

Hysteresis from Employer Subsidies*

By Emmanuel Saez, Benjamin Schoefer, and David Seim

May 24, 2021

Abstract

This paper uses administrative data to analyze a large and 8-year long employer payroll tax rate cut in Sweden for young workers aged 26 or less. We replicate previous results documenting that during the earlier years of the reform, it raised youth employment among the treated workers, driven by labor demand (as workers' take-home wages did not respond). First, drawing on additional years of data, this paper then documents that the longer-run effects during the reform are twice as large as the medium-run effects. Second, we document novel labor-demand-driven "hysteresis" from this policy – i.e. persistent employment effects even after the subsidy no longer applies – along two dimensions. *Over the lifecycle*, employment effects persist even after workers age out of eligibility. Three years *after the repeal*, employment remains elevated at the maximal reform level in the formerly subsidized ages. These hysteresis effects more than double the direct employment effects of the reform. Discrimination against young workers in job posting fell during the reform and does not bounce back after repeal, potentially explaining our results.

*Saez: University of California, 530 Evans Hall #3880, Berkeley, CA 94720 (e-mail: saez@econ.berkeley.edu); Schoefer: University of California, 530 Evans Hall #3880, Berkeley, CA 94720 (e-mail: schoefer@berkeley.edu); Seim: Stockholm University, Universitetsvagen 10 A, 106 91 Stockholm, Sweden (e-mail: david.seim@ne.su.se). We thank co-editor Alex Gelber, David Card, Peter Fredriksson, Patrick Kline, and three anonymous referees for useful comments. Jesper Ekman, Clancy Green, Linn Hansen, and Katie Donnelly Moran provided outstanding research assistance. We acknowledge financial support from NSF Grant SES-1559014; the Center for Equitable Growth at UC Berkeley and FORTE Grant 2015-00490.

1 Introduction

Governments use an array of active labor market policies, in particular subsidies, to improve employment of disadvantaged groups. Oftentimes, the policies are presented as a one-time push to lift individuals onto better employment trajectories. Yet, evidence for active labor market policies to entail such persistent effects on employment is scarce (see e.g., Kluve et al. 2007).¹ On the labor supply side, employment effects from welfare reform or in-work subsidies quickly vanish after the policies end with no lasting gains (see e.g., Card and Hyslop 2005). This elusive policy persistence is puzzling in light of the large body of evidence for employment hysteresis — i.e. persistent employment shifts even after the original cause has disappeared (Blanchard and Summers 1986) — from non-policy labor market shocks such as job loss (Davis and Von Wachter, 2011), recession shocks (Blanchard and Katz 1992, Yagan 2019), trade shocks (Autor et al. 2014). Perhaps most related to our analysis, there is a literature showing that early labor market experiences such as graduating in a bad economy have long-lasting effects on career experiences and earnings (Kahn 2010 and Oreopoulos, Von Wachter, and Heisz 2012, Altonji, Kahn, and Speer 2016).

Our paper breaks new ground by showing that employment subsidies affecting the labor demand side can deliver employment hysteresis. To do so, we study a large and long-lasting employer-borne payroll tax cut for young workers aged 26 or less in Sweden that started in 2007 but was then suddenly repealed in 2015. This setting is uniquely suited for a comprehensive and credibly identified test for hysteresis from labor demand policies in three ways. First, the tax differential by age translates fully into a labor cost differential with no differential in net wages. Specifically, the labor cost of young workers fell by 12% during the tax cut and went up correspondingly by 14% when the tax cut was repealed (or workers age out), allowing us to study employment effects driven by labor demand responses. Second, it features sharp age cutoffs of eligibility, permitting us to study *lifecycle hysteresis*, i.e. whether previously treated workers' employment biographies are affected after they have aged out of the subsidy. Third, the abolition of the subsidy after having been in place for multiple years allows us to additionally measure persistent *aggregate hysteresis* effects on Swedish youth overall. We find first that the tax cut caused large and *growing* employment effects on the directly treated youth aged 26 and

¹The widely cited survey by Calmfors, Forslund, and Hemstrom (2002) for Sweden concludes “Our survey also questions the use of large-scale youth programmes ... as it is unclear whether there are any positive employment effects for the participating individuals.” For a policy-focused analysis of a broader set of tools boosting youth employment including, e.g., training, and an international perspective, see OECD (2017), ILO (2020), and OECD and ILO (2020).

below. The estimated employment effects are around 2 points three to five years into the reform – replicating the findings from previous work studying this reform.² Strikingly, these effects *double*, to 4–5 points, by the sixth and seventh year into the reform. This multi-year build-up of the treatment effect can be driven either by a slow adjustment in labor demand, such as under gradual accumulation of youth-specific capital, employer learning, or erosion of labor market discrimination against the young. It also suggests the possibility of hysteresis-like effects.

Next, we turn to a sharp test to cleanly isolate hysteresis effects, studying the evolution of employment as workers age out of the policy, as well as after the repeal of the reform. We document substantial employment hysteresis on both margins in response to the employer subsidies.

We start with lifecycle hysteresis. Here, we find clear positive and growing hysteresis employment effects emerging among young workers previously exposed to the tax cut after they age out and hence their wages are no longer subsidized. The lifecycle hysteresis suggests that the positive youth employment gains caused by the reform have consequences for the subsequent careers of initially directly treated workers.

The multi-year lead-up required to see these effects explains why they were missed by the existing short-run analyses (and perhaps may never have emerged in shorter-run policies). This finding of lifecycle hysteresis was not possible to detect in our previous evaluation of the reform (Saez, Schoefer and Seim, 2019), which studies the earlier years of the reform through 2013. In that earlier analysis, a slightly higher treatment effect is indeed visible in 2012-13 compared to 2009-11 (Figure 4), which is only briefly mentioned but not interpreted (p. 1734), and the estimation of employment effects pools the entire post-period (2009-13). This paper by contrast adds five more years to the data, through 2018.

Importantly, the lifecycle hysteresis patterns are consistent with an earlier cohort case study of the 1982 cohort from 2006 to 2011 included as one intriguing additional empirical check within the broader study of the tax cut in Egebark and Kaunitz (2018). The 1982 cohort was treated in 2007 but immediately ages out in 2008. Egebark and Kaunitz (2018, page 172) document a persistent effect in 2008 but not in 2009, 2010 or 2011. Like Egebark and Kaunitz (2018), our estimates for 2010 (which we pool through 2013) do not show significant lifecycle hysteresis,

²Earlier work had found significant but modest positive effects on youth employment rates (and no net-wage incidence) studying up to the first five years of the reform, Skedinger (2014), Egebark and Kaunitz (2013, 2018), and Saez, Schoefer, and Seim (2019). These studies have investigated neither the later years of the reform nor its repeal, and hence have not been able to document the hysteresis patterns we document emerge in those later years. Egebark and Kaunitz (2018) do include an intriguing and important early check for lifecycle hysteresis during the first five years of the reform for one cohort, which we summarize below.

which only start becoming large and statistically significant in the long run (2014-15). As we, Egebark and Kaunitz (2018) worry about the 2009 downturn and its aftermath masking such effects, so we, thanks to additional sample years, drop that year, but we can study the longer run effects. Egebark and Kaunitz (2018) presciently note the “open question whether the effect would have persisted longer, had the economic downturn not taken place in 2009” (page 172). While we cannot distinguish the downturn and lifecycle hysteresis taking longer to show up generally as an explanation for the growing lifecycle hysteresis effects in the late years of the reform, our study drawing on data after 2011 (which marked the end of the Egebark and Kaunitz (2018) sample period) confirms their conjecture.

Finally, we turn to aggregate hysteresis in the time series, studying the aftermath of the repeal in 2015. Strikingly, despite the associated large increase in youth labor costs, in the first three years after the repeal for which we have data, we do not see *any* reduction in youth employment. We view this as simple and compelling evidence for market-level hysteresis effects due to the clean and sharp difference-in-difference nature of our quasi-experimental policy variation and its sharp deactivation at repeal, compared to, e.g., trade shocks or recessions.

We recognize that young workers’ unemployment rates are more sensitive to the business cycle than overall unemployment rates. Therefore, the business cycle might confound the effects we find. To address this issue, we construct adjusted employment rates by age that control for the business cycle using the pre-reform period. These adjusted series also display excess youth employment—during and after the reform—confirming that our employment results are robust to controlling for business cycle effects. These controlled effects show that employment effects grow during the reform and remain stable after the repeal at the level of the last years of the reform (while they actually grow somewhat post-repeal in our baseline results).

What mechanism could explain our results? The persistent youth employment effects could be explained by the time it takes for employers to factor in the repeal into their personnel decisions. Alternatively, it could reflect a persistent change in hiring decisions and organization. For example, firms might have developed youth labor intensive technologies that cannot be reversed quickly. It is also conceivable that the payroll tax cut reduced discrimination against the young in hiring decisions, and that such discrimination does not come back after the repeal. Consistent with this explanation, we find a sharp decrease in the fraction of job postings with minimum age requirements after the reform starts with no bounce back after the reform ends (in contrast, job postings with gender requirements stay flat). These patterns echo evidence in the anti-discrimination literature that affirmative action policies can have long-term effects on minority employment even after the policies end (Miller and Segal 2012), and possibly driven

by permanent changes in screening methods for potential hires (Miller 2017).

One implication of our three employment findings combined is that overall employment effects of the reform in latter years are much larger (and fiscally cheaper) than the effects estimated by previous work on the first five years of the reform and excluding hysteresis (Egebark and Kaunitz 2013, 2018, Skedinger 2014, and Saez, Schoefer, Seim 2019).

Future years will show whether the youth employment gains start to disappear. Interestingly, during the COVID crisis, Sweden reintroduced a new temporary payroll tax cut for the young from January 2021 to March 2023 as a way to address the surge in unemployment particularly acute among the young.

Our paper is connected to other literatures on hysteresis effects. First, there is a large literature that studies policies targeting disadvantaged areas (see Kline and Moretti 2014 for a recent survey). For example, the Tennessee Valley Authority program to develop the US South during the Great Depression recently analyzed by Kline and Moretti (2013) had permanent effects on economic development. More recent policies, such as the US empowerment zones seem to have more modest effects (see e.g., Busso, Gregory, and Kline 2013). Second, World War II was a huge labor demand shock for American women. While female labor force participation fell back shortly after the war (Goldin 1991), geographical variation shows that there was some long-run persistence (Acemoglu, Autor, Lyle 2004 and Goldin and Olivetti 2013). Third, a number of studies in behavioral economics have shown persistent effects of temporary policies on various individual behaviors such as energy use (Allcott and Rogers 2014), exercise (Charness and Gneezy 2009), smoking (Giné, Karlan, and Zinman 2010), and voting (Fujiwara, Meng, and Vogl 2016).

2 Institutional Setting and Data

In this section, we first discuss the institutional setting of the payroll tax in Sweden and the payroll tax cut reform and its subsequent repeal. Next, we present the data we use for the analysis.³

2.1 Payroll Tax Cut for Young Workers in Sweden

Swedish payroll tax. In Sweden, the entirety of the payroll tax on earnings is nominally paid by employers and the tax is proportional to wage earnings with no exemption and no cap.

³More complete details can be found in Egebark and Kaunitz (2013, 2018), Skedinger (2014), or Saez, Schoefer, Seim (2019).

The payroll tax rate is uniform across industrial sectors and covers all employers public and private. The top series in the solid line in Figure 1 depicts the normal payroll tax rate from 2004 to 2021. The normal tax rate has been quite stable at around 31-32 percent over this period.

Young workers' payroll tax cuts. The second series in the dashed line in Figure 1 depicts the preferential payroll tax rate for young workers. In 2007-9, a new center-right coalition government implemented a payroll tax cut targeted toward young workers. The explicit aim of this reform was to fight youth unemployment, which had risen in previous years, and was perceived to be excessively high. It was enacted as a permanent tax change (with no change in benefits). By January 1st, 2009, the payroll tax rate was 15.49 percent (less than half the normal rate of 31.4%) for all workers turning 26 or less during the calendar year. To be precise, in 2009, the payroll tax cut applied to all workers born in 1983 or later on the totality of their 2009 earnings; in 2010, the payroll tax cut applied to all workers born in 1984 or later, etc. Hence, a worker's *only* determinant of eligibility for a full calendar year is year of birth (and not actual age when the earnings are received). Correspondingly, in our analysis, *age* is always defined as year of observation minus birth year (regardless of birthday date within the year). The tax cut is directly administered through the payroll tax software used by employers where individual earnings and year of birth are reported on a monthly basis by employers. Therefore, take-up is close to 100 percent.

Repeal in 2015-16. The left-wing opposition parties were against this payroll tax cut from the start. They lost the 2010 election but narrowly won the 2014 election on September 14. Therefore, in 2015, the new center-left government abolished the payroll tax cut for young workers. The new government worried that the number of new jobs were too small in relation to these costs. The lower payroll tax rate for the young expired in steps in 2015-6 as depicted and explained in Figure 1. The bill was passed on March 25, 2015 following a proposal put forward on October 7, 2014, just after the election. On August 1, 2015, all workers turning 25 and less in 2015 had their taxes increased to 25.46 percent (63% of the tax rate gap was closed). On June 1, 2016, the normal payroll tax rate was restored for all workers. The payroll tax cut lasted 9 years, and 6.5 years in its strongest form.

New COVID payroll tax cut. As part of the response to the COVID crisis, a new and temporary youth payroll tax cut started on January 1st, 2021 and is set to expire after 2.25 years on March 31, 2023. For workers turning 19-23 in the calendar year, the payroll tax is cut

to 19.73% as depicted in the figure. For workers turning 15-18, the payroll tax was cut even more, down to 10.21%. The 2021 payroll tax cut applies only up to a cap of 25,000 SEK of monthly earnings (approximately \$3,000). Data to study this new payroll tax cut are not yet available and hence this analysis is left for future work.

Other Contemporary Reforms. The 2007 payroll tax cut was accompanied with other reforms aimed at improving employment: a new hiring subsidy (in the form of temporary payroll tax cuts) for people unemployed or disabled for at least one year, an extension of the maximum duration of temporary labor contracts from one to two years, an earned income tax credit, and a value added tax cut in the restaurant sector (see Skedinger 2014 for a detailed presentation of these reforms). Younger workers are relatively more likely to benefit from these reforms but, in contrast to the payroll tax cut we focus on, they do not create a sharp discontinuity by age.⁴ Saez, Schoefer, Seim (2019, online appendix A.3) provide detailed analysis on the New Start Jobs and the Reform of Temporary Contracts. They show that participating in the New Start Jobs is low (around 2%) and only slightly decreasing by age ruling out confounding effects (Figure A22(b)). They also show that excluding all workers benefitting from the New Start Jobs does not affect the employment results at all (Figure A22(c)). For temporary contracts, they show that there is a *smooth* gradient by age in the likelihood of being in a temporary contract from 75 percent for workers aged 20 down to 15 percent at age 35 but with no discontinuity at age 25 (Figure A23(d)). Furthermore, the share of temporary contracts has stayed stable from the pre-reform period (2002-2006) and into the post-reform period (2009-2013). The value-added-tax cut for restaurant workers and the earned income tax credit do not have an explicit discontinuity at age 25 and hence should not generate a discontinuity in any employment effects they might generate. Therefore, none of these contemporary reforms could explain the specific discontinuities by age that we obtain here.

Wage Setting in Sweden. The vast majority of employees in Sweden are covered by collective bargaining agreements. Fredriksson and Topel (2010) show that 36 percent of all employees are covered by agreements where wages are bilaterally bargained between employer and employee and 57 percent are covered by agreements in which increases in total labor costs are only set at the firm level and local negotiations then set the distribution of increases within the firm. Therefore, there is scope for bargaining at the individual level for many workers but it is

⁴The only exception is the New Start Jobs, which eliminates payroll taxes for people out of work for at least 12 months. The young aged 20-25 only need to have been unemployed for 6 months to qualify.

important to note that the importance of union bargaining in Sweden could have a large impact on our results, hereby potentially affecting the external validity of our results for settings where unions have less power as for example in the United States. Saez, Schoefer, and Seim (2019) find that the minimum wage floors, set at the industry, occupation, tenure, and sometimes age level in Sweden, are too low to explain any of the employment and wage incidence results they obtain.

2.2 Administrative Data from Statistics Sweden

Our analysis is based on the full population of all Swedish residents (as of December 31 each year) aged 16 and above for years 1990-2018. To study employment, we obtain annual earnings and employment spells from matched employer-employee records covering 1985-2018. For each spell, these data record annual wage payments and months worked (used to administer the social security and income tax systems). Additionally, from the Income Tax Register, we retrieve total wage earnings. From the Integrated Database for Labour Market Research (various administrative records compiled by Statistics Sweden), we obtain the unemployment history (days registered with the unemployment insurance agency and benefits receipts), gender, year and month of birth.

To study wages, we also link to this baseline population a matched employer-employee annual data set: the Structure of Earnings Survey, which covers worker-level wages, occupational codes and hours of work, for a very large sample of firms. The data set covers all industrial sectors, and specifically all public sector employees and around 50 percent of private sector workers.⁵ The information is collected during a measurement week (in September-November, hence the 2016 data covers post-abolition wages) for all workers employed for at least one hour during that week. The wage concept is the full-time equivalent monthly wage prevailing in the given month, including all fixed wage components, piece-rate and performance pay and fringe benefits. Our wage concept is defined as the full-time equivalent contracted monthly wage. We also adjust wages for inflation (CPI base-year 2018) and convert to US dollars using an exchange rate of 9.8 SEK/USD (as of 9/30/2019). We use this wage concept to study the incidence of the payroll tax cut on market wages in Section 3.1.

⁵The sample is a stratified random sample of firms, with larger weights on larger firms. All firms with more than 500 employees are included. Our wage results are robust to reweighting the wage sample to match the industry- and the firm-size distribution of the total population of employees.

3 Empirical Analysis: Direct Employment Effects, and Hysteresis

We first analyze the effects of the payroll tax reform on cohort-specific wages to determine the incidence of the payroll tax, and then move to employment effects and its dynamics. Our strategy follows the methodology adopted by previous work (Egebark and Kaunitz 2013, 2018, Skedinger 2014, Saez, Schoefer, and Seim 2019) and extends it to the last years of the reform and the three years after the repeal. Naturally, the results we obtain for the first years of the reform replicate these earlier analyses. The methodology follows a basic difference-in-differences (DiD) strategy. This requires the parallel trend assumption between the treatment and control groups (in the counterfactual absent the reform). This assumption is more demanding in the longer run than in the medium run analyzed by earlier work. To address potential concerns and particularly business cycle effects, we provide a robustness analysis using a much wider range of years pre-reform in Section 3.3.

3.1 Wage Incidence: Effects on Net and Gross Wages

Figure 2 depicts the average monthly wage in Sweden by age for different time periods for all employees in the Structure of Earnings Survey data. The top panel depicts *net wages* defined as monthly wage earnings net of payroll taxes. The bottom panel depicts *gross wages* defined as monthly wage earnings gross of payroll taxes, i.e. the labor costs the employer pays including payroll taxes. Age is defined as the age reached during the calendar year (the relevant concept for the payroll tax cut). 2003-2006 are pre-tax cut years (red square series). 2007-2015 are tax cut years (blue circle series). 2016-2018 are post-repeal years (green triangle series). The reform applies to ages up to 26 in 2009-2014 and up to age 25 in 2007-8 and 2015 (as depicted by the two dashed vertical lines, see Figure 1 for details). The size of the tax cut is less in 2007-8 and 2015 as the tax cut is phased-in and phased-out.

The top panel shows that net wages are continuous at the age thresholds, and the bottom panel shows that gross wages are discontinuous. In the bottom panel, the discontinuity happens at age 25 to 26 in years 2007-2008 and 2015 when the tax cut did not cover workers aged 26. In years, 2009-2014, when the reform applied to workers aged 26 as well, the discontinuity in gross wages is between ages 26 and 27. As a result, the labor cost of young workers went down by about 12% during the tax cut and then went up correspondingly by 14% when the tax cut was repealed. (The tax rate for young workers is 15.49 percent while the normal rate is 31.42 percent, hence a reduction of labor costs of $(31.42 - 15.49)/(100 + 31.42) = 12.1\%$.)

This implies that employers do not adjust wages by age in response to the payroll tax cut and its subsequent repeal, and therefore absorb fully the tax discontinuity created by the age and time specific payroll tax reform.⁶ These results replicate the earlier analyses of the payroll tax introduction, which found no or very little effect on net wages (Egebark and Kaunitz 2013, 2018, Skedinger 2014, Saez, Schoefer, and Seim 2019).⁷ Our new analysis shows that the absence of effects on net wages continues into the last years of the reform and in the post-repeal period as well. This complete absence of incidence on net wages also goes against the standard view in public economics that employer payroll taxes are borne by employees. Our results imply that the tax changes translate fully into changes in the labor costs of young workers relative to slightly older untreated workers, permitting us to study employment effects of the reform-induced cost reduction of youth labor next.

3.2 Effects on Employment: During and After the Subsidy

Our wage results imply that an employer would save 12% of labor costs by replacing an ineligible older worker (say, aged slightly above 26) by an eligible young worker (aged 26 or less), given the lack of net wage incidence. In Sweden, CVs typically include age and job applications often include social security IDs that reveal birth dates. Hence employers can generally observe age before hiring. As these two groups of workers should be close substitutes, profit maximizing firms should want to hire more eligible workers or put more effort in retaining eligible workers (relative to ineligible workers). Indeed, this labor demand increase triggers the wage pressure that eventually equalizes gross wages across treated subsidized and control groups in the standard competitive model.

Methodology. To analyze such employment effects, we examine the employment rate in the labor force by age group and over time using the individual annual earnings data (see Section 2 for details). The employment rate is defined as the ratio of employment to the labor force. The employment numerator is defined as all employed residents with annual wage earnings above a small annual threshold. The small annual threshold is equal to \$4,490 in 2012 (in 2019 USD)

⁶Perhaps wages do not appear to respond in ongoing jobs due to wage rigidity or implicit contracts. To test this, we repeat Figure 2 but limiting the sample to new hires in Appendix Figure A1, and even for this subsample, we do not see any discontinuity in net wages either when the reform is put in place or when it is repealed.

⁷The absence of effects on net wages in these earlier studies is pervasive and present even when restricting the sample to new hires or to high turnover industries, and cannot be explained by the minimum wage either (Saez, Schoefer, and Seim 2019).

(and adjusted for median wage growth in other years).⁸

The labor force denominator is defined as all residents who are either (i) employed with annual wage earnings above the small annual threshold or (ii) unemployed (defined as having registered with the Unemployment Office at any point during the year).

Figure 3a depicts the *unadjusted* employment rate by age and time periods. Series in red squares are before the payroll tax cut is in place. Series in blue circles are when the reform is in place in three periods 2010-11, 2012-13, 2014-15 with growing circle sizes over time. We exclude 2007-09 because these years are strongly affected by the business cycle (2007 boom followed by the Great recession years of 2008-2009). We show in Section 3.3 that youth employment is more sensitive to the business cycle than employment at older ages and show that our results are robust to including all years and controlling for business cycle effects. Series in green triangles are after the repeal i.e. 2016-18 – and this post-repeal period serves as our additional time series for a DiD strategy to test persistence and aggregate hysteresis effects in all age groups.

Figure 3b presents an *adjusted* difference-in-difference version of these results, by measuring employment rates by age and single years relative to 2006, which controls for overall employment effects across years. The normalization is made by aligning the unemployment rate (one minus the employment rate) for ages 35-40 to the 2006 level for each year (multiplicatively) and then taking the difference in employment rates with 2006.⁹ That adjusted series can therefore be read as the deviation of employment rates by age and year (relative to 2006) expressed in points of employment and controlling for the overall level of unemployment across years. The 35-40 workers hence serve as our “far-away” control group never directly affected by the reform that permits us to investigate spillovers onto the “nearby” control group (workers aged 27 and above), as would arise from lifecycle hysteresis effects.¹⁰ (Because employment rates in the top panel for ages 35-40 are very close in all periods (within 1 percentage point), results are robust to the normalization choice.)

We have three sets of results on the dynamics of the employment effects of the payroll tax cut

⁸This earnings threshold is a social security benchmark against which benefits are indexed, and corresponds approximately to working at 20 percent of full-time a full year at the minimum wage in the restaurant sector. We have checked that our results are robust to using alternative thresholds. Appendix Figure A2 reproduces Figure 3 using a threshold that is twice as high (\$8,980) and displays comparable employment effects, perhaps even slightly larger, implying that our employment effects are not concentrated solely at the very bottom of the earnings distribution.

⁹Formally, if e_{at} is the employment rate at age a in year t , and $u_{at} = 1 - e_{at}$ is the unemployment rate in year t at age a , the normalized unemployment rates are $\hat{u}_{at} = u_{at} \cdot (u_{35-40,2006} / u_{35-40,t})$ and hence the normalized employment rates are $\hat{e}_{at} = 1 - \hat{u}_{at} = 1 - (1 - e_{at}) \cdot (u_{35-40,2006} / u_{35-40,t})$. We plot the series: $\hat{e}_{at} - \hat{e}_{a,2006} = (e_{at} - e_{a,2006}) \cdot (u_{35-40,2006} / u_{35-40,t})$.

¹⁰The only exceptions are the cohorts aged 35 in 2017 and 35 in 2018. The first cohort of these was treated for one and a half years and the second for half a year, as the reform started in July 2007.

and the aftermath of the repeal.

Long- vs. medium-run effects while the reform is active. First, the data reveal a clear increase in youth employment rates from pre-reform periods (2003-4 and 2005-6) to early years of the reform (2010-11 and 2012-13) as documented in Saez, Schoefer, and Seim (2019). In these early years, the employment effects are concentrated at ages 22 to 24 with smaller effects for workers close to the age threshold (25 and especially 26 year old). It is possible that effects are smaller for the oldest treated workers because employers understand that such workers will age out of the reform quickly (recall that employers can observe age on job applicants' CVs or applications in most cases, and otherwise infer it from their education biography). Effects are also smaller for very young workers (aged 20 or 21), perhaps because such workers are very young relative to other employees making it more challenging for employers to find them, evaluate them, and incorporate them into their workforce.

Interestingly however, in 2014-2015, there is a clear further increase in youth employment while the employment rates of older workers aged 30-40 remain stable. Therefore, the effect of the reform appears much stronger in 2014-2015 than in earlier years. By year 2015, the employment effect is about three times as large as in 2010. In fact, as the reform matures, the employment effects both deepen and widen: the effects are larger at all ages but especially the very young. It is striking to see how monotonically the age gradients line up across years.

Lifecycle hysteresis and spillovers on control ages. Second, workers slightly older than 26 and hence no longer eligible for the tax cut appear to have a higher employment rate in 2014-2015 than in earlier years. These workers were exposed to the reform, suggesting that exposure to the reform has hysteresis effects on employment (or that spillover effects on the group level may take longer to materialize). Both panels show clearly that the effect of the reform spills over gradually across slightly older groups. The spillover is almost perfectly monotonic, providing compelling evidence that it is reform-driven. Consistent with this evidence, Saez, Schoefer, and Seim (2019, Figure A13) had shown that the employment effects of the young during the reform are driven primarily by retaining young workers more rather than hiring more young workers.

Post-repeal hysteresis for all young workers. Third and perhaps most striking, the employment effects of the young if anything keep increasing during the years 2016-2018 – *after the tax cut is repealed*. This suggests that the positive employment effects of the reform do not vanish after the payroll tax cut ends – a clear indication of hysteresis effect at this group level.

Moreover, the lifecycle hysteresis, or spillover, effects also continue after the reform, mirroring the reform-period shape (even appearing to be further increasing). This is particular clear in the normalized bottom panel, where we disaggregate by individual years.

Placebo test: 2004-5. Finally, the bottom panel includes two pre-reform years 2004 and 2005, which can serve as a placebo test (recall that the bottom panel is normalized relative to 2006). For these years, we do not detect any employment effects at any age, nor any pre-trend upward in youth employment rates (if anything the reverse), further lending credence to our identification assumption.

Regression evidence. Table 1A provides the corresponding estimates using a basic difference-in-differences regression based solely on the aggregate unadjusted cohort-year time series as depicted in Figure 3a. We use data at the year \times age level with 21 age categories a (20 to 40), and 13 years t (2003 to 2018, excluding 2007-2009 when the reform was partially phased-in and also the global recession year), yielding a total of $13 \times 21 = 273$ year \times age level observations. To estimate treatment effects and spillovers, we also create 6 coarser age groups denoted by A , comprising directly treated ages 20-26, and various spillover groups namely 27-28, 29-30, 31-32, and 33-34, and finally our baseline control age group 35-40. Following Figure 3, we also group years into 5 periods denoted by T , comprising 2003-4 (pre-period), 2005-6 (base period against which effects are estimated), 2010-3 (medium run i.e. combining the two nearly identical gradients in Figure 3a), 2014-5 (long run), 2016-8 (post-repeal). We run the following basic difference-in-differences regression at the year \times age level:

$$e_{at} = \alpha_0 + \alpha_a + \alpha_t + \gamma_{AT} + \varepsilon_{at}. \quad (1)$$

The regression includes the set of treatment age fixed effects α_a , and year fixed effects α_t . The coefficients of interest are the interaction effects $\gamma_{AT} = \sum_T \sum_A \gamma_{AT} \cdot \mathbb{1}(a \in A \wedge t \in T)$, and are reported in the table for each T and A (interactions exclude the 2005-6 baseline control period). These coefficients correspond to the average of the employment rates by age group and (now grouped) periods depicted in Figure 3b (except that the figure separates out 2005 and that we run the regression on unadjusted gradients reported in Figure 3a). Robust standard errors are reported.

The table confirms and quantifies the visual impression in Figure 3b. In 2010-13, the medium run, the effect on employment for the treated young (20-26) is 2.3 points. In 2014-15, the longer run, this effect jumps to 4.4. This shows that the effect doubles from the early years to the late

years of the reform. In 2016-18, after the repeal, the effect increases further to 6.3 points. The effect for ages 27-28 is small and insignificant in the early years 2010-13 but becomes significant for years 2014-15 (1.3 point effect) and even larger after the repeal (2.5 points). At ages 29-30, we find a significant effect of 1.2 points after the repeal (recall that these workers were exposed to the tax cut when they were younger). The placebo years 2003-4 do not show any effect either for the treatment group or for the other ages, supporting the validity of our research design.

Heterogeneity. Figure 4 explores heterogeneity by gender (top panels) and by local youth unemployment rate (as of 2006). The methodology is otherwise the same as in Figure 3, bottom panel (raw employment rates are presented in Appendix Figure A3). The y-axis scales are the same for comparison.

Gender. The top panels show the employment rate effects by age and year relative to 2006 for females (left) and for males (right). The figure shows a similar effect of the reform by gender in early years (2010-11, 12-13), but the effect on young female workers grows much more in subsequent years and remains much higher for females after the tax cut is repealed (although hysteresis effects are present for both genders).

Local youth unemployment rate. The bottom panels show the employment rates by age and year relative to 2006, separately for regions in the top quintile of their 2006 youth unemployment (left) and for bottom quintile regions (right). In 2006, just before the reform, there was wide variation across Sweden's 21 regions in youth unemployment. Regions in the bottom quintile had rates in the range of 10.5-12.4% compared to 20-23.3%, i.e. about twice as high, in regions in the top quintile.

The figure reveals a much larger effect of the payroll tax cut in high youth unemployment regions, already in the medium run. Strikingly, in the longer run, the acceleration of the treatment effect the national analysis indicated is particularly pronounced in the high youth unemployment regions (an effect not driven by mean reversion as the flat pre-period gradients indicate). In neither group do the gains immediately disappear as workers age out of the policy, since employment effects appear among the older ineligible workers, particularly so in the later years of the policy. Most strikingly however, in the high youth unemployment regions the additional gains in employment are not more quickly undone, but the larger treatment effect persists even after the subsidies are abolished. As a result, even three years after the policy was repealed, the regions maintain the substantial gains towards convergence in youth unemployment, as illustrated in Appendix Figure A3, which plots the unadjusted employment

rate levels.

3.3 Controlling for Business Cycle Effects

In our employment effects analysis presented so far, we focus on years 2003-06 and 2010-18 because all these years have remarkable stability in the employment rate of control workers aged 35-40, who are never affected by the reform. This is because all these years have fairly stable employment rates for the overall population implying no large business cycle movement.¹¹ However, it is well known that young workers' unemployment rates are more sensitive to the business cycle than overall unemployment rates (see, e.g., Jaimovich and Siu 2009 and Hoynes, Miller, and Schaller 2012). Therefore, even a modest business cycle movement over the years we consider might affect young workers and confound the effects we find. We address this issue in this section.¹²

Figure 5 plots the raw employment rates (defined exactly as in Figure 3, panel (a)) by age groups from 1993 to 2018. The young groups (ages 20-22, 23-24, 25-26), treated during the reform, are depicted in red; the hysteresis groups (aged 27-28, 29-30, 31-32, 33-34), who age out of the reform in our time frame, are depicted in green; and the older control groups (ages 35-36, 37-38, 39-40), who are never treated, are depicted in blue. Four points are worth noting. First, for each year, employment rates increase monotonically with age, and all move in tandem over the business cycle with a sharp downturn in the 1990s, a boom from 1997 to 2001, a milder downturn from 2001 to 2005, a short boom from 2005 to 2008, the Great Recession, which is pretty mild in Sweden from 2008 to 2010, followed by a slow recovery. Second, the young do display much larger fluctuations in employment rates over the business cycle with a clearly visible gradient by age. Third and consistent with our results above, the period 2010 to 2018 stands out. Employment rates for the old are almost constant but they improve considerably for the young shortly after the reform, and this improvement continues even after the reform. Fourth, the hysteresis effect is visible in the form of a compression of employment rates across ages that starts in the middle of the reform and continues after the reform. By 2018, employment rates are almost identical from age 23 to age 37, which is unprecedented over the full period we depict.

The excess sensitivity of youth employment rates to the business cycle motivates the construction of adjusted employment rates by age that control for the business cycle. We can then

¹¹As mentioned above, we excluded the early years of the reform 2007-2009 because they were subject to much larger business cycle fluctuations invalidating our basic approach of comparing workers by age.

¹²We thank three reviewers and the co-editor in charge for encouraging us to develop a strategy assessing this important concern.

test whether our excess youth employment effects (during and after the reform) are robust to using such business cycle adjusted employment rates. We proceed as follows.

To construct the normal employment level, we estimate a reduced-form Engel curve of age-specific employment to the average employment rate E_t across all ages 20-40 ($E_t = \sum_a e_{at} / 20$), which is our business cycle indicator. We do so over the pre-reform period 1993-2006 so that our estimation is not confounded with reform effects. That is, we run the following regression:

$$e_{at} = \alpha_t + \alpha_a + \sum_{a'=20}^{39} \phi_{a'} \cdot 1(a = a') \cdot E_t + \varepsilon_{at}, \quad (2)$$

where a is age 20,...,40, t is year from 1993 to 2006, α_t denote year dummies, α_a denote age dummies. The coefficients of interest are the ϕ_a on age dummies interacted with E_t . We omit the interaction with age $a = 40$ so that $\phi_{a=40}$ is effectively normalized to zero. Panel (a) of Figure 6 depicts the sensitivity of employment rates to the business cycle by age over the pre-reform period 1993-2006 by depicting these estimated coefficients $\hat{\phi}_a$ at each age. Consistent with Figure 5, there is indeed excess sensitivity of youth employment rates to the business cycle which is monotonically decreasing with age.

Next, we test whether we find excess youth employment during and after the reform adjusting for exposure to the business cycle according to these age-patterns. We include two groups: the young aged 20-26 (the treated group) and the old aged 35-40 (the control group never treated by the reform). We conduct a standard difference-in-differences regression, where the dependent variable is the employment rate and the independent variables are year dummies, age dummies and year dummies interacted with the treated group. We exclude the interaction for year 2006 so that the coefficient is normalized to zero in this year. In addition, we control linearly for the $\hat{\phi}_a$ -coefficients estimated above and allow the coefficients on the $\hat{\phi}_a$ -coefficients to vary by year. This step implies that the employment rates are effectively the residual ones controlling for the excess sensitivity of the business cycle across the age distribution, where we permit flexible year-specific coefficients on the $\hat{\phi}_a$. To account for this two-step procedure when making inference, we estimate standard errors using a bootstrap method which draws observations with replacement 100 times. Panel (b) of Figure 6 reports the coefficients on these interactions by year. Two results are worth noting.

First, none of the coefficients are significant in the pre-reform period from 1993 to 2005, showing that our method is indeed successful to control for business cycle effects of youth employment. Second, the coefficients increase and become significant after the reform starts. The coefficients grow over years from around 1.5 percentage points (and still insignificant) in

2007 to around 2.5 points in the early years of the reform (2008-2012), and to about 3 to 3.5 points (and significant) in 2013-2015, the last 3 years of the reform. Furthermore, the strong positive effects persist above this 3 points level (and significant) in the three years after the reform (2016-2018). These effects are quantitatively similar in magnitude to the ones we documented in our unadjusted method above (Table 1A). The only discrepancy is that the direct method finds even larger effects post-repeal while the adjusted method finds effects that stabilize post-repeal (at the level for the last years of the reform). Overall, this analysis shows that the business cycle is not driving our findings of positive effects on youth employment.

As a caveat, our adjusted unemployment rates are based on a relatively short time series from 1993-2006. None of the recessions in that time period exactly mimics the 2008-9 recession so that the adjusted series require making the identification assumption that, absent the payroll tax cut, the excess sensitivity of youth unemployment to the business cycle is stable over time.¹³

4 Implications of Hysteresis for Policy Effectiveness

We now present a quantitative assessment of how the hysteresis-like employment dynamics we have uncovered affect the policy evaluation of the reform. We proceed in two parts, first studying job creation effects, and then the fiscal costs of the policy.¹⁴

Jobs created. First, in Table 1B, we calculate the count of job creation from the policy, separating the contributions of the direct effects and the hysteresis effects. Table 1B is organized as our regression Table 1A, with rows denoting periods and columns denoting age groups. The job-years are defined as 1 more employed person in a given year (i.e. with annual earnings above the aforementioned threshold). These effects are computed as the product of the age-year-specific treatment effect from Table 1A, multiplied by the age group’s pre-period (2006) labor force count (detailed in the table note). The additional last entry of each row and column additionally includes a total count for a given age group or period as well as the share of that entry contributing to total job creation (itself reported in the bottom right corner).

Out of the total 229,228 job-years created (bottom right entry), 91% were created among

¹³We have also experimented with estimating the age-specific employment cyclicity depicted in Panel (a) of Figure 6 in regional panel data, using the 21 counties, to obtain broader employment variation. We have found very similar results for the analogs of both panels in Figure 6.

¹⁴Our difference-in-difference design cannot detect negative spillovers between treated and untreated cohorts (e.g., Crépon et al. 2013, Lalive, Landais, Zweimuller 2015). In that case, we may overestimate the employment effects and hence underestimate the fiscal costs per job. However, our raw time series of employment in Figure 3(a) did not exhibit an obvious dent in the plausibly closest substitutes plausibly most subject to such negative spillovers, i.e., the slightly younger workers.

the directly treated 20-26 olds. The dynamics of job creation reveal, first, that 23% of these jobs were created in the final two years of the policy (2014-15), as many as within the four-year medium run period (2010-13, 21%).

Second and perhaps most strikingly, half (51%) of these jobs among ages 20-26 were created *after* the repeal of the policy (2016-18), hence from hysteresis effects. Across all age groups, these three post-repeal years account for 128,886 (56%) of all jobs created .

Lifecycle hysteresis effects are strongest among the 27-28 year olds, accounting for 19,237 jobs. Spillover effects on the even older groups are smaller while the policy is active, but strikingly start showing up even among the 29-30 year olds *after* the repeal.

Payroll tax costs per job created. Second, we provide simple evaluation of a widely used policy number, namely fiscal costs per job created, here in terms of payroll tax revenue foregone, but excluding revenue gains from the larger tax base. Our comprehensive analysis dramatically raises this jobs per dollar measure not just because overall job creation is larger (as reported in Table 1B), thereby raising the numerator – but additionally because any additional jobs from hysteresis come for free as they are no longer actually subsidized – i.e. they do not enter in the denominator (the fiscal cost).

We calculate the baseline (2006) payroll tax base for ages 20-26 (labor income) and multiply it with the payroll tax subsidy rate (15.9 percentage points, namely the gap between the regular and lowest tax rate, 31.4% vs. 15.4%), yielding lost payroll tax revenue of \$1,496,753,536 (in 2019 USD) per year. Dividing by the age 20-26, medium-run job creation (per year), the per-job year cost was \$113,943 during the medium run. In the longer-run period 2014-15 considered on its own, the per-job cost among the 20-26 year old drops by nearly *half*, to \$60,648. Our full analysis of all policy years plus the “free” post-repeal hysteresis effects, implies an *average* per-job cost among ages 20-26 of \$64,569 even when including the four expensive medium-run years. When we furthermore include the spillover treatment effects the slighter older workers exhibit, the comprehensive measure implies another reduction in costs to \$58,766. This is a reduction by *half* (48%) compared to the \$113,943 effectiveness measures if drawing from narrower treatment effect estimates restricted to the initial reform years such as in Saez, Schoefer, Seim (2019).

5 Hysteresis Mechanism: The End of Youth Discrimination?

We conclude with a discussion of the mechanisms that may drive the hysteresis in youth employment. The employment response we have uncovered is likely due to labor demand effects rather than labor supply effects because the net wage of eligible young workers does not increase (Section 3.1), and hence labor costs were lowered one to one (and in light of the high youth unemployment). Sluggish adjustment from attention to the tax reversal may explain the findings. Alternatively, it could reflect a persistent or even permanent changes in labor demand, for example if firms switched to youth intensive technologies that cannot be reversed quickly.

Conceivably, the payroll tax cut also removed discrimination against the young in hiring (such as posting job ads requiring prior years of experience or some minimum age), and such discrimination did not come back after the tax cut repeal as employers had learned about employing young workers more effectively and perhaps revised beliefs about their productivity. Below, we provide a suggestive empirical evaluation of this plausible mechanism.

We investigate this potential mechanism using job vacancy postings from the Swedish Public Employment Service (PES). Employers post vacancies free of charge on the PES platform (“Platsbanken”), which covers a large share of all posted vacancies in Sweden, around 30-50% between 2006-2013 (Arbetsförmedlingen 2014). Any job seeker can register with the portal at no cost, and UI recipients must enroll. The open-source data span 2006-17, and cover 4.9 million job ads corresponding to around 9 million jobs (due to multi-job listings). The listings include the job description and desired characteristics of the hire.

Using text search for discriminatory key phrases, we classify these listings by whether they discriminate against youth (our focus) and also by gender (a control demographic dimension). The main step was two external researchers manually reviewing and classifying a random sample of 3,000 ads for discriminatory phrases; we then classify the universe of vacancies using text search following these phrases, by youth and by gender phrases. Anti-youth discrimination manifests itself largely as a minimum age cutoff in the description. Appendix A and Appendix Figure A4 report the full list of the age and gender phrases, and their time series.

Figure 7 shows the time series of this share of anti-youth phrases in job listings (solid blue lines with circles). Since the data start in 2006, we are unable to test a parallel pre-trend assumption. But the time series of youth discrimination starts off high in the pre-reform years in 2006 and 2007, reaches its peak in 2007 (the year the reform starts getting phased-in), and then steadily declines steadily over time, continuing to fall and reaching its the lowest levels in

the post-repeal years, consistent with our hysteresis effects.

As a control time series, Figure 7 also plots the share of listings containing phrases of gender discrimination (dashed red line with squares). This control time series is very stable, with only a small trend downwards. Importantly, the two series reach nearly the same level in the end of the time period, whereas the youth discrimination share was initially three times as high. The pattern (plotted for private-sector jobs) is robust to including public-sector jobs, restricting to permanent or full-time positions, and weighting each ad by its underlying number of open positions. The fraction of ads that are discriminatory against young falls by about .4 percentage points over the period. This is quantitatively small relative to the employment effects we have documented (around 2-3 percentage points). Discrimination could still explain the results if only the tip of discrimination iceberg comes to view through the ads.

One concrete confounding driver we cannot definitively rule out is the introduction of a broad-based anti-discrimination law, which came into effect in 2009 (after being ratified in June 2008). The scope of the law included many categories besides age, and covered not just the labor market but various areas of society (such as education, public policies, medical care, social insurance, or housing), although age was explicitly exempted from most areas (but included working life). An interesting possibility is that the law curbed age-based discrimination but less so gender-based discrimination, which had been illegal earlier. To our knowledge, the effects of the law on employment outcomes by age group have not been studied empirically. Since the law had no age cutoff, it is unlikely to explain our main results on the employment effects of the cohorts treated by the payroll tax reform.

We tentatively conclude that a persistent reduction in discrimination against young workers could explain our results, but we cannot provide a definite test of this mechanism. Our findings suggest that a dedicated empirical study isolating youth discrimination in the labor market may be promising.¹⁵

6 Conclusion

Our paper has shown that a large payroll tax cut for the young had positive and growing effects on youth employment. We have also documented novel labor-demand-driven “hysteresis” from

¹⁵In Sweden, discrimination in hiring decisions has been documented in field studies in the context of *older* workers in a sample that excluded workers younger than 35 (Carlsson and Eriksson 2019). Indirectly related but consistent with anti-youth discrimination, field experiments and resume studies have found evidence for discrimination of online platform workers with fewer evaluations (Pallais 2014), and applicants with longer unemployment durations (Eriksson and Rooth 2014 for Sweden, Kroft, Lange and Notowidigdo 2016 for the United States).

this policy – i.e. persistent employment effects even after the subsidy no longer applies – along two dimensions. First, over the lifecycle, employment effects persist even after workers age out of eligibility. Second, after the repeal, employment remains elevated at the maximal reform level in the formerly subsidized ages. These hysteresis effects more than double the direct employment effects of the reform. A persistent reduction in discrimination against young workers could possibly explain our results.

An important question beyond the scope of our analysis of the unique introduction and abolition of the age-specific payroll tax cut in Sweden is external validity. For instance, it is possible that—formal or informal—institutional factors such as firing restrictions or implicit contracts have affected the degree of hysteresis, e.g., by precluding layoffs of workers aging out of eligibility. Similarly, the erosion of youth discrimination naturally depends on the degree of pre-existing youth discrimination. Whether our findings carry over to other disadvantaged worker groups is an open question. Finally, since the initial hysteresis effect on untreated workers took seven years to show up, firms’ adjustment to the repeal may also take time.

Future years will show whether the youth employment gains start to disappear. As we mentioned, during the COVID crisis, Sweden introduced a new and temporary payroll tax cut for the young from January 2021 to March 2023 as a way to address the surge in unemployment that was particularly acute among the young.¹⁶ This new payroll tax cut applies only to the quite young (age 23 or less) and is even larger for the very young (age 18 or less). Furthermore, it applies only to monthly wages up to about \$3,000. This reform offers promising new variation for a future empirical analysis.

¹⁶Labor force survey statistics from Sweden show that the (seasonally adjusted) unemployment rate among the young (aged 15-24) surged by 3.6 points from 20.5% in January 2020 pre-COVID to 24.1% in January 2021 (when the new payroll tax cut starts) while overall unemployment (age 15-74) increased by 1.9 points from 7.1% to 9.0% during the same period. In percent term though, the increase is actually lower for the young (17.6%) than overall (26.8%). See online statistics at <https://www.ekonomifakta.se/Fakta/Arbetsmarknad/Arbetsloshet/Arbetsloshet/>.


References -

- Acemoglu, Daron, David H. Autor, and David Lyle.** 2004. "Women, War, and Wages: The Effect of Female Labor Supply on the Wage Structure at Midcentury." *Journal of Political Economy* 112(3): 497-551.
- Allcott, Hunt, and Todd Rogers.** 2014. "The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation." *American Economic Review* 104(10): 3003-3037.
- Altonji, Joseph G., Lisa B. Kahn, and Jamin D. Speer.** 2016. "Cashier or Consultant? Entry Labor Market Conditions, Field of Study, and Career Success." *Journal of Labor Economics* 34(S1): S361-S401.
- Arbetsförmedlingen.** 2014. "Arbetsmarknadsrapport 2014".
- Autor, David, David Dorn, Gordon Hanson, and Jae Song.** 2014. "Trade Adjustment: Worker-Level Evidence." *Quarterly Journal of Economics* 129(4): 1799-1860.
- Blanchard, Olivier and Lawrence Summers.** 1986. "Hysteresis and the European Unemployment Problem. *National Bureau of Economic Research Macroeconomics Annual* 1, 15-78.
- Busso, Matias, Jesse Gregory, and Patrick Kline.** 2013. "Assessing the Incidence and Efficiency of a Prominent Place Based Policy." *American Economic Review* 103(2): 897-947.
- Calmfors, Lars and Forslund, Anders and Maria Hemstrom.** 2002. "Does Active Labour Market Policy Work? Lessons from the Swedish Experiences." CESifo Working Paper Series No. 675.
- Carlsson, Magnus, and Stefan Eriksson.** 2019. "Age Discrimination in Hiring Decisions: Evidence from a Field Experiment in the Labor Market." *Labour Economics* 59, 173-183
- Card, David, and Dean R. Hyslop.** 2005. "Estimating the Effects of a Time-Limited Earnings Subsidy for Welfare-Leavers." *Econometrica* 73(6): 1723-1770.
- Charness, Gary, and Uri Gneezy.** 2009. "Incentives to Exercise." *Econometrica* 77(3): 909-931.
- Crépon, Bruno, Esther Duflo, Marc Gurgand, Roland Rathelot, and Philippe Zamora.** 2013. "Do Labor Market Policies Have Displacement Effects? Evidence from a Clustered Randomized Experiment." *Quarterly Journal of Economics* 128(2): 531-580.
- Davis, Steven, and Till Von Wachter.** 2011. "Recessions and the Costs of Job Loss." *Brookings Papers on Economic Activity*, 1-55.
- Egebark, Johan and Niklas Kaunitz.** 2013. "Do Payroll Tax Cuts Raise Youth Employment?", IFAU Working Paper 2013:27.
- Egebark, Johan and Niklas Kaunitz.** 2018. "Payroll Taxes and Youth Labor Demand." *Labour Economics* 55: 163-177.
- Eriksson, Stefan and Dan-Olof Rooth.** 2014. "Do Employers Use Unemployment as a Sorting Criterion when Hiring? Evidence from a Field Experiment." *American Economic Review* 104(3): 1014-39.
- Fredriksson, Peter, and Robert H. Topel.** 2010. "Wage Determination and Employment in Sweden since the Early 1990s: Wage Formation in a New Setting." In R. Freeman, B. Swedenborg and R. Topel eds., *Reforming the Welfare State: Recovery and Beyond in Sweden*. (Chicago: University of Chicago Press), 83-126.
- Fujiwara, Thomas, Kyle Meng, and Tom Vogl.** 2016. "Estimating Habit Formation in Voting." *American Economic Journal: Applied Economics* 8(4): 160-188.
- Giné, Xavier, Dean Karlan, and Jonathan Zinman.** 2010. "Put your Money where

- your Butt is: A Commitment Contract for Smoking Cessation.” *American Economic Journal: Applied Economics* 213-235.
- Goldin, Claudia.** 1991. “The Role of World War II in the Rise of Women’s Employment.” *American Economic Review* 81(4), 741-756.
- Goldin, Claudia, and Claudia Olivetti.** 2013. “Shocking Labor Supply: A Reassessment of the Role of World War II on Women’s Labor Supply.” *American Economic Review* 103 (3): 257-62.
- Hoynes, Hilary, Douglas Miller, and Jessamyn Schaller.** 2012. “Who Suffers During Recessions?” *Journal of Economic Perspectives* 26(3): 27–48.
- Jaimovich, Nir and Henry Siu.** 2009. “The Young, the Old, and the Restless: Demographics and Business Cycle Volatility.” *American Economic Review* 99(3): 804–26.
- ILO.** 2020. *Global Employment Trends for Youth 2020*. (Geneva: International Labour Organization).
- Kahn, Lisa.** 2010. “The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy.” *Labour Economics* 17 (2), 303-316.
- Kline, Patrick, and Enrico Moretti.** 2013. “Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority.” *Quarterly Journal of Economics* 129(1), 275-331.
- Kline, Patrick, and Enrico Moretti.** 2014. “People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs.” *Annual Review of Economics* 6(1), 629-662.
- Kluve, Jochen, David Card, Michael Fertig, Marek Gora, Lena Jacobi, Peter Jensen, Reelika Leetmaa, Leonhard Nima, Eleonora Patacchini, Sandra Schaffner, Christoph M. Schmidt, Bas van der Klaauw, Andrea Weber.** 2007. *Active labor market policies in Europe: Performance and perspectives*. Springer Science and Business Media.
- Kroft, Kory and Lange, Fabian and Notowidigdo, Matthew.** 2013. “Duration Dependence and Labor Market Conditions: Evidence from a Field Experiment.” *The Quarterly Journal of Economics* 128 (3): 1123-67.
- Lalive, Rafael, Camille Landais, and Josef Zweimuller.** 2015. “Market Externalities of Large Unemployment Insurance Extension Programs.” *American Economic Review* 105(12): 3564-96.
- Miller, Amalia R., and Carmit Segal.** 2012. “Does Temporary Affirmative Action Produce Persistent Effects? A Study of Black and Female Employment in Law Enforcement.” *Review of Economics and Statistics* 94 (4): 1107-25.
- Miller, Conrad.** 2017. “The Persistent Effect of Temporary Affirmative Action.” *American Economic Journal: Applied Economics* 9 (3): 152-90.
- OECD.** 2017. *Employment Database-Labour Market Policies and Institutions*. (Paris: Organisation for Economic Cooperation and Development).
- OECD and ILO.** 2020. *Helping Disadvantaged Youth: Progress and Policy Action towards the Antalya G20 Youth Goal*. (Paris: Organisation for Economic Cooperation and Development).
- Oreopoulos, Philip, Till Von Wachter, and Andrew Heisz.** 2012. “The Short- and Long-Term Career Effects of Graduating in a Recession.” *American Economic Journal: Applied Economics* 4 (1): 1-29.
- Pallais, Amanda.** 2014. “Inefficient Hiring in Entry-Level Labor Markets.” *American Economic Review* 104 (11): 3565-99.

- Saez, Emmanuel, Benjamin Schoefer, and David Seim.** 2019. "Payroll Taxes, Firm Behavior, and Rent Sharing: Evidence from a Young Workers' Tax Cut in Sweden." *American Economic Review* 109(5), 1717-1763.
- Skedinger, Per.** 2014. "Effects of Payroll Tax Cuts for Young Workers." in eds. Michael Rosholm and Michael Svarer, *Nordic Economic Policy Review: Consequences of Youth Unemployment and Effectiveness of Policy Interventions* 1, Norden.
- Yagan, Danny.** 2019. "Employment Hysteresis from the Great Recession." *Journal of Political Economy* 127(5): 2505-2558.


Figure 1: Preferential Payroll Tax rate for Young Workers




Notes: The figure depicts the normal payroll tax rate (solid line) and the lower rate for young workers (dashed line) in Sweden over time. The payroll tax rate in Sweden applies to the totality of earnings and is nominally fully paid by employers. The first reform lowered the payroll tax rate for earnings received on or after July 1, 2007 for all workers turning 19 to 25 during the calendar year down to 21.3 percent. The second reform further lowered the tax rate down to 15.5 percent for earnings received on or after January 1, 2009 for all workers turning 26 or less during the calendar year. The reform was repealed in three steps on May 1, 2015, August 1, 2015, and on June 1, 2016. In the first step on May 1, 2015 (not depicted on the figure), which lasted 3 months only, the tax cut was repealed for workers turning 26 in 2015 (and was actually further lowered for workers turning 23 or less to be budget neutral). In the second step on August 1, 2015, which lasted 10 months, all workers turning 25 and less in the calendar year had their taxes increased to 25.46 percent (63% of the tax rate gap was closed). In the third step on June 1, 2016, the normal payroll tax rate applies to all workers. As part of the response to the COVID crisis, a new and temporary youth payroll tax cut starts on January 1st, 2021 (set to expire after 2.25 years on March 31, 2023). For workers turning 19-23 in the calendar year, the payroll tax is cut down to 19.73 percent as depicted on the figure. For workers turning 15-18, the payroll tax is cut even more down to 10.21 percent. The 2021 payroll tax cut applies only up to a cap of 25,000 SEK of monthly earnings (approximately \$3,000).

Figure 2: The Effect of the Payroll Tax Cut on Wages

(a) Monthly net wage (wage earnings net of the payroll tax)

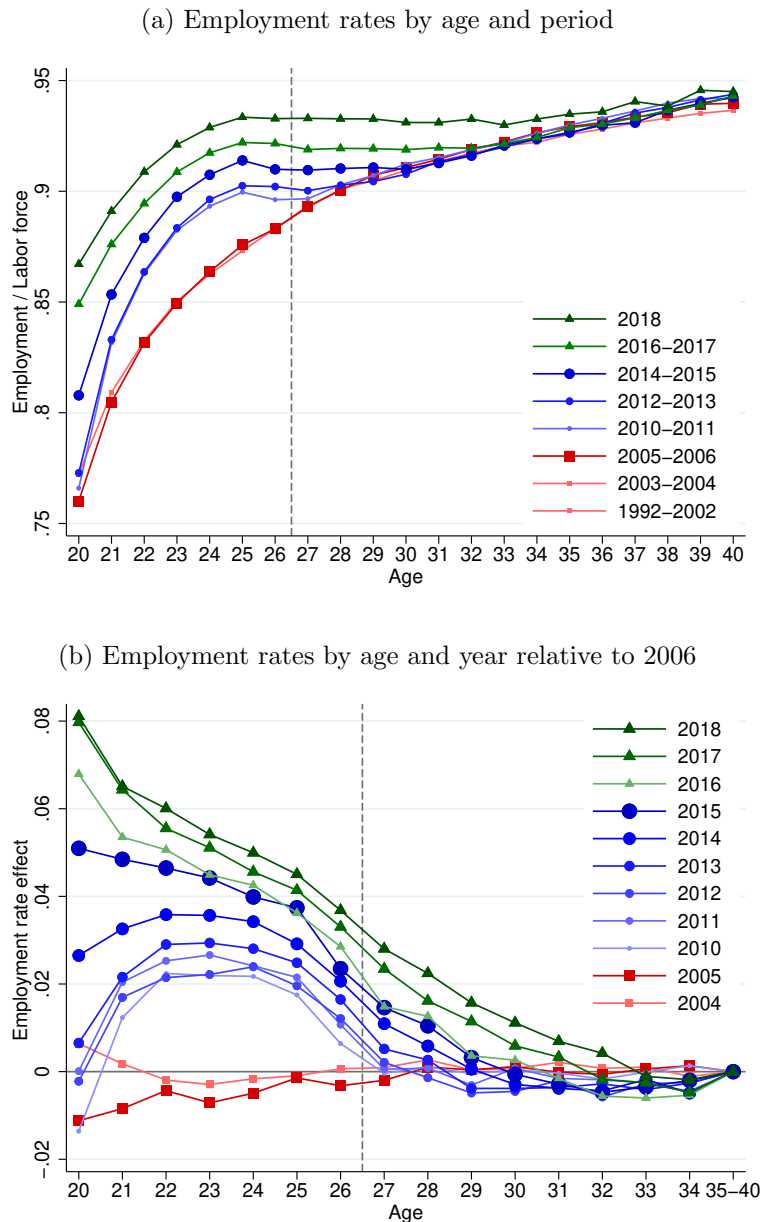


(b) Monthly gross wage (wage earnings gross of the payroll tax)



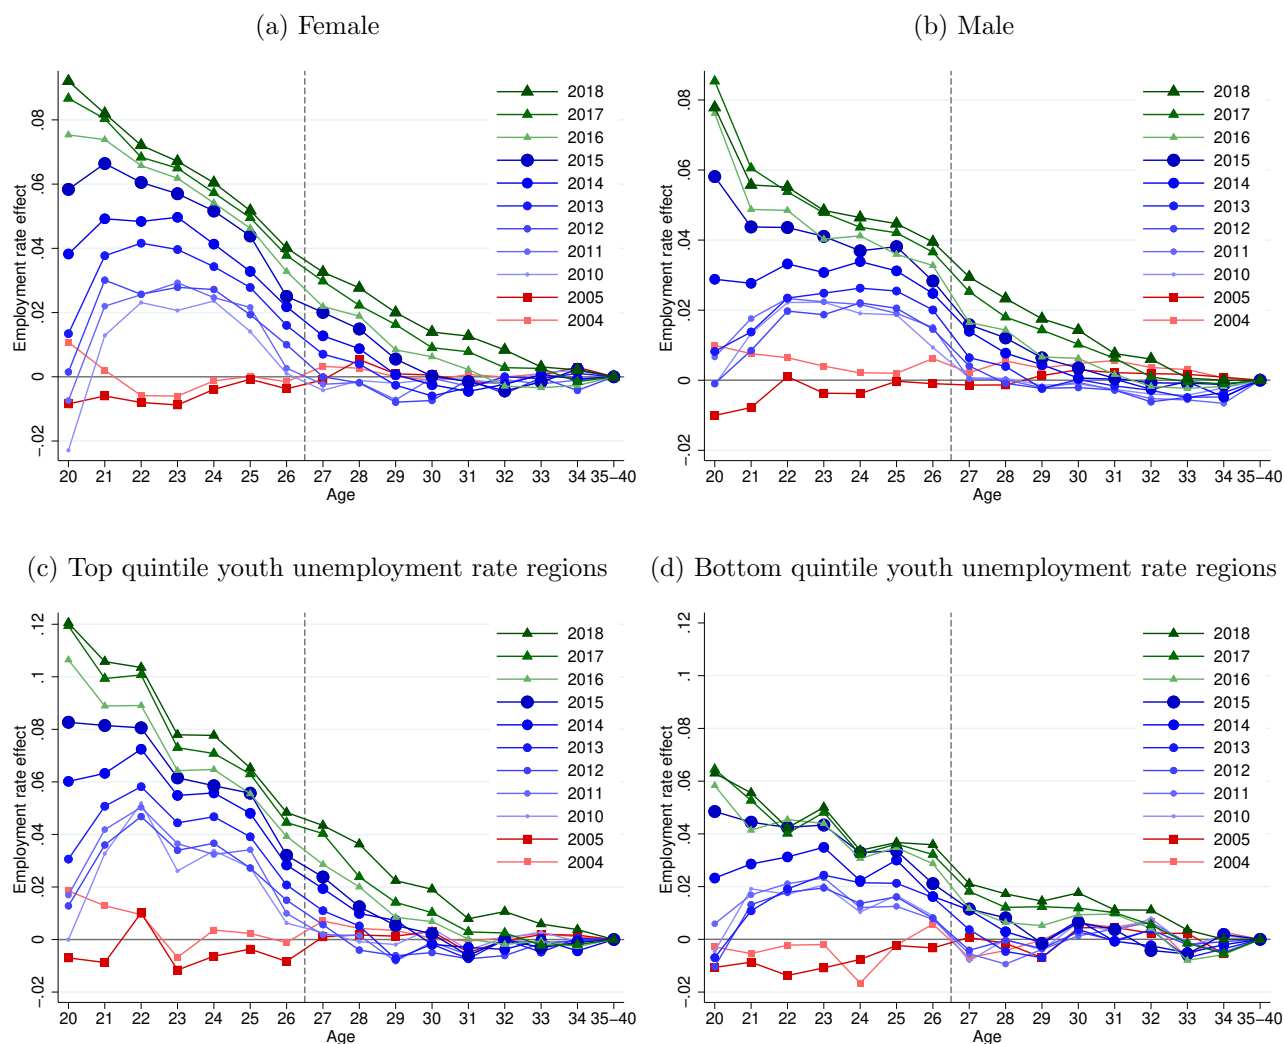
Notes: This figure depicts the average monthly wage in Sweden by age for different time periods using the Structure of Earnings Survey data. The survey measures wages mostly for the month of September (with some measurements in October and November). The top panel depicts *net wages* defined as monthly wage earnings net of payroll taxes. The bottom panel depicts *gross wages* defined as monthly wage earnings gross of payroll taxes. The wage is defined as the full-time equivalent contracted monthly wage. It is adjusted for inflation (CPI base-year 2018) and converted to US dollars using an exchange rate of 9.8 SEK/USD (as of 9/30/2019). Age is defined as the age reached during the calendar year, which is the relevant concept for the payroll tax cut. 2003-6 are pre-tax cut years (red square series). 2007-15 are tax cut years (blue circle series). 2016-2018 are post-repeal years (green triangle series). The reform applies up to age 26 in 2009-14 and up to age 25 in 2007-8 and 2015 (as depicted by the two dashed vertical lines, see Figure 1 for details). The sample includes all employed individuals in the Structure of Earnings Survey, which covers all industrial sectors (see Section 2.2). The top panel shows that net wages are continuous at the age thresholds and the bottom panel shows that gross wages are discontinuous. This implies that employers do not adjust wages by age in response to the payroll tax cut and its subsequent repeal. Hence, young workers become cheaper to employers during the reform.

Figure 3: The Effect of the Payroll Tax Cut on Employment




Notes: The top panel depicts the employment rate by age and time periods. The employment rate is the employment to labor force ratio. The employment numerator is all residents employed with annual wage earnings above a small annual threshold (\$4,490 in 2012 and adjusted for median wage growth in other years). The labor force denominator is defined as all residents who are either (i) employed as just defined for the numerator; (ii) unemployed defined as having registered with the Unemployment Office at any point during the year. The bottom panel shows the employment rate by age and single years relative to 2006. The normalization (detailed in Footnote 9) is made by aligning multiplicatively the unemployment rate (one minus the employment rate) for ages 35-40 to the 2006 level for each year and then taking the difference in employment rates with 2006. The figure shows a strong and increasing positive effect of the reform on the employment rate of young targeted workers. The effect does not diminish after the reform is repealed (2016-2018). The figure also shows an increase in the employment rate of workers exposed to the reform after they aged out consistent with a hysteresis effect.

Figure 4: Employment Rate Effects by Gender and Local Unemployment Rate




Notes: This figure shows the employment rates by age and year relative to 2006 by gender (top panels) and by local youth unemployment rate (bottom panels). The methodology is the same as in Figure 3, bottom panel, but splitting the sample by gender (top panels) and by regions based on their youth unemployment rates in 2006 (bottom panels). For bottom panels, the 21 Swedish regions are divided into quintiles weighted by the number of young in the labor force in 2006 so that each quintile includes roughly the same number of people. The normalization (detailed in Footnote 9) is made by aligning multiplicatively the unemployment rate (one minus the employment rate) for ages 35-40 to the 2006 level for each year and then taking the difference in employment rates with 2006 (raw employment rate series by gender and local unemployment areas are presented in Appendix Figure A3). The scale on the female vs. male panels (and for low vs. high youth unemployment rate) is the same for comparison. Panels (a) and (b) show a similar effect of the reform by gender in early years (2010-2012), but the effect on young female workers grows much more in subsequent years and remains much higher for females after the tax cut is repealed. Panels (c) and (d) show a much larger effect of the reform in high youth unemployment areas both during the reform and after the repeal, and also for people who age out of the reform. Hence, hysteresis effects appear much stronger in high unemployment areas.

Figure 5: Employment Rate By Age over Business Cycles




Notes: This figure depicts the employment rates by age groups from 1993 to 2018. The employment rate is defined as the employment to labor force ratio as in Figure 3. The young groups (ages 20-22, 23-24, 25-26) treated during the reform are depicted in shades of red from lighter to darker, the hysteresis groups (aged 27-28, 29-30, 31-32, 33-34), who age out of the reform in our time frame, are depicted in shades of green from lighter to darker, and the older control groups (ages 35-36, 37-38, 39-40), who are never treated, are depicted in shades of blue from lighter to darker. The figure shows that employment rates for the young fluctuate more over the business cycle.

Figure 6: Employment Effects Correcting for the Business Cycle



Notes: Panel (a) depicts the sensitivity of employment rates to the business cycle by age over the pre-reform period 1993-2006. The depicted coefficients $\hat{\phi}_a$ by age a are obtained from a regression of employment rate e_{at} at the age $a \times$ year t level on age dummies interacted with E_t (the overall employment rate across all ages 20-40, which is our business cycle indicator) and also including year and age dummies. The interaction for age 40 is omitted, which normalizes this coefficient to zero. Conventional OLS standard errors are reported. Youth employment is much more sensitive to the business cycle, consistent with Figure 5. Panel (b) depicts the business cycle adjusted youth employment rate effects. We compare the young (aged 20-26) to the control group (aged 35-40). We regress the employment rate on age dummies, year dummies, a linear control for the $\hat{\phi}_a$ coefficients depicted in Panel (a) interacted with year dummies, and year dummies interacted with an indicator for being in the treated age group (20-26). We estimate standard errors using a bootstrap method described in the main text. The figure shows excess (business cycle adjusted) youth employment after the reform starts, which persists after the reform is repealed. Therefore, the employment effects during the reform and after its repeal are robust to reform-unrelated business cycle effects on youth employment.

Figure 7: Share of Youth and Gender Discriminatory Job Ads



Notes: This figure shows the share of all (4.9 million) vacancies posted on the Platsbanken platform of the Swedish Employment Service, that are discriminating either against the young or by gender using a text search for phrases associated with discrimination by youth or gender (see Appendix for details, and Appendix Figure A4 for time series by phrase). We restrict attention to private sector jobs (around 65% of all ads), but our results are robust to this restriction. Youth discrimination falls during the reform and does not bounce back after repeal (while gender discrimination, as a placebo test, is stable).

Table 1: Direct and Indirect Effects of Payroll Tax Cut on Employment

Periods	Age Groups					
	<i>Directly Treated</i>		<i>Spillovers</i>			
	20-26	27-28	29-30	31-32	33-34	20-34
Panel A: Treatment Effects on Employment Rate in Percentage Points						
2003-4 <i>Placebo</i>	0.543 (0.395)	0.300 (0.238)	0.116 (0.242)	0.171 (0.286)	0.011 (0.315)	
2010-13 <i>Medium Run</i>	2.316 (0.339)	0.204 (0.183)	-0.292 (0.215)	-0.251 (0.201)	-0.321 (0.263)	
2014-15 <i>Long Run</i>	4.352 (0.346)	1.340 (0.218)	0.161 (0.234)	-0.196 (0.237)	-0.173 (0.284)	
2016-18 <i>Post-Repeal</i>	6.272 (0.421)	2.473 (0.230)	1.213 (0.227)	0.484 (0.210)	-0.077 (0.251)	
Panel B: Decomposition: Job-Year Counts (and Share of Total)						
2010-13 <i>Medium Run</i>	52,544	1,436	-2,135	-2,014	-2,662	47,169 (20.6%)
2014-15 <i>Long Run</i>	49,359	4,724	591	-783	-716	53,173 (23.2%)
2016-18 <i>Post-Repeal</i>	106,723	13,078	6,657	2,907	-479	128,886 (56.2%)
All Years	208,625 (91.0%)	19,237 (8.4%)	5,113 (2.2%)	110 (0.0%)	-3,857 (-1.7%)	229,228 (100%)

Notes: Panel A presents the effects of the payroll tax cut on employment rates in percentage points by periods (by row) and by age groups (by column) using the aggregated times series by age and year displayed in Figure 3 bottom panel. We omit the phase-in period and the global recession year (2007-9), so that we have thirteen years (2003-6 and 2010-18) and 21 ages (20-40), rendering 273 observations in total. We first divide our “treated” ages into groups as follows: 20-26; 27-28; 29-30; 31-32 and 33-34. We then regress the outcome variable (employment) on year dummies, age dummies and interactions between the five treatment-age-groups and time periods, defined as in the rows (excluding 2005-6 as the omitted category). The table reports coefficients on the interactions. We report robust OLS standard errors. The treatment effect provides an average effect by age groups and periods using the age distribution of the labor force in 2006. Labor force (LF) is defined as all residents who are either (i) employed with annual wage earnings above a small annual threshold (\$4,490 in 2012 and adjusted for median wage growth in other years); or (ii) unemployed (defined as having registered with the Unemployment Office at any point during the year). Panel B computes the reform-induced creation of job-years. We take the age-group \times period treatment effects γ_{AT} from DiD regression model (1) reported in Panel A. We multiply each γ_{AT} with the labor force body count of age group A in the pre-period 2006, i.e. $\Delta E_{AT} = \gamma_{AT} LF_{A,2006}$. In the bottom line of each entry, we report the share of a given entry contributing to total job creation is $\Delta E_{AT} / \Delta E_T$. The last row contains age-specific job creation over all years $\Delta E_A = \sum_T \Delta E_{AT}$. The rightmost column presents total time-specific job creation over all cohorts $\Delta E_T = \sum_A \Delta E_{AT}$. The bottom right entry reports total overall employment effects $\Delta E = \sum_A \sum_T \Delta E_{AT}$. Since our time periods combine multiple years, we also multiply a given time period’s treatment effect by its length in years (4,2,3).

Online Appendix of: Hysteresis from Employer Subsidies

By Emmanuel Saez, Benjamin Schoefer, and David Seim

A Classifying Job Listings by Hiring Preferences

Text search procedure. From the full vacancy data set, which includes around 4.9 million posted vacancies over the period 2006-2017, we extracted a random sample of 3,000 ads. Two research assistants independently classified these 3,000 ads as discriminatory against young, and by gender (men or women). From the set of ads marked by at least one reader as discriminatory, we determined discriminatory phrases. We then text-searched the full set of vacancies for these phrases, and reviewed a random sample of 1,000 ads classified as discriminatory to determine false positives. We iterated this procedure until such false positives became negligible. The final round of this procedure yielded our partitioning of the full sample of vacancies into discriminatory (either by gender, against youth, or both) and non-discriminatory.

Youth phrases. We exclude nondiscretionary age restrictions imposed by legal constraints (e.g., taxi drivers with minimum legal ages). The set of phrases (and their English translations) used to classify jobs as discriminatory against youth (where “\d” denotes digits) are, in phrase-specific time series, depicted in Appendix Figure A3b) and listed below:

- “ha(r) fyllt 2[1-9] år” – english translation: “must be above 2[1-9] years of age”.
- “mellan (1[8-9][2-5][0-9]6[0-4]) och \d \d år” – “between (1[8-9][2-5][0-9]6[0-4]) and \d \d years”.
- “(M|m)inst 2\d år” but not “i minst 2 \d år” – “At least 2\d years” but not “for at least 2 \d years”, to avoid strings such as “You have worked for at least 20 years”.
- “(D|d)u (som söker)?är (man|kvinna)?i åldern 2 \d(- \d \d)?” or “(N|n)ågon i åldern 2 \d(- \d \d)?” – “You (who are applying) are (a man/woman) in ages 2 \d (- \d \d)” or “Someone in ages 2\d (- \d \d)”.
- “(Ä|ä)ldre sökande” but not “(Y|y)ngre” – “Older applicants” but not “Younger”, to avoid strings such as “We are looking equally for younger and older applicants”.
- “är [2-6] \d \d+” – “are [2-6] \d \d+”.
- “(M|m)inimiålder [2-5] \d år” – “Minimum age [2-5] \d years”.
- “2 \d-[2-6] \d årsåldern” – “Agerange 2 \d-[2-6] \d”.

- “2 \d år fyllda” – “Age 2 \d or more”.
- “2[1-9] - [2-6] \d år” – “2[1-9] - [2-6] \d years”.

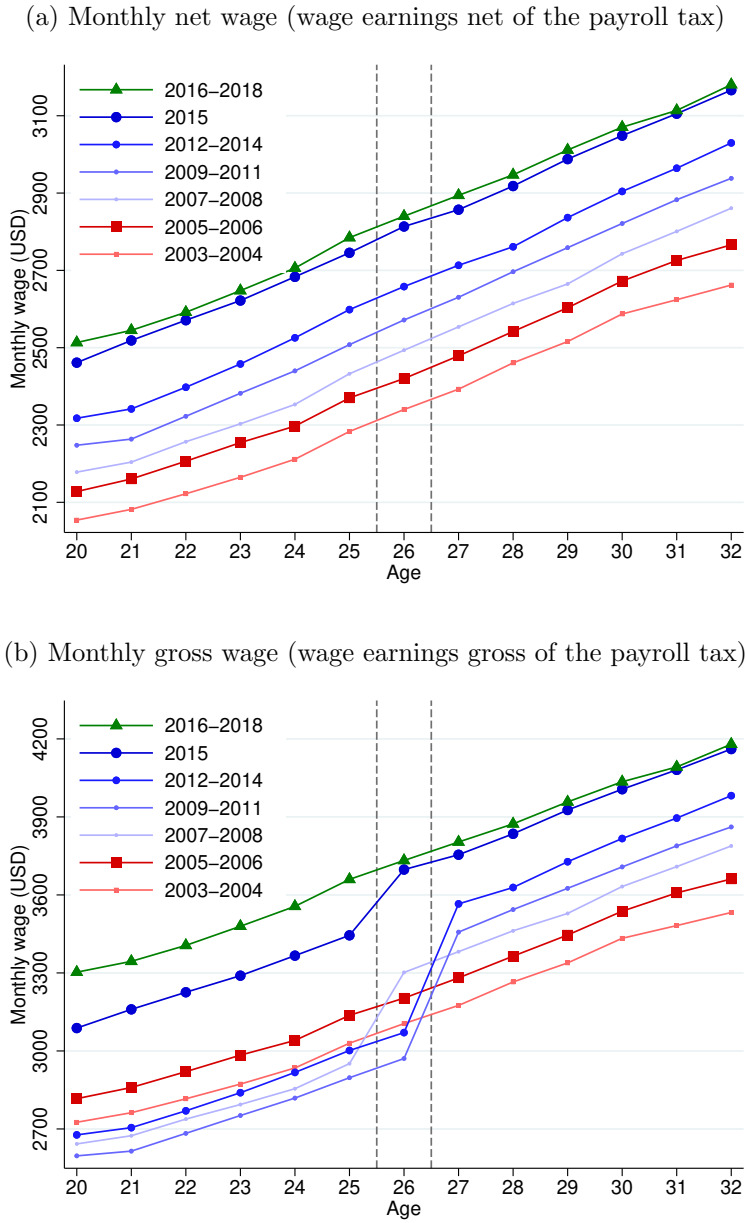
Gender phrases. For gender discrimination, we search for ads that are directed particularly to men or women. The list of terms (phrase-specific time series are depicted in Appendix Figure A3b) is below:

- “(G|g)ärna (kille|tjej)” – “Preferably man/woman”.
- “(D|d)u är (en)?(kvinna|man|kille|tjej)” but does not contain “((K|k)vinna|(M|m)an|(K|k)ille|(T|t)tjej) eller (kvinna|man|kille|tjej)” – “You are a man/woman/guy/girl” but not “man/woman/guy/girl or a man/woman/guy/girl”, to avoid strings such as “You are a man or a woman”.
- “(S|s)öker (en)?(kvinna|kvinnlig|man|kille|tjej)” but does not contain “((K|k)vinna|(M|m)an|(K|k)ille|(T|t)tjej) eller (kvinna|man|kille|tjej)” – “Looking for a man/woman/guy/girl” but not “man/woman/guy/girl or a man/woman/guy/girl”.
- “händig (kvinna|man|kille|tjej)” – “Handy man/woman/guy/girl”.

However, many ads state that they seek to hire a man/woman to make the gender distribution more balanced. We do not label those ads as discriminatory and identify them by including the strings “jämlighet” – “equality”; “jämställdhet” – “equality”; “Könsfördelning” – “Gender distribution”; “Jämn fördelning” – “Equal distribution”; “Mångfald” – “Diversity”; “Mansdominerad”/“Kvinnodominerad” – “Male dominated”/“Female dominated”; “Dominerad av män”/“Dominerad av kvinnor” – “Dominated by men”/“Dominated by women”.

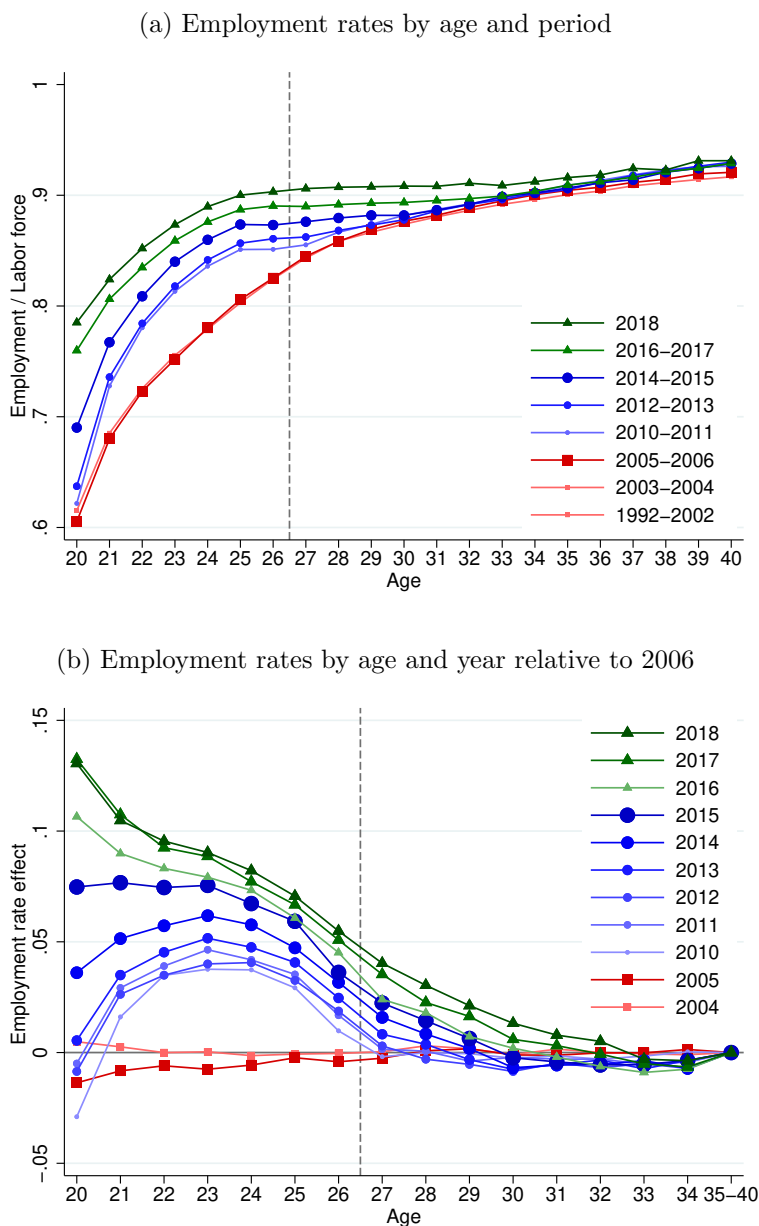
Additional Figures

Figure A1: The Effects on Wage for New Hires/Job Switchers



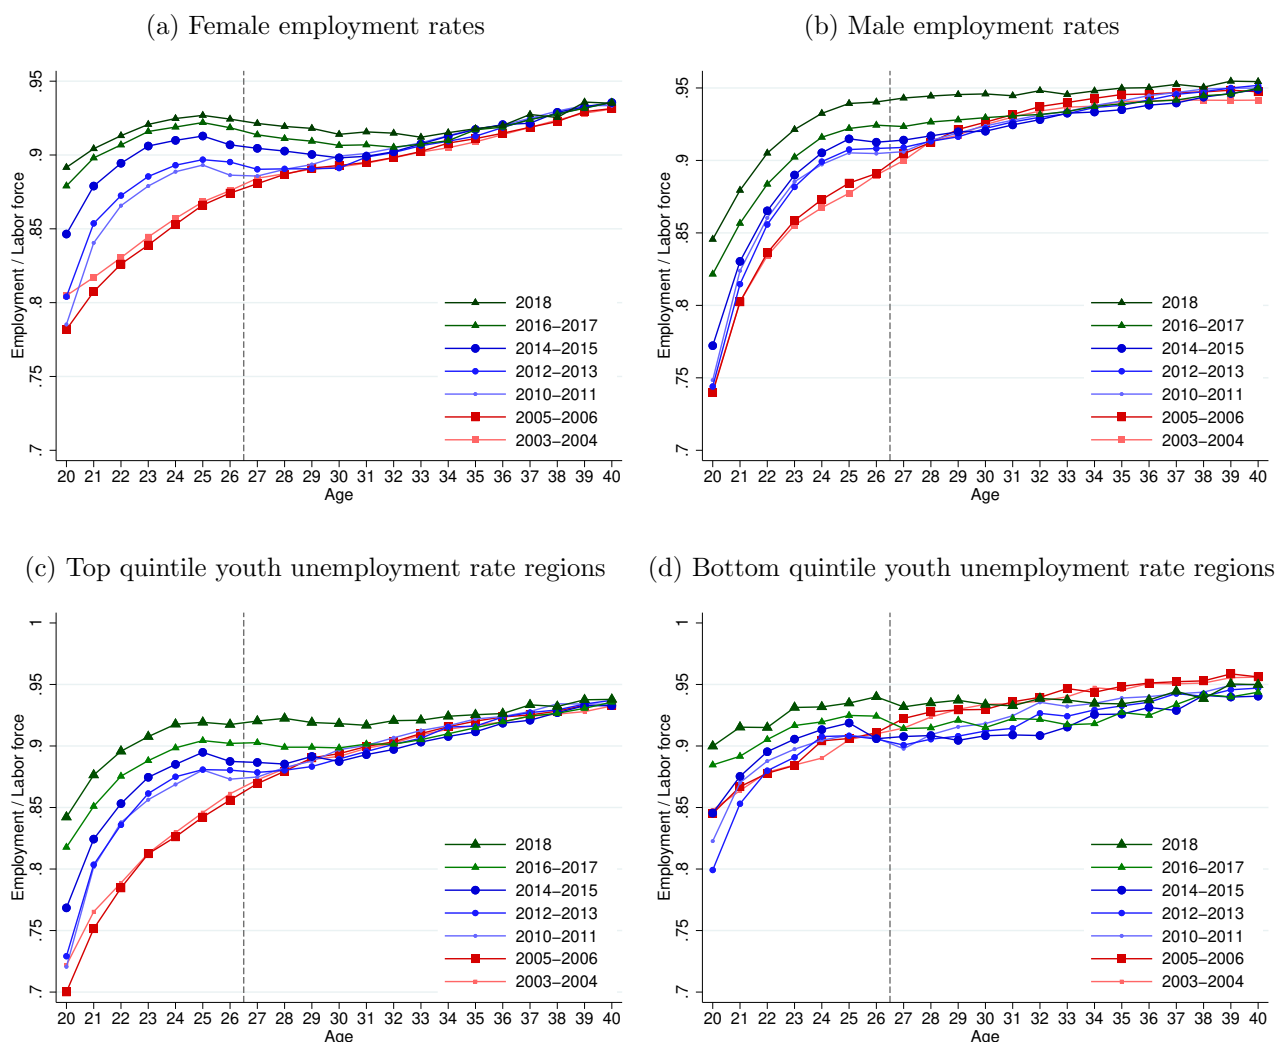
Notes: This figure repeats the wage statistics displayed in Figure 2 but limiting the sample to new hires or job switchers, defined as having a new firm identifier as the main (i.e., highest paying) employer relative to September of the previous year. It includes both job-to-job transitions as well as new hires among the nonemployed. As we found in Figure 2, there is no discontinuity in net wages (top panel) and a corresponding discontinuity in gross wages due to the tax differentials. This implies that employers do not pass on the payroll tax cut or its abolition to young workers (relative to older workers).

Figure A2: Robustness of Employment Effects to using a Twice Higher Earnings Threshold





Notes: This figure reproduces Figure 3 but using an annual earnings threshold twice as high in our definition of employment. The top panel depicts the employment rate by age and time periods. The employment rate is the employment to labor force ratio. The employment numerator is all residents employed with annual wage earnings above an annual threshold set twice as high as in our main text Figure 3 (\$8,980 in 2012 and adjusted for median wage growth in other years). The labor force denominator is defined as all residents who are either (i) employed as just defined for the numerator; (ii) unemployed defined as having registered with the Unemployment Office at any point during the year. The bottom panel shows the employment rate by age and single years relative to 2006. The normalization (detailed in Footnote 9) is made by aligning multiplicatively the unemployment rate (one minus the employment rate) for ages 35-40 to the 2006 level for each year and then taking the difference in employment rates with 2006. Like Figure 3, this figure shows a strong and increasing positive effect of the reform on the employment rate of young targeted workers. The effect does not diminish after the reform is repealed (2016-2018). The figure also shows an increase in the employment rate of workers exposed to the reform after they aged out consistent with a hysteresis effect.

Figure A3: Employment Rates by Gender and Local Youth Unemployment Rates



Notes: The figure depicts the employment rate by age and time periods by gender (top panels) and by local youth unemployment rate (bottom panels). The methodology is the same as in Figure 3, top panel, but splitting the sample by gender (top panels) and by regions based on their youth unemployment rates in 2006 (bottom panels). In 2006, just before the reform, there was wide variation across Sweden's 21 regions in youth unemployment. Regions in the lowest quintile of youth unemployment rates (left panel) had rates in the range of 10.5-12.4% while regions in the highest quintile (right panel) had youth unemployment rates in the range of 20-23.3%, i.e. about twice as high. Regions are divided into quintiles weighted by the number of young in the labor force in 2006 so that each quintile includes roughly the same number of individuals. The employment rate is the employment to labor force ratio. The employment numerator is all residents employed with annual wage earnings above a small annual threshold (\$4,490 in 2012 and adjusted for median wage growth in other years). The labor force denominator is defined as all residents who are either (i) employed as just defined for the numerator; (ii) unemployed defined as having registered with the Unemployment Office at any point during the year. The top panels show a strong and positive effect of the reform on the employment rate of young targeted workers. The effect appears much stronger for females than for males especially in later years. The panels also show an increase in the employment rate of workers exposed to the reform after they age out consistent with a hysteresis effect for both genders. The bottom panels show a much larger effect of the reform in high youth unemployment regions both during the reform and after the repeal, and also for people who age out of the reform.

Figure A4: Discrimination in Posted Vacancies Over Time



Notes: This figure shows the share of all vacancies posted on the website Platsbanken, which is the online job portal provided by the Swedish Employment Service, that are discriminating against the young (Panel (a)) or on the basis of gender (Panel (b)). Discrimination is defined as follows (detailed in Appendix A). Out of the 4.9 million job ads that were posted during the years 2006-2017, two research assistants independently read and coded a random sample of 3,000 ads as discriminatory against young or against men or women. We then matched the two samples of discriminatory ads, resolved inconsistencies and selected discriminatory phrases. We then extracted a random sample of 1,000 ads contained the phrases, and refined the phrases until the share of false positives was negligible. This graph shows the share of ads that includes a given phrase. “\d” denotes digits for age-related phrases.