Help Really Wanted?
The Impact of Age Stereotypes in Job Ads on Applications from Older Workers

Broad research agenda: policies that promote employment and economic self-sufficiency

- Redistribution policies that promote work and human capital investment, rather than the opposite
  - Most notably, minimum wage vs. EITC
- Addressing spatial and other impediments to employment
  - Place-based policies
  - Hiring credits
- Reducing barriers to work
  - Studying (and testing for) discrimination
  - Effects of anti-discrimination policies
- Other workforce challenges (labor market networks, opioids, health, workplace injuries)
- Next up: Low hanging fruit (?) of increasing employment among those who find (but lose) jobs
  - Different focus than on the long-term non-employed

My approach

- Primary goal is to generate rigorous and reliable evidence, focused on policy questions
- Not "constrained" by what theory says can/can't happen, but I consider implications of theory for thinking about the evidence, e.g.:
  - Minimum wages:
    - I didn't "reject the possibility that Card/Krueger could be right" (as some urged); I worked to obtain better data
    - Tests of richer implications than just "job loss"
  - Discrimination:
    - Not constrained by Becker's results that discrimination can't persist (nepotism, statistical disc., employee disc., customer disc.)
    - But informed Becker, keenly attuned to non-discriminatory explanations
- Look for and create opportunities to interact with policymakers about research evidence

Help Really Wanted?
The Impact of Age Stereotypes in Job Ads on Applications from Older Workers

Ian Burn
Daniel Firoozi
Daniel Ladd
David Neumark

Motivation – population aging and employment of older workers

- Aging populations lead to rising dependency ratios, slower economic output/growth – hence policy imperative to increase employment of older individuals
- Supply-side incentives (Soc. Sec. reforms) enacted to encourage employment of older workers beyond current retirement ages
- Demand-side barriers – age discrimination – may frustrate these efforts
- Interest not limited to retirement ages
  - Substantial employment declines among less-skilled men at older, but pre-retirement, age
  - Lower employment at these ages can impact longer-term employment/retirement

The importance of hiring

- Hiring of older individuals likely an important part of the solution
  - "Bridge" jobs, "partial retirement," and "unretirement" common (nearly 50%) before transitioning to complete retirement (e.g., Johnson et al., 2009; Maestas, 2010)
- Discrimination laws (ADEA in this case) may be ineffective at reducing hiring discrimination
  - Anti-discrimination enforcement largely driven by private attorneys who share in damages
  - Damages from not being hired may be low at the individual level
  - Much harder to identify a "class," which is a damage multiplier, compared to, e.g., a group of current employees promoted less or terminated at a higher rate

Culmination of a set of papers on age discrimination

- Large-scale field experiment (correspondence study) on age discrimination
  - Strong evidence of age discrimination (especially for women) (Neumark et al., *JPE*, 2019)
- Machine learning/computational linguistics study of job ads from field experiment
  - Discriminatory employers were more likely to use ageist stereotypes in job ads (Burn et al., *JoLE*, 2022)
- This paper (+ an intermediate one): Field experiment manipulating job-ad language to see if ageist stereotypes discourage older job seekers from applying
  - "Independent" motivation: "[d]espite protections by the Age Discrimination in Employment Act of 1967…, employers have gotten cleverer in masking what is age discrimination" by using ageist phrases in job ads (AARP, Brenoff, 2019)

Overview of paper/findings

- Create bank of job ads for three occupations using language from real job ads
- Random variation across job ads in use of age-related stereotypes related to communications skills, physical ability, and technology skills (plus one treatment with more blatant language)
- Post on job board in 14 cities, over approximately 16 months, collect responses from 2,646 job applicants
- Findings:
  - Statistically significant evidence that ageist stereotypes reduce the likelihood that older workers apply; e.g., when all three used, average age across cities lowered by about 2.5 years (mean = 32.7), and proportion 40+ lowered by 12 percentage points (mean = 20%)
  - Effect on hiring of older workers is roughly as large as effect of hiring discrimination

Why are the results surprising/interesting?

- Not surprising if blatant ageist language is perceived as ageist
  - E.g., one of our "benchmarking" treatments uses language suggested by AARP that is sufficiently blatant (e.g., "digital native")
- Two reasons our evidence is much more interesting (and surprising) than we would get from such examples
  - Our evidence speaks to real-world job-ads and phrases with language that is much more subtle
    - Constructed from ~14,000 job ads from our correspondence study
  - The phrases we study predict discrimination by employers, as measured experimentally in our earlier correspondence study
  - Core finding plus other evidence suggest that discriminating employers use this language intentionally to discourage applications from older workers

Potential policy implications (I)

- Lowering applications from older job seekers can have same impact as direct age discrimination in hiring, but we focus more on the latter
- Explicit language ("over 40 need not apply") is illegal, and we don't see it in job ads
- CFR also defines as illegal less explicit phrases:
  - E.g.: "Notices or advertisements that contain terms such as age 25 to 35, young, college student, recent college graduate, boy, girl, or others of a similar nature violate the Act unless one of the statutory exceptions applies" (§1625.4)
- Our evidence implies much more subtle language can act as a form of age discrimination

Potential policy implications (II)

- EEOC can issue stronger guidance to employers to use age-stereotyped language in job ads
- EEOC might investigate firms that use age-stereotyped language in their job ads
  - Discrimination may be occurring even in the absence of shortfalls between the share of older applicants hired and the share of older workers who apply for jobs
  - If share of older applicants declines due to job-ad language, then age shortfalls in hiring out of the applicant pool – the most definitive evidence of hiring discrimination – can be obscured
    - The applicant pool can be "tainted"
- Current enforcement mechanisms may be missing an important channel of age discrimination

Past experimental evidence on age discrimination in hiring

- Most credible evidence comes from field experiments – resume-correspondence studies (Fix and Struyk, 1993; Gaddis, 2018; Neumark, 2018)
- Have been many on race, ethnicity, sex, and in a few cases, age
- Applied to age
  - Create fictitious but realistic job applicants who are on average equivalent except for age, which is signaled through school graduation year(s)
  - Apply for real job openings
  - Hiring discrimination measured by comparing interview request rates ("callbacks")
  - Previous studies uniformly point to substantial age discrimination in hiring (e.g., Bendick et al., 1997, 1999; Lahey, 2008; Farber, Silverman, and Von Wachter, 2017)

Recent experimental evidence on age discrimination in hiring (Neumark et al., 2019)

- 40,223 applications (resumes) to 13,371 job positions in 12 cities (in 11 states)
  - By far the largest resume correspondence study
  - Large number of job ads included in the study is critical to the methods we use in the present paper
- Occupations that, based on CPS data, older individuals often take as new jobs (hence likely bridge jobs/unretirement): administrative assistant and retail sales for women, and retail sales, security, and janitor for men (positions also common for younger workers)
- 3 applications per position: one younger applicant, and two older applicants of different ages (49-51 or 64-66) or with different work experience histories
- Study addresses potential biases in past studies, including treatment of experience, differences in variance of unobservables

Experimental evidence on age discrimination in hiring (Neumark et al., 2019)

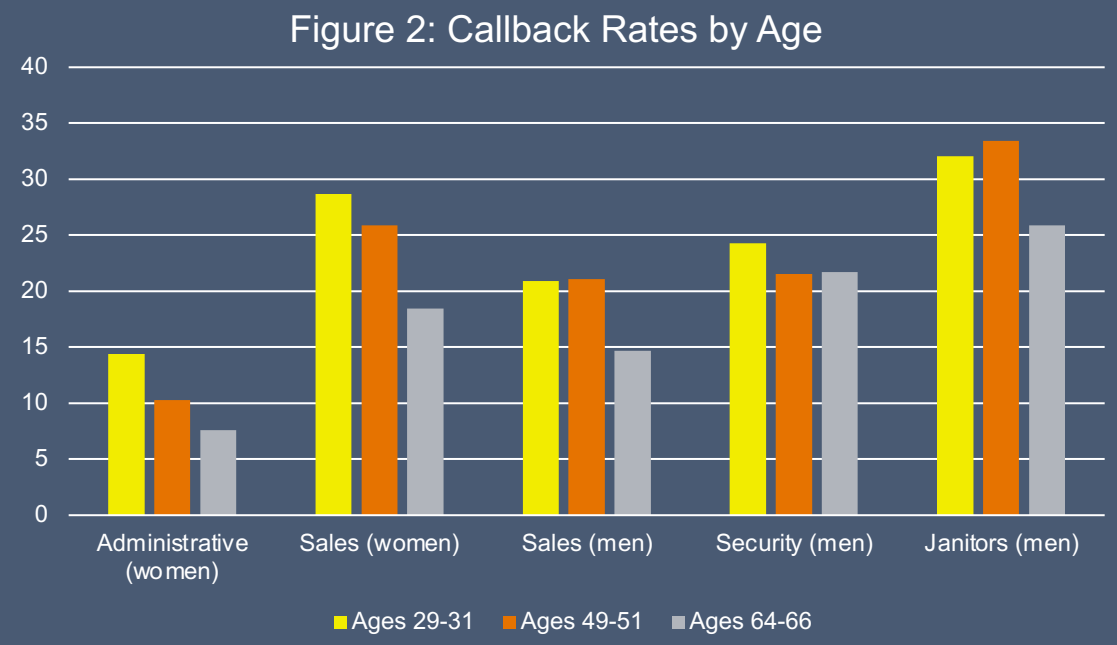

Source: Neumark et al. (*JPE*, 2019)

Ageist stereotypes and discrimination

- Research proposal from correspondence study:
  - "We will retain the text of the job ads to which we apply to see if we can detect ageist stereotypes in the ads that might predict employer behavior."
  - Good prediction for research project; bad prediction for how complicated it was to do this

Does ageist job-ad language predict age discrimination in experiment? (Burn et al., 2022)

- Why would discriminatory employers use ageist language in job ads?
- Hypo. 1: Intentional discrimination: use stereotyped language to shape applicant pool, reduce detection of discriminatory behavior
  - Could be attributable to taste or statistical discrimination
- Hypo. 2: "Innocuous?" statistical discrimination: jobs have different requirements, stated in job ads; employers hold stereotypes about older workers' ability to meet them, and act on them when older workers apply
- Both illegal, although second nuanced, depends on whether requirement is for RFOA
  - "… a non-age factor that is objectively reasonable when viewed from the position of a prudent employer mindful of its responsibilities under the ADEA under like circumstances"
    - E.g., Hodgson vs. Greyhound, upholding a maximum hiring age (going beyond decisions based on a factor related to age)
  - I'll come back to why – based on our new paper – we think intentional discrimination is most likely explanation

Steps in analysis of ageist stereotypes and discrimination

- Identify common age stereotypes from industrial psychology and related literature
- Scrape text of job ads from experiment, and use machine learning/computational linguistics to identify relationships between job-ad language and age stereotypes
- Use all phrases in each job ad to identify ads that have high "semantic similarity" with each of the age-related stereotypes
- Estimate regression models to identify which age stereotypes in job ads predict age discrimination as measured by our experiment

Step 1. Stereotypes from industrial psych literature

- Detailed review of literature
- Compiled list, including different phrasings
- 17 stereotypes
  - 11 negative
  - 6 positive
    - Some are ambiguous: e.g., “worse communication skills” and “better communication skills”

All 17 stereotypes based on industrial psychology lit., alternative phrasings

Health
Less Attractive
Hard of Hearing
Worse Memory
Less Physically Able

Personality
Less Adaptable
Careful
Less Creative
Dependable
Negative Personality
Warm Personality

Skills
Lower Ability to Learn
Better Communication Skills
Worse Communication Skills
More Experienced
More Productive
Less Productive
Worse with Technology

Step 2: Matching stereotypes to words and phrases in job ads

- Problem: we do not expect stereotypes in job ads to be expressed exactly as they are in research literature – "string matching" impractical (and also ineffective and subject to data mining)
- Strategy: use computational linguistics/machine learning to determine semantic similarity between *all* phrases in job ads, and age-related stereotypes from research literature
- Steps:
  - 1. Train model using text from corpus of Wikipedia
  - 2. Apply weights that allow measuring similarity of words to phrases (3-word trigrams in ads, and age stereotypes)
  - 3. Compute semantic similarity score between every 3-word trigram in ads and all of the stereotypes
    - Ranges from −1 (words never appear in same inputs) to 1 (identical, always appear in same inputs)

Sample ad (selected text)

Part Time Insides Sales/Customer Service Rep

The ideal candidate will be operationally minded and have excellent sales abilities. Candidate must be reliable, energetic, customer-friendly, detail-oriented, and possess excellent time management skills.

Our Mission is to provide an unparalleled experience by fostering an environment that inspires teamwork and unrivaled passion for exceptional service.

Essential duties and responsibilities include:

- Provide excellent customer service.
- Perform order and data entry with precision in a given time frame.

Requirements:

- Must possess excellent telephone etiquette, communication and organizational skills
- Candidate must be self motivated and be able to work independently.
- Effective time management
- Must be computer savvy; experience with MS Outlook and Office required
- 2-4 years of inside sales/customer service experience preferred.

Our company is a stable, employee oriented organization, providing a fun, team oriented atmosphere.

Example of distribution of CS scores

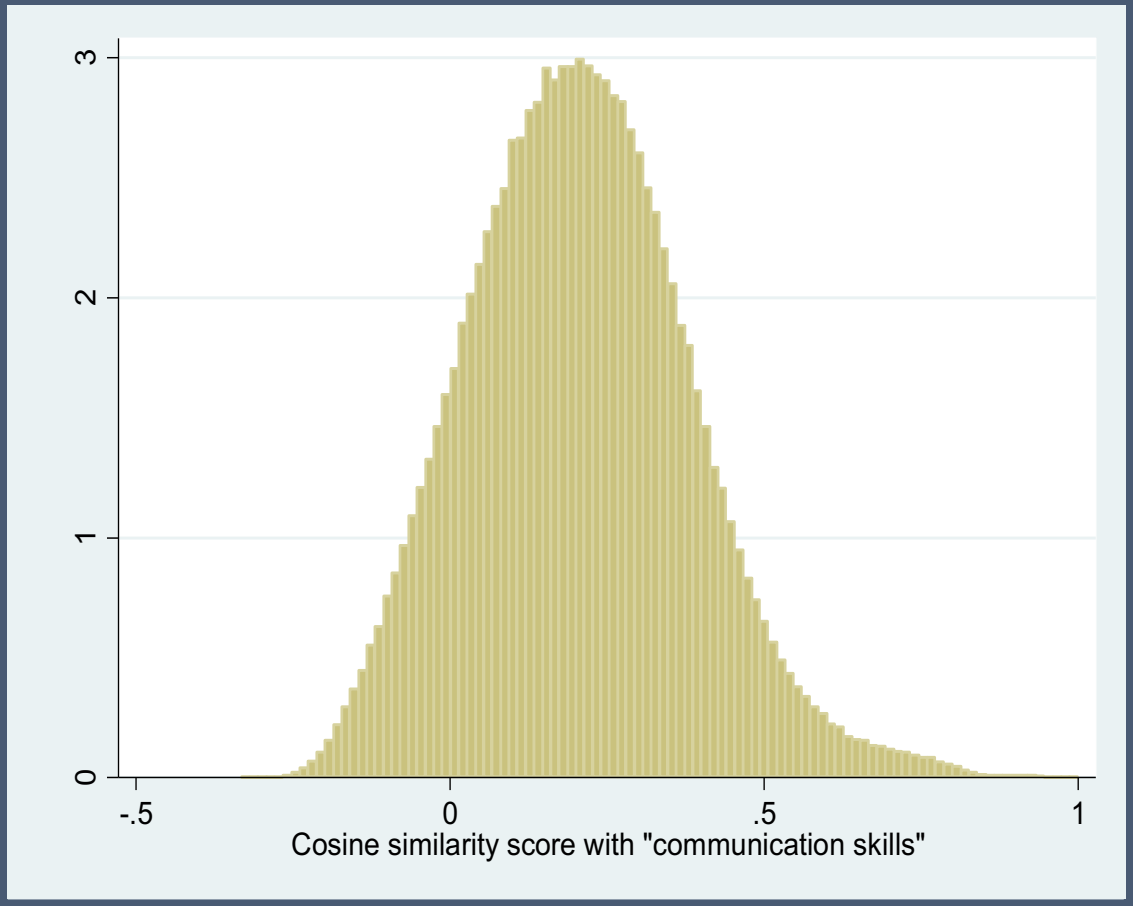

- Centered above zero, since we are looking at text from job ads
- Examples:
  - Near −0.3: “Christmas season near,” “hotel near Seattle”
  - Near 0: “every Sunday pm,” “work year round”
  - 1.0: “excellent communication skills,” “prioritizing skill communication”

Step 3: Characterize language in each job ad

- Take an individual job ad, and calculate distribution of CS scores for each phrase in ad with each
- Use 95th percentile of scores for each stereotype to "characterize" how stereotyped that ad is in relation to that stereotype
  - Idea is to capture whether ad has a small number of phrases strongly related to the stereotype, and to distinguish ads based on their most stereotyped phrases

Step 4: Identify which stereotypes in job ads predict lower callbacks to older applicants

- Core hypothesis: job ads with negative age stereotypes will have relatively lower callback rates for older applicants
- Data set is all responses to triplet of job applications sent in response to each ad, which could be matched to employer and ad
  - 34,260 job applications to 11,420 job ads, corresponding to 22,840 observations on middle-aged or older applicants
- Outcome D_{ij}: younger applicant called back, middle/older applicant not called back (0/1) – "experimental measure of discrimination"
  - Estimate probit model
    - $Pr[D_{ij} = 1] = \alpha + \sum_s \beta_s P_{js}^{95} + X_{ij}\delta + \varepsilon_{ij}$

How does the job-ad language vary?

		Trigrams closest to 1 standard deviation above mean				
	Mean 95th percentile	2 trigrams below	1 trigram below	1 standard deviation above the mean	1 trigram above	2 trigrams above
	(1)	(2)	(3)	(4)	(5)	(6)
Physically able	assistant position available	fast paced fun	experience preferred necessary	work preferred necessary	required flexibility required	anything required help
Adaptable	making learning agility	must reliable apply	duties needed excellent	dedicated providing well	individual fast learner	moment respectful willing
Careful	computer skills required	basis important able	utmost professionalism integrity	processing payments must	requirements need really	assistant good verbal
Creative	position energetic detail	support vp marketing	office experience ability	involves spearheading independent	professional personal presentation	experience skills essential
Dependable	media reputation management	reliable transportation needed	service skills critical	hours retirees welcome	professional attitude well	experience excellent telephone
Ability to learn	word excel must	sales skills managing	communication skills competency	career using cashiering	excellent people communication	abilities ability complete
Communication skills	politeness absolutely fundamental	backend customers required	skills abilities ability	client relation skills	filing typing computer	communication organizational skills
Experienced	evenings weekends requires	computer knowledge exciting	duties qualifications experience	level gain exposure	quickly become one	full time receptionist
Productive	fun eventful one	banker seeking ambitious	computer literate interested	tasks experience necessary	organization efficiency must	player good communication
Technology	paced integrative medical	ability learn new	willingness learn new	social media google	communication skills proficiency	computer skills proficiency

Results

- Job-ad language highly related to ageist stereotypes is associated with experimental measures of hiring discrimination by age from correspondence study
  - For men, age stereotypes about all three categories we consider – health, personality, and skills – predict age discrimination
  - Most evidence in direction predicted by IP literature
- Weaker evidence of this for women, but field experiment detected more age discrimination against women
  - Stereotypes about older women may be harder to express in job-ad language
  - Or language may be used less because age discrimination laws less effective for older women (intersectional claims?)

Current paper

- Field experiment to test whether ageist stereotypes in job ads deter older workers from applying
- Utilizing ageist language to deter older applicants may be rational for employers
  - Avoid detection of age discrimination (taste discrimination)
  - Express job requirements that are correlated with age
    - Could be rational, but illegal
  - Lower search costs for employers who don't want to hire older workers by signaling ageism

Most closely related research

- Tilcsik (2011) identifies words in job ads related to masculine stereotypes and links them to hiring outcomes in a correspondence study of discrimination against gay men
- Card et al. (2021) study a policy change in Austria where the government launched a campaign to inform employers and newspapers that gender preferences in jobs ads were illegal, and relates to hiring of women and men in different jobs
- Delfino (2021) manipulates photos of workers in job ads to convey the gender composition of the workforce and studies effects on women and men applying
- Kuhn and Shen (2021) study the impact of removing preferred gender fields from ads on a Chinese job board on applicant responses and study applications and callbacks to women and men
- Last 2 are experiments in the flavor of ours; ours is unique in studying subtle shifts in language rather than explicit statements (e.g., diversity statements), in focusing on age, and in tying results to experimental estimates of discrimination from earlier study

Experimental variation/design

- Chose 3 of the IP stereotypes :
  - Commonly expressed in job-ad language about ideal/preferred candidate or attributes
  - Evidence from prior study of link to measured employer discrimination
  - Older workers aware that employers hold stereotypes (from, e.g., AARP reports)
- Communication skills, physical ability, and technological skills
- 3 of the 4 four occupations from original study
  - Retail sales, administrative assistants, security guards

Cities

- 14 cities selected due to large size, geographic distribution, and different population age distributions
  - Two more than in initial field experiment PAP because of large presence on on-line job board
  - Size corresponds to # of applicants

Key study design element: designing job ads

- Created job-ad templates for each city-occupation combination, based on actual ads from correspondence study, supplemented with additional ads from job board
  - Used similar format
  - Rewrote to give enough details about company and job to be realistic, but not as specific to company as actual ad
  - Modified job requirements to reduce number of applicants potentially excluded
    - Flexible hours, competitive pay, PT and FT
    - Randomized ½ to require HS degree

Stereotyped job requirements

- Design of primary treatment(s):
  - Calculate semantic similarity of thousands of job ad phrases drawn from actual ads
  - Insert into job ad sentences/phrases
  - Iteratively edit sentences/phrases (including treatment and control phrases), and entire ads, to create high semantic similarity with targeted stereotype, but not with the other two (and not the others from our longer list)
  - Try to have related “skills” in treatment and control text

Administrative asst example

Psychiatric office is in need of a full or part time Administrative Assistant to assist in front/back office general clerical duties. This individual will work on a several tasks and stay on course at all times. The Administrative Assistant we hire will be trained in various duties that cover the entire office.

This individual MUST possess the foll

-Exceptional customer service backgr
patients, answer phones, schedule ap
-Can multitask.
-High School diploma or GED.
-Professional attitude.
-*Communication Skill Requirement*.
-*Technology Requirement*
-*Physical Requirement*
-Available for flexible hours.
(Schedule hours and days will alternate every other week)

Please email us a CV or resume and put “full-time” or “part-time” in the subject line.

Vary as treatment or control phrases, with 4 treatments:

--Single stereotyped phrase for one of these

--All 3 stereotyped phrases

Treatment and control sentences by stereotype (selected)

Occupation	Stereotype	Control	Machine Learning Treatment
(1)	(2)	(3)	(4)
Administrative Assistants	Communication skills	You must be good at working without supervision (CSS = 0.20)	You must have good communication and teamwork on tasks (CSS = 0.48)
Administrative Assistants	Technological skills	You must produce and distribute documents such as correspondence memos, faxes and forms (CSS = 0.08)	You must use accounting software systems like Netsuite, Freshbook, and QuickBooks (CSS = 0.29)
Retail sales	Physical ability	You must enter bills and keep track of invoices (CSS = 0.11)	You must be able to lift 40 pounds (CSS = 0.41)
Retail sales	Technological skills	You must help to clean and organize the store (CSS = 0.09)	You must use software such as Microsoft Office/Excel or Google Sheets (CSS = 0.27)
Security guard	Communication skills	You must follow instruction from supervisors (CSS = 0.21)	You must maintain communication about tasks with supervisors (CSS = 0.38)
Security guard	Physical ability	You need to carry a flashlight (CSS = 0.20)	You must be able to lift 50 pounds (CSS = 0.41)
Security guard	Technological skills	You must write patrol records in journal notebook (CSS = 0.03)	You must type patrol entries into a journal application on a computer system (CSS = 0.24)

Locations of treatment (dashed) and control (solid) phrases in CSS distribution of all job ad phrases

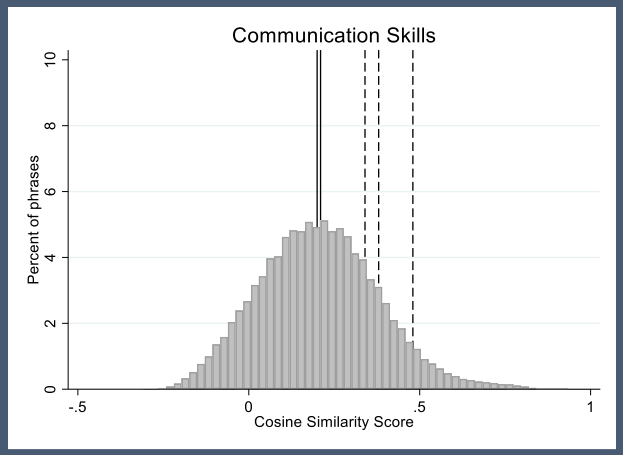

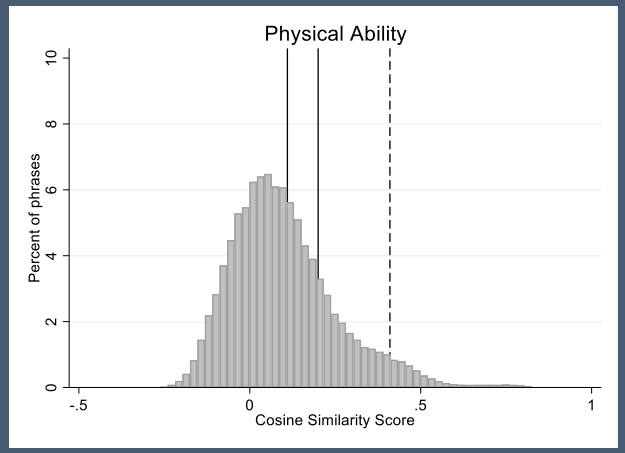

Control phrases near median and wouldn't be regarded as ageist

Treatment phrases near 75th percentile – so ageist, but not blatant

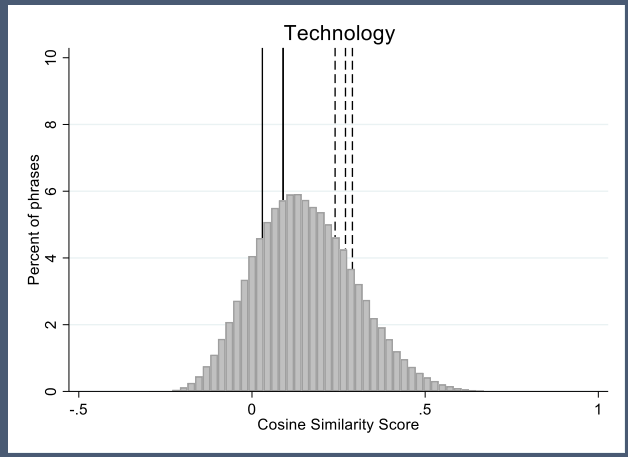

How well do our job ads convey intended stereotypes?

- Would like treatment ads to have high "sensitivity" (conveying ageist stereotypes) and "specificity" (conveying information about the specific ageist stereotype intended)
- First type of evidence:
  - Take all three-word phrases from each ad
  - For each ad compute the median, mean, and each percentile up to 99th of CS score with stereotype
  - Average these "moments" across all ads with different treatments

How well do our job ads convey intended stereotypes?

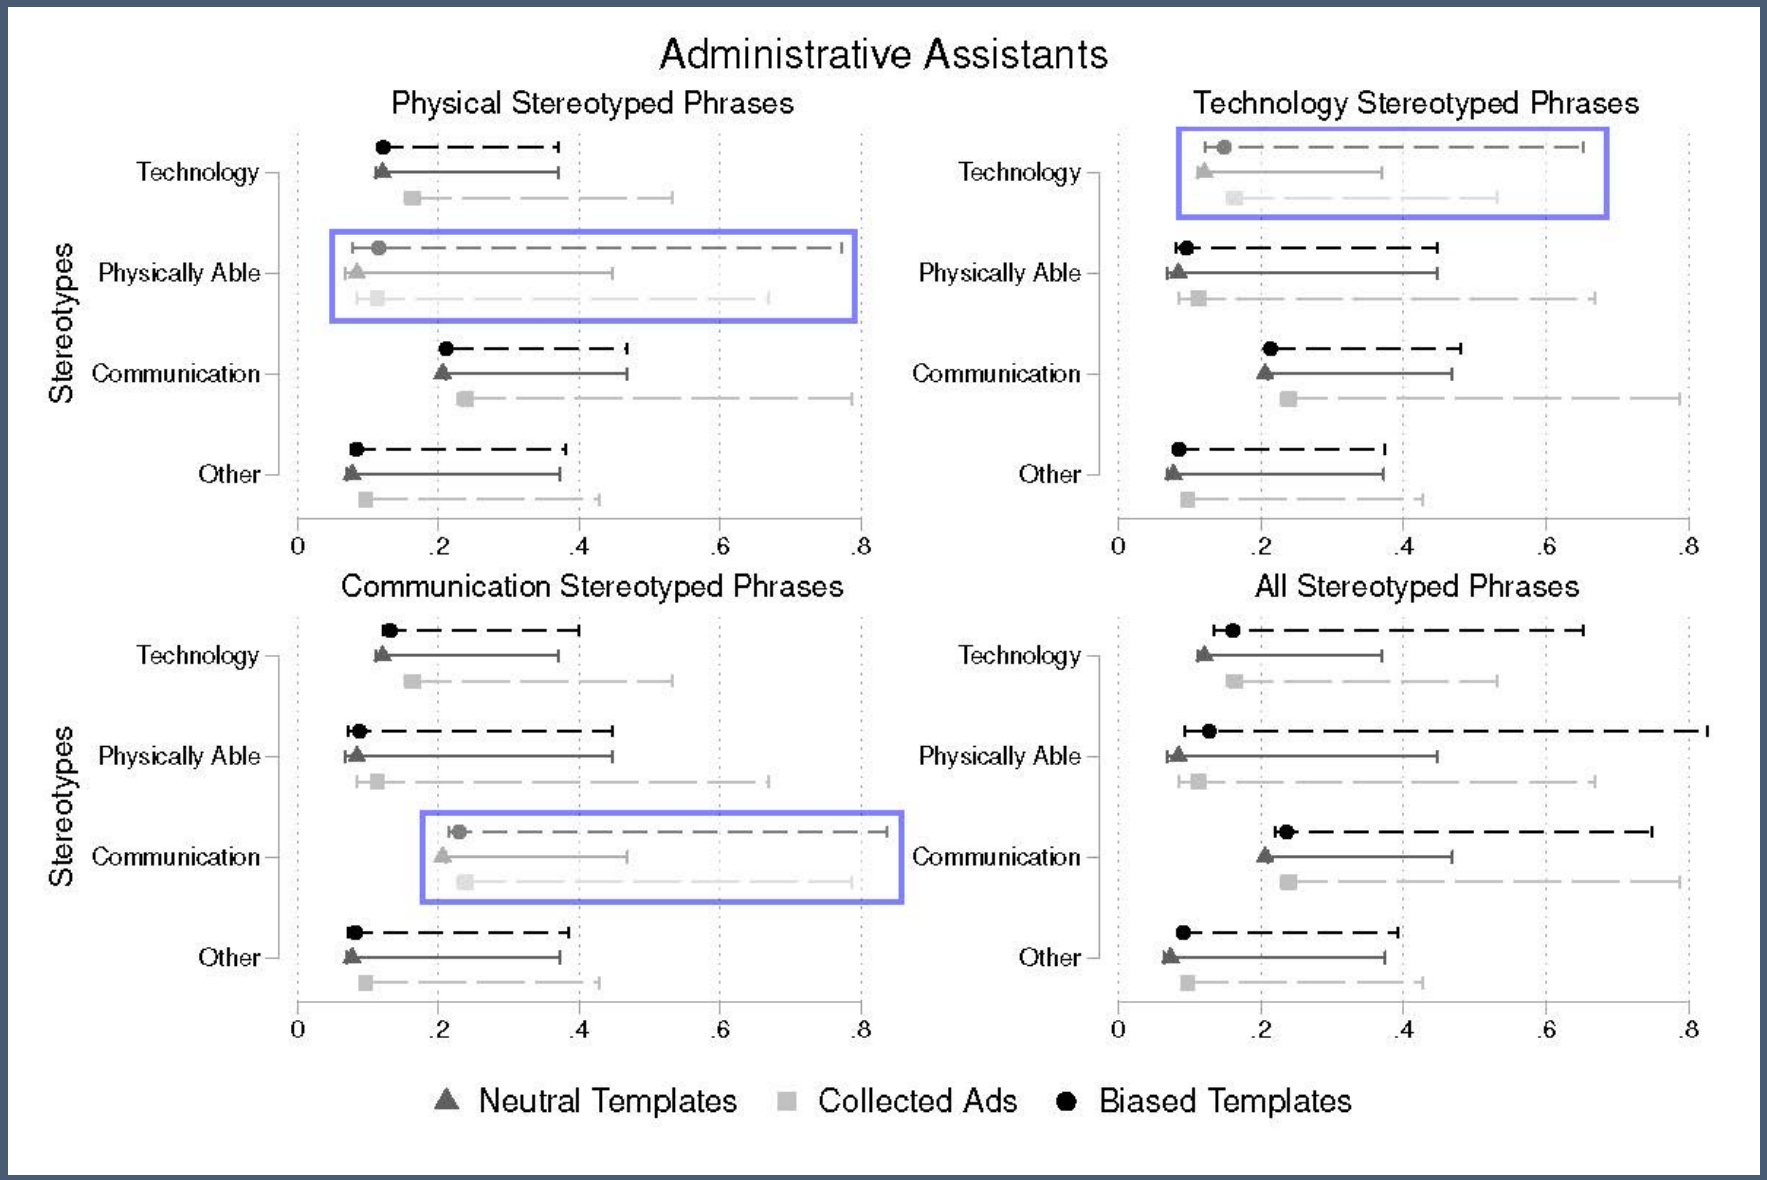

How well do our job ads convey intended stereotypes?

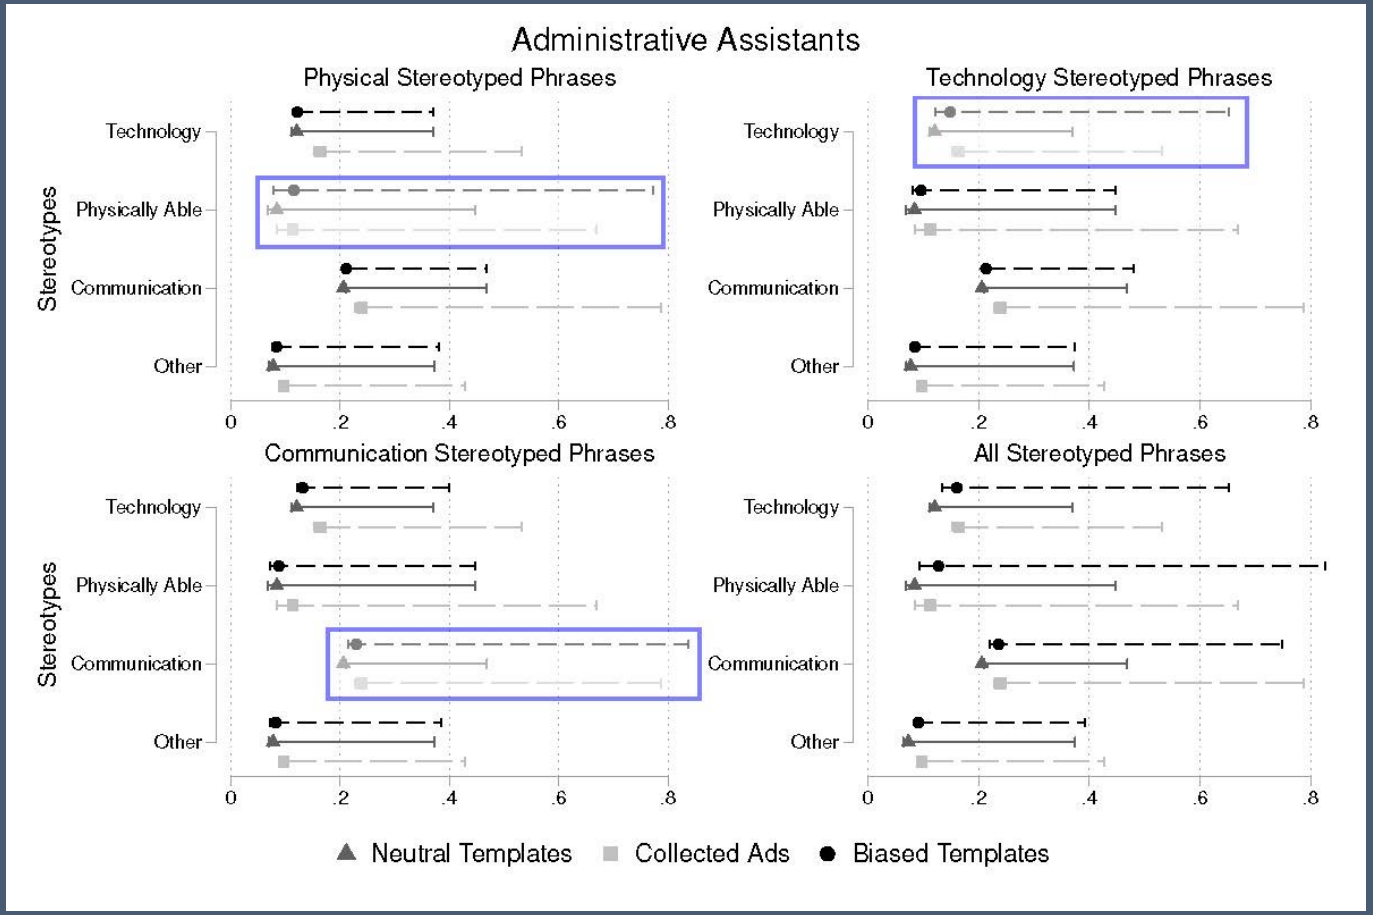

Biased treatments have higher upper percentiles and means (vs. controls)

True for "intended/activated" stereotype, but not others

(Note collected ads more similar to treated ads – *these stereotypes get used in ads*)

Secondary treatment

- Phrases identified as ageist by AARP, among those closer to our stereotypes
  - "Cultural fit," "energetic person," and "digital native"
  - Adapted to fit job ads and created three treatment sentences
    - "You must be up-to-date with current industry jargon and communicate with a dynamic workforce" – communication skills (cultural fit)
    - "You must be a fit and energetic person" – physical ability (energetic person)
    - "You must be a digital native and have a background in social media" – technological skills (digital native)
  - More blatant, designed to tell us whether experiment could be informative, in case we found no effect of our experimental treatments using more subtle phrases ("upper bound" of effect)

Are the treatment phrases perceived as ageist relative to the control phrases?

- If not, absence of evidence of effects in our experiment may be uninformative
- Ran experiment/validation exercise on MTurk, asking about perceived bias against workers 50+ for treatment and control phrases (Burn et al., 2023)
- First elicited respondents' self-assessments, then asked others to predict those responses
  - Avoid social desirability biases (which we found)
  - Eliciting predictions allows us to incentivize responses

How well do our job ads convey intended stereotypes?

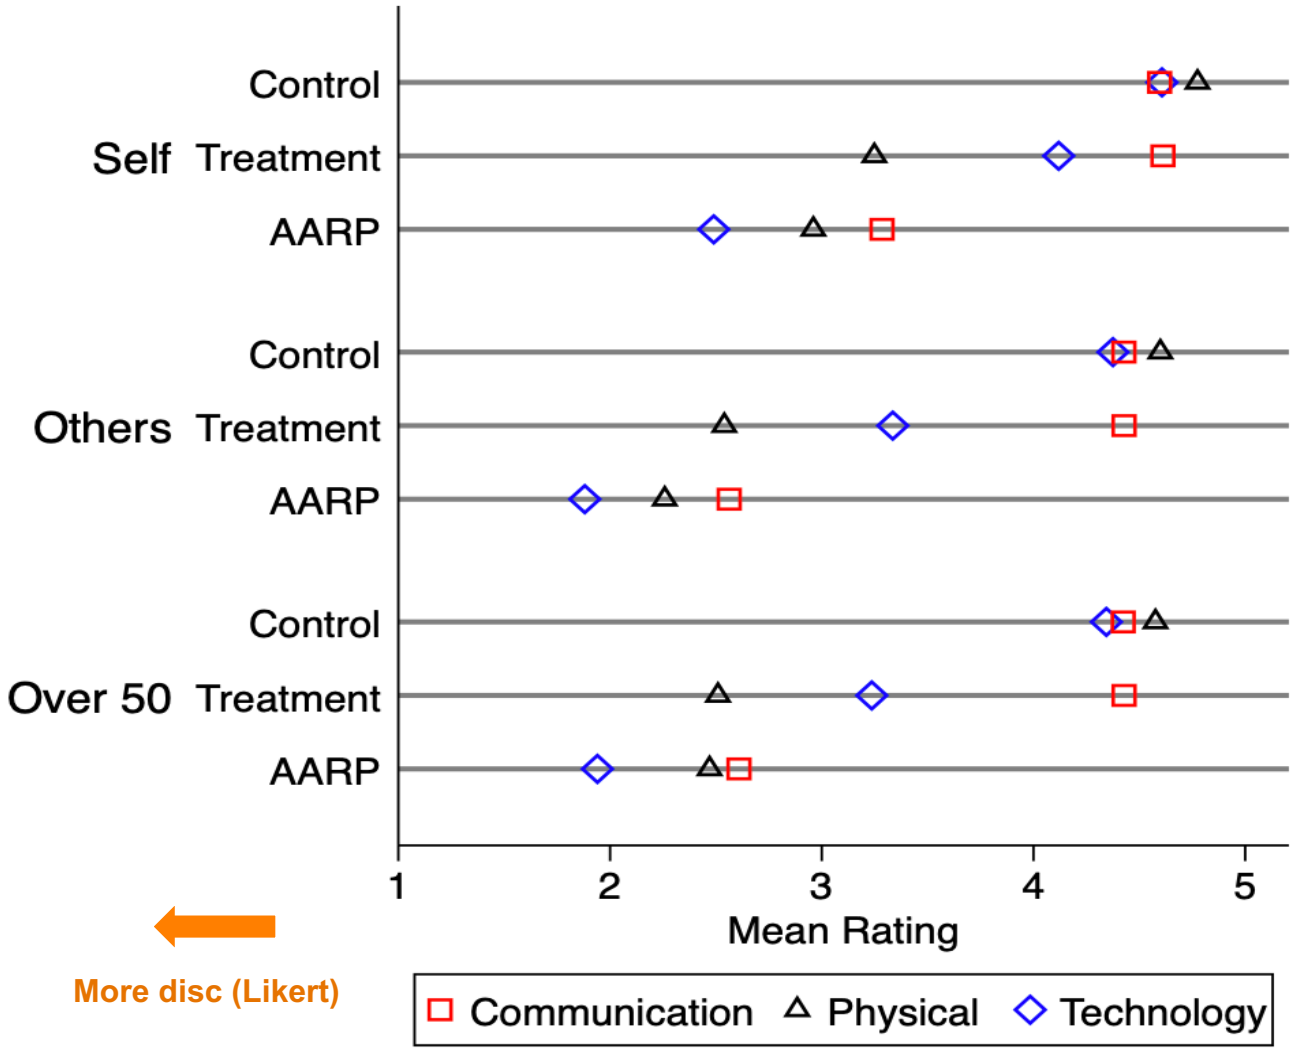

Is job requirement indicating "treatment/control phrase" biased against workers over the age of 50?

"Strongly Agree" -- 1
"Somewhat Agree"– 2
"Neither agree nor disagree" – 3
"Somewhat Disagree" – 4
"Strongly Disagree" 5

Summary of MTurk findings

- ML/CL treatment phrases were perceived as more ageist than the control phrases
  - AARP phrases were perceived as most ageist (biased against workers over age 50)
  - Less apparent for communications skills, but that is not an unambiguously negative stereotype in the IP literature
  - (Results “weaker” for self, indicating social desirability biases and importance of avoiding them)
- This evidence indicates our experiment should be informative

Plan for posting the job ads

- 18 ads (6 arms x 3 occupations) per city
- Staggered posting of ads with timing to capture most responses and create long period between posting for same occupation (different arm) in same city
  - Never had two ads up for same occupation in same city
- To avoid p-hacking, specified these numbers in advance, with schedule hence taking 54 weeks, with ads posted randomly on different days of the week (M-F)

IRB approved and pre-analysis plan registered

- UCI Office of Research Institutional Review Board on October 18, 2019: HS# 2015-2017, modification application #26404
  - To minimize cost to participants, informed within 24 hours that we had decided to hire a different candidate
  - Informed consent not possible, so applicants notified and given option to have data excluded from study
    - Only a handful requested exclusion of their data
- Pre-Analysis Plan for this project was registered on Open Science Framework on December 31, 2020 (now public)
  - "Plan B": Anticipated some challenges, built into PAP
  - Some additional analyses added based on feedback; clearly delineated

Implementation of job posting was challenging

- Job board makes money from fees for posting job ads, and hence tries to screen out fake ads used for phishing, etc.
- Human "checkers" for each city on the job board monitor for highly similar ads or ads that appear to be from fictitious companies
- Problems encountered if we tried to use the same credit card to pay for ads in different cities, or used the same IP address for posting ads in different cities
  - Used large number of gift cards to avoid repeated use
    - Learned that gift cards used by money launderers, and credit card thieves, so many didn't work
  - Needed large number of IP addresses as gift card companies don't allow registration of many from same IP address
    - Hence purchased large numbers of cell phones and SIM cards, with PAYG plans that randomized IP addresses

“If the FBI comes to my apartment, I’m in big trouble!”

Calculating applicant age

- 1: based on year of HS graduation (assumed at age 18)
- 2: based on date of earliest work experience (assumed at age 16)
  - Minimal possible age, assuming didn't start work before age 16
- 3: earliest non-work year listed on resume (assumed at age 18, rare)
- Assigned oldest age based on these
  - Avoids assigning too young an age from someone who leaves off earlier job market experience, unless they don't list year of HS graduation
- Used age or birthday if explicitly noted, but that was rare

Bias from omitted or manipulated age information?

- Older applicants may hide age, especially in response to ageist ads
  - No a priori check – had to roll the dice
  - Tested for effect of ads on reporting age information – no impact
  - Some applicants (400) applied to more than one ad – no impact of treatment on reporting lower age
- Suggests applicants don't manipulate age-related information when applying in response to job-ad language

CDFs based on individual data (any treatment)

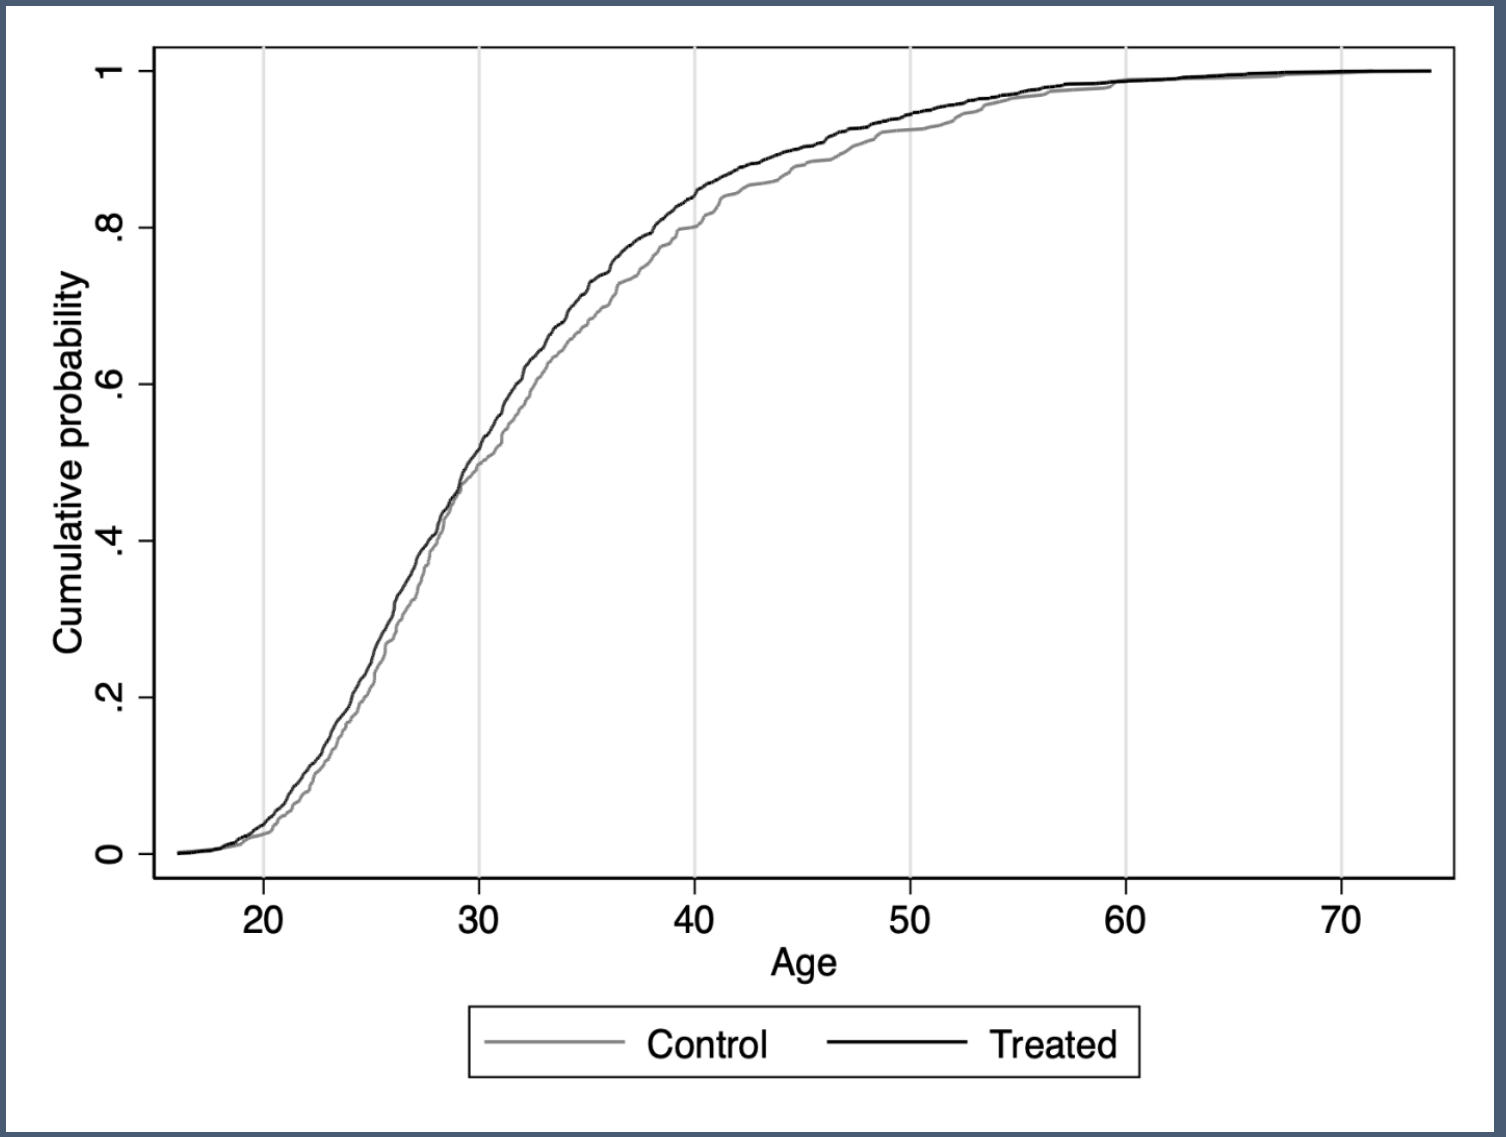

CDFs based on individual data

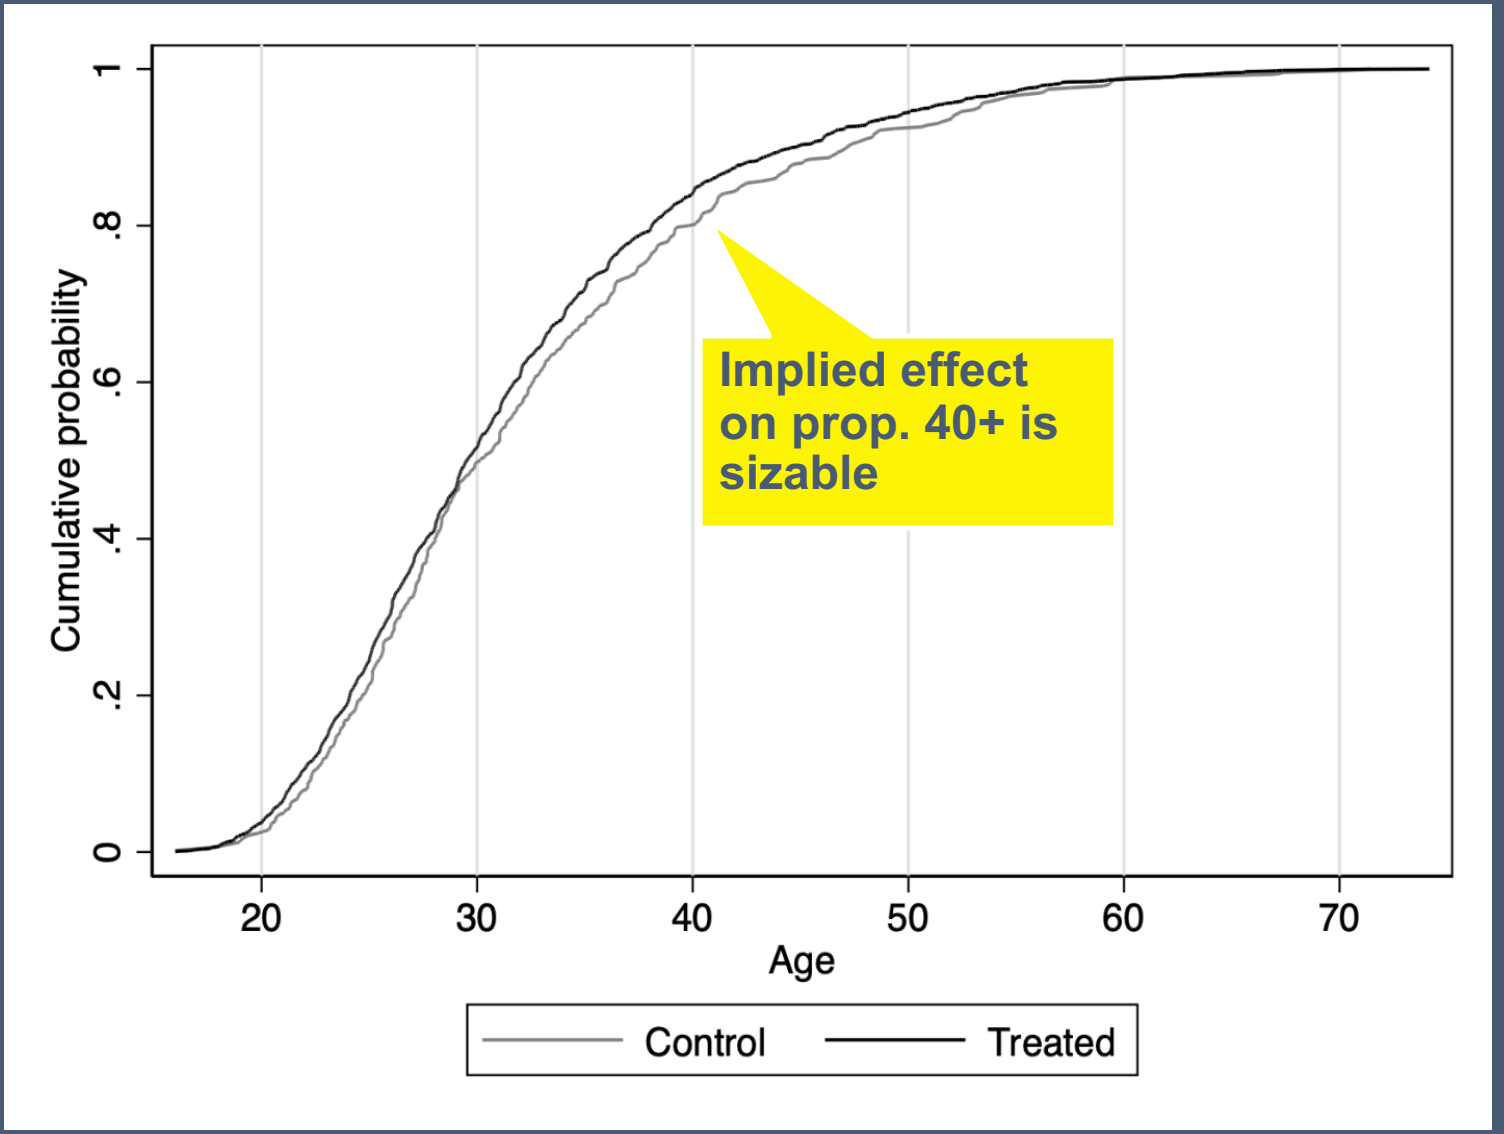

CDFs based on individual data

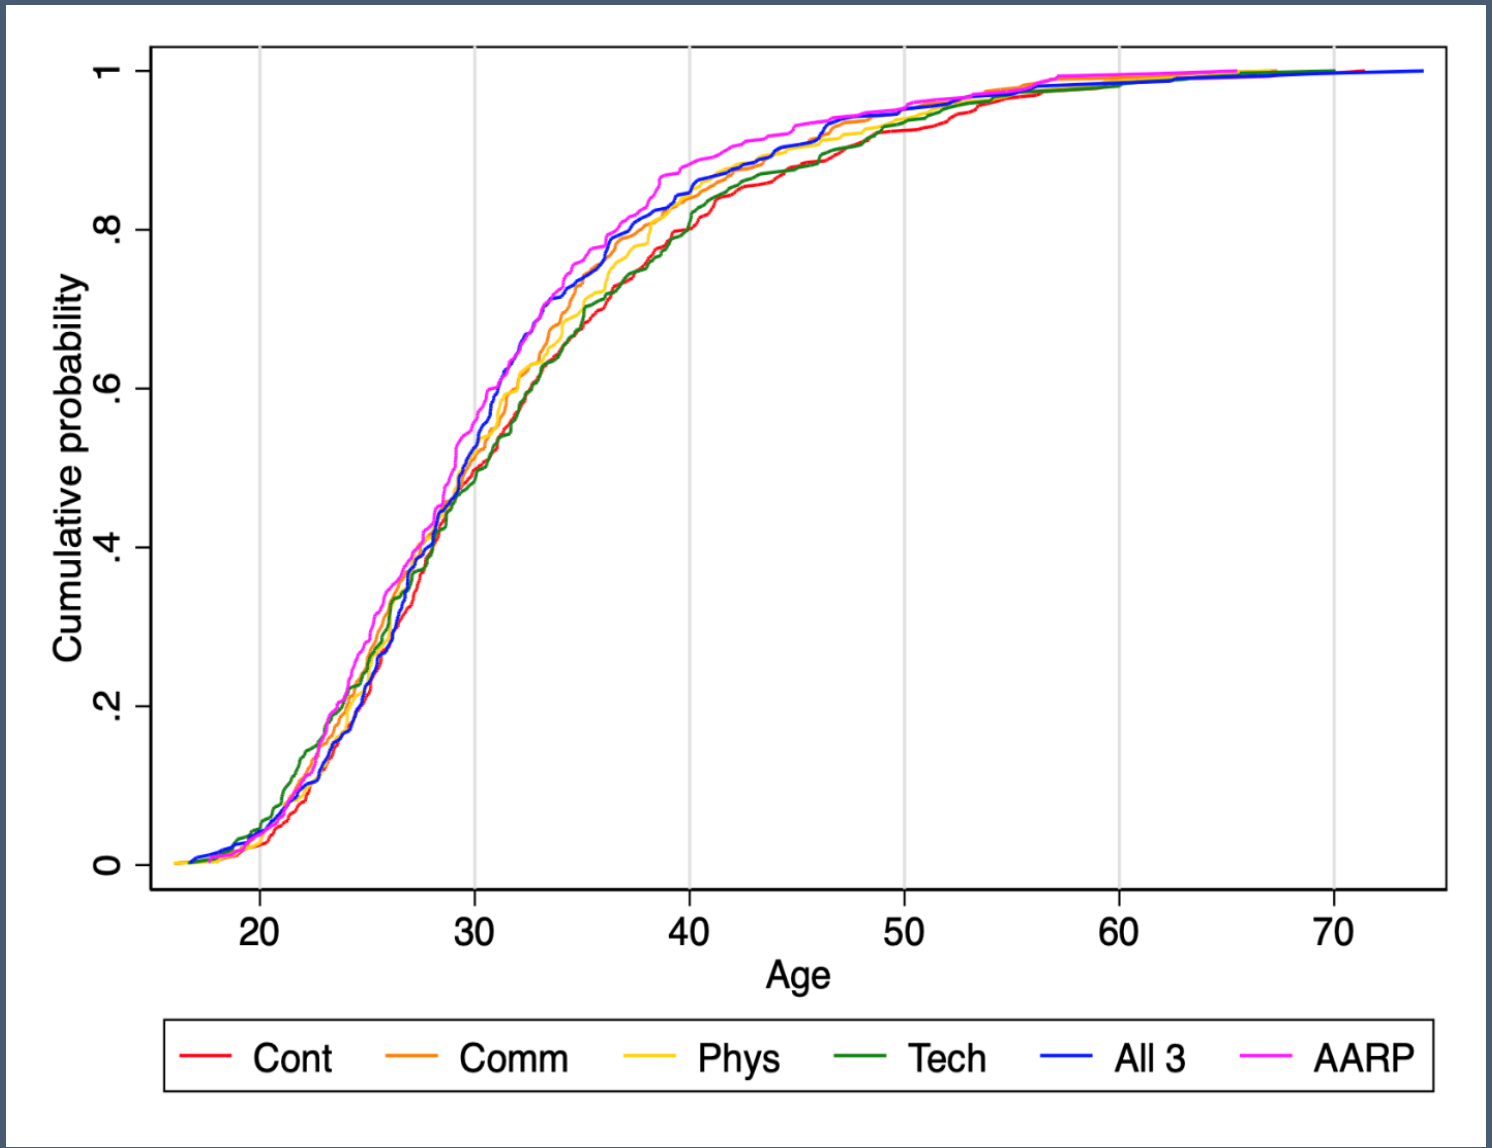

CDFs based on individual data

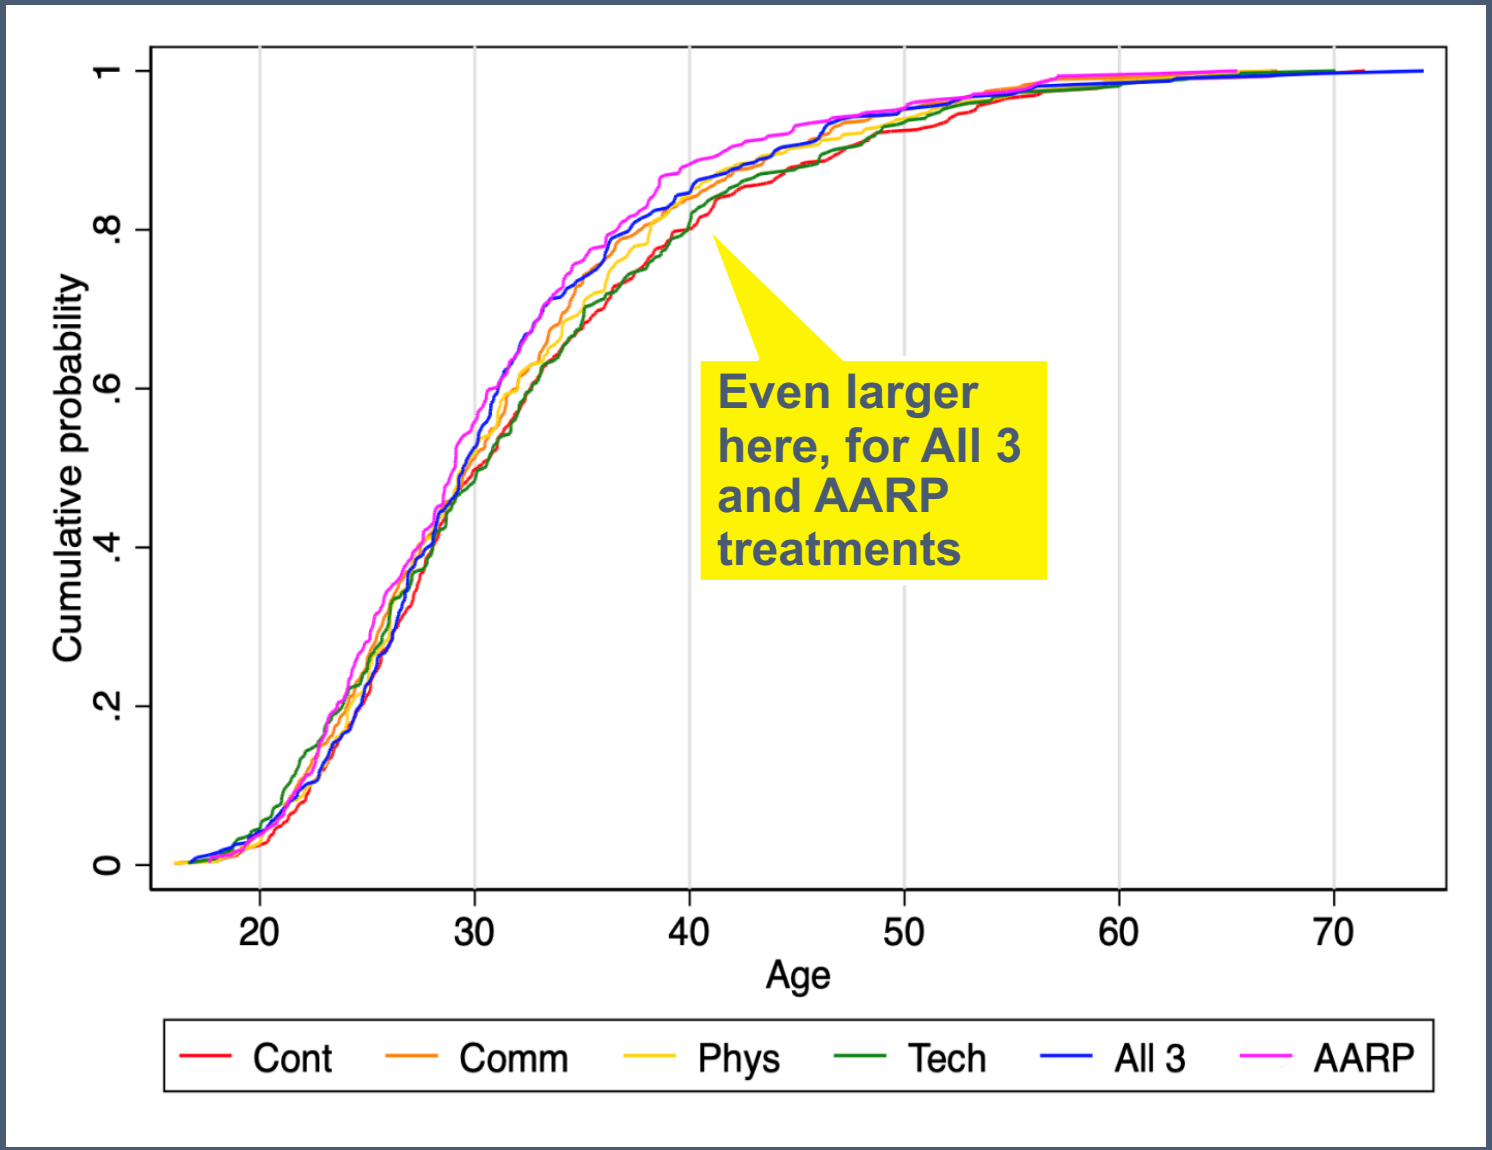

Empirical density functions show hollowing out above 40 (to about 60)

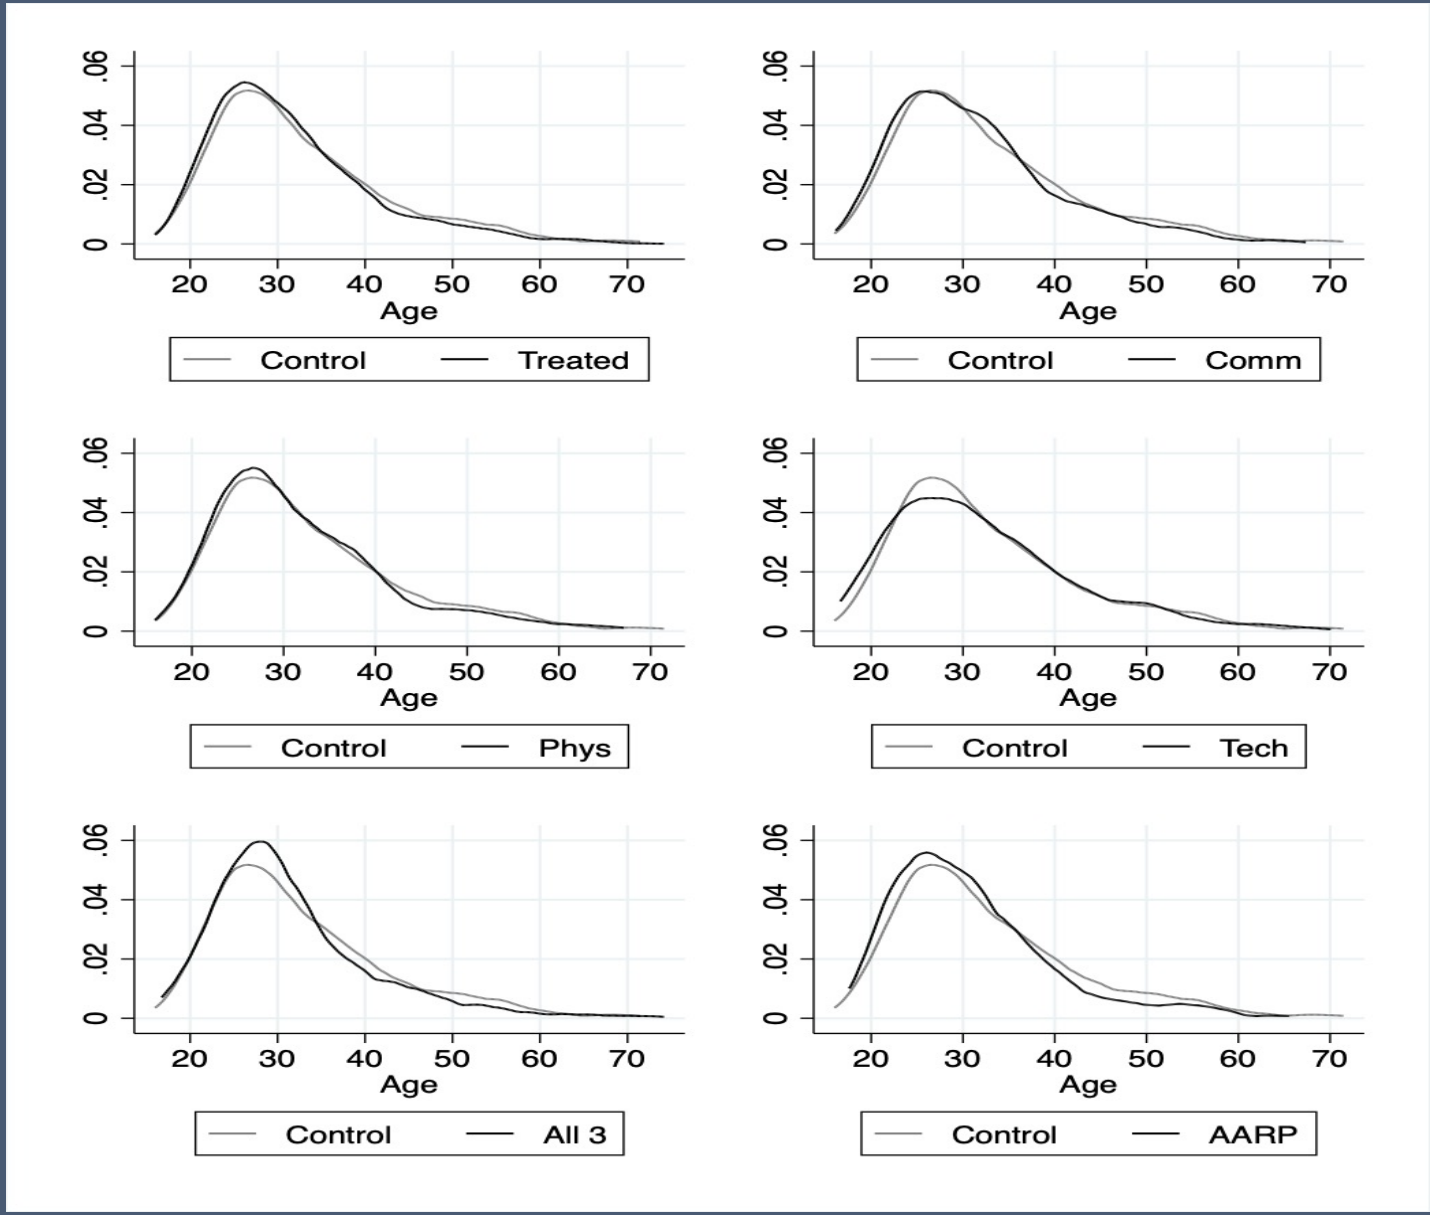

Regression results for single stereotypes

	Average Age	Median Age	75th Percentile	Over 40
Communication	-2.632**††	-2.948*††	-2.583	-0.070†
	(1.232)	(1.460)	(2.043)	(0.050)
N	79	79	79	79

	Average Age	Median Age	75th Percentile	Over 40
Physical	-1.828†	-2.062†	-2.454	-0.082†
	(1.403)	(1.394)	(2.134)	(0.050)
N	79	79	79	79

	Average Age	Median Age	75th Percentile	Over 40
Technology	-1.930†	-2.105†	-0.728	-0.044
	(1.221)	(1.285)	(1.945)	(0.042)
N	79	79	79	79

Effect on *every* age measure is negative

Strong for communications (in contrast to MTurk evidence); so job applicants perceive such language as negative

Regression results for all treatment arms

	Average Age	Median Age	75th Percentile	Over 40	No Age Information
Communication	*-2.683***††	*-2.986***††	-2.704†	-0.075†	0.014
	(1.147)	(1.350)	(1.898)	(0.046)	(0.039)
Physical	-1.879†	-2.071†	-2.385	-0.083*††	0.034
	(1.288)	(1.276)	(2.000)	(0.045)	(0.051)
Technology	-1.889†	-2.002†	-0.707	-0.041	0.056
	(1.165)	(1.222)	(1.822)	(0.041)	(0.045)
All 3	*-2.516***††	*-2.504***††	*-4.156***††	**-0.117*****†††	0.016
	(1.122)	(1.104)	(1.754)	(0.041)	(0.048)
AARP	**-4.559*****†††	**-3.888*****†††	**-5.896*****†††	**-0.156*****†††	0.027
	(1.222)	(1.212)	(1.799)	(0.038)	(0.057)
N	228	228	228	228	237

Italicized (bold): sign at 10% (5%) in multiple testing (FDR: probability that at least some rejected hypotheses are false)

No age information column implies no manipulation of age; confirmed from repeat applicants

Only some columns point to "dosage" response

Even a single stereotyped phrase triggers a response among job applicants

Regression results for average cosine similarity score or average MTurk Likert scale (reversed)

	Average Age	Median Age	75th Percentile	Over 40	No Age Information
CSS	**-1.722****†††	**-1.534****††	**-3.139*****†††	**-0.083*****†††	0.008
	(0.652)	(0.637)	(0.972)	(0.023)	(0.031)
N	228	228	228	228	237
Likert score (perceived age bias)	**-2.048*****†††	**-1.524****†††	**-3.354*****†††	**-0.085*****†††	0.004
	(0.652)	(0.637)	(0.972)	(0.023)	(0.031)
N	228	228	228	228	237

Variables are standardized.
Italicized (bold): sign at 10% (5%) in multiple testing

Regression results for all treatment arms, other age cutoffs

	Over 40	Over 50	Over 65
Communication	-0.075†	-0.073*††	-0.005
	(0.046)	(0.037)	(0.010)
Physical	-0.083*††	-0.042	-0.009
	(0.045)	(0.040)	(0.007)
Technology	-0.041	-0.038	-0.004
	(0.041)	(0.036)	(0.009)
All 3	-0.117***†††	-0.081**††	-0.006
	(0.041)	(0.034)	(0.007)
AARP	-0.156***†††	-0.090**†††	-0.008
	(0.038)	(0.035)	(0.008)
N	228	228	228

Smaller and insignificant effect at oldest ages consistent with what we saw in descriptive figures – few observations at those ages, and main impact was hollowing out from about 40 to 60 – but % changes actually large

Regression results for any treatment arms, many alternative age cutoffs

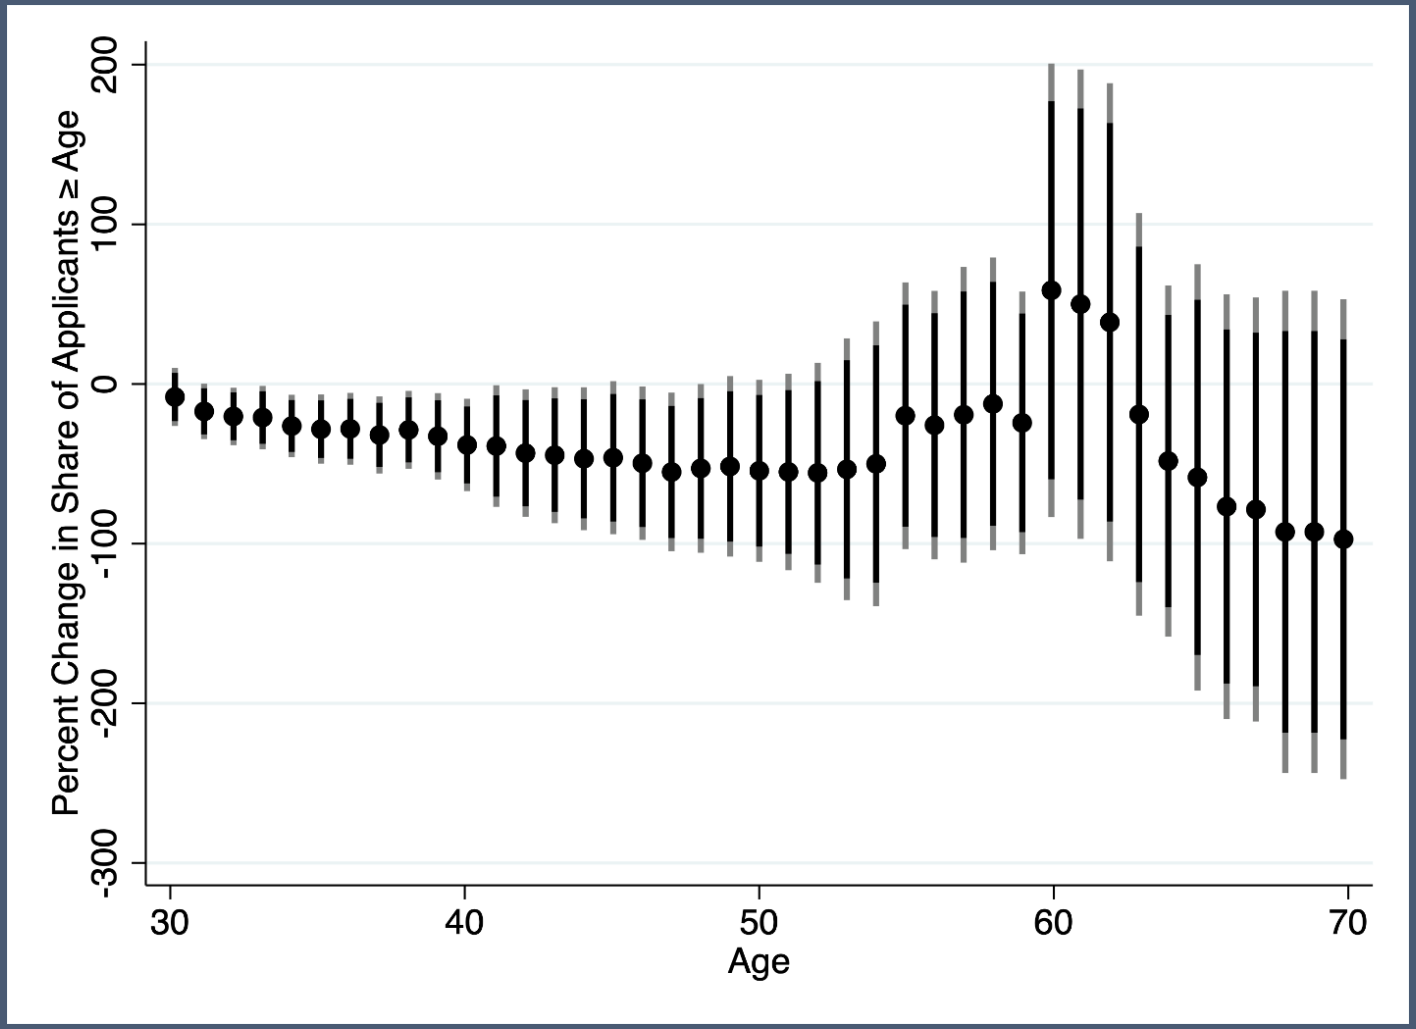

Smaller and insignificant effect at oldest ages consistent with what we saw in descriptive figures – few observations at those ages, and main impact was hollowing out from about 40 (or a bit younger) to 60 – but % changes actually larger

Results driven by discouraging applications from older job seekers

- Decline in average age of applicants or share of applicants over age 40 could be generated from increases in the number of people under age 40 who apply for jobs
  - Could happen because the ageist stereotypes are presented in a positive rather than a negative light
- Key result: Applications for younger job seekers do not increase, so our evidence implies that job-ad language related to ageist stereotypes discourages applications from older job seekers
  - Experimental treatments generate declines in applications from younger and older applicants (because treatment phrases reflect skill demands), but relative decline for older applicants is much larger – and the decline is significant only for older applicants

Direct vs. "indirect" age discrimination

- Can compare estimated effects on share of older job seekers "hired" (called back) from the discouragement of older applicants by job-ad language, to direct impact of age discrimination in hiring in correspondence study
- Share of applicants over 40 in control group is 20.00%
- This experiment: Discouragement of older applicants reduces % of over 40 among applicants from 20% to 16.31%
- Correspondence study: Age discrimination in hiring would reduce % over 40 among hires (callbacks) from 20% to 15.56%
- If these numbers roughly generalize to actual labor market, implies that by focusing only on hiring shortfalls we miss potentially half the problem by focusing on hiring relative to applicant pool
  - Critical implications for anti-discrimination enforcement

Do stereotypes in ads matter?

- Can older workers just avoid these, and move on to next job ad, at little/no cost?
- Not likely
  - Not that many ads – by city and occupation median/25th/10th percentiles are 123/39/12, and fewer outside administrative assistant jobs
  - Recall that ageist phrases were only at ~75th percentile of distribution, so many ads have similar phrases (or worse)

Interpretation: Other explanations?

- Natural interpretation is job searchers perceive ageist job-ad language as discriminatory and are deterred from applying
- Alternative: job ads signal job requirements older workers are less likely to fulfill or workers think firms believe this (or older workers prefer less)
  - Statistical discrimination by employers based on older workers' skills/qualifications
  - Worker sorting
    - Workers believe employers make assumptions about older workers' skills/qualifications, and workers know this
    - Or older workers have lower preferences for jobs signaled by treatment arms

Alternative to discrimination unlikely

- Employers who used physical and technology age stereotypes discriminated against older men in hiring in the original field experiment, and that experiment did a lot to rule out statistical discrimination
- Both older and younger respondents perceive ML job requirements as biased against older workers
- Changes in age composition occurred between ages 40 and 60, before ages at which age might more likely make one unqualified
  - Lifting 40 lbs.
  - Computer programs that are generally over 30 years old
  - Communications skills?

Other evidence generally doesn't support skill declines over these age ranges

- Some declines in "fluid" intelligence (memory, processing speed, etc.), but stable or perhaps increasing "crystallized" intelligence (knowledge acquired over time, vocabulary, etc.)
- Detailed studies of age and performance for a truck factory and a service firm find productivity increases until age 65 (when everyone retires), or productivity declines for routine and undemanding jobs but not for more complex jobs, and increases in the most challenging jobs (Börsch-Supan and Weiss, 2016; Börsch-Supan et al., 2021)
- Recent study by Quinby et al. (2023) studies relationship between profitability and age composition of workforce, instrumenting with age composition of commuting zone workforce (since declining firms will have older workers), and finds no clear relationship with profitability
- Research indicates there is no good reason to assume that productivity declines over the age ranges we study (National Academies of Sciences, Engineering, and Medicine: Committee on Understanding the Aging Workforce and Employment at Older Ages, 2022)
- Evidence further undermines statistical discrimination against older workers – at least based on correct stereotypes

Similar conclusion from resumes that skills/qualifications don't decline with age

- We assess changes in the skills/qualifications of applicants by taking control group resumes, and computing CS scores of the resumes (rather than the job ads) with the stereotypes
  - We use the 95th percentile of the CSS for each resume with each stereotype
- These are quite flat by age (as is previous occupational experience)

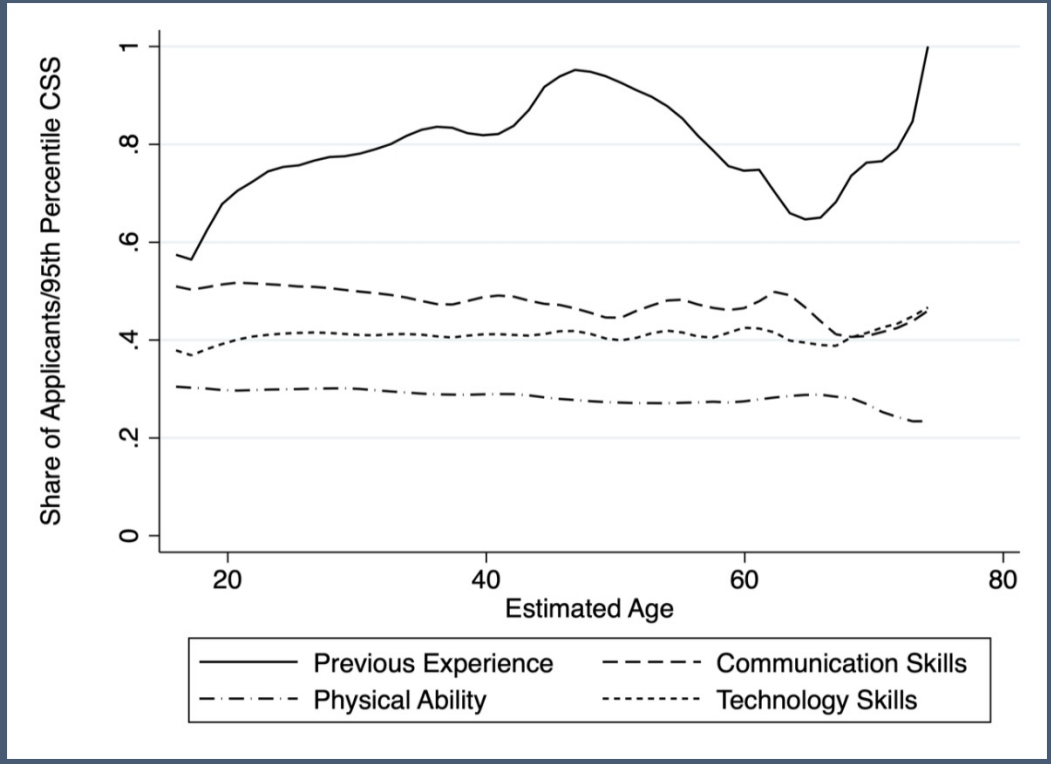

Results inconsistent with sorting

- Finally, we use the same computations to ask whether older workers "select out" of applying based on their skills or qualifications or preferences in response to job-ad language
- Steepening age gradients of resumes' CS scores with stereotypes would be evidence of older workers with resumes less consistent with stereotypes selecting out of the applicant pool in response to treatment
  - We find no evidence of this

Results inconsistent with sorting

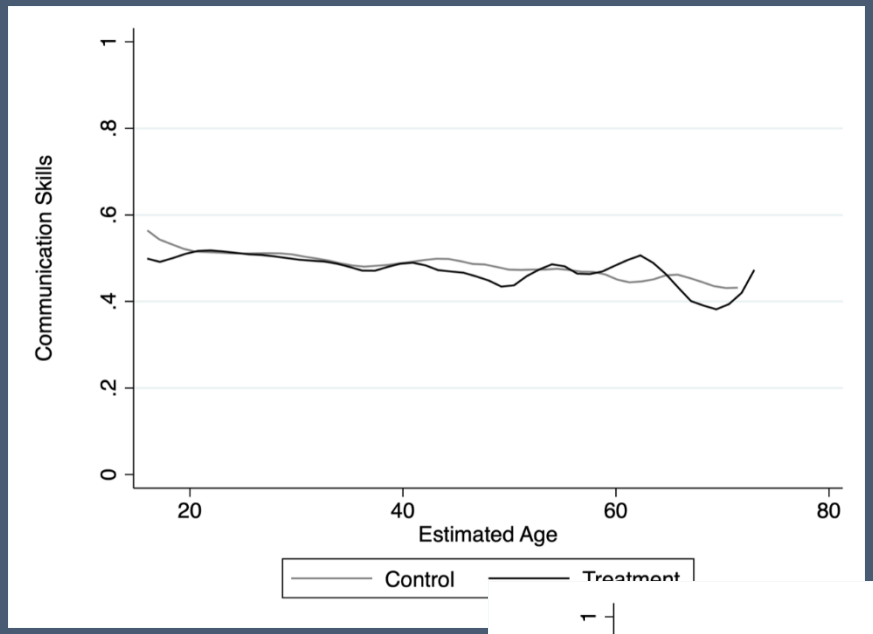

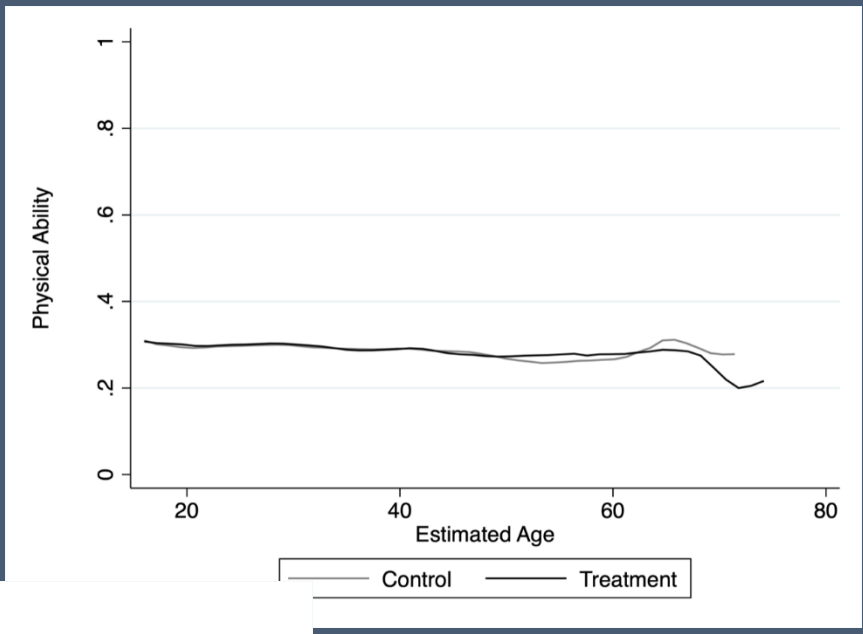

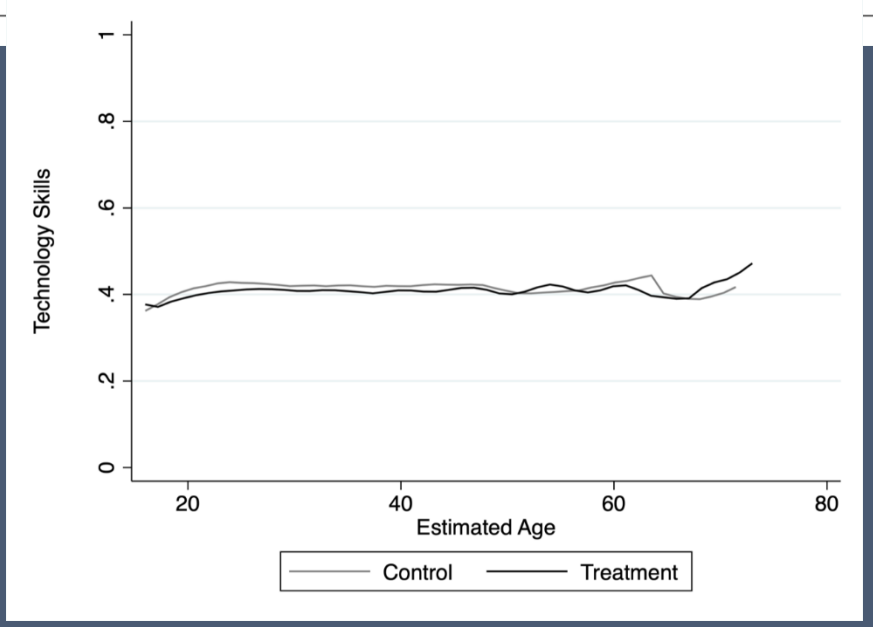

Results inconsistent with sorting

- Finally, we use the same calculation to ask whether older workers "select out" of applying based on their skills or qualifications or preferences in response to job-ad language
- Steepening age gradients of resumes' CS scores with stereotypes would be evidence of older workers with resumes less consistent with stereotypes selecting out of the applicant pool in response to treatment
  - We find no evidence of this
- And intercept doesn't shift down
  - Slight exception for technology skills, which is perhaps perceived as signaling a skill fewer workers have
- Instead, workers – older workers in particular – respond mainly to the ageist "cues" in job-ad language rather than the specific skill or other requirements expressed in the stereotyped treatments

Summary and conclusions

- First field experiment on how older job seekers respond to ageist language in job ads
- Use ageist language verified as linguistically related to specific age stereotypes, and perceived as age biased
- Find strong evidence that ageist language in job ads deters older job seekers
  - Even happens for single stereotyped phrase, in many cases
- Core results are for subtle ageist phrases (even stronger for blatant AARP phrases)
  - Ageist language does not have to be explicit to have pernicious effects on older workers and possibly facilitate age discrimination
- Intentional age discrimination is the most likely explanation of our findings

Potential policy implications

- Explicit language ("over 40 need not apply," "recent college graduate") is illegal (per CFR), and we really don't see it in job ads
- Our evidence implies much more subtle language can act as a form of age discrimination
  - EEOC can issue stronger guidance to employers to use age-stereotyped language in job ads
  - EEOC might investigate firms that use age-stereotyped language in their job ads
- "Shortfalls" of older workers in hiring relative to applicants likely understate effects of discrimination
  - Even possible that discrimination is occurring when there is no shortfall between the share of older workers hired and the share of older workers who apply for jobs
- Methods could be applied to discrimination against other groups, and implications for job-ad language might be similar

Thank you!