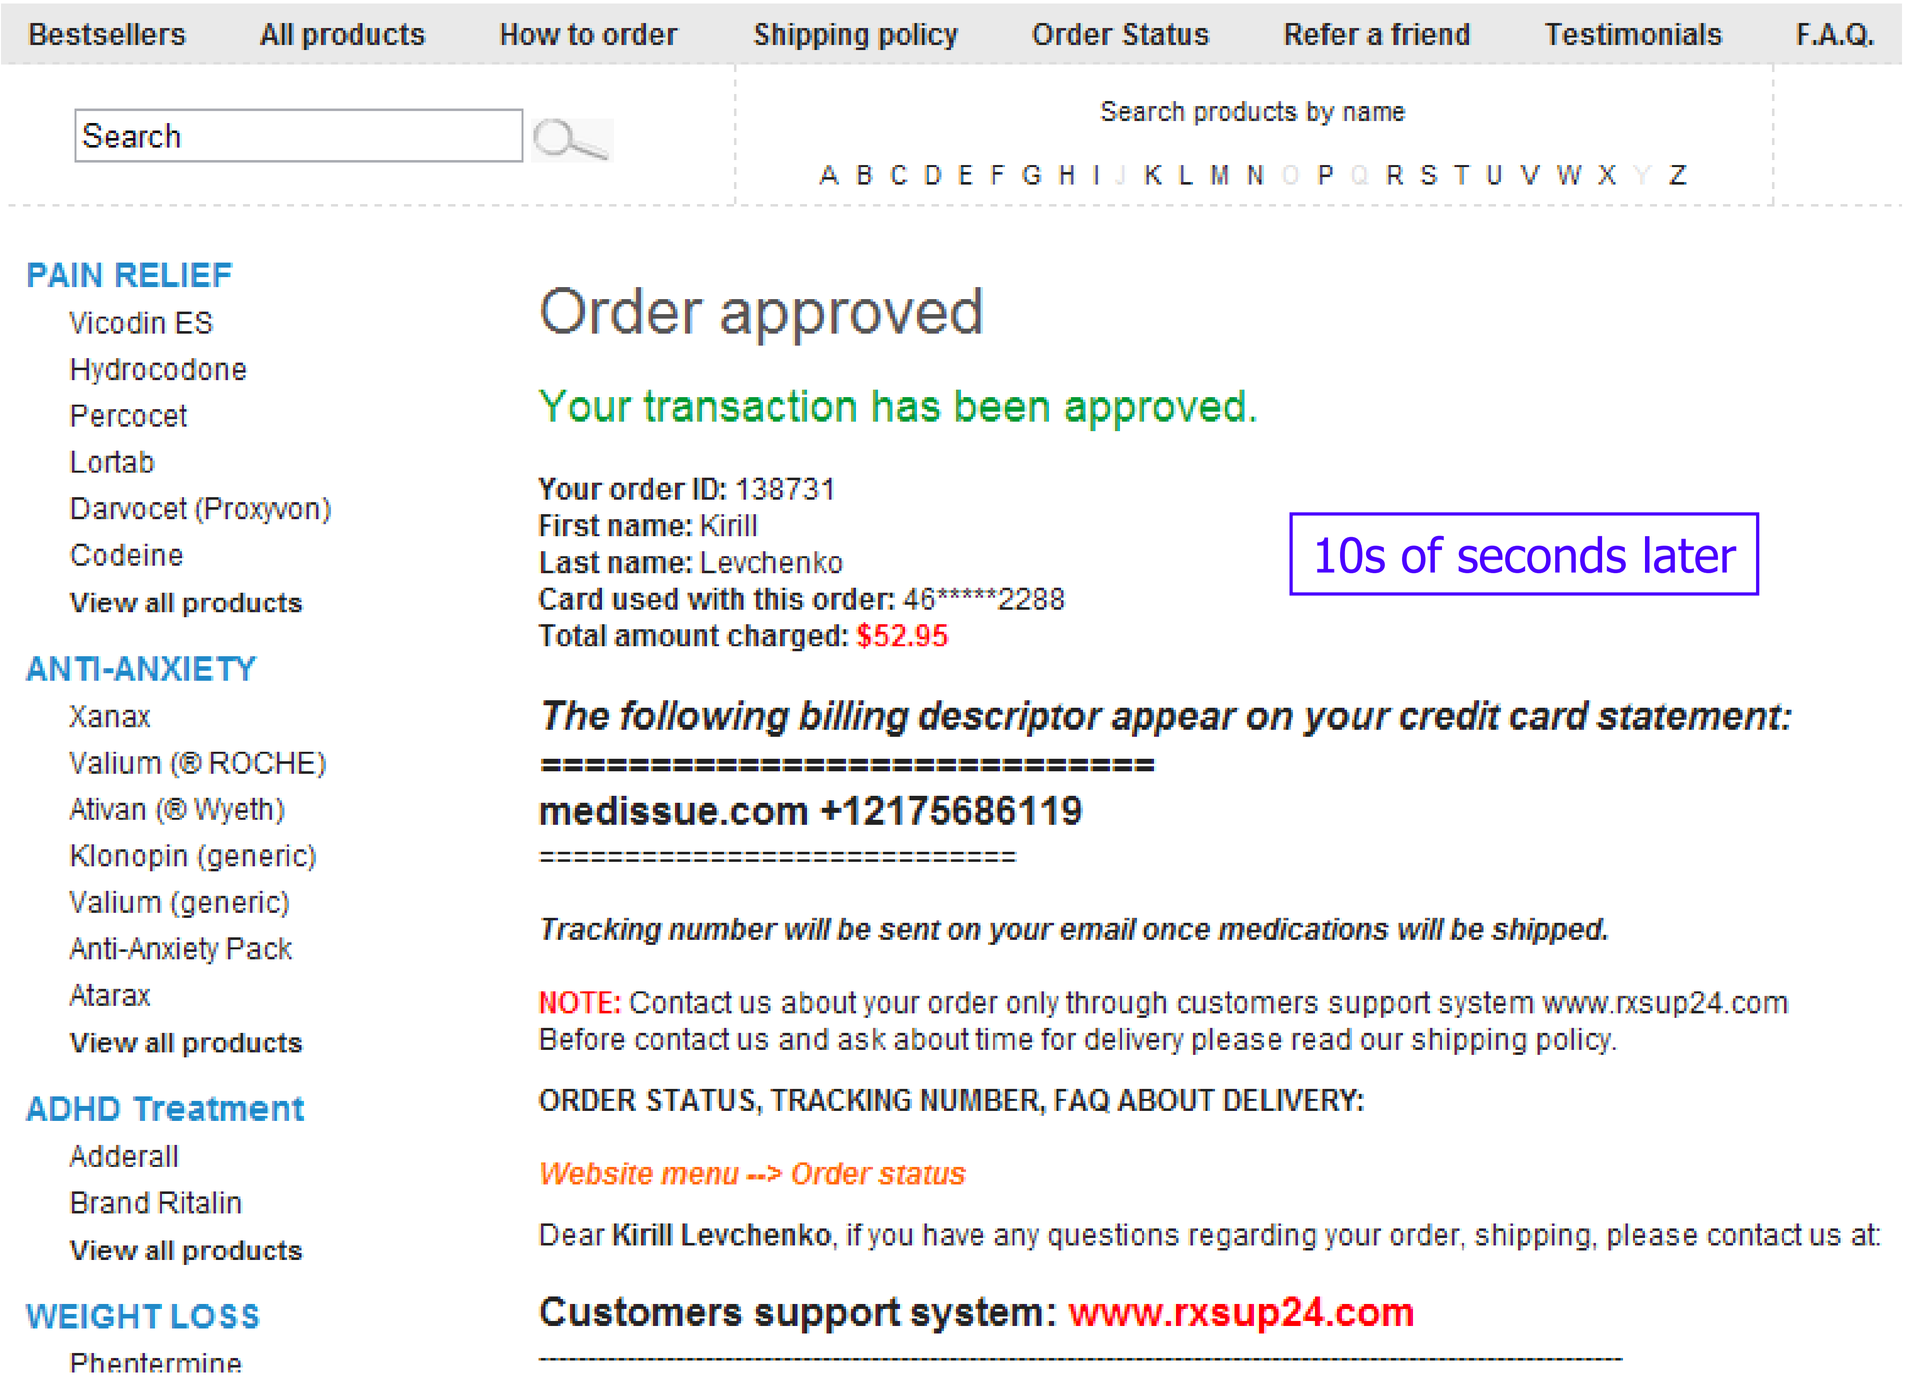

Order approved

Your transaction has been approved.

Your order ID: 138731
First name: Kirill
Last name: Levchenko
Card used with this order: 46*****2288
Total amount charged: $52.95

10s of seconds later

The following billing descriptor appear on your credit card statement:

================================

medissue.com +12175686119

============================

Tracking number will be sent on your email once medications will be shipped.

NOTE: Contact us about your order only through customers support system www.rxsup24.com
Before contact us and ask about time for delivery please read our shipping policy.

ORDER STATUS, TRACKING NUMBER, FAQ ABOUT DELIVERY:

Website menu --> Order status

Dear **Kirill Levchenko**, if you have any questions regarding your order, shipping, please contact us at:

Customers support system: www.rxsup24.com

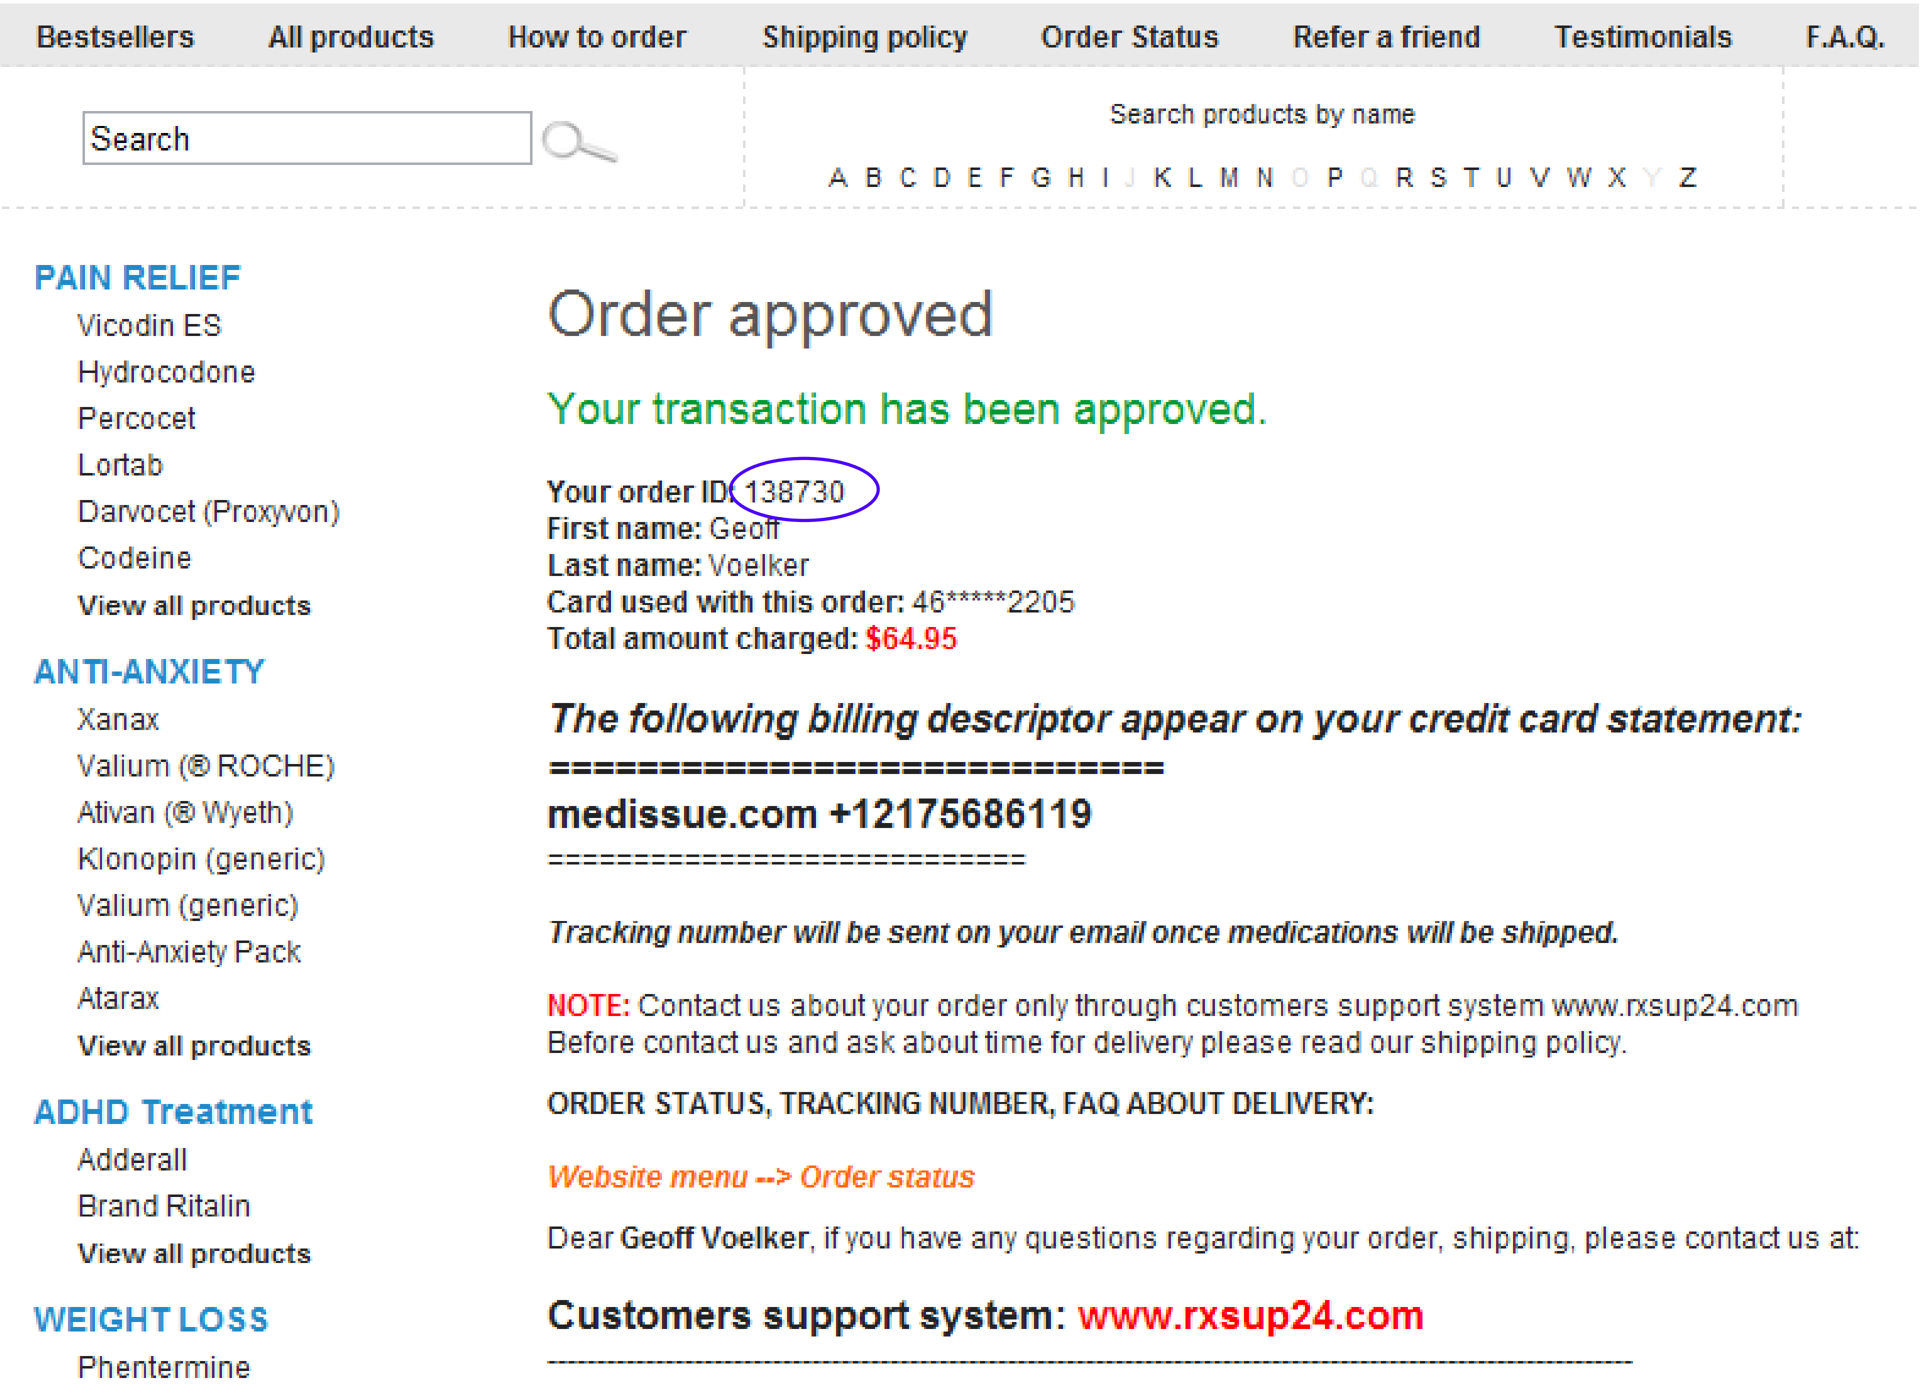

Order approved

Your transaction has been approved.

Your order ID: 138730
First name: Geoff
Last name: Voelker
Card used with this order: 46*****2205
Total amount charged: $64.95

The following billing descriptor appear on your credit card statement:

==============================

medissue.com +12175686119

============================

Tracking number will be sent on your email once medications will be shipped.

NOTE: Contact us about your order only through customers support system www.rxsup24.com
Before contact us and ask about time for delivery please read our shipping policy.

ORDER STATUS, TRACKING NUMBER, FAQ ABOUT DELIVERY:

Website menu --> Order status

Dear **Geoff Voelker**, if you have any questions regarding your order, shipping, please contact us at:

Customers support system: www.rxsup24.com

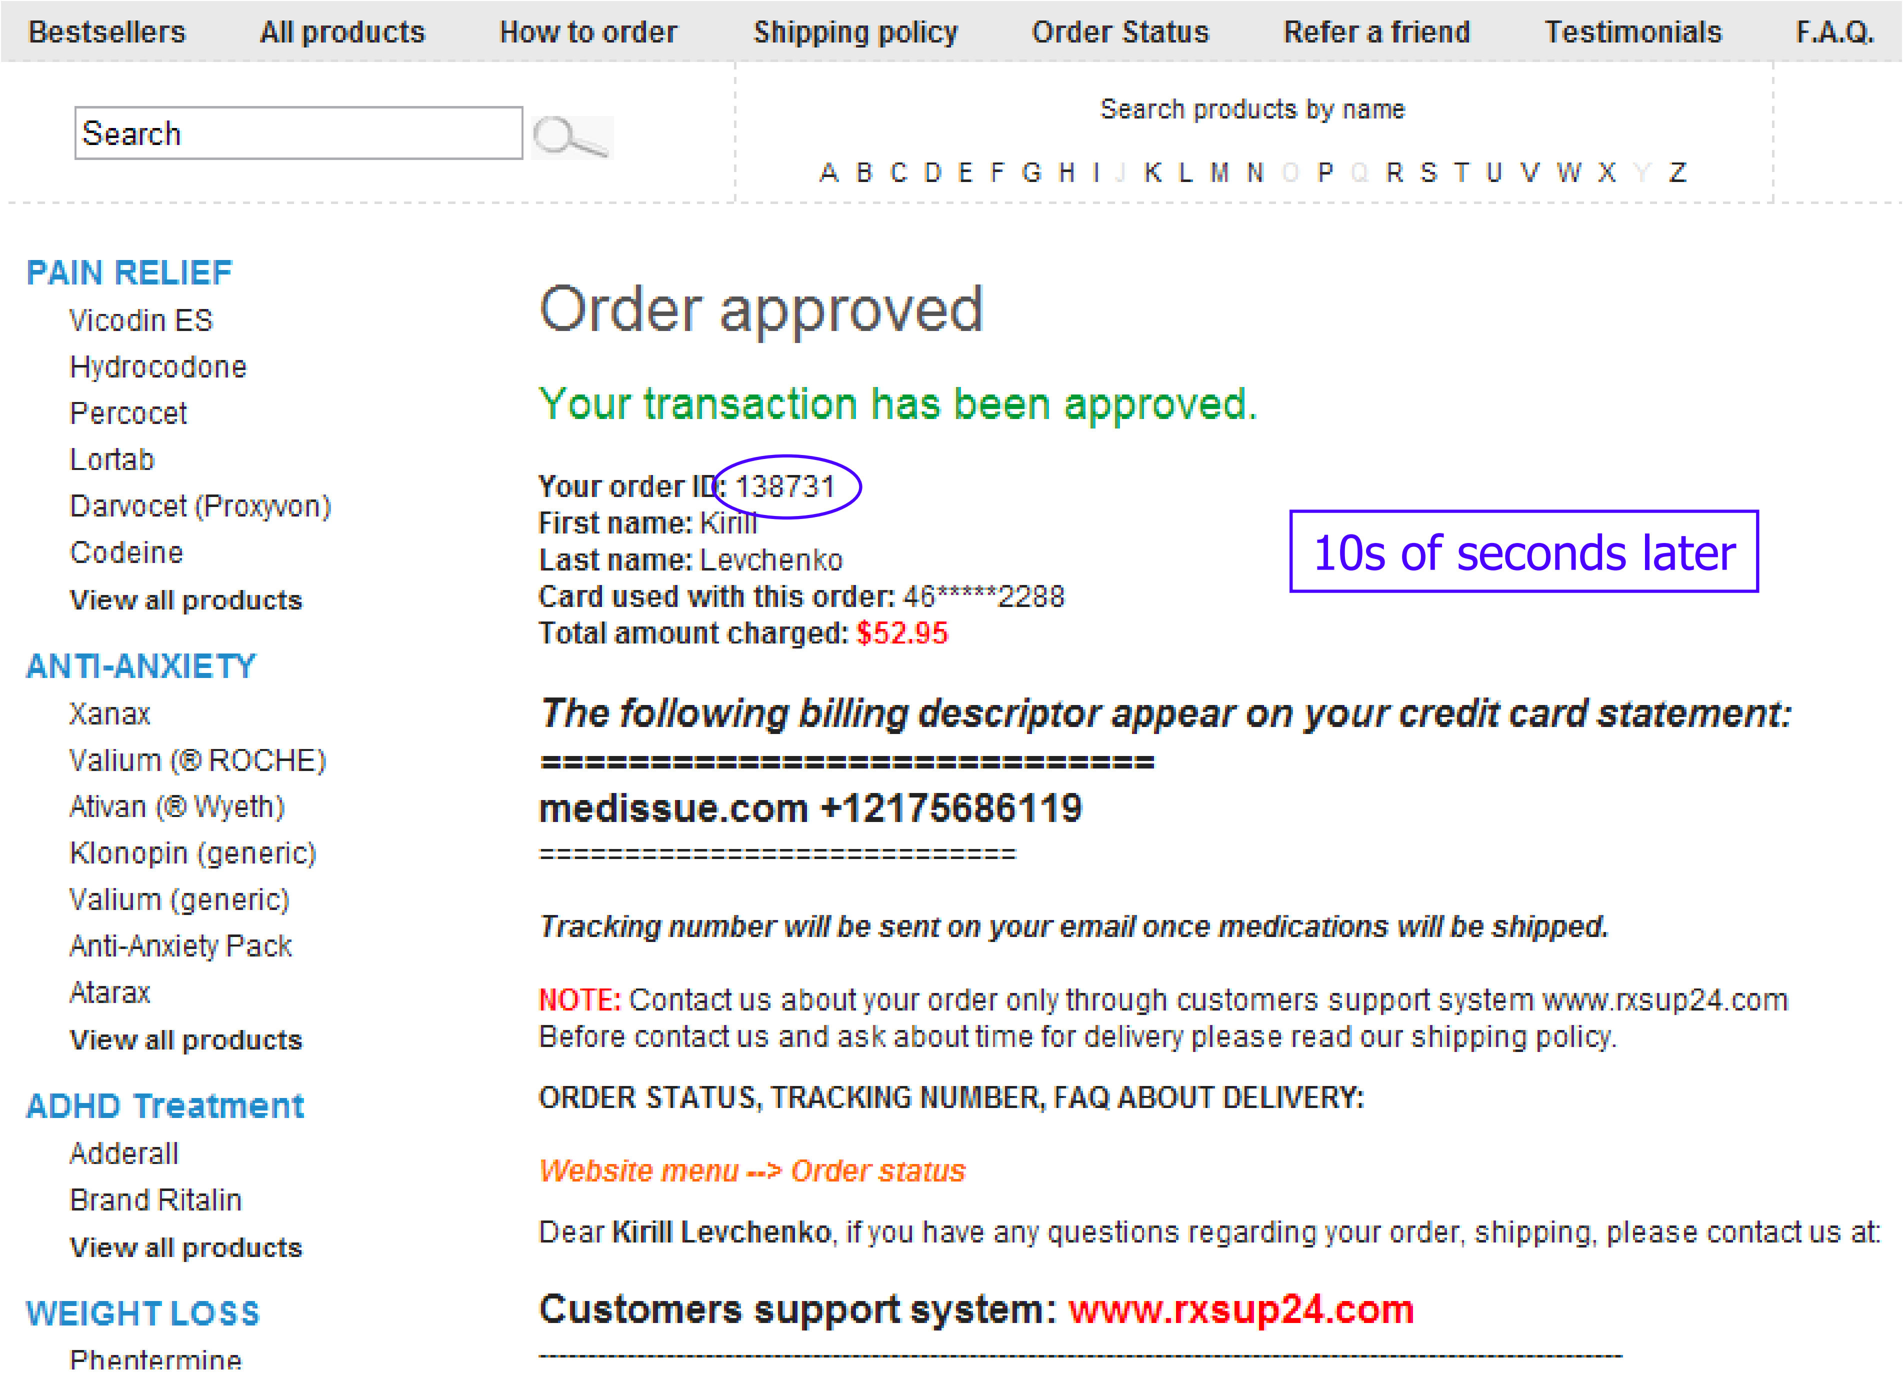

Order approved

Your transaction has been approved.

Your order ID: 138731
First name: Kirill
Last name: Levchenko
Card used with this order: 46*****2288
Total amount charged: $52.95

10s of seconds later

The following billing descriptor appear on your credit card statement:

================================

medissue.com +12175686119

============================

Tracking number will be sent on your email once medications will be shipped.

NOTE: Contact us about your order only through customers support system www.rxsup24.com
Before contact us and ask about time for delivery please read our shipping policy.

ORDER STATUS, TRACKING NUMBER, FAQ ABOUT DELIVERY:

Website menu --> Order status

Dear **Kirill Levchenko**, if you have any questions regarding your order, shipping, please contact us at:

Customers support system: www.rxsup24.com

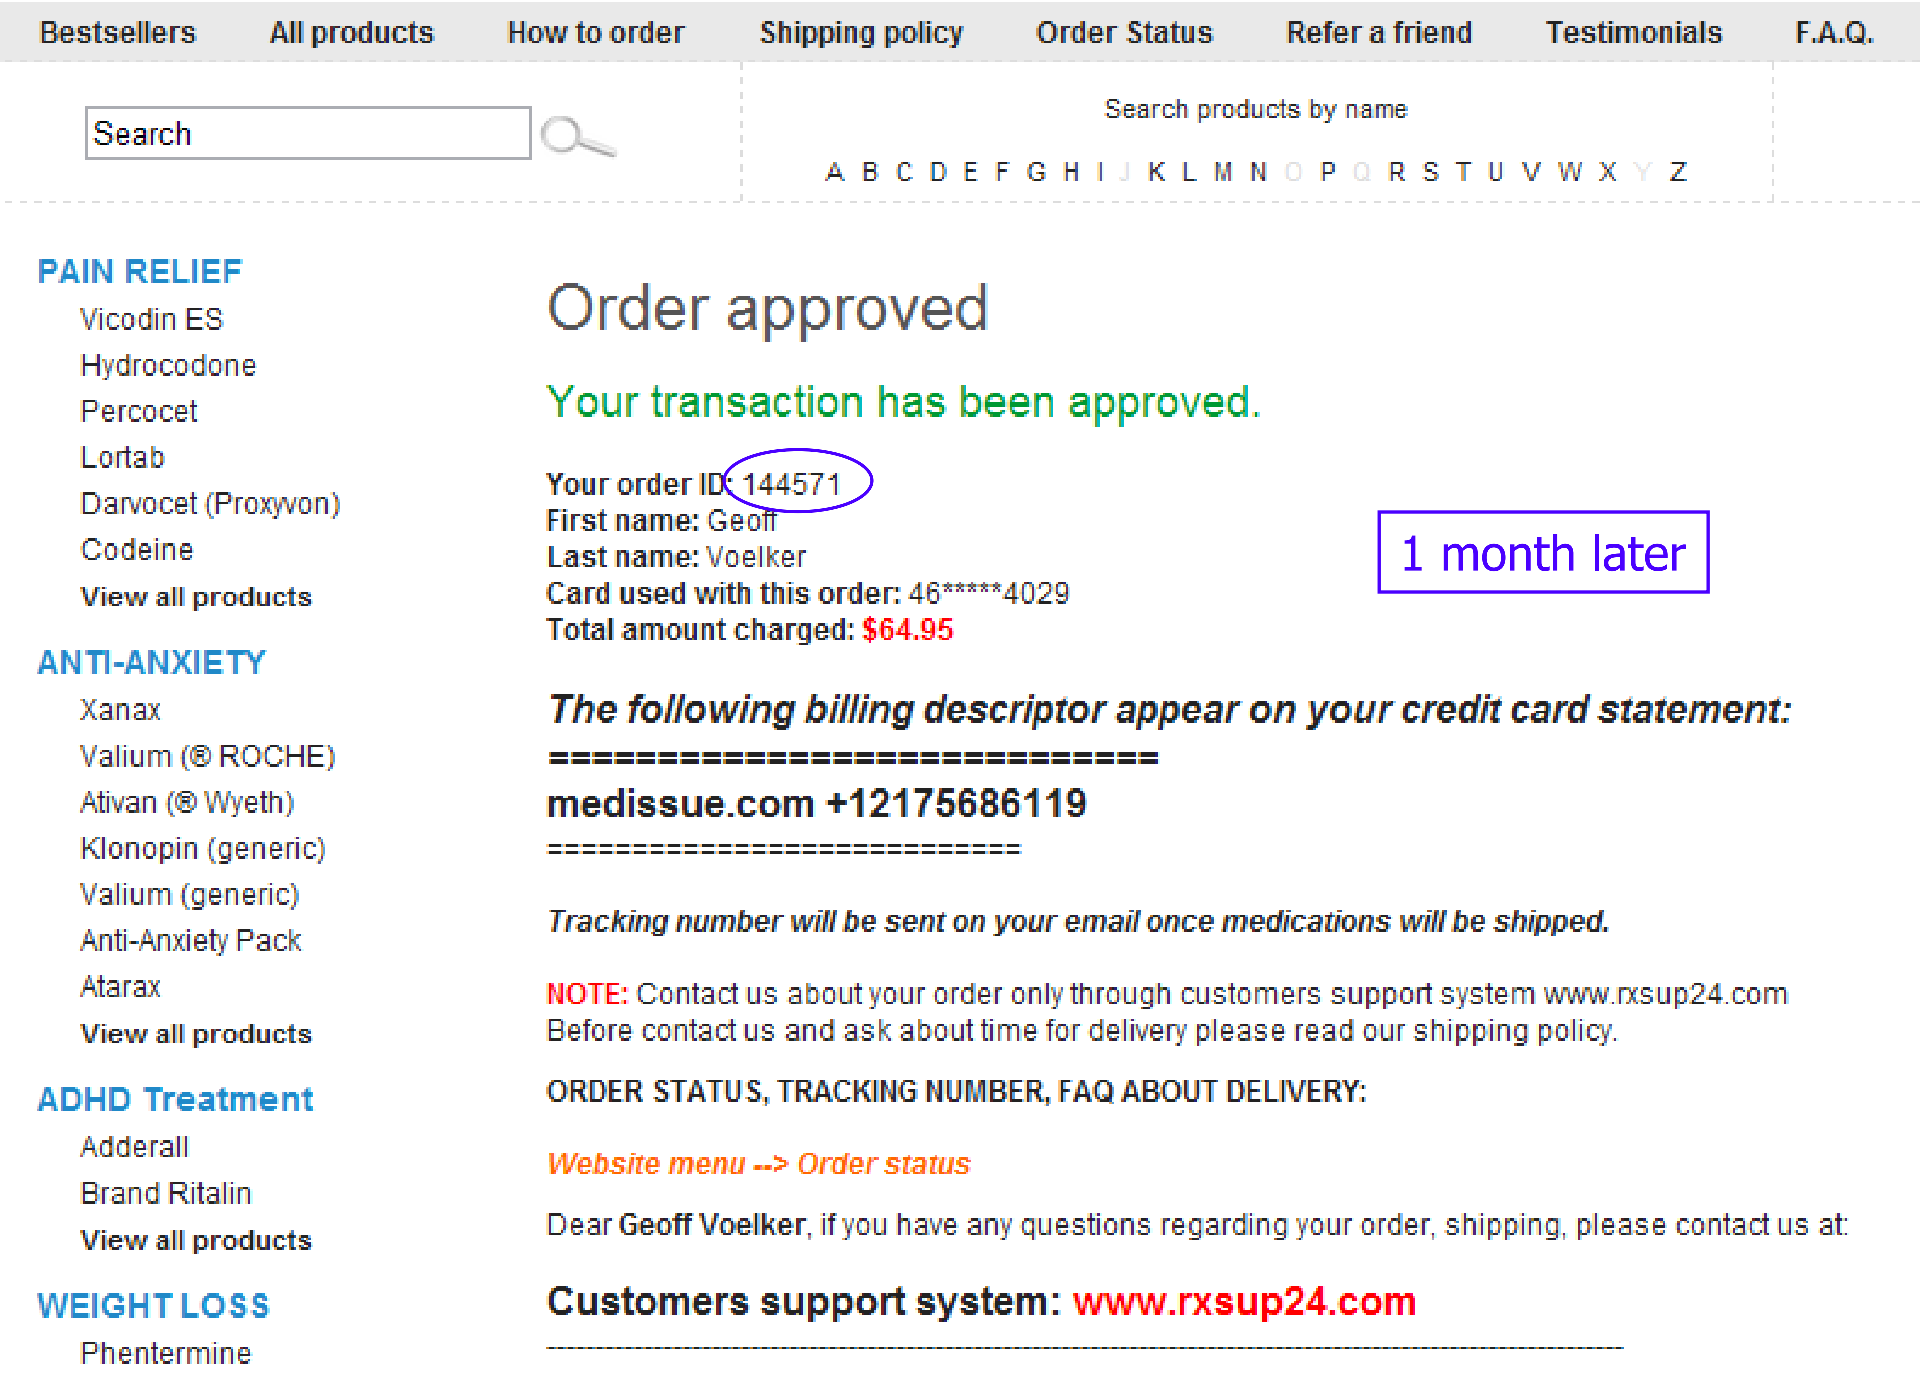

Order approved

Your transaction has been approved.

1 month later

Your order ID: 144571
First name: Geoff
Last name: Voelker
Card used with this order: 46*****4029
Total amount charged: $64.95

The following billing descriptor appear on your credit card statement:

================================

medissue.com +12175686119

============================

Tracking number will be sent on your email once medications will be shipped.

NOTE: Contact us about your order only through customers support system www.rxsup24.com
Before contact us and ask about time for delivery please read our shipping policy.

ORDER STATUS, TRACKING NUMBER, FAQ ABOUT DELIVERY:

Website menu --> Order status

Dear **Geoff Voelker**, if you have any questions regarding your order, shipping, please contact us at:

Customers support system: www.rxsup24.com

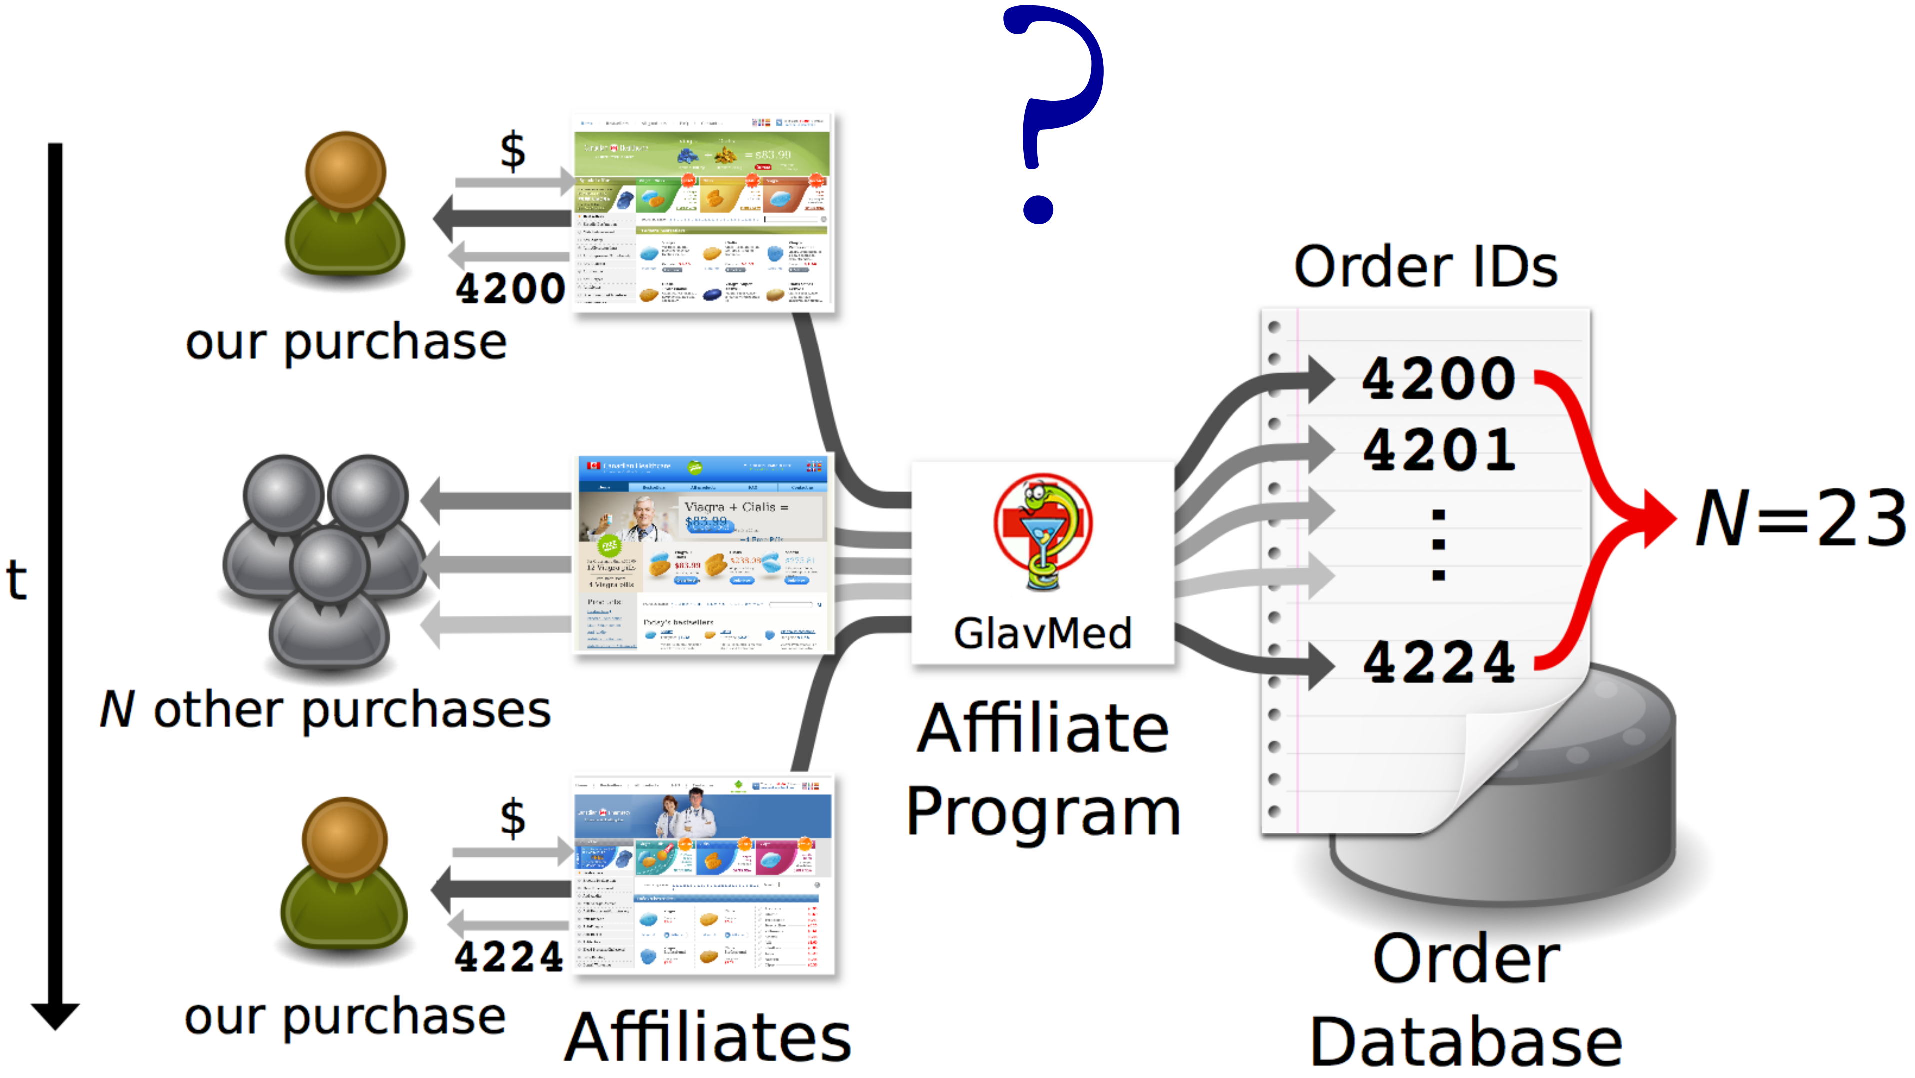

t

$

4200

our purchase

N other purchases

$

4224

our purchase

Affiliates

?

GlavMed

Affiliate Program

Order IDs

4200

4201

⋮

4224

N=23

Order Database

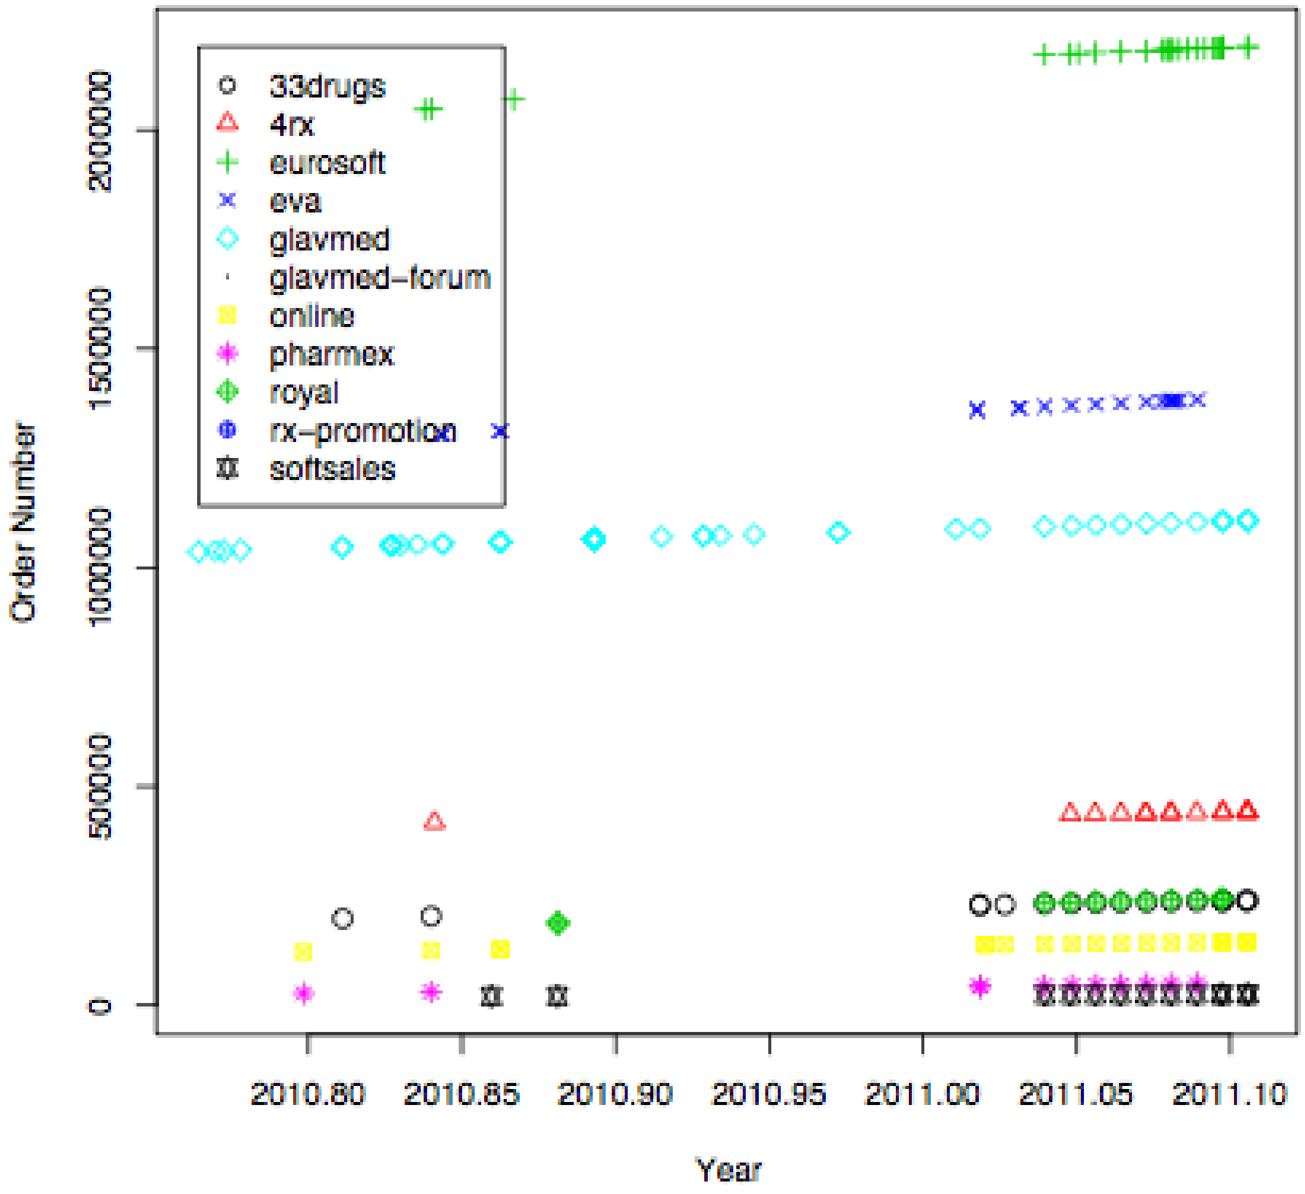
33drugs
4rx
eurosoft
eva
glavmed
glavmed–forum
online
pharmex
royal
rx–promotion
softsales
Order Number
0
500000
1000000
1500000
2000000
2010.80
2010.85
2010.90
2010.95
2011.00
2011.05
2011.10
Year

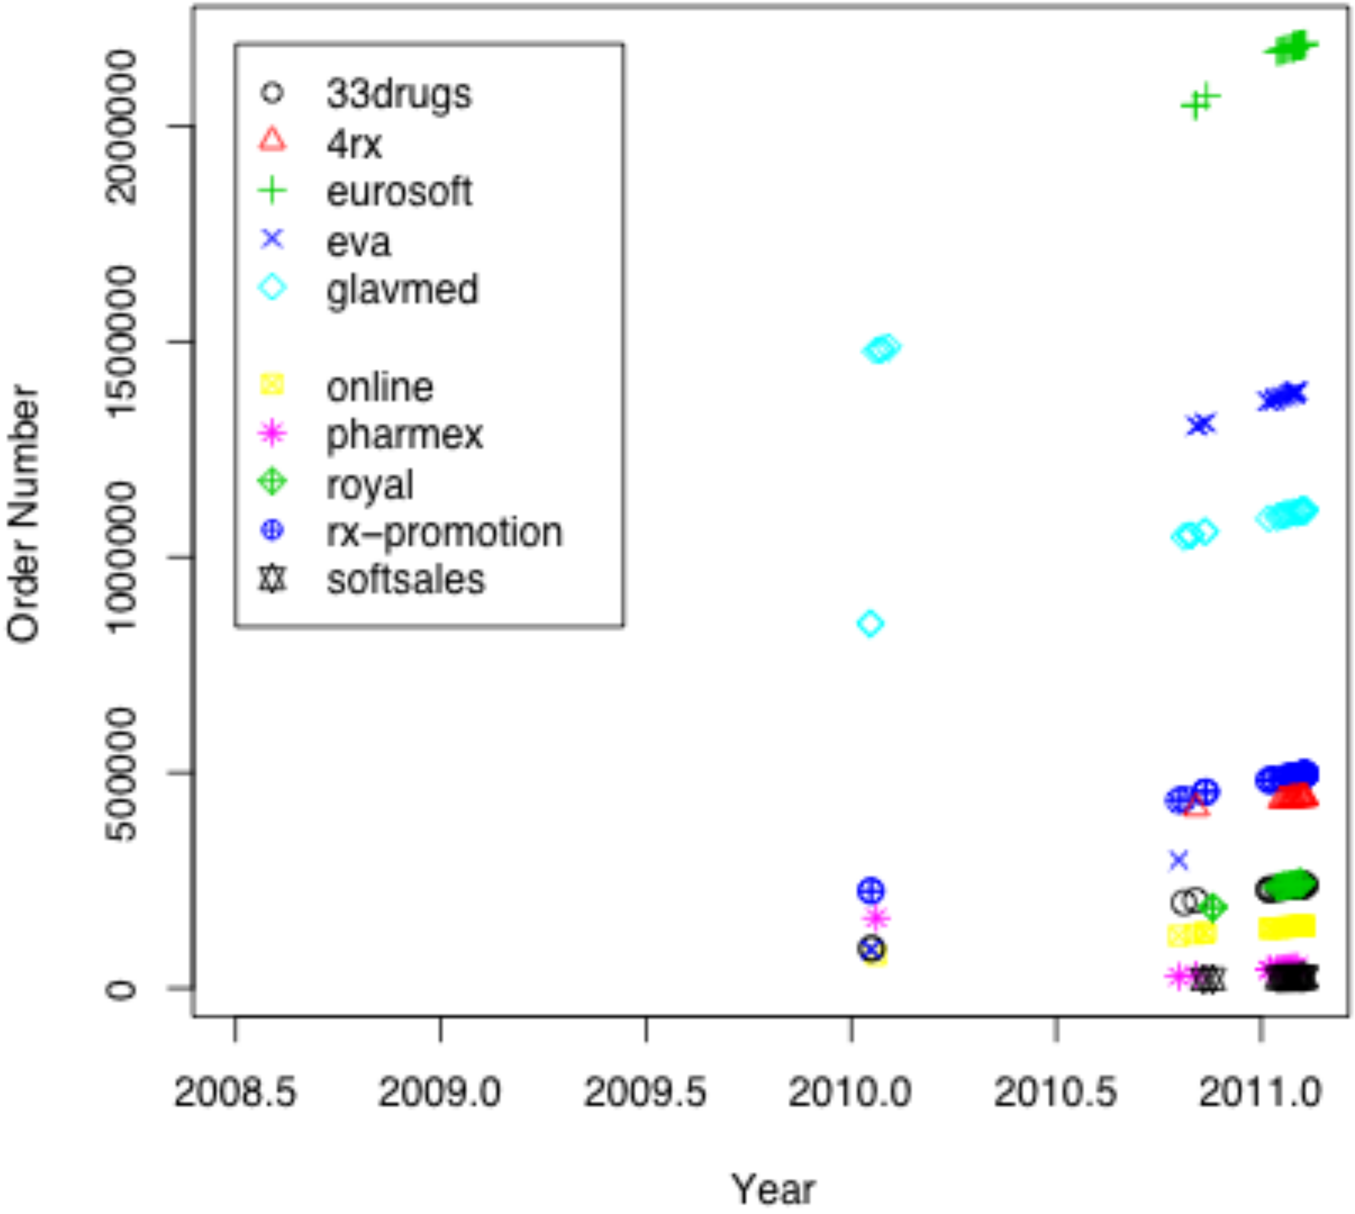

○ 33drugs
△ 4rx
+ eurosoft
× eva
◇ glavmed
⊠ online
✳ pharmex
◇ royal
⊕ rx–promotion
✡ softsales
Order Number
0
500000
1000000
1500000
2000000
2008.5
2009.0
2009.5
2010.0
2010.5
2011.0
Year

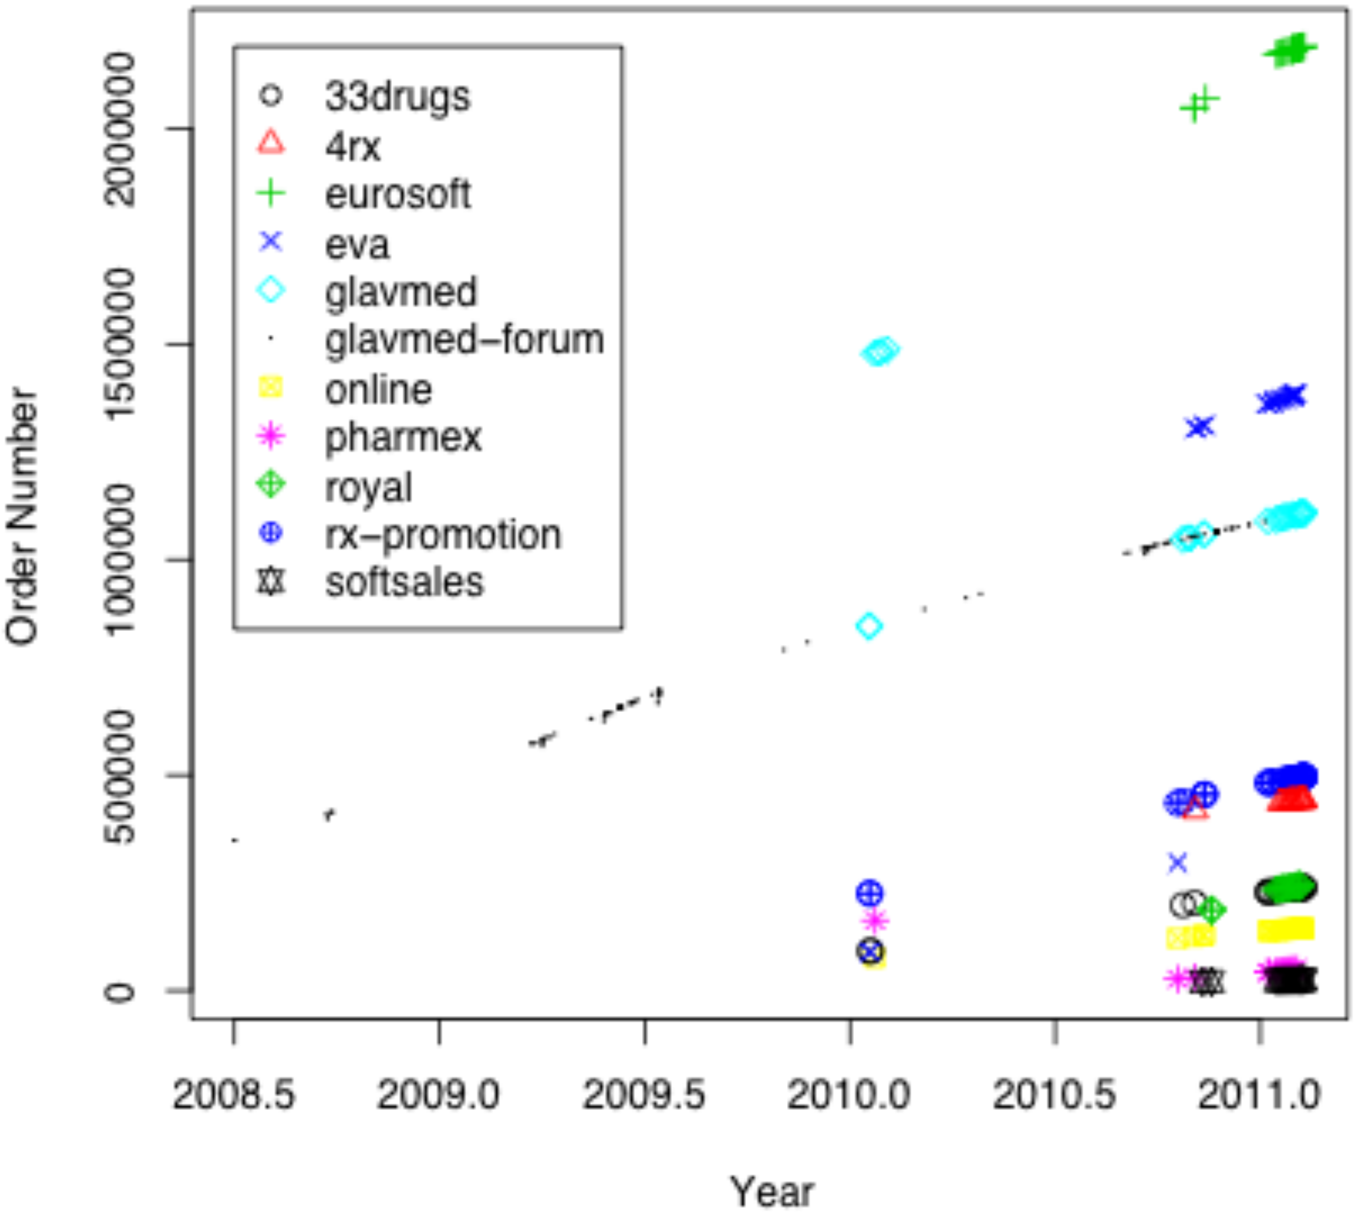
33drugs
4rx
eurosoft
eva
glavmed
glavmed–forum
online
pharmex
royal
rx–promotion
softsales
Order Number
0
500000
1000000
1500000
2000000
2008.5
2009.0
2009.5
2010.0
2010.5
2011.0
Year

Order-Driven or Timer-Driven?

- How to definitively establish whether IP-ID inference trick works for pharma orders?
- Alternative hypothesis also fitting with monotonicity: order #s are timer-based
  - Perhaps in batches, hence nearby ones just go up by one rather than by time interval
- Approach: take a month of data
  - Compute *least-squares fit* of order # vs. time
  - For each order, look at its offset from the fit
    - Reflects either measurement error (we didn't record time right) …
    - … or required batching granularity for consistent measurement

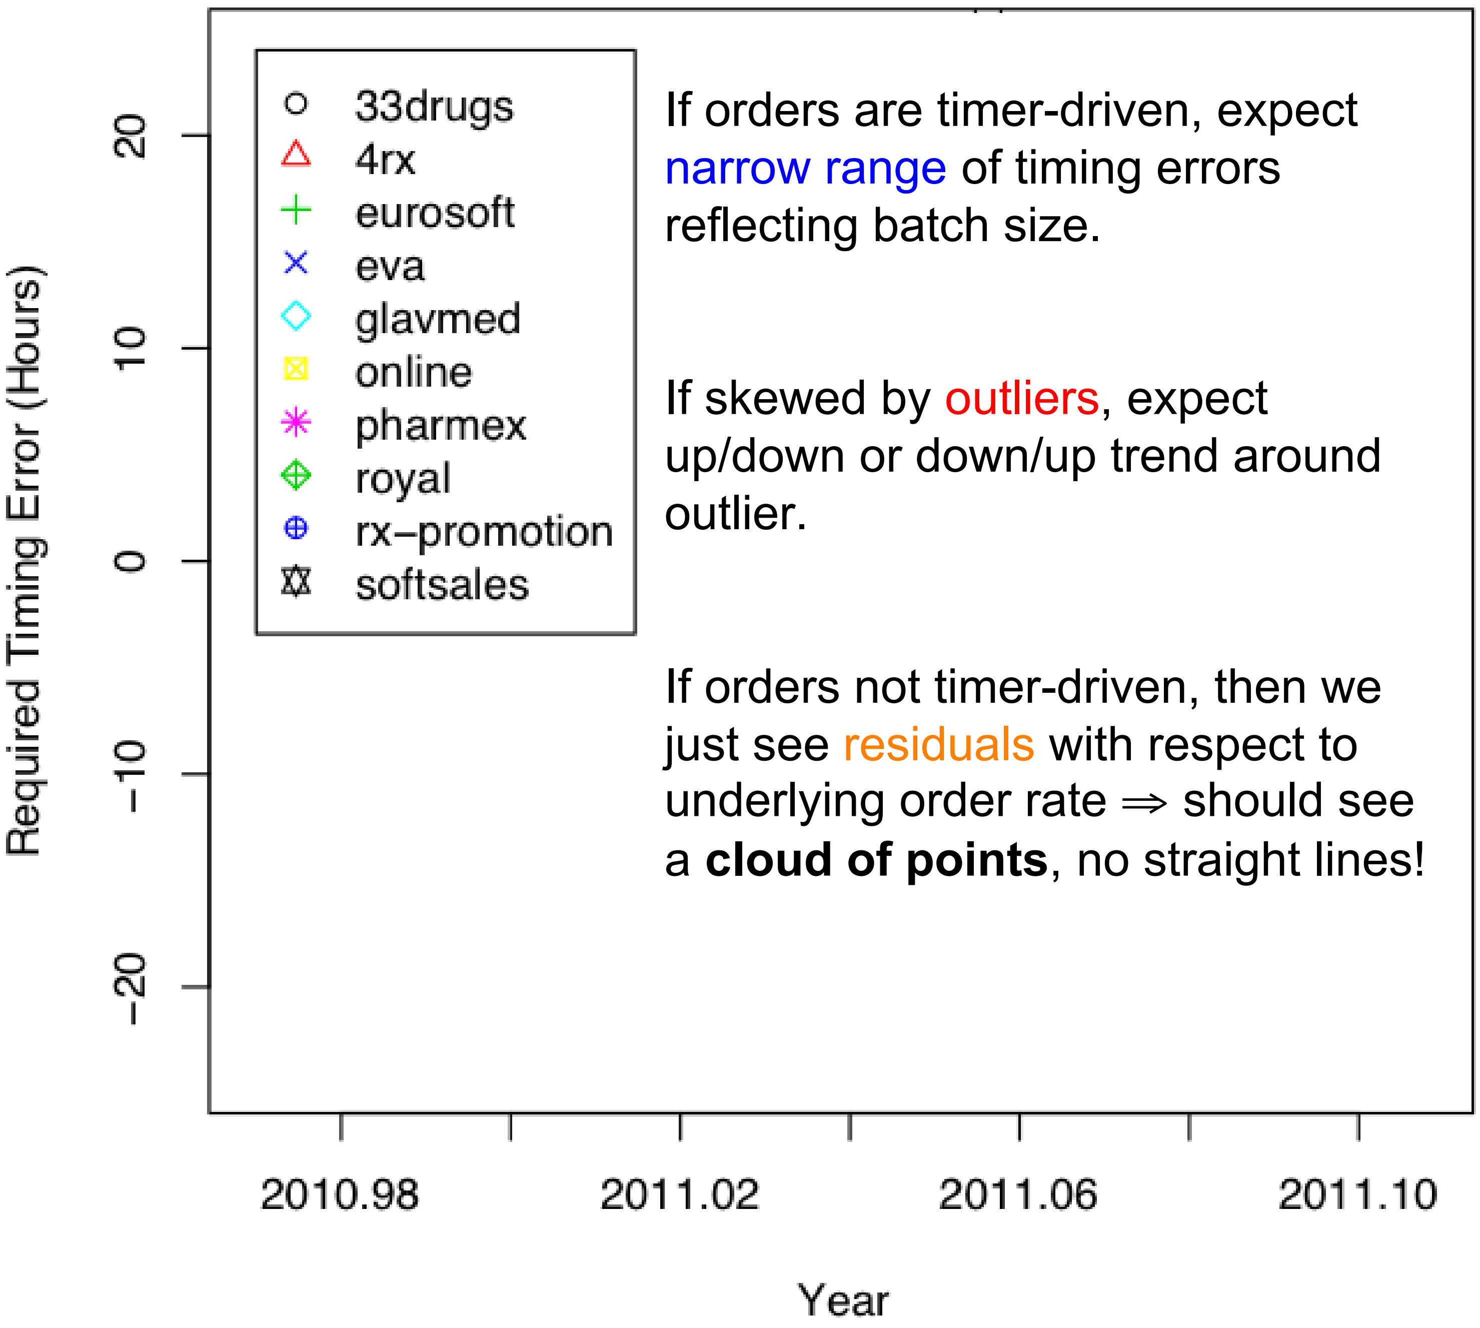

If orders are timer-driven, expect narrow range of timing errors reflecting batch size.

If skewed by outliers, expect up/down or down/up trend around outlier.

If orders not timer-driven, then we just see residuals with respect to underlying order rate ⇒ should see a **cloud of points**, no straight lines!

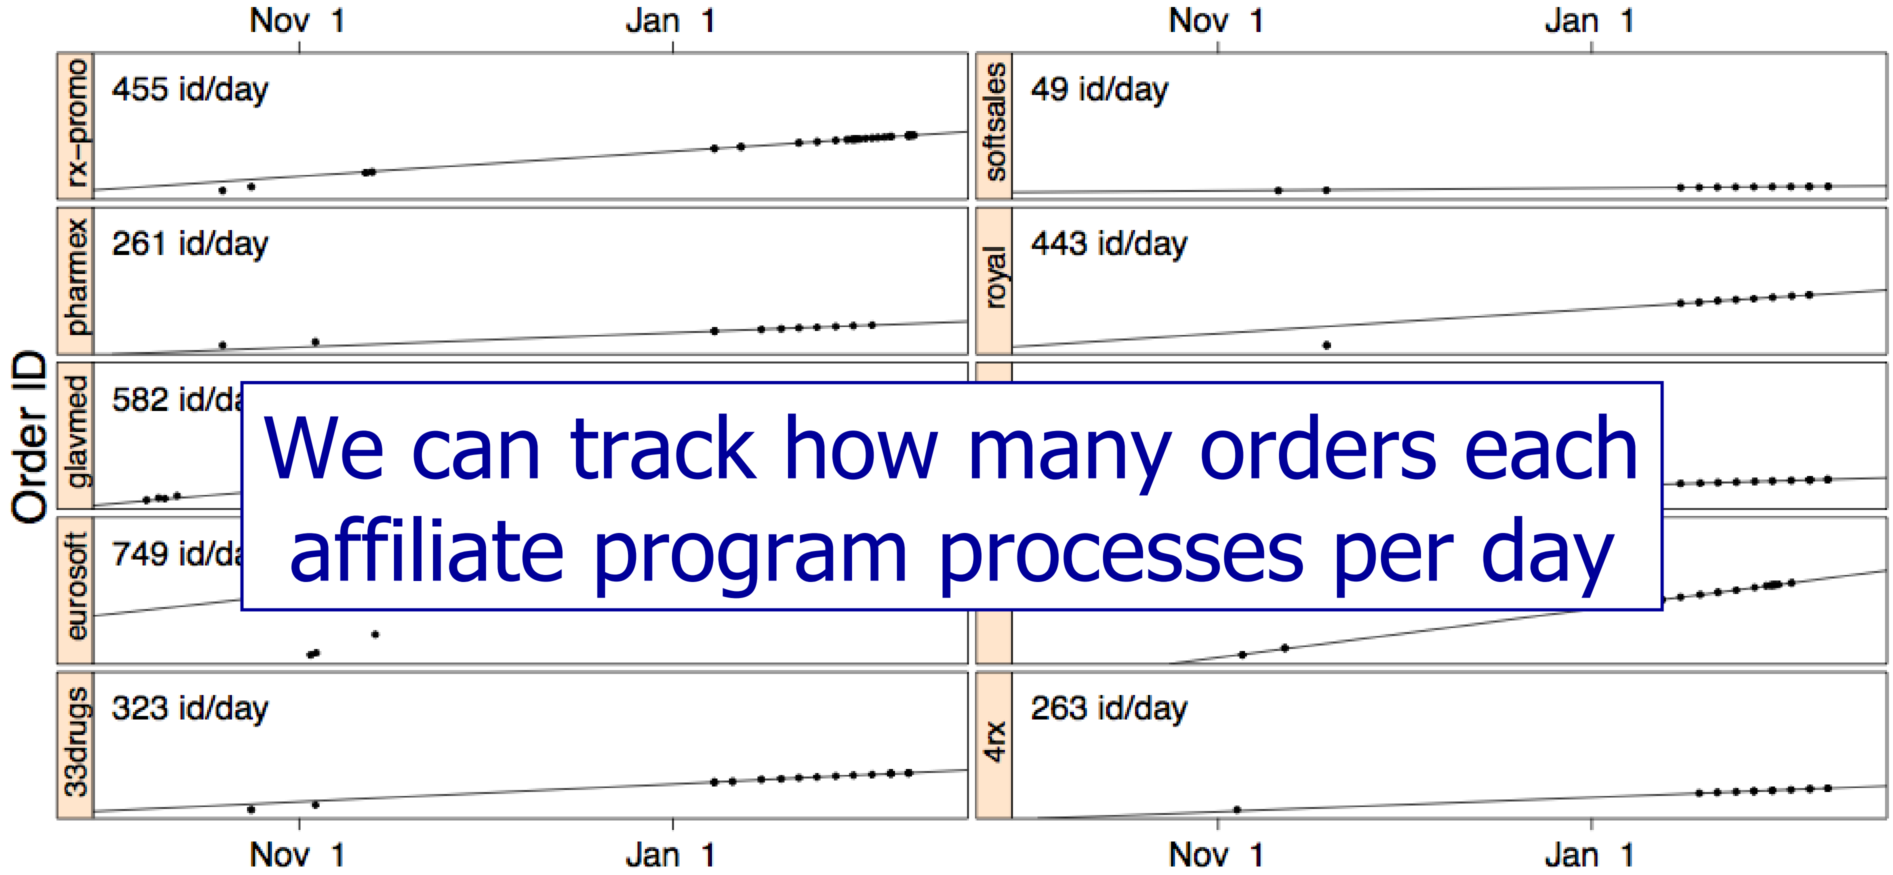

We can track how many orders each affiliate program processes per day

Coupled with a (separate) *structural inference* ⇒ entire spam "pharma" revenue ≈ $50-100M/year

Part III:

Pretty Pictures & Moneyshots

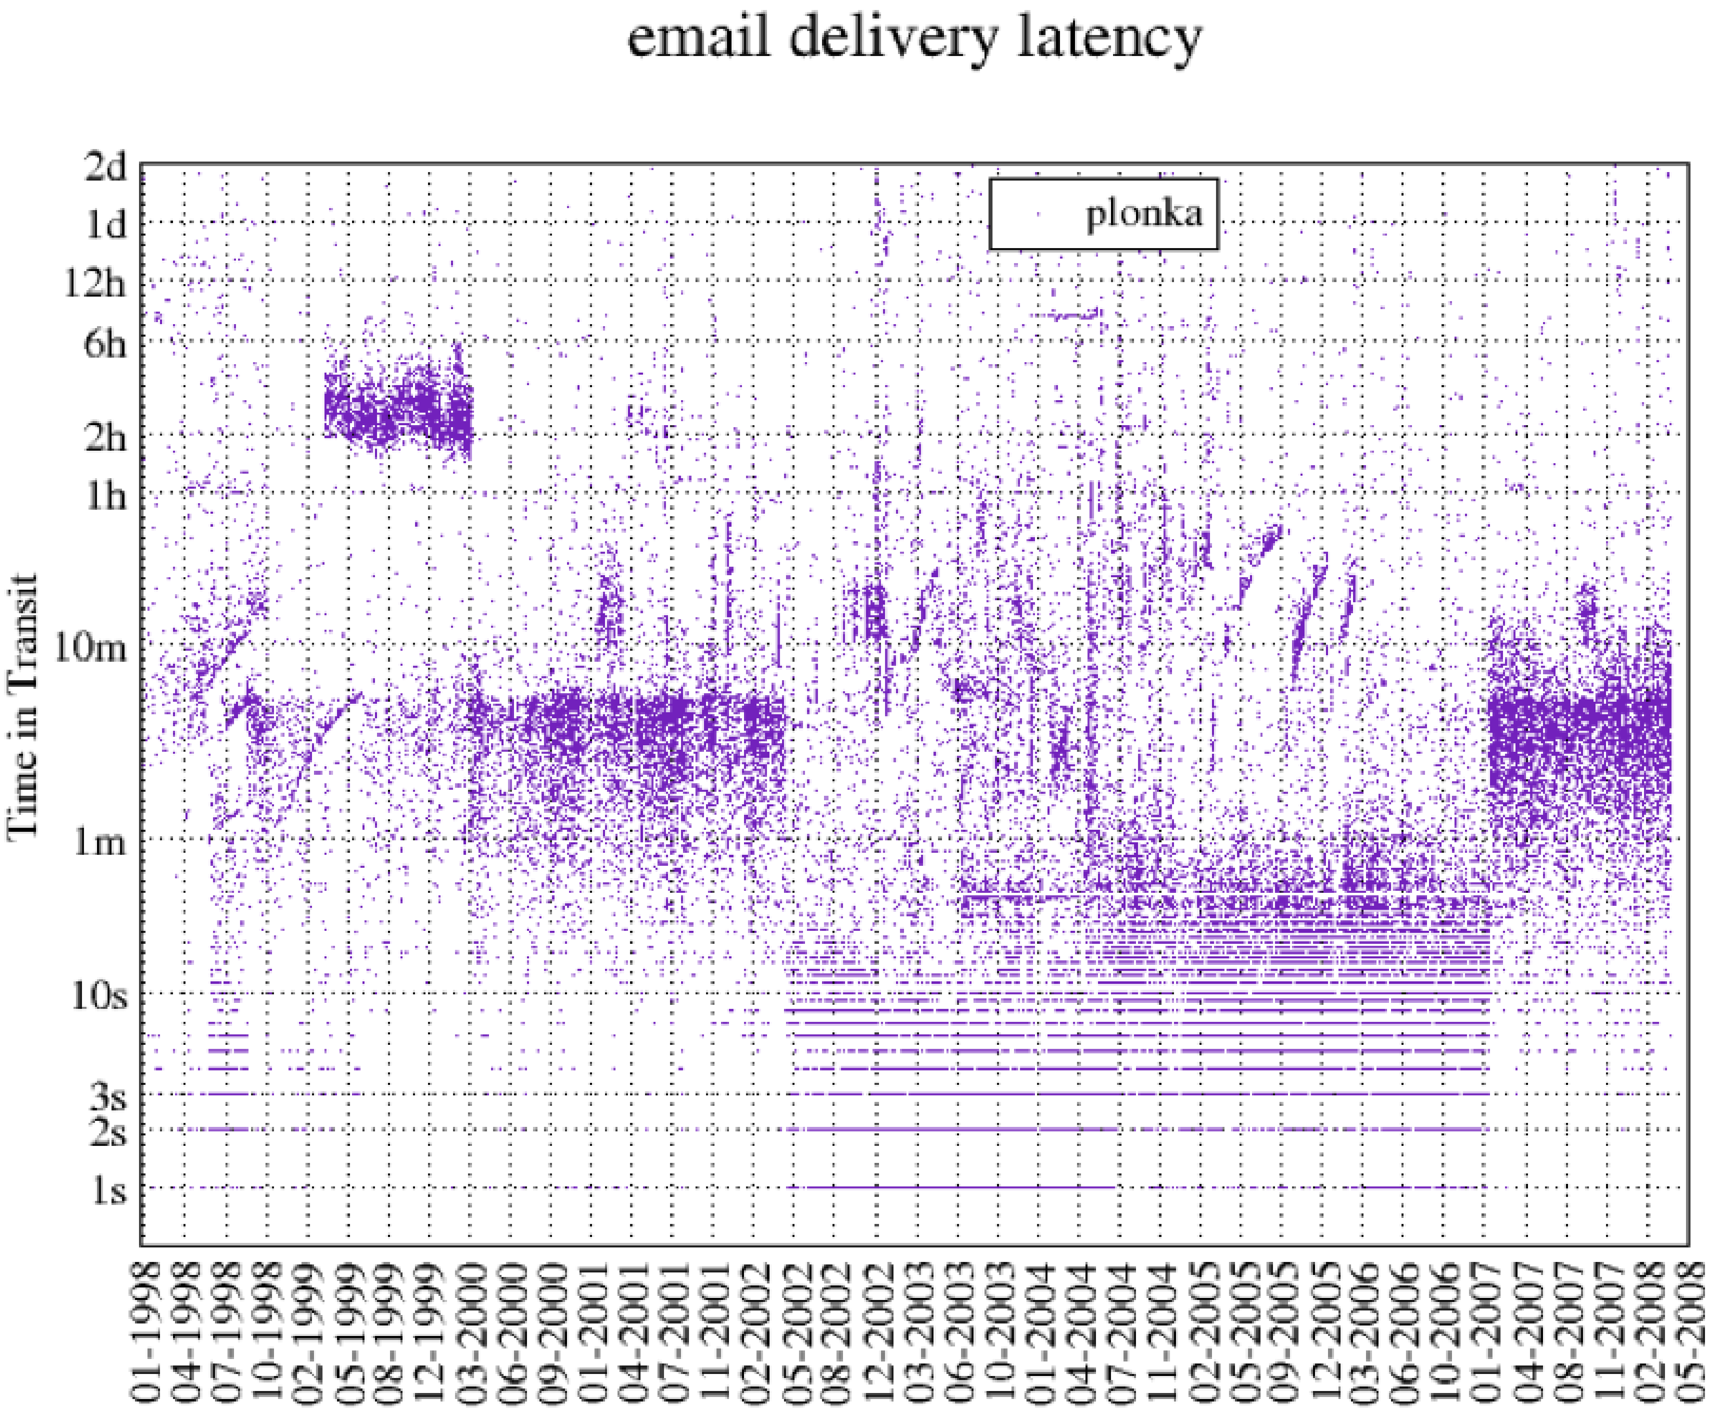
email delivery latency
plonka
Time in Transit
2d
1d
12h
6h
2h
1h
10m
1m
10s
3s
2s
1s
01-1998
04-1998
07-1998
10-1998
02-1999
05-1999
08-1999
12-1999
03-2000
06-2000
09-2000
01-2001
04-2001
07-2001
11-2001
02-2002
05-2002
08-2002
12-2002
03-2003
06-2003
10-2003
01-2004
04-2004
07-2004
11-2004
02-2005
05-2005
09-2005
12-2005
03-2006
06-2006
10-2006
01-2007
04-2007
08-2007
11-2007
02-2008
05-2008

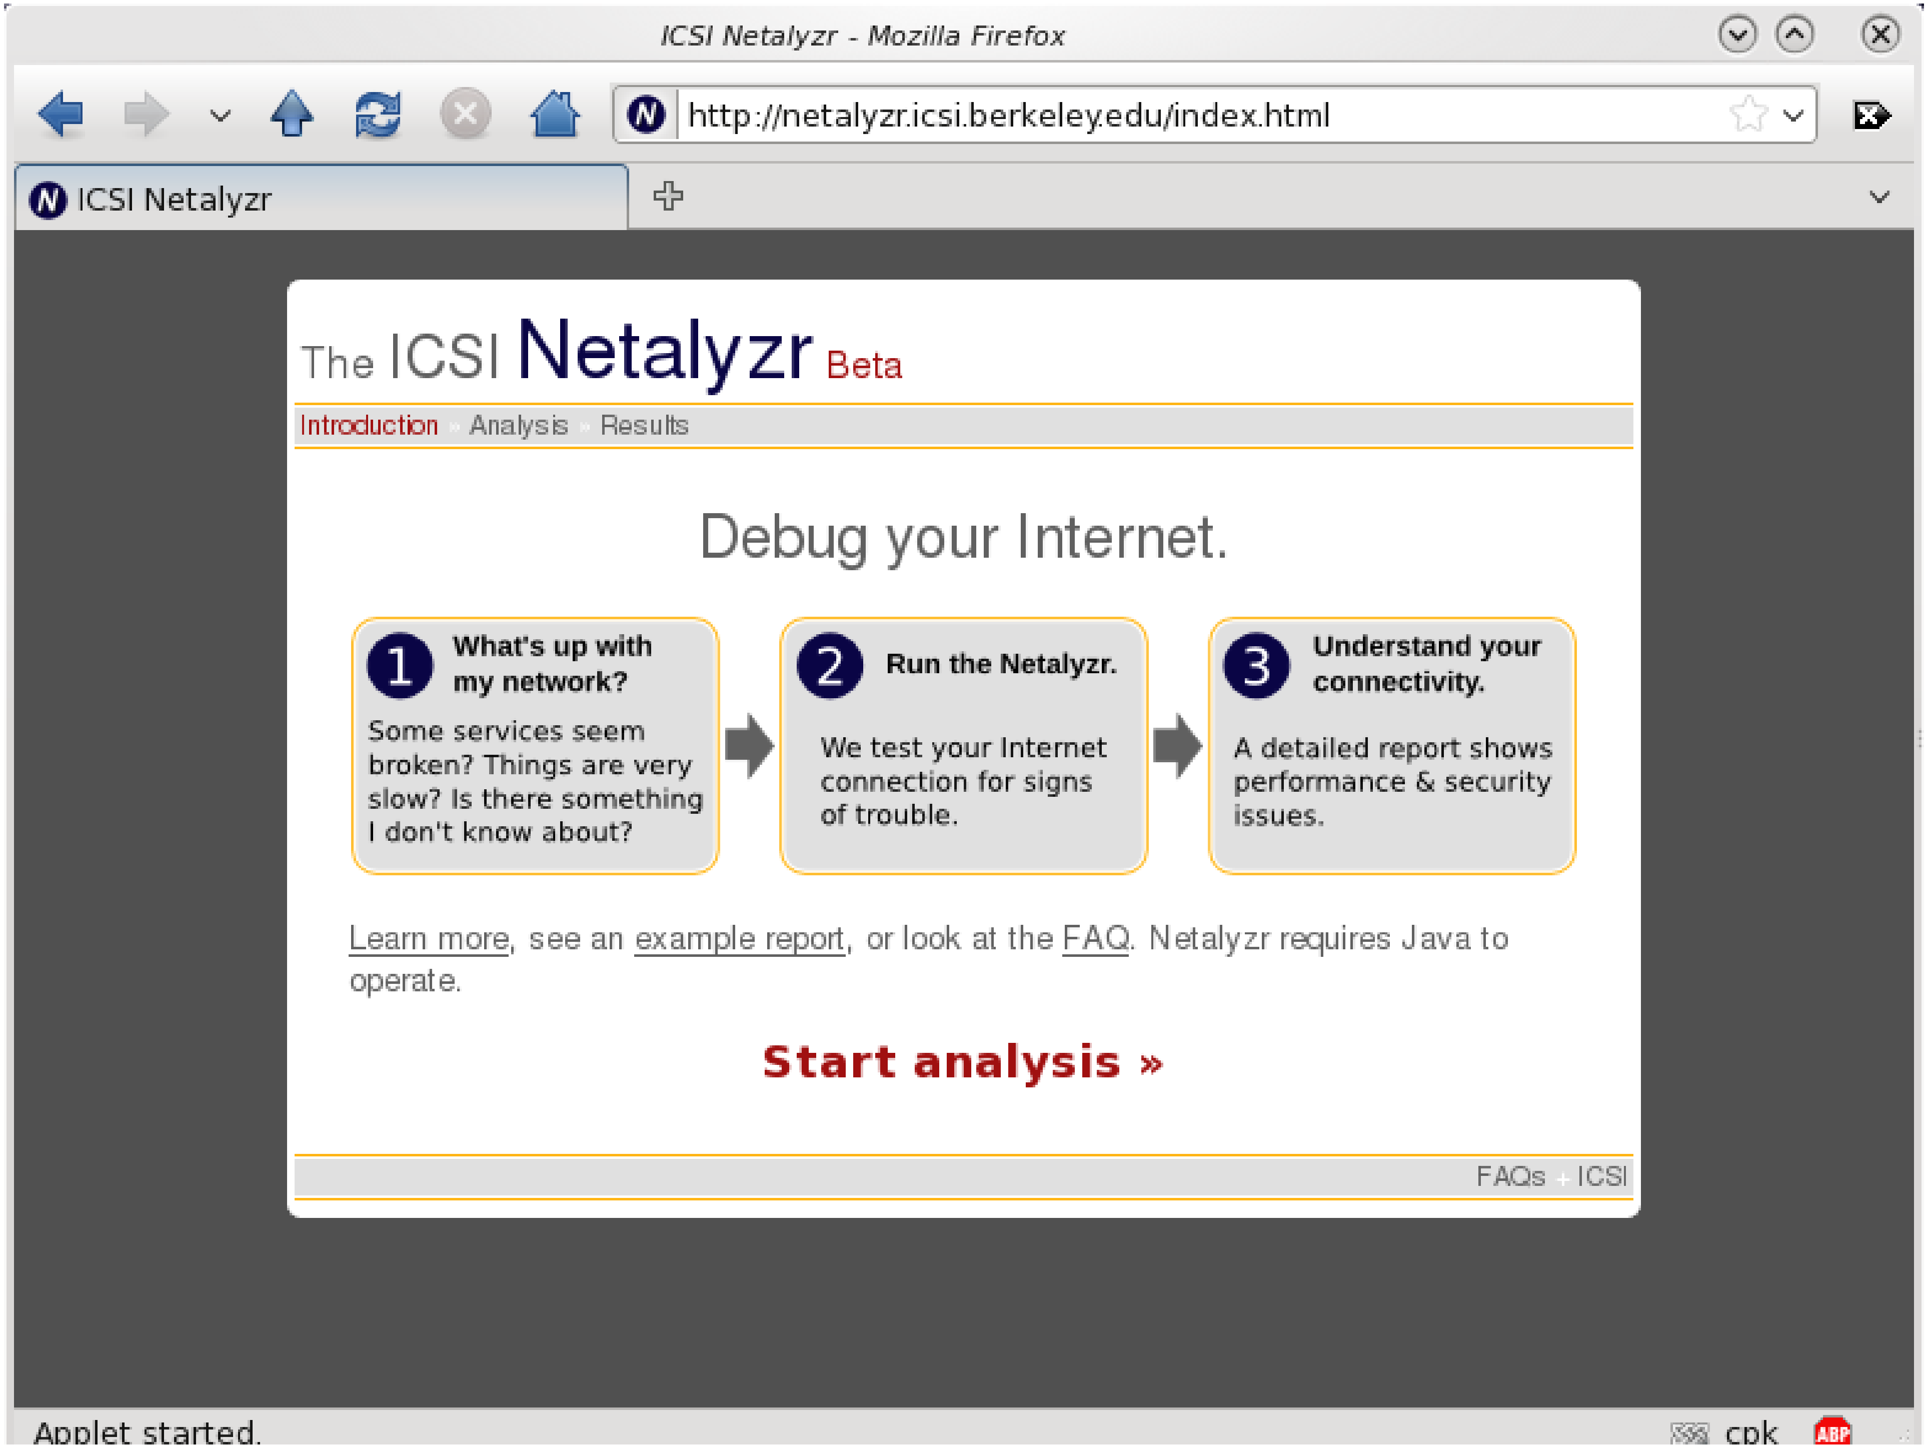
ICSI Netalyzr - Mozilla Firefox
http://netalyzr.icsi.berkeley.edu/index.html
ICSI Netalyzr

The ICSI **Netalyzr** Beta

Introduction » Analysis » Results

Debug your Internet.

1 **What's up with my network?**

Some services seem broken? Things are very slow? Is there something I don't know about?

2 **Run the Netalyzr.**

We test your Internet connection for signs of trouble.

3 **Understand your connectivity.**

A detailed report shows performance & security issues.

Learn more, see an example report, or look at the FAQ. Netalyzr requires Java to operate.

Start analysis »

Applet started.

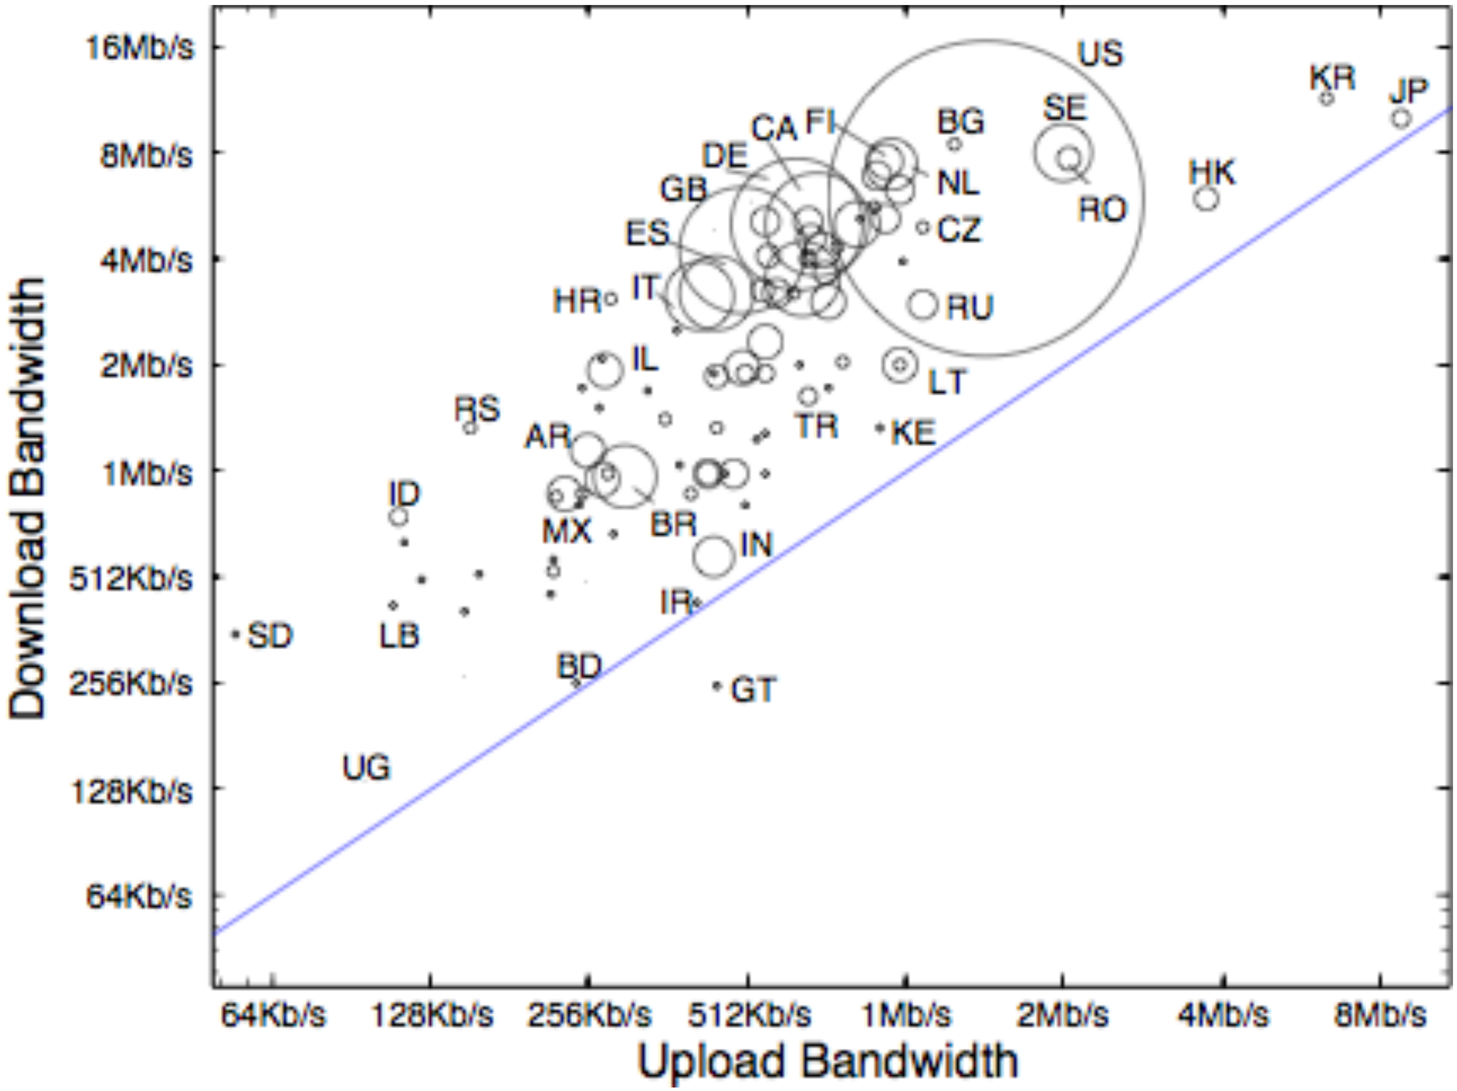

Download Bandwidth
Upload Bandwidth
16Mb/s
8Mb/s
4Mb/s
2Mb/s
1Mb/s
512Kb/s
256Kb/s
128Kb/s
64Kb/s
64Kb/s
128Kb/s
256Kb/s
512Kb/s
1Mb/s
2Mb/s
4Mb/s
8Mb/s
US
KR
JP
SE
BG
FI
CA
DE
NL
RO
HK
GB
ES
CZ
IT
HR
RU
IL
LT
RS
AR
TR
KE
ID
MX
BR
IN
IR
SD
LB
BD
GT
UG

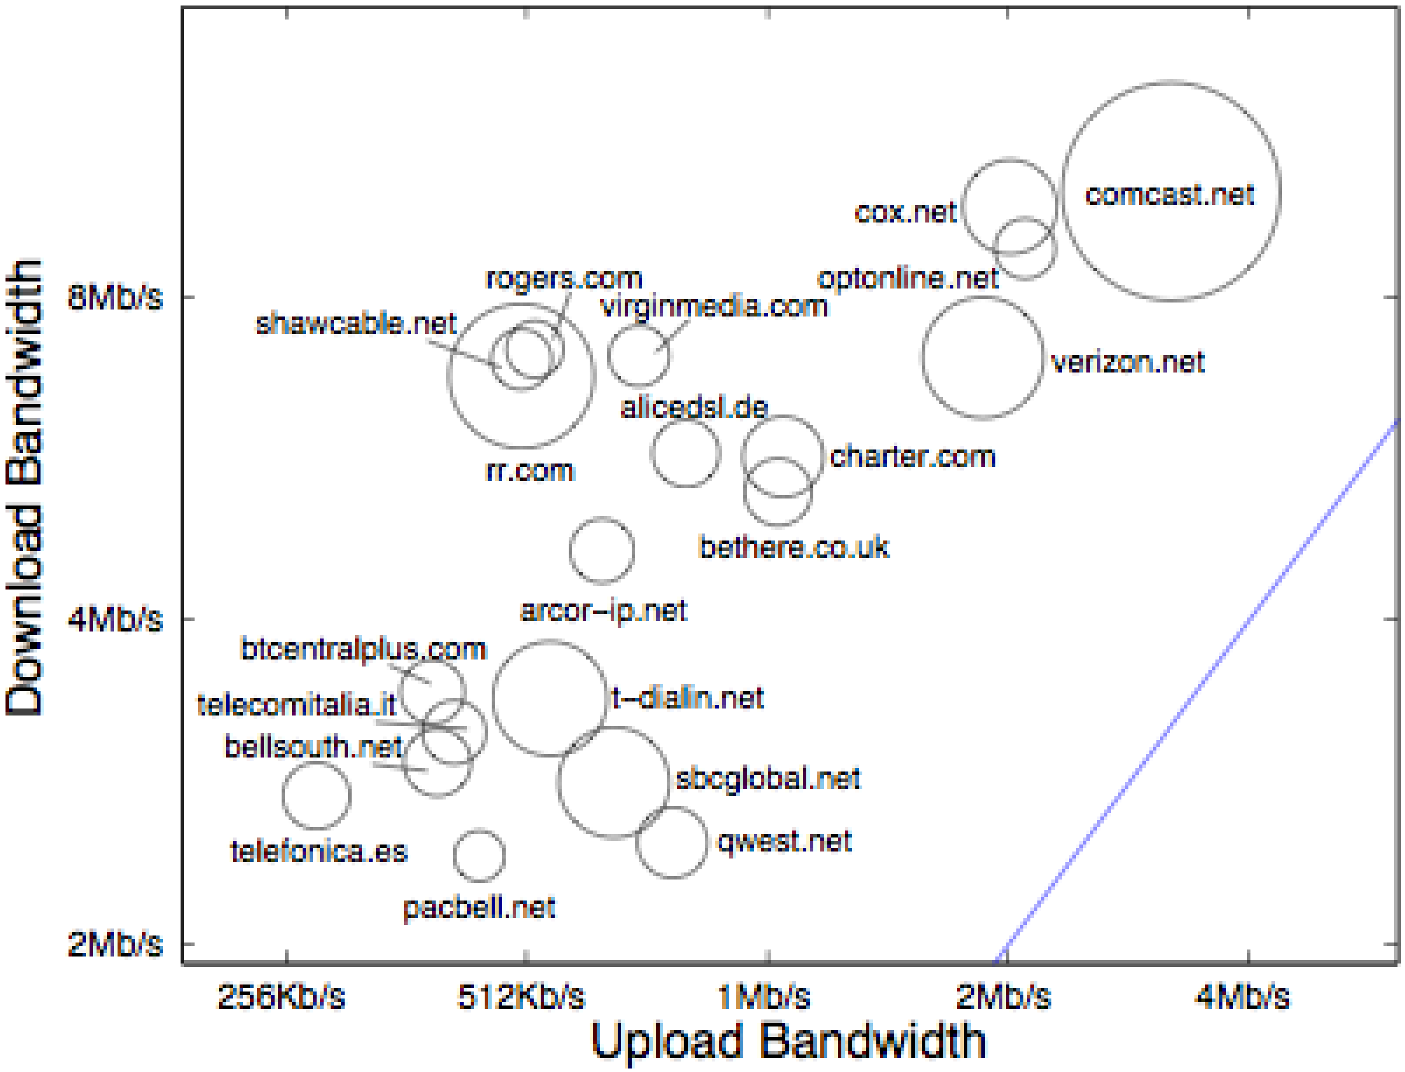
Download Bandwidth
Upload Bandwidth
2Mb/s
4Mb/s
8Mb/s
256Kb/s
512Kb/s
1Mb/s
2Mb/s
4Mb/s
comcast.net
cox.net
optonline.net
verizon.net
rogers.com
virginmedia.com
shawcable.net
rr.com
alicedsl.de
charter.com
bethere.co.uk
arcor–ip.net
btcentralplus.com
telecomitalia.it
bellsouth.net
telefonica.es
pacbell.net
t–dialin.net
sbcglobal.net
qwest.net

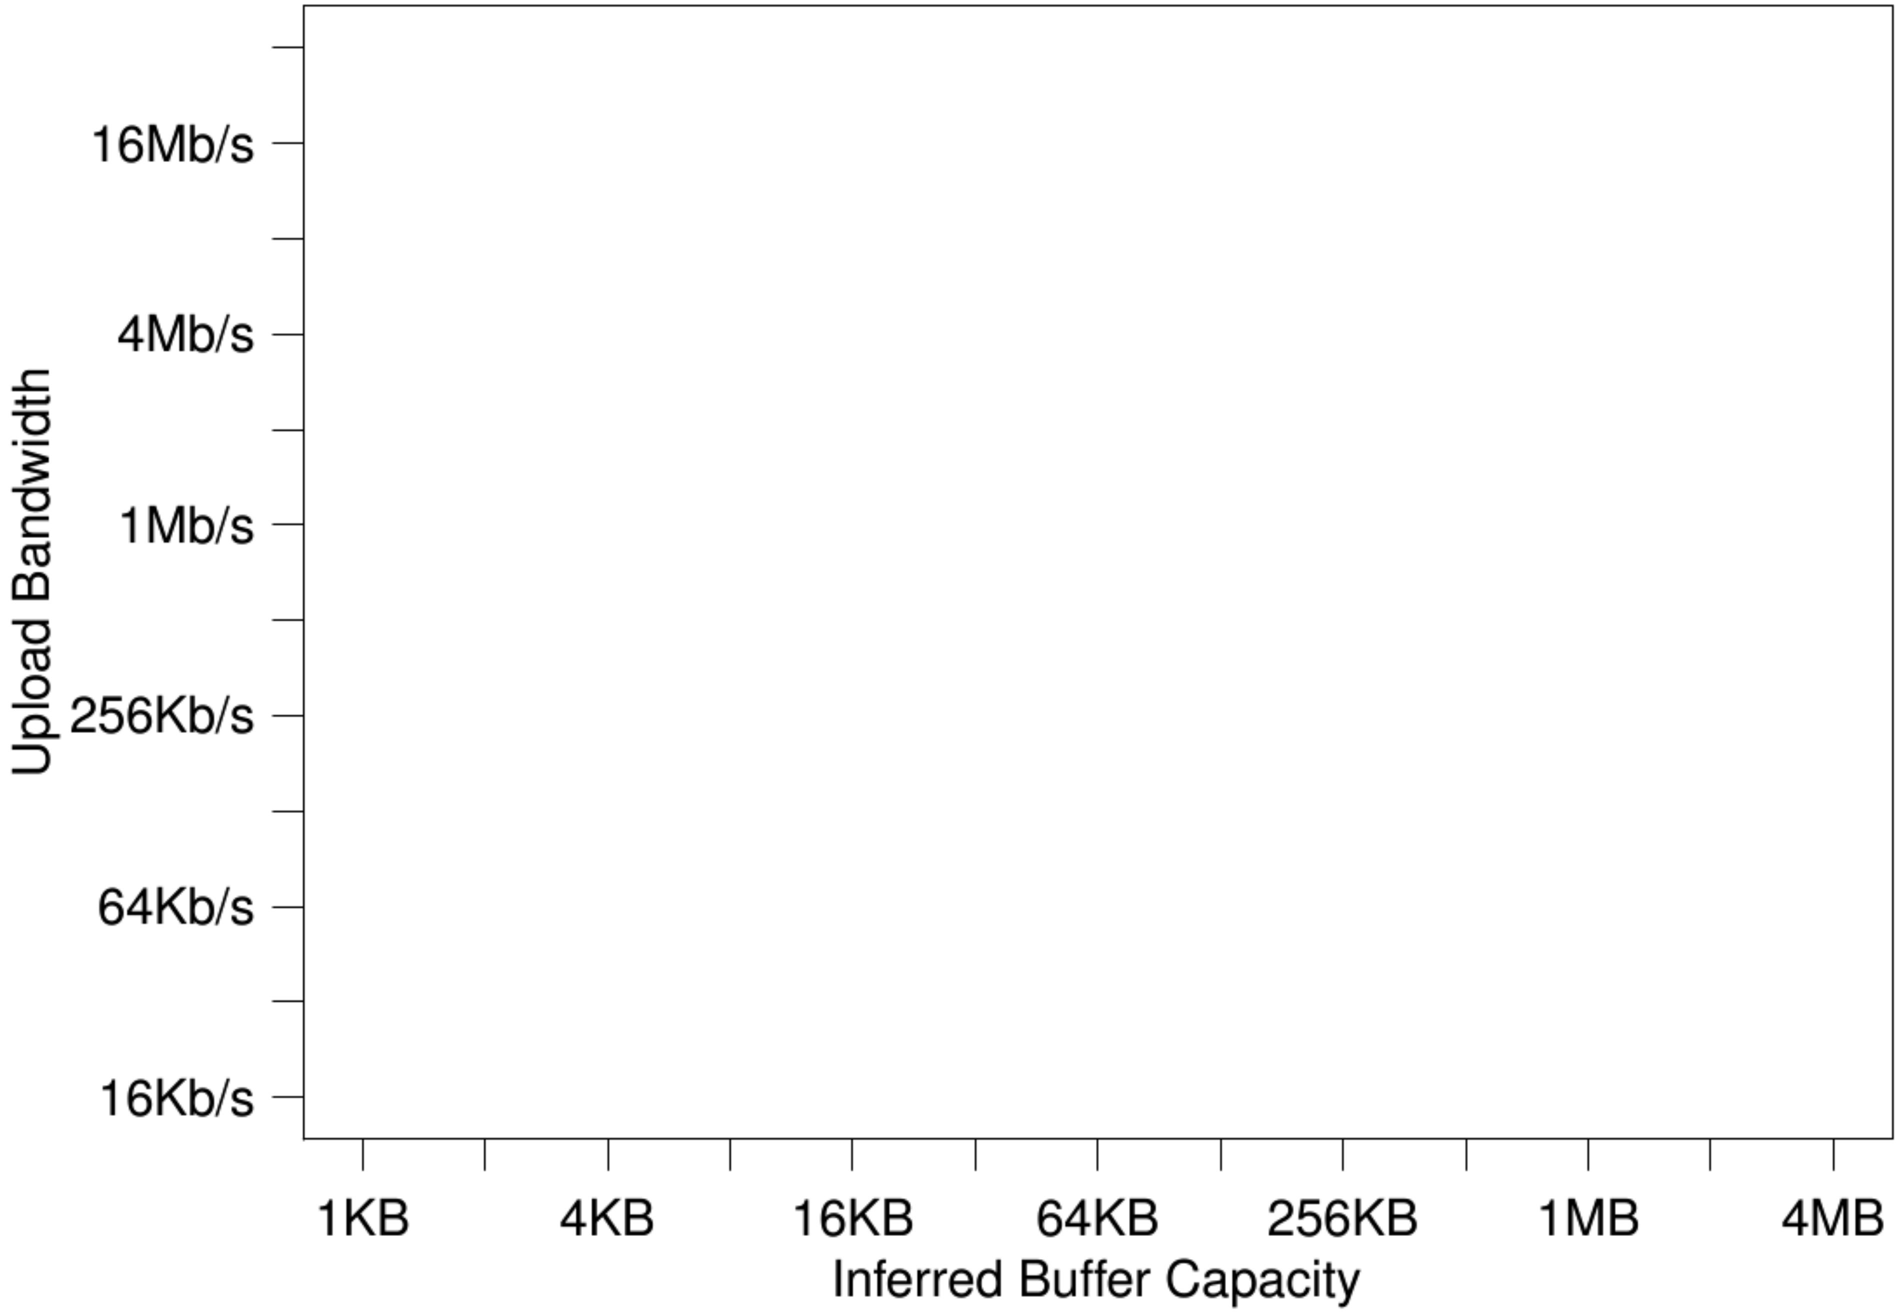
16Mb/s
4Mb/s
1Mb/s
256Kb/s
64Kb/s
16Kb/s
Upload Bandwidth
1KB
4KB
16KB
64KB
256KB
1MB
4MB
Inferred Buffer Capacity

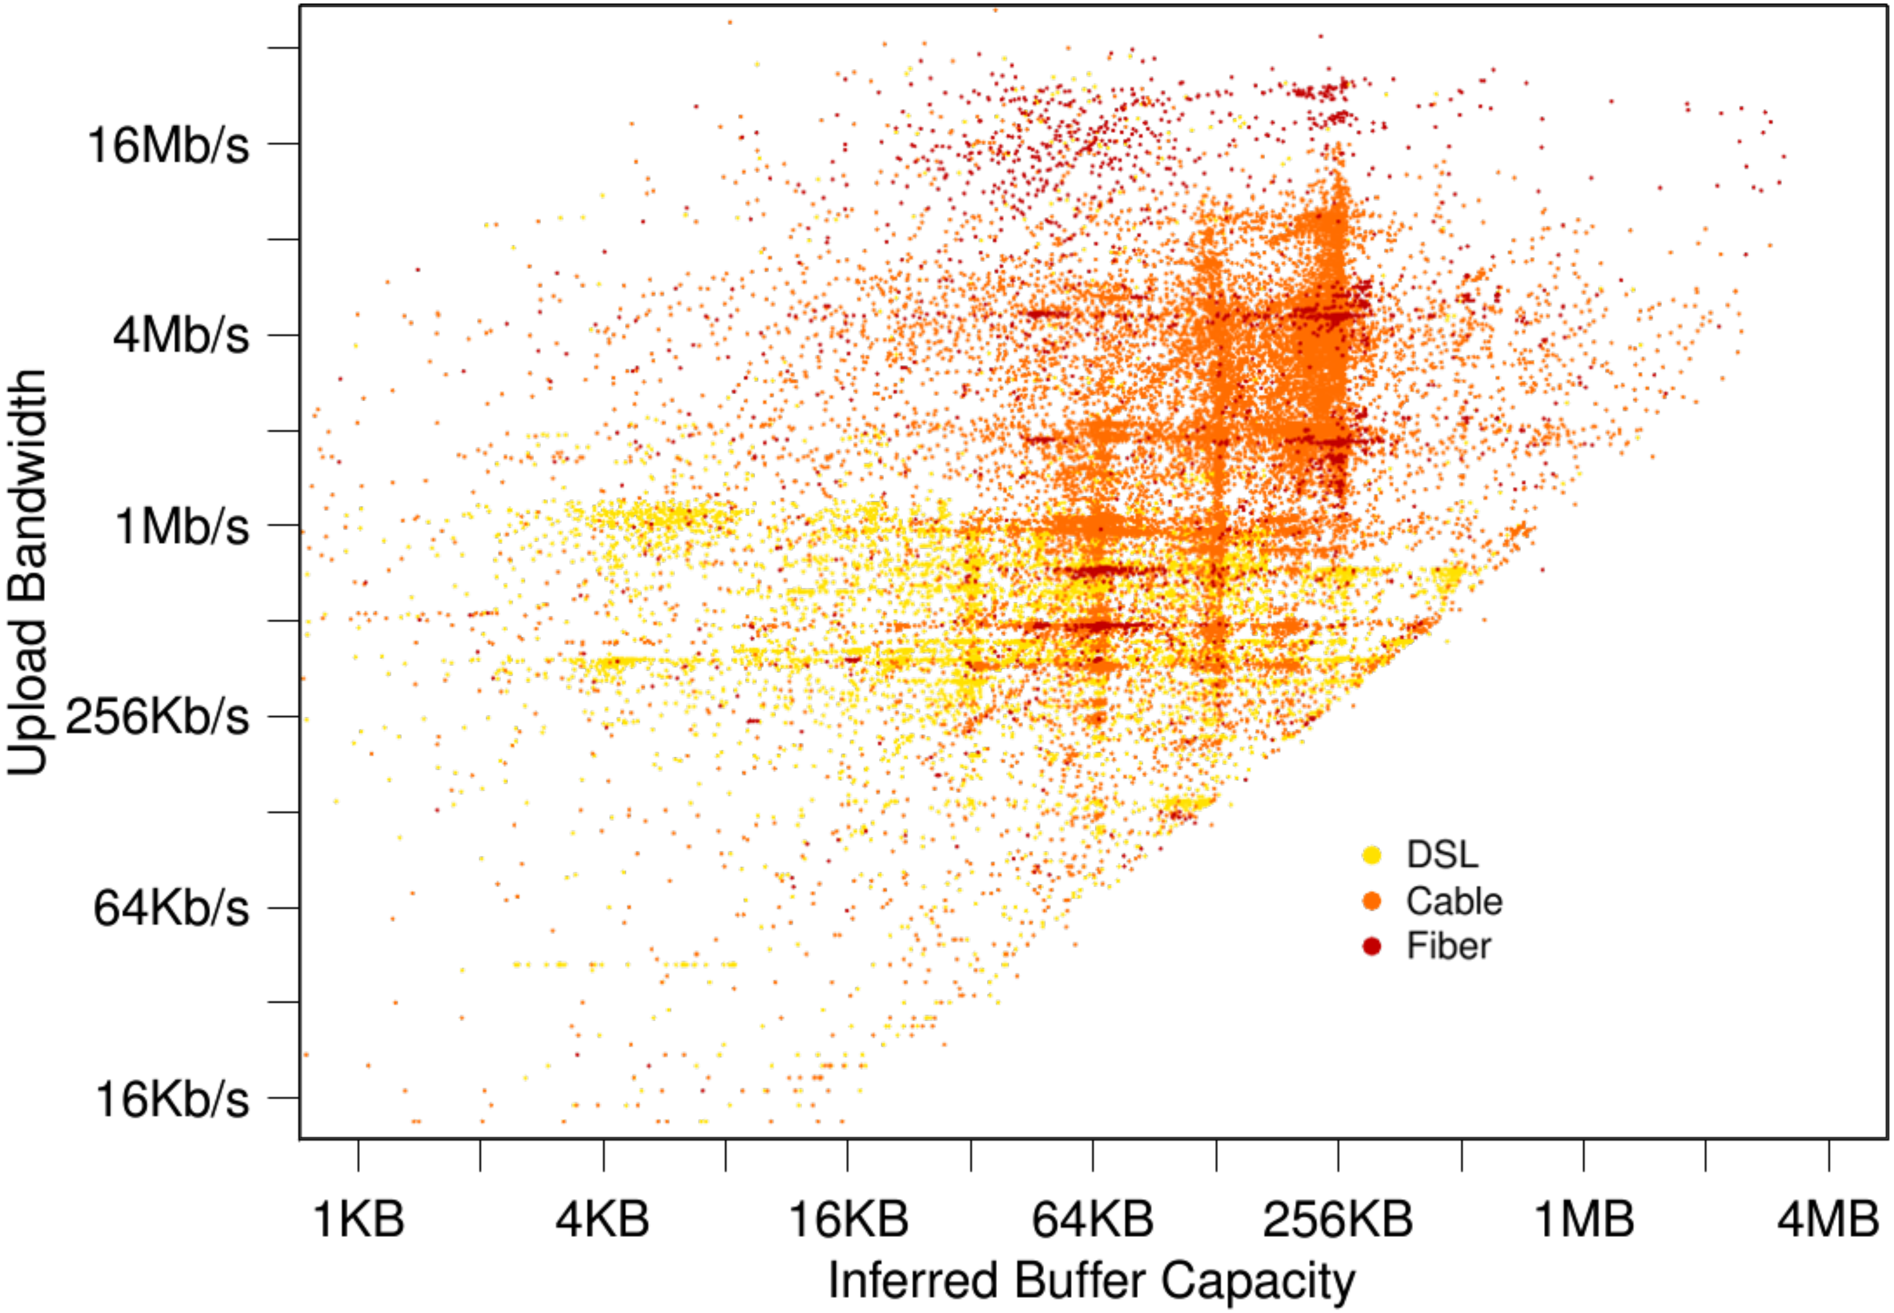
Upload Bandwidth
16Mb/s
4Mb/s
1Mb/s
256Kb/s
64Kb/s
16Kb/s
1KB
4KB
16KB
64KB
256KB
1MB
4MB
Inferred Buffer Capacity
DSL
Cable
Fiber

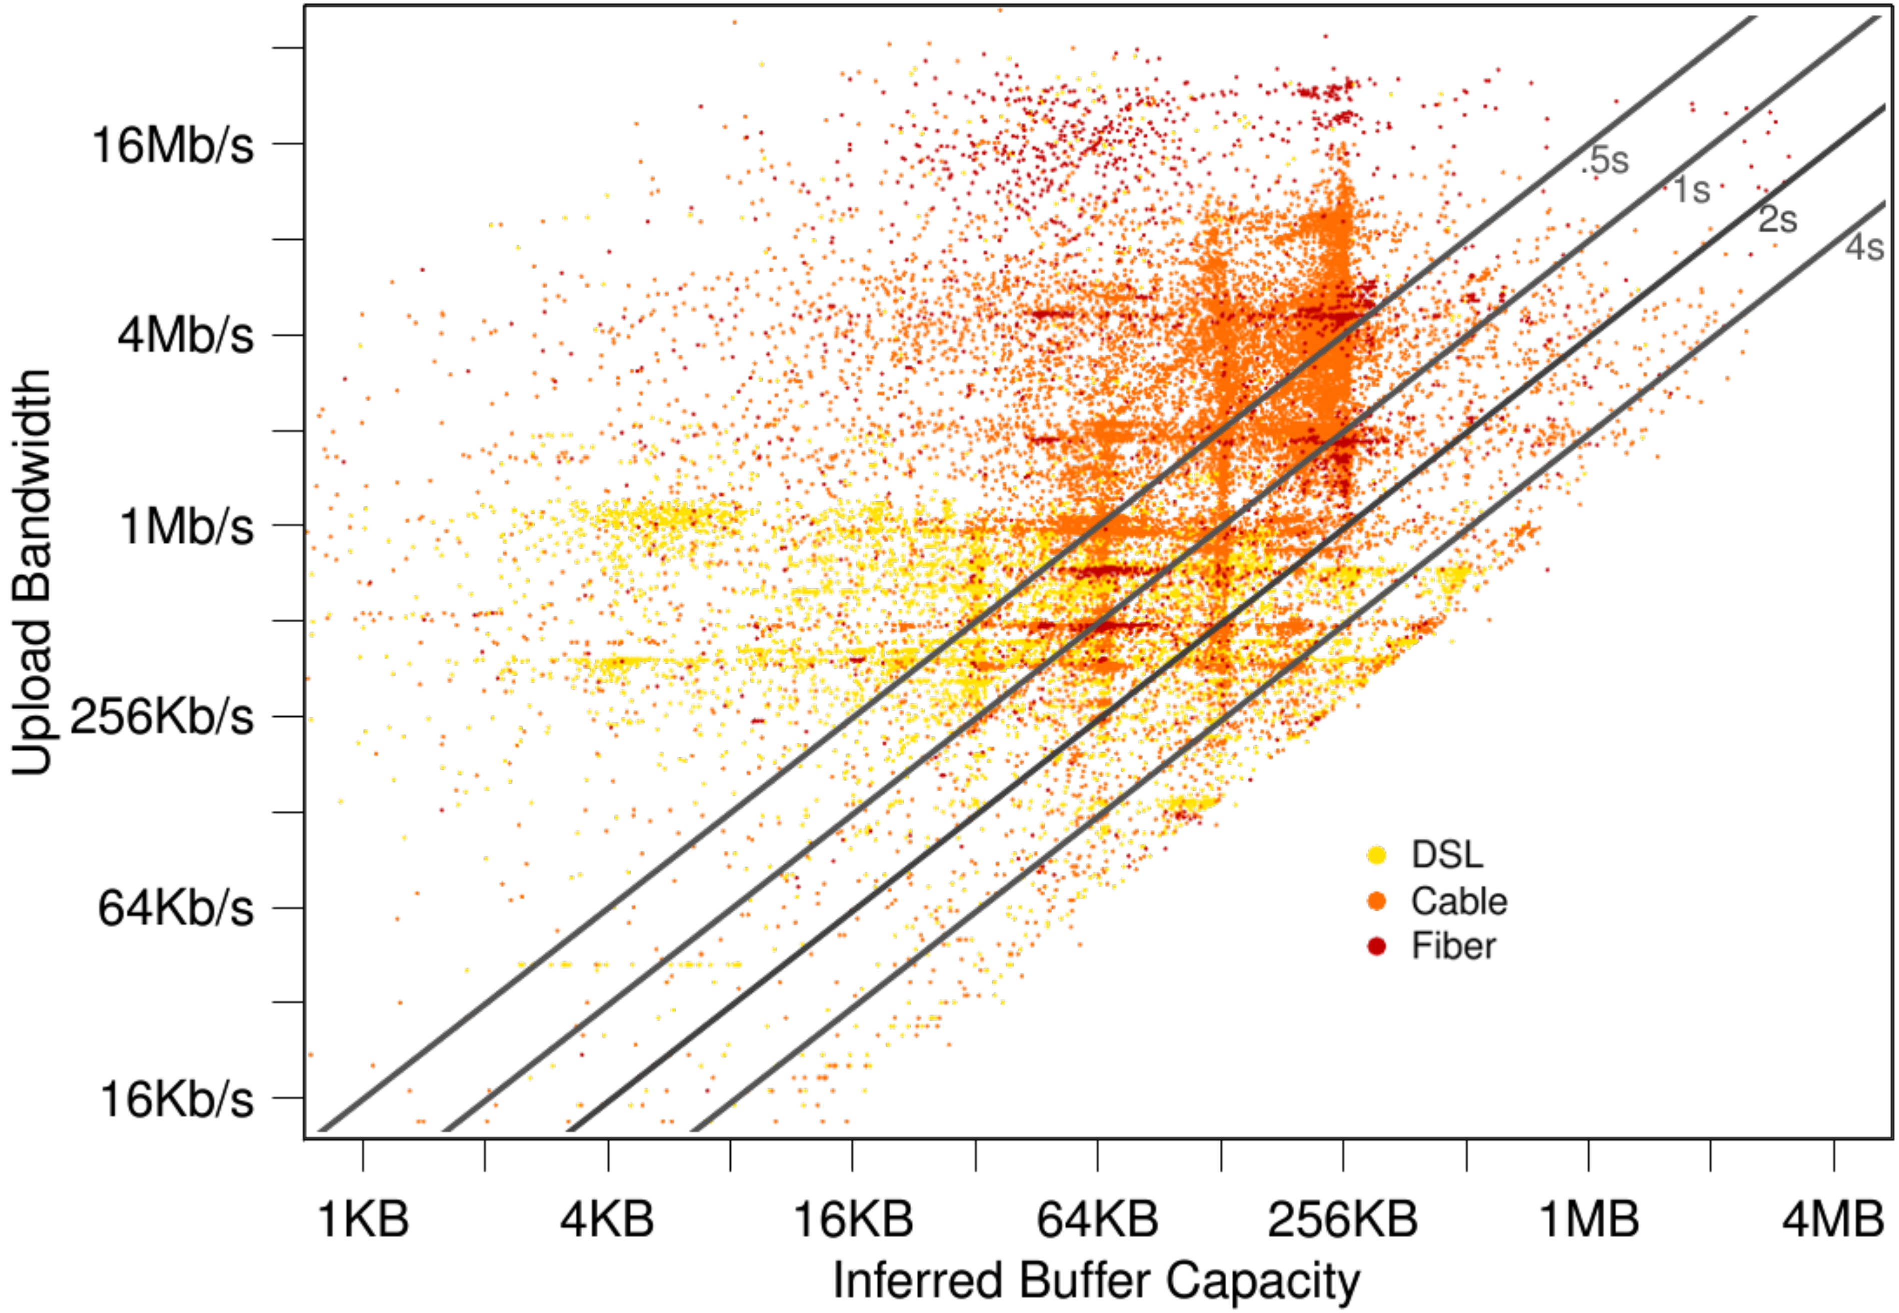
Upload Bandwidth
16Mb/s
4Mb/s
1Mb/s
256Kb/s
64Kb/s
16Kb/s
1KB
4KB
16KB
64KB
256KB
1MB
4MB
Inferred Buffer Capacity
.5s
1s
2s
4s
● DSL
● Cable
● Fiber

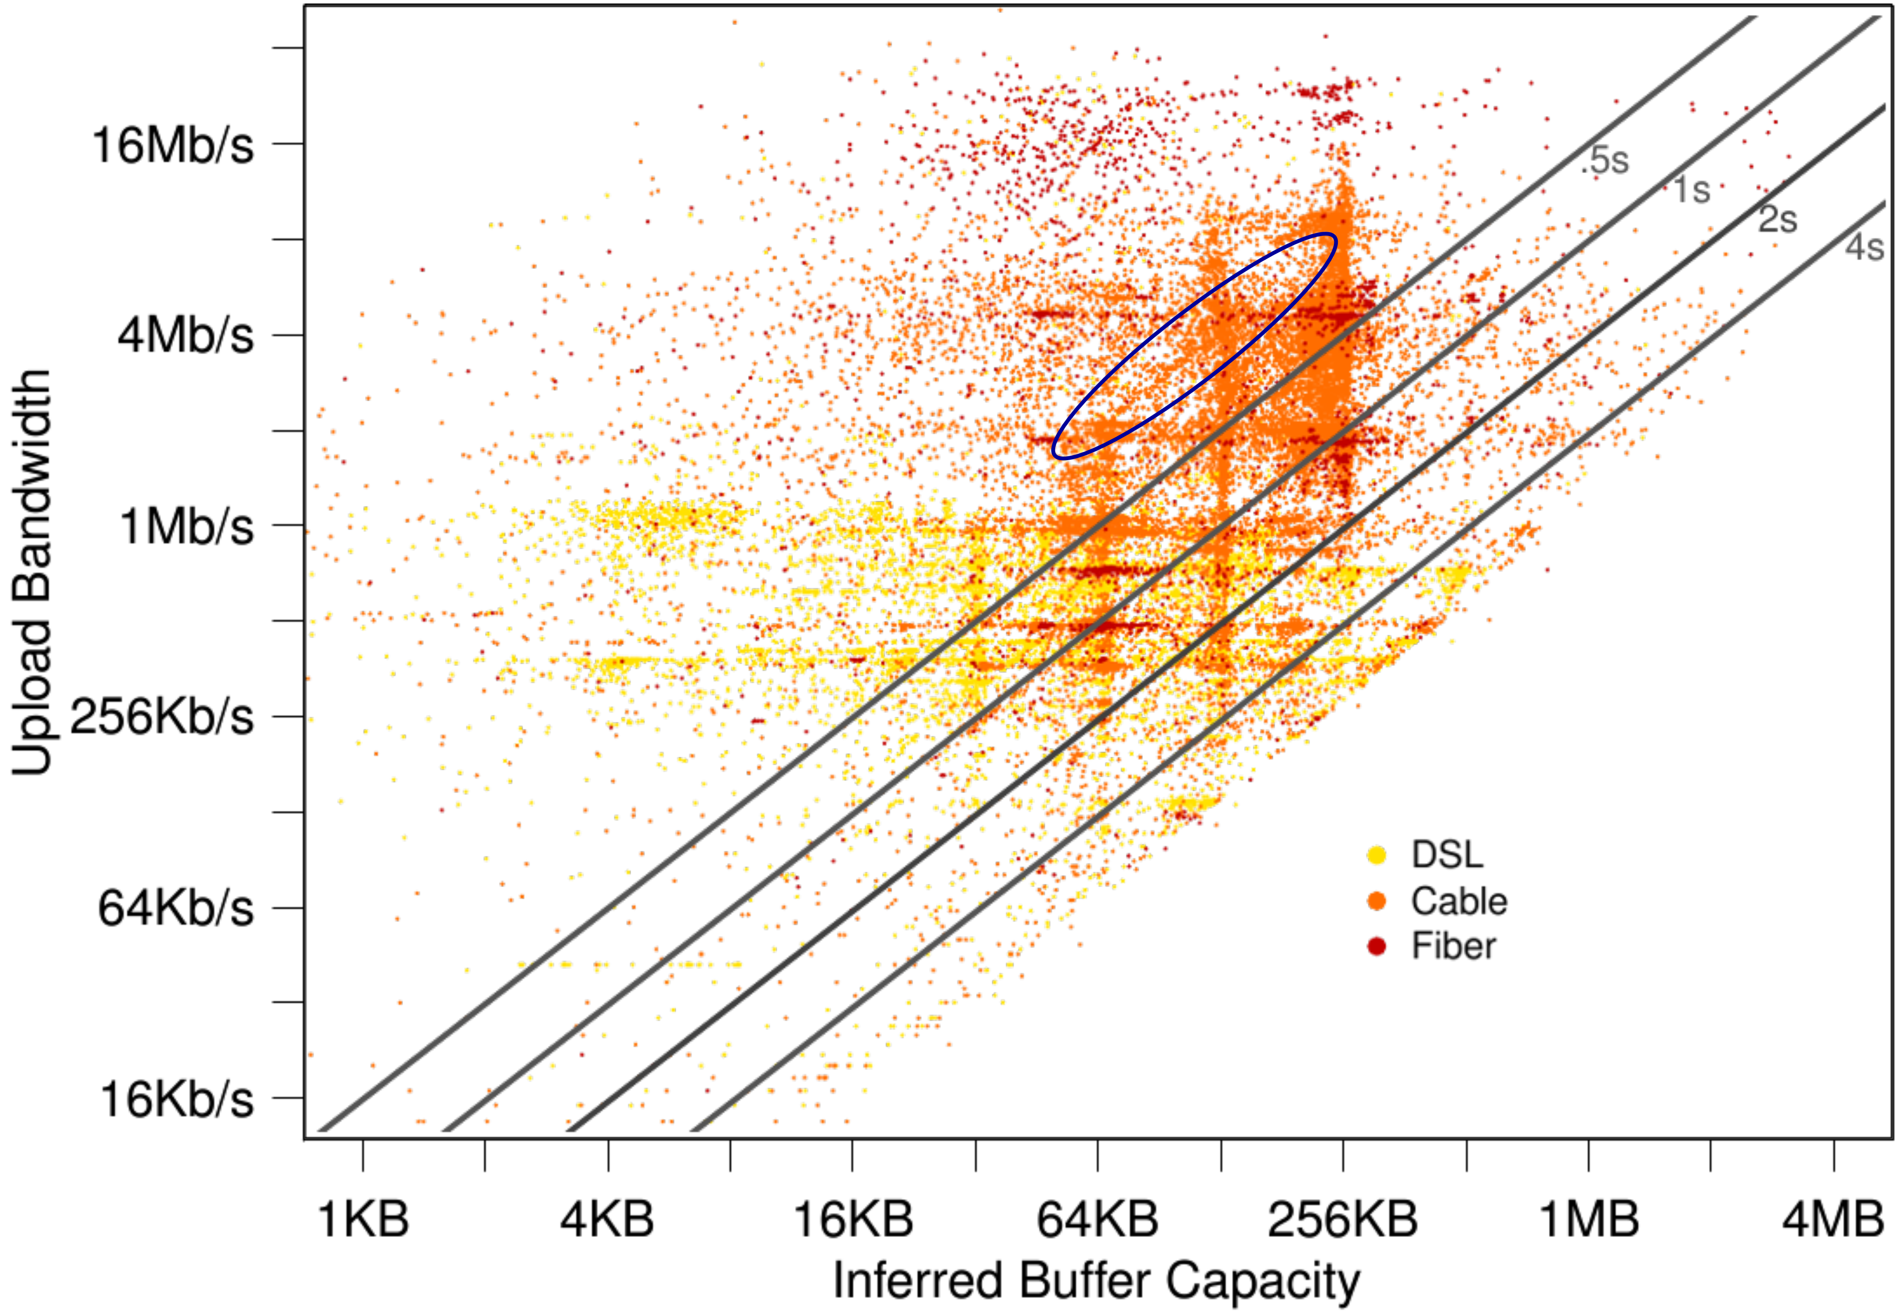
Upload Bandwidth
16Mb/s
4Mb/s
1Mb/s
256Kb/s
64Kb/s
16Kb/s
Inferred Buffer Capacity
1KB
4KB
16KB
64KB
256KB
1MB
4MB
.5s
1s
2s
4s
DSL
Cable
Fiber

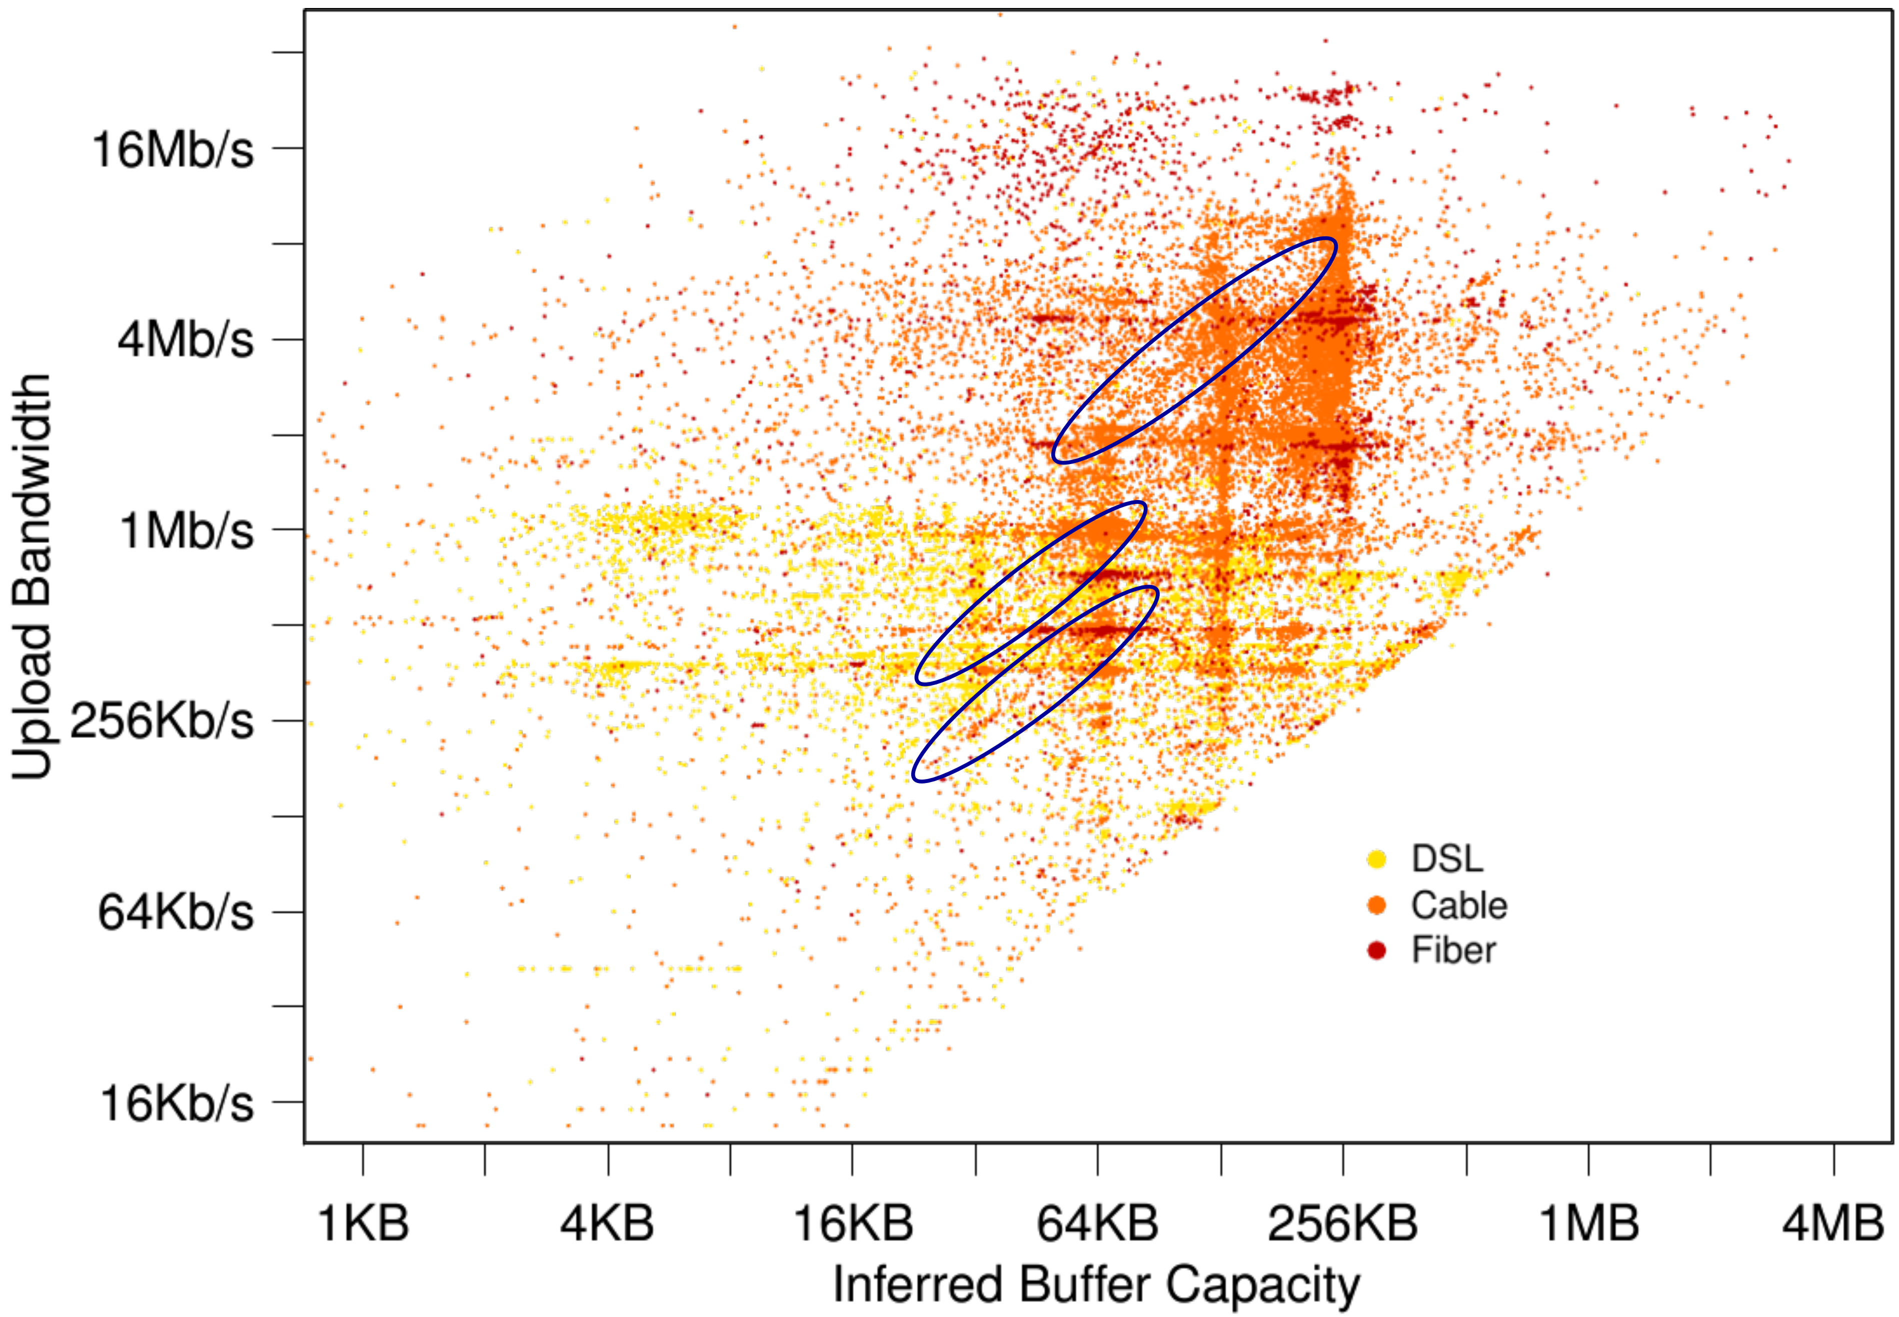
16Mb/s
4Mb/s
1Mb/s
256Kb/s
64Kb/s
16Kb/s
Upload Bandwidth
1KB
4KB
16KB
64KB
256KB
1MB
4MB
Inferred Buffer Capacity
● DSL
● Cable
● Fiber

Part IV:

Measuring Measurement

The Role of Measurement in Network Research

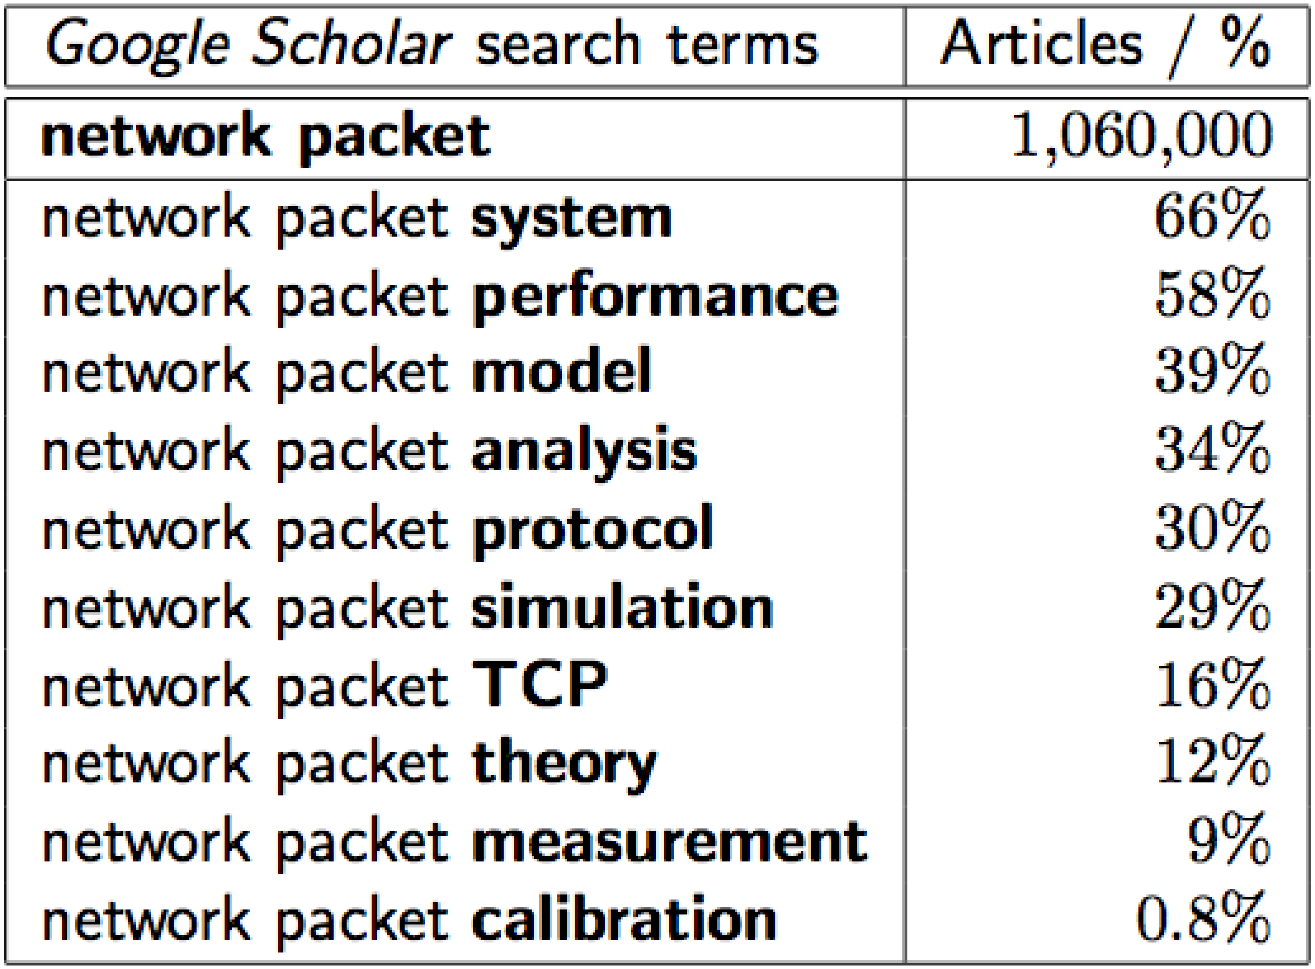

Google Scholar search terms	Articles / %
network packet	1,060,000
network packet **system**	66%
network packet **performance**	58%
network packet **model**	39%
network packet **analysis**	34%
network packet **protocol**	30%
network packet **simulation**	29%
network packet **TCP**	16%
network packet **theory**	12%
network packet **measurement**	9%
network packet **calibration**	0.8%

Thinking About *Network Science*

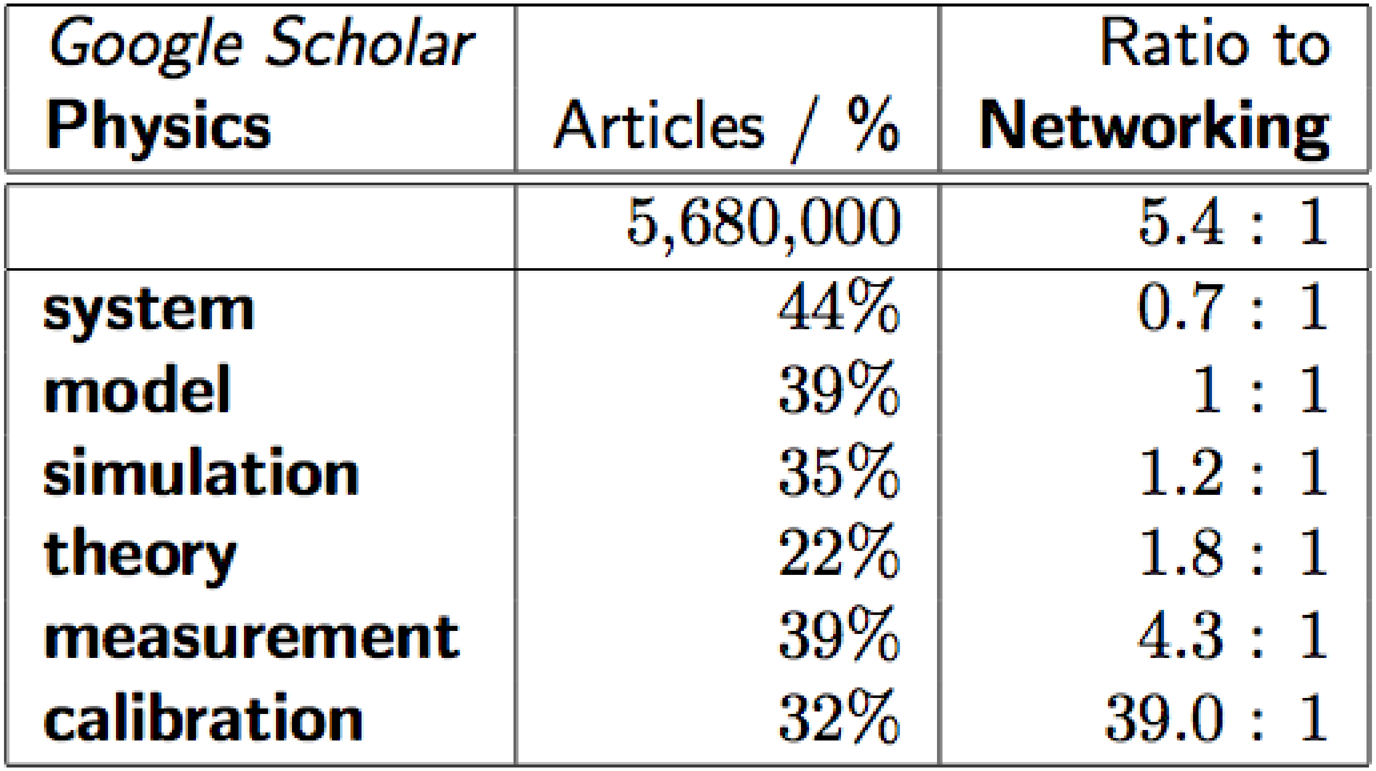

Google Scholar		Ratio to
Physics	Articles / %	**Networking**
	5,680,000	5.4 : 1
system	44%	0.7 : 1
model	39%	1 : 1
simulation	35%	1.2 : 1
theory	22%	1.8 : 1
measurement	39%	4.3 : 1
calibration	32%	39.0 : 1

What I’ve Tried To Convey

- What makes measurement practitioners tick?
  - A passion for knowing how things really work
    - Not: how we guess that they work
    - Not: how we’d like them to work
  - Indeed, term “measurement” is unfortunate -- detail-oriented
    - More accurate: “*empirical analysis*”
- Students: follow what really jazzes you
  - That’s what brings out excellence
  - Take heart when the path seems obscure

What I've Tried To Convey, con't

- Measurement folks: analyzing Modern Badness is a blast
  - All sorts of compelling problems & surprising possibilities
  - BUT: technically a big headache
    - Can require developing a lot of disparate elements
    - Plus significant issues regarding ethics & legality
- Research folks: measurement can fundamentally change our understanding & perception
  - Seems we should figure out how to do more of it
  - But: this is not at all easy for a whole bunch of reasons