Essays on

The Selectiveness of Firms' Hiring: Determinants and Measurement

Inauguraldissertation
zur
Erlangung des Doktorgrades
der

Wirtschafts-und Sozialwissentschaftlichen Fakultät
der
Universität zu Köln

2019

vorgelegt
von

Sebastian Butschek, MSc
aus
Innsbruck

Referent:	Prof. Dr. Dirk Sliwka
Korreferent:	Prof. Dr. Erik Lindqvist
Tag der Promotion:	03.09.2019

Acknowledgements

First I want to thank Dirk Sliwka, my supervisor. His curiosity, intrinsic motivation, rigor and humility have been an inspiration. In addition, he has given me great freedom and support for pursuing my research ideas. He has also been most generous with his time—both in terms of discussing my work and in terms of making my teaching load very manageable.

Next I want to express my gratitude to Jan Sauermann. Since our chance encounter in Mannheim three years ago he has been a fantastic co-author, mentor and friend. His generosity in sharing both the data access and the research grant he had won has enabled me to repeatedly visit Stockholm and work with some of the richest administrative data in the world. I have very much enjoyed working with and learning from Jan and greatly appreciate his openness to my ideas.

For financial support I thank the German Research Foundation, which has provided funding through priority program DFG SSP 1764. I benefited from the support of the team at IAB's research data center in Nuremberg, who helped me obtain access to some non-standard data. In addition, I am grateful to ZEW Mannheim, where I worked before coming to Cologne in 2016. ZEW has continued to host me on many occasions and has allowed me to use its fantastic IT infrastructure.

My thanks also go to my colleagues, both at the University of Cologne and at ZEW Mannheim, for making me feel welcome and at home and for their support at many stages over the last years.

Finally, I am deeply indebted to my partner Lili, my family and my friends for always having my back. I have been able to talk to them about my research when necessary, and, even more importantly, have experienced awesome things with them that had nothing to do with my work. In this way they have helped me both to get through the tougher bits and to enjoy the pleasant parts of the PhD journey.

Contents

Introduction

Firms spend considerable time and resources on choosing their workers. For instance, data from the Linked Personnel Panel (LPP) on 1,783 German private-sector establishments with 50 employees or more shows that at the average employer, a successful applicant for a non-managerial position experiences three hours of pre-hire screening (LPP, waves 1-4, 2012-2018).

Some theoretical work in labor economics takes screening into account when modelling the hiring process and shows that doing so makes a difference (e.g., Villena-Roldán, 2012; Wolthoff, 2017). There is also some empirical evidence describing determinants of the screening or recruitment process that are related to job or firm characteristics as well as some work exploring the correlation of screening/recruitment with match quality (see Oyer and Schaefer, 2010, for an overview). However, serious attempts to identify causal relationships are rare (e.g., Pellizzari, 2011).

On the effect of labor market regulation on firms' worker selection, hardly any evidence exists. That is, we know very little about whether firms' hiring becomes more or less selective when labor market institutions change. In part, the dearth of empirical evidence on the link between labor market policy and screening may be a data issue: studying the selectiveness of firms' hiring requires information about the ability of new hires, which is rarely observed. The three articles constituting this thesis try to reduce this gap in the literature in two ways. First, Chapters 1 and 2 provide some of the first evidence on the causal effect of minimum wages and employment protection, respectively, on the selectiveness of firms' hiring. Second, Chapter 3 offers an empirical assessment of a widely known estimated ability proxy. Named after their originators in Abowd et al. (1999), AKM worker effects are time-invariant individual wage premia estimated from labor market data and are used as a measure of worker productivity by a growing list of empirical studies.

Why is it important to know if firms respond to policy changes by hiring more or less selectively? The primary benefit of evidence on this causal link is that it may inform the design of labor market policy. For example, a tightening of regulation may have the unintended side effect of making firms hire more selectively. Then, in addition to benefiting a targeted group of workers, the policy change may have negative distributional effects on a group of workers that is particularly vulnerable. That is, the least productive workers may find it even harder to get hired following the reform-induced increase in firms' hiring standards and may get trapped in unemployment.

Finding out how well estimated ability proxies work, on the other hand, is primarily intended to help researchers. In particular, this includes empirical economists who need a measure of individual productivity and would like to know the pros and cons of using estimated AKM worker effects. Quantifying their performance as an ability measure appears especially useful in light of a number of recent contributions questioning AKM worker effects' validity as a productivity proxy both on theoretical and on empirical grounds (e.g., Eeckhout and Kircher, 2011; Abowd et al., 2018; Bonhomme et al., 2018).

Chapter 1 investigates whether the 2015 introduction of a statutory minimum wage in Germany made a sample of 1,491 private-sector employers hire more selectively. My difference-in-differences analysis compares the evolution of realized hiring standards across establishments treated and untreated by the reform. I proxy hiring standards by establishments' minimum hire quality. As a measure of hire quality I use worker effects from two-way fixed-effect log wage regressions estimated before my period of analysis, which capture both observed and unobserved ability. To classify employers as treated I rely on their pre-reform wage structure: whether or not they employed sub-minimum wage workers. I find that in response to the minimum wage introduction, treated establishments hired more productive workers: the relative increase in new hires' minimum hire quality is 18.9% of a standard deviation. Using pre-reform survey information on screening intensity I show that the effect is increasing in the importance of screening to the establishment's hiring process. This strengthens my interpretation of the effect on minimum hire quality as a change in employers' hiring standards.

Chapter 2 estimates the causal effect of employment protection on firms' worker selection. Together with Jan Sauermann I study a 2001 policy change in Sweden that reduced dismissal costs for firms with less than ten employees. These small firms employ over a tenth of Sweden's workforce. Our difference-in-differences analysis of firms' hiring uses individual ability measures including estimated AKM worker effects and cognitive test scores from the military draft. Comparing treated small firms to slightly larger, untreated ones we find that the reform reduced minimum hire quality by five per cent of a standard deviation, half of which we can attribute to firms' hiring becoming less selective. Consistent with this, we show that employers' reliance on network-based hiring also goes down as firms react to the reduction of dismissal costs. Our results help discriminate between existing theories, supporting the prediction that firms shift their hiring standards in response to changes in dismissal costs.

Chapter 3 is a descriptive study that assesses the performance of estimated worker effects as an ability proxy. It is also joint work with Jan Sauermann. We focus on worker effects from two-way fixed-effect log-wage regressions as introduced by Abowd et al. (1999). Using Swedish register data, we study the correlation between such worker effect estimates and standardized cognitive test scores from the military draft. For the period from 2007 to 2013 we find a correlation of roughly .4. We show that this is more than three quarters of the predictive performance achieved by machine-learning algorithms. However, we find that the worker effects' prediction performance varies with observed worker and firm characteristics—for example, across industries. Moreover, we identify biases—for instance, the worker effects underestimate cognitive test scores for blue-collar workers. Our analysis of the worker effects' determinants shows that they "contain" a significant amount of both skill-related and non-skill attributes such as education and cognitive ability as well as age. Finally, our exploration of the importance of estimation ingredients for prediction performance broadly supports the estimation choices of existing applications.

While Chapters 1 and 2 both study labor market regulation as a determinant of firms' worker selection, they focus on very different policy changes: a tightening versus a loosening of regulation and the creation of a uniform wage floor versus a modification of dismissal costs. To the extent, however, that the two reforms can be compared, my estimates of their effects are consistent with each other. In a broad sense both minimum wages and employment protection provisions change the cost of labor; Chapter 1 suggests an exogenous *increase* in the cost of labor makes firms' hiring *more* selective and Chapter 2 indicates that an exogenous *decrease* in the cost of labor makes firms' hiring *less* selective. The studies appear in less agreement about the size of the respective reforms' effects: the increase in minimum hire quality in Chapter 1 is nearly eight times as large as the decrease in minimum hire quality in Chapter 2. One potential explanation for this difference is that the immediate costs imposed by minimum wages are more salient than the potential future costs of employment protection. Another potential explanation is that the 2001 Swedish employment protection reform represented a small reduction in dismissal costs: the principle was maintained that dismissals are allowed only on the grounds of proven economic necessity (not out of mere dissatisfaction with a worker's ability).

In Chapter 1 I choose a binary measure for establishments' minimum wage affectedness, arguing that employers' incentives to raise their hiring thresholds does not depend on the share of pre-reform sub-minimum wage workers, only on whether or not they employ very low-paid workers at all. In Chapter 2 we also work with a binary measure for firms' exposure to the employment protection reform, but we show that the effect is stronger for smaller firms and argue that this is due to the reform's greater bite—the share of workers for whom employment protection is lifted by the reform is much larger for small firms. This may appear inconsistent. The reason it is not a contradiction is that the Swedish reform actually reduced dismissal costs more for smaller firms than for bigger ones while the German reform did not introduce a higher minimum wage for employers with more low-wage workers. Establishments with a larger share of sub-minimum wage workers would only have to apply their new hiring threshold to a larger number of new hires, not raise their hiring threshold further.

By and large Chapter 3 provides an empirical validation of the preceding chapters' use of AKM worker effects as an individual ability proxy. The most relevant caveat that emerges is that the worker effects contain a substantial age component. In Chapter 1, where I find that a minimum wage introduction increases the selectiveness of hiring, the age bias of AKM worker effects might exaggerate the reform's true effect on worker selection. This would be the case if employers affected by the minimum wage started hiring older workers in response to the policy change. I show that this is not the case. In Chapter 2, where we identify a downward shift in firms' hiring standards caused by a reduction of employment protection, the age bias in AKM worker effects might also overstate the policy change's effect on worker selection. This would happen if treated firms responded to the reform by hiring younger workers. While this is theoretically plausible, we find the opposite.

There is a small body of evidence on the link between minimum wages and firms' hiring standards, most of it descriptive (Fairris and Bujanda, 2008; Hirsch et al., 2015; Gürtzgen et al., 2016). My findings in Chapter 1 confirm and extend these results by providing causal evidence that employers indeed hire more selectively in response to minimum wages. In a contemporaneous working paper, Clemens et al. (2018) also find that minimum wages raise hiring standards by looking at job postings. Chapter 1 differs from their paper in its focus on realized hires using worker inflow data and its focus on unobserved worker ability.

Three existing papers present evidence consistent with an effect of employment protection on firms' worker selection (Kugler and Saint-Paul, 2004; Marinescu, 2009; Bjuggren and Skedinger, 2018). Our findings in Chapter 2 are consistent with the results of these papers. They go further than them by providing evidence on new hires' ability and by discriminating between potential channels driving the effect.

To my knowledge, there is no paper that uses an established measure of individual ability to assess the validity of AKM worker effects as an ability proxy. That doing so is nonetheless interesting is underscored by several methodological contributions, some of them recent, that call into question whether AKM worker effects are correlated with the individual's true productivity (e.g., Eeckhout and Kircher, 2011; Hagedorn et al., 2017; Abowd et al., 2018; Bonhomme et al., 2018).

Chapters 1-3 that follow are structured like regular journal articles, each with its own introduction, conclusion, list of references and appendix. A concluding chapter briefly considers some of their limitations and implications. An estimate of my contribution to each of the chapters is found in the appendix to the dissertation.

References

Abowd, John M, Francis Kramarz, and David N Margolis (1999): "High wage workers and high wage firms," *Econometrica*, 67(2): 251–333.

Abowd, John M, Kevin L McKinney, and Ian M Schmutte (2018): "Modeling endogenous mobility in earnings determination," *Journal of Business & Economic Statistics*, 1–14.

Bjuggren, Carl Magnus and Per Skedinger (2018): "Does job security hamper employment prospects?" IFN Working Paper 1255, Research Institute of Industrial Economics.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa (2018): "A distributional framework for matched employer employee data," *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2019-15).

Clemens, Jeffrey, Lisa Kahn, and Jonathan Meer (2018): "Dropouts need not apply: The minimum wage and skill upgrading," Tech. rep., Working Paper.

Eeckhout, Jan and Philipp Kircher (2011): "Identifying sorting—in theory," *Review of Economic Studies*, 78(3): 872–906.

Fairris, David and Leon Fernandez Bujanda (2008): "The dissipation of minimum wage gains for workers through labor-labor substitution: evidence from the Los Angeles living wage ordinance," *Southern Economic Journal*, 473–496.

Gürtzgen, Nicole, Alexander Kubis, Martina Rebien, and Enzo Weber (2016): "Neueinstellungen auf Mindestlohnniveau: Anforderungen und Besetzungsschwierigkeiten gestiegen," Tech. rep., IAB-Kurzbericht.

Hagedorn, Marcus, Tzuo Hann Law, and Iourii Manovskii (2017): "Identifying equilibrium models of labor market sorting," *Econometrica*, 85(1): 29–65.

Hirsch, Barry T, Bruce E Kaufman, and Tetyana Zelenska (2015): "Minimum wage channels of adjustment," *Industrial Relations: A Journal of Economy and Society*, 54(2): 199–239.

Kugler, Adriana D. and Gilles Saint-Paul (2004): "How do firing costs affect worker flows in a world with adverse selection?" *Journal of Labor Economics*, 22(3): 553–584.

Marinescu, Ioana (2009): "Job security legislation and job duration: evidence from the United Kingdom," *Journal of Labor Economics*, 27(3): 465–486.

Oyer, Paul and Scott Schaefer (2010): "Personnel economics: hiring and incentives," Tech. rep., National Bureau of Economic Research.

Pellizzari, Michele (2011): "Employers' search and the efficiency of matching," *British journal of industrial relations*, 49(1): 25–53.

Villena-Roldán, Benjamin (2012): "Aggregate implications of employer search and recruiting selection," unpublished manuscript, University of Chile.

Wolthoff, Ronald (2017): "Applications and Interviews: Firms' Recruiting Decisions in a Frictional Labour Market," *The Review of Economic Studies*.

Chapter 1

Raising the bar: minimum wages and employers' hiring standards

Author: Sebastian Butschek

Abstract: This paper investigates whether the 2015 introduction of a statutory minimum wage in Germany made a sample of 1,491 private-sector employers hire more selectively. My difference-in-differences analysis compares the evolution of realized hiring standards across establishments treated and untreated by the reform. I proxy hiring standards by establishments' minimum hire quality. As a measure of hire quality I use worker effects from two-way fixed-effect log wage regressions estimated before my period of analysis, which capture both observed and unobserved ability. To classify employers as treated I rely on their pre-reform wage structure: whether or not they employed sub-minimum wage workers. I find that in response to the minimum wage introduction, treated establishments hired more productive workers: the relative increase in new hires' minimum hire quality is 18.9% of a standard deviation. Using pre-reform survey information on screening intensity I show that the effect is increasing in the importance of screening to the establishment's hiring process. This strengthens my interpretation of the effect on minimum hire quality as a change in employers' hiring standards.

Keywords: worker selection, hiring standard, screening, minimum wage, labor-labor substitution.

JEL Codes: J23, J24, J31, J38, M51.

1.1 Introduction

The effects of minimum wages on hiring have been studied by labor economists for decades. Scores of papers on the employment effects of minimum wages notwithstanding, very little is known about how policy-mandated wage floors affect the composition of hires. Yet it is plausible that a minimum wage could affect how selectively firms hire. If it has any bite, a new minimum wage will exogenously make low-wage workers more expensive. For employers of sub-minimum wage workers, this may raise hiring standards: to ensure new hires remain profitable for them, firms may require greater productivity of applicants for low-wage jobs. Despite the scarcity of evidence, it is important to understand if worker selection is a channel of employer adjustment to minimum wages. Firms hiring more selectively in response to a minimum wage may have significant distributional consequences for job-seekers. Workers with the lowest productivity may experience an increased likelihood of getting stuck in unemployment, even in the absence of negative aggregate employment effects of minimum wages.

In this paper I provide some of the first causal evidence on the effect of a minimum wage on employers' hiring standards by studying the 2015 introduction of a statutory minimum wage (SMW) in Germany. Exploiting variation in pre-reform wage structure, I compare establishments treated by the SMW introduction to untreated ones. Using a differences-in-differences (DiD) framework, I estimate the effect of the reform on minimum hire quality, my proxy for employers' hiring standards. Minimum hire quality is based on estimated person fixed effects, a measure of time-invariant productivity that captures both observed and unobserved skills. My results suggest that the hiring of German establishments became more selective in response to the SMW introduction: the reform increased minimum hire quality by 18.9% of a standard deviation.

My empirical analysis uses administrative data for 2010-2016 on the entire workforces of 1,491 establishments representative of the German private sector (excluding establishments with fewer than 50 workers). To assign establishments to the treated or control group I construct a measure for their exposure to the minimum wage introduction. I classify those employers as treated that employed sub-minimum wage workers in 2010-2013, before the SMW introduction could have been anticipated. My preferred measure for SMW exposure is binary for two reasons. First, establishments employing many minimum-wage workers do not have an incentive to raise hiring standards *further* than employers with just a few minimum-wage workers. Second, with a binary measure measurement error in SMW exposure should bias my results towards zero whereas it may bias them away from zero with a continuous minimum wage bite measure.

As a proxy for the productivity of new hires I utilize worker effects from a two-way fixed-effects log-wage regression as introduced by Abowd et al. (1999)(henceforth AKM). I use the AKM worker effects estimated for West Germany by Card et al. (2013) for the period 2002-2009, ensuring that ability proxies are pre-determined for any new hires considered in my analysis. AKM worker effects capture those components of time-invariant individual productivity that are reflected in workers' wage earnings. This includes education or vocational qualifications but extends to unobserved cognitive and non-cognitive ability such as motivation or conscientiousness (Butschek and Sauermann, 2019b). For the purpose of estimating the effect of an SMW introduction on hiring standards I view AKM worker effects as a more appropriate productivity measure than observed proxies (such as age or years of schooling). The first reason is that AKM worker effects vary continuously while the level of schooling is measured in coarse categories—possibly too coarse to capture the adjustment of hiring standards to a minimum wage. The second reason is that AKM worker effects should be more informative of productivity in low-wage jobs than age or schooling: it is not clear that older or more highly educated individuals have higher performance behind the cashier or on a lawnmower.

I first show that my measure of SMW exposure is relevant: the 2015 SMW introduction increased bottom starting wages by EUR 6.56 per day. Next I estimate the DiD effect of the SMW introduction on minimum hire quality, finding an increase of 18.9% of a standard deviation. I then estimate an event-study version of the DiD specification and find that the positive effect on hiring standards is already present in 2014, when the SMW was widely anticipated, and persists through 2016.

There are some threats to identification, which I address in turn. First, if the reform reduced the share of new hires whose ability I can measure this may spuriously produce the pattern I attribute to a shift of hiring standards. A similar argument applies to groups of workers whose ability may be systematically underestimated by AKM worker effects, such as women and young workers. I show that the reform did not significantly affect the share of hires for whom I have an AKM worker effect, nor the share of women or the age of hires. Second, despite parallel pre-reform trends minimum hire quality in treated and control establishments may have evolved differently post-reform even in the absence of the SMW introduction. As treated employers pay lower wages than untreated ones I run placebo tests that compare *untreated* establishments with lower average wages to *untreated* establishments with higher average wages. There are no significant placebo effects of the reform, providing suggestive evidence for parallel counter-factual post-reform trends.

To strengthen my interpretation that the SMW effect on minimum hire quality reflects a raising of employers' hiring standards I confront two alternative interpretations of the result. First, the increase in minimum hire quality may be a mechanical consequence of a drop in establishments' hiring rate. If employers hired significantly fewer workers each year, this would reduce the spread of the ability distribution of new hires. An implication is that maximum hire quality should drop, which I do not find. In addition, the SMW introduction did not significantly reduce the hiring rate in my sample.

Second, rather than reflecting an adjustment of qualitative labor demand, the increase in minimum hire quality may be driven by better workers sorting into treated establishments. I look at voluntary movers coming from high-wage employers to test for changes in worker self-selection and find no evidence for it. Moreover, I show that the SMW effect on minimum hire quality is more pronounced for establishments where survey evidence suggests that pre-hire screening plays a bigger role. That employers for whom screening is more feasible respond more strongly supports my interpretation that the SMW effect on minimum hire quality is due to hiring becoming more selective.

The two defining features of this paper are its focus on the minimum wage's effect on hiring standards and on a worker quality measure that captures unobserved productivity. There are three existing papers that also address the link between minimum wages and employers' hiring standards. First, using a survey of restaurant managers, Hirsch et al. (2015) find that a federal minimum wage hike increased self-reported intentions of hiring older and more experienced workers (though no positive effect on hire age materialized). Second, Gürtzgen et al. (2016) descriptively compare survey information on firms' hiring criteria from before and after the German SMW introduction I also study, finding that these criteria became stricter.[1] Unlike these two papers, my analysis does not use management-reported information but relies on inferring employers' hiring standards from the ability of realized hires. This makes my proxy more informative of effective changes in hiring standards than the intentions stated in a survey. Third, a recent working paper by Clemens et al. (2018) pursues an alternative way of providing revealed-preference evidence on hiring standards: they use data from job postings to show that employers' advertised hiring criteria became stricter in response to federal minimum wage increases in the US. In further contrast to Clemens et al. (2018) I employ an ability measure that varies continuously and captures not just observed human capital but also unobserved cognitive and non-cognitive ability. Reassuringly, our independent studies produce consistent and complementary results. Though based on different measures of worker skill and strategies for identifying employer hiring standards, both find a positive minimum wage effect on the selectiveness of hiring.

More generally, this paper is related to a small empirical literature on labor-labor substitution as a channel of firm adjustment to minimum wages. That literature almost exclusively uses observed characteristics such as age or schooling as proxies for worker skill.[2] It produces mixed evidence. The findings of some papers are consistent with substitution toward higher-skilled workers: Neumark and Wascher (1995) find that state-level minimum wage increases in the US were associated with teenage transition patterns between school, work and unemployment hinting at substitution toward higher-quality teenage labor. The analysis of a living wage ordinance for Los Angeles city contractors in Fairris and Bujanda (2008) shows that more workers hired after the wage mandate had pre-hire formal training than incumbents hired before. Bernini and Riley (2016) present evidence that increases of the UK National Minimum Wage were associated with occupational shifts among new hires toward more skill-intensive production. This suggestive evidence is strengthened by

[1] Fairris and Bujanda (2008) also have a direct survey item on whether employers raised hiring standards in response to a Los Angeles living-wage ordinance. However, they cannot use it as a regression outcome as all surveyed firms were equally affected by the wage floor.

[2] An exception is Fairris and Bujanda (2008). While they cannot provide a causal estimate, their descriptive case study of a living wage and labor-labor substitution looks at past wages conditional on observables as a proxy for unobserved worker ability. Conceptually, using past wages as an outcome while controlling for worker and past employer characteristics is similar to the AKM method.

Clemens et al. (2018), who find that state-level minimum wage hikes in the US led to substitution away from high-school dropouts and younger workers.

More papers, though using similar skill proxies, find no evidence of substitution toward higher-skilled workers along observed dimensions. Several do not support substitution away from the young (Portugal and Cardoso, 2006; Fairris and Bujanda, 2008; Giuliano, 2013; Hirsch et al., 2015; Cengiz et al., 2019). Some also look at schooling, where results do not support substitution away from those with low education (Fairris and Bujanda, 2008; Hirsch et al., 2015; Cengiz et al., 2019). In addition, Fairris and Bujanda (2008) find no indication of substitution toward more experienced workers. While labor-labor substitution along observed dimensions is not the focus of this paper, my results on this are in line with the latter group of papers. I find no significant effect of the German SMW on the share of hires with low education or out of unemployment nor on the age of hires.[3]

My empirical analysis is also relevant to theoretical macro work on the role of firm search for the formation of matches in the labor market. For example, several papers show that including screening in matching models alters the predictions for aggregate labor market behaviour with respect to cyclicality or in response to regulation (Villena-Roldán, 2012; Wolthoff, 2017). My paper lends support to these papers by providing evidence that screening plays a significant role in match formation. Other theoretical studies have explicitly modelled the impact of regulatory changes in a matching framework (Pries and Rogerson, 2005; Brochu and Green, 2013; Sengul, 2017). Their predictions for minimum wage effects on hiring standards are closely related to the hypothesis I test in this paper.[4]

One potential implication of my results is that the SMW makes it more difficult for the least productive job-seekers to find employment. Given that boosting the incomes of low-wage workers is a key objective behind introducing an SMW, this is an undesired side effect. It would be misguided however to conclude from this that the minimum wage does more damage than good. Rather, my finding points to a potential need to provide extra support to the losers from the minimum wage's effect on worker selection.

1.2 Theoretical framework

In this section I look at ways of formalizing the intuition behind the paper. My first step is to model a firm's decision whether or not to hire a worker in the most simplified set-up possible. I then review the predictions of existing papers that analyze the effect

[3]In this, my results differ from those of Clemens et al. (2018). One potential explanation is that being a high-school drop-out in the US is a stronger negative signal than having neither an *Abitur* (the high school diploma required for university entry in Germany) nor a vocational qualification is in Germany. Another explanation is that the effect on hiring standards with respect to observed characteristics is too small for me to detect with the statistical power my sample affords. Finally, note that even recent US evidence is ambiguous (e.g., Clemens et al., 2018; Cengiz et al., 2019).

[4]In Section 1.2 I use a toy model to formalize the intuition that minimum wages may raise hiring standards and discuss predictions from existing theoretical work.

of minimum wages on firms' hiring thresholds in search-and-matching frameworks, though their predictions refer to match productivity rather than worker ability.

1.2.1 A toy model

Consider a firm that uses only labor to produce a good, which it sells at price $p = 1$. The firm's output depends solely on the ability a of its single worker, where $a \sim N(m_a, \sigma_a^2)$. Before production starts, the firm receives an application from one worker. If it hires the worker, it has to pay her the market wage w, which is exogenously given and independent of ability. The firm then produces and sells all of its output. If the firm does not hire the worker, it produces zero output and incurs no costs. Thus the firm's profit is given by $\Pi(a, w) = a - w$ if it hires the applicant and by $\Pi(a, w) = 0$ otherwise.

When ability is perfectly observed, the firm will hire the worker if $a \geq w$. Similarly, when a is unobserved but its distribution is known, the firm will make a hire if $m_a \geq w$.

Now consider the case when a is imperfectly observed and the firm receives a noisy ability signal $z = a + \epsilon$, where $\epsilon \sim N(0, \sigma_\epsilon^2)$ and ϵ is uncorrelated with ability. The firm will hire the worker if, conditional on the ability signal, expected profits are non-negative:

$$\mathrm{E}[\Pi|z] = m_a + \frac{\sigma_a^2}{\sigma_a^2 + \sigma_\epsilon^2}(z - m_a) - w \geq 0 \tag{1.1}$$

or, equivalently, if the ability signal is weakly greater than the hiring threshold t:

$$z \geq \frac{1}{\sigma_a^2}[(\sigma_a^2 + \sigma_\epsilon^2)w - \sigma_\epsilon^2 m_a] = t. \tag{1.2}$$

This threshold goes up when the wage is increased (such as by a minimum wage introduction that bites):

$$\frac{dt}{dw} = \frac{\sigma_a^2 + \sigma_\epsilon^2}{\sigma_a^2} > 0. \tag{1.3}$$

This toy model formalizes the intuition that when workers become more expensive, firms' hiring becomes more selective. It yields the prediction tested in this paper that the introduction of a biting minimum wage will increase firms' hiring thresholds.[5] See Section 1.7.1 in the appendix for the proofs.

[5] In my empirical analyses I proxy the hiring threshold (the lowest ability signal the firm would tolerate) with my ability measure for the least productive new hire, which is not exactly the same. Note, however, that the expected ability of a new hire with $z = t$ is simply the wage and thus will also go up when a minimum wage is introduced.

1.2.2 Theoretical predictions from the literature

Pries and Rogerson (2005) introduce screening to the hiring process. They consider a matching model with employer learning that treats the quality of the worker-firm match as both an inspection and an experience good. In their framework, employers hire workers when the match productivity signal they receive is above a certain threshold. Studying the impact of various types of regulation in this set-up they predict that a minimum wage increase will raise firms' hiring thresholds. Brochu and Green (2013) also model screening in a matching model with endogenous separations. They, too, obtain the analytical prediction that, under some reasonable conditions, minimum wages will increase the hiring standard. Both Pries and Rogerson (2005) and Brochu and Green (2013) model screening as costless to the firm. Sengul (2017) relaxes the assumption that screening is costless and explicitly models firms' choice of screening. She also gets the analytical result that increases in minimum wages increase the threshold value of the match productivity signal. The mapping of these models' predictions to the hypothesis tested in this paper is imperfect in the sense that they all view workers as ex-ante homogeneous. That is, ex-post match productivity differs across worker-firm matches but ex-ante worker ability does not. It is reassuring that the set-up I consider above—which is highly simplified but uses heterogeneous workers—yields a similar prediction for the ability signal threshold.

1.3 Institutional setting and data

1.3.1 Policy shock: statutory minimum wage introduction

Germany's statutory minimum wage (SMW) was set at EUR 8.50 per hour (gross) and came into effect on 1 January 2015. It applied to all of Germany's federal states (*Länder*) equally, irrespective of prior regional differences in wage levels.[6] Certain groups of workers were temporarily exempted:

- Apprentices for the duration of their vocational training, teenagers up to their 18th birthday and interns for a maximum duration of three months.
- The long-term unemployed for the first six months of their new job.[7]
- Employees subject to a pre-existing sectoral minimum wage set below the new SMW, for a one-year transition period ending on 31 December 2015.

[6] How many workers would benefit from the SMW was the subject of considerable debate in 2014. Its bite in my sample of employers will be discussed in Section 1.5.

[7] The law adopted a strict definition of long-term unemployment: it applied only to individuals entering employment who had been unemployed for at least a year without interruption.

Figure 1.1: Timing of minimum wage introduction

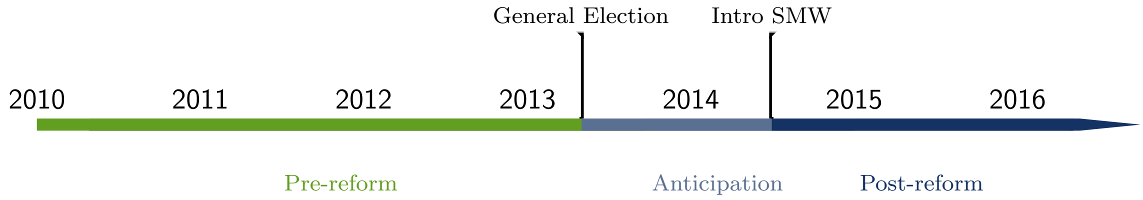

Anticipation The timing of the SMW introduction is sketched in Figure 1.1. Though the statutory minimum wage was written into law in August 2014, its introduction was widely anticipated earlier, most likely from the end of 2013 on. By December 2013, the general election of 22 September 2013 had resulted in a grand coalition between the Conservatives (CDU/CSU) and the Social Democrats (SPD). The SPD—the junior coalition partner—had managed to push into the coalition agreement its election manifesto pledge of a statutory minimum wage of EUR 8.50.[8] In this the SPD had triumphed over the Conservatives, who in their election manifesto had rejected a statutory minimum wage, instead promoting additional sector- and region-specific minimum wages agreed by collective bargaining. As the Conservatives had been favoured to win the general election, it was difficult to anticipate the introduction of a statutory minimum wage prior to the general election, i.e., earlier than 22 September 2013.[9] This view is supported by Figure 1.2, which shows an index of Google web searches for "*Mindestlohn*" (minimum wage) in Germany over time: the graph remains broadly flat from 2010 through the general election of September 2013, when it starts rising.

Concurrent policy changes The timing of Germany's SMW introduction on 1 January 2015 coincided with only one other labor market reform: the *Pflegestärkungsgesetz* granted workers the right to ten days of unpaid vacation if they were unexpectedly forced to care for a relative. During this time they would receive a government transfer (*Pflegeunterstützungsgeld*). This is a rather marginal policy change unlikely to have interacted with the SMW introduction. Policy changes that became effective at other times in 2014 and 2015 made pensions more generous (July 2014), increased the social transfers to asylum seekers (March 2015) and ended the requirement that asylum seekers only be considered for jobs if no EU nationals are available. In addition, an option of doubling the duration of

[8] The result of the SPD's internal vote on the coalition agreement was announced on 14 December 2013.

[9] Newspapers reported at the end of May 2013 that Chancellor Merkel had reaffirmed her rejection of an SMW (https://www.handelsblatt.com/politik/deutschland/tagung-in-muenster-merkel-bekraeftigt-nein-zum-gesetzlichen-mindestlohn/8253832.html, http://www.spiegel.de/politik/deutschland/union-zum-gesetzlichen-mindestlohn-merkel-lehnt-forderung-der-cda-ab-a-901892.html). Only at the end of October were there reports that if a Grand Coalition came to pass it would likely introduce an SMW (see, e.g., https://www.zeit.de/wirtschaft/2013-10/mindestlohn-auswirkungen-arbeitsmarkt).

Figure 1.2: Timing of minimum wage introduction

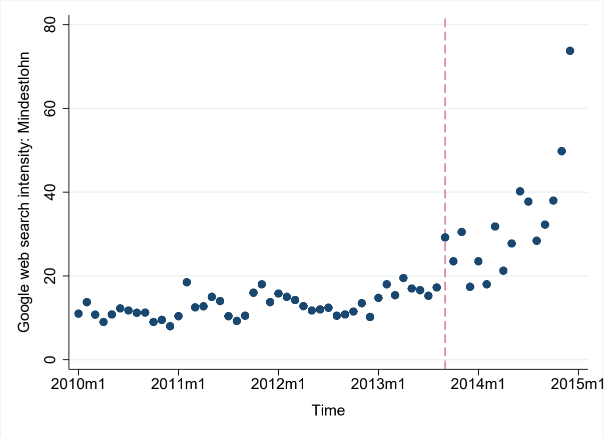

paid parental leave by combining it with part-time work was introduced (July 2015). The only one of these reforms that may threaten to affect the quality of hires is the reduction of employment barriers for asylum seekers. However, it is unlikely that this loosening will confound the results of the minimum wage as it made it easier for all employers to hire (potentially more qualified) asylum seekers, not just for those establishments affected by the minimum wage.[10]

1.3.2 Data source

My main data source is the Integrated Employment Biographies (IEB). These data are provided by FDZ, the research data centre of the German Federal Employment Agency's research institute (IAB). The IEB are longitudinal individual-level spell data that are assembled from social security records. They contain information on workers' age, gender, education, qualification, occupation, industry, the timing of entry into and exit from employment, and total (annual) earnings from each employment spell. There are person and establishment identifiers allowing me to follow individuals over time and to observe who is employed at the same establishment. The data do not contain firm identifiers. As a consequence, several observed establishments may be part of the same company but I will not know.

The data underlying my empirical analysis are a flow sample of the workforces of 1,520 German establishments for the period from 2010 through 2016. That is, they contain information on all individuals ever formally employed at these establishments during those seven years, even if only for a day. The sample of establishments is representative for the

[10]In practice, a bias from the liberalization of the German labor market for asylum seekers is ruled out by my measure of individual productivity: AKM worker effects are only available for individuals who have been in regular employment in Germany before. See Section 1.4.3 for details.

German private sector by size, industry and region strata, but excluding establishments that employ fewer than 50 workers. As the empirical analysis is at the establishment level I aggregate individual-level spell information in establishment-year cells.

I use the employment records of full-time, part-time and marginal workers and exclude the remaining groups such as interns, apprentices and people in a part-time transition to retirement. The resulting sample includes 4,913,135 observations for 918,056 workers at 1,520 establishments. Most relevant for my analysis of firm's selectiveness in hiring is the sample of worker inflows during 2010-16. This hire sample includes 448,013 observations for 440,813 workers at 1,512 establishments.

The final analysis sample is at the establishment-year level and includes the 1,491 establishments for which I can measure pre-reform minimum wage exposure and hire quality in at least one year.[11]

1.4 Method

1.4.1 Identification strategy

To estimate how employers adjusted the selectiveness of their hiring to the introduction of the minimum wage I look for a group of establishments that were untreated by the policy change but similar in their hiring behaviour. I use the following difference-in-differences (DiD) specification:

$$y_{jt} = \alpha + \beta TR_j * POST_t + \gamma_t + \delta_j + \epsilon_{jt}, \tag{1.4}$$

where I assign establishment j to the treated group ($TR_j = 1$) if it pays one or more workers strictly less than the future minimum wage in the pre-reform years (before policy anticipation is possible). The control group consists of all establishments that do not employ any minimum-wage workers in the pre-reform period. The pre-reform period includes the years 2010-2013; the post-reform period, $POST_t = 1$, is 2015-2016. I drop observations for the year 2014 from this DiD specification: in this year, the minimum wage is not yet effective but there may already be an anticipation effect on hiring behaviour. There are fixed effects for years, γ_t, and establishments, δ_j. Beyond using establishment fixed effects to control for time-invariant heterogeneity I do not include co-variates. y_{jt} measures establishments' hiring standards (see Section 1.4.3 fur further detail). The DiD coefficient $\hat{\beta}$ provides an estimate of the effect of the minimum wage introduction on establishments' worker selection. Standard errors are clustered at the establishment level.[12]

[11] The hire quality data I use are also provided by IAB. See Section 1.4.3 for details.

[12] Some papers applying a similar identification strategy weight regressions with log employment (see, e.g., Harasztosi and Lindner, 2018). I show in the appendix that doing so leaves my results qualitatively unchanged (see Figure 1.6).

The key identifying assumption is $\mathbb{E}\left[\epsilon_{jt}|TR_j * POST_t\right] = 0$: the selectiveness of hiring in treated and untreated establishments follows parallel counter-factual trends, i.e., would have been parallel in the entire analysis period if the minimum wage had not been introduced.[13] I can directly test whether this holds only for the pre-reform years by considering an event-study version of the DiD specification:

$$y_{jt} = \alpha + \sum_{t}(\beta_t TR_j * \gamma_t) + \gamma_t + \delta_j + \epsilon_{jt}. \tag{1.5}$$

Now β_t is a year-specific DiD coefficient and 2013 is the reference year. Parallel pre-reform trends in outcomes would mean that $\hat{\beta}_t$ is not significantly different from zero for 2010-2012. Anticipation effects of the minimum wage would show up as $\hat{\beta}_{t=2014}$ different from zero. $\hat{\beta}_t$ for 2015-2016 provides estimates of the minimum wage's short-run effect. While it is impossible to test for parallel counter-factual trends in the post-reform period a placebo analysis in Section 1.6.1 provides suggestive evidence in their support.

1.4.2 Measuring establishment minimum wage exposure

My empirical strategy compares establishments treated by the minimum wage introduction to untreated establishments. I use establishments' pre-reform wage structure to determine their exposure to the SMW. My measure for SMW exposure is an indicator for whether any of the establishment's workers earned less than the future SMW (deflated using the CPI) in the years 2010-2013, when the minimum wage was not anticipated. As such, this dummy varies at the establishment level (but not over time). I follow Harasztosi and Lindner (2018) in averaging SMW exposure over the pre-reform period.[14] There is a theoretical motivation for my focus on a binary measure for SMW exposure rather than, e.g., the fraction of sub-SMW workers or a wage gap measure: I am looking for a shift in employers' hiring thresholds. If an establishment's lowest-paid workers exogenously become more expensive and this raises its hiring standard, recruitment should avoid workers with a productivity below the threshold - irrespective of whether the minimum wage bites for just a handful or many of its workers.[15]

The minimum wage sets a EUR 8.50 floor for hourly wages. The IEB data contain average daily wages but not hours worked. I therefore approximate hourly wages by focusing on establishments' full-time workers and dividing their daily wages by 8, the length of a standard working day in Germany. This means that my SMW exposure measure does not consider establishments' part-time workers and mini jobbers. This may be problematic if

[13]See Section 1.7.2 in the appendix for a discussion of a further, more subtle, assumption.

[14]Column (4) of Table 1.8 in the appendix shows measuring SMW exposure only in July 2013 produces qualitatively similar results.

[15]In fact, if there are frictions in the hiring process, it is more difficult to shift the hiring standard up for establishments using more sub-MW workers. While a continuous MW exposure measure is thus theoretically less appealing it produces broadly similar results—see Column (2) of Table 1.8 in the appendix.

SMW exposure is measured with more error at establishments more heavily affected by the minimum wage, as may be the case if these establishments also employ a higher share of part-timers or mini jobbers. Indeed there is evidence that the share of minimum wage earners is particularly high among these groups, highlighting this concern (Machin et al., 2003). I next discuss how I deal with the potential bias arising from this.

Limiting non-classical measurement error In this section I argue that focusing on a binary, extensive-margin measure for SMW exposure helps me to limit potential biases from non-classical measurement error. First consider treated establishments. The share of part-timers and mini jobbers is higher at treated than at untreated establishments, as Table 1.1 shows. Existing evidence suggests that a greater share of part-timers and mini jobbers than of full-timers earn below the minimum wage (Machin et al., 2003). At the intensive margin, this would make me underestimate the SMW bite more at treated establishments. With the extensive-margin exposure measure, however, treated establishments are correctly classified even if I ignore part-timers and mini jobbers. I will mis-classify some establishments as untreated - namely the ones where full-timers earn above the minimum wage but part-timers or mini jobbers earn below it. Note, however, that falsely assigning some treated employers to the control group in this way should bias my estimates towards zero. This is how using an extensive-margin measure for SMW exposure should help me to avoid overestimating the SMW effect.

I attempt to go one step further in reducing measurement error by using supplementary information on individual hours for a sub-set of my data. These data come from the German accident insurance. The hours data have hardly been used for scientific purposes because of a number of reasons: (1) they are available only for the years 2010-2014; (2) they contain a substantial share of missing values; (3) they contain either contractual, actual or estimated working hours, but it is unknown which of the three. When investigating questions at the individual level, the third point can be addressed by applying establishment fixed effects and assuming that reporting differences within employers are negligible, making within-establishment comparisons valid. In my setting, however, it is precisely the differences (in hourly wages) across establishments that I am interested in. Relying on hours data that report different things across establishments will thus almost certainly introduce additional noise, which is why I view it as a robustness check. Column (3) of Table 1.8 in the appendix shows that my main results remain qualitatively similar when I use the supplementary hours data to compute hourly wages for the construction of an SMW exposure measure.

1.4.3 Measuring employers' hiring standards

I measure establishment j's hiring standard using the ability of its least productive new hire i in a given year t, $\min_{j,t}\{ability_i\}$. The idea behind focusing on the bottom of the new hire productivity distribution is to be able to detect the shift of employers' hiring thresholds predicted by my toy model in Section 1.2. In addition, the minimum wage has the largest bite for the lowest-earning workers, so that I expect its effect on worker selection to be concentrated in the left tail of the new hire ability distribution.[16]

Underlying $\min_{j,t}\{ability_i\}$ is a measure of individual worker productivity.[17] For this I use estimated worker effects from a two-way fixed-effects log-wage regression as introduced by Abowd et al. (1999).[18] These AKM worker effects capture the portion of workers' time-invariant productivity that is reflected in their wage earnings, including not only observable human capital but also typically unobserved aspects of productivity. Among Swedish workers in the 2000s, for example, AKM worker effects detect significant variation in both non-cognitive and cognitive ability even conditional on the level of education (Butschek and Sauermann, 2019b). For the purpose of measuring employers' hiring standards it is an advantage that AKM worker effects (a) vary continuously and (b) can distinguish between the ability of workers with the same level of schooling. Educational qualifications alone may be too coarse to differentiate between the productivity of very low-paid workers. They may also be less relevant for performance in a minimum wage job than, e.g., conscientiousness, reliability or common sense.

I use the AKM worker effects estimated by Card et al. (2013) for the interval 2002-2009. CHK implement the following AKM specification:

$$\ln(w_{ijt}) = \alpha_i + \psi_j + \gamma_t + x'_{it}\beta + r_{ijt}, \tag{1.6}$$

where $\ln(w_{ijt})$ is the natural logarithm of individual i's wage at establishment j in year t. There are additive fixed effects for individuals (α_i) and establishments (ψ_j) as well as a vector of time-varying individual-level controls (x_{it}), including age squared and age cubed as well as education categories interacted with the year dummies, age squared and age cubed. Using the largest connected set of establishments in West Germany and the universe of workers employed there during the interval from 2002 to 2009, CHK estimate Equation (1.6) separately for men and women.

Matching CHK's worker effect estimates with my sample of 2010-16 hires gives me an individual ability measure for 38.86% of hired workers.[19] However, I observe an ability measure for at least one hire at 99.01% of establishments. Measuring ability of only a

[16] Despite the theoretical appeal of using minimum hire quality as a proxy for hiring standards one may be worried that the minimum makes the analysis vulnerable to outliers. In Section 1.7.3 of the appendix I also look at the SMW's effect on P10, P25 and the median. Reassuringly, the effect is also positive for these aggregations but becomes smaller and turns insignificant as one moves up the hire quality distribution—see Table 1.9.

[17] I use the words ability and productivity interchangeably.

[18] Estimated individual fixed effects have gained some popularity as a broad measure of worker quality, see the review in Butschek and Sauermann (2019b).

[19] One reason for the low individual-level coverage is that CHK's estimates are only for West Germany. In a robustness check I use a preliminary update of CHK's AKM worker effects that includes East German workers. This expands individual coverage to 51.58% of hired workers and produces similar results (see Figure 1.7 in the appendix).

subset of hires would be problematic if this selectively affected treated employers after the reform. I explore this issue in Section 1.6.1.

1.4.4 Summary statistics

Table 1.1: Firm characteristics by treatment status

A: Continuous characteristics	Treated		Control	
	mean	sd	mean	sd
Head count (1 Jan)	362.629	1477.994	154.849	184.199
Establishment age (yrs)	21.977	11.682	21.080	11.605
Female worker share	0.349	0.251	0.259	0.217
Mean worker age	42.285	3.920	43.930	3.296
Mean wage	94.481	31.903	117.158	25.494
Part-time worker share	0.181	0.220	0.117	0.167
Mini-job worker share	0.064	0.111	0.030	0.070
Mean daily starting wage (EUR)	79.307	26.309	99.785	23.438
Hire rate	0.255	0.414	0.147	0.302
Hire share with AKM FE	0.387	0.200	0.431	0.236
Observations	1,147		344	

B: Binary characteristics	Treated		Control	
	Share Yes	Frequency	Share Yes	Frequency
Manufacturing	0.295	338	0.358	123
Metal, electrical, automotive	0.251	288	0.328	113
Trade, transport, news	0.155	178	0.157	54
Business/financial services	0.177	203	0.087	30
Information, communication	0.122	140	0.070	24
Former East Germany	0.307	352	0.291	100
Observations	1,147		344	

Note: This table summarizes treatment and control establishments' average pre-reform characteristics for the main estimation sample. Panel A provides mean and standard deviation of continuous variables. Panel B gives means and frequencies for dummy variables.

Table 1.1 provides descriptive statistics from 2013 for the estimation sample separately for treated and control establishments. Panel A reports means and standard deviations for continuous characteristics. The average treated establishment is larger and employs a greater share of women, part-time workers and mini jobbers. At treated establishments the mean wage is 19.3% lower, the mean starting wage is 20.5% lower and the hiring rate is larger. Panel B gives the share and number of establishments that display certain binary characteristics. The share in manufacturing as well as in metal, electrical and automotives is smaller among treated establishments. Their share in business or financial services and information and communication is bigger. Their share in trade, transport and news is similar, as is their share located in former East Germany. Note that those establishment characteristics that are not endogenous to the reform—i.e., mostly industry and location—are relatively time-invariant and will be captured by the establishment fixed effects in the regression specifications.

1.5 Main results

I start with a plausibility test of my minimum wage exposure measure. That is, I test whether the SMW introduction had any bite for those establishments I identify as treated. If so, it should increase the wages establishments pay their lowest-earning new hires. Figure 1.3 presents the yearly DiD coefficients $\hat{\beta}_t$ from estimating the event-study regressions in Equation (1.5) with establishments' minimum starting wages as an outcome. 2013 is the reference year. The DiD estimates for 2010-2012 are indistinguishable from zero, confirming the validity of the parallel-trends assumption for the pre-reform period. The 2014 coefficient shows that there is no significant anticipation of the minimum wage introduction in the bottom starting wages employers pay. Reassuringly, starting wages react positively to the introduction of the minimum wage at treated establishments.

Figure 1.3: Minimum wage effect on minimum starting wages: event study

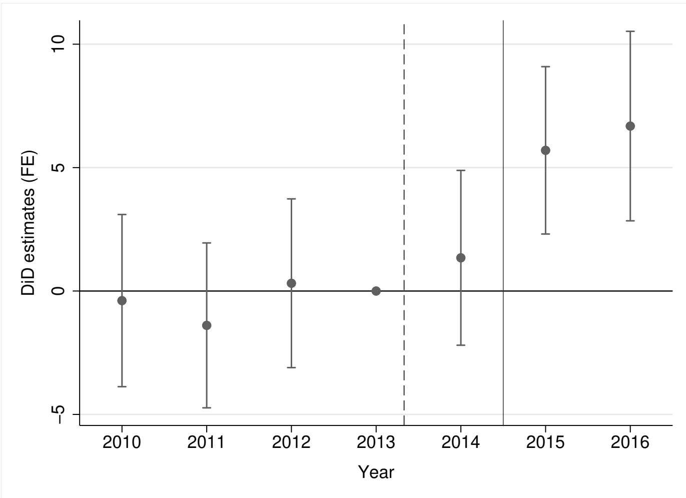

Note: This figure shows yearly DiD estimates for the effect of the minimum wage introduction on minimum starting wages (daily, in EUR). Estimates are from an establishment fixed-effects specification without covariates other than year dummies. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the establishment level.

Having established that the minimum wage had a bite for those establishments that I assign to the treated group I move on to the central question whether the minimum wage introduction made employers hire more selectively. To estimate the dynamic effects on minimum hire quality I again start with the fixed-effects event-study framework from Equation (1.5). Figure 1.4 presents the estimation results, revealing a persistent positive effect of the 2015 minimum wage introduction of nearly 20 per cent of a standard deviation. It also shows a significant anticipation effect in 2014. This suggests that employers already factored higher future wages into their recruitment decisions the year before the minimum wage was binding. For the pre-reform period, the DiD estimates are closely centred around zero, providing evidence in support of parallel pre-reform trends.

Figure 1.4: Minimum wage effect on minimum hire quality: event study

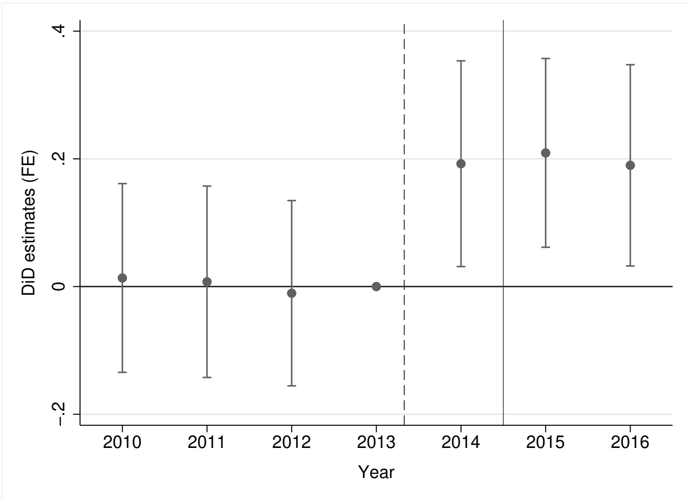

Note: This figure shows yearly DiD estimates for the effect of the minimum wage introduction on minimum hire quality. Estimates are from an establishment fixed-effects specification without covariates other than year dummies. Hire quality is measured by CHK's individual AKM worker effects estimated 2002-2009. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the establishment level.

Table 1.2 condenses the results from Figures 1.3 and 1.4 by reporting before-vs-after DiD estimates (see Equation (1.4)). Here I exclude data from the year 2014 to make sure my estimates do not capture anticipation effects. Column (1) shows that the statutory minimum wage increases new hires' minimum daily pay by about EUR 6.60. As illustrated by Column (2), minimum hire quality also increased, by about 18.9 per cent of a standard deviation. The remainder of this paper investigates whether this increase in minimum hire quality can indeed be attributed to employers hiring more selectively.

1.6 Discussion

1.6.1 Threats to identification

Systematic ability mis-measurement The individual ability measure used in this paper, estimated AKM worker fixed effects, doubtlessly contains substantial measurement error. Thanks to the DiD-strategy, however, even non-classical measurement error will only lead to an overestimation of the effect of the minimum wage on worker selection

Table 1.2: Minimum wage effect on starting wages and hire quality

	(1) Starting wages	(2) Hire quality
DiD estimate (Treated*Post)=1	6.5550***	0.1888***
	(1.2219)	(0.0496)
Observations	8,288	7,591
Establishments	1,491	1,491
Adjusted R^2	0.0234	0.0150

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the establishment level. Dependent variable is minimum starting wage (1) and minimum hire quality (2), where hire quality is measured by CHK's individual AKM worker effects estimated 2002-2009. Estimates are from establishment fixed-effects specifications without covariates other than year dummies. The year 2014 is excluded to rule out anticipation effects.

under very specific conditions. This would be the case if pre-determined ability had been systematically overestimated for new hires entering treated establishments after the reform, i.e., if the bias of the individual ability estimate were correlated with where and when the individual is eventually hired. This is highly unlikely: the AKM worker effect used in the 2010-2016 analysis period is estimated using a non-overlapping earlier period (2002-2009) and is necessarily based on employment at other establishments than the one eventually hiring the worker.

Instead a potential threat to identification stems from the aggregation of individual ability measures to the firm level. As detailed in Section 1.4.3 I do not have an AKM worker effect estimate for all new hires. If a smaller share of hires were observed at treated establishments after the reform, I might mechanically overestimate minimum hire quality. Such selective attrition may produce the same pattern as a change in the selectiveness of hiring. I test for the presence of selective attrition by using the share of new hires for whom an estimated AKM worker effect is available as an outcome. Column (1) of Table 1.3 suggests that the reform had no significant effect on the proportion of new hires for whom I have an ability measure, allaying concerns about selective attrition.

Table 1.3: Minimum wage effect on availability of hire quality measure

	(1) AKM coverage	(2) Female hires	(3) Hire age
DiD estimate (Treated*Post)=1	0.0020	-0.0105	0.2691
	(0.0117)	(0.0093)	(0.2838)
Observations	8,473	8,473	8,473
Establishments	1,491	1,491	1,491
Adjusted R^2	0.0952	0.0003	0.0112

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the establishment level. In all columns the dependent variable reflects certain characteristics of establishments' new hires: in (1) the share of hires for whom an AKM worker effect is available; in (2) the female hire share; and in (3) new hires' average age. DiD estimates are from establishment fixed-effects specifications without covariates other than year dummies. The year 2014 is excluded to rule out anticipation effects.

A related threat arises from the possibility that AKM worker effects systematically under-estimate individual ability for certain groups and that the minimum wage leads to fewer members of these groups being hired. This combination of a shift in the composition of hires and systematic mis-measurement for the newly under-represented groups may drive observed minimum hire quality up at treated establishments without the increase in hiring standards I am attempting to identify. I can test this possibility for two groups: women and younger workers. First, to the extent that there is still some discrimination against women in the labor market AKM worker effects may underestimate ability for women relative to men. It is therefore reassuring that the minimum wage did not significantly reduce the share of women hired (see Column (2) of Table 1.3). Second, there is evidence that AKM worker effects underestimate ability for younger workers (Butschek and Sauermann, 2019b). Column (3) of Table 1.3 suggests, however, that this is unlikely to drive my results as the minimum wage did not significantly increase hires' average age.

Non-parallel counter-factual trends While the pre-reform trends in outcomes are parallel in the treatment and control groups, it is possible that minimum hire quality of treated and untreated establishments would have diverged (or converged) in the absence of the minimum wage introduction. The identifying assumption of parallel counter-factual trends itself is untestable. It is possible, however, to use a placebo test as indirect evidence for its plausibility. The variation in minimum wage exposure I use in this paper comes from differences in establishment wage structures. I stick to this idea in constructing a placebo treatment: using only employers untreated by the minimum wage, I split them into two roughly equal-sized groups based on their average wages. Those establishments with average wages below the median are assigned the placebo treatment and those with average wages at the median or greater make up the placebo control group. Figure 1.5 reports the results from estimating event-study regressions of minimum hire quality on the placebo treatment. There is no clear pattern to the yearly coefficient estimates $\hat{\beta}_t$ and none of them is significantly different from zero. I interpret this as indicative evidence for the validity of the assumption of parallel counter-factual trends during the post-reform period.

1.6.2 Alternative interpretations

Mechanical effect of fewer hires An alternative to the firm selectiveness explanation for the minimum wage effect on minimum hire quality is that a reduced number of hires mechanically increased minimum hire quality. Suppose that employers hire randomly, not screening workers at all. Then, hires into a given establishment are a random draw from the overall ability distribution of available workers. If the hiring rate goes down, this will reduce the ability dispersion of new hires, raising hire quality at the bottom of the establishment hire quality distribution. But it will symmetrically reduce quality at the top.

Figure 1.5: Placebo effect on minimum hire quality: event study

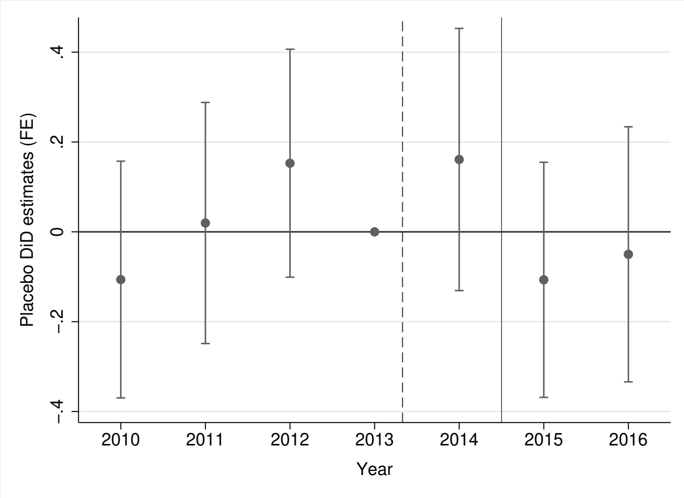

Note: This figure shows yearly DiD estimates for the effect of a placebo treatment on minimum hire quality. Estimates are from an establishment fixed-effects specification without covariates other than year dummies. Establishments affected by the actual treatment (minimum wage) are excluded. The placebo treatment is assigned to those untreated establishments with below-median average wages. Hire quality is measured by CHK's individual AKM worker effects estimated 2002-2009. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the establishment level.

Thus, if the minimum wage substantially reduced the number of hires, this may have mechanically increased minimum hire quality even with random hiring, delivering the empirical pattern I have hitherto interpreted as evidence of raised hiring standards. In addition, if it operated mechanically through the number of hires, the minimum wage should have reduced maximum hire quality. I test these predictions by estimating Equation (1.4) with the number of hires and maximum hire quality as outcome variables. As Table 1.4 illustrates, I neither find a significant negative effect of the minimum wage on the number of hires nor a significant reduction of maximum hire quality. This suggests that if the mechanical effect of fewer hires is at play it is too small to be detectable with the available statistical power. It is therefore an unlikely explanation for the effect of the minimum wage on minimum hire quality.

Demand or supply So far in this paper I have interpreted the changes in minimum hire quality induced by the introduction of the minimum wage as evidence of shifts in firm hiring standards, i.e., a change in the demand for the quality of labor. In principle, however, they could also have been driven by changes in self-selection of workers—i.e., the supply of the quality of labor. Due to the minimum wage, a greater number of productive workers may apply to employers treated by the minimum wage as these are now more attractive—their wages (at least at the bottom) improve relative to those of untreated establishments.

Table 1.4: Minimum wage effect on number of hires and maximum hire quality

	(1) Number of hires	(2) Maximum hire quality
DiD estimate (Treated*Post)=1	-1.3134	-0.0135
	(2.6061)	(0.0480)
Observations	8,473	7,591
Establishments	1,491	1,491
Adjusted R^2	0.0033	0.0054

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the establishment level. Dependent variable is the number of hires (1) and maximum hire quality (2), where hire quality is measured by CHK's individual AKM worker effects estimated 2002-2009. Estimates are from establishment fixed-effects specifications without covariates other than year dummies. The year 2014 is excluded to rule out anticipation effects.

A first piece of evidence in favour of the labor demand or screening interpretation comes from the literature. In a descriptive study of the evolution of establishmens' hiring criteria and job-filling success before and after the introduction of the German minimum wage, Gürtzgen et al. (2016) find that it becomes harder to fill low-wage vacancies. This contradicts the possibility that higher-productivity workers more readily self-select into formerly minimum-wage establishments.

To go further than this I look to empirical evidence. First I test for indications of a change in worker self-selection. Ideally, one would have data on applicants rather than hires.[20] Lacking such information I consider a subset of hires for whom it is likely that the workers self-selected into their new establishments—voluntary job changers—and check whether the share of these coming from high-wage establishments goes up.[21] I then use the share of voluntary transitions from a high-wage establishment among new hires as a proxy for establishment attractiveness in the regression framework from Equation (1.4). As Column (1) of Table 1.5 shows, the introduction of the minimum wage does not make treated establishments significantly more attractive by this metric. This is a further indication that worker self-selection does not drive the increase in minimum hire quality.

Second I attempt to identify more direct evidence for the screening channel. To this end I distinguish employers by how feasible it is to screen workers. Depending on occupation, production technology and task structure it will be more or less possible to ex-ante assess worker productivity. Assuming that employers' screening intensity in the pre-reform period contains information about the feasibility of pre-hire screening, I use information from the 2012 wave of the Linked Personnel Panel (LPP), an employer survey on human resource management, to split the sample into low- and high-screening feasibility establishments.[22]

[20] It is tempting to use the LPP, an IAB survey data set on employer management practices, to look at the number of applications. Unfortunately, this information is available only for skilled workers, who earn above the minimum wage.

[21] To empirically identify voluntary job changes I use employment-to-employment (E-E) transitions with an employment gap of less than thirty days. I define those as coming from a high-wage employer whose previous employer's estimated AKM establishment effect was in ventile 11 or higher.

[22] The idea is that when screening is (nearly) impossible, employers do not bother to invest much time, deferring most worker assessment to on-the-job screening. If, on the other hand, it is possible to learn

Table 1.5: Minimum wage effect on hire quality: worker self-selection and screening feasibility

	(1) Estab. attractiveness	(2) Low feasibility	(3) High feasibility
DiD estimate (Treated*Post)=1	-0.0023	0.0999	0.3065***
	(0.0023)	(0.0692)	(0.0919)
Observations	8,473	3,608	2,233
Establishments	1,491	723	434
Adjusted R^2	0.0008	0.0133	0.0195

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the establishment level. The dependent variable in (1) is the share of hires transitioning directly from employment at a high-wage firm (above-median AKM firm effect). In (2) and (3) the dependent variable is minimum hire quality, with CHK's individual AKM worker effects estimated 2002-2009 used as a hire quality measure. (2) reports the estimated effect of the minimum wage introduction on establishments with below-median screenability of new hires; (3) does the same for establishments above the median. The proxy for screening feasibility is self-reported screening intensity for non-management hires in hours, measured in 2012. DiD estimates are from establishment fixed-effects specifications without covariates other than year dummies. The year 2014 is excluded to rule out anticipation effects.

If the reform-induced increase in minimum hire quality is due to employers raising their hiring standards then this response should be more pronounced at establishments where pre-hire screening is more feasible. As Columns (2) and (3) of Table 1.5 show, the estimated effect of the minimum wage on minimum hire quality is indeed bigger at establishments where screening plays a bigger role.[23] This is a final piece of evidence suggesting that the change in minimum hire quality is a demand-side phenomenon, driven by a shift of employers' hiring standards.

1.6.3 Observed worker characteristics and statistical discrimination

The evidence presented so far suggests that the SMW introduction in Germany raised establishments' hiring standards. A potential implication of increased employer selectiveness is that vulnerable groups in the labor market find it harder to get hired. This would be worrying, suggesting that some of the weakest labor market participants are hurt by a measure intended to improve their outcomes. To check if employers statistically discriminate against vulnerable groups I estimate the SMW's effect on the hire share of unemployed and low-skill workers, where I follow CHK's definition of low skill (education level lower than *Abitur* and no vocational qualification). Table 1.6 presents the results of estimating the DiD specification in Equation (1.4) for these outcomes.

There is no evidence that the minimum wage makes employers less likely to hire disadvantaged labor market participants along these observed dimensions. It may still be though

something substantial about the worker's productivity before hiring them, it makes sense to ex-ante screen for several hours and more. Based on this reasoning those with above-median screening intensity in 2012 are classified as employers where pre-hire screening is more feasible (see Figure 1.8 in the appendix for a pre-hire screening histogram).

[23] Figure 1.9 in the appendix provides further support for this interpretation. Splitting the sample of establishments at the pre-reform screening intensity quartiles, it suggests that the minimum wage effect on hiring standards increases monotonically in the feasibility of screening.

Table 1.6: Minimum wage effect on observable hire quality

	(1) Unemployed	(2) Low-skill
DiD estimate (Treated*Post)=1	-0.0009	0.0047
	(0.0090)	(0.0059)
Observations	8,473	8,473
Establishments	1,491	1,491
Adjusted R^2	0.0098	0.0402

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the establishment level. Dependent variable is (1) the share among new hires of formerly unemployed workers and (2) workers that have neither a vocational qualification nor a secondary school diploma that qualifies them for university entry (*Abitur*). DiD estimates are from establishment fixed-effects specifications without covariates other than year dummies. The year 2014 is excluded to rule out anticipation effects.

that in terms of the unobserved characteristics captured by estimated AKM worker fixed effects the least productive workers find it harder to find a job as the minimum wage makes employers more selective.[24]

Is this a contradiction, given that unemployment and a lack of formal qualifications are likely to be correlated with the ability captured in estimated AKM worker fixed effects? I argue that it is not. It is plausible that among the very low-pay workers who may be close to employers' hiring standards observed characteristics are less useful for distinguishing productivity on the job. This may be because, for example, educational qualifications are largely irrelevant for productivity in a cleaning job. In addition, the observed characteristics may be too coarse to differentiate between more and less able workers. One piece of evidence in support of this is that the partial correlations between estimated AKM worker effects and observed characteristics are lower among sub-minimum wage workers than among those higher up the wage distribution, especially for the low-education dummy; the share of the variance in AKM worker effects explained by low-skill and entry from unemployment is also lower. This is illustrated by Table 1.11 in the appendix, which compares a regression of AKM worker effects on observed characteristics between workers earning below and above the SMW in the pre-reform period (2010-13).

1.7 Conclusion

In this paper I study the effect of the introduction of a statutory minimum wage in Germany on the hiring standards of 1,491 private-sector employers. My empirical strategy is a difference-in-differences comparison of treated and untreated establishments, classified by whether their wage structure prior to the reform exposed them to the introduction of the

[24]For consistency with my main analysis, I use the binary SMW exposure measure as described in Section 1.4.2. However, given that the outcome here is not the hiring threshold but the share of a particular group among new hires, a continuous measure of SMW exposure is arguably more appropriate. Table 1.10 in the appendix reproduces the analysis using a continuous SMW bite measure (the pre-reform share of sub-MW workers). It also does not show a statistically significant effect of the SMW introduction on the share of hires from unemployment or with low education.

SMW. To identify employers' hiring standards I do not rely on manager-reported intentions or hiring criteria but on the quality of establishments' realized hires, using the ability of the least productive new hire in a year as a proxy for the hiring standard. The worker-level data I use allows me to observe every single worker inflow during a period of four years pre-reform and two years post-reform. Matching these with the AKM worker effects estimated by Card et al. (2013) gives me an individual-level measure of time-invariant productivity. Using AKM worker effects as ability proxies implies that I measure employers' hiring standards with respect to both observed and unobserved worker productivity.

My DiD estimate of the SMW introduction's effect on establishments' minimum hire quality is an increase of 18.9% of a standard deviation. That this result identifies a causal effect is primarily threatened by two scenarios. The first is selective attrition or measurement, where the availability or accuracy of my worker ability proxy is systematically different for treated establishments after the reform. I show that this is not the case. The second potential threat is a deviation from parallel counter-factual trends after the reform. I provide suggestive evidence that this is unlikely: a placebo test shows that trends are roughly parallel between two groups of untreated employers that, like the treated and control groups, differ in their pre-reform wage structure.

I also explore the possibility that while the reform's effect on minimum hire quality is real, it does not reflect an increase of employers' hiring standards. There are two alternative explanations I address: the first is that a reform-induced reduction in the hiring rate mechanically increased establishments' minimum hire quality by reducing the sample they draw from the worker ability distribution. Such a mechanism requires a drop in the hiring rate and implies a reduction in maximum hire quality. I find neither. The second alternative interpretation is that a change in workers' self-selection improved the applicant pool of treated establishments. Thus a shift in the quality of labor supplied, rather than raised hiring standards, could be driving the increase in minimum hire quality. I offer two pieces of evidence supporting the interpretation that demand for a higher quality of labor is the primary driver of my result. First, the share of voluntary job changers coming from high-wage establishments did not increase at treated establishments, suggesting they did not become more attractive for workers. Second, the effect of the SMW introduction on minimum hire quality is concentrated in those establishments that report more pre-hire screening in the pre-reform period, i.e., where it is more feasible for employers to distinguish worker quality before hiring.

The main contribution of this paper is providing some of the first causal evidence that a minimum wage raises employers' hiring standards. This is consistent with the conclusion of the only other causal study on the question, written contemporaneously and independently (Clemens et al., 2018), though my paper emphasizes hiring standards with respect to unobserved productivity. Some disagreement remains between our studies whether increased hiring standards also materialize in new hires' observed characteristics such as age and education; like the majority of papers looking at such labor-labor substitution, I do not find evidence for substitution toward older or more highly educated workers. The most likely explanation for this is that these observed characteristics do not sufficiently differentiate between the productivity of the relevant low-wage workers.

A potential implication of my findings for policy-makers is that even without negative aggregate employment effects there may be some losers from the introduction of a minimum wage—namely those who are deemed too unproductive to be hired at the new wage floor. Such unemployed workers may benefit more from productivity-enhancing interventions than from measures aimed at improving the matching process (such as application training). Another implication pointed out by Fairris and Bujanda (2008) is that the wage-boosting effect of the minimum wage for low earners may be "dissipated" by increased hiring standards: the more productive the workers who are hired at or just above the new minimum wage, the higher their earnings potential would have been elsewhere, too. This would dampen the redistributive power of the minimum wage.

References

Abowd, John M, Francis Kramarz, and David N Margolis (1999): "High wage workers and high wage firms," *Econometrica*, 67(2): 251–333.

Bernini, M and R Riley (2016): "Exploring the Relationship between the NMW and Productivity," *London: LPC*.

Brochu, Pierre and David A. Green (2013): "The Impact of Minimum Wages on Labour Market Transitions," *Economic Journal*, 123(573): 1203–1235.

Butschek, Sebastian and Jan Sauermann (2019b): "Do individual fixed effects capture worker ability?" unpublished manuscript, Stockholm University.

Card, David, Jörg Heining, and Patrick Kline (2013): "Workplace heterogeneity and the rise of West German wage inequality," *The Quarterly Journal of Economics*, 128(3): 967–1015.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer (2019): "The Effect of Minimum Wages on Low-Wage Jobs," *The Quarterly Journal of Economics*.

Clemens, Jeffrey, Lisa Kahn, and Jonathan Meer (2018): "Dropouts need not apply: The minimum wage and skill upgrading," Tech. rep., Working Paper.

Fairris, David and Leon Fernandez Bujanda (2008): "The dissipation of minimum wage gains for workers through labor-labor substitution: evidence from the Los Angeles living wage ordinance," *Southern Economic Journal*, 473–496.

Giuliano, Laura (2013): "Minimum wage effects on employment, substitution, and the teenage labor supply: Evidence from personnel data," *Journal of Labor Economics*, 31(1): 155–194.

Gürtzgen, Nicole, Alexander Kubis, Martina Rebien, and Enzo Weber (2016): "Neueinstellungen auf Mindestlohnniveau: Anforderungen und Besetzungsschwierigkeiten gestiegen," Tech. rep., IAB-Kurzbericht.

Harasztosi, Péter and Attila Lindner (2018): "Who Pays for the Minimum Wage?" *American Economic Review*, forthcoming.

Hirsch, Barry T, Bruce E Kaufman, and Tetyana Zelenska (2015): "Minimum wage channels of adjustment," *Industrial Relations: A Journal of Economy and Society*, 54(2): 199–239.

Machin, Stephen, Alan Manning, and Lupin Rahman (2003): "Where the minimum wage bites hard: Introduction of minimum wages to a low wage sector," *Journal of the European Economic Association*, 1(1): 154–180.

Neumark, David and William Wascher (1995): "Minimum-wage effects on school and work transitions of teenagers," *The American Economic Review*, 85(2): 244–249.

Portugal, Pedro and Ana Rute Cardoso (2006): "Disentangling the minimum wage puzzle: an analysis of worker accessions and separations," *Journal of the European Economic Association*, 4(5): 988–1013.

Pries, Michael and Richard Rogerson (2005): "Hiring policies, labor market institutions, and labor market flows," *Journal of Political Economy*, 113(4): 811–839.

Sengul, Gonul (2017): "Learning about match quality: Information flows and labor market outcomes," *Labour Economics*, 46: 118–130.

Villena-Roldán, Benjamin (2012): "Aggregate implications of employer search and recruiting selection," unpublished manuscript, University of Chile.

Wolthoff, Ronald (2017): "Applications and Interviews: Firms' Recruiting Decisions in a Frictional Labour Market," *The Review of Economic Studies*.

Appendix to Chapter 1

1.7.1 Proofs of theoretical claims

Ad Equation (1.1):

$$\begin{aligned}
\mathrm{E}[\Pi|z] &= \mathrm{E}[a - w|z] \\
&= \mathrm{E}[a|z] - w \\
&= \mathrm{E}[a] + \frac{\mathrm{Cov}[a,z]}{\mathrm{Var}[z]}(z - \mathrm{E}[z]) - w \\
&= \mathrm{E}[a] + \frac{\mathrm{Cov}[a,a+\epsilon]}{\mathrm{Var}[a+\epsilon]}(z - \mathrm{E}[a+\epsilon]) - w \\
&= m_a + \frac{\sigma_a^2}{\sigma_a^2 + \sigma_\epsilon^2}(z - m_a) - w.
\end{aligned}$$

Ad Equation (1.2):

$$\begin{aligned}
\mathrm{E}[\Pi|z] &\geq 0 \\
\Leftrightarrow m_a + \frac{\sigma_a^2}{\sigma_a^2 + \sigma_\epsilon^2}(z - m_a) - w &\geq 0 \\
\Leftrightarrow \frac{\sigma_\epsilon^2}{\sigma_a^2 + \sigma_\epsilon^2}m_a + \frac{\sigma_a^2}{\sigma_a^2 + \sigma_\epsilon^2}z - w &\geq 0 \\
\Leftrightarrow \frac{\sigma_a^2}{\sigma_a^2 + \sigma_\epsilon^2}z &\geq w - \frac{\sigma_\epsilon^2}{\sigma_a^2 + \sigma_\epsilon^2}m_a \\
\Leftrightarrow z &\geq \frac{\sigma_a^2 + \sigma_\epsilon^2}{\sigma_a^2}w - \frac{\sigma_\epsilon^2}{\sigma_a^2}m_a \\
\Leftrightarrow z &\geq \frac{1}{\sigma_a^2}[(\sigma_a^2 + \sigma_\epsilon^2)w - \sigma_\epsilon^2 m_a].
\end{aligned}$$

Ad Equation (1.3):

$$\begin{aligned}
\frac{dt}{dw} &= \frac{d}{dw}\frac{1}{\sigma_a^2}[(\sigma_a^2 + \sigma_\epsilon^2)w - \sigma_\epsilon^2 m_a] \\
&= \frac{\sigma_a^2 + \sigma_\epsilon^2}{\sigma_a^2} > 0.
\end{aligned}$$

1.7.2 Stable treatment unit value assumption

In addition to the parallel-trends assumption discussed in Section 1.4.1, identifying a causal effect in a DiD setting requires the stable unit treatment value assumption (SUTVA) to hold. SUTVA maintains that units do not interfere with each other. More specifically, SUTVA says that one unit receiving the treatment must not influence the outcome of another establishment (particularly a control group one). Here I provide some indicative correlation evidence for the validity of SUTVA in Section 1.6.1.

In my setting, SUTVA would be violated if, by hiring more selectively, establishments treated by the minimum wage "poached" productive workers that would else have been hired by establishments untreated by the minimum wage. Such a hire stealing effect may reduce minimum hire quality at control group establishments, causing me to overestimate the effect of the minimum wage on employers' worker selection.

To address this threat I identify circumstances where hire stealing effects are more and less likely and test whether heterogeneity in the effect of the minimum wage is consistent with the SUTVA violation. That a treated establishment's recruitment directly affects untreated establishments by depriving them of a good potential hire is more likely in tight labor markets where there are many vacancies for each unemployed worker, causing employers to compete for hires. If the hire stealing effect biased my findings, the effect of the minimum wage on worker selection would be less pronounced in slack labor markets - where one establishment's hiring has little effect on other establishments' choices from the pool of unemployed - than in tight labor markets.

Table 1.7 looks at effect heterogeneity by labor market tightness, considering variation in the vacancy-to-unemployment ratio across German federal states. Column (1) estimates Equation (1.4) for establishments in states where the average pre-reform vacancy-to-unemployment ratio is smaller than the median, i.e., for slacker labor markets. Column (2) does the same for establishments in tighter labor markets. The estimated minimum wage effect is bigger in slack labor markets. This is contrary to the prediction of the SUTVA violation considered. One may interpret this as a piece of suggestive evidence that the hire-stealing effect is unlikely to perceptibly bias my findings. It is important to remember, however, that this comparison is far from ceteris-paribus: German federal states with different vacancy-to-unemployment ratios are likely to vary in many unobserved dimensions that may also influence the effect of the minimum wage on firms' worker selection. Moreover, German federal states do not coincide with local or regional labor markets.

While I cannot present a conclusive empirical test to rule out the hire-stealing effect the latter appears theoretically unlikely. It is not plausible in practice that employers formerly paying below the minimum wage would suddenly, by paying the minimum wage, lure workers who would otherwise have ended up in higher-wage establishments in any systematic way.

1.7.3 Robustness checks and additional results

Table 1.7: Effect on hire quality by state labor market tightness

	(1) Slack labor markets	(2) Tight labor markets
DiD estimate (Treated*Post)=1	0.2484***	0.1332**
	(0.0729)	(0.0663)
Observations	3,624	3,967
Establishments	764	731
Adjusted R^2	0.0110	0.0210

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the establishment level. The dependent variable is minimum hire quality, with CHK's individual AKM worker effects estimated 2002-2009 used as a hire quality measure. (1) reports the estimated effect of the minimum wage introduction on establishments in federal states (*Länder*) with below-median average pre-reform labor market tightness (V/U ratio); (2) does the same for establishments in federal states with a vacancy-unemployment ratio above the median. DiD estimates are from establishment fixed-effects specifications without covariates other than year dummies. The year 2014 is excluded to rule out anticipation effects.

Table 1.8: Minimum wage effect on hire quality: MW bite measures

	(1)	(2)	(3)	(4)
D_treat_post=1	0.1888***		0.0986**	0.1040**
	(0.0496)		(0.0497)	(0.0415)
MW_post		0.2743**		
		(0.1183)		
Observations	7,591	7,591	7,579	7,489
Establishments	1,491	1,491	1,489	1,444
Adjusted R^2	0.0150	0.0134	0.0130	0.0140

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the establishment level. Dependent variable is minimum hire quality, where hire quality is measured by CHK's individual AKM worker effects estimated 2002-2009. In (1) and (2) minimum wage exposure (MW) is computed based on full-time workers' wages. In (3) MW is calculated using additional data on hours worked. In (4) MW is measured only on 1 July 2013 (rather than averaged over the pre-reform period). DiD estimates are from establishment fixed-effects specifications without covariates other than year dummies. The year 2014 is excluded to rule out anticipation effects.

Table 1.9: Minimum wage effect on hire quality: aggregation methods

	(1) Minimum	(2) P10	(3) P25	(4) P50
DiD estimate (Treated*Post)=1	0.1888***	0.1604***	0.0730	0.0595
	(0.0496)	(0.0584)	(0.0557)	(0.0550)
Observations	7,591	7,591	7,591	7,591
Establishments	1,491	1,491	1,491	1,491
Adjusted R^2	0.0150	0.0040	0.0029	0.0018

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the establishment level. Dependent variable varies according to the establishment-level aggregation of individual hire quality: minimum (1), first decile (2), first quartile (3) and median (4). Hire quality is measured by CHK's individual AKM worker effects estimated 2002-2009. (Continuous) DiD estimates are from establishment fixed-effects specifications without covariates other than year dummies. The year 2014 is excluded to rule out anticipation effects.

Figure 1.6: Minimum wage effect on minimum hire quality (log employment weighting)

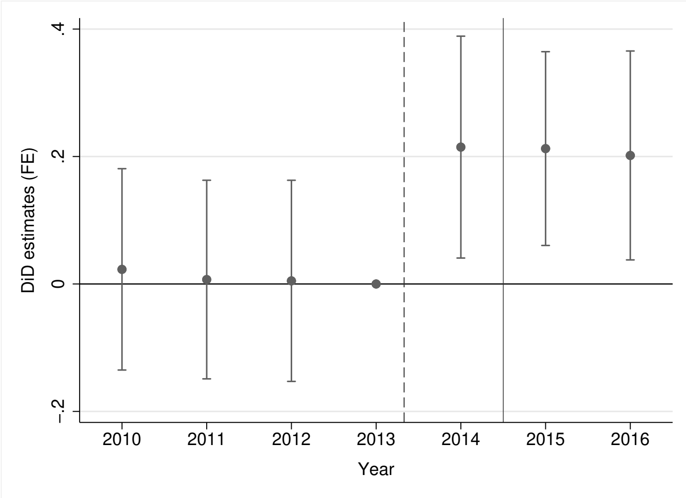

Note: This figure shows yearly DiD estimates for the effect of the minimum wage introduction on minimum hire quality. Estimates are from an establishment fixed-effects specification without covariates other than year dummies and with log(employment) weights. Hire quality is measured by CHK's individual AKM worker effects estimated 2002-2009. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the establishment level.

Figure 1.7: Minimum wage effect on minimum hire quality (preliminary CHK update)

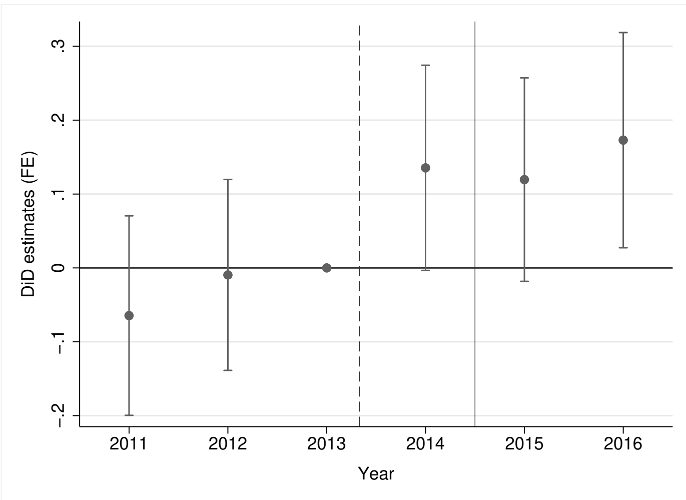

Note: This figure shows yearly DiD estimates for the effect of the minimum wage introduction on minimum hire quality. Estimates are from an establishment fixed-effects specification without covariates other than year dummies. Hire quality is measured by a preliminary update of CHK's individual AKM worker effects. These are estimated 2003-2010 (hence the shortened period of analysis) and include workers from both West and East Germany. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the establishment level.

Figure 1.8: 2012 screening intensity distribution

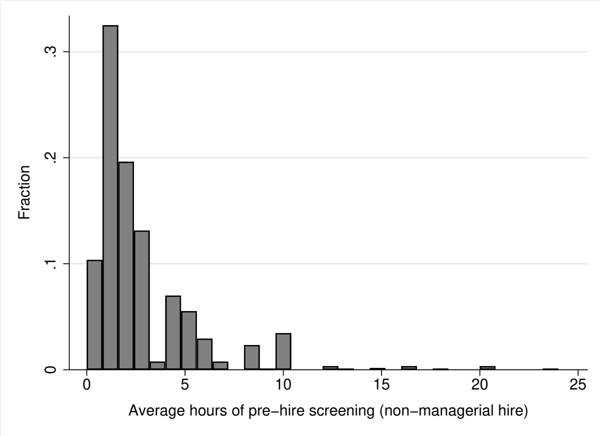

Note: This figure shows a histogram for screening intensity (in hours) for an average successful applicant in 2012. The screening intensity measure is from a survey of HR executives (LPP, 2012).

Figure 1.9: Minimum wage effect by screening feasibility

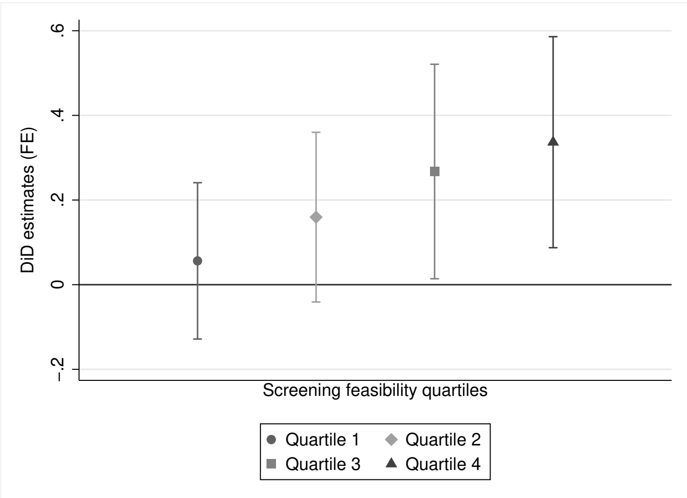

Note: This figure shows DiD estimates for the effect of the minimum wage introduction on minimum hire quality. Estimates are from establishment fixed-effects specifications without covariates other than year dummies. The year 2014 is excluded to rule out anticipation effects. Hire quality is measured by CHK's individual AKM worker effects estimated 2002-2009. Screening feasibility is measured by establishments' self-reported pre-reform screening intensity (LPP 2012). Vertical bars denote 95% confidence intervals. Standard errors are clustered at the establishment level.

Table 1.10: Minimum wage effect on observable hire quality

	(1) Unemployed	(2) Low-skill	(3) Age
Continuous DiD estimate (Bite*Post)	0.0042	0.0019	0.7845
	(0.0197)	(0.0176)	(0.6754)
Observations	8,473	8,473	8,473
Establishments	1,491	1,491	1,491
Adjusted R^2	0.0098	0.0401	0.0113

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the establishment level. Dependent variable is (1) the share among new hires of formerly unemployed workers; (2) the share of new hires that have neither a vocational qualification nor a secondary school diploma that qualifies them for university entry (*Abitur*); (3) new hires' mean age. DiD estimates are from establishment fixed-effects specifications without covariates other than year dummies. The year 2014 is excluded to rule out anticipation effects.

Table 1.11: Relationship between AKM worker effects and observed worker ability

	(1) Wage < SMW	(2) Wage ≥ SMW
	(1)	(2)
Establishment entry from UE dummy	0.0579***	-0.0592***
	(0.0046)	(0.0030)
Less than Abitur or equiv, no voc quali	-0.0649***	-0.2295***
	(0.0048)	(0.0039)
Constant	3.4779***	3.8197***
	(0.0034)	(0.0015)
AKM person effect SD	.2912	.3702
Observations	16025	79081
Adjusted R^2	0.0189	0.0261

Note: *** p<0.01, ** p<0.05, * p<0.1, with heteroskedasticity-robust standard errors. Coefficients are from a cross-sectional OLS regression of AKM worker effects on observed worker characteristics. The table compares two groups of new hires in the pre-reform period (2010-13): Column (1) reports results for the sub-sample of workers with starting wages below the minimum wage and (2) focuses on the workers with starting wages at or above the minimum wage.

Figure 1.10: Minimum wage effect on new hires' AKM coverage

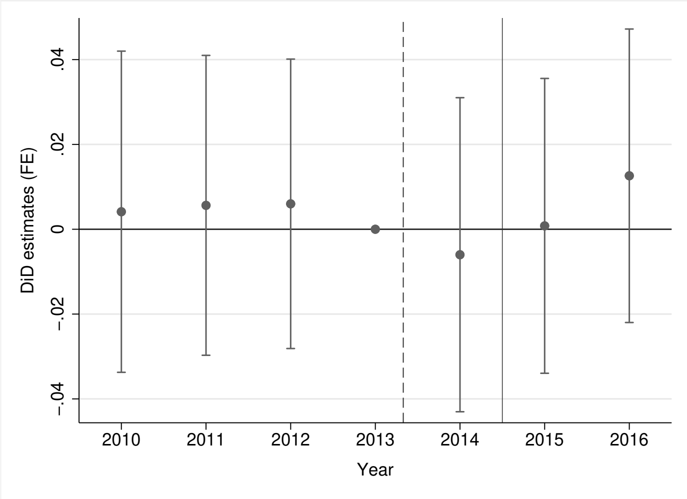

Note: This figure shows yearly DiD estimates for the effect of the minimum wage introduction on the share of new hires for whom an AKM worker effect estimate from CHK is available. Estimates are from an establishment fixed-effects specification without covariates other than year dummies. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the establishment level.

Figure 1.11: Minimum wage effect on women's hire share

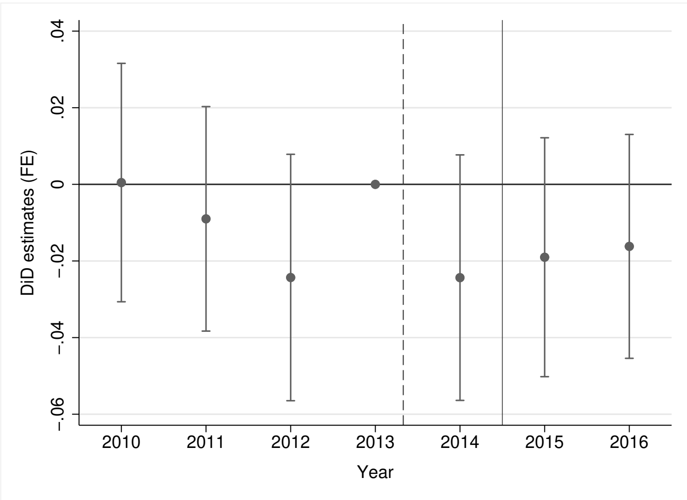

Note: This figure shows yearly DiD estimates for the effect of the minimum wage introduction on the female share of new hires. Estimates are from an establishment fixed-effects specification without covariates other than year dummies. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the establishment level.

Figure 1.12: Minimum wage effect on new hires' average age

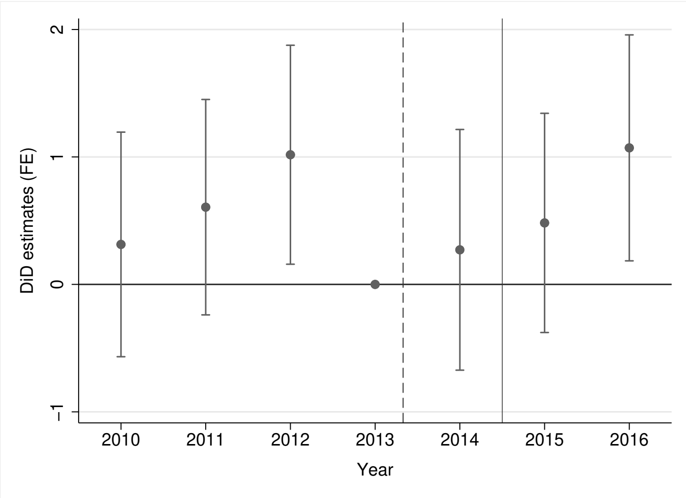

Note: This figure shows yearly DiD estimates for the effect of the minimum wage introduction on new hires' mean age. Estimates are from an establishment fixed-effects specification without covariates other than year dummies. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the establishment level.

Figure 1.13: Minimum wage effect on the number of hires

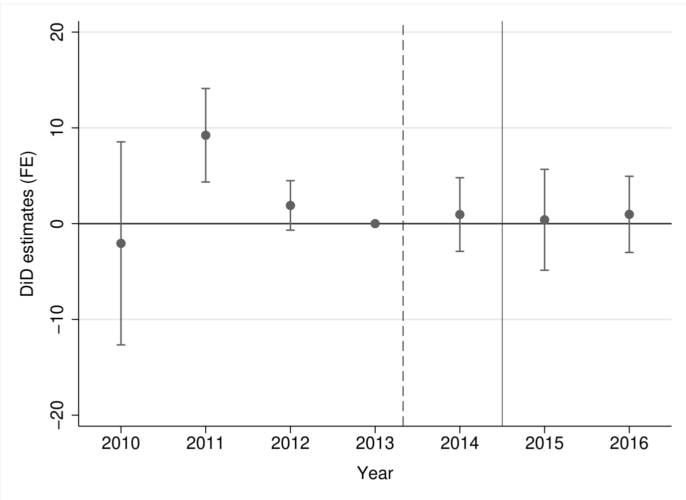

Note: This figure shows yearly DiD estimates for the effect of the minimum wage introduction on the number of hires. Estimates are from an establishment fixed-effects specification without covariates other than year dummies. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the establishment level.

Figure 1.14: Minimum wage effect on maximum hire quality

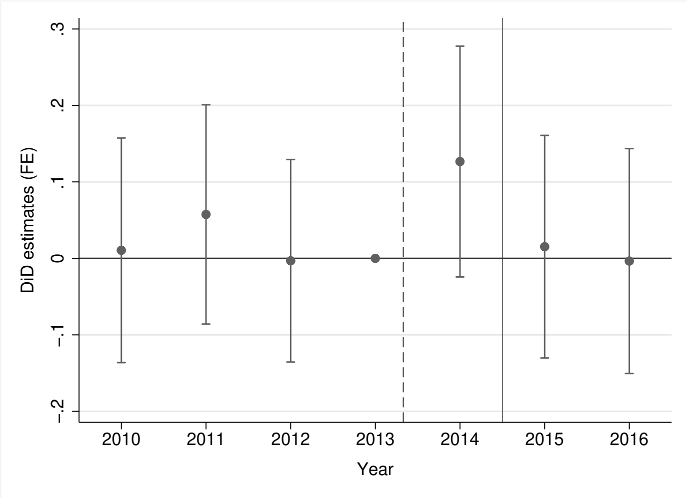

Note: This figure shows yearly DiD estimates for the effect of the minimum wage introduction on maximum hire quality. Estimates are from an establishment fixed-effects specification without covariates other than year dummies. Hire quality is measured by CHK's individual AKM worker effects estimated 2002-2009. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the establishment level.

Figure 1.15: Minimum wage effect on establishment attractiveness

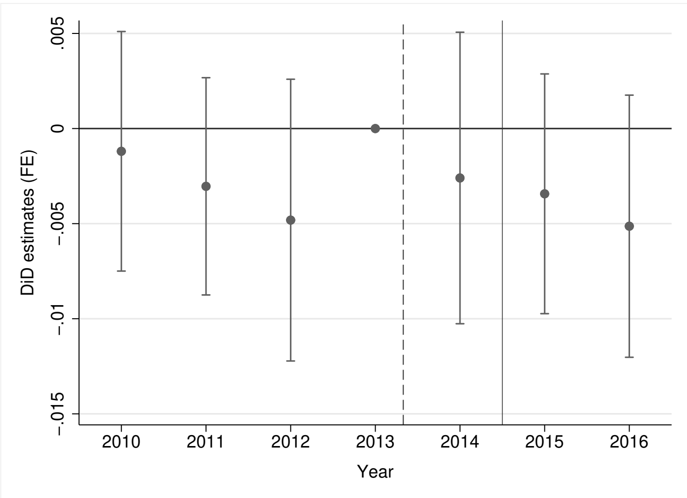

Note: This figure shows yearly DiD estimates for the effect of the minimum wage introduction on the share of hires coming directly from high-wage establishments. Estimates are from an establishment fixed-effects specification without covariates other than year dummies. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the establishment level.

Figure 1.16: Minimum wage effect heterogeneity by screening feasibility

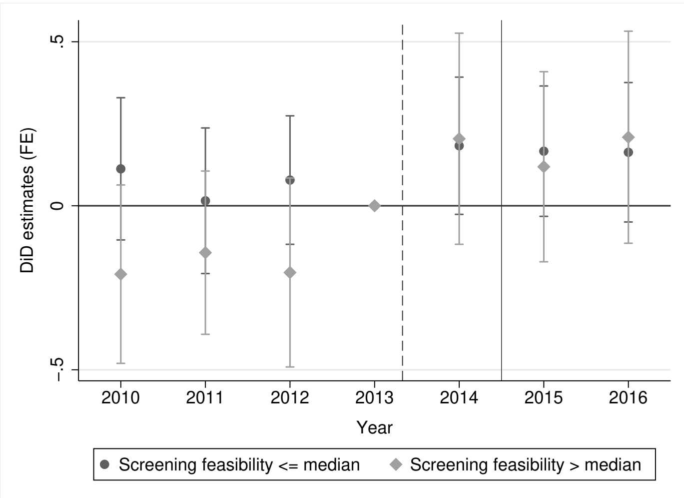

Note: This figure shows yearly DiD estimates for the effect of the minimum wage introduction on minimum hire quality separately for below- and above-median screening feasibility. Estimates are from an establishment fixed-effects specification without covariates other than year dummies. Hire quality is measured by CHK's individual AKM worker effects estimated 2002-2009. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the establishment level.

Chapter 2

The effect of employment protection on firms' worker selection

Authors: Sebastian Butschek and Jan Sauermann

Abstract: To estimate the causal effect of employment protection on firms' worker selection, we study a 2001 policy change in Sweden that reduced dismissal costs for firms with less than ten employees. These small firms employ over a tenth of Sweden's workforce. Our difference-in-differences analysis of firms' hiring uses individual ability measures including estimated AKM worker effects and cognitive test scores from the military draft. Comparing treated small firms to slightly larger, untreated ones we find that the reform reduced minimum hire quality by five per cent of a standard deviation, half of which we can attribute to firms' hiring becoming less selective. Consistent with this, we show that employers' reliance on network-based hiring also goes down as firms react to the reduction of dismissal costs. Our results help discriminate between existing theories, supporting the prediction that firms shift their hiring standards in response to changes in dismissal costs.

Keywords: worker selection, screening, hiring standard, employment protection, dismissal costs.

JEL Codes: M51, D22, J24, J38.

2.1 Introduction

Economists have long studied the effect of labor market regulation on *how many* workers are hired. Far less is known about whether regulation also affects *who* gets a job. Does labor market policy influence firms' worker selection? This question has consequences for the efficiency of labor market outcomes, as firms' selectiveness directly influences the allocation of workers. Moreover, the question has distributional implications: changes in firms' worker selection will alter the employment prospects of the most disadvantaged labor market participants, such as the long-term unemployed. Firms' worker selection may thus be an overlooked pathway from labor market regulation to labor productivity and to the employment dynamics of vulnerable workers.

Several theoretical papers on the hiring process agree that there should be an effect of labor market regulation on the selectiveness of firms' hiring. The existing models suggest different mechanisms for the link from regulation to firms' worker selection though. For example, the model in Lazear (1995) implies that tighter regulation makes hiring more selective as firms hire less "risky" workers, i.e., firms reduce the spread of their hires' ability distribution. By contrast, the matching model in Pries and Rogerson (2005) predicts that stricter regulation increases firm selectiveness by causing employers to no longer hire "bad" workers, i.e., to raise their hiring standard. Do firms reduce the variance of hires' ability at the bottom and the top or do firms primarily avoid the bottom tail of the ability distribution? These are, roughly, the testable predictions of the risky-hires framework formulated by Lazear (1995) and the matching model variant in Pries and Rogerson (2005). Empirically studying the effect of regulation on firms' selectiveness offers the prospect of enhancing our understanding of the hiring process by discriminating between these suggested mechanisms.

In this paper, we investigate *if* and *how* a relaxation of employment protection legislation (EPL) reduces firms' selectivity in hiring. To estimate the effect of weaker EPL on firms' worker selection we exploit a Swedish reform that made it easier for small firms to lay off workers. The policy change generated exogenous variation in dismissal costs across firms and over time, whose effects we analyze in a difference-in-differences (DiD) framework. Administrative data on the universe of Swedish workers and firms allow us to precisely measure firms' affectedness by the reform and to credibly proxy for the ability of new hires. We find that the reduction in dismissal costs directly reduces firms' selectiveness by about 2.5% of a standard deviation. We further conclude that this is driven by an adjustment of hiring thresholds rather than by making riskier hires.

The main contribution of this paper is to show that employment protection legislation directly affects *whom* firms hire, and that this results from firms lowering their hiring threshold. Regarding the first point, we confirm and extend the findings of two other papers that provide evidence consistent with EPL affecting worker selection (Marinescu, 2009; Bjuggren and Skedinger, 2018). In this study we go further by establishing a causal effect on the *ability* of new hires and by providing direct evidence for a change in screening intensity. With respect to the second point, no evidence exists on *how* dismissal costs affect the selectiveness of firms' hiring. This paper is the first to study potential mechanisms underlying the link between EPL and firms' worker selection, enabling us to discriminate between the theoretical predictions of Lazear (1995) and Pries and Rogerson (2005).

Our first step in investigating whether EPL changes worker selection is to estimate the overall effect of Sweden's 2001 EPL reform on firms' minimum hire quality, i.e., the ability of the least productive hire a firm makes in a year. As we discuss below, this overall effect may contain indirect effects other than changes in firms' selectiveness. Our preferred proxy for worker productivity is an estimated worker effect (AKM) from two-way-fixed effect log-wage regressions in the spirit of Abowd et al. (1999). Using a DiD specification that exploits the size cut-off in firms' eligibility for the EPL reduction, we estimate that the reform reduced firms' minimum hire quality by 5% of a standard deviation. Our alternative ability measures are military draft test scores for cognitive and psychological aptitude from age 18 and grade point averages (GPAs) from age 15. We prefer AKMs as the alternatives cover either only men or only young people. Nonetheless, the alternative ability proxies produce remarkably similar effect estimates.

We then isolate the direct effect of Sweden's EPL reform on firms' worker selection. To do so we quantify the role of a mechanism that may lower firms' minimum hire quality but *does not* go through the selectiveness of firms' hiring: as reduced dismissal costs make firms hire more workers, they may—by chance—employ lower-ability workers without actively adjusting their selectiveness.[1] We estimate the importance of this "number-of-hires" channel by simulating a counter-factual scenario designed to preclude changes in firm selectiveness. Our findings suggest the number-of-hires channel accounts for about half of the overall effect of the EPL reform on firms' minimum hire quality. This leads to our estimated direct effect of Sweden's EPL reform on firms' selectiveness of about 2.5% of a standard deviation.

Having established that lower dismissal costs reduce firms' selectiveness we proceed to investigate two mechanisms suggested by the theoretical literature. The first potential channel is based on the insight that firms may like hiring risky workers as long as workers are easy to fire: these hires provide option value as they may turn out more productive than expected (Lazear, 1995). If dismissal costs go down, firms should then intensify their hiring of risky workers. In terms of unobserved ability, the variance of new hires' ability distribution will go up as the firm now hires workers that will either turn out less able than earlier cohorts or prove to be more productive than expected. In our setting, this would imply that lower EPL both reduces minimum hire quality and increases maximum hire quality: under reduced dismissal costs, hires should be better than expected as often as they are worse than expected. Adjusting for the number-of-hires channel, we find no evidence that firms symmetrically increase the riskiness of their hiring. That is, our evidence does not support the mechanism suggested by the risky-hires framework in Lazear (1995).

The second suggested mechanism for the link between EPL and worker selection works through firms' hiring standards. The idea is that there is a productivity threshold below which a firm does not hire applicants. In their matching model with employer search,

[1] Pries and Rogerson (2005) predict a direct effect of EPL on both firms' hiring threshold and on worker turnover. There is also empirical evidence on an effect of EPL on hiring rates both internationally and in Sweden (Kugler and Pica, 2008; von Below and Skogman Thoursie, 2010).

Pries and Rogerson (2005) show that dismissal costs shift this threshold, implying that a reduction in EPL should shift firms' hiring standards downward. For our setting this results in the prediction that the EPL relaxation should lower firms' minimum hire quality without increasing maximum hire quality. Consistent with this prediction, we find that the reform reduced firms' selectiveness only at the bottom of the hire ability distribution. An additional implication of firms lowering their hiring threshold in response to the EPL reform is that employers may relax their pre-hire screening of applicants. To obtain a direct measure of such screening, we build on the finding in Hensvik and Skans (2016) that firms use incumbent workers' networks of past co-workers to hire selectively. Using the number of co-worker links from new hires to incumbents as a proxy, we find that reduced EPL causes firms to rely less on such referral-based hiring. Both pieces of evidence suggest that firms' worker selection responded to the EPL relaxation by becoming less stringent with respect to low ability. This favors the hiring threshold mechanism of Pries and Rogerson (2005) as an explanation of the effect of dismissal costs on firms' worker selection.

Our paper is the first to provide empirical evidence on the mechanisms driving the effect of EPL on firms' worker selection. In showing that there is a link between EPL and firms' selectiveness it extends scarce existing evidence. Three studies present evidence consistent with regulation affecting firms' worker selection. Marinescu (2009) studies the effect of an increase of EPL in the UK on job durations. Her main result is that shortening the qualification period for claiming unfair dismissal decreases the hazard of termination not only for directly affected workers, but also for unaffected workers hired after the policy change. Although her data do not allow her to provide direct evidence on the quality of hires, she shows that the pattern is consistent with EPL inducing more selective hiring. We confirm and complement her suggested explanations by making use of explicit information on worker ability.[2]

In a recent working paper written independently from and contemporaneously to ours, Bjuggren and Skedinger (2018) analyze the effect of the same 2001 EPL reform on the share of workers firms hire from unemployment and from active labor market programs. They find that the EPL reduction increases the share hired from these sources and argue that this is consistent with firms reducing their screening. Their paper and ours arrive at similar conclusions using different pieces of evidence: their argument is that unemployed workers are harder to screen, so that intensified hiring of unemployed workers is an indication that firms are happy to screen less. Our argument goes through hire ability and through firms' network-based hiring, both of which point to a relaxation of screening. The findings

[2]The finding in Kugler and Saint-Paul (2004) that stronger EPL decreases the chances of unemployed individuals is also similar to our conclusion. In their setting with adverse selection, unemployed workers are more negatively selected than employed workers because firms first lay off the lemons. When firing costs increase, firms rely on statistical discrimination to avoid unproductive hires—they shy away from hiring unemployed job-seekers. In our setting, by contrast, firms screen and directly identify less productive workers before hiring them.

of Bjuggren and Skedinger (2018) and ours strengthen each other: they bring different evidence to bear on the question whether EPL affects firms' worker selection and come to consistent and complementary answers.

Also relevant to our study are the findings of von Below and Skogman Thoursie (2010) and Bjuggren (2018), who study the effects of the EPL reduction on labour turnover and firm productivity, respectively. At first glance our finding that the 2001 EPL reform made firms' hiring less selective seems at odds with Bjuggren (2018), who finds positive effects of the same EPL relaxation on labor productivity. Looking at the theoretical motivation of his analysis, however, suggests that there does not need to be a contradiction: in standard models of reduced employment protection, it is not a change in who is hired in the first place that improves the allocation of labor, but increased labor market churning (Hopenhayn and Rogerson, 1993). Indeed, the reform has been shown to have increased both firms' hiring and their firing without a change in net employment (von Below and Skogman Thoursie, 2010). We show that the reform increased the dispersion of hire quality. It is possible that firms increased worker productivity by first hiring less selectively but then identifying and retaining the abler workers. This would be consistent with both higher churn and a larger ability range of workers entering.

By studying the effects of dismissal costs on worker-firm matching, this paper also expands the scope of the literature studying the effects of Sweden's 2001 EPL reform. Labor market outcomes considered by existing papers include sickness absence (Lindbeck et al., 2006; Olsson, 2009) and parental leave take-up (Olsson, 2017).[3]

At a more general level, this study contributes to the literature looking at determinants of selectivity in worker-firm matching. Examples include the analysis of foreign ownership and hiring selectivity (Balsvik and Haller, 2015), the relation between management quality and the quality of worker flows in and out of the establishment (Bender et al., 2018), the role of networks in hiring (Hensvik and Skans, 2016), and skill sorting across firms over time (Håkanson et al., 2015). By analyzing the effect of dismissal costs on worker selection, our paper not only looks at a very different determinant of worker-firm matching than these studies, but it also offers exogenous variation in the determinant of interest, allowing us to identify a causal relationship.

To support the causal interpretation of our results, we present evidence addressing a number of threats to causal identification. These include concurrent changes in labor market legislation, workers' differential self-selection into firms, non-classical measurement error in worker ability, selective attrition and time-invariant firm heterogeneity. Our results are robust to a broad range of tests.

[3] As there is practically no evidence on the effects of EPL on worker-firm matching internationally, our paper generally broadens the literature on the effects of employment protection legislation—see, e.g., Bastgen and Holzner (2017) for a recent review.

2.2 Institutional background

The 2001 reform of Sweden's employment protection provisions created variation in firms' recruitment incentives that we want to exploit. In this section we briefly describe the institutional setting of our quasi-experiment before discussing selected features of Swedish EPL and its 2001 reform that are relevant to our analysis.[4]

2.2.1 Employment protection legislation in Sweden

Swedish EPL is relatively strict: dismissals are only allowed in case of misconduct and redundancy. In the latter case, i.e., when a firm's business is underperforming and it wants to lay off a worker, legislation from 1974 further dictates whom it may fire. A last-in-first-out (LIFO) rule stipulates that the establishment's most recent hire has to be dismissed first. The LIFO rule constrains firms' ability to lay off its least productive workers in a downturn, instead forcing them to dismiss those with the lowest seniority (SFS, 1982).

The 2001 reform of Sweden's EPL did not change the basic idea of protecting higher-tenure workers but allowed small firms to exempt up to two of its workers from the LIFO rule. This meant that out of their three most recent hires, firms with ten or fewer workers became able to choose whom to lay off.[5] The degree to which this relaxed employment protection varies by firm size (see below for details).

Given that small firms (with up to ten workers) make up a sizeable share of the workforce (11.2% in 1999), these changes constitute a substantial liberalization of labor market regulation. Notwithstanding its importance the 2001 EPL reform came about quickly and rather unexpectedly: two possible versions of it were proposed in February 2000 and the second, after some modification, was written into law in October 2000 to become effective on 1 January 2001. That the reform was unexpected is best illustrated by who originated it: an atypical coalition of opposition parties (Liberal-Conservatives and Greens) pressed the unwilling minority government to make a proposal, which the governing Social Democrats sought to derail even after presenting it (Lindbeck et al., 2006).

Most changes to labor market regulation in the 1990s did not coincide with the EPL reform in either timing or targeting. One was a temporary provision allowing *all* firms to exempt two workers from the LIFO tenure ranking, effective only during 1994 (SFS, 1993; Skedinger, 2008). Second, there was a 1997 reform that made it easier to hire workers under fixed-term contracts without a particular reason, though for at most 5 individuals and with individual employment on a temporary contract capped at 12 months in a three-

[4]See Lindbeck et al. (2006), von Below and Skogman Thoursie (2010) and Bjuggren (2018) for other information on the 2001 EPL reform. Skedinger (2008) and Böckerman et al. (2018) contain useful information on EPL legislation in Sweden generally.

[5]One implication is that on-the-job screening may have been made more feasible by the reform: in addition to the probation period affected firms now had time to fire a recent hire as long as there were no more than two more recent hires.

years period (Sjöberg, 2009; Holmlund and Storrie, 2002). Unlike the LIFO reform of 2001 these policy changes applied to firms of all sizes equally. They also were not introduced at the same time as the change in EPL. The only exception is an amendment to the Gender Equality Act, which we discuss in detail in Section 2.6.

2.2.2 Institutional detail of the last-in-first-out rule

The change in EPL applied to small firms of all sectors except private households with employed persons, as the LIFO rules do not apply to these (SFS, 1982). Defining "small", the reform law makes it clear that the threshold of ten workers is in heads (not full-time equivalents) and includes both part-time and full-time workers as well as those on permanent and temporary contracts (von Below and Skogman Thoursie, 2010). The head count, however, excludes owners and their family members, (leading) managers and workers on employment subsidies (SFS, 1982). What the law is less explicit on is the time frame during which LIFO-relevant firm size is measured (more in Section 2.3.3).

The policy change did not alter the fact that firms (both small and large) had to define a tenure ranking of eligible workers in case of dismissals.[6] What did change was the proportion of workers in the tenure ranking who enjoyed LIFO protection: before the reform, a downsizing firm had to lay off the person at the bottom of the relevant tenure ranking. After the reform, the firm was able to choose between the bottom three workers. As a consequence, the bite of the 2001 EPL reform was largest for the smallest firms. For firms with 2 or 3 workers, employment protection was practically abolished. For firms with 4-10 workers, the share of protected workers fell by between 67% and 22%. For larger firms, EPL did not change (cf. Table 2.1).

Finally, there is a degree of fuzziness to the reform. Some firms too large to be eligible will, through other means, have managed to relax or avoid LIFO's strict rules. For example, Böckerman et al. (2018) point out that LIFO-relevant tenure rankings are defined separately for each establishment, blue-collar and white-collar workers and may even be split up again by worker skill.[7] Conversely, a certain fraction of eligible firms will have retained their previous EPL provisions, as employers and trade unions may negotiate modifications to the LIFO rules in local agreements. Nonetheless, the evidence on the LIFO reform's effects suggests that the EPL provisions from 1974 are a binding constraint on firms' decisions (e.g., von Below and Skogman Thoursie, 2010; Bjuggren, 2018).

[6] Both before and after the reform, the protection afforded by the LIFO rule did not cover workers on temporary contracts (even though these workers entered in the head count determining whether the firm was affected by the reform).

[7] To avoid confusion, we reiterate the role of the *firm* and the *establishment* in Swedish EPL: it is firm size that determines whether or not an employer is eligible for the relaxation of EPL effected by the 2001 reform. The establishment is only relevant for determining the LIFO ranking, which must be respected in laying off individual workers. As will become clear below this implies that for us the firm is the relevant unit of analysis.

Table 2.1: Degree of EPL before and after the 2001 reform, by firm size

# workers	Pre-reform		Post-reform		Δ protected
	Protected	Unprotected	Protected	Unprotected	
2	1	1	0	2	-100%
3	2	1	0	3	-100%
4	3	1	1	3	-67%
5	4	1	2	3	-50%
6	5	1	3	3	-40%
7	6	1	4	3	-34%
8	7	1	5	3	-29%
9	8	1	6	3	-25%
10	9	1	7	3	-22%
11+	10+	1	10+	1	0%

Note: Calculated levels and changes in EPL ignore exemptions, such as workers on fixed-term contracts.

2.3 Data and method

2.3.1 Data

To be able to analyze firms' hiring behavior, we need to characterize firms' new hires. For this we would like to capture the timing and nature of all inflows (not just the change in net employment between a given day one year and the same day the next year). To do so we use individual-level data at monthly frequency on the universe of workers from Statistics Sweden (SCB). We extract and infer information about when workers enter new firms and about worker characteristics, most importantly their ability. We also use the firm identifiers in the individual-level data to infer the existence and size of firms at any point in time. We then combine this information to obtain a yearly firm-level panel that tracks all Swedish firms of a certain size over time and characterizes their worker inflows.

Data sources We construct information on worker flows from *Jobbregistret* (JOBB), a register containing the universe of individual employment spells in Sweden from 1986 onwards. The spell data allow us to both identify new hires and measure firm size at monthly resolution.[8] In addition, we use JOBB to obtain information on firms' industry, ownership (private/public), location and age. Finally, JOBB is our main source of information on individuals' labor market history, including where and with whom they worked and how much they earned. We use this information to estimate individual ability (see Section 2.3.5 for detail on our estimation of AKM worker effects) as well as co-worker links (see Section 2.5.3). Drawing on other registers, we obtain further individual characteristics, such as workers' gender, age, highest education and GPA from grade 9. For a majority of male

[8]LISA, a multi-purpose register containing similar information, would only have given us individuals' employment in November of each year.

hires we can also match military draft test scores to our sample. Table 2.9 in the appendix provides more detail on variables and data sources used in this study.

Time period Our period of analysis is 1993-2004. This gives us seven years (1993-1999) before the EPL reform we use was discussed and could have been anticipated (in 2000). We then have four years (2001-2004) during which we can study the short- and medium-run effects of the change in EPL. We do not extend our window of observation further into the past than 1993 as we use the period from 1986-1992 to estimate individual ability (further detail in Section 2.3.5).

Sample restrictions In our main analyses we focus on small firms employing between two and 15 workers in 1999. To ensure we have the relevant firm size measure for all, we restrict our sample to firms that existed in 1999. There are 275,731 firms in 1999 and we keep the 129,187 firms that were between two and 15 in size that year (see Section 2.3.2).[9] We then remove the 1 private household with employed persons, as LIFO rules do not apply to such "firms". Finally, we drop extreme outliers in terms of employment growth, i.e., the 0.5% of firms that experienced the largest inflows of workers.[10]

2.3.2 Identification strategy

We want to estimate the effect of a change in EPL on the selectiveness of firms' hiring. The 2001 reform made EPL less strict for small firms. Because until 1999 this policy change was unexpected, we view it as exogenous variation in EPL over time for small firms (our treated group). The empirical strategy we employ addresses two key requirements for identifying a causal effect: first, we find a control group of similar but unaffected firms to obtain an estimate of the counter-factual hiring behavior of the treated firms in absence of the reform. Second, we recognize that, from 2000 onwards, well-managed firms may have downsized strategically to benefit from the reform. To ensure that such selection is not captured by our effect estimates we assign firms to the treated and control groups before the reform and keep them in that group. For assignment we use the latest possible time pre-reform when we can be reasonably sure that firms could not have anticipated the reform. As Section 2.2 illustrates, such anticipation cannot be ruled out in 2000 but is highly unlikely in 1999. We therefore use 1999 firm size to assign firms to the treated or control groups.

[9] Note that firm size may be below two or above 15 before or after 1999.

[10] We show in Table 2.10 in the appendix that this restriction does not drive our results.

More formally, to estimate the effect of the 2001 relaxation of Sweden's LIFO rule on small firms' hiring standards, we use the following difference-in-differences (DiD) specification:

$$y_{jt} = \alpha + \beta TR_j * POST_t + \gamma_t + \delta TR_j + \epsilon_{jt}, \tag{2.1}$$

where firm j is defined as treated ($TR_j = 1$) if it has 10 or fewer employees in 1999. Firms with 11 to 15 workers are used as the control group. Observations between 2001 and 2004 make up the post-reform period, $POST_t = 1$. Observations from 1993 through 1999 are defined as the pre-reform period. The year 2000, when anticipation behavior is plausible but the reform was not yet in effect, is dropped from this DiD specification. γ_t are year fixed effects. Our main estimates do not include control variables or firm fixed effects, though we show that controlling for time-invariant heterogeneity leaves results qualitatively unchanged.[11] We cluster standard errors at the firm level. The outcome y_{jt} is a measure for firms' selectiveness in hiring. In our main analyses, it is the minimum ability of firm j's hires in year t (more detail in Section 2.3.4).

The DiD estimate of the reform's effect on hiring standards is given by $\hat{\beta}$. As we use a pre-determined and time-invariant measure of firm size, our DiD specification yields an intention-to-treat (ITT) estimate of β. This has the benefit of immunity to endogenous selection of firms into treatment: firms remain classified as control (treated) if they employed up to (more than) ten workers in 1999, even if they strategically downsize (naively expand). A potential drawback of the ITT approach is that it may bias our estimates toward zero to the extent that firms non-selectively shrink into or grow out of the treatment group.

Our key identifying assumption is $\mathbb{E}\left[\epsilon_{jt} | TR_j * POST_t\right] = 0$: absent the EPL reform, treated firms' worker selection would have followed a trend parallel to that of control firms'. While it is impossible to test this assumption for the reform years, we can check the extent to which it holds in the pre-reform period. To do so, we consider an event-study framework of the following form:

$$y_{jt} = \alpha + \sum_t (\beta_t TR_j * \gamma_t) + \gamma_t + \delta TR_j + \epsilon_{jt}, \tag{2.2}$$

where the DiD coefficient β_t is year-specific and 1999 is used as the reference year. Parallel trends in the outcome variable during the pre-reform period would imply estimates of β_t that are indistinguishable from zero for the years 1993-1998. Non-zero $\hat{\beta}_{t=2000}$ would constitute evidence for anticipation effects. The yearly effect estimates for the years 2001-2004 make it possible to learn whether the reform's effects, if any, are sustained in the medium run or die out over time.

[11]Not including fixed effects allows us to look at effect heterogeneity by firm size, which is time invariant in our empirical approach.

There are potential threats to our identifying assumption that are not addressed by comparing pre-reform trends. Our identifying assumption would be violated if trends in outcomes, while parallel in the pre-reform period, would have deviated from their parallel paths in the post-reform period even in the absence of the EPL reform. We return to this issue in Section 2.6.7.

2.3.3 EPL reform bite measure

We want to compare firms potentially affected by the EPL reform to unaffected firms. This makes 1999 firm size a central measure. A key issue in defining firm size is whom to include in the headcount that determines whether the firm is small enough to benefit from the LIFO relaxation. As described in Section 2.2, owners/managers are excluded from this headcount. We follow Bjuggren (2018) in arguing that each firm will have at least one owner/manager and subtract one from the raw headcount.[12]

There is another important issue in defining firm size: is LIFO-relevant firm size the average number of employees in a year or the snapshot headcount at the time the firm wants to lay off a worker? Here the law makes provides no explicit information. We argue that this ambiguity is relatively unproblematic because we are most interested in how the *expected* ease of firing affects firms' hiring. When forming expectations about what rules apply to them in present and future, firms, like us, will need to rely on some approximation of their headcount. We choose the average number of employees in a year as a simple proxy that also takes account of seasonal fluctuations. That is, we use the employment spell data to determine firms' monthly number of individuals in dependent employment and take the average for each year. We then round this average to the nearest integer.

2.3.4 Measuring the selectiveness of firms' hiring

As a proxy for the selectiveness of firms' hiring we use their minimum hire quality in a given year, $\min_{j,t}\{ability_i\}$. This allows us to capture changes in firms' worker selection driven by both the risky-hiring mechanism (Lazear, 1995), which predicts a change in hire quality spread, and the hiring-threshold channel (Pries and Rogerson, 2005), which hypothesizes a shifting of hiring standards at the bottom. Measuring firms' time-varying minimum hire quality requires individual-level ability proxies.

[12] In principle, the data allow us to account for most qualifications detailed in Section 2.2: they identify owners/entrepreneurs and we can use family linkages to find out whether their parents, siblings or children also work at the firm. However, because in 85.4% of all firm-level observations for 1999, none of the individuals linked to a firm is categorized as either owner or manager, we ignore this partial information. Von Below and Skogman Thoursie (2010) choose a different approach, excluding owners/managers from the count where such information is available and reducing the head count by one where no manager/owner appears in the employment spell data under that firm ID. We prefer the Bjuggren (2018) approach because it requires fewer assumptions. Table 2.11 in the appendix shows that our main results are robust to the firm size definition suggested by von Below and Skogman Thoursie (2010).

Worker ability measures Thanks to the wealth of Swedish register data there are three main options for measuring worker ability: first, average school grades at age 15 (GPA); second, military draft test scores at age 18 for cognitive ability (COG) and psychological aptitude (NON-COG) (Lindqvist and Vestman, 2011); and third, estimated worker fixed effects (AKM) from a two-way fixed-effects log wage regression with person and firm effects (Abowd et al., 1999). GPA is dominated by the other two measures: it is available only for some of the young birth cohorts in our data (born between 1973 and 1982).[13] While draft test scores are arguably more precise measures of individual ability than AKM estimates, the former were collected only for men. AKM, on the other hand, will capture individual ability only to the extent that it is reflected in people's gross monthly earnings. Some of Sweden's elaborate wage-setting institutions will no doubt cause individual wages to depart from a worker's marginal product of labor. Nonetheless, AKMs will contain information about those aspects of time-constant individual productivity that are remunerated in the labor market. The key advantage of AKMs is that they can be obtained for both men and women. For this reason we use AKMs as our main ability measure and rely on draft test scores and GPA to explore the robustness of our results.[14]

2.3.5 Estimating AKM worker effects

To estimate time-invariant individual productivity we use individual-level spell data on employment from JOBB. We obtain full time-equivalent (FTE) monthly wages from the Wage Survey Statistics (WSS). For the AKM estimation we use data on the period from 1986 (the beginning of records) through 1992 (the year before the beginning of our analysis period). We want our ability estimates to be pre-determined from the point of view of our analysis and so avoid an overlap with our analysis period.

In our AKM estimation we include individuals aged 18-65 with a November spell for whom we have some FTE wage data.[15] We deflate wages using the CPI and winsorize at 0.5% and 99.5% of the annual real monthly FTE wage distribution. Information on years of schooling/education is available only from 1990; for earlier years, we impute it from the first year it is available.[16]

[13]AKM estimates are available for birth cohorts 1922-1976 and draft test scores males born 1951 -1991.

[14]Butschek and Sauermann (2019b) compare different individual ability measures and find a correlation between estimated AKM worker effects and cognitive test scores of around 29% and GPAs at age 15 and cognitive test scores of around 63%.

[15]We have a measure of FTE monthly wages for 45.4% of person-year observations, or at least one monthly FTE wage observation for 63.7% of individuals. Coverage is partial because WSS is a survey-based register that only covers a stratified random sample of smaller firms.

[16]We do not have reliable information on whether individuals are primarily studying. We therefore omit this information.

Following Abowd et al. (1999) and Card et al. (2013) we estimate a two-way fixed-effects regression:

$$\ln(w_{ijt}) = \alpha_i + \psi_j + \gamma_t + x'_{it}\beta + r_{ijt}, \tag{2.3}$$

where $\ln(w_{ijt})$ is the natural logarithm of individual i's hourly wage at firm j in year t. Moreover, there are additive fixed effects for individuals (α_i) and firms (ψ_j) as well as a set of year dummies (γ_t) and a vector of time-varying individual-level controls (x_{it}). Controls include age squared and age cubed as well as education categories interacted with the year dummies, age squared and age cubed.[17] We deviate from Card et al. (2013) by estimating the two-way fixed-effects regression for men and women together so $\hat{\alpha}_i$ is comparable across gender. We obtain individual fixed-effect estimates $\hat{\alpha}_i$ for 78.7% of the workers for whom we observe a monthly FTE wage and estimates of ψ_j for the 46.6% of firms that have multiple workers and are connected by worker mobility.[18]

2.3.6 Summary statistics

Table 2.2 provides summary statistics for our main estimation sample, separately for treated and control firms. Panel A summarizes continuous variables. The average treatment group firm employs 4.9 workers, compared to 12.7 workers for the average control group firm. Treated firms have existed for a bit shorter, employ older workers and slightly more women and they pay a little less. Means are virtually identical across groups for workers' years of schooling and estimated AKM person and firm fixed effects. Panel B shows means for dummy variables. These are very similar across groups, with two exceptions: treated firms are less likely to be manufacturing firms and are less likely to have expanded since the previous year.

2.4 Results

2.4.1 EPL effect on firms' minimum hire quality

We first consider the results from the event-study estimation framework described in Equation 2.2. Recall that the DiD estimates for 2001-2004 show the reform's (dynamic) effect on minimum hire quality $\min_{j,t}\{\hat{\phi}_i\}$, where $\hat{\phi}_i$ are estimated AKM worker effects. DiD estimates for 1993-1998 provide a test of the parallel-trends assumption for the pre-reform period. Figure 2.1 displays these DiD estimates over time. With the exception of 1996 pre-reform DiD estimates are statistically indistinguishable from zero. While this is a deviation from parallel pre-reform trends in outcomes we argue that overall the assumption

[17] Age squared and age cubed are included to avoid capturing the effect of experience on wages in $\hat{\alpha}_i$.

[18] To estimate Equation 2.3, we use `reghdfe` (Correia, 2016), which drops singletons (22.3% of individuals and 38.9% of firms).

Table 2.2: Firm characteristics by treatment status

A: Continuous characteristics	Treated		Control	
	mean	sd	mean	sd
Head count (rounded)	4.857	2.453	12.688	1.388
Firm age (years)	8.171	4.548	8.859	4.407
Mean worker age (years)	40.158	9.311	38.644	8.532
Female worker share	0.409	0.297	0.373	0.267
Mean worker years of schooling	11.418	1.450	11.356	1.222
Mean worker monthly wage (100 SEK 1980)	45.080	29.996	45.506	29.199
Mean worker AKM person effect (std, 1986-92)	-0.007	0.998	-0.022	0.817
Estimated AKM firm effect (std, 1986-92)	-0.007	0.992	0.038	0.995
Observations	87,441		13,947	

B: Binary characteristics	Treated		Control	
	Share Yes	Frequency	Share Yes	Frequency
Privately owned	0.991	86,651	0.986	13,750
Manufacturing firm	0.090	7,881	0.132	1,845
Expanding firm	0.430	35,124	0.581	7,786
Downsizing firm	0.286	23,316	0.287	3,847
Stockholm county	0.227	19,845	0.230	3,212
Malmo (Skane county)	0.123	10,745	0.115	1,606
Gothemburg (Vastra county)	0.163	14,261	0.175	2,435
Observations	87,441		13,947	

Note: This table summarizes treatment and control firms' characteristics in 1999 for the main estimation sample. Panel A provides mean and standard deviation of continuous variables. Panel B gives means and frequencies for dummy variables.

of parallel counter-factual trends is likely to hold. We present specific evidence supporting this claim in Section 2.6.7, where we perform placebo tests. For the years 2001-2004 Figure 2.1 shows a negative and persistent effect of the EPL reform on minimum hire quality. The event-study framework suggests that there was strong anticipation in 2000, the year that the LIFO reform was debated and passed in parliament. While the size of the anticipation effect may be somewhat surprising it is intuitive that, if they respond, firms may adjust their worker selection once they anticipate that firing workers becomes easier in the near future.[19]

We next compare the EPL reform's estimated effect on minimum hire quality across our four different worker ability measures. Table 2.3 presents estimated DiD coefficients for the specification in Equation 2.1. The difference between the columns is the worker ability proxy underlying firms' yearly minimum hire quality: estimated AKM worker effects (1), military draft cognitive test scores (2), draft psychological test scores (3) and GPA at age 15 (4). The outcome variables are standardized to make them comparable across scores. The DiD results are not only qualitatively consistent across measures but also

[19]This is especially true given that the reform facilitates firing of recent hires only once they are no longer the most recent hire.

Figure 2.1: Dynamic EPL effect on minimum hire quality

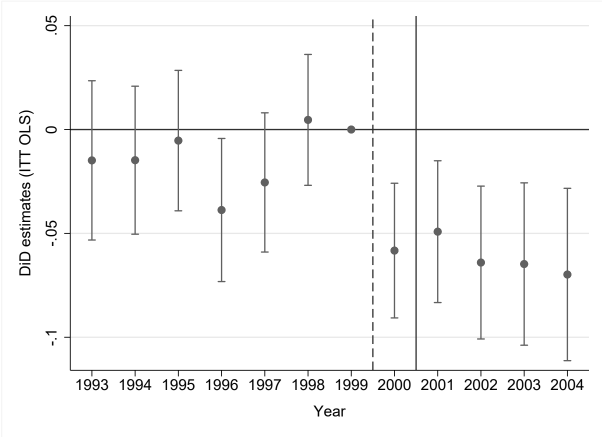

Note: This figure shows yearly DiD estimates for the effect of the EPL reform on minimum hire quality from an OLS (ITT) specification without covariates other than year dummies. Hire quality is measured by AKM worker effects estimated 1986-1992. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the firm level.

quantitatively comparable. This impression is reinforced by looking at the event-study graphs for the alternative ability measures: Figures 2.5-2.7 in the appendix reproduce the event-study DiD graphs for cognitive ability, psychological aptitude and GPA. They show similar dynamics as Figure 2.1. A somewhat more pronounced effect on minimum hire quality as measured by GPAs (available for younger cohorts) may hint at a relaxation of firms' screening: it is plausible that relatively young hires, about whom firms may know little but their school grades, may benefit disproportionately from laxer hiring standards.

Table 2.3: EPL effect on different minimum hire quality measures

	(1) AKM	(2) COG	(3) NON-COG	(4) GPA
DiD estimate (Treated*Post)=1	-0.0484***	-0.0603***	-0.0488***	-0.0724***
	(0.0096)	(0.0126)	(0.0124)	(0.0150)
Observations	314,144	193,025	189,923	132,281
Firms	101,388	79,923	79,210	63,909
Adjusted R^2	0.0027	0.0062	0.0058	0.0071

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the firm level. Dependent variable is one of three minimum hire quality measures, based on: AKM worker effects estimated 1986-1992 (1), military draft cognitive test scores (2), military draft psychological test scores (3), GPA at age 15 (4). Estimates are from OLS (ITT) specifications without covariates other than year dummies and are all based on the sample of firms from (1). The year 2000 is excluded to rule out anticipation effects.

2.4.2 Effect heterogeneity by reform bite

The EPL reform did not affect all treated firms equally. As we discuss in some detail in Section 2.2, firms with two and three relevant workers experienced a near abolition of employment protection, with the degree of EPL relaxation diminishing monotonically up to firms with ten workers. The light gray diamonds in Figure 2.2 illustrate this variation in reform bite. The figure also shows results from a DiD specification exploiting this heterogeneity. It splits up the overall DiD term into separate DiD terms for each firm size group (and uses firms with 14 workers as the reference group). The results, shown in dark grey dots, are noisy, particularly further up the firm size distribution where there are fewer observations. Still, the results in Figure 2.2 strengthen the case that firms' minimum hire quality responded to the reform. The absolute magnitude of the DiD estimates is greatest for the smallest firms in our sample and near-monotonically falls with firm size.

Figure 2.2: EPL effect heterogeneity by firm size

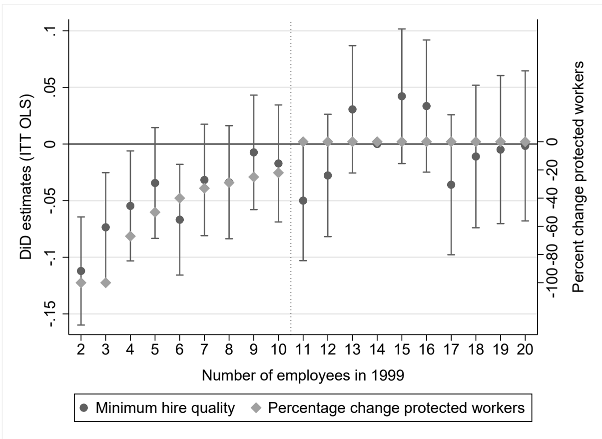

Note: The dark gray dots show firm size-specific DiD estimates for the effect of the EPL reform on minimum hire quality from an OLS (ITT) specification without covariates other than year dummies. The year 2000 is excluded to rule out anticipation effects. Hire quality is measured by AKM worker effects estimated 1986-1992. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the firm level. The light gray diamonds illustrate the bite of the reform in terms of the percentage change of workers protected by the seniority rule.

2.5 Mechanisms

2.5.1 Number of hires channel

Von Below and Skogman Thoursie (2010) find that the 2001 EPL reform increased worker flows both at the hiring and the firing margin. Is the drop in minimum hire quality at treated firms driven by a reform-induced increase in the hiring rate, mechanically lowering minimum hire quality? To test this, we simulate a scenario where firms hire the same type of workers before and after the reform, not changing their selectiveness. We implement this by randomly re-shuffling new hires over time within firm size groups. For the re-shuffle we use the sub-period from 1996-2004 and exclude the year 2000.[20] This ensures that the pre-and post-reform periods are of the same length so that there is approximately the same mass of pre- and post-reform hires in the simulation data.

In our simulation the type of hires is the same before and after the reform by design. If in reality, too, firms' selectiveness had been unchanged by the reform, the counter-factual exercise should leave the reform's estimated effect on minimum hire quality the same as the effect based on actual data. If, on the other hand, the estimated effect is attenuated by the simulation, this gives us an indication that firms hired systematically different workers after the reform. Figure 2.3 shows that in the counter-factual scenario, the reform's effect is less pronounced. Reassuringly, pre-reform coefficients are nearly identical but simulated reform effects are attenuated. Point estimates for the actual and counter-factual reform effects on minimum hire quality are given in Column (1) and (2) of Table 2.4. The estimated actual effect is nearly twice as large in absolute value as the counter-factual one. This suggests that a substantial part of the overall effect is driven not by firms hiring more workers but by firms otherwise changing their hiring behavior. We argue below that the likeliest candidate for a shift in firm hiring behavior is a change in firm selectiveness.

2.5.2 Risky hires channel

Theoretically, changes in firm selectiveness may come in at least two guises. One is that firms are more willing to hire risky workers. That is, firms could refrain from screening out candidates whose ability is so difficult to judge a priori that it may be either so low the firm would not want to employ the worker or so high the worker would really boost productivity at the firm. Before the reform, Sweden's six-month probation period may not have been long enough to give them confidence they would learn the new hire's productivity in time to fire the worker without extra hiring costs. After the reform, however, treated firms may want to use their extra firing flexibility to harness the option value of risky workers explored by Lazear (1995). That is, these firms may want to hire workers whose ability is difficult to tell precisely because they can try to retain the workers that turn out to be unusually productive and lay off those that are revealed to be unproductive. One piece of evidence in favor of this mechanism is that the reform not only lowered minimum but also increased maximum hire quality at affected firms - see column (3) of Table 2.4. As above, though, this result may come about without a change in selectiveness as firms hire more workers, thus drawing larger samples from the ability distribution. We therefore

[20] Potential anticipation effects make it difficult to unambiguously assign the year 2000 to before or after the reform.

Figure 2.3: EPL effect and random hiring simulation

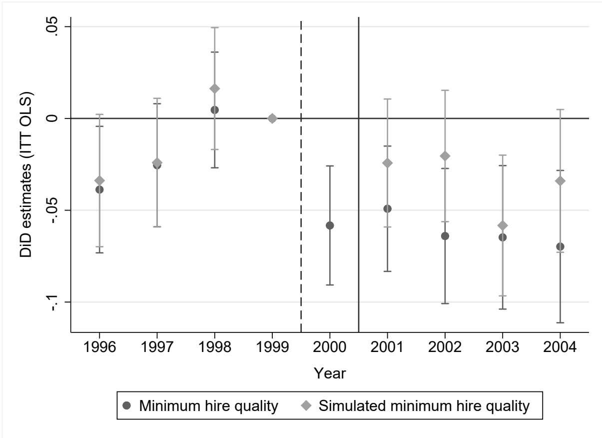

Note: The light gray diamonds are yearly DiD estimates for a counter-factual effect of the EPL reform on minimum hire quality from an OLS (ITT) specification. The data have been modified, leaving each firm's actual number of hires unchanged but randomly re-shuffling hired individuals across time within a firm size group. (The year 2000, when anticipation is possible, is omitted from the re-shuffle). For comparison, the DiD estimates for actual minimum hire quality are included as dark gray dots. Estimates are from OLS (ITT) specifications without covariates other than year dummies and are all based on the sample of firms from (1). Hire quality is measured by AKM worker effects estimated 1986-1992. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the firm level.

use the same simulated data as above and estimate the reform's counter-factual effect on maximum hire quality in a scenario where, by design, firms cannot have changed their selectiveness. Column (4) of Table 2.4 shows that unlike the drop in minimum hire quality, the increase in maximum hire quality is entirely explained by the number of hires effect in our simulation. The evidence therefore does not support the interpretation that firms' intensified hiring of risky workers is the mechanism behind firms' reduced selectiveness.

2.5.3 Lowered hiring threshold channel

An alternative mechanism, suggested by Pries and Rogerson (2005), is that loosened job protection makes firms more willing to give a chance to workers who at first do not appear very productive. The reform-induced drop in minimum hire quality may thus be driven by firms lowering their hiring standard. For more direct evidence of a lowering of hiring thresholds we look to firms' network-based hiring. Hensvik and Skans (2016) provide evidence that firms use their employees' past co-worker networks for screening purposes. We build on their result and use new hires' past links with incumbents as a screening proxy. The intuition is the following: if affected firms after the reform are happy to hire workers of potentially lower ability, it is less important for them to obtain reliable signals of new hires' productivity, e.g., through employee referrals. If so, we would expect to see

Table 2.4: EPL effect mechanisms

	(1)	(2)	(3)	(4)
	Minimum		Maximum	
	Actual	Simulated	Actual	Simulated
DiD estimate (Treated*Post)=1	-0.0478***	-0.0253**	0.0309***	0.0390***
	(0.0102)	(0.0099)	(0.0104)	(0.0104)
Observations	232,707	244,459	232,707	244,459
Firms	92,164	95,473	92,164	95,473
Adjusted R^2	0.0028	0.0017	0.0026	0.0023

Note: *** $p<0.01$, ** $p<0.05$, * $p<0.1$, with standard errors clustered at the firm level. Dependent variable: actual minimum hire quality (1), counter-factual minimum hire quality (2), actual maximum hire quality (3), counter-factual maximum hire quality (4). For the counter-factual effect estimate, the data have been modified, leaving each firm's actual number of hires unchanged but randomly re-shuffling hired individuals across time within a firm size group. Hire quality is measured by AKM worker effects estimated 1986-1992. All estimates are from OLS (ITT) specifications without covariates other than year dummies. The year 2000 is excluded to rule out anticipation effects.

the intensity of network-based hiring decrease at firms affected by the reform.[21] As Figure 2.4 shows, this is indeed the case. The dynamics of the effect are very similar to those of the effect on minimum hire quality.[22] We interpret this as evidence that relaxed screening in the form of a lowered hiring threshold is a key driver of the reform-induced drop in minimum hire quality.

There are two caveats to this result. First, unlike our other results the estimates in Figure 2.4 are from a specification adding firm fixed effects to Equation 2.2. Without controlling for time-invariant firm heterogeneity in this way, pre-trends are not parallel, as Figure 2.8 in the appendix shows. Second, including firm fixed effects results in parallel pre-reform trends only back to 1994, as is clear from Figure 2.4. Barring a data issue we are unaware of this suggests that some specific event differentially affected the network-based hiring of treatment and control firms in 1993. Note, however, that trends are parallel again as we go back further in time, as Figure 2.9 in the appendix illustrates.

2.6 Threats to identification

2.6.1 The Gender Equality Act

As pointed out in Section 2.2 there was an amendment to the gender equality act (GEA) that roughly coincided with the 2001 LIFO reform. GEA, in force from 1992, called for firms of ten employees or more to annually publish two documents specifying measures to

[21] We follow Hensvik and Skans (2016) in our definition of new hires' co-worker links by counting the number of incumbents who previously worked at the same firm and establishment as the new hire. To ensure this measure is not mechanically related to firm size we divide new hires' average number of co-worker links to incumbents by the number of incumbents, giving us the share of incumbents the new hire is linked to.

[22] Column (1) of Table 2.12 in the appendix shows that the effect size is also comparable to FE estimates of the reform's effect on minimum hire quality given in Table 2.7.

Figure 2.4: EPL effect on network-based hiring

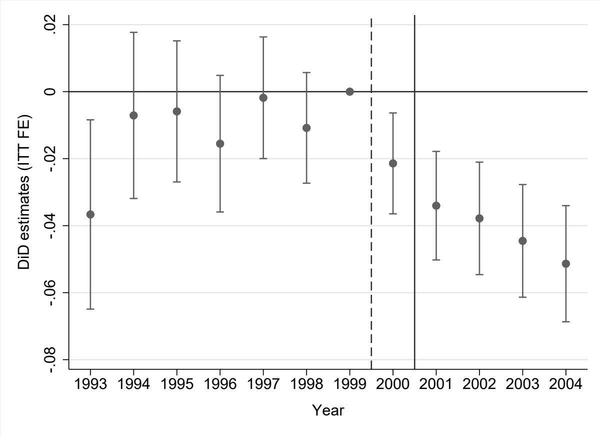

Note: This figure shows yearly DiD estimates for the effect of the EPL reform on network-based hiring from a firm fixed effects specification without covariates other than year dummies. Network-based hiring is measured by new hires' average share of incumbents with whom they previously worked at another firm. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the firm level.

promote gender equality.[23] Effective in 2001, the wording of these requirements was made more concrete. Importantly, however, it did not fundamentally change their content (SFS, 1991, 2000).[24] In light of the modest legal changes it is unlikely that the GEA's effect on firm's hiring abruptly increased in 2001. This view is reinforced by the fact that compliance with the GEA's provisions was neither monitored nor enforced through sanctions. Still, the 2001 re-phrasing of the legal text may have reflected growing awareness of gender inequality and rising pressure to do something about it. In this case firms with 10 or more workers (mostly contained in our control group) may have felt more obliged to promote gender equality in their hiring choices. It appears likely that they would have responded by marginally favoring women in hiring, thus relaxing their hiring standard. In difference-in-difference terms, this is equivalent to a tightening of the hiring standard of the unaffected small firms (our treated group). As a consequence, if the 2001 strengthening of GEA had a perceptible effect on larger firms' hiring standards for women, it should bias our estimated EPL reform effects towards zero. To provide suggestive evidence whether this is the case we separately estimate the effect of the EPL reform on minimum female and male hire quality. Columns (1) and (2) of Table 2.5 are consistent with such an interpretation: the point estimate for women is less pronounced than that for men.

[23]The mandate for an annual gender equality action plan was introduced in 1992 and called for, e.g., measures to promote women in male-dominated positions through training and recruitment. The mandate to publish an equal pay action plan was added in 1994 and asked firms to report on gender differences in pay and ways of removing them.

[24]The requirement for an annual gender action plan was largely re-stated and the equal pay action plan mandate was re-worded to specify in more concrete terms what comparing women's and men's pay should entail (SFS, 1991, 2000).

Table 2.5: EPL effect on minimum hire quality by gender and on firm attractiveness

	(1) Women	(2) Men	(3) E-E hires
DiD estimate (Treated*Post)=1	-0.0353***	-0.0522***	0.0109
	(0.0128)	(0.0122)	(0.0084)
Observations	180,194	192,387	314,144
Firms	76,453	78,853	101,388
Adjusted R^2	0.0012	0.0019	0.0022

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the firm level. Dependent variable: female hires' minimum AKM worker effect (1), male hires' minimum AKM worker effect (2), share of hires transitioning directly from employment at a high-wage firm (above-median estimated AKM firm effect) (3). Estimates are from OLS (ITT) specifications without covariates other than year dummies and are based on the sample of firms from Column (1) in Table 2.3. The year 2000 is excluded to rule out anticipation effects.

2.6.2 Demand vs. supply

We have been using an equilibrium outcome (observed hires) to infer something about firms' (qualitative) labor demand. Implicit in this analysis is the assumption that labor supply remained constant, i.e., that workers' self-selection into firms was not systematically affected by the LIFO reform. To bias our results away from zero, the reform would have had to make affected firms less attractive employers. Theoretically the reform's value for new hires is ambiguous: in the short run it made treated firms more attractive as the lowest-tenure worker need no longer be laid off first. In the medium run, however, the reform reduced the value of seniority capital as second- and third-lowest tenure no longer afforded protection from layoffs. Beyond this, attractiveness was unchanged. A conclusive empirical answer to this question would require data on applicants, which we lack. We can, however, look at job changes that were likely voluntary and therefore contain information about firms' attractiveness as employers. As a proxy for voluntary job changes we use the share of new hires who worked at a high-wage firm right up to starting their new job.[25] As Column (3) of Table 2.5 shows, the LIFO reform has no significant effect on voluntary employment-to-employment transitions. We view this as suggestive evidence that the reform did not make treated firms less attractive to workers and that our assumption of unchanged (qualitative) labor supply is valid.

2.6.3 Non-classical measurement error in worker ability

It is likely that our estimated AKM worker effects measure worker ability with substantial error. Those instances of such measurement error are unproblematic that affect our control and treated firms in the same way. As we estimate individual ability before our analysis period and in other firms, this is plausible in most cases. Take an example: to the extent that there is any wage discrimination in the labor market, women's wages (and thus AKM)

[25] We define as high-wage firms those with a firm fixed effect estimate (from our AKM specification) above the median, irrespective of size.

will systematically understate their ability. However, thanks to our DiD-setting, this will not be a problem unless we underestimate to a particularly large extent the ability of women hired into treated firms from 2001 on, which is unlikely. However, a related problem would be if the reform increased the share of women hired: in this case, the share of underestimated ability would have increased and that might bias minimum hire quality downward. We test this possibility in Column (1) of Table 2.6. There is no detectable effect of the 2001 reform on the female hire share, allaying this concern.[26]

Table 2.6: EPL effect on female hire share and AKM coverage

	(1) Share Women	(2) Share AKM
DiD estimate (Treated*Post)=1	-0.0040	-0.0196**
	(0.0085)	(0.0089)
Observations	314,144	314,144
Firms	101,388	101,388
Adjusted R^2	0.0014	0.0292

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the firm level. Dependent variable: female hire share (1), share of hires for whom an estimated AKM worker effect is observed (2). Estimates are from OLS (ITT) specifications without covariates other than year dummies and are based on the sample of firms from Column (1) in Table 2.3. The year 2000 is excluded to rule out anticipation effects.

2.6.4 Selective worker ability measurement

We do not observe worker ability for all new hires. In 1999, at the average firm in our estimation sample we have an estimated AKM worker effect for 55.4% of new hires. Our DiD setting and ability measurement ensure that this is only problematic if the share of new hires whose ability we can measure is positively affected by the reform. Similarly to the number of hires effect described in Section 2.5.1, this may mechanically drive down minimum hire quality even if firms' selectiveness remains unchanged. To test for this in our sample, Column (2) of Table 2.6 reports results from using firms' share of new hires with an AKM worker effect estimate as an outcome. It shows we do not observe the ability of relatively more hires at treated firms. In fact, the point estimate is *negative* and significant, implying that, if anything, selective ability measurement will bias our results towards zero.

[26] An analogous source of bias would be that (i) the EPL reform made firms hire younger workers and (ii) AKM underestimates young workers' ability. Table 2.13 in the appendix suggests this is not the case, as mean age of new hires increased in response to the reform.

2.6.5 Time-invariant firm heterogeneity

Consider a scenario in which the composition of small firms deteriorates over time, so that small firms observed after the reform are negatively selected as compared to those observed before. Suppose that the composition of large firms remains unchanged. A credit boom with lots of new entrants, for example, may produce such a pattern. This scenario could theoretically deliver similar results as our study in the absence of any EPL-induced change of selectiveness in hiring. The danger of such a confounding scenario is somewhat reduced by the fact that we condition on firms existing in 1999. However, it can be ruled out by controlling for unobserved time-invariant firm heterogeneity. Column (2) of Table 2.7 reports results from a fixed-effects version of our DiD specification. While this reduces the size of our estimated effect by roughly a quarter, the reform's negative effect on minimum hire quality remains strongly significant.[27]

Table 2.7: EPL effect on minimum hire quality: FE and IV

	(1) OLS	(2) FE	(3) IV
DiD estimate (Treated*Post)=1	-0.0484***	-0.0359***	-0.1162***
	(0.0096)	(0.0104)	(0.0207)
Observations	314,144	314,144	314,144
Firms	101,388	101,388	101,388
Adjusted R^2	0.0027	0.0010	.

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the firm level. (1) OLS (ITT); (2) FE (ITT) (3) 2SLS (LATE). All specifications include only year dummies as covariates. Dependent variable: minimum hire quality. Hire quality is measured by AKM worker effects estimated 1986-1992. The year 2000 is excluded to rule out anticipation effects. The first-stage F statistic for (3) is 5563.93.

2.6.6 Endogenous firm size and overcorrection

Column (3) of Table 2.7 reports an estimate of the reform's local average treatment effect (LATE). The two-stage least squares specification uses size in 1999 as an instrument for time-varying actual firm size. It zooms in on the reform's effect on those firms that, based on their 1999 size, were likely to also be small enough to benefit from the reform from 2001 onward. The LATE addresses the possibility that the ITT effect used in our main specifications is overly conservative as it classifies firms no longer benefiting from the reform as treated and vice versa. As expected, the LATE results are more pronounced than our main (ITT) results.

[27] Figures 2.10-2.13 in the appendix show that adding firm fixed effects to Equation 2.2 yields results that are qualitatively the same as OLS, though noisier.

2.6.7 Non-parallel counter-factual trends: firm size

Our results may be spurious if minimum hire quality in different firm size groups had deviated from parallel trends even in the absence of the reform. While this is impossible to ascertain for the firm sizes actually affected by the reform, placebo estimates comparing firms further up the firm size distribution provide a sense of how likely non-parallel counter-factual trends in minimum hire quality are. To this end we define as placebo treated and placebo control those firms that are 10, 20 and 30 workers bigger than our actual treated and untreated firms. This gives us placebo firm size cut-offs at 20, 30 and 40 workers. Table 2.8 gives the results from estimating Equation 2.1 on these placebo treatments. Reassuringly, none of the DiD estimates of the placebo effects are statistically significant.

Table 2.8: EPL effect on minimum hire quality: placebo tests

	(1)	(2)	(3)
Placebo firm size cut-off	20	30	40
DiD estimate (Treated*Post)=1	0.0016	-0.0061	-0.0012
	(0.0029)	(0.0040)	(0.0050)
Observations	109,735	54,044	33,130
Firms	22,331	9,063	4,902
Adjusted R^2	0.0051	0.0064	0.0088

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the firm level. Placebo tests. (1) placebo TR: firms with 12-20 workers, placebo CT: 21-25. (2) placebo TR: 22-30, placebo CT: 31-35. (3) placebo TR: 32-40, placebo CT: 41-45. All specifications include only year dummies as covariates. Dependent variable: minimum hire quality. Hire quality is measured by AKM worker effects estimated 1986-1992. The year 2000 is excluded to rule out anticipation effects.

2.7 Conclusion

In this paper we study the effect of reduced employment protection on firms' worker selection. By doing so, we provide empirical evidence on a less obvious way in which labor market regulation affects both efficiency (through the quality of worker-firm matches) and equity (through the opportunities of the least productive workers).

To obtain quasi-experimental variation in employment protection we use the 2001 relaxation of Sweden's seniority rule. This allows us to test *if* reduced dismissal costs make firms less stringent in their worker selection and if so, *how*. Using individual-level data for the universe of Swedish workers and several ability measures, we find that the EPL reform lowers minimum hire quality at affected firms by about 5% of a standard deviation. This result is robust to addressing various scenarios that threaten the validity of our identification strategy.

To isolate the portion of the EPL effect on minimum hire quality due to a change in firms' selectiveness we implement a counter-factual simulation. We find that about half of the overall EPL effect is explained by an increase in the hiring rate, which mechanically lowers minimum hire quality, whereas the other half represents a change in the selectiveness of firms' hiring. This is our first main finding: employment protection legislation directly affects *who* gets hired.

We devise tests to investigate two mechanisms that may be driving the effect on firms' hiring selectivity. The first possible mechanism is that firms hire riskier workers as dismissal costs fall (Lazear, 1995), implying a widening of the spread of hire quality both at the bottom and the top. To test for it, we look at the response of maximum hire quality to the EPL reform. We do not find evidence for this channel.

The second possible mechanism is that employers respond to looser employment protection by lowering their hiring standard, i.e., the productivity threshold above which they hire applicants (Pries and Rogerson, 2005). We build on the literature on network-based hiring, in particular, Hensvik and Skans (2016), to obtain direct evidence for a change in firms' screening of workers. Our analysis of new hires' past co-worker links with incumbents at their new firm suggests that the reform makes employers rely less on their employees' networks to screen potential recruits. We conclude that a lowering of firms' hiring standard drives the EPL effect on firms' selectiveness in hiring. Our second key result is thus that EPL affects firms' worker selection by shifting their hiring threshold as suggested by Pries and Rogerson (2005).

Our finding that employment protection directly affects firms' worker selection is consistent with the scarce evidence that already exists (Marinescu, 2009; Bjuggren and Skedinger, 2018). We go beyond these papers by establishing a causal effect on the *ability* of new hires and by providing direct evidence for a change in screening intensity.

This is the first study to provide evidence on the mechanisms by which EPL influences firms' worker selection, i.e., *how* dismissal costs affect the selectiveness of firms' hiring. By empirically assessing the predictions of different theoretical models we hope to inform the modelling choices of future theoretical work on worker-firm matching.

Following the 2018 general election in Sweden, the Swedish government re-opened debate on a possible loosening of employment protection legislation that would exempt more firms from the LIFO rule. The implication of our results for this policy discussion is that a further reduction of employment protection may make it yet easier for the least productive job-seekers to get a foot in the door. This points to an intuitive trade-off in the design of labor market policy between safeguarding the job security of insiders and making work attainable for the weakest outsiders.

References

Abowd, John M, Francis Kramarz, and David N Margolis (1999): "High wage workers and high wage firms," *Econometrica*, 67(2): 251–333.

Balsvik, Ragnhild and Stefanie Haller (2015): "Ownership Change and its Implications for the Match between the Plant and its Workers." *NHH Dept. of Economics Discussion Paper*, (10).

Bastgen, Andreas and Christian L. Holzner (2017): "Employment protection and the market for innovations," *Labour Economics*, 46: 77–93.

Bender, Stefan, Nicholas Bloom, David Card, John Van Reenen, and Stefanie Wolter (2018): "Management practices, workforce selection, and productivity," *Journal of Labor Economics*, 36(S1): S371–S409.

Bjuggren, Carl Magnus (2018): "Employment protection and labor productivity," *Journal of Public Economics*, 157: 138–157.

Bjuggren, Carl Magnus and Per Skedinger (2018): "Does job security hamper employment prospects?" IFN Working Paper 1255, Research Institute of Industrial Economics.

Böckerman, Petri, Per Skedinger, and Roope Uusitalo (2018): "Seniority rules, worker mobility and wages: Evidence from multi-country linked employer-employee data," *Labour Economics*, 51: 48–62.

Butschek, Sebastian and Jan Sauermann (2019b): "Do individual fixed effects capture worker ability?" unpublished manuscript, Stockholm University.

Card, David, Jörg Heining, and Patrick Kline (2013): "Workplace heterogeneity and the rise of West German wage inequality," *The Quarterly Journal of Economics*, 128(3): 967–1015.

Correia, Sergio (2016): "Linear Models with High-Dimensional Fixed Effects: An Efficient and Feasible Estimator," Tech. rep., working Paper.

Håkanson, Christina, Erik Lindqvist, and Jonas Vlachos (2015): "Firms and skills: the evolution of worker sorting," Working Paper Series 2015:9, IFAU Institute for Evaluation of Labour Market and Education Policy.

Hensvik, Lena and Oskar Nordström Skans (2016): "Social networks, employee selection, and labor market outcomes," *Journal of Labor Economics*, 34(4): 825–867.

Holmlund, Bertil and Donald Storrie (2002): "Temporary work in turbulent times: the Swedish experience," *Economic Journal*, 112(480): F245–F269.

Hopenhayn, Hugo and Richard Rogerson (1993): "Job turnover and policy evaluation: A general equilibrium analysis," *Journal of Political Economy*, 101(5): 915–938.

Kugler, Adriana D. and Giovanni Pica (2008): "Effects of employment protection on worker and job flows: evidence from the 1990 Italian reform," *Labour Economics*, 15(1): 78–95.

Kugler, Adriana D. and Gilles Saint-Paul (2004): "How do firing costs affect worker flows in a world with adverse selection?" *Journal of Labor Economics*, 22(3): 553–584.

Lazear, Edward (1995): "Hiring risky workers," NBER Working Paper 5334, National Bureau of Economic Research.

Lindbeck, Assar, Mårten Palme, and Mats Persson (2006): "Job Security and Work Absence: Evidence from a Natural Experiment," CESifo Working Paper Series 1687, CESifo Group Munich.

Lindqvist, Erik and Roine Vestman (2011): "The Labor Market Returns to Cognitive and Noncognitive Ability: Evidence from the Swedish Enlistment," *American Economic Journal: Applied Economics*, 3(1): 101–128.

Marinescu, Ioana (2009): "Job security legislation and job duration: evidence from the United Kingdom," *Journal of Labor Economics*, 27(3): 465–486.

Olsson, Martin (2009): "Employment protection and sickness absence," *Labour Economics*, 16(2): 208–214.

——— (2017): "Direct and Cross Effects of Employment Protection: The Case of Parental Childcare," *Scandinavian Journal of Economics*, 119(4): 1105–1128.

Pries, Michael and Richard Rogerson (2005): "Hiring policies, labor market institutions, and labor market flows," *Journal of Political Economy*, 113(4): 811–839.

SFS (1982): *80, Lag om anställningsskydd*, Statens författningssamling.

——— (1991): *433, Jämställdhetslag*, Statens författningssamling.

——— (1993): *1496, Lag om anställningsskydd*, Statens författningssamling.

——— (2000): *733, Jämställdhetslag*, Statens författningssamling.

Sjöberg, Ola (2009): "Temporary work, labour market careers and first births in Sweden," Working Paper 2009:7, SPaDE, Stockholm University, Stockholm.

Skedinger, Per (2008): *Effekter av anställningsskydd: vad säger forskningen?*, SNS Förlag.

von Below, David and Peter Skogman Thoursie (2010): "Last in, first out?: Estimating the effect of seniority rules in Sweden," *Labour Economics*, 17(6): 987–997.

Appendix

Table 2.9: Data sources

Data source	Period used	Description	Used for measurement of	Variables used/generated
JOBB	1986-2004	Employment spells	Worker flows	Timing and identity of hir
JOBB	1993-2004	Employment spells	Firm size	Monthly head count
JOBB	1993-2004	Employment spells	Other firm characteristics	Industry, age, location, owner
JOBB	1986-2004	Employment spells	Worker labor market history	Co-workers
WSS	1986-1992	Wage survey statistics	Wages	Full-time equivalent monthly w
LISA	1990-2004	Various information	Worker characteristics	Highest educational qualifica
BAKGRUND	n.a.	Birth register	Worker characteristics	Year of birth, gender
KRIGSARKIVET	1970-2004	Military draft information	Worker ability	Cognitive, psychological test s
ÅRSKURS9	1988-1997	Grade 9 school grades	Worker ability	GPA at age 15

Table 2.10: EPL effect on minimum hire quality: mechanisms, with outliers

	(1)	(2)	(3)	(4)
	Minimum		Maximum	
	Actual	Simulated	Actual	Simulated
DiD estimate (Treated*Post)=1	-0.0473***	-0.0249**	0.0341***	0.0375***
	(0.0101)	(0.0099)	(0.0104)	(0.0104)
Observations	235,187	245,211	235,187	245,211
Firms	92,774	95,689	92,774	95,689
Adjusted R^2	0.0027	0.0017	0.0026	0.0024

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the firm level. Data include worker inflow outliers. Dependent variable: actual minimum hire AKM (1), counter-factual minimum hire AKM (2), actual maximum hire AKM (3), counter-factual maximum hire AKM (4). For the counter-factual effect estimate, the data have been modified, leaving each firm's actual number of hires unchanged but randomly re-shuffling hired individuals across time within a firm size group. All estimates are from OLS (ITT) specifications without covariates other than year dummies. The year 2000 is excluded to rule out anticipation effects.

Table 2.11: EPL effect on minimum hire quality: firm size corrections

	(1)	(2)	(3)
DiD estimate (Treated*Post)=1	-0.0484***	-0.0469***	-0.0621***
	(0.0096)	(0.0096)	(0.0089)
Observations	314,144	312,152	347,798
Firms	101,388	99,873	124,328
Adjusted R^2	0.0027	0.0027	0.0028

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the firm level. Different versions of reform bite measure: (1) head count -1 all firms; (2) head count - managers and family where known, -1 otherwise; (3) unadjusted head count. All estimates are from OLS (ITT) specifications without covariates other than year dummies. Dependent variable: minimum hire quality. Hire quality is measured by AKM worker effects estimated 1986-1992. The year 2000 is excluded to rule out anticipation effects.

Table 2.12: EPL effect on network-based hiring

	(1)	(2)
DiD estimate (Treated*Post)=1	-0.0347***	-0.0311***
	(0.0059)	(0.0059)
Observations	618,260	659,637
Firms	121,242	121,701
Adjusted R^2	0.0071	0.0078

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the firm level. Estimates are from firm fixed-effects specifications without covariates other than year dummies. Dependent variable: network-based hiring, measured by new hires' average share of incumbents with whom they previously worked at another firm. (1) is for the period 1994-2004. (2) is for 1993-2004. The year 2000 is excluded to rule out anticipation effects.

Table 2.13: EPL effect on age of hires

	(1)
	Mean hire age
DiD estimate (Treated*Post)=1	0.5032***
	(0.1018)
Observations	314,144
Firms	101,388
Adjusted R^2	0.0264

Note: *** p<0.01, ** p<0.05, * p<0.1, with standard errors clustered at the firm level. Dependent variable: age of new hires. Estimates are from OLS (ITT) specifications without covariates other than year dummies and are based on the sample of firms from Column (1) in Table 2.3. The year 2000 is excluded to rule out anticipation effects.

Figure 2.5: EPL effect on minimum hire cognitive test score

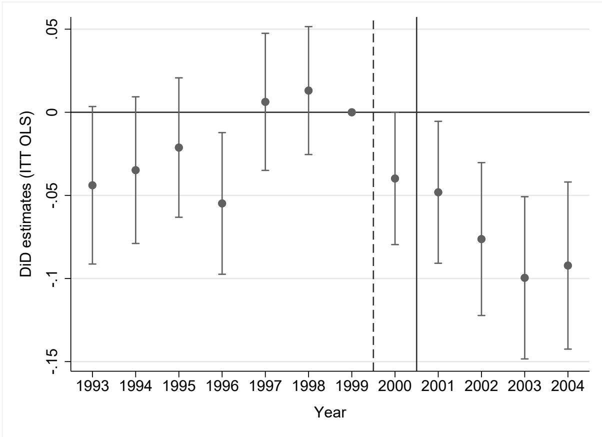

Note: This figure shows yearly DiD estimates for the effect of the EPL reform on minimum hire quality from an OLS (ITT) specification without covariates other than year dummies. The sample of firms is the same as in Figure 2.1. Hire quality is measured by cognitive test scores from the military draft. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the firm level.

Figure 2.6: EPL effect on minimum hire psychological test score

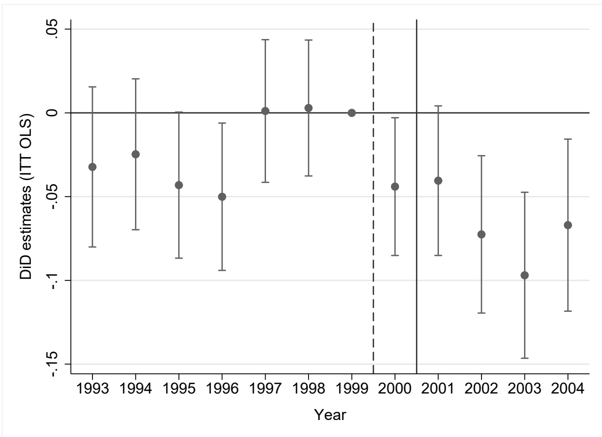

Note: This figure shows yearly DiD estimates for the effect of the EPL reform on minimum hire quality from an OLS (ITT) specification without covariates other than year dummies. The sample of firms is the same as in Figure 2.1. Hire quality is measured by psychological test scores from the military draft. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the firm level.

Figure 2.7: EPL effect on minimum hire grade point average

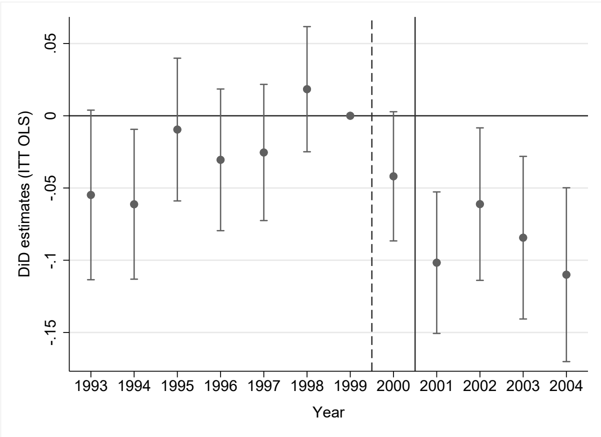

Note: This figure shows yearly DiD estimates for the effect of the EPL reform on minimum hire quality from an OLS (ITT) specification without covariates other than year dummies. The sample of firms is the same as in Figure 2.1. Hire quality is measured by GPA at age 15. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the firm level.

Figure 2.8: EPL effect on network-based hiring

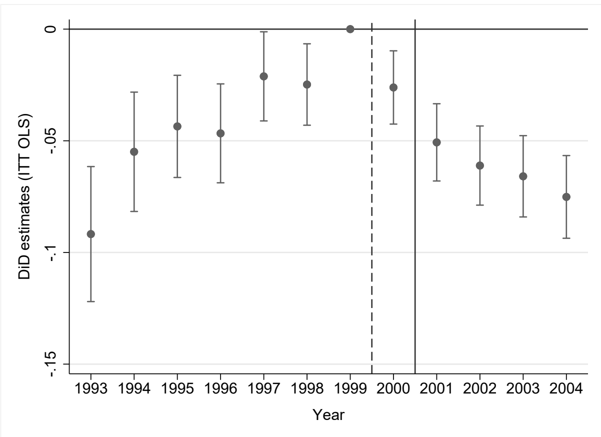

Note: This figure shows yearly DiD estimates for the effect of the EPL reform on network-based hiring from an OLS specification without covariates other than year dummies. Network-based hiring is measured by new hires' average share of incumbents with whom they previously worked at another firm. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the firm level.

Figure 2.9: EPL effect on network-based hiring (FE)

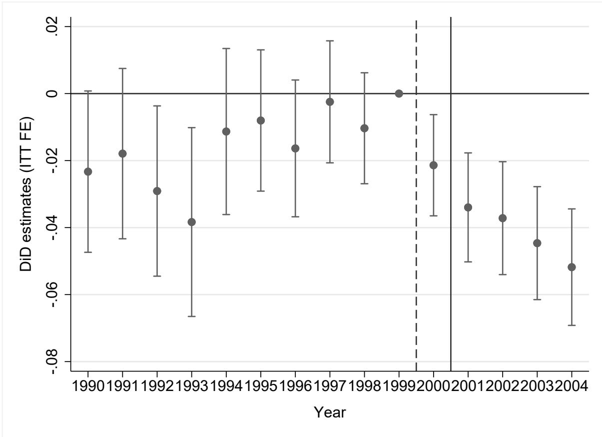

Note: This figure shows yearly DiD estimates for the effect of the EPL reform on network-based hiring from a firm fixed-effects specification without covariates other than year dummies. Network-based hiring is measured by new hires' average share of incumbents with whom they previously worked at another firm. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the firm level.

Figure 2.10: EPL effect on minimum hire quality (FE)

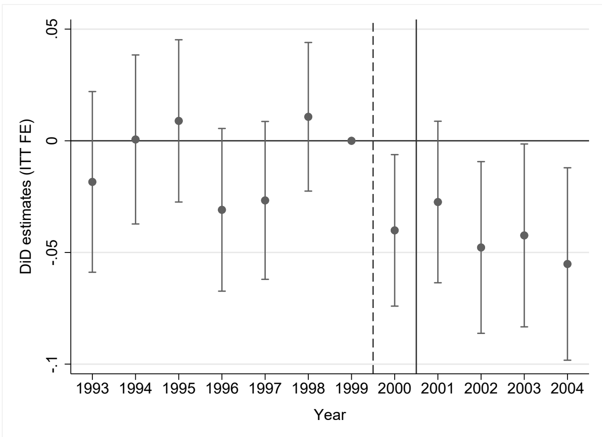

Note: This figure shows yearly DiD estimates for the effect of the EPL reform on minimum hire quality from a firm fixed-effects specification without covariates other than year dummies. Hire quality is measured by AKM worker effects estimated 1986-1992. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the firm level.

Figure 2.11: EPL effect on minimum hire cognitive test score (FE)

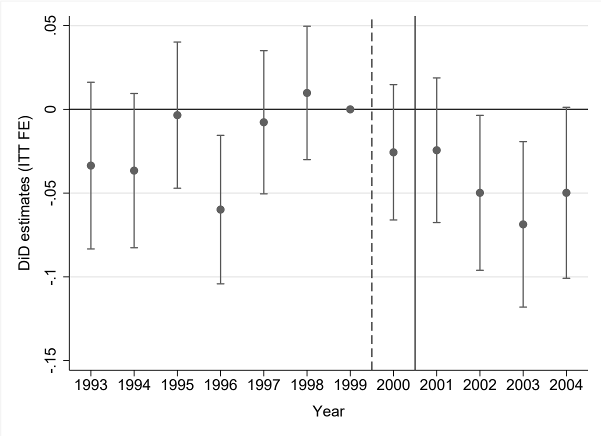

Note: This figure shows yearly DiD estimates for the effect of the EPL reform on minimum hire quality from a firm fixed-effects specification without covariates other than year dummies. The sample of firms is the same as in Figure 2.1. Hire quality is measured by cognitive test scores from the military draft. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the firm level.

Figure 2.12: EPL effect on minimum hire psychological test score (FE)

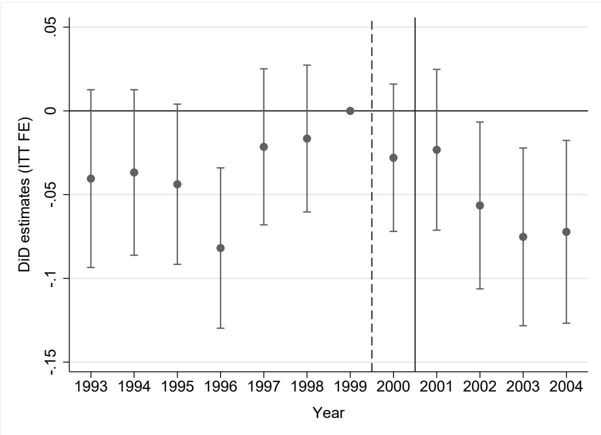

Note: This figure shows yearly DiD estimates for the effect of the EPL reform on minimum hire quality from a firm fixed-effects specification without covariates other than year dummies. The sample of firms is the same as in Figure 2.1. Hire quality is measured by psychological test scores from the military draft. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the firm level.

Figure 2.13: EPL effect on minimum hire grade point average (FE)

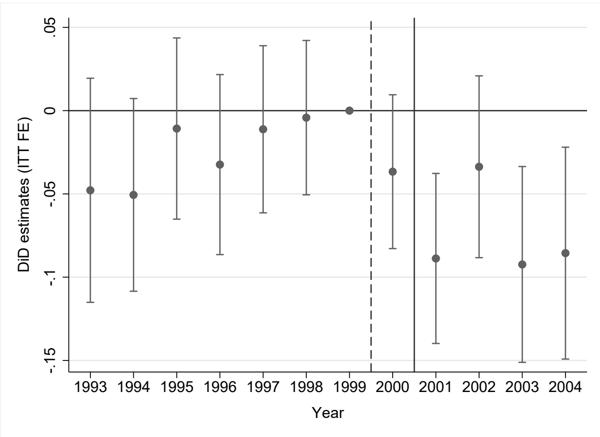

Note: This figure shows yearly DiD estimates for the effect of the EPL reform on minimum hire quality from a firm fixed-effects specification without covariates other than year dummies. The sample of firms is the same as in Figure 2.1. Hire quality is measured by GPA at age 15. Vertical bars denote 95% confidence intervals. Standard errors are clustered at the firm level.

Chapter 3

Do estimated individual fixed effects capture worker ability?

Authors: Sebastian Butschek and Jan Sauermann

Abstract: In this paper we assess the performance of estimated worker effects as an ability proxy. We focus on worker effects from two-way fixed-effect log-wage regressions as introduced by Abowd et al. (1999). Using Swedish register data, we study the correlation between such worker effect estimates and standardized cognitive test scores from the military draft. For the period from 2007 to 2013 we find a correlation of roughly .4. We show that this is more than three quarters of the predictive performance achieved by machine-learning algorithms. However, we find that the worker effects' prediction performance varies with observed worker and firm characteristics—for example, across industries. Moreover, we identify biases—for instance, the worker effects underestimate cognitive test scores for blue-collar workers. Our analysis of the worker effects' determinants shows that they "contain" a significant amount of both skill-related and non-skill attributes such as education and cognitive ability as well as age. Finally, our exploration of the importance of estimation ingredients for prediction performance broadly supports the estimation choices of existing applications.

Keywords: cognitive ability, test scores, unobserved productivity, AKM.

JEL Codes: J00, J24, J31.

3.1 Introduction

As direct measures of individual ability are scarce, a growing body of empirical work in applied microeconomics uses estimated productivity proxies. Following Abowd et al. (1999, henceforth AKM), a range of papers interpret estimated worker effects from two-way fixed-effects log wage regressions as a broad measure of time-invariant worker ability. Opposing this tendency, there is an increasing theoretical consensus, backed up by some empirical evidence, that there are problems with the interpretation of AKM worker effects as a productivity proxy. These criticisms notwithstanding, AKM worker effects continue to be used to approximate individual ability—often, because they are relatively easy to obtain or because of a lack of ready alternatives. Moreover, while several papers have made a strong theoretical case for what AKM effects *are not*, it remains unclear what AKM effects *are*. AKM worker effects, in particular, are a blackbox. That is, we do not know what characteristics or estimation biases they reflect.

In this paper we bring empirical evidence to bear on these questions. Leveraging the availability of a high-quality measure of cognitive ability in Swedish register data, we (i) assess estimated AKM worker effects' performance as an ability proxy; (ii) study how this performance varies across worker and firm observables; (iii) explore AKM worker effects' potential bias; (iv) describe their empirical contents, i.e., their determinants; and (v) investigate the consequence of estimating AKM worker effects under certain data constraints. To answer these questions we make descriptive comparisons between standardized test scores for cognitive ability (COG) and estimated AKM worker effects.

We find an overall AKM-COG correlation of nearly .4, which our machine-learning benchmarks suggest is sizeable. Our results further show that observed firm and worker characteristics predict both the size of the AKM-COG correlation and whether AKM worker effects systematically over- or underestimate COG. Our attempt to open the AKM blackbox suggests that AKM worker effects "contain" both skill—such as education, cognitive and non-cognitive ability—and other worker or job attributes that command wage premia—such as age, occupation type and industry. Finally, we are able to provide some practical guidelines for AKM estimation when the goal is to use approximate worker ability. For instance, we find that an estimation interval of six years' length harnesses nearly all the information on individual ability that AKM worker effects are able to extract in an interval as long as 14 years.

The AKM method uses a two-way fixed-effects regression to decompose log wages into a firm component related to the time-invariant characteristics of the employer and a worker component related to the time-constant characteristics of the individual. This approach assumes that firm and worker pay premia are additively separable, so that the firm pay premium is independent of the characteristics of the worker: regardless of the precise worker hired, wages are higher at more productive firms. Moreover, AKM estimations assume that conditional on time-invariant worker and firm characteristics and time-varying observables, worker mobility is exogenous. The original application of the AKM method is a study of sorting in the French labor market. AKM use the correlation between estimated AKM worker and firm effects (as proxies of fixed characteristics) to measure the degree of assortative matching (Abowd et al., 1999). Their approach has been replicated using data on other countries, sometimes with modifications or extensions (Abowd et al., 2002; Iranzo et al., 2008; Gruetter and Lalive, 2009; Bagger et al., 2013; Card et al., 2013, 2015; Alvarez et al., 2018; Barth et al., 2018; Lopes de Melo, 2018; Song et al., 2018).

Interest in labor market sorting is not only the source of a string of replications but also supplies a key critique of the AKM method. Equilibrium sorting models argue that if there are complementarities between worker and firm productivity, wages will not be monotonic in firm productivity (e.g., Eeckhout and Kircher, 2011; Hagedorn et al., 2017). For example, a bad worker may earn less at a good firm than at a bad firm as the good firm will only hire him at a very low wage to compensate it for foregoing a more productive hire. As a consequence the AKM firm effects will not be correlated with firm productivity (Eeckhout and Kircher, 2011). In the presence of search frictions, their model implies that wages may also not be monotonic in worker productivity. A good worker in a bad match may earn less than a slightly less productive worker in a good match. This may result in AKM worker effects not being correlated with worker productivity.

A second major critique of the AKM method concerns the assumption on (conditionally) exogenous mobility, which rules out specific instances of endogenous mobility (e.g., Abowd et al., 2018; Bonhomme et al., 2018).[1] One example is mobility to high-wage firms of high-productivity workers whose ability is revealed only gradually in the pre-move job. Such a pattern would attribute too little of the wage premium to the worker and too much to the firm (Card et al., 2013; Abowd et al., 2018). A second example is mobility driven by specific labor demand shocks, which will also overstate the firm pay premium (Card et al., 2013).[2] A final example of endogenous mobility ruled out by AKM is that the current wage realization (e.g., from a match-specific productivity component) influences future mobility. There is empirical evidence for the US and Sweden that current match-specific heterogeneity in wages affects whether and to what type of firm a worker moves in future (Abowd et al., 2018; Bonhomme et al., 2018). The resulting biases of estimated AKM worker and firm effects may undermine the use of AKM effects as measures of worker skill and firm productivity (Abowd et al., 2018).

Despite these potential or actual problems there are numerous applications that interpret AKM worker effects as a measure of individual ability. This may in part be due to the attention received by Card et al. (2013, henceforth CHK), who use the AKM method to study the contribution of employer pay heterogeneity and assortative matching to a rise in West German wage inequality. Both pre-dating CHK and citing their work, a string of papers takes the AKM method beyond the study of labor market sorting and uses AKM worker effects as an explicit proxy for time-invariant worker characteristics (e.g., Fox and

[1] This is different from the critique that identification of AKM firm effects in practice is undermined by limited mobility bias, i.e., that too few of firms' workers are movers (Abowd et al., 2004; Andrews et al., 2008, 2012; Borovičková and Shimer, 2017).

[2] Consider a worker that regularly moves between two industries. Industry A experiences strong cyclical fluctuations so that it pays high wages during an upswing but jobs are quickly destroyed in a downswing. Industry B offers stable jobs at lower wages. If the worker moves to an A firm during an A upswing and moves back to a B firm when the A boom ends, AKM firm effects will exaggerate the firm-specific wage component. This is because the estimation attributes wage changes accompanying worker movements to time-invariant firm premia although they are partly driven by labor demand shocks.

Smeets, 2011; Abowd et al., 2012; Davidson et al., 2014; Krishna et al., 2014; Balsvik and Haller, 2015; Schmutte, 2015; Carlsson et al., 2016; Macis and Schivardi, 2016; Cornelissen et al., 2017; Roca and Puga, 2017; Butschek, 2017; Bender et al., 2018; Gathmann et al., 2018; Bombardini et al., 2019; Butschek and Sauermann, 2019a).[3] Another explanation for the use of AKM worker effects as ability proxies is the scarcity of alternatives. High-quality direct measures of cognitive ability are available in only a few data sets. The best-known example is the Armed Services Vocational Aptitude Battery (ASVAB) contained in the National Longitudinal Survey of Youth (NLSY79). Also well-known are military draft test scores for men's cognitive ability from the Swedish Armed Forces.[4]

Motivated by the continued popularity of AKM worker effects as a proxy for individual ability and the dearth of substantive justification for doing so, we attempt to provide direct empirical evidence for the strengths and weaknesses of AKM worker effects as an ability proxy. To this end we estimate the type of AKM specification implemented by CHK on a sample of Swedish men for the seven years from 2007 to 2013. We then compare the estimated AKM worker effects to cognitive test scores from the Swedish military draft in simple descriptive analyses to answer five sub-questions.

First, how well do AKM worker effects approximate cognitive test scores? We use the correlation between the two measures as our main performance indicator. We find an AKM-COG correlation of nearly .4. To help assess whether this makes AKM worker effects a good ability proxy we benchmark it with machine-learning predictions of cognitive test scores based on the same labor market data. While the machine-learning algorithms do better, AKM worker effects achieve three quarters of their predictive performance. This finding will not end the debate on the merits of AKM worker effects as an ability proxy, but it is the first one that puts a simple number on AKM worker effects' predictive performance.

Second, does the quality of AKM worker effects as a cognitive ability proxy vary across worker and firm characteristics? Comparing the AKM-COG correlation across observables, we find considerable heterogeneity. In particular, the correlation improves with age, is higher for white-collar than blue-collar occupations and varies strongly across industries. It is similar across educational categories and regions. The main use of these findings is practical—to help identify settings in which the use of AKM worker effects is more and less justified.

Third, are AKM worker effects biased? To identify systematic exaggeration or understatement of cognitive ability we regress AKM prediction error on the worker and employer characteristics we observe. Our findings suggest that AKM worker effects underestimate

[3] AKM person effects have also been used as a proxy for other dimensions of time-invariant individual heterogeneity (e.g., Finkelstein et al., 2016). Analogously to AKM, education economists have used two-way fixed-effects regressions with test scores as an outcome to estimate teacher or school value added (e.g., Rothstein, 2010). In this literature fixed effects are also interpreted as capturing ability, though the focus is on teachers or schools.

[4] Test scores from the military draft have been used, amongst others, by Lindqvist and Vestman (2011); Håkanson et al. (2015); Fredriksson et al. (2018).

cognitive ability for younger workers, blue-collar workers, more educated workers and educators. These results do not directly map onto the theoretical debate on sources of bias in the AKM method. By contrast, our finding that the (downward) bias of AKM worker effects for firm movers as compared to stayers is negligible in size is relevant to the literature on AKM mobility biases, though we cannot distinguish types of mobility so as to identify potential countervailing effects.

Fourth, what individual characteristics do AKM worker effects contain? Regressing estimated AKM worker effects on COG, non-cognitive test scores and typically observed worker and firm characteristics we find that ceteris paribus, most variation in AKM worker effects is explained by education and age. This is followed by occupation type, non-cognitive test scores and COG, all of which are also contained in AKM worker effects in economically significant quantities. The contribution of our analysis is to improve our empirical handle on what AKM worker effects actually are—which may have both theoretical and practical relevance.

Fifth, how is the quality of AKM worker effects as an ability proxy affected by the way they are estimated? We show that the AKM-COG correlation varies only slightly across four non-overlapping 7-year estimation intervals; that hourly wages are a key estimation ingredient unless part-time workers can be excluded from the AKM estimation; that the performance of AKM worker effects as ability proxies is reduced, but still acceptable, when (employed) workers are not observed for the entire estimation period; that longer estimation intervals increase the predictive performance of AKM worker effects, but only up to six years; that including dummies for the level of education reduces the AKM-COG correlation only slightly; and that estimating AKM specifications for men and women jointly does not undermine AKM worker effects' performance as an ability proxy for men. These findings may facilitate estimation decisions for researchers approximating ability with AKM worker effects.

3.2 Method

3.2.1 AKM estimation

The starting point of this paper is the two-way log-wage fixed-effects specification introduced by Abowd et al. (1999). The estimates of time-invariant individual wage premia this yields are widely referred to as AKM worker effects. To obtain AKM worker effects for our analysis we closely follow CHK by estimating the following specification:

$$\ln(w)_{ijt} = \alpha_i + \psi_j + \gamma_t + x'_{it}\beta + r_{ijt}, \tag{3.1}$$

where $\ln(w)_{ijt}$ is individual i's log wage at firm j in year t. α_i denotes additive fixed effects for individuals i and ψ_j denotes additive fixed effects for firms j. γ_t is a set of year dummies and x_{it} is a vector of time-varying individual-level controls: age squared, age cubed and interactions of education categories with year dummies, age squared and age cubed. This allows time effects and second- and third-degree age effects to vary by educational category.[5] For comparability we also follow CHK in estimating the AKM specification separately for men.[6] We only deviate from CHK in our wage measure: they use wages unadjusted for hours worked but exclude part-time workers. Our data allow us to make an approximate distinction between full-time and part-time workers for a subset of workers; for this subset, however, we can do even better by using full time-equivalent (FTE) wages, which is what we opt for in our main analyses.[7]

To be able to estimate α_i for a worker it is necessary to have at least two observations for this individual. Similarly, ψ_j is not identified for firms that only have a single employee and for firms whose workers are never employed elsewhere, i.e., those that are not connected to other firms by worker movements. We obtain AKM worker effect estimates $\hat{\alpha}_i$ for 80.13% of the workers for whom we observe a monthly FTE wage and estimates of ψ_j for 42.68% of firms.[8] While this may seem low in comparison to, e.g., CHK, the main driver of the difference is that our register data contain self-employed individuals, whereas the social security records underlying CHK's estimations do not.

3.2.2 AKM approximation quality and bias

Throughout the paper, we use standardized versions of AKM worker effects and COG. To evaluate the performance of estimated AKM worker effects as predictors of cognitive test scores, we focus on the Pearson sample correlation coefficient between AKM worker effects and COG, the AKM-COG correlation $r = \frac{\sum_{i=1}^{N}(AKM-\mu_{AKM})(COG-\mu_{COG})}{\sqrt{\sum_{i=1}^{N}(AKM-\mu_{AKM})^2}\sqrt{\sum_{i=1}^{N}(COG-\mu_{COG})^2}}$, where μ_{AKM} and μ_{COG} are AKM worker effect and COG means.

The correlation coefficient is an easily interpretable measure for the predictive performance of AKM worker effects but it is blind to the direction in which AKM worker effects go wrong when approximating COG. To get a sense of whether AKM worker effects systematically under- or overestimate COG we consider prediction error $PE = AKM - COG$ as a measure of the bias of AKM worker effects.

[5]Note that Equation (3.1) does not include non-interacted education category dummies. This follows the formulation in CHK; their reason for excluding education levels is not given. Presumably, one reason is that for most people education does not vary over time. Perhaps a more important reason is that CHK *want* to absorb time-invariant educational human capital in the worker fixed effect.

[6]We explore the effects on AKM worker effects' predictive performance of including women in the AKM estimation in Section 3.4.5.

[7]We explore the value added from hourly wages in Section 3.4.5.

[8]We rely on `reghdfe` (Correia, 2016) for our estimation of Equation (3.1). This module drops singletons (19.41% of individuals and 47.52% of firms).

3.2.3 Machine learning as a benchmark

It is difficult to know what correlation constitutes good predictive performance. Both cognitive test scores and the labor market/demographic data underlying AKM worker effects have stochastic components. Because of this randomness even an algorithm that fully captures the data-generating process will not perfectly predict cognitive ability. This implies that the correlation between any cognitive ability predictor and cognitive test scores will never be one—there is an irreducible prediction error.

Estimated AKM worker effects were not specifically designed to approximate cognitive ability; that they may be used as a (cognitive) ability proxy is a beneficial side effect. When predicting cognitive ability is the explicit goal, it is likely that one can do significantly better than estimating AKM worker effects by training an algorithm specifically for the task of predicting cognitive ability—for example, by using machine learning.[9]

Why might one still rely on AKM worker effects to approximate worker ability? The key advantage is that an AKM regression model does not need training and testing, making it feasible in settings where no explicit information on cognitive ability is available. By contrast, ML algorithms are trained by minimizing their mean squared error (MSE) in predicting observed cognitive ability.[10] Model assessment also requires information on cognitive ability.

In light of this, the point of comparing the predictive performance of AKM worker effects and ML algorithms is not to help decide which one to use. Instead, the idea is to use machine learning as a benchmark—to obtain an (upper-bound) estimate of the size of the irreducible error of a cognitive ability prediction based on labor market data. The performance of AKM worker effects as an ability proxy can then be gauged against the correlation a state-of-the-art prediction algorithm achieves based on the same data.

We consider three ML algorithms: a gradient boosting model, a bootstrap-aggregated decision tree (bagging) model and a random forest model.[11] These are so-called ensemble methods that differ in how they combine several decision trees. A decision tree is an algorithm that partitions the feature space into regions and uses the mean outcome in each region to make predictions. The splits chosen by decision trees are often sensitive to particular observations, implying that decision tree predictions may vary a lot depending on the particular sample the trees have been trained on. Gradient boosting, bagging and random forest models are different tools for reducing the variance of decision trees—by training several trees on derivatives or sub-sets of the training data and aggregating their predictions. A gradient boosting model trains decision trees sequentially (each on the

[9] Given input-output pairs and one or more ways of making predictions, supervised machine learning uses chooses a range of parameters to optimize the prediction of outputs based on the inputs it observes.

[10] $MSE = \frac{1}{N}\sum_{i=1}^{N}(AKM - COG)^2$.

[11] Admittedly, this selection is arbitrary. Our main reason for choosing gradient boosting, bagging and random forests was that these models are both powerful and have been implemented in open-source software.

prediction residuals from the previous one) and adds their predictions. A bagging model randomly draws several samples of the same size as the training set from the training data with replacement, trains a decision tree on each of them, obtains a prediction from all trees and averages them. A random forest proceeds like a bagging model but increases tree differentiation by using only a random sub-set of features each time it looks for an optimal split of the feature space.[12]

For training, model assessment and prediction we utilize the ML algorithm implementations in Python's SciKitLearn. We use data on 2000-2006 for ML training and model assessment, and data on 2007-2013 for benchmarking the predictive performance of AKM worker effects. We focus on the balanced panel in each data set, i.e., the subset of workers we observe for the full seven years (2000-06 or 2007-13).[13] The data we use for training, model assessment and prediction include cognitive test scores and the information (at yearly frequency) that underlies the AKM estimation: log wages, firm ID, education categories, age squared and age cubed (see Equation (3.1)). We standardize the continuous explanatory variables other than firm ID. We randomly split the 2000-2006 data in half, creating a training and a validation set. Having trained the algorithms, we compare their predictive performance in the validation set, choosing the one with the lowest MSE. We then use the best algorithm's prediction performance for cognitive ability in 2007-2013 as a benchmark for AKM worker effects for 2007-13.

3.3 Data

Labor market information We have individual-level register data from Statistics Sweden (SCB) for the period 1986-2013. In the bulk of our analyses we focus on AKM worker effect estimates for the years 2007-2013. The estimation sample includes men aged 18-65 who earn non-zero wages from paid employment in November of a given year.[14] To estimate our main AKM worker effects we use full time-equivalent (FTE) contractual wages from the Wage Survey Statistics (WSS). WSS only covers the workers of a subset of firms (45.76% of observations for our main estimation period).[15] We use the CPI to deflate wages and winsorize at .5% and 99.5% of the annual wage distribution to exclude extreme outliers. We obtain firm identifiers, industry, location and worker occupation from *Job-*

[12]See, e.g., James et al. (2013) for an introduction to machine learning.

[13]SciKitLearn's ML tools best handle panel data in wide format and cannot deal with missing values.

[14]Women are excluded as we have no military draft test scores for almost all of them.

[15]WSS wage data's partial coverage of firms and workers may reduce the quality of AKM worker effects. Reassuringly, this suggests the AKM-COG correlation we find is a lower-bound estimate of AKM worker effects' prediction performance.

bregistret (JOBB), a data base of individual-level employment spells.[16] Information on age and education is drawn from other registers—see Table 3.10 in the appendix for details.[17]

We also estimate Equation (3.1) on the earlier years in three intervals: 1987-1992, 1993-1999 and 2000-2006. The age distribution of workers we observe in the WSS data looks similar across periods apart from showing particularly big birth cohorts in the mid-1940s and the late 1960s (see Panel (a) of Figure 3.13). Having AKM worker effect estimates from non-overlapping estimation intervals allows us to study differences in the predictive performance with increasing worker age (see Section 3.4.2) and across time periods (see Section 3.4.5). [18]

Cognitive test scores Our cognitive ability measure comes from Sweden's military enlistment, which was compulsory from 1969 to 2005 for male Swedes aged 18-19. Draft testing took place over two days and included assessment of cognitive ability, physical fitness, health status, and psychological aptitude for the demands of war.[19] Between 1970 and 1996, well over 90% of Swedish-born males were tested at age 18 or 19. Between 1997 and 2005 this gradually fell to about 70%, before rapidly declining (see Figure 3.14 in the appendix). As a consequence of this timing, our AKM estimation intervals differ in terms of the age groups for whom we have military draft test scores. We primarily focus on the years 2007-2013 because the age distribution of workers with cognitive test scores has the broadest support in this interval (see Panel (b) of Figure 3.13 in the appendix).

The cognitive ability scores we use are based on performance in four dimensions: verbal reasoning; logical induction; metal folding; and technical comprehension. Each sub-test contains 40 questions, is graded on a scale from 1 to 9, and is then transformed to a discrete, "stanine" distribution by the enlistment agency. We use the average of these four sub-scores and standardize it within cohort to give it mean zero and unit variance (Lindqvist and Vestman, 2011; Håkanson et al., 2015) and again across cohorts to make it comparable over time.[20]

While the draft test scores, relatively speaking, are a high-quality measure of cognitive ability, they are not without limitations. Although they are based on performance in four dimensions, they cannot hope to capture all aspects of cognitive ability. Moreover, they are surely measured with error - not just because they reflect performance on one or two single days but also because not everyone participating will have tried their best to impress

[16]JOBB has total monthly wage earnings unadjusted for the number of hours worked and no information on labor supply. We only use these unadjusted wages in Section 3.4.5 to explore how the prediction performance of AKM worker effects is affected by the quality of the wage measure.

[17]LISA, the register that contains education information, is available only from 1990. We impute the level of education for 1986-1989 from 1990 (or the first time an individual is observed in LISA).

[18]We also use the raw 2000-2006 data to train our ML algorithms.

[19]For more detailed information on the military enlistment, see, e.g., Lindqvist and Vestman (2011).

[20]For the test years 1969 to 1994, the cognitive test scores were kept unchanged; they were subject to slight revisions after that (Håkanson et al., 2015).

the enlistment officers. Finally, additional measurement error may arise as cognitive ability increases or decreases over time. This is not unreasonable in the space of half a lifetime, though fluid intelligence is viewed as relatively stable over time.

3.4 Results

3.4.1 Correlation AKM worker effects and COG

We first assess the closeness of the link between AKM worker effects (estimated 2007-13) and military draft test scores for cognitive ability. Figure 3.1 provides a non-parametric illustration of the AKM-COG relationship. The binned scatter plot considers equal-sized bins of standardized cognitive test scores and shows the averages of the standardized AKM worker effect estimates for each of them. Figure 3.1 suggests the relationship is positive and close to linear.

Figure 3.1: Non-parametric relationship AKM worker effects and cognitive ability

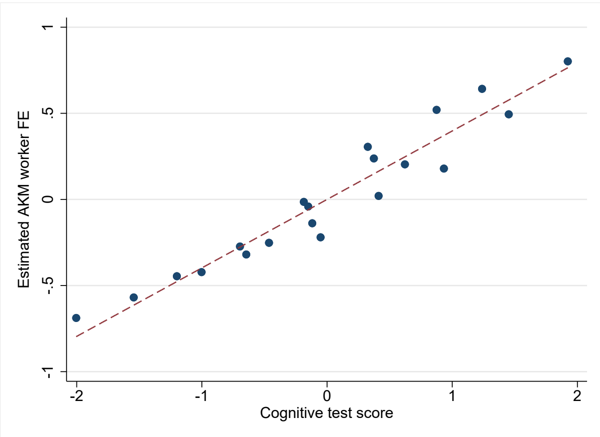

Note: This figure gives a binned scatter plot of standardized AKM worker effects (estimated 2007-2013) and cognitive test scores (standardized separately for each year and overall).

The Pearson sample AKM-COG correlation is .398 among the 881,490 male workers for whom we have both measures (see Column (1) of Table 3.1).[21] As a first gauge for whether this makes AKM worker effects a good proxy for cognitive ability we check how frequently AKM worker effects and COG are in the same quartile. Column (2) of Table 3.1 shows this is the case in 35.9% of cases. To put this into perspective: if AKM worker effects held no information about cognitive ability, AKM worker effect and COG quartiles should

[21] Spearman's ρ, a rank correlation coefficient, is very similar, see Table 3.7 in the appendix.

coincide in roughly 25% of cases. Column (2) of Table 3.1 implies that as a proxy for COG quartiles AKM worker effect quartiles outperform chance by 10.9 percentage points or by 43.6%.

Table 3.1: Estimated AKM worker effects and cognitive ability

	(1) Correlation r	(2) Share same quartile
AKM worker effects and COG	.398	.359
Observations	881,490	881,490

Note: Estimated AKM worker effects and cognitive test scores are standardized to mean zero and unit variance. (1) displays their correlation coefficient; (2) shows the share of individuals for whom the estimated AKM worker effect and the cognitive test score fall in the same quartile.

Benchmarking the prediction performance of AKM worker effects A further way to judge whether an AKM-COG correlation of .398 makes AKM worker effects a good cognitive ability proxy is to ask how high the correlation could possibly be given the data the AKM method uses—after all both cognitive test scores and the labor market/demographic information underlying AKM worker effects have a random component. To get a sense of how much information about cognitive ability *can* be extracted from the data used by AKM estimations we look at the correlation with cognitive test scores achieved by the prediction from a state-of-the-art machine-learning algorithm.

We use data on the years 2000-2006 to train three ML algorithms and choose between them. Out of the three tree-based algorithms we consider, the gradient boosting model marginally outperforms the bagging and the random forest models.[22]

Table 3.2 compares the predictive performance of AKM worker effects and gradient boosting. Column (1) reports our baseline AKM-COG correlation for 2007-13 as in Table 3.1, though for the sub-set of workers who are observed in the labor market throughout 2007-2013 (our ML algorithms require a balanced panel for making a prediction). Column (2) reports the correlation of the gradient boosting model's cognitive ability prediction for 2007-13 with actual COG. The table shows that the AKM-COG correlation is lower than the ML-COG correlation. However, estimated AKM worker effects achieve 77.99% of the predictive performance of the top ML algorithm we trained.

3.4.2 Size of correlation varies with worker and firm characteristics

We explore the heterogeneity of the link between AKM worker effects and COG by looking at their correlation separately by various observed worker and firm characteristics.

[22]Mean squared error in the validation set is smallest for the gradient boosting model—see Table 3.8 in the appendix.

Table 3.2: Benchmarks for AKM worker effects as ability proxies

	(1) AKM worker effects	(2) Boosting
Correlation r with COG	.450	.577
Observations	356,958	356,958

Note: AKM worker effects, out-of-sample ML predictions and cognitive test scores have been standardized to mean zero and unit variance. Columns (1) and (2) compare the prediction performance of AKM worker effects and a gradient boosting model. Both correlations are computed for the sample of workers observed throughout 2007-13.

Worker age We have two pieces of evidence that AKM worker effects are a better ability proxy for older than for younger workers. One is given by Panel (a) of Figure 3.2, which takes a group of workers and compares the correlation between their (time-invariant) COG and their AKM worker effects, estimated in four non-overlapping time periods. It shows that the correlation is .387 for 1986-92, at the start of which workers are 27-33 years old. This correlation increases gradually to .457 for 2007-13 (in 2007 workers are 48-54 years old). A drawback of this comparison is that it conflates the effect of age with that of any longer-run macro-economic changes, such as business cycles.[23] The second piece of evidence is a comparison of AKM-COG correlations between different age groups within the same AKM estimation interval.[24] This avoids capturing time effects at the cost of making inter-cohort comparisons. As Figure 3.2 Panel (b) shows, this comparison confirms that the AKM-COG correlation increases with age.[25] It additionally suggests that the correlation plateaus around .47 after age 40.

Worker education and occupation Our second heterogeneity result is that the AKM-COG correlation varies non-monotonically with education and is higher for white-collar than for blue-collar workers. Moreover, the correlation is lower conditional on either occupation or education than overall. Figure 3.3 Panel (a) compares the AKM-COG correlation across educational categories. It shows a correlation of around .26 for most categories. For workers with a PhD, however, it is markedly lower (r=.140)[26] and for those with only some post-secondary education, it is somewhat lower (r=.219).That AKM worker effects are a worse proxy for cognitive ability *within* educational categories implies that some of the overall correlation between AKM worker effects and COG comes from differences in both

[23] We further explore the effect of the time period in Section 3.4.5.

[24] Age groups 2-5 in Figure 3.2 Panel (b) are chosen to match the beginning-of-period ages in Figure 3.2 Panel (a).

[25] The variation in AKM worker effects is much lower than that in COG for younger workers (see Table 3.9 in the appendix). This would be consistent with higher-cognitive ability individuals entering the labor market later and/or displaying experience wage profiles with a similar intercept and a steeper slope.

[26] PhD holders have more variation in their AKM worker effects than in their COG (see Table 3.9 in the appendix), which may explain some of their lower AKM-COG correlation.

Figure 3.2: AKM worker effect prediction performance by worker age

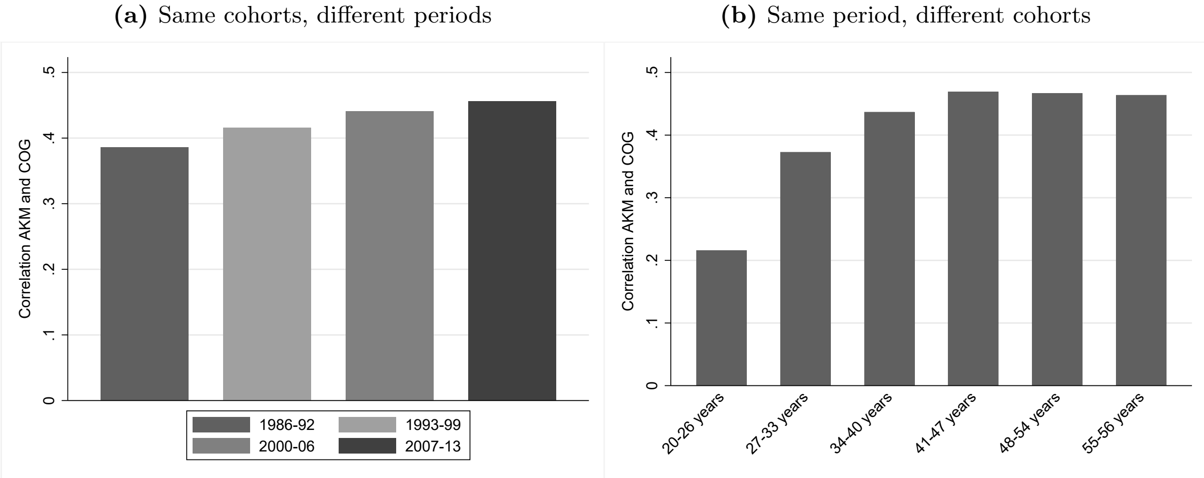

Note: Panel (a) shows correlations between estimated AKM worker effects and cognitive test scores for four non-overlapping AKM estimation intervals. The same seven cohorts (born 1953-1959) are followed across the four intervals. Panel (b) shows correlations between estimated AKM worker effects and cognitive test scores for different age groups, all from the same AKM estimation interval (2007-2013).

average cognitive ability and AKM worker effects *between* educational groups.[27] Figure 3.3 Panel (b) shows a comparison between occupation types. AKM worker effect prediction performance is .248 for white-collar workers, .226 for intermediate occupations and .145 for blue-collar workers. Again, much of the overall AKM-COG correlation seems to be coming from between-occupation variation in cognitive test scores and AKM worker effects. Among white-collar workers, AKM worker effects are a less good proxy for cognitive ability than overall, but still nearly twice as good as for blue-collar workers. That the AKM-COG correlation is so low for blue-collar workers is mirrored by much less variation in their AKM worker effects than their COG (see Table 3.9 in the appendix). This may in turn be due to greater wage compression and lower prevalence of variable pay in blue-collar occupations.

Employer characteristics: industry and location Our third heterogeneity result is that the correlation between AKM worker effects and COG varies considerably across industries but hardly across regions. Figure 3.4 Panel (a) compares AKM-COG correlations across industries, assigning workers employed in several industries to the one they spent most time in. AKM worker effects work best as cognitive ability proxies in manufacturing (r=.449) and worst in the hospitality sector (r=.200) and extra-territorial organizations (there are very few observations in the last). We do not see a clear pattern in the het-

[27]Note that the AKM specification adopted from CHK (see Equation (3.1)) contains interactions of education dummies but not education dummies by themselves. This explains why the AKM worker effect estimates capture considerable between-education variation in cognitive ability.

Figure 3.3: AKM prediction performance by worker education and occupation

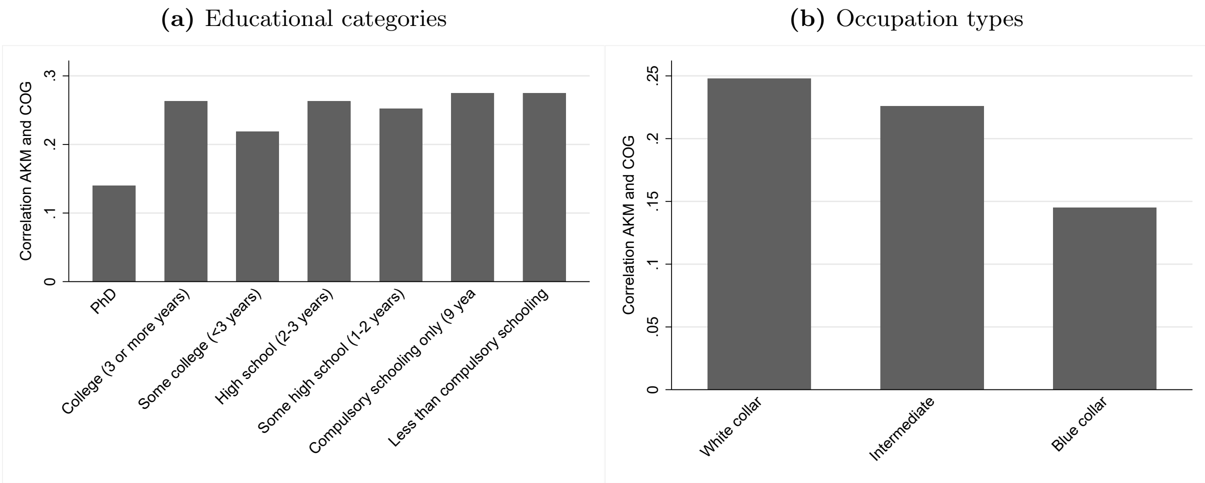

Note: Panel (a) shows correlations between estimated AKM worker effects and cognitive test scores by the worker's highest educational category during the estimation interval (2007-2013). Panel (b) gives correlations between estimated AKM worker effects and cognitive test scores by workers' most frequent (modal) type of occupation during the estimation interval (2007-2013).

erogeneity between industries. It is noteworthy though that the AKM-COG correlation conditional on industry is higher than that conditional on education or occupation. Figure 3.4 Panel (b) presents region-specific correlations between AKM worker effects and COG. Workers moving between counties during the AKM estimation interval are again assigned to the modal county. There is little variation across counties in AKM worker effect prediction performance, with the AKM-COG correlation ranging between .350 and .435 for all but two counties. AKM-COG correlations conditional on place of work are thus similar to the unconditional correlation.

3.4.3 Bias of AKM worker effects

We now ask whether AKM worker effects are simply a noisier proxy for COG in some circumstances than in others or whether AKM worker effects systematically under- or overestimate COG in a way that is related to characteristics we observe. AKM prediction error, our bias measure, subtracts standardized cognitive test scores from standardized estimated AKM worker effects. It is roughly normally distributed and centered around zero, as Figure 3.15 in the appendix shows.

Worker ability We first regress AKM prediction error (for COG) on COG and non-cognitive test scores (NONCOG) without control variables.[28] Table 3.3 shows that, hold-

[28]NONCOG are also from the military draft, evaluated by a certified psychologist on the basis of a semi-structured 25-minute interview of the conscript. Designed to assess the individual's ability to cope with the

Figure 3.4: AKM prediction performance by employer characteristics

(a) Industry (SNI 92, sections) **(b)** Place of work (county)

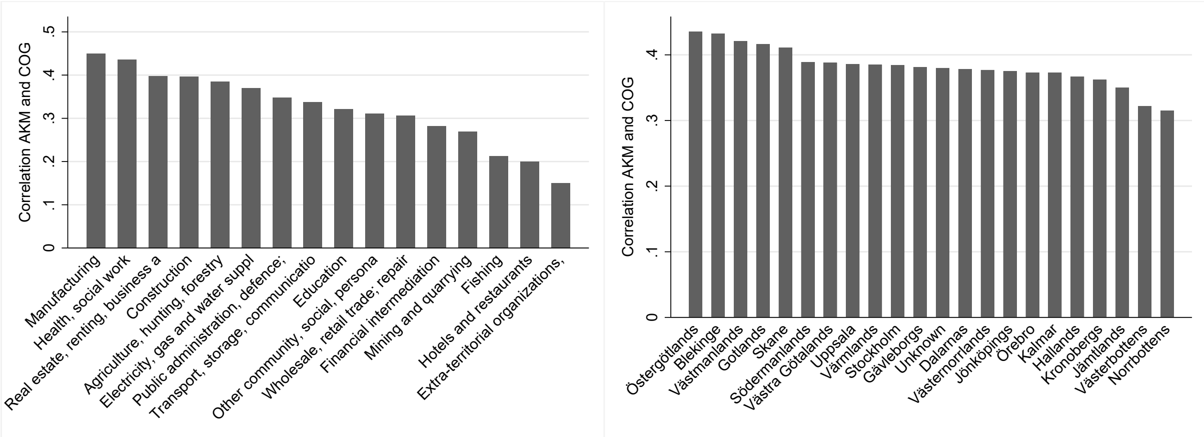

Note: Panel (a) shows correlations between estimated AKM worker effects and cognitive test scores by the worker's most frequent (modal) industry during the estimation interval (2007-2013). Panel (b) gives correlations between estimated AKM worker effects and cognitive test scores by workers' modal place of work (county) during the estimation interval (2007-2013).

ing NONCOG constant, a one standard deviation (SD) higher COG is associated with a prediction error that is .616 SDs lower. This means that AKM worker effects tend to overestimate COG in the lower half of the COG distribution and underestimate it in the upper half. This finding is consistent with Figure 3.1, where the fitted line has a slope smaller than one. That AKM worker effects "underestimate" COG for high levels of COG need not, however, imply that AKM worker effects are biased as a proxy for productivity more broadly. It would also be sufficient for this result that (i) COG does not perfectly determine productivity in the labor market and that (ii) productivity differences are not fully captured by wages, both of which seem plausible.

The coefficient on NONCOG is .187, less than a third as large in absolute magnitude. Its positive sign implies that, conditional on COG, AKM worker effects overestimate COG more the higher the level of NONCOG. This result can be consistent with the positive correlation between COG and NONCOG (r=.362 in our sample) if COG and NONCOG are also positively associated with AKM worker effects. That AKM worker effects show attenuation bias as a cognitive ability proxy may reflect bias with respect to (more readily) observed characteristics, too. For which workers at what type of firms do AKM worker effects over- or underestimate cognitive ability? To address this question we regress AKM prediction error on age, education, occupation, movement between firms, industry and

psychological requirements of military service and war, the test score reflects the following character traits: willingness to assume responsibility, independence, an outgoing character, persistence, emotional stability, and power of initiative (Lindqvist and Vestman, 2011).

Table 3.3: AKM bias with respect to worker ability

	AKM prediction error
Military draft cognitive test score (std)	-0.6162***
	(0.0010)
Military draft non-cognitive test score (std)	0.1870***
	(0.0010)
Observations	803,822
Adjusted R^2	0.3313

Note: *** p<0.01, ** p<0.05, * p<0.1, with heteroskedasticity-robust standard errors. Dependent variable is AKM prediction error, i.e., AKM worker effects minus COG. AKM worker effects estimated 2007-2013. The regression includes no control variables.

place of work. The results are given by Figures 3.5-3.7, which report different groups of OLS coefficients from the same regression. For comparison, they have the same x-axis range.

Worker characteristics Figure 3.5 reports coefficients on the year-of-birth dummies (with the 1974 birth cohort serving as reference category). The estimates reveal that, conditional on the other observables, bias is positively and monotonically related to age: AKM prediction error is nearly 1 standard deviation greater for the oldest workers (born 1951) than for the youngest (born 1992). The primary reason for this is likely to be that wages tend to increase with age (while our cognitive ability measure is time-invariant) and that the specification in Equation (3.1) cannot account for worker age because it includes year dummies.

Figure 3.5: AKM bias with respect to worker cohort

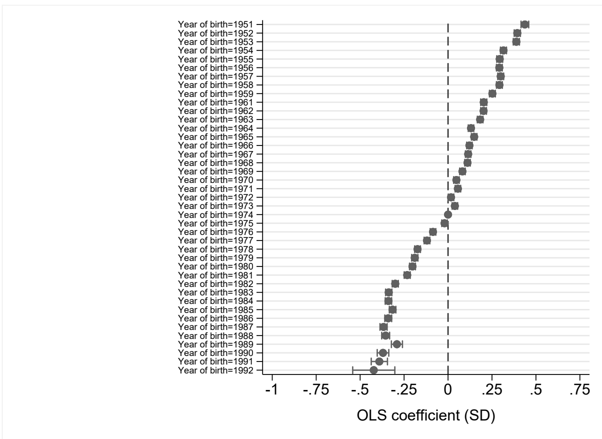

Note: This figure shows estimates from a regression of standardized AKM prediction error (std. AKM worker effects-std. COG) on worker and firm characteristics. The OLS coefficients shown are for the worker's year of birth. 95% confidence intervals are displayed (based on heteroskedasticity-robust standard errors).

Figure 3.6 shows point estimates for education categories, occupation type and for whether the worker switched firms. The estimated coefficients suggest, first, that bias is negatively related to the level of schooling. That is, AKM worker effects overestimate COG more for workers with lower education than for workers with higher education. Given that average COG is higher at higher levels of schooling[29] this is consistent with the finding above that AKM prediction error is negatively related to cognitive ability. The magnitude of the association between education and AKM prediction error is sizeable: compared to workers with some college, a college degree (reference category) or a doctorate, AKM prediction error is a quarter to a third of a standard deviation greater for those with completed or some high school, and one half to two thirds of a standard deviation higher for those with only compulsory schooling or less. Note that this is conditional on other observables. For instance, less than compulsory schooling will mostly be found among older workers but this positive link between age and bias is not captured by the education coefficients.

Figure 3.6: AKM bias with respect to education, occupation and movement between firms

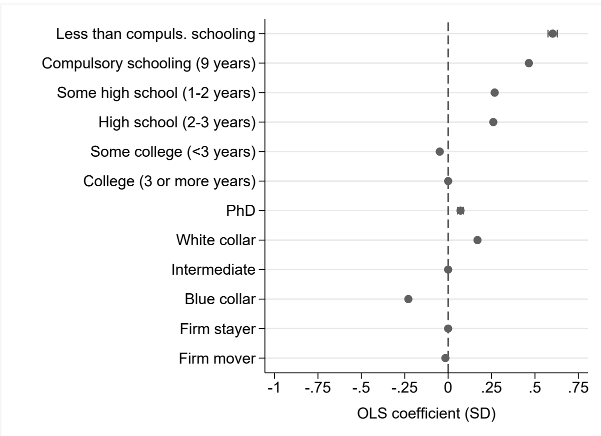

Note: This figure shows estimates from a regression of standardized AKM prediction error (std. AKM worker effects- std. COG) on worker and firm characteristics. The OLS coefficients shown are for educational categories, type of occupation and a dummy for movement between firms during the estimation period (2007-2013). 95% confidence intervals are displayed (based on heteroskedasticity-robust standard errors).

Figure 3.6 further shows that, compared to intermediate occupation types, AKM worker effects overestimate COG for white-collar workers and underestimates it for blue-collar workers. The gap is large, with AKM prediction error about two fifths of a standard

[29]The correlation between COG and a rough measure for years of schooling (imputed from the highest educational qualification) is .486.

deviation larger for white- than for blue-collar workers. This is an interesting result in that it goes in the opposite direction than the findings for education. While wages (the central ingredient of AKM worker effects) seem to dampen differences in cognitive ability across educational categories (conditional on occupation type), they appear to exaggerate difference in cognitive ability across occupation types (holding education constant). Finally, Figure 3.6 shows that AKM bias is negligibly different between firm movers and firm stayers. This is notable in that endogenous mobility ruled out by the AKM method has been identified as a potential source of bias (e.g., Abowd et al., 2018; Bonhomme et al., 2018). As we do not distinguish different types of mobility we cannot rule out that AKM worker effects systematically over- or underestimate cognitive ability for specific types of movers. However, our results do suggest that while the negative AKM bias for firm movers is statistically significant, the point estimate is very small compared to other sources of bias.

Figure 3.7: AKM bias with respect to employer characteristics

(a) Industry (SNI 92, sections) **(b)** Place of work (county)

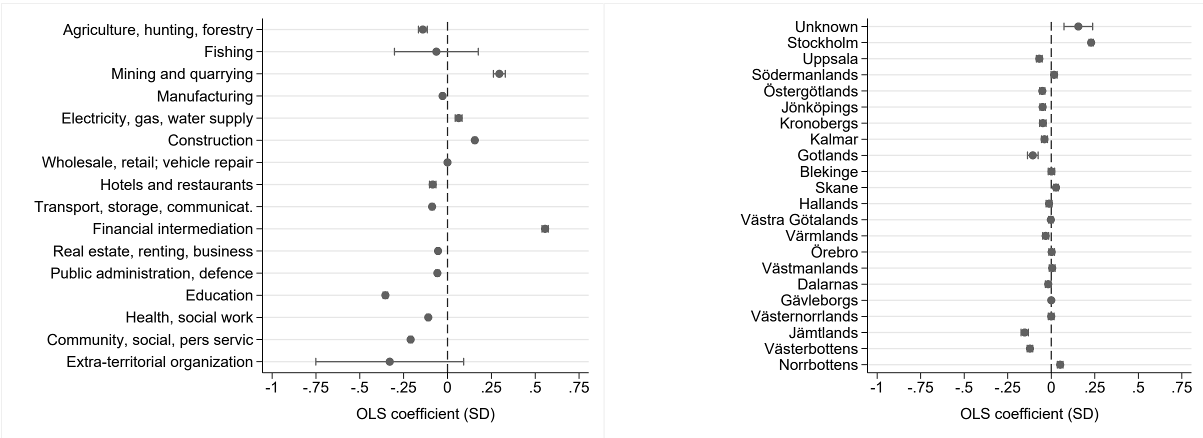

Note: Both panels show estimates from the same regression of standardized AKM prediction error (std. AKM worker effects-std. COG) on worker and firm characteristics (AKM worker effects estimated 2007-2013). They include 95% confidence intervals based on heteroskedasticity-robust standard errors. Panel (a) reports OLS coefficients for the worker's modal industry during the estimation period. Panel (b) reports OLS coefficients for the worker's modal place of work during the estimation period.

Firm characteristics Figure 3.7 Panel (a) reports coefficient estimates for the industry dummies. Wholesale and retail trade/repair of vehicles and household goods serves as the reference category. Though there is some heterogeneity across industries in AKM bias this heterogeneity lacks an obvious pattern. The three industries where AKM prediction error is largest in absolute magnitude are mining, education and financial intermediation. For mining, AKM prediction error is .296 standard deviations greater than for the reference category, for financial intermediation it is .557 standard deviations greater and for education it is .354 standard deviations smaller.[30] Heterogeneity in AKM bias by employer location is relatively limited other than for Stockholm, where prediction error is nearly a quarter of a standard deviation larger than in the reference county (Gävleborg)—see Figure 3.7 Panel (b).

3.4.4 Determinants of AKM worker effects

So far we have been exploring the validity of AKM worker effects as a cognitive ability proxy. We have shown that AKM worker effect estimates do correlate with cognitive ability, though this link displays both heterogeneity and biases. In this section we widen the focus and explore the determinants of AKM worker effects more generally. To study the empirical contents of estimated AKM worker effects we simply regress them on the worker and firm characteristics also considered above. In doing so we want to not only know what information AKM worker effects contain about workers and their jobs but also learn about the relative importance of these determinants. Figures 3.8-3.10 present estimated OLS coefficients from a regression of AKM worker effects on worker age, education, occupation type and movement between firms as well as employer industry and location. Their x-axes have the same range for comparability.

Worker characteristics Figure 3.8 reports the coefficient estimates for birth cohort dummies. The reference year is 1973. Very simply put the graph shows that age is a major determinant of AKM worker effects. As discussed in Section 3.4.3, this makes sense given that Equation (3.1) cannot control for age. It is possible of course that some of the age component of AKM worker effects is a productivity increase as workers acquire human capital in the course of their career. However, it seems unlikely that age-related productivity gains make up much of the overall age effect: first, the age premium is flat or even decreasing for the youngest twelve cohorts. If productivity improvements were driving the age effect we would expect the premium to grow fastest in the early years. Second, the age premium keeps rising right up to our second oldest cohort whose workers are 61 in our last estimation year. We would not expect productivity to keep rising so quickly up until a few years before retirement. As a consequence we consider most of the age effect a statistical artefact of the AKM estimation procedure.

Figure 3.9 is the centerpiece of this section. It reports the coefficients on dummies for cognitive and non-cognitive ability quartiles, educational categories, occupation types and

[30]Academic economists may find it consoling that (i) their wages do not do justice to their cognitive ability and (ii) that the size of this understatement is second only to the size of the overstatement enjoyed by bankers.

Figure 3.8: Worker age as a conditional AKM determinant

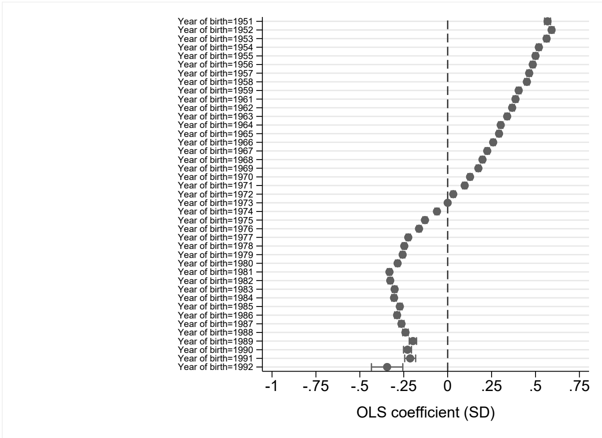

Note: This figure shows estimates from a regression of standardized AKM worker effects (estimated 2007-2013) on worker and firm characteristics. The OLS coefficients shown are for worker year of birth dummies. 95% confidence intervals are displayed (based on heteroskedasticity-robust standard errors).

movement between firms. With the exception of the firm mover dummy all of these worker characteristics are economically significant determinants of AKM worker effects. Educational categories are the strongest determinants, with the estimated average difference in AKM worker effects between the lowest-educated and PhD holders equal to 1.16 standard deviations. There appears to be little difference in AKM worker effects between workers with compulsory schooling and some high school (both around .26 SDs below the reference category). Similarly, there is only a small difference in AKM worker effects between workers with some post-secondary education (the reference category) and those with only high school. A regular college degree is associated with about a quarter of a SD higher AKM worker effects.

Among the worker characteristics shown in Figure 3.9 occupation types are second most important for explaining differences in AKM worker effects: white-collar workers' AKM worker effects are nearly .8 standard deviations higher than those of blue-collar workers, with intermediate occupations, the reference category, closer to white-collar workers. Next are non-cognitive abilities as proxied by the military draft test scores: the average difference in AKM worker effects between the top and bottom test score quartiles is .353 standard deviations. Finally cognitive test scores also explain a considerable amount of variation in AKM worker effects: the interquartile range is .225 SDs. Importantly, these are all ceteris-paribus statements. That is, for instance, the difference in AKM worker effects between the top and bottom cognitive test score quartiles of .225 SD is for a comparison of workers with similar levels of education and in the same type of occupation. This is likely to be part of the explanation for why non-cognitive test scores explain a bigger share of variation in AKM worker effects than cognitive test scores *when holding education constant*: cognitive

Figure 3.9: Ability, education, occupation and movement between firms as conditional AKM determinants

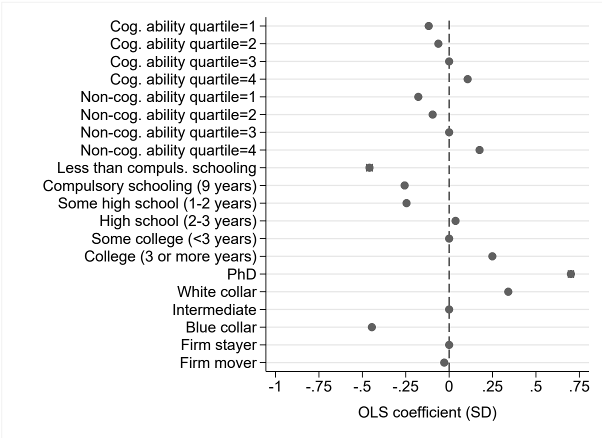

Note: This figure shows estimates from a regression of standardized AKM worker effects (estimated 2007-2013) on worker and firm characteristics. The OLS coefficients shown are for cognitive and non-cognitive ability quartiles, educational categories, type of occupation and a dummy for movement between firms during the AKM estimation period. 95% confidence intervals are displayed (based on heteroskedasticity-robust standard errors).

test scores vary more systematically with the level of education than do non-cognitive test scores.[31]

Firm characteristics Figure 3.10 presents coefficient estimates for firm characteristics. Panel (a) shows considerable heterogeneity in average AKM worker effects across industries with financial intermediation at the top and education at the bottom of a range of 1.07 standard deviations (extra-territorial organizations contains too few observations for reliable interpretation). AKM worker effect averages also vary across counties, though this variation is less pronounced (see Panel (b) of Figure 3.10). Stockholm stands out for AKM worker effects .307 standard deviations higher than Uppsala, the reference county, and Jämtland is at the bottom of the range at .185 standard deviations below Uppsala. The upshot from looking at firm-level determinants is that AKM worker effects contain considerable industry premia and non-negligible location premia. This is surprising: as firms do not frequently change industries or locations, we would have expected the firm fixed effects in the AKM specification in Equation (3.1) to absorb these differences. It is possible, though, that these premia are unequally distributed across jobs, workers or units within firms. One example of how this could come about is labor shortages in certain region-industry-occupation cells that drive up wages (and thus AKM worker effects) for only some workers in a firm so that the resulting wage premia are not fully captured by firm fixed effects.

[31] The correlation of a rough years of schooling measure is .486 with COG and .271 with NONCOG.

Figure 3.10: Employer characteristics as conditional AKM determinants

(a) Industry (SNI 92, sections) **(b)** Place of work (county)

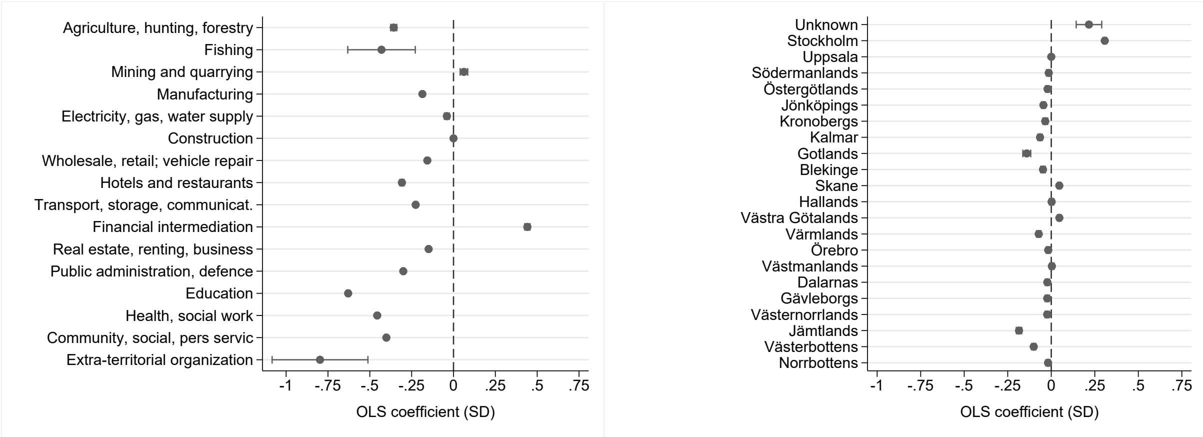

Note: Both panels show estimates from the same regression of standardized AKM worker effects on worker and employer characteristics (AKM worker effects estimated 2007-2013). They include 95% confidence intervals based on heteroskedasticity-robust standard errors. Panel (a) reports OLS coefficients for the worker's modal industry during the estimation period. Panel (b) reports OLS coefficients for the worker's modal place of work during the estimation period.

3.4.5 Size of correlation varies with estimation parameters

In this section we investigate how the correlation between AKM worker effects and COG depends on how the former are estimated.

Wage measure The first estimation ingredient we look at is the quality of the wage measure. Our main estimations use full time-equivalent (FTE) monthly wages. Many administrative data sources, while providing wage information, do not contain hours worked or the information necessary to infer FTE/hourly wages. How much the lack of hourly wage information undermines the ability of AKM worker effects to approximate cognitive test scores will depend on the degree of variation in labor supply and on the selection of workers into part-time or overtime. To quantify the importance of FTE or hourly wages for AKM worker effects as a cognitive ability proxy we compare two sets of AKM worker effect estimates, with each set estimated on the same sample: one set that use total monthly wages (unadjusted for hours) and one set that use FTE wages. We first do this for all workers and then separately for part-time and full-time workers.

Table 3.4 presents the results. As Column (1) shows, the AKM-COG correlation is nearly halved when hours worked are not adjusted for and all workers are included in the AKM estimation. We next consider the case when AKM worker effects are estimated only for full-time workers (defined as working at least 35 hours a week). The idea behind this is that some administrative data sources contain only rough hours information. For

Table 3.4: AKM prediction performance by availability of working hours info

Sample	(1) All workers	(2) FT workers (≥35 hrs.)	(3) PT workers (<35 hours)
r without hours info	.219	.394	.127
r with hours info	.398	.415	.354
Observations	881,490	756,078	213,218

Note: This table compares the correlation between estimated AKM worker effects (2007-2013) and cognitive test scores for different samples depending on whether information on working hours is used/available. Column (1) looks at all workers; (2) looks at full-time workers with 35 or more contractual weekly hours; (3) looks at part-time workers with less than 35 weekly hours.

instance, CHK use the information available in the German social security records to drop part-time workers and marginally employed workers and estimate their AKM specification only on full-time workers. Column (2) shows that the AKM-COG correlation for full-time workers is only marginally higher when hours-adjusted wage data are used, suggesting that limiting the AKM estimation to full-time workers is a valid strategy when this is possible. Column (3) looks at part-time workers. As expected, the AKM-COG correlation is reduced markedly when wages are not adjusted for hours worked, though the correlation remains positive.

Panel completeness The second estimation issue we address concerns panel completeness. Because our wage data come from a survey that covers only a subset of firms we have an unbalanced panel even for those workers that are actually continuously employed. We limit our attention to those workers who, according to the complete labor market records, are employed throughout our 7-year estimation window.[32] Among these workers there is still variation in the number of years during which they are observed in the FTE wage data and thus variation in the number of years on which our AKM worker effect estimate for them is based. Figure 3.11 explores how this affects the AKM-COG correlation. Pearson's r increases from .334 for workers observed for only 2 years to .450 for workers observed throughout the period.[33]

Estimation period Boom-and-bust cycles may affect the ability of AKM worker effects to capture individual ability. Also, CHK find that the correlation between estimated worker and firm effects changes over time in Germany, presumably reflecting structural changes in the labour market. Here we ask whether macro changes have an impact on how much information AKM worker effects contain about workers' cognitive ability. Table 3.5 compares the correlation between AKM worker effects and cognitive test scores across four

[32] Conditioning on continuous employment ensures that we do not capture selection into inactivity/unemployment or composition effects (such as workers entering the labor market partway through the estimation period).

[33] Half of workers are observed throughout, with the rest approximately evenly distributed across the shorter durations, see right y-axis of Figure 3.11.

Figure 3.11: AKM prediction performance by number of years worker is observed

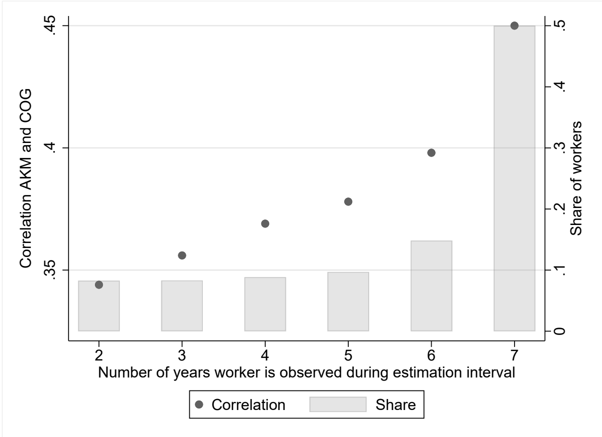

Note: This figure shows correlations between estimated AKM worker effects (from 2007-2013) and cognitive test scores by the number of years the worker is observed in the WSS data during the estimation interval.

seven-year intervals among workers aged 27-33. We look at the same age groups in each estimation interval to avoid capturing age effects. The correlations fluctuate between .37 and .41 across estimation intervals without a monotonic trend. The pattern may, however, reflect the business cycle: Sweden experienced a financial crisis with high unemployment between 1992 and 1998. Most of this falls into our second estimation interval. The Great Recession, contained in our final estimation interval, also affected Sweden, with unemployment peaking in 2009 and 2010. The AKM-COG correlation is lower in the crisis estimation intervals, but the difference is not large.

Table 3.5: AKM prediction performance by estimation time period

	(1) 1986-1992	(2) 1993-1999	(3) 2000-2006	(4) 2007-2013
r	.392	.382	.409	.372
Observations	159,432	158,562	198,140	165,957

Note: This table compares the correlation between AKM worker effects estimated in four different intervals and cognitive test scores. In each interval, the correlations are computed for workers aged 27-33.

Length of estimation interval Another choice (or constraint) faced by researchers applying the AKM method to approximate worker ability is the length of the panel used for AKM estimation. We use non-overlapping seven-year estimation intervals because this allows us to split our observation period into four windows of equal length. CHK, addressing a different question, choose overlapping intervals of seven to eight years. Here we look at how the length of the estimation interval affects the quality of AKM worker effects as an ability proxy. Figure 3.12 compares the AKM-COG correlation across AKM estimations that differ in the length of the underlying panel. To avoid composition effects, it only considers individuals who are observed throughout and who are of a similar age. The results suggest that estimation intervals longer than six years do not improve the performance of AKM worker effects as ability proxies: the AKM-COG correlation is between .38 and .39 for all estimation intervals that are at least six years long. Using a shorter estimation period does reduce the AKM-COG correlation, though it is still surprisingly high (over .33) for AKM worker effects estimated on a two-year panel.

Figure 3.12: AKM prediction performance by length of estimation interval

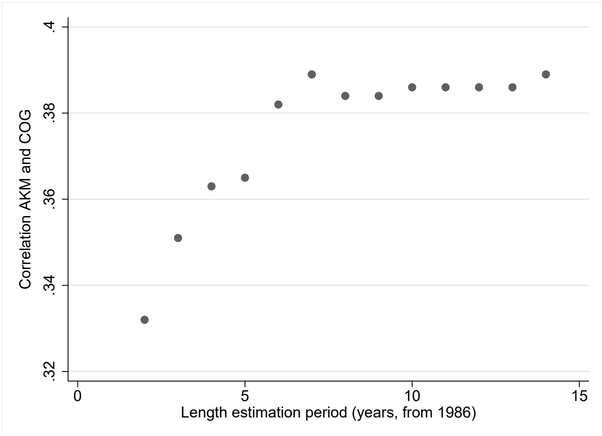

Note: This figure shows correlations between estimated AKM worker effects and cognitive test scores for estimation intervals of different lengths. AKM-COG correlations are computed only for those workers aged 27-33 in 1987 (birth cohorts 1954-1960) who are observed in the labor market throughout 1987-1999.

Deviating from CHK Finally we consider selected deviations from the AKM specification chosen by CHK and followed in this paper. We focus on two variations that may be of particular interest to applications seeking to use estimated AKM worker effects as proxies for individual productivity. We ask how the AKM-COG correlation is affected when (i) education is explicitly controlled for and (ii) when women are included in the AKM specification. Table 3.6 reports the results with Column (1) replicating the AKM-COG correlation for the AKM specification used throughout this paper. Column (2) confirms the finding from Section 3.4.4 that AKM worker effects capture variation in cognitive ability over and above between-education levels differences in COG: adding dummies for levels of education to the AKM estimation reduces the AKM-COG correlation only slightly. Whether it makes sense in applications to include education dummies in the AKM specification is another matter. It may, for instance, depend on whether the researcher wants to focus on differences in cognitive ability not reflected by the level of education or is looking for a very broad measure of human capital.

Column (3) suggests that including women into the AKM estimation has a negligible effect on the AKM-COG correlation. This means that estimating AKM worker effects for men and women jointly does not reduce the predictive performance of AKM worker effects for men. It does not, by itself, constitute evidence that AKM worker effects serve similarly well as cognitive ability proxies for women. We refrain from making AKM-COG comparisons for the small number of women for whom we have military test scores as these are likely to be a highly selected sample.

Table 3.6: AKM prediction performance—deviations from CHK

	(1) Baseline	(2) Education dummies	(3) Men and women
r	.398	.384	.394
Observations	881,490	881,490	882,633

Note: This table compares the correlation between estimated AKM worker effects (2007-2013) and cognitive test scores for different AKM specifications that compare to the specification in Equation (3.1) and CHK as follows: (1) is unchanged (used throughout the paper); (2) includes education category dummies; (3) includes women.

3.5 Conclusion

In this paper we provide an empirical assessment of AKM worker effects as an ability proxy. We use Swedish register data that includes both detailed labor market information, allowing us to estimate AKM regressions, and standardized test scores for cognitive ability at age 18. Our approach to assessing approximation quality is simple: we take COG as a measure of true ability and look at the correlation between AKM worker effects and COG. We find an AKM-COG correlation of about .4 and use machine-learning algorithms trained to predict COG as a benchmark. They outperform AKM worker effects, but AKM worker effects achieve over three quarters of their prediction performance. We are inclined to conclude from this that overall AKM worker effects *are* a valid proxy for individual ability.

The remainder of the paper explores qualifications to this assessment. We first uncover substantial heterogeneity in the prediction performance of AKM worker effects: they "work" better for older workers, white-collar workers and certain industries. Our attempt to discern whether the prediction deficit for some groups reflects systematic mis-estimation or noise suggests that there are biases: for example, AKM worker effects underestimate the cognitive ability of young workers, blue-collar workers and highly educated workers.

By contrast, we do not find evidence for mobility bias on average, though we cannot directly investigate AKM worker effects' potential biases associated with specific patterns of endogenous mobility. As a consequence the lack of an overall mobility bias may well mask countervailing biases for sub-groups of firm movers. When we open the AKM worker effect blackbox and regress them on a set of worker and firm characteristics we find that on the one hand, AKM worker effects contain time-invariant skill-related attributes one would like them to—such as the level of education, cognitive ability and non-cognitive ability. On the other hand, they also contain potentially time-varying features less closely related to individual skill, such as age or industry; while the age premium is a plausible mechanical consequence of the estimation strategy, one might have expected firm fixed effects to absorb most industry differences in pay.

One recommendation arising from these findings is the following: when AKM worker effects suggest ability differences, it is important to rule out that these are driven by some of the characteristics shown here to bias them, such as age or occupation type. The final sub-section of the paper offers further practical implications; it is explicitly addressed to researchers wishing to use AKM worker effects as ability proxies. Our analyses show that having information on either hourly pay or full-time/part-time status is key for reasonable prediction performance; that longer estimation intervals are better than shorter ones but in a strongly concave way; that balanced panels are superior but unbalanced ones are not much worse; that in Sweden, the prediction performance of AKM worker effects does not vary much across four non-overlapping estimation intervals; and that estimating AKM worker effects jointly for men and women does not perceptibly reduce their prediction performance for men.

We hope these findings will provide some input to the methodological debate on the relationship between AKM worker effects and underlying worker productivity. We also hope that they offer useful guidance on the circumstances in which AKM worker effects are reasonable ability proxies.

References

Abowd, John M, Robert H Creecy, Francis Kramarz, et al. (2002): "Computing person and firm effects using linked longitudinal employer-employee data," Tech. rep., Center for Economic Studies, US Census Bureau.

Abowd, John M, Francis Kramarz, Paul Lengermann, Kevin L McKinney, and Sébastien Roux (2012): "Persistent inter-industry wage differences: rent sharing and opportunity costs," *IZA Journal of Labor Economics*, 1(1): 1–25.

Abowd, John M, Francis Kramarz, Paul Lengermann, and Sébastien Pérez-Duarte (2004): "Are good workers employed by good firms? A test of a simple assortative matching model for France and the United States," *Unpublished Manuscript*.

Abowd, John M, Francis Kramarz, and David N Margolis (1999): "High wage workers and high wage firms," *Econometrica*, 67(2): 251–333.

Abowd, John M, Kevin L McKinney, and Ian M Schmutte (2018): "Modeling endogenous mobility in earnings determination," *Journal of Business & Economic Statistics*, 1–14.

Alvarez, Jorge, Felipe Benguria, Niklas Engbom, and Christian Moser (2018): "Firms and the decline in earnings inequality in Brazil," *American Economic Journal: Macroeconomics*, 10(1): 149–89.

Andrews, Martyn J, Len Gill, Thorsten Schank, and Richard Upward (2008): "High wage workers and low wage firms: negative assortative matching or limited mobility bias?" *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 171(3): 673–697.

Andrews, Martyn J, Leonard Gill, Thorsten Schank, and Richard Upward (2012): "High wage workers match with high wage firms: Clear evidence of the effects of limited mobility bias," *Economics Letters*, 117(3): 824–827.

Bagger, Jesper, Kenneth L Sørensen, and Rune Vejlin (2013): "Wage sorting trends," *Economics Letters*, 118(1): 63–67.

Balsvik, Ragnhild and Stefanie Haller (2015): "Ownership Change and its Implications for the Match between the Plant and its Workers." *NHH Dept. of Economics Discussion Paper*, (10).

Barth, Erling, James Davis, and Richard B Freeman (2018): "Augmenting the human capital earnings equation with measures of where people work," *Journal of Labor Economics*, 36(S1): S71–S97.

Bender, Stefan, Nicholas Bloom, David Card, John Van Reenen, and Stefanie Wolter (2018): "Management practices, workforce selection, and productivity," *Journal of Labor Economics*, 36(S1): S371–S409.

Bombardini, Matilde, Gianluca Orefice, and Maria D Tito (2019): "Does exporting improve matching? Evidence from French employer-employee data," *Journal of International Economics*, 117: 229–241.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa (2018): "A distributional framework for matched employer employee data," *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2019-15).

Borovičková, Katarína and Robert Shimer (2017): "High wage workers work for high wage firms," Tech. rep., National Bureau of Economic Research.

Butschek, Sebastian (2017): "Raising the bar: the effect of wage cost shocks on worker selection," unpublished manuscript, University of Cologne.

Butschek, Sebastian and Jan Sauermann (2019a): "The effect of employment protection legislation on firms' worker selection," unpublished manuscript, Stockholm University.

Card, David, Ana Rute Cardoso, and Patrick Kline (2015): "Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women," *Quarterly Journal of Economics*, 131(2): 633–686.

Card, David, Jörg Heining, and Patrick Kline (2013): "Workplace heterogeneity and the rise of West German wage inequality," *The Quarterly Journal of Economics*, 128(3): 967–1015.

Carlsson, Mikael, Julián Messina, and Oskar Nordström Skans (2016): "Wage adjustment and productivity shocks," *Economic Journal*, 126(595): 1739–1773.

Cornelissen, Thomas, Christian Dustmann, and Uta Schönberg (2017): "Peer Effects in the Workplace," *American Economic Review*, 107(2): 425–56.

Correia, Sergio (2016): "Linear Models with High-Dimensional Fixed Effects: An Efficient and Feasible Estimator," Tech. rep., working Paper.

Davidson, Carl, Fredrik Heyman, Steven Matusz, Fredrik Sjöholm, and Susan Chun Zhu (2014): "Globalization and imperfect labor market sorting," *Journal of International Economics*, 94(2): 177–194.

Eeckhout, Jan and Philipp Kircher (2011): "Identifying sorting—in theory," *Review of Economic Studies*, 78(3): 872–906.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams (2016): "Sources of geographic variation in health care: Evidence from patient migration," *Quarterly Journal of Economics*, 131(4): 1681–1726.

Fox, Jeremy T. and Valerie Smeets (2011): "Does Input Quality Drive Measured Differences In Firm Productivity?" *International Economic Review*, 52(4): 961–989.

Fredriksson, Peter, Lena Hensvik, and Oskar Nordström Skans (2018): "Mismatch of Talent Evidence on Match Quality, Entry Wages, and Job Mobility," *American Economic Review*, forthcoming.

Gathmann, Christina, Ines Helm, and Uta Schönberg (2018): "Spillover effects of mass layoffs," *Journal of the European Economic Association.*

Gruetter, Max and Rafael Lalive (2009): "The importance of firms in wage determination," *Labour Economics*, 16(2): 149–160.

Håkanson, Christina, Erik Lindqvist, and Jonas Vlachos (2015): "Firms and skills: the evolution of worker sorting," Working Paper Series 2015:9, IFAU Institute for Evaluation of Labour Market and Education Policy.

Hagedorn, Marcus, Tzuo Hann Law, and Iourii Manovskii (2017): "Identifying equilibrium models of labor market sorting," *Econometrica*, 85(1): 29–65.

Iranzo, Susana, Fabiano Schivardi, and Elisa Tosetti (2008): "Skill dispersion and firm productivity: An analysis with employer-employee matched data," *Journal of Labor Economics*, 26(2): 247–285.

James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani (2013): *An introduction to statistical learning*, vol. 112, Springer.

Krishna, Pravin, Jennifer P Poole, and Mine Zeynep Senses (2014): "Wage effects of trade reform with endogenous worker mobility," *Journal of International Economics*, 93(2): 239–252.

Lindqvist, Erik and Roine Vestman (2011): "The Labor Market Returns to Cognitive and Noncognitive Ability: Evidence from the Swedish Enlistment," *American Economic Journal: Applied Economics*, 3(1): 101–128.

Lopes de Melo, Rafael (2018): "Firm wage differentials and labor market sorting: Reconciling theory and evidence," *Journal of Political Economy*, 126(1): 313–346.

Macis, Mario and Fabiano Schivardi (2016): "Exports and wages: Rent sharing, workforce composition, or returns to skills?" *Journal of Labor Economics*, 34(4): 945–978.

Roca, Jorge De La and Diego Puga (2017): "Learning by working in big cities," *Review of Economic Studies*, 84(1): 106–142.

Rothstein, Jesse (2010): "Teacher quality in educational production: Tracking, decay, and student achievement," *Quarterly Journal of Economics*, 125(1): 175–214.

Schmutte, Ian M (2015): "Job referral networks and the determination of earnings in local labor markets," *Journal of Labor Economics*, 33(1): 1–32.

Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter (2018): "Firming up inequality," *Quarterly Journal of Economics*, 134(1): 1–50.

Appendix

Figure 3.13: Worker age distribution by estimation interval

(a) Male workers with AKM **(b)** Male workers with both COG and AKM

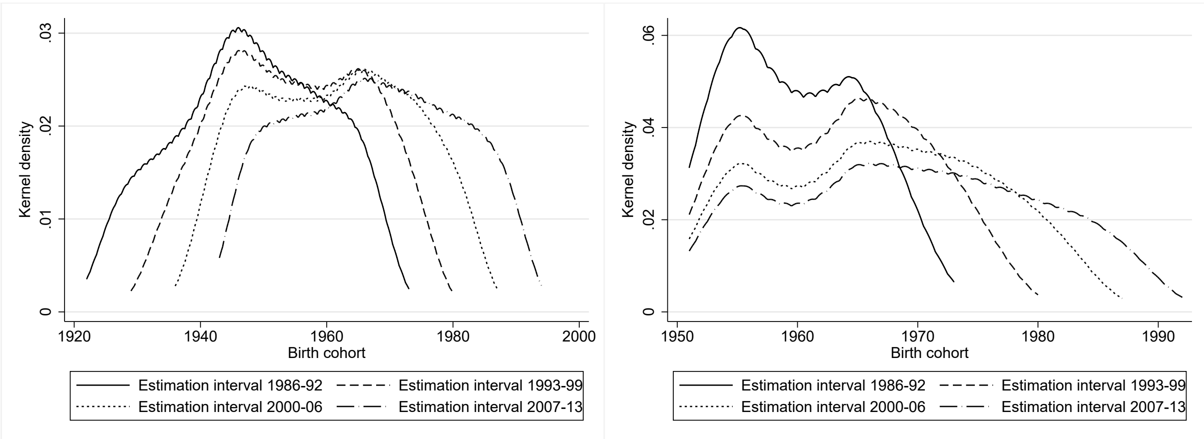

Note: Both panels show estimated kernel densities for the age distribution within each AKM estimation interval. Panel (a) shows the age distributions of those workers with an estimated AKM worker effect and Panel (b) gives the age distributions of the subset of workers for whom both a cognitive military test score and an estimated AKM worker effect is available.

Table 3.7: Correlations between estimated AKM worker effects and cognitive ability

	(1) Pearson's r	(2) Spearman's ρ
AKM and COG	.398	.396
Observations	881,490	881,490

Note: Estimated AKM worker effects and cognitive test scores are standardized to mean zero and unit variance. The columns give different sample correlation coefficients: (1) is Pearson's r; (2) is Spearman's ρ.

Figure 3.14: Cognitive test score coverage

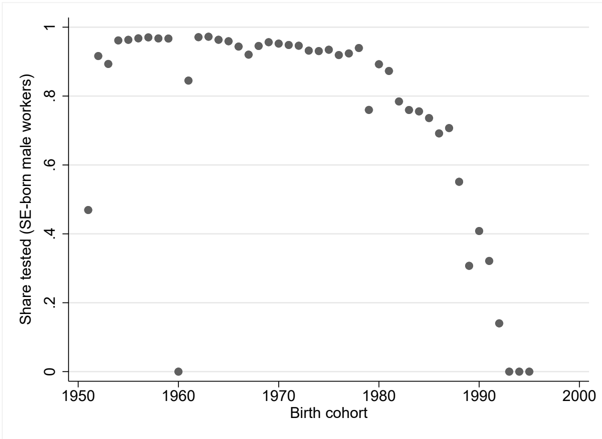

Note: This figure gives the share of Swedish-born males in the 2007-2013 WSS data for whom cognitive test scores are available.

Table 3.8: Benchmarks for AKM worker effects as ability proxies

	(1) AKM	(2) Boosting	(3) Bagging	(4) Random forest
Validation set MSE	-	.639	646	.653
r with COG	.450	.577	.576	.580
Observations	356,958	356,958	356,958	356,958

Note: Predictions and cognitive test scores have been standardized to mean zero and unit variance. r is their correlation coefficient. The columns compare the prediction performance of AKM worker effects to those of three tree-based machine-learning algorithms trained on 2000-2006 data to predict cognitive test scores. Columns (1)-(4) all use the same sample of male workers observed throughout 2007-13. (1) uses the AKM worker effects also used in Figure 3.1 and Table 3.1; (2) is from a gradient boosting model; (3) is from a bagging model; (4) is from arandom forest model.

Figure 3.15: AKM prediction error distribution

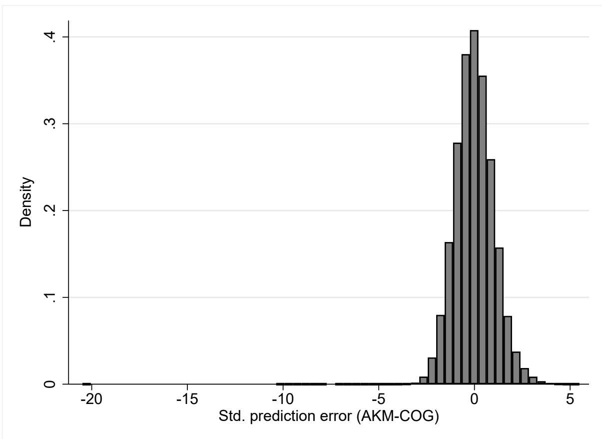

Note: This figure shows a histogram for AKM prediction error, defined as standardized AKM worker effects (from 2007-2013) minus standardized cognitive test scores.

Table 3.9: Summary statistics for AKM and COG, by subgroups

	AKM worker effects		Cognitive test scores	
	mean	sd	mean	sd
All workers	0.000	1.000	-0.000	1.000
20-26 years	-0.424	0.555	-0.021	0.993
27-33 years	-0.167	0.799	0.008	0.994
34-40 years	0.048	0.992	0.013	1.000
41-47 years	0.160	1.104	0.018	1.012
48-54 years	0.303	1.165	0.008	1.004
55-56 years	0.383	1.183	-0.040	1.006
White collar	0.622	1.027	0.476	0.872
Intermediate	0.176	1.012	0.128	0.874
Blue collar	-0.592	0.488	-0.453	0.926
Less than compulsory schooling	-0.546	0.603	-1.129	0.839
Compulsory schooling only (9 years)	-0.488	0.684	-0.786	0.894
Some high school (1-2 years)	-0.366	0.725	-0.489	0.855
High school (2-3 years)	-0.286	0.805	-0.224	0.884
Some college (¡3 years)	0.194	0.980	0.342	0.851
College (3 or more years)	0.534	1.114	0.605	0.842
PhD	1.167	1.127	1.101	0.739
Firm stayer	0.016	1.016	-0.020	1.004
Firm mover	-0.034	0.965	0.043	0.989
Observations (all)	881,490		881,490	

Note: This table provides summary statistics for AKM worker effects and military test scores for cognitive ability, separately for sub-groups defined by observed characteristics.

Table 3.10: Data sources

Data source	Period used	Description	Used for measurement of	Variables used/generated
JOBB	1986-2013	Employment spells	Employment	Firm ID
JOBB	1986-2013	Employment spells	Worker characteristics	Occupation type
JOBB	1986-2013	Employment spells	Firm characteristics	Industry, age, location
WSS	1986-2013	Wage survey statistics	Wages	Full-time equivalent monthly wages
LISA	1990-2013	Various information	Worker characteristics	Highest educational qualification
BAKGRUND	n.a.	Birth register	Worker characteristics	Year of birth, gender
KRIGSARKIVET	1970-2004	Military draft information	Worker ability	Cognitive, psychological test scores

Conclusion

Chapters 1 and 2 of this dissertation attempt to increase our understanding of the role of labor market regulation in shaping firms' hiring behaviour. Their contribution is to provide some of the first evidence on how minimum wages and employment protection affect *which* workers get hired. The direct conclusion emerging from both studies is that firms do adjust the selectiveness of their hiring to these policy changes: they become less willing to employ less able workers when minimum wages drive up labor costs at the bottom of the wage distribution and more willing to do so when dismissal costs sink. Chapter 3 strengthens this conclusion by confirming that estimated AKM worker effects, the productivity proxy used in Chapters 1 and 2, is informative about workers' cognitive ability. This final chapter of the dissertation discusses some limitations of Chapters 1-3 and attempts to synthesize their implications.

Limitations One weakness of Chapter 1 is that my sample includes only establishments with 50 or more workers. As smaller employers tend to pay lower wages, they are an important group to consider in studying the minimum wage. Having nothing to say on smaller firms thus undermines the representativeness of my findings. A disadvantage of relying on pre-determined AKM worker effects as an ability proxy is that I do not have a productivity measure for all new hires. In particular, workers too young to have held a job in the estimation interval (2002-2009) do not enter my analysis of the minimum wage's effect on the selectiveness of hiring. Again, this implies that I partly omit from my effect estimates a group for which the minimum wage is important. Another limitation concerns my minimum wage affectedness measure: in my main analysis, I approximate hourly wages by looking only at full-time workers. This means that part-time workers and mini-jobbers do not enter my classification of establishments as affected or unaffected by the future minimum wage. While this is likely to bias my results toward zero and I can broadly confirm my findings using additional hours data, my classification of employers' exposure to the minimum wage remains imperfect. Finally, I am unable to completely rule out that systematic self-selection of workers causes me to overestimate the effect of the minimum wage on employers' worker selection.

In Chapter 2, larger and richer data as well as a different quasi experiment allow us to overcome some of Chapter 1's limitations. For example, we observe firms of all sizes, we need not rely exclusively on AKM worker effects as an ability measure (but have several productivity proxies that all point in the same direction), and the data allow us to precisely measure exposure to the employment protection reform. A limitation Chapters 1 and 2 have in common, however, is that while we present indicative evidence that worker self-selection plays no significant role, we are unable to conclusively reject the possibility it biases our estimates (though the direction is unclear in Chapter 2). Unlike in Chapter 1, a mechanical effect of the reform on hire quality through the number of hires makes up a significant fraction of the overall effect in Chapter 2. Both a simulation exercise we use to purge our estimates of the mechanical effect and our analysis of new hires' links with incumbents from past employment suggest there is a substantial direct effect of employment protection on firms' worker selection; it is difficult, however, to make precise statements about the size of this direct effect.

Chapter 3's key limitation is that we are only able to say something about the performance of AKM worker effects as an ability proxy for men as nearly all workers who underwent military draft testing are male. A further limitation of Chapter 3 is that our cognitive ability measure is far from capturing everything that makes a worker productive. This is true in terms of scope—traits as diverse as communication skills, physical strength, charisma, patience or conscientiousness may contribute to individual productivity. It is also true in terms of time and measurement error—we do not have true contemporaneous cognitive ability but a noisy test score from age 18. Yet, it is difficult to imagine data better suited to our purpose: maybe one would prefer incentivized standardized tests along multiple dimensions that are repeated at regular intervals. However, such a measure is unlikely to ever become available for a sufficiently large sample of workers.

Policy implications The findings of Chapters 1 and 2 strongly suggest that labor market regulation affects employers' hiring standards. The implications of this for policy-makers have already been discussed briefly in the Chapters' respective conclusions: there may be distributional consequences of minimum wages and employment protection that run counter to the policies' intended effect.

To the extent that a minimum wage is supposed to boost the income of low-wage workers, for instance, it may have an unintended side effect by making it harder for the least productive to find a job. This implication is different from the old contention that minimum wages may destroy low-wage jobs. The difference is that my findings in Chapter 1 suggest that regardless of whether there are net employment effects, low-wage jobs may be re-distributed away from the least productive workers, increasing the likelihood of being unemployed for those with the lowest earnings power. Fairris and Bujanda (2008) have aptly described this as dissipation of the minimum wage-induced gains for low earners.

The claim that employment protection may benefit insiders (incumbents) but hurt outsiders (job seekers) has also been around for a long time. Indeed, one study has already drawn attention to the distributional aspect of dismissal costs: Kugler and Saint-Paul (2004) show that an increase in employment protection favors employed job-seekers against the unemployed. Our findings from Chapter 2 provide evidence on an additional distributional dimension: by affecting firms' worker selection, employment protection may make it likelier that the lowest-ability workers end up as the outsiders, stacking the distribution of jobs and job security against them.

What then could policy-makers do with this information? It would be wrong-headed to ignore or disproportionately discount the distributional benefits of minimum wages and employment protection on the basis that, even without aggregate employment losses, there might be losers from these regulations. Instead, public employment services could allow for the possibility that it is actually increasingly difficult for those ending up in unemployment to be matched with jobs—partly because minimum wages and employment protection raise hiring thresholds. This realization may (further) shift the focus of active labor market policies for unemployed workers with low productivity from instruments targeting the matching process (such as job search and application coaching) to interventions attempting to improve workers' productive capabilities (such as training or re-training).

Potential future research Based on the evidence presented in this dissertation and available elsewhere, the policy implications discussed above remain somewhat speculative. One important step toward providing a more solid evidence base would be to directly test whether minimum wages and employment protection cause the pool of unemployed low-skill workers to become more negatively selected.

Another follow-up analysis putting the policy implications on a firmer footing would be to relax the implicit assumption of Chapters 1 and 2 that workers' self-selection into firms remained unchanged by the reforms. This assumption makes it possible to infer a change in firms' hiring thresholds from the equilibrium outcomes of labor supply and demand. I see two possible ways of relaxing it. One is working with data on job requirements from posted vacancies as in Deming and Kahn (2018). In the course of reviewing the chapters for inclusion in this dissertation I realized that there is an independent and contemporanous working paper, Clemens et al. (2018), that does just that and confirms the positive effect of minimum wages on firms' demand for the quality of labor.

A second way would be to pursue a similar strategy as Chapters 1 and 2 (identifying hiring standards from actual worker ability) but explicitly studying the ability of applicants, not of realized hires. Presumably, this would only be feasible with company personnel records. Obtaining such data from enough firms that there is variation in the exposure to a policy change in question appears difficult. A more promising approach might be to work with records from an HR service provider that recruits workers for other employers.

Chapter 3 does not have any direct policy implications – its target audience are applied economists in search of valid individual ability proxies. There are, however, worthwhile ways in which the analysis could be extended in future. In particular, it would be instructive to also assess the performance as ability proxies of the alternatives to estimated AKM worker effects proposed in the recent literature (e.g., Hagedorn et al., 2017; Abowd et al., 2018; Bonhomme et al., 2018).

References

Abowd, John M, Kevin L McKinney, and Ian M Schmutte (2018): "Modeling endogenous mobility in earnings determination," *Journal of Business & Economic Statistics*, 1–14.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa (2018): "A distributional framework for matched employer employee data," *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2019-15).

Clemens, Jeffrey, Lisa Kahn, and Jonathan Meer (2018): "Dropouts need not apply: The minimum wage and skill upgrading," Tech. rep., Working Paper.

Deming, David and Lisa B. Kahn (2018): "Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals," *Journal of Labor Economics*, 36(S1): S337–S369.

Fairris, David and Leon Fernandez Bujanda (2008): "The dissipation of minimum wage gains for workers through labor-labor substitution: evidence from the Los Angeles living wage ordinance," *Southern Economic Journal*, 473–496.

Hagedorn, Marcus, Tzuo Hann Law, and Iourii Manovskii (2017): "Identifying equilibrium models of labor market sorting," *Econometrica*, 85(1): 29–65.

Kugler, Adriana D. and Gilles Saint-Paul (2004): "How do firing costs affect worker flows in a world with adverse selection?" *Journal of Labor Economics*, 22(3): 553–584.

Appendix to dissertation

Declaration of contribution to the chapters of the cumulative dissertation

My co-author Jan Sauermann and I estimated my respective contributions to Chapters 2 and 3 together.

Raising the bar: minimum wages and employers' hiring standards (Chapter 1)

Author: Sebastian Butschek (single author)

Contribution:

Funding and data acquisition:	5%
Idea and study design:	100%
Data preparation and analysis:	100%
Drafting and development of manuscript:	100%

The effect of employment protection on firms' worker selection (Chapter 2)

Authors: Sebastian Butschek and Jan Sauermann

Contribution:

Funding and data acquisition:	0%
Idea and study design:	75%
Data preparation and analysis:	50%
Drafting and development of manuscript:	66%

Do estimated individual effects capture worker ability? (Chapter 3)

Authors: Sebastian Butschek and Jan Sauermann

Contribution:

Funding and data acquisition:	0%
Idea and study design:	75%
Data preparation and analysis:	60%
Drafting and development of manuscript:	80%

Curriculum vitae

Name:	Sebastian Butschek
Birth:	7 March 1985, Innsbruck, Austria
Secondary school:	IB Diploma, Lester B. Pearson UWC, Victoria, Canada, 2003
Diploma:	Physiotherapy, AZW, Innsbruck, Austria, 2008
Bachelor:	BA History and Economics, University of Oxford, UK, 2011
Master:	MSc Economics, BGSE, Barcelona, Spain, 2012
Professional:	Researcher, ZEW, Mannheim, Germany, 2012-2016
Doctorate:	Universität zu Köln, Cologne, Germany, expected: 2019