How does labor mobility affect business adoption of a GPT? The case of machine learning[*]

Ruyu Chen
Dyson School of Applied Economics and Management
Cornell University
137 Reservoir Avenue
Ithaca, NY 14853
USA
rc723@cornell.edu

Natarajan Balasubramanian
Whitman School of Management
Syracuse University
721 University Ave
Syracuse, NY 13244
USA
nabalasu@syr.edu

Chris Forman
Dyson School of Applied Economics and Management
Cornell University
137 Reservoir Avenue
Ithaca, NY 14853
USA
chris.forman@cornell.edu

Abstract: We investigate how worker mobility influences the adoption of a new general-purpose technology (GPT). Using data from over 153,000 establishments between 2010 and 2018, we observe establishment decisions to adopt machine learning. Taking advantage of state-level changes to the enforceability of noncompete agreements as an exogenous shock to worker mobility, we find that changes that facilitate worker movements are associated with a significant decline in the likelihood of adoption. Moreover, the magnitude of establishment response depends upon characteristics of the establishment, the location in which it resides and its industry—in particular, establishment size and number of large establishments in the same industry-location and the level of experimentation with analytics technology. These results are consistent with the view that increases in worker mobility lead to greater risks for establishments that are contemplating adoption of a new GPT that involves significant downstream innovation.

Keywords: ML adoption, labor mobility, GPT, co-invention

[*] We thank Sara Bana, Kenneth Huang, Aija Leiponen, Matt Marx, Evan Starr, Jonathan Timmis, Pengxiang Zhang, and seminar participants at Cornell University; the Workshop on Information Systems and Economics; the Temple-CMU-NYU 2020 Conference on Artificial Intelligence, Machine Learning, and Business Analytics; the ISB 2nd AI & Strategy Consortium; Wharton Innovation Doctoral Symposium (WINDS); 19th ZEW Conference on the Economics of Information and Communication Technologies; and the 2021 Academy of Management Annual Meeting (AOM) for helpful comments and suggestions. We thank Aberdeen Computer Intelligence for supplying data. All opinions and errors are ours alone.

How does labor mobility affect business adoption of a GPT? The case of machine learning

INTRODUCTION

Recent research has highlighted the potential benefits of adopting advanced analytics and data-driven decision-making for the productivity of firms (e.g., Brynjolfsson and McElheran 2016; Brynjolfsson, Jin, and McElheran 2021; Tambe 2014; Wu, Hitt, and Lou 2020). However, it is well known that adoption of new IT systems such as advanced analytics requires complementary investments to achieve productivity gains (e.g., Bloom et al. 2012; Bresnahan et al. 2002; Brynjolfsson, Rock, and Syverson 2021). In particular, adoption of IT systems requires firms to invest in complementary business process innovation to adapt general-purpose systems to the idiosyncratic needs of firms (e.g., Bresnahan and Greenstein 1996). Implementing such complementary innovation requires significant human capital, both as part of the implementation process as well as part of changes to organizational processes, strategy, and structure that have been documented in the literature (see Brynjolfsson and Milgrom 2013 for a recent review).

An important way for firms to acquire the necessary human capital is by hiring workers from other firms. This type of human capital acquisition from other firms has been shown to have a significant impact on hiring firms' productivity (Tambe and Hitt 2014; Wu et al. 2018), particularly for human capital related to new technologies (Tambe 2014). Hiring workers has also been highlighted as an important channel for human capital acquisition and knowledge transfer in the literature on IT spillovers (e.g., Chang and Gurbaxani 2012a, b; Cheng and Nault 2007, 2012; Tambe and Hitt 2014). Thus, labor mobility across firms can facilitate the diffusion of new technologies by acting as a conduit through which firms can obtain the human capital needed to deploy general-purpose IT systems.

However, labor mobility across firms can also have negative effects on technology diffusion. In particular, the risk of workers leaving the firm, potentially with valuable knowledge, can depress incentives for firms to make investments in new technologies that will engender gains to worker skills, through either formal training or on-the-job learning. These risks are likely to be greatest among new general-purpose technologies (GPTs). This is because GPTs cannot be used productively without significant downstream innovation (Bresnahan and Trajtenberg 1995) and investments in worker knowledge, which in turn make workers particularly valuable to a firm's competitors.

We take a first look at examining the empirical salience of these competing effects of labor mobility on adoption of machine learning (ML). Following recent work, we will focus on ML as a prediction technology that is a particular subfield of artificial intelligence, which has been argued to be a new GPT (Cockburn et al. 2019; Goldfarb et al. 2020; Trajtenberg 2019). We use changes in state-level enforceability of noncompete agreements (NCAs) as a plausibly exogenous source of variation in labor mobility (Balasubramanian et al. 2020; Ewens and Marx 2018; Marx et al. 2009). In particular, we study how state-level changes in the strength of NCA enforceability between 2010 and 2018 influence the adoption of machine learning within firms over the same period. Our measure of machine learning adoption is based on data from over 150,000 establishments with 50 or more employees in the Aberdeen Computer Intelligence database over this period. An important advantage of this database is that it includes establishments that vary significantly with respect to their size, industry, and location in ways that influence the (net) benefits to technology adoption and the implications of NCA enforceability for their behavior.

Using a two-period difference model, we find that increased labor mobility, as measured by a loosening of enforceability of NCAs, is associated with a 0.6 percentage point decline in the likelihood of machine learning adoption by establishments in our sample. Given an average adoption rate in 2018 of 9.7%, this translates to a roughly 6.2% decline in the likelihood of adoption on average. We probe and confirm the robustness of this result in a number of ways, including a variety of different controls, subsamples, and measurement strategies. We also explore a falsification exercise to an alternative technology—touchscreen tablets—that is likely to require less downstream innovation and so should be less sensitive to the potential risks of within-industry labor mobility. We find that changes in NCAs have no significant effect on this alternative technology.

Motivated by prior research in economics and information systems, we then examine heterogeneity in our main results. The costs of adopting new technologies are higher in larger organizations, often requiring significant co-invention by users to adapt GPTs to idiosyncratic firm needs, processes, structures, and technologies (Bresnahan and Greenstein 1996; Forman et al. 2005; Ito 1995). Further, the potential risks of labor mobility to technology adoption are likely to be greater in larger regions and in those with a greater concentration of workers in the same industry (Forman et al. 2005; Greenwood et al. 2019; Tambe 2014). These facts suggest that the impact of labor mobility on ML adoption is likely to be greater for larger establishments and establishments that are in locations with a large number of establishments from the same industry. We further examine heterogeneity in our results based on industry-level differences in experimentation with analytics technology.

Our findings about the impact of labor mobility on technology adoption are consistent with these predictions. Establishments with more than 100 employees see a 0.9 percentage point decline in the likelihood of ML adoption when NCAs change in a way that benefits employees (or 7.4% of the average adoption rate in this set of establishments, and statistically significant at the 5% level), while those with fewer than 100 employees see no economically or statistically significant change. We also find that the effects of NCAs on technology adoption are greater in large urban areas; consistent with NCAs limiting mobility to a firm's competitors, these results are explained by a stronger effect of NCAs when the focal establishment is in a location with a greater number of large establishments in its own industry. Together, our results add new empirical evidence about an understudied effect of labor mobility on the adoption of GPTs.

THEORETICAL FRAMEWORK

In this section we describe the theoretical motivation behind our empirical tests. We begin by characterizing the nature of investments required to adopt (and receive benefits from adopting) ML technology. We then describe how the costs of these investments will be shaped by worker mobility and how they interact with establishment and local characteristics.

Machine Learning and Business Process Innovation

As noted above, ML facilitates prediction of events in contexts where traditional algorithms perform poorly (e.g., Agrawal et al. 2018, 2019; Cockburn et al. 2019; Goldfarb et al. 2020; Trajtenberg 2019). For example, in a business setting, ML can help predict why sales are down in a particular region and where attrition will occur within an organization (LePlante 2020).

The diffusion of ML technology in business has depended in part on upstream innovation in algorithms that have been documented in earlier work (Cockburn et al. 2019; Nilsson 2010).

However, as with all GPTs, it also requires downstream innovation that occurs through a combination of work in application firms (Bresnahan and Gambardella 1998; Rosenberg 1982) and by user firms themselves.

Our focus in this paper is across a range of downstream industries and firm sizes, rather than on digital firms for whom ML technology is central to the product or service provided. In this setting, firms adopting ML have the option of either building tools themselves or adapting existing application software to unique circumstances. In earlier generations of general-purpose information technology, firms often adopted packaged software solutions and then engaged in complementary business process innovation to adapt such solutions to firm needs (Bresnahan and Greenstein 1996).

Therefore, in our analysis we focus on the adoption of ML that has been incorporated into enterprise application software, specifically business analytics software, which facilitates organizational decision-making by identifying patterns in data (LePlante 2020). Sometimes known as augmented analytics, such incorporation of ML represents an advance over analytics software that incorporates descriptive visualization technologies and diagnostic tools such as online analytical processing (LePlante 2020). Although acquiring ML-enabled business application software—rather than building such software internally—will reduce some of the necessary human capital investments, applying ML to specific downstream applications will still require significant investments by downstream firms in the human capital of workers to adapt general-purpose algorithms to specific applications.

In particular, knowledge of how to implement ML models must frequently be combined with domain knowledge (Tambe 2019), through some combination of formal training and on-the-job learning. One example that sheds light on the nature of investments required by adopters is a recent collaboration between C3 and 3M (Siegel and Makinen 2018). C3 used a combination of connected devices and advanced analytics enabled by ML. C3 worked in a variety of specialized industries, from mining to national defense, but it did not employ specialists in those domains. 3M deployed C3 to help 3M's health care division build a new software system for hospitals and health care providers that used predictive analytics and ML for clinical documentation, computer-assisted coding, and health analytics. Employees at 3M who used the system trained on C3 at C3's headquarters for several weeks, passed an exam, and then worked with a C3 team to build out the application over a 90-day period (Siegel and Makinen 2018). Thus, in this example, workers at 3M gained valuable industry- and domain-specific skills and knowledge as a result of these investments made by 3M during the implementation of C3.

Mobility and Machine Learning Adoption

As noted, significant investments are required to adapt ML-enabled enterprise software to business needs. In particular, ML adoption will require significant complementary human capital investments in the firm's workers related to the functionality of ML and its application to business processes. Some of these investments will be firm- and industry-specific. In this subsection, we detail how the costs and benefits of these investments can depend on labor mobility, and in turn influence ML adoption.

When employees move to a new employer, they are able to apply the experiences they have learned to new environments, lowering the implementation costs and increasing the productivity of investments in their new firms (Tambe 2014; Tambe and Hitt 2014; Wu et al. 2018). However, at the same time, the possibility of employee movements out of a firm will increase the risks (to the firm) of making the necessary human capital investments by reducing future benefits (e.g., through the additional costs of training new workers). These latter risks may loom particularly large for investment in GPTs like ML, when initial investments are particularly large and the stream of payoffs accrues only after significant periods of time (Brynjolfsson et al. 2021).

As a stylized example, consider a business establishment evaluating a potential investment in ML. As a result of such investment, workers will accumulate human capital that will increase their value to other firms. This, in turn, will increase their outside options for alternative wage offers from other firms in the same region.

Reducing the costs of workers moving within a state (by loosening the legal enforceability of NCAs) will have competing effects on the likelihood of adoption across establishments within that state. First, a decline in the costs of mobility to the workers will increase the likelihood that workers who have accumulated human capital related to ML will move to other firms. These increased risks of workers leaving the establishment will decrease the establishment's net benefits from making the investments necessary for ML adoption, other things equal. As noted, investments in ML produce a stream of payoffs over time. However, to obtain these benefits, firms require skilled workers who have been appropriately trained. If workers leave the firm after receiving training, the firm will need to train new workers to receive the benefits from AI adoption. If training costs are sufficiently high, this can reduce the benefits of making ML investments and dissuade adoption.

However, if some establishments continue to make investments in ML, increased worker mobility will lower the costs of adoption and complementary investments for other firms in the same region. The reason is that IT workers tend to disproportionately change jobs within the same region (Neffke et al. 2017; Tambe and Hitt 2014). Workers who change jobs can share their experiences of deploying ML in their earlier firm. This is particularly true if they change jobs within the same industry, because of the significant industry- and process-specific human capital that is necessary for deploying ML systems. In this way, increases in mobility could increase adoption in a region.

In sum, increased worker mobility can increase or decrease the net benefits of adopting ML. The net effect of these two factors on the likelihood of ML adoption for a particular establishment can in turn depend upon other additional factors, related to characteristics of the establishment and its location.

Establishment size. The benefits and costs to business process innovation enabled by adopting enterprise software systems, and consequently the impact of increasing worker mobility, vary with size. For several reasons, increases in worker mobility are likely to have a more negative effect on adoption behavior in large establishments. First, large plants are more likely to have related technological complements and more mature processes (Bresnahan and Greenstein 1996; Brynjolfsson and McElheran 2016) that will increase the value of adopting new enterprise software systems. However, integrating new systems within an installed base is also costlier and can delay adoption (Bresnahan and Greenstein 1996; Forman 2005). This deterrence is likely to be amplified by the risks associated with worker mobility. Second, to the extent process innovation reduces unit costs, its benefits will be proportional to the output of the firm (Klepper 1996). This in turn implies that any reduction in unit cost benefits due to higher worker mobility is also proportional to the output of the firm. Thus, in the presence of high worker mobility, this effect is likely to deter larger establishments from adopting. Third, business process innovation that is embedded within complex business processes (associated with large establishments) may be more difficult for competitors to replicate (Mata et al. 1995; Melville et al. 2004), further increasing the benefits to large establishments. However, some of these very characteristics that make the benefits of IT adoption more valuable in large establishments make the risks of mobility highest to such establishments. Workers who are involved in ML adoption in large establishments will acquire more human capital than workers employed elsewhere, and this human capital will include a significant industry-specific component, making them particularly valuable to other firms in the same industry. This leads to a prediction that increases in worker mobility will have a more negative effect on adoption behavior in large establishments.

Location size. The effects of worker mobility will similarly be magnified in large locations. In large locations, the outside options faced by workers who have accumulated ML-related human capital are higher, depressing incentives to make investments in ML that will train workers. However, in large regions the benefits to firms that are seeking to hire workers who have obtained the requisite skills elsewhere will also be larger. Specifically, in large regions, the larger number of workers and firms will make matching between worker skills and employer needs easier, thus increasing the benefits to firms hiring workers from elsewhere.

In particular, the effects of worker mobility will be larger in locations with more firms from the same industry. The implementation of ML requires a significant amount of industry- and process-specific knowledge. The experience of early adopters' attempts to adapt general-purpose systems to the unique needs of firms (Bresnahan and Greenstein 1996) can potentially be shared within firms in related industries. As a result, the training and on-the-job experience obtained by IT workers will be most valuable to firms within the same industry. Furthermore, restrictions related to NCAs apply only to movement of workers to competitors. So, any heterogeneity in the effects of NCAs are most likely to be dependent on the same-industry scale, and they will be less influenced by the scale of other industries in the same location. This leads to the prediction that increases in worker mobility will have a more negative effect on adoption behavior when the establishment is in a location with a large number of establishments in the same industry.

Industry experience with related technologies. We also explore heterogeneity in the effects of worker mobility based upon industry use of related technologies. In particular, we examine differences in industry-level experimentation with predictive analytics during our sample period, based upon measures reported in Brynjolfsson, Jin, and McElheran (2021). The outside option will be higher for workers in industries that are actively using predictive analytics, increasing the risks to making investments in ML that will train workers who may go to competing firms in the same industry. Industry-level differences may also capture establishment level differences in the net benefits to adoption, which may also influence sensitivity to changes in NCA on adoption behavior.

MOBILITY AND NONCOMPETE AGREEMENTS

The prior section established how changes in worker mobility can shape business process innovation related to ML and the adoption of ML. However, measuring the mobility of workers engaged in such business process innovation and estimating a causal relationship between their movement and adoption is difficult. First, public and private data available from the Census Bureau such as Quarterly Workforce Indicators (QWI) measure aggregate mobility within a region, but do not identify occupation (Balasubramanian et al. 2020), making it impossible to identify the set of workers engaged in business process innovation and their movements. Although private data sets employed in recent studies (Tambe 2014; Tambe and Hitt 2014; Wu et al. 2018) do identify occupation, identifying a causal relationship between their movements and adoption is difficult as the benefits and costs of local labor mobility will be correlated with other local factors that influence the benefits of adoption. As a result, determining the impact of mobility on adoption requires an exogenous change to mobility across locations. In this study, we use plausibly exogenous state-level changes in the enforceability of NCAs to capture changes in the ease of worker mobility on adoption of ML.

NCAs are agreements between employers and employees that restrict employees from joining or starting a competing firm for a period of time, commonly around one or two years, after they leave their employer. The intent of these agreements, where enforceable, is to protect valuable investments made by a firm that may spill out to its competitors through worker mobility, and thus eventually encourage such investments by firms. The enforceability of these agreements in the United States varies across states; while most states allow "reasonable" restrictions, they are mostly or completely unenforceable in three states (California, North Dakota, and Oklahoma). Importantly, the enforceability of these agreements has changed over time in a number of states, either through legislative action or through judgments in courts.

We use changes in NCA enforceability as a measure of the state-level changes in the costs and benefits of labor mobility for several reasons. First, because they are shaped by legal and legislative changes, changes in the strength of NCAs are not likely to be correlated with IT investment. Second, NCAs are important restrictions that cover a significant proportion of the U.S. workforce. Starr (2019), using a survey of 11,000 workers, found that 38% workers had signed an NCA and that 19% were subject to an NCA by the time of the survey. Moreover, workers in knowledge-intensive positions are more likely to be subject to NCAs, such as workers in architecture, computers, and engineering (Starr 2019); CEOs (Garmaisse 2011); and inventors (Marx et al. 2009). Third, the enforceability of NCAs is a matter of state law (rather than federal law), which allows us to compare changes in adoption behavior in establishments that are in states that altered NCA enforceability with the corresponding changes among establishments in states that did not alter NCA enforceability.

Most importantly, existing studies across a range of settings have provided strong and compelling evidence that changes in the enforceability of NCAs influence worker mobility. Using U.S. patent data, Marx et al. (2009) show that strengthening the enforceability of NCAs reduces inventor mobility, particularly among those with firm-specific skills or in narrow technology fields. Balasubramanian et al. (2020), who looked at a recent ban on NCAs for technology workers in Hawaii, show that banning NCAs led to an 11% increase in mobility (and a 4% increase in new-hire wages) among technology workers (defined by industry) relative to other workers. Garmaise (2011) and Jeffers (2019) find a similar mobility-hindering effect among executives and knowledge workers, respectively. The latter study uses a similar set of legal changes as we adopt in this study and finds that following an increase in NCA enforceability, the total departure rate of employees drops by around 9%. Hence, we use changes in the legal enforceability of NCAs to proxy for changes in the ease of worker mobility. The next section provides further details on the specific changes considered.

DATA

Our primary source of data is the Aberdeen Computer Intelligence Technology Database (hereafter CI database). It contains information on establishment- and firm-level characteristics such as the number of employees, installations of IT software and hardware, and industry classification, among others. As one of the most comprehensive sources of micro-level IT investment, this dataset has been used by many researchers to study the adoption and economic implications of IT investments (e.g., Bloom et al. 2012; Bresnahan et al. 2002; Bresnahan and Greenstein 1996; Forman et al. 2005, 2012; Nagle 2019).

Historically, the CI data were collected by interview teams that surveyed establishments throughout the calendar year. However, beginning in 2017 the data collection methodology changed to one based upon evidence of technology usage and topical queries recorded online. For example, one source indicates that data are captured from over 1,000 websites that host content concerning technologies, including job boards, forums, tutorials, and educational sites (Levy 2019). For example, if a user at a company lists experience with a technology on her resume, or if a user is active on a tutorial or user forum associated with a technology, this is considered evidence of usage of the technology at the organization. While the CI database includes information on evidence of both current and expected use of technologies, we use only evidence on current usage. In this way, the approach is similar to that used by recent authors to detect IT investment and use based on online employee resumes and job advertisements (Goldfarb et al 2019; Tambe and Hitt 2014; Tambe et al. 2019).

Prior researchers have compared the set of establishments in the CI database that have more than 100 employees and demonstrated that it is broadly representative of the U.S. economy (Forman et al. 2002), with a slight oversampling of large establishments and technology-intensive industries. We similarly compared our establishment-level sample to U.S. Census County Business Patterns (CBP) data and found it similar to the U.S. CBP in terms of industry and geography, though like prior versions of the database, it seems to have slightly oversampled large establishments (Table B1).

The focus of our research is to understand how changes in the ease of mobility for workers, as proxied by changes to state-level enforceability of NCAs, influences the diffusion of ML technology among businesses. Thus, we require data on adoption of ML over time. Our decision regarding the sample period is shaped by several factors. First, since NCA changes occur over a period of years, we must use a sample period that allows sufficient time for these changes to occur and for them to influence adoption. Similarly, our sample must end sufficiently late so that we can observe adoption of ML: for example, advanced data analytics enabled by ML entered the market starting only around 2015 (see, for example, Sallam et al. 2017). Third, as noted above, Aberdeen changed its data collection strategy recently, making the intervening years difficult to compare. Given these constraints, in our analysis we examine adoption over two distinct years, 2010 and 2018, studying the adoption decisions in 2018 from a base of zero (no adoption in 2010).

Given our research design, we require establishments that appear both in 2010 and 2018 in the data. We identify 686,878 such establishments. We exclude government, military, nonprofit organizations (including elementary and high school education and libraries), and agriculture because the relationship between ML adoption and labor mobility for these organizations is likely to be different than for nonfarm businesses.[1] Prior research that has used establishment-level CI data to examine adoption and economic implications of frontier IT adoption has focused on larger establishments because of low rates of adoption of frontier IT among small establishments and also because of potential measurement error (e.g., Forman et al. 2005, 2008, 2012). Following that research, we also exclude small establishments (fewer than 50 employees) because of low adoption rates and the risk that we may be unable to correctly observe ML adoption given Aberdeen's data collection methodology. The final baseline sample contains 306,208 establishment-year observations.

Dependent Variable

[1] The observations we excluded are: Public administration (SIC 90-99); Agriculture, forest, and fishing (SIC 01-09), Elementary and secondary schools (SIC 8211); Colleges and universities (8221), Junior Colleges and Technical Institutes (8222), and Libraries (8231); and some establishments affiliated with a county-, city-, or state-level government.

Machine Learning Analytic Software Adoption

Our interest is primarily in understanding the implications of worker mobility on the downstream innovation required to deploy a GPT. Thus, to isolate the implications of this downstream innovation from other types of innovation that could occur when building new software, we focus our analysis on the adoption of packaged software that incorporates ML technology.

Specifically, we measure ML adoption based on whether an establishment adopts enterprise data analytics software that incorporates ML technology. Analytics software incorporating ML functionality enables new applications by facilitating prediction (Agrawal et al. 2018, 2019). The adoption of such advanced analytic tools focused on predictive analysis is thus different from traditional data analytics tools that focus on descriptive analysis.

We identify ML adoption based on the functionality of the software applications adopted by establishments. The CI data report the vendor and product name installed by the establishment. Based on intensive research on the functionality of the application packages, we identify 31 packages as incorporating ML technology. We create a dummy variable for *ML adoption* at the establishment level that is equal to 1 if the establishment has adopted one of these packages and 0 otherwise. The overall ML adoption rate in our sample is 9.7% in 2018. Details of our coding approach are provided in Appendix A.

Independent Variables

Changes in Noncompete Enforceability

We identified 14 states that experienced a significant change in NCA enforceability based on the following four sources: (1) Ewens and Marx (2018), which provides a list of significant state-level changes used in their study period ending in 2014; (2) Beck Reed Riden LLP (2010, 2013, 2019), which provides a state-by-state snapshot of key aspects of noncompete enforceability such as whether they are permitted, whether there are any exemptions, etc.; (3) Malsberger et al. (2017), which contains the most comprehensive treatment available of noncompete enforceability; and (4) Jeffers (2019), which provides a list of nine state Supreme Court decisions between 2009 and 2013 that changed the enforceability of NCAs.

As we assume that changes would take some time to show any effects on technology adoption, we focus on significant changes during the period 2010 to 2017, a window that ends one year prior to the end of our analysis sample. We reviewed each of the four sources independently to identify relevant state-level changes. We classified each potential change into two categories, those that favored employers and those that favored workers (details of each change are provided in the Appendix). We then compared across sources to confirm the direction and significance of each of the identified changes. In a few cases where there appeared to be contradictions among sources, we relied on Malsberger et al. (2017), given its comprehensive treatment. The state-level changes developed using this process are shown in Table 1. If a state had changes in both directions or was inconsistent in some other way, we treated it as no change in the baseline estimation and performed robustness checks.

Other Controls

We include additional variables as controls to address potential heterogeneity in establishment-level, firm-level, and local factors that can influence adoption.

Firm and Establishment Characteristics. The CI data include a range of information about the focal establishment and the firm that they belong to. Establishment size is measured using the number of employees; firm size is measured using the total number of establishments in the firm. Because of changes in measurement in these variables and since changes over time can be correlated with our dependent variable, we use the base-year values of these characteristics and interact with a 2018 dummy to identify their effects on adoption. We do not include technology controls in our regressions because in the base year of our data the CI database did not survey a significant percentage of establishments (they surveyed only 57%) on their software installations. However, we estimated the robustness of our results to using this smaller sample with technology controls and found our results to be qualitatively similar.

Local characteristics. To control for the intensity with which establishments in other technology-intensive industries are collocated with the focal establishment, we collect county-level high-tech industry employment data from the Quarterly Census of Employment and Wages (QCEW). We calculate the fraction of employment in high-tech industries at the county level. The definition of high-tech industries come from the U.S. Bureau of Labor Statistics.[2] Specifically, we identify the total employment in high-tech industries based on four-digit NAICS and compute a dummy to indicate whether the county ranks in the top quartile in the United States.

We also obtain information from the 2010 and 2018 American Community Survey to control for state-level demographic and economic factors that may affect IT adoption cost, including the percentages of the population that are aged 15 to 64, aged over 65, Black, and female. We also include controls for college or graduate school attendance rate among adults aged 18 to 24, the logarithm of state population, state GDP, and median household income.

We also identify other state-level policies and laws that might affect the employment and labor flows among firms in the observed period. Based on this, our controls include: whether the state has adopted public policy, implied contract, and good-faith exceptions to at-will

---

[2] https://www.bls.gov/opub/btn/volume-7/high-tech-industries-an-analysis-of-employment-wages-and-output.htm?view_full, Retrieved May 2020

employment (Autor et al. 2006); whether the state has adopted right-to-work laws (Starr et al. 2018); and the state-level top corporate income tax rates (Seegert 2012).

Table 2 presents descriptive statistics across the entire set of sample establishments, and Table 3 presents these results based on whether the state in which the establishment resides experienced a change in the strength of NCA enforceability during our sample period. All statistics use base year (2010) values.

IDENTIFICATION STRATEGY

Estimating Average NCA Effects on IT Adoption

Our primary specification examines changes to establishment-level adoption decisions of ML technology between 2010 and 2018. Our baseline empirical model takes the following form:

$$Adoption_{isj(18)} - Adoption_{isj(10)} = \beta_0 + \beta_1 NCA_{\mathrm{s}} + \beta_2 X_{isj} + \beta_3 Z_{isj} + \varepsilon_{isj} \qquad (1)$$

where $Adoption_{isjt}$ is the binary variable equal to 1 if establishment i in state s from industry j adopts ML at time t, and 0 otherwise. ML software only became widely available in commercial software in 2015 (Sallam et al 2017), so $Adoption_{isj(10)} = 0$. X_{isj} is a vector of controls at the establishment level. Z_{isj} is a vector of controls for local factors that influence the net benefits to IT adoption. Standard errors are robust and clustered at the state level.

Our estimation equation can be obtained from an underlying adoption model P($y^*_{ijst} > 0$) $= \beta_0 + \beta_1 NCA_{\mathrm{st}} + \beta_2 X_{ijst} + \beta_3 Z_{ist} + \omega_i + \gamma_{jt} + \varepsilon_{ijst}$, where y^* is some unobserved return function based on which firms decide to adopt. Here, the industry-year fixed effects control for time-varying factors at the industry level such as changes in output and input prices. Importantly, they include the price of the technology, which is very high initially so that there is no adoption in the first period. Taking the difference between the two periods (i.e.,$P(y^*_{is18} > 0) - P(y^*_{is10} > 0)$), taking the linear approximation of the underlying latent variable model, and allowing for slopes on X_{isj} and Z_{isj} to be flexible by year gets our estimating equation. Using this approach differences out the establishment fixed effect ω_i.

Consistent with prior literature that has examined the implications of changes in NCA enforceability, in our baseline model, we code changes in NCA enforceability using three levels (e.g., Ewens and Marx 2018; Garmaise 2011; Jeffers 2019): NCA_s represents the changes in NCA enforceability in state s, which is coded as 1 for a decrease (favoring employees), 0 for no change, and –1 for an increase (favoring employers). β_1 is our main coefficient of interest, which denotes the difference between the IT adoption rate in establishments that were exposed to a legal change in NCA enforceability relative to those located in a place without any NCA-related changes, after controlling for other factors. Specifically, if β_1 is positive (negative), it indicates that a loosening in NCA enforceability, which favors workers, is associated with an increase (decrease) in the likelihood of ML adoption in 2018. As noted in Table 1, our changes to NCA enforceability occur over a period of time between 2011 and 2016. This approach reflects potential delays in the effects of NCA on worker mobility, and in turn how changes to worker mobility would influence firm adoption decisions.

Our primary identification assumption is that there are no state- or local-level unobservables that are correlated with the incidence of changes in the strength of NCA enforceability and adoption of ML. In addition to controlling for a variety of state-level features, we probe the validity of this assumption in a number of ways. First, we estimate models that predict the likelihood of changes in NCA enforceability (Appendix Table B2) and find no correlation between common factors that we expect might be correlated with adoption and the likelihood of NCA enforceability. Second, we explore the effects of a pseudo-treatment analysis that randomly assigns NCA treatment to a similar-sized group of states as in the original regression, and we find that this alternative measure has no effect on adoption. As a falsification exercise, we examine the implications of NCA-related changes on the adoption of another technology that diffused over the same period but for which industry-specific human capital investments are likely to be lower and for which NCA-related changes are likely to be less important (touchscreen tablets) and show that changes in NCA enforceability have no significant effect on the adoption of this alternative technology.

We examine how the effects of NCA vary based upon the size of the establishment, the size of the local geographic region, the IT-intensity of the establishment's industry, and the presence of large establishments in the same region. These additional tests will provide additional confidence in a causal interpretation of our results if they are consistent with the framework developed in the prior section. In particular, if our results are influenced by unobserved heterogeneity, then the source of unobserved heterogeneity must act in a way that is consistent with our predictions related to how establishment size, industry, and location interact with NCA changes to shape adoption.

A common additional test in data such as ours would be to examine the effects of changes in NCA laws in periods prior to treatment. Several aspects of our data complicate the use of such an approach. First, as shown in Table 1, the overwhelming majority of our states make changes to NCA enforceability during or prior to 2015, the first year in which ML software is widely available in the marketplace (Sallam et al 2017). Only two states, Utah and Oregon, make changes to NCA enforceability later, in 2016. In this environment, any such "pretrend" test will have no power as no establishments will be at risk of adopting prior to the initiation of the policy change. Second, as noted earlier, there is a change in the data collection strategy in the CI database in 2017, which further complicates any attempts to conduct a pretrend analysis through the use of panel data.

RESULTS

We first investigate whether changes in NCA enforceability are associated with a higher likelihood of ML adoption. We then investigate differences in our results based upon the establishment size and the characteristics of the location in which the establishment is situated. We then explore robustness with respect to sample and specification. We finally explore robustness to two separate falsification exercises.

Baseline Results

Baseline Results of NCA Effects on ML Adoption

In Table 4, we present the baseline results, successively adding controls in each column so that column (6) presents the estimates from the two-period adoption model with the full set of controls as specified in equation (1). Our specifications include an extensive set of controls, including all of those listed in Panel A of Table 2. To conserve space, however, we include only a subset of control variables in our tables. The results in Table 4 show that changes in NCA enforceability in favor of workers have a negative effect on adoption. Based on the baseline specification in column (6), a decrease in NCA enforceability (a change favoring workers) is associated with a 0.6 percentage point decline in the likelihood of ML adoption, which translates into a 6.2% percentage decrease (the adoption rate in 2018 is 9.7%).

Heterogeneous Effects of NCA on Adoption by Employment and Location Size

Table 5 presents split-sample results that estimate different NCA coefficients for large establishments (with more than 100 employees) versus small establishments (50–100 employees). We estimate separate regressions for the two groups to allow for unrestricted effects of covariates in differently scaled firms. Column (1) shows that NCA enforceability changes have no statistically or economically significant effect on adoption behavior for establishments with fewer than 100 employees. In contrast, column (2) shows that the magnitude of the NCA enforceability effect is larger for establishments with more than 100 employees; these establishments experience a 0.9 percentage point decline in adoption when changes in NCA enforceability favor workers (which translates to a 12.2% decline when compared with an adoption rate of 7.4%), an effect that is statistically significant at the 5% level. A test, based on seemingly unrelated regressions (Zellner 1962), of the difference between the estimates in column (1) and (2) shows they are statistically different at the 10% level. Appendix Table B2 examines the robustness of these results using additional establishment size categories. The results show that the negative effects of loosening NCA enforceability increase as the sample is increasingly restricted to establishments of larger size. For establishments with more than 800 employees, the NCA coefficient suggests that weakening NCA enforceability in favor of workers is associated with a 3.6 percentage point decline in the likelihood of ML adoption (or a 16.3% change). In short, the effects of changes in NCA enforceability appear to be greater for large establishments, in line with our predictions.

In Table 6, we compare the effects of changes in NCA enforceability for establishments located in large metropolitan statistical areas (MSAs, defined as populations over 1 million) with establishments in other locations. Column (1) shows that a change in NCA enforceability in favor of workers has a –0.8 percentage point effect on the likelihood of adoption for establishments located in large MSAs, which is statistically significant at the 5% level. This translates into a 7.1% decrease in the likelihood of adopting ML by establishments in those locations. In column (2), the coefficient of NCA enforceability changes in locations outside of a major MSA is almost zero and neither statistically nor economically significant. These estimates in columns (1) and (2) are statistically different from one another at the 10% level.

Table 6 provides evidence that the effects of changes in NCA enforceability is strongest for establishments located in large MSAs. However, our earlier discussion offers a sharper prediction: that the effects of NCA enforceability will be strongest in locations with a large number of large establishments in the same industry. These are the locations for which the risks of labor mobility to potential adopters of ML will be greatest. In Table 7 we explore this hypothesis in further detail, examining heterogeneity in our effects based on the number of large (more than 100 employees) establishments in the same four-digit SIC code and in the same MSA.

The number of large same-industry establishments may be correlated with other local characteristics—in particular, whether the focal establishment is in a large MSA. Accordingly, we control for location in a large MSA in many of our specifications. Because we are seeking to separately discern the effects of multiple local characteristics—including location size as well as the number of small and large establishments in the focal establishment's own industry—in this table we capture their simultaneous effects by using interaction terms rather than split samples.

Columns (1) and (2) explore the effects of interacting our NCA variable with (log of) number of establishments in the same location and four-digit SIC, providing a baseline for the effects of an increasing number of same-industry establishments. Increases in the total number of establishments has no economically or statistically significant effects on the marginal effect of a change in NCA enforceability, whether or not the regression specification controls for location in a large MSA. Column (3) presents a similar specification using the number of small establishments, and similarly shows that an increase in the number of small establishments similarly has no impact on the marginal effect of NCA enforceability.

Columns (4)–(6) represent the regression results that show the implications of increases in the number of large establishments in the same industry-MSA, progressively adding controls for the number of small establishments in the industry-MSA and a large MSA dummy (and their respective interactions with our NCA variable). These results show that increases in the number of large establishments strengthen the effects of NCA changes on adoption. For example, the results in column (6) (row: NCA Post x Log number of large establishments by MSA-SIC4 industry) suggest that a 10% increase in the number of large establishments will lead to an additional 3.1% decline in the likelihood of adoption when there is a change in NCA enforceability that favors workers. In short, the effects of NCA are stronger for establishments in locations where there are a large number of large establishments in the same industry-location. In Table B3 we show that all these results are robust to using an alternative threshold for large establishments (more than 50 employees).

In Table 8 we examine heterogeneity in the effects of NCA changes based upon the industry-level differences in the propensity to adopt predictive analytics technology. To do this, we identify lead user industries of predictive analytics based upon survey evidence of manufacturing establishments from the U.S. Census Bureau 2015 Management and Organization Practice Survey (MOPS) reported in Brynjolfsson, Jin, and McElheran (2021). We identify lead user industries as those that have average predictive analytics adoption greater than or equal that 0.75, based upon Figure 2 in that paper. As shown in Table 8, this roughly divides our sample of manufacturing establishments in half. The advantage of this measure is that it provides an industry-level measure of adoption of a closely related technology; one disadvantage is that it is available for manufacturing only. The evidence in table 8 shows that the effects of NCA changes on adoption are stronger in industries that are lead users of predictive analysis; establishments in other industry do not show statistically or economically significant changes in behavior in response to NCA.

Robustness Checks

A primary assumption necessary for identification of our model is that there are no unobservable factors that are changing over time in states experiencing NCA changes and that are simultaneously correlated with ML adoption. In our analyses in the prior section, we showed the robustness of our results to adding a variety of local controls. The robustness of our results to the addition of these controls suggests that the impact of unobservables would need to be large relative to observables to explain our results (Altonji et al. 2005). We probe this assumption further through additional analysis.

First, we explore the impact of NCA changes on adoption of a different technology that requires fewer industry-specific downstream investments and so should be less influenced by changes in NCA. To do this, we explore the effects of NCA changes on the adoption of tablets (touchscreens) within enterprises. Adoption of tablets should require fewer industry-specific investments to be used productively (Zolas et al. 2020) and diffused among businesses around the same time as ML. Apple launched the first-generation iPad in 2010, so we can assume that adoption among enterprises is zero as we do for ML. Column (2) of Table 9 shows that changes in NCA enforceability has little effect on adoption of tablets, consistent with our expectations.

We conduct a separate falsification exercise by examining the robustness of our results to randomization inference (Hess 2017). We follow Hess (2017) and randomly assign NCA-related treatment to a similar-sized group of states as that in our original regression, and collect estimates using equation (1) from 500 replications. The two-sided *p*-value of the test was 0.022, which suggests that it is extremely unlikely that our observed treatment effect is purely due to random chance.

We further probe whether state-level unobservables could influence our results by estimating a linear probability model to explain the factors influencing the likelihood of a change in NCA enforcement during our sample period. In Appendix Table B2 we run a descriptive regression of changes in NCA enforceability on state-level features such as population, GDP, and propensity that adult-aged workers are in the labor force, among others. These regressions show that the only variable that is consistently correlated with a change in NCA enforceability is the log of median household income. Therefore, we include this variable as a control in our regressions. In short, the results appear supportive of our identification assumption that there do not exist unobserved state-level factors that are correlated with NCA-related changes and ML adoption.

We also explore the robustness of our results to changes in specification and alternative samples. Following prior literature (e.g., Ewens and Marx 2018; Garmaise 2011; Jeffers 2019), our baseline results impose symmetry on the effects of NCA enforceability changes on adoption, assuming that strengthening and weakening NCA enforceability will have effects that are of similar magnitude but opposite signs. In Appendix Table B4 we probe the validity of this assumption and include two separate treatment variables: one for changes favoring workers and another for changes favoring employers. The results are in general consistent with those shown in Tables 4–6, in both direction and magnitudes—the size of effects for strengthening and weakening NCA enforceability are quite similar in the entire sample and when splitting our estimates by location size. However, we do observe some asymmetry in effects when splitting the sample by establishment size. In particular, for large establishments, the effects of NCA enforceability changes favoring employers are much larger (in absolute value) than for those favoring workers. The seemingly unrelated regressions (SUR) show that the two groups are statistically different in terms of the effects of NCA-related changes on ML adoption.

We explore the robustness of our results for different samples. One potential concern in our establishment-level analysis is that firm decisions to adopt ML in multi-establishment firms may be insensitive to local labor market factors. This could be the case if adoption depends on human capital located elsewhere in the firm (Forman et al. 2008) or if decisions to adopt new technology are made elsewhere in the organization (McElheran 2014). While these factors would likely depress our estimates of the effects of NCA enforceability on local adoption, we probe their salience further by estimating our baseline model on standalone establishments. While this cuts our sample size almost in half, column (3) of Table 9 shows that the point estimates are almost identical to those in the baseline sample. The implied percentage change in adoption likelihood becomes more negative (–19.5%) since the adoption rate is lower among standalone establishments. In Table B5 we also show that the results in Table 5 are robust to the use of different size thresholds; it further shows that the marginal effect of changes in NCA enforceability are stronger for very large establishments than for the 100-plus employees threshold that we use in Table 5.

As noted above, our identification of state-level changes to the enforceability of NCAs was based on recent research completed in law, finance, strategy, and economics. We probe the robustness of our results to changes for which there is some ambiguity in the direction of change. This occurred when, for example, changes to NCA enforceability within a state could affect workers both positively and negatively, or when there were multiple changes in a state during our sample period. Appendix C describes these changes in more detail, and Table 10 re-estimates our models when we make changes to our coding of the changes in states for which there is some ambiguity. Table 10 shows that our results are robust to these changes.

To examine whether our results are influenced disproportionately by changes in a specific state, we re-estimated the model in equation (1) excluding each of the states in the United States in turn. To examine whether our results are influenced by a specific industry, in Appendix Table B6 we estimate our results for manufacturing and non-manufacturing separately. We also re-estimated the baseline model excluding one industry at a time (these results are available from the authors upon request). All of our results remain robust to these changes.

DISCUSSION AND CONCLUSIONS

In this paper we explore the implications of labor mobility for adoption of a frontier GPT. Prior research in the information systems literature has focused upon the implications of how labor mobility into firms can increase their productivity and the benefits to new IT investments (Tambe 2014; Tambe and Hitt 2014; Wu et al. 2018). We highlight a second mechanism through which the presence of labor mobility influences the benefits of new IT. Specifically, we provide a framework for demonstrating how the presence of labor mobility can reduce the (net) benefits from investing in a new GPT.

We then bring this framework to the data by exploring business adoption of ML. Using exogenous changes in NCA enforceability as a shifter for labor mobility, we demonstrate that the likelihood of adopting ML will decline when NCA enforceability changes in a direction that will increase the incidence of labor mobility. These effects are strongest for establishments where the potential costs of worker mobility are highest; namely, larger establishments and those located in regions with many other large establishments from the same industry. In sum, by providing evidence on how labor mobility can negatively influence the benefits to a new GPT, we further our understanding of this important mechanism for the diffusion of new GPTs and the value that firms obtain from them.

Viewed from another perspective, our results also add to a literature that has explored how the benefits to new IT systems vary based on the size of the location (Dranove et al. 2014; Forman et al. 2005, 2012; Tambe 2014;). The contrast between these papers is informative in showing the additional nuance that our paper adds to the findings of this prior work. As was the case with this prior research, our work similarly suggests that adoption rates of a new technology are higher in large urban areas (see, for example, the Sizable MSA x Year results in Table 7), likely owing to the range of complementary assets that are available. However, the presence of many establishments from the same industry increases the risks of labor mobility. Viewed within the context of the urban economics literature, these results highlight a limitation through which Marshallian (same-industry) agglomeration economies can influence the benefits to new IT adoption. Viewed from a managerial perspective, our results point to potential risks for establishments that seek to invest in new IT and that are in locations with many other firms from the same industry.

Our research also contributes to a small but growing literature that has recently examined the implications of NCA enforceability for firm outcomes such as investment in R&D (Conti 2014), physical capital (Garmaise 2011; Jeffers 2019), and worker training (Starr 2019). While those papers similarly show that NCA enforceability can influence investment, we extend that literature by exploring a different and important context, *viz.* investments in early-stage GPTs. Perhaps closest to our context is Rock (2019). While not a primary focus of that paper, Rock (2019) does not find evidence that changes to NCA enforceability are associated with changes in the number or wages of engineering employees in his sample of firms. This may be because his firm-level data use the location of the headquarters establishment to measure the implications of NCA changes. Moreover, changes in NCA enforceability are expected to directly affect the mobility of workers, not necessarily the equilibrium number of employees in a firm. In general, our focus on establishment-level data and the characteristics of the location in which the establishment resides allows a more direct measurement between NCA enforceability and outcomes and allows us to test hypotheses not investigated in prior research.

As with other prior papers that have studied NCA enforceability, our paper also has policy implications, particularly with regard to adoption and diffusion of advanced technologies. Our results suggest that worker mobility affects the adoption and diffusion of GPTs that require significant investments in downstream innovation, especially during the early stages of the technology. Thus, our study offers another important dimension that policy makers must consider when evaluating the societal impact of NCA enforceability.

We also advance recent research that has sought to better understand the factors influencing the diffusion of artificial intelligence and ML technology. Our approach of measuring ML adoption through the adoption of ML-enabled business applications software complements other approaches that use alternative data sources such as the job postings in Burning Glass or member profile information from LinkedIn (Goldfarb et al. 2020; Rock 2019) or through direct survey of the use of ML/AI through confidential Census data (Zolas et al. 2020). Given the continuing policy interest in the diffusion of AI-related technologies among businesses (e.g., United States Executive Office of the President, 2016), our strategy for measuring ML adoption across a large sample of establishments is of independent interest. Further, while ML in many ways represents a transformative GPT (Cockburn et al. 2018;

Goldfarb et al. 2019; Trajtenberg 2019), our results, while we do not emphasize them, show that at least some of the factors influencing adoption bear similarities to other forms of IT, including the effects of establishment and firm size, industry, and location in which the organization is situated (Fichman 2000; Forman et al. 2005).

Limitations

Our results also have limitations that suggest avenues for future research. For one, we rely on a recent literature that has demonstrated how NCA enforceability influences worker mobility (Garmaise 2011; Marx et al. 2009)—in particular, recent work that has explored how recent changes to NCA enforceability influence mobility of technology and knowledge workers (Balasubramanian et al. 2020; Jeffers 2019). However, unlike some of the studies in that literature (e.g., Marx et al. 2009), we are unable to measure worker movements in our sample directly. Though this limits our ability to understand which workers may play a role in driving our results, in this regard, our study is similar to recent research that has explored the implications of NCA enforceability on firm-level outcomes such as R&D investment and investment in equipment (Conti 2014; Garmaise 2011; Jeffers 2019). Future research can potentially combine information on workers with information on ML adoption to develop a more fine-grained view of the impact of labor mobility.

Although informative, our measurement of adoption also has some limitations. For example, we can capture a product's adoption but know little about its intensity of usage. For example, we do not observe the number of users or the degree to which insights derived from ML influence business operations. If these quantities are differentially affected by labor mobility relative to adoption, then the insights that we can draw from our analysis will be limited.

Evaluating the impact of labor mobility on intensity of usage can be a fruitful avenue for future research.

As noted, a critical assumption in our research is that there exist no unobserved state-level factors that are changing in ways that are correlated with NCA enforceability changes and technology adoption. We have probed the credibility of this assumption through the exploration of different controls, subsamples, and measurement strategies and through falsification exercises. Our results have proven robust to all these changes, increasing confidence in the results and in our interpretation. However, we leave it to other work to further probe the robustness of our findings.

Our research starts from the premise that investments in AI require complementary inputs to be deployed successfully. It seeks to unpack how mobility in these complementary human capital inputs shapes adoption of AI in the short run. In the long run, the implications of labor mobility may be different, as the number of adopters increases, the need for downstream innovation in the deployment of ML declines, and the value of the human capital accrued during new ML implementation decreases. One avenue for research would be to examine if the interplay between local labor market features and investment in ML is different in the short run than in the long run.

More broadly, it is worthwhile noting that the focus of our paper is on how changes in labor mobility influence short run ML adoption decisions, not the implications of these decisions for labor demand. While it is possible that adoption of ML may lead to organizational changes that could reduce the demand for labor in the long run, we do not advance recent work that seeks to examine the labor demand implications of ML investments (e.g., Acemoglu and Restrepo

2018, 2019; Acemoglu et al 2010; Frey and Osbourne 2017; Felten, Raj, and Seamans 2018; Brynjolfsson, Mitchell, and Rock 2020).

Our research has taken a first step toward understanding how labor mobility can depress incentives to invest in a new GPT. However, there remain many ways to build on these results. There are opportunities to study where workers of GPT adopters come from and move to, as well as to study how the patterns we have investigated evolve over time, just to name a couple. We hope our research encourages additional work in this important area.

References

Acemoglu, Daron, David Autor, Jonathan Hazell, and Pascual Restrepo. 2020. AI and Jobs: Evidence from Online Vacancies. Working Paper.

Acemoglu, Daron and Pascual Restrepo. 2018. The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. American Economic Review 108(6): 1488-1542.

Acemoglu, Daron and Pascual Restrepo. 2019. Artificial Intelligence, Automation, and Work. In The Economics of Artificial Intelligence: An Agenda, eds. Ajay Agrawal, Joshua Gans, and Avi Goldfarb. Chicago: University of Chicago Press, p. 197-236.

Agrawal, A., Gans, J., and Goldfarb, A. 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence.* Cambridge, MA: Harvard Business Press.

Agrawal, A., Gans, J. S., and Goldfarb, A. 2019. "Exploring the Impact of Artificial Intelligence: Prediction Versus Judgment," *Information Economics and Policy* (47), pp. 1-6.

Altonji, J. G., Elder, T. E., and Taber, C. R. 2005. "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools," *Journal of Political Economy* (113:1), pp.151-184.

Angrist, J. D., and Pischke, J. S. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion.* Princeton, NJ: Princeton University Press.

Autor, D. H., Donohue III, J. J., and Schwab, S. J. 2006. "The Costs of Wrongful-Discharge Laws," *Review of Economics and Statistics* (88:2), pp. 211-231.

Balasubramanian, N., Chang, J. W., Sakakibara, M., Sivadasan, J., and Starr, E. 2020. "Locked In? The Enforceability of Covenants Not to Compete and the Careers of High-Tech Workers," *Journal of Human Resources,* pp. 1218-9931R1.

Beck Reed Riden LLP. 2010. *Employee Noncompetes: A State by State Survey*.

Beck Reed Riden LLP. 2013. *Employee Noncompetes: A State by State Survey*.

Beck Reed Riden LLP. 2019. *Employee Noncompetes: A State by State Survey*.

Bloom, N., Sadun, R. and Van Reenen, J. 2012. "Americans Do IT Better: US Multinationals and the Productivity Miracle," *American Economic Review*, (102:1), pp.167-201.

Bresnahan, T. F., Brynjolfsson, E., and Hitt, L. M. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence," *Quarterly Journal of Economics* (117:1), pp. 339-376.

Bresnahan, T., and A. Gambardella. 1998. "The Division of Inventive Labor and the Extent of the Market," in E. Helman (ed.), *General Purpose Technologies and Economic Growth*, Cambridge, MA: MIT Press, pp. 253-282.

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

Bresnahan, T. F., and Trajtenberg, M. 1995. "General Purpose Technologies 'Engines of Growth'?" *Journal of Econometrics* (65:1), pp. 83-108.

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

Brynjolfsson, E., Jin, W., and McElheran, K. 2021. "The Power of Prediction: Predictive Analytics, Workplace Complements, and Business Performance. Available at SSRN: https://ssrn.com/abstract=3849716.

Brynjolfsson, E., and McElheran, K. 2016. "The Rapid Adoption of Data-Driven Decision-Making," *American Economic Review* (106:5), pp. 133-39.

Brynjolfsson, E., and Milgrom, P. 2013. "Complementarity in Organizations," *Handbook of Organizational Economics*, R. Gibbons and J. Roberts (eds.), Princeton, NJ: Princeton University Press, pp. 11-55.

Brynjolfsson, E., Mitchell, T., and Rock, D. 2018. "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" *AEA Papers and Proceedings* (108), pp. 43-47).

Brynjolfsson, E., Rock, D., and Syverson, C. 2021. "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," *American Economic Journal: Macroeconomics* (13:1), pp. 333-372.

Chang, Y. B., and Gurbaxani, V. 2012a. "The Impact of IT-Related Spillovers on Long-Run Productivity: An Empirical Analysis," *Information Systems Research* (23:3-part-2), pp. 868-886.

Chang, Y. B., and Gurbaxani, V. 2012b. "Information Technology Outsourcing, Knowledge Transfer, and Firm Productivity: An Empirical Analysis," *MIS Quarterly* (36:4), pp. 1043-1063.

Cheng, Z., and Nault, B. R. 2007. "Industry Level Supplier-Driven IT Spillovers," *Management Science* (53:8), pp. 1199-1216.

Cheng, Z., and Nault, B. R. 2012. "Relative Industry Concentration and Customer-Driven IT Spillovers," *Information Systems Research* (23:2), pp. 340-355.

Cockburn, I. M., Henderson, R., and Stern, S. 2019. "The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis," in *Economics of Artificial Intelligence: An Agenda*, A. Agrawal, J. Gans, and A. Goldfarb (eds.), Chicago, IL; University of Chicago Press, pp. 115-146.

Conti, R. 2014. "Do Non-competition Agreements Lead Firms to Pursue Risky R&D Projects?" *Strategic Management Journal,* (35:8), pp.1230-1248.

Conti, R., Gambardella, A., and Novelli, E. 2019. "Specializing in General Purpose Technologies as a Firm Long-Term Strategy," *Industrial and Corporate Change* (28:2), pp. 351-364.

Dranove, D., Forman, C., Goldfarb, A., and Greenstein, S. 2014. "The Trillion Dollar Conundrum: Complementarities and Health Information Technology," *American Economic Journal: Economic Policy* (6:4), pp. 239-270.

Ewens, M., and Marx, M. 2018. "Founder Replacement and Startup Performance," *Review of Financial Studies* (31:4), pp. 1532-1565.

Felten, E.W., Raj, M. and Seamans, R., 2018, May. A method to link advances in artificial intelligence to occupational abilities. In *AEA Papers and Proceedings* (Vol. 108, pp. 54-57).

Fichman, R. G. 2000. "The Diffusion and Assimilation of Information Technology Innovations," in *Framing the Domains of IT Management: Projecting the Future Through the Past,* Zmud, R. (ed.) Pinnaflex Educational Resources, pp. 105-128.

Forman, C., Goldfarb, A., and Greenstein, S. 2002. "Digital Dispersion: An Industrial and Geographic Census of Commercial Internet Use," National Bureau of Economic Research (No. w9287).

Forman, C., Goldfarb, A., and Greenstein, S. 2005. "How Did Location Affect Adoption of the Commercial Internet? Global Village vs. Urban Leadership," *Journal of Urban Economics* (58:3), pp. 389-420.

Forman, C., Goldfarb, A., and Greenstein, S. 2008. “Understanding the Inputs into Innovation: Do Cities Substitute for Internal Firm Resources?” *Journal of Economics & Management Strategy* (17:2), pp. 295-316.
Forman, C., Goldfarb, A., and Greenstein, S. 2012. “The Internet and Local Wages: A Puzzle,” *American Economic Review* (102:1), pp. 556-75.
Frey, C.B. and Osborne, M.A., 2017. The future of employment: How susceptible are jobs to computerisation?. *Technological forecasting and social change*, *114*, pp.254-280.
Garmaise, M. J. 2011. “Ties That Truly Bind: Noncompetition Agreements, Executive Compensation, and Firm Investment,” *Journal of Law, Economics, and Organization* (27:2), pp. 376-425.
Goldfarb, A., Taska, B., and Teodoridis, F. 2019. “Could Machine Learning Be a General-Purpose Technology? Evidence from Online Job Postings,” available at SSRN: https://ssrn.com/abstract=3468822
Goldfarb, A., Taska, B., and Teodoridis, F. 2020. “Artificial Intelligence in Health Care? Evidence from Online Job Postings,” *AEA Papers and Proceedings* (110), pp. 400-404.
Greenwood, B., Ganju, K., and Angst, C. 2019. “How Does Implementation of Enterprise Information Systems Affect a Professional’s Mobility? An Empirical Study,” *Information Systems Research* (30:2), pp. 563-594.
Hess, S. 2017. “Randomization Inference with Stata: A Guide and Software,” *Stata Journal* (17:3), pp. 630-651.
Ito, H. 1995. “The Structure of Adjustment Costs in Mainframe Computer Investment,” Working Paper, Stanford University.
Jeffers, J. 2019. “The Impact of Restricting Labor Mobility on Corporate Investment and Entrepreneurship,” available at SSRN 3040393.
Klepper, S. 1996. Entry, Exit, Growth, and Innovation over the Product Life Cycle. *The American Economic Review* (86:3), pp. 562-583.
LePlante, A. 2020. “What is Augmented Analytics? Powering Your Data with AI,” O’Reilly Media.
Levy, M. 2019. Aberdeen Behavioral Technographics. retrieved Nov 2020, from https://gzconsulting.org/2019/10/07/aberdeen-behavioral-technographics/
Malsberger, B., Carr, D. J., Pedowitz, A. H., and Tate, E. A. 2017. *Covenants Not to Compete: A State-by-State Survey,* 11th edition, Arlington, VA: Bloomberg BNA.
Marx, M., Strumsky, D., and Fleming, L. 2009. “Mobility, Skills, and the Michigan Non-Compete Experiment,” *Management Science* (55:6), pp. 875-889.
Mata, F.J., Fuerst, W.L. and Barney, J.B. 1995. “Information Technology and Sustained Competitive Advantage: A Resource-based Analysis,” *MIS Quarterly* (19:4), pp. 487-505.
McElheran, K. 2014. “Delegation in Multi-establishment Firms: Evidence From IT Purchasing,” *Journal of Economics & Management Strategy*, (23:2), pp. 225-258.
Melville, N., Kraemer, K. and Gurbaxani, V. 2004. “Information technology and organizational performance: An integrative model of IT business value,” *MIS Quarterly* (28:2)*,* pp. 283-322.
Nagle, F. 2018. “Learning by Contributing: Gaining Competitive Advantage Through Contribution to Crowdsourced Public Goods,” *Organization Science* (29:4), pp. 569-587.
Neffke, F. M., Otto, A., and Weyh, A. 2017. “Inter-industry Labor Flows,” *Journal of Economic Behavior & Organization* (142), pp. 275-292.

Nilsson, N. 2010. *The Quest for Artificial Intelligence: A History of Ideas and Achievements*, Cambridge: Cambridge University Press.
Rock, D. 2019. "Engineering Value: The Returns to Technological Talent and Investments in Artificial Intelligence," available at SSRN 3427412.
Rosenberg, N. 1982. *Inside the Black Box: Technology and Economics*, Cambridge: Cambridge University Press.
Sallam, R., Howson, C., Idoine. 2017. Augmented Analytics Is the Future of Data and Analytics. *Gartner Report.*
Seegert, N. 2012. "Optimal Taxation with Volatility: A Theoretical and Empirical Decomposition," working paper, University of Michigan, Ann Arbor.
Siegel, R. E., and Makinen, J. 2018. "C3 IoT: Enabling Digital Industrial Transformation," Stanford Business School, Case Reference no. SM307.
Starr, E. 2019. "Consider This: Training, Wages, and the Enforceability of Covenants Not to Compete," *ILR Review* (72:4), pp. 783-817.
Starr, E., Balasubramanian, N., and Sakakibara, M. 2018. "Screening Spinouts? How Noncompete Enforceability Affects the Creation, Growth, and Survival of New Firms," *Management Science* (64:2), pp. 552-572.
Tambe, P. 2014. "Big Data Investment, Skills, and Firm Value," *Management Science* (60:6), pp. 1452-1469.
Tambe, P. 2019. "From Data to Decisions: Domain Knowledge and the Machine Learning Workforce." Working Paper, University of Pennsylvania.
Tambe, P., and Hitt, L. M. 2014. "Job Hopping, Information Technology Spillovers, and Productivity Growth," *Management Science* (60:2), pp. 338-355.
Tambe, P., Hitt, L. M., Rock, D., and Brynjolfsson, E. 2019. "IT, AI and the Growth of Intangible Capital," available at SSRN 3416289.
Trajtenberg, M. 2019. "Artificial Intelligence as the Next GPT," in *Economics of Artificial Intelligence: An Agenda*, A. Agrawal, J. Gans, and A. Goldfarb (eds.), Chicago, IL; University of Chicago Press, pp. 175-186.
United States Executive Office of the President. 2016. *Artificial Intelligence, Automation, and the Economy.*
Wu, L., Hitt, L., and Lou, B. 2020. "Data Analytics, Innovation, and Firm Productivity," *Management Science* (66:5), pp. 1783-2290.
Wu, L., Jin, F., and Hitt, L. M. 2018. "Are All Spillovers Created Equal? A Network Perspective on Information Technology Labor Movements," *Management Science* (64:7), pp. 3168-3186.
Zellner, A. 1962. "An Efficient Method of Estimating Seemingly Unrelated Regressions and Tests for Aggregation Bias," *Journal of the American Statistical Association* (57:298), pp. 348-368.
Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D. N., Buffington, C., Goldschlag, N., Foster, L., and Dinlersoz, E. 2020. "Advanced Technologies Adoption and Use by US Firms: Evidence from the Annual Business Survey," National Bureau of Economic Research (No. 28290).

Table 1. Changes in NCA enforceability

State	Case/Code	Effective Date /Decision Date
Changes favoring employers (coded as -1)		
Arkansas	Ark. Code 4-75-101	7/22/2015
Colorado	*Lucht's Concrete Pumping, Inc. v. Horner*	5/31/2011
Georgia	Restrictive Covenants Act	5/1/2011
Texas	*Marsh v. Cook*	12/16/2011
Virginia	*Assurance Data Inc. v. Malyevac*	9/12/2013
Wisconsin	*Runzheimer International v. Friedlen*	4/30/2015
Changes favoring workers (coded as 1)		
Hawaii	Haw. Rev. Stat. Sec. 480-4(d)	7/1/2015
Kentucky	*Creech v. Brown*	6/9/2014
Montana	*Wrigg v. Junkermier, Clark, Campanella, Stevens*	11/22/2011
New Hampshire	N.H. Rev. Stat. Ann. Sec. 275-70	7/12/2012, 7/28/2014
New York	*Brown & Brown v. Johnson*	6/11/2015
Oregon	ORS 653.295	1/1/2016
Pennsylvania	*Socko v. Mid-Atlantic Systems*	11/18/2015
South Carolina	*Poynter v. Century Builders*	5/24/2010
Utah	Utah Codes 34-51-101 to 34-51-301	5/10/2016

Table 2. Descriptive statistics

Variable	Obs.	Mean	SD	Min	Max
Panel A: Summary statistics of variables					
Machine learning adoption in 2018	153,104	0.097	0.297	0.000	1.000
Log number of site employees	153,104	4.719	0.797	3.932	10.309
Log number of sites in the enterprise	153,104	2.183	2.072	0.693	10.390
Top quartile county high-tech employment fraction	153,104	0.751	0.433	0.000	1.000
Dummy for having public policy exceptions to at-will employment	153,104	0.824	0.381	0.000	1.000
Dummy for having implied contract exceptions to at-will employment	153,104	0.784	0.412	0.000	1.000
Dummy for having good-faith exceptions to at-will employment	153,104	0.244	0.429	0.000	1.000
Right-to-work law	153,104	0.374	0.484	0.000	1.000
Top corporate tax rate	153,104	6.755	2.719	0.000	12.000
State log total employment in private sectors	153,104	15.145	0.860	12.510	16.484
State log # of establishments in private sectors	153,104	12.473	0.903	9.964	14.103
State log total wages in private sectors	153,104	25.882	0.942	23.155	27.367
State log GDP	153,104	13.022	0.937	10.244	14.537
State log population	153,104	16.021	0.897	13.244	17.435
State percent age 65+	153,104	0.131	0.017	0.077	0.173
State percent age 15-64	153,104	0.671	0.012	0.641	0.746
State percent Black	153,104	0.129	0.082	0.004	0.516
State percent female	153,104	0.509	0.005	0.480	0.528
State log medium household income	153,104	10.825	0.142	10.515	11.140
State percent age 18-24 enrolled in college	153,104	0.433	0.043	0.275	0.573
Panel B: Machine learning adoption by NCA group					
NCA changes favoring employers	26,960	0.109	0.312	0.00	1.00
No change	100,669	0.096	0.295	0.00	1.00
NCA changes favoring workers	25,475	0.089	0.285	0.00	1.00

Note: Unless otherwise indicated, in Panel A all values are from 2010.

Table 3. Mean comparison by NCA group

Variable	NCA change favoring employers (-1)	No NCA change (0)	NCA change favoring workers (1)
Percent machine learning adoption in 2018	10.9%	9.6%	8.9%
Log number of site employees	4.724	4.713	4.740
Log number of sites in the enterprise	2.320	2.163	2.117
Top quartile county high-tech employment fraction	0.739	0.757	0.739
Dummy for having public policy exceptions to at-will employment	0.843	0.863	0.647
Dummy for having implied contract exceptions to at-will employment	0.843	0.786	0.714
Dummy for having good-faith exceptions to at-will employment	0.046	0.355	0.016
Right-to-work law	0.759	0.333	0.130
Top corporate tax rate	3.717	7.345	7.642
State log total employment in private sectors	15.381	15.084	15.137
State log # of establishments in private sectors	12.613	12.443	12.447
State log total wages in private sectors	26.101	25.816	25.914
State log GDP	13.248	12.959	13.032
State log population	16.257	15.964	16.000
State percent age 65+	0.114	0.134	0.139
State percent age 15-64	0.674	0.670	0.673
State percent Black	0.148	0.125	0.125
State percent female	0.506	0.509	0.511
State log medium household income	10.817	10.828	10.821
State percent age 18-24 enrolled in college	0.399	0.438	0.450
Number of establishments	26,960	100,669	25,475

Note: Unless otherwise indicated, all values are from 2010. We conducted one-way analysis of variance to compare the means across the three NCA groups, and all variables listed in the table are significantly different across the three groups.

Table 4. Baseline results of NCA effects on ML adoption

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
NCA(-1,0,1) x Post	-0.008**	-0.009**	-0.004*	-0.004**	-0.003***	-0.006**
	(0.003)	(0.003)	(0.002)	(0.002)	(0.001)	(0.002)
Log number of site employees in 2010 x Year		0.041***	0.021***	0.021***	0.021***	0.021***
		(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Log number of sites in the enterprise in 2010 x Year			0.050***	0.050***	0.049***	0.049***
			(0.001)	(0.001)	(0.001)	(0.001)
Top quartile county high-tech employment fraction x Year				0.016***	0.016***	0.016***
				(0.001)	(0.001)	(0.001)
Establishments	153,090	153,090	153,090	153,090	153,090	153,090
R^2	0.588	0.593	0.638	0.638	0.638	0.638
Other laws	N	N	N	N	Y	Y
Demographic controls	N	N	N	N	N	Y
Economic controls	N	N	N	N	N	Y

Notes: Columns 5 and 6 include controls for other laws, including a dummy for having public policy exceptions to at-will employment, a dummy for having implied contract exceptions to at-will employment, a dummy for having good-faith exceptions to at-will employment and right-to-work laws. Column 6 includes demographic controls for state-level log population, percent age 65+, percent age 15-64, percent Black, log of medium household income, and percent age 18-24 enrolled in college. Robust standard errors clustered by state are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 5. Effects of NCA enforceability on ML adoption by establishment employment size

VARIABLES	(1) 50–99 employees	(2) 100+ employees	Test of differences (SUR) Chi-square	p-value
NCA(-1,0,1) x Post	-0.002	-0.009**	3.262*	0.071
	(0.002)	(0.004)		
Log number of site employees in 2010 x Year	0.008*	0.028***	14.890***	0.000
	(0.004)	(0.003)		
Log number of sites in the enterprise in 2010 x Year	0.050***	0.049***	0.948	0.330
	(0.001)	(0.001)		
Top quartile county high-tech employment fraction x Year	0.013***	0.017***	2.067	0.151
	(0.001)	(0.002)		
Establishments	78,082	74,974		
R^2	0.637	0.644		
Mean adoption rate in 2018	0.0745	0.1210		

Notes: All regressions include controls listed for column (6) of Table 4. Robust standard errors clustered by state are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 6. Heterogeneous effects of NCA enforceability on ML adoption by geographical location size

VARIABLES	(1) Sizable MSA (with over 1m population)	(2) Other locations	Test of differences (SUR) Chi-square	p-value
NCA(-1,0,1) x Post	-0.008**	-0.000	3.680*	0.055
	(0.003)	(0.002)		
Log number of site employees in 2010 x Year	0.025***	0.013***	15.932***	0.000
	(0.003)	(0.002)		
Log number of sites in the enterprise in 2010 x Year	0.051***	0.046***	20.744***	0.000
	(0.001)	(0.001)		
Top quartile county high-tech employment fraction x Year	0.014***	0.012***	0.133	0.715
	(0.002)	(0.002)		
Establishments	88,175	64,873		
R^2	0.641	0.642		
Mean adoption rate in 2018	0.1093	0.0812		

Notes: All regressions include controls listed for column (6) of Table 4. Robust standard errors clustered by state are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 7. Heterogeneous effects of NCA on adoption by industry-location size

	(1)	(2)	(3)	(4)	(5)	(6)
Location (MSA) heterogeneity	\	Sizable MSA	\	\	\	Sizable MSA
Industry by location (MSA) heterogeneity	Total establishments	Total establishments	Small establishments	Large establishments	Small & large establishments	Small & large establishments
NCA(-1,0,1) x Post	-0.004*	-0.003	-0.004*	-0.003	-0.005**	-0.004
	(0.003)	(0.003)	(0.003)	(0.002)	(0.003)	(0.003)
NCA Post x Log number of establishments by MSA-SIC4 industry	-0.000	-0.000				
	(0.001)	(0.001)				
Log number of establishments by MSA-SIC4 industry x Year	0.001	0.000				
	(0.001)	(0.001)				
NCA Post x Log number of small establishments by MSA-SIC4 industry			-0.000		0.001	0.001
			(0.001)		(0.001)	(0.001)
Log number of small establishments by MSA-SIC4 industry x Year			0.001*		0.001	0.001
			(0.001)		(0.001)	(0.001)
NCA Post x Log number of large establishments by MSA-SIC4 industry				-0.002*	-0.003**	-0.003**
				(0.001)	(0.001)	(0.001)
Log number of large establishments by MSA-SIC4 industry x Year				0.001	-0.000	-0.001
				(0.001)	(0.002)	(0.002)
NCA Post x Sizable MSA (with over 1m population)		-0.003				-0.003
		(0.004)				(0.004)
Sizable MSA x Year		0.007***				0.008***
		(0.002)				(0.002)
Establishments	153,090	153,090	153,090	153,090	153,090	153,090
R^2	0.638	0.638	0.638	0.638	0.638	0.638

Notes: All regressions include controls listed for column (6) of Table 4. Robust standard errors clustered by state are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 8. Heterogeneous effects of NCA enforceability on ML adoption by industry predictive analytics (PA) adoption intensity

VARIABLES	(1) Industry PA adoption rate >= 0.75	(2) Industry PA adoption rate < 0.75	Test of differences (SUR) Chi-square	*p*-value
NCA(-1,0,1) x Post	-0.009**	0.003	6.49**	0.0108
	(0.004)	(0.003)		
Log number of site employees in 2010 x Year	0.018***	0.016***	0.37	0.5455
	(0.004)	(0.003)		
Log number of sites in the enterprise in 2010 x Year	0.056***	0.028***	111.68***	0.000
	(0.002)	(0.002)		
Top quartile county high-tech employment fraction x Year	0.011***	0.001	3.51*	0.0611
	(0.005)	(0.002)		
Establishments	36,840	33,916		
R^2	0.633	0.589		
Mean adoption rate in 2018	0.1102	0.0403		

Notes: All regressions include controls listed for column (6) of Table 4. Robust standard errors clustered by state are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 9. Additional robustness checks: Alternative technologies and standalone establishments

VARIABLES	(1) Baseline – ML	(3) Adoption of tablets	(2) Standalone establishments – ML
NCA(-1,0,1) x Post	-0.006**	-0.003	-0.006***
	(0.002)	(0.003)	(0.002)
Log number of site employees in 2010 x Year	0.021***	0.258***	0.024***
	(0.002)	(0.003)	(0.001)
Log number of sites in the enterprise in 2010 x Year	0.049***	0.001*	
	(0.001)	(0.001)	
Top quartile county high-tech employment fraction x year	0.016***	0.010***	0.013***
	(0.001)	(0.003)	(0.001)
Establishments	153,090	153,090	74,081
R^2	0.638	0.738	0.549
Mean adoption rate in 2018	0.097	0.420	0.0308

Notes: All regressions include controls listed for column (6) of Table 4. Robust standard errors clustered by state are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 10. Robustness checks: Alternative NCA measures

NCA Measures	(1) Original measures	(2) AL = -1, (baseline =0)	(3) AL =1, (baseline =0)	(4) ID=-1, (baseline =0)
NCA x Post	-0.006**	-0.005**	-0.006**	-0.005**
	(0.002)	(0.002)	(0.002)	(0.002)
Log number of site employees in 2010 x Year	0.021***	0.021***	0.021***	0.021***
	(0.002)	(0.002)	(0.002)	(0.002)
Log number of sites in the enterprise in 2010 x Year	0.049***	0.049***	0.049***	0.049***
	(0.001)	(0.001)	(0.001)	(0.001)
Top quartile county high-tech employment fraction x Year	0.016***	0.016***	0.016***	0.016***
	(0.001)	(0.001)	(0.001)	(0.001)
Establishments	153,090	153,090	153,090	153,090
R^2	0.638	0.638	0.638	0.638

NCA Measures	(6) IL=1, (baseline =0)	(7) NV=1, (baseline =0)	(8) NY=0, (baseline =1)	(5) IL=-1, (baseline =0)
NCA x Post	-0.005**	-0.006***	-0.004**	-0.005**
	(0.002)	(0.002)	(0.002)	(0.002)
Log number of site employees in 2010 x Year	0.021***	0.021***	0.021***	0.021***
	(0.002)	(0.002)	(0.002)	(0.002)
Log number of sites in the enterprise in 2010 x Year	0.049***	0.049***	0.049***	0.049***
	(0.001)	(0.001)	(0.001)	(0.001)
Top quartile county high-tech employment fraction x Year	0.016***	0.016***	0.016***	0.015***
	(0.001)	(0.001)	(0.001)	(0.001)
Establishments	153,090	153,090	153,090	153,090
R^2	0.638	0.638	0.638	0.638

Notes: This table reports the effects of NCA changes using alternative NCA measures (descriptions are available in Appendix C). All regressions include controls listed for column (6) of Table 4. Robust standard errors clustered by state are in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.

APPENDIX

Appendix A. Construction of the dependent variable—ML adoption

Our process of identifying ML adoption in 2018 includes the following steps:

- *Match product name from CI data to vendor website:* Our first step was to match the product name in the CI database with that which appears in the vendor's website. Our goal was to identify whether the identified products incorporated ML functionality, according to the vendor's product description on its website. There were several special cases that we needed to consider. In particular, some products in the database did not list the exact product name (for example, in some cases the product was listed as "BI"). In these cases, we identified the product from the vendor for which the primary functionality was business intelligence or data analytics and used that product. Some other products experienced name changes over time from the time when the data were created (2018) to when we identified the product details on the web (2019). These name changes could be due to product upgrades or mergers and acquisitions of vendors. We thus assumed that products with older names were upgraded by vendors to the latest versions, since tech vendors regularly end premier support of their older version products to promote upgrades to the latest version. After these adjustments for each vendor-product, we were able to identify the associated product description during our web search in 2019.
- *Identifying products with ML functionality*: We searched product functionality from descriptions and product manuals on the vendor's website and conducted string matching using ML keywords. Our keywords were motivated by a similar procedure to identify ML in other settings such as patents (see, e.g., Cockburn el al. 2019) and included ML, neural networks, reinforcement learning, unsupervised learning, and machine intelligence. We also confirmed that the keywords were used to describe product features and were not used as context in the product description or manual (e.g., an example of using keywords as context would be noting that new technologies like ML were changing business but not describing a specific functionality of the product).

Appendix B: Supplementary Tables

Table B1: Distributions of CI data vs. the U.S. Census County Business Patterns data, 2010

	CI 2010 full	CBP 2010 full	CI 2010>100	CBP 2010 >100
Number of establishments	4,370,901	7,403,197	273,072	171,632
% MSA	85.2	93.8	91.8	95.2
% > 100 employees / % > 500 employees given have 100 employees	6.2	2.3	16.2	13.9
% Northeast	18.2	19.4	14.7	19.6
% Midwest	21.4	21.9	24.3	23.9
% South	37.7	35.2	23.9	35.6
% West	22.7	23.5	37.0	20.9
% Agriculture, Forestry, Fishing, and Hunting (NAICS = 11)	0.7	0.3	0.4	0.1
% Mining (NAICS = 21)	0.4	0.4	1.8	0.6
% Utilities (NAICS = 22)	0.5	0.2	0.6	0.8
% Construction (NAICS = 23)	4.5	9.2	2.7	3.7
% Manufacturing (NAICS = 31, 32, 33)	6.7	4.1	12.9	13.8
% Wholesale Trade (NAICS = 42)	4.1	5.6	3.7	4.6
% Retail Trade (NAICS = 44, 45)	7.7	14.4	25.3	15.2
% Transportation & Warehousing (NAICS = 48, 49)	3.3	2.8	2.8	4.1
% Media, Telecommunications, and Data Processing (NAICS = 51)	3.0	1.8	2.9	3.4
% Finance and Insurance (NAICS = 52)	7.5	6.4	3.0	4.7
% Real Estate and Rental and Leasing (NAICS = 53)	2.7	4.7	0.9	1.0
% Professional, Scientific, and Technical Services (NAICS = 54)	10.0	11.5	8.3	5.9
% Management of Companies and Enterprises (NAICS = 55)	0.8	0.7	0.2	3.3
% Administrative and Support and Waste Management and Remediation Services (NAICS = 56)	5.3	5.2	3.0	9.3
% Educational Services (NAICS = 61)	6.1	1.2	8.0	2.8
% Health Care and Social Assistance (NAICS = 62)	23.8	11.0	10.0	17.0
% Arts, Entertainment, and Recreation (NAICS = 71)	1.0	1.7	1.1	2.1
% Accommodation and Food Services (NAICS = 72)	1.7	8.7	2.0	5.4
% Other Services (except Public Administration) (NAICS = 81)	4.1	9.8	4.7	2.2
% Government (NAICS = 92)	6.0	0.3	5.9	0.0

Table B2. Predicting changes in NCA enforceability

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	NCA Enf. Up NCA changes favoring employers (NCAPost = –1)			NCA Enf. Down NCA changes favoring employers (NCAPost = 1)		
State Republicans to Democrats ratio		-0.040	-0.036		0.045	0.054
		(0.084)	(0.085)		(0.100)	(0.100)
State Labor Force (rate)	0.015	0.021	0.022	-0.036*	-0.038	-0.035
	(0.018)	(0.019)	(0.020)	(0.021)	(0.023)	(0.023)
State Unemployment (rate)	-0.017	-0.027	-0.020	-0.057	-0.071	-0.057
	(0.036)	(0.038)	(0.040)	(0.043)	(0.046)	(0.048)
Uniform Trade Secrets Act (UTSA)			-0.130			-0.287
			(0.238)			(0.281)
State log median household income	-0.248	-0.148	-0.112	0.913*	1.349*	1.429**
	(0.454)	(0.569)	(0.577)	(0.542)	(0.679)	(0.683)
State log population	0.131	0.417	0.479	0.374	0.844	0.980*
	(0.243)	(0.448)	(0.465)	(0.290)	(0.534)	(0.550)
State log GDP	-0.012	-0.283	-0.359	-0.400	-0.852	-1.021*
	(0.251)	(0.441)	(0.467)	(0.299)	(0.527)	(0.552)
Observations	51	49	49	51	49	49
R^2	0.116	0.140	0.146	0.100	0.123	0.144

Notes: Dependent variable NCA Enf. Up (Down) is an indicator variable equal to 1 if state experienced an increase (decrease) in NCA enforceability between 2010 and 2018. Washington, DC, and Nebraska are excluded when including controls of Republicans to Democrats ratio.

Table B3. Heterogeneous effects of NCA on adoption by industry-location size: Alternative threshold (50 employees) for large vs. small establishments

VARIABLES	(1)	(2)	(3)	(4)
NCA(-1,0,1) x Post	-0.004*	-0.003	-0.004*	-0.003
	(0.002)	(0.002)	(0.002)	(0.002)
NCA Post x Log number of small establishments by MSA-SIC4 industry	-0.000		0.001	0.001
	(0.001)		(0.001)	(0.001)
Log number of small establishments by MSA-SIC4 industry x Year	0.001*		0.002	0.001
	(0.001)		(0.001)	(0.002)
NCA Post x Log number of large establishments by MSA-SIC4 industry		-0.001	-0.002**	-0.002**
		(0.001)	(0.001)	(0.001)
Log number of large establishments by MSA-SIC4 industry x Year		0.001	-0.001	-0.002
		(0.001)	(0.002)	(0.002)
NCA Post x Sizable MSA (with over 1m population)				-0.003
				(0.004)
Sizable MSA x Year				0.008***
				(0.002)
Establishments	153,090	153,090	153,090	153,090
R^2	0.638	0.638	0.638	0.638

Notes: Notes: All regressions include controls listed for column (6) of Table 4. Robust standard errors clustered by state are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table B4. Robustness checks: Asymmetric NCA coefficient

	(1)	(2)	(3)	(4)	(5)
		By employment size		By location	
VARIABLES	All	Below 100 employees	Above 100 employees	Sizable MSA (with over 1m population)	Other locations
NCA changes favoring employers	0.005*	-0.002	0.011**	0.010**	-0.003
	(0.003)	(0.003)	(0.004)	(0.004)	(0.003)
NCA changes favoring workers	-0.006**	-0.005**	-0.007	-0.006*	-0.003
	(0.003)	(0.002)	(0.005)	(0.004)	(0.003)
Log number of site employees in 2010 x Year	0.021***	0.008*	0.028***	0.025***	0.013***
	(0.002)	(0.004)	(0.003)	(0.003)	(0.002)
Log number of sites in the enterprise in 2010	0.049***	0.050***	0.049***	0.051***	0.046***
x Year	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Top quartile county high-tech employment fraction	0.016***	0.013***	0.017***	0.014***	0.012***
x Year	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)
Establishments	153,090	78,082	74,974	88,175	64,873
R^2	0.638	0.637	0.644	0.641	0.642

Notes: All regressions include controls listed for column (6) of Table 4. Robust standard errors clustered by state are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table B5. Robustness checks of effects of NCA enforceability by employment size

VARIABLES	(1) All (50+ employees)	(2) 100+ employees	(3) 200+ employees	(4) 400+ employees	(5) 800+ employees
NCA(-1,0,1) x Post	-0.006**	-0.009**	-0.011*	-0.023***	-0.036**
	(0.002)	(0.004)	(0.006)	(0.009)	(0.017)
Log number of site employees in 2010 x Year	0.021***	0.028***	0.030***	0.027***	0.035***
	(0.002)	(0.003)	(0.004)	(0.006)	(0.011)
Log number of sites in the enterprise in 2010	0.049***	0.049***	0.048***	0.048***	0.044***
x Year	(0.001)	(0.001)	(0.002)	(0.003)	(0.004)
Top quartile county high-tech employment fraction x year	0.016***	0.017***	0.017***	0.022***	0.042***
	(0.001)	(0.002)	(0.004)	(0.007)	(0.011)
Establishments	153,090	74,974	31,719	12,267	4,895
R^2	0.638	0.644	0.656	0.683	0.707
Mean adoption rate in 2018	0.097	0.121	0.155	0.202	0.221

Notes: All regressions include controls listed for column (6) of Table 4. Robust standard errors clustered by state are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table B6. Robustness checks: Manufacturing and non-manufacturing

VARIABLES	(1) Manufacturing	(2) Non-manufacturing
NCA(-1,0,1) x Post	-0.005***	-0.006**
	(0.001)	(0.003)
Log number of site employees in 2010 x Year	0.022***	0.022***
	(0.003)	(0.002)
Log number of sites in the enterprise in 2010 x Year	0.047***	0.050***
	(0.001)	(0.001)
Top quartile county high-tech employment fraction x Year	0.005	0.019***
	(0.003)	(0.001)
Establishments	37,059	116,031
R^2	0.613	0.644
Mean adoption rate in 2018	0.0772	0.1039

Notes: All regressions include controls listed for column (6) of Table 4. Robust standard errors clustered by state are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Appendix C: NCA Enforceability Changes Considered Only in Robustness Checks

Alabama (2016): Effective 1/1/2016, Ala. Code 8-1-193 was amended to permit judicial reformation of covenants overbroad as written (Malsberger et al. 2017, pp. 45, 1299). However, this was part of Ala. Code 8-1-190 to 8-1-197, which replaced the old Ala. Code 8-1-1. These codes seem to have some worker-favorable features such as consideration (Malsberger et al. 2017, p. 1318) and presumptions of reasonableness (Malsberger et al. 2017, p. 1321). Moreover, it appears that judicial reformation was the norm prior to the repeal (Malsberger et al. 2017, p. 1333). Given that the changes had features favorable and unfavorable to workers, we set the baseline to ***0*** and performed a robustness check with values –1 and +1.

Idaho (2016) *as favoring employers:* In 2016, the Idaho legislature passed law HB 487, which adjusted Idaho's noncompete laws to say that if a "key employee…is in breach of an agreement, a rebuttable presumption of irreparable harm has been established." This effectively put the onus on the employee to prove they did not cause irreparable harm to the employer. However, shortly thereafter, and following some controversy, SB 1287 was introduced in 2018 to eliminate the language that was added through HB 487.[3] Furthermore, there was a decision in 2008 that favored employers (Ewens and Marx 2018), the effects of which could have lingered in the early years of the sample. Hence, we set the baseline to ***0*** and considered a robustness check as the change favoring employers.

Illinois (2011) *as favoring employers*: In *Reliable Fire Equipment Co v. Arredondo,* the state supreme court ruled that the enforceability of the employees' covenant not to compete should be judged by the three-prong test of reasonableness, of which the employer's legitimate business interest continues to be a part, and which looks to the totality of all of the circumstances, rather

[3] https://idahofreedom.org/sb-1287-non-compete-contracts/, retrieved Oct. 28, 2019.

than focusing on named specific factors (*Reliable Fire Equipment Co. v. Arredondo, 2011 IL 111871*). Thus, it possibly expanded the scope of legitimate business interest. In contrast to this, subsequently, in *Fifield v. Premier Dealer Services, Inc., 2013 IL App (1st) 120327*, the appellate court set a "bright line rule" that said a minimum of two years of continued employment is necessary to establish adequate consideration. However, this bright line rule does not appear to have been universally adopted. For instance, in *R.J. O'Brien & Associates, LLC v. Williamson*, the United States District Court for the Northern District of Illinois Eastern Division observed, "Indeed, some Illinois courts have adopted a two year bright line rule" but that "[o]ther courts, however, have rejected the two year bright line rule in favor of considering other factors in determining whether sufficient consideration was given to enforce a restrictive covenant." Hence, we set the baseline to ***0*** and considered robustness checks as the change favoring employers and favoring workers.

Nevada (2016) *as favoring workers:* In *Golden Rd v. Islam,* the state supreme court affirmed that if even one provision were invalid, the whole contract would be invalid. This would favor workers since employers would be hesitant to write overly broad contracts. However, this was superseded by Assembly Bill 276 (signed into law on 6/3/2017), which amended the law to allow courts to modify any unreasonable or overbroad restrictions.[4] Hence, we set the baseline to ***0*** and considered a robustness check as the change favoring workers.

New York (6/11/2015): In *Brown & Brown v. Johnson*, the court of appeals held that Florida law on restrictive covenants would violate New York public policy.[5] It also dismissed an overbroad restriction that prohibited the worker from working with any of the employer's customers,

[4] https://www.jacksonlewis.com/publication/new-law-brings-changes-nevada-s-non-compete-law, retrieved Oct. 28, 2019.
[5] https://law.justia.com/cases/new-york/court-of-appeals/2015/92.html, retrieved Oct. 28, 2019.

regardless of whether she had met them. Malsberger et al. (2017, p. 4063) note, "Following *BDO Seidman* [1999], and consistent with…*Brown & Brown*, NY courts have declined to partially enforce an overly broad noncompete provision." Our research and inputs from lawyers suggest that the impact of *Brown & Brown* was mainly clarificatory and marginal. Hence, we also test for robustness to treating this state as a "no change."