Internet Initiative Japan

Hochschule für Angewandte Wissenschaften Hamburg

Hamburg University of Applied Sciences

PEERING
The BGP Testbed

Measuring Adoption of RPKI Route Validation and Filtering

Andreas Reuter (andreas.reuter@fu-berlin.de)

**Joint work with Randy Bush,
Ethan Katz-Bassett, Italo Cunha,
Thomas C. Schmidt, and Matthias Wählisch**

Once upon a time ... someone incorrectly announced an IP prefix.

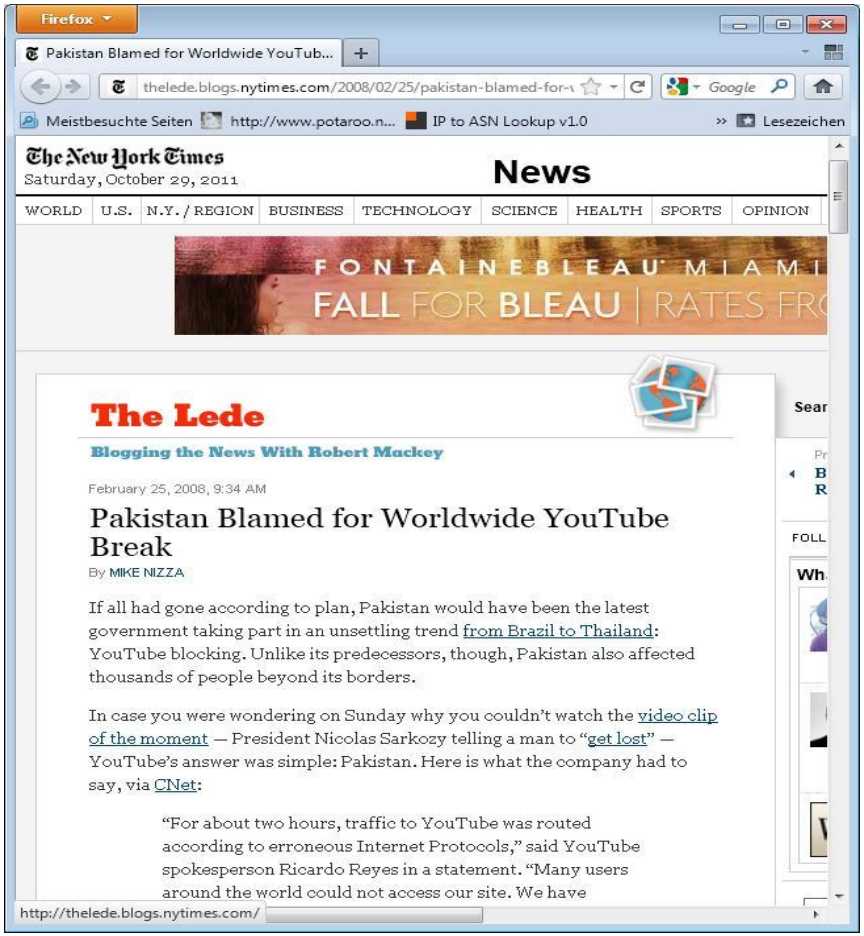

Once upon a time ... someone incorrectly announced an IP prefix.

For about 18 minutes on April 8, 2010, China Telecom advertised erroneous network traffic routes that instructed U.S. and other foreign Internet traffic to travel through Chinese servers.* Other serv-

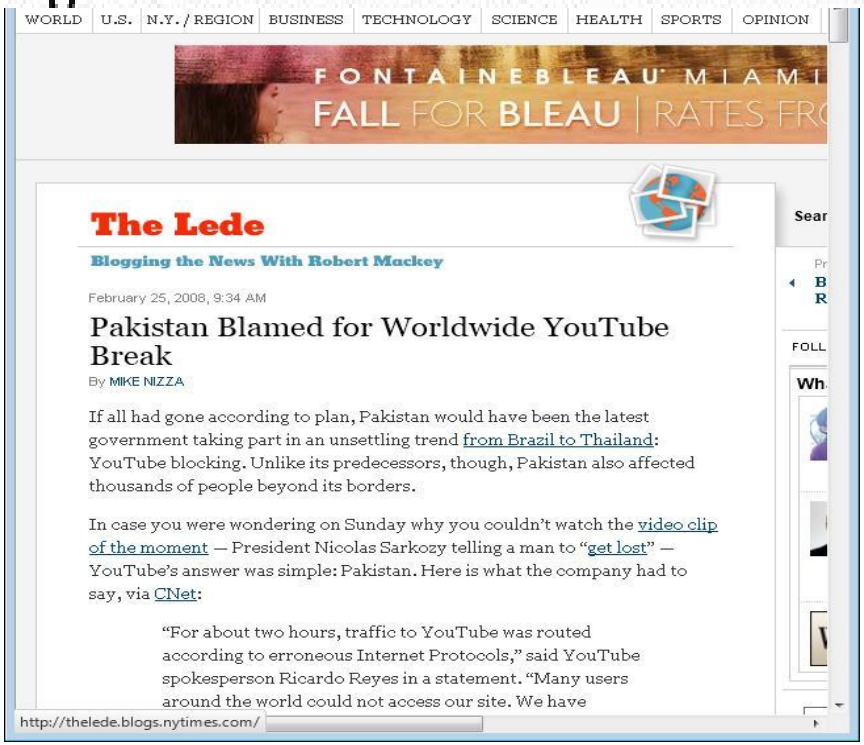

Once upon a time ... someone incorrectly announced an IP prefix.

For about 18 minutes on April 8, 2010, China Telecom advertised erroneous network traffic routes that instructed U.S. and other foreign Internet traffic to travel through Chinese servers.* Other serv-

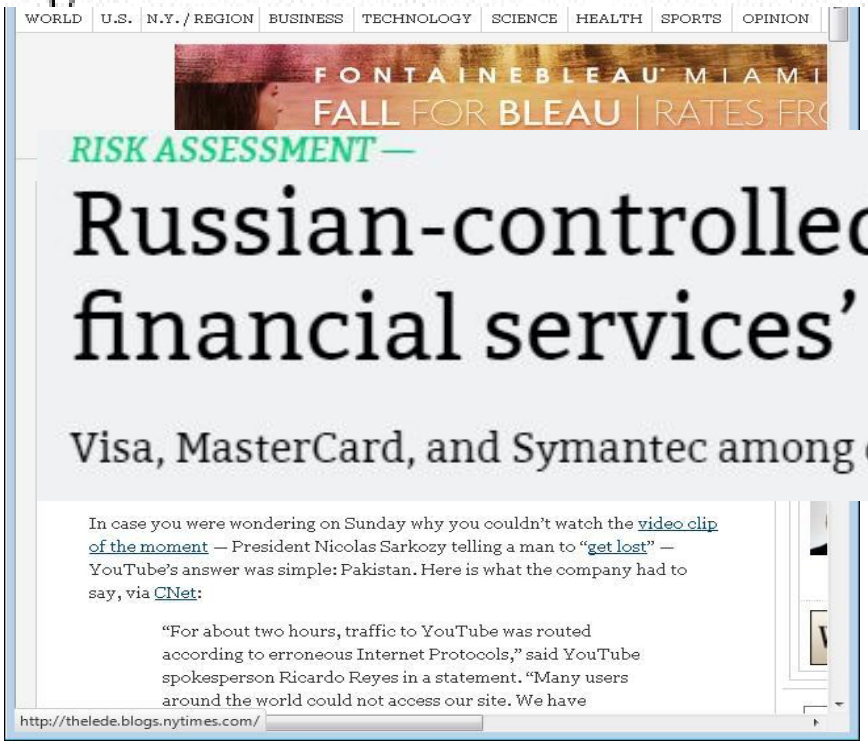

Enter RPKI

Prefix hijacking prevention using Resource Public Key Infrastructure

Enter RPKI

Prefix hijacking prevention using Resource Public Key Infrastructure

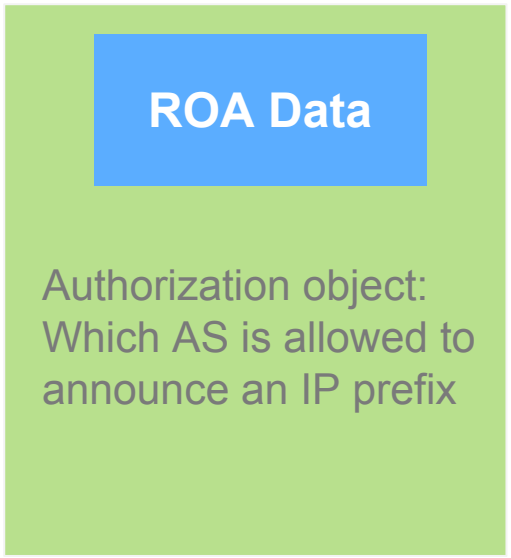

Enter RPKI

Prefix hijacking prevention using Resource Public Key Infrastructure

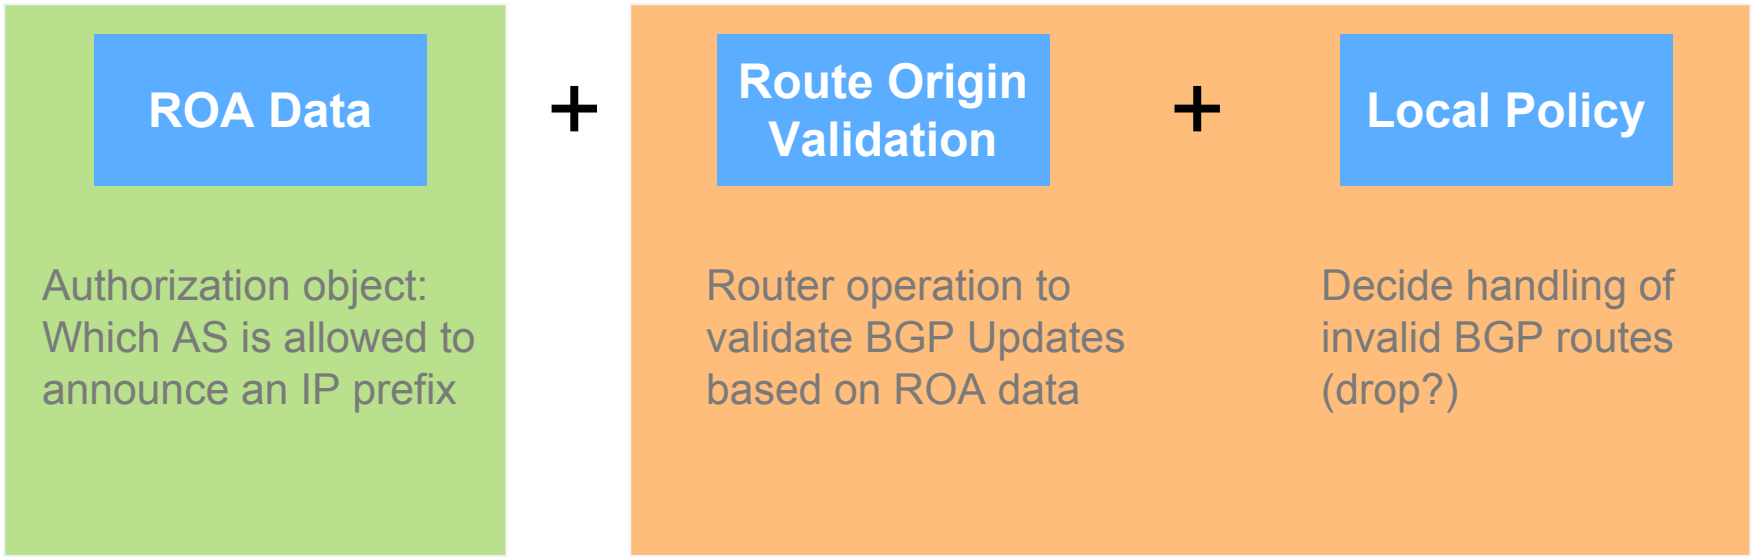

Enter RPKI

Prefix hijacking prevention using Resource Public Key Infrastructure

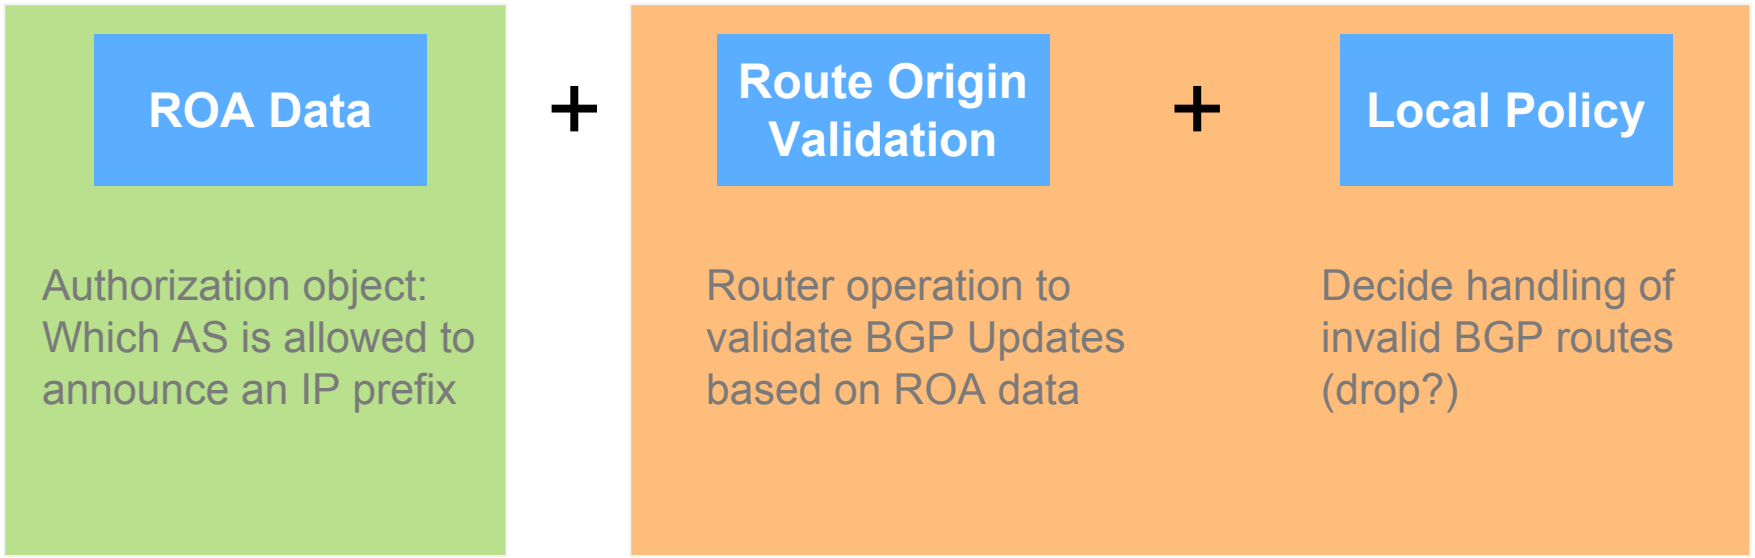

ROA: 10.20.0.0/16-24 AS100

BGP: 10.20.0.0/16 AS100 ✔ Accept
BGP: 10.20.0.0/16 AS666 ✖ Reject

Research Problem

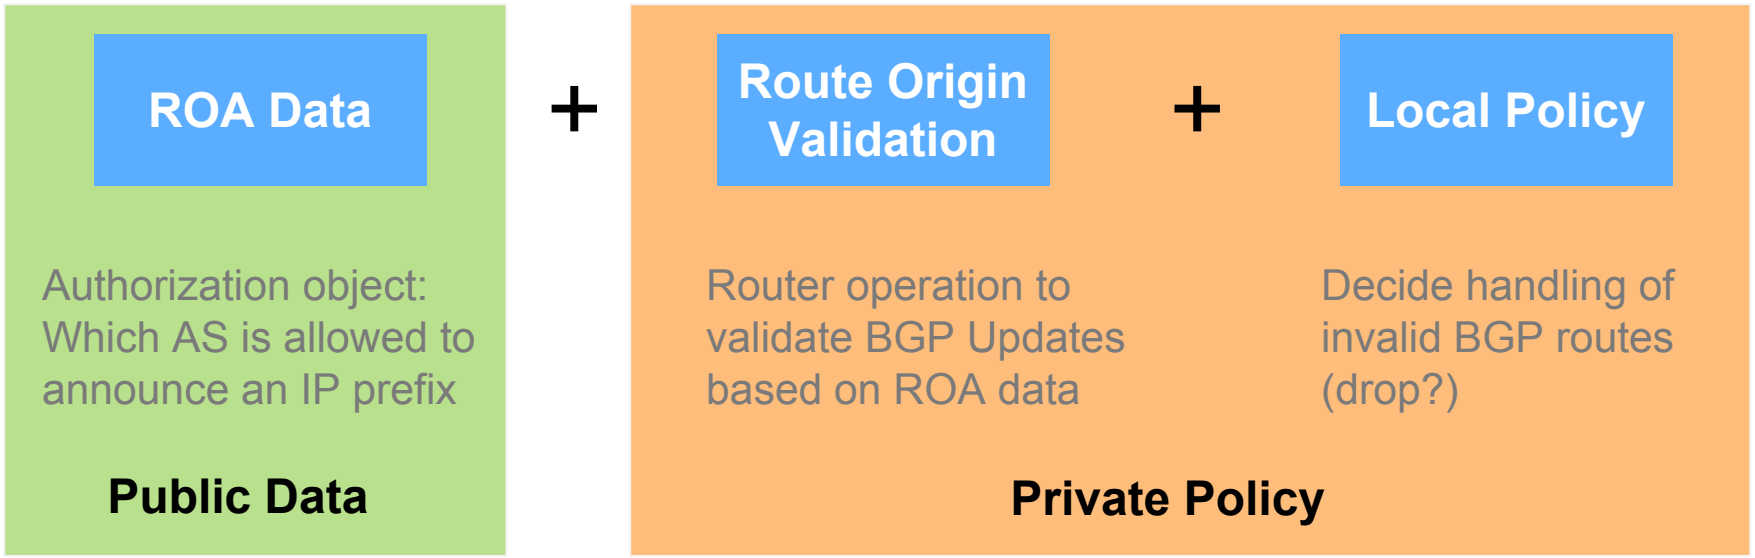

Measure the adoption of RPKI-based filter policies.

Research Challenge

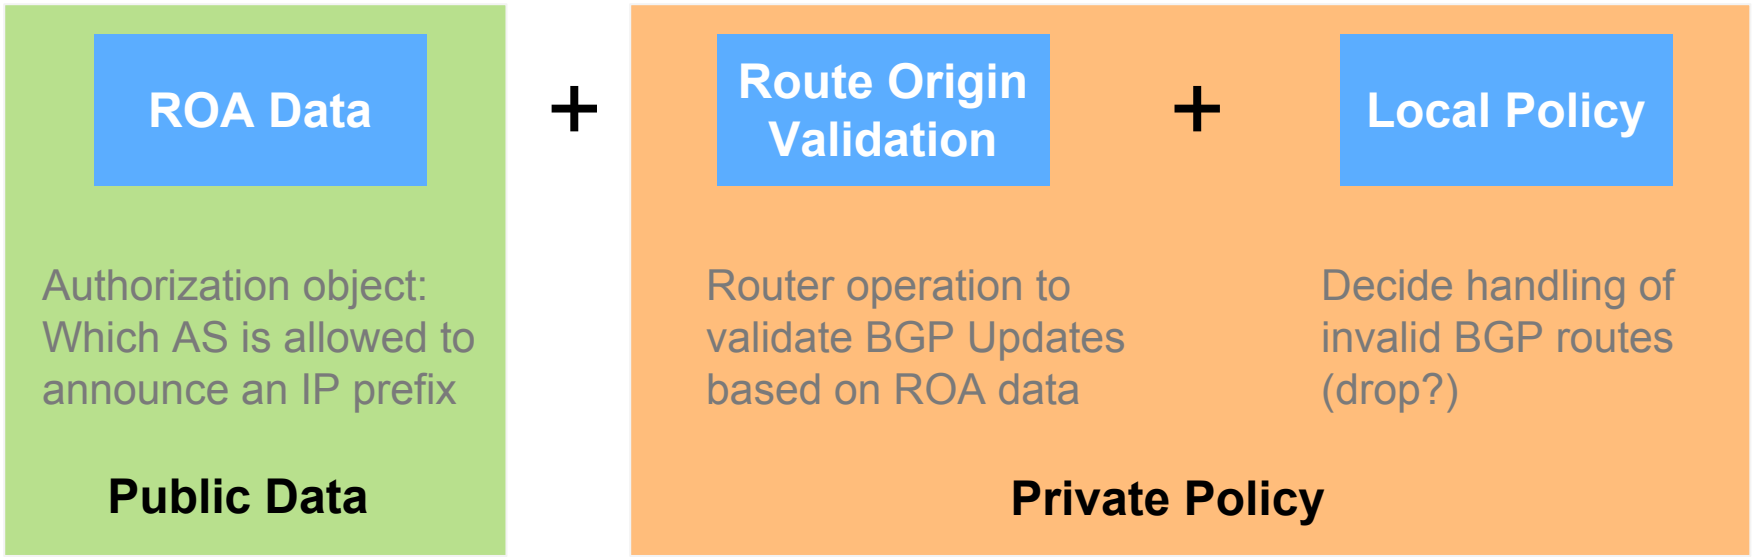

Measure the adoption of RPKI-based filter policies.

Challenge: Private policies must be inferred from measurements.

Two principle approaches

Uncontrolled experiments

Analysing existing BGP data and ROAs, trying to infer who is filtering.

- ➔ **Fast**
- ➔ **Easy**

Controlled experiments

Actively announcing BGP Updates and dynamically creating ROAs

Analyse resulting BGP data to infer who is filtering.

- ➔ **Slow**
- ➔ **Needs experimental facilities**

Uncontrolled Experiments: The Basic Idea

➔ **Leverage divergence between AS paths of invalid and non-invalid routes to infer if an AS is filtering**

Uncontrolled Experiments: The Basic Idea

➔ Leverage divergence between AS paths of invalid and non-invalid routes to infer if an AS is filtering

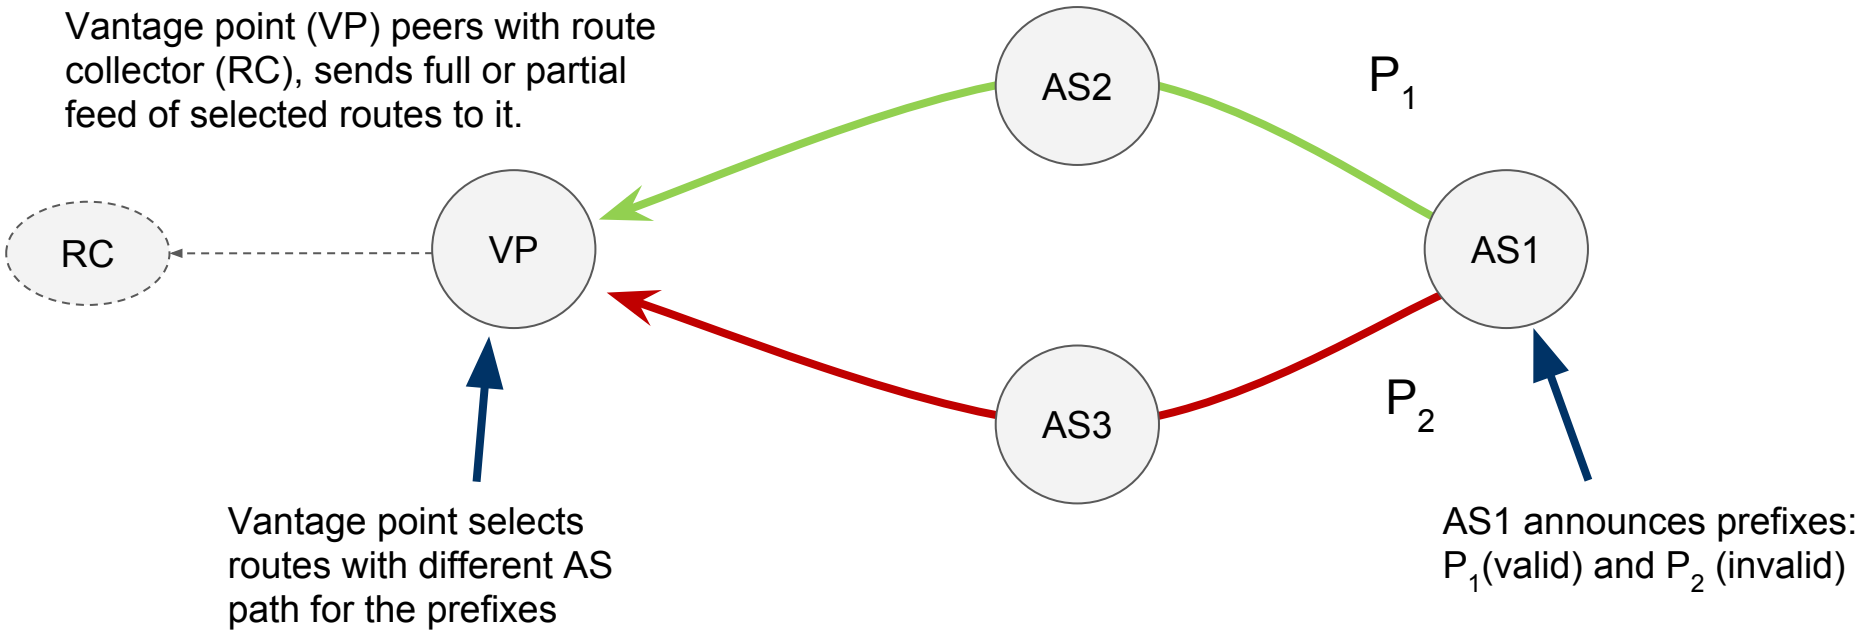

Uncontrolled Experiments: The Basic Idea

➔ Leverage divergence between AS paths of invalid and non-invalid routes to infer if an AS is filtering

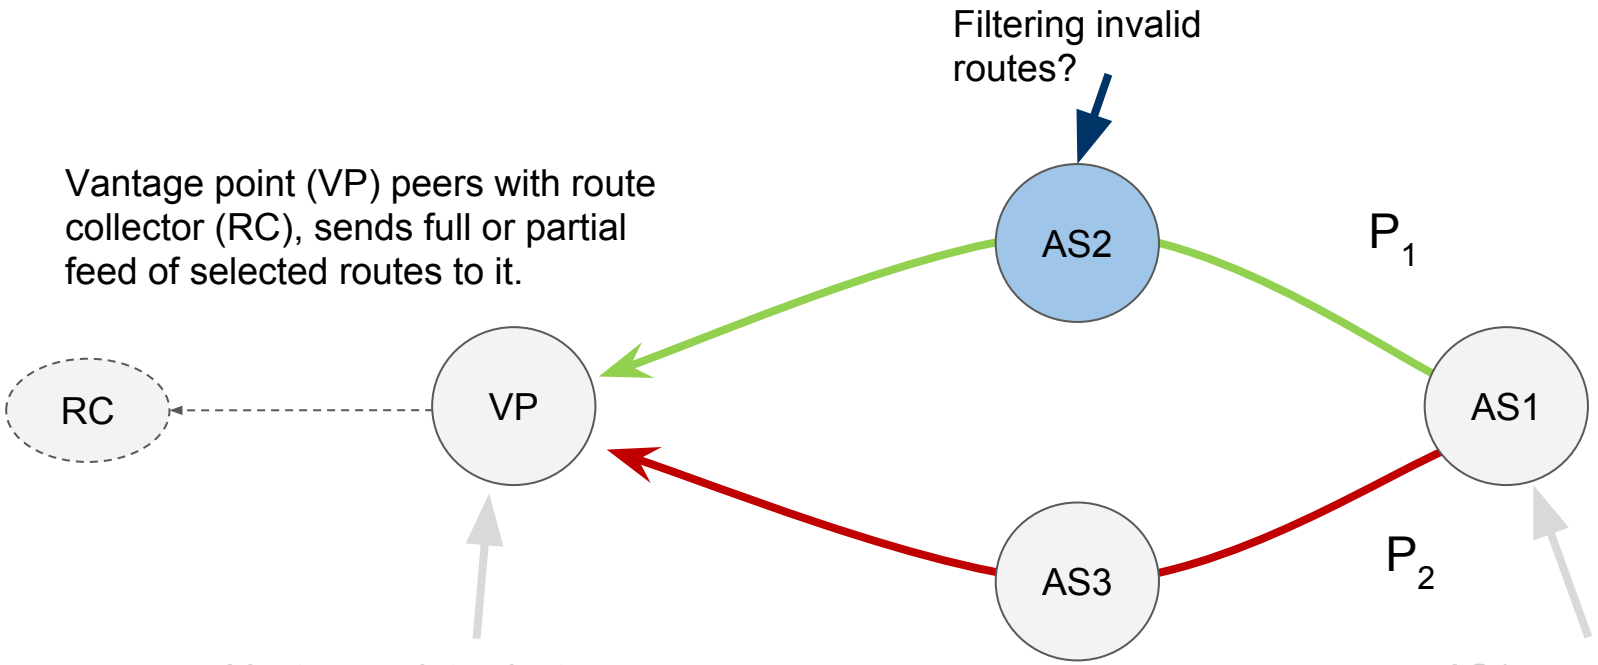

Uncontrolled Experiments: Problems

Uncontrolled Experiments: Problems

➔ Limited Control

Uncontrolled Experiments: Problems

- ➔ Limited Control
  - ◆ Do not know origin AS policy. Traffic engineering might look like RPKI-based filtering.

Uncontrolled Experiments: Limited Control

Origin Policy

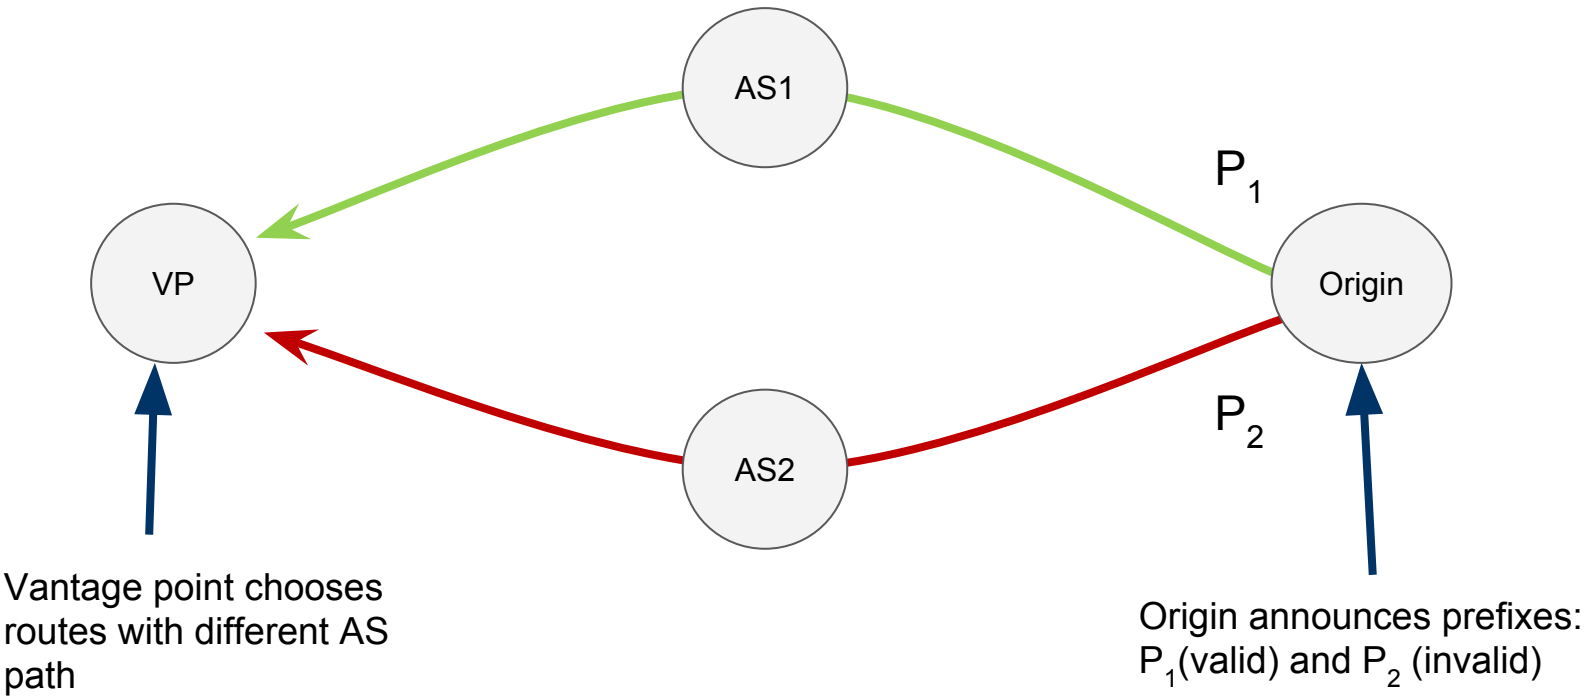

Uncontrolled Experiments: Limited Control

Origin Policy

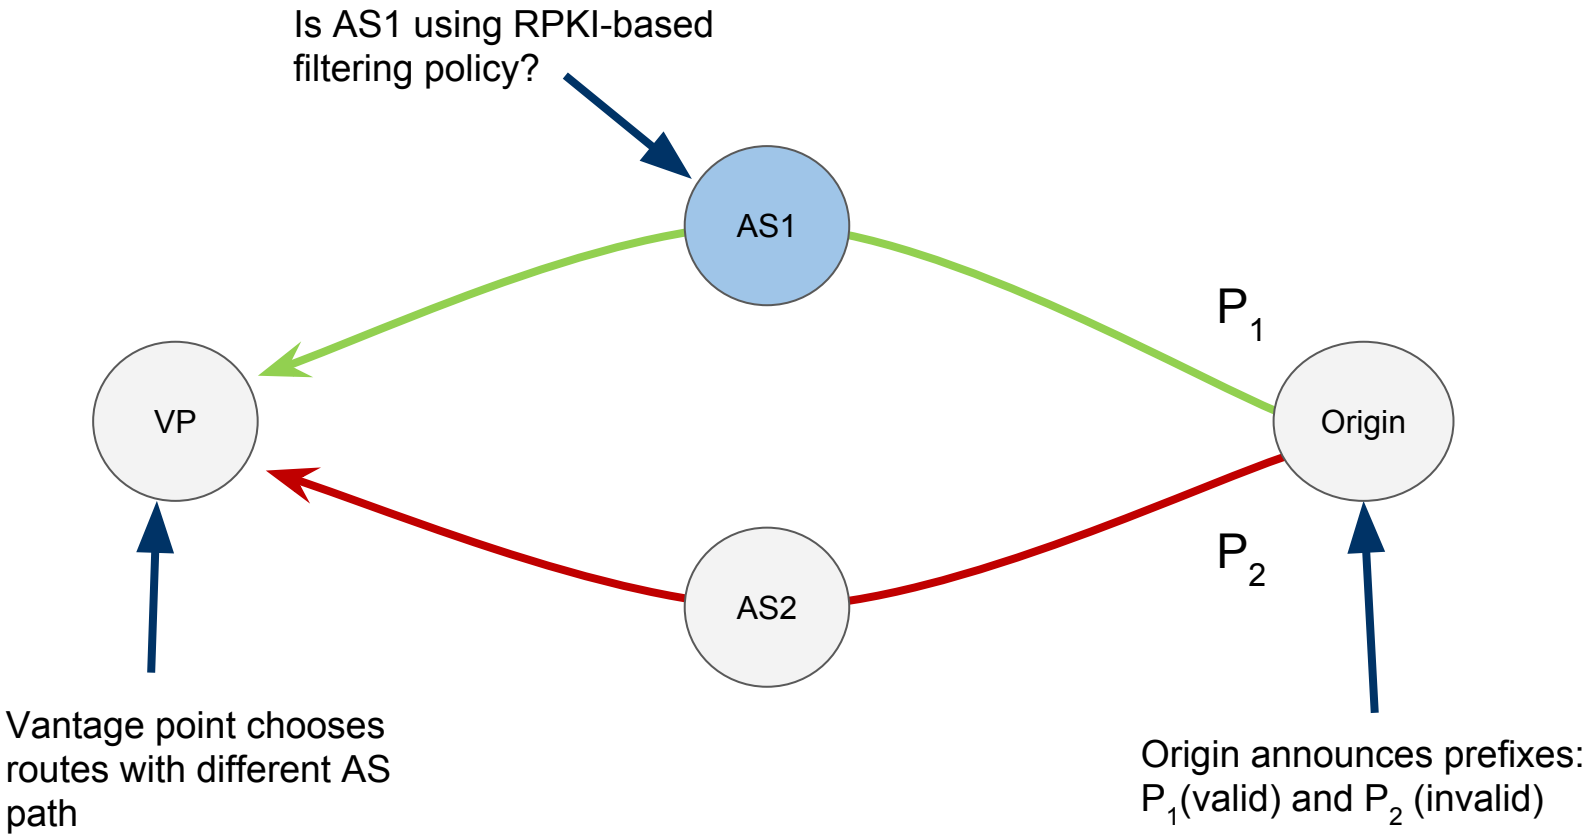

Uncontrolled Experiments: Limited Control

Origin Policy

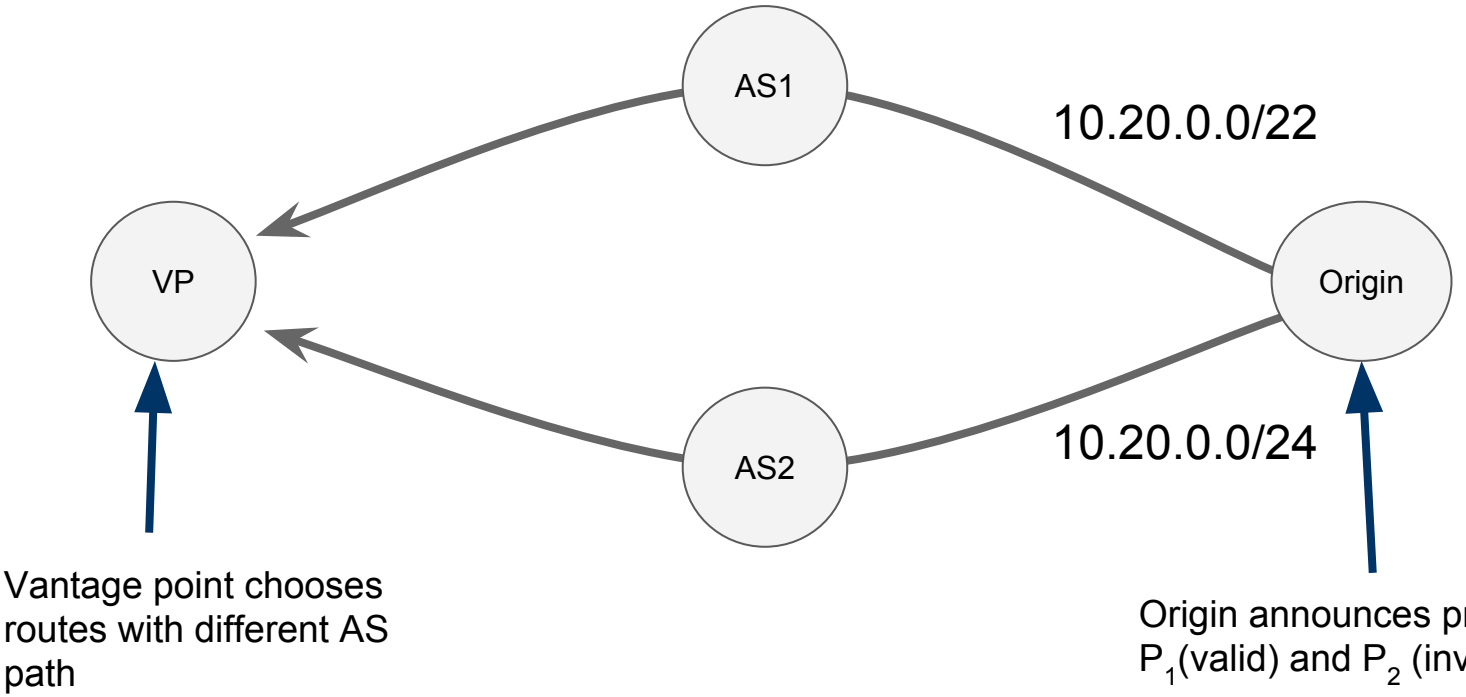

Uncontrolled Experiments: Limited Control

Origin Policy

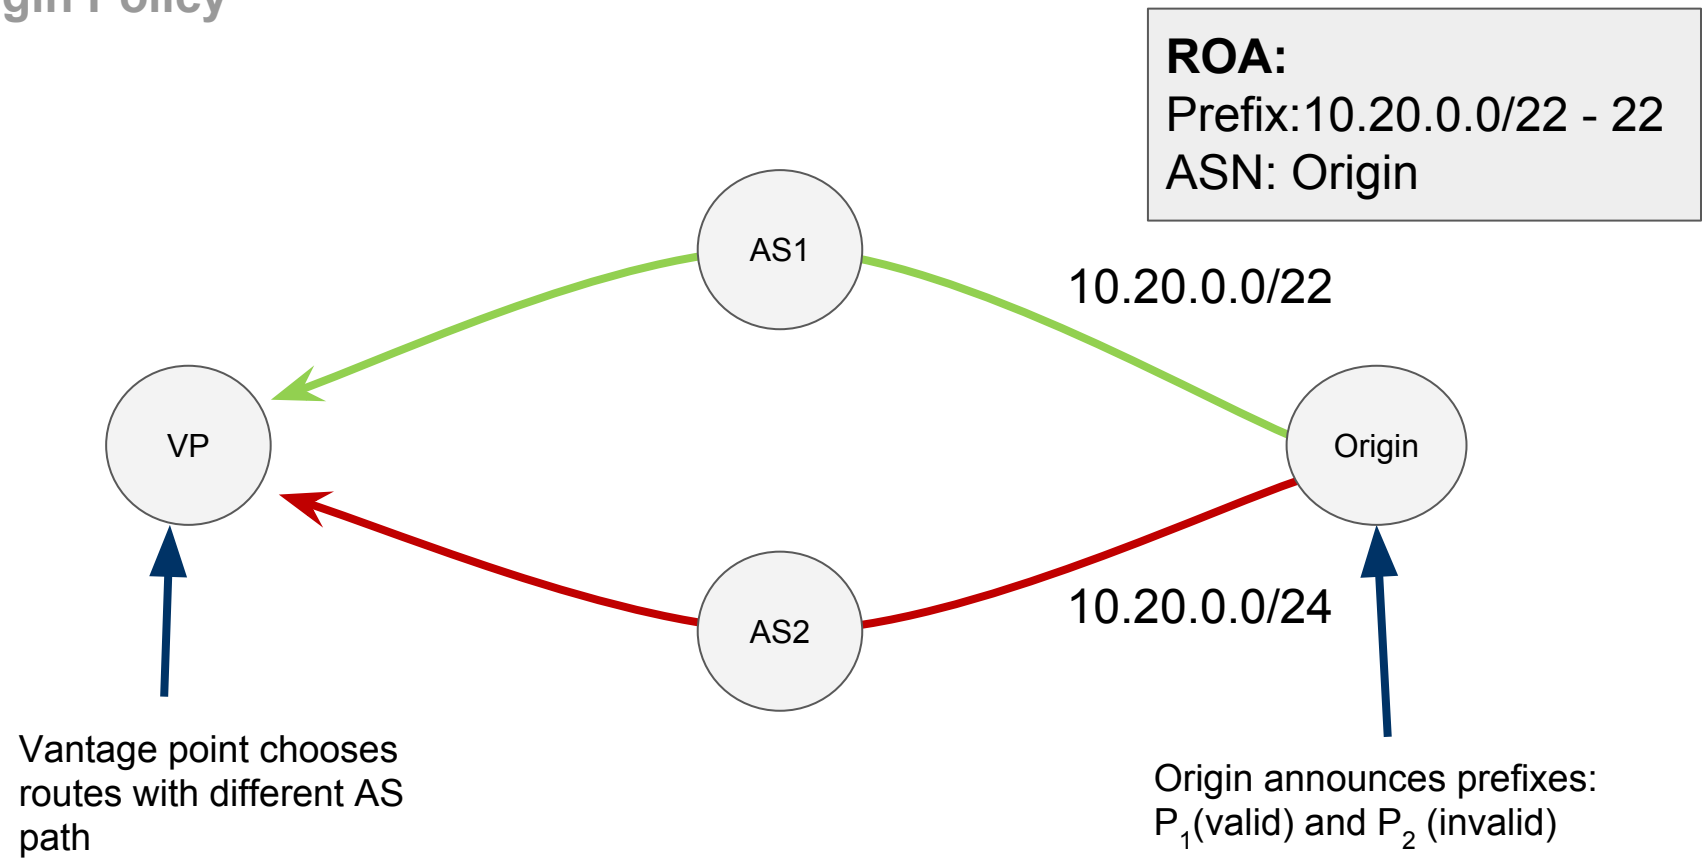

Uncontrolled Experiments: Limited Control

Origin Policy

Is AS1 using RPKI-based filtering policy?

Path divergence at first hop is more likely to be the result of traffic engineering at origin.

Vantage point chooses routes with different AS path

Origin announces prefixes: P_1(valid) and P_2 (invalid)

Path Divergence

Divergence between AS paths of routes with the same origin

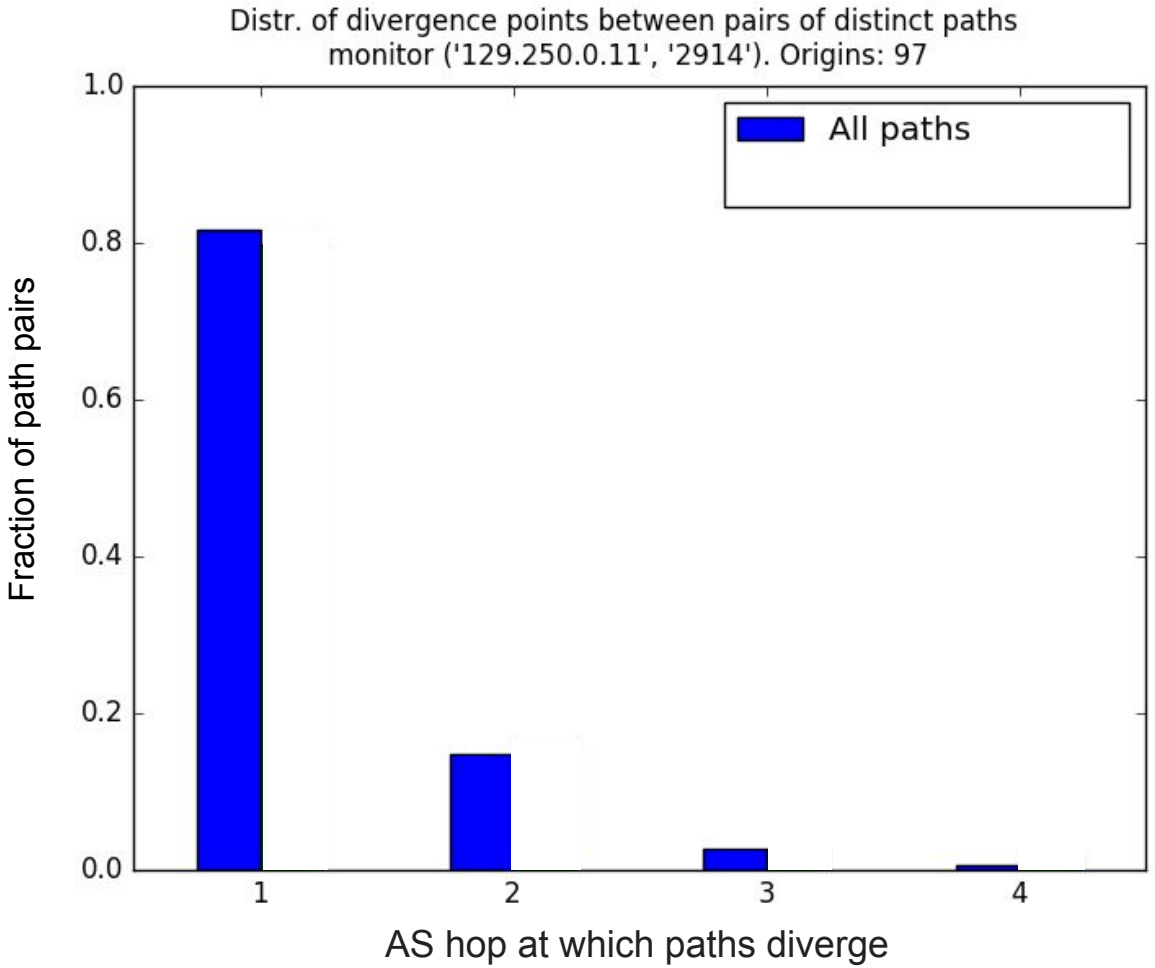

Path Divergence

Divergence between AS paths of routes with the same origin

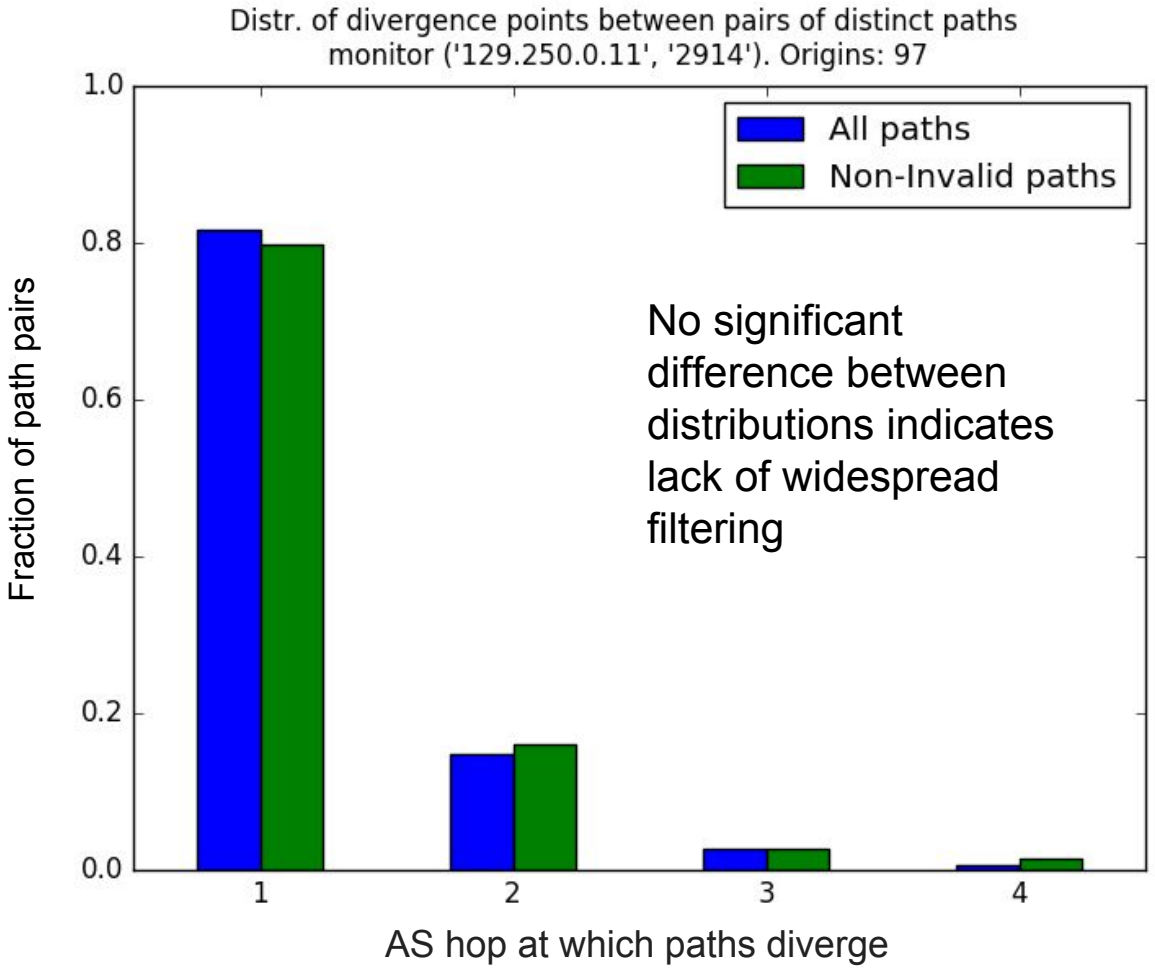

➔ Invalid routes (probably) have different AS paths for non-RPKI reasons

Uncontrolled Experiments: Problems

➔ Limited Control

- Do not know origin AS policy. Traffic engineering might look like RPKI-based filtering.
- Cannot distinguish between filtering based on RPKI vs. filtering based on other attributes

Uncontrolled Experiments: Limited Control

Real World Example

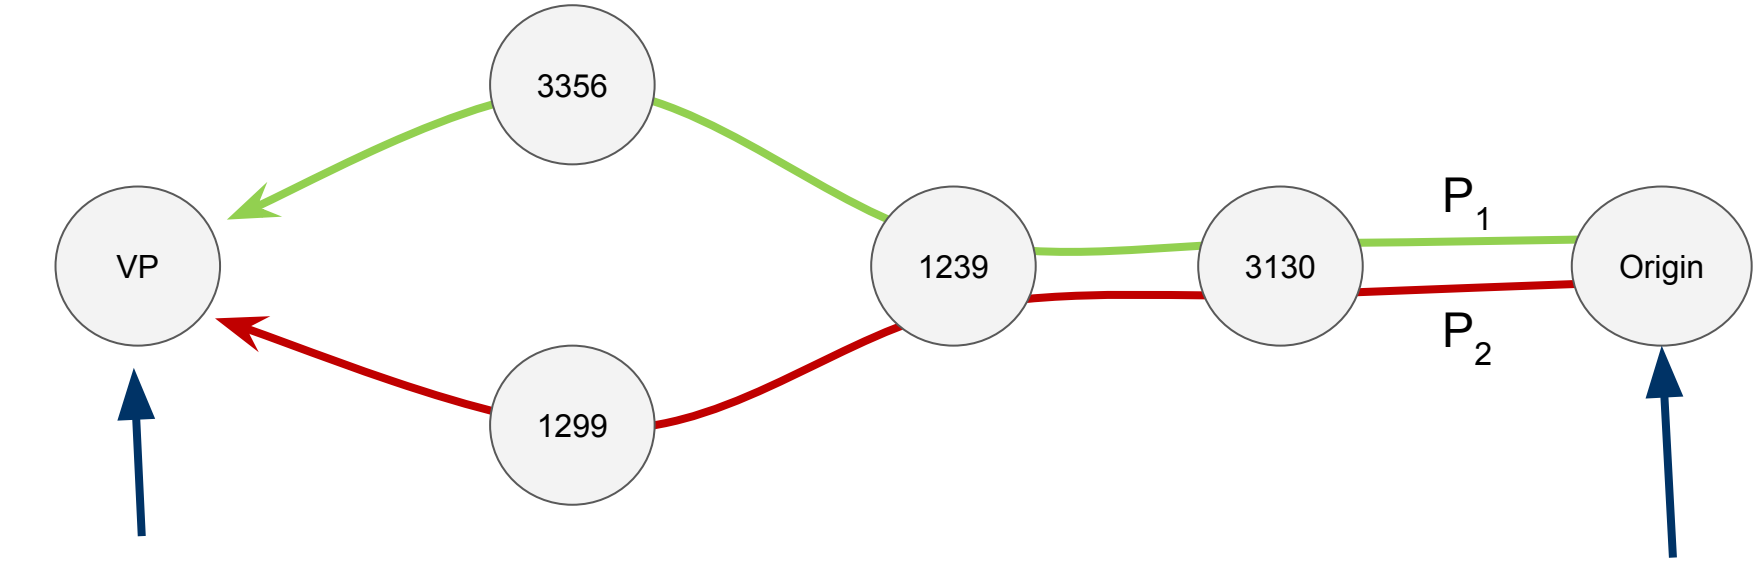

Uncontrolled Experiments: Limited Control

Real World Example

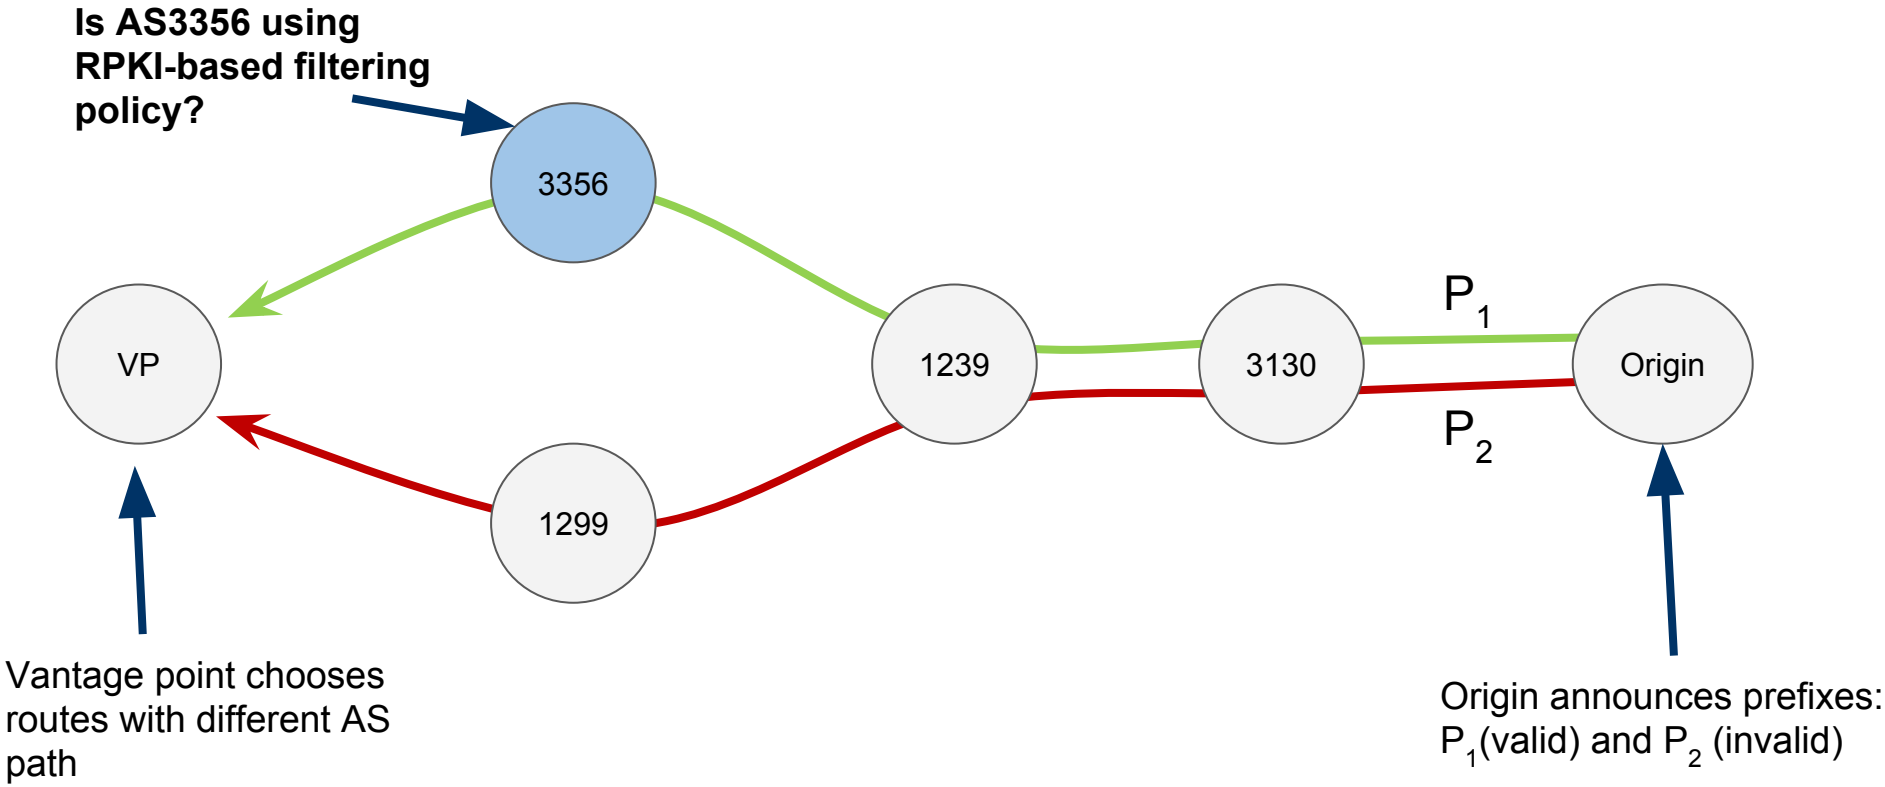

Uncontrolled Experiments: Limited Control

Real World Example

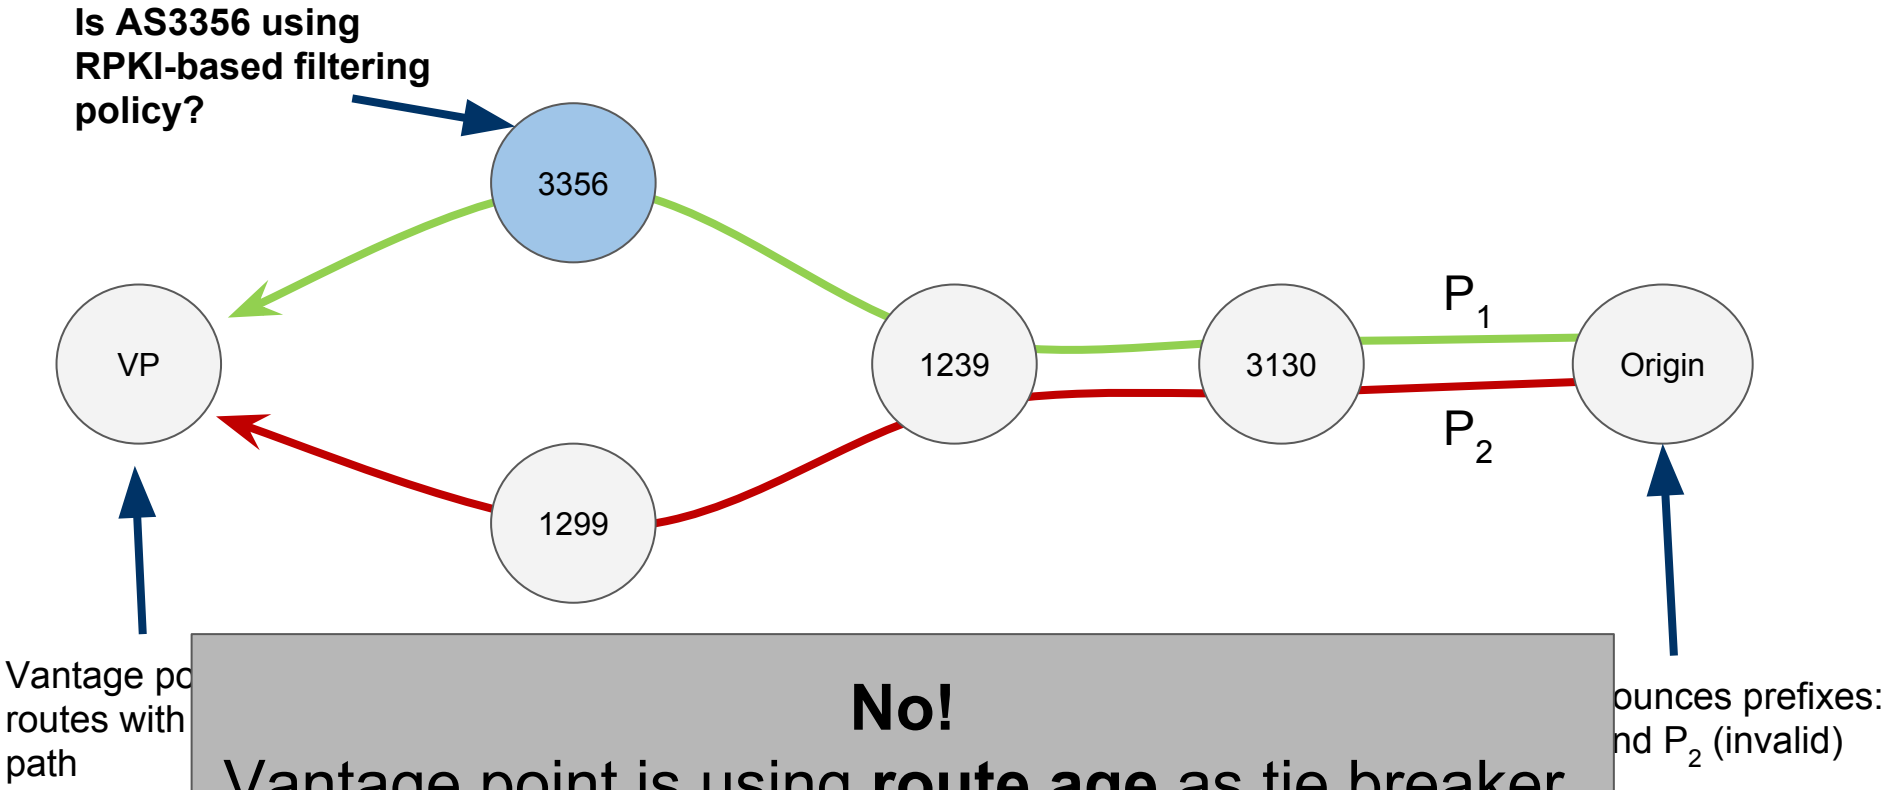

No!

Vantage point is using **route age** as tie breaker.

Uncontrolled Experiments: Problems

➔ Limited Control

- Do not know origin AS policy. Traffic engineering might look like RPKI-based filtering.
- Cannot distinguish between filtering based on RPKI vs. filtering based on other attributes

Uncontrolled Experiments: Problems

- ➔ Limited Control
  - ◆ Do not know origin AS policy. Traffic engineering might look like RPKI-based filtering.
  - ◆ Cannot distinguish between filtering based on RPKI vs. filtering based on other attributes
- ➔ Limited Visibility can lead to misclassification

Uncontrolled Experiments: Limited Visibility

➔ Analysing data from different sets of vantage points can yield different classifications

Uncontrolled Experiments: Limited Visibility

➔ Analysing data from different sets of vantage points can yield different classifications

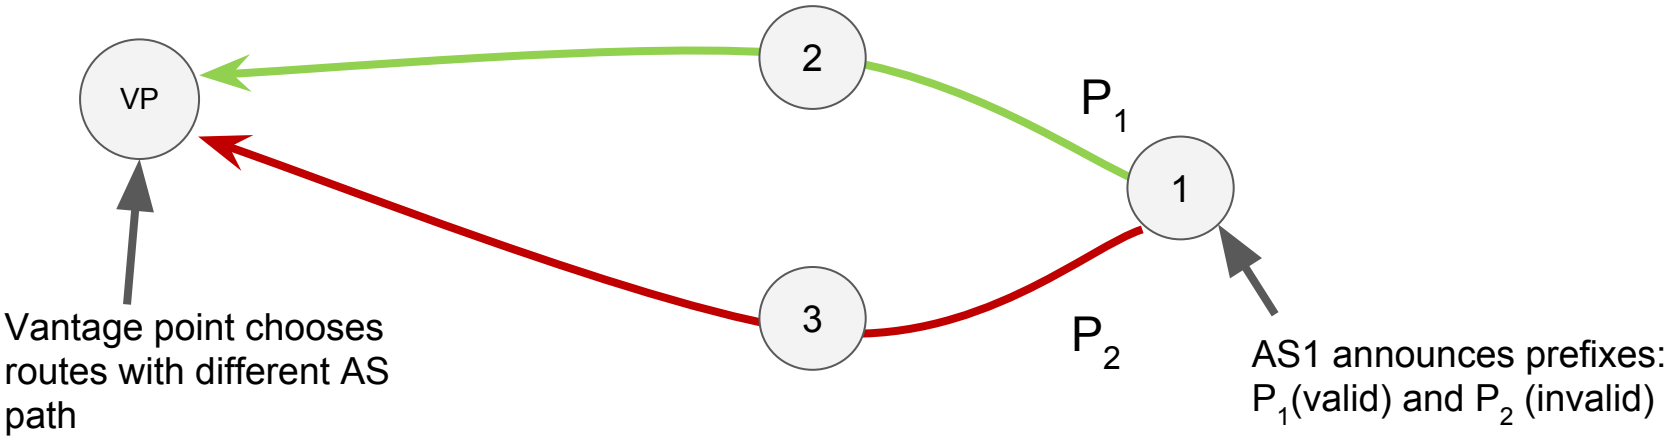

Uncontrolled Experiments: Limited Visibility

➔ Analysing data from different sets of vantage points can yield different classifications

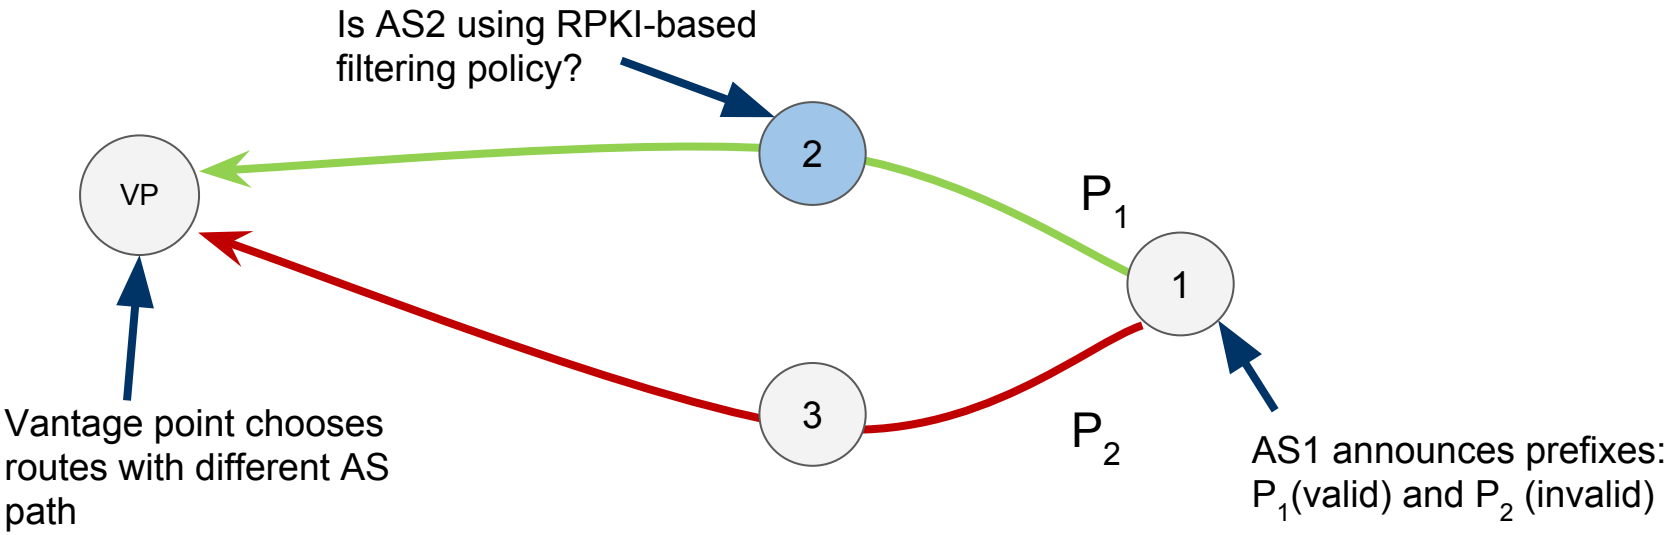

Uncontrolled Experiments: Limited Visibility

➔ Analysing data from different sets of vantage points can yield different classifications

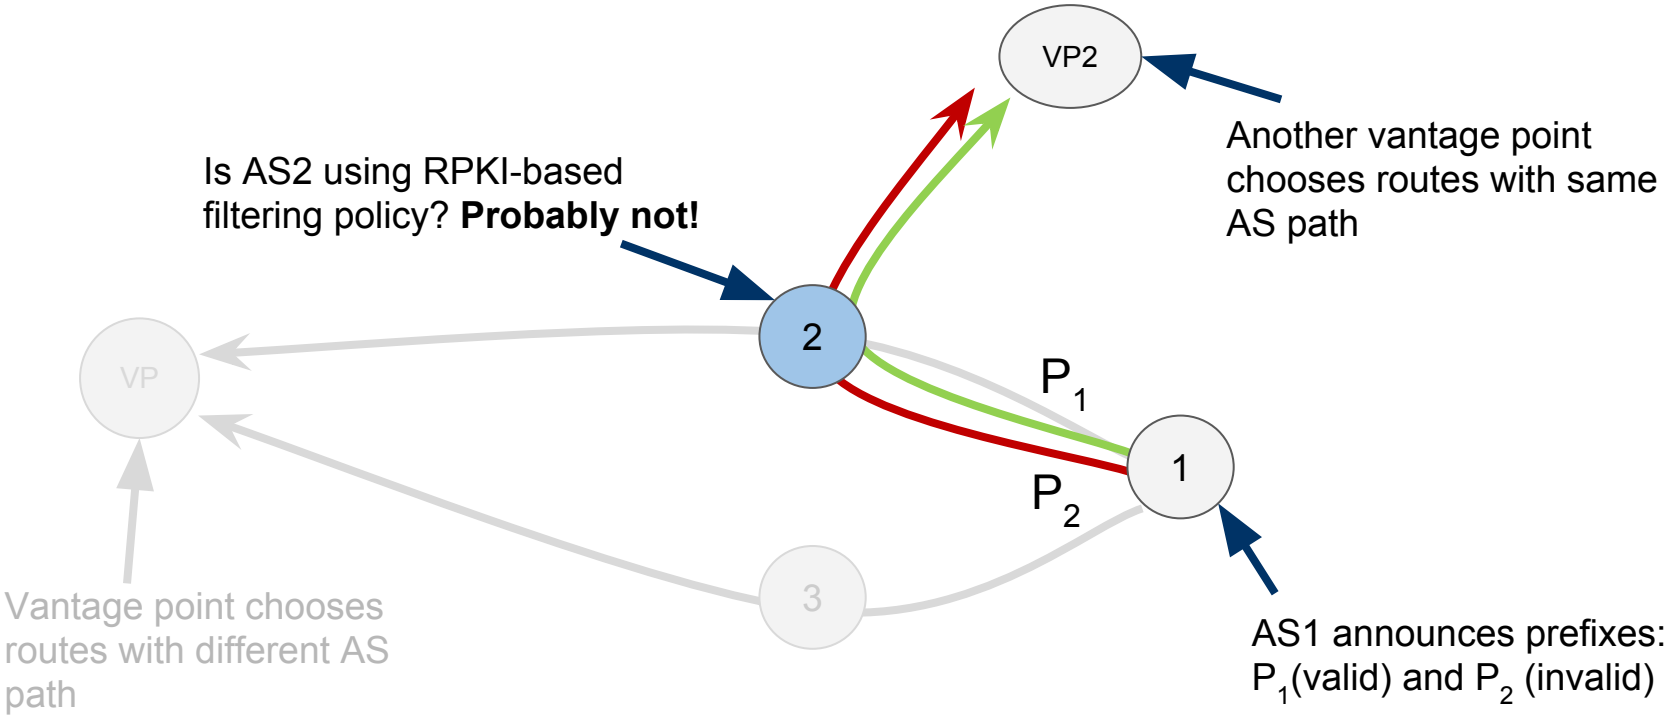

Uncontrolled Experiments: Limited Visibility

➔ Analysing data from different sets of vantage points can yield different classifications

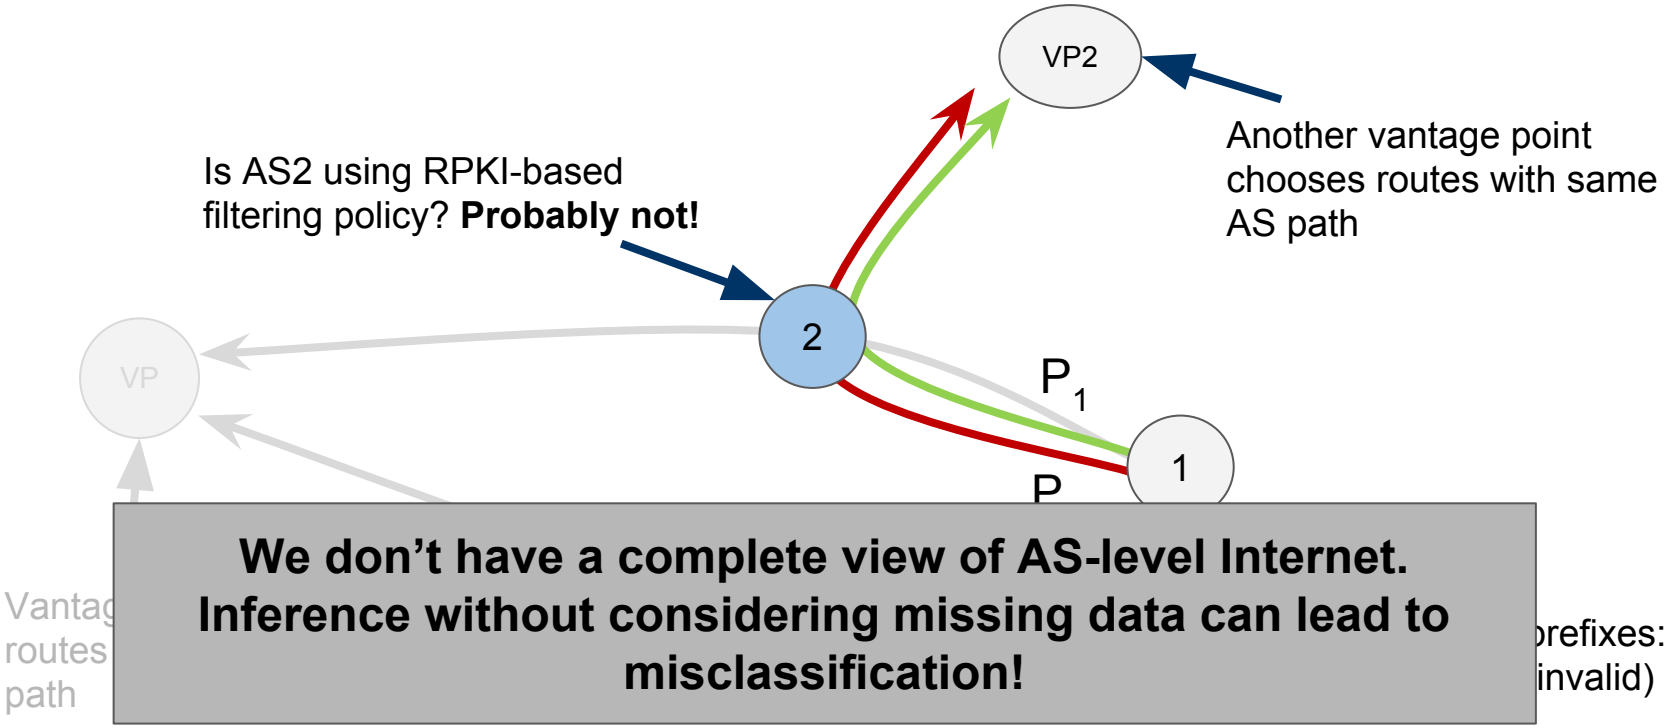

Uncontrolled Experiments: Problems

- ➔ Limited Control
  - ◆ Do not know origin AS policy. Traffic engineering might look like RPKI-based filtering.
  - ◆ Cannot distinguish between filtering based on RPKI vs. filtering based on other attributes
- ➔ Limited Visibility can lead to misclassification

Uncontrolled Experiments: Problems

- ➔ Limited Control
  - ◆ Do not know origin AS policy. Traffic engineering might look like RPKI-based filtering.
  - ◆ Cannot distinguish between filtering based on RPKI vs. filtering based on other attributes
- ➔ Limited Visibility can lead to misclassification
- ➔ Not possible to reproduce

Uncontrolled Experiments: Problems

- Limited Control
  - D… …ering m…
  - C… …on R…
- Limit…
- Not p…

Inferring if a specific AS is using RPKI-based filtering on the basis of uncontrolled experiments is prone to misclassification!

Controlled Experiments

Controlled Experiments

Hand-crafted ROAs *and* BGP Updates

Controlled Experiments: Advantages

Hand-crafted ROAs *and* BGP Updates

➔ Limited Control

- We know the routing policy of origin AS

Controlled Experiments: Advantages

Hand-crafted ROAs *and* BGP Updates

➔ Limited Control

- We know the routing policy of origin AS
- Can distinguish between RPKI-based filtering vs. filtering based on other attributes by changing ROAs/Updates

Controlled Experiments: Advantages

Hand-crafted ROAs *and* BGP Updates

- ➔ Limited Control
  - ◆ We know the routing policy of origin AS
  - ◆ Can distinguish between RPKI-based filtering vs. filtering based on other attributes by changing ROAs/Updates
- ➔ Limited Visibility is less of an issue, we only care about our prefixes

Controlled Experiments: Advantages

Hand-crafted ROAs *and* BGP Updates

- ➔ Limited Control
  - ◆ We know the routing policy of origin AS
  - ◆ Can distinguish between RPKI-based filtering vs. filtering based on other attributes by changing ROAs/Updates
- ➔ Limited Visibility is less of an issue, we only care about our prefixes
- ➔ Can repeat experiments and target specific AS.

Controlled Experiments: Our Setup

BGP

Announce prefixes P_A (Anchor) and P_E (Experiment)

+ Same RIR DB route object
+ Same length
+ Minimal bit difference
+ Announced at the same time
+ Announced from same origin AS
+ Announced to same peers

RPKI

Issue ROAs for both prefixes

Periodically change ROA for experiment prefix

➔ Flips announcement from VALID to INVALID to VALID once a day

(Yes, we operate a grandchild RPKI CA ;))

Controlled Experiments: Observations

Situation: Origin and vantage point peer directly

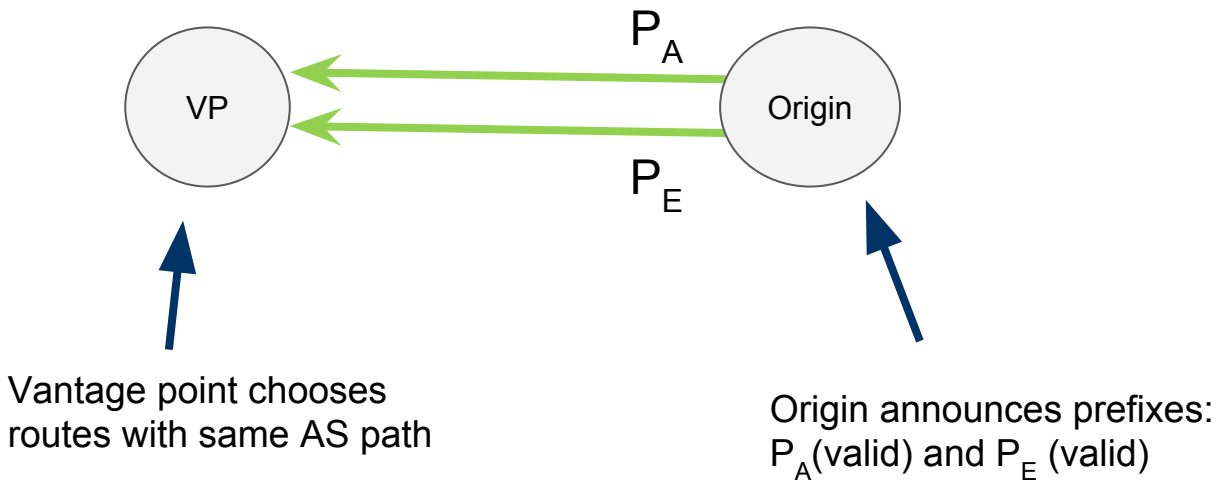

Controlled Experiments: Observations

Situation: Origin and vantage point peer directly

Vantage point chooses routes with same AS path

Origin announces prefixes: P_A(valid) and P_E (valid)

Controlled Experiments: Observations

Situation: Origin and vantage point peer directly

Observation 1: VP has no route for P_E now that it's announcement is invalid

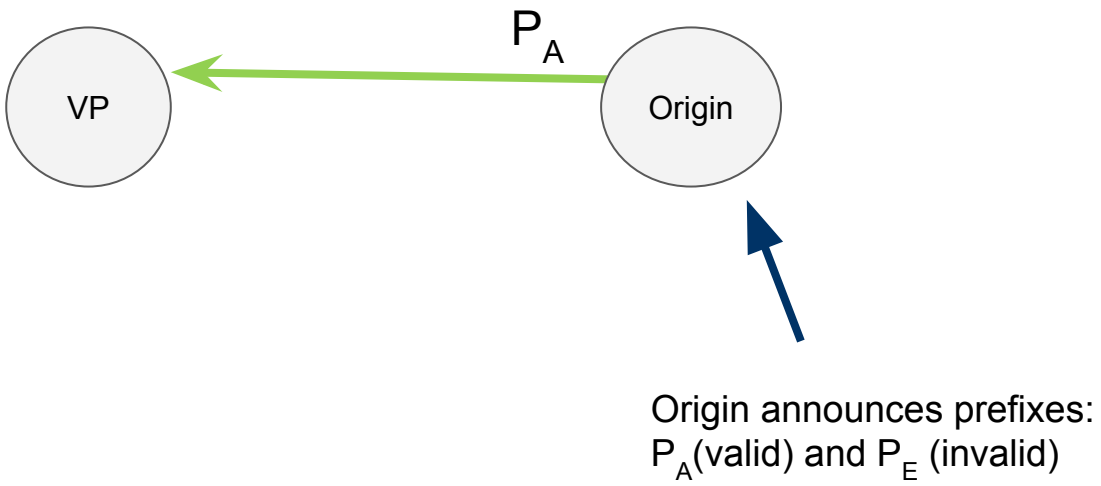

Conclusion: VP is using RPKI-based filtering.

Controlled Experiments: Observations

Situation: Origin and vantage point peer directly

Observation 2: VP has route via AS X for P_E now that it's announcement is invalid

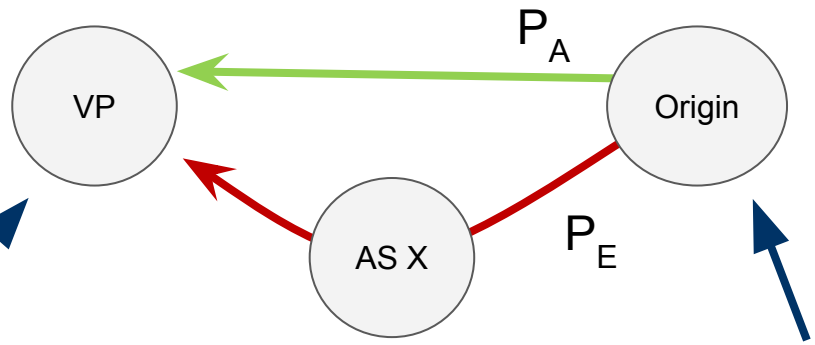

Vantage point chooses routes with different AS path

Origin announces prefixes: P_A(valid) and P_E (invalid)

Conclusion: VP uses RPKI-based filtering **selectively**.

Controlled Experiments: Observations

Situation: Origin and vantage point do not peer directly, other AS on path

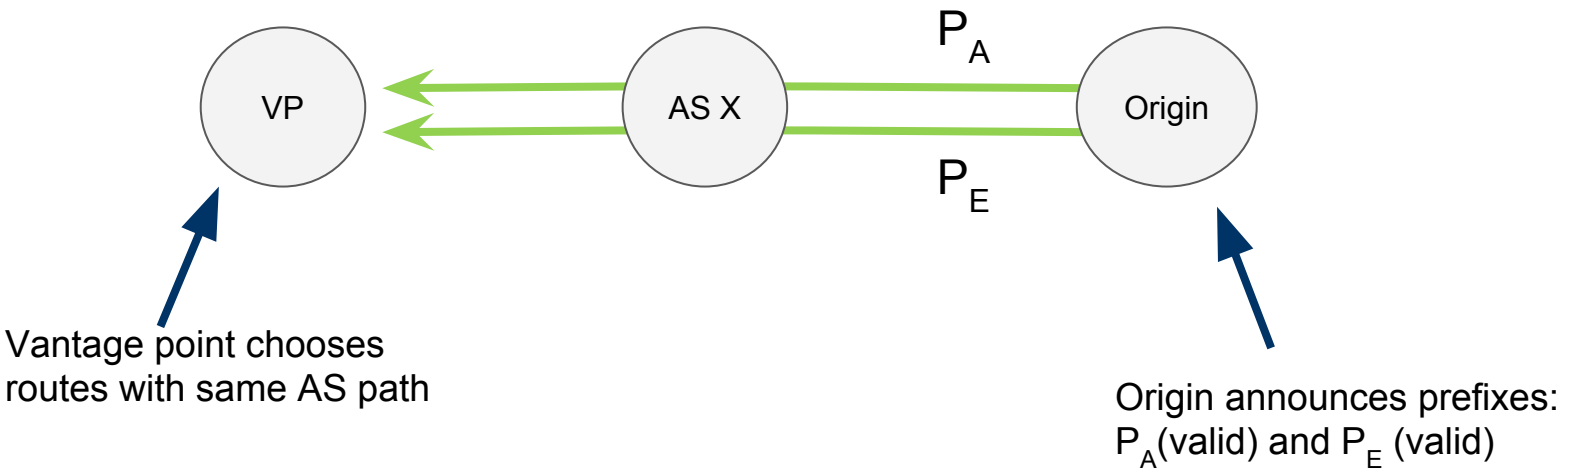

Controlled Experiments: Observations

Situation: Origin and vantage point do not peer directly, other AS on path

Controlled Experiments: Observations

Situation: Origin and vantage point do not peer directly, other AS on path

Observation 1: VP has no route for P_E now that it's announcement is invalid

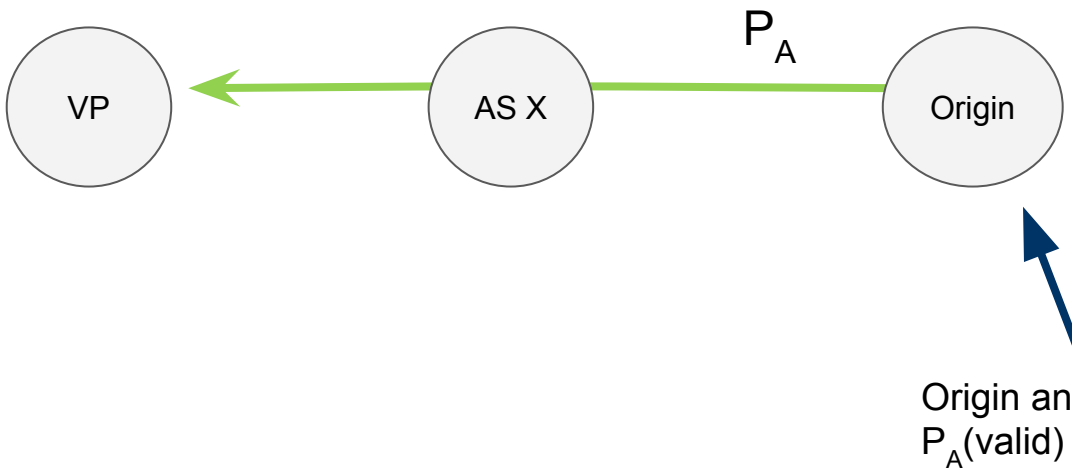

Conclusion: VP or AS X (or both) are using RPKI-based filtering.

Controlled Experiments: Observations

Situation: Origin and vantage point do not peer directly, other AS on path

Observation 2: VP has different route for P_E now that it's announcement is invalid

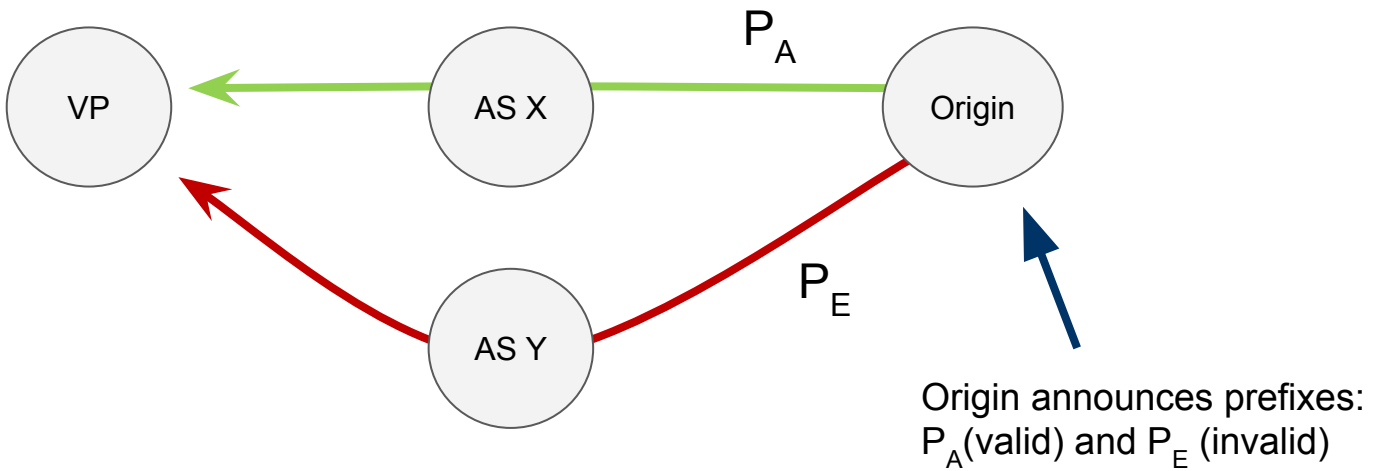

Conclusion: VP or AS X (or both) are using RPKI-based filtering.

Controlled Experiments: Observations

Situation: Origin and vantage point do not peer directly, other AS on path

Observation 2: VP has different route for P_E now that it's announcement is invalid

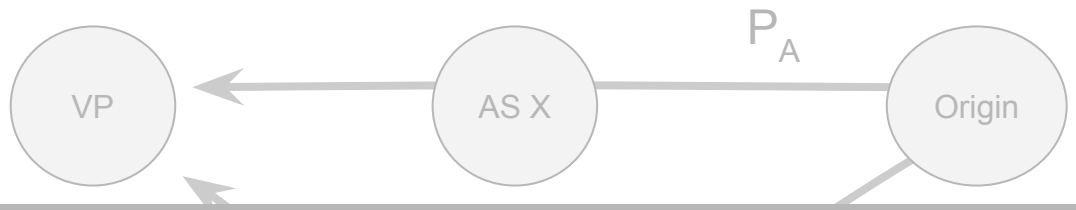

Resolve ambiguity by:

➔ Establishing direct peering with VP

Controlled Experiments: Observations

Situation: Origin and vantage point do not peer directly, other AS on path

Observation 2: VP has different route for P_E now that it's announcement is invalid

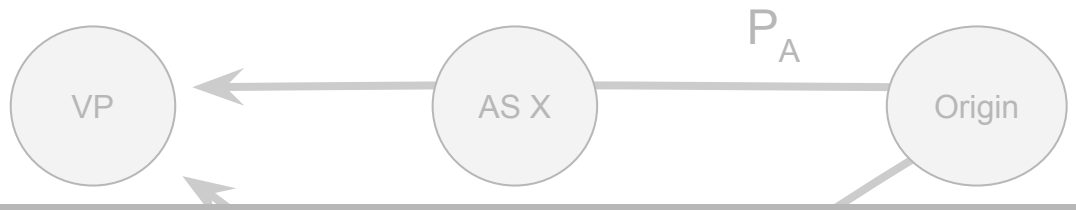

Resolve ambiguity by:

➔ Establishing direct peering with VP

➔ Checking if AS X has a vantage point

Results

Results

We found at least 3 AS that deployed RPKI-based filtering!

None of them are large providers ...

2 AS filtered all invalid routes

1 AS filtered selectively

Another measurement study found other results.

Results

We found at least 3 AS that deployed RPKI-based filtering!

None

2

Confirmed by repeated experiments and talking to operators.

Another

Conclusion

➔ There are ASes that do RPKI-based filtering.
Not many, not the big ones, but at least some (>3).

➔ Uncontrolled experiments are unsuited to infer RPKI-based filtering policies

➔ Controlled experiments are crucial to measuring adoption of RPKI-based filtering policies

Internet infrastructure requires proper monitoring.

Next Steps

➔ We will extend our measurement methodology.

➔ We will establish a live monitoring system with public access.

BGP monitoring is based on collaboration!

➔ Please, establish direct peering with PEERING testbed.
- ◆ https://peering.usc.edu/peering/

➔ Please, peer with public route collectors.

Next Steps

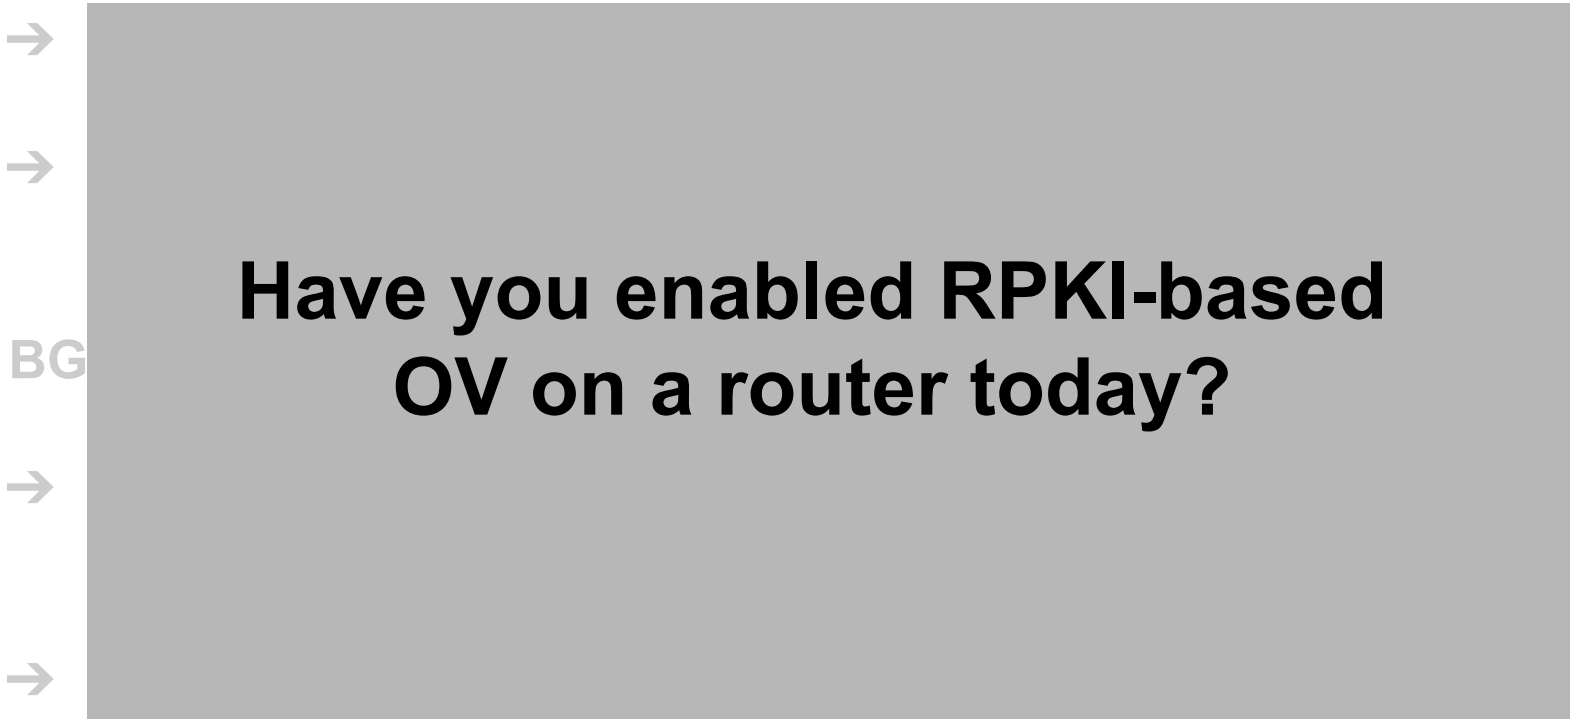

Backup

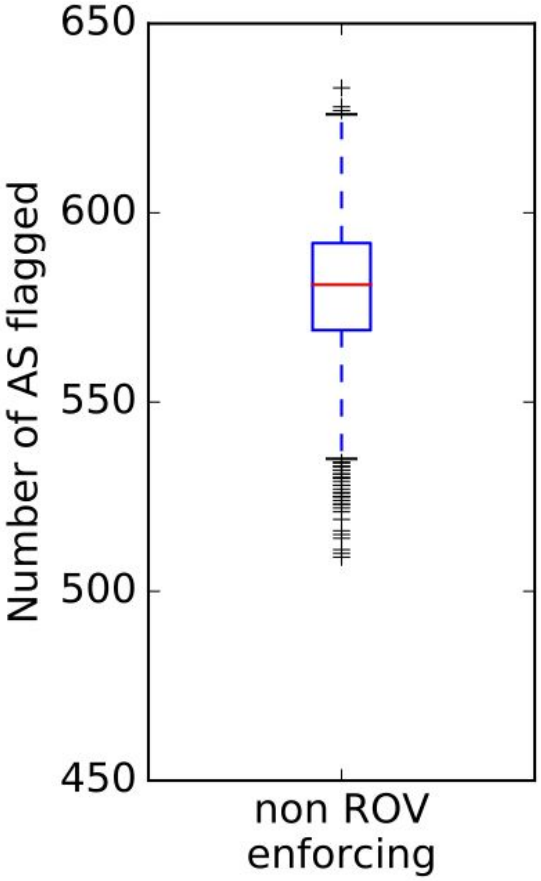
650
600
550
500
450
Number of AS flagged
non ROV
enforcing

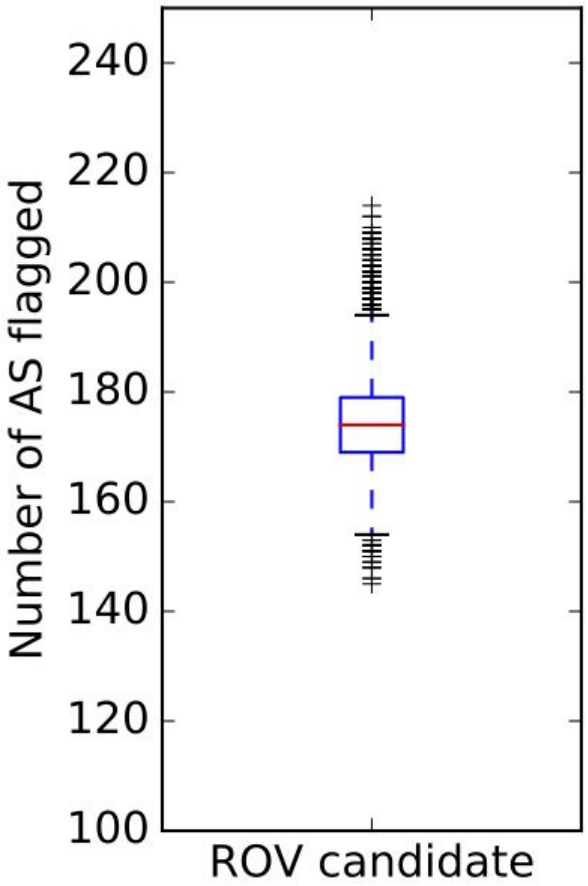
240
220
200
180
160
140
120
100
Number of AS flagged
ROV candidate

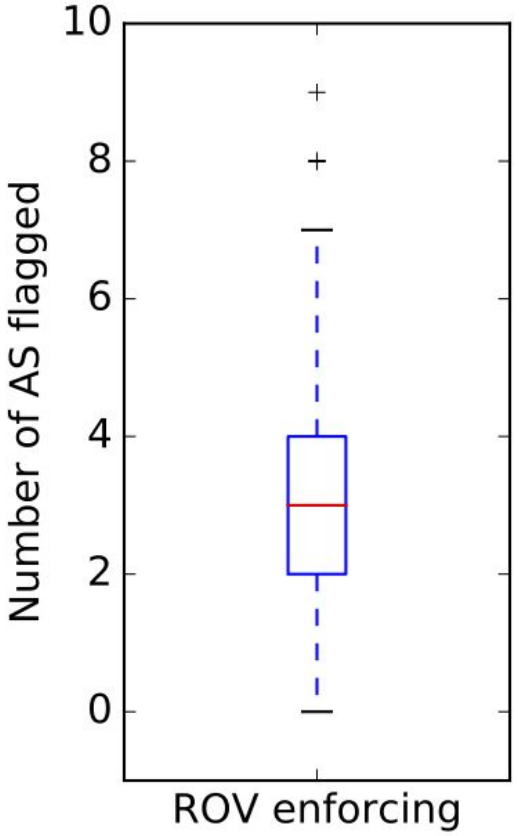
10
8
6
4
2
0
Number of AS flagged
ROV enforcing

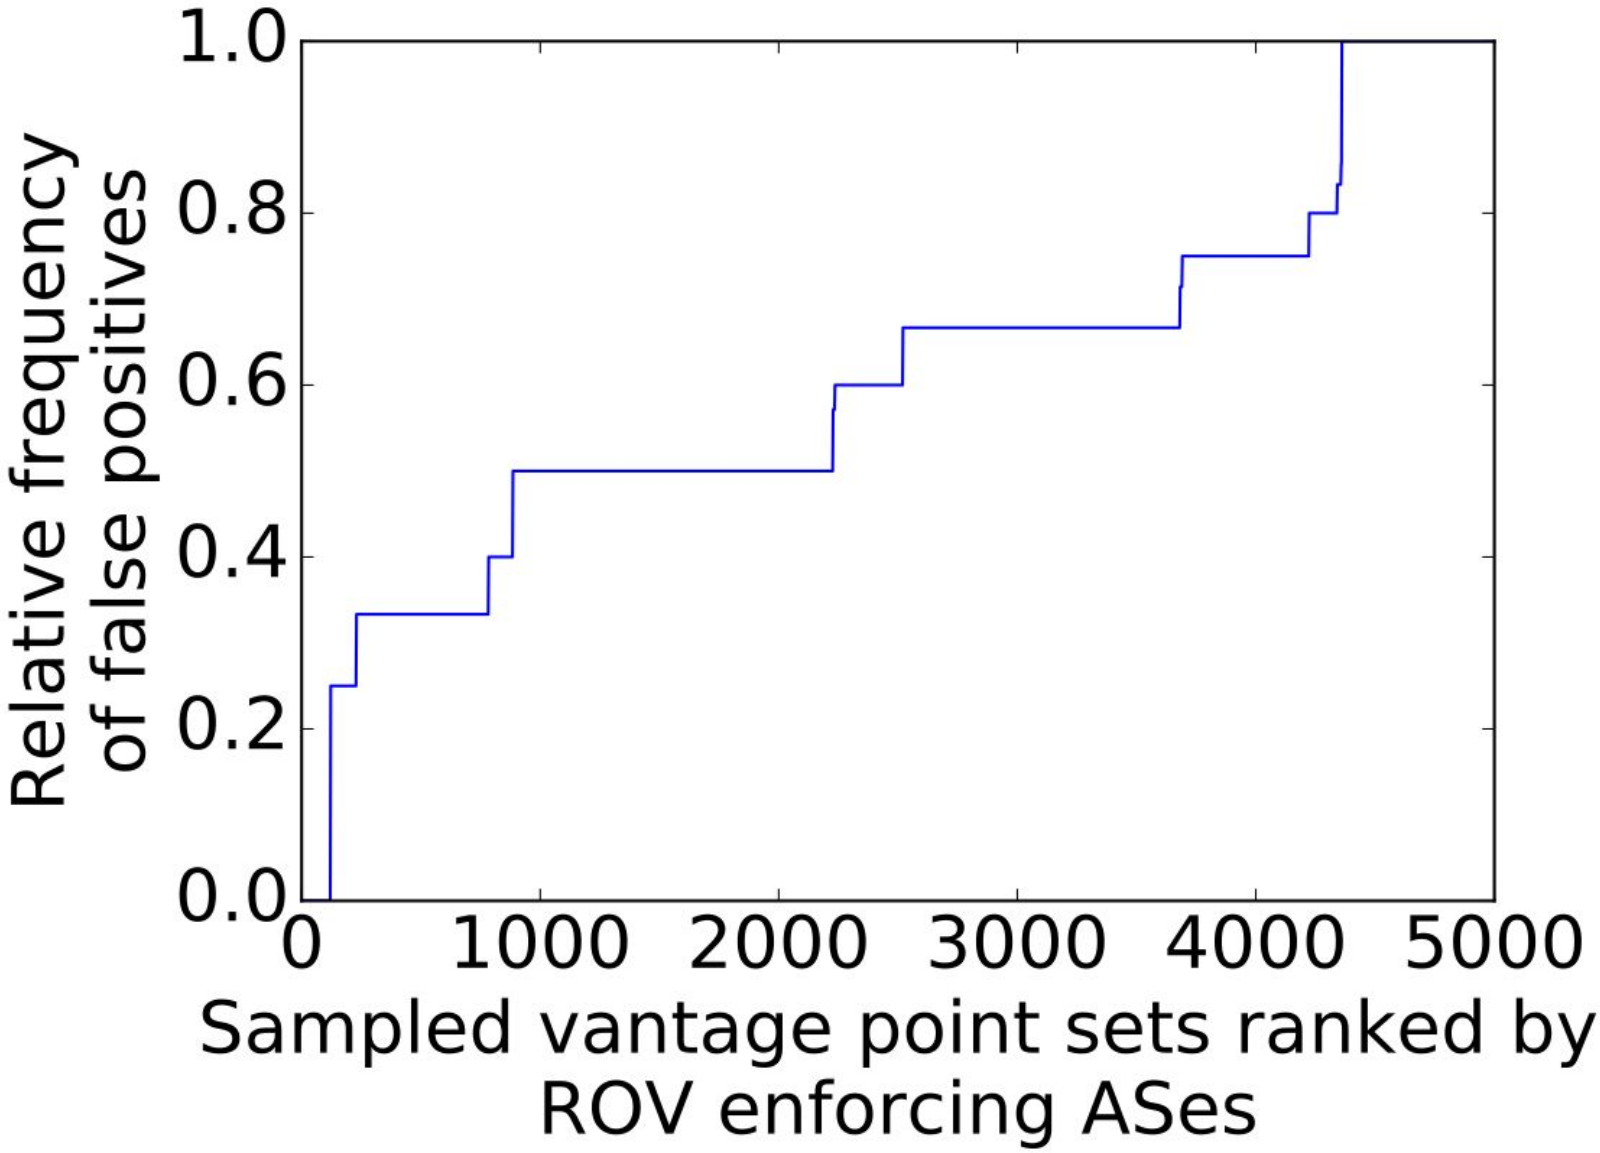
1.0
0.8
0.6
0.4
0.2
0.0
Relative frequency of false positives
0
1000
2000
3000
4000
5000
Sampled vantage point sets ranked by ROV enforcing ASes

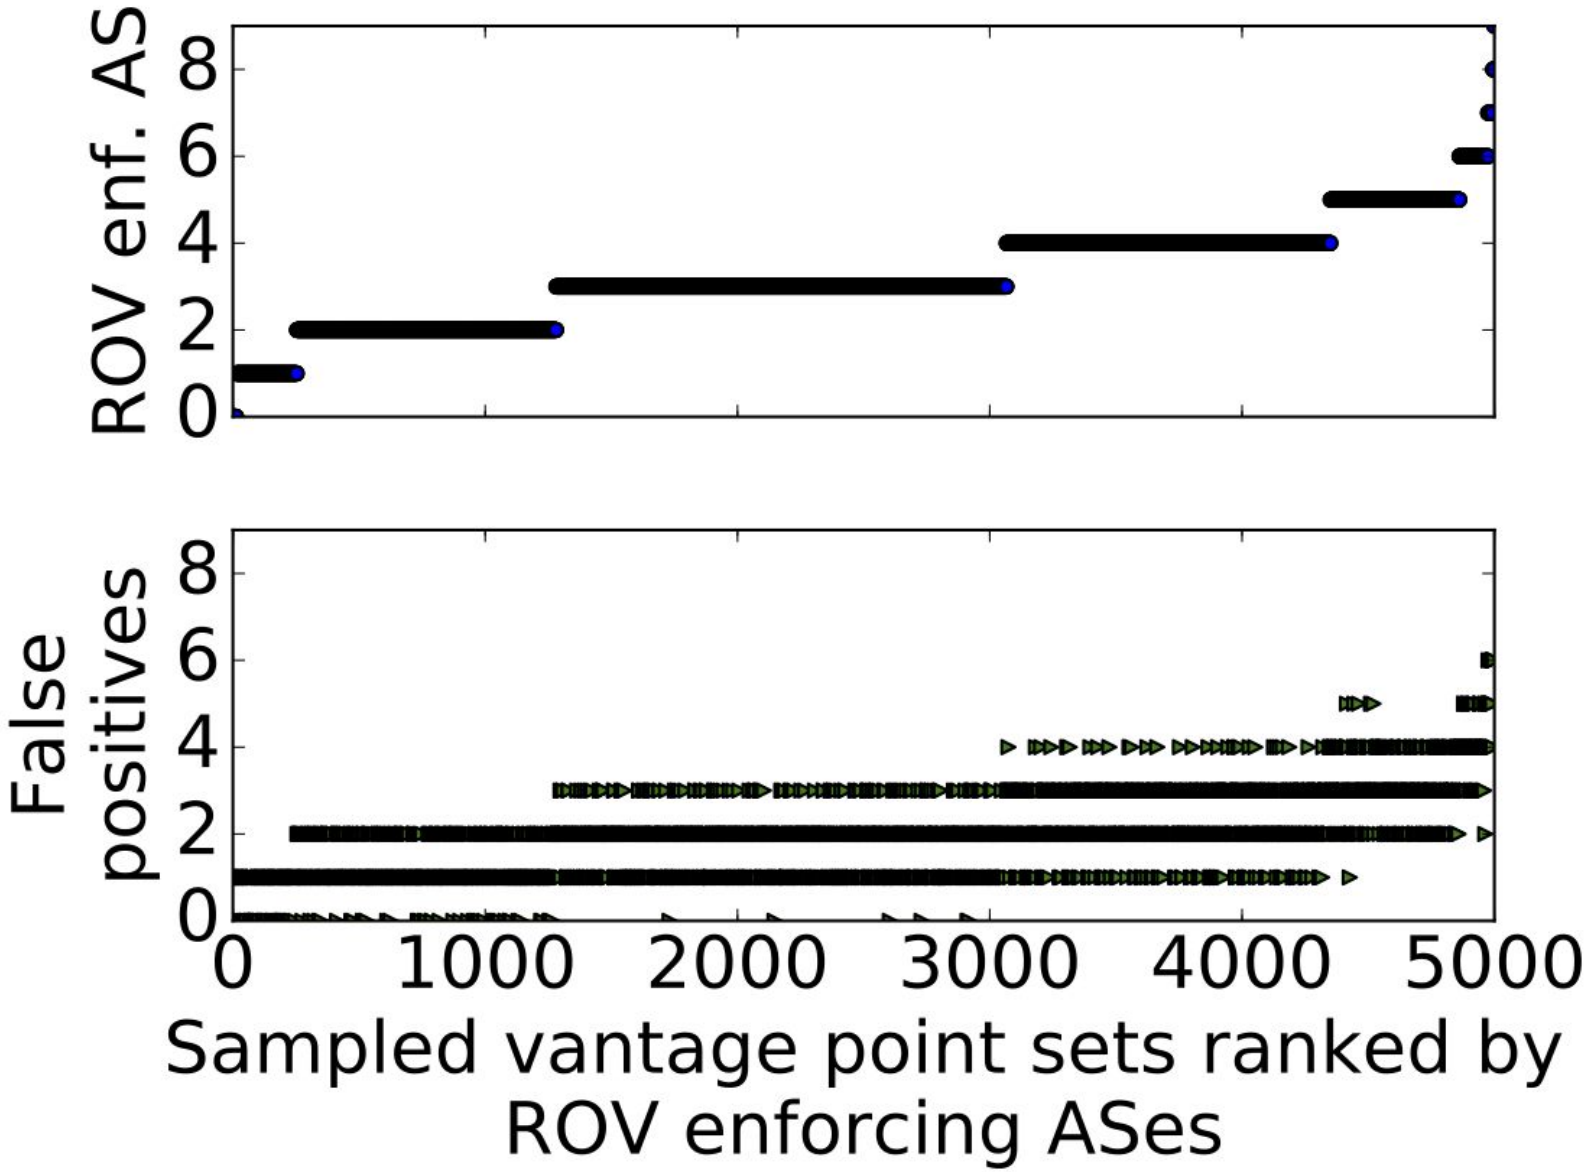
ROV enf. AS
0
2
4
6
8
False positives
0
2
4
6
8
0
1000
2000
3000
4000
5000
Sampled vantage point sets ranked by ROV enforcing ASes

Path Diversity

Path Diversity Distribution of a single vantage point

For ~50% of origins, there is exactly one distinct AS path

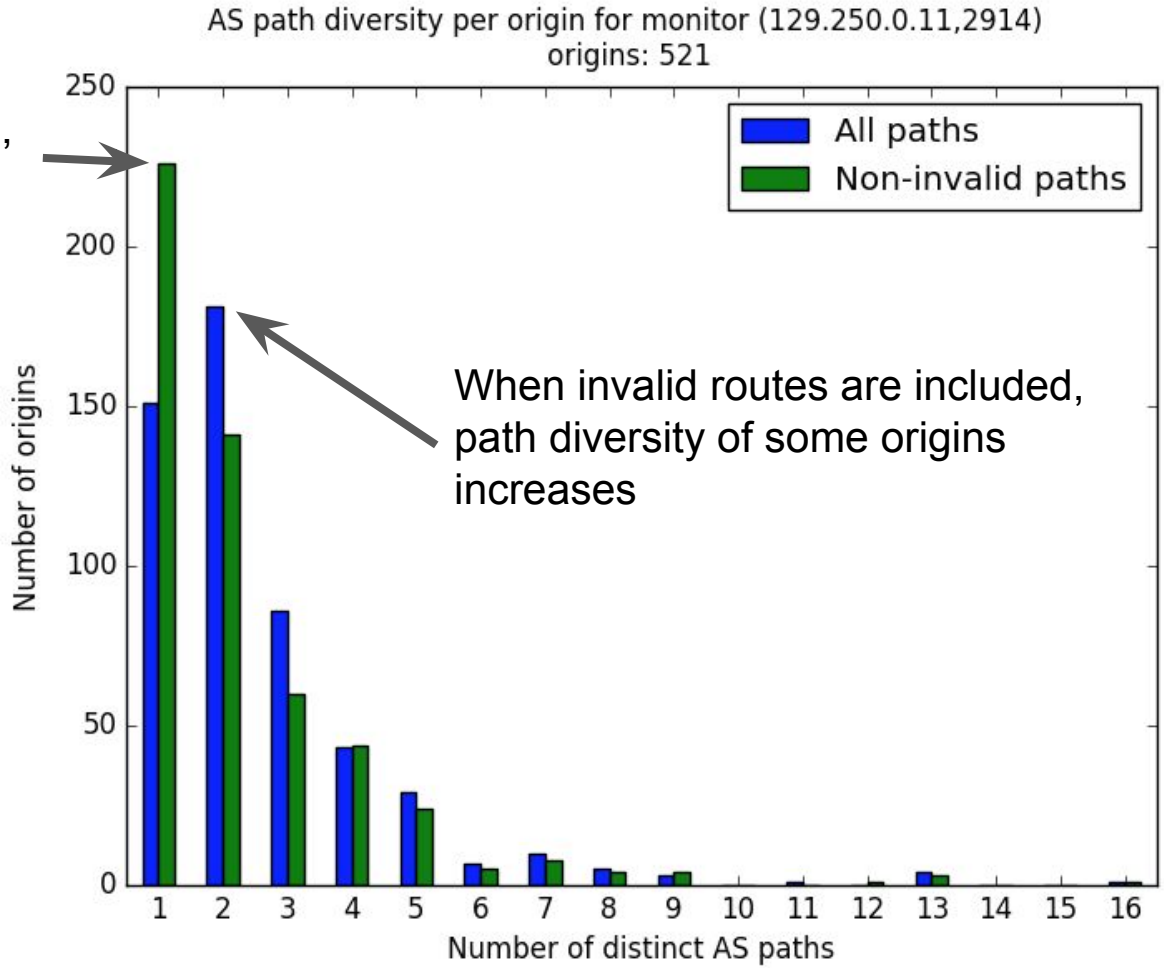

➔ Invalid routes tend to have different AS paths than non-invalid routes

Vantage Point Visibility Matters

Prefixes and their Origins

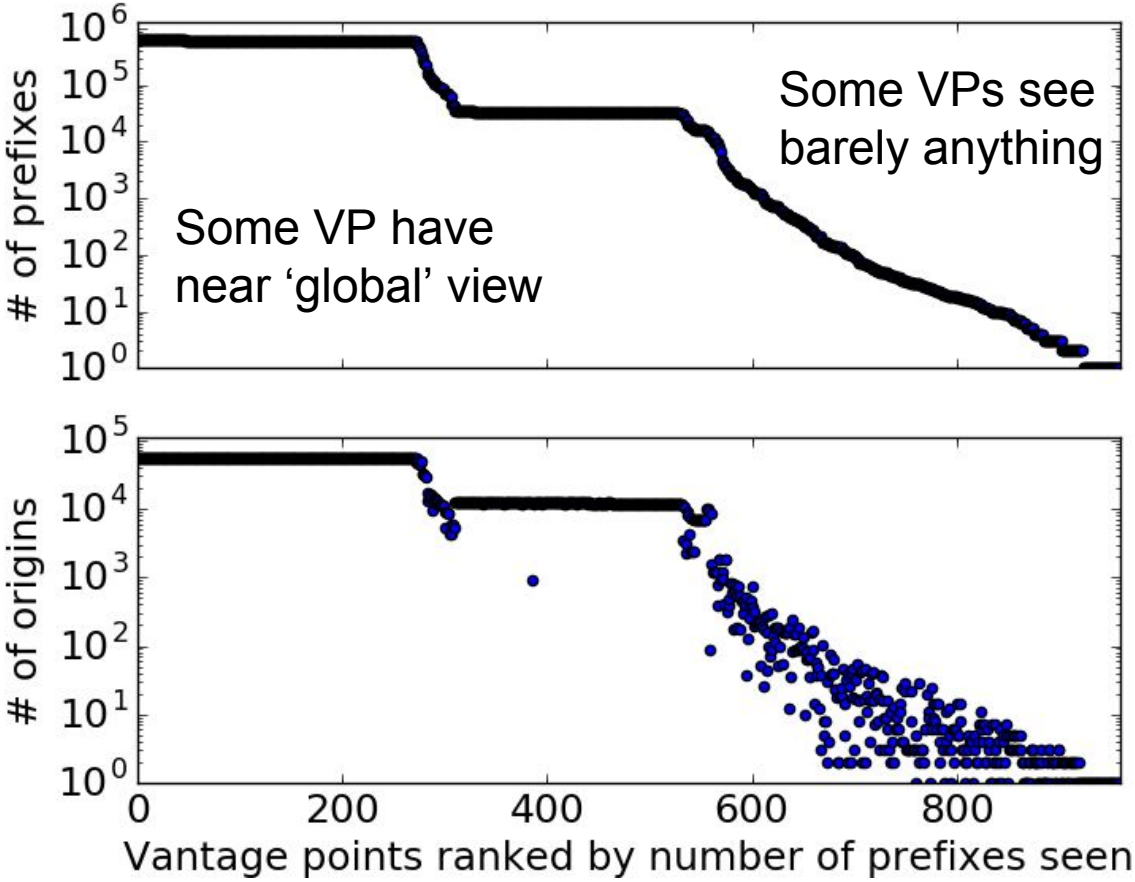

Vantage Point Visibility Matters

Prefixes of invalid routes and their reasons for invalidity

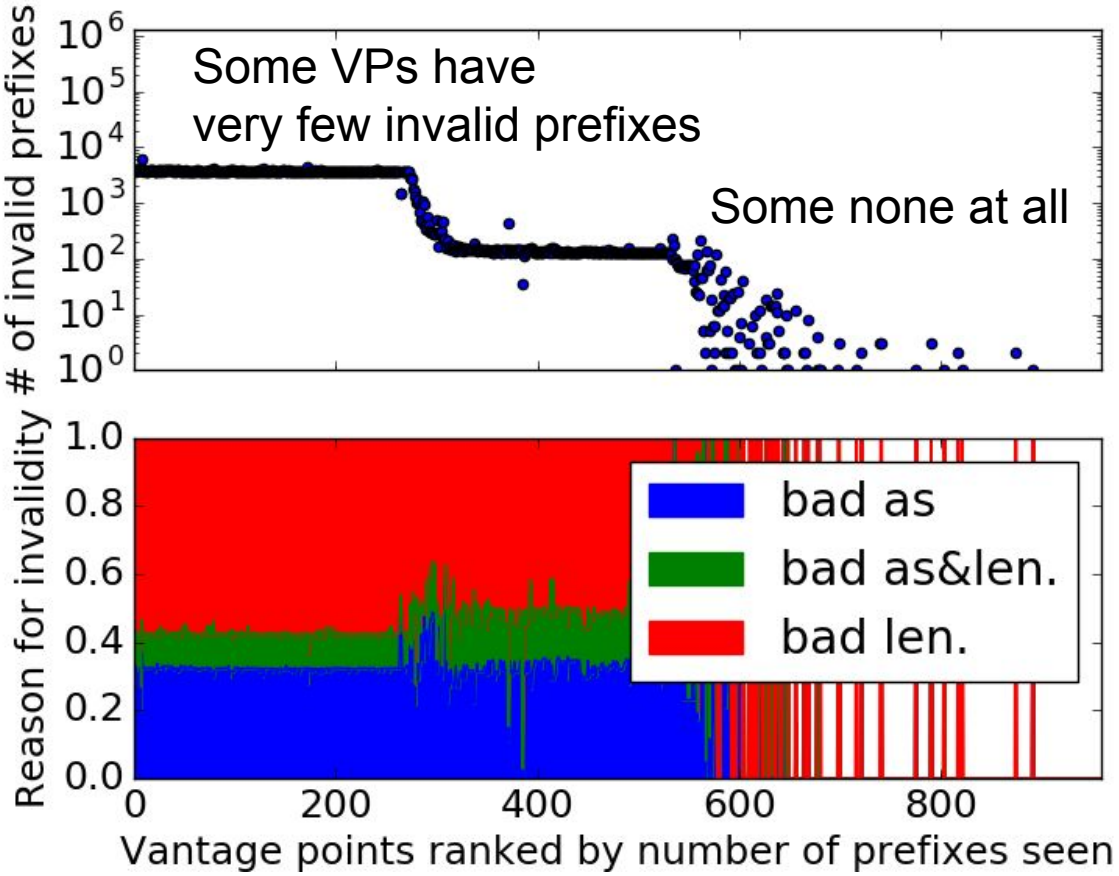

Vantage Point Visibility Matters

Per-Origin Prefix Visibility

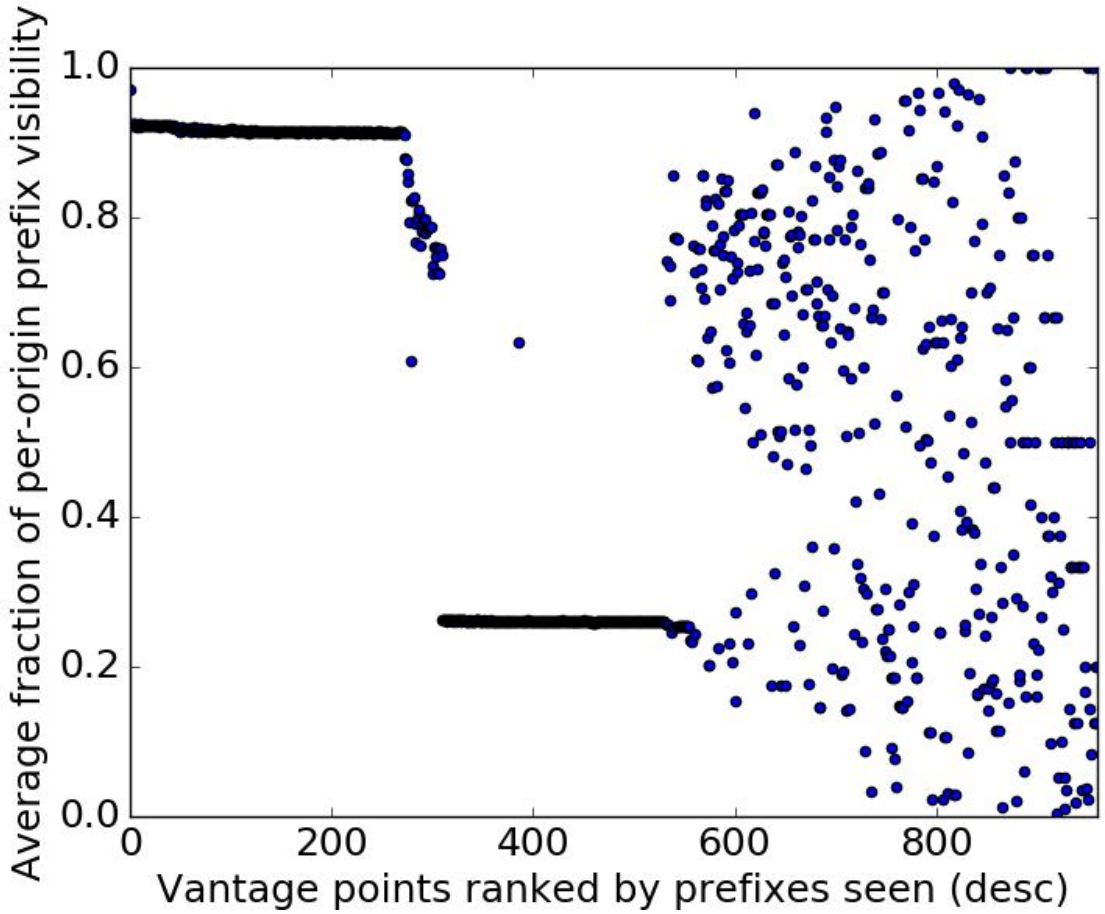

➔ Virtually all VPs have some origin AS they only ‘see’ incompletely. Oops!

Invalid Announcements: Path Diversity