Similarities and Differences in the Adoption of General Purpose Technologies[1*]

Ajay Agrawal, Joshua S. Gans and Avi Goldfarb

Rotman School of Management, University of Toronto and NBER

Ajay Agrawal is the Geoffrey Taber Chair in Entrepreneurship and Innovation and Professor of Strategic Management, Joshua S. Gans is Jeffrey Skoll Chair in Technical Innovation and Entrepreneurship and Professor of Strategic Management, and Avi Goldfarb is Rotman Chair in Artificial Intelligence and Healthcare and Professor of Marketing, all at the Rotman School of Management, University of Toronto, Toronto, Canada. Their email addresses are ajay.agrawal@rotman.utoronto.ca, joshua.gans@gmail.com, and agoldfarb@rotman.utoronto.ca.

[1*] We thank the editors for helpful comments and Yan Zhou for research assistance.

Economists recognize the central role of productivity in economic growth. At the same time, economic models provide little insight into when the next big idea and its associated productivity dividend will come along. Once an important technology is identified, however, the next steps are more within the domain of economic analysis. The economist's toolkit can provide an understanding when firms will adopt a new technology and for what purpose.

Bresnahan and Trajtenberg (1995) emphasize one category of important technologies, which they label "General Purpose Technologies" or GPTs. These technologies have the potential to generate sustained productivity growth through a positive feedback loop of innovation in producing and using industries. They note that, "[m]ost GPTs play the role of 'enabling technologies', opening up new opportunities rather than offering complete, final solutions. For example, the productivity gains associated with the introduction of electric motors in manufacturing were not limited to a reduction in energy costs."

Since Bresnahan and Trajtenberg's article appeared, a growing literature has explored the diffusion and impact of general purpose technologies. Some of this literature takes a historical perspective, for example examining how the steam engine and electricity diffused (e.g. Lipsey, Carlaw, and Bekar 2005; Moser and Nicholas 2004). Much of the literature focuses on more recent technologies, particularly computers, the internet, and artificial intelligence (e.g. Bresnahan, Brynjolfsson, and Hitt 2002; Bresnahan and Hitt 2003; Forman, Goldfarb, and Greenstein 2005; Cockburn, Henderson, and Stern 2019).

Regardless of the technology studied, the focus of the literature has been on commonalities across each type of GPT. This focus is natural, given that the goal of the literature has been to identify generalizable insights across technologies. The main commonality has been the role of complementary innovation, or co-invention. For example, Bresnahan and Greenstein (1996) examined the diffusion of client/server computing systems in large companies. They emphasize the role of co-invention, the invention of new technologies and processes that enable the technology to generate growth. Looking at a different information technology, enterprise resource planning software, Aral, Brynjolfsson, and Wu (2012) emphasize complementary innovations related to organizational structures. Broadly, this literature emphasizes heterogeneity in co-invention costs. Organizations with lower co-invention costs, whether because of internal or external advantages, are more likely to adopt.

Each GPT, however, provides a distinct benefit. Steam provided a new power source. The internet facilitated communication. Recent advances in artificial intelligence reduce the cost of prediction (Agrawal, Gans, and Goldfarb 2018). The differences between GPTs are important for understanding adoption patterns. Co-invention costs do matter, but so do the benefits. The distinct benefits of GPTs determine why a particular GPT might be useful to a particular firm. Therefore, in order to understand how GPTs diffuse, and how they impact economic growth, it is necessary to consider the barriers to adoption common to all GPTs, as related to co-invention costs. It is also necessary to consider the benefits from adoption, which tend to be different with each GPT. These benefits often define the particular type of co-invention needed, as well as providing an understanding of which firms and industries will gain most from adoption.

Using the examples of the internet and artificial intelligence, we will discuss how both co-invention costs and distinct benefits determine the adoption of technology. The discussion of the internet will be grounded in the empirical literature on internet adoption, particularly a re-framing of Forman, Goldfarb, and Greenstein's (2005) *Journal of Urban Economics* paper. The discussion of AI will build on a theory literature and an extrapolation of Mullainathan and Obermeyer's (2021) study of AI in emergency departments. For both technologies, we demonstrate that discussions of the impact of a GPT on productivity and growth need to emphasize the benefits as well as the costs. The goal of this paper is therefore to link the literature on co-invention costs with an understanding of the distinct benefits of each GPT.

General Purpose Technologies

There has been a renewed interest in GPTs over the last five years. Figure 1 tracks citations to Bresnahan and Trajtenberg (1995) over time. Citations grew steadily until 2010, flattened through 2016, and then grew again. Many of these more recent citations discuss the potential of artificial intelligence to be a GPT.

Figure 1: Citations to Bresnahan & Trajtenberg (1996) in Web of Science

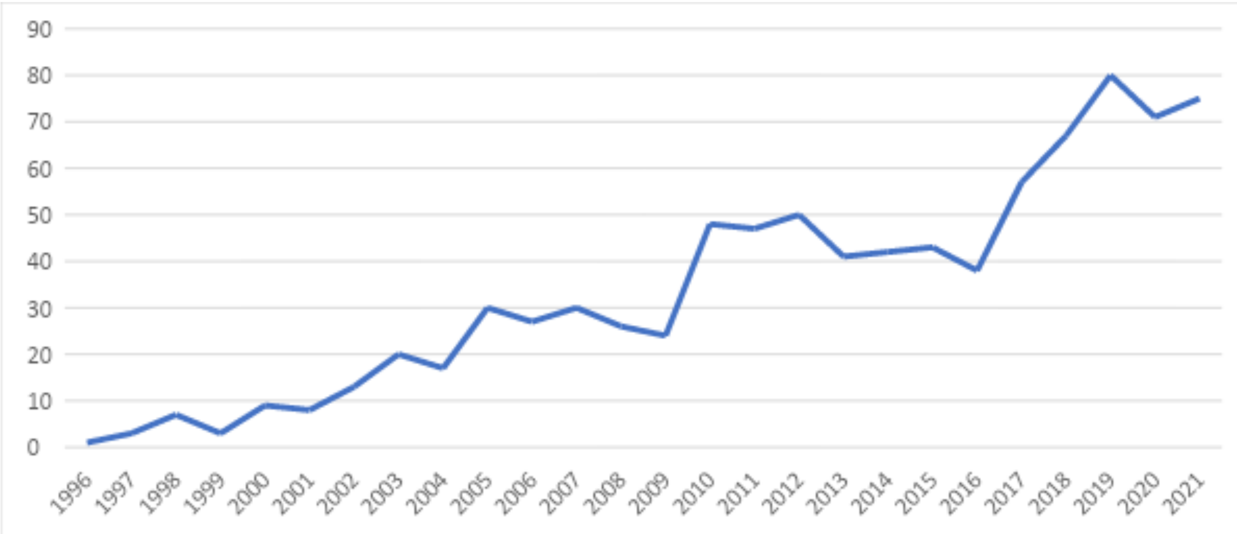

Bresnahan and Trajtenberg (1995) note that "[w]hole eras of technical progress and growth appear to be driven by a few 'General Purpose Technologies.'" These GPTs generate follow-on innovations through what they label "innovational complementaries". This perspective mirrors a larger literature on the economics of technology that emphasizes the importance of a handful of innovations in driving economic growth. Such innovations have received different labels, with varying definitions, elsewhere. Mokyr (1990) emphasized the role of "macroinventions" in generating the potential for growth and innovation, and the role of follow-on "microinventions" in catalyzing the growth. This literature builds on studies that examine technological change in specific industries, such as David's (1990) examination of electrification and the reorganization of factories and Rosenberg's (1963) documentation of the breadth of innovations underlying the productivity increases in the 19th century US machine tools industry. Here we emphasize the particular framework provided in Bresnahan and Trajtenberg (1995) and Bresnahan's (2010) review paper.

GPTs are defined by their “potential for pervasive use in a wide range of sectors and by their technological dynamism” (Bresnahan and Trajtenberg 1995). The technological dynamism is not limited to the industry or sector that produces the GPT. GPTs have an outsized impact on long term economic growth because they also generate innovation in using, or application, industries. The initial innovation in one sector generates productivity in the R&D in a downstream sector. This, in turn, generates further innovation in the upstream producing industry. This creates a positive feedback loop that pushes against decreasing returns to innovation. This positive feedback loop between producing and using industries defines GPTs as distinct from other useful technologies.

Bresnahan and Trajtenberg attribute the increase in downstream R&D productivity to “innovational complementarities.” These complementarities generate what Bresnahan and Greenstein (1996) label “co-invention”. Gans (1995) and Aral, Brynjolfsson, and Wu (2012) note that these innovations can also occur in management and organization, and so they label them “organizational complementarities.”

An empirical literature on information technology adoption has emphasized the costs associated with co-invention. Bresnahan, Brynjolfsson, and Hitt (2002) demonstrate that it takes several years for the productivity benefits of information technology to appear. They argue that this delay is caused by the challenges in developing new processes. Forman, Goldfarb, and Greenstein (2008) show that early adoption of advanced internet technologies was more likely in large firms or in large cities. They argue that this was driven by differences in co-invention costs. It is easier to undertake co-invention in firms with more resources, and in cities with more expertise. Dranove et al (2014) further explore this hypothesis in the context of electronic medical records. They find that larger hospitals, and hospitals in cities with a large pool of health information technology workers, are more likely to gain from adopting electronic medical records, as measured by a reduction in hospital costs. As with Bresnahan, Brynjolfsson, and Hitt (2002), this cost reduction appeared after several years. Dranove et al (2014) also show that many adopters fail. Even six years after adoption, hospitals that adopted but did not have access to a large local pool of health information technology workers, experienced a substantial increase in costs with little benefit. Brynjolfsson, Rock, and Syverson (2021) take a more macroeconomic approach. They measure the relationship between IT investments and firm productivity in a cross-industry firm-level panel data set. They demonstrate a "productivity J-curve" in which measured productivity falls in the early years after IT investments but eventually rises. Their evidence suggests it is investments in intangible assets that lead to the short-term decrease in measured productivity. These investments generate complementary innovation.

Broadly, this empirical literature on IT adoption and its consequences has emphasized the common thread of co-invention costs leading to a lag in the productivity benefit to adoption and higher levels of adoption in firms that could adapt their processes efficiently.

The empirical literature on GPTs has emphasized commonalities in the costs of adopting different GPTs. In contrast, there has been little discussion of the benefits. That might be because the benefits are, at the surface, different across GPTs. Steam engines and electricity enabled new power sources. Interchangeable parts enabled production at scale. Railroads and the internal combustion engine enabled new forms of transportation. The internet enabled new forms of communication. In each case, these benefits are different.

Models simplifying these benefits provide insight into which firms will benefit the most. There is literature on new technologies, separate from the GPT literature, that emphasizes how technology reduces economic frictions. Shapiro and Varian (1999) emphasized the internet as a technology for cheap communication, reducing the costs of copying and of search. Goldfarb and Tucker (2019) noted how the internet also reduced the cost of transporting information and of tracking behavior. More broadly, Nordhaus (2007) linked computing advances to a drop in the cost of computation or arithmetic. For AI, Agrawal, Gans, and Goldfarb (2018) emphasized that it can be seen as a drop in the cost of prediction.[2]

Adding this perspective on a common model of the benefits of a GPT provides insight into which firms will adopt early and experience a productivity gain. It also provides insight into the types of co-invention needed. Put differently, co-invention is necessary, and is the source of the positive feedback loop that generates sustained productivity growth. Nevertheless, that productivity growth depends on adoption and co-invention in the industries that use the technology. As a technology makes a particular process cheaper, whether computation, search, or prediction, then that process will be used more broadly. How much more broadly depends on the elasticity of demand. Thomas J. Watson, then-CEO of IBM, allegedly declared global demand for five computers. That implies

[2] This framing can apply to a variety of technologies, even those that are not widely seen as GPTs. For example, for blockchain, Catalini and Gans (2020) saw it through the lens of a fall in verification costs.

a very inelastic model of demand. As computers reduced the price of machine arithmetic, the number of applications grew rapidly.

In the remainder of this paper, we attempt to connect these two streams of literature. This attempt builds on work in economic history that emphasizes the particular benefits of a technology along with the challenges in co-invention. For example, David (1990) discusses how electrification provided a new way to power factories, but substantial co-invention was required for the new way of providing power to affect aggregate productivity. Similarly, Mokyr (1990, p. 82) emphasizes that during the Industrial Revolution, "a clustering of macroinventions occurred…thus creating a complementary flow of microinventions." For the macroinvention of the steam engine, Mokyr (1990) describes several microinventions. These microinventions were in the producing industry, making the engines more efficient. They also occurred downstream, applying the engines to provide power in new industries such as shipping (e.g. Fulton). Mokyr details innovation in a competing power source, water, for comparison. Understanding the diffusion and impact of the technology requires understanding the challenges to innovation, but also the benefits.

We use examples from relatively recent technologies in order to demonstrate how the impact of a GPT is best understood if the costs of co-invention and the benefits with respect to reducing a particular economic friction are taken into account. To do so, we first examine the diffusion of the commercial internet. Specifically, we re-interpret Forman, Goldfarb, and Greenstein's (2005) paper in internet adoption using this framework. We demonstrate how an emphasis on both co-invention costs and the benefits of the technology provides a useful lens for interpreting the results. Second, we describe emerging research on the usefulness of artificial intelligence. We combine the theoretical framework in Bresnahan (2020) and in Agrawal, Gans, and Goldfarb (2021) on AI systems with the empirical setting in Mullainathan and Obermeyer (2021). In doing so, we demonstrate a feedback loop between the benefits of prediction technology and the need for a particular type of co-invention.

Internet

Forman, Goldfarb, and Greenstein (2005) examine internet adoption patterns by 86,879 US business establishments, measured in the year 2000. The paper takes an economic geography perspective, examining whether the internet was adopted more by rural or urban business establishments. The main result of the paper is that establishments in larger cities were more likely to adopt relatively advanced internet technologies such as enterprise resource planning (ERP) and customer relationship management (CRM), but they were somewhat less likely to adopt the basic technologies such as email and web browsing.

The paper emphasizes that the result for advanced internet—that urban firms are more likely to adopt—supports the hypothesis that co-invention costs are important drivers of adoption. The other main result, that basic internet is more likely to be adopted in rural areas and smaller cities, received less focus in the paper and in the literature that followed. This is the result about the benefit of the technology, that the internet reduced communication costs and therefore has a distinct benefit for rural firms compared to urban firms. The authors wrote two additional papers with the same dataset that only looked at the advanced internet, emphasizing co-invention cost differences in urban and rural areas (Forman, Goldfarb, and Greenstein 2008; Forman, Goldfarb, and Greenstein 2012).

Table 7 (replicated below) of the paper provides the most important results for understanding the difference between co-invention costs and the particular benefits of internet technology for facilitating communication. To do so, the paper further categorizes basic and advanced internet into those that involve communication outside the establishment and those that facilitate only within-establishment communication. This leads to four broad types of internet technologies:

1) Basic cross-establishment internet: These are applications of the internet that involve communication between establishments and do not require much expertise to use. Examples include email and web browsing.

2) Advanced cross-establishment internet: These are applications that involve communication between establishments and require considerable expertise to implement. Examples include using the internet for commercial transactions between firms, or between firms and customers.

3) Basic within-establishment internet: These are applications that primarily involve communication within an establishment and that do not require much expertise to use. Examples include internal web pages and networking services run through an intranet.

4) Advanced within-establishment internet: These are applications that primarily involve communication within an establishment and require considerable expertise to implement. Examples include ERP and CRM systems.

Basic cross-establishment internet is more likely to be adopted in less populated areas (panel A columns 4-6). Advanced within-establishment internet is more likely to be adopted in bigger cities (panel C columns 1-3). The other two categories are in the middle, with no consistent significant difference between urban and rural firms for advanced cross-establishment internet and basic within-establishment internet. Together, as noted in the paper, these results are consistent with evidence of co-invention: advanced technologies are more likely to be adopted in urban areas. These results also show that less populated areas are more likely to adopt cross-establishment internet, conditional on whether the technology is basic or advanced. Even for advanced technologies, small cities are as likely to adopt as large ones. Adoption is determined by more than co-invention costs. It is also determined by who benefits most from being able to communicate more efficiently across establishments.

Table 7
Effect of population size and density on adoption of WEI firm and CEI adoption (standard errors in parentheses)

	WEI			CEI		
	(1)	(2)	(3)	(4)	(5)	(6)
Participation						
A. Coefficients from (weighted) probit regressions						
Small MSA	−0.000307	−0.00472	0.0318	0.019097	0.0209	0.0334
	(0.00222)	(0.0173)	(0.0294)	(0.0317)	(0.0200)	(0.0417)
Medium MSA	0.0138	0.00497	0.0303	−0.0356	−0.0294	−0.0244
	(0.0160)	(0.0138)	(0.0215)	(0.0200)+	(0.0153)+	(0.0260)
Large MSA	−0.00678	−0.00747	−0.0013	−0.0492	−0.0208	−0.0572
	(0.0127)	(0.0119)	(0.0151)	(0.0169)**	(0.0133)	(0.0238)*
Log likelihood	−30431.8	−52357.1	−17680.6	−15356.3	−34342.7	−8625.6
Observations	53231	86879	30260	53231	86879	30119
B. Marginal effects from (weighted) probit regressions						
Small MSA	−0.000102	−0.00181	0.0107	0.00280	0.00489	0.0048
	(0.00736)	(0.00661)	(0.00978)	(0.00459)	(0.00462)	(0.00587)
Medium MSA	0.00457	0.00190	0.0102	−0.00536	−0.00700	−0.0036
	(0.00528)	(0.00527)	(0.00717)	(0.00308)+	(0.00369)+	(0.00393)
Large MSA	−0.00225	−0.00285	−0.000431	−0.00727	−0.00490	−0.0084
	(0.00423)	(0.00457)	(0.00513)	(0.00256)**	(0.00316)	(0.00361)**
Enhancement						
C. Coefficients from (weighted) probit regressions						
Small MSA	0.0801	0.0724	0.114	0.0330	0.0379	0.0183
	(0.0236)**	(0.0222)**	(0.0330)**	(0.0194)+	(0.0159)*	(0.0268)
Medium MSA	0.0841	0.0931	0.0992	0.0462	0.0172	0.0370
	(0.0163)**	(0.0154)**	(0.0230)**	(0.0174)**	(0.0131)	(0.0192)+
Large MSA	0.130	0.147	0.125	0.0334	0.00703	0.0013
	(0.0135)**	(0.0125)**	(0.0164)**	(0.0109)**	(0.00920)	(0.0138)
Log likelihood	−24608.8	−27269.8	−12499.4	−29910.4	−46272.2	−16510.5
Observations	53227	86872	30260	53227	86872	30265
D. Marginal effects from (weighted) probit regressions						
Small MSA	0.0217	0.0116	0.0274	0.0108	0.0117	0.00570
	(0.00662)**	(0.00372)**	(0.00841)**	(0.00642)+	(0.00498)*	(0.00848)
Medium MSA	0.0226	0.0148	0.0235	0.0151	0.00530	0.0116
	(0.00456)**	(0.00261)**	(0.00570)**	(0.00577)**	(0.00404)	(.00611)+
Large MSA	0.0340	0.0224	0.0285	0.0108	0.00215	0.0004
	(0.00367)**	(0.00208)**	(0.00393)**	(0.00356)**	(0.00282)	(0.00433)

Notes: All regressions include dummy variables for three-digits NAICS, whether it was a multi-establishment firm, employment and employment squared as controls. Robust standard errors, clustered by MSA, are in parentheses. Non-MSA is the base for these regressions.

(1) & (4) include only establishments that received the supplementary Harte Hanks survey.

(2) & (5) include entire sample.

(3) & (6) include only establishments from single-establishment firms who received the supplementary Harte Hanks survey.

\+ Significant at the 90% confidence level.

* Idem., 95%.

** Idem., 99%.

Taken together, these results demonstrate that the internet is a GPT, but of a particular type. There is evidence that adoption is more likely in cities, but there is also evidence that the applications most related to communication technology were also likely to be adopted in rural areas.

This gives a hint of what kind of co-invention was necessary and which types of firms benefited most. While co-invention is easier in cities, a key benefit of the technology is cross-firm communication. Several papers by Chris Forman examine the co-invention needed for this cross-firm communication, although the papers do not frame it as such. Forman and Gron (2011) shows that internet adoption was faster for insurers that were vertically integrated with their agents, and therefore non-vertically integrated adopters need to adapt their systems to avoid "opportunistic behavior", such as agents selling competitor products to consumers acquired through an insurer-funded electronic commerce system. The benefit of the communication technology is clear, but the incentives to adopt may not be aligned. Langer, Forman, Kekre, and Sun (2012) show that the successful addition of an internet distribution channel requires a change in product offerings that appeal to the more diverse set of customers that become feasible through digital communication. Forman and Van Zeebroeck (2012, 2019) examine within-firm knowledge flows after internet adoption. They show that internet connections increase collaboration and citations across inventors at different locations within a firm, but primarily for researchers with overlapping interests. They then suggest changes to the geography of the research group within the firm.

Each of these papers documents how the internet changed communication patterns. This, in turn, meant that firms needed to change their processes to take advantage of the change. While not labeled as such, this is co-invention. Put differently, the literature on internet adoption by businesses has emphasized co-invention costs (as in Forman, Goldfarb, and Greenstein's papers) and it has discussed how easier communication benefits certain types of businesses (as in Forman's other work). Our goal in this paper is to connect the two and point out that understanding co-invention and how a GPT might have an impact on productivity and growth requires an understanding of the benefits of the particular technology.

Artificial Intelligence

We proceed by assuming that AI is a GPT (whether part of a broader GPT or distinct), in the sense that AI technologies are defined by their "potential for pervasive use in a wide range of sectors and by their technological dynamism" (Bresnahan and Trajtenberg 1995). There has also been substantial innovation in both AI-producing and AI-using industries over the past ten years, suggesting potential for a positive feedback loop. AI patenting has growth rapidly since 2010 (Bloom et al 2018). AI is frequently mentioned in earnings calls and job postings in a variety of industries (Bloom et al 2021; Goldfarb, Taska, and Teodoridis 2021).[3]

In our prior work (Agrawal, Gans, and Goldfarb 2018, 2019), we have emphasized that it is useful to see AI as prediction technology, and so advances in AI can be understood as reductions in the cost of prediction, where prediction is defined in the statistical sense. AI makes it easier to fill in

[3] Artificial intelligence technologies are still diffusing, and it remains an open question as to whether they represent a new GPT (Agrawal, Gans, and Goldfarb 2019), an extension of the information technology GPT that has been diffusing since the 1950s, an aspect of data science as a GPT, or not a GPT at all (Goldfarb, Taska, and Teodoridis 2021).

missing information. Just as the internet is communication technology, steam provides an energy source, and the railway is transportation technology, AI enables prediction.

AI has been developed to address a number of prediction problems. It underlies Google's search engine, predicting which website a user wants in response to a query. It underlies Facebook's news feed, predicting which information a user is most likely to engage with. It underlies Amazon's recommendation engine, predicting which products people are likely to purchase. It is also an important component of business processes in a variety of other industries. It helps banks, credit card companies, and other financial institutions detect fraud. It helps insurance companies determine risks. It helps healthcare providers with diagnosis.

Bresnahan (2020) argues that only the information technology companies have received a substantial productivity benefit from AI so far. Google, Facebook, and Amazon had processes that could adapt to the opportunities in AI relatively quickly. They had data tools in place to enable good prediction, and they had the ability to adapt workflows to accommodate the predictions. Bresnahan emphasizes that this requires a system-level approach to using AI. A system-level approach involves changing workflows, adding new products, and providing new forms of value to customers. In other words, co-invention around the technology was easier. Therefore, IT companies adopted the technology early and with success.

In contrast, many other industries that have adopted AI have found relatively narrow applications within existing workflows. Financial services have used AI for fraud detection. As a result, fraud detection has improved, but the overall workflow has changed little. There has been relatively little innovation, either in terms of algorithms or organizational structure.

In Agrawal, Gans, and Goldfarb (2021), we formalized Bresnahan's insight, focusing on AI as a prediction technology. Prediction is useful because it improves decision-making. A system of decisions can involve multiple decision-makers who may have different information and whose incentives may not be aligned. Focusing on differences in information, our model showed that a productivity-enhancing AI might not be adopted if coordination between decision-makers is difficult. In other words, as a prediction technology, AI helps with decision-making. When there are multiple decisions and multiple decision-makers, AI can face resistance due to alignment. Even if it improves some decisions, if they cannot be coordinated with other decisions, then adopting AI could make things worse. We, therefore, hypothesize that co-invention for AI will require the development of systems to overcome coordination challenges.

Coordination of decisions in medical care: An application of AI as a GPT

In this section, we provide an example of an AI that improves medical diagnosis and discuss the necessary coordination of decisions that would lead to a substantial productivity boost. Specifically, we take the AI for heart attack diagnosis in Mullainathan and Obermeyer (2021) and examine what changes the hospital would need to undertake in order to get substantial value out of the tool.

When a patient arrives in the emergency department of a hospital, the medical staff assesses the likelihood of various ailments and recommends further testing. Mullainathan and Obermeyer focus on physician decisions to recommend testing for a heart attack. This testing can be invasive and costly. At the same time, a missed heart attack has severe consequences for the patient, as well as bottom line implications for the hospital. Using information available to physicians at the time they make a decision about testing, Mullainathan and Obermeyer developed an AI that performed substantially better than the doctors. The AI is built from data on 246,265 emergency visits at a large top-ranked hospital. Approximately 15% of patients required treatment. The paper documents that the AI predicted better than doctors. Doctors sometimes overtested, providing invasive tests to patients who were unlikely to be having a heart attack. Doctors also undertested, failing to order tests to many patients who were likely to need treatment. If the number of tests were fixed, the AI could re-allocate them from low risk to high risk patients, leading to better outcomes. It could also achieve the same outcomes with fewer tests. The focus of their paper as published is on understanding the nature of physician mistakes, showing that physicians use an overly simplistic model of the heart attacks.

The paper also contains an online appendix that explores how hospitals might change their processes in response to the algorithm. Next, we take the results of that appendix and build on them to model the potential for disagreement between physicians and hospital administrators. Specifically, in Appendix 3, they discuss two different tests for heart attacks. The definitive test is cardiac catheterization. This involves a cardiologist inserting an instrument into the coronary arteries. If a blockage is found, a stent is inserted during the same procedure. Alternatively, physicians can recommend that the patient first receive a stress test. These tests are less expensive and less invasive, but also less accurate. Furthermore, if positive, the patient still needs to have the catheterization in order to insert the stent.

While the main paper does not distinguish between stress tests and testing through catheterization, it is an important decision that can cost the hospital thousands of dollars and can affect the patient's long-term outcomes. They estimate that the cost of catheterization is about $28,000 while the cost of a stress test is about $4,000. Patient outcomes and hospital profits are largely aligned. If hospitals identify patients who are having a heart attack, then the patient's life is saved and the hospital generates substantial revenue from caring for the patient.[4] Mullainathan and Obermeyer (appendix page 6) note, "it is intuitive and efficient for physicians to begin with stress testing for lower-risk patients…Higher-risk patients, on the other hand, should go straight to catheterization without incurring the additional cost and likelihood of false negative from stress testing."

Currently, stress testing serves a useful purpose. While Mullainathan and Obermeyer note that there is an active medical debate about whether stress tests should be used at all, stress tests continue to be used. They give doctors an outlet when they are highly uncertain whether a patient is at risk.

The question we explore next is what happens as the AI prediction gets better. Better prediction means less uncertainty. The risk will be more accurately measured. More patients will be clearly high risk, and more patients will be clearly such low risk that even a stress test is not cost effective. The simulations in Mullainathan and Obermeyer use their predictions of risk and demonstrate that stress tests should be used less when those predictions are more accurate.

[4] Details in the appendix are supplemented by an interview of Ziad Obermeyer on January 25, 2022.

Health economics often emphasizes differences in incentives between different players in the healthcare system, whether patients, doctors, insurers, or hospital administrators. For example, Bhattacharya, Hyde, and Tu's (2014) textbook, *Health Economics*, discusses the model in Harris (1977) that hospitals should be thought of as two separate economic actors: physicians and administration. The physicians decide on testing and diagnosis for each patient. The administration decides on which inputs should be available to physicians. With this division of decision-making, doctors have little incentive to control hospital costs (p. 103). In terms of testing, in other contexts widespread screening leads to better patient outcomes but at higher cost (p. 281). Doctors and administrators may be at odds as to which tests to do.

Table 2 provides back-of-the-envelope estimates of the benefits of allowing stress testing to doctors and to administrators with today's 15% rate of diagnosis accuracy upon arrival into the emergency department, an improved AI with 50% accuracy, and an almost-perfect AI with 99% accuracy. While loosely based on the values in Mullainathan and Obermeyer's appendix and elsewhere in the medical literature, these numbers are for demonstration purposes only.[5]

[5] Estimates of revenue and profit vary. We use $100,000 in revenue and $20,000 in margin, based on Urbich et al's (2020) two-year revenue estimates meta-analysis and a 20% margin (RANGE 20k to 200k).

Table 2

AI accuracy if positive	**1%**	**15%**	**50%**	**99%**
Doctor payoff				
Send patient home	$148,500	$127,500	$75,000	$1,500
Stress test first	$140,101	$135,915	$125,450	$110,799
Direct to catheterization	$122,000	$122,000	$122,000	$122,000
Administrative payoff				
Send patient home	$0	$0	$0	$0
Stress test first	$-892.15	$1,807.75	$8,557.5	$18,007.15
Direct to catheterization	$-1285	$1,725	$9250	$19,785

Calculations: Doctor: Benefit to doctor/patient of a healthy patient $150,000; cost of stress test to doctor $4,000; Cost of catheterization $28,000; false positive stress test 20%; false negative stress test 5%. Administration: Profit from detecting heart attack $20,000; profit from sending home $0; loss from stress test $800; loss from unnecessary catheterization $1,500; false positive stress test 20%; false negative stress test 5%. Numbers are back-of-the envelope and are not meant to represent accurate doctor or administration values.

The second column shows that under a reasonable set of assumptions, doctors and administration are currently aligned. If a patient is predicted as having a 15% chance of a heart attack, both doctors and administrators will first send that patient for a stress test. If that stress test is positive, both doctors and administrators will move to catheterization.

The first and fourth columns show that the doctors and administration will also be aligned with a very accurate AI. The first column shows that if a patient has a 1% chance of having a heart attack or less, then the doctors and administrators agree that the patient should be sent home untested. The downside of the test is too high relative to the benefit of detecting the heart attack. The fourth column shows that if a patient is predicted as 99% likely to be having a heart attack, both doctors and administrators agree that the stress test should be skipped and the patient should go direct to catheterization.

However, the third column shows that if a patient is predicted to have a 50% chance of having a heart attack, then doctors and administrators might disagree. The source of disagreement relates to different estimates of the downside of an unnecessary catheterization, and to different estimates of the cost of undertaking a stress test. We have assumed that administrators perceive a larger cost to unnecessary testing than doctors.

Given the allocation of decision rights, as long as stress tests are available, doctors will choose to do a stress test first if they predict the patient has a 50% chance of having a heart attack. While the administration would prefer that the patient go straight to catheterization, the decision is the doctor's. The administrator, however, has the choice of whether to allow stress testing at all. If the administration provides a doctor with an AI that allocates patients to, say, either a 0.1% chance of having a heart attack or a 50% chance of having a heart attack, the doctor will then provide stress tests to anyone that the AI predicts at 50%. The administration should then decide that stress tests are not an option. And so the doctor would need to either send patients home, or to catheterization. The doctor will choose catheterization.

As modeled, the solution is straightforward. The administration should not offer the stress test. The doctor will choose catheterization. As practiced, however, it is unlikely to be so straightforward. The doctor will push back. A regulatory body may be called in. Patients' rights will be discussed. When the decision-makers are no longer aligned, then the structure of the game as modeled may change. In other words, the system as described has administration making a take-it-or-leave-it offer to the doctors, and then the doctors making the best decision given the circumstances. The doctors, however, can change the game. If the administration wants the doctors on board, it may need to change how doctors perceive the payoffs of unnecessary testing. This might involve new billing procedures, different liability protection, or other changes that suggest a change to the way the hospital operates.

This, in turn, might make the administrator hesitant to adopt the AI in the first place. The AI changes the decision alignment. The organization then needs to be more specific about its objective function and coordinate on whose payoffs determine the decision. Even though the AI improves patient outcomes, it creates conflict. Depending on the costs of that conflict, the administration might decide to forego the benefits of better prediction.[6] Predictions can change decision alignment, and this lack of alignment becomes a barrier to AI adoption.[7]

We provide this example to demonstrate how AI can improve outcomes, but the specific nature of AI as prediction technology also generates challenges to its adoption. As noted in the previous section, our goal in this paper is to connect the literature on co-invention to an understanding of the particular benefits of each GPT. For AI, these are connected. The benefits relate to improved decision-making. Most organizations, however, have multiple decision-makers. The technology can lead to misalignment. Successful adoption may require negotiation, new processes, and a change in how payoffs are allocated.

[6] Another point is that if the AI isn't deployed early, then it may never gather enough data in order to improve the predictions to the point where the doctors and the administration are aligned give the predictions.

[7] Gans (2022) provides a general treatment of when internal conflicts can make the adoption of radical technologies more difficult.

Conclusion

This paper has connected two different streams in the literature on the diffusion and consequences of GPTs. One stream focuses on the barriers to adoption, and the need for co-invention to enable productivity growth. A second stream focuses on the benefits of each generation of technology. Using the examples of the internet and artificial intelligence, we show that understanding the consequences of GPT diffusion, and its associated challenges, requires understanding commonalities with prior GPTs in terms of co-invention, while recognizing that each technology brings with it specific benefits that suggest which firms will adopt and the direction of the necessary co-invention.

References

Agrawal, Ajay, Joshua Gans, Avi Goldfarb. 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Review Press, Boston MA.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2019. Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction. *Journal of Economic Perspectives,* 33(2): 31-50.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2021. AI Adoption and System-Wide Change. *National Bureau of Economic Research, DOI:* 10.3386/w28811.

Aral, Sinan, Erik Brynjolfsson, Lynn Wu. 2012. Three-Way Complementarities: Performance Pay, Human Resource Analytics, and Information Technology. *Management Science* 58(5), 913-931.

Bhattacharya, Hyde, and Tu 2014. *Health Economics.* Red Globe Press.

Bresnahan, T. 2020. Artificial intelligence technologies and aggregate growth prospects. Working paper, Stanford University.

Bresnahan, Timothy. 2010. General Purpose Technologies. In Bronwyn Hall, Nathan Rosenberg Eds. *Handbook of the Economics of Innovation*. Chapter 18, 761-791.

Bresnahan, Timothy, Erik Brynjolfsson, Lorin Hitt. 2002. Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence, *The Quarterly Journal of Economics*, 117(1), 339–376.

Bresnahan, T., S. Greenstein. 1996. Technical Progress and Co-invention in Computing and in the Uses of Computers. *Brookings Papers on Economic Activity, Microeconomics* 1996; 1-83.

Bresnahan, T., M. Trajtenberg. 1995. General Purpose Technologies 'Engines of Growth'? *Journal of Econometrics*. 65, 83-108.

Brynjolfsson, Erik, and Lorin M. Hitt. 2003. "Computing Productivity: Firm-Level Evidence." *Review of Economics and Statistics* 85 (4): 793–808.

Brynjolfsson, Erik, Daniel Rock, Chad Syverson. 2019. Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics. In Agrawal, Gans, Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press, 23-57.

Brynjolfsson, Erik, Daniel Rock, Chad Syverson. 2021. Productivity J Curve. AEJ Macroeconomics.

Catalini, Christian and Joshua S. Gans (2020), "Some Simple Economics of the Blockchain," *Communications of the ACM*, 63 (7), pp.80-90.

Cockburn, Iain, Rebecca Henderson, Scott Stern. 2019. The Impact of Artificial Intelligence on Innovation. In Agrawal, Gans, Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press, 115-146.

David, Paul. 1990. The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox. *American Economic Review Papers & Proceedings* 80(2), 355-361.

Forman, Chris, Avi Goldfarb, and Shane Greenstein. 2005. How Did Location Affect Adoption of the Commercial Internet: Global Village vs. Urban Leadership. *Journal of Urban Economics,* 58 (3): 389-420.

Forman, Chris, Avi Goldfarb, and Shane Greenstein. 2008. Understanding the Inputs into Innovation: Do Cities Substitute for Internal Firm Resources? *Journal of Economics and Management Strategy*, 17 (2).

Forman, Chris, Avi Goldfarb, and Shane Greenstein. 2012. The Internet and Local Wages: A Puzzle. *American Economic Review,* 102 (1): 556-575.

Forman, Chris, and Anne Gron. 2011. Vertical Integration and Information Technology Investment in the Insurance Industry, *The Journal of Law, Economics, and Organization* 27(1), 180–218.

Forman, C. and N. van Zeebroeck. 2012. From wires to partners: How the internet has fostered R&D collaborations within firms. *Management Science* 58 (8), 1549–1568.

Forman, C., and N. van Zeebroeck. 2019. Digital technology adoption and knowledge flows within firms: Can the Internet overcome geographic and technological distance? *Research Policy* 48 (8): 103697.

Gans, Joshua S. 1995. "Inside the Black Box: A Look at the Container," *Prometheus*, 13 (2) pp.169-183.

Gans, Joshua S. 2022. "A Theory of Visionary Disruption," *mimeo*., Toronto.

Goldfarb, Avi, Bledi Taska, and Florenta Teodoridis. 2021. "Could Machine Learning Be a General Purpose Technology? Evidence from Online Job Postings," Working paper, University of Toronto.

Goldfarb, Avi, and Catherine Tucker. 2019. "Digital Economics." *Journal of Economic Literature* 57(1): 3-43.

Jovanovic, Boyan, Peter Rousseau. 2005. General Purpose Technologies. In Philippe Aghion, Steven N. Durlauf Eds. *Handbook of Economic Growth* Volume 1B, Chapter 18, 1181-1224.

Lipsey, Richard, Kenneth Carlaw, Clifford Bekar. 2005. *Economic Transformations: General Purpose Technologies and Economic Growth*. Oxford University Press, Oxford UK.

Langer N, Forman C, Kekre S, Sun B. 2012. Ushering buyers into electronic channels: An empirical analysis. Inform. Systems Res. 23(4):1212–1231.

Mokyr, Joel. 1990. *The Lever of Riches,* Oxford University Press.

Moser, Petra, Tom Nicholas. 2004. Was electricity a general purpose technology? Evidence from historical patent citations. *American Economic Review Papers & Proceedings* 94(2), 388-394.

Mullainathan, Sendhil, and Ziad Obermeyer. 2021. Diagnosing Physician Error: A Machine Learning Approach to Low-Value Health Care, *The Quarterly Journal of Economics*.

Rosenberg, Nathan. 1963. Technological Change in the Machine Tool Industry, 1840-1910. *Journal of Economic History*, 414-443.

Shapiro and Varian. 1999. *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business Review Press.

Urbich M, Globe G, Pantiri K, et al. 2020. A Systematic Review of Medical Costs Associated with Heart Failure in the USA (2014-2020). *Pharmacoeconomics*. 38(11):1219-1236.