




Cloud Computing: Co-Invention for the Masses

Cloud computing is a big and growing business. Revenue [estimates](#) for 2017 range from just under \$150 billion to about \$180 billion.



In some respects, cloud computing is simply an incremental improvement over traditional client/server computing. Users work at terminals that are typically personal computers, while the processing and storage are located elsewhere. In other respects, however, cloud is a monumental shift in computing. Unlike traditional c/s computing, where access is specialized and available only within a given organization, anybody with access to a broadband connection can access some of the most powerful computing power available at rapidly decreasing prices.

One would expect cloud, therefore, to have two economic effects. First, it is likely to become increasingly cost-ineffective for firms to house and maintain their own computing resources. This process may be similar to firms' evolution from using mainframes to minicomputers, c/s architecture, and personal computers. Second, the process of co-invention, [examined in detail](#) by Tim Bresnahan and Shane Greenstein more than two decades ago, where the technology becomes more valuable as users adapt it and themselves in unexpected ways, may occur more quickly and ultimately have a larger economic impact than other high-end computing advancements due to the larger scale and scope of cloud.

This TPI Insight discusses the evolution of cloud, its similarities and difference to earlier computing innovations, its likely economic effects, and sets the stage for research into the economics of cloud computing.

Cloud Computing: Decreasing Costs and the Evolving C/S Architecture

A large part of the story of cloud computing—where access to mainframe computing power and massive storage

Economics Staff

Scott Wallsten, PhD

(202) 828-4405

swallsten@techpolicyinstitute.org

Robert Hahn, PhD

(202) 828-4405

rhahn@techpolicyinstitute.org

Thomas Lenard, PhD

(202) 828-4405

tlenard@techpolicyinstitute.org

Sarah Oh, JD, PhD

(202) 425-7725

soh@techpolicyinstitute.org

Lindsay Poss, MS

(202)-828-4405

lposs@techpolicyinstitute.org

Nathaniel Lovin

(202) 828-4405

nlovin@techpolicyinstitute.org

For Press Inquiries

David Fish

(571) 389-4446

dfish@techpolicyinstitute.org

Technology Policy Institute

409 12th Street, SW


Suite 700


Washington, D.C. 20024

capacity is available to anyone—is decreasing costs. The necessary inputs into cloud computing have allowed it to scale and set prices that make it widely available, while the costs of using cloud services make it accessible to almost anyone, not just big businesses.

Decreases in the Cost of Cloud Inputs

At the risk of grossly oversimplifying, cloud computing requires three primary inputs: processing power, data storage, and an ability to access both of those at a reasonable cost. The price of all three inputs has been falling steadily and precipitously for decades, as the figures below demonstrate. Processing and storage prices have fallen so much that figures must use log scales to avoid making prices look like they fell to nothing years ago. Similarly, data transit prices fell from about \$1200 per Mbps in 1998 to \$0.02 per Mbps in 2017.





Affordable Massive Computing Power and Storage Available to Everyone, Everywhere

The real prices of personal computers have been falling for decades, but the low cost of processing power, storage, and transit made it possible for anyone to access computers with significantly more capabilities. Relatively suddenly, high-end computing became available and affordable to small firms and entrepreneurs, not just to firms large enough to support their own mainframe computer systems.


Turning Fixed Costs into Marginal Costs

Rather than investing significant sums into high-end computing equipment, cloud services make it possible to buy storage and processing a la carte, with a large number of available combinations of features. Amazon's AWS, for example, lists (in addition to a free tier) three prices per GB for standard storage based on the amount of data stored plus three additional tiers for data needed only infrequently. It then has a menu of prices for related services, like data transfers, queries, and tagging. And all of those [differ by geographic region](#). Google cloud [explains](#) how the user can create "virtual machines" that use as much computing power as the user needs, ranging from 1 vCPU (virtual CPU) to 64 vCPUs, each with up to 6.5 GB of memory. The more processing power the user needs, the more the user pays.


In short, users can buy the processing power and storage they need at any given moment rather than sink money into expensive hardware. Cloud has turned what used to be a fixed cost to many firms into a marginal cost.

Competition and Prices in Cloud Services

The cloud "industry" itself is comprised of a fairly large number of firms and rapidly decreasing prices. The largest providers of cloud services in the U.S. are Amazon (AWS), Google (Google Cloud), and Microsoft (Azure), with the largest by revenues being AWS. But many other firms offer cloud services, sometimes targeted at particular uses.



Comparing cloud prices is complicated due to the detailed a la carte nature of pricing discussed above. Nevertheless, [Byrne et al \(2017\)](#) created price indices of the various cloud services offered. They estimate that cloud processing prices fell by about half between 2009 and the end of 2016, while storage prices fell between 70 and 80 percent during that time period.



Cloud Computing Also a Throwback to Early Computing

In addition to similarities to traditional c/s computing, cloud also has certain similarities to computing even prior to the adoption of c/s. In particular, cloud computing is creating a new type of worker trained in specific cloud skills and, although making powerful computing services available at low prices, also increases the costs of programming

errors.

Data Scientist as an Occupation

Operating a mainframe computer required certain skills, and operating any given company's computing infrastructure required additional specialized training. Similarly, using cloud resources directly also requires specialized skills—generally knowledge of coding in C++ and the ability to run complex SQL queries at a minimum.¹ There is sufficient demand for cloud-applicable skills that universities and other educational institutions increasingly offer courses and degrees in “data science.” While the phrase has been used for decade and even today has no single accepted definition, generally a data scientist is trained to work with data in the cloud. In 2018, the U.S. Bureau of Economic Analysis added “Data Scientist” to its [official list of occupations](#).² Thus, just as mainframe computing required specialized training, so, too, does use of cloud computing.

Increase in Cost of Programming Errors

In the era of mainframe computing, programming errors could be costly in ways that programming in a local environment could not. Similarly, programming errors in the cloud can entail large financial costs.

From the 1930s until the early 1980s computer programs were written and run using Hollerith punch cards. Writing programs was labor-intensive and errors were costly, since error-correction required replacing punch cards. Decreases in the cost of computer memory, which allowed programs to be stored on tape and held in the computer's memory while running, eventually made punch cards obsolete.³ Eliminating punch cards reduced the cost of programming errors since they could generally be fixed quickly.

Cloud computing continues to reduce normalized prices to users of every relevant metric. However, its a la carte nature means that errors can be costly in ways many of today's computer users have never experienced. Sloppy coding for queries on a massive dataset can result in charges significantly higher than a more carefully-designed query. One user on Stack Overflow [wrote about](#) how a Google BigQuery program ended up costing over \$4,000 when a more careful implementation based on only a sample of the data would have cost just a few dollars.⁴

Some fields, like those involving genomics, protein modeling, and rendering special effects animation, have always had to consider computing costs. But other fields, like empirical social sciences and many business applications, have been able to treat the cost of each additional computer operation as zero—the cost of the machine and any relevant software were fixed and the only measurable marginal cost was time. With new types of data available for analysis, it is conceivable that the nature of programming itself may change as a result of cloud.

Coinvention

As mentioned earlier, cloud computing itself is already big business. But the real economic effects will come from the applications and services that rely on the cloud to work—services we already have but run more efficiently and ones that we cannot even imagine. As Bresnahan and Greenstein (1996) explain, technological improvements become truly economically valuable as people discover new uses: “Users, through their own experimentation and discovery, make technology more valuable. We call this activity co-invention to distinguish it from original invention.” (p.1). In the context of their 1996 paper, they discuss how the real benefits of client-server computing relative to traditional mainframe computing did not materialize immediately with the introduction of c/s computing because of the time required for co-invention. Somewhat foreshadowing the current obsession with platforms, they note that they “prefer a broad definition of platforms that includes the co-invented parts: users' knowledge about how to best configure, program, and utilize the system. What a platform can do depends not only on the technology invented by vendors and the market supply of complementary goods, but also on human capital at user establishments.” (B&G p.7)

The BG hypothesis suggests that the biggest benefits of cloud computing will not happen immediately. It will take time for people, firms, and other institutions to fully integrate this new approach to computing into their various activities. However, cloud computing is available to everyone rather than just select companies and institutions.

¹We are not referring here to products built on the cloud.

²Under the 2018 Standard Occupational Classification System:

“15-2051 Data Scientists. Develop and implement a set of techniques or analytics applications to transform raw data into meaningful information using data-oriented programming languages and visualization software. Apply data mining, data modeling, natural language processing, and machine learning to extract and analyze information from large structured and unstructured datasets. Visualize, interpret, and report data findings. May create dynamic data reports. Excludes ‘Statisticians’ (15-2041), ‘Cartographers and Photogrammetrists’ (17-1021), and ‘Health Information Technologists and Medical Registrars’ (29-9021).

Illustrative examples: Business Intelligence Developer, Data Analytics Specialist, Data Mining Analyst, Data Visualization Developer.”

³Although they did have an impressive second life as notecards for more than a decade after their last use in computing until the world's supply was extinguished.

⁴The post is from 2013, and cloud computing prices have come down significantly since then. Still, the point remains correct.

This mass availability radically increases the size of the population that can innovate with these new tools and, presumably, decrease the amount of time required to generate large benefits.

Conclusion

Cloud computing represents an upward leap in computing resources available to society, largely due to the confluence of decreases in processing power, storage, and data transit. The large number of firms and entrepreneurs who can make use of these tools suggests that it will generate significant economic benefits. History suggests that the biggest economic effects are yet to come.

Recent TPI Insights

Econometrics in the Cloud: Robust Standard Errors in BigQuery ML (Dec 10, 2019)

Econometrics in the Cloud: Extending Google BigQuery ML (Nov 6, 2019)

Economics, Experts, and Federalism in *Mozilla v. FCC* (Oct 4, 2019)

The Law and Economics of *Apple Inc. v. Pepper* (Dec. 20, 2018)

Stay tuned for more economic and legal analysis from Washington, D.C. in *TPI Insights*. Contact Ashley Benjamin at (202) 828-4405 for more information

Technology Policy Institute

The Technology Policy Institute is a non-profit research and educational organization that focuses on the economics of innovation, technological change, and related regulation. More information is available at www.techpolicyinstitute.org.