

# Internet Performance: Who Cares and by How Much?

Ramesh K. Sitaraman

College of Information and Computer Science

UMass, Amherst

&


Chief Consulting Scientist

Akamai Technologies



Why

~~How~~ Performance?

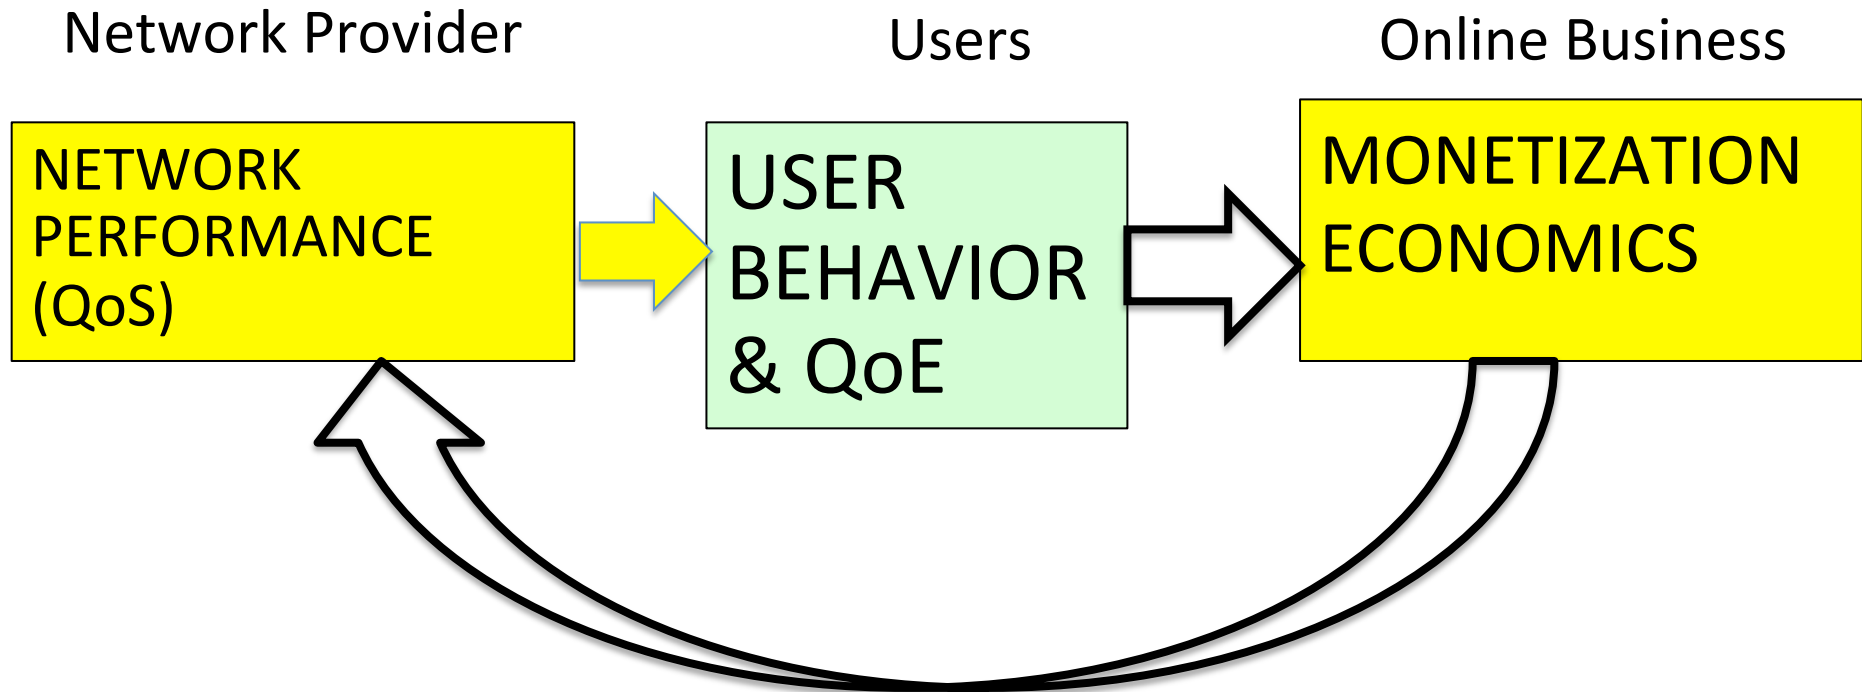


The Economic Value of Performance

# Economic Value of Web Performance?



# Holistic Network Design: A New Science in the Making



Every real-world network has a virtuous cycle!

# Impact of Network Performance on QoE & User Behavior

## Case Study: Rigorous Experimental Design


# Lets talk video performance...

If a video that you want to watch freezes in middle, would you watch less of the video?



Can we answer the question *causally* and *quantitatively*?

# Basic Tool: Correlate Performance and Behaviour



Hypothesis: Video rebuffering causes viewers to watch less

Strong negative rank correlation. Kendall correlation = -0.421.

# Correlation $\neq$ Causality

Correlation: A and B “move together”.

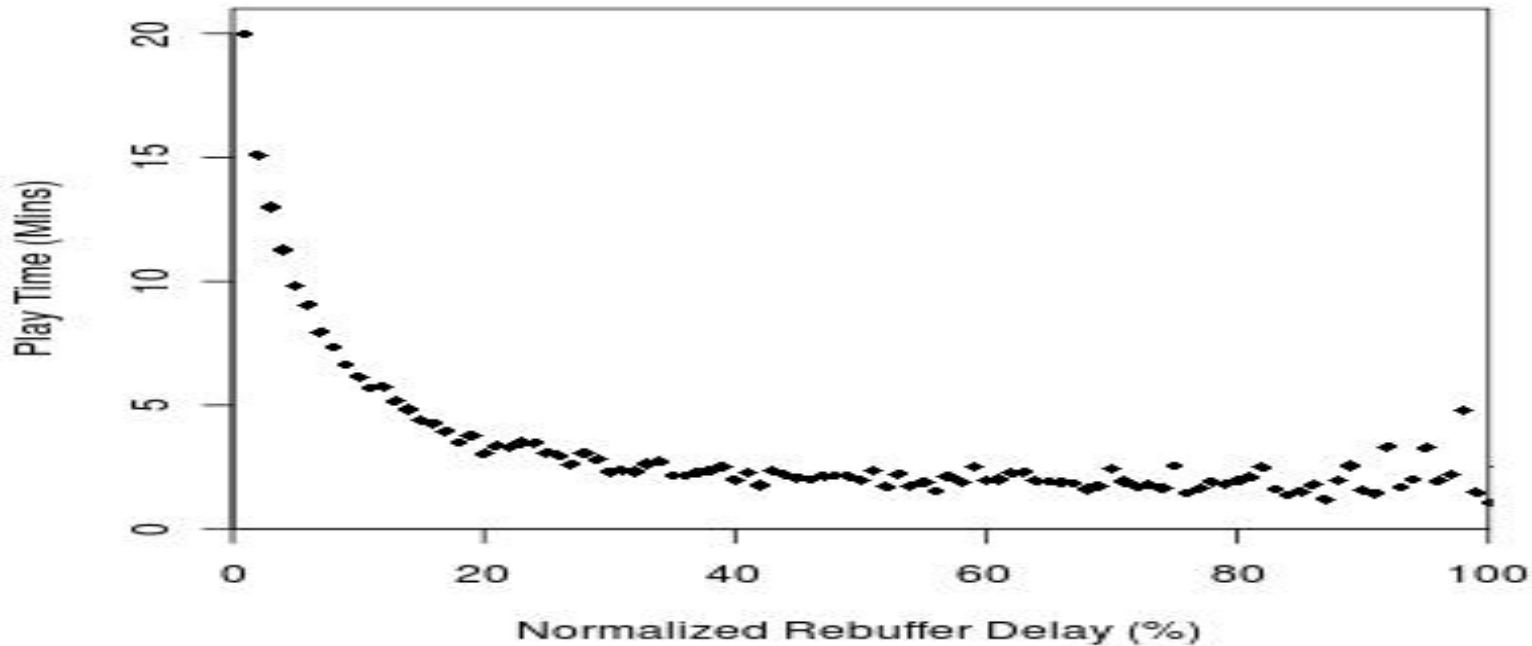
versus

Causality: A causes B to occur.

Threats to Causality: **Confounding variables** that could account for both A and B.




# Threats to Causality



Possible Threat: Users who are better off can afford better network **connectivity**, resulting in less rebuffering. They can also afford access to more interesting **content** causing them to watch longer.

Confounding variables: **Connectivity, Content, Geography.**

# Threats to Causality



Possible Threat: Users watching videos during office hours watch less, so as not to catch the eye of the boss 😊. Also rebuffers more likely due to peak hour Internet traffic jams.

Confounding variables: Time (day of week, time of day), geography.

# Challenges to a Causal Analysis

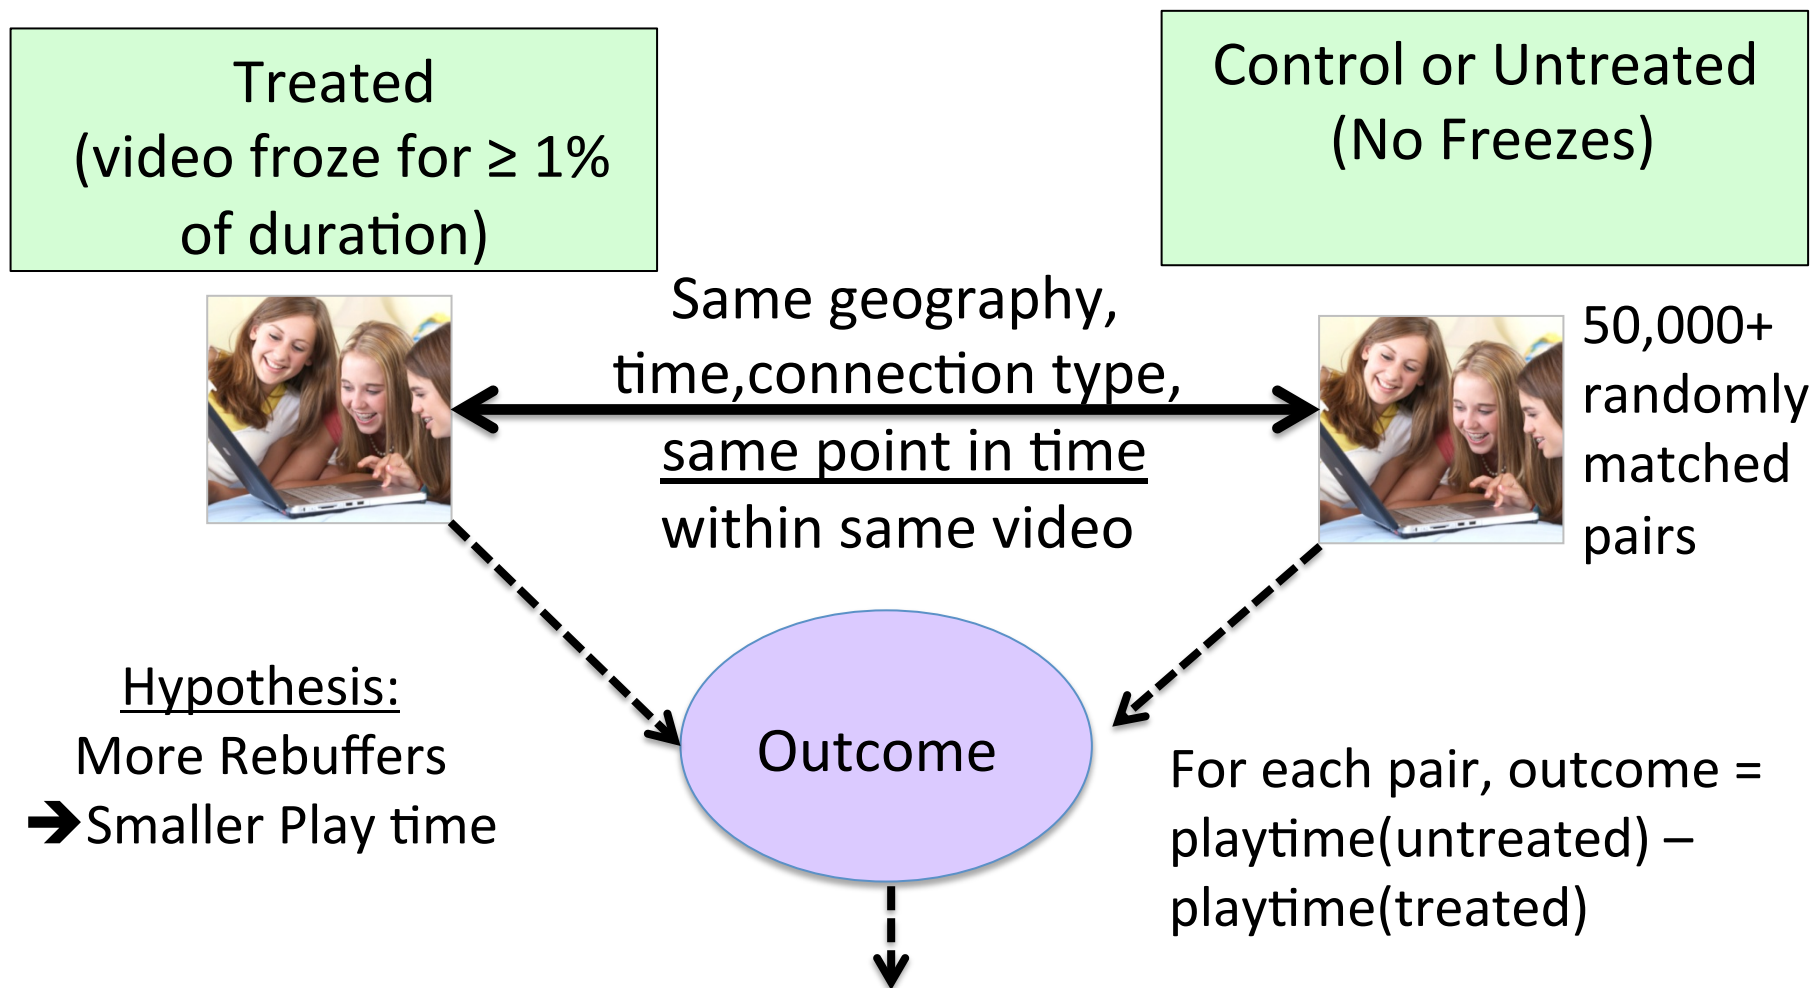
## Pro:

- Plenty of data!

## Cons:

- Randomized (controlled) experiments not possible.
- Diversity of devices, connectivity, geographies, and content to capture within a lab setting.

# Our approach: Quasi-Experimental Design

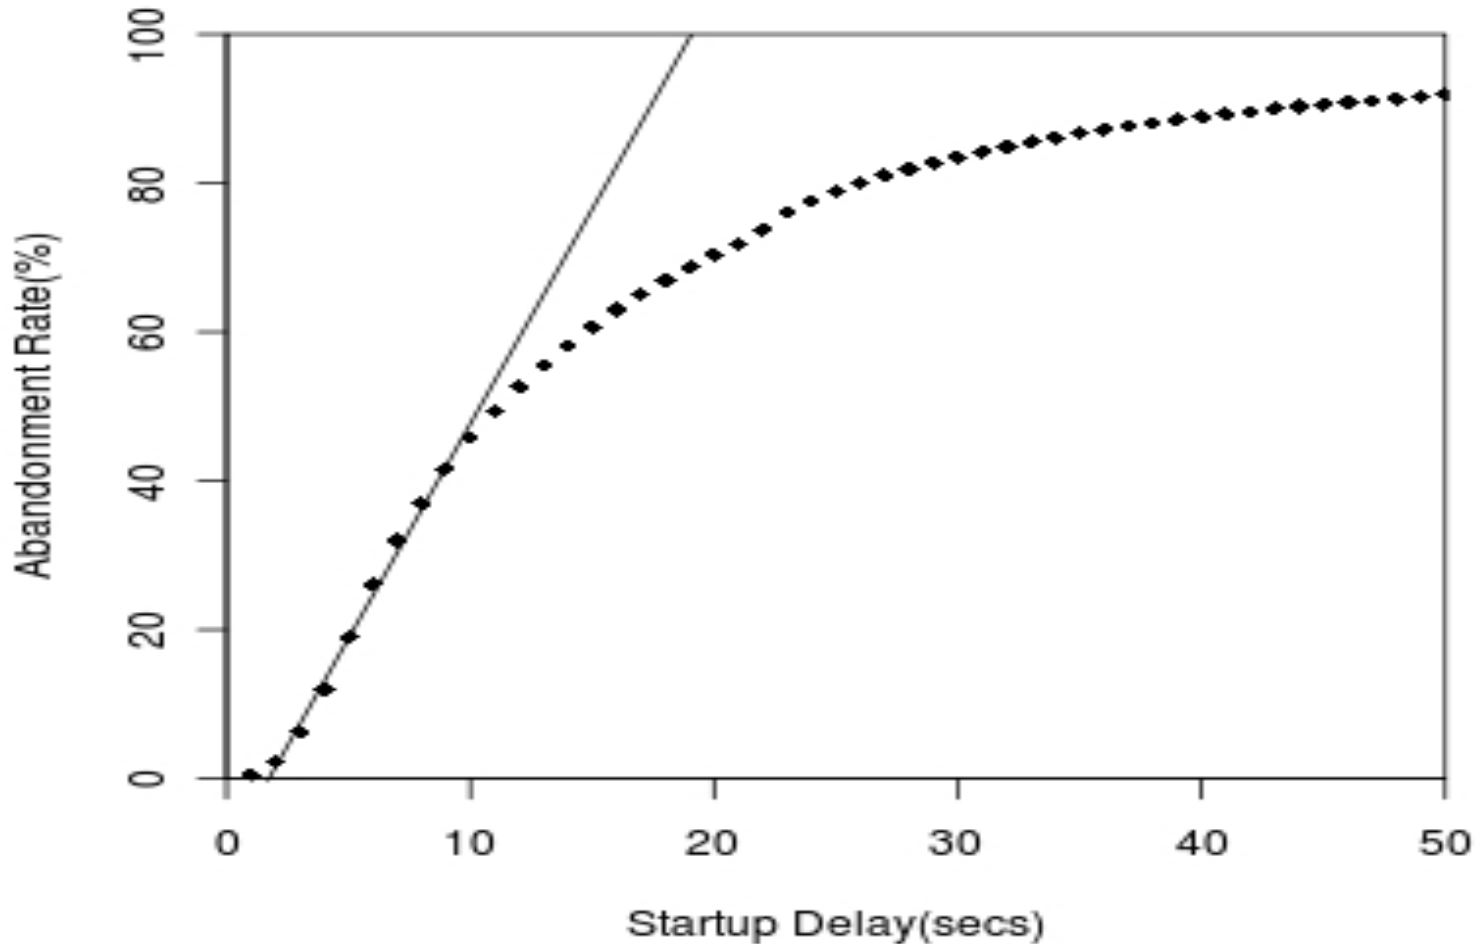


1% rebuffering causes 5% decreasing in play time

# Understanding Human Psychology in the Virtual World


## Case Study: Viewer Patience

# How long will viewers wait for a video to startup?



Abandonment rate = %viewers who loose patience and abandon before the video starts up

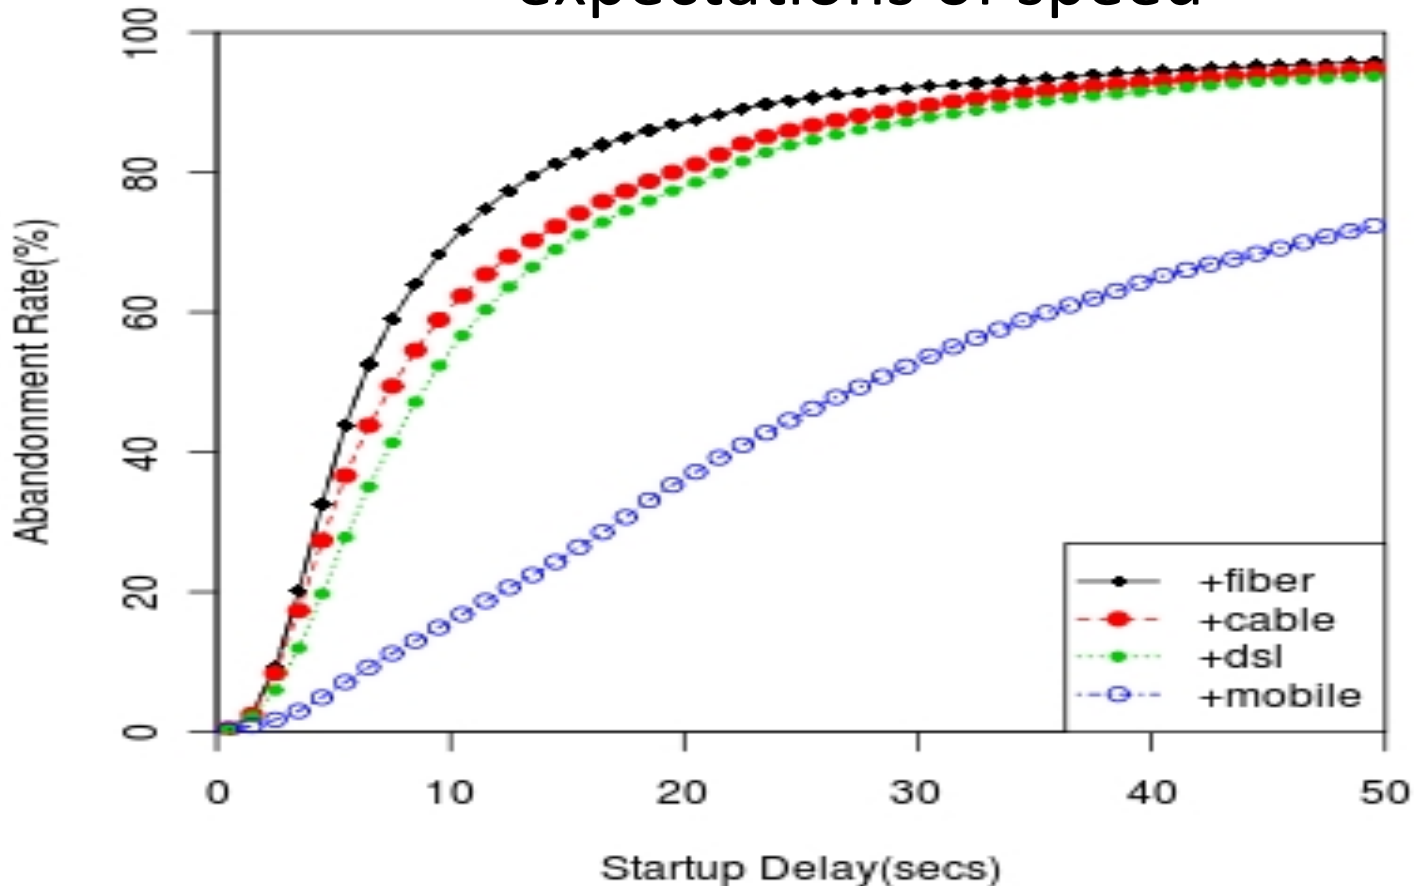
# How long will viewers wait for a video to startup?



## 2-Second Rule:

- a) Viewers start to abandon if startup delay exceeds 2 secs.
- b) At 5 seconds, ~ 25% abandoned. At 10 seconds, ~ 50% gone.


Time is relative: viewer patience is influenced by expectations of speed



- Viewers with better connectivity have less patience for startup delay and abandon sooner.
- Key implications for the Net Neutrality debate.

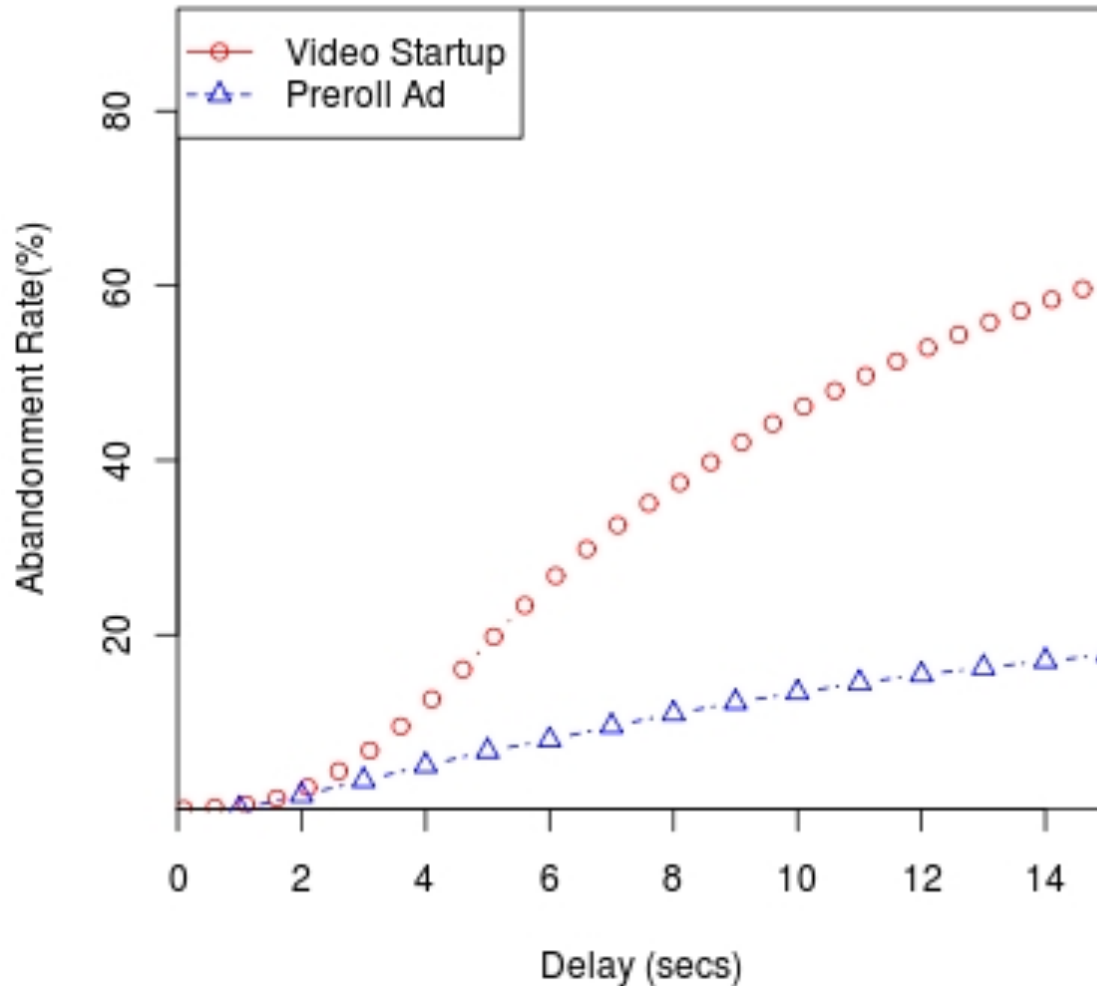


# Viewer patience is related to perceived value



Viewers are less tolerant of startup delay for short videos (e.g, news clips) in comparison to longer videos (e.g, movies)

# Viewers are more patient with ads than performance glitches



Abandonment rate 3X more for slow-loading video than pre-roll ad.

# SUMMARY

- **Virtuous Cycles:** Understand the *virtuous cycle* that links QoS, QoE, user behavior, and economics.
- **Rigorous Experimental Design:** Take a leaf from the social and medical sciences.
- Understand **online user psychology** for informed network design, policy, and governance.

# Questions?

*“Video Stream Quality Impacts Viewer Behavior: Inferring Causality using Quasi-Experimental Designs”.*

Krishnan and Sitaraman, ACM Internet Measurement Conference (IMC), 2012. (Also, IEEE/ACM TON, 2013)

*“Understanding the Effectiveness of Video Ads: A Measurement Study.”*

Krishnan and Sitaraman, ACM Internet Measurement Conference (IMC), 2013.