

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: Economic Analysis of the Digital Economy

Volume Author/Editor: Avi Goldfarb, Shane M. Greenstein, and Catherine E. Tucker, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-20684-X; 978-0-226-20684-4

Volume URL: <http://www.nber.org/books/gree13-1>

Conference Date: June 6–7, 2013

Publication Date: April 2015

Chapter Title: Comment on "Modularity and the Evolution of the Internet"

Chapter Author(s): Timothy F. Bresnahan

Chapter URL: <http://www.nber.org/chapters/c13056>

Chapter pages in book: (p. 47 – 53)

- Rysman, M., and T. Simcoe. 2008. "Patents and the Performance of Voluntary Standard Setting Organizations." *Management Science* 54 (11): 1920–34.
- Saltzer, J. H., D. P. Reed, and D. D. Clark. 1984. "End-to-End Arguments in System Design." *ACM Transactions on Computer Systems* 2 (4): 277–88.
- Sanchez, R., and J. T. Mahoney. 1996. "Modularity, Flexibility, and Knowledge Management in Product and Organization Design." *Strategic Management Journal* 17:63–76.
- Simcoe, T. 2012. "Standard Setting Committees: Consensus Governance for Shared Technology Platforms." *American Economic Review* 102 (1): 305–36.
- Simon, H. A. 1962. "The Architecture of Complexity." *Proceedings of the American Philosophical Society* 106 (6): 467–82.
- Smith, A. 1776. *Wealth of Nations*, vol. 10, Harvard Classics, edited by C. J. Bullock. New York: P. F. Collier & Son.
- Stigler, G., and R. Sherwin. 1985. "The Extent of the Market." *Journal of Law and Economics* 28 (3): 555–85.

Comment Timothy F. Bresnahan

In "Modularity and the Evolution of the Internet" Tim Simcoe brings valuable empirical evidence to bear on the structure and governance of the Internet's more technical, less customer-facing, layers. His main empirical results are about the Internet's protocol stack, that is, the structure of the technical layers' modular architecture and of the division of labor in invention of improvements.

To organize my discussion, I will follow Simcoe's main results. There are, however, three distinctions that I want to draw before proceeding: (1) modularity is not the same as openness; (2) one can say that an architecture is modular (or open), which is not the same as saying the process by which the architecture changes is modular (or open); and (3) the Internet, like most ICT platforms, includes both purely technical standards and de facto standards in customer-facing products.

1. Modularity is related to, but not the same as, openness. Modularity is an engineering design concept. A large, complex problem can be broken up into pieces, and engineers working on one piece need know only a small amount about all the other pieces. They *do* need to know how their piece can interact with the other pieces—for which they (ideally) need know only the information contained in the interface standards described in the IETF (and preceding) and W3C documents analyzed by Simcoe. In contrast, openness is an economic organization concept. It refers to the availability and control

Timothy F. Bresnahan is the Landau Professor in Technology and the Economy at Stanford University and a member of the board of directors of the National Bureau of Economic Research.

For acknowledgments, sources of research support, and disclosure of the author's material financial relationships, if any, please see <http://www.nber.org/chapters/c13056.ack>.

of information about interface standards and to the role of a platform sponsor as a gatekeeper. In a closed (or proprietary) architecture, a GPT sponsor controls certain interface standards, and access to information about those standards flows to other firms through contracting with the sponsor. The sponsor can compel others to contract either because it only has the interface information or because it controls access to distribution to customers or both. Modularity makes openness feasible, but many proprietary architectures are quite modular.

2. Modularity is most precisely used as a modifier of an architecture at a moment in time. Modularity in this sense means that the boundaries between layers exist and “local” inventive effort can proceed. An architecture can remain modular over time, however, either by respecting the old boundaries (a part of “backward compatibility”) or by moving them in light of new technical or market developments. As we move to this dynamic viewpoint, an important element of openness is that outsiders can define new general-purpose layers and add them to the stack.

3. The Internet, like most multilayered GPTs, has both technical layers and user-facing layers among its general-purpose components. Simcoe focuses on technical layers and the interfaces between them. He does not focus on the commercial layers that connect the Internet to customers. Search, from Google or Microsoft, is an important general-purpose layer in the Internet for both users and advertisers. So, too, is product search inside Amazon or eBay or other storefronts, for both merchants and consumers. For a long time, the Internet index created by Yahoo appeared to be a general-purpose component. Other examples abound. The key point is that not all of the general components associated with the Internet fall within the organized standard setting of the IETF or the W3C. Some are, instead, set in markets or by dominant firms in some layer.

A Great Transformation as New Uses Are Found

Simcoe usefully notes that the time-series pattern of the count of Internet documents (RFCs and W3C publications) corresponds to the role of the Internet as a GPT, or more precisely, a GPT for which important applications were discovered after a lag. If we interpret the count of documents as an indicator of the amount of inventive activity, there is a burst of invention in the 1970s, comparatively less until the 1990s, and a steady growth from the mid-1990s through the present day. This corresponds broadly to the two main eras of the application of the Internet. From its invention until the commercialization of the Internet in the early 1990s, the Internet largely connected technical users in military and academic labs. While there was steady invention throughout this period, Simcoe shows that the architecture of the Internet, at least as measured by the count of documents, needed to

be invented to support this technical-user era but, once invented, did not need radical expansion in capabilities.

The second main era in the application of the Internet is its widespread use for commercial and mass market electronic communication, commerce, and content, hereafter EC³. The commercial portion of this begins in the early to mid-1990s, and, famously, the mass market part of this in the mid- to late 1990s. As Simcoe shows, the ongoing explosion in the range of applications of the Internet that began then and continues to the present has been associated with a dramatic expansion in the number of Internet documents. His interpretation, which is clearly right, is that the wider range of applications elicited new improvements in the general-purpose components. This pulls together a familiar and an unfamiliar aspect of GPT economics. Familiarly, important applications of a GPT can lag years behind its original invention. Less familiarly, new applications, particularly if they involve much larger demand for the GPT than earlier ones, can call for changes in the technical capabilities of the general-purpose components themselves.

Surprising Persistence of Openness

As Simcoe suggests, this transformation involves at least two surprising and very positive developments: commercialization without proprietization and expansion by outsiders. Both are related to modularity and openness.

Most commercial computing and communications platforms are proprietary.¹ The IBM 360 family was proprietary from the get-go, though an essential feature of the family was its modular architecture. The personal computer (PC) began as an open system, but is now the proprietary Microsoft Windows platform, even though there is a great deal of modularity in its architecture. The Oracle or SAP software platforms of the present are at once modular and proprietary. In each case, a single-firm GPT sponsor maintains control over the GPT and, in particular, either controls or commodifies supply of general-purpose layers. The Internet moved from being mostly a technical-uses GPT to being mostly a commercial-uses GPT without (yet) becoming a proprietary platform with a dominant sponsor firm, and with continued openness. This is a borderline miracle.

How the miracle of commercialization without proprietization was achieved is partly reflected in Simcoe's tables. Within the technical layers there continues to be an open architecture, and he shows this. Still, our best understanding of how and why this miracle occurred comes from detailed

1. As Bresnahan and Greenstein (1999) point out, this tendency is less marked for technical platforms such as minicomputers. Thus, the distinction between the technical layers of the Internet and the commercial GPTs running "on top of" them is economically important.

examinations of the important historical epochs at which there was a risk of some or all of the Internet becoming proprietary. Shane Greenstein (forthcoming) writes with compelling depth and understanding of the exit of the NSF from Internet funding, the “commercialization of the Internet.” At that stage, it could easily have transited to being an IBM technology—only a very thoughtful exit by the NSF prevented this. Another moment when the Internet might have become proprietary was after Microsoft won the browser war. Faced with substantial scope diseconomies between the businesses offering Windows and the Internet (Bresnahan, Greenstein, and Henderson 2012), the firm ultimately focused on maintaining control of the Windows standard for mass market computing and chose not to use command of the browser to proprietize the Internet.

These important historical transitions illustrate an important theme about causation. The technical layers of the Internet stack studied by Simcoe have remained open and modular in part because of their governance, as Simcoe suggests. Equally important, however, has been the absence of a takeover of standards setting by the firm supplying a complementary commercial layer.

Outsider Innovation

The second surprising and very positive development is expansion of the set of open, modular, general-purpose layers of the Internet by outsiders. An important pair of examples is the World Wide Web (WWW) and the web browser. These inventions transformed the Internet into a mass medium. Today, if you ask most consumers what the Internet is, they will answer in terms of the WWW viewed through a browser. Both the WWW and the web browser were new layers in the stack. Economically, they are complements to the preexisting layers of the Internet.

The openness of the Internet architecture meant that the WWW could be invented without getting the permission of any suppliers of existing Internet components or engaging in contracts with them. Instead, the WWW could be defined in a way that it “runs on top of” the Internet; that is, that it interacts with the other layers through open interface standards. This is, as Shane Greenstein (forthcoming) has emphasized, an important element of open organization. In turn, the outsiders who invented and (some of whom) later commercialized the web browser did not need to get the permission of the inventors of the WWW or engage in contracts with them. This would have gone badly if it were required, since Tim Berners-Lee, inventor of the Web, strongly disapproved of the web browser once it became commercialized at Netscape. This is an important example of uncontrolled, uncontracted for, invention by outsiders permitted by open systems, for the series of events culminating in the commercialization of the web browser is one of the top ten economic growth innovations of the twentieth century.

Decomposability, Division of Labor, and Diffusion

Simcoe uses citations—from later Internet documents and from patents—to Internet documents to examine the structure of Internet innovation, both organizationally and technically, and the diffusion of new applications of the Internet. This is an extremely valuable undertaking and we can learn much from it. Of course, it also suffers from the difficulties of citations analysis generally.

Simcoe's analysis of the division of innovative labor seems to me to be a particularly successful deployment of citations methods. The Internet is largely modular in its different technical layers, and firms that work on a layer also tend to patent inventions that are related to that layer. As he points out, considerable gains have been made by having multiple firms inventing and supplying general-purpose components.

The study of the diffusion of new applications for the Internet is a difficult one, and particularly so from a technical-layer-centric perspective. This is, of course, not particularly a weakness of Simcoe's chapter. Data sets on new technologies generally emphasize the technical rather than application. One cautionary note, however, is what the measurable perspective of an "application" is here. Most of the "applications" studied by Simcoe are themselves GPTs, which connect to the Internet and to which, in turn, many specific applications are connected. This is not a small point. A list of things that are *not* applications from the perspective of the citations used in this chapter includes Google Search, Facebook social networking, and Apple media and applications sales in the iTunes store. My interpretation would be that there is no doubt that the enormous transformation of the uses of the Internet to the commercial realm and then to mass market EC³ is behind these tables, but that it is less obvious that the timing or breadth of the spread of applications can be seen in these tables. A difficulty for patent citations is that patent policy is changing over the relevant time period, so that it is not obvious whether the quantitative growth lies in the breadth of applications or in the tendency to patent inventions. The Internet document citations difficulty is that they are, by their nature, from within the standardized GPT layers of the Internet, not from applications. Only insofar as new applications lead to a change in the GPT layers will an expansion of applications be reflected there.

The Framework

Ultimately, the most interesting thing about Simcoe's chapter is the perspective it takes on the analysis. We have two very different literatures on coordination between suppliers of general-purpose components and applications. These are sufficiently different, especially in their treatment of the optimal form of coordination, that much confusion has arisen.

The first literature, typically writing about “two-sided markets” or “platform economics,” is concerned mostly with the coordination of production and prices.² The literature takes a contractual approach to the coordination of applications supply with platform (GPT) supply. To facilitate the contractual approach, the most common assumption is that the general-purpose components are supplied by a single firm. By that I mean each platform or GPT cluster has a single supplier of general-purpose components at its center, and that this firm contracts with, or offers incentives to, suppliers of applications. Sometimes there is competition to be (or to become) the dominant platform or GPT, so that there are competing central sponsors, each offering contracts or incentives to an atomless distribution of applications developers.

While the second literature, typically calling itself “GPT” or “Recombination,”³ treats the same industries, it emphasizes very different phenomena and modeling elements. First, this literature is concerned with the problem of invention, especially repeated rounds of invention, much more than pricing and production. This arises because the practical GPT literature has had to deal with the phenomenon—so emphasized by Simcoe—of general-purpose components supplied by many firms. The “layered” architecture of systems like the Internet involves competition within each layer (rather than competition between whole systems), but complementary invention of improvements across layers. An important general point of this literature is that explicit contracts to coordinate innovation may be impossible so that “softer” governance structures such as the one described by Simcoe are optimal.

Why might the softer governance structures work? Are they optimal only because the governance structure we would really like, explicit contracts among complementary suppliers, is impossible? There are several important points to make here. The most important point concerns the possibility of unforeseen and perhaps unforeseeable change. Sometimes after a period of exploitation of a general-purpose technology, new demands or new inventions call for improvements in the general-purpose components. This is a moment at which not drawing too sharp a distinction between “applications” and general-purpose components can be valuable. A system that is open to the invention of new applications (in the strong sense that they do not need to contract with anyone) will have low barriers to entry. If an application is very widely used and itself becomes a general purpose input into new applications, then the platform is transformed.

In Simcoe’s chapter, as in other studies, we see the value of uncoordinated (or only loosely coordinated) innovation for this kind of *ex post* flexibility.

2. See Jullien (2011) or Rysman (2009). An important exception is Tirole and Weyl (2010), which attempts to extend this framework to invention.

3. See Bresnahan and Trajtenberg (1995).

Modularity and openness permit flexible innovation ex post. They permit flexibility not only in reconfiguration of the platform's general-purpose components but also in allowing an ex post opportunity for multiple heterogeneous innovators to undertake differentiated efforts to improve the general-purpose components of the same GPT. Elsewhere (Bresnahan 2011) I have argued that it was the modularity and openness of the Internet that made it the winner in a multiway race to be the general-purpose technology underlying the enormous EC³ breakthroughs of the last two decades. Simcoe offers us a fascinating glimpse into the workings of that modularity and openness underlying flexible improvements in the Internet's GPT components.

References

- Bresnahan, T. 2011. "General Purpose Technologies." In *Handbook of the Economics of Innovation*, edited by Bronwyn Hall and Nathan Rosenberg. North Holland: Elsevier.
- Bresnahan, T., and S. Greenstein. 1999. "Technological Competition and the Structure of the Computer Industry." *Journal of Industrial Economics* 47 (1): 1–40.
- Bresnahan, T., S. Greenstein, and R. Henderson. 2012. "Schumpeterian Competition and Diseconomies of Scope: Illustrations from the Histories of Microsoft and IBM." In *The Rate and Direction of Inventive Activity Revisited*, edited by Josh Lerner and Scott Stern. Chicago: University of Chicago Press.
- Bresnahan, T., and Manuel Trajtenberg. 1995. "General Purpose Technologies: 'Engines of Growth'?" *Journal of Econometrics* special issue 65 (1): 83–108.
- Greenstein, S. Forthcoming. *Innovation from the Edges*. Princeton, NJ: Princeton University Press.
- Jullien, B. 2011. "Competition in Multi-Sided Markets: Divide-and-Conquer." *American Economic Journal: Microeconomics* 3 (4): 1–35.
- Rysman, M. 2009. "The Economics of Two-Sided Markets." *Journal of Economic Perspectives* 23:125–44.
- Tirole, Jean, and Glen Weyl. 2010. "Materialistic Genius and Market Power: Uncovering the Best Innovations." IDEI Working Paper no. 629. Institut d'Économie Industrielle (IDEI), Toulouse, France.