The basic economics of the internet

Shane Greenstein

Professor, Harvard Business School

June 22, 2021

Connectivity, Competition & Collaboration

The 3 Cs of the Digital Trade

Thank you for the opportunity to speak

This talk draws on material from the article, “The Basic Economics of Internet Infrastructure.” If you would like to know more about the history of the commercialization of the internet, then please read my book, “How the Internet Become Commercial.”.

Journal of Economic Perspectives—Volume 34, Number 2—Spring 2020—Pages 192–214

The Basic Economics of Internet Infrastructure

Shane Greenstein

This internet barely existed in a commercial sense 25 years ago. In the mid-1990s, when the data packets travelled to users over dial-up, the main internet traffic consisted of email, file transfer, and a few web applications. For such content, users typically could tolerate delays. Of course, the internet today is a vast and interconnected system of software applications and computing devices, which society uses to exchange information and services to support business, shopping, and leisure. Not only does data traffic for streaming, video, and gaming applications comprise the majority of traffic for internet service providers and reach users primarily through broadband lines, but typically those users would not tolerate delays in these applications (for usage statistics, see Nevo, Turner, and Williams 2016; McManus et al. 2018; Huston 2017). In recent years, the rise of smartphones and Wi-Fi access has supported growth of an enormous range of new businesses in the "sharing economy" (like, Uber, Lyft, and Airbnb), in mobile information services (like, social media, ticketing, and messaging), and in many other applications. More than 80 percent of US households own at least one smartphone, rising from virtually zero in 2007 (available at the Pew Research Center 2019 Mobile Fact Sheet). More than 86 percent of homes with access to broadband internet employ some form of Wi-Fi for accessing applications (Internet and Television Association 2018).

It seems likely that standard procedures for GDP accounting underestimate the output of the internet, including the output affiliated with "free" goods and the restructuring of economic activity wrought by changes in the composition of firms who use advertising (for discussion, see Nakamura, Samuels, and Soloveichik

■ *Shane Greenstein is the Martin Marshall Professor of Business Administration, Harvard Business School, Boston, Massachusetts. His email address is sgreenstein@hbs.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.34.2.192.

Outline

- Introduction
- Who pays for what?
- The frontier
- Summary

Journal of Economic Perspectives—Volume 34, Number 2—Spring 2020—Pages 192–214

The Basic Economics of Internet Infrastructure

Shane Greenstein

This internet barely existed in a commercial sense 25 years ago. In the mid-1990s, when the data packets travelled to users over dial-up, the main internet traffic consisted of email, file transfer, and a few web applications. For such content, users typically could tolerate delays. Of course, the internet today is a vast and interconnected system of software applications and computing devices, which society uses to exchange information and services to support business, shopping, and leisure. Not only does data traffic for streaming, video, and gaming applications comprise the majority of traffic for internet service providers and reach users primarily through broadband lines, but typically these users would not tolerate delays in these applications (for usage statistics, see Nevo, Turner, and Williams 2016; McManus et al. 2018; Huston 2017). In recent years, the rise of smartphones and Wi-Fi access has supported growth of an enormous range of new businesses in the "sharing economy" (like, Uber, Lyft, and Airbnb), in mobile information services (like, social media, ticketing, and messaging), and in many other applications. More than 80 percent of US households own at least one smartphone, rising from virtually zero in 2007 (available at the Pew Research Center 2019 Mobile Fact Sheet). More than 86 percent of homes with access to broadband internet employ some form of Wi-Fi for accessing applications (Internet and Television Association 2018).

It seems likely that standard procedures for GDP accounting underestimate the output of the internet, including the output affiliated with "free" goods and the restructuring of economic activity wrought by changes in the composition of firms who use advertising (for discussion, see Nakamura, Samuels, and Soloveichik

■ *Shane Greenstein is the Martin Marshall Professor of Business Administration, Harvard Business School, Boston, Massachusetts. His email address is sgreenstein@hbs.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.34.2.192.

Outline

- **Introduction**
- Who pays for what?
- The frontier
- Summary

Journal of Economic Perspectives—Volume 34, Number 2—Spring 2020—Pages 192–214

The Basic Economics of Internet Infrastructure

Shane Greenstein

This internet barely existed in a commercial sense 25 years ago. In the mid-1990s, when the data packets travelled to users over dial-up, the main internet traffic consisted of email, file transfer, and a few web applications. For such content, users typically could tolerate delays. Of course, the internet today is a vast and interconnected system of software applications and computing devices, which society uses to exchange information and services to support business, shopping, and leisure. Not only does data traffic for streaming, video, and gaming applications comprise the majority of traffic for internet service providers and reach users primarily through broadband lines, but typically these users would not tolerate delays in these applications (for usage statistics, see Nevo, Turner, and Williams 2016; McManus et al. 2018; Huston 2017). In recent years, the rise of smartphones and Wi-Fi access has supported growth of an enormous range of new businesses in the "sharing economy" (like, Uber, Lyft, and Airbnb), in mobile information services (like, social media, ticketing, and messaging), and in many other applications. More than 80 percent of US households own at least one smartphone, rising from virtually zero in 2007 (available at the Pew Research Center 2019 Mobile Fact Sheet). More than 86 percent of homes with access to broadband internet employ some form of Wi-Fi for accessing applications (Internet and Television Association 2018).

It seems likely that standard procedures for GDP accounting underestimate the output of the internet, including the output affiliated with "free" goods and the restructuring of economic activity wrought by changes in the composition of firms who use advertising (for discussion, see Nakamura, Samuels, and Soloveichik

■ *Shane Greenstein is the Martin Marshall Professor of Business Administration, Harvard Business School, Boston, Massachusetts. His email address is sgreenstein@hbs.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.34.2.192.

Questions

Process. How does the internet work? What are its main economic elements?

Supporting frontier. How has the network improved in the last two decades? What are the economic incentives to improve the frontier?

Orientation. Analysis to inform policy, not to advocate any particular policy. Plain words.

Motivation: why it is worthwhile to understand how the internet works and improves.

Deployment of internet restructured private/public boundary for communications in the last three decades.

Invisible. Hard to appreciate how the economics of the modern internet works. The mechanisms are not easily visible or measurable.

Big and important. Part 1. $105.9 billion in revenue (advertising) among "Internet Publishing and Broadcasting and Web Search Portals (NAICS 519130)." Grew 250% since 2012.

Big and important. Part 2. $545 B in 2017 for "Electronic Shopping & Mail Order (NAICS 4541)." Grew 65% from 2012.

The answer: Who pays who for what?

- Most observers already understand two types of economic transactions, subscriptions and ad revenue.
- *Subscription:* **users** pay **content providers** for content.
- *Ad revenue:* **advertisers** pay **content providers** for access to users.

- This presentation will focus on the less visible economic transactions that support the system. Will address the questions:
  - *Who pays for access and what do they receive for those payments?*
  - *Who pays to reduce latency experienced by users?*
  - *How does paying for collocation reduce latency?*
  - *Who pays for transit, and how does that result in moving data?*

Outline

- Introduction
- **Who pays for what?**
- The frontier
- Summary

Journal of Economic Perspectives—Volume 34, Number 2—Spring 2020—Pages 192–214

The Basic Economics of Internet Infrastructure

Shane Greenstein

This internet barely existed in a commercial sense 25 years ago. In the mid-1990s, when the data packets travelled to users over dial-up, the main internet traffic consisted of email, file transfer, and a few web applications. For such content, users typically could tolerate delays. Of course, the internet today is a vast and interconnected system of software applications and computing devices, which society uses to exchange information and services to support business, shopping, and leisure. Not only does data traffic for streaming, video, and gaming applications comprise the majority of traffic for internet service providers and reach users primarily through broadband lines, but typically these users would not tolerate delays in these applications (for usage statistics, see Nevo, Turner, and Williams 2016; McManus et al. 2018; Huston 2017). In recent years, the rise of smartphones and Wi-Fi access has supported growth of an enormous range of new businesses in the "sharing economy" (like, Uber, Lyft, and Airbnb), in mobile information services (like, social media, ticketing, and messaging), and in many other applications. More than 80 percent of US households own at least one smartphone, rising from virtually zero in 2007 (available at the Pew Research Center 2019 Mobile Fact Sheet). More than 86 percent of homes with access to broadband internet employ some form of Wi-Fi for accessing applications (Internet and Television Association 2018).

It seems likely that standard procedures for GDP accounting underestimate the output of the internet, including the output affiliated with "free" goods and the restructuring of economic activity wrought by changes in the composition of firms who use advertising (for discussion, see Nakamura, Samuels, and Soloveichik

■ *Shane Greenstein is the Martin Marshall Professor of Business Administration, Harvard Business School, Boston, Massachusetts. His email address is sgreenstein@hbs.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.34.2.192.

A brief history of origins

- **DARPA & NSF funding**
  - DARPA. Packet switching. Prototypes.
  - NSF funded operations 1985 – 1990s. Funded pragmatic operations → scaling.
- **World Wide Web**
  - Hypertext created by Tim Berners-Lee at CERN & first deployed in 1991
  - Consortium established at MIT in 1994.
- **Internet privatized in 1995.**
  - Commercial browser introduced in 2/95.
  - Want to know more? Please see book.

- **Architecture and use both evolved during and after the 1990s. But changes not widely appreciated.**
  - **Scale** of access and participation are much larger today. Frontier of the network has evolved to encourage and accommodate that scale.
  - **Composition** of applications much different. Email & web browsing → interactive, streaming, & video.
  - Traffic increasingly "**asymmetric**" – i.e., households receive more data than send. Internet changed with accordingly.

The mechanics – e.g., user looks up Wikipedia

- If the user & Wikipedia **use the same ISP**, then straightforward.
  - Web browser calls Wikipedia.
  - Browser can be on any device.
  - ISP sends request to Wikipedia & Wikipedia's servers receive it.
  - Wikipedia's servers send data to ISP, who sends data to user.
- When does traffic **use same ISP**?
  - Email tends to be with another local person → stays on same ISP.
  - Not much else.

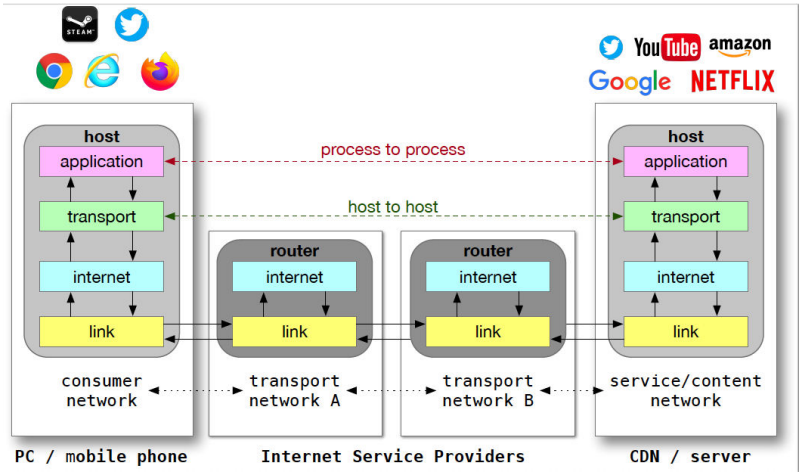

Most traffic **does NOT use same ISP**.

What if the user and Wikipedia do not use the same ISP?

Data travel b/w ISPs. How? Part 1 (of 3)

- **ISP moves data b/w Content Delivery Network (CDN) & user.**
  - How? Wikipedia arranged with CDN to make some content available.
  - User requested something on CDN.
- **Minimize distance:** reduce latency.
- **Most common approaches.**
  - Content provider hires third-party CDNs – e.g., Akamai, Cloudflare, or limelight.
  - **Large firms operate their own CDNs** – e.g., Alphabet, Apple, Facebook, Amazon, Microsoft, Netflix

An example of a CDN network.

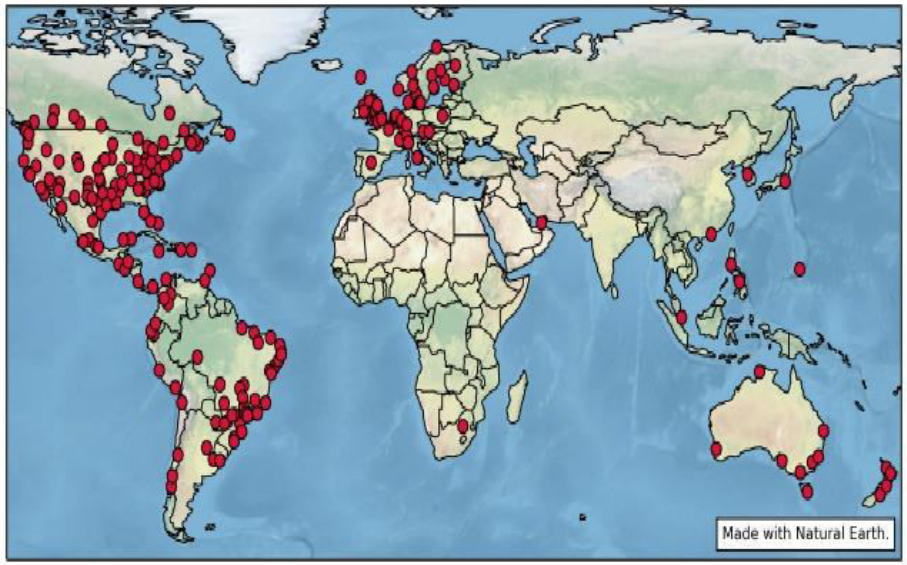

A map of Netflix's CDNs – comprised of 4600+ servers in 223 locations, according to Bottger et al (2018).

Data travel b/w ISPs. How? Part 2 (of 3)

- **Server to user.** Data travels directly from servers at Wikipedia's HQ to user's browser over multiple ISPs.
  - How things worked before CDNs were invented. Still an option.
- **Why uncommon?** It is slow.
  - **Why use this approach?** User requested less popular content stored on servers at HQ and content provider puts only popular content on CDN to save on expenses.
  - **Why else?** Content provider is tiny & cannot afford to pay CDN.

- **Data takes the least congested pathway** on the network.
  - When routing is working the way it is supposed to work.

Lines that connect the internet, with their interconnection points in red.

Data travels between ISPs. Part 3.

Contracting

- **Private peering with symmetry**: ISPs exchange data of equal volumes → no money changes hands.
- **Private peering with asymmetry**: The ISP who sends more data pays the ISP who receives it.
- **Transit:** If ISP A gives data to ISP B who gives to ISP C, then ISP B's network gets paid for "transit."
- **Increasing asymmetry:** content providers send more data in one direction (to households) than receive.

Physical contact

- **Bilateral.** Two networks build a line and meet. *(Modern practice).*
- **Bilateral at internet exchange point (IXP)** that acts as "carrier hotel." Pay a fee. ISPs also make a contract to peer. *(Modern practice)*
- **Incentive to build bigger pipes & interconnection:** It is better to receive data than give it, so it pays to be bigger and have more points of interconnection → receive more data & provide more transit.

Summary: Who pays who for what?

- *Access:* **User** and **content provider** pays the **ISP** for access to the internet.
- *Hosting to reduce latency:* **Content provider** pays **CDN** to host content closer to user to reduce latency.
- *Collocation:* some **CDNs** pay **ISPs** to collocate in the ISP & some **CDNs** pay **IXPs** to collocate in the IXP.
- *Transit:* **ISPs** pay other **ISPs** to carry data between ISPs.
- *Subscription:* **users** pay **content providers** for content.
- *Ad revenue:* **advertisers** pay **content providers** for access to users.

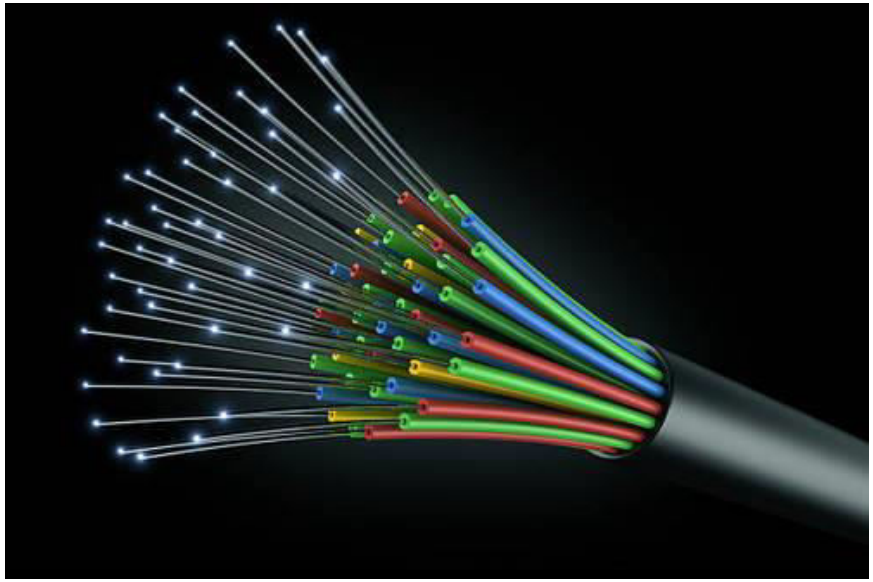

Outline

- Introduction
- Who pays for what?
- **The frontier**
- Summary

Journal of Economic Perspectives—Volume 34, Number 2—Spring 2020—Pages 192–214

The Basic Economics of Internet Infrastructure

Shane Greenstein

This internet barely existed in a commercial sense 25 years ago. In the mid-1990s, when the data packets travelled to users over dial-up, the main internet traffic consisted of email, file transfer, and a few web applications. For such content, users typically could tolerate delays. Of course, the internet today is a vast and interconnected system of software applications and computing devices, which society uses to exchange information and services to support business, shopping, and leisure. Not only does data traffic for streaming, video, and gaming applications comprise the majority of traffic for internet service providers and reach users primarily through broadband lines, but typically these users would not tolerate delays in these applications (for usage statistics, see Nevo, Turner, and Williams 2016; McManus et al. 2018; Huston 2017). In recent years, the rise of smartphones and Wi-Fi access has supported growth of an enormous range of new businesses in the "sharing economy" (like, Uber, Lyft, and Airbnb), in mobile information services (like, social media, ticketing, and messaging), and in many other applications. More than 80 percent of US households own at least one smartphone, rising from virtually zero in 2007 (available at the Pew Research Center 2019 Mobile Fact Sheet). More than 86 percent of homes with access to broadband internet employ some form of Wi-Fi for accessing applications (Internet and Television Association 2018).

It seems likely that standard procedures for GDP accounting underestimate the output of the internet, including the output affiliated with "free" goods and the restructuring of economic activity wrought by changes in the composition of firms who use advertising (for discussion, see Nakamura, Samuels, and Soloveichik

■ *Shane Greenstein is the Martin Marshall Professor of Business Administration, Harvard Business School, Boston, Massachusetts. His email address is sgreenstein@hbs.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.34.2.192.

New frontier: household wirelines and wireless broadband follows classic s-curve

Well documented: Broadband (cable/DSL) replaced dial-up.

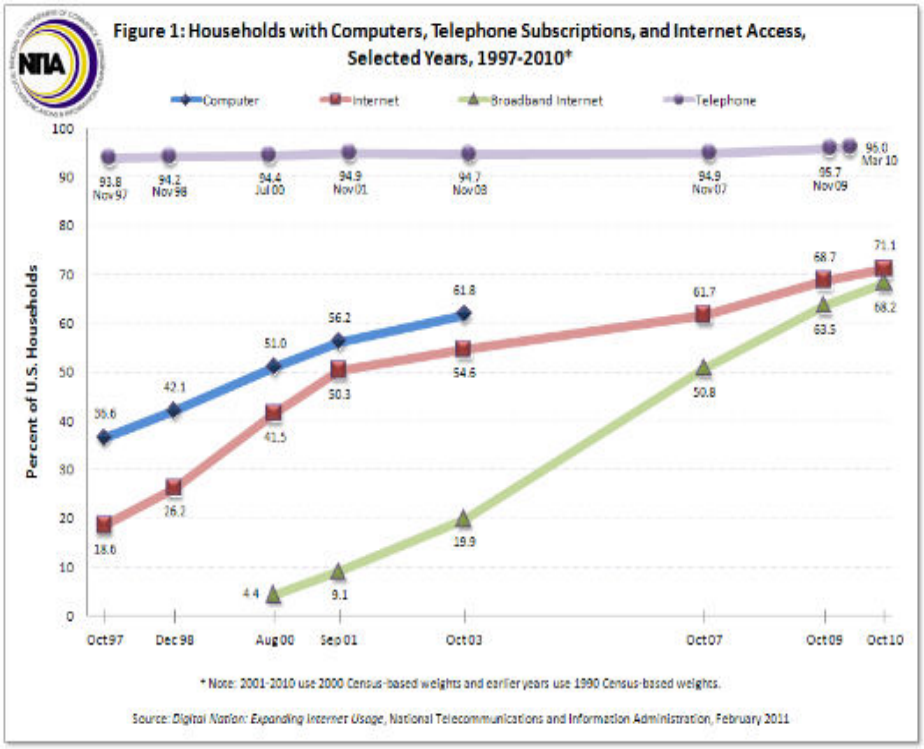

- **Dial-up** encouraged by universal available telephone system.
- **Broadband** comes in three forms: DSL, cable, fiber-to-home. Largely replaced dial-up.
- Today approaching 80% **adoption** by US households.
- Wireless in **Wifi** at homes & business. Near ubiquity.
- Wireless in **3G or 4G** with smart phone. **5G buildout** will generate a new generation of diffusion.

New frontier: Data Centers & Cloud

- **Data Centers.** Support computing & storage. Locate near user & good connectivity.
- **Cost determinants:** electricity, land, cooling requirements, government tax abatements.
- **Recent.** NYSE electronic exchanges moved to a data center in New Jersey in 2010.
- Google's first data center in 2003 (GA), then 2006 (OR), & 2007 (SC, IA, NC), & more….
- **Cloud is even more recent.** AWS begins 2006, $8B by 2015, $35B by 2019.
- **Frontier:** Hyperscale to support cloud services, & distributed architectures to reduce latency to users.

Ashburn, VA contains the largest collection of cloud & data centers in US.

Map of the US data centers (pre-pandemic)

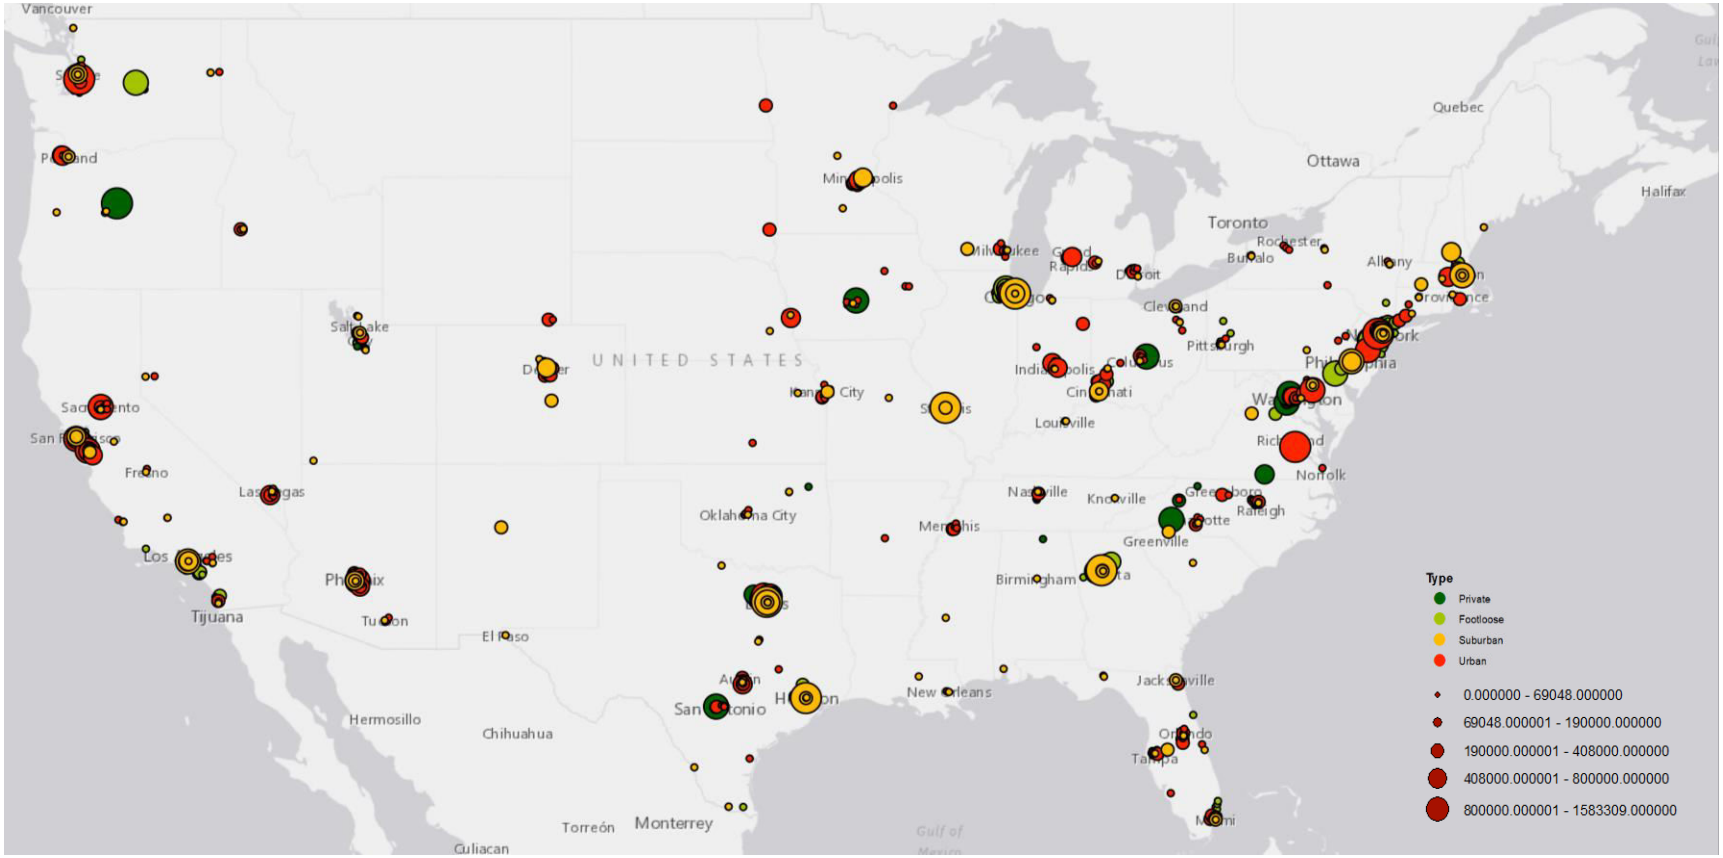

Third party data centers locate in the biggest cities near their biggest users

Optimistic v. pessimistic outlook about frontier

Optimists

- The frontier must show up somewhere first.
- Early use goes where lead users locate – i.e., the tech-hub cities.
- Eventually things become cheaper, more reliable & widely available.
- Policy should be patient, and at most, should try to speed up diffusion processes.

Pessimists

- Frontier stays in same places & spreads slowly.
- Privileged access for coastal clusters, rich urbanites, elites.
- Inequality of access/use breeds resentment & hinders delivery of public services.
- Policy should intervene, and at minimum, overcome drawbacks of uneven distribution.

Outline

- Introduction
- Who pays for what?
- The frontier
- **Summary**

Journal of Economic Perspectives—Volume 34, Number 2—Spring 2020—Pages 192–214

The Basic Economics of Internet Infrastructure

Shane Greenstein

This internet barely existed in a commercial sense 25 years ago. In the mid-1990s, when the data packets travelled to users over dial-up, the main internet traffic consisted of email, file transfer, and a few web applications. For such content, users typically could tolerate delays. Of course, the internet today is a vast and interconnected system of software applications and computing devices, which society uses to exchange information and services to support business, shopping, and leisure. Not only does data traffic for streaming, video, and gaming applications comprise the majority of traffic for internet service providers and reach users primarily through broadband lines, but typically these users would not tolerate delays in these applications (for usage statistics, see Nevo, Turner, and Williams 2016; McManus et al. 2018; Huston 2017). In recent years, the rise of smartphones and Wi-Fi access has supported growth of an enormous range of new businesses in the "sharing economy" (like, Uber, Lyft, and Airbnb), in mobile information services (like, social media, ticketing, and messaging), and in many other applications. More than 80 percent of US households own at least one smartphone, rising from virtually zero in 2007 (available at the Pew Research Center 2019 Mobile Fact Sheet). More than 86 percent of homes with access to broadband internet employ some form of Wi-Fi for accessing applications (Internet and Television Association 2018).

It seems likely that standard procedures for GDP accounting underestimate the output of the internet, including the output affiliated with "free" goods and the restructuring of economic activity wrought by changes in the composition of firms who use advertising (for discussion, see Nakamura, Samuels, and Soloveichik

■ *Shane Greenstein is the Martin Marshall Professor of Business Administration, Harvard Business School, Boston, Massachusetts. His email address is sgreenstein@hbs.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.34.2.192.

Summary

- **Architecture and use both evolved during last two decades.**
  - **Scale** of access and participation are much larger today. **Frontier** of the network has evolved.
  - **Composition** of applications much different. Email & web browsing → interactive, streaming, & video.
  - Traffic increasing "**asymmetric**" – i.e., households receive more data than send.
- **Open research questions:**
  - **Demand induced innovation & externalities.** Competitive incentives to improve elements with dispersed benefits?
  - **Value of substitutes to user?** Willingness to substitute across options? Willingness to pay for moving targets.
  - **Building & accumulation.** Local tax abatements for building → what activities follow, if any? Do business and household gains accrue?

Thanks for your attention

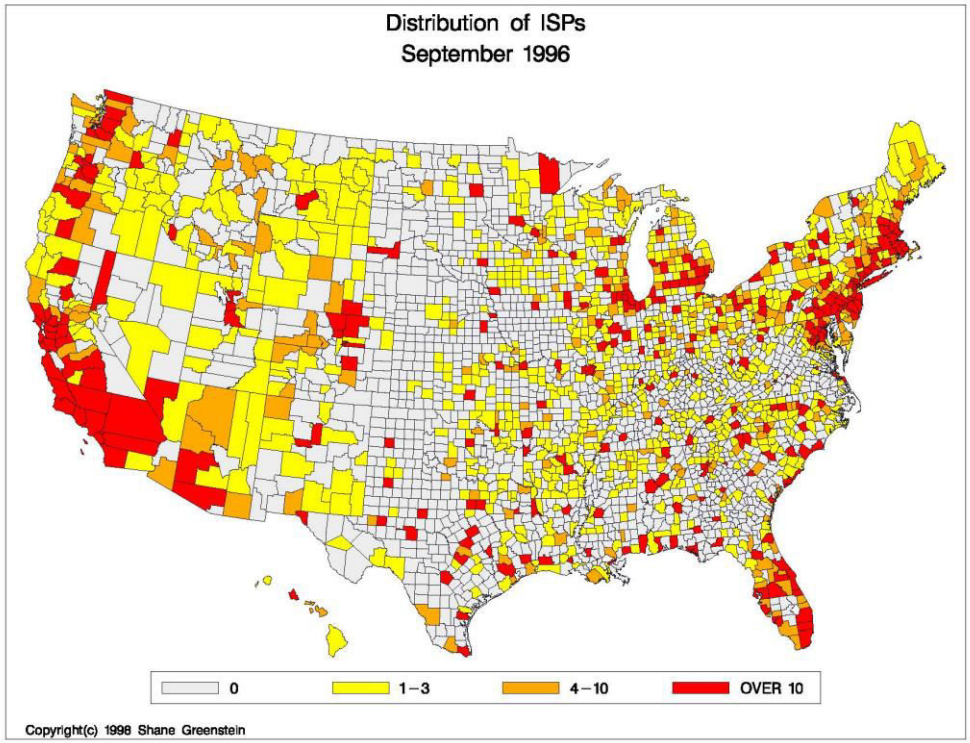

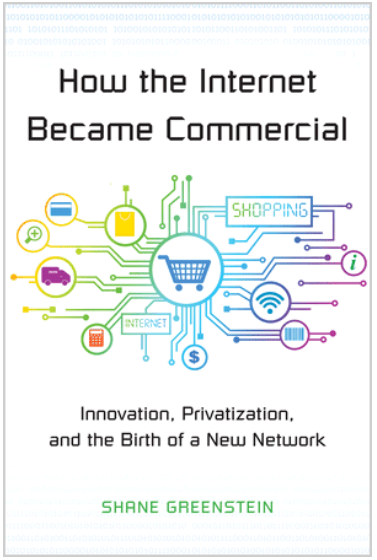

Extra slides

Variance in broadband use/speed. Why does this arise? What impact does it have, if any?

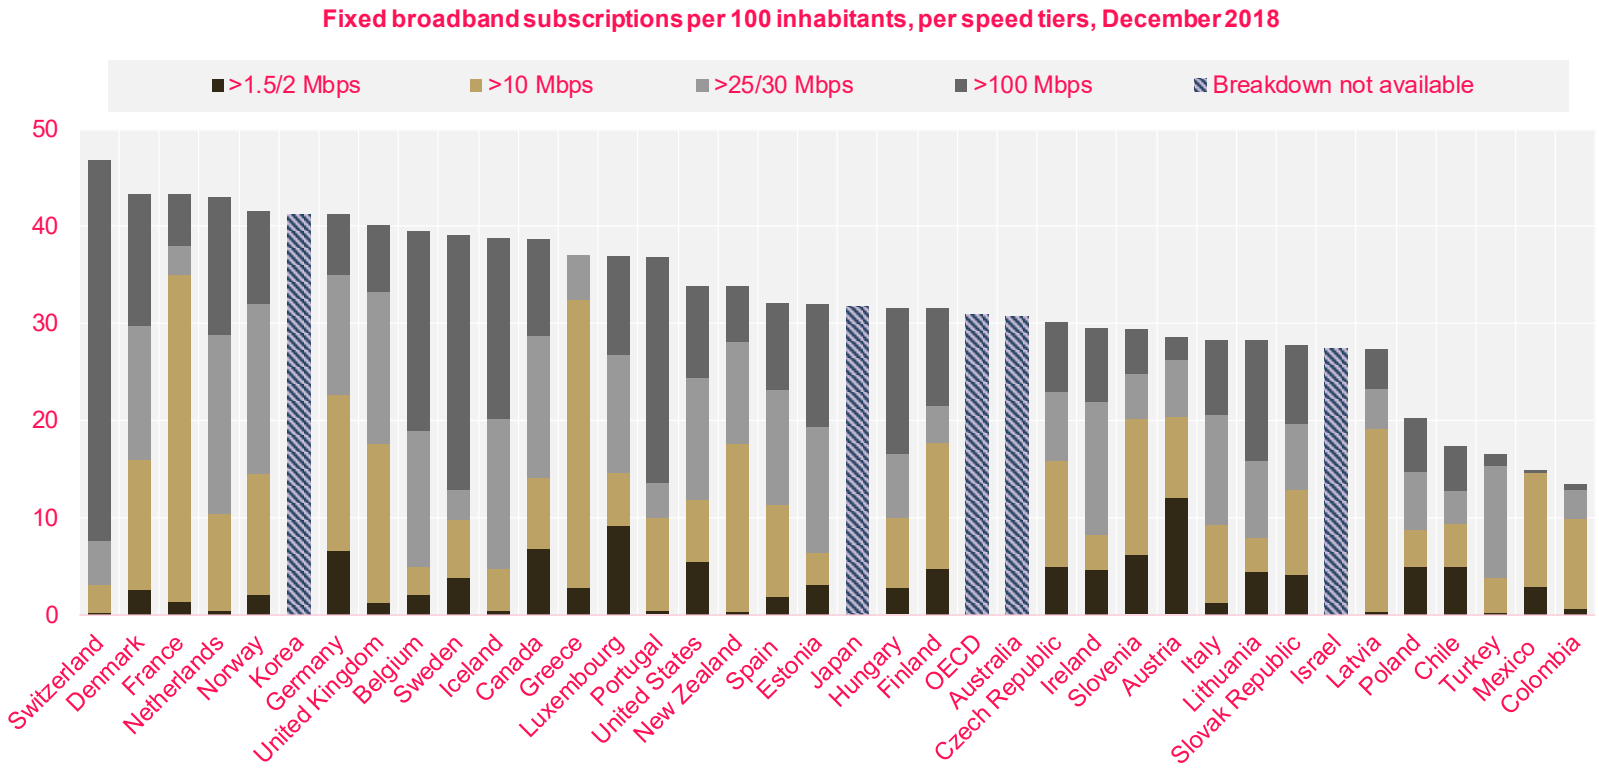

The OECD keeps lots of country-wide data. Room for x-country analysis….
Wallsten and Riso, 2010, good first step on pricing….but much less research on this topic than you would have thought…

Information superhighway? A flawed comparison

Like a highway? Rights of way shared with roads. Physical resemblance.

Shared capital. Like a new road, reduces frictions bringing assets into contact, reducing frictions for new services (Goldfarb and Tucker, 2018).

Continued improvement. Frequently restructured architecture that no designer could anticipate. New innovative applications built atop frontier infrastructure. Unlike any other infrastructure in this respect.

Privately financed. Excludable. Discretion to supplier over territory, pricing, and quality. Minimal subsidy. Few universal service regulations.