Older Workers and Employer-Sponsored Health Insurance Premiums: Evidence From Linked Administrative and Survey Data

Owen F. Davis*

December 1, 2023

The role of employee age in employers' decisions surrounding health insurance is a topic of long-standing public-policy concern. This study quantifies the impact of workforce age structure on health insurance premiums by linking administrative earnings records from the Longitudinal Employer-Household Dynamics data set to rich plan-level data from private employers surveyed by the Medical Expenditure Panel Survey–Insurance Component. The baseline results indicate that a 10 percentage-point increase in the share of a firm's workforce ages 50 and older is associated with 4.3% higher single-coverage premiums, an estimate that varies significantly by firm size, managed-care provision, and plan funding mechanism. I further explore how the age-premium relationship changes under community rating policies that limit the degree to which insurers can vary small-group premiums according to the age and health of the workforce. Using the implementation of the Affordable Care Act's community ratings provisions in 2014 as a quasi-experimental exogenous shock, I find that the relationship between workforce age and small-group premiums is mitigated under age- and health-focused ratings policies. Finally, I examine the effects on health insurance premiums of Medicare acting as first payer for the smallest firms. Results suggest that the effect of age on premiums is reduced by 25% when Medicare acts as first payer.

1 Introduction

The aging of the U.S. workforce makes the employment prospects of older workers a top public-policy concern. One suspected barrier to hiring and retaining older workers is the cost

*This paper is part of a larger research project with Joelle Abramowitz and Alice Zawacki; I am grateful to them for their collaboration and guidance. This research was supported by the Dissertation Fellowship in Retirement Research awarded by the Center for Retirement Research at Boston College and sponsored by the Social Security Administration. This research does not represent the views of the Social Security Administration. The Census Bureau has ensured appropriate access and use of confidential data and has reviewed these results for disclosure avoidance protection (Project 7527352: CBDRB-FY24-CED006-0005). The views expressed in this chapter are those of the author and do not necessarily reflect the official policy or position of the Census Bureau.

of health insurance. Although the higher health costs incurred by older workers may be passed through to their own wages, these costs are also borne to some degree by their employers. The salience of older workers' health costs has grown in recent decades as the working population has skewed older and as expenditures on employer-provided health benefits have exploded, with costs now equivalent to about to 12% of total wages (Bureau of Labor Statistics, 2021). This all combines to potentially disadvantage older job-seekers; while older workers are perceived to offer positive qualities such as know-how and people skills, employers also consider them more costly to employ and age discrimination is well documented.

The expense and burdens associated with obtaining insurance for an older and harder-to-insure workforce, particularly among small employers, have drawn attention from policy-makers for decades. U.S. states implemented a wave of small-group insurance reforms in the 1990s and thereafter, culminating in the 2010 Patient Protection and Affordable Care Act or ACA. These policies included ratings reforms, which limit how much insurers can adjust premiums based on health conditions and workforce demographic characteristics such as age and gender. These polices promised to bring stable and affordable health insurance options into reach for smaller employers with older or less healthy employee pools while mitigating one of the barriers older job-seekers may face. Yet ratings reforms and attendant policies also entailed the potentially undesirable trade-offs: cross-subsidization of high-risk participants by younger and healthier workers leading to changes in insurance offer decisions and lower insurance take-up in low-risk segments.

This study examines the association between the age of a firm's workforce and the cost of its health insurance premiums, as well as the impact of state and federal policies on these associations. To accomplish this, I link administrative employment and demographic data with rich survey data from the Medical Expenditure Panel Survey–Insurance Component (MEPS-IC), building on the work of Abramowitz et al. (2018). The resulting data set captures the age structure of MEPS-IC firms far more accurately than the comparable measures collected in the survey—namely, the share of workers 50 and older—allowing for more precise estimates of relationships between employees' age and the characteristics of employer-sponsored health insurance plans.

The first set of results reported in this paper regress log single health premiums on older age share and a wide range of covariates. As expected, the association of age and premiums is significant, with a 10 percentage-point increase in the 50-and-older share of a workforce corresponding to a 4.3% rise in single premiums. The relationship exhibits substantial heterogeneities across firm size, funding arrangement, and presence of a managed care plan (e.g., a health maintenance organization or HMO). Yet these estimates cannot be treated as causal because premiums and workforce age are endogenous to a number of other unobserved factors relating to employers' decisions to offer health insurance and their choice of plans. Beyond the decision to offer health insurance at all, these endogenous factors include the offer of multiple plans, variation in plan generosity, and changes to employee eligibility and work hours. Employee age also correlates with factors like income and seniority that are associated with more generous and thus pricier health plans (Tilipman, 2022).

To help assess the degree to which the main results reflect workforce age and not other factors endogenous to age, the second part of the empirical section exploits variation over time and between states in ratings policies that limit the degree to which insurers can vary premiums based on the age profile of small-group insurance buyers. The more the age-premium relationship observed in the baseline results is driven by confounders, the less of an effect we would expect to see from small-group insurance policies specifically targeted around workforce age. Consistent with the aims of these policies, I find that the age-premium relationship is 71% smaller in states with the most stringent age rating policies relative to states without these policies. Increasing the share of workers 50 and older in a firm by 10 percentage points is associated with 1.8% higher premiums in the presence of the most stringent restrictions, compared to 6.4% where no restrictions exist.

Due to endogeneity concerns arising from the potential correlation of age rating restrictions with unobserved state-level factors, an additional test exploits the nationwide implementation of the ACA's adjusted community rating policies limiting premium variation by age to a 3:1 ratio. Although pre-existing state ratings regulations may suffer some degree of policy endogeneity, the timing of the ACA's implementation represents a plausibly exogenous shock to small-group markets, with exposure to pre-existing ratings reforms providing a source of identifying variation. Results of difference-in-differences regressions using the implementation of ACA show a reduction in the age-premium relationship only for those states that lacked age ratings restrictions prior to the ACA's implementation. I fail to find clear and consistent evidence of reduced health insurance offer or an increase in self-funding among firms in the those states most affected by the ACA's age rating restrictions, though adjustments along other margins still need to be investigated.[1]

In a final empirical application, I explore the role played by Medicare in providing primary insurance to employees at the smallest firms, those with fewer than 20 employees. The results indicate that the increase in premiums associated with a higher share of workers 65 and older is 25% lower for firms where Medicare likely acts as first payer relative to slightly larger firms in the small-group market. This evidence is consistent with Medicare's first-payer status reducing the cost burden for smaller firms that employ workers 65 and older.

This study contributes to the literature exploring how the health insurance decisions of employers respond to employee demographics and particularly workforce age (Bundorf, 2002; Lahey, 2007; Bundorf, Levin and Mahoney, 2012; Tilipman, 2022). It also informs questions over

[1]When firms self-insure plans rather than purchase health insurance on the small-group market, those plans need not abide by age rating restrictions or state benefit mandates. It is plausible that firms with younger and healthier workforces would choose to leave the small-group market upon the introduction of age rating restrictions, since they now bear some of the costs incurred by older and less healthy groups (Park, 2000; Simon, 2005).

small-group health insurance policy design at both the state and federal levels, where adjusted community ratings policies aim to stabilize health insurance premiums for small firms subject to sharp and sudden changes in costs arising from the volatility of employee health expenditures and the limited size of the risk pools. While it is no surprise that health insurance premiums rise with workforce age, quantifying this relationship should help inform future research on the intersection of workforce aging and hiring and compensation outcomes. The additional benefits costs associated with hiring older workers plays an important though still only partially explored role in the demand for their labor (Allen, 2019).

2 Background

Employers are the primary source of health insurance in the U.S., with more than half of the population insured through their workplace or that of a family member (Kaiser Family Foundation [KFF], 2021). Employers typically pay the bulk of insurance premiums with employees covering a portion out of their own pay. The average total premium for single coverage at establishments offering health insurance rose from $3,287 in 1996 to $7,149 in 2020, in 2020 dollars. The average employee contribution rose over this period from 17% to 21% of single-coverage premiums (Agency for Healthcare Research and Quality, 2022).

The availability of health benefits and their method of funding vary considerably across firms and industries. The offer rate of employer-based coverage increases with firm size, ranging from an average of 60% among employers with fewer than 100 workers to 92% for employers of 500 or more (Bureau of Labor Statistics, 2021). Employers can fund health benefits through different arrangements: fully insured plans, self-insured plans, or a mix of the two. Firms that fully insure health benefits purchase coverage through insurance providers in the small- or large-group markets. Firms that self-insure reimburse participating employees' medical claims directly out of company funds, often using stop-loss insurance coverage to insure against catastrophic costs; these firms also often contract with large insurers for administrative services. Larger firms are more likely to self-insure; 27% of covered workers in the 50-199 size range belonged to self-insured plans in 2020, compared to 87% in firms with employment of 1,000 or more (KFF, 2021).

The relationship between employee age and employers' hiring, retention and insurance offer decisions has important implications for the older workers' labor market prospects (Lahey, 2007). Data on health care expenditures and employer reimbursements confirm the intuition that older workers cost employers more. Burtless and Koepcke (2018) used household data from the Medical Expenditure Panel Survey–Household Component to calculate average health reimbursements net of employee premium contributions by age. They found that, compared to workers ages 45-49, net reimbursements were 35% higher for those ages 55-59 and 79% higher for those 60-64 after controlling for other worker characteristics. The considerable costs of childbirth and raising dependents, concentrated among younger age groups, do not change the overall pattern that health costs rise sharply beginning in a worker's 50s.

In the absence of institutional and legal restrictions, employers could pass the additional health costs of older workers onto these employees in the form of diminished wages. Though inconclusive, prior empirical research finds evidence that older workers do indeed absorb some of the costs of rising insurance costs (Sheiner, 1999; Bailey, 2014), in line with a much broader literature on the incidence of health insurance benefits (Gruber, 1994; Baicker and Chandra, 2006; Bhattacharya and Bundorf, 2009; Anand, 2017; Arnold and Whaley, 2020). Federal law prohibits employers from discriminating on the basis of age in pay and requires that employee contributions be uniform within any employer-provided plan. Yet in cases where salaries are negotiable and are not standardized across the firm, older workers may accept lower wage offers in order to obtain health insurance that would be costlier in the individual market.

Employers deciding whether to offer insurance and the types of plan or plans to offer must balance expected costs with employee preferences and competitive pressures. Employers are barred from charging plan enrollees different premiums according to age, health, or any other factors, giving rise to complex choices as to plan pricing and generosity (Bundorf, Levin and Mahoney, 2012). At a high level, employers appear responsive to employee preferences regarding insurance coverage options (Pauly and Herring, 2007; Peele et al., 2000). Employers may also reap benefits from encouraging a healthy and insured workforce, such as lower turnover (perhaps due to job lock) and higher productivity (O'Brien, 2003; Decressin et al., 2005). Among the ways that employers tailor their insurance offerings to reconcile employee demand with a need to minimize expected costs is by offering multiple types of plans: for instance, a high-premium plan with generous coverage and one that is cheaper and less generous for low-risk employees. Employers are more likely to offer a choice in plans at firms with higher-paid workers and those with greater variance in employee health risk (Bundorf, 2002; Tilipman, 2022). Employers may also steer workers towards higher-deductible plans or those with narrower networks in order to reduce costs.

Plan premiums are a function of plan generosity and expected health costs among other factors.[2] A positive association between workforce age and premiums emerges from simple statistical associations (e.g., Kaiser Family Foundation 2021: Fig 1.6; Lurie and Miller 2023) as well as in multivariate regression analyses that use age-related variables as controls (Cutler, 1994; Bundorf, 2010). While this association hardly comes as a surprise, it bears further investigation to determine how much of the age-premium relationship is driven by confounders such as income, union coverage, and industry, which may be associated with both workforce age and health plan generosity.

In an effort to increase the affordability of health insurance and reduce the volatility of costs

[2] One important contributor to plan costs is the competitiveness of the health insurance market in the area where a firm is located (Dafny, Duggan and Ramanarayanan, 2012). In the empirical section I am unable to control for this influence on premiums outside of using state fixed effects.

for smaller employers, most U.S. states enacted some form of ratings reforms by the 1990s and early 2000s through legislation that also often included companion policies such as guaranteed issue and renewal. These regulations set limits on the degree to which small-group insurance providers could vary premiums based on the perceived risk of the group. Ratings reforms came in a few varieties, the most common of which employs bands over which employers can vary premiums according to a set of employee or firm characteristics, including age, health, gender, industry, and whether employees smoke (Hall, 2000). This system, known as adjusted or modified community rating, was made federal policy by the ACA, which eliminated risk rating based on health (or pre-existing conditions) and enforces a 3:1 age rating band (i.e., insurers can rate the oldest employees as no more than three times more expensive than the youngest employees). Modified community rating stands in contrast to so-called pure community rating, in which all groups in an area pay the same premiums for a given plan. New York and Vermont are the only states that enforce such a policy; Massachusetts also has more stringent ratings restrictions than those imposed by the ACA.[3]

Prior research examining ratings reforms has found only modest effects on coverage rates and small-firm insurance decisions. Buchmueller and DiNardo (2002) compared outcomes after the pure community rating introduction in New York to concurrent trends in the neighboring states of Connecticut, which imposed modified community rating, and Pennsylvania, which had none. They found a slight shift toward coverage of older individuals and increased HMO penetration in New York, though no "death spiral" in insurance coverage as some had feared. Park (2000) used cross-sectional data to document higher rates of self-insurance among small employers in states with stronger small-group ratings reforms. Simon (2005) found that such reforms were associated with higher premiums and lower coverage among the lowest-risk segment of the workforce, never-married young men, consistent with cross-subsidization of higher-risk insurance consumers by lower-risk group members. Trish and Herring (2018) documented an increase in self-insurance among small firms in states with more stringent ratings regulations. Studies that have examined the small-group market after the ACA have found null to moderate changes in offer rates and take-up (Abraham, Royalty and Drake, 2016; Kattih, Mansour and Mixon, 2019).

An open question is how small-group ratings policies affect older workers. By spreading the health risks associated with older workers over a larger group, these policies have the potential to encourage the hiring and retention of older workers. Exploring this hypothesis using Current Population Survey data, Rutledge and Crawford (2016) found no evidence for increased employment of older workers in small firms under stronger ratings regulations. While I do not address the question of older-worker employment stemming from ratings regulations, the estimates I present help to quantify the degree to which such policies reduce the costs of employing older

[3]The ACA brought other changes to the small-group market at the same time as ratings restrictions were implemented. These include the requirement that plans offer a set of minimum essential benefits and caps on out-of-pocket spending and deductibles. Pre-existing plans could be granted "grandfathered" status in order to circumvent these requirements.

workers at small firms.

One final institutional influence on the health insurance and employment of older workers is the existence of Medicare as first payer of most insurance claims for seniors employed by firms with fewer than 20 workers. These smallest of employers represent a carve-out from the Medicare as Secondary Payer rules passed by Congress in 1982, which required employers offering health plans to cover Medicare-eligible workers and put employer-provided plans first in line to pay out benefits. The requirement that Medicare act only as second payer for most employers acts as an implicit tax on seniors who work after 65 (Goda, Shoven and Slavov, 2007) and creates a sharp discontinuity between the smallest firms and all other firms in the costs associated with employing seniors.

3 Data

The data set used in this study matches establishment-level survey data on health insurance offerings to employee-employer linked administrative records capturing workforce demographic characteristics. The survey data comes from the Medical Expenditure Panel Survey–Insurance Component (MEPS-IC), which surveys roughly 40,000 establishments annually, providing detailed information on insurance offerings, premiums, employee contributions, deductibles, copayments, and many other plan and establishment details. The sampling frame for the MEPS-IC is the U.S. Census Bureau's Business Register (BR) database of active businesses, which provides the survey with accurate data on firm and establishment characteristics including firm size, firm age, and industry. It also provides a means of linking MEPS-IC data to other Census Bureau data sets. Although the MEPS-IC includes some data on employee demographics, including the share 50 and older, the self-reported data contains a substantial share of missing values and measurement error (as will be seen below). The survey does not ask about the share of younger workers.

Individual data used to create firm demographic statistics come from the Longitudinal Employer-Household Dynamics (LEHD), an administrative data set overseen by the U.S. Census Bureau. The LEHD consists of matched employer-employee records constructed from state unemployment insurance records, providing quarterly data on workers' place of employment and earnings. Establishments in the LEHD can be linked to other Census data products using the BR, which crosswalks between establishments and parent firms as well as Employer Identification Numbers (EINs) used in tax filings. Individual workers in the LEHD can be linked to individual-level Census data in the Individual Characteristics File recording demographic characteristics including age, gender, and race. States entered the LEHD program in a staggered fashion over time, with 45 participating by 2000 (the start of the sample used in this study) and all states online by 2010; four states have dropped out of the voluntary data-sharing program in recent years (Alaska, Arkansas, Michigan and Mississippi).

Inconsistencies in defining the boundaries of an employer present a challenge to using LEHD data to provide employer demographic information for MEPS-IC establishments. While firms identified for MEPS-IC establishments using BR data can span multiple states, LEHD employers are identified by state-specific EINs or SEINs. For this reason, I define firms as a combination of state and firmid, the Census-defined firm identifier. First I identify the firmid associated with the MEPS-IC establishment using the BR identifiers provided in that data set. I then identify all EINs associated with that firm for each state. These state-EIN combinations can then be associated with LEHD SEIN identifiers. These are grouped together by state–firmid and it is on these state–firmid entities that I calculate firm-level demographic characteristics.[4]

Calculation of firm-by-state demographic characteristics follows typical procedures in analyzing LEHD data. For a worker to be included in the calculation of these statistics in quarter t, I require that they have earnings at the same SEIN in both periods t and $t-1$; this corresponds to the Census Bureau's definition of stable employment. I then sum stable employment counts by demographic group over the calendar year and across SEINs associated with a state-firmid and calculate the share of workers in each of three age groups (age 25 or younger, ages 26-49, age 50 and older) and the share who are female. In some analyses, I also use variables denoting finer-grained older-workers shares (ages 50-64 and 65 and up). The resulting data set provides relatively accurate firm-by-state–level demographic data for each MEPS-IC establishment that successfully links to the LEHD.

The linking process yields an overall match rate of 88.8%, rising from 77.9% in 2000 (when not all states had joined the LEHD) to more than 90% for most of the period 2004–2016, then falling slightly to 87.0% as the four states noted above dropped out of the LEHD program; see Appendix Table 3.A1. Weighted by employment, the total match rate is 93.1%. Appendix Table 3.A2 reports the results of a linear probability model regressing a successful match on observable establishment- and firm-level covariates in the MEPS-IC. Of particular note are the statistically significant negative coefficient on the older age-share variable—indicating that MEPS-IC establishments with older reported workforces were less likely to match to LEHD records—as well as the significant coefficients associated with part-time, union and female shares. To help ensure that selection into the matched sample does not bias the results, I reweight observations by the inverse match probability as calculated in this model.

The baseline sample consists of all health insurance plans 2000-2020[5] at matching establishments that offer insurance and whose parent firms record state LEHD stable employment of at least 10 employees within a calendar year. While this size cutoff is not atypical in the literature, it has an additional benefit in this context due to the extreme variance in age share measures for the smallest firms. It is not unusual for tiny firms with just one employee or a handful of workers

[4]At first glance it may appear suboptimal that MEPS-IC measures are defined at the establishment level while the LEHD-derived demographic measures are defined at the state-firm level. Yet employers of multi-establishment firms often offer the same insurance plans to employees at all establishments. Thus the demographics of the firm as a whole (within that state) are likely more relevant to insurance offer decisions that those of the establishment.

[5]Due to a change in survey format between 2007 and 2008, the MEPS-IC records no data for the year 2007.

to record age 50+ shares of 100% or 0%, while such values are rare to nonexistent higher up the firm size distribution (this follows simply from the law of large numbers). Note that multiple plans can be recorded belonging to the same establishment and multiple establishments in any sample year could belong to the same state-firm. In addition to this baseline sample, I define a narrower plan-level sample consisting of plans at smaller firms, 50 employers or fewer, that purchase insurance rather than self-insuring. These small-firm, fully-insured plans correspond to the small-group market that has been the subject of extensive state and federal policy. Table 1 reports summary statistics for the broader sample as well as the narrower small-group sample.

Table 1: Summary statistics

	Broad sample		Narrow sample	
	Mean	Std. Dev.	Mean	Std. Dev.
Establishment size	44.2	212.3	19.7	11.1
Firm size (LEHD)	1591	5222	25.0	106.9
Firm age $<$ 5 years	0.05	0.22	0.11	0.32
Firm age 20+ years	0.73	0.44	0.49	0.50
Union presence	0.09	0.28	0.03	0.17
Nonprofit	0.10	0.30	0.08	0.28
Share female	0.50	0.26	0.44	0.28
Share age $\leq$ 25	0.18	0.17	0.14	0.15
Share age 50+	0.27	0.16	0.31	0.18
Share part-time	0.19	0.26	0.17	0.24
Count establishments	214000		46500	
Plan enrollment	13.9	85.3	8.7	8.2
Single premium (2020)	5720	2276	5820	2580
HMO plan	0.24	0.43	0.32	0.47
PPO plan	0.09	0.29	0.16	0.37
Worker cost share	0.25	0.20	0.24	0.24
Self-funded	0.50	0.50		
Count plans	359000		56000	

Note: MEPS-IC plans linked to LEHD earnings records, 2000–2020. Broad sample includes all firms with LEHD employment of at least 10 and without imputed premiums. The narrow sample further restricts the sample to fully insured plans at firms with 50 or fewer employees. The top panel records establishment-level statistics and the bottom records plan-level statistics. Counts are rounded to adhere to Census disclosure guidelines.

The primary benefit of linking LEHD-derived firm demographics to the MEPS-IC is to improve on the latter's demographic measures. To this end, Figure 1 depicts kernel density estimates of the share of workers 50 and older at MEPS-IC firms, comparing the survey measures with those derived from the administrative LEHD data.[6] Although the broad pattern is similar, the smoother profile of the LEHD-derived estimates reflects an absence of round-number bias

[6]The sample for the survey-derived measure is all MEPS-IC establishments 2000-2020 with non-missing responses to the age 50+ question. The sample for the LEHD-derived measure is all MEPS-IC establishments matching to the LEHD over that period.

that afflicts the survey-derived measures. By reducing measurement error, the LEHD-derived demographic measures have the potential to provide substantially more precise estimates of the relationships of interest.

Figure 1: Distribution of linked establishments by share age 50+, survey data vs administrative records

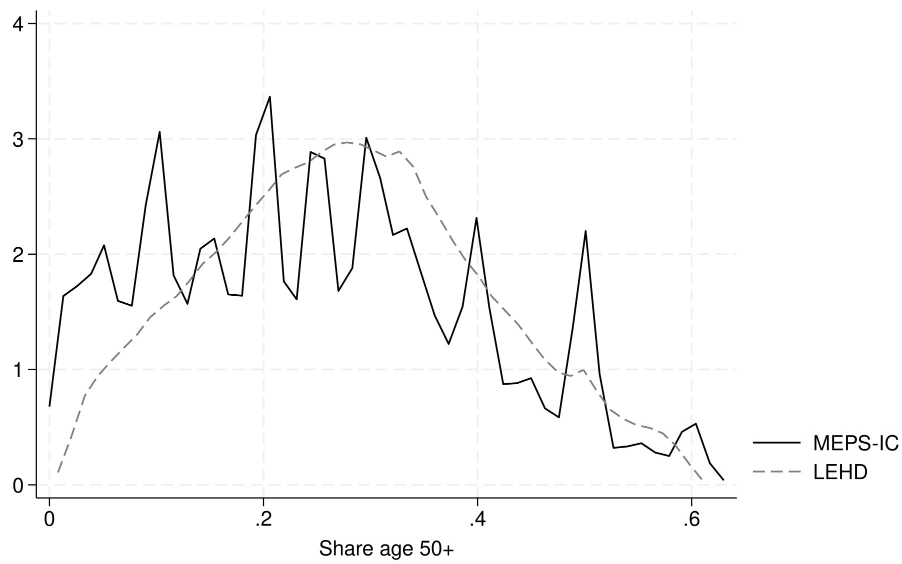

Note: MEPS-IC establishments linked to LEHD earnings records, 2000–2020. Kernel density estimates weighted by MEPS-IC establishment weights. The top and bottom 5% excluded to adhere to Census disclosure guidelines.

I draw a number of other variables from the MEPS-IC and linked administrative data. These include firm multi-unit status, firm age (binned), establishment industry sector, nonprofit status, and binned firm and establishment size. Additional survey-derived establishment-level variables include the share of workers who work part-time, the share who are covered by a union contract, the share who are high- and low-paid (cutoffs for these definitions vary over time according to the MEPS-IC survey instrument). I also use plan-level variables including indicators for whether the plan provider is a health management organization (HMO), a point-of-service plan (POS), a preferred provider organization (PPO) or else not a managed-care provider; and an indicator for the presence of a deductible and its inflation-adjusted level.

The dependent variable in the main regressions is log real premiums for single coverage plans as reported by the employer in MEPS-IC or imputed using logical edits by AHRQ statisticians. I exclude from the samples plans whose premiums are imputed using MEPS-IC imputation models. Because my analysis makes use of state policy differences it is important that the premium measures reflect variation specific to each state. The methodology used to impute missing premiums in MEPS-IC, however, employs a hot-deck process that fills in missing values using other observations that are similar among a number of dimensions including geography, provider type, firm size, and industry (Sommers, 2000). In this process, premiums may be imputed using donor observations residing in other states (though likely in the same Census division).[7] Given the importance of state variation to the analysis below, I exclude observations with imputed premiums.

Data reflecting state-level ratings restrictions was generously provided by Rutledge and Crawford (2016). I examine both age and health ratings restrictions in the models reported below. In both cases I group states into three levels of stringency of their ratings restrictions. The categories are defined according to the ratio by which states allow premiums to vary according to age or health.[8] Appendix Table 3.A3 below reports the number of plan observations by year for each category. The implementation of the ACA moved many states from the least stringent category of age ratings restrictions—where the majority of plans were located before 2014—to the moderately stringent bin. The ACA moved all states to the most stringent category of health ratings restrictions, a category that less than half of the plans in the sample fell into prior to 2014.

4 Methodology

Identifying the impact of workforce age on the premiums paid on health insurance plans presents a number of conceptual challenges. Numerous decisions regarding health care provision are likely made simultaneously, including the decision to offer, the type of provider chosen, whether multiple plans are offered and the generosity of the plan. Each of these is endogenous to other decisions and the premiums charged. For this reason, empirical work on employer-provided health insurance sometimes employs structural models incorporating several aspects of decision-making into a multi-part model (e.g., Bundorf, Levin and Mahoney, 2012; Tilipman, 2022).

Endogeneity also arises between the share of older workers at a firm and the characteristics of its health insurance plan(s). An older workforce, and particularly older workers with more seniority, might push an employer to offer more generous and higher-deductible plans. Conversely, an employer already offering generous health benefits might attract and retain a greater share of older workers. Absent any sources of exogenous variation in the age of the workforce, it is difficult to overcome this source of endogeneity. Thus, since this study does not exploit any shocks to workforce age or other identifying sources of exogenous variation, the conditional associations it reports between the age of the workforce and employer-sponsored health insurance

[7]A similar issue is that imputed premiums may also be drawn from firms of a difference broad size class; i.e., premiums for a 40-employee firm may be imputed using the value for a 100-employee firm. Since the recipient firm would be subject to small-group regulations while the donor firm is not, including such imputed premiums could give rise to additional measurement error.

[8]For health ratings, the categories are least stringent (none to 2.3), moderately stringent (2.08 to 1.67), and the most stringent (1.5 to 1). For age ratings the categories are least stringent (none), moderately stringent (5 to 3), and the most stringent (2.8 to 1).

premiums cannot be interpreted causally.

The baseline results report the coefficients associated with the share of workers age 50+ in an ordinary least squares model regressing log single-coverage premiums on age share and other covariates:

$$\log Prm_{ijst} = \alpha + \beta Age50_j + \mathbf{X}_i\Gamma_1 + \mathbf{Y}_j\Gamma_2 + \phi_s + \phi_t + \varepsilon_{ijst} \tag{1}$$

In Equation 1, log premiums for plan i at firm j in state s and in year t are regressed on a constant, the share age 50 and older at that firm, a vector of plan-level covariates $\mathbf{X_i}$, a vector of firm- and establishment-level covariates $\mathbf{Y_j}$, and fixed effects ϕ_s and ϕ_t absorbing unobserved variation at the state and year level. Standard errors are clustered at the establishment and state levels. Premiums are inflated to 2020 prices using CPI-U.

I use a wide range of controls to account for potential confounders of the age-premium relationship. Because premiums could vary by industry in some state prior to the ACA—and more broadly because certain health conditions may be associated with particular industries—I include indicators for two-digit NAICS sector at the establishment level. I also control for non-profit status at the establishment level. Firm-level controls include firm age, an indicator for multi-unit status, and an interaction of binned firm size with binned establishment size. I also include a number of controls capturing workforce characteristics. The MEPS-IC survey instrument records variables reflecting the share of workers who are part-time, the share covered by union contracts, and the shares with high and low hourly pay (according to cutoffs that vary over time).[9] From the LEHD linkage I also include controls for the share of workers who are women and the share who are 25 years old or younger. Finally, I control for several dimensions of plan characteristics as a way of accounting for plan quality and generosity. I include an indicator for whether the plan has a deductible and separately control for the level of the deductible (in 1000s of 2020 dollars). I also include indicators for a plan's managed-care arrangement: HMO, PPO, POS, or none of these.

I use multiple functional forms to gauge the effect of ratings policies on the age-premium relationship. The first simply interacts the age 50+ variable with indicators for whether the policy in effect in a given state-year falls within each of the three levels of ratings restrictions, r. I run this specification separately for age and health ratings in a sample that includes only fully insured plans belonging to firms with 50 or fewer employees measured using BR data (these plans are the target group of ratings reforms):

$$\log Prm_{ijst} = \alpha + \sum_{r=1}^{3} \beta_r \cdot \mathbb{1}(ratings_{s,t} = r) \cdot Age50_j + \mathbf{X}_i\Gamma_1 + \mathbf{Y}_j\Gamma_2 + \phi_s + \phi_t \tag{2}$$

This specification is intended to capture heterogeneity in the effect of the older-worker share on premiums across different ratings policy regimes, β_r. Yet it could suffer from endogeneity

[9] I use dummy variables to indicate nonresponse to the survey questions underlying these variables so that missing values, here set to zero, do not need to be dropped.

if these policies correlate with separate state regulations that affect the small-group market or other state-level trends concurrent with policy shifts.

To mitigate these concerns (if not entirely eliminate them), a second specification uses the 2014 implementation of the ACA's small-group policies as an exogenous shock to ratings policies whose impact was felt differentially across states according to pre-existing ratings restrictions. As Table 3.A3 indicates, the ACA led to a sharp and permanent shift in ratings policies for a large part of the sample, making age and health ratings more restrictive in states that previously lacked ratings reforms or had less stringent policies. To conduct this analysis I repeat the model in Equation 2 but with a triple-interaction term that interacts the age 50+ share with an indicator for ACA implementation ($year \geq 2014$) and a categorical indicator for pre-ACA state ratings policies.[10] This results in six estimates of $\hat{\beta}_{a,r,}$ corresponding to each of the two periods pre- and post-2014 $a \in (pre, post)$ and each of the three pre-ACA ratings policies $r \in (1, 2, 3)$. The sample for this analysis is limited to 2008–2020 so that the lengths of the pre- and post-ACA implementation periods are roughly equivalent. The expectation is that the age 50+ coefficient $\hat{\beta}_{a,r}$ changes from pre- to post-ACA only for those states whose ratings restrictions were made exogenously more restrictive by implementation of the ACA. This implies $\beta_{pre,r} = \beta_{post,r}$ for those states unaffected by ACA implementation, which are those in age ratings categories 2 or 3 (somewhat stringent or most stringent).

For these estimates to be considered causal requires that both the implementation of the ACA in 2014 and underlying policies are conditionally independent of the changes in the age-premium relationship over the sample period by type of policy. This implies the assumption that the conditional pre-ACA age-premium relationship provides a good counterfactual for how that relationship would have evolved post-ACA within each category of ratings policies. This may not be the case, however, if implementation of the ACA interacted with other state policy measure to affect small-group insurance markets. For instance, for a state in the least-stringent age ratings category in 2013, the ACA's implementation mandated a more stringent age rating band, 3:1. Yet it also imposed minimum essential benefits rules on insurance offer. If states without age ratings restrictions in 2013 also had lower essential benefits requirements pre-ACA, then part of the effect attributed to the new age ratings rules could instead stem from the change in the essential benefits policy. Another possibility is that states imposed age ratings policies pre-ACA due to some unobserved underlying demographic trend which means that the pre-ACA relationship for these states does not provide a good counterfactual fro the post-ACA relationship between age and premiums. Finally, states had four years to anticipate implementation of the ACA's ratings restrictions, a time during which they may have implemented other anticipatory policies affecting the age-premium relationship.[11]

[10] Using an indicator for state policies in 2013 requires that ratings restrictions are relatively stable in the sample period. Over the years 2008–2013, only one state moves between ratings categories (Maryland, in 2009).

[11] As noted above, few states altered their ratings policies between 2008–2013 and none moved across broader ratings categories between the 2010 passage of the ACA and its 2014 implementation.

Research conducted since ACA implementation provides some assurance that the markets affected by the law avoided dramatic shifts in outcomes such as the rate of offer, though some shifts have been observed. Abraham, Royalty and Drake (2016) examined changes in insurance offer between 2013 and 2014 and found no substantial patterns, even for states for whom the ACA introduced adjusted community rating. Kattih, Mansour and Mixon (2019) found a moderate reduction in health insurance take-up following 2014 ACA implementation among small firms relative to larger firms. Yet, counterintuitively, they found increased takeup among small firms with a higher share of adults younger than 26. This is the opposite of what would be expected given the expansion of adjusted community rating, which shifts costs from older and less healthy workers to younger ones.

5 Results

In Table 2 I report results building up to the baseline specification. These models progressively add three groups of controls: covariates at the firm and establishment level, workforce demographic indicators, and controls for plan characteristics. Across specifications the $Age50$ coefficient is positive and significant as expected. Although the size of the coefficient is reduced when additional firm and worker controls are added, the full specification indicates that a 10 percentage-point increase in the 50-and-older share is associated with 4.3% higher single-coverage premiums, equivalent to an increase in per-employee annual premiums of $246 a year at the sample mean. Although these baseline results indicate a statistically significant effect of workforce age on premiums conditional on a range of other observables, variation in LEHD age structure accounts for only 6% of the variation in log premiums in the sample.

How much do the age measures derived administrative records improve over the MEPS-IC collected demographic measures? In column 5, I replace the LEHD-based age share variable with the equivalent measure in MEPS-IC. Though it is still positive and significant, the magnitude of the $Age50$ coefficient estimate is reduced by more than half and is less precisely estimated, consistent with substantial measurement error in the collected MEPS-IC data on workforce age. Caution in warranted in making a direct comparison in these estimates, however, since the MEPS-IC measure of share 50+ includes missing values that may not be missing completely at random, giving rise to the potential for biased estimates.[12]

Next I explore heterogeneities in the age-premium relationship by firm size, funding arrangement, and type of plan. It might be expected that larger firms are better able to moderate the costs stemming from an older workforce because they can cover the costs associated with more complex plan administration or because they have greater bargaining power with insurers. Column 1 of Table 3 interacts the $Age50$ variable with binned firm size indicators (10–24,

[12] I deal with these missing values by assigning missing values of the MEPS-IC collected age share measure $Age50 = 0$ and including a dummy variable for missingness. This strategy has the benefit of keeping the sample consistent across models, though it may introduce bias.

Table 2: Effect of workforce age structure and other factors on log single premiums

	(1)	(2)	(3)	(4)	(5)
Age 50+	0.637***	0.503***	0.515***	0.428***	0.193***
	(0.0187)	(0.0186)	(0.0196)	(0.0224)	(0.0166)
Share in union		0.190***	0.184***	0.170***	0.179***
		(0.0116)	(0.0130)	(0.0131)	(0.0122)
Age $\leq$ 25				-0.139***	-0.312***
				(0.0136)	(0.0112)
Share female				0.146***	0.101***
				(0.00984)	(0.00951)
Share part-time				0.00218	0.0207**
				(0.00689)	(0.00647)
Firm and estab characteristics		X	X	X	X
Plan characteristics			X	X	X
Other workforce demographics				X	X
Original MEPS-IC demographics					X
Observations	359000	359000	359000	359000	359000
R-squared	0.210	0.234	0.250	0.258	0.250

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: MEPS-IC plans linked to LEHD earnings records, 2000–2020. Firm and establishment controls include indicators for multi-unit status, firm age (binned), establishment industry sector, nonprofit status, share of the workforce in union, and the interaction of binned firm and establishment sizes. Plan controls are indicators for HMO, POS, and PPO, an indicator for the presence of a deductible and its inflation-adjusted level. Workforce controls include share of workers who are female, share 25 or younger, share part-time, and shares low- and high-paid. In column (5), MEPS-IC collected measures of share 50+ and share female are substituted for the corresponding LEHD measures. Outcome is log single-coverage health insurance premiums in 2020 dollars. Standard errors clustered by establishment and state.

25–99, 100–999 and 1,000+). Only for the largest firm size does the age-premium relationship differ significantly from that of other firm sizes, with a statistically significant 17% reduction in the $Age50$ coefficient at firms with 1,000+ employees relative to plans at firms with 100–999 employees. Column 2 explores heterogeneity in the age-premium relationship by plan type, comparing PPO, HMO, and POS plans to those without any specific provider arrangements. The age-premium relationship is significantly lower among managed-care plans, though differences between POS, PPO and HMO in the age-premium relationship are minimal. This suggests that managed-care plan arrangements succeed in restraining health costs for older workers. Finally, column 3 of Table 3 explores heterogeneity by plan funding arrangement, self-insured or fully insured. For self-insured plans, the age-premium gradient is cut nearly in half. This may reflect a generally greater administrative capacity for firms that are able to self-insure, though further investigation is warranted.

Table 3: Effect of workforce age structure on log single premiums: heterogeneity by firm and plan type

	(1)	(2)	(3)
Firm 10–24 × age 50+	0.461***		
	(0.0468)		
Firm 25–99 × age 50+	0.432***		
	(0.0362)		
Firm 100-999 × age 50+	0.453***		
	(0.0224)		
Firm 1000+ × age 50+	0.377***		
	(0.0238)		
Any providers × age 50+		0.736***	
		(0.0507)	
POS × age 50+		0.399***	
		(0.0195)	
PPO × age 50+		0.393***	
		(0.0448)	
HMO × age 50+		0.407***	
		(0.0359)	
Fully insured × age 50+			0.522***
			(0.0331)
Self-insured × age 50+			0.266***
			(0.0190)
Establishment characteristics	X	X	X
Firm characteristics	X	X	X
Other workforce demographics	X	X	X
Observations	359000	359000	359000
R-squared	0.258	0.259	0.261

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: MEPS-IC plans linked to LEHD earnings records, 2000–2020. Full set of covariates included in each model. Outcome is log single-coverage health insurance premiums in 2020 dollars. Standard errors clustered by establishment and state.

5.1 State and federal policies

Having established a robust relationship between age and premiums conditional on covariates, this section moves on to explore how state and federal adjusted community ratings policies moderate this relationship. In examining these policies, I further limit the sample to firms likely to be in the small-group market, having 50 or fewer employees.[13] To examine the role of state-level small-group policies in mediating the age-premium relationship, I first interact $Age50$ with

[13]This sample restriction does not perfectly align with every small-group market since definitions vary over time and across states as to what constitutes a small group. Moreover, the small-group definition typically depends on the number of eligible or full-time-equivalent employees, which often differs from the total number of employees with taxable earnings at the firm (which may include part-time workers). The cutoff of 50 used here is a conservative one in that it likely excludes more employers eligible for the small-group market than it includes large-group employers.

indicators for the three levels of age rating regulatory stringency, as indicated in Equation 2. This model also includes these indicators as separate controls to capture the spillover effects of ratings restrictions. Appendix Table 3.A4 reports the results and Figure 2 depicts the estimates associated with the coefficients of interest for both age and health ratings (these are examined in separate regressions).

Figure 2: Coefficient estimates on interactions of age 50+ share with state age ratings policy category

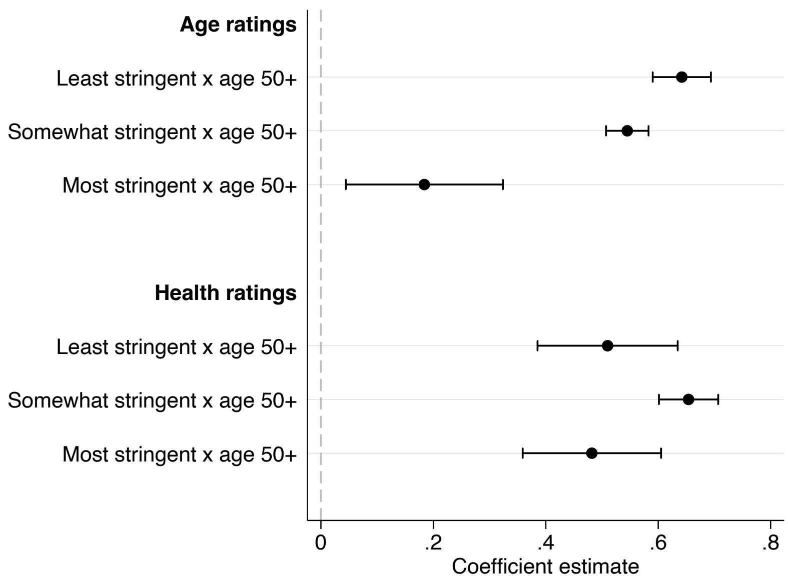

Note: MEPS-IC plans linked to LEHD earnings records, 2000–2020; fully insured plans at small firms. Plots depict coefficient estimates and 95% confidence intervals associated with ratings restriction categories interacted with age 50+ share. Age and health ratings regressions run separately. Full set of covariates included in each model. Outcome is log single-coverage health insurance premiums in 2020 dollars. Standard errors clustered by establishment and state.

In line with expectations, the presence of limits on age rating moderates the relationship between the older-worker share and premiums. Moving from a policy regime without age ratings restrictions to one with the most stringent restrictions reduces the age 50+ coefficient from 0.64 to 0.18, or by 71%. The coefficients on the stringency dummies reported in Appendix Table 3.A4 indicate that the presence of ratings reforms is associated with higher overall premiums only for the most stringent age rating restriction policies; single premiums are 14.4% higher in states with the most restrictive age ratings rules. This is consistent with cross-subsidization: the higher costs of older groups, who can no longer be fully risk-rated, are passed onto younger groups. Yet this conclusion should be considered as merely suggestive, since there might be other, unmeasured policy or state-level confounders driving this difference. The results for the interactions of health ratings restrictions and age share are less clear. There is virtually no difference in the age-premium relationship between the most and least stringent health ratings policies. Nor is the presence of the most stringent health ratings restrictions associated with higher premiums overall (see Appendix Table 3.A4, column 2). This suggests that, despite the correlation of age and health costs, workforce age only accounts for part of the variation in expected health costs (even after controlling for industry sector, share female, and other health-cost-related covariates).

To mitigate the aforementioned concerns over policy endogeneity—if not eliminate them entirely—I leverage the implementation of the ACA's small-group provisions in 2014 as a source of variation whose timing is largely exogenous to state policy changes. These regressions include the triple interaction $Age50 \times ACA_t \times ratings_{s,2013}$ where ACA_t indicates whether $year \geq 2014$ and $ratings_{s,2013}$ records the ratings policy category for state s in 2013, the year before ACA implementation. The expectation is that the impact of the ACA on the $Age50$ coefficient will manifest only for those states with weak or nonexistent ratings restrictions prior to the ACA's implementation. Appendix Table 3.A5 reports the results of regressions with this triple interaction and Figure 3 plots the coefficient estimates for the triple interactions of age by ACA by age ratings policy. The top panel in the figure reports the simpler interaction $Age50 \times ACA_t$ as a reference point.

The results indicate that only those states most affected by the implementation of the ACA—those without any small-group age ratings policies prior to 2014—experience a change in the $Age50$ coefficient between the pre- and post-ACA periods. This reduction is statically significant, with the post-ACA coefficient on $Age50$ falling by 17.5%. This change is comparable to moving from the lowest age ratings category (no age ratings restrictions) to the middle category in the models reported in Figure 2, which is indeed the change that took place for states that had no age ratings policies in 2013. Looking at interactions with health ratings policy categories in Appendix Figure 3.A1 reveals a different pattern. At the implementation of the ACA, the age-premium relationship is reduced across all three policy groups, with the largest reduction occurring for those with the least stringent health ratings restrictions prior to ACA implementation. This is consistent with age and health status both independently affecting group health insurance premiums.

One potential confounder in these results is that states with different pre-ACA age ratings policies experienced divergent pre-trends in the relationship between age and premiums. This could be driven, for instance, by different rates of population aging within the age 50-and-up group between states with different policy regimes (indeed, the timing of ratings policy enactment could be affected by state-level population aging). To help assure that this does not drive the results reported above, I present results of two models exploring the age-premium relationship over time. The first reports interactions between $Age50$ and year dummies for all fully insured small-firm plans. The second model interacts $Age50$ with year dummies as well as a binary indicator for age ratings policies: 0 for no age rating restrictions in 2013, 1 for any age rating restrictions in 2013.

Figure 3: Coefficient estimates on interactions of age 50+ share with ACA and ACA by 2013 state age ratings policy category

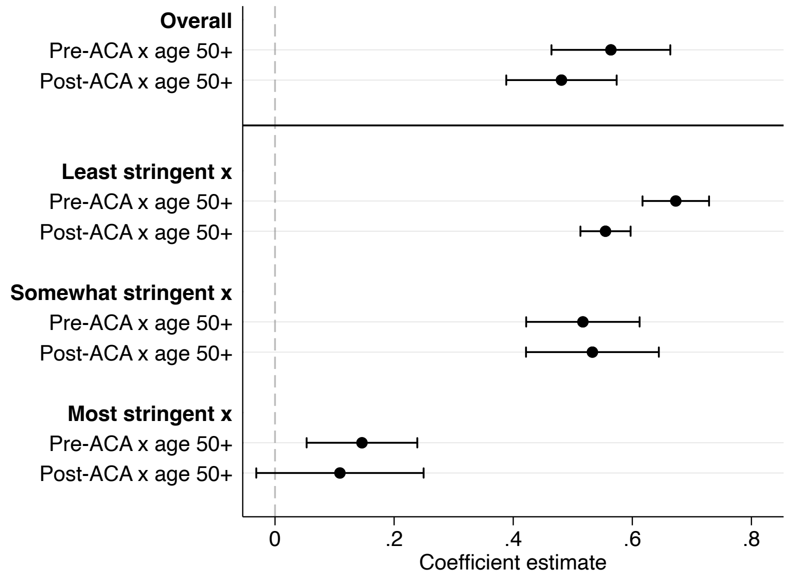

Note: MEPS-IC plans linked to LEHD earnings records, 2008–2020; fully insured plans at small firms. Plots depict coefficient estimates and 95% confidence intervals associated with age ratings restriction categories in 2013 interacted with ACA implementation by age 50+ share. Full set of covariates included in each model. Outcome is log single-coverage health insurance premiums in 2020 dollars. Standard errors clustered by establishment and state.

The estimated coefficients associated with these terms are plotted in Figure 4. For ease of analysis, each panel also includes the mean linear combinations of the coefficients by policy type both for pre-ACA years 2008–2013 and for post-ACA years 2014–2020. Two conclusions follow. First, although there is variation in the year-to-year coefficient estimates, they do not exhibit any clear trend in the years prior to the implementation of the ACA, either for the basic two-way interactions or for the policy-by-year interactions. Second, the coefficients on $Age50$ interactions fall to a statistically significant degree after ACA implementation, but *only* for the group of states with no age ratings restrictions prior to 2014. These results underscore and further illustrate the conclusions reached above: the effects of the ACA on the age-premium relationship were concentrated in those states without strong age rating restrictions prior to 2014. These effects do not appear to be driven by state-level trends that are correlated with policy stances.

A final open question is whether the implementation of adjusted community ratings policies drove spillover effects such as health insurance offer or plan funding arrangements due to cross-subsidization pressures (Park, 2000; Trish and Herring, 2018). One possibility is that small-group employers with a lower-risk workforce might opt not to offer coverage when premiums rise due to cross-subsidization of costs from higher-risk groups. The same process could drive these

Figure 4: Trends in the age-premium relationship over time, small firms

Panel A: Year × age 50+ interactions

Panel B: Year × age ratings group (2013) × age 50+ interactions

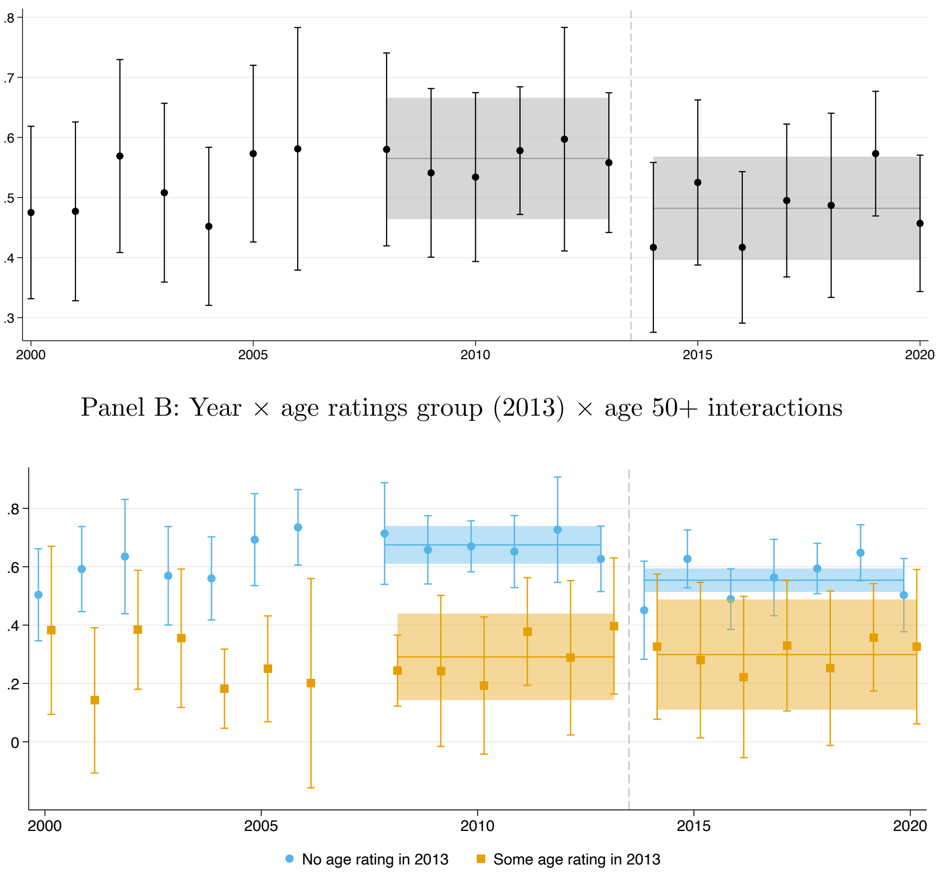

Note: MEPS-IC plans linked to LEHD earnings records, 2000–2020; fully insured plans at small firms. The year 2007 is missing from the MEPS-IC due to a change in survey format between 2007 and 2008. Plots depict coefficient estimates and 95% confidence intervals associated with binary age ratings restriction categories in 2013 (none vs any restrictions) interacted with year by age 50+ share. Full set of covariates included in each model. Outcome is log single-coverage health insurance premiums in 2020 dollars. Standard errors clustered by establishment and state.

employers to shift to self-insuring their plans.

To explore these potential spillover effects, I run a series of models where the coefficients of interest are interactions using different age group shares. The younger measure $Age25$ reflects the share of a firm's workforce that is aged 25 years or younger and the mid-career measure $AgeMid$ reflects the share ages 26–49. I follow the same approach as outlined in Equation 2, employing the interactions $Age \times ratings_{s,t}$ to probe whether offer or self-insurance outcomes vary by state policy regimes. I run these models separately for each of the three age share measures used independently in the interaction. Examining these outcomes requires expanding the samples from those used in the rest of the small-group regressions. For the models with insurance offer as the outcome, the unit of analysis is the establishment and the sample expands the previously used small-group sample to include establishments not offering insurance. For the models with self-insurance as the outcome, the unit of analysis is the plan and the sample expands the previous small-group sample to include self-insured plans.

Table 4: Effect of age structure and policy environment on insurance offer and self-insurance

	(1) Offer	(2) Offer	(3) Offer	(4) Self	(5) Self	(6) Self
Age ratings group 2	0.00109	-0.0118	-0.0200	0.0327	0.0371	0.0303
	(0.0137)	(0.0138)	(0.0196)	(0.0173)	(0.0192)	(0.0225)
Age ratings group 3	0.00125	0.127***	0.0197	-0.00901	0.0172	0.00456
	(0.0269)	(0.0263)	(0.0247)	(0.0111)	(0.0284)	(0.0173)
Age ratings group 1 × age ≤ 25	-0.159***			0.0347		
	(0.0194)			(0.0258)		
Age ratings group 2 × age ≤ 25	-0.229***			0.0385		
	(0.0275)			(0.0345)		
Age ratings group 3 × age ≤ 25	-0.0364			0.0980*		
	(0.0687)			(0.0478)		
Age ratings group 1 × age 26–49		0.0598**			0.0133	
		(0.0174)			(0.0131)	
Age ratings group 2 × age 26–49		0.0704**			0.00885	
		(0.0247)			(0.0212)	
Age ratings group 3 × age 26–49		-0.121***			-0.0139	
		(0.0159)			(0.0389)	
Age ratings group 1 × age 50+			0.0801***			-0.0313
			(0.0136)			(0.0203)
Age ratings group 2 × age 50+			0.113***			-0.0229
			(0.0246)			(0.0202)
Age ratings group 3 × age 50+			0.110			-0.0456**
			(0.0615)			(0.0154)
Observations	94000	94000	94000	62500	62500	62500
R-squared	0.314	0.312	0.312	0.020	0.019	0.019

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: MEPS-IC plans linked to LEHD earnings records, 2000–2020; fully insured plans at small firms. Linear probability model. Full set of covariates included in each model. Sample for model using offer as outcome (columns 1–3) includes all establishments of small firms regardless of health insurance offer, funding type and premium imputation. Sample for model with self-insurance as the outcome (columns 4–6) is restricted to plans with non-imputed premiums at small firms. Standard errors clustered by establishment and state.

The results of the models described above are reported in Table 4. Looking first at insurance offer as the outcome (columns 1–3), the results fail to pick up any clear shift away from offering insurance to the youngest (and therefore lower-risk) groups under the most stringent age ratings restrictions, as the coefficients on the $Age25$ interactions show in column 1. The pattern of coefficients related to the $AgeMid$ group, however, is consistent with some degree of shifting away from offering health insurance to relatively younger groups under the most stringent age ratings restrictions—the coefficient on the $AgeMid$ interactions shifts from positive in less restrictive environments to significantly negative under more restrictive policies. Looking at self-insurance as the outcome (columns 4–6), there is some evidence of a shift towards self-insuring the youngest groups under the most stringent policies (column 4), but not for the age 26–49 group interactions (column 5). Under the most restrictive age ratings policies, an older workforce is associated with a reduced tendency to self-insure (column 6), in line with what might be expected if self insurance is a way of minimizing age-related health costs.

As before, however, policy endogeneity might confound these results. Thus I again use the implementation of the ACA as a semi-exogenous shock to small-group health insurance markets. Appendix Table 3.A6 reports results of interactions $Age \times ACA_t \times ratings_{s,2013}$ for the three different age groups to test whether, in states undergoing the most substantial ACA-driven policy changes, lower-risk groups became any less likely to offer insurance or any more likely to self-insure. These patterns are not evident in any of the estimates, either for insurance offer or self-insurance as the outcome of interest.

Taken together, these results suggest that the small-group adjusted community ratings provisions may induce shifts away from health coverage and towards self-insured plans by the types of firms one would expect to be most impacted by cross-subsidizing higher-risk groups. Yet the evidence is mixed and hardly conclusive. The implementation of the ACA does not appear to have given rise to any spillovers related to its age ratings policies in the models reported here. These results should be understood as provisional and merely suggestive. It may be the case that while ratings provisions had small or nonexistent spillover effects on offer rates by employers, individual low-risk workers shifted out of coverage due to cost increases. Firms may also have opted to defray higher costs by other means, for instance by moving to higher-deductible plans.

5.2 Medicare and small firms

Another policy that has a significant impact on small firms is the exemption from Medicare as Secondary Payer rules for firms with fewer than 20 employees. For these smallest firms, which I label as "tiny," Medicare acts as the first payer for medical claims, reducing the potential costs of employing seniors 65 and older. It has yet to be estimated, however, to what degree this federal policy design affects health insurance premiums by tiny firms employing seniors.

Using the fully insured small-firm sample 2000-2020 defined above, I regress log premiums on age share variables interacted with an indicator for tiny-firm status (employment less than 20). I run this same regression for three age share variables: the share 50 and older, the share 50-64, and the share 65 and older. Since Medicare eligibility begins at age 65, the expectation is that the difference in the age-premium relationship between tiny firms and those with 20–50 employees is greatest when the age 65+ variable is used. As Table 5 indicates, this expectation is borne out in the results. For this sample, the additional premium costs associated with employing seniors—conditional on all the covariates—are a statistically significant 25% lower (0.690 vs 0.520) for tiny firms relative to slightly larger firms in the small-group market.

Table 5: Effect of workforce age structure on log single premiums around Medicare-as-second-payer firm size cutoff

	(1)	(2)	(3)
Tiny firm	-0.0190	-0.0178	-0.0131
	(0.0189)	(0.0204)	(0.0188)
Age 50+	0.528***		
	(0.0448)		
Tiny × age 50+	-0.0148		
	(0.0252)		
Age 50–64		0.535***	
		(0.0553)	
Tiny × age 50–64		-0.0210	
		(0.0304)	
Age 65+			0.690***
			(0.0554)
Tiny × age 65+			-0.170**
			(0.0580)
Observations	56000	56000	56000
R-squared	0.257	0.250	0.237

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: MEPS-IC plans linked to LEHD earnings records, 2000–2020; fully insured plans at small firms. Full set of covariates included in each model. Outcome is log single-coverage health insurance premiums in 2020 dollars. "Tiny" indicates firms with fewer than 20 employees. Standard errors clustered by establishment and state.

6 Conclusion

This paper used a new linkage between administrative payroll records and survey data on employer health insurance offerings to estimate the relationship between workforce age and health insurance premiums at U.S. employers. Focusing on small employers, the paper examined how policies impact this relationship, including state and federal adjusted community ratings policies and the exemption to the Medicare as Secondary Payer program for firms with fewer than 20 workers. The baseline results indicate that premiums are strongly and positive associated with the share of a firm's workforce that is 50 years of age or older, with an additional 10% of a firm's workforce associated with 4.3% higher single-coverage premiums. Looking at heterogeneities, I found that this relationship is somewhat dampened among the largest firms and significantly reduced among managed-care plans and in self-insured plans.

Results exploring the impacts of policies that reduce the extent to which insurers can rate premiums by employee age and health status are consistent with state-level restrictions on age rating having their intended effects among small-group fully insured plans. States with the most stringent age ratings restrictions exhibit an age-premium relationship that is 71% smaller than states with no age ratings restrictions, while their overall premiums are somewhat higher, consistent with cross-subsidization. Using the implementation of the ACA's adjusted community ratings rules in 2014 as an alternate source of variation, the results showed a pattern largely consistent with the prior results: the age-premium relationship was reduced after 2014, but only for those states with the weakest age ratings restrictions prior to the ACA. Turning to the analysis of potential spillovers, I found evidence consistent with only limited, if any, knock-on effects of age ratings restrictions on insurance offer and self-insurance. The results fail to show any spillovers related to the ACA's implementation of age ratings restrictions in the small-group market, though this analysis should be regarded as provisional.

This study makes several contributions. First, although it is no surprise that workforce age has a strong positive association with employer-sponsored health insurance premiums, carefully quantifying this relationship and its heterogeneities may help researchers and policy makers better analyze the barriers older job-seekers face as well as the impacts of policies aimed at older workers and the small-group market. Second, the analysis of the effects of adjusted community ratings policies, while not strictly causal, provides compelling evidence that these policies do indeed reduce the costs associated with hiring older workers. These results have implications for future policies that may aim to make the offer of health insurance more manageable among small employers.

This research also paves the way for future work exploring the interactions of workforce demographics and benefits. A natural question is to what degree the implementation of adjusted community ratings policies, particularly those of the ACA, may have affected small-firm hiring and retention of older workers. The results reported here point toward an opportunity for research that uses the ACA or similar community ratings policy changes to explore how the hiring and retention of older workers react to shifts in the costs associated with hiring older workers. More research is also needed to examine potential spillovers brought about by the cross-subsidization of older and high-risk workers by younger, lower-risk workers. This could involve lower insurance take-up or changes to particular plan types and offerings, outcomes not explicitly explored in this paper.

References

Abraham, Jean, Anne B. Royalty, and Coleman Drake. 2016. "Employer-Sponsored Insurance Offers: Largely Stable In 2014 Following ACA Implementation." *Health Affairs*, 35(11): 2133–2137.

Abramowitz, Joelle, Kristin McCue, Britni Wilcher, and Alice Zawacki. 2018. "How is Employer-sponsored Health Insurance Changing in Response to an Aging Workforce?" Unpublished manuscript.

Agency for Healthcare Research and Quality. 2022. "Medical Expenditure Panel Survey–Insurance Compenent: Table Series I. Private-Sector Data by Firm Size and Selected Characteristics for the U.S." Center for Financing, Access and Cost Trends.

Allen, Steven G. 2019. "Demand for Older Workers: What Do Economists Think? What Are Firms Doing?" National Bureau of Economic Research. NBER Working Paper No. 26597.

Anand, Priyanka. 2017. "Health insurance costs and employee compensation: Evidence from the national compensation survey." *Health Economics*, 26: 1601–1616.

Arnold, Daniel, and Christopher M. Whaley. 2020. "Who Pays for Health Care Costs? The Effects of Health Care Prices on Wages." RAND Corporation. working paper WR-A621-2.

Baicker, Katherine, and Amitabh Chandra. 2006. "The labor market effects of rising health insurance premiums." *Journal of Labor Economics*, 24.

Bailey, James. 2014. "Who pays the high health costs of older workers? Evidence from prostate cancer screening mandates." *Applied Economics*, 46: 3931–3941.

Bhattacharya, Jay, and M. Kate Bundorf. 2009. "The incidence of the healthcare costs of obesity." *Journal of Health Economics*, 28(3): 649–658.

Buchmueller, Thomas, and John DiNardo. 2002. "Did Community Rating Induce an Adverse Selection Death Spiral? Evidence from New York, Pennsylvania, and Connecticut." *The American Economic Review*, 92(1): 280–294.

Bundorf, M. Kate. 2002. "Employee demand for health insurance and employer health plan choices." *Journal of Health Economics*, 21: 65–88.

Bundorf, M. Kate. 2010. "The Effects Of Offering Health Plan Choice Within Employment-Based Purchasing Groups." *Journal of Risk & Insurance*, 77(1): 105–127.

Bundorf, M. Kate, Jonathan Levin, and Neale Mahoney. 2012. "Pricing and Welfare in Health Plan Choice." *American Economic Review*, 102(7): 3214–48.

Bureau of Labor Statistics. 2021. "National Compensation Survey: Employee Benefits in the United States."

Burtless, Gary, and Eric Koepcke. 2018. "Age Related Health Costs and Job Prospects of Older Workers." Paper presented at Brookings Institution seminar, "Reducing Employer Bias against Older Workers: Public Policies to Shrink Age-Related Costs of Employee Compensation." Washington, DC, November 2, 2018.

Cutler, David M. 1994. "Market Failure in Small Group Health Insurance." National Bureau of Economic Research. NBER Working Paper No. 4879.

Dafny, Leemore, Mark Duggan, and Subramaniam Ramanarayanan. 2012. "Dafny, Duggan and Ramanarayanan (2012). Paying a Premium on Your Premium? Consolidation in the US Health Insurance Industry." *American Economic Review*, 102(2): 1161–1185.

Decressin, Anja, Julia Lane, Kristin McCue, and Martha Stinson. 2005. "Employer-Provided Benefit Plans, Workforce Composition and Firm Outcomes." Center for Economic Studies, U.S. Census Bureau. LEHD Technical Paper No. TP-2005-01.

Goda, Gopi Shah, John B Shoven, and Sita Slavov. 2007. "A Tax on Work for the Elderly: Medicare as a Secondary Payer."

Gruber, Jonathan. 1994. "The Efficiency of a Group-Specific Mandated Benefit: Evidence from Health Insurance Benefits for Maternity." National Bureau of Economic Research. NBER Working Paper No. 4157.

Hall, Mark A. 2000. "The Structure and Enforcement of Health Insurance Rating Reforms." *Inquiry*, 37(4): 376–388.

Kaiser Family Foundation. 2021. "Employer Health Benefits: 2021 Annual Survey."

Kattih, Nour, Fady Mansour, and Franklin G. Mixon. 2019. "Small Firms, Bigger Changes: Health Insurance Coverage Take-up Rates in Small Firms After the ACA." *Applied Economics*, 51(54): 5878–5889.

Lahey, Joanna N. 2007. "Does Health Insurance Affect the Employment of Older Workers?" Sloan Center on Aging & Work at Boston College.

Lurie, Ithai Z., and Corbin L. Miller. 2023. "Employer-Sponsored Health Insurance Premiums and Income in U.S. Tax Data." *Journal of Public Economics.*

O'Brien, Ellen. 2003. "Employers' Benefits from Workers' Health Insurance." *The Milbank Quarterly*, 81(1): 5–43.

Park, Christina H. 2000. "Prevalence of Employer Self-Insured Health Benefits: National and State Variation." *Medical Care Research and Review*, 57(3): 340–360.

Pauly, Mark V., and Bradley Herring. 2007. "The Demand for Health Insurance in the Group Setting: Can You Always Get What You Want?" *Journal of Risk & Insurance*, 74(1): 115–140.

Peele, Pamela B., Judith R. Lave, Jeanne T. Black, and John Evans. 2000. "Employer-Sponsored Health Insurance: Are Employers Good Agents for Their Employees?" *Milbank Quarterly*, 78(1): 5–21.

Rutledge, Matthew S., and Caroline Crawford. 2016. "Would Reducing The Price Of Employing An Older Worker Improve Labor Market Outcomes By Socioeconomic Status? Evidence From Health Insurance Premium Restrictions." Center for Retirement Research at Boston College. CRR WP 2016-17.

Sheiner, Louise. 1999. "Health Care Costs, Wages and Aging." Federal Reserve Board of Governors. Finance and Economics Discussion Series 1999-19.

Simon, Kosali Ilayperuma. 2005. "Adverse Selection in Health Insurance Markets? Evidence from State Small-Group Health Insurance Reforms." *Journal of Public Economics*, 89(9-10): 1865–1877.

Sommers, John Paul. 2000. "Imputation of employer information for the 1996 Medical Expenditure Panel Survey Insurance Component." Agency for Healthcare Research and Quality. MEPS Methodology Report No. 10. AHRQ Pub. No. 00-0039.

Tilipman, Nicholas. 2022. "Employer Incentives and Distortions in Health Insurance Design: Implications for Welfare and Costs." *American Economic Review*, 112(3): 998–1037.

Trish, Erin, and Bradley Herring. 2018. "Does Limiting Allowable Rating Variation in the Small Group Health Insurance Market Affect Employer Self-Insurance?" *Journal of Risk and Insurance*, 85(3): 607–633.

A Additional Figures and Tables

Figure 3.A1: Coefficient estimates on interactions of age 50+ share with ACA and ACA by 2013 state health ratings policy category

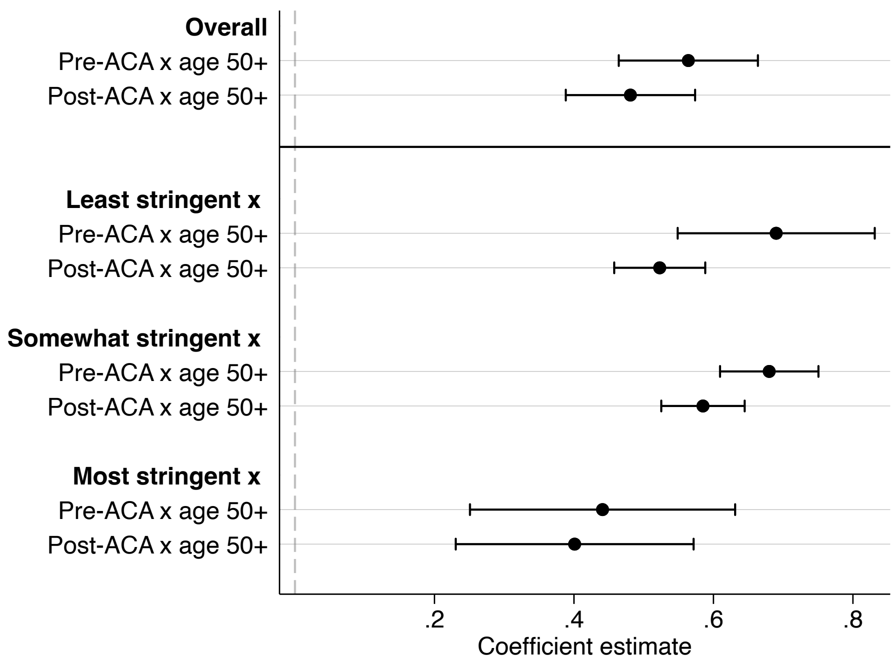

Note: MEPS-IC plans linked to LEHD earnings records, 2008–2020; fully insured plans at small firms. Plots depict coefficient estimates and 95% confidence intervals associated with health ratings restriction categories in 2013 interacted with ACA implementation by age 50+ share. Full set of covariates included in each model. Outcome is log single-coverage health insurance premiums in 2020 dollars. Standard errors clustered by establishment and state.

Table 3.A1: Match percentages by year

Year	Raw	Weighted
2000	77.9	87.6
2001	81.4	89.4
2002	86.6	92.9
2003	89.5	93.9
2004	90.3	94.1
2005	90.8	94.5
2006	91.2	94.9
2008	88.8	92.5
2009	89.9	93.0
2010	90.0	93.4
2011	90.3	93.6
2012	90.5	93.6
2013	90.6	93.7
2014	90.3	94.1
2015	90.7	94.2
2016	91.9	94.1
2017	89.8	93.5
2018	89.7	93.8
2019	87.0	92.4
2020	87.0	92.7
Total	88.8	93.1

Note: Table reports the percent of MEPS-IC establishments that can be matched to LEHD earnings records by year and overall. The weighted column reports percentages weighted by employment.

Table 3.A2: Results of match probability model

	(1)
Offers health insurance	0.00194
	(0.00123)
Multi-unit	0.00864***
	(0.00186)
Share part-time	-0.0396***
	(0.00173)
Share union	0.0135***
	(0.00373)
Share female	0.0231***
	(0.00187)
Age 50+	-0.0662***
	(0.00187)
Observations	578000
R-squared	0.175

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: MEPS-IC establishments 2000–2020. Unweighted linear probability model with robust standard errors. Model also controls for binned firm age, binned firm size interacted with binned establishments size, two-digit NAICS sector, nonprofit status, and state and year fixed effects. Census disclosure requirements limit the number of categorical regressors that can be simultaneously reported.

Table 3.A3: Distribution of plans across ratings restriction categories

	Age ratings group			Health ratings group		
Year	1	2	3	1	2	3
2000	1800	200	300	200	1200	900
2001	1900	300	400	300	1300	950
2002	2000	250	500	400	1300	1100
2003	2300	300	450	450	1500	1100
2004	1900	300	400	400	1200	900
2005	1900	250	450	400	1100	1000
2006	1900	300	450	350	1200	1100
2008	2200	350	450	500	1300	1200
2009	2400	550	450	550	1400	1500
2010	2400	550	500	550	1400	1400
2011	2400	600	450	600	1400	1400
2012	2100	500	400	450	1300	1300
2013	2000	500	400	450	1200	1200
2014	0	2300	300	0	0	2600
2015	0	2100	300	0	0	2400
2016	0	2500	350	0	0	2800
2017	0	2400	350	0	0	2800
2018	0	2600	300	0	0	2900
2019	0	2400	350	0	0	2800
2020	0	1900	250	0	0	2200

Note: MEPS-IC plans linked to LEHD earnings records, 2000–2020. Ratings group range from least stringent (1) to most stringent (3) with cutoffs given in text. Counts are rounded to adhere to Census disclosure guidelines.

Table 3.A4: Effect of workforce age structure on log single premiums: Heterogeneity by state policy

	(1)	(2)
Age ratings group 2	0.000925	
	(0.0328)	
Age ratings group 3	0.144***	
	(0.0317)	
Age ratings group 1 × age 50+	0.642***	
	(0.0264)	
Age ratings group 2 × age 50+	0.545***	
	(0.0193)	
Age ratings group 3 × age 50+	0.184*	
	(0.0713)	
Health ratings group 2		-0.0739
		(0.0435)
Health ratings group 3		-0.0257
		(0.0504)
Health ratings group 1 × age 50+		0.510***
		(0.0636)
Health ratings group 2 × age 50+		0.654***
		(0.0269)
Health ratings group 3 × age 50+		0.482***
		(0.0628)
Observations	56000	56000
	0.261	0.258

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: MEPS-IC plans linked to LEHD earnings records, 2000–2020; fully insured plans at small firms. Full set of covariates included in each model. Outcome is log single-coverage health insurance premiums in 2020 dollars. Ratings restrictions groups range from least stringent (1) to most stringent (3). Standard errors clustered by establishment and state.

Table 3.A5: Effect of workforce age structure on log single premiums: Heterogeneity by state policy in 2013 and ACA implementation

	(1)	(2)	(3)
Pre-ACA × age 50+	0.564***		
	(0.0509)		
Post-ACA × age 50+	0.481***		
	(0.0473)		
Age ratings group 1 × pre-ACA × age 50+		0.673***	
		(0.0285)	
Age ratings group 1 × post-ACA × age 50+		0.555***	
		(0.0215)	
Age ratings group 2 × pre-ACA × age 50+		0.517***	
		(0.0486)	
Age ratings group 2 × post-ACA × age 50+		0.533***	
		(0.0569)	
Age ratings group 3 × pre-ACA × age 50+		0.146**	
		(0.0474)	
Age ratings group 3 × post-ACA × age 50+		0.109	
		(0.0717)	
Health ratings group 1 × pre-ACA × age 50+			0.690***
			(0.0721)
Health ratings group 1 × post-ACA × age 50+			0.523***
			(0.0333)
Health ratings group 2 × pre-ACA × age 50+			0.680***
			(0.0360)
Health ratings group 2 × post-ACA × age 50+			0.585***
			(0.0305)
Health ratings group 3 × pre-ACA × age 50+			0.441***
			(0.0970)
Health ratings group 3 × post-ACA × age 50+			0.401***
			(0.0870)
Observations	37500	37500	37500
R-squared	0.205	0.212	0.207

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: MEPS-IC plans linked to LEHD earnings records, 2008–2020; fully insured plans at small firms. Full set of covariates included in each model. Outcome is log single-coverage health insurance premiums in 2020 dollars. Ratings restrictions groups range from least stringent (1) to most stringent (3). Standard errors clustered by establishment and state.

Table 3.A6: Effect of age structure on insurance offer and self-insurance: ACA and state policy interactions

	(1) Offer	(2) Offer	(3) Offer	(4) Self	(5) Self	(6) Self
Age ratings group 1 × pre-ACA × age ≤ 25	-0.205***			0.0219		
	(0.0218)			(0.0304)		
Age ratings group 1 × post-ACA × age ≤ 25	-0.226***			0.0548		
	(0.0278)			(0.0487)		
Age ratings group 2 × pre-ACA × age ≤ 25	-0.168***			0.0449		
	(0.0390)			(0.0390)		
Age ratings group 2 × post-ACA × age ≤ 25	-0.250***			0.0262		
	(0.0197)			(0.0513)		
Age ratings group 3 × pre-ACA × age ≤ 25	0.00632			0.0633		
	(0.0787)			(0.0364)		
Age ratings group 3 × post-ACA × age ≤ 25	-0.0548			0.126		
	(0.0847)			(0.131)		
Age ratings group 1 × pre-ACA × age 26–49		0.0501*			0.0147	
		(0.0190)			(0.0149)	
Age ratings group 1 × post-ACA × age 26–49		0.0473			0.00972	
		(0.0299)			(0.0248)	
Age ratings group 2 × pre-ACA × age 26–49		0.0790			-0.0143	
		(0.0710)			(0.0201)	
Age ratings group 2 × post-ACA × age 26–49		0.00907			-0.0152	
		(0.0368)			(0.0426)	
Age ratings group 3 × pre-ACA × age 26–49		-0.119***			0.0396	
		(0.0200)			(0.0308)	
Age ratings group 3 × post-ACA × age 26–49		-0.205***			-0.0281	
		(0.0226)			(0.0725)	
Age ratings group 1 × pre-ACA × age 50+			0.120***			-0.0194
			(0.0218)			(0.0190)
Age ratings group 1 × post-ACA × age 50+			0.119***			-0.0301
			(0.0292)			(0.0218)
Age ratings group 2 × pre-ACA × age 50+			0.0616			-0.0109
			(0.0438)			(0.0268)
Age ratings group 2 × post-ACA × age 50+			0.170***			0.00429
			(0.0373)			(0.0510)
Age ratings group 3 × pre-ACA × age 50+			0.0586			-0.0657*
			(0.0624)			(0.0323)
Age ratings group 3 × post-ACA × age 50+			0.171*			-0.0405
			(0.0753)			(0.0206)
Observations	63500	63500	63500	42000	42000	42000
R-squared	0.322	0.318	0.319	0.021	0.020	0.020

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: MEPS-IC plans linked to LEHD earnings records, 2008–2020; fully insured plans at small firms. Linear probability model. Full set of covariates included in each model. Sample for model using offer as outcome (columns 1–3) includes all establishments of small firms regardless of health insurance offer, funding type and premium imputation. Sample for model with self-insurance as the outcome (columns 4–6) is restricted to plans with non-imputed premiums at small firms. Standard errors clustered by establishment and state.