

Computer Mediated Transactions

Hal R. Varian*

March 6, 2010

Every now and then a set of technologies becomes available that sets off a period of “combinatorial innovation.” Think of standardized mechanical parts in the 1800s, the gasoline engine in the early 1900s, electronics in the 1920s, integrated circuits in the 1970s, and the internet in the last decade or so.

The component parts of these technologies can be combined and recombined by innovators to create new devices and applications. Since these innovators are working in parallel with similar components, it is common to see simultaneous invention. There are many well-known examples, such as the electric light, the airplane, the automobile, and the telephone. Many scholars have described such periods of innovation, using terms such as “recombinant growth,” “general purpose technologies,” “cumulative synthesis” and “clusters of innovation.”¹

The internet and the web are wonderful examples of combinatorial innovation. In the last 15 years we have seen a huge proliferation of web applications, all built from a basic set of component technologies.

The internet itself was a rather unlikely innovation; I like to describe it as a “lab experiment that got loose.” Since the internet arose from the research community rather than the private sector, it had no obvious business model. Other public computer networks, such as AOL, CompuServe, and Minitel, generally used a subscription models, but were centrally controlled and offered little scope for innovation at the user level. The internet won out over these alternatives, precisely because it offered a flexible set of component technologies which encouraged combinatorial innovation.

*UC Berkeley and Google. hal@ischool.berkeley.edu

¹See, for example, Weitzman (1998), Bresnahan and Trajtenberg (1995), Bresnahan (2010), Rosenberg (1976), Usher (1998), and Schumpeter (2000).

The earlier waves of combinatorial innovation required decades or more to play out. For example, Hounshell (1984) argues that the utopian vision of interchangeable parts took well over a century to be realized. The web was invented in the early 1990s, but didn't really become widely used till the mid-1990s. Since then we have seen a huge number of novel applications from web browsers, to search engines, to social networks, to mention just a few examples. As with the internet, the web initially had no real business model, but offered a fertile ground for combinatorial innovation.

Why was innovation so rapid on the internet? The reason is that the component parts were all bits. They were programming languages, protocols, standards, software libraries, productivity tools and the like. There was no time to manufacture, no inventory management, and no shipping delay. You never run out of HTML, just like you never run out of email. New tools could be sent around the world in seconds and innovators could combine and recombine these bits to create new web applications.

This parallel invention has led to a burst of global innovation in web applications. It is true that the internet was an American innovation, but the web was invented by an Englishman living in Switzerland. Linux, the most used operating system on the web, came from Finland, as did MySQL, a widely used database for web applications. Skype, which uses the internet for voice communication, came from Estonia.

Of course there were many other technologies with worldwide innovation such as automobiles, airplane, photography, and incandescent lighting to name just a few. However applications for the internet, which is inherently a communications technology, could be developed everywhere in the world in parallel, leading to the very rapid innovation we have observed.

My interest in this lecture is in the economic aspects of these technological developments. I start with a point so mundane and obvious, it barely seems worth mentioning.

1 Computer mediated transactions

Nowadays, most economic transactions involve a computer. Sometimes the computer takes the form of a smart cash register, sometimes it is part of a sophisticated point of sale system, and sometimes it is a web site. In each of these cases, the computer creates a record of the transaction.

The record-keeping role was the original motivation for adding the com-

puter to the transaction. Creating a record of transactions is the first step in building an accounting system, thereby enabling a firms to understand how its business is doing.

But now that these computers are in place, they can be used for many other purposes. In this lecture I would like to explore some of the ways that computer mediation can affect economic transactions. I argue that these computer mediated transactions have enabled significant improvements in the way transactions are carried out and will continue to impact the economy for the foreseeable future.

I classify impact of computer mediated transactions into four main categories.

- Facilitate new forms of contract;
- Facilitate data extraction and analysis;
- Facilitate controlled experimentation;
- Facilitate personalization and customization;

2 Enable new forms of contract

Contracts are fundamental to commerce. The simplest commercial contract says “I will do X if you do Y,” as in “I will give you \$1 if you give me a cup of coffee.” Of course, this requires that the actions be verifiable. Just because I ask for coffee doesn’t mean that I will get some. As Abraham Lincoln supposedly remarked, “If this is coffee, please bring me some tea; but if this is tea, please bring me some coffee.”

A computer in the middle of a transaction can observe and verify many aspects of a transaction. The record produced by the computer can allow the contracting parties to condition the contract on terms that were previously unobservable, thereby allowing for more efficient transactions.

I am not claiming that increased observability will *necessarily* lead to more efficient contracts. There are counterexamples to the claim that “more information is better” such as the famous Hirshleifer (1971) example. I am only claiming that additional information *allows* for more efficient contracts.

Of course, the study of contracts is a highly developed field in economics and it is hardly novel to suggest that contractual form depends on what is ob-

servable. What is interesting, I think, is the way that progress in information technology enables new contractual forms.

Consider, for example, a rental-car agency that buys insurance based on accident rates, and that accident rates in turn depend on the speed at which a vehicle is operated. All renters would prefer to drive within the speed limit if they were compensated with a lower rental fee. However, if there is no way to monitor the speed of the rental car, such a contractual provision is unenforceable. Putting a computer transmitter in the trunk of the car that records the vehicle's speed makes the contract enforceable and offers a Pareto improvement on the original arrangements.²

The transportation sector has capitalized on the availability of computerized transmitters to create more efficient contracts in a number of areas.

- Car dealers are selling cars with “starter interrupt” devices that inhibit operations if car payments are missed. (Press (2006))
- Similar interrupt devices attached to breath analyzers are mandated for drunk driving offenders in many states.
- Parents can buy a device known as “MyKey” which allows them to limit auto speed, cap the volume on the radio, require seat belt use and encourage other safe-driving habits for teenage drivers. (Bunkley and Vlastic (2008))
- In the economics literature, Hubbard (2000) and Baker and Hubbard (2004) examine a variety of ways that vehicular monitoring systems have impacted the trucking industry.

There are many other examples of computer mediated contracts. One nice example is the work of Dana and Spier (2001) and Mortimer (2008) that describes the efficiency gains resulting from revenue sharing in the video tape rental industry.

Video tapes were originally purchased by retail stores from distributors for about \$65 per tape. Since the videos were so expensive, stores bought a only a few. As a result, the popular videos quickly disappeared from the shelves, making everyone unhappy.

²This is a particularly simple case. If drivers have heterogeneous preferences, those who prefer to speed may be made worse off by the availability of such a device.

In 1998, retailers and distributors adopted a new business model, a revenue sharing arrangement in which the stores paid a small upfront fee of \$3 to \$8 but split the revenue when the video was rented, with 40% to 60% going to the retailer. Stores no longer had an incentive to economize on purchase, and all parties to the transaction—retailers, distributors, and customers—were made better off.

Sharing revenue at point of sale requires that both parties be able to monitor the transaction. The technological innovations of bar code scanning, the computerized cash register, and dial-up computer networks were the technologies that enabled revenue sharing arrangements.

Of course, when a transaction takes place online, revenue sharing is much easier. Online advertising is a case in point where revenue from an advertiser for an ad impression or click may be split among publishers, ad exchanges, ad networks, affiliates and other parties based on contractual formulas.

I have emphasized the benefits from computers offering *more* information to contracting parties, but there are also cases where computers can be used to improve contractual performance by hiding information, by using cryptographic methods. A picturesque example is the “cocaine auction protocol” which describes an auction mechanism designed to hide as much information as possible. (Stajano and Anderson (1999).)

Finally, I should mention “algorithmic game theory,” an exciting hybrid of computer science and economic theory. This subject brings computational considerations to game theory (how a particular solution can be computed) and strategic considerations to algorithm design (is a particular algorithm actually incentive compatible?) See Nisan, Roughgarden, Tardos and Vazirani, eds (2007) for a comprehensive collection of articles and Varian (1995) for an early contribution.

2.1 Some history of monitoring technologies

Though I have emphasized *computer* mediated transactions, the definition of computer can be considered to be quite broad. The earliest example I have been able to find for an accounting technology that enabled new forms of contract involves Mediterranean shipping circa 3300 BC.

The challenge was how to write a “bill of lading” for long distance trade in societies that were pre-literate and pre-numerate. The brilliant solution was to introduce small clay tokens, known as “bullae,” which were small representations of the material being transported. As each barrel of olive oil

was loaded onto a ship, a barrel-shaped token was placed in clay envelope. After the loading was completed, the envelope was baked in a kiln and given to the ship’s captain. At the other end of the voyage, the envelope was broken open and the tokens were compared to the barrels of oil on the ship as they were unloaded. If the numbers matched, the contract was verified. Later marks were scratched on the outside of the envelope that indicated the number of tokens inside, and some authors believe that this innovation led to the invention of writing between 3400 and 3300 BC. (Glassner, Bahrani and de Miero (2005))

A somewhat more recent example is the invention of the cash register in 1883 by James Ritty. Ritty was a saloon owner who discovered that his employees were stealing money. He hit upon the idea of building a device which would record each transaction on a paper tape, an invention that he patented under the name of “the incorruptible cashier.” (Patent 271,368.) Ritty’s machine formed the basis for the National Cash Register company founded in 1884. The NCR device added a cash drawer and a bell that sounded “ka-ching” whenever the drawer was opened which alerted the owner to the transaction, thereby discouraging pilfering. It is thought that this improved monitoring technology made retailers willing to hire employees outside the immediate family, leading to larger and more efficient establishments. See Yates (2000) for a more detailed account of the role of office machinery in the development of commercial enterprises.

2.2 Online advertising

Online advertising serves as a poster child for algorithmic mechanism design. A Pasadena company called GoTo came up with the idea of ranking search results using an auction. Users did not find this particular form of search attractive, so GoTo switched to using their auction to rank advertisements. In the original auction, ads were ranked by “bid per click” and advertisers paid the amount they bid. After consultation with auction theorists, GoTo moved to second-price auction: an advertiser paid a price per click determined by the bid of the advertiser in the next lower position. See Battelle (2005) and Levy (2009) for accounts of the development of these auctions.

There is a fundamental divergence of incentives in advertising. The publisher (the content provider) has space on its the web page for an ad and it wants to sell these ad impressions to the highest bidders. The advertiser doesn’t care directly about ad impressions but does care about visitors to its

website, and ultimately about sales of its products. So the publisher wants to sell impressions, but the advertiser wants to buy clicks.

This is like an international trade transaction where the buyer wants to pay in euros but the seller wants to receive dollars. The solution in both cases is the same: an exchange rate. In the context of online advertising the exchange rate is the predicted clickthrough rate, an estimate of how many clicks a particular ad impression will receive. This allows one to convert the advertiser's offered bid per click to an equivalent bid per impression, allowing the publisher to sell each impression to the highest bidder.

This mechanism aligns the interests of the buyer and seller, but creates other problems. If the advertiser only pays for clicks, then it has no direct incentive to economize on impressions. But excessive impressions impose an attention cost on users, so some further attention to ad quality is important in order to ensure that ad impressions remain relevant to users.

Nowadays, the major providers of search engine advertising all estimate clickthrough rates along with other measures of ad quality and use auctions to sell ads. Economists have applied the tools of game theory and mechanism design to analyze the properties of these auctions. See, for example, Athey and Ellison (2007), Edelman, Ostrovsky and Schwartz (2007), Varian (2007), and Varian (2009).

3 Enable data extraction and analysis

The data from computer mediated transactions can be analyzed and used to improve the performance of future transactions.

An illustrative example is Sabre air passenger reservation system offered by American Airlines. The original conception in 1953 was simply to automate the creation of an airline reservation. However, by the time the system was released in 1960, it had become apparent that such a system could also be used to study patterns in the airline reservation process: the acronym Sabre stands for Semi-Automatic Business Research Environment. (Sabre (2009).)

The existence of airline reservation systems enabled sophisticated differential pricing (also known as “yield management”) in the transportation business. See Smith, Leimkuhler and Darrow (1992) for the history of yield management in the airline industry and Talluri and van Ryzin (2004) for a textbook treatment.

Many firms have built data warehouses based on transaction level data which can then be used as input to analytic models for customer behavior. A prominent example is supermarket scanner data which has been widely used in economic analysis (see Nevo and Wolfram (2002), Hendel and Nevo (2006) for just two examples.) Scanner data has also been useful in constructing price indexes (Feenstra and Shapiro, eds (2003), Farm Foundation (2003)), since it allows for much more direct and timely access to prices. The fact that the data is timely is worth emphasizing, since it allows for real time analysis and intervention at both the business and policy level.

Choi and Varian (2009b) and Choi and Varian (2009a) use real-time publicly available search engine data to predict the current level of economic activity for automobile, real estate, retail trade, travel, and unemployment indicators. There are many other sources of real-time data such as credit card data, package delivery data, and financial data. Clements and Hendry (2003) and Castle and Hendry (2009) have coined the term “nowcasting” to describe the use of real-time data to estimate the current state of the economy. They use a variety of econometric techniques to deal with the problems of variable selection, gaps, lags, structural changes and so on. One promising development is that much of the real-time data is also available at state and city levels, allowing for regional macroeconomic analysis.

In the last 20 years or so, the field of machine learning has made tremendous strides in “data mining.” This term was once pejorative (at least among econometricians) but now enjoys a somewhat better reputation due to the exciting applications developed by computer scientists and statisticians; see Hastie, Friedman and Tibshirani (2009) for a technical overview. One of the big problems with data mining is over-fitting, but various sorts of cross-validation techniques have been developed that mitigate this problem. Econometricians have only begun to utilize these techniques; the previously mentioned work by Castle and Hendry (2009) is noteworthy in this respect.

4 Enable experimentation

As econometricians have observed, “if you torture the data long enough it will confess to anything.” It is difficult to establish causality from retrospective data analysis, so it is noteworthy that computer mediation allows one not only to measure economic activity but also to conduct controlled experiments.

In particular, it is relatively easy to implement experiments on web-based

systems. Such experiments can be conducted at the query level, the user level, or the geographic level.

In 2008, Google ran 6,000 experiments involving web search which resulted in 450–500 changes in the system. Some of these were experiments with the user interface, some were basic changes to the algorithm (Hoff (2009)). The ads team at Google ran a similar number of experiments, tweaking everything from the background color of the ads, to the spacing between the ads and search results, to the underlying ranking algorithm.

In the 1980s, Japanese manufacturers touted their “kaizen” system that allowed for “continuous improvement” of the production process. In a well-designed web-based business, one can have continuous improvement of the product itself—the website.

Google and other search engines also offers various experimental platforms to advertisers and publishers, such as “Ad Rotation,” which rotates ad creatives among various alternatives to choose the one that performs best and “Website Optimizer,” a system that allows websites to try out different designs or layouts and determine which performs best.

Building a system that allows for experimentation is critical for future improvement, but it is all too often left out of initial implementation. This is a shame, since it is the early versions of a system that are often in need of the most improvement.

Cloud computing, which I will discuss later in the lecture, offers a model for “software as service,” which typically means software which is hosted in a remote data center and accessed via a web interface. There are numerous advantages to this architecture but one that is not sufficiently appreciated is the fact that it allows for controlled experiments which can in turn lead to continuous improvement of the system. Alternatives, such as packaged software, make experimentation much more difficult.

Ideally, experiments lead to understanding of causal relations that can then be modeled. In case of web applications there are typically two “economic agents:” the users and the applications. The applications are already modeled via the source code that is used to implement them, so all that is necessary is to model the user behavior. The resulting model will often take the form of a computer simulation which can be used to understand how the the system works.

A nice example of this are the Bid Simulator and Bid Forecasting tools offered by Google and Yahoo. These tools give an estimate of the cost and clicks associated with possible bids. The cost per click is determined by the

rules of the auction and can be calculated directly; the clicks are part of user behavior and must be estimated statistically. Putting them together gives a model of the auction outcomes.

4.1 How experiments change business

The fact that computer mediation drastically reduces the cost of experimentation changes the role of management. As Kohavi, Longbotham, Sommerfield and Henne (2008) have emphasized, decisions should be based on carefully controlled experiments rather than “the Highest Paid Person’s Opinion (HiPPO).”

If experiments are costly, expert opinion by management is a plausible way to make decisions. But when experiments are cheap, they are likely provide more reliable answers than opinions—even opinions from highly paid people. Furthermore, even when experienced managers have better-than-average opinions, it is likely that there are more useful things for them to do than sit around a table debating about which background colors will appeal to web users. The right response from managers to such questions should be “run an experiment.”

Businesses have always engaged in experimentation in one form or another. But the availability of computer mediated transactions makes these experiments much more inexpensive and flexible than they have been in the past.

5 Customization and personalization

Finally, computer mediated transactions allow for customization and personalization of the interactions by basing current transactions on earlier transactions or other relevant information.

Instead of a “one size fits all” model, the web offers a “market of one.” Amazon, for example, makes suggestions of things to buy based on your previous purchases, or on purchases of consumers like you. These suggestions can be based on “recommender systems” of various sorts (Resnick and Varian (1997)).

Not only content, but prices may also be personalized, leading to various forms of differential pricing. What are the welfare effects of such personalized pricing? Acquisiti and Varian (2005) examine a model in which firms can

condition prices on past history. They find that the ability of firms to extract surplus is quite limited when consumers are sophisticated. In fact, firms have to offer some sort of “enhanced services” in order to justify higher prices.

Varian (2005b) suggests that there is a “third welfare theorem” that applies to (admittedly extreme) cases with perfect price discrimination and free entry: perfect price discrimination results in the optimal amount of output being sold while free entry pushes profits to zero, conferring all benefits to the consumers. See Ulph and Vulkan (2007) for a theoretical analysis of first-degree price discrimination.

The same sort of personalization can occur in advertising. Search engine advertising is inherently customized since ads are shown based on the user’s query. Google and Yahoo offer services that allow users to specify their areas of interest and then see ads related to those interests. It is also relatively common for advertisers use various forms of “re-targeting” that allows them to show ads based on previous responses of users to related ads.

6 Transactions among workers

My emphasis so far has been on transactions among buyers, sellers and advertisers. But computers can also mediate transactions among workers. The resulting improvements in communication and co-ordination can lead to productivity gains, as documented in the large literature on the impact of computers on productivity.

In a series of works, Paul David has drawn an extended analogy between the productivity impact of electricity at the end of the nineteenth century and the productivity impact of computing at the end of the twentieth century.³ Originally factories were powered by waterwheels which drove a shaft and all the machines in the factory had to connect to this central shaft. The manufacturing process involved moving the piece being assembled from station to station during assembly.

The power source evolved from waterwheels to steam engines to electric motors. Eventually electric motors were attached to each machine, which allowed more flexibility in how the machines were arranged within the factory. However, factories still stuck to the time-honored arrangements, grouping the same sort of machines in the same location—all the lathes in one place, saws in another, and drills in yet another.

³See David (1990), David (1991b), and David (1991a).

It wasn't until Henry Ford invented the assembly line in the first decade of the twentieth century that the flexibility offered by electric motors was really appreciated.⁴ As David (1990) shows, the productivity impact of the assembly line was huge and over the last century manufacturing has become far more efficient.⁵

I want to extend David (1990)'s assembly line analogy to examine "knowledge worker productivity" (Drucker (1999)). Prior to the widespread use of the personal computer, office documents were produced via a laborious process. A memo was dictated to a stenographer who later typed the document, making a half-dozen carbon copies. The typed manuscript was corrected by the author and circulated for comments. Just as with pre-assembly line production, the partially produced product was carried around from station to station for modification. When the comments all came back, the document was retyped, reproduced and recirculated.

In the latter half of the twentieth century there were some productivity enhancements for this basic process, such as White-Out, Post It Notes, and photocopy machines. But the basic production process remained the same for a century.

When the personal computer came along, editing became much easier, and the process of collaborative document production involved handing floppy disks back and forth. The advent of email allowed one to eliminate the floppy disk and simply mail attachments from person to person.

All of these effects contributed to improving the quantity and quality of collaborative document production. However, they all mimicked the same physical process: circulating a document from person to person for comment. Editing, version control, tracking changes, circulation of the documents and other tasks remained difficult.

Nowadays, we have a new model for document production enabled by "cloud computing." (Armbrust, Fox, Griffith, Joseph, Katz, Konwinski, Lee, Patterson, Rabkin, Stoica and Zaharia (2009), Wikipedia (2009b)). In this model, documents live "in the cloud," that is, in some data center on the

⁴Ford (1923) suggests that the inspiration for the assembly line came from observing the meatpacking plants in Chicago, where an animal carcass was hung on hooks and moved down a line where workers carved off different pieces. If you could use this process to dis-assemble a cow, Ford figured you could use it to assemble a car.

⁵I do not mean to imply that the only benefit from electric motors came from improved factory layout. Motors were also more efficient than drive belts and the building construction was simpler.

internet. The documents can be accessed any time, from anywhere, on any device, by any authorized user.

Cloud computing dramatically changes the production process for knowledge work. Now there is a single master copy that can be viewed and edited by all relevant parties, with version control, check points and document restore built in. All sorts of collaboration, including collaboration across time and space, have become far easier.

Instead of carrying the document around from collaborator to collaborator, a single master copy of the document can be edited by all interested parties. By allowing workflow to be re-organized, cloud computing changes knowledge worker productivity the same way that electricity changed the productivity of physical labor.

7 Deployment of applications

As mentioned earlier, cloud computing offers “software as service.” This architecture reduces support costs and makes it easier to update and improve applications.

But cloud computing doesn’t just offer “software as service”, it also offers “platform as service,” which means that software developers can deploy new applications using the cloud infrastructure.

Nowadays, it is possible for a small company to purchase data storage, hosting services, an applications development environment, and internet connectivity “off the shelf” from vendors such as Amazon, Google, IBM, Microsoft, Sun and others.

The “platform as service” model turns what used to be a fixed cost for small web applications into a variable cost, dramatically reducing entry costs. Computer engineers can not only explore the combinatorial possibilities of generic components to create new inventions—they can actually purchase standardized services in the market in order to deploy those innovations.

This development is analogous to what happened in the book publishing business. At one time publishers owned facilities for printing and binding books. Due to the strong economies of scale in this process, most publishers have outsourced the actual production process to a handful of specialized book production facilities.

Similarly, in the future it is likely that there will be number of cloud computing vendors that will offer computing on a utility-based model. This

production model dramatically reduces the entry costs of offering online services, and will likely lead to a significant increase in businesses that provide such specialized services. (Armbrust et al. (2009).)

The hallmarks of modern manufacturing are routinization, modularization, standardization, continuous production, and miniaturization. These practices have had a dramatic impact on manufacturing productivity in the twentieth century. The same practices can be applied to knowledge work in the twenty-first century.

Computers, for example, can automate routine tasks, such as spell-checking, or data retrieval. Communications technology allows tasks to be modularized and routed to the workers best able to perform those tasks. Just as the miniaturization of the electric motor allowed physical production to be rearranged in 1910, the miniaturization the computer—from the mainframe, to the workstation, to the PC, to the laptop, to the mobile phone—allows knowledge production to be rearranged at both a local and a global scale.

8 Micro-multinationals

One interesting implication of computer mediated transactions among knowledge workers is that interactions are no longer constrained by time or distance.

Email and other tools allow for asynchronous communication over a distance, which allows for optimization of tasks on a global basis. Knowledge work can be subdivided into tasks, much like the physical work in Adam Smith’s pin factory. But even more, those tasks can be exported around the world to where they can most effectively be performed.

For example, consultants at McKinsey routinely send their PowerPoint slides to Bangalore for beautification. There are many other cognitive tasks of this sort that can be outsourced, including translation, proofreading, document research and so on. Amazon’s Mechanical Turk (Wikipedia (2009a)) is a intriguing example of how computers can aid in matching up workers and tasks. As of March 2007 there were reportedly over 100,000 workers from 100 countries who were providing services via the Mechanical Turk (Pontin (2007)).

The dramatic drop in communications costs in the last decade has led to the emergence of what I call “micro-multinationals.” (Varian (2005a)) Nowadays, a 10 or 12 person company can have communications capabilities

that only the largest multinationals could afford 15-years ago. Using tools like email, web pages, wikis, voice over IP, and video conferencing, tiny companies can co-ordinate workflow on a global basis. By handing work off from one time zone to the next, these companies can effectively work around the clock, giving them a potential competitive advantage over firms that are restricted to one time zone.

Many micro-multinationals have a common history: a student comes to the US for graduate school. They use the internet and the collaborative tools available in scientific workgroups. Some may get bitten by the start-up bug. They draw on their friends and colleagues back home, who have other contacts living abroad. The collaborative technologies mentioned above allow such loose groups to collaborate on producing computer code which may end up as a working product.

As Saxenian (2006) has pointed out, “emigration” means something quite different now than it did 30 years ago. As she puts it “brain drain” has been replaced by “brain circulation.” Now we have email, web pages, wikis, voice over IP, and a host of other collaborative technologies that allow an immigrant to maintain ties to his social and professional network in his home country.

I began this essay with a discussion of combinatorial innovation, and pointed out that innovation has been so rapid in the last decade because innovators around the world can work in parallel, exploring novel combinations of software components. When the innovations are sufficiently developed to be deployed, they can be hosted using cloud computing technology, and managed by global teams—even by tiny companies. I believe that these capabilities will offer a huge boost to knowledge worker productivity in the future.

References

- Acquisiti, Alessandro and Hal R. Varian, “Conditioning Prices on Purchase History,” *Marketing Science*, 2005, 24 (3), 367–381.
- Armbrust, Michael, Armando Fox, Rean Griffith, Anthony D. Joseph, Randy H. Katz, Andrew Konwinski, Gunho Lee, David A. Patterson, Ariel Rabkin, Ion Stoica, and Matei Za-

- haria**, “Above the Clouds: A Berkeley View of Cloud Computing,” Technical Report, UC Berkeley 2009.
- Athey, Susan and Glenn Ellison**, “Position Auctions with Consumer Search,” Technical Report, Harvard 2007.
- Baker, George and Thomas N. Hubbard**, “Contractibility and Asset Ownership: On-Board Computers and Governance in US Trucking,” *Quarterly Journal of Economics*, 2004, *119* (2), 1443–1479.
- Battelle, John**, *The Search*, Portfolio Hardcover, 2005.
- Bresnahan, Timothy**, “General Purpose Technologies,” in Bronwyn Hall and Nathan Rosenberg, eds., *Handbook of the Economics of Innovation*, Elsevier-North Holland, 2010.
- and **M. Trajtenberg**, “General Purpose Technologies: Engines of Growth?,” *Journal of Econometrics*, 1995, *65* (1), 83–108.
- Bunkley, Nick and Bill Vlasic**, “Ensuring Junior Goes for a Mild Ride,” *New York Times*, October 7 2008.
- Castle, Jennifer L. and David Hendry**, “Nowcasting from Disaggregates in the Face of Location Shifts,” Technical Report, Oxford University 2009.
- Choi, Hyunyoung and Hal R. Varian**, “Predicting Initial Claims for Unemployment Benefits,” Technical Report, Google 2009.
- and —, “Predicting the Present with Google Trends,” Technical Report, Google 2009.
- Clements, M. P. and David F. Hendry**, “Forecasting in the National Accounts at the Office for National Statistics,” Technical Report, Statistics Commission 2003.
- Dana, James D. and Kathryn E. Spier**, “Revenue Sharing and Vertical Control in the Video Rental Industry,” *Quarterly Journal of Economics*, 2001, *XLIX* (3), 223–245.

- David, Paul**, “The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox,” *American Economic Review*, May 1990, 80 (2), 355–61.
- , “Computer and the Dynamo: The Modern Productivity Paradox in the Not-Too-Distant Mirror,” in “Technology and Productivity: The Challenge for Economic Policy,” OECD, 1991, pp. 315–348.
- , “General Purpose Engines, Investment, and Productivity Growth: From the Dynamo Revolution to the Computer Revolution,” in E. Deiac, E. Hornel, and G. Vickery, eds., *Technology and Investment: Crucial Issues for the 90s*, London: Pinter Publishers, 1991.
- Drucker, Peter F.**, “Knowledge-worker productivity: the biggest challenge,” *California Management Review*, 1999, 41 (2), 79–94.
- Edelman, Benjamin, Michael Ostrovsky, and Michael Schwartz**, “Internet Advertising and the Generalized Second Price Auction,” *American Economic Review*, March 2007, 97 (1), 242–259.
- Farm Foundation**, “Food CPI, Prices, and Expenditures: A Workshop on the Use of Scanner Data in Policy Analysis,” June 2003.
- Feenstra, Robert C. and Matthew Shapiro, eds**, *Scanner Data and Price Indexes*, Vol. 64, NBER and University of Chicago Press, 2003.
- Ford, Henry**, *My Life and Work*, Doubleday, Page & Company, 1923.
- Glassner, Jean-Jacques, Zainab Bahrani, and Marc Van de Miero**, *The invention of cuneiform: writing in Sumer*, Baltimore: Johns Hopkins University Press, 2005.
- Hastie, Trevor, Jerome Friedman, and Robert Tibshirani**, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2 ed., Springer, 2009.
- Hendel, Igal and Aviv Nevo**, “Measuring the Implications of Sales and Consumer Inventory Behavior,” *Econometrica*, 2006, 74 (6), 1637–1673.
- Hirshleifer, Jack**, “The Private and Social Value of Information and the Reward to Inventive Activity,” *American Economic Review*, 1971, 61 (4), 561–574.

- Hoff, Rob**, “Google Search Guru Singhal: We Will Try Outlandish Ideas,” *Business Week*, October 2009.
- Hounshell, David A.**, *From the American System to Mass Production, 1800-1932*, Baltimore: Johns Hopkins University Press, 1984.
- Hubbard, Thomas N.**, “The Demand of Monitoring Technologies: The Case for Trucking,” *Quarterly Journal of Economics*, 2000, 115 (2), 533–560.
- Kohavi, Ron, Roger Longbotham, Dan Sommerfield, and Randal M. Henne**, “Controlled experiments on the web: survey and practical guide,” *Data Mining and Knowledge Discovery*, 2008, 19 (1), 140–181.
- Levy, Steve**, “Secret of Googlenomics: Data-Fueled Recipe Brews Profitability,” *Wired*, 2009, 17 (6).
- Mortimer, Julie H.**, “Vertical Contracts in the Video Rental Industry,” *Review of Economic Studies*, 2008, 75, 165–199.
- Nevo, Aviv and Catherin Wolfram**, “Why Do Manufacturers Issue Coupons? An Empirical Analysis of Breakfast Cereals,” *The RAND Journal of Economics*, 2002, 22 (2), 319–339.
- Nisan, Noam, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, eds**, *Algorithmic Game Theory*, Cambridge, England: Cambridge University Press, 2007.
- Pontin, Jason**, “Artificial Intelligence, With Help From the Humans,” *New York Times*, March 25 2007.
- Press, Associated**, “For Some High-Risk Auto Buyers, Repo Man Is a High-Tech Gadget,” *LA Times*, Nov 18 2006.
- Resnick, Paul and Hal R. Varian**, “Recommender Systems,” *Communications of the Association for Computer Machinery*, March 1997, 3, 56–58.
- Rosenberg, Nathan**, “Technological Change in the Machine Tool Industry,” in “Perspectives in Technology,” Cambridge, UK: Cambridge University Press, 1976, pp. 9–31.

- Sabre**, “Sabre History Website,” November 2009.
- Saxenian, Annalee**, *The New Argonauts: Regional Advantage in a Global Economy*, Cambridge, MA: Harvard University Press, 2006.
- Schumpeter, Joseph A.**, “The Analysis of Economic Change,” in Richard V. Clemence, ed., *Essays on Entrepreneurs, Innovations, Business Cycles and the Evolution of Capitalism*, New Brunswick: Transactions Publishers, 2000, pp. 134–149. Originally published in *Review of Economic Statistics*, May 1935.
- Smith, Barry C., John F. Leimkuhler, and Ross M. Darrow**, “Yield Management at American Airlines,” *Interfaces*, 1992, 22 (1), 8–31.
- Stajano, Frank and Ross Anderson**, “The Cocaine Auction Protocol: On The Power Of Anonymous Broadcast,” in “Proceedings of Information Hiding Workshop” Lecture Notes in Computer Science 1999.
- Talluri, Kalyan T. and Garrett J. van Ryzin**, *The theory and practice of revenue management*, Kluwer Academic Publishers, 2004.
- Ulph, David and Nir Vulkan**, “Electronic Commerce, Price Discrimination, and Mass Customisation,” Technical Report, Said Business School, Oxford 2007. <http://vulkan.worc.ox.ac.uk/wp-content/images/combined-paper.pdf>.
- Usher, Abbott Payson**, *A History of Mechanical Invention*, revised ed., Dover Publications, 1998.
- Varian, Hal**, “Technology Levels the Business Playing Field,” *New York Times*, August 25 2005.
- Varian, Hal R.**, “Economic Mechanism Design for Computerized Agents,” in “USENIX workshop on Electronic Commerce” 1995, pp. 13–21.
- , “Competition and Market Power,” in Joseph Farrell, Carl Shapiro, and Hal R. Varian, eds., *The Economics of Information Technology: an Introduction*, Cambridge University Press, 2005, pp. 1–46.
- , “Position Auctions,” *International Journal of Industrial Organization*, 2007, 25 (7), 1163–1178.

———, “Online Ad Auctions,” *American Economic Review*, 2009, 99 (2), 430–434.

Weitzman, Martin, “Recombinant Growth,” *Quarterly Journal of Economics*, 1998, 113, 331–360.

Wikipedia, “Amazon Mechanical Turk,” *Wikipedia*, 2009.

———, “Cloud Computing,” *Wikipedia*, 2009.

Yates, JoAnne, “Business Use of Information and Technology from 1880–1950,” in Alfred D. Chandler and James Cortada, eds., *A Nation Transformed by Information: How Information Has Shaped the United States from Colonial Times to the Present*, Oxford University Press, 2000, pp. 107–135.