

THE TRILLION DOLLAR CONUNDRUM: COMPLEMENTARITIES AND HEALTH INFORMATION TECHNOLOGY*

David Dranove, Northwestern University
Chris Forman, Georgia Institute of Technology
Avi Goldfarb, University of Toronto
Shane Greenstein, Northwestern University

November 2012

Abstract

We examine the heterogeneous relationship between the adoption of electronic medical records (EMR) and hospital operating costs at thousands of US hospitals between 1996 and 2009. Combining data from multiple sources, we first identify a puzzle that has been seen in prior studies: Adoption of EMR is generally associated with a slight increase in costs. We draw on the literature on information technology and productivity to analyze why this average effect arises, and explain why it masks important differences over time, across locations, and across hospitals. We find evidence consistent with this approach, namely, that: (1) EMR adoption is initially associated with a rise in costs; (2) EMR adoption at hospitals in favorable conditions – such as urban locations – leads to a decrease in costs after three years; and (3) Hospitals in unfavorable conditions experience a sharp increase in costs even after six years.

JEL Code: I10, L30

*Corresponding author: Avi Goldfarb, Rotman School of Management, University of Toronto, 105 St George St, Toronto ON, M5S 3E6, CANADA, phone 416-946-8604, fax 416-978-5433, agoldfarb@rotman.utoronto.ca. Bingyang Li provided excellent research assistance. We thank seminar participants at Georgia Tech, Harvard Business School, MIT, the University of Arizona, the University of Georgia, the University of Helsinki, and the University of Zurich for helpful comments, and we thank Dr. David Artz, Leemore Dafny and anonymous reviewers for useful comments and suggestions. We also thank Harte Hanks Market Intelligence for supplying data. Avi Goldfarb thanks SSHRC for research support. We are responsible for all opinions and errors.

I. Introduction

More than a quarter century ago economists engaged in a vigorous debate about the benefits from investment in information technology (IT) in manufacturing and services. That debate was encapsulated in the Solow ‘Productivity Paradox’—“You can see the computer age everywhere but in the productivity statistics” (Solow 1987). That debate eventually faded from view, in part because the data began to reject it. Over time it was found that firms achieved productivity benefits from IT, just with a lag. Moreover, explanations for the lag emerged from considerable work on IT use in enterprises. The challenges to productivity benefits were due to the costly adaptations required for the successful implementation of new IT. In time, it was found that the firms realizing benefits from their IT investments were those that had made complementary investments in areas such as worker skills and organizational decision rights.¹

A new manifestation of this debate has surfaced around the use of electronic medical records (EMR). A small sampling of research from the last half dozen years provides a sense of the uncertainty about the productivity benefits from these investments. The Congressional Budget Office states: “No aspect of health information technology (i.e., EMR) entails as much uncertainty as the magnitude of its potential benefits” (Congressional Budget Office 2008). A widely cited 2005 report by the RAND Corporation, published in the leading policy journal *Health Affairs*, estimates that widespread adoption of EMR by hospitals and doctors could reduce annual health spending by as much as \$81 billion while simultaneously leading to better outcomes (Hillestad et al. 2005). Jaan Sidorov, a medical director with the Geisinger Health Plan, an early adopter of EMR, published a response to the RAND report in *Health Affairs*. Sidorov (2006) highlights the high cost of adoption and cites evidence that EMR leads to greater health spending and lower productivity. Other recent studies, cited below, also fail to find consistent evidence that EMR savings offset adoption costs.

What impact does EMR have on a key determinant of an existing organization’s productivity, such as its operating costs? We view EMR as a type of business process innovation, one that involves not only

¹ See, for example, Bresnahan, Brynjolfsson, and Hitt (2002) and Bloom, Sadun, and Van Reenen (2012). Several other explanations have also been highlighted for these empirical findings, including mismeasurement of IT capital or output. For further details on these issues, see Triplett (1999).

investments in IT but also changes in the operational practices within the adopting organization. If EMR is viewed in this way, how does that change the understanding of EMR's impact? We argue that prior literature has missed important features of EMR by not building on prior research studying the adoption and productivity benefits of large scale enterprise IT and "insider econometrics" studies of IT adoption.²

Building on this prior research, we stress the complementary assets that reduce the costs associated with "co-invention," which is the process of adapting an innovation to unique circumstances and turning the overall change into a net benefit to the enterprise. These complementary assets come from several sources. *Local* resources may be available as market services, such as expertise in implementing similar business process innovations, or widespread spillovers in how to use the IT. Resources available *internally to the enterprise*, such as expertise with other business processes inside the enterprise, may also help implement the business process innovation and often cannot be purchased from markets in the short run. To summarize, variance in local and internal resources provides an explanation for why the payoff from EMR may be delayed, and for why we observe variance in the returns to investments in enterprise IT across locations.

We conduct an empirical examination of the impact of EMR adoption on hospital operating costs during the period 1996 to 2009, using the prior research on enterprise IT adoption to frame the analysis and results. The data come from several sources linking hospital costs to EMR adoption and the potential for complementarities. Our main analysis regresses logged operating costs on EMR adoption, hospital fixed effects, and a large number of controls. We focus on whether the impact of EMR is greater for hospitals that are positioned to exploit available complementarities. Thus, our key independent variable is the interaction between EMR adoption and the presence of local complements, as measured by the IT-intensity of local industry. Our key identification assumption is that EMR adoption is not correlated with unobservable cost factors that are differentially trending in hospitals with locally available complementary inputs relative to hospitals that lack these inputs. We explain below why we believe that this is a reasonable

² For examples of multi-industry studies that examine the adoption of enterprise IT and the accompanying organizational adaptations, see for example Bresnahan and Greenstein (1996), Bresnahan, Brynjolfsson, and Hitt (2002), Forman, Goldfarb and Greenstein (2005), and Bloom et al. (2009). For examples of single-industry "insider econometrics" studies, see Ichinowski and Shaw (2003) and Bartel, Ichinowski, and Shaw (2007).

assumption; even so, we show robustness to instrumenting for EMR adoption using hospital proximity to EMR vendors and EMR adoption in alliance systems and geographically linked markets.

We find the evidence consistent with our reframing. The timing of cost savings is consistent with what we would expect given the literature on the productivity paradox in IT. For the average hospital, the gains from EMR adoption appear with some delay. Moreover, there is significant heterogeneity in the gains achieved that depend upon the local availability of complementary factors such IT workers. We focus on costs because much of the political discussion has emphasized cost savings, and because the multi-product nature of hospitals makes it easier to measure the implications of EMR for costs than for productivity. However, our focus on costs means that we cannot use our data to rule out the possibility that our results are paralleled by opposite results on clinical benefits, though the prior literature, including Agha (2012), Miller and Tucker (2011), and McCullough et al (2010), have found clinical benefits on average to be small.³

We find that hospitals that adopted EMR between 1996 and 2009 did not experience a statistically significant decrease in costs on average. In fact, under many specifications, costs rose after EMR adoption, particularly for the more advanced EMR systems. However, this effect is mediated by measures of the availability of technology skills in the local labor market. Specifically, in strong IT locations, costs can fall sharply after the first year of adoption to below pre-adoption levels. In weak IT locations, costs remain above pre-adoption levels indefinitely. Overall, hospitals in IT-intensive markets enjoyed a statistically significant 3.4 percent decrease in costs from three years after adoption of basic EMR and a marginally significant 2.2 percent decrease in costs from three years after adoption of advanced EMR. These results are significantly better than the up to 4 percent *increase* in costs after adoption by hospitals in other markets.

Figure 1 displays these general patterns in the raw data, comparing hospitals that adopt basic and advanced EMR before the adoption period, during the adoption period, and after the adoption period. For basic EMR, costs do not fall until three years after adoption. For non-IT-intensive locations, costs rise sharply in the year of adoption, and then fall back. For IT-intensive locations, costs fall with adoption, and

³ Thus far, most significant evidence of significant clinical benefits has been found primarily in mortality risks of high-risk patients (e.g., Miller and Tucker 2011 and McCullough et al. 2011, 2012).

are substantially lower three years after adoption. For advanced EMR, the patterns are similar: costs rise in the period of adoption for non-IT-intensive locations and fall over time for the other hospitals.

We provide evidence that the benefits of strong IT locations arise in part from an agglomeration of IT employment in (other) IT hospitals. Hospitals in locations with strong HIT employment enjoyed a statistically significant 6.9 percent decrease in costs from three years after adoption of basic EMR and a statistically significant 7.3 percent decrease in costs from three years after adoption of advanced EMR. However, concentration of IT employment in other industries is not associated with greater benefits from adopting basic or advanced EMR. Controlling for strong HIT employment, costs still fall more rapidly in strong IT locations than in weak ones. In short, one benefit of strong IT locations is a thicker labor market for HIT workers, though other benefits persist as well.

We also show results suggesting that complementary skills can be found internally in the hospital. For advanced EMR, the initial increase in costs is mitigated substantially if hospitals already have substantial software experience. Hospitals without experience are hurt in the short run for the most sophisticated technologies. We do find, however, that within a short time inexperienced hospitals can make up the difference. Specifically, the difference in costs after adoption for hospitals with and without internal expertise disappears within three years. This suggests that, in contrast to complementary assets that depend on a location with favorable agglomeration economies, some complementary assets to business process innovation can be acquired relatively quickly.

These findings have several implications. As annual U.S. healthcare expenditures climb towards \$3 trillion and with spending forecast to exceed \$4.5 trillion by 2020, many analysts hope that electronic medical records (EMR) can stem the tide (Centers for Medicare & Medicaid Services). For example, David Cutler and Melinda Beeuwkes Buntin make EMR the centerpiece of their “Two Trillion Dollar” solution for modernizing the health care system (Buntin and Cutler 2009). While some are confident in EMR, others remain cautious, especially due to EMR’s sluggish diffusion. As of 2009, only about 30 percent of America’s hospitals have adopted any advanced elements of EMR.⁴ This may have been due, in part, to the

⁴ Source: Authors’ calculations based on data supplied by HIMSS.

lack of consistent evidence of cost savings.

In order to spur EMR adoption, Congress in 2009 passed the Health Information Technology for Economic and Clinical Health Act (HITECH Act), which provides \$20 billion in subsidies for providers who adopt EMR. Two thirds of hospitals said they planned to enroll in the first stage of HITECH subsidy programs by the end of 2012 (US Department of Health and Human Services 2011). The 2010 Patient Protection and Affordable Care Act also contains provisions promoting EMR adoption. Despite these legislative actions, many remain unconvinced of the benefits of EMR. Our findings also may help resolve the ongoing debate. Supporters and detractors both seem to treat EMR as if its economic impact is independent of other environmental factors, as if it either works or it doesn't. This creates a conundrum for both sides. If EMR is going to save hundreds of billions of dollars or more, as its supporters claim, why isn't it working in obvious ways? If it costs more than it saves, as the skeptics argue, why are policy makers so keen to expand adoption? Our results suggest that the debate about EMR should be reframed by drawing on the general literature on business adoption of IT, where it is very common for successful technology adoption to require complementary changes in business processes that rely on specific labor and information inputs. It is also common for new enterprise IT to be more productive when companies have access to these inputs in their local market. Using this experience, it is not surprising that EMR can simultaneously have the potential to generate substantial savings but demonstrate mixed results in practice.

We proceed as follows. Sections II and III describe the institutional setting for EMR, and some of the prior evidence about its effects on hospitals. This motivates a comparison in Section IV between EMR and the adoption of IT inside other organizations, which leads to a reframing of several key hypotheses. Sections V and VI present data and results. Section VII concludes.

II. What is EMR?

EMR is a catchall expression used to characterize a wide range of information technologies used by hospitals to keep track of utilization, costs, outcomes, and billings. In practice, EMR includes, but is not limited to:

- A *Clinical Data Repository (CDR)* is a real time database that combines disparate information

about patients into a single file. This information may include test results, drug utilization, pathology reports, patient demographics, and discharge summaries.

- *Clinical Decision Support Systems (CDSS)* use clinical information to help providers diagnose patients and develop treatment plans.
- *Order Entry* provides electronic forms to streamline hospital operations (replacing faxes and paper forms).
- *Computerized Provider Order Entry (CPOE)* is a more sophisticated type of electronic order entry and involves physician entry of orders into the computer network to medical staff and to departments such as pharmacy or radiology. CPOE systems typically include patient information and clinical guidelines, and can flag potential adverse drug reactions.
- *Physician Documentation* helps physicians use clinical information to generate diagnostic codes that are meaningful for other practitioners and valid for reimbursement

As this list shows, there is no single technology associated with EMR, and different EMR technologies may perform overlapping tasks.

Nearly all of the information collected by EMR already resides in hospital billing and medical records departments and in physicians' offices. EMR automates the collection and reporting of this information, including all diagnostic information, test results, and services and medications received by the patient. EMR can also link this information to administrative data such as insurance information, billing, and basic demographics. EMR can reduce the costs and improve the accuracy of this data collection. Two components of EMR, Clinical Decision Support Systems and Computerized Provider Order Entry, use clinical data to support clinical decision making (Agha (2012) refers to this as a distinct category labeled Clinical Decision Support or CDS). If implemented in ideal conditions and executed according to the highest standards, EMR can reduce personnel costs while facilitating more accurate diagnoses, fewer unnecessary and duplicative tests, and superior outcomes with fewer costly complications.

Despite these potential savings, EMR adoption has been uneven. Table 1 reports hospital adoption rates for the five components of EMR described above. The data is taken from HIMSS Analytics, which we

describe in more detail in Section V. Clinical Data Repository, Clinical Decision Support, and Order Entry are older technologies that were present in many hospitals in the 1990s. Even for these older technologies, adoption rates range from 75 to 85 percent in 2009. The remaining applications emerge in the early to mid-2000s. Adoption rates for these are below 25 percent.

While informative, Table 1 lacks several crucial pieces of information. It lacks comparable data on physician adoption of EMR, for example, which is much lower than hospital adoption (Callaway and Ghosal 2012). Our data do not tell us about intensity of use by physicians and staff within hospitals, about the details of the installation, or on how close operations come to ideal conditions. Interviews with hospital administrators suggest that adoption can be uneven within hospitals, with some departments enthusiastically embracing change while others do not. Although beyond the scope of this study, compatibility issues may shape the success of EMR at a regional level, and this too is missing from the table. There are many different EMR vendors and their systems do not easily interoperate. As a result, independent providers cannot always exchange information, which defeats some of the purpose of EMR adoption (Miller and Tucker 2009).

III. Evidence on the Potential Savings from EMR

Has adoption of EMR reduced hospital costs? This section reviews prior evidence, stressing the absence of work focusing on operational savings, lack of emphasis on complementarities with the labor market, and the absence of accounting for the functional heterogeneity of EMR's components. This discussion will motivate our concerns and our approach to framing the study of EMR's impact on productivity using past research that emphasized enterprise IT as a business process innovation.

Every EMR study remarks on the expense. One prominent estimate, from the Congressional Budget Office (CBO 2008), estimates that the cost of adopting EMR for office-based physicians is between \$25,000 and \$45,000 per physician, with annual maintenance costs of \$3000 to \$9000. For a typical urban hospital, these figures range from \$3-\$9 million for adoption and \$700,000-\$1.35 million for maintenance. These costs are quite significant: If the adoption costs are amortized over ten years, EMR can account for about 1 percent of total provider costs. It would be no surprise, therefore, if research suggested that EMR

does not pay for itself, let alone generate hundreds of millions of dollars in savings.

In their review of 257 studies of EMR effectiveness, Chaudry et al. (2006) note that few studies focus on cost savings, providing, at best, indirect evidence of productivity gains. Most of the studies they review focus on quality of care. Ten studies examine the effects of EMR on utilization of various services. Eight studies show significant reductions of 8.5-24 percent, mainly in laboratory and radiology testing. While fifteen studies contained some data on costs, none offered reliable estimates of cost savings.

Hillestad et al. (2005; the widely cited RAND study mentioned in our introduction) uses results from prior studies of EMR and medical utilization and extrapolates the potential cost savings net of adoption costs. They estimate that if 90 percent of U.S. hospitals were to adopt EMR, total savings in the first year would equal \$41.8 billion, rising to \$77.4 billion after fifteen years. They also predict that EMR adoption could eliminate several million adverse drug events annually, and save tens of thousands of lives through improved chronic disease management.

Sidorov (2006) challenges these findings, arguing that the projected savings are based on unrealistic assumptions. For example, the RAND study appears to assume that EMR would entirely replace a physician's clerical staff. Sidorov argues that providers who adopt EMR tend to reassign staff rather than replace them. To take another example, EMR is supposed to eliminate duplicate tests, while it is just as likely that EMR may allow providers to justify ordering additional tests.⁵ Buntin et al. (2011) review 73 studies of the impact of EMR on medical utilization. EMR is associated with a significant reduction in utilization in 51 (70 percent) of these studies. They do not identify any studies of EMR and costs.

Indeed, we have identified only three focused cost studies. Borzekowski (2009) uses fixed effects regression to examine whether early versions of financial and clinical information technology systems generated significant savings between 1987 and 1994. He finds that hospitals adopting the most thoroughly automated versions of EMR realize up to 5 percent savings within five years of adoption. He also finds that hospitals that adopt less automated versions of EMR experience an increase in costs. His conclusions mirror the popular discussion: there appears to be the potential for savings but there is little understanding

⁵ McCormick et al (2012) document that physicians with EMR tend to order more diagnostic tests, though they do not address the role of omitted variables in driving this result.

of the drivers of the heterogeneity across hospitals. Second, Furukawa, Raghu, and Shao (2010) study the effect of EMR adoption on overall costs among hospitals in California for the period 1998-2007. Also using fixed effects regression, they find that EMR adoption is associated with 6-10 percent higher costs per discharge in medical-surgical acute units, in large part because nursing hours per patient day increased by 15-26 percent. This is plausible because nurse use of EMR can be very time consuming. Third, Agha (2012) uses variation in hospitals' adoption status over time, analyzing 2.5 million inpatient admissions across 3900 hospitals between the years 1998-2005. Health IT is associated with an initial 1.3 percent increase in billed charges. She finds no evidence of cost savings, even five years after adoption. Additionally, adoption appears to have little impact on the quality of care, measured by patient mortality, medical complication rates, adverse drug events, and readmission rates. While not directly about costs, Lee, McCullough, and Town (2012) document small positive effects of hospital IT on productivity.

None of the studies frame EMR in the context of the prior literature on enterprise IT. In other words, there is no examination of factors that shape availability of complementary components such as local expertise or prior experience with related technology. This may be due to a lack of familiarity with the theoretical frameworks that would suggest such differential effects. In the next section, we offer such a framework, based on research on the productivity of large scale IT projects in enterprises, and develop some specific implications for the deployment of EMR.

IV. Information Technology and Complementarities

The existing literature on effective implementation of IT within businesses has emphasized the view of IT as a business process innovation.⁶ Such innovations alter organizational practices, generally with the intent of improving services, reducing operational costs, and taking advantage of new opportunities to match new services to new operational practices. Typically this type of innovation involves changes in the discretion given to employees, changes to the knowledge and information that employees are expected to

⁶ Specifically, Attewell (1992), Bresnahan and Greenstein (1996), Black and Lynch (2001), Bresnahan, Brynjolfsson, and Hitt (2002), Brynjolfsson and Hitt (2003), Hubbard (2003), Forman, Goldfarb, and Greenstein (2005), Bloom, Garicano, Sadun, and Van Reenen (2009), and Bloom, Sadun, and Van Reenen (2012). Forman and Goldfarb (2005) summarize the earlier literature. The literature on "insider econometrics" has touched on related themes, particularly stressing channels of communication, the influence of hierarchy on communications, and changes in hierarchy as a result of the deployment of new information technology (e.g. Ichinowski and Shaw 2003; Bartel, Ichinowski, and Shaw 2007).

retain and employ, and changes to the patterns of communications between employees and administrators within an organization.

Because important innovation in enterprise IT occurs on a large scale, it typically involves a range of investments, both in computing hardware and software, and in communications hardware and software. It also involves retraining employees and redesigning organizational architecture, such as hierarchy, lines of control, compensation patterns, and oversight norms. In the discussion below, we draw on a wide literature to explain a number of common misunderstandings about business process innovations.

For example, there is a misperception that new IT hardware or software yields the vast majority of productivity gains without the need for adaptation by the firm. Prior research has shown that each generation of IT is not readily interchangeable with older products or processes, meaning that the initial investment often does not generate a substantial productivity gain until after complementary investments, adaptations, and organizational changes (e.g. Bresnahan and Greenstein 1996; Bresnahan, Brynjolfsson, and Hitt 2002; Bartel, Ichniowski, and Shaw 2007; Bloom, Sadun, and Van Reenen 2012). Many of these necessary changes are made long after the initial adoption. Hence, it is common for IT investments to have no or negative returns in the short run before yielding positive returns. Among the functions mentioned in EMR, for example, CPOE generates many changes to routine processes. These changes often take time to make, and their productivity gains can come long after the initial rollout.

Thus, business process innovation is not equivalent to installing shrink-wrap software for a PC that works instantly, or merely after training of staff. Instead, prior studies stress the importance of *co-invention*, the post-adoption invention of complementary business processes and adaptations aimed at making IT adoption useful (Bresnahan and Greenstein 1996). The initial investment in IT is not sufficient for ensuring productivity gains. Those gains depends on whether the employees of the adopting organization—in the case of hospitals, administrative staff, doctors, and nurses—find new uses to take advantage of the new capabilities, and/or invent new processes for many unanticipated problems. For example, at one ophthalmology unit at a teaching hospital, the physicians could not find a way to put their traditional “hand-drawings” into the new formats. They found that the new electronic formats sometimes

reduced the richness of the information they could record.

This relates to another common misunderstanding: expectations that the entire cost of investment is incurred as monetary expense. Non-monetary costs comprise a substantial risk from installing a business process innovation. Prior studies emphasize the cost of delays, for example. Delays can arise from non-convexities in investment (e.g., all the wiring must be installed before the communications routines can be tested), the technical necessity to invest in one stage of a project only after another is completed (e.g., the client cannot be modified until the servers work as designed), lack of interoperability during upgrades (which some software handles better than others), and cognitive limits (e.g., staff does not anticipate idiosyncratic issues until a new process is at their fingertips). Moreover, interruptions to ongoing operations generate large opportunity costs in foregone services that can be substantially mediated with internal resources (e.g., development of middleware by in-house IT staff) for which there may be no market price or, for that matter, no potential for resale.⁷

Thus, planning is another common difficulty in IT adoption. Though the installation of any substantial business process innovation requires planning – i.e., administrative effort by the enterprise in advance of installation – such planning alone rarely ends the administrative tasks required to generate productivity gains. Administrative effort does not cease after installation, and continues throughout implementation. Hiring and training personnel generates use of new hardware, software, and procedures. New users in new settings then notice unanticipated problems, which generates new insight about unexpected issues.

As an example of the necessary adjustments and co-invention required for successful EMR investment, consider one large teaching hospital that supported a diverse and geographically dispersed affiliate network.⁸ The IT staff configured the records for patients to suit the needs of physicians treating severe medical issues. The central hospital saw many severely ill patients, since it was a major trauma center for its region. After rolling out this new system, the doctors at the satellite campuses complained of wasting time “cleaning up the records” because no responsibility was assigned for updating the records after an

⁷ Private communication with David Artz, Medical Director of Information Systems, Memorial Sloan-Kettering Cancer Center, in August 2012.

⁸ These observations were obtained from field work that was conducted by one of the authors at a hospital whose management wished to remain anonymous.

urgent event. The doctors at satellite campuses also frequently found themselves wading through many screens when routine issues did not require it, and while the patient was present, diminishing the patient experience. It also lengthened the physician's day, as they spent time updating records. In this simple example, gaining the maximal productivity gains required tailoring the software to the specific types of users and the specific setting, as well as implementing procedures to keep it updated.

Two key empirical implications arise from this discussion. First, given that there was considerable heterogeneity across US locations in the availability of complementary factors, such as skilled labor and knowledge spillovers (Forman, Goldfarb, and Greenstein 2005, 2012), third-party software support and service (Arora and Forman 2007), and infrastructure (Greenstein 2005, Greenstein and McDevitt 2011), there should have been a visible relationship between investment in health IT and local conditions in a limited metropolitan geographic area. Large cities had thicker labor markets for complementary services or for specialized skills. We expect that thicker markets lowered the (quality-adjusted) price of obtaining IT services such as contract programming and of hiring workers to develop in-house functions.⁹ Such locations may also have had better availability of complementary information technology infrastructure, such as broadband services. Increases in each of these factors may have increased the (net) benefits of adopting complex technologies in some cities and not others, other things being equal. Non-IT employees in advantaged locations also may have adapted more easily to EMR. Overall, the presence of thicker labor markets for technical talent, greater input sharing of complex IT processes, and greater knowledge spillovers in cities should have increased the benefits to adoption of frontier technologies in big cities relative to other locations (Henderson 2003; Forman, Goldfarb, and Greenstein 2008).

A few examples help to illustrate. One example is El Camino Hospital in Mountain View, California (e.g., near Silicon Valley). This hospital is an otherwise small community hospital that would normally be a laggard in advanced IT, but not due to its location. It could hire plenty of sophisticated administrators to implement the components of EMR, and IT was "in the air", so the hospital was able to gain worker

⁹ For example, see this quote: "There's a lot of dedicated health care professionals out there in the universe," said Josh Lee, a doctor at and chief medical information officer for the University of California San Diego Medical Center. "There's a lot of dedicated IT professionals. But it's a much narrower band where you have people that can live in both of those worlds." http://www.pcworld.com/businesscenter/article/229071/hospitals_compete_for_it_talent_with_funding_at_stake.html, accessed September 16, 2012.

acceptance and adapt to its advanced EMR system. This type of spillover is not unique to Northern California. It is also related to hospitals in the Milwaukee area, which is not a location normally regarded as one with a thick market for IT talent. However, the largest EMR provider in the country, Epic, has its headquarters in Madison, Wisconsin, less than a 90 minute drive from Milwaukee. This proximity resulted in early and extensive support for hospitals in the Milwaukee area, giving them more experience and, hence, greater success with advanced EMR services than otherwise would be expected.¹⁰

This framework has a second implication, which is less surprising, as it mimics long-standing economic work on learning curves inside organizations. Enterprises with existing IT facilities should expect lower co-invention costs than establishments without extensive operations, and that should shape costs around the time of adoption. Prior IT projects may reduce development costs if on-staff programmers are able to transfer lessons learned from one project to another.¹¹ Prior work on other IT projects may create learning economies and spillovers that decrease the costs of adapting general purpose IT to organizational needs, reducing the importance of external consultants and local spillovers. For example, many major medical centers in the US – such as Duke, Vanderbilt, Hopkins, UPMC, Yale, or Washington University in St. Louis – invested in advanced IT in order to remain competitive, and those centers initially built their EMR with in-house software instead of packages, using internal expertise during every additional investment. That internal expertise proved valuable when the hospitals later adopted packages and customized them to their organizations.

In summary, if the productivity impact of EMR follows patterns seen with other types of IT, then it should come with a lag. Furthermore, the productivity impact of EMR should depend on factors that shape the supply conditions for complements, such as the experience of a hospital's IT staff, as well as the local labor market for skilled labor and third-party software and support.

¹⁰ Private communication with Dr. David Artz, Medical Director of Information Systems, Memorial Sloan-Kettering Cancer Center, August, 2012. Epic's location arose from the preferences of the founder.

¹¹ For example, software developers may be able to share common tools for design, development, and testing (Banker and Slaughter 1997), or they may be able to reuse code (Barnes and Bollinger 1991). Software development may also have learning economies (Attewell 1992) that through experience reduce the unit costs of new IT projects. Much prior research in the costs of innovative activity has also presumed experience with related projects lowers the costs of innovation (Cohen and Levinthal 1990).

V. Data

We use a variety of data sources to examine the relationship between EMR adoption and costs. In particular, the data for this study matches data on EMR adoption from a well-known private data source on health IT investments (HIMSS Analytics) with cost data from the Medicare Hospital Cost Report. We add data from the American Hospital Association’s (AHA) Annual Survey of Hospitals. We obtain regional controls and information on local complementary factors from the decennial U.S. Census and from U.S. County Business Patterns data. We supplement the sources above with information on lagged hospital-level IT experience and local IT workforce from another private source on IT investment, the Harte Hanks Computer Intelligence Database. Our data are organized as an unbalanced panel, with data available every year from 1996 to 2009. Table 2 provides descriptive statistics.¹²

EMR adoption. Information about EMR adoption comes from the Healthcare Information and Management Systems Society (HIMSS) Analytics data base. The HIMSS Annual Study collects information systems data related to software and hardware inventory and reports the current status of EMR implementation in more than 5300 healthcare providers nationwide, including well over 3000 community hospitals.¹³ Because most organizations tend to participate for a long period of time, the HIMSS Analytics data closely approximates panel data and can be used for fixed effects regression.

HIMSS reports adoption of 99 different technologies in 18 categories. Examples include Emergency Department Information Systems, Financial Modeling for Financial Decision Support, and a Laboratory Information System. Following most other studies, we restrict attention to five applications in the category Electronic Medical Records, which we listed above. These closely represent the kind of EMR applications that the RAND study and others believe will lead to dramatic cost savings and quality enhancements.

We aggregate the five EMR applications into two broad categories that we label “basic” and “advanced” EMR, similar to our prior framework for enterprise IT (e.g. Forman, Goldfarb, and Greenstein

¹² The number of observations column in table 2 shows a key challenge within and across data sources: missing data. There is considerable variation across hospitals and years for each of the variables. We simply drop observations with any missing data from our main specifications, though we document that results are robust to some alternatives.

¹³ Community hospitals provide treatments for a wide range of diseases and have relatively short (less than 30 day) average lengths of stay. There are approximately 5000 community hospitals in the United States. HIMSS hospitals are more likely than average to be privately owned and tend to be larger than non-reporting hospitals.

2012). Applications within each of these categories involve similar costs of adoption and require similar types of co-invention to be used successfully. We say that a hospital has basic EMR if it has adopted a clinical data repository (CDR), clinical decision support systems (CDSS), or order entry/communication. We say that a hospital has advanced EMR if it has adopted either computerized practitioner order entry (CPOE) or physician documentation, applications that are more difficult to implement and more difficult to operate successfully due to the need for physician training and involvement. Analyses of health IT adoption, such as the HIMSS Forecasting Model, consider advanced EMR applications to represent the final stage of EMR adoption (HIMSS Analytics 2011).

Table 1 shows sharp increases in adoption of all of these technologies over the sample period. By 2009, at least 70 percent of responding hospitals had adopted each of the basic EMR technologies and at least 20 percent had adopted the advanced technologies.

Our estimation sample is based on the set of hospitals that replied to the HIMSS survey. Thus, we may exclude hospitals that systematically invest little in information systems and have little incentive to reply to the HIMSS survey. Missing data about specific technologies (and to a lesser extent about covariates) mean that our regressions involve 2214 to 3653 hospitals observed an average of 10 to 13 years. A comparison of hospitals that report and do not report data on adoption of basic EMR in 1996 reveals that hospitals who report basic EMR have similar costs per admit (\$9138 versus \$9497 for non-reporters) but are substantially smaller, with 47.2 percent fewer beds. Furthermore, while ownership structures are similar, hospitals that do not report data are less likely to be located in metropolitan statistical areas and are less likely to be teaching hospitals. The Appendix provides further details on the comparison.

Hospital costs. Our primary dependent variable is equal to total hospital operating expenses per admission. There are several reasons why we study the impact of EMR on costs and not productivity. From a policy perspective, the debate on EMR focuses on two dimensions, costs and outcomes. From an econometric perspective, hospitals are multi-product firms. It may be easier to specify cost as the dependent variable and include *ad hoc* controls for product mix than to try to define output on a uniform scale. This may explain why there are many published studies of hospital cost functions but few published studies of

hospital production functions.

We collect data on hospital costs from Medicare Cost Reports. Hospitals are required to report costs to Medicare so that Medicare can compute national reimbursement rates. While these cost data are not audited, hospitals have little incentive to report inaccurately. The cost measure that we use includes the fully amortized operating costs across the entire hospital. These will include the costs of property, plant, and equipment depreciation, but exclude costs of services such as parking garages and public cafeterias.¹⁴ Physician salaries are generally excluded from this measure, except for in-house staff with administrative roles, or those whose billing contributes to hospital revenue, such as radiologists. While our primary measure is total operating expenses per admission, we also show robustness to using total expenses and a case-mix weight on admissions.¹⁵ In some years Cost Report data are missing; in our estimation sample 11 percent of hospital observations are missing cost data. We interpolate values for these missing cost data using the geometric mean of adjacent year costs though results are robust to excluding these observations. Table 2 shows that, on average, costs rise considerably over the sample period (from 9.065 to 9.885 in logged values) but there is a great deal of variation across hospitals.

Because operating expenses include amortized investments in EMR, it would be unsurprising to see increases in costs around the time EMR adoption, reflecting increased purchases of computing and networking hardware and software. We do not expect any such mechanical bump in costs to be correlated with the local IT environment.

We emphasize results on aggregate costs to the hospital, and show robustness to labor costs. While the Cost Report does provide a breakdown of costs by department (e.g. diagnostic radiology, housekeeping, pharmacy), we do not believe these data are reliable reflections of the costs in those areas. In particular, there appears to be substantial accounting discretion in how costs are allocated within the hospital. For example, housekeeping costs range from -\$16 per admit to \$70,000 per admit with a mean of \$196 and a

¹⁴ Depreciation rules are standardized across hospitals. For further details on these rules, see the documentation for the cost report data available at https://www.cms.gov/CostReports/02_HospitalCostReport.asp#TopOfPage.

¹⁵ We obtained annual data on the case mix of Medicare patients for 87% of the sample available from the CMS website (<https://www.cms.gov/Medicare/Medicare-Fee-for-Service-Payment/AcuteInpatientPPS/Historical-Impact-Files-for-FY-1994-through-Present.html>). While Medicare case mix does not match actual case mix, we still document that our results are robust to (i) including the Medicare case mix as a control and (ii) normalizing the cost per admit by the Medicare case mix index. Our main specification assumes that the case mix does not change simultaneously with EMR adoption for reasons other than EMR adoption.

standard deviation of \$523. Thus, our understanding of these data is that the aggregate costs (and aggregate labor costs) are reliable, but the cost breakdown by department is subject to substantial variation in accounting norms across hospitals.

Hospital characteristics. We obtain hospital characteristics from the American Hospital Association Annual Survey. The survey contains details about hospital ownership, service offerings, and financials. We match AHA, Cost Report, and HIMSS data using the hospital Medicare ID and retain only matching hospitals. Specifically, we start with 4915 hospitals in the HIMSS data, drop to 4819 when merging with the Cost Report, and drop to 4493 when merging with the AHA data. We use information from the AHA data and the Medicare Cost Report to exclude several types of hospitals whose costs might be affected by unobservable and/or idiosyncratic factors unrelated to EMR adoption. In particular, we exclude federal hospitals, as well as hospitals that are not defined as short-term general medical and surgical hospitals. (The hospitals that we exclude are not usually considered to be “community” hospitals). Finally, we dropped a small number of hospitals that report very low total costs (less than \$100,000) over one or more years in our sample period. After dropping these, the minimum cost is \$1.2 million and the average cost is \$61 million.

Our final data set contains 4231 hospitals for which we have data on basic EMR adoption for 2228 and data on advanced EMR adoption for 3306. We observe 96 percent of the hospitals in all fourteen years of the data (though both cost and EMR adoption information may not be available in all years).

We use the AHA data to compute the following covariates¹⁶:

- Hospital size: We include number of outpatient visits and number of inpatient days. We also include 1996 values of number of beds and total number of admissions.
- Integration with physicians: We include indicators of whether the hospital operates an independent practice association or a management service organization hospital (as of 1996).
- Hospital ownership: We include for-profit ownership, non-secular non-profit ownership, non-profit church ownership, equity model hospital, or foundation hospital (in 1996).

¹⁶ In a small number of cases, specific pieces of the AHA data are missing for a hospital in a given year but available in other years. In these cases, we impute the missing value using the other years though results are robust to dropping these observations.

- Other characteristics: We include whether the hospital is a teaching hospital (defined as having a residency program or being a member of the Council of Teaching Hospitals), number of births, total costs per admit (to control for different trends on the base level of the dependent variable), and the number of Medicare and Medicaid discharges (all values are from 1996).¹⁷

In our regressions, we interact the 1996 values with a time trend. We emphasize the 1996 baseline to avoid potential changes in hospital characteristics that are driven by the EMR adoption. Results are robust to allowing the characteristics to change over time, but we prefer the simpler specification as a baseline.

Local features. We use U.S. Census data to identify location-level factors that might affect costs independent of IT and to measure complementary factors that might facilitate process innovation. We focus on cross-sectional values to facilitate interpretation (so that locations do not switch status), though results are robust to allowing these values to change over time. For controls, we obtain the following variables from the 2000 decennial U.S. Census and match on county: population, percent Black, percent age 65+ and percent age 25-64, percent university education, and median household income. In our regressions, these are interacted with a time trend to allow different locations to have different cost trends.

To measure the availability of local complementary factors, we use three measures from the Census and two measures based on the Harte Hanks Market Intelligence Computer Intelligence Technology Database (hereafter CI database).

Our main measure of complementary factors is the percentage of local firms that are in IT-using and IT-producing industries. We measure the fraction of firms in IT-using and IT-producing industries in the county as of 1995 from the US Census County Business Patterns data. National aggregate data shows that such industries have unusually high returns from investment in IT in the 1990s. We define these industries using the classification reported in Jorgenson, Ho, and Stiroh (2005, p. 93).¹⁸ Table 2 shows that 43 percent

¹⁷ We only have discharge information for approximately half the hospitals in our sample. When the information is missing, we assume it does not change over time by setting it at average levels and allowing the hospital fixed effects to absorb differences across hospitals.

¹⁸ These industries are Communications (SIC 48), Business Services (73), Wholesales Trade (50-51), Finance (60-62, 67), Printing and Publishing (27), Legal Services (81), Instruments and Miscellaneous Manufacturing (38-39), Insurance (63-64), Industrial Machinery and Computing Equipment (35), Gas Utilities (492, 496, and parts of 493), Professional and Social Services (832-839), Other Transportation Equipment (372-379), Other Electrical Machinery (36, ex. 366-267), Communications Equipment (SIC 366), and Electronic Components (367).

of the hospitals in our data are in counties in the top quartile in IT intensity.

Our second measure of local complementary factors combines county-level income, education, population, and IT intensity. This measure, also used in Forman, Goldfarb, and Greenstein (2012), defines “high all factors” counties as those with over 150,000 population that are in the top quartile in income, education, and IT intensity; 23 percent of the hospitals in our sample fit this criteria.

Third, we include a dummy for whether the hospital is located in an MSA. Urban locations will benefit from additional supply of complementary factors, including thicker labor markets, third party services firms, and better infrastructure. In particular, urban location has been shown to be correlated with lower costs for enterprise IT adoption in a variety of industries (Forman, Goldfarb, and Greenstein 2005, 2008).

To measure local availability of IT skills, we create two measures of local employment of programmers using the CI database: the number of programmers working in hospitals and the number of programmers working outside of hospitals. The CI database contains establishment- and firm-level data on characteristics such as the number of employees, personal computers per employee, number of programmers, and the use of specific software applications. We use data from the 1996 and 2002 data releases. A number of researchers have used this data previously to study adoption of IT (e.g., Bresnahan and Greenstein 1996) and the productivity implications of IT investment (e.g., Bresnahan, Brynjolfsson, and Hitt 2002, Brynjolfsson and Hitt 2003, Bloom, Sadun, and Van Reenen 2012). As has been discussed elsewhere (e.g., Forman, Goldfarb, and Greenstein 2005), this data set represents one of the most comprehensive sources of information on the IT investments of private firms available: for example, as of 2000, it comprised roughly one half of all US establishments over 100 employees

For each establishment in the CI database, we take the midpoint of the (categorical) measure of the number of programmers. For each county-industry, we add up the number of programmers.¹⁹ We weight these by comparing the county-industry employment in our sample to total county-industry employment, where the latter is estimated from the US Census County Business Patterns data. Our measure total non-

¹⁹ Industries are defined by SIC in 1996 and NAICS in 2002. For both years, we use 2-digit industries (SIC in 1996, NAICS in 2002) plus a separate industry for hospitals (SIC 806 in 1996, NAICS 622 in 2002).

health programmers is the weighted sum across all non-health industries.²⁰ For our measure of health programmers, we simply compute the (weighted) employment for hospitals, differencing out the focal hospital. We show results using both the 1996 and 2002 CI databases (and the County Business Patterns from the same year). As we do for strong IT locations, we compute the top quartile for both health and non-health programmers.

IT capabilities. We also use the CI database to obtain measures of historical hospital-level internal IT capabilities. For capabilities of hospitals, we gather data from 1996 on the number of computer programmers and numbers of business and clinical software applications at the hospital, which we interpret as measures of hospital experience with IT. We merge this information into our main data set using hospital names. Unfortunately, because the CI database is itself a sample from a broader population of establishments, there is a significant loss of data from merging these two data sources: the number of hospitals in our sample falls by more than half in the regressions that use the CI database to measure internal hospital experience in software in 1996.

V. Empirical strategy and results

We perform linear regression with hospital and year fixed effects on an unbalanced panel of hospitals observed annually from 1996 to 2009. We proceed in four stages. First, we regress logged costs per patient on different measures of EMR adoption. We document that costs appear to rise on average after adoption. Second, we decompose the rise in costs by years since adoption and show that the rise is largest in the first year of adoption. Third, we examine different margins of complementarity, and show that the results are much stronger for location than for internal IT experience. In particular, we show that the results are strongest in locations with a large number of HIT workers. These results provide suggestive evidence of a difference between complementarities related to available internal expertise and complementarities related to agglomeration economies. Finally, we examine robustness, identification, and plausibility with a variety of further tests.

²⁰ The county-industry weights are equal to (total county-industry employment in County Business Patterns Data)/(total county-industry employment in the CI database). That is, if our data undersamples a given two-digit industry related to the Census it is given more weight in our estimates. We use the health industries described in the previous footnote.

Overall effects: We begin by examining the relationship between (the log of) total administrative costs per admit (c_{it}) and EMR:

$$(1) \text{Log}(c_{it}) = \alpha X_{it} + \beta_t X_i + \gamma_t Z_i + \theta \text{EMR}_{it} + \tau_t + \mu_i + \varepsilon_{it},$$

Here, τ_t captures average changes to costs over time; μ_i is a hospital-specific fixed effect that gets differenced out in the estimation; and EMR_{it} is a discrete variable for whether hospital i had adopted a particular EMR technology by time t . Thus, θ identifies our main effect of interest. We assume that ε_{it} is a normal i.i.d. variable and calculate heteroskedasticity-robust standard errors that are clustered by hospital.

We include three categories of controls. First, X_{it} are controls for hospital characteristics that change over time: inpatient days and outpatient visits. We choose to allow inpatient days and outpatient visits to vary over time to be consistent with prior work on hospital costs that specified these with a translog function (e.g., Dranove and Lindrooth 2003). Second, X_i are all other controls for hospital characteristics. These include beds, type of hospital, ownership status, and discharges. We are concerned that EMR adoption may drive changes in these variables, so including contemporaneous values would be an error. We take their 1996 values and interact them with a linear time trend. Third, Z_i are controls for county-specific characteristics (such as population and income) that do not vary sufficiently over time for changes in their values to have much identifying power. However, the location-level characteristics do have power to identify cost trends. Therefore, we interact these local characteristics with a linear time trend.

For this part of our analysis, our identification relies on the assumption that any systematic changes in hospital costs after EMR adoption are captured by the changes in the hospital-level controls over time and the time trends for the locations.²¹ Put another way, adoption of EMR is uncorrelated with unobservable cost trends that were experienced differentially by adopting hospitals.

Table 3 shows the results of this regression. For columns 1 to 7, the dependent variable is total operating costs per admission, as defined in the AHA data. For column 8 and 9, we use total hospital operating costs (i.e., we do not divide by admissions). Columns 1 to 3 use the specific EMR technologies

²¹ As in Athey and Stern (2002), Hubbard (2003), Bloom et al (2009), Agha (2012), and Forman, Goldfarb, and Greenstein (2012) we initially treat the diffusion of a new technology as an exogenous factor that leads to a change in economic outcomes, and then examine the consequences of the exogeneity assumption.

that together we label “basic EMR”; columns 4 and 8 use the aggregated basic EMR measure (which is equal to one when the hospital has adopted any of the three technologies); column 5 and 6 use the EMR technologies that make up “advanced EMR”; columns 7 and 9 use the aggregated advanced EMR measure.

The results suggest that, on average, EMR does not reduce costs. Instead, in many specifications, EMR is associated with a positive and significant increase in costs of about one to two percent.

Effects by time since adoption: As discussed above, a rich literature on IT productivity has documented that IT adoption affects productivity with a lag. Table 4 examines the extent to which the increase in costs is driven by initial adoption costs such as co-invention and learning new processes. Specifically, Table 4 splits the *EMR* variable into seven pieces, based on time since adoption:

$$(2) \text{Log}(c_{it}) = \alpha X_{it} + \beta_t X_i + \gamma_t Z_i + \sum_{L=0 \dots 6} \theta_L \text{EMR}_{it+L} + \tau_t + \mu_i + \varepsilon_{it},$$

We therefore identify separate coefficients for the first year observed after adoption and for each of the six subsequent years. For the dummy for the sixth year, we use “adopt at least six years earlier.” The hospital fixed effects mean these coefficients should be interpreted relative to the period before adoption.

Costs often rise significantly immediately after adoption, increasing 1.9 percent for advanced EMR. After the first period, costs gradually return to the pre-adoption levels. Generally, the costs return to the pre-adoption levels faster for the basic EMR technologies than for advanced EMR. This is consistent with the prior literature on enterprise IT: initial adoption costs are high because of disruptions to established processes, over time these disruptions diminish, and more complicated technologies take more time to be effectively implemented. It is also consistent with Agha (2012) who finds a transitory increase in total medical expenditures upon adoption but that this increase goes away over time to yield no essentially no change in costs.

Table 5 sets up the sparser specifications that are used in the remainder of the paper to facilitate interpretation. In particular, it focuses on the aggregate measures of basic and advanced EMR and it combines individual years into two variables: “adopt in previous three year period” and “adopt at least three years earlier”. As expected, the results are similar to Table 4.

Effects by location: The literature on enterprise IT has emphasized that efficient use of IT requires the

availability of complementary factors such as skilled labor, third-party software support and service, and infrastructure. To explore this hypothesis, we interact EMR adoption measures with the IT-intensity of a location:

$$(3) \text{Log}(c_{it}) = \alpha X_{it} + \beta_t X_i + \gamma_i Z_i + \theta_1 EMR_{it} + \theta_2 EMR_{it-3+} \\ + \varphi_1 IT_INTENSE_i \times EMR_{it} + \varphi_2 IT_INTENSE_i \times EMR_{it-3+} + \tau_t + \mu_i + \varepsilon_{it},$$

where EMR_{it-3} is a dummy variable for whether the hospital adopted EMR at least three years earlier and $IT_INTENSE_i$ is a measure of whether the location is IT-intensive.

Table 6 examines three distinct measures of IT-intensity: (i) a dummy variable for whether the hospital is in a county that is in the top quartile in terms of IT-using and IT-producing industry, (ii) a dummy variable for whether the hospitals is in a county with high population, income, education and IT-intensive industry (labeled “high all factors” in Forman, Goldfarb, and Greenstein (2012)), and (iii) a dummy variable for whether the hospital is in an MSA. For these estimates, we add a control for these measures interacted with a time trend.

Recall that in the previous analysis our identification assumption was that adoption of EMR was uncorrelated with unobservable cost trends that were experienced differentially by adopting hospitals. In this analysis, which is central to our study, our identifying assumption is weaker. We do not need to assume that adopters and non-adopters experience the same trends in unobservables. Rather, we need to assume that there is no difference in unobservable cost trends around the time of IT adoption in high IT-intensity markets versus low IT-intensity markets; i.e., there is no differential selection on trends in unobservables. Although we can think of no obvious economic reason why this assumption would be violated, in later specifications we will instrument for adoption. Another identification assumption that we require is that hospitals do not relocate to respond to lower EMR adoption costs, and that hospitals cannot easily hire to overcome local IT deficiencies. That is, we assume that an IT-intensive environment requires sufficient local scale, and that hospitals will be a small part of a local IT environment.²²

The first two rows show that costs per admission do not fall in non-IT-intensive counties. In contrast,

²² Our data are consistent with this assumption: for the average county in our data, less than 1% of IT workers are employed in hospitals.

for advanced EMR, costs per admission appear to rise substantially in such locations. The differences between IT-intensive locations and other locations increase after the initial adoption period. For basic EMR, after three years, costs fall a statistically significant 3.4 percent in IT-intensive counties while the coefficient is positive but insignificant in all other counties. For advanced EMR, after three years, costs fall a marginally significant (p-value is 0.09) 2.2 percent in IT-intensive counties while costs rise 3.8 percent in other counties.

Taking the point estimates in columns 1 and 2 of Table 6 at face value, a hospital that installed basic EMR in a favorable location had an average cost reduction of 3.4 percent starting three years after installation, while an installation of advanced EMR in the same location experienced a cost reduction of 2.2 percent. In contrast, a hospital in a poor location would experience an (insignificant) rise in costs of 1.3 percent from three years after adoption of basic EMR and a strongly significant rise of 3.8 percent after adoption of advanced EMR. With average annual operating costs in the tens of millions, these differences are substantial.

Effects by local IT workers: We next explore whether the difference between IT-intensive locations and other locations is driven by workers in the hospital IT sector, or by IT workers of any kind. Specifically, as described above, we use measures of the number of IT programmers (from 1996 and 2002) in hospitals and in all other (non-health) industries in the county:

$$(4) \text{Log}(c_{it}) = \alpha X_{it} + \beta_i X_i + \gamma_i Z_i + \theta_1 EMR_{it} + \theta_2 EMR_{it-3} + \\ + \psi_1 HIT_EMPLOYMENT_i \times EMR_{it} + \psi_2 HIT_EMPLOYMENT_i \times EMR_{it-3} + \\ + \phi_1 OTH_IT_EMPLOYMENT_i \times EMR_{it} + \phi_2 OTH_IT_EMPLOYMENT_i \times EMR_{it-3} + \tau_t + \mu_i + \varepsilon_{it},$$

where $HIT_EMPLOYMENT_i$ is a dummy indicating if the hospital is in a county where HIT employment is in the top quartile, while $OTH_IT_EMPLOYMENT_i$ indicates if the hospital is in a top quartile county by non-health IT employment.

This analysis therefore compares whether the correlations between EMR adoption, IT intensity, and costs are more strongly driven by local hospital IT expertise or by IT expertise across other sectors. As noted above, HIT shares many features with other kinds of enterprise IT, so it is possible that labor skills and knowledge will transfer across industries. Table 7 columns 1 and 2 show the results using the 1996

measures of hospital IT employment and other IT employment and columns 3 and 4 use the 2002 measure. The results suggest that it is local hospital IT employment rather than local IT employment in other sectors that drives the results of the previous subsection. Costs appear to be substantially lower for hospitals located in the top quartile of counties in terms of hospital IT employment. This does not seem to be true of hospitals in the top quartile of counties in non-health IT employment, controlling for hospital IT employment.

Columns 5 through 8 of table 7, however, suggest that local hospital IT employment may not explain the entire result on IT-intensive locations found in table 6. In particular, we include both local hospital IT employment and our measure of IT-intensive location in the regression:

$$(5) \text{Log}(c_{it}) = \alpha X_{it} + \beta_t X_i + \gamma_t Z_i + \theta_1 \text{EMR}_{it} + \theta_2 \text{EMR}_{it-3+} \\ + \psi_1 \text{HIT_EMPLOYMENT}_i \times \text{EMR}_{it} + \psi_2 \text{HIT_EMPLOYMENT}_i \times \text{EMR}_{it-3+} \\ + \phi_1 \text{IT_INTENSE}_i \times \text{EMR}_{it} + \phi_2 \text{IT_INTENSE}_i \times \text{EMR}_{it-3+} + \tau_t + \mu_i + \varepsilon_{it},$$

Consistent with columns 1 to 4 of Table 7, we find that local hospital IT employment is strongly correlated with reduced costs. In addition, we find a weakly persistent relationship between IT-intensive locations and reduced costs, even controlling for hospital IT employment. We interpret this to suggest that local expertise in hospital IT is particularly important but it may not explain all of the difference between IT-intensive counties and other counties.

Effects by hospital IT experience: Internal expertise also can mitigate the costs of adoption of a new process innovation. Importantly, unlike local factors, a hospital may be able to overcome some of these issues by hiring outside expertise. Table 8 examines the interaction in the following format:

$$(6) \text{Log}(c_{it}) = \alpha X_{it} + \beta_t X_i + \gamma_t Z_i + \theta_1 \text{EMR}_{it} + \theta_2 \text{EMR}_{it-3+} \\ + \phi_1 \text{HIT_EXPERIENCE}_i \times \text{EMR}_{it} + \phi_2 \text{HIT_EXPERIENCE}_i \times \text{EMR}_{it-3+} + \tau_t + \mu_i + \varepsilon_{it},$$

As measures of hospital IT experience, we examine business software applications, clinical software applications, and programmers employed (all measured in 1996, at the beginning of the sample). These can be seen as measures of whether the hospital had prior experience in managing software. Given that the sample is reduced by more than half when we merge in the CI database that contains experience information, the additional insight imposes a significant cost on the analysis.

Still, Table 8 suggests a striking contrast to the effects of local IT-intensity. Internal expertise appears

to have little impact on the relationship between basic EMR and costs. It does appear related to reduced costs for hospitals that adopt advanced EMR, but only in the first period after adoption. For each of the three measures, a one standard deviation change yields a 2.5 to 3.9 percent decrease in costs per admittance in the initial three years after advanced EMR adoption. Internal expertise therefore seems particularly important for the most advanced applications that might involve a great deal of co-invention to be successfully employed but any cost disadvantages from a lack of expertise are quickly overcome. We speculate that this might be because it is not difficult for the hospital to hire the expertise from outside. Broadly, the main message of Tables 6, 7, and 8 is consistent with this study's framing, using the results of prior literature on enterprise IT to understand EMR adoption.

Robustness, identification, and plausibility: Next, we explore the degree to which we can claim our main results in table 6 are causal and general. There are four potential types of concerns. First, there might be an omitted variable correlated with EMR adoption and with costs. Second, and related to this, it is possible that unobservable changes in cost drivers are associated with EMR adoption differentially in high and low IT intensity markets. Third, it is possible that anticipated changes in costs drive EMR adoption (rather than EMR adoption driving changes in costs). Fourth, the large amount of missing data may mean that our sample is not representative.

In anticipation of these concerns, we included in our previous analyses hospital and time fixed effects as well as a very large set of covariates as controls. In order to address additional concerns we conduct three types of analyses, examining the timing of the relationship between EMR adoption and cost changes, instrumental variables analysis, and robustness to alternative specifications such as alternative treatments of the missing variables.

In Figure 2, we examine the timing of the relationship between EMR adoption and changes in costs. Specifically, we focus on eventual adopters and exploit variation across hospitals in year of adoption. We run the equation (2) above, but add variables for 1 year before adoption, 2 years before adoption, 3 years before adoption, and 4 or more years before adoption. Figures 2a and 2b demonstrate distinct effects for IT-intensive and non-IT-intensive locations, defined by the top quartile of counties in terms of IT-intensive

industry. Figure 2a examines basic EMR adoption and Figure 2b examines advanced EMR adoption. Prior to adoption, the costs follow similar patterns. During and after the initial adoption, however, the costs in non-IT-intensive locations rise while the costs in IT-intensive locations fall substantially. The coefficients for these regressions are shown in Appendix Table A.2. The timing of the impact of EMR displayed in Figure 2 suggests that there is not a noticeable omitted variable driving the estimates. Similarly, there is no evidence of differential time trends between IT-intensive and non-IT-intensive locations prior to EMR adoption.

In Table 9, we apply instruments for EMR adoption to explore concerns about the direction of causality; one possible concern is that hospitals (especially those in IT-intensive locations) anticipating a reduction in costs will buy an EMR system. While the results are, at best, weakly significant, the signs are consistent with the results of the main specifications. We emphasize three instruments that have some power in the first stage (shown in the Appendix) and, under certain assumptions, may not directly impact costs. First, we use geographic variation in hospitals that belong to multi-location hospital systems and use EMR adoption by competing hospitals in other counties within the same systems as an instrument. This instrument is similar to the one used in Forman, Goldfarb, and Greenstein (2008) to examine the impact of internal expertise in IT on advanced internet adoption by US businesses. The identification assumption is that adoption by competing hospitals in other geographic markets will increase the likelihood of EMR adoption by hospitals within the same system but in those other geographic markets. This will decrease the costs of EMR adoption by the focal hospital but should not affect its other costs.

Second, we use the distance from the hospital to the closest EMR vendor as of 1996 (and interact this with a time trend). The identification assumption is that hospitals near EMR vendor offices will have lower costs for learning about EMR systems.

Third, we use information on hospital alliances and use adoption by other hospitals in the same alliance as an instrument for own adoption. The identification assumption is that adoption by other hospitals in the alliance might lead to lower EMR adoption costs or better information about EMR but will not be coincident with trends in costs for other operating procedures. Given that we think the third

instrument requires the strongest assumptions, we show results for all three and for just the first two. In the Appendix, we show results for each instrument separately.

Because the instruments are at the hospital level rather than the hospital-year level, we focus on one covariate for EMR adoption: whether the hospital adopted EMR at least three years earlier. There is variation in the power of the instruments in the first stage, though they are generally weak with first stage F-statistics ranging from 2.00 to 23.23. The competing hospitals instrument is the most powerful; it works best for basic adoption but also for advanced adoption. The distance to nearest vendor instrument is quite weak. The hospital alliance instrument has some power for advanced adoption but little for basic adoption. For the second stage, Hausman tests show that the coefficient values are not significantly different from the main results, though this is driven more by high standard errors than similar coefficient values. The p-values of the overidentification tests range from 0.28 to 0.75.

Column 1 of Table 9 shows that the coefficients on the main effect of basic EMR adoption turn negative when we instrument for adoption. The main effect is positive (but not significant) for advanced EMR adoption in column 2. Perhaps more importantly, columns 3 and 4 show that the signs of the results on the difference between IT-intensive counties and other counties hold, and the results are weakly significant for basic EMR. Columns 5 through 8 provide nearly identical qualitative results with just two instruments. While not conclusive, we view the instrumental variables analysis as suggestive that our result on the difference between IT intensive locations and other locations is unlikely to be driven by anticipated changes in costs leading to more adoption.

In the Appendix, we address other concerns regarding specification. We show that results are robust to dropping all controls, to adding controls for time-varying hospital characteristics and the Medicare case mix index, and to changing the dependent variable to labor costs per admit, direct costs per admit, or total costs per admit weighted by the Medicare case mix index. We also show that the results are robust to including only those hospitals observed in all years (a balanced panel) and to dropping locations in which hospital IT workers make up a substantial fraction of IT employment. We also document strong similarity between the coefficients on the controls for our main sample and for a sample that excludes hospitals that never adopt,

suggesting that the control group of hospitals has a similar cost function to the treatment group.

VI. Conclusion

Drawing on a variety of data sources on IT, EMR, local demographics, and hospital characteristics, this study shows that the impact of EMR adoption is consistent with the view of EMR as a business process innovation. While EMR adoption appears to be associated with an increase in costs on average, there is important heterogeneity over time, across technologies, across locations, and across hospitals. Both basic and advanced EMR adoption are initially associated with a rise in costs, and this initial increase in costs is mitigated in hospitals with some internal information technology expertise. After three years, hospitals in IT-intensive locations experience a (significant) 3.4 percent decrease in costs after adopting basic EMR, and a marginally significant 2.2 percent decrease in costs after adopting advanced EMR. These benefits are greatest in locations with a large number of HIT workers, though the benefits of IT-intensive locations appear to extend beyond local expertise in hospital IT. In contrast, hospitals in other locations experience an increase in costs, even after several years.

As with any empirical work, our analysis has a number of limitations. First, we observe only a subset of the medical providers in the United States. Doctors' offices, outpatient clinics, nursing homes, and other medical practices may have had a different experience. While we believe it is likely that the general principles of the prior literature on IT would apply broadly, our evidence is specific to hospitals. Second, we focus on a particular set of EMR technologies over a particular time period. It is possible that the technologies that have arisen since 2009 may be more effective and easier to implement. Third, a key assumption is that hospitals represent a small fraction of local IT expertise and employment. If this is not the case, then our explanation based on complementarities related to co-invention costs is hard to justify.

This study also leaves open questions such as why hospitals adopt if their costs do not fall. It might be due to the pursuit of a societal ideal in spite of the cost, or perhaps misconceptions, expected benefits that we do not measure, the difference between *ex ante* aspirations and *ex post* experience, or something else. We have tried to address the endogeneity of this adoption through various techniques, but we cannot completely rule out the possibility that adopting hospitals in IT-intensive locations adopt because they

expect their costs to fall for some reason other than the complementarities of the local IT environment. Relatedly, though the evidence in the literature is mixed on whether hospitals accrue benefits, such as improved clinical outcomes or reduced errors, it is possible that hospitals outside IT-intensive locations experience a sharp increase in benefits such as clinical outcomes and reduced errors. In that case our findings on reduced costs only tell part of the story.

Despite these limitations, we believe our results help inform the discussion on the “trillion dollar conundrum,” providing the (perhaps missing) link between healthcare IT and healthcare costs. Indeed, our results can be restated as a possible resolution to the trillion dollar conundrum. EMR may succeed when the necessary complements are present and the complementary components are in place. Until then, the results of EMR implementation, at best, can be only mixed. While EMR’s past mixed performance is no guarantee of a future result, the past experience also is no guarantee of future failure. Over time, complementary IT skills are expected to become more widely available, and the various components more widely deployed. If so, more hospitals will enjoy the benefits of EMR and it may yet fulfill its promise.

References

- Agha, L. 2012. The Effects of Health Information Technology on the Costs and Quality of Medical Care. Working paper, Boston University.
- Allison, P. 2002. *Missing Data*, Sage Publications: Thousand Oaks, CA.
- Arora, A. & C. Forman. 2007. Proximity and Information Technology Outsourcing: How Local are IT Services Markets. *Journal of Management Information System* 24(2): 73-102.
- Athey, S., & S. Stern. 2002. The Impact of Information Technology on Emergency Health Care Outcomes. *RAND Journal of Economics* 33: 399-432.
- Attewell, P. 1992. Technology Diffusion and Organizational Learning: The Case of Business Computing. *Organizational Science* 3: 1-19.
- Autor, D., F. Levy, & R. J. Murnane. 2003. The Skill Content of Recent Technological Change: An Empirical Exploration. *Quarterly Journal of Economics* 118(4): 1279-334.
- Banker, R. & S. Slaughter. 1997. A Field Study of Scale Economies in Software Maintenance. *Management Science* 43: 1709-1725.
- Barnes, B.H. & T.B. Bollinger. 1991. Making reuse cost-effective. *IEEE Software* 8(1): 13-24.

- Bartel, A. C. Ichniowski, and K. Shaw. 2007. How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills. *Quarterly Journal of Economics* 122(4): 1721-1758.
- Beaudry, Paul, Mark Doms, & Ethan Lewis. 2006. Endogenous Skill Bias in Technology Adoption: City-Level Evidence from the IT Revolution. Federal Reserve Bank of San Francisco. Working Paper #06-24
- Black, S. E., & L. M. Lynch. 2001. How To Compete: The Impact of Workplace Practices and Information Technology on Productivity. *Review of Economics and Statistics* 83: 434-445.
- Bloom, N., L. Garicano, R. Sadun, & J. Van Reenen. 2009. The Distinct Effects of Information Technology and Communication Technology on Firm Organization. NBER Working Paper #14975.
- Bloom, N., R. Sadun, & J. Van Reenen. 2012. Americans Do IT Better: US Multinationals and the Productivity Miracle. *American Economic Review* 102(1): 167-201.
- Borzekowski, R. 2009. Measuring the Cost Impact of Hospital Information Systems: 1987-1994 *Journal of Health Economics* 28: 938-49.
- Bresnahan, T. F., E. Brynjolfsson, & L. M. Hitt. 2002. Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence. *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.
- Brynjolfsson, E. & L. Hitt. 2003. Computing Productivity: Firm-Level Evidence. *The Review of Economics and Statistics* 85(4): 793-808.
- Buntin, M. & Cutler, D. 2009. The Two Trillion Dollar Solution. Center for American Progress. http://www.americanprogress.org/issues/2009/06/pdf/2trillion_solution.pdf.
- Buntin, M., M. Burke, M. Hoagline, & D. Blumenthal. 2011. The Benefits of Health Information Technology: A Review of the Recent Literature Shows Predominately Positive Results. *Health Affairs* 30(3): 464-471.
- Callaway, Brant, and Vivek Ghosal. 2012. The Adoption and Diffusion of Health Information Technology: The Case of Primary Care Clinics. Mimeo, Georgia Institute of Technology.
- Centers for Medicare & Medicaid Services. 2010. National Health Expenditure Projections 2010-2020. At <https://www.cms.gov/NationalHealthExpendData/downloads/proj2010.pdf> . Accessed Feb. 14, 2012.
- Chaudhry, B., J. Wang, S. Wu, M. Maglione, W. Mojica, E. Roth, S. Morton, & P. Shekelle. 2006. Systematic Review: Impact of Health Information Technology on Quality, Efficiency, and Costs of Medical Care. *Annals of Internal Medicine* 144(10): 742-752.
- Cohen, W. & D. Levinthal. 1990. Absorptive Capacity: A New Perspective on Learning and Innovation. *Administrative Science Quarterly* 35: 128-152.
- Congressional Budget Office. 2008. Evidences on the Costs and Benefits of Healthcare Information Technology. Available at <http://www.cbo.gov/ftpdocs/91xx/doc9168/maintext.3.1.shtml>. Accessed on February 14, 2012.
- Dranove, David, & Richard Lindrooth. 2003. Hospital consolidation and costs: another look at the evidence. *Journal of Health Economics* 22: 983-997.

- Fonkych, K. & R. Taylor. 2005. *The State and Pattern of Health Information Technology Adoption*. Santa Monica, RAND Corporation
- Forman, C., & A. Goldfarb. 2006. Diffusion of Information and Communication Technologies to Businesses. In *Handbook of Information Systems, Volume 1: Economics and Information Systems*, ed. Terrence Hendershott. pp.1–52. Amsterdam: Elsevier.
- Forman, C., A. Goldfarb, & S. Greenstein. 2005. How Did Location Affect the Adoption of the Commercial Internet? Global Village vs Urban Density. *Journal of Urban Economics*, 58(3): 389-420.
- Forman, C., A. Goldfarb, & S. Greenstein. 2008. Understanding Inputs into Innovation: Do Cities Substitute for Internal Firm Resources? *Journal of Economics and Management Strategy*. 295-316.
- Forman, C., A. Goldfarb, & S. Greenstein. 2012. The Internet and Local Wages: A Puzzle. *American Economic Review*, 102(1) 556-575.
- Furukawa, M., Raghu, T. & B. Shao. 2010. Electronic Medical Records, Nurse Staffing, and Nurse-Sensitive Patient Outcomes: Evidence from California Hospitals, 1998-2007. *Health Services Research* 45(4): 941-962.
- Greenstien, S. 2005. The Economic Geography of Internet Infrastructure in the United States. In *Handbook of Telecommunication Economics, Volume II* eds. Martin Cave, Sumit Majumdar, and Ingo Vogelsang. pp. 289-364. Amsterdam: Elsevier.
- Greenstein, S., & R. McDevitt. 2011, The Broadband Bonus: Estimating Broadband Internet's Economic Value. *Telecommunications Policy* 35: 617-632.
- Hillestad, R., J. Bigelow, A. Bower, F. Girosi, R. Meili, R. Scoville, & R. Taylor. 2005. Can Electronic Medical Record Systems Transform Healthcare? An Assessment of Potential Health Benefits, Savings, and Costs. *Health Affairs* 24(5): 1103-17.
- HIMSS Analytics. 2011. Essentials of the U.S. Hospital IT Market, 6th Edition.
- Henderson, V. 2003. Marshall's Scale Economies. *Journal of Urban Economics* 53(1): 1-28.
- Hubbard, T. N. 2003. Information, Decisions, and Productivity: On-Board Computers and Capacity Utilization in Trucking. *American Economic Review* 93: 1328-1353.
- Ichniowski, C., and K. Shaw. 2003. Beyond Incentive Pay: Insiders' Estimates of the Value of Complementary Human Resource Management Practices. *Journal of Economic Perspectives* 17(1): 155-180.
- Jorgenson, Dale W., Mun S. Ho, & Kevin J. Stiroh. 2005. *Productivity, Volume 3: Information Technology and the American Growth Resurgence*. Cambridge, MA: MIT Press.
- Kazley, A. & Y. Ozcan. 2007. Organizational and Environmental Determinants of Hospital EMR Adoption: A National Study. *Journal of Medical Systems* 31: 375-84.
- Kolko, J. 2002. Silicon Mountains, Silicon Molehills: Geographic concentration and convergence of internet industries in the US. *Information Economics and Policy* 14(2): 211–32.
- Lee, Jinhyung, J. S. McCullough, & R.J. Town. 2012. The Impact of Health Information Technology on Hospital Productivity. NBER Working Paper #18025.
- Manski, C. 1993. Identification of Endogenous Social Effects: The Reflection Problem. *Review of Economic Studies* 60(3): 531-42

- McCormick, D., D. H. Bor, S. Woolhandler, & D. U. Himmelstein. 2012. Giving Office-Based Physicians Electronic Access to Patients' Prior Imaging And Lab Results Did Not Deter Ordering Of Tests. *Health Affairs* 31(3): 488-496.
- McCullough, J. 2008. The Adoption of Hospital Information Systems. *Health Economics* 17: 649-64.
- McCullough, J., M. Casey, I. Moscovice, & S. Prasad. 2010. The effect of health information technology on quality in US hospitals. *Health Affairs* 29: 647-654.
- McCullough, J., S. T. Parente, & R.J. Town. 2011. The Effect of Health Information Technology on Patient Outcomes. Working Paper, University of Minnesota.
- McCullough, J., S. T. Parente, & R.J. Town. 2012. Health Information Technology and Quality: The Role of Organizational and Informational Complementarities. Working Paper, University of Minnesota.
- Miller, A. & C. Tucker. 2009. Privacy Protection and Technology Diffusion: The Case of Electronic Medical Records. *Management Science* 55(7): 1077-93.
- Miller, A. & C. Tucker. 2011. Can Healthcare Information Technology Save Babies? *Journal of Political Economy* 119(2): 289-324.
- Parente, S. & L. Van Horn. 2006. Valuing Hospital Investment in Information Technology: Does Governance Make a Difference?" *Health Care Financing Review* 28(2): 31-43
- Schaffer, M.E., 2010. xtivreg2: Stata module to perform extended IV/2SLS, GMM and AC/HAC, LIML and k-class regression for panel data models. <http://ideas.repec.org/c/boc/bocode/s456501.html>
- Sidorov, J., 2006, "It Ain't Necessarily So: The Electronic Health Record and the Unlikely Prospect of Reducing Health Care Costs." *Health Affairs* 25(4): 1079-85.
- Solow, R. 1987. We'd better watch out. *New York Times* (July 12), Book Review, 36.
- Triplett, J. 1999. The Solow Productivity Paradox: What Do Computers Do to Productivity? *Canadian Journal of Economics* 32(2): 309-334.
- U.S. Department of Health and Human Services. 2011. Surveys Show Significant Proportions of Hospitals and Doctors Already Plan to Adopt Electronic Health Records and Qualify for Federal Incentive Payments. HHS Press Release, 1/13/ 2011.
- Wager, Karen A., Frances Wickham Lee, & John P. Glaser. 2009. Health Care Information Systems. Second Edition. Wiley/Jossey-Bass, San Francisco.