Does Employer-Provided Health Insurance Exacerbate Age Discrimination in Labor Markets?

Principal Investigator: Kurt Lavetti, Ohio State University and IZA *[Website]* *[CV]**
Co-Investigator: Jonathan Holmes, NBER *[Website]* *[CV]*

Overview

In 2019, the average employer in the US paid $14,561 per employee for family health insurance coverage, which is more than 30% of the median annual earnings of full-time workers. The high cost of insurance relative to earnings gives US employers (that offer health insurance) strong incentives to hire workers with lower expected healthcare costs, all else equal. A worker's age is one of the strongest and most easily observed predictors of healthcare costs available to employers at the hiring stage. In 2014, an average person aged 60-64 in the US spent almost four times as much on healthcare as a 25-29 year old, and 2.4 times as much as a 40-44 year old. Although it is illegal in the US to discriminate against a worker on the basis of their age or disabilities, most employers face clear incentives to find legal reasons to prefer younger workers with lower average healthcare costs.

In this project we will quantify the component of age-based discrimination in US labor markets that operates through health channels. Although age-based discrimination can take many forms, the rapidly growing share of compensation devoted to health insurance costs has potentially made health status a dominant motivation for preferring to hire younger workers. To measure how health costs are related to employment, we use a unique new database that links individual UI-derived earnings records to administrative health insurance claims for all workers in Utah. We will use these data to implement two distinct research designs. The first uses aging into Medicare eligibility combined with discontinuous firm size cutoffs that mandate employees in small businesses drop their employer provided insurance and enroll in Medicare. The second uses firm size cutoffs in access to ACA small business exchange marketplaces, which allow firms to purchase insurance at prices that are independent of their employees' health status. With each of these research designs we answer the following questions:

1) What are the effects of the age-based component of healthcare costs on hiring rates, separation rates, earnings levels, relative earnings growth rates within firms?

2) To what extent do these effects on labor market outcomes depend on individual health status, particularly among elderly workers?

3) Among elderly workers eligible for Medicare (65+), does enrolling in Medicare instead of employer-sponsored insurance attenuate the magnitude of these labor market effects?

Significance

Since the mid-1980s labor force participation among the elderly (65+) in the US nearly doubled (prior to Covid-19). But older workers face a more divergent set of labor market conditions than younger workers on average. While older workers are substantially less likely to be displaced, those that are displaced face greater difficulty finding new jobs and are more likely to have long unemployment spells. These patterns are consistent with a combination of competing forces: older workers are on average more protected by their relative seniority within firms, but once displaced they face resistance to being rehired into comparable jobs at pre-displacement earnings levels. As a result average wage losses associated with job displacement are more than ten times larger for older workers (62+) than they are work workers aged 25-34 (Johnson & Mommaerts, 2011).

Understanding the labor market experiences of older workers depends heavily on understanding what causes older workers to be treated differently than other workers at both the separation and hiring margins. On the hiring margin, Lahey (2008) finds that resumes with lower randomly assigned ages were more than 40% more likely to lead to interview requests. This age disparity is more extreme among larger employers, which tend to offer more generous health insurance and other benefits. It remains unclear, however, what mechanism drives these results. They could, for example, be explained by productivity differences combined with wage rigidities, statistical discrimination, or perceived animus by employers and or customers to elder workers. This project aims to help clarify the mechanisms behind these patterns by testing whether difference in average insurance-based compensation between large and small employers contribute to age-based hiring disparities and help explain why elderly workers are over-represented at small businesses.

On the separation margin, Johnson and Mommaerts (2011) find that although older workers are unconditionally less likely to be displaced, they are actually significantly more likely be displaced conditional on job tenure. By bundling health insurance with jobs and making it illegal for wages to adjust with the cost of insurance for different workers, the employer-based healthcare system in the US creates a uniquely large incentive to discriminate based on age.

Discriminatory behavior in labor markets has long been a topic of interest among economists, but identifying evidence of discrimination and its causes using observational data on labor market outcomes is notoriously difficult. Studies of this type generally test whether wage gaps across groups remain after attempting to control for differences in marginal productivity, and interpret residual gaps as evidence of discrimination. A pervasive concern is that measuring marginal productivity accurately is very difficult. If productivity is imperfectly measured, and if groups differ in characteristics that are correlated with productivity, then including any of these characteristics in the model of equilibrium wage determination introduces endogeneity. Field-based experimental studies, often structured as audit studies or correspondence studies, circumvent this problem by experimentally varying labor

supply characteristics to evaluate aspects of the hiring process.[1] These types of experimental approaches have clear advantages in identifying discriminatory behavior, but by their nature in focusing on pre-employment applicant selection they yield less information about the implications of potential discrimination on a broad set of equilibrium labor market outcomes.

Fewer studies have used natural experiments to quantify discrimination. For example, Goldin and Rouse (2000) study the consequences of switching to blind auditions for orchestras on gender discrimination, and Price and Wolfers (2010) use variation in the racial composition of NBA officiating crews to test for racial discrimination in foul calls. Our proposal fits into this category in that we will use policy-driven variation that creates discontinuities in firms' cost functions associated with hiring older workers. Rather than attempting to control for the age-productivity profile of workers, we will compare workers of the same age who are employed at comparable firms that face different health insurance cost functions, where the cost functions are being shifted nonlinearly by policy variation.[2]

The study will consist of two distinct components with different research designs. The two approaches address similar questions and use the same data and study populations, but rely on different policies that shift the cost of employer provided health insurance in different ways, and for workers at different age margins. In *Research Design One* we use variation caused by federal regulations that mandate workers employed at firms with 20 or fewer workers must enroll in Medicare when they turn 65, in contrast to workers at large firms who can choose between Medicare or employer-provided insurance. In *Research Design Two* we use nonlinear variation and kinks in the age-premium function for health insurance caused by the creation of small business exchanges in the Affordable Care Act, along with firm size discontinuities in access to this marketplace.

Data

This project will use the Utah All-Payer Claims Dataset (APCD) from 2013-2019 linked to administrative earnings records constructed from Utah UI records. We currently have access to a short panel of this database, from 2013-2015, and this grant will be used in part to extend the panel to 2019, which will provide several important benefits to the research designs that we discuss below.

The APCD contains information from nearly all commercially insured individuals in Utah. Insurers are required by state law to submit their administrative enrollment records and medical, drug, and dental claims to the state, who constructs this database for research purposes. The APCD includes administrative health insurance enrollment records, demographics (gender, dates of birth, location of residence), and household structures. For individuals

[1] See Neumark (2018) and Bertrand and Duflo (2017) for a review of the literature on these field experiments.

[2] To be clear, we use the term discrimination with respect to health status or healthcare costs in a legal sense. Unlike taste-based discrimination on dimensions that have no impact on profit functions, this type of discrimination may be viewed differently. Still, the use of the term discrimination in this context is consistent with the literature on gender wage gaps, which typically estimate discrimination effects on wages that include average gender differences in employer-sponsored healthcare costs as a component.

(and dependents) whose insurance coverage comes from employers, we observe employer identifiers, group policy numbers, group sizes, whether the employer self-insures, and whether the individual has a health savings account. In most cases we can also impute the employer menu of insurance plan choices and plan design features, such as deductibles and cost-sharing rules. The records also include all medical claims, diagnoses, prescriptions filled, costs, member liabilities.

The APCD is linked by anonymized patient identifiers to a customized dataset containing earnings records that was produced by the Utah Department of Workforce Services. This file contains information on quarterly earnings for every individual in the state in every quarter between 2002-2019, as well as information about the labor market network structure, such as which sets of workers are employed by the same firm. To our knowledge this is the first large-scale administrative database available to researchers in the U.S. that combines detailed administrative health records with earnings and employment records.

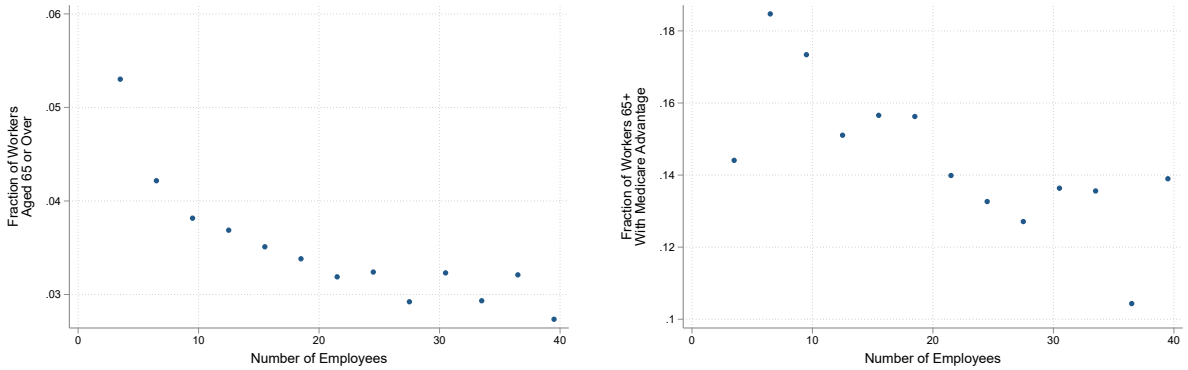
Utah is a clean case study for both natural experiments because of its light regulatory policies. Prior to the ACA, the price of health insurance for small businesses was minimally regulated, and premiums could vary substantially across small businesses based on employee health status. The ACA had a more pronounced effect on the small employer insurance market in Utah relative to states with stronger pre-ACA price controls or community rating.

Although we have a few years of APCD data already, one limitation in working with a short panel is that we can only observe the ages of workers who appear in the APCD at some point. Workers who have periods of uninsurance may appear in the earnings records but not in the APCD. Most of these workers are either part-time, or they are employed at firms that do not offer insurance. This latter group is important for our analyses, since older workers, especially those eligible for Medicare, may be more likely to be employed in firms that don't offer insurance. To include these individuals in our analyses we need to observe them at least once in the APCD. A longer panel will greatly improve our ability to do this, especially as older workers eventually join Medicare plans and appear in the APCD. Additional years of data will also improve the match quality in the linkage between earnings and health records. Although the match rate is high, there are individuals in each database that do not appear in the other between 2013-2015.

Research Designs

Research Design One: Our first design measures the impact of Medicare eligibility on the employment outcomes of aging workers. Upon turning 65, workers become eligible to enroll in Medicare, but employees at large firms also have the option to stay in their employer-sponsored plan while they are working. Employees that do switch to Medicare generate substantial savings for their employers: single-worker health benefits represent 21% of the median salary of a 55-64 year-old worker, and family plans represent a larger share. In this section, we ask whether

Figure 1: Preliminary Regression Discontinuity Figures



(a) Fraction of Employees Age 65+

(b) Fraction of Employees 65+ in Medicare Advantage

firms are more willing to hire and retain elderly workers once they are no longer required to provide them health insurance.

Our research design uses differences in how firms are treated depending on how many employees they have. By law, all employees at firms with fewer than 20 workers (including part-time workers) must enroll in Medicare as their primary insurer as soon as they turn 65 (Medicare has sizable financial penalties for failing to enroll after turning 65). In contrast, workers at firms with 20 or more workers can choose between keeping their employer-provided insurance or switching to Medicare. Most workers prefer to keep their employer-provided insurance, which typically provides more generous coverage at lower out of pocket costs than Medicare. Empirically, among workers aged 65-74 who are offered employer health insurance 67% choose to keep their employer insurance (Burless, 2017).

Our preliminary analyses are consistent with differences in the age and insurance composition of employees for firms above and below 20 employees. Figure 1a shows that the share of workers who are elderly is higher in small firms. As firms grow larger they tend to hire younger workers, decreasing the share of elderly workers (future analyses will examine worker flows directly). Figure 1b shows elderly workers in firms with fewer than 20 employees are more likely be enrolled in Medicare Advantage plans than are elderly workers in larger firms.[3] In these preliminary analyses using our current short panel we were able to link ages from health records to roughly 75% of the individuals in the earnings data. We expect that by lengthening the panel this match rate will improve substantially.

We use the 20 employee threshold to measure differences in labor market outcomes with respect to hiring, retention, salary, and health insurance provision, comparing similar firms above and below the 20-employee threshold. Methodologically we will use a regression discontinuity design such as the following:

$$Y_{ft} = f_1(e_{ft} \mid e_{ft} \leq 20) + f_2(e_{ft} \mid e_{ft} > 20) + \beta \mathbb{I}(e_{ft} > 20) + \varepsilon_{ft} \tag{1}$$

[3]Further data work is required to differentiate primary and secondary insurers to improve the coding of Medicare enrollees.

where Y_{ft} includes outcomes like the share of new hires at firm f in period t who are 65+, the share of separations that are 65+, and the average wage gap between elderly (65+) and near-elderly (60-64) workers. $f_1(e_{ft} \mid e_{ft} \leq 20)$ is a flexible continuous function of employment levels in small firms, $f_2(e_{ft} \mid e_{ft} > 20)$ is a flexible continuous function of employment levels in larger firms, and β identifies the discontinuous jump in outcomes associated with crossing the 20 worker threshold. This model can also be extend to a difference in RD design by interacting every term in the model with an indicator for whether the firm offers health insurance, and including in the sample firms that never offer insurance.

A separate estimation approach that we will also test uses age 65 as a discontinuity. In this approach we define Y_{fta} to be the fraction of employees at firm f in period t who are in age group a (or hiring rates, separation rates, and average employee earnings) and estimate:

$$Y_{fta} = f_1(a \mid a < 65, e_{ft} \leq 20) + f_2(a \mid a \geq 65, e_{ft} \leq 20) + f_3(a \mid a < 65, e_{ft} > 20) + f_4(a \mid a \geq 65, e_{ft} > 20) + \beta_1 \mathbb{I}(a \geq 65, e_{ft} \leq 20) + \beta_2 \mathbb{I}(a \geq 65, e_{ft} > 20) + \varepsilon_{fta} \quad (2)$$

where $f_1 - f_4$ are flexible functions of age that may differ in small and large firms, and above or below the age 65 discontinuity. β_1 identifies the discontinuous change when workers age into Medicare eligibility in small firms, β_2 identifies a similar discontinuity in large firms, and $\beta_1 - \beta_2$ is the DRD estimate. We hypothesize that $\beta_1 - \beta_2$ will be positive for wage outcomes and hiring rates and negative for separation rates.

The richness of our dataset allows for additional robustness checks. For example, a subset of firms in our data will cross the 20-employee threshold, and we can look at hiring and retention of older workers within these firms both before and after they cross this threshold. We can also cut the analyses by health status at age 64 (and/or dependent health status) to use intensive margin variation in individual healthcare costs.

Research Design Two: The Affordable Care Act (ACA) gave firms with 2-50 employees the opportunity to purchase insurance coverage for their workers from small employer marketplaces, in which premiums are community-rated. Prior to the enactment of the ACA, most private insurance plans were risk-rated, meaning that premiums varied across firms depending on the average healthcare spending of their employees. The ACA marketplaces, which began operating in 2014, have pricing regulations that prevent insurers from charging higher premiums to firms with sicker workers.

Insurance premiums in the ACA small business exchanges are allowed to vary depending on geography, age, household composition (number and ages of dependents), and tobacco use. Importantly, they cannot vary based on health status or pre-existing conditions. The age component of premiums is also regulated and follows a federally mandated relative age-premium function. This function sets the premium charged to a 64 year old to be exactly 3 times the premium charged to a 21 year old. However, after age 64 the mandated premium function has a kink,

Figure 2: Adverse Selection into the Small Firm ACA Market, 2014

	Small Firms		Large Firms
	ACA-Compliant	Grandfathered	
Average Age	43.6	42.5	43.3
Fraction Male	0.51	0.53	0.51
Predicted ACG	1.33	1.08	1.14
Top 10% ACG in 2013	0.11	0.072	0.082
Enrollees	25542	34425	323937

Notes: This table measures average enrollee characteristics in Utah's employer-sponsored insurance market in 2014. Small firms are those with under 50 full-time equivalent workers. Among small firms, ACA compliant firms are those that purchase community-rated coverage subject to the ACA price regulations, and grandfathered plans are those that continue purchasing experienced-rated plans. ACG scores measure average predicted health spending according to the John Hopkins ACG System, where an ACG value of 1 is associated with a person with average health spending among insured individuals in Utah.

and premiums may not increase further with age. Researchers (Orsini & Tebaldi, 2017) have shown that the federal pricing function causes cross-subsidization that inflates premiums for younger workers and reduces premiums for older workers, relative to the distribution of healthcare spending. The policy change therefore reduced the relative cost to small employers of hiring older workers or other workers with high average health expenditures. In addition, it fully eliminated age-based differences in health insurance costs within the subset of workers aged 64+ employed at small firms.

We will use the 50 employee cutoff in access to community-rated insurance to estimate a differences-in-differences model to test whether decoupling premiums from employee health status:

(1) Increased the hiring and retention of older workers and workers with disabilities.

(2) Increased the wages of older workers relative to younger workers.

(3) Increased the number of older workers who receive health insurance at their place of work.

We will also explore whether there is sufficient power to conduct similar tests using the age 64 pricing kink, and using a regression discontinuity design around the 50 employee cutoff for access to the exchanges.

Specifically, we will estimate the following firm-level DiD model in which Y_{ft} is an outcome of interest (e.g. average age, age of new hires, age of separations, etc.) at firm f in quarter t:

$$Y_{ft} = \alpha_f + \gamma_t + \sum_{t \neq 2013Q4} \beta_t^S S_f + \epsilon_{ft}$$

where α_f is a fixed employer effect, γ_t are fixed quarter-year effects, and S_f is an indicator for having 50 or fewer employees. β_t^S are a set of event-time dummies measuring changes in the small-large employer age gap relative to the omitted quarter, 2013Q4. These coefficients are intent-to-treat estimates of allowing small employers to access community-rated insurance markets. We will also explore specifications that replace fixed employer effects with employer controls, such as industry effects.

In Figure 2, we conduct a preliminary analysis of selection of employees using 2013-2015 data. These summary statistics show clear evidence of adverse selection into ACA plans: employees at small firms that opted to enroll

on the ACA marketplace were 2.6% older and had 23% higher risk scores on average. In future work we will test whether this pattern is driven by selection of employers with high-cost employees into ACA plans, or whether employers in these plans also begin hiring older employees.

References

Bertrand, M., & Duflo, E. (2017). Chapter 8 - field experiments on discrimination. In A. V. Banerjee & E. Duflo (Eds.), *Handbook of field experiments* (Vol. 1, p. 309-393). North-Holland.

Burless, G. (2017). *Age related health costs and job prospects of older workers* (Tech. Rep.). Stanford Institute for Economic Policy Research. Retrieved from https://siepr.stanford.edu/system/files/BURTLESS_Age-Related-Health-Costs_1st-Draft_Oct-2017.pdf

Johnson, R., & Mommaerts, C. (2011). Age differences in job loss, job search, and reemployment. *Urban Institute Discussion Paper 11-01*.

Lahey, J. N. (2008). Age, women, and hiring: An experimental study. *The Journal of Human Resources*, *43*(1), 30–56. Retrieved from http://www.jstor.org/stable/40057338

Neumark, D. (2018). Experimental research on labor market discrimination. *Journal of Economic Literature*, *56*(3), 799-866.

Orsini, J., & Tebaldi, P. (2017). Regulated age-based pricing under the affordable care act. *Working Paper*.

Budget and Justification

1) Update to Utah All-Payer Claims Database (APCD) license: $50,000
We currently have access to APCD data covering 2013-2015, which is linked to earnings records from 2001-2019. To complete this project we will add APCD data from 2016-2019 to create a longer panel in which we can observe both healthcare claims and individual earnings. The cost structure for the Utah APCD is described at: http://stats.health.utah.gov/about-the-data/data-series/. The base price for a new application for one calendar year of data is $20,000 and each additional year of data costs $10,000 per year. The total cost for four years of data is $50,000. This amount is tentatively assumed to be excluded from a subaward with indirect costs.

2) Graduate Research Assistant: $17,827
This will cover the cost of hiring one PhD student at OSU to work on the project for one semester. This amount is based on the minimum monthly stipend rate of $2364.44 plus 11.5% fringe benefit rate for graduate students plus $5964 in-state tuition fees.

3) ACG Software License: $2,000
We will purchase a two-year license for the Johns Hopkins ACG software (www.hopkinsacg.org). This is specialized software that is designed to be run on healthcare claims records, and produces important summary variables such as patient risk scores and predicted risk scores using retrospective claims. This software is used by many commercial insurers to set risk-adjusted prices, and it provides an important enhancement to the claims data alone.

Total Direct Costs Requested: $69,827
Indirect Costs (on GRA): $2,139
Total Direct and Indirect Costs Requested: $71,966