Received: 18 January 2022 | Revised: 19 January 2024 | Accepted: 27 January 2024

DOI: 10.1002/smj.3595

RESEARCH ARTICLE

SMS | Strategic Management Journal

WILEY

How does worker mobility affect business adoption of a new technology? The case of machine learning

Ruyu Chen[1] | **Natarajan Balasubramanian**[2] | **Chris Forman**[3]

[1]Stanford Institute for Human-Centered Artificial Intelligence, The Digital Economy Lab, Stanford University, Stanford, California, USA

[2]Whitman School of Management Syracuse University, Syracuse, New York, USA

[3]Dyson School of Applied Economics and Management, Cornell SC Johnson College of Business, Cornell University, Ithaca, New York, USA

Correspondence
Chris Forman, Dyson School of Applied Economics and Management, Cornell SC Johnson College of Business, Cornell University, Ithaca, NY 14853, USA.
Email: chris.forman@cornell.edu

Abstract

Research Summary: We investigate how worker mobility influences the adoption of a new technology using state-level changes to the enforceability of non-compete agreements as an exogenous shock to worker mobility. Using data on over 153,000 establishments from 2010 and 2018, we find that changes that facilitate worker movements are associated with a significant decline in the likelihood of adoption of machine learning. Moreover, we find that the magnitude of decline depends upon the size of the establishment, the extent of predictive analytics adoption in its industry, and the number of large establishments in the same industry-location. These results are consistent with the view that increases in outward worker mobility increase costs for adoption of a new technology that involves significant downstream investments in the early years of its diffusion.

Managerial Summary: Successful business adoption of new technologies such as machine learning requires skilled workers with experience in implementing those technologies. In the early years of technology diffusion workers in early adopting businesses typically acquire these skills through on-the-job learning that is paid for by the adopter. So, if such early adopters face an increased risk of those skilled workers quitting, then their incentives to adopt the technology decrease. We examine this possibility using changes in noncompete enforceability as a proxy for changes in worker mobility and find that the likelihood of adopting machine learning decreases as the risk of worker mobility increases, particularly for larger establishments, establishments in industries where adoption may be more beneficial and in locations with many large competing establishments.

KEYWORDS

business process innovation, human capital, machine learning, noncompete agreements, worker mobility

1 | INTRODUCTION

Recent research has highlighted the increasing diffusion of machine learning (ML) technologies among firms (Goldfarb et al., 2023, 2020; Zolas et al., 2020). This has sparked a growth in research on ML (Felten et al., 2021); especially, an interest in understanding how the value of such technologies is influenced by firm-level choices and the availability of firm-level complements (Bughin et al., 2019; Cao & Iansiti, 2021; Raj & Seamans, 2019). One topic of particular interest has been the link between the value of ML adoption and human capital (Allen & Choudhury, 2022; Brynjolfsson, Jin, & McElheran, 2021; Choudhury et al., 2020; McElheran et al., 2023).

Human capital investments in ML, and more broadly in many new information technologies, have two features with implications both for the costs of investments and for worker mobility. First, a large part of these human capital investments is vintage-specific (Barth et al., 2023; Chari & Hopenhayn, 1991; Choudhury et al., 2020) and is developed through on-the-job learning by workers. Hence, the value to firms from adopting the new technology depends in part on the availability of workers with the associated vintage-specific human capital ("VSHC") (Chari & Hopenhayn, 1991; Choudhury et al., 2020). Second, like prior vintages of business information technology (IT), adopting firms need to make large investments in complementary business process innovation ("BPI") (Bresnahan et al., 2002; Brynjolfsson et al., 2019; Brynjolfsson, Jin, & McElheran, 2021; Tambe, 2014). By their very nature, such innovations tend to be domain-specific and involve extensive on-the-job learning that is typically paid for by adopting firms.

These features of human capital investments associated with adoption of technologies such as ML create risks for firms that are considering adopting them. In particular, investments in adoption efforts enable workers involved in those efforts to gain human capital through on-the-job learning that is valuable to other firms in the industry, thus increasing the risk of outward worker mobility. In this regard, recent work has found that the risk of workers leaving the firm with valuable knowledge can depress incentives for firms to make investments that give rise to workers developing such capital (Agarwal et al., 2009; Conti, 2014; Kang & Lee, 2022; Starr, 2019). Thus, these studies would suggest that increased worker mobility may reduce the incentives to adopt new technologies that involve costly investments and require on-the-job learning. This difficulty is likely to be more severe for early adopters of new technologies such as ML when few experienced workers are available. To our knowledge, however, no studies have investigated this potential investment-dampening risk of mobility on technology adoption.

In this study, we develop a parsimonious model of adoption that incorporates these key features of technologies such as ML. Using this model, we predict that increasing mobility is likely to dissuade adoption, particularly for larger establishments that likely benefit from economies of scale, in industries where benefits to adoption are likely higher, and in larger locations where workers likely have lower costs to move.

Empirically, we focus on adoption of ML incorporated in enterprise business analytics software and measure its adoption in 2018 based on data in the Aberdeen Computer Intelligence (CI) database from over 150,000 establishments with 50 or more employees. Rather than study the movements of workers directly, we use changes in state-level enforceability of noncompete agreements (NCAs) as a plausibly exogenous source of variation in worker mobility (Ewens & Marx, 2018; Marx et al., 2009). We then use a two-period long differences model to study how state-level changes in the strength of NCA enforceability between 2010 and 2018 influence the adoption of ML within establishments over the same period. Importantly, since ML adoption was not widespread by 2018, our analyses focus on short-run effects when the potential benefits of inward mobility to adopters may be small relative to the costs of outward mobility.

In line with our predictions, we find that increased worker mobility, as measured by a loosening of enforceability of NCAs, is associated with a 0.6 percentage point decline in the likelihood of ML adoption by establishments in our sample. Given an average adoption rate of 9.7% in 2018, this translates to a 6.2% decline in the likelihood of adoption on average. Our results are stronger among larger establishments and in industries that have been lead users of prediction technology, both of which are environments in which the benefits of ML adoption are likely to be high and so where establishments will be more affected by worker mobility. We also find that the effects of NCAs on technology adoption are greater in large urban areas, especially when the focal establishment is in a location with a greater number of large establishments in its own industry, consistent with NCAs deterring mobility to competitors.

Together, our theoretical framework and empirical results shed new light on an understudied effect of worker mobility on adoption of ML in the short run. We suggest that worker mobility influences ML adoption by affecting the costs of VSHC and complementary BPI. These costs characterize a range of present and past business IT, and so have implications for the diffusion of a broad range of technologies beyond ML. We add to recent studies of the relationship between information technologies and human capital at the firm level (e.g., Bresnahan et al., 2002; Nagle, 2018), particularly those related to ML (Alekseeva et al., 2020; Allen & Choudhury, 2022; Rock, 2021).

We contribute to research on the causes and implications of worker mobility. While worker mobility has been understood to shape product innovation (e.g., Almeida & Kogut, 1999; Conti, 2014; Singh & Agrawal, 2011), how it shapes the benefits and costs of BPI, especially innovation that accompanies adoption of software technologies, is only now beginning to be appreciated (Tambe, 2014; Tambe & Hitt, 2014; Wu et al., 2018). In this regard, while the focus has largely been on inward mobility, our study highlights the role of outward mobility. This focus gives rise to some distinct implications and relates our study to broader work on mobility and competition. Earlier work has suggested that increased mobility can erode a firm's first-mover advantages from technological leadership (Agarwal & Gort, 2001) through faster imitative entry. Our research highlights an additional implication of mobility: that it can dissuade early movers from making the investments that give rise to their technological leadership, in turn potentially slowing down such entry.

Finally, we contribute to a small but growing literature that on how NCA enforceability affect firm outcomes, such as investment in R&D (Conti, 2014), physical capital (Garmaise, 2011; Jeffers, 2024), and worker training (Starr, 2019). Our findings also have implications for firm choices regarding the use of NCAs. Among others, they suggest (indirectly) that large, early adopters in lead-user industries may find it more beneficial to adopt NCAs for their workers.

2 | THEORETICAL FRAMEWORK

We motivate our hypotheses using a stylized model of adoption. Although our exposition focuses on ML for the sake of simplicity and its match to our empirical context, the key features modeled here are found in many other new technologies, making the model broadly applicable beyond ML adoption. We first characterize the human capital investments required to adopt ML, and then discuss of the role of worker mobility before turning to the model.

2.1 | ML and human capital investments

Our focus in this study is on downstream industries and firms, rather than on digital firms for which ML is central to the product or service provided. Like earlier generations of IT, downstream firms adopting ML are likely to adopt packaged software solutions and then adapt such solutions to firm needs (Bresnahan & Greenstein, 1996). Hence, 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 (LaPlante, 2020).

Innovations in business IT such as ML have two distinguishing features. First, they involve both technological innovation and a process of innovation in downstream application sectors to develop new uses for existing technologies. In this vein, a vibrant literature has documented the complementary business adaptations necessary to adopt and obtain value from enterprise IT (Bresnahan et al., 2002; Bresnahan & Greenstein, 1996; Brynjolfsson, Rock, & Syverson, 2021), including the creation of new business processes, management practices, organizational structures, and investments in new human capital. These often take time and involve a substantial period of foregone productivity during which firm resources are committed to such adaptations (Brynjolfsson, Rock, & Syverson, 2021).

The second distinguishing feature, relevant for workers, is that the VSHC to implement such technologies cannot be easily acquired through academic or technical training and will typically need to be acquired on the job. Relatedly, for reasons noted above, learning at a "purely technological level" (Rosenberg, 1963) is less important than application-specific downstream innovation (Bresnahan, 2019). While workers in adopting firms may accumulate human capital related to a new technology, such capital will not be easily translated to new contexts, including to subsequent vintages of the technology, without again undergoing a similar process of innovation that is specific to a new downstream application sector, thus inducing a degree of industry- and vintage-specificity in worker human capital (Chari & Hopenhayn, 1991; Choudhury et al., 2020). Importantly, the vintage-specificity implies that there will be few workers with the relevant on-the-job experience during the early stages of technology adoption

(Chari & Hopenhayn, 1991), while industry-specificity means that sharing of technical advances through cross-industry mobility is likely to play a secondary role (Bresnahan, 2019).

Given the worker accumulation of VSHC and industry-specific human capital through on-the-job learning, when workers of the adopting firm move to new employers, particularly in the same sector, they can apply their experiences to those firms. While those destination firms benefit from the movement of these workers as highlighted in past studies (e.g., Marx et al., 2009), the possibility of employee movements out of their original firm will increase the original adopting firm's costs. Moreover, ML as a technology has yet to see widespread adoption in firms. Hence, adopters in this setting will typically be lead users who must make complementary investments in workers to deploy ML successfully and will face difficulty in acquiring skilled workers themselves. Thus, the role of mobility costs is likely to be particularly salient to such firms, at least in the short run when there are few trained workers. The costs of such departures will be particularly high when it is difficult to replace and retrain workers—for example, when departures involve team leaders who have worked on multiple aspects of the project and can create linkages across the project team (Davenport, 2000).

2.2 | Model setup

Having motivated how worker mobility can influence the costs of ML adoption, we turn to the formal model. Our model simplifies aspects of the adoption decision to focus on the key issues at stake in our study. While our model incorporates some important factors that have been emphasized in shaping adoption of new technology, such as firm and establishment size and industry as well as location characteristics (e.g., Forman et al., 2005; Griliches, 1957; Hannan & McDowell, 1984), it excludes others, such as the roles of network effects and competition (Karshenas & Stoneman, 1993; Kretschmer et al., 2012). Further, the model does not explicitly incorporate how the benefits and costs of adopting new IT such as ML evolve within firms over time (Brynjolfsson et al., 2019; Brynjolfsson, Rock, & Syverson, 2021). Rather, our model parsimoniously highlights the key assumptions necessary for mobility to influence ML adoption and how these assumptions interact with location and firm characteristics.

Consider the following model of a decision to adopt ML technology within a firm establishment's production processes[1]:

$$bq - \frac{m}{\theta}(1-t) \geq F \tag{1}$$

where $F \geq 0$ is some fixed cost of adopting ML, $m > 0$ is the cost associated with expected worker mobility out of the firm,[2] $\theta > 0$ is a scaling parameter that captures factors that influence the costs of mobility and how they decline over time, and b is a parameter that influences the relationship between size and the benefits to adoption. $q \geq 0$ is size (e.g., establishment output), and t is time, assumed to be continuous and between 0 and 1.

[1]The transition to ML-assisted production occurs at the production system level (Agrawal et al., 2021; Bresnahan, 2020). However, because it is difficult to observe underlying production systems within firms and because firms often include many production systems, we follow recent work and focus on the adoption decision at the establishment (or plant) level (Brynjolfsson & McElheran, 2016; Goldfarb et al., 2020).

[2]Note that we are agnostic about net mobility, that is, whether outward mobility is higher or lower than inward mobility. We only require that the net cost of outward mobility be positive (i.e., $m > 0$).

The first term (bq) is the net benefit of adopting ML, which is assumed to be increasing in size. There are several reasons for this assumption. Prior literature has shown that large establishments are more likely to have related technological complements and more mature processes that will increase the value of adopting enterprise software systems (Bresnahan & Greenstein, 1996; Brynjolfsson & McElheran, 2016). In the case of ML, large establishments may also have data assets that give them a scale advantage (Farboodi et al., 2019). Recent work finds that artificial intelligence facilitates scale advantages for the largest and most productive firms (Babina et al., 2024). Also, if process innovation through ML reduces unit costs, net benefits will be proportional to firm output (Klepper, 1996). Moreover, BPI that is embedded within complex business processes (associated with large establishments) may be more difficult for competitors to replicate (Mata et al., 1995), thus increasing payoffs to large establishments.

While the benefits to using ML may be increasing in scale, costs may rise as well. ML adoption in large establishments may require significant adjustment costs to integrate legacy systems (Bresnahan & Greenstein, 1996). Further, as noted above, adoption of ML will require integration with existing business processes. These adjustments will involve costly training of workers, which prior work suggests may be particularly costly in large firms (McElheran, 2015).

Although both benefits and costs are likely to increase with size, recent productivity studies have generally been supportive of the assumption that the *net* benefits of IT investments are increasing in size. Using firm-level data, Tambe and Hitt (2012) show that the returns to IT are smaller in midsize firms than those in the Fortune 500; however, they appear more slowly in large firms. More recently, Tambe et al. (2019) measure the intangible capital value of the complementary organizational adjustments and training described above. They find that the quantity of this digital capital is greatest in the highest decile by firm market value, suggesting that these investments are most highly valued in the largest firms.

The parameter b denotes the extent of net benefits that the technology offers for a given firm size. This may vary across industries. For instance, the net benefits of ML may be greater in contexts with digital outputs and so can be scaled at a low marginal cost. Thus, the benefits of scale for ML will be particularly large for firms in digital industries (Farboodi et al., 2019). Even for products and processes that are available or delivered physically, ML will be valuable when involved in processes in which the outcomes of ML prediction can be applied with little increase in marginal costs. For example, the benefits of a recommendation engine can be applied across many physical goods with little incremental cost (Bresnahan, 2020).

Beyond scalability, recent work has characterized settings where it is more difficult to apply ML to existing processes. For example, it may be more difficult to apply ML in settings where the costs of a prediction error are particularly high (Bresnahan, 2020). Further, interdependence of production processes may influence the costs of the organizational adjustments necessary for deploying ML (Agrawal et al., 2021; Bresnahan, 2020). These differences in the characteristics of the product or process are likely to translate to differences at the broader industry level.[3] Indeed, at the industry level, it is a well-known fact that adoption of new technologies often varies significantly across industries, as do the productivity benefits of technology adoption (e.g., Stiroh, 2002).

The sizes of b, q, and F provide motivation for why there may be adopters of ML even early in the diffusion of ML when (as we will detail below) mobility costs are high. For example, the

[3]We could directly hypothesize about the underlying product and process characteristics rather than about industries. However, in addition to limited data availability on the underlying characteristics, it is difficult to ex ante describe and specify the underlying characteristics of processes that give rise to the issues we describe.

early diffusion of ML technologies in marketing and consumer-facing applications has been argued to be an example of capital deepening of existing investments in web and mobile technologies (Bresnahan, 2020). In the context of our model, a large firm (high q) in a high b industry may find it optimal to adopt ML even when the costs are high.

The second term in Equation (1), $\left(\frac{m}{\theta}(1-t)\right)$, denotes the costs to acquire and retain the human capital necessary for implementing ML. Given the discussion above, we model these as a function of the costs of potential outward mobility and assume they will be positive, reducing the net benefits of ML adoption. In keeping with the vintage-specific nature of the human capital investments required to adopt ML (Chari & Hopenhayn, 1991; Choudhury et al., 2020), we posit that the cost of replacing departing workers declines over time as the number of workers with the requisite experience increases. Indeed, it is possible that after the technology diffuses further and trained workers are more widely available, mobility increases the net benefits of ML adoption, by lowering the costs of acquiring new skills (this can be reflected in our model by multiplying t by an additional parameter, $\omega > 0$; see Online Appendix D).

The parameter θ allows for potential differences across regions in labor market thickness that can affect the ease with which local employer-worker matches can be made. In large regions, thicker labor markets mean that it will be easier to find workers with the requisite skills elsewhere, and the costs of replacing departing workers will be lower (Forman et al., 2005, 2008). In the model, this is reflected in such regions having a higher θ.

2.3 | Hypotheses

In our model, the time at which a firm adopts ML is determined by when the benefits of adoption exceed the costs. Formally, based on Equation (1), this is given by:

$$t \geq t^* = 1 + \frac{\theta}{m}(F - bq) \tag{2}$$

To proceed, we make a simplifying assumption that all firms eventually adopt ML (i.e., at $t = 1$). This is merely to simplify the exposition here and can be relaxed without changing the inferences (see Online Appendix D). Then, from Equation (1), at $t = 1$, we must have $bq > F$ or, equivalently, $F - bq < 0$ for all q. We further assume that for all q, $bq < F + \frac{m}{\theta}$. Intuitively, this means that when $t = 0$, there are no firms that adopt instantaneously.

Taking the derivative of t^* with respect to m, we get $\frac{\partial t^*}{\partial m} = -\frac{\theta(F - bq)}{m^2} > 0$ since $F - bq < 0$. Thus, as worker mobility increases, adoption is delayed. Intuitively, as noted, this happens because at any given time, keeping other things constant, increasing mobility decreases the expected benefit from adopting. We state this as our first baseline hypothesis and then explore how this baseline prediction varies with size, industry, and location of the establishment:

Hypothesis 1. Increasing worker mobility will increase the time to ML adoption.

2.3.1 | Establishment size (q)

From Equation (2), $\frac{\partial t^*}{\partial q} = -\frac{b\theta}{m} < 0$. Thus, large establishments (i.e., those with higher q) are likely to adopt ML earlier. This follows straightforwardly from the assumption that benefits are increasing in size. However, the characteristics that make the benefits of ML adoption more valuable for larger establishments also make the costs of mobility higher.

To see this, consider the cross-partial derivative $\frac{\partial^2 t^*}{\partial m \partial q} = \frac{b\theta}{m^2} > 0$. Thus, increases in mobility reduce the adoption time to a greater extent for larger establishments. This arises because larger establishments are at greater risk of adopting ML, and so changes in mobility costs will influence their behavior more on the margin. We state this as the second hypothesis:

Hypothesis 2. Increasing worker mobility will increase the time to ML adoption more in larger establishments than in smaller establishments.

2.3.2 | Expected benefits of adoption (b)

The relationship between the scale of output q and the benefits of adoption may vary across contexts, depending upon the characteristics of the product or process. These ideas are expressed formally in our model, where $\frac{\partial t^*}{\partial b} = -\frac{q\theta}{m} < 0$ so that industries where the benefits are higher (i.e., those with higher b) tend to adopt ML earlier. The effects of worker mobility will be magnified in such industries. Formally, $\frac{\partial^2 t^*}{\partial m \partial b} = \frac{q\theta}{m^2} > 0$. To see the underlying intuition, consider an industry where the benefits to ML adoption are so low that few firms find it optimal to adopt; changing worker mobility in that industry will yield almost no effect on adoption. In contrast, mobility changes will have a larger impact in an industry where many firms find it optimal to adopt. In sum:

Hypothesis 3. Increasing worker mobility will increase the time to ML adoption more in industries where benefits of adoption are likely to be higher.

2.3.3 | Location size (θ)

As noted earlier, in large regions the costs of hiring workers who have obtained the requisite skills will be lower. In the model, this is reflected in such regions having a higher θ. Based on this, the direct effect (benefit) of location size can be seen by considering the partial derivative $\frac{\partial t^*}{\partial \theta} = \frac{(F - bq)}{m} < 0$ since $F - bq < 0$.

However, the effects of worker mobility will be magnified in larger locations. To see this formally, note that $\frac{\partial^2 t^*}{\partial m \partial \theta} = -\frac{(F - bq)}{m^2} > 0$ since $F - bq < 0$. Intuitively, this arises because in such regions, the larger number of workers and firms will make matching between worker skills and employer needs easier, thus reducing the costs to firms hiring workers from elsewhere. As a result, there are more outside options for workers who have accumulated ML-related human capital, depressing incentives to make investments in ML that will train workers. Thus, the costs associated with mobility are higher in large regions (i.e., those with higher θ).[4]

Hypothesis 4. Increasing worker mobility will increase the time to ML adoption more in larger locations than in smaller locations.

[4]In our empirical analysis we will probe this hypothesis further. The relevant labor market for acquiring workers may be jobs that exist within the same industry. Therefore, we will investigate whether it is increases in industry size, rather than location size per se, that most directly influence the relationship between mobility and adoption.

3 | MOBILITY AND NCAs

The prior section established how changes in worker mobility can shape adoption of ML. However, measuring the mobility of workers engaged in BPI and estimating a causal relationship between their movement and adoption is difficult as the benefits and costs of local worker mobility will be correlated with other local factors that influence the benefits of ML adoption. As a result, determining the impact of mobility on ML 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 ML adoption.

NCAs are agreements between employers and employees that restrict employees from joining or starting a competing firm for some time after they leave their employer. These agreements, where enforceable, intend to protect investments made by a firm that may spill out to its competitors through worker mobility, and thus encourage 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 (CA, ND, and OK). Importantly, the enforceability of these agreements has changed over time in many states, either through legislative action or through court judgments.

We use changes in NCA enforceability as a measure of the state-level changes in the costs and benefits of worker 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% of 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 (Garmaise, 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, 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. (2022), 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 (2024) 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.

That said, it is important to keep in mind that there are other links in the chain linking enforceability to mobility, including the strategic decision to use (and agree to) NCAs. Those links likely mute the relationship between enforceability and mobility by introducing noise. For instance, the propensity to use NCAs is higher among larger firms and in lead-user industries such as information; finance; professional, scientific, and technical services; and manufacturing (Starr et al., 2021). Such variations in NCA use may interact with changes in NCA enforceability to influence mobility. Moreover, prior research has shown that in addition to increasing mobility, decreasing NCA enforceability increases wages, although this effect is neither uniform nor theoretically unambiguous (e.g., Balasubramanian et al., 2022). However, such a wage increase is unlikely to occur without any accompanying increase in worker mobility, and thus NCAs are unlikely to affect ML adoption through wage increases alone.

4 | DATA

Our primary source of data is the Aberdeen CI Technology Database. It contains information on establishment- and firm-level characteristics such as 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 to study the adoption and implications of IT investments (e.g., Bloom et al., 2012; Bresnahan et al., 2002; Bresnahan & Greenstein, 1996; Forman et al., 2005, 2012; Nagle, 2018, 2019).[5]

Historically, the CI data were collected by interview teams that surveyed establishments. However, beginning in 2017 the data collection became based upon evidence of technology usage and topical queries recorded online. One source indicates that data are captured from over 1000 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 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 recent methods to detect IT investment and use based on online employee resumes and job advertisements (Goldfarb et al., 2023; Tambe et al., 2019; Tambe & Hitt, 2014).

We want to understand how changes in the ease of mobility for workers, as proxied by changes to state-level enforceability of NCAs, influence the diffusion of ML technology among businesses. Thus, we require data on adoption of ML over time. Our choice of sample period is shaped by several factors. First, since NCA changes occur over a period of years, we must use a 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, e.g., Sallam et al., 2017). Third, as noted, Aberdeen changed its data collection strategy recently, making the intervening years difficult to compare. Given these constraints, we examine adoption over two distinct years, 2010 and 2018, studying the adoption decisions in 2018 from a base of zero (i.e., no adoption in 2010). Similarly, in our empirical models, we use only 1 year (2010) to control for establishment characteristics. This mitigates the effects that changes in data collection strategy might have on our results.

[5]The database was previously known as the Harte Hanks Market Intelligence CI Technology database. It has since become known as the Aberdeen CI Technology Data Set and, more recently, the Aberdeen Technology Data Cloud. We retain the CI database label because it was used during part of our sample period and because it is the name that has been referenced in prior research projects.

Our research design requires establishments to appear in both 2010 and 2018 in the data. We identify 1,046,523 such establishments. We exclude government, military, nonprofit, and agriculture organizations (∼178,000 establishments) because the relationship between ML adoption and worker mobility for these organizations is likely to be different. Prior-related research using establishment-level CI data have 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 exclude small establishments with fewer than 50 employees (∼681,000 establishments) because of low adoption rates and the risk that we cannot accurately observe ML adoption given Aberdeen's recent data collection methodology that relies on the use of online signals to infer technology use. However, as we note later, our results are robust to including them. The final baseline sample contains 153,090 establishments. Online Appendix Table B1 provides details of the sample construction. Further, in Online Appendix Table B2, we compare our data to U.S. Census County Business Patterns data and find it similar in terms of industry and geography, though like prior versions of the database, the CI data seem to have slightly oversampled large establishments (e.g., Forman et al., 2002).

4.1 | Dependent variable: ML analytic software adoption

Our interest is in understanding the implications of worker mobility through the complementary innovation required to deploy a new technology. Thus, to isolate the implications of this innovation from other types of innovation that could occur when developing and implementing new software (e.g., development of algorithms), 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 predictive analysis (Agrawal et al., 2018, 2019) and hence is different from traditional data analytics tools that focus on descriptive analysis. To assess whether a software incorporates ML, we examined the functionality of that software and identified 31 packages that incorporate 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 and the list of 31 packages are provided in Online Appendix A.

4.2 | Key independent variable: 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 (2019), which provides a state-by-state snapshot of key aspects of noncompete enforceability, such as whether NCAs are permitted, whether there are any exemptions, and so forth; (3) Malsberger et al. (2017), which contains the most comprehensive treatment available of NCA enforceability; and (4) Jeffers (2024), 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 from 2010 to 2017, a window that ends 1 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 Online 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.

4.3 | Other controls

4.3.1 | 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 methods in these variables and since changes in these variables over time can be affected by our dependent variable (i.e., they would be "bad controls" per Cinelli et al., 2020), we use the

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

base-year values of these characteristics in our estimation. We also control for industry-level differences using four-digit SIC dummies.

We do not include technology controls in our baseline regressions because in the base year of our data (2010), the CI database surveyed only 43% of establishments on their software installations (e.g., database software) that are likely to directly influence the net benefits to ML adoption. However, we examined the robustness of our results to a model including controls for IT hardware investments and IT workers, which will capture differences in the general IT intensity of the establishment. Our results are robust to these additions.

4.3.2 | Local characteristics

We include several controls related to local characteristics. Because changes in these local factors during our sample period may influence adoption decisions and can potentially be misattributed to changes in NCA enforceability, these controls are entered as differences between 2010 and 2018. 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 and use this to indicate whether the county ranks in the top quartile in the United States.[6]

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 15–64 years, aged over 65, Black, and female. We include controls for college or graduate school attendance rate among adults aged 18–24 years, the logarithm of state population, state GDP, and median household income.

We control for other state-level policies and laws that might affect worker mobility among firms. Specifically, we control for whether the state has adopted public policy, implied contract, and good-faith exceptions to at-will employment (Autor et al., 2006); whether the state has adopted right-to-work laws (Starr et al., 2018); and the state top corporate income tax rate (Seegert, 2012).

Table 2 presents the descriptive statistics across the entire set of sample establishments. These include statistics related to adoption and establishment characteristics; additional descriptive statistics across our sample and based on whether the state in which the establishment resides experienced a change in the strength of NCA enforceability are included in Tables B3 and B4 in the Online Appendix. All statistics in these tables use base year (2010) values.

5 | IDENTIFICATION STRATEGY

5.1 | Estimating average NCA effects on IT adoption

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

[6]Specifically, we identify high-tech industries from Roberts and Wolf (2018) and compute the fraction of employment in such industries based on the four-digit NAICS.

TABLE 2 Descriptive statistics.

Variable	Obs.	Mean	SD	Min	Max
Panel a: Summary statistics of firm variables					
Machine learning adoption in 2018	153,090	0.097	0.297	0.000	1.000
Log number of site employees	153,090	4.719	0.797	3.932	10.309
Log number of sites in the enterprise	153,090	2.183	2.072	0.693	10.390
Panel b: Machine learning adoption by NCA group					
NCA changes favoring employers	26,959	0.109	0.312	0.00	1.00
No change	100,658	0.096	0.295	0.00	1.00
NCA changes favoring workers	25,473	0.089	0.285	0.00	1.00

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

$$\Delta Adoption_{isj} = \beta_0 + \beta_1 \Delta NCA_s + \beta_2 X_{isj(10)} + \beta_3 \Delta Z_{isj} + \varepsilon_{isj} \tag{3}$$

where $\Delta Adoption_{isj} = Adoption_{isj(18)} - Adoption_{isj(10)}$.[7] Here, $Adoption_{isj(t)}$ is equal to 1 if establishment i in state s from industry j adopts ML at time t, and 0 otherwise. ML only became widely available in commercial software in 2015 (Sallam et al., 2017), so $Adoption_{isj(10)} = 0$.ΔNCA_s reflects the change in NCA enforceability (described below). $X_{isj(10)}$ is a vector of establishment-level controls fixed at 2010 values, as noted earlier. ΔZ_{isj} is a vector of controls for local factors that influence the net benefits of ML adoption, entered as differences between 2010 and 2018. Standard errors are robust and clustered by state.

Our hypotheses predict how changes in NCA enforceability and other factors influence the time to adoption. However, we observe adoption only in 2018. Therefore, as an equivalent approach, we test our hypotheses by examining how mobility and its interaction with establishment and local factors influence adoption by 2018: when changes in these factors reduce the likelihood of adoption by 2018, we interpret this as an increase in the time to adoption.

We focus on establishment-level adoption because, as noted, it is closer to the underlying process or system over which ML adoption decisions will be made. Our focus on establishment-level data is consistent with recent attempts to measure adoption of ML and analytics (Brynjolfsson & McElheran, 2016; Goldfarb et al., 2020) as well as prior studies of enterprise software adoption using the CI database and other sources of business data available at the establishment level, such as U.S. Census data (e.g., Bresnahan & Greenstein, 1996; Forman et al., 2012; McElheran, 2015).

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 & Marx, 2018; Garmaise, 2011; Jeffers, 2024). 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 ML adoption rate in establishments that were

[7]Our estimation equation can be obtained from a fixed effects model in which adoption is a (linear) function of the level of NCA strength, establishment-level and local factors, and establishment and industry-year fixed effects that control for time-varying factors at the industry level such as changes in output and input prices. Importantly, the latter include the price of the technology, which is very high initially so that there is no adoption in the first period.

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. While this baseline specification assumes that the effects of enforceability increases and enforceability decreases are symmetric, in Online Appendix B we provide results in which we allow for these effects to be asymmetric and show that our qualitative results remain unchanged.

A key aspect of our specification is that it is a two-period difference model in which the treated states (i.e., where NCA laws changed between 2010 and 2018) are compared only to untreated states (i.e., those with no such changes) and not to other treated states. A common test with a method such as ours would be a "pre-trend analysis," which examines if the observed effects of changes in NCA laws on ML adoption manifest before these changes occur. Several aspects of our data preclude the estimation of pre-trends as well as estimating the implications of NCA changes on ML adoption in all years between 2010 and 2018. First, as shown in Table 1, the majority of our states made changes to NCA enforceability during or prior to 2015, the first year in which ML software was widely available in the marketplace (Sallam et al., 2017). Only two states, Utah and Oregon, made changes to NCA enforceability later, in 2016. In this environment, any pre-trend test would have little power, as few establishments would have risked adopting ML prior to the policy change. Second, as noted earlier, there was a change in the data collection strategy in the CI database in 2017. Hence, we focus our analysis on two-period difference models and probe robustness through other means.

This modeling approach also has implications for the use of recent innovations in the difference-in-difference (DiD) literature (e.g., Callaway & Sant'Anna, 2021; De Chaisemartin & D'Haultfoeuille, 2020; Goodman-Bacon, 2021). While the timing of treatments (i.e., changes in NCA enforceability) may vary across states, in our setup this does not influence the comparisons made to identify the DiD parameters: the identifying variation only compares treated states to never-treated states. Hence, we do not compare groups experiencing treatment with other groups that will or had already experienced treatment earlier in the sample (and are not experiencing it now) (Goodman-Bacon, 2021).

After establishing the baseline effect (Equation (1); Hypothesis 1), we examine how the effects vary based on establishment, industry, and geography characteristics (Hypotheses 2–4). Beyond verifying our predictions, these additional tests provide further 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.

6 | RESULTS

6.1 | Baseline results of NCA effects on ML adoption

In Table 3, we present the baseline results, successively adding controls in each column so that column (6) presents the estimates from the two-period difference 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 Table 2 and in Table B3 in the Online Appendix. To conserve space, we present only a subset of control variables in our tables. The results in Table 3 show that changes in NCA enforceability in favor of workers have a negative effect on adoption. Based on the

TABLE 3 Baseline results of NCA effects on ML adoption.

	(1)	(2)	(3)	(4)	(5)	(6)
NCA (−1,0,1)	−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		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			0.050	0.050	0.049	0.049
			(0.001)	(0.001)	(0.001)	(0.001)
Top quartile county high-tech employment fraction				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	.131	.141	.237	.237	.237	.237
Other laws	N	N	N	N	Y	Y
Demographic controls	N	N	N	N	N	Y
Economic controls	N	N	N	N	N	Y

Note: 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, and a dummy for having good-faith exceptions to at-will employment and right-to-work laws. Column (6) includes changes in demographic controls for changes in state-level log population, % age 65+, % age 15–64, % Black, log of medium household income, and % age 18–24 enrolled in college. Robust standard errors clustered by state in parentheses.

TABLE 4 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)	−0.002 (0.002)	−0.009 (0.004)	3.262	.071
Log number of site employees in 2010	0.008 (0.004)	0.028 (0.003)	14.890	.000
Log number of sites in the enterprise in 2010	0.050 (0.001)	0.049 (0.001)	0.948	.330
Top quartile county high-tech employment fraction	0.013 (0.001)	0.017 (0.002)	2.067	.151
Establishments	78,082	74,974		
R^2	.245	.240		
Mean adoption rate in 2018	0.0745	0.1210		

Note: All regressions include controls listed in column (6), Table 3. Robust standard errors clustered by state in parentheses.

baseline specification in column (6), a decrease in NCA enforceability (a change favoring workers) is associated with a 0.6 percentage point decline (p-value .015) in the likelihood of ML adoption, which translates into a 6.2% decrease (the adoption rate in 2018 is 9.7%). Thus, the results support Hypothesis 1.

While these results assume that increasing and decreasing NCA enforceability have symmetric effects (following the literature), we tested for asymmetric effects and found that the magnitude of the coefficients for changes favoring workers and employers are roughly equal (Online Appendix Table B5). Hence, and given its ease of interpretation, we use the symmetric specification for our baseline analyses and discuss asymmetric effects as robustness checks.

6.2 | Heterogeneous effects of NCA on adoption

Table 4 presents split-sample results that estimate different NCA coefficients for large establishments (with more than 100 employees) versus small establishments (with 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 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 show a 0.9 percentage point decline in adoption (p-value .027) when changes in NCA enforceability favor workers (which translates to a 7.4% decline when compared with an adoption rate of 12.1%). A test, based on seemingly unrelated regressions (Zellner, 1962), of the difference between the estimates in columns (1) and (2) shows they are different (p-value .071).

In Table 5, we test Hypothesis 3. We proxy the industry-level net benefits to ML adoption with the propensity to adopt predictive analytics technology. We identify industries that are lead users of predictive analytics based on survey evidence reported in Brynjolfsson, Jin, and McElheran (2021). We identify lead user industries as those that have average predictive

TABLE 5 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)	−0.009	0.003	6.49	.0108
	(0.004)	(0.003)		
Log number of site employees in 2010	0.018	0.016	0.37	.5455
	(0.004)	(0.003)		
Log number of sites in the enterprise in 2010	0.056	0.028	111.68	.000
	(0.002)	(0.002)		
Top quartile county high-tech employment fraction	0.011	0.001	3.51	.0611
	(0.005)	(0.002)		
Establishments	18,420	16,958		
R^2	.220	.160		
Mean adoption rate in 2018	0.1102	0.0403		

Note: Split on PA adoption rate based on Brynjolfsson, Jin, and McElheran (2021), fig. 2. All regressions include controls listed for column (6) of Table 3. Robust standard errors clustered by state are in parentheses.

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

VARIABLES	(1) Sizable MSA (with over 1 m population)	(2) Other locations	Test of differences (SUR) Chi-square	*p*-Value
NCA (−1,0,1)	−0.008	−0.000	3.680	.055
	(0.003)	(0.002)		
Log number of site employees in 2010	0.025	0.013	15.932	.000
	(0.003)	(0.002)		
Log number of sites in the enterprise in 2010	0.051	0.046	20.744	.000
	(0.001)	(0.001)		
Top quartile county high-tech employment fraction	0.014	0.012	0.133	.715
	(0.002)	(0.002)		
Establishments	88,175	64,873		
R^2	.230	.254		
Mean adoption rate in 2018	0.1093	0.0812		

Note: All regressions include controls listed in column (6), Table 3. Robust standard errors clustered by state in parentheses.

analytics adoption greater than or equal to 0.75, based on fig. 2 in that paper. As shown in Table 5, this divides our sample of manufacturing establishments roughly 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 only for manufacturing. Table 5 shows that the effects of NCA changes on adoption are stronger in industries that are lead users of predictive analysis and so is consistent with Hypothesis 3 (the *p*-value of the difference in coefficient estimates between lead user industries and others is .011).

In Table 6, we compare the effects of changes in NCA enforceability for establishments located in large metropolitan statistical areas (MSAs, large defined as populations over 1 million) with establishments in other locations. Column (1) shows that a change in NCA enforceability in favor of workers decreases the likelihood of ML adoption for establishments located in large MSAs by 0.8 percentage points (*p*-value .016). This translates to a 7.1% decrease in the likelihood of ML adoption by establishments in those locations. In column (2), the coefficient of NCA enforceability changes in locations outside of a major MSA is almost 0. These estimates in columns (1) and (2) are different from one another (*p*-value .055).

In our theoretical motivation, we argued that establishments in large regions will be more sensitive to NCA changes because of easier matching between worker skills and employer needs. This matching will be easier still when the large region includes many large establishments (relative to the case where large regions include many smaller establishments). Further, because employer-worker matching is easiest for worker movements within the same industry, and because NCA laws restrict worker movement within the same industry, the effects of NCA enforceability will be strongest in locations with many large establishments in the *same industry*.

In Table 7, we explore this prediction further, 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 local characteristics—in particular, whether the focal establishment is in a large MSA. Accordingly, we control for location in a large MSA in many specifications. Because we want 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 industry—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 have no economically 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 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 ML adoption. For example, the results in column (6) (row: NCA × Log number of large establishments by MSA-SIC4 industry, highlighted in bold) suggest that a 10% increase in the number of large establishments will lead to an additional 3.1% decline in the likelihood of ML adoption when there is a change in NCA enforceability that favors workers (*p*-value .018). In short, the effects of NCA are stronger for establishments in locations where there are more large establishments in the same industry-location.

6.3 | Robustness checks

We follow prior literature (e.g., Ewens & Marx, 2018; Garmaise, 2011; Jeffers, 2024) by imposing symmetry on the effects of NCA enforceability changes on ML adoption. Our tests in Online

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 and large establishments**	**Small and large establishments**
NCA (−1,0,1)	−0.004 (0.003)	−0.003 (0.003)	−0.004 (0.003)	−0.003 (0.002)	−0.005 (0.003)	−0.004 (0.003)
NCA × Log number of establishments by MSA-SIC4 industry	−0.000 (0.001)	−0.000 (0.001)				
Log number of establishments by MSA-SIC4 industry	0.001 (0.001)	0.000 (0.001)				
NCA × 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			0.001 (0.001)		0.001 (0.001)	0.001 (0.001)
NCA × Log number of large establishments by MSA-SIC4 industry				**−0.002** **(0.001)**	**−0.003** **(0.001)**	**−0.003** **(0.001)**
Log number of large establishments by MSA-SIC4 industry				0.001 (0.001)	−0.000 (0.002)	−0.001 (0.002)
NCA × Sizable MSA (with over 1 m population)		−0.003 (0.004)				−0.003 (0.004)
Sizable MSA		0.007 (0.002)				0.008 (0.002)
Establishments	153,090	153,090	153,090	153,090	153,090	153,090
R^2	.237	.237	.237	.237	.237	.237

Note: All regressions include controls listed for column (6) of Table 3. Robust standard errors clustered by state are in parentheses.

Appendix Table B5 support our approach, showing that the magnitudes for changes favoring workers and employers are roughly equal. However, for robustness, we also reestimate all the models in Tables 4–7 in which we allow for heterogeneity in establishment size, industry IT intensity, location size, and own industry size (Tables B6–B9). In general, the results are consistent with the directional predictions of our model, though in some cases the parameters are imprecisely estimated, and in one case the results are inconsistent with model predictions (the coefficient estimate in Table B7 for changes favoring workers in industries with PA adoption rate <0.75 is positive rather than negative). However, given the small number of states underlying some of the coefficients, we urge caution when interpreting results from these tables.

We then probe the robustness of our results to different strategies for measuring changes in the strength of NCA enforceability. We examine cases with 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. Online Appendix C describes these changes in more detail. The results in Table B10 show that our results are robust when we reestimate our models after making changes to our coding for the states in which there is some ambiguity.

Like many prior studies in this literature (e.g., Ewens & Marx, 2018), our baseline analyses do not consider differences in the intensity of change in NCA enforceability. We attempted a rough exploration based on the strength of NCA change. Table B11 presents the results of these explorations, showing that our baseline results are directionally robust to allowing parameter estimates for NCA changes to vary by our measure of change strength.

We now probe our identification assumptions. In Table B12, we examine the robustness of our baseline results to including controls for IT hardware as well as the fraction of IT employees and IT developers out of total establishment employment. Our results are robust to including these controls. In Table B13, we run a descriptive regression of changes in NCA enforceability on state-level features. The results appear supportive of our assumption.

As another way to assess the validity of our theoretical arguments, we explore the impact of NCA changes on adoption of a different technology that requires lesser 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, one of the technologies studied by Zolas et al. (2020) and one for which we have data. Adoption of tablets should require less VSHC and lower 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 (1) of Table B14 in the Online Appendix shows that changes in NCA enforceability have little effect on adoption of tablets, consistent with our expectations. We similarly examined the implications of changes in NCA on the changes in the number of IT workers and show in column (2) of Table B14 that changes in NCA enforceability in favor of workers are associated with a 4.0% decline in the number of IT workers.

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 .022, which suggests that it is unlikely that our observed treatment effect is due to random chance.

We explore the robustness of our results for different thresholds and samples. In Table B15, we include all establishments (column (1)) and find that the NCA coefficient is negative and larger in magnitude than our baseline result. In that table, we also provide a split-sample analysis by different size classes. Except for the smallest size class (which likely is the most susceptible to measurement errors), the magnitude of the effects is generally increasing in establishment size. Tables B16 and B17 show that the heterogeneous effects by establishment and industry-location size are robust to alternative size thresholds.

In Table B18, we split the sample into manufacturing and non-manufacturing and find the coefficients to be similar across these sectors. In Table B19, we restricted our sample to single-establishment firms only, and the results are qualitatively similar. We also confirmed the robustness of our baseline model to excluding one state at a time, and excluding one industry at a time (results available upon request). Last, we retested Hypotheses 3 and 4 using interactions rather than split samples and find qualitatively similar results (Table B20).

6.4 | Survey evidence

To shed additional light on the theoretical arguments underlying our model, we conducted a Qualtrics survey of about 200 "key team members" (i.e., project manager, technical lead, business lead, or a similar role) in large IT projects at for-profit companies. Of these, 80% had experience in new technologies (that companies in an industry have begun using only in the past 5 years), and another 12% were familiar with them. Broadly, we attempted to understand the (i) kinds of human capital relevant for such individuals; (ii) likelihood, drivers, and consequences of key team members voluntarily departing during the implementation of a project; and (iii) use and potential impact of NCAs. In addition, regarding the first two aspects, we attempted to understand the difference between new technologies and existing technologies. Below, we briefly discuss of the key findings and include details in Appendices E and F.

Focusing first on the types of human capital required, our theoretical arguments highlight the importance of on-the-job learning developed through working on projects, especially for new technologies. In this regard, and most directly, about 41% of respondents rated on-the-job learning as being more important for new technologies, relative to 25% who rated it as being more important for existing technologies. In contrast, completing a college degree, a more general form of human capital, was viewed as equally important for both technologies.

Turning to the risk of mobility, about 40% of respondents said that in their industry, key team members quit during a project very often (i.e., in more than 50% of projects) or often (26–50% of projects). Furthermore, 51% of respondents said that mobility risk was higher for projects with new technologies (compared with about 17% who said that it was higher for projects with existing technologies). Together, these support our theoretical arguments about mobility risk being higher for new technologies.

Consistent with industry-specific human capital being an important driver of mobility risk, 55% of key team members who quit went to a competitor. Symmetrically, the most common external source for replacing key team members who quit was competitors (26%) (replacements were most likely to come from within the company (53%)).

Focusing on costs of mobility, 54% said the costs of finding, training, and onboarding a replacement were higher for new technologies (compared with 20% for existing technologies). The mean expected time to replace a departing key team member was 5.8 months for new technology projects (vs. 4.6 months for existing technologies). The mean expected cost increase was also higher for new technologies (45.5% vs. 38.2%). Together, this evidence supports our thesis that mobility costs are likely to be higher for new technologies.

Finally, the survey evidence suggests that NCAs are widely used for key project members and is consistent with the thesis that NCAs hinder mobility. Then, 57% had signed an NCA and another 10% stated that they had probably signed an NCA. In addition, 66% of respondents said that their company requires NCAs for key project members, and a little over 69% said that NCA use was somewhat common (i.e., 25–50% of people) or very common (i.e., >50% of people) in their local market. NCAs had made switching jobs difficult for about 6% of respondents and hindered hiring key team members for about 34% of respondents. Together, the survey evidence seems broadly consistent with the key assumptions underlying our theoretical arguments.

7 | DISCUSSION AND CONCLUSIONS

In this article, we explore the implications of worker mobility for adoption of a new technology. We provide a framework to demonstrate how worker mobility can reduce the (net) benefits from investing in a new technology that requires investments in VSHC (Chari & Hopenhayn, 1991; Choudhury et al., 2020) and complementary BPI (Brynjolfsson et al., 2019; Brynjolfsson, Jin, & McElheran, 2021; Tambe, 2014). We then bring this framework to the data by exploring business adoption of ML, which provides evidence on how worker mobility can negatively influence the benefits of adopting a new technology. In so doing, we advance understanding of the diffusion of new technologies and the value that firms obtain from them (e.g., McElheran, 2015; Seamans, 2012).

Our research advances recent work that has sought to better understand the factors influencing the diffusion of artificial intelligence and ML by highlighting a critical role for worker mobility. Specifically, our study adds to work on the link between the value of ML adoption and human capital (e.g., Alekseeva et al., 2020; Alekseeva et al., 2021; Allen & Choudhury, 2022; Cao & Iansiti, 2021; Choudhury et al., 2020; Rock, 2021). These studies provide insights into the complementarities between human capital and ML when human capital is fixed, or they focus on the implications of ML-related human capital investments. In contrast, by focusing on the drivers of ML adoption, we explore factors that shape the accumulation of ML-related human capital, which naturally highlights a role for worker mobility in adoption decisions.

By focusing on ML and worker mobility, our study is also related to recent research that examines IT worker skills and has shown the benefits of IT worker mobility for the productivity of firms (Tambe, 2014; Tambe & Hitt, 2014; Wu et al., 2018). Research in this area has focused on the potential benefits of inward mobility; our research design takes a contrasting perspective and measures the net effects of mobility caused by policy changes, thereby providing evidence that outward mobility can shape decisions to adopt a new technology.

Our research adds to studies that examine the implications of NCA enforceability for firm outcomes (Conti, 2014; Garmaise, 2011; Jeffers, 2024; Rock, 2021; Starr, 2019) by adding new technology adoption as a possible firm outcome that may be influenced by NCA enforceability.[8] Our results differ from earlier work in the NCA literature that has emphasized the benefits of

[8]While not a focus of that paper, Rock (2021) does not find evidence that changes to NCA enforceability are associated with changes in the number or wages of engineering employees. This may be because his firm-level data use the location of the headquarters to measure the implications of NCA changes. 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.

weaker NCA enforceability for innovation and local economic growth (Marx et al., 2009), and is closer to more recent work that has highlighted the potential benefits of stronger NCA enforceability (e.g., Conti, 2014; Kang & Lee, 2022). We conjecture that this difference in results may arise from our focus on early-stage investments in VSHC (Chari & Hopenhayn, 1991; Choudhury et al., 2020). Further, our results indirectly suggest that early adopters may have greater incentives to adopt NCAs for their workers since such adopters are the most affected by outward worker mobility.

Our research has implications for understanding the relationships between mobility, market structure, and innovation. Broadly, increases in workforce mobility can help to erode entry barriers that arise with technology leadership and facilitate imitative entry (Agarwal & Gort, 2001). Within the context of IT investments, recent work has found evidence that such investments can lead to increases in industry concentration (Autor et al., 2020; Bessen, 2020; Brynjolfsson et al., 2023), and argued that worker mobility has the potential to weaken the IT-based advantages of dominant firms (Bessen, 2022). Our research suggests that these relationships may evolve over time when the use of new technology requires VSHC investments (Chari & Hopenhayn, 1991). During the early stages of technology diffusion, workforce mobility may dissuade firms from making investments in innovation that could give rise to technology leadership, and thus in turn potentially slow down imitative entry.

Finally, our approach of measuring ML adoption through ML-enabled business applications software complements other approaches that use alternative data sources such as job postings in Burning Glass or information from LinkedIn (Alekseeva et al., 2020; Alekseeva et al., 2021; Goldfarb et al., 2020; Rock, 2021) or 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), this strategy for measuring ML adoption across a large sample of establishments is likely to be of independent interest.

7.1 | Limitations

Our results have limitations that suggest avenues for future research. For one, we rely on a recent literature that has demonstrated how NCA enforceability influences technology and knowledge worker mobility (Balasubramanian et al., 2022; Jeffers, 2019; Marx et al., 2009). However, unlike some studies (e.g., Marx et al., 2009), we are unable to measure worker movements. Future research can use information on workers to develop a more fine-grained view of the impact of outward and inward worker mobility.

Our research seeks to unpack how mobility in complementary human capital inputs shapes adoption of new technologies such as ML in the short run. Because of the vintage-specific nature of these investments (Chari & Hopenhayn, 1991; Choudhury et al., 2020), the long-run implications may be different. With a longer panel, researchers could measure how the dynamics of VSHC shape technology diffusion. For example, as the number of adopters increases, the need for downstream innovation and the value of the human capital accrued during a new implementation decrease. Research could examine if the interplay between human capital and investment in new technologies such as ML is different in the short run than in the long run. Relatedly, because of data limitations we examine adoption in a single year, 2018. Researchers can leverage data from the years before and after 2018 to provide a more comprehensive look.

We have studied ML, a particular type of technology where the role of complementary BPI is very important and the cost of outward mobility is likely high. This is not likely to be true for other technologies. While our exploration of the implications of mobility for tablet adoption represent a first step in this direction, future research can explore the generalizability of our results by studying how the role of worker mobility might vary in environments where VSHC and BPI are less important.

It is worthwhile noting issues that we do not address surrounding the interplay between ML adoption, mobility, and its implications. First, we have viewed ML as an enabler of process innovation. However, ML can also enable product innovation. That may be affected differently by worker mobility than process innovation and can be examined. Further, the focus of our paper is on how changes in worker mobility influence short-run ML adoption decisions, not on the implications of these decisions for labor demand. As a result, we do not advance recent work that seeks to examine the labor demand implications of ML investments (e.g., Brynjolfsson et al., 2018; Felten et al., 2018).

Relatedly, we have also not delved into the performance or welfare implications of ML adoption. On the one hand, as ML diffuses and becomes more widely available, the associated increases in productivity (Brynjolfsson, Jin, & McElheran, 2021) and innovation (Babina et al., 2024) can lead to additional market entry (Bennett & Hall, 2020). On the other hand, recent research has highlighted some negative consequences for society (e.g., Acemoglu, 2021), and as an anonymous reviewer noted, slower adoption may allow the development of guardrails to prevent such consequences. Research can assess such trade-offs associated with faster adoption.

A critical assumption in our research is that there are no unobserved state-level factors that are changing in ways correlated with NCA enforceability changes and technology adoption. We have probed this assumption through the exploration of different controls, subsamples, and measurement strategies and through falsification exercises. Our results have proven robust across all these, increasing confidence in the results and in our interpretation. However, we leave it to other work to further probe the robustness of our findings. For example, additional work could investigate our hypotheses using different data and measurement strategies for key constructs such as firm benefits from adoption of ML. A related area of research could be a deeper investigation of the potential asymmetric effects between increasing NCA enforceability (that benefits employers) and decreasing enforceability (that benefit employees). We cannot rule out the null hypothesis of symmetry in our baseline and heterogeneity analyses (Tables B5–B8). However, this could be a result of limited statistical power, and future studies can examine this further. In this regard, it may also be useful to incorporate NCA use by firms into the analyses, as research has shown that the NCA use varies by firm size and industry (Starr et al., 2021).

In conclusion, our research has taken a first step toward understanding how worker mobility can depress incentives to invest in a new technology. However, there remain many ways to build on these results. We invite our readers to do so.

ACKNOWLEDGMENTS

The authors thank Sara Bana, Erik Brynjolfsson, Kenneth Huang, Aija Leiponen, Matt Marx, Evan Starr, Jonathan Timmis, Pengxiang Zhang, the editor (Rajshree Agarwal), and two anonymous referees, as well as seminar participants at Cornell University; Stanford University; the 2020 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; the 2021 Academy of

Management Annual Meeting (AOM), and the 2021 NBER Economics of Artificial Intelligence Conference for helpful comments and suggestions. The authors also thank Aberdeen Computer Intelligence for supplying data. All opinions and errors are their own.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from Aberdeen Group, LLC. Restrictions apply to the availability of these data, which were used under license for this study. Data are available from the authors with the permission of Aberdeen Group, LLC.

ORCID

Ruyu Chen https://orcid.org/0009-0005-4104-7942
Natarajan Balasubramanian https://orcid.org/0000-0002-3513-5090
Chris Forman https://orcid.org/0000-0003-3055-1607

REFERENCES

Acemoglu, D. (2021). Harms of AI. (NBER Working Paper No. 31065). National Bureau of Economic Research.

Agarwal, R., Ganco, M., & Ziedonis, R. H. (2009). Reputations for toughness in patent enforcement: Implications for knowledge spillovers via inventor mobility. *Strategic Management Journal*, *30*(13), 1349–1374.

Agarwal, R., & Gort, M. (2001). First-mover advantage and the speed of competitive entry, 1887–1986. *Journal of Law and Economics*, *44*(1), 161–177.

Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Press.

Agrawal, A., Gans, J., & Goldfarb, A. (2021). AI adoption and system-wide change. (NBER Working Paper No. 28811) National Bureau of Economic Research.

Agrawal, A., Gans, J. S., & Goldfarb, A. (2019). Exploring the impact of artificial intelligence: Prediction versus judgment. *Information Economics and Policy*, *47*, 1–6.

Alekseeva, L., Azar, J., Giné, M., Samila, S., & Taska, B. (2021). The demand for AI skills in the labor market. *Labour Economics*, *71*, 1–27.

Alekseeva, L., Giné, M., Samila, S., & Taska, B. (2020). AI adoption and firm performance: Management versus IT. Working Paper.

Allen, R., & Choudhury, P. (2022). Algorithm-augmented work and domain experience: The countervailing forces of ability and aversion. *Organization Science*, *33*(1), 149–169.

Almeida, P., & Kogut, B. (1999). The localization of knowledge and the mobility of engineers in regional networks. *Management Science*, *45*(7), 905–917.

Autor, D., Dorn, D., Katz, L. F., Patterson, C., & Van Reenen, J. (2020). The fall of the labor share and the rise of superstar firms. *Quarterly Journal of Economics*, *135*(2), 645–709.

Autor, D. H., Donohue, J. J., III, & Schwab, S. J. (2006). The costs of wrongful-discharge laws. *Review of Economics and Statistics*, *88*(2), 211–231.

Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, *151*, 103745.

Balasubramanian, N., Chang, J. W., Sakakibara, M., Sivadasan, J., & Starr, E. (2022). Locked in? The enforceability of covenants not to compete and the careers of high-tech workers. *Journal of Human Resources*, *57*(1), S349–S396.

Barth, E., Davis, J. C., Freeman, R. B., & McElheran, K. (2023). Twisting the demand curve: Digitalization and the older workforce. *Journal of Econometrics*, *233*(2), 443–467.

Beck Reed Riden LLP. (2019). Employee noncompetes: A state by state survey.

Bennett, V. M., & Hall, T. A. (2020). Software availability and entry. *Strategic Management Journal*, *41*, 950–962.

Bessen, J. (2020). Industry concentration and information technology. *Journal of Law and Economics*, *63*(3), 531–555.

Bessen, J. (2022). *The new goliaths: How corporations use software to dominate industries, kill innovation, and undermine regulation*. Yale University Press.

Bloom, N., Sadun, R., & Van Reenen, J. (2012). Americans do IT better: US multinationals and the productivity miracle. *American Economic Review*, *102*(1), 167–201.

Bresnahan, T. (2019). Technological change in ICT in light of ideas first learned about in the machine tool industry. *Industrial and Corporate Change*, *28*(2), 331–349.

Bresnahan, T. (2020). Artificial intelligence technologies and aggregate growth prospects. Working Paper.

Bresnahan, T., & Greenstein, S. (1996). Technical progress and co-invention in computing and in the uses of computers. *Brookings Papers on Economic Activity*, *1996*, 1–83.

Bresnahan, T. F., Brynjolfsson, E., & Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *Quarterly Journal of Economics*, *117*(1), 339–376.

Brynjolfsson, E., Jin, W., & McElheran, K. (2021). The power of prediction: Predictive analytics, workplace complements, and business performance. *Business Economics*, *56*(4), 217–239.

Brynjolfsson, E., Jin, W., & Wang, X. (2023). Information technology, firm size, and industrial concentration (NBER Working Paper No. 31065). National Bureau of Economic Research.

Brynjolfsson, E., & McElheran, K. (2016). The rapid adoption of data-driven decision-making. *American Economic Review*, *106*(5), 133–139.

Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What can machines learn and what does it mean for occupations and the economy? *American Economic Review Papers and Proceedings*, *108*, 43–47.

Brynjolfsson, E., Rock, D., & Syverson, C. (2019). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence* (pp. 23–57). University of Chicago Press.

Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, *13*(1), 333–372.

Bughin, J., Kretschmer, T., & van Zeebroeck, N. (2019). Experimentation, learning, and stress: The role of digital technologies in strategy change. Working Paper.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230.

Cao, R., & Iansiti, M. (2021). Growth, transformation and digital capital: The importance of technological and organizational architecture (Harvard Business School Research Paper Series No. 21-122). Harvard Business School.

Chari, V. V., & Hopenhayn, H. (1991). Vintage human capital, growth, and the diffusion of new technology. *Journal of Political Economy*, *99*(6), 1142–1165.

Choudhury, P., Starr, E., & Agarwal, R. (2020). Machine learning and human capital complementarities: Experimental evidence on bias mitigation. *Strategic Management Journal*, *41*, 1381–1411.

Cinelli, C., Forney, A., & Pearl, J. (2020). A crash course in good and bad controls. *Sociological Methods & Research*, 00491241221099552. https://doi.org/10.1177/00491241221099552

Conti, R. (2014). Do non-competition agreements lead firms to pursue risky R&D projects? *Strategic Management Journal*, *35*(8), 1230–1248.

Davenport, T. (2000). *Mission critical: Realizing the promise of enterprise systems*. Harvard Business School Press.

De Chaisemartin, C., & D'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, *110*(9), 2964–2996.

Ewens, M., & Marx, M. (2018). Founder replacement and startup performance. *Review of Financial Studies*, *31*(4), 1532–1565.

Farboodi, M., Mihet, R., Philippon, T., & Veldkamp, L. (2019). Big data and firm dynamics. *American Economic Review*, *109*, 38–42.

Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, *42*(12), 2195–2217.

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

Forman, C., Goldfarb, A., & Greenstein, S. (2002). Digital dispersion: An industrial and geographic census of commercial internet use (NBER Working Paper No. w9287). National Bureau of Economic Research.

Forman, C., Goldfarb, A., & Greenstein, S. (2005). How did location affect adoption of the commercial internet? Global village vs. urban leadership. *Journal of Urban Economics*, *58*(3), 389–420.

Forman, C., Goldfarb, A., & Greenstein, S. (2008). Understanding the inputs into innovation: Do cities substitute for internal firm resources? *Journal of Economics & Management Strategy*, *17*(2), 295–316.

Forman, C., Goldfarb, A., & Greenstein, S. (2012). The internet and local wages: A puzzle. *American Economic Review*, *102*(1), 556–575.

Garmaise, M. J. (2011). Ties that truly bind: Noncompetition agreements, executive compensation, and firm investment. *Journal of Law, Economics, and Organization*, *27*(2), 376–425.

Goldfarb, A., Taska, B., & Teodoridis, F. (2023). Could machine learning be a general purpose technology? A comparison of emerging technologies using data from online job postings. *Research Policy*, *52*(1), 104653.

Goldfarb, A., Taska, B., & Teodoridis, F. (2020). Artificial intelligence in health care? Evidence from online job postings. *AEA Papers and Proceedings*, *110*, 400–404.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, *225*(2), 254–277.

Griliches, Z. (1957). Hybrid corn: An exploration in the economics of technological change. *Econometrica*, *25*(4), 501–522.

Hannan, T. H., & McDowell, J. M. (1984). The determinants of technology adoption: The case of the banking firm. *RAND Journal of Economics*, *15*(3), 328–335.

Hess, S. (2017). Randomization inference with Stata: A guide and software. *Stata Journal*, *17*(3), 630–651.

Jeffers, J. S. (2024). The impact of restricting labor mobility on corporate investment and entrepreneurship. *The Review of Financial Studies*, *37*(1), 1–44.

Kang, H., & Lee, W. (2022). How innovating firms manage knowledge leakage: A natural experiment on the threat of worker departure. *Strategic Management Journal*, *43*(10), 1961–1982.

Karshenas, M., & Stoneman, P. L. (1993). Rank, stock, order, and epidemic effects in the diffusion of new process technologies: An empirical model. *RAND Journal of Economics*, *24*(4), 503–528.

Klepper, S. (1996). Entry, exit, growth, and innovation over the product life cycle. *American Economic Review*, *86*(3), 562–583.

Kretschmer, T., Miravete, E. J., & Pernías, J. C. (2012). Competitive pressure and the adoption of complementary innovations. *American Economic Review*, *102*(4), 1540–1570.

LaPlante, A. (2020). *What is augmented analytics? Powering your data with AI*. O'Reilly Media.

Levy, M. (2019). Aberdeen behavioral technographics. [Online resource]. Retrieved from https://gzconsulting.org/2019/10/07/aberdeen-behavioral-technographics/

Malsberger, B., Carr, D. J., Pedowitz, A. H., & Tate, E. A. (2017). *Covenants not to compete: A state-by-state survey* (11th ed.). Bloomberg BNA.

Marx, M., Strumsky, D., & Fleming, L. (2009). Mobility, skills, and the Michigan non-compete experiment. *Management Science*, *55*(6), 875–889.

Mata, F. J., Fuerst, W. L., & Barney, J. B. (1995). Information technology and sustained competitive advantage: A resource-based analysis. *MIS Quarterly*, *19*(4), 487–505.

McElheran, K. (2015). Do market leaders lead in BPI? The case(s) of e-business adoption. *Management Science*, *61*(6), 1197–1216.

McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L. S., & Zolas, N. (2023). AI adoption in America: Who, what, and where. (NBER Working Paper No. 31788). National Bureau of Economic Research.

Nagle, F. (2018). Learning by contributing: Gaining competitive advantage through contribution to crowdsourced public goods. *Organization Science*, *29*(4), 569–587.

Nagle, F. (2019). Open source software and firm productivity. *Management Science*, *65*(3), 1191–1215.

Raj, M., & Seamans, R. (2019). Primer on artificial intelligence and robotics. *Journal of Organizational Design*, *8*(11), 1–14.

Roberts, B., & Wolf, M. (2018). High-tech industries: An analysis of employment, wages, and output. *Beyond the Numbers*, *7*(7), 1–11.

Rock, D. (2021). Engineering value: The returns to technological talent and investments in artificial intelligence. *SSRN Electronic Journal* Retrieved from https://ssrn.com/abstract=3427412

Rosenberg, N. (1963). Technological change in the machine tool industry, 1840–1910. *Journal of Economic History*, *23*(4), 414–443.

Sallam, R., Howson, C., & Idoine, C. (2017). Augmented analytics is the future of data and analytics. [Gartner Report].

Seamans, R. (2012). Fighting city hall: Entry deterrence and technology upgrades in cable TV markets. *Management Science*, *58*(3), 461–475.

Seegert, N. (2012). Optimal taxation with volatility: A theoretical and empirical decomposition. Working Paper, University of Michigan, Ann Arbor.

Singh, J., & Agrawal, A. (2011). Recruiting for ideas: How firms exploit the prior inventions of new hires. *Management Science*, *57*(1), 129–150.

Starr, E. (2019). Consider this: Training, wages, and the enforceability of covenants not to compete. *ILR Review*, *72*(4), 783–817.

Starr, E., Balasubramanian, N., & Sakakibara, M. (2018). Screening spinouts? How noncompete enforceability affects the creation, growth, and survival of new firms. *Management Science*, *64*(2), 552–572.

Starr, E., Prescott, J. J., & Bishara, N. (2021). Noncompetes in the US labor force. *Journal of Law and Economics*, *64*, 53–84.

Stiroh, K. (2002). Information technology and the U.S. productivity revival: What do the industry data say? *American Economic Review*, *92*(5), 1559–1576.

Tambe, P. (2014). Big data investment, skills, and firm value. *Management Science*, *60*(6), 1452–1469.

Tambe, P., & Hitt, L. M. (2012). The productivity of information technology investments: New evidence from IT labor data. *Information Systems Research*, *23*(3), 599–617.

Tambe, P., & Hitt, L. M. (2014). Job hopping, information technology spillovers, and productivity growth. *Management Science*, *60*(2), 338–355.

Tambe, P., Hitt, L. M., Rock, D., & Brynjolfsson, E. (2019). IT, AI and the growth of intangible capital. *SSRN Electronic Journal* Retrieved from https://ssrn.com/abstract=3416289

United States Executive Office of the President. (2016). Artificial intelligence, automation, and the economy. Retrieved from https://obamawhitehouse.archives.gov/sites/whitehouse.gov/files/documents/Artificial-Intelligence-Automation-Economy.PDF

Wu, L., Jin, F., & Hitt, L. M. (2018). Are all spillovers created equal? A network perspective on information technology labor movements. *Management Science*, *64*(7), 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), 348–368.

Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D. N., Buffington, C., Goldschlag, N., Foster, L., & Dinlersoz, E. (2020). Advanced technologies adoption and use by US firms: Evidence from the annual business survey. (NBER Working Paper No. 28290). National Bureau of Economic Research.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Chen, R., Balasubramanian, N., & Forman, C. (2024). How does worker mobility affect business adoption of a new technology? The case of machine learning. *Strategic Management Journal*, 1–29. https://doi.org/10.1002/smj.3595