AS Topology Visibility - You Can't Get There from Here

2010.05.05 / RIPE Praha

Randy Bush <randy@psg.com>
Olaf Maennel <olaf@maennel.net>
Matt Roughan <matthew.roughan@adelaide.edu.au>
Steve Uhlig <steve.uhlig@gmail.com>

<http://archive.psg.com/100505.ripe-visibility.pdf>

We Study *Visibility*

- What is the real routing graph of the Internet?
- What is the AS topology of BGP routing?
- How do we debug our network?
  - Are ping and traceroute the best we can do?
- How biased is our methodology?

RIPE-RIS & Route Views

- RIPE RIS/RouteViews were designed for operators
- Researchers discovered them – most **without consideration of limitations**

Google Scholar search for papers mentioning the term "Route Views"

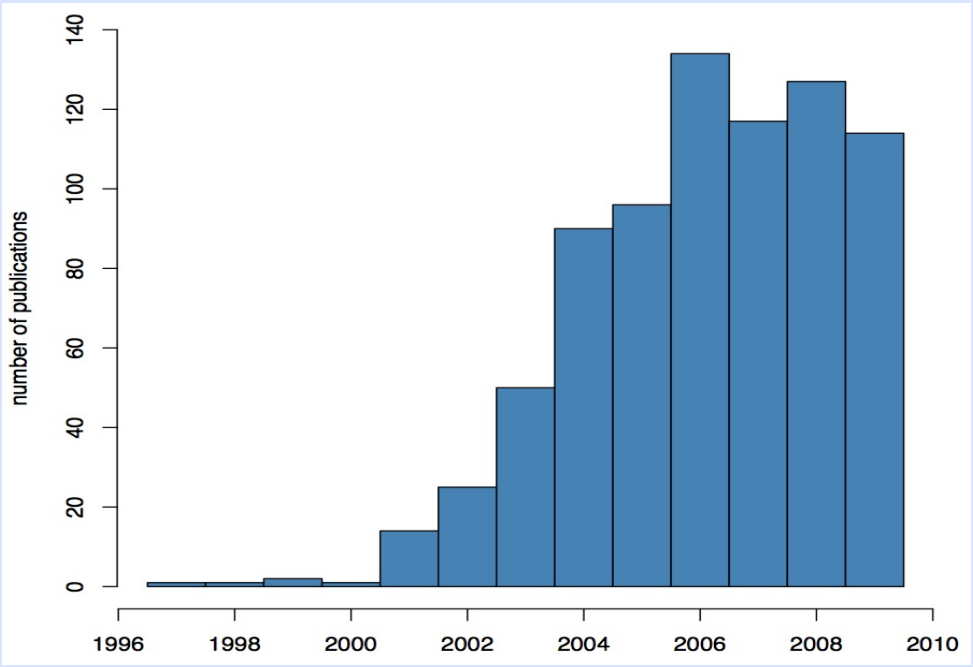

Bogon Diagnosis Work

- R&D for ARIN to enable them to diagnose what ASs were filtering newly allocated address space. See 2007 SIGCOMM NetMgt Workshop.
- Though ARIN never deployed, we continued to measure to see how long it takes to get filters removed.
- Bored, we turned the tool to other use

Announcing a /25

- We announced a /25 to NTT Global
- They passed it only to customers
- RV/RIS/... showed 15 ASs could see it

Whoops!

- We used ping from the /25 to 'all' ASs
- 1024 ASs could get packets back to the /25 source!
- So Route-Views and RIS were off by a FACTOR OF 60!
- And one was as good/bad as another, adding more views did not help.

/25 AS Hops

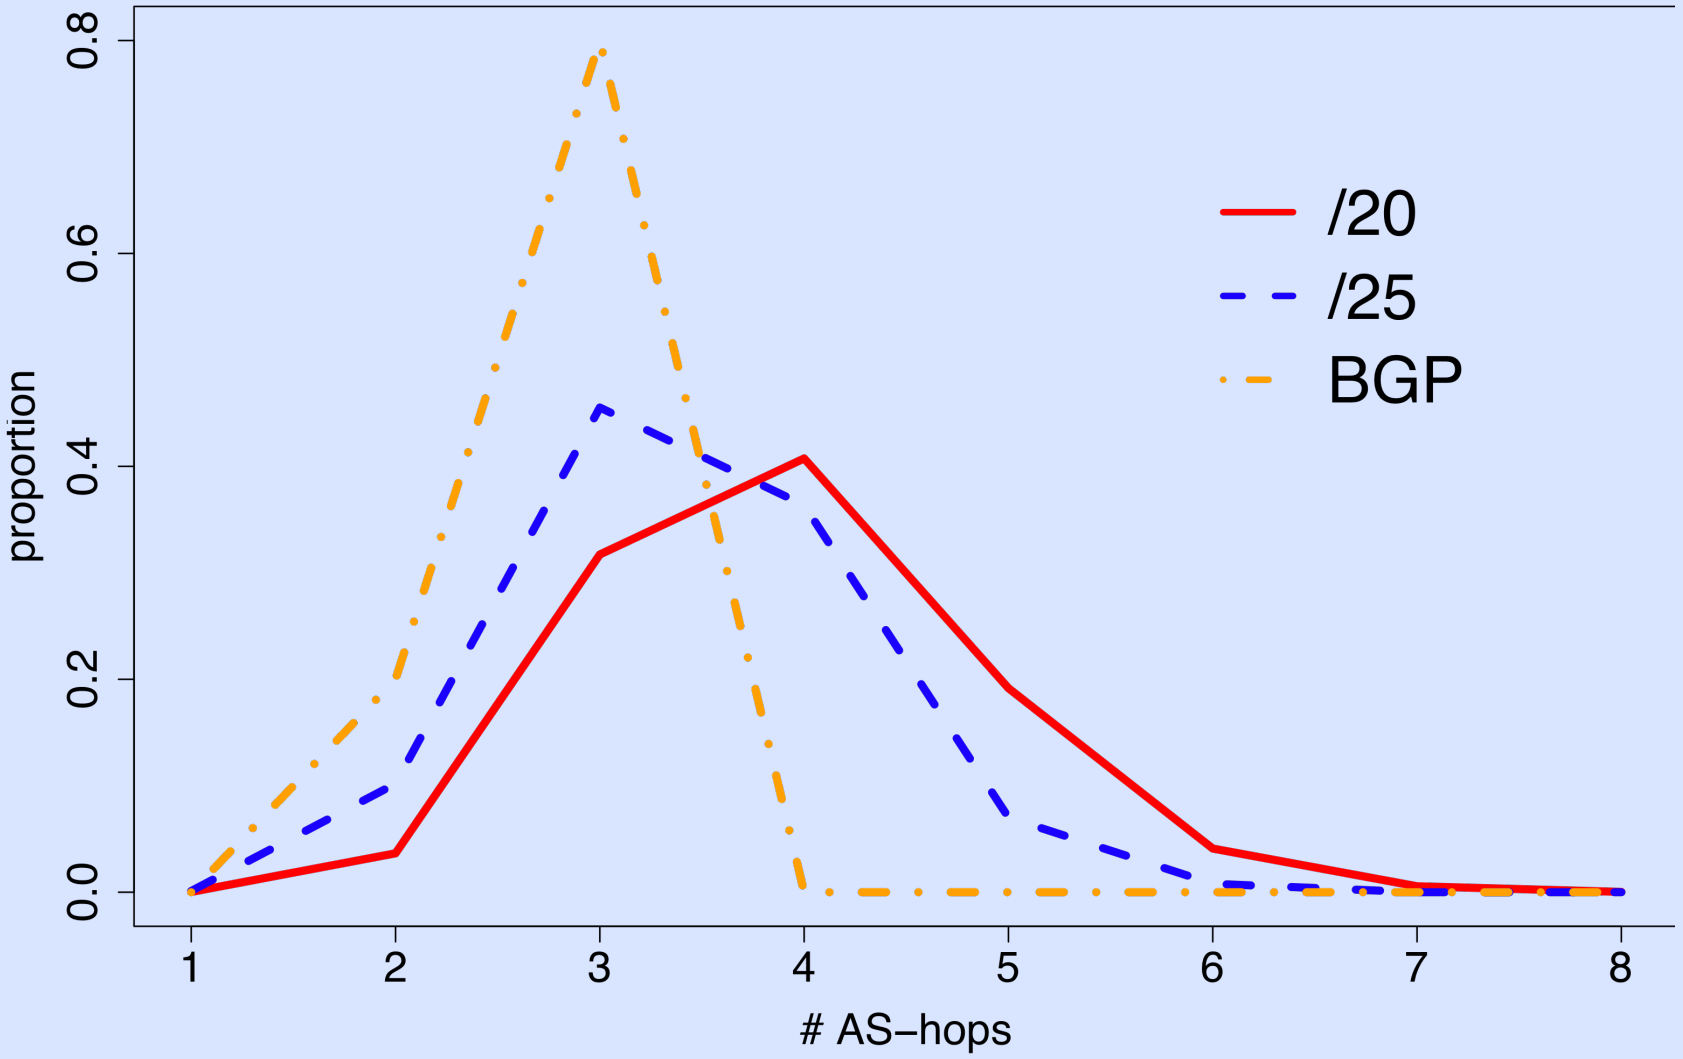

How Much of This was Due to Default as Opposed to Poor BGP Visibility?

Default Detection

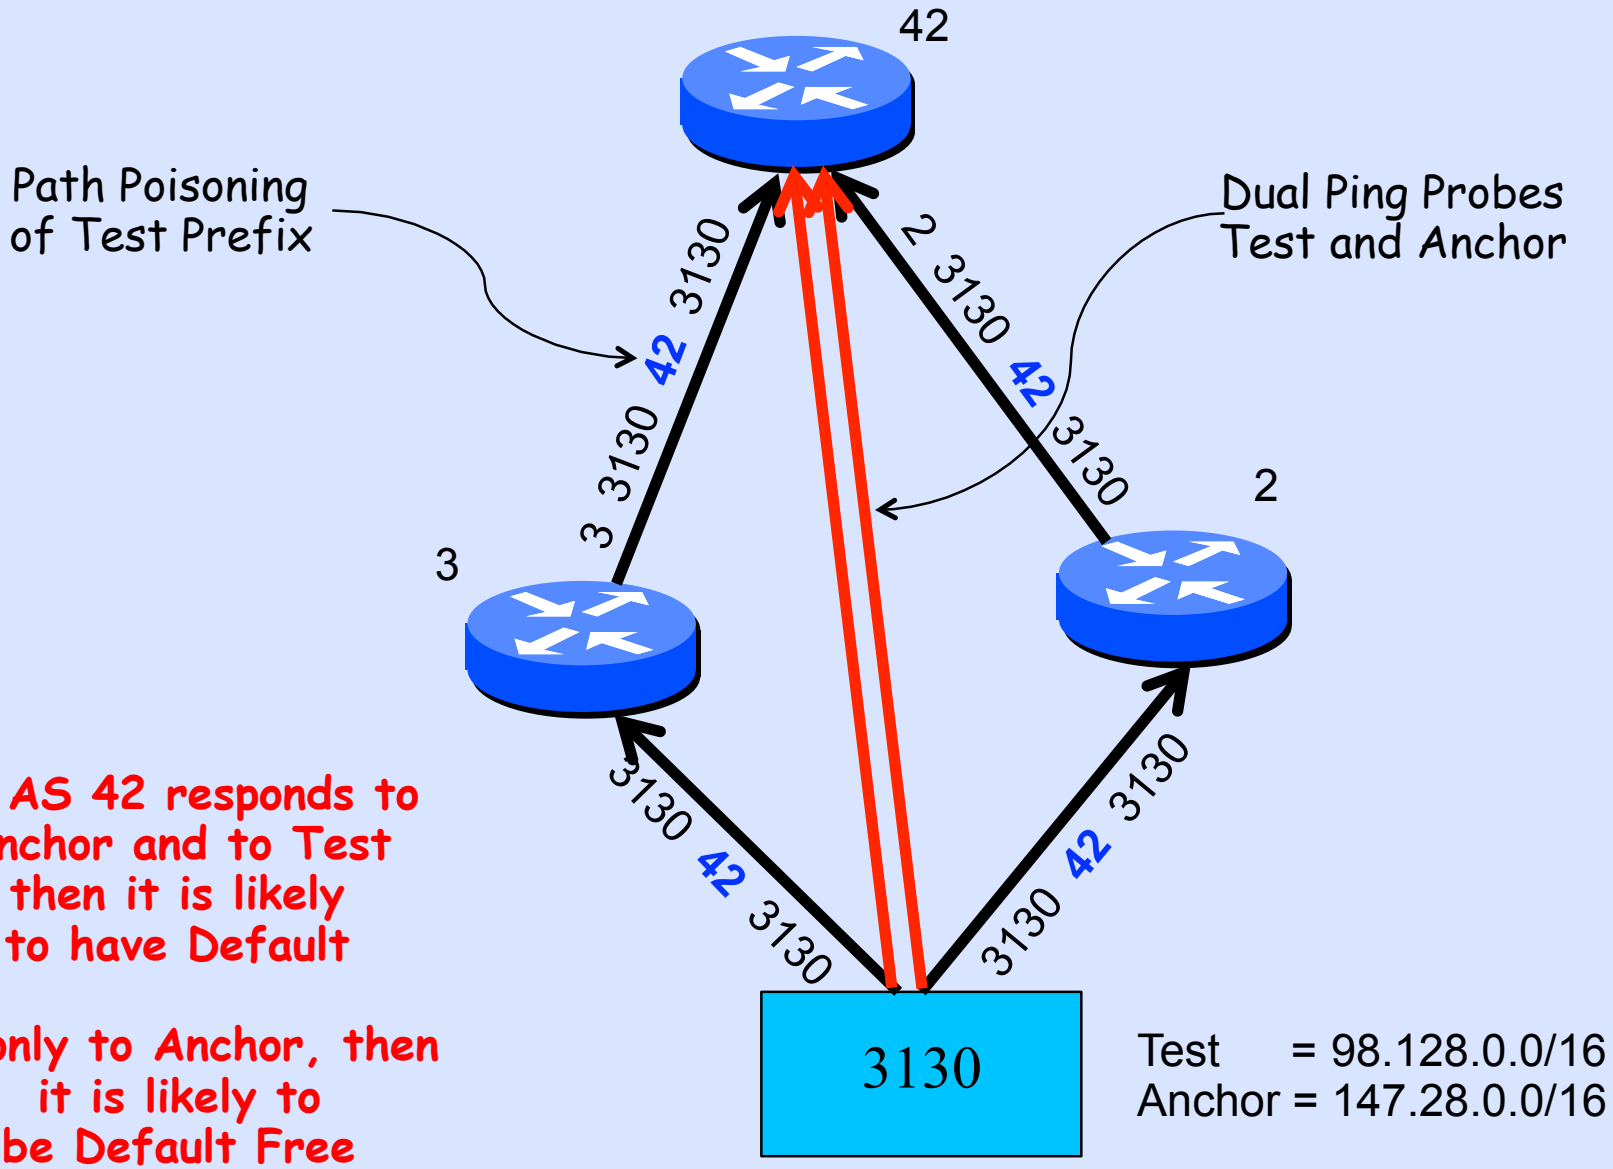

If AS 42 responds to Anchor and to Test then it is likely to have Default

If only to Anchor, then it is likely to be Default Free

Use of Default toward /25

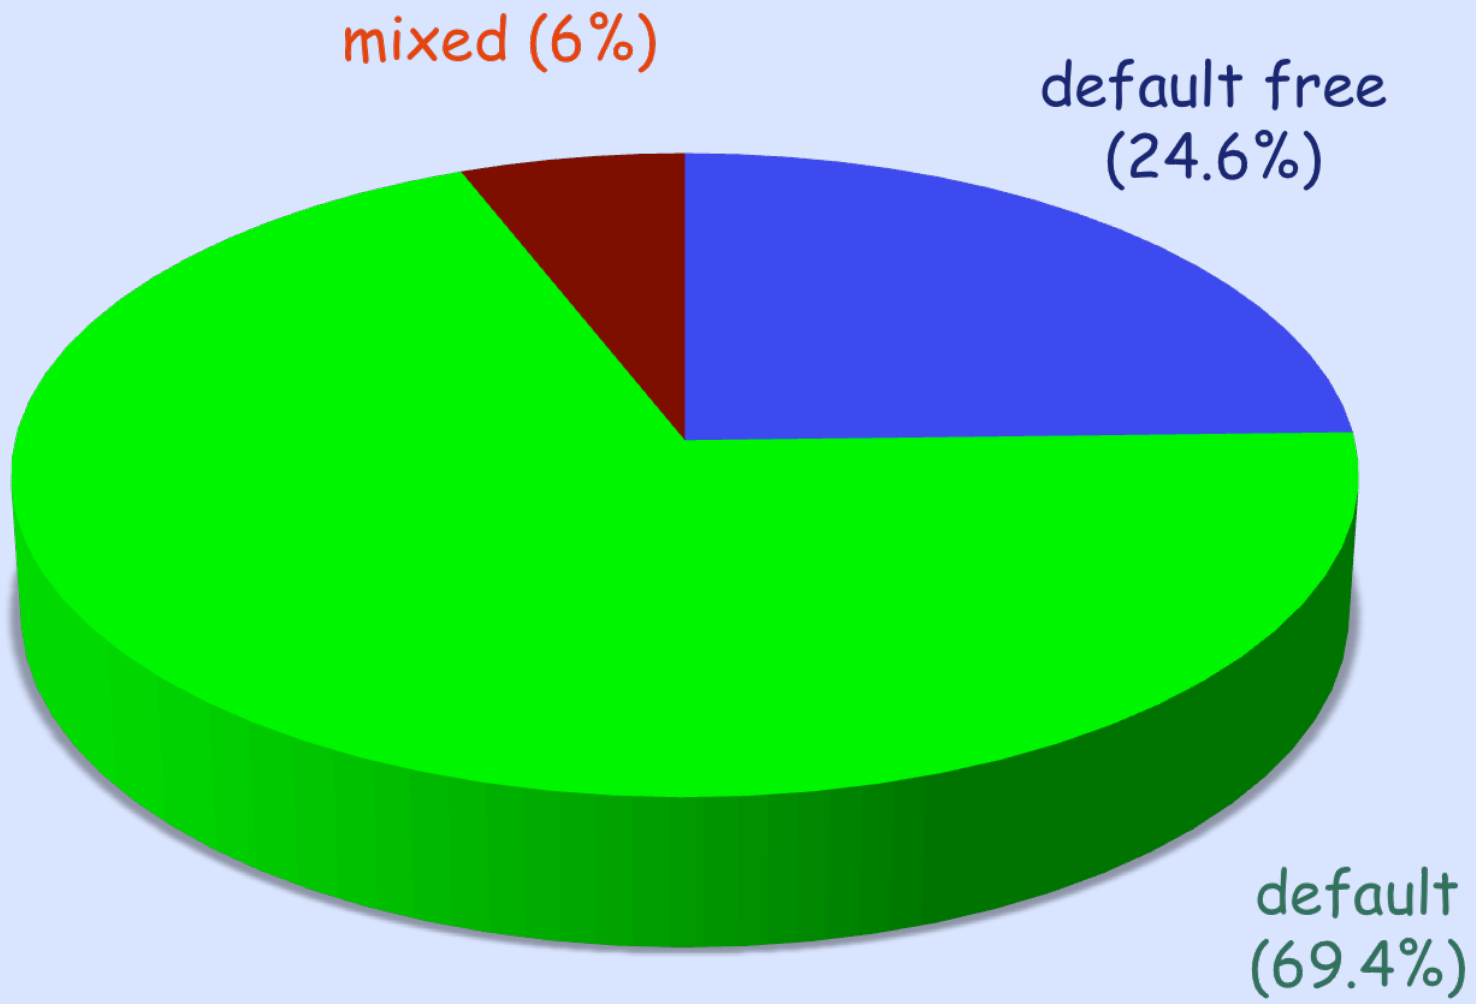

Defaults in /25-Experiment

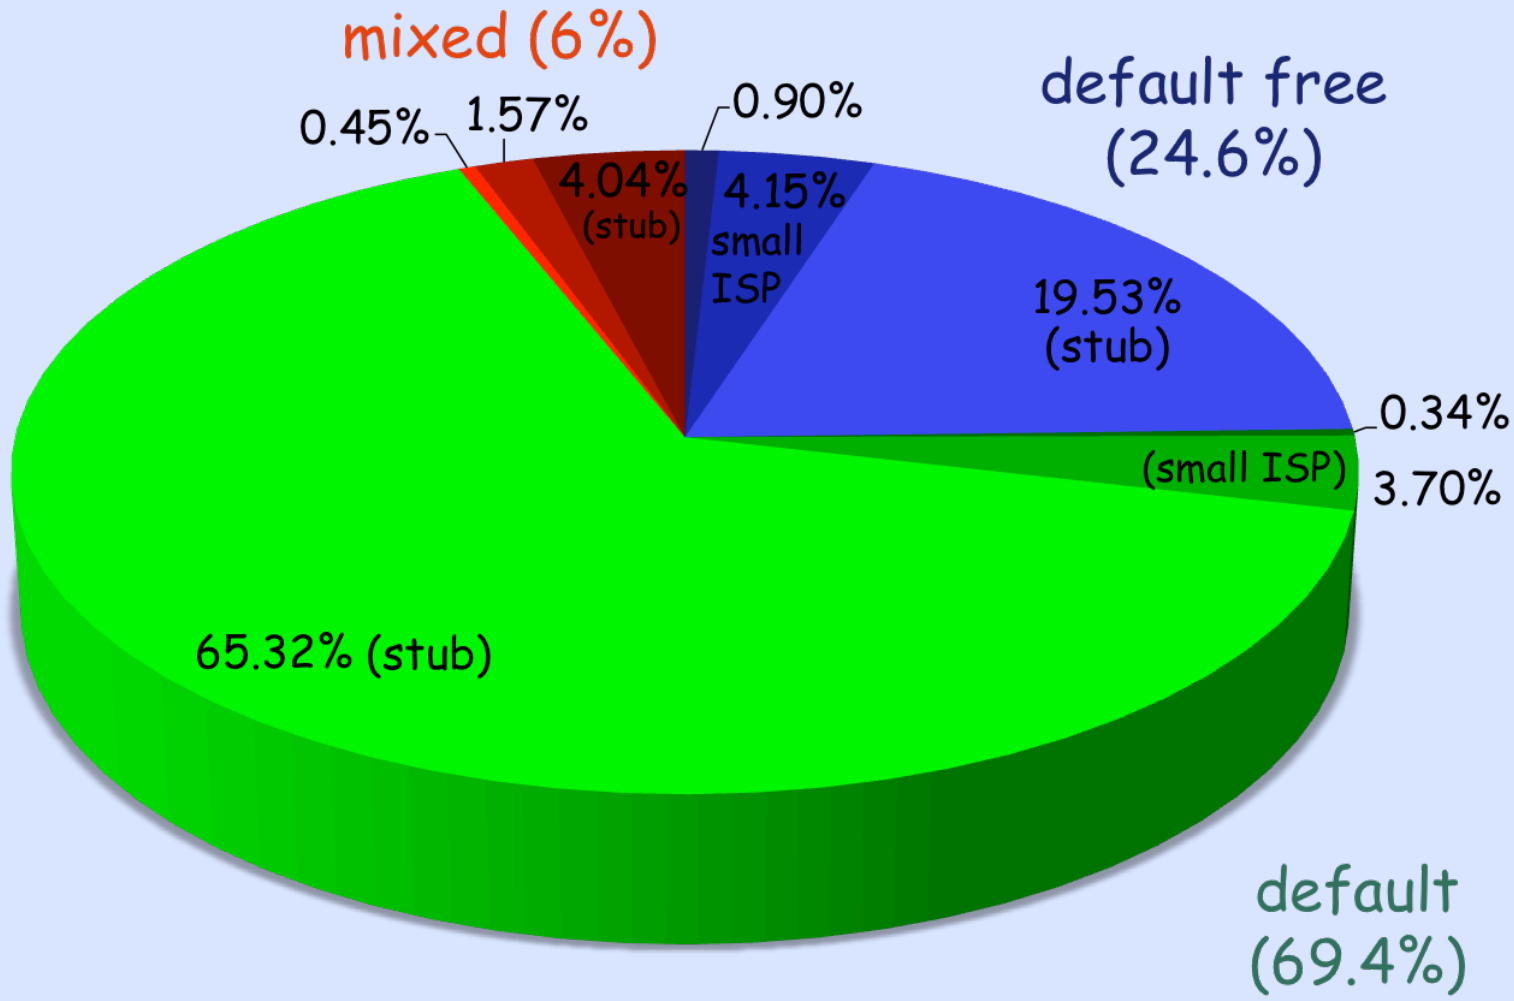

Default Free Zone?

Not Really!

Testing Most ASes

UCLA taxa	tested/total	default	default-free	mixed
stub	24,224/31,517	77.1%	19.3%	3.6%
small ISP	1,307/1,361	44.5%	42.2%	13.3%
large ISP	246/255	17.1%	60.6%	22.3%

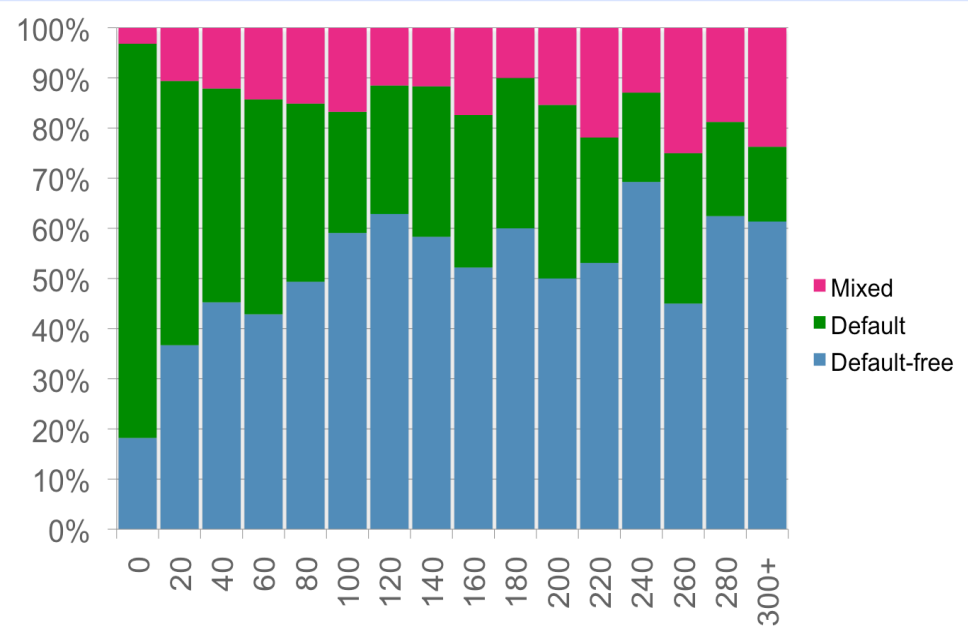

Default routing use as a function of AS out-degree

ASes with out-degree ≥ 300 are combined in the last value.

But Japan is Different

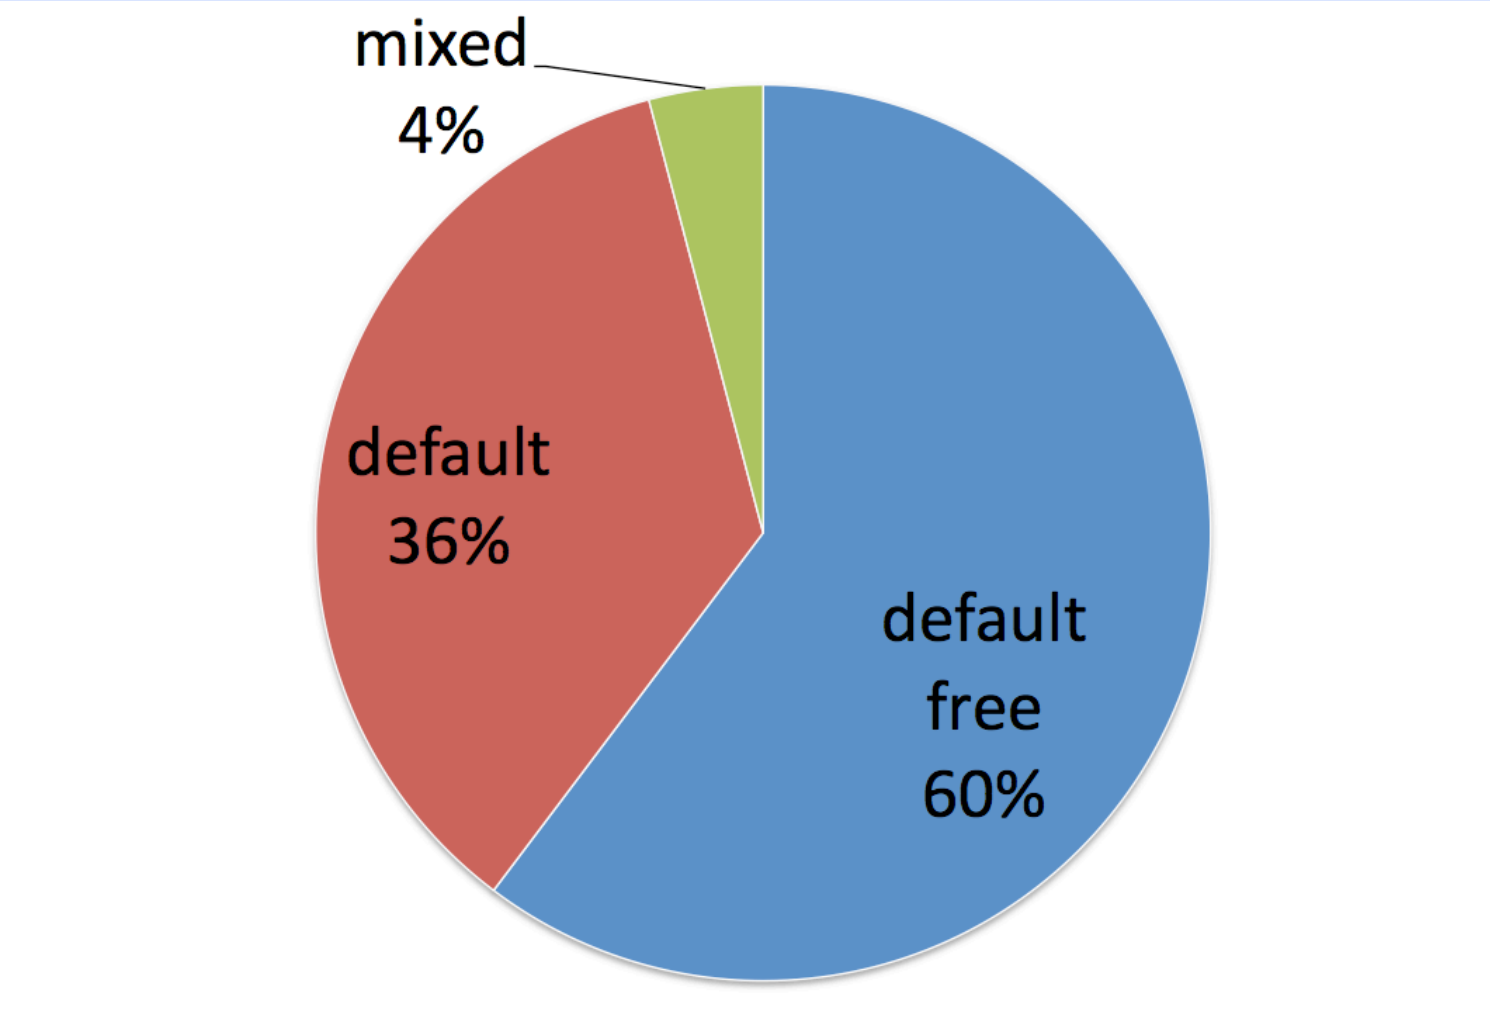

Tomoya Yoshida <yoshida@nttv6.jp>

Asia Varies Widely

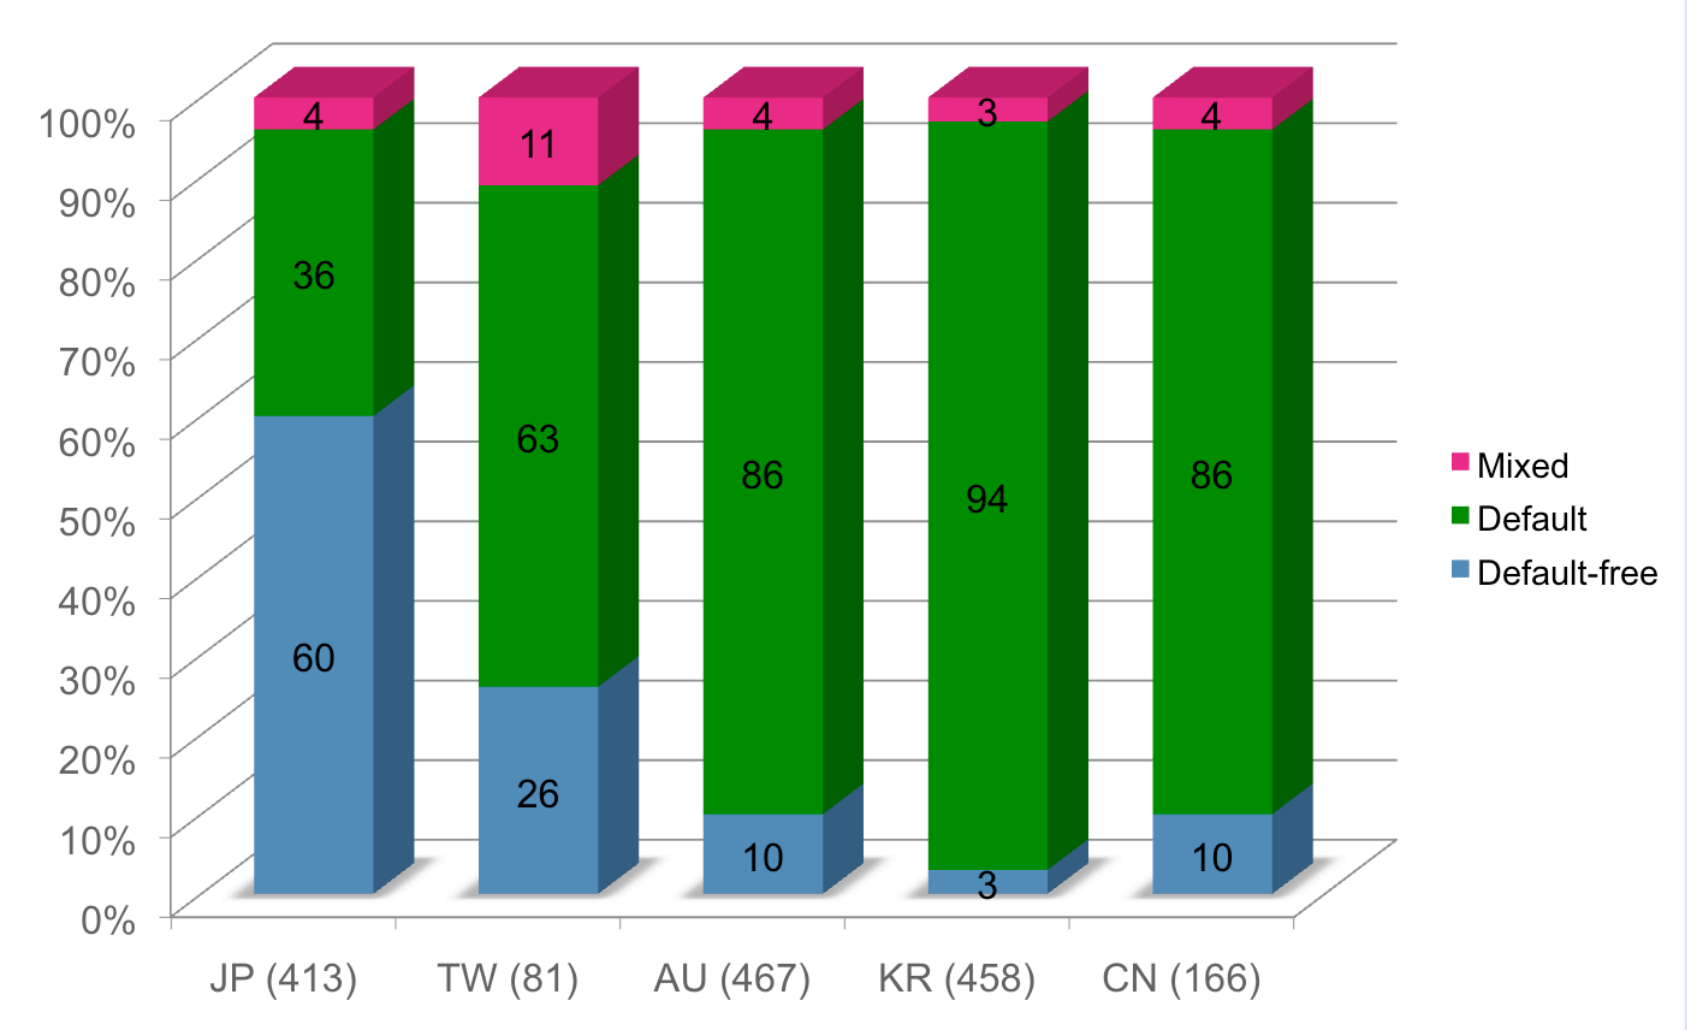

Tomoya Yoshida <yoshida@nttv6.jp>

Validation – We Asked

- 216 operators answered,
- 172 (79.6%) said "correct",
- 21 (9.7%) "almost" correct (e.g., correctly measured, but network is more complex),
- 10 (4.6%) believed we were right (did not recheck),
- 8 (3.7%) we measured wrongly (e.g., AS address space from different provider),
- 5 (2.3%) said we must be wrong ☺

R-Views / RIPE-RIS

- 1024 path-poisoned ASs could reach the test prefix
- Assume 70% used default
- The other 30%, or 307, had a path
- Only 15 of them showed up in RV/RIS
- **RV/RIS was off by a factor of 20**
- **And that is a lower bound!**

Policy Interactions – the "fun" of BGP research… ;-)

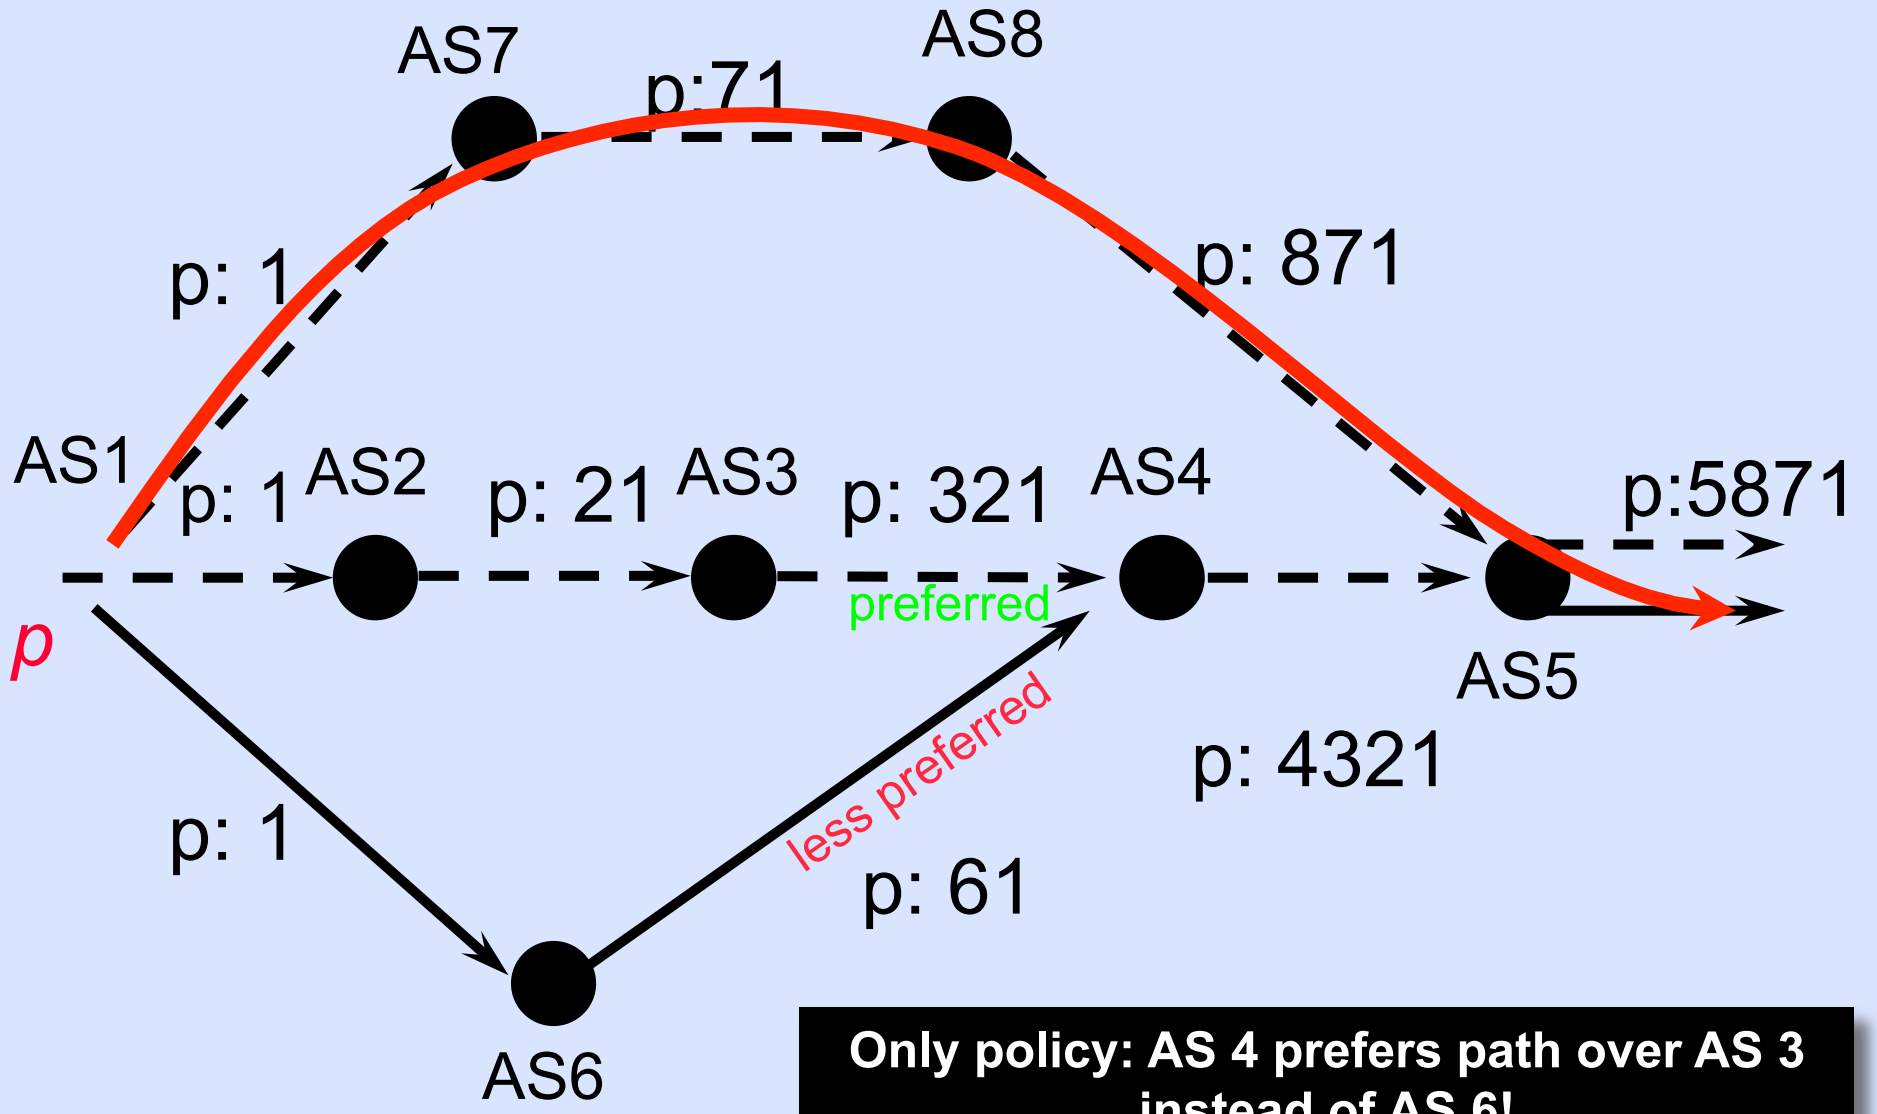

Only policy: AS 4 prefers path over AS 3 instead of AS 6!

Policy Interactions – the "fun" of BGP research… ;-)

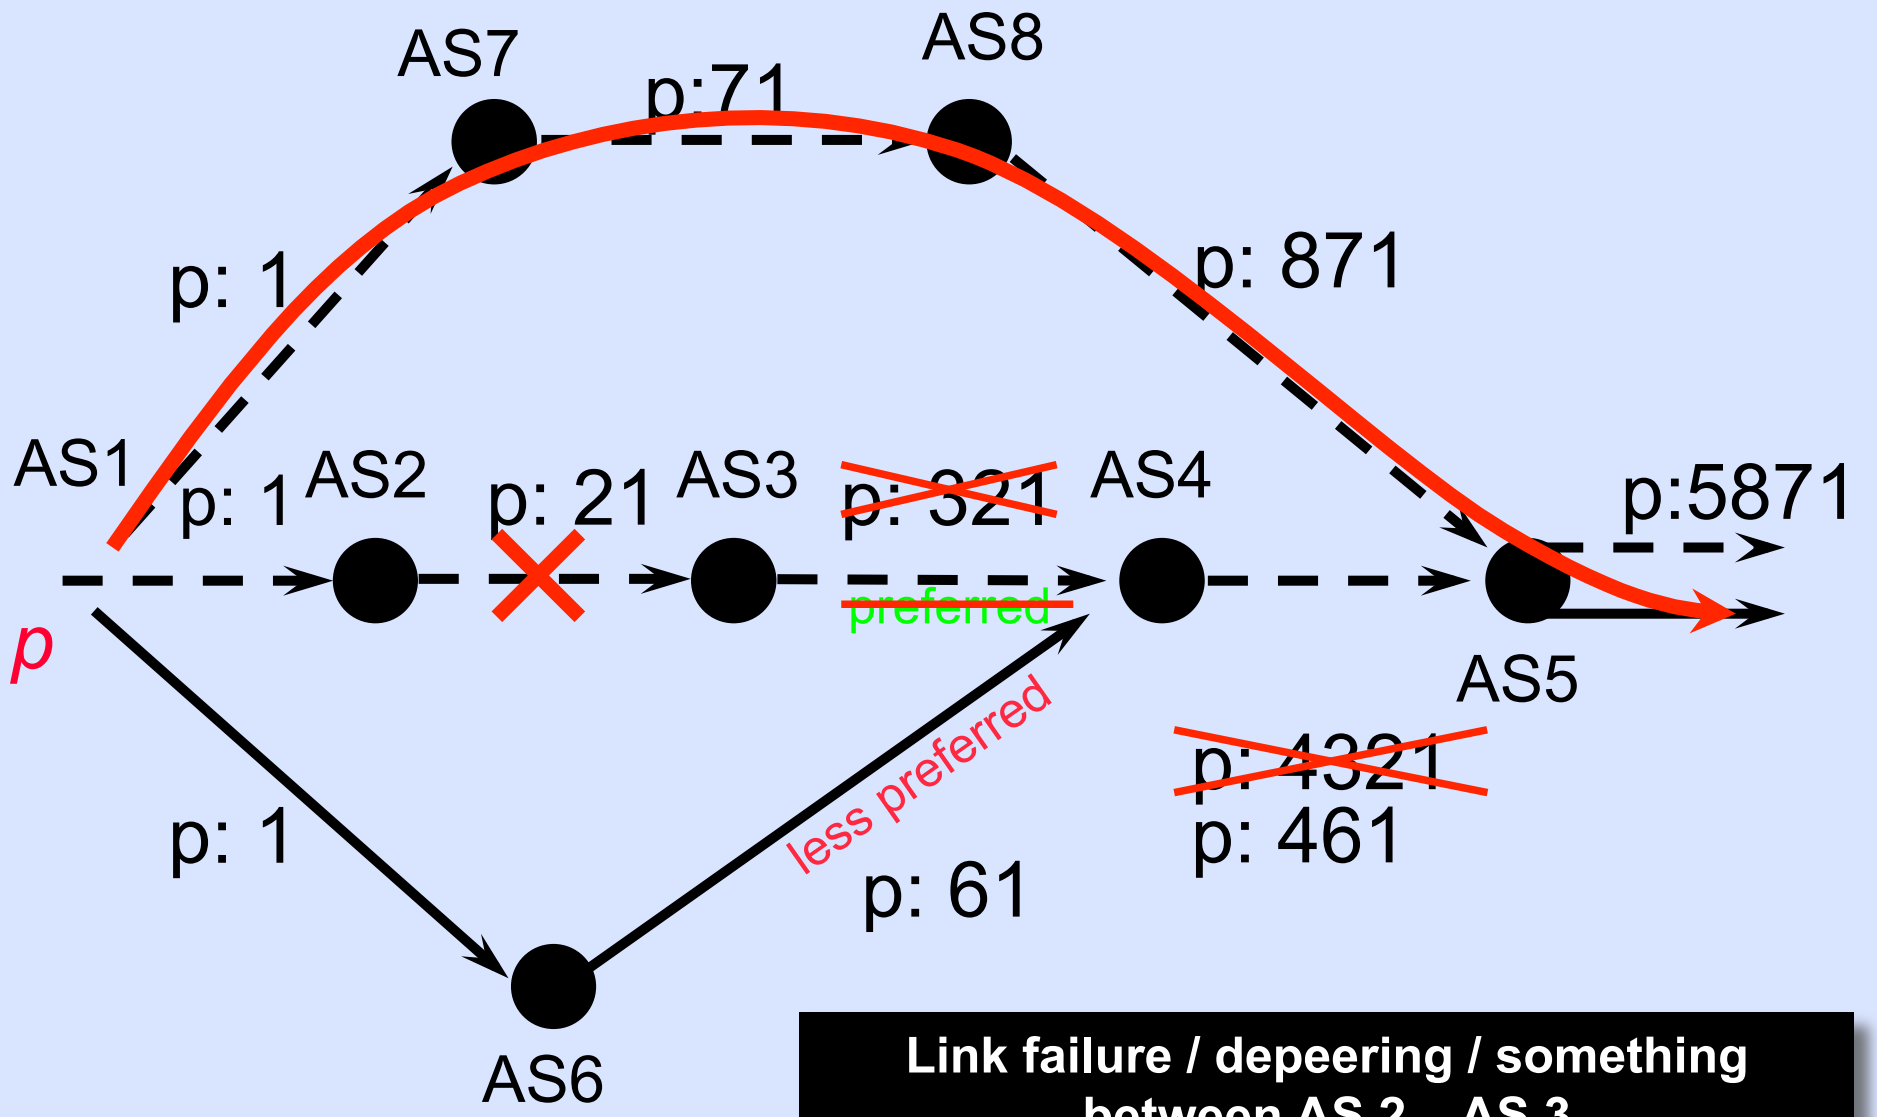

Link failure / depeering / something between AS 2 – AS 3

Policy Interactions – the "fun" of BGP research… ;-)

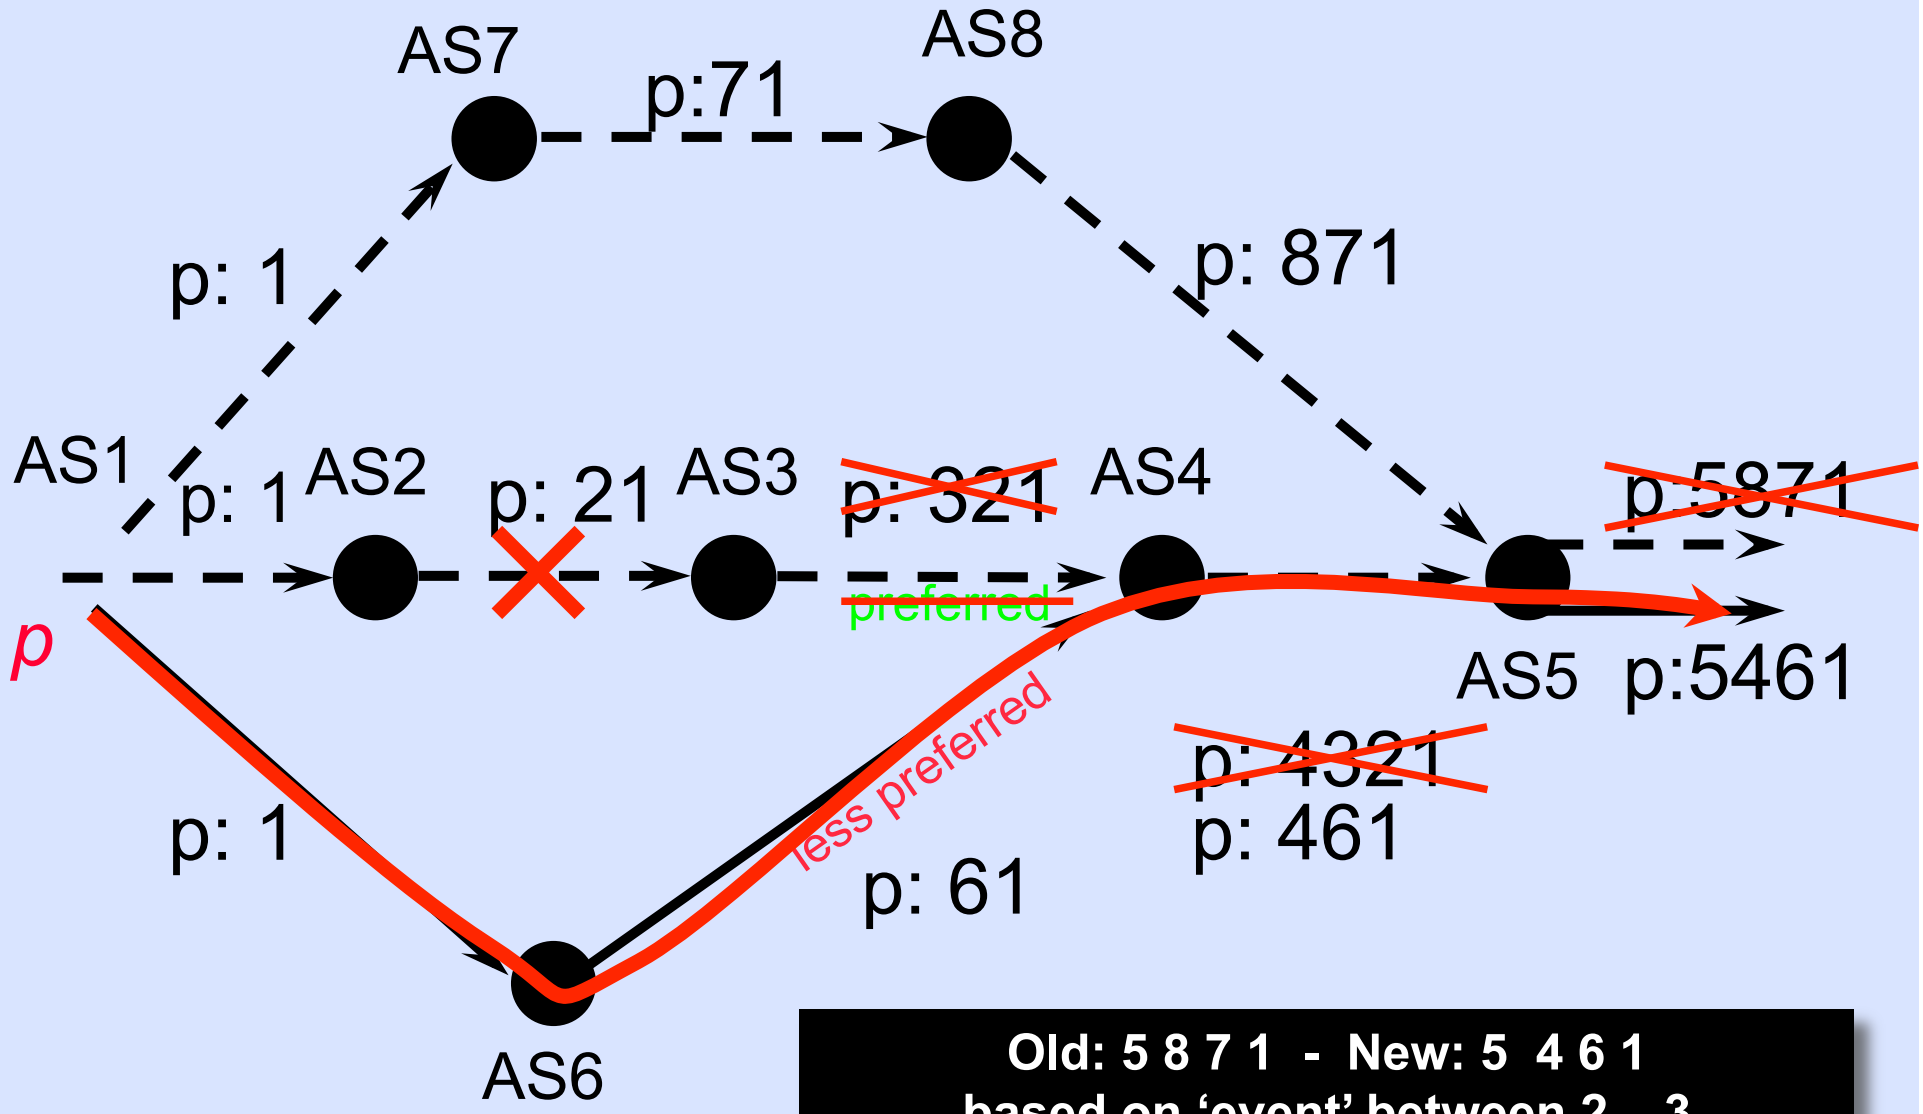

Old: 5 8 7 1 - New: 5 4 6 1
based on 'event' between 2 – 3

How do you know

if this is what happened?

-

Not good for BGP-based

formal Root Cause Analysis

Our Glasses are Broken

- Looking in RV/RIS/... does not tell you if they can reach you
- Looking just in RV or RIS is as good (well bad) as hundreds of BGP feeds
- Researchers should be very wary of using RV/RIS data for many classes of analysis, e.g. AS topology, traffic
- Are Renesys-style presos bogus?

Work Supported By

- **Cisco**

  SUPPORTED IN PART BY A CISCO UNIVERSITY RESEARCH PROJECT GIFT VIA KEIO UNIVERSITY

- **Internet Initiative Japan**

- **Google, NTT, Equinix**