Risk, Norms, and Soft Skills: Non-Standard Drivers of Labor Market Outcomes

by

Jennifer Kades

Department of Economics
Duke University

Defense Date: March 25, 2025

Approved:

Robert Garlick, Supervisor

Peter Arcidiacono, Supervisor

Pat Bayer

Matt Masten

Pengpeng Xiao

Dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Economics in the Graduate School of Duke University
2025

ABSTRACT

Risk, Norms, and Soft Skills: Non-Standard Drivers of Labor Market Outcomes

by

Jennifer Kades

Department of Economics
Duke University

Defense Date: March 25, 2025

Approved:

Robert Garlick, Supervisor

Peter Arcidiacono, Supervisor

Pat Bayer

Matt Masten

Pengpeng Xiao

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Economics in the Graduate School of Duke University
2025

Abstract

How individuals decide whether and in what jobs to work, how firms decide whom to hire, and the aggregate wage and employment implications of these decisions are of central importance to labor economics and labor market policy. This dissertation explores the role of non-standard variables such as risk aversion, soft skills, and social norms in shaping these labor market outcomes.

Chapter 2 examines job transition decisions and their wage growth implications over workers' careers. Large transitions into higher-skill jobs can boost wage growth, but are rare. Whereas existing work largely focuses on the (mis)match between worker skills and job skill requirements to explain this pattern, this chapter presents evidence that, while skills matter, larger transitions are also riskier. Since workers are risk averse, this risk also contributes to their hesitance to make bigger career moves. These results carry important policy implications for the relative efficacy of skills-focused (e.g., training) compared to risk-focused (e.g., job search assistance) active labor market policies.

Chapter 3 provides experimental evidence on how using information about job applicants' soft skills in firms' hiring decisions affects both firm and workseeker outcomes. Partnering with the largest recruitment agency in South Africa, the experiment randomizes the criteria used to shortlist job applicants for job listings at partner firms. Results suggest that including measures of soft skills in candidate ranking leads to more efficient hiring and better firm-worker matches.

Chapter 4 explores the relationship between social norms around work and labor market outcomes such as labor supply and wages. The analysis employs a novel measure of work norms in the U.S. over roughly the past decade, constructed using large-scale social media text data and machine learning methods. Results suggest that there is a meaningful relationship between social norms and labor supply, particularly for men.

Contents

List of Tables

List of Figures

Acknowledgements

As Thomas Wolfe put it, "each of us is all of the sums he has not counted." Below is my chaotic and entirely non-exhaustive attempt at an accounting—thank you to everyone in it, and everyone I am surely forgetting.

To my committee, Peter Arcidiacono, Pat Bayer, Matt Masten, Pengpeng Xiao, and most especially my advisor, Rob Garlick: you were all extremely kind and helpful, much more so than anyone should reasonably expect from economists. Thank you also to Erica Field, for your generous support in my sixth year.

To the countless other Duke faculty and staff who have helped me over the past six years, especially Christyn Klinck and the rest of the Economics Department team, David McDade in Grants & Contracts, and Tom Milledge the Duke Compute Cluster wizard: as someone whose affairs are perpetually in disorder, I truly could not have done this without you all.

To my co-authors, Lukas Hensel, Kate Orkin, and Rob again: I've learned so much from you and so enjoyed working with you. Thank you for the opportunity to contribute to an incredibly meaningful research portfolio. Thank you also to all of the other wonderful people I've worked with on these projects across J-PAL Africa and our research partners.

To J-PAL, the National Science Foundation, and the Weiss Fund for Research in Development Economics: thank you for your generous support of various components of this research.

To the outstanding economics faculty at Davidson College, especially Mark Foley, Shyam Gouri Suresh, and Vikram Kumar: your passion and thoughtful pedagogy made a famously "dismal" subject feel fun and enlightening. I am certain you're all to thank (or blame) for the strong Davidson-to-Duke-Econ-PhD pipeline.

To the indie & pop girls I formed parasocial relationships with, especially Samia the superlative sad girl, Sabrina Carpenter (who inspired the name of one these chapters), and Sidney Gish, who somehow makes debilitating levels of neuroticism funny instead of sad (or at least funny in addition to sad): thank you for allowing me to briefly disassociate from

the bleak reality of PhD life now and then.

Last but certainly not least, to all the friends and family who supported me in myriad ways throughout this odyssey, especially my parents, whose love and encouragement have been consistently unflagging and unconditional, my fellow masochists Alison, Chelsea, Erin, Nikita, and Rosie, my *other* fellow masochists the Bull City Beer Runners, and Beth, Elissa, Heather, Karen, Moe, and Olivia: each of you, at one point or another, kept me going at a moment I didn't think I could.

1. Introduction

This dissertation explores the role of non-standard variables such as risk aversion, soft skills, and social norms in shaping labor market outcomes. The first essay studies job transitions and resulting wage growth trajectories, exploring whether the rarity of big career jumps is driven by worker risk aversion, in addition to more traditional skills-focused explanations. The second essay studies hiring criteria, presenting evidence that including measures of applicants' soft skills can improve outcomes. The final essay studies social norms around work, proposing a novel measurement approach and demonstrating a relationship between work norms and labor supply.

This work is not an indictment of existing labor market models that exclude these variables. Simplifying assumptions are not only necessary but often quite helpful for any model of the world, crystallizing the most important components and mechanisms of the question at hand. Rather, this dissertation builds on and contributes to these existing models, shedding light on when and how we should expect the non-standard variables in question to matter.

Chapter 2 examines job transitions and wage growth for young workers using detailed panel data on work histories, heterogeneous risk preferences, and a task-based measure of skill distance. It provides new evidence on how "steep" transitions—moves to jobs for which workers are relatively underskilled—affect short- and long-term wage outcomes. These transitions yield immediate wage growth premiums, contrary to the expectation that skill mismatch penalizes wage growth, suggesting productivity benefits that overshadow mismatch costs. However, it appears that workers struggle to target strictly upward moves (moves to higher-skill jobs), prompting a proposed framework in which larger transitions are riskier due to information frictions. Supporting this, risk-tolerant workers are observed to make more frequent and larger transitions, though many of these transitions are downward. A model grounded in these findings quantifies the joint implications of risk aversion and skill alignment concerns in job mobility, showing that the degree to which workers' hesitance to make larger occupational transitions is due to risk preferences versus skill deficiency has

important policy implications for the efficacy of training and re-skilling programs.

Chapter 3 presents the results of a field experiment testing the use of "soft skills" in the hiring process.[1] Firms often struggle to find suitable workers despite large pools of applicants. In the absence of other information, they rely on traditional signals—education and experience—or referrals, potentially excluding suitable candidates without signals of their skills. This chapter provides experimental evidence on how broadening the information in firms' hiring decisions affects both firm and workseeker outcomes. Partnering with the largest recruitment agency in South Africa, the experiment randomizes the criteria used to shortlist job applicants for job listings at partner firms. In vacancies in the control group, firms receive shortlists of candidates ranked on education and experience. In vacancies in the treatment group, shortlists rank candidates on these criteria and, in addition, measures of soft skills correlated with labor market outcomes. Broadening firms' information set alters which candidates are hired: in treated vacancies, hired applicants have on average higher soft skills and less work experience than in control vacancies. It also boosts hiring: in treated vacancies, firms issue more interview invitations and make more actual hires. Individuals who would have been shortlisted for the job posting if it had been assigned to a different treatment are not negatively affected, other than by not getting the one job in the posting. Findings are in line with models where higher match quality raises labor demand.

Chapter 4 explores how work norms may shape the labor market. Media narratives and existing research suggest such norms are important drivers of worker behavior and, ultimately, macroeconomic outcomes like labor supply and wages. However, this idea has not been tested empirically due to the difficulty of measuring norms. This chapter develops a novel measure of work norms using a large set of unstructured text data from social media, combined with a custom-built machine learning model to label the text. The measure construction approach produces time-varying, geographically granular measures of work norms in the U.S. over roughly the past decade, and could be easily generalized to study other norms. The resulting measure is used to examine the relationship between work norms

[1] Chapter 3 reflects collaborative work with Robert Garlick, Lukas Hensel, and Kate Orkin.

and labor market outcomes such as labor force participation, employment, and wages. There appears to be a meaningful relationship between norms and attitudes toward work and labor supply, suggesting that any study of labor market outcomes and policies should carefully consider the social context in which they are operating.

2. Short n' Steep: Wage Growth Maximizing Job Ladders with Heterogeneous Skills and Risk Preferences

2.1 Introduction

It is well established that job transition patterns play an important role in determining wage growth over the life cycle, often jointly and in complex interplay with human capital—a relationship that search and matching models as far back as Ben-Porath (1967) and Burdett (1978) have been trying to disentangle. At the same time, job transitions are typically characterized by substantial uncertainty, and hence risk. Therefore, variation in the riskiness of transitions, combined with heterogeneity in worker risk preferences, is a potentially meaningful driver of career trajectories and wage dispersion. This paper uses a longitudinal dataset of detailed labor market histories to document novel empirical patterns consistent with this idea, leveraging the intuitive proposition that making bigger occupational jumps (in terms of differences in skill requirements between the source and destination job) is riskier. I then propose a model of worker mobility decisions incorporating the worker's current stock of skills and heterogeneous risk appetite that can explain patterns in the data that are difficult to square with existing models of skill-based matching.

Worker human capital—or skills—impacts hiring and wages. This is clear from direct observation of job advertisements and surveys of employers, as well as a large body of work exploring the returns to increasingly specific formulations of human capital, starting with general ability or education level to industry-, occupation, or firm-specific human capital to granular skill bundles based on detailed job tasks.[1] Thus, while early matching models assumed workers were identical in terms of human capital (e.g., Burdett (1978) and Mortensen and Pissarides (1994)), this restriction was quickly relaxed (e.g., J. J. Heckman

[1] Worker human capital impacts hiring and wages as found in, e.g., Brunello and Wruuck (2021), Cappelli (2015), D. Deming and Kahn (2018), Hershbein et al. (2018), and Ziegler et al. (2020). A small sample of the work exploring reruns to human capital (by level of specificity) includes: J. J. Heckman et al. (2006), Ingram and Neumann (2006), and Lindqvist and Vestman (2011) (general ability or education level); Hashimoto (1981), Kambourov and Manovskii (2009), Lazear (2009), Neal (1995), and Sullivan (2010) (industry-, occupation-, or firm-specific human capital); Autor and Handel (2013) and Gibbons and Waldman (2004) (granular skills based on detailed job tasks).

and Sedlacek (1985) and Keane and Wolpin (1997)). Canonical models incorporating human capital generally posit that workers sacrifice wage growth early in their career in order to attain more productive matches later in their career, either through investment in human capital accumulation through training or on-the-job learning or through an "exploration" or "job shopping" period that enables them to find the jobs that are the best match for their (initially uncertain) skills.[2] These models align with several stylized empirical facts, like convex wage growth, high rates of occupational switching early in one's career, and wage penalties for industry or occupation switching (Eckardt, 2019; Goldsmith and Veum, 2002; Rubinstein and Weiss, 2006).

This paper examines job transitions by "skill distance" and "direction," revealing novel empirical results that both build upon and add nuance to this existing work. Using work history panel data from the National Longitudinal Survey of Youth ("NLSY") and detailed occupational skill requirements from the Occupational Information Network ("O*NET"), the skill distance of a job transition describes the absolute difference between the worker's current skill bundle at the time of transition and the skill requirements of their destination job (Guvenen et al. (2020), Lise and Postel-Vinay (2020), and Blair et al. (2021) use similar skill distance measures). To capture the direction of the transition—i.e., whether the move is up or down the skill ladder—I develop additional distance measures that identify the extent to which the worker is under- or overskilled for their new job.

Considering job transitions through the lens of these skill distance and direction measures yields five key results. First, as these existing models would predict, skill mismatch appears to matter for productivity: workers moving into jobs they are initially underskilled for face wage *level* penalties relative to occupation means and are more likely to receive firm-provided training compared to well-matched or overskilled peers. Second, however, steep upward transitions yield immediate and sustained wage *growth* premiums, suggesting that job quality overshadows skill mismatch in determining wage growth. This is somewhat

[2] See, e.g., Bagger et al. (2014), Barron et al. (1989), and Engbom (2022) for explorations of the trade-offs between short-term wage growth and human capital accumulation, and Antonovics and Golan (2012), Light (2005), and Neal (1995) for "exploration" or "job shopping" periods.

in contrast to the results discussed above, many of which would imply that, from a wage growth perspective, the best-case scenario is to start your career in an occupation you are already well-matched with and stay in it, or else make only small upward transitions to jobs with similar skill profiles, thus minimizing skill mismatch penalties. Third, these large upward transitions are rare, despite being extremely beneficial. While this is likely to reflect selection to some extent, the main results are robust to the inclusion of a large set of demographic and transition context controls, and remain significant after including worker fixed effects. Fourth, many workers make transitions down the skill ladder, which do not appear to be driven by compensating differentials, as upward moves correlate with higher job satisfaction levels and growth after controlling for wages. This suggests that larger job transitions may be riskier due to information frictions that prevent workers from perfectly targeting higher-skill jobs.[3] Consistent with this idea, risk aversion significantly predicts the frequency and size of transitions, and larger moves carry more risk in terms of realized job outcomes.

Motivated by these empirical patterns, I introduce a model of worker transition decisions that includes heterogeneous risk preferences alongside heterogeneous skills. This model accounts for the observed wage growth premium associated with steep upward transitions and the rarity of such moves by emphasizing how information frictions and workers' risk tolerance may influence their job mobility strategies. Expanding on existing models that focus solely on skill alignment, this approach provides a more comprehensive explanation of job transition patterns, revealing the potentially important role that risk aversion plays in shaping career trajectories.

In particular, I propose a model of worker mobility decisions in which workers are heterogeneous in terms of risk aversion, skills, and education level, and larger occupational transitions are riskier due to a lack of information about skill level and quality for jobs that

[3] It is well-documented that various information frictions exist in the labor market and significantly impact job search behavior and outcomes. For recent examples documenting the presence and implications of such frictions see, for example: Abebe et al. (2023), Banerjee and Sequeira (2023), Conlon et al. (2018), Mueller and Spinnewijn (2023), and C. H. Wheeler (2008).

workers are less familiar with. In the model, workers choose an absolute skill distance range in which to search for their next job, but are not able to precisely target upward moves. Intuitively, they decide whether to look for a job very similar to their current job, very different, or somewhere in between. Given a choice of search strategy, they draw a new job from a known distribution centered around their current job's skill level, with higher-skill-distance search strategies implying higher variance. The model is able to match observed transition patterns and outcomes in the data, and estimated parameters are consistent with the idea that skill match may be less important than worker characteristics and job quality in determining wages. Counterfactual exercises exploring the impact of two different active labor market policies (skills training and job search assistance) show that considering the role of risk tolerance in transition decisions has important implications for expected policy outcomes. In particular, to the extent that workers avoid larger career moves due to risk aversion rather than skill misalignment, skills training programs may be less effective than anticipated.

In sum, this paper provides the first empirical examination of the joint relationships between task-based skills, risk preferences, job transitions, and wage growth. I show that while existing models capture some elements of these dynamics, it is difficult to fully explain the nuanced interplay observed in real-world data without considering risk preferences. By incorporating risk aversion and skill mismatch into a new model of worker transition decisions, I demonstrate how these factors may jointly influence career trajectories and wage outcomes. This contribution is crucial for understanding how workers navigate complex labor markets and for designing policies that support effective job transitions and sustained wage growth.

This project contributes to a vast corpus of empirical and theoretical work on patterns of occupational mobility and wage growth over the life cycle. Exploring the wage implications of job transitions along the vectors of both *skill distance* and *skill direction*, as well as relating transition types to risk preferences, provides a new layer of insight to empirical work in this area, building on research such as Parent (2000), Light (2005), C. H. Wheeler (2008),

Munasinghe et al. (2008), Barlevy (2001), L. B. Kahn (2010), and Hoffmann et al. (2016) that uses U.S. survey data to carefully catalog patterns in the nature and consequences of job transitions in various contexts. From a theoretical standpoint, the model feature of workers adopting particular search or job mobility strategies in response to uncertainty, which then have downstream implications for future mobility, human capital accumulation, and wage growth, is in the spirit of work such as Neal (1999), Antonovics and Golan (2012), Pallais (2014), or more recently D. J. Deming and Noray (2020), Barbulescu and Bonet (2024), and Macaluso and Borovičková (2024). In fact, the proposed model maps quite well onto Neal's two-stage model of early-career transitions, though with significant additional complexity due to the inclusion of job skills and heterogeneous worker risk aversion. Additionally, workers in this model choose a job search strategy based on the degree of dissimilarity with their current occupation—a novel but intuitive targeting mechanism.

As touched upon above, this paper also adds to the growing body of work concerned with a task-based approach to defining jobs, human capital, and the degree of (mis)match between them (e.g., Macaluso and Borovičková (2024)). In particular, Guvenen et al. (2020) and Lise and Postel-Vinay (2020) use similar concepts of skill distance, though are more concentrated on carefully accounting for the dynamics of skill accumulation in various dimensions and the implication of these distinct skill paths for skill mismatch and resulting productivity, which is not the main focus of this paper. Both of their models imply meaningful and persistent wage penalties for skill mismatch. This is consistent with some of the empirical patterns I observe but contrasts with the lack of wage growth penalty (and in some cases wage growth premium) I find for higher levels of underqualification. This paper is also closely related to Blair et al. (2021), which explores the impact of educational credentials on a worker's ability to make larger transitions. Blair et al. conclude that having a bachelor's degree makes it easier to move into jobs one is more underqualified for, a result I confirm here. While they use a cross-sectional sample of job transitions, I build upon their work by exploiting longitudinal data to capture longer-term transition patterns and outcomes as well as by adding risk preferences as a source of variation in the kinds of

transitions workers make.

Finally, I provide new empirical evidence on the relationship between risk preferences, the types of transitions workers make, and wage growth, which motivates the inclusion of heterogeneous risk preferences in a model of worker mobility decisions. This presents, to my knowledge, the first such evidence on the relation between risk aversion and the kinds of transitions workers make (in terms of similarity to one's current job and wage changes). The closest existing empirical work is van Huizen and Alessie (2019), which explores the impact of risk tolerance on binary worker mobility. The authors confirm that there is almost no other research in this area, pointing to Argaw et al. (2017), which uses the German Socio-Economic Panel Survey to investigate the relation between risk attitudes and job mobility, as the only other known example. Huizen and Alessi also offer an excellent review of adjacent research on the role of risk aversion in labor market outcomes, including in the context of education and occupational choices, migration decisions, entrepreneurship, the behavior of unemployed work seekers, and wage growth, which I will not reproduce here other than citing perhaps the most closely related of this subset—Dillon (2018)—which documents risk premiums in earnings for occupations with higher employment and wage risk.

In terms of theory, risk aversion tends to be assumed constant across all workers, or else absent altogether, presumably for analytical tractability. Exceptions include an early model of job search due to Pissarides that incorporates constant risk aversion, which generates a wage premium for occupations with higher wage variance (Pissarides, 1974). In a similar vein, Cozzi (2012) uses a model incorporating heterogeneous risk tolerance and income risk in an attempt to match observed wealth inequality in U.S. data. Perhaps the most relevant theoretical work is Chade and Lindenlaub (2022), which presents a model of job matching with risk-averse agents and risky up-front skill investment and uses it to decompose rising U.S. wage inequality into components due to risk, heterogeneity, and technology. However, risk is still assumed to be constant across individuals in this model, and the environment and mechanisms are quite different from the framework presented here, in which risk preferences

influence the *size* of the transitions that workers target.

Section 2 presents the data and measurement approaches used in the empirical analysis. Section 3 presents the main empirical findings. Section 4 presents the model and estimation results, with counterfactual exercises. Section 5 concludes.

2.2 Data and Measurement

Sections 2.1 through 2.3 describe the data and methodology underlying the measures of risk tolerance and skill distance used throughout the paper. Sections 2.4 and 2.5 present an overview of the skill distance and dynamics of job transitions in the sample, showing that large occupational transitions are rare and workers do not tend to return to previously-held occupations after switching.

2.2.1 Data

Data on job transitions is constructed using labor market histories for a sample of roughly 9,000 individuals from the 1997 cohort of the National Longitudinal Survey of Youth (NLSY). This is the second cohort of the well-known NLSY survey, which tracks a nationally representative sample of U.S. workers and is described in detail in many earlier papers, e.g., Light (2005). Respondents range from ages 14 to 18 years old at the start of the survey, and are surveyed at one- or two-year intervals from that point forward, providing detailed weekly employment and wage histories. The NLSY also contains a rich set of variables on worker demographics, educational background, personality, and attitudes, including direct measures of worker general ability (as measured by ASVAB scores) and risk aversion.

I use the 1997 cohort for the primary analysis in order to avoid data quality issues identified in early rounds of the 1979 survey, as well as to simplify occupational classification.[4] However, I confirm that patterns are robust to the inclusion of the 1979 cohort, and in fact that patterns for this cohort are strikingly similar to those of the younger cohort.

[4] Training data in the NLSY79 are incomplete prior to 1988, as discussed in e.g. Kaymak (2014). Additionally, occupations are classified using 1970 Census codes before 2000, but switch to 1990 Census coding thereafter, complicating the mapping of occupations to O*NET skills.

The main unit of analysis throughout is job transitions, where a transition is identified as a change in the worker's primary employer.[5] I apply standard restrictions to the sample of transitions by (i) limiting the set of job spells to those that start after the respondent earned their highest degree, (ii) limiting the sample to jobs averaging at least 20 hours per week and lasting at least four weeks, and (iii) excluding military jobs and self-employment (following, e.g., Farber and Gibbons (1996)). Table 1 presents summary statistics for the resulting sample of job transitions. As can be seen from the table, job transitions are a relatively rare occurrence—the median worker has made only four job transitions as of the last survey round. Additionally, the majority of transitions are among workers with less than a college degree, who also make up a majority of the worker sample. About a quarter of transitions are involuntary, and while transitions result in small gains in wages and job satisfaction on average, variance in these outcomes is substantial. In other words, transitions appear to be risky in terms of realized changes in wage and job satisfaction.

Data from the Occupational Information Network ("O*NET") furnishes skill requirements for each three-digit occupation code in the sample. O*NET is a database maintained by the US Department of Labor that offers detailed information on the skills, activities, and educational requirements of approximately 1,000 occupational categories. The skill requirements of each occupation are captured by skill requirement ratings (from one to five) in 30 skill components, as assessed by expert occupational analysts (Tsacoumis & Willison, 2010).[6]

2.2.2 Risk Attitudes in the NLSY

The measure of heterogeneous risk preferences used throughout the analysis is taken directly from the NLSY. Specifically, all NLSY97 respondents were asked to assess their level of risk tolerance on a 0-10 scale with the following question:

[5] Because, in the NLSY, job spell information is gathered by asking about the employer, not the job, employer and job transitions are not generally distinguishable.

[6] Skill ratings are merged into the NLSY work history data using a crosswalk between Census and SOC codes, as found on the Census website: https://www.census.gov/topics/employment/industry-occupation/guidance/code-lists.html.

Table 2.1: Summary Statistics for Job Transitions

Variable	Category	Min	Mean	Std Dev	Median	Max
Ability (AFQT)		0.00	32.54	29.53	25.90	100.00
Age at Transition		14.00	26.36	5.74	25.00	42.00
Education Level	Less than HS		10.60			
	HS graduate		69.73			
	Some college		4.71			
	College graduate		11.97			
	More than college		2.99			
Race or Ethnicity	Black		28.94			
	Hispanic		20.75			
	Other		0.90			
	White (Non-Hispanic)		49.41			
Risk Tolerance		0.00	5.78	2.67	5.00	10.00
Male			0.55			
Number of Transitions per R[1]		1.00	5.02	3.83	4.0	30.00
Total Work Experience (wks)[2]		5	357.59	255.29	296.00	1118
Transition Year Category	<2000		3.08			
	[2000, 2005)		27.95			
	[2005, 2010)		31.55			
	[2010, 2015)		19.02			
	>2015		18.39			
Spell Length (wks)[2]		5	70.45	96.13	35.00	1118
Involuntary Departure			0.28			
Unemp. or OOLF Between Jobs			0.53			
Length Non-work Between Jobs (wks)		0	16.91	43.87	0.00	211
Change in Real Hourly Pay		-7.05	0.48	3.11	0.26	10.44
Change in Job Satisfaction		-4.00	0.43	1.62	0.00	4.00
Transition to Same Occupation			0.17			
Skill Distance of Transition		0.00	2.94	1.31	2.84	8.58
N=33,848						

[1] Cumulative total number of job transitions, for workers with at least one transition observed in the sample.

[2] Total work experience and spell length reflect the accumulated work experience at the time of transition and the spell length in the job preceding the transition, respectively.

> *Are you generally a person who is fully prepared to take risks or do you try to avoid taking risks? Rate yourself from 0 to 10, where 0 means "unwilling to take any risks" and 10 means "fully prepared to take risks."*

Each respondent answered the question at most once, either in survey year 2010 or 2011. Risk tolerance is roughly normally distributed across the sample, with the distribution centered around five and with a standard deviation of about 2.6, as shown in Table 2.1. In other words, there is substantial variation in risk tolerance among respondents.

Although this risk attitudes scale is self-assessed and not incentivized, existing research demonstrates that the risk measures in the NLSY appear to accurately reflect individu-

als' true preferences, in the sense that they are predictive of risky behaviors such as entrepreneurship and divorce (Chekmasova (2020) and Light and Ahn (2010)). More broadly, research has shown that survey-elicited risk measures align well with incentivized lottery choices and predict real-world behaviors (Dohmen et al. (2011) and Falk et al. (2015)).

2.2.3 Measuring Skill Distance and Direction

Similar to approaches in Guvenen et al. (2020) and Lise and Postel-Vinay (2020), and Blair et al. (2021), I use O*NET data to define the skill requirements of each job as identified by its three-digit occupation code. Workers' skill bundles are then dynamically constructed based on their full labor market histories up to that point, such that at any given time t, a worker's current skill level for any given skill component k can be calculated as follows:

$$S_{k,t} = \begin{cases} S_{k,T} + (S_{j'} - S_{k,T}) \left(1 - e^{-r_u \cdot (t-T)}\right), & \text{if } S_{j'} > S_{k,T}, \\ S_{k,T} + (S_{j'} - S_{k,T}) \left(1 - e^{-r_d \cdot (t-T)}\right), & \text{if } S_{j'} < S_{k,T}, \\ S_{k,T}, & \text{if } S_{j'} = S_{k,T} \end{cases}$$

where $S_{i,T}$ denotes the worker's skill level in skill i at the time T that they departed from their prior job, $S_{i,j'}$ denotes the required skill level in skill i for their current job j', identified from O*NET ratings, and r_u and r_d represent the rate of skill updating and skill depreciation, respectively.[7] The main empirical results are robust to alternative appreciation and depreciation rates, and in particular are robust to the extreme case of zero skill acquisition before the most recent job, as shown in Appendix A.2.

Intuitively, this equation captures the gradual adjustment of the worker's skills to those required by the job. Skills in which the worker is initially underskilled will adjust upward until they meet the new job's requirement, while skills in which the job's requirement is lower than the worker's starting level will depreciate over the course of the worker's tenure

[7] As estimated in. e.g., Lise and Postel-Vinay (2020) and Dinerstein et al. (2022), updating and depreciation (r_u and r_d) are assumed to be asymmetrical, with skill acquisition happening much faster than depreciation. In particular, for the empirical analyses, the skill updating rate is set to $r_u = 0.11552$, which reflects the idea that skills update quickly to meet new job requirements. The skill depreciation rate is set to $r_d = 0.00444$, consistent with depreciation rate of ∼5% as estimated in Dinerstein et al. (2022).

in the destination job.

With worker skill bundles and occupational skill requirements thus specified, I define the *absolute skill distance* of an occupational transition for worker i from job j to j' as the Euclidean distance between the worker's vector of current skills at the time they leave their source job and the vector of skill requirements in their destination job, as follows:

$$\text{Absolute Skill Distance} = \sqrt{\sum_{k=1}^{N} \left(S_{k,j'} - S_{k,i,T}\right)^2}$$

Where $S_{k,i,T}$ captures the worker's level of skill k at time T of transition, and the sum is taken over the 35 job skills identified in the O*NET database.

To capture the *direction* of transitions—i.e., whether they reflect moves up or down the skill ladder, in aggregate—I further decompose skill distance into components in which the worker is underskilled or overskilled. In particular, define underskill distance as:

$$\text{Underskill} = \sqrt{\sum_{k=1}^{N} \mathbb{1}[S_{k,j'} > S_{k,i,T}] \left(S_{k,j'} - S_{k,i,T}\right)^2}$$

as the sum of skill distances for *only the skill components in which the worker is underskilled*. Overskill distance is defined analogously.

The difference between these identifies the worker's relative skill deficit, in terms of their current skill bundle, for the job they are transitioning into:

$$\text{Relative Skill Distance} = \text{Underskill} - \text{Overskill}$$

where positive values indicate that the worker is relatively underskilled for the destination job, while negative values indicate they are relatively overskilled. Broadly speaking, relative skill distance captures the extent to which the worker is moving up the skill ladder, in aggregate.

To build intuition, Table 2.2 provides illustrative examples of job transitions and their associated distances. Assume for these examples that each worker's skills have completely adjusted to their source job's skill requirements, so the worker's skill bundle is identical to the source job's required skill bundle.

Transition Example 1: Small absolute and relative distance, upward move.
Ben resigned in disgrace from his job as a public auditor due to a minor bribery scandal, but nevertheless managed to leverage his accounting skills to land a role as the financial manager of a charitable trust for a large candy company. While both jobs have a similar skill portfolio, the financial manager job requires slightly higher levels of many skills, especially management-related ones such as Complex Problem Solving and Monitoring, leading to an underskill score of about 1.32. At the same time, there are a few, more execution-oriented, skills for which the auditor role has higher requirements (e.g., Math and Writing), leading to an overskill score of about 0.86. Combining these results in a moderately positive relative skill distance of about 0.47 and a small absolute skill distance (less than the 25th percentile of the absolute skill distance distribution over all transitions).

Transition Example 2: Larger absolute and relative distance, upward move.
Andy was working as a shoeshiner in City Hall while awaiting his rock band's "big break," but networked his way into a job as a personal assistant to a bureaucratic employee, helping her launch a campaign for public office. These jobs require very different skill sets, and, in terms of direction, being a personal assistant requires higher levels of most skills. This results in a moderately large absolute skill distance (about median), the majority of which is underskill, as well as a large and positive (about 75th percentile) relative skill distance.

Transition Example 3: Large absolute distance with smaller relative distance, downward move. Ann, a nurse, had grown tired of acquaintances and even strangers soliciting her medical advice, often with upsetting pictures. So, when the opportunity arose to serve as the public relations manager for her city's public health department, she happily accepted.

These jobs are quite different and require substantially different skill sets, but it isn't immediately clear which job is higher-skill. This is reflected in the skill distance measures—the absolute skill difference between the two occupations is large (above median), but this is driven by both underskill (e.g., PR management requires higher levels of sales-type skills such as Negotiation than nursing) and overskill (e.g., nursing requires higher skills in Science than PR management). The result is a negative, though small, relative skill distance—nursing has a slightly higher skill requirement than PR management, in aggregate.

Table 2.2: Illustrative Job Transition Examples

Name	Source Occupation	Destination Occupation	Absolute Skill Dist.	Underskill	Overskill	Relative Skill Dist.
(1) Ben	Auditor *(800)*	Financial Manager *(120)*	1.577	1.324	0.858	0.466
(2) Andy	Shoeshiner[1] *(4520)*	Personal Assistant[2] *(5700)*	2.498	2.393	0.715	1.678
(3) Ann	Nurse *(3130)*	PR Manager *(60)*	2.917	1.679	2.385	-0.707

2002 Census codes corresponding to each occupation in parentheses.

[1] Shoeshiner is mapped to Census code 4520: Miscellaneous personal appearance workers.

[2] Personal Assistant is mapped to Census code 5700: Secretaries and administrative assistants.

2.2.4 Distribution of Skill Distance

Figure 2.1 shows the sample distribution of each skill distance measure for (a) the distance between the worker's skill bundle at the time of transition and the skill requirements of their destination job, and (b) the distance between the source job and destination job skill requirements. Because skills are assumed to update faster than they depreciate, workers tend to be more overskilled than underskilled when considering their full labor market histories, though this is not by construction (e.g., this would not be the case if workers moved into progressively higher-skill jobs throughout their careers on average—see Appendix A.1 for simulated distributions with random transitions for comparison). On the other hand, when comparing the skill requirements of source and destination jobs, underskill and overskill are roughly symmetric, resulting in a distribution of relative skill distance that is roughly centered at zero. In other words, *workers are about as likely to move down the skill ladder as up the skill ladder on any given transition*, relative to their source job, suggesting that workers may struggle to target strictly higher-skill jobs (and this remains true when restricting the sample to only voluntary transitions—see Appendix A.2 for details).

The distributions are generally right-skewed, reflecting the fact that larger transitions are relatively rare. In particular, when using the worker's skill bundle based on their full job history, large underskill is quite rare, and in fact the modal transition has zero underskill distance. This results in an asymmetric distribution of relative skill distance, in which the majority of transitions have negative relative distance (workers have a higher aggregate skill level than the requirements of the jobs to which they are transitioning). Additionally, when comparing source and destination job skills only, we see that the modal transition is zero-distance; that is, the modal transition is to a job in the same occupation with a different employer.

The empirical analyses shown throughout the remainder of the paper use skill distance based on the worker's skill bundle, incorporating their full work history as laid out in section 2.3. Since skills are assumed to appreciate faster than they depreciate, this measure

(a) Distances between worker skills and destination job skills

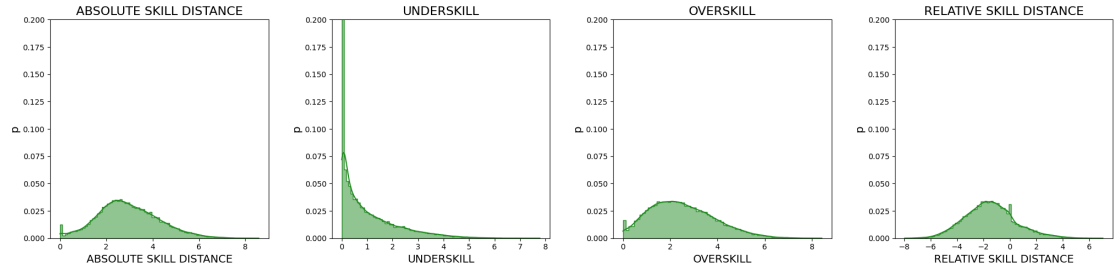

(b) Distances between source and destination job skills

FIGURE 2.1: Skill Distances of Job Transitions

Note: Reflects the distribution of each skill distance measure across all job transitions observed in the sample. The first row reflects distances between the skill requirements of the worker's new job and their current skill bundle at the time of the transition, including (potentially depreciated) skills from prior jobs. The second row reflects distances between the skill requirements of the worker's new job and the skill requirements of their old (source) job.

provides a rough upper bound on each worker's skills, while the distance between source and destination jobs can be thought of as a lower bound (as it assumes workers possess only the skills acquired in their most recent job). Therefore, estimates of the effect of underskill using the first measure are likely somewhat conservative. In practice, however, results are quite similar when using skill distances between the source and destination jobs (see Appendix C).

2.2.5 Transition Dynamics

Figure 2.2 documents the dynamics of transitions as workers progress in their careers. Figures reflect averages across all workers' nth job transitions, for up to the first ten transitions of each worker's career. Thus, each point on the x-axis reflects the sample average *conditional on having made that many transitions.* The top left figure shows that workers are less likely to change occupations as they progress in their careers, with the probability of occupation switching dropping from 84% on a worker's first transition to 76% on a worker's tenth transition. The top right figure shows that the absolute skill distance between a worker's destination job and their nth most recent job at any given transition is increasing in n—in other words, workers tend to move monotonically further away from their older jobs in terms of skills, on average. The bottom two figures show that it is relatively rare for workers to return to prior occupations after making a switch. In particular, the probability of returning to the $n-1$th, conditional on switching occupations at the nth transition, is only about 10%, regardless of transition number. Additionally, while the probability of moving back to *any* previously-held occupation is monotonically increasing (since a higher transition number means a larger set of occupations to potentially move back to), it never exceeds 35%, and is only about 20% for the fourth transition (the maximum for the median worker in the sample).

2.3 Empirical Patterns

The following sections describe the empirical patterns that motivate the model of occupational transition decisions in Section 4.

Sections 3.1 and 3.2 present an empirical puzzle: as shown in 3.1, skill mismatch appears to matter—large underskill distances are rare, and being underskilled predicts employer training and wage *level* penalties compared to workers in the same occupation—but as shown in 3.2, greater skill mismatch is associated with better short- and long-term transition outcomes *as long as the direction of the move is positive (moving up the skill ladder).* In other words, in addition to yielding better long-term outcomes, being underskilled actually

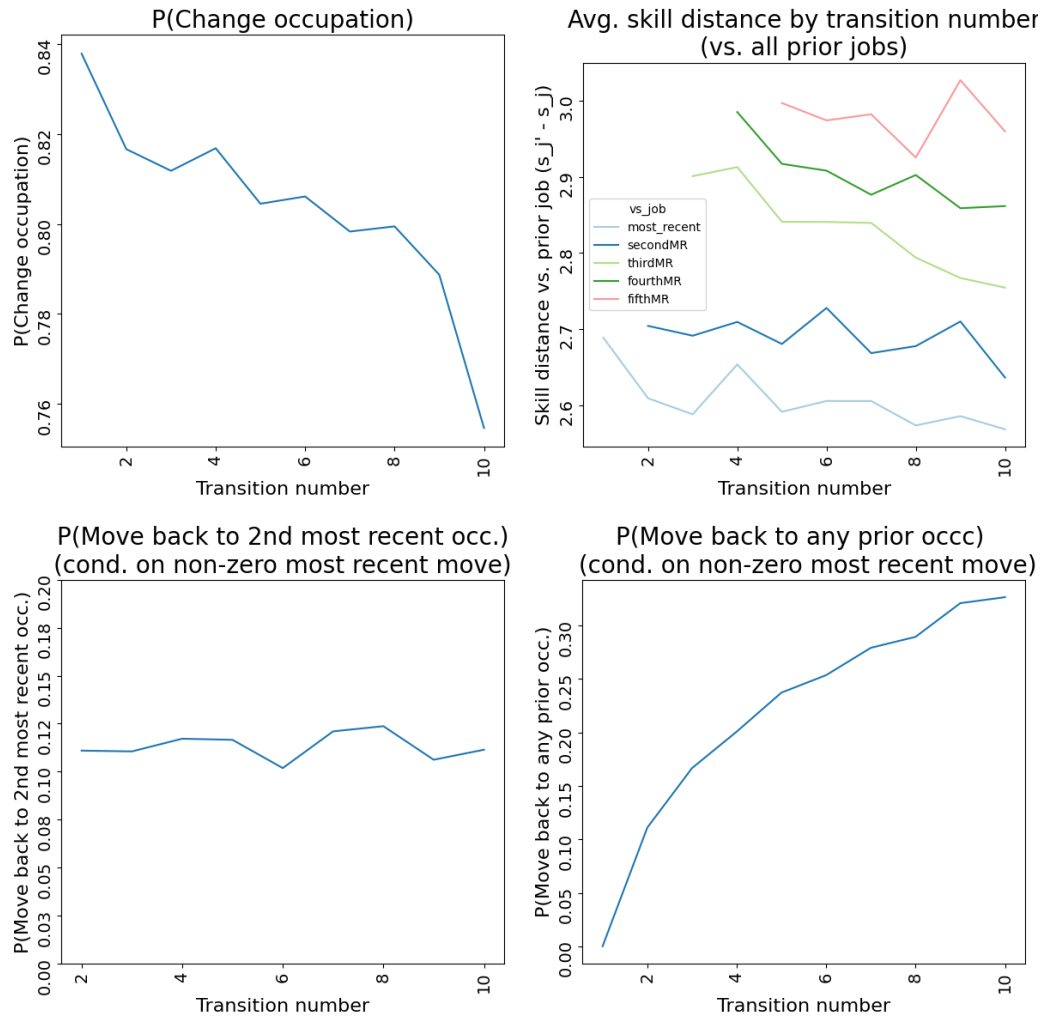

FIGURE 2.2: Transition Dynamics

carries immediate wage *growth* premiums compared to transitions in which the worker is closely matched or overskilled compared to their destination job. This raises the question: if moving up the skill ladder is so beneficial, and does not carry short-term wage *growth* penalties, why aren't more workers making these steep transitions?

Section 3.3 offers risk as a potential explanation. Specifically, workers may be hesitant to make larger transitions because they are riskier. I hypothesize that higher absolute skill distance is associated with higher risk due to information frictions—workers are well-informed about jobs that are similar to their own, but know increasingly less about jobs that are further afield. This makes larger moves risky because workers don't necessarily

know whether they are moving into overall higher-skill (and thus likely higher-quality) jobs or lower-skill ones. While it may at first seem implausible that workers don't know ex-ante whether they are moving up or down the skill ladder, recall from Figure 2.1 that workers are about as likely to move down the skill ladder as up compared to their source job, suggesting workers are not able to precisely target higher-skill jobs. Further, the multi-dimensionality of skills makes aggregate comparisons non-trivial, particularly for jobs that are very different (e.g., Ann's transition from nurse to PR manager as shown in Table 2.2 is a downward move, but this is not obvious). Supporting this hypothesis, section 3.3 demonstrates that higher risk tolerance is associated with both transition probability and the absolute skill distance of transitions, and confirms that higher absolute skill distance is indeed associated with higher risk, in the form of higher variance in realized job quality.

2.3.1 Worker-job Skill Match Matters

This section presents results that are consistent with models of skill mismatch in which task-specific skills are important for productivity. Specifically, larger job transitions, and particularly transitions with large underskill distance, are relatively rare, and underskill is predictive of training receipt in the destination job, suggesting that workers and employers care about skill match, and underskill in particular. Further, absolute skill distance is associated with wage growth penalties and negative transition conditions such as involuntary departure and longer non-work spells between jobs, suggesting that workers may tend to be "pushed" rather than "pulled" into making larger absolute skill distance moves.

First, large transitions are rare. As shown in figure 2.1, all skill distance distributions are right-skewed, and the modal transition is within occupations (i.e., the worker moves into the same occupation with a different employer).[8] This is consistent with an environment in which it is difficult or costly to be hired into or perform jobs that require very different skills from your current skill bundle. Further, relative skill distance is less than zero on

[8] Note that within-occupation transitions need not, and generally will not, have zero skill distance between the worker's skill bundle and the destination job's skill requirements, as the worker likely has accumulated other skills from prior jobs that have not fully depreciated.

average—i.e., workers are more likely to move to jobs they are relatively overskilled, rather than underskilled, for. While this is not surprising given that skills are assumed to update faster than they depreciate, it is not fully by construction[9], and this also could indicate a state of the world in which employers care about skills and are reluctant to hire workers with insufficient skills.

Secondly, absolute skill distance is associated with a higher incidence of employer-provided training receipt in the first year of the new job, and this is driven completely by underskill. Specifically, as shown in Table 2.3, an increase of one standard deviation in underskill increases the probability of receiving employer-provided training in the first year of the destination job by about 1 percentage point, or about a 20% increase from the sample average of 4.2%. This suggests that workers must adjust their skills when their current level of a given skill is too low to meet the requirements of the job, sometimes with assistance from employers. This relationship holds even after controlling for a broad array of destination job characteristics such as wages, hours, and occupation and industry categories, which implies that the observed training provision is truly driven, at least in part, by skill mismatch, rather than job complexity or job quality.

Additionally, to the extent that skill mismatch negatively impacts productivity, we should expect workers with skill mismatch to earn less than their counterparts with well-matched skills, at least initially. The data confirms this. As shown in Table 2.4, a worker's real wage *level* in the first year of a job, controlling for occupation and industry fixed effects, is negatively impacted by skill mismatch. That is, given two workers in the first year of a new job in the same occupation, the one with skills that are a closer match to the occupation's skill requirements will earn more. This is again driven more by underskill: an increase of one standard deviation in underskill is associated with a 4.7% lower real wage rate in the first year of the new job, relative to other workers in the same occupation. Interestingly, however, overskill also incurs a small but significant wage level penalty.

[9] This would not be true if, e.g., workers always moved into jobs with higher skill requirements in every skill component relative to their prior job.

Table 2.3: Any Employer-Provided Training in First Year of New Job

	(1)	(2)	(3)
Absolute skill distance	0.004***		
	(0.001)		
Underskill		0.008***	
		(0.002)	
Overskill		0.002	
		(0.001)	
Relative skill distance			0.004**
			(0.002)
Controls:			
Demographics	✓	✓	✓
Job & transition vars.	✓	✓	✓
Year, occupation & industry FEs	✓	✓	✓
Outcome variable:			
Mean	0.042	0.042	0.042
Std. Deviation	0.200	0.200	0.200
R2	0.065	0.066	0.065
N	33483	33483	33483

Robust standard errors in parentheses (* p<0.1, ** p<0.05, ***p<0.01)

Reflects a binary indicator for whether the respondent reported any employer-provided training in the first year of their new job, regressed on skill distance and controls for worker demographics, characteristics of the source and destination job and of the transition, year fixed effects, and fixed effects for the occupation and industry of the new job.

Finally, there is a meaningful relationship between worker and transition characteristics and skill distance. In particular, general human capital (in the form of both education and general ability as measured by ASVAB scores) appears to be positively associated with skill distance, suggesting that broader aptitude may be somewhat substitutable for specific skills in some cases. In particular, as shown in Table 2.5, having a college degree is strongly associated with underskill and relative skill distance, consistent with the idea that higher education may mitigate skill mismatch costs (e.g., by allowing workers to learn new skills more quickly). Additionally, there is a positive relationship between more negative transition conditions (involuntary departures, a higher unemployment rate at the time of transition) and larger transitions in terms of absolute skill distance. This could suggest that workers would generally prefer to move to jobs that are a good fit with their skills, and

Table 2.4: Log Wage Level vs. Occupation Mean in First Year of New Job

	(1)	(2)	(3)
Absolute skill distance	-0.029***		
	(0.002)		
Underskill		-0.047***	
		(0.002)	
Overskill		-0.012***	
		(0.002)	
Relative skill distance			-0.022***
			(0.002)
Controls:			
Demographics	✓	✓	✓
Job & transition vars.	✓	✓	✓
Year, occupation, & industry FEs	✓	✓	✓
Log wage level:			
Mean	6.277	6.277	6.277
Std. Deviation	0.493	0.493	0.493
R2	0.565	0.569	0.563
N	32658	32658	32658

Robust standard errors in parentheses (* p<0.1, ** p<0.05, ***p<0.01)

Reflects logged real wage level in the worker's first year of the job regressed on skill distance and controls for worker demographics, characteristics of the job and of the transition into the job, year fixed effects, and fixed effects for the occupation and industry of the job.

are more likely to make bigger moves under duress or when a close skill match is difficult to find.

Table 2.5: Correlates of Skill Distance

	Absolute skill distance	Underskill	Overskill	Relative skill distance
Demographics:				
Ability (ASVAB)	0.158***	0.112***	0.110***	-0.012
	(0.026)	(0.026)	(0.026)	(0.026)
High school grad	0.183***	0.123***	0.134***	-0.022
	(0.018)	(0.018)	(0.018)	(0.018)
Some college	0.202***	0.148***	0.142***	-0.013
	(0.033)	(0.031)	(0.033)	(0.032)
College grad	0.185***	0.291***	0.059**	0.129***
	(0.027)	(0.027)	(0.026)	(0.026)
More than college	-0.080*	0.137***	-0.131***	0.173***
	(0.041)	(0.037)	(0.040)	(0.036)
Age	0.114***	-0.147***	0.198***	-0.226***
	(0.011)	(0.010)	(0.010)	(0.010)
Age^2	-0.002***	0.002***	-0.003***	0.004***
	(0.000)	(0.000)	(0.000)	(0.000)
Male	0.260***	0.085***	0.228***	-0.111***
	(0.011)	(0.011)	(0.011)	(0.011)
Black	-0.043***	-0.018	-0.027**	0.008
	(0.013)	(0.013)	(0.013)	(0.013)
Hispanic	0.016	0.071***	-0.018	0.054***
	(0.015)	(0.015)	(0.015)	(0.015)
Other Race or Ethn.	-0.143**	-0.039	-0.113**	0.057
	(0.057)	(0.058)	(0.055)	(0.054)
Total experience (yrs)	0.001	-0.011***	0.007***	-0.011***
	(0.002)	(0.002)	(0.002)	(0.002)
Job and transition variables:				
Real wage (old job)	-0.163***	-0.068***	-0.130***	0.052***
	(0.014)	(0.014)	(0.014)	(0.014)
Hrs per wk (old job)	-0.000	-0.005***	0.002***	-0.005***
	(0.001)	(0.001)	(0.001)	(0.001)
Involuntary depart.	0.032**	-0.088***	0.082***	-0.110***
	(0.013)	(0.012)	(0.013)	(0.013)
Time between jobs (wks)	0.000***	0.001***	0.000	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)
Training between jobs	0.012	-0.003	0.013	-0.011
	(0.034)	(0.037)	(0.034)	(0.036)
Unemployment rate	0.008**	0.000	0.005	-0.003
	(0.004)	(0.004)	(0.004)	(0.004)
R2	0.034	0.046	0.051	0.062
N	32707	32707	32707	32707

Robust standard errors in parentheses (* p<0.1, ** p<0.05, ***p<0.01)

Reflects the skill distance of each observed job transition regressed on worker demographic variables and characteristics of the source job and of the transition.

2.3.2 Job Skill Level Matters More for Transition Outcomes

Although skill match is important, the skill *level* of the destination job appears to be more important for determining the outcome of job transitions. Conditional on making a job transition, higher levels of underskill and relative skill distance are associated with better transition outcomes, both immediately and in the long term. This holds true for wages as well as other measures of job quality such as job satisfaction and tenure. Further, I find that making larger upward occupational transitions is associated with sustained wage gains and higher mid-career wages conditional on the overall number of job moves, respectively.

In other words, it appears that the optimal job ladder is short (comprised of a low number of transitions) but steep (in the sense that the higher the skill gradient of each transition, the better). While it makes sense that moving into higher-skill jobs ultimately results in higher wages, it is somewhat surprising that these benefits appear to be immediate and sustained, rather than involving an initial trade-off between human capital accumulation or learning and future wage growth.

Throughout this section, results for relative skill distance, which captures the extent to which the job requires higher aggregate skills compared to the worker's current skill bundle, are highlighted. Relative skill distance is a parsimonious indicator of how far up or down the skill ladder the worker is moving. Results for underskill and overskill are also shown throughout to provide intuition for how these are each contributing to the total effect.

As shown in Table 2.6, higher levels of relative skill distance are strongly positively associated with the wage change upon transition. Specifically, an increase of one standard deviation in relative skill distance is associated with wage growth that is 5.7 percentage points higher than for the average transition, almost doubling the sample mean of roughly 6.6%. This is in stark contrast to the large and significant wage growth penalty for absolute skill distance, which is driven by overskill. Clearly, the skill *direction* of the move is an important dimension to consider when analyzing occupational transitions, as absolute distance alone fails to capture this divergence. The immediate and marked wage boost

from making a larger move into a job one is relatively underskilled for is surprising in the context of a job-specific skills model, many of which would predict at least a temporary wage growth penalty. Instead, it seems that the productivity gains from these types of larger upward jumps overshadow any skill mismatch costs.

This could be consistent with positive selection on unobservable worker or match characteristics, or with a basic job ladder model in which output is mostly driven by job quality, rather than attributes of the particular worker-job match. Indeed, selection is almost surely part of the story. As shown in Table 2.5, larger relative skill distance is associated with positive observable characteristics such as higher education (college degree or higher) and voluntary departures. However, the wage growth effects persist after controlling for all standard observables (as shown in Table 2.6, and even after including individual fixed effects (as shown in Appendix B). This suggests that the wage growth premium for steeper upward transitions is not entirely due to the selection of workers with higher unobserved productivity; rather, there appears to be some degree of luck involved in determining who moves up the skill ladder.

In addition to higher pay, higher relative skill distance is positively associated with changes in job satisfaction as well as tenure in the destination job, as shown in 2.7. In particular, a one standard deviation increase in relative skill distance is associated with an almost 25% larger increase in job satisfaction compared to the sample mean, as well as about three months of additional tenure in the new job on average. These patterns reinforce the idea that steep upward transitions result in higher-quality or more productive matches.

In particular, it is interesting to note that higher relative skill distance is associated with larger gains in job satisfaction *even after controlling for wages* (and other job attributes). One possible explanation of this pattern is that, as workers gain labor market experience, they learn private information about attributes not observable in the data (and potentially not observable to employers) that impact their productivity in different occupations (e.g., preferences over non-wage job attributes), which they use to target better-matched jobs. However, the fact that job satisfaction, job skill level, and wages tend to co-move makes

Table 2.6: Wage Growth from Transitions

	(1)	(2)	(3)
Absolute skill distance	-0.026***		
	(0.002)		
Underskill		0.031***	
		(0.002)	
Overskill		-0.042***	
		(0.002)	
Relative skill distance			0.057***
			(0.002)
Controls:			
Demographics	✓	✓	✓
Job & transition vars.	✓	✓	✓
Year FEs	✓	✓	✓
Change in log wage:			
Mean	0.066	0.066	0.066
Std. Deviation	0.419	0.419	0.419
R2	0.256	0.270	0.270
N	32031	32031	32031

Robust standard errors in parentheses (* p<0.1, ** p<0.05, ***p<0.01)

Reflects the change in logged real wage level from the end of the worker's old job to the start of their new job regressed on skill distance and controls for worker demographics, characteristics of the source and destination jobs and of the transition, and transition year fixed effects.

the fact that so many workers make downward transitions more difficult to rationalize, as it does not appear to be the case that workers are sacrificing wages for non-wage amenities (which should presumably show up in job satisfaction). This result further indicates that workers may struggle to understand which jobs are higher-skill in aggregate or to effectively target these jobs when searching.

Turning to longer-term post-transition outcomes, the wage benefits of moving up the skill ladder are immediate and sustained, based on wage paths pre- and post-transition (using each respondent's *largest* transition in terms of absolute skill distance, compared to non-movers in the same periods), as shown in Figure 2.3. While making an above-median transition in terms of *absolute* skill distance is associated with wage stagnation or even a

Table 2.7: Other Transition Outcomes

	Change in Job Sat.			Tenure in New Job		
	(1)	(2)	(3)	(4)	(5)	(6)
Absolute skill distance	0.009			-0.001		
	(0.008)			(0.012)		
Underskill		0.094***			0.231***	
		(0.008)			(0.014)	
Overskill		-0.029***			-0.097***	
		(0.008)			(0.011)	
Relative skill distance			0.091***			0.245***
			(0.008)			(0.013)
Controls:						
Demographics	✓	✓	✓	✓	✓	✓
Job & transition vars.	✓	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓	✓
Outcome variable:						
Mean	0.401	0.401	0.401	1.883	1.883	1.883
Std. Deviation	1.459	1.459	1.459	2.681	2.681	2.681
R2	0.025	0.029	0.028	0.050	0.068	0.065
N	32707	32707	32707	28250	28250	28250

Robust standard errors in parentheses (* p<0.1, ** p<0.05, ***p<0.01)

Reflects the job satisfaction from the end of the worker's old job to the start of their new job (columns (1)-(3)) or the tenure in years in the new job (columns (4)-(6)) regressed on skill distance and controls for worker demographics, characteristics of the source and destination jobs and of the transition, and transition year fixed effects. The sample for tenure regressions is limited to transitions in which the destination job spell is complete (i.e., not truncated by the end of the survey).

slight wage penalty relative to non-movers, making an above-median transition in terms of *relative* skill distance is associated with immediate and continued wage gains. It is also interesting to note that workers making larger upward transitions do not appear to be positively selected in terms of pre-transition wages or wage growth.

Finally, Table 2.8 presents the results of a Mincer-style wage decomposition of age 35 wages in the sample (roughly the highest age that all survey respondents have reached by the most recent available survey round). Making larger upward transitions on average significantly predicts mid-career wage levels in the sample, even after controlling for starting wage, total work experience, the total number of job transitions the worker has made, and

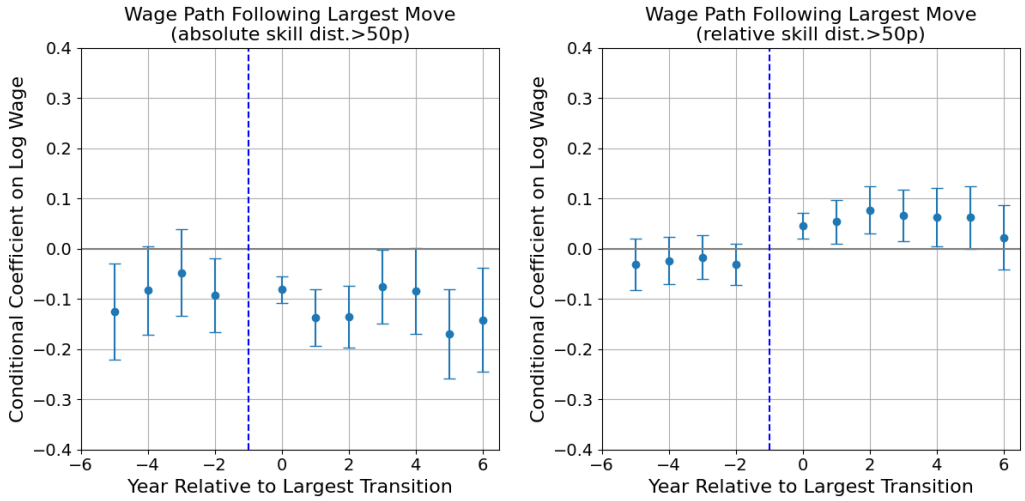

FIGURE 2.3: Wage Paths Following Transitions

a rich set of controls. In particular, a one-unit increase in the relative skill distance of the worker's *average* transition is associated with the same wage gain as over a year of additional work experience. On the other hand, there is a negative relationship between the raw count of job transitions and wages, suggesting that fewer, but steeper, upward moves (relative to the worker's current skills) are associated with the best wage trajectories.

Taken together these patterns suggest that, conditional on being able to make such a move, the job or match quality effects of moving into a job one is relatively underskilled for outweigh the effects of skill mismatch. Which raises the question: who is making these steep upward transitions, and why aren't more workers making similar transitions (or at least moving strictly upward)? As discussed in Section 3.1, there does appear to be some degree of selection. Workers with more education (specifically, at least a college degree) tend to make larger upward moves, while involuntary transitions and higher unemployment are associated with downward moves. This could align with either selection (lower-quality workers are both more likely to be fired and less able to match with highskill jobs) or a simple job ladder in which laid-off or unemployed workers have lower reservation wages. Stepping back, the R-squared for the regressions in Table 2.5 (predicting skill distance) are quite low, suggesting that, while worker characteristics and the circumstances of the

Table 2.8: Log Wage Level at Age 35

	(1)	(2)	(3)	(4)
Cumulative avg., absolute skill dist.	-0.010			
	(0.007)			
Cumulative avg., underskill		0.059***		
		(0.011)		
Cumulative avg., overskill			-0.030***	
			(0.007)	
Cumulative avg., relative skill dist.				0.043***
				(0.006)
Real wage in first job	0.292***	0.312***	0.290***	0.299***
	(0.021)	(0.021)	(0.021)	(0.021)
Total experience (yrs)	0.037***	0.036***	0.037***	0.036***
	(0.002)	(0.002)	(0.002)	(0.002)
Total number of transitions	-0.010***	-0.007***	-0.008***	-0.005**
	(0.002)	(0.002)	(0.002)	(0.002)
Controls:				
demographics	✓	✓	✓	✓
year FEs	✓	✓	✓	✓
Log wage level at age 35:				
Mean	6.722	6.722	6.722	6.722
Std. Deviation	0.572	0.572	0.572	0.572
R2	0.358	0.363	0.361	0.366
N	3827	3827	3827	3827

Robust standard errors in parentheses (* p<0.1, ** p<0.05, ***p<0.01)

Reflects the logged real wage level of each worker at age 35 (for workers with a wage observation available) regressed on the average skill distance across all the transitions the worker has made up to age 35, along with controls for the worker's wage level in their first observed job, their total work experience, the total number of job transitions they've made, worker demographics, and year fixed effects.

transition do play a meaningful role in the career moves workers choose to or are able to make, there may also be a good deal of randomness. To the extent that the degree of randomness—or risk—is correlated with the size of the transition (in terms of absolute distance), heterogeneous risk tolerance may help explain these patterns. The following section explores this idea.

2.3.3 Ladders or Chutes: the Role of Risk

Given the preceding results, it is difficult to understand why more workers don't make larger upward transitions, or why so many transitions are actually *down* the skill ladder. I propose one explanation in which transitions are risky because workers have imperfect information about job skill levels ex ante, and cannot precisely target higher-skill jobs. In other words, they may be uncertain whether a particular transition would be a "ladder" (upward move) or a "chute" (downward move). Additionally, larger skill distance transitions are riskier due to higher uncertainty—workers know less about jobs that are less similar to their current work. Thus, workers may be hesitant to make larger moves not (or not only) because their skills are mismatched, but also because they are risk averse and larger transitions carry more risk. Supporting this hypothesis, I show that higher risk tolerance is associated not only with the probability of making a job transition, but also with absolute skill distance—i.e., more risk-tolerant workers make larger moves. Additionally, higher absolute distance is in fact riskier in terms of realized destination job quality.

As one might expect, higher risk tolerance is correlated with a higher probability of making *any* job transition. However, the size of the effect is relatively small. As shown in Table 2.9, moving from the lowest possible level of risk tolerance to the highest is associated with an increase of roughly 1.5 percentage points in the probability of making a transition, which is small relative to the sample average of about 28%. Though small, the relationship between risk tolerance and transition probability is robust to a large set of demographic and source job controls.

Beyond this relationship between risk tolerance and binary transition probability, risk tolerance also predicts the *size* of the transition. As shown in Table 2.10, moving from the lowest possible level of risk tolerance to the highest is associated with a significant increase in absolute skill distance—about double the effect of involuntary departure (from Table 2.5). There is generally not a significant relationship between risk tolerance and the other skill distance measures, which capture direction. This makes sense in a world where workers

Table 2.9: Probability of Any Job Transition in a Given Year

	(1)	(2)	(3)	(4)
Risk tolerance	0.014***	0.014***	0.015***	0.013***
	(0.004)	(0.004)	(0.004)	(0.004)
Controls:				
Demographics	✓	✓	✓	✓
Job vars.		✓	✓	✓
Year FEs			✓	✓
Ooccupation & industry category FEs				✓
Transition probability:				
Mean	0.280	0.280	0.280	0.280
R-squared	0.102	0.133	0.143	0.145
N	148490	148490	148490	148490

Robust standard errors in parentheses (* p<0.1, ** p<0.05, ***p<0.01)

Reflects a binary variable reflecting whether each worker made a job transition in each year in which they are observed in the data, regressed on risk tolerance (normalized to a 0 to 1 scale) and (depending on the specification) controls for worker demographics, characteristics of the worker's current job, year fixed effects, and occupation and industry fixed effects for the worker's current job.

are not able to precisely identify higher-skill jobs. Indeed, this is potentially the source of the risk. While more risk-tolerant workers are more willing to move into jobs that are more different from their own, they don't seem to be able to move into strictly higher-skill jobs.

Table 2.10: Skill Distance and Risk Tolerance

	Absolute skill distance	Underskill	Overskill	Relative skill distance
Risk tolerance	0.062***	0.029	0.039*	-0.006
	(0.021)	(0.020)	(0.020)	(0.020)
Controls:				
Demographics	✓	✓	✓	✓
Job & transition vars.	✓	✓	✓	✓
Year & occupation, industry cat. FEs	✓	✓	✓	✓
R2	0.063	0.094	0.139	0.162
N	30347	30488	30488	30488

Robust standard errors in parentheses (* p<0.1, ** p<0.05, ***p<0.01)

Reflects the skill distance of each observed job transition regressed on risk tolerance (normalized to a 0 to 1 scale) and controls for worker demographics, characteristics of the worker's current job, year fixed effects, and occupation and industry fixed effects for the worker's current job.

Finally, bigger skill distance moves *are*, in fact, riskier. Proxying occupation quality

with the mean real wage of each occupation in the sample, Figure 2.4 shows that higher absolute skill distance is associated with greater variance in the change in job quality following a transition. Specifically, the dispersion of outcomes is monotonically increasing in the absolutes skill distance quartile of the transition. In other words, making a larger transition in terms of absolute skill distance can result in moves to much higher-quality jobs *or* much lower-paying jobs. Notably, this is true even when limiting the sample to voluntary transitions; while the distributions for involuntary departures are less favorable, moves to lower-quality occupations are not driven only by these involuntary transitions.

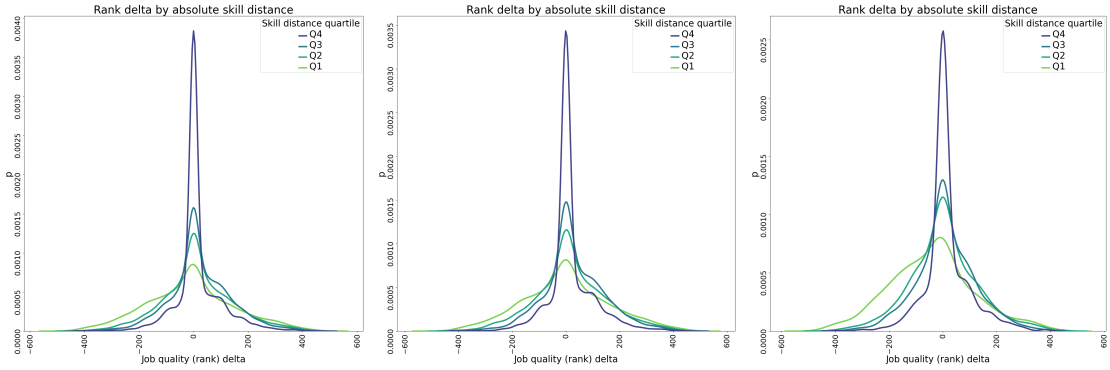

(a) All Transitions (b) Voluntary Transitions (c) Involuntary Transitions

FIGURE 2.4: Change in Job Quality by Absolute Skill Distance Quartile

Taken all together, these patterns are consistent with a world in which big career moves are risky because job skill level and quality are difficult to observe during job search, with higher uncertainty about jobs that are more distant from your current field. Due to this uncertainty, more risk-averse workers avoid making larger transitions. Thus, even though large *upward* transitions are so advantageous, we do not observe many in practice. Further, since workers are not able to perfectly observe aggregate job skill levels ex ante, we see workers making transitions both up and down the skill ladder. The next section presents a model exploring the implications of these hypothesized conditions.

2.4 Model

2.4.1 Model Set-up and Intuition

Workers are characterized by their education level α_i (less than college or college graduate), skill level κ_i, and heterogeneous risk aversion γ_i. Jobs are characterized by their skill requirement κ_j and job quality q_j, which is assumed to be imperfectly correlated with skill requirements (as observed in the data).

There are two periods. In period 1, workers choose a job search strategy s defined by an *absolute* skill distance range. Workers are only able to choose the absolute skill distance over which they are searching for jobs, reflecting the idea that they are uncertain about aggregate job skill levels ex-ante, and thus cannot precisely target strictly higher-skill jobs. Workers then draw a new job from a known distribution that depends on their chosen strategy s. They receive a wage in the new job determined by their education and skill levels and the new job's skill requirements and quality. Thus, workers choose a search strategy s to maximize expected utility in period 2 according to:

$$\max_s \mathbb{E}[u(w_{ij'}(\alpha_i, d_{ij'}, q_{j'}), \gamma)|s, \alpha_i, \kappa_i] \tag{2.1}$$

where $w_{ij'}$ is the wage in the destination job (j'), α_i is the worker's education level, $d_{ij'} = \kappa_{j'} - \kappa_i$ is the distance between the job's skill requirement and the worker's skill level, and $q_{j'}$ is the job's quality.

To provide intuition about the mechanisms through which skill mismatch and risk tolerance influence workers' job search decisions in this setting, the remainder of the section solves a *simplified* version of the model. The model specification used for estimation is presented in the following section.

Solution to simplified model: Start by specifying the wage function as follows:

$$w_{ij} = \beta_e \alpha_i + \beta_q \kappa_j + \beta_s(\kappa_i - \kappa_j) \tag{2.2}$$

Where β_e captures the effect of education on the wage level, β_q captures the effect of job

quality (assumed to be equal to skill level in this simplified version of the model) and β_s captures the effect of skill distance, or skill mismatch.

The job search strategy determines the *variance* of the distribution of skill levels from which the worker is drawing their new job. This captures the idea that workers can influence the riskiness of their job search strategy. The specification used in estimation introduces more complex distributions that approximate discontinuous skill distance ranges, reflecting the idea that workers can choose to look "closer" or "further" from their current job, but these discontinuities make a closed-form solution intractable. So, for this example specify the distribution of new job skill level $\kappa_{j'}$ as a function of search strategy s as follows:

$$\kappa_{j'} \sim N(\kappa_i + ds, \sigma^2(s)) \tag{2.3}$$

Where κ_i is the worker's current skill level (i.e., search is centered around the worker's current skill), $d \geqslant 0$ is a risk premium implying that workers who choose riskier search strategies are able to (imprecisely) target somewhat higher skill jobs by shifting the center of the distribution upwards, and $\sigma^2(s)$ is the variance of the distribution, which is assumed to be monotonically increasing in s. In other words, choosing a higher level of s results in higher risk.

Finally, worker utility over wages follows a simple CARA form:

$$U(w_{ij}) = -\exp(-\gamma_i w_{ij}) \tag{2.4}$$

where $\gamma_i > 0$ is worker i's coefficient of absolute risk aversion.

Then, workers choose s to maximize expected utility in their new job j':

$$\max_s E[U(w_{ij'})] = -\exp\left(-\gamma E[w_{ij'}] + \tfrac{1}{2}\gamma^2 \operatorname{Var}[w_{ij'}]\right)$$

Taking the expectation of wages over the distribution of new job skill level depending on search strategy s yields expected wage:

$$E[w_{ij'}] = \beta_e \alpha_i + \beta_q \kappa_i + \beta ds \tag{2.5}$$

with variance: $\text{Var}[w_{ij'}] = \beta^2\sigma^2(s)$ (where $\beta = \beta_q - \beta_s$).

Substituting this into the utility function gives expected utility:

$$E[U(w_{ij'})] = -\exp\left(-\gamma(\beta_e\alpha_i + \beta_q\kappa_i + \beta ds) + \tfrac{1}{2}\gamma^2\beta^2\sigma^2(s)\right)$$

Collecting only the terms dependent on s into Ψ yields the following first order condition:

$$\frac{d\Psi}{ds} = -\gamma\beta d + \tfrac{1}{2}\gamma^2\beta^2\frac{d\sigma^2(s)}{ds} = 0$$

Rearranging terms, we can derive the following condition for the optimal search strategy:

$$\frac{1}{2}\gamma^2\beta^2\frac{d\sigma^2(s^*)}{ds} = \gamma\beta d \implies \frac{d\sigma^2(s^*)}{ds} = \frac{2d}{\gamma\beta}$$

The above implies that, as long as $\beta_q > \beta_s$, more risk-tolerant workers choose higher s because they are willing to accept higher variance for higher expected wages. In other words, as long as the skill mismatch penalty is sufficiently low, the wage benefits of moving up the skill ladder outweigh the wage penalty of skill mismatch, and workers will accept some degree of risk for the chance to move up the skill ladder, depending on their risk tolerance.

2.4.2 Estimation

The model specification used for estimation is more complex, and does not permit a closed-form solution. It includes a richer wage equation, CRRA utility, and discontinuous distributions of new job skill level that capture the idea of workers searching in absolute skill distance "ranges."

Wage determined based on worker education and skill level and job skill level and quality as follows:

$$w_ij = \beta_1 q_j + \beta_2 d_{ij} + \beta_3((1-\alpha_i) * d_{ij}) + \beta_4(\alpha_i q_i) \tag{2.6}$$

where q_j is job quality, $d_{ij} = \kappa_j - \kappa_i$ captures relative skill distance, and α_i is the worker's education level. For computational tractability, skill levels and job skill requirements are

reduced to uni-dimensional variables capturing aggregate skill levels. There are two education (α) groups, comprised of workers without a college degree and workers with a college degree.

Worker utility is given by:

$$U(w_{ij}) = \frac{w_{ij}^{(1-\gamma_i)}}{(1-\gamma_i)} \tag{2.7}$$

For estimation, workers take one of four possible values of risk aversion γ_i, determined based on their self-reported risk aversion and mapped into values ranging from 0 (no risk aversion) to 5 (highest possible risk aversion).

Workers choose one of four possible skill distance strategies ($s = 0$ to 3), where 0 represents the special case of a zero-distance move (stay in the same occupation) and 3 targets the largest absolute skill distance. This richer distributional specification is meant to capture the idea that workers are able to tell roughly how different a job is from their current one (absolute skill distance), but not necessarily whether it's higher- or lower-skill (relative distance). In this version, workers choose a search strategy that determines the distribution of *absolute* skill distance from which they are drawing their new job as follows:

$$D(s) \sim N(\mu_D 2^{s-1}, \sigma_D^2 \exp(1.1s)) \tag{2.8}$$

where μ_D and σ_D are distributional scaling parameters to be estimated. Intuitively, choosing a higher s means a larger mean absolute skill distance as well as higher variance in the distribution. The skill level of the new job is assumed to be centered around the worker's current job skill level (as suggested by the symmetry of relative skill distance around the worker's source job as shown in Figure 2.1).

Whether *realized* relative skill distance $d_{ij'}$ is positive or negative (upward or downward move) then depends on a direction parameter p_{up} (reflecting the idea that workers don't observe new job skill level $\kappa_{j'}$ or quality $q_{j'}$ ex-ante) and a targeting parameter τ. To provide some intuition for what these distributions look like in practice, Figure 2.5 shows the distribution of new job skill levels $k_{j'}$ for each non-zero search strategy (recall that for

$s = 0$, workers simply remain in the same occupation, and thus get the same skill level). This example shows distributions for a worker with current job skill level $\kappa_j = 0.5$ (hence the distributions are centered around 0.5). This specification reflects the idea that workers may target their job search in approximate ranges of absolute skill distance ("I want a job similar to what I'm doing now" vs. "I want to do something totally different").

FIGURE 2.5: Distribution of New Job Skill by Search Strategy

Due to the CRRA utility specification and discontinuities in the skill level distributions, this version of the model does not permit a closed-form solution for the worker's optimal search strategy. Therefore, I use simulated method of moments to estimate the model parameters. I estimate eight parameters in total: the parameters of the wage equation (capturing the wage effects of skill mismatch (β_2), the interactions of education level*skill mismatch (β_3), and education level*job quality (β_4); β_1 is held fixed at 1 for scaling), the scaling parameters of the new job skill distribution δ_μ and δ_σ, and the directional parameters of the new job skill distribution (τ, which captures the extent to which workers are able to target higher-skill jobs, and the probability of moving up vs. down the skill ladder *by education category*: $p_{up,non-col}$ and $p_{up,col}$). Note that the probability of making an upward

move is permitted to vary with education level, reflecting patterns observed in the data.

I estimated the model by matching moments relating to the absolute skill distance and wage outcomes of each respondent's largest (in terms of absolute skill distance) transition in the sample. In particular, the estimation targets the mean and standard deviation of absolute skill distance for each education-by-risk tolerance group, as well as the mean and standard deviation of wage growth for each of these groups. As there are two education levels and four levels of risk tolerance, this results in 32 targeted moments.

2.4.3 Parameter Estimates and Fit

The model is able to replicate empirical patterns of interest and closely matches data moments. In particular, as shown in Figure 2.6, we see that the model generates a negative relationship between absolute skill distance of transitions and risk aversion, as observed in the data. Additionally, realized wage growth is higher for moves up the skill ladder, and higher for college-educated compared to non-college-educated workers.

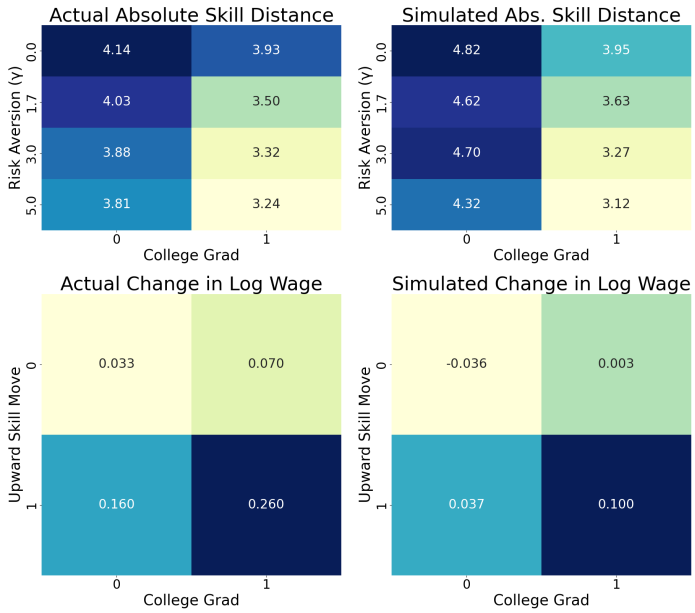

FIGURE 2.6: Moment Match

The parameter estimates also align with the empirical patterns documented previously. In particular, as shown in Table 2.11, the effect of skill distance on wages (β_2 and β_3) is

small compared to the impact of job quality and education level (β_1 and β_4) (recall that β_1 is fixed at 1 for scaling). Additionally, college-educated workers ($p_{up,col}$) are much more likely to move up the skill ladder than non-college-educated workers ($p_{up,non-col}$), which is consistent with the empirical result that a college degree predicts relative skill distance. Notably, this was not a targeted moment. Finally, the positive estimate for ($\tau > 0$) indicates that workers are able to noisily target higher-skill jobs to some extent.

Table 2.11: Parameter Estimates

Parameter	Value
β_1	1.0000
β_2	-0.0331
β_3	-0.0233
β_4	1.3091
δ_μ	0.0083
δ_σ	0.0162
τ	0.1616
$p_{\text{up, non-col}}$	0.2272
$p_{\text{up, col}}$	0.6786

2.4.4 Counterfactual: Reducing Riskiness of Job Search

We can use the model to explore the impact of risk on worker's transition decisions. In particular, to the extent that larger transitions are riskier due to higher variance in outcomes, risk-averse workers may be sacrificing potential wage gains from larger transitions. To explore this idea, I run a counterfactual simulation with much lower variance (dividing the estimated scaling parameter δ_σ by 10). Results are shown in Figure 2.7. Distributions of skill distance and wage change using the actual estimated parameters are shown in blue, while distributions generated using the lower δ_σ value are shown in green. As anticipated, reducing the variance component of risk (δ_σ) increases skill distance. However, the effect on wages is mixed—the distribution of wage changes becomes more skewed relative to the baseline, with a higher median but lower average. Notably, the probability of moving up or down the skill ladder was unchanged. Simply reducing the variance in outcomes of larger transitions, while preserving the downside risk, induces workers to make larger transitions.

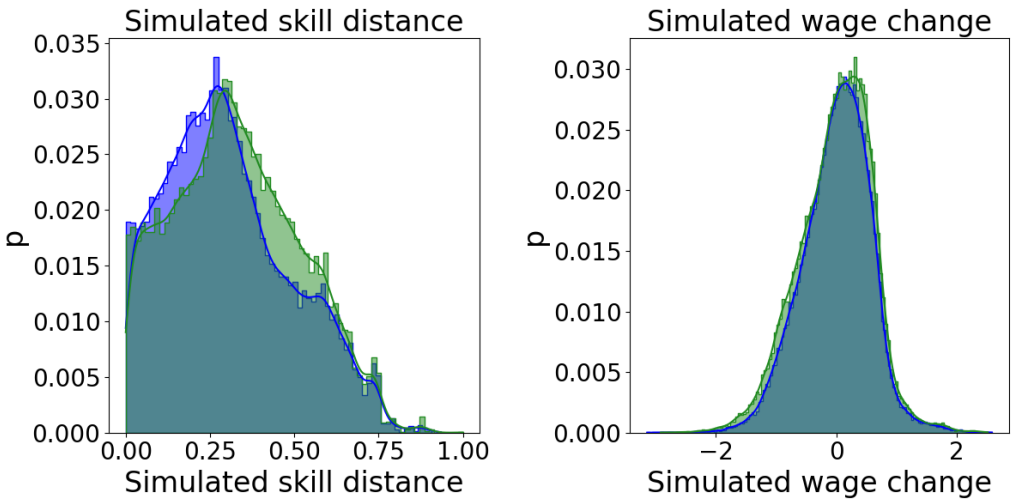

FIGURE 2.7: Counterfactual Simulation: Reducing Risk

2.4.5 Counterfactual: Comparing Skills Training and Job Search Assistance

The relative importance of skill mismatch compared to risk aversion in worker transition decisions has important implications for the effectiveness of different active labor market policies. I use the model to present a stylized example of this dynamic. In particular, to the extent that workers are avoiding larger transitions due more to risk aversion than a lack of the necessary skills, job search assistance that tempers the riskiness of larger moves will be more effective at improving wage outcomes than skill training programs.

I first simulate the effects of a "skills training" intervention that increases each worker's current skill level s_i by a moderate amount (based on the skill updating rate over a one-year period). As shown in the top column of Figure 2.8, this intervention does not shift worker search strategies or wage growth outcomes. This is due to the fact that the estimated skill mismatch penalty ($\beta_2 + \beta_3$) is small compared to the effect of job quality, and most workers are already relatively highly skilled compared to the distribution of jobs they are drawing from. Note that this is likely an overly restrictive characterization of most skills training interventions, as it assumes that the intervention does not change the distribution of jobs the worker is drawing from (whereas many skills training interventions could feasibly be

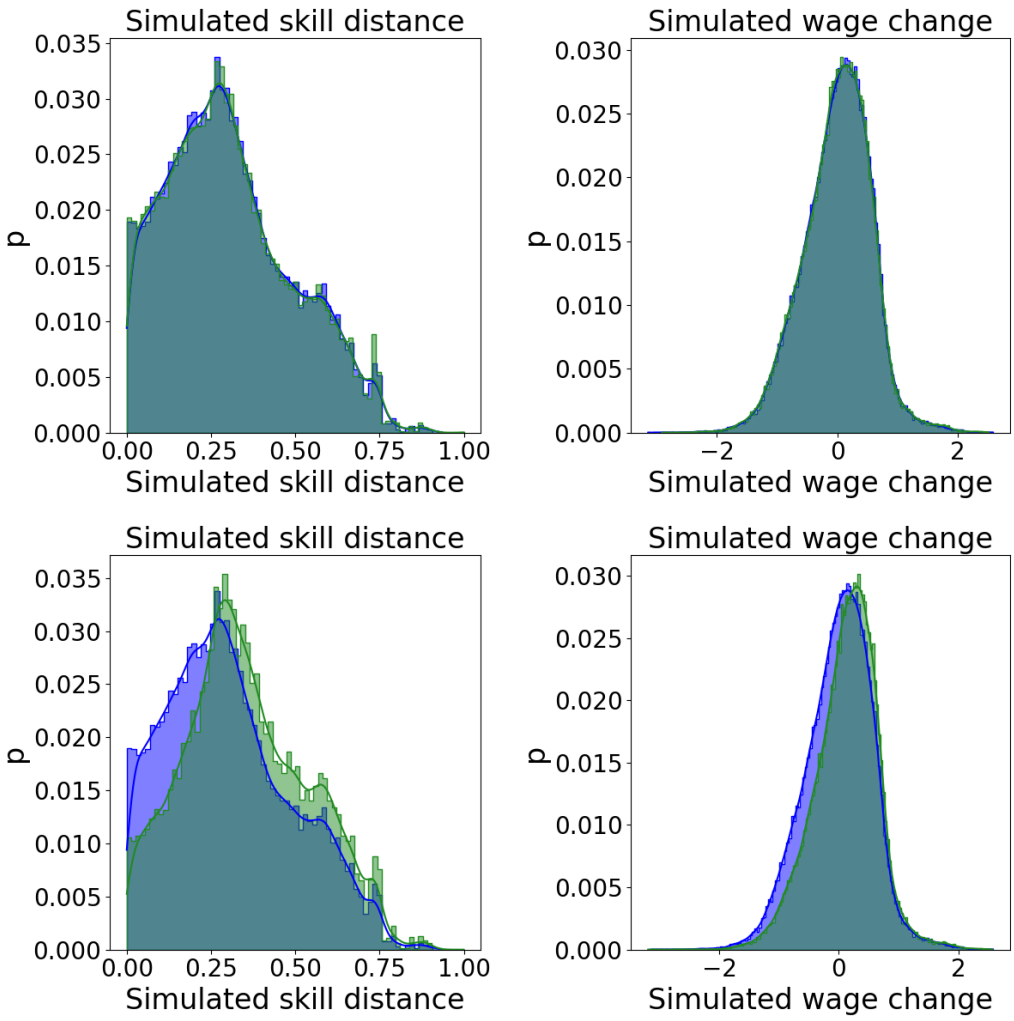

FIGURE 2.8: Counterfactual Simulations: Skill Training vs. Job Search assistance

signaled to employers, thus presumably enhancing workers' access to higher-skill jobs).

I next simulate the effect of a "job search assistance" intervention that increases the targeting parameter τ by $\approx 1/3$ of the std. deviation of absolute skill distance. In other words, this intervention helps workers to better target higher-skill jobs, shifting the distributions they are drawing from to the right. As expected, many workers choose to make larger transitions when facing more favorable distributions. This can be seen in the bottom lefthand graph in Figure 2.8, as the distribution of realized absolute skill distances of transitions shifts rightward, with a particular redistribution of mass from the lower values of absolute skill distance. This combination of larger transitions and better job targeting results in

higher wage growth for workers on average, as it allows them to move into higher-skill jobs.

While these simplified counterfactuals abstract from much of the real-world complexity of these policies, they capture an important idea. To the extent that workers are not making large transitions because of risk aversion rather than insufficient skills, skills training will be less effective than we would expect without accounting for risk aversion. For example, much of the discourse around displaced manufacturing workers impacted by the "China shock" has been focused on retraining and skills updating. However, such policies may fail to induce labor reallocation if risk aversion is also at play.

2.5 Conclusion

The results presented above show that large, steep transitions up the job ladder yield large and immediate wage growth premiums compared to other transitions. That is, workers moving into jobs for which they are meaningfully underskilled realize the highest wage growth upon transition, even compared to workers whose skills are well-matched to their new job. However, such transitions are rare, which is often attributed to skill mismatch costs. This paper presents risk as an additional explanation. In particular, to the extent that larger occupational transitions are riskier and workers are risk averse, this risk may discourage big career jumps. Empirical patterns support this hypothesis, and a simple model of worker transition decisions incorporating this idea is able to match observed transition outcomes quite well. Directions for future research include exploring the role of selection in making larger upward moves, which is partially addressed in this analysis with worker fixed effects, but is difficult to fully account for without firm-side data. The role of worker risk tolerance compared to skill mismatch in transition decisions has potentially important implications for active labor market policies, such as weighing the benefits of job search assistance compared to retraining.

3. Soft Skills and Hiring

3.1 Introduction

Developing country labor markets often suffer from inefficiencies that contribute to high turnover and persistently high unemployment, especially among youth (Donovan et al., 2023). A significant factor is information frictions: firms have limited information about prospective workers' skills, particularly "soft" skills that are difficult to observe but correlate with productivity (Adhvaryu et al., 2023; Autor & Scarborough, 2008). This lack of information can lead to lower productivity, wages, and higher worker turnover (Blattman & Dercon, 2018). At the same time, these frictions can make firms wary of hiring younger, less experienced workers, making it difficult for young people to get a foothold in the formal labor market and stunting human capital accumulation.

In South Africa, these information frictions are highly relevant. The unemployment rate exceeds 30% (Africa, 2023), and firms express a willingness to pay for information about applicants' soft skills (Abel et al., 2020; Carranza et al., 2022). High turnover rates (Benner et al., 2007; Zizzamia & Ranchhod, 2019) further exacerbate the problem. These issues are not unique to South Africa—many other developing-country labor markets are also plagued by persistent unemployment along with indicators of information frictions and frequently low match quality (Abebe, Caria, & Ortiz-Ospina, 2021; Bassi & Nansamba, 2022; Hardy & McCasland, 2023). Thus, finding ways to reduce these frictions could be extremely impactful given the scale and generalizability of the problem.

In an effort to address these frictions, we provide experimental evidence on how incorporating information about job applicants' soft skills in firms' hiring decisions affects both firm and workseeker outcomes. The soft skills measure we utilize captures concepts like initiative, tenacity, and willingness to learn over time, which have been shown to meaningfully predict important behaviors and outcomes (Duckworth & Quinn, 2009; Dweck, 2006; Gable et al., 2000). Partnering with the largest recruitment agency in South Africa, we randomize the criteria used to shortlist job applicants for job listings at partner firms. We

test whether including measures of soft skills in candidate ranking leads to better firm-worker matches, more inclusive hiring, and more hiring overall. We are also able to test for negative spillovers for jobseekers who are randomly displaced from shortlists due to our treatment, whom we term "counterfactual" applicants. Our results are largely positive, suggesting that including soft skills information leads to increased overall hiring with limited negative spillover effects.

Our experiment randomizes roughly 1,000 online job listings on our partner's job search and matching platform. All employers who list a job on the platform receive a ranked shortlist of applicants who are considered to be the best potential match for the vacancy. Our treatment alters the algorithm used to rank applicants for this shortlist. While applicants to job listings in the control group are ranked using the platform's current default algorithm, which is based on traditional criteria such as educational credentials and formal work experience, applicants to job listings in the treatment group are ranked using our updated algorithm, which incorporates information about applicants' soft skills in addition to these original criteria. Notably, the experiment is double-blind in the sense that neither jobseekers nor employers know which algorithm a given job listing is using—all listings look the same to applicants, and all shortlists look the same to employers. After the randomization, shortlisting, and hiring for a given vacancy has been completed, we collect data on hiring outcomes for the specific vacancy as well as broader job search and employment outcomes, for both shortlisted and non-shortlisted applicants.[1]

Our first set of results document the direct effect of our treatment on the hiring process and outcomes, broadly suggesting that incorporating soft skills boosts hiring throughout the recruitment funnel. In particular, we find that employers in the treatment group issue more interview invitations and make more hires compared to those in the control group. We also observe directional evidence of more inclusive hiring—treated vacancies tend to hire a greater number of younger workers with less formal labor market experience—consistent

[1] We are also in the process of collecting longer-term follow-up data from firms and workseekers in order to study outcomes such as attrition, job and firm performance, and career trajectories over a longer time horizon. We plan to include these longer-term results in a subsequent paper.

with the idea that including information about soft skills in the hiring process may help identify strong candidates from disadvantaged backgrounds who may lack formal credentials.

At the same time, our analysis of "counterfactual" applicants suggests that the intervention does not lead to negative spillover effects that reduce employment prospects for untreated workers. We find that counterfactually shortlisted applicants (those who would have been included on the shortlist under the alternative ranking algorithm to the one actual applied for each vacancy) are less likely to be employed at the time of our short-term follow-up survey, but this effect is entirely driven by the direct displacement mechanism. In other words, these applicants are less likely to receive offers for jobs in our experiment because they were not included on one final shortlist given to employers, which is the specific treatment we implement. However, when we adjust for this direct displacement effect, we find no other evidence that untreated jobseekers suffer broader employment losses. Specifically, the probability of being employed outside our experimental vacancies is unchanged, and we find no evidence of discouragement effects in the job search behavior post-treatment of counterfactual list applicants. This suggests that workers who miss out on a job in our experiment do not face lasting disadvantages in the broader labor market.

This paper contributes to two broad areas of existing literature. First, it adds to a large body of existing work evaluating the effectiveness of different labor market intermediation services in developing countries. Our experiment provides novel insights in this area by (1) exploiting our unique access to both firm- and worker-side data, which allows us to better understand the potential mechanisms of our treatment effects and also directly study the effect of treatment on the composition of hired workers, and (2) using follow-up data on "counterfactual" applicants (those displaced on shortlists by our treatment) to explore the potential spillover effects of labor market policies that shift hiring criteria—a question with critical and timely policy relevance and little conclusive evidence. Second, by testing the inclusion of soft skills measures as one specific potential route to improving labor market intermediation, we add to a relatively new but growing literature on soft skills in hiring.

Providing intermediation services to reduce labor market frictions has long been viewed as a promising direction for policy and research but has yielded mixed results in practice, particularly in developing countries. However, recent developments in the accessibility of large online job-matching platforms such as the one used in our study offer new opportunities for intervention design and scale (Behaghel et al., 2024; Belot et al., 2019; Ben Dhia et al., 2022; Fernando et al., 2023; Kelley et al., 2024; L. Wheeler et al., 2022). In high-income countries (e.g. the U.S. and Europe), public and private job-matching services have been widely implemented and evaluated. Meta-analyses of active labor market programs find that job search assistance and placement services generally have modest but positive short-run effects on employment, with larger gains in the medium and longer run, especially for disadvantaged jobseekers, though results vary by context (Card et al., 2018). In developing countries, however, labor market intermediation experiments have often found disappointingly small effects on employment, consistent with severe structural frictions. More traditional, in-person interventions like transport assistance or job fairs generally have not led to significant employment gains (Abebe et al., 2023; Bassi & Nansamba, 2022; Groh et al., 2015). At the same time, low-touch online interventions, like encouraging unemployed youth to sign up on job portals or providing information about vacancies, also often show no effect (Afridi et al., 2023; Chakravorty et al., 2023; Jones & Sen, 2022). One interpretation is that when frictions are primarily informational, simply expanding access or contact is not enough—the content of information must improve. Our results suggest that including soft skills in hiring may help achieve this objective.

Our first novel contribution in this space is studying the impact of our labor market intermediation with data from both sides of the labor market. In particular, since our experiment is conducted on a centralized job search and matching platform, we can observe rich data on applicant demographics and job search and application behavior, as well as detailed data on each job listing and the employer posting it. By contrast, much of the existing literature in this area focuses on only one side of the market. For example, a large number of papers study worker-facing interventions that improve the information available

to jobseekers (e.g., about their skills or labor market prospects), but are unable to study any resulting changes in firm hiring directly (Busso et al., 2023; Heller & Kessler, 2024). On the firm side, existing work testing alternate hiring processes is generally unable to shed light on the effect of treatment for applicants, especially those who aren't hired (Autor & Scarborough, 2008; Hoffman et al., 2018). Thus, our combination of worker and firm data allows us to provide a novel set of results that captures both firms-side outcomes, such as how many interviews and job offers are extended per vacancy, as well as outcomes for all applicants (whether shortlisted or hired or not) such as employment (whether with a firm in the sample or not) and subsequent job search behavior.

Our second novel contribution to work studying labor market intermediation services is our "counterfactual" analysis, which allows us to study the spillover effects of our treatment. This is crucial to understanding whether an intervention will generate net positive employment effects in a labor market in aggregate after accounting for impacts on untreated workers. As such, it has increasingly been a focus of research in this area. Perhaps most prominently, Crépon et al. (2013) find that positive effects of intensive, in-person job search advice in France are almost completely offset by negative spillovers on untreated jobseekers. However, evidence from other contexts paints a more mixed picture. Cheung et al. (2025) find that negative spillovers offset some but not all of the positive effects of more frequent meetings between jobseekers and job search advisors in Sweden. Gautier et al. (2018) similarly find small but positive effects of mandatory job search advice and applications in Denmark once they take spillovers into account. Ferracci et al. (2014) show that training for unemployed jobseekers in France produces negative spillovers on other jobseekers. Field et al. (2023) find no evidence of negative spillovers on untreated jobseekers of an intervention that encourages additional job applications in Pakistan. We add a piece to the puzzle by providing a new approach for analyzing spillover effects, as well as evidence pointing to relatively small spillovers in our context.

Given the specific intervention we are testing, we also contribute to a growing literature on soft skills in hiring. While several papers study the effects of varying hiring information

or criteria, our results reflect a unique combination of focus and experimental design. For example, many existing papers investigating hiring criteria alter several aspects of the criteria or hiring process at once, making it difficult to judge the efficacy of any single criterion (Autor & Scarborough, 2008; Hoffman et al., 2018). Additionally, among interventions that study soft skills exclusively, ours is the only one that is double-blind, allowing us to learn more about the mechanisms through which incorporating soft skills into the hiring process improves outcomes. In particular, the closest existing work in this area studies the inclusion of soft skills in hiring through formal testing and certifications distributed to either the worker or the firm (Bassi & Nansamba, 2022; Laajaj & Macours, 2021), potentially complicating interpretation. In our case, on the other hand, we can rule out mechanisms such as workers updating their beliefs about their own soft skills or firms receiving explicit, formal signals of these skills. Instead, it appears that firms already screen for soft skills to some extent (likely implicitly), and thus the control ranking algorithm, by ignoring soft skills, wasn't delivering applicants that were best-matched with employer preferences on average. This result adds important nuance to existing evidence on soft skills in hiring.

Section 2 describes the context of our experiment, while Section 3 lays out the experimental design. Section 4 discusses the data sources and sample used in our main analysis. Section 5 provides results, and Section 6 concludes.

3.2 Context

This section describes the research context of our experiment. It provides additional detail on the economic environment in which the study takes place, as well as a description of our research partner and the online job matching platform through which the experiment is implemented. This context is important for interpreting the experiment and results, as well as for understanding the potential for generalizability and scalability of the intervention.

3.2.1 Economic Environment

Our study takes place in South Africa, a middle-income country that faces a blend of labor market challenges. Like many developing countries, South Africa's economy is

characterized by high unemployment, particularly among youth. As of the first quarter of 2023 (roughly when our study began), the overall unemployment rate stood at over 30%, while the unemployment rate for those under 34 years old stood at over 40% (Africa, 2023). Wage labor is the primary source of income for most workers and most employment is within formal firms, though not always covered by formal contracts. Similar to many developed economies, entry-level employment is mainly in service industries (especially retail and restaurants) or, to a lesser degree, manufacturing (Africa, 2023). Given the limited viability of entrepreneurship or self-employment, helping improve the hiring efficiency and match productivity of vacancies in formal firms is of critical importance for improving labor market outcomes in this setting.

At the same time, firms struggle to identify and hire workers who are a good match with the job. Despite such high rates of unemployment, there are also high rates of turnover and absenteeism among workers, suggesting low match quality (Benner et al., 2007; Zizzamia & Ranchhod, 2019). This reflects in part the difficulty of screening and hiring suitable workers due to information frictions. In particular, employers lack credible information about workers' skills, as indicated by existing research and our own conversations with employers. In similar contexts, workseekers with credible signals of their skills or past performance are more likely to be employed, and firms are willing to pay for information about applicants' skills (Abel et al., 2020; Carranza et al., 2022). Employers in our study—those using our research partner's job matching platform—face similar challenges. Based on data from the platform, only 66% of observed vacancies make a hire, despite large applicant pools (median 83 applicants per vacancy). Employers on the platform have also directly communicated that certain worker skills and attributes, like soft skills, are crucial for many roles but difficult to observe—a major impetus for our study.

3.2.2 Research Partner

To implement our experiment, we partnered with the largest youth unemployment services provider in South Africa. This non-profit organization works closely with the South

African government to design and implement programs aimed at reducing unemployment through job search support and partnerships with employers, primarily targeting disadvantaged workseekers aged 18 to 35. In particular, they maintain a large online job search and matching platform that provides an ideal setting for our study.

Our partner's online job search and matching platform captures a large and dense swath of the South African labor market. Given the organization's focus on youth unemployment, usage of the platform is generally limited to workseekers under the age of 35 (as verified through a South African ID number or similar documentation upon sign-up) from disadvantaged backgrounds. Within these parameters, the site has over four million users, over a quarter of whom were active on the platform in the past year. Importantly, the site is "zero-rated," meaning that interacting with it through a cellular network does not incur cell data charges (a major expense for many South Africans), boosting utilization compared to comparable for-profit job search websites. As the majority of these workseekers have little or no formal work experience, the platform primarily aims to match users with entry-level jobs. The platform directly partners with over 2500 employers spanning the private and public sectors, and also aggregates job opportunities from other job search websites, such that in practice the site contains the near-universe of entry-level formal job openings in South Africa at any given time.

For employers that directly post job vacancies on the platform, applicants are ranked and ordered according to an algorithm (described in more detail in Section 3.3) that attempts to capture their fit for the job based on their education, work history, and geographic proximity to the job. These rankings are automatically provided for all opportunities created on the platform, and employers can view and use them at their discretion. However, some employers choose not to interact directly with the platform, but instead rely on our research partner to manage their vacancy posting and provide them with a ranked shortlist of applicants. We use this subset for our experiment as we can exercise full control over the applicant shortlist (and ranking of applicants on each shortlist) from which employers hire for these vacancies. As shown in Section 3.4.2, these vacancies are broadly distributed

geographically and in terms of industry. That said, the sample is concentrated in the urban centers of Johannesburg and Capetown, and in the retail, insurance, and other business services sectors, reflecting patterns in the full set of vacancies on the platform (and in the South African labor market more broadly) (Africa, 2023).

3.3 Experimental Design

To test the impact of including information about soft skills in hiring, we partner with a South Africa-based non-profit social enterprise (described in more detail in Section 3.2.2). Our research partner operates an online job search and matching platform that helps match young workseekers to entry-level work and training opportunities. In particular, the platform uses an algorithm based on hiring criteria such as education, work experience, and distance from the posted job to rank applicants to each vacancy. Employers then select candidates to interview and hire based on this ranked list. Using a sample of job postings from the platform's partner firms, we explore the effects of including information about job applicants' soft skills in the candidate ranking algorithm.

The sections below describe the control and treatment ranking algorithms, our randomization protocol and experiment timing, and how we measure soft skills.

3.3.1 Ranking Algorithm

The platform's current ranking algorithm serves as the control in our experiment. Our alternative ranking algorithm, which incorporates a measure of applicants' soft skills *in addition to* the criteria used in the current ranking algorithm, is the treatment. We also tested, on a smaller subset of vacancies, a second alternative ranking algorithm that *replaced* some of the criteria in the current ranking algorithm with the soft skills measure, though this subset is not included in the primary results.[2]

For each applicant to a given vacancy, the ranking algorithm first calculates an inter-

[2] We ultimately dropped this second treatment group in order to maximize statistical power, due to sample size constraints. However, this smaller "soft skills only" treatment sample allows us to directionally explore non-linearities in the *intensity* of the emphasis placed on soft skills (relative to other criteria) in the applicant ranking algorithm.

mediate score that is based on the applicant's formal education and work history as well as the job type. In particular, the platform sorts vacancies into one of four job "quadrants," and ranking score formulas vary slightly between quadrants (mostly in order to capture the relevance of the applicant's past work experience to the job type in question). The four quadrants reflect the categorical product of (customer-facing or non-customer-facing) x (low-complexity or high-complexity), such that each quadrant captures the relative degree of customer interaction required as well as the relative complexity of the role. For example, a job as a server at a restaurant would be sorted into the "customer-facing low-complexity" quadrant, while a dishwasher job at the same restaurant would fall into the "non-customer-facing low-complexity" quadrant. Once the intermediate score has been determined, the final rank is calculated by adjusting this intermediate score to (very slightly) boost inclusivity and to prioritize applicants who are located nearer to the job.

The intermediate (excluding geography) ranking score used for applicants to *control group* vacancies (the current default on the platform) is based on their education and formal work experience relative to the job quadrant and is calculated as follows:

$$\begin{aligned}\text{Rank}_{\text{Control}} = \max\Big(&[\beta_1 \text{Quintile} + \beta_2 \text{Maths Lit} + \beta_3(\text{Pure Maths} \\ &+ \beta_4 \text{Has Pure Maths}) + \beta_5 \text{Has Degree} + \beta_6 \text{Has Diploma} \\ &+ \beta_7 \text{Has Exp. in Quadrant} \\ &+ \beta_8 \text{Has Exp. in Other Quadrant}]/C_1, C_2\Big)\end{aligned}$$

where "Quintile" captures the quintile rank of the secondary school the applicant attended (workseekers from lower-ranked schools get a boost), "Maths Lit," "Pure Maths," and "Has Pure Maths" capture the mathematics track the applicant chose in secondary school and their score on the relevant test[3], "Has Degree" captures whether the applicant obtained their matric degree (roughly equivalent to a high school diploma in the U.S.), "Has Diploma" captures whether the applicant has obtained any post-secondary diploma, and "Has Exp.

[3] Specifically, "Maths Literacy" and "Pure Maths" are the two core math options South African secondary school students may choose from. "Pure Maths" is generally considered more technical, so applicants who chose this track get a boost.

in [Other] Quadrant" captures whether the applicant has any formal work experience in a job in the same [any other] quadrant as the job in question. The coefficients, as well as constants C_1 and C_2, are quadrant-specific scaling parameters. The maximization ensures that all intermediate rank scores are bound below by a constant floor (C_2), accounting for, e.g., missing data in some score components.

The intermediate ranking score for applicants to *treatment group* vacancies is then given by:

$$\text{Rank}_{\text{Treat}} = \text{Rank}_{\text{Control}} * \text{Soft Skill Score}$$

where the Soft Skill Score is a standardized sum of scores for the three soft skill measures used, as described in Section 3.3.3 below. In other words, the applicant's soft skills score is incorporated multiplicatively *on top of* the baseline rank score.

Then the final score used to rank applicants, depending on treatment assignment T, incorporates a small inclusivity boost as well as geographic proximity to the job, and is given by:

$$\text{Final Rank}_T = \text{Inclusivity} * \text{Geo} * \text{Rank}_T$$

where "Inclusivity" reflects a small boost of 1.01 for women applicants, while "Geo" is based on the distance between the worker's reported address and the job location, calculated using a bounded geometric decay function.

Once the Final Rank has been calculated, applicants are sorted and displayed according to their Final Rank score, with higher-scoring applicants appearing higher on the list.

3.3.2 Randomization and Timing

Overall, we randomized a sample of over 1,000 job vacancies on the platform. Assignment of vacancies to treatment and control groups was conducted in real time, based on pre-generated sequences of random assignments stratified within blocks. The timing of job postings depends on employer demand and was therefore not uniform or precisely predictable. The chain of events after a vacancy is randomized, up until an applicant is hired

and starts work, is laid out in Figure 3.1. The randomization and treatment stage of the experiment started in November 2022 and finished in the latter half of 2024. Outcome data collection (as described below and in more detail in Section 3.4.1) started in January 2023 and is planned to finish by the end of 2025.

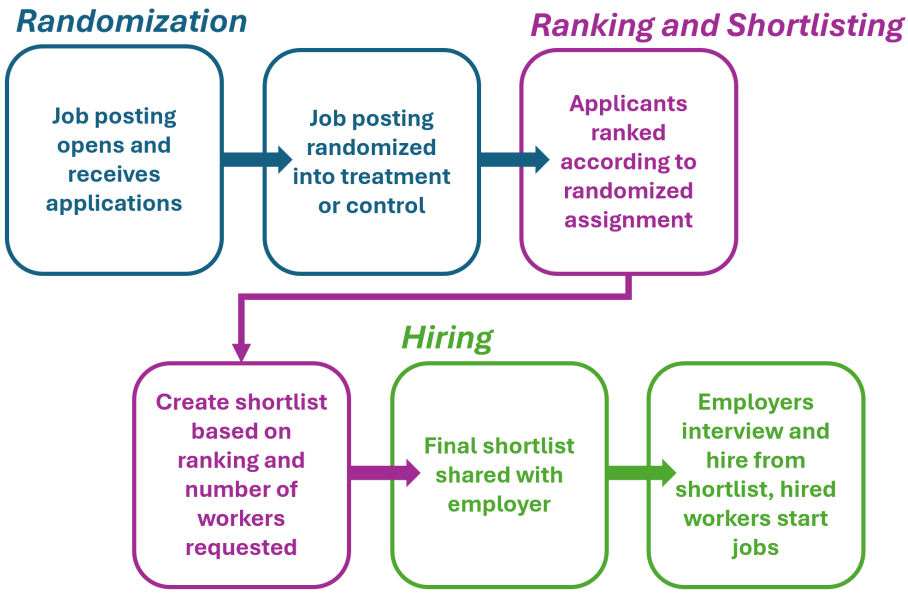

FIGURE 3.1: Experiment Chain of Events

The randomization and data collection process for each job posting is as follows:

1. Platform staff members share a pre-filtered applicant list with the research team. This list includes only the applicants who meet baseline requirement for the job as specified by the employer (e.g., located within 10 miles of the job).
2. The job posting is randomly assigned to either the control group or the treatment group.
3. The research team ranks the candidate list according to the algorithm specified by the assignment from step 2.
   (a) Candidate lists for job postings assigned to the control group are ranked using the platform's traditional algorithm, which is based on education, work experience, and distance from the job.
   (b) Candidate lists for job postings assigned to the treatment group are ranked using

an algorithm that is based on the criteria listed above as well as a measure of the candidate's soft skills.

4. The research team shortens the ranked list to either the top 20 applicants by rank or the top n*4 applicants by rank (whichever is greater), where n is the target number of hires listed by the employer (e.g., n would be equal to four if a retailer posted a job ad looking to hire four cashiers).
5. This ranked shortlist is sent to the employer, with a note at the top of the list explaining that the applicants are ranked according to their expected fit with the job.
6. After the shortlist is shared with the employer, we conduct a short-term survey with the shortlisted applicants, as well as applicants who would have been shortlisted, to measure interviews, offers, and employment at the job for which they applied, as well as employment, wages, promotions, job satisfaction, turnover, and job search more broadly.
7. We also survey the employer about their hiring decisions for the job they posted, as well as post-hiring outcomes for any hired applicants, subsequent hiring, and branch-level performance.
8. Finally, we conduct a second round of follow-up interviews with the applicants surveyed in step 6 to measure longer-term labor market outcomes.

The results presented in this paper primarily utilize data from our short-term workseeker follow-up survey, as outlined in Step 6 above and in more detail in Section 3.4.1. Data collection for the firm surveys and long-term workseeker follow-up surveys is still in progress.

3.3.3 Measuring Soft Skills

we construct a measure of workers' soft skills using a weighted average of Behavioural Activation (Gable et al., 2000), Grit (Duckworth & Quinn, 2009), and Growth Mindset (Dweck, 2006) scales. Motivated by existing research in psychology and behavioral economics, these scales are intended to capture initiative, tenacity, and willingness to learn over time—traits that are valued by employers but difficult to observe. The Behavioural

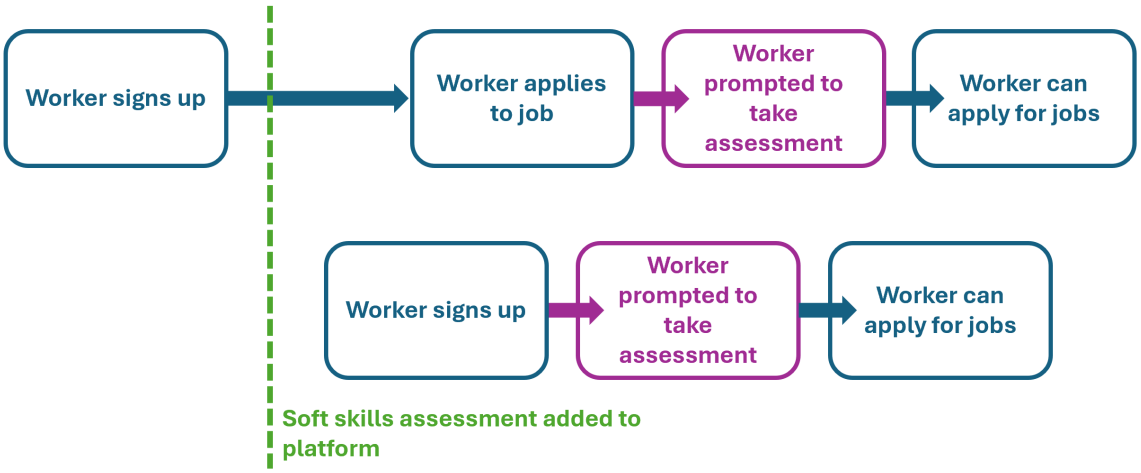

FIGURE 3.2: Soft Skills Assessment Timeline

Activation scale assesses intiative by gauging respondents' degree of alignment with statements such as "In the last week, I was active and achieved the goals I set out to do" (Gable et al., 2000). The Grit scale assesses tenacity similarly, with statements such as "I have overcome a setback to achieve an important goal" (Duckworth & Quinn, 2009). The Growth Mindset scale assesses the degree to which respondents feel they can exercise control over their own development and learning, using agreement with statements such as "No matter who you are, you can change your skills a lot" (Dweck, 2006). We construct the aggregate soft skill score used in the ranking algorithm by first summing the standardized scales and then rescaling the sum to be between 0 and 1 (using the distribution of scores among all workseekers on the platform).

These scales are self-administered on the platform and must be completed by applicants in order to apply to the job postings included in the study. Workseekers already registered on the platform were prompted to complete the assessment upon their first job application after the experiment start date, while workseekers joining the platform after the experiment start were prompted to take the assessment upon registration, as shown in Figure 3.2.

These soft skill assessments as originally developed were administered in-person by our research partner during job search workshops. Before moving forward with the self-administered, online version of the measure, we conducted a pilot study to confirm that

the assessments remained valid in this format. Our pilot found that the online version maintained strong psychometric properties, as shown in Appendix D. Further, our pilot found no evidence of workseekers "gaming" the assessments or giving deliberately untruthful answers, even under conditions emphasizing the potential importance of the assessments (Appendix D).

Table 3.1: Properties of Soft Skill Measures in the Full Sample

	Correlation coefficients			Cronbach's α
	Grit	GM	BA	
Grit	1	0.38	0.53	0.73
Growth Mindset		1	0.27	0.74
Behavioral Activation			1	0.87

Examining the properties of the soft skill measures in our final sample increases our confidence in their validity and high signal content. First of all, measures for each of the scales are highly (but not perfectly) correlated, as shown in Table 3.1, suggesting distinct but related concepts as we should expect from prior research. Second, the distribution of the final aggregate soft skills score, as used in the treatment ranking algorithm, displays a good deal of variation and a roughly normal (right-truncated) distribution (Figure 3.3). This again suggests that workseekers are not responding disingenuously to inflate their scores. Finally, it appears that our soft skills measure is meaningfully predictive of behavior and resulting outcomes. In the full sample, the aggregate soft skills score predicts planned search and work outcomes in the *control* group (i.e., workseekers who applied to vacancies that did *not* use soft skills in the ranking algorithm), as shown in Table 3.2.

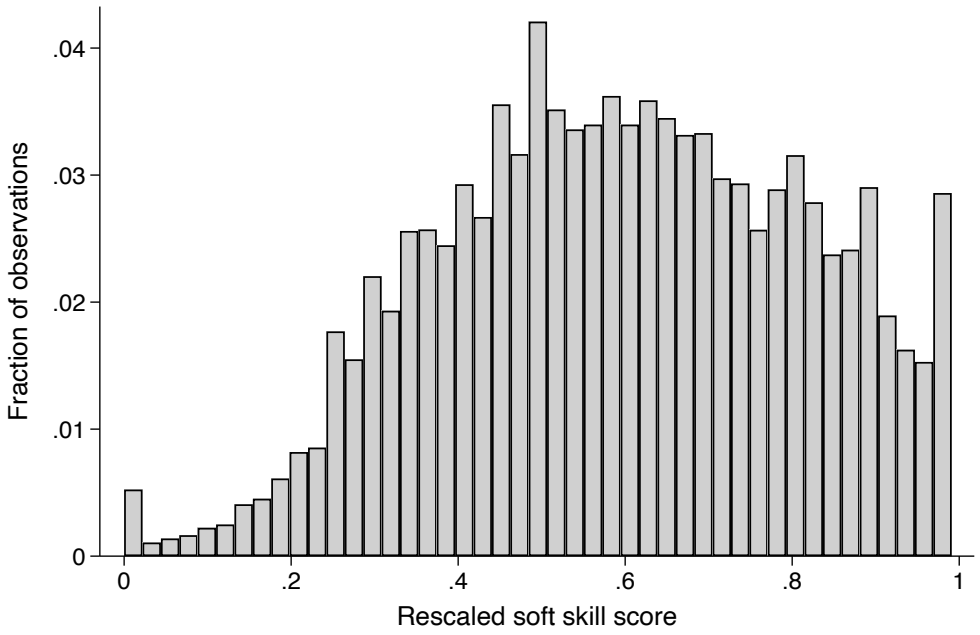

FIGURE 3.3: Distribution of Aggregate Soft Skill Measure

Table 3.2: Soft Skills Predict Planned Search and Work Outcomes in the Control Group

	Planned search			Work		
	(1) Hours	(2) Expend	(3) Apps	(4) Any	(5) Hours	(6) Earnings
Soft skill score (z)	3.455***	24.070***	3.735***	0.014***	0.636***	195.027***
	(0.507)	(6.583)	(0.293)	(0.005)	(0.177)	(45.319)
Any work experience	-1.854*	-26.383*	1.101*	0.008	-0.419	13.148
	(1.010)	(13.511)	(0.584)	(0.010)	(0.359)	(88.005)
Math score (z)	-0.411	-23.578***	-0.414	-0.009*	-0.367**	-61.529
	(0.496)	(6.463)	(0.293)	(0.005)	(0.176)	(45.158)
English score (z)	0.781	-0.351	1.857***	0.010**	0.525***	90.289*
	(0.505)	(6.544)	(0.273)	(0.005)	(0.186)	(50.497)
Control mean	46.772	571.251	31.127	0.394	9.102	2012.282
Observations	9903	9886	9879	10100	10088	10061

3.4 Data and Sample

3.4.1 Data Sources

Our primary data source for the analyses presented in this paper is short-term follow-up surveys with jobseekers who applied to the job listings in our sample. We supplement this with backend data from our partner's online job matching platform.

Our short-term follow-up surveys are administered within six months following randomization and shortlisting (the process illustrated in Figure 3.1), and collect data on labor market outcomes, beliefs, and job search. After the ranked shortlist of applicants is shared with the employer, we conduct a survey with the shortlisted applicants, as well as "counterfactual" applicants (those who would have been shortlisted under the alternate treatment assignment). The survey asks applicants about interviews, offers, and employment at the job for which they applied, as well as employment, wages, promotions, job satisfaction, turnover, and job search more broadly.

The specific outcomes we collect in this survey are as follows. In terms of hiring outcomes for the sample job listing, we ask whether the applicant was invited to interview, whether they attended interview or not (and if not why not), whether they were offered the job or not, whether they accepted job or not if offered (and if not why not), whether they started the job yet or not, and the start date of employment. In terms of job attributes (if hired), we ask about the type of employment agreement (informal, formal with verbal contract, formal with written contract), hours and days worked per week, and compensation levels and types (e.g., salary vs. commission). If they are hired, we also ask about attrition and tenure, including whether they are still in the job (if they have started work), why they left (if they left), how long they anticipate remaining at the job, and their likelihood of quitting in the next six months (if they haven't left). In terms of job and match quality, we ask workers for their reservation wage for the surveyed job and (if hired) a measure of worker surplus. In terms of job search, we ask all respondents how much time and money they are spending on job search, how many applications they are submitting, and their beliefs

about their job-finding rate conditional on their self-reported job search behavior. Finally, we also collect data on employment and earnings in out-of-sample jobs.

While survey data always carries some risk of response bias, our experimental design ensures that there is minimal risk of response bias in our main outcome (employment) that is correlated with treatment. Since the jobseekers in our sample do not know the treatment status of the job they applied to, response bias or over-reporting of job offers or hires for treated job listings due to, e.g., experimenter demand effects is unlikely. Further, the surveys are administered either fully online or over the phone by field officers who do not know the treatment status of the surveyed job opportunity, minimizing the risk of surveyor-induced bias.

We supplement this survey data with platform data from our research partner's job search and matching website. We are able to collect from the website, for all registered users, standard demographic variables, geographic data, education history, standardized test scores, and formal work history. We are also able to observe users' on-platform job search behavior, including all applications they start and submit. Our research partner also collects hiring data from the employers on their website directly. While this hiring data is not suitable as a primary data source due to low response rate, we will use it to verify our survey data.

Finally, there are several data sources we are currently in the progress of collecting and incorporating into our analysis, which will be included in future work. In particular, we are in the process of surveying the employers in our sample about their hiring decisions for the job they posted, as well as post-hiring outcomes for any hired applicants, subsequent labor demand and hiring, and branch-level performance. We are also currently conducting longer-term follow-up surveys with workseekers (roughly a year to 18 months post-treatment) in order to collect data on attrition, post-hiring outcomes, and career and earnings trajectories. These longer-term outcomes will allow us to say more about the potential changes in match quality and labor demand induced by our treatment, as well as potential longer-term spillover effects and career ladder effects for treated and counterfactual jobseekers.

3.4.2 Sample

The sample used in our main analysis is made up of roughly 1,000 job listings from our partner's platform. Some of the employers who post job listings choose not to interact directly with the platform, but instead rely on our research partner to manage their listings and provide them with a ranked shortlist of applicants. We use job listings from this subset of employers for our experiment as we are able to fully control the shortlists, and the ranking of applicants on these shortlists, used in the hiring process. Employers generally do not schedule job listings far in advance, but submit them as needed to fill vacancies. As such, we add job listings to our sample in real time as they are posted to the platform, randomizing them into the control or treatment group using pre-specified block randomization. Our randomization ran from roughly the start of 2024 to the end of 2025.

Table 3.3: Job Listing and Applicant Sample

	(1) Mean	(2) 5^{th} perc.	(3) 25^{th} perc.	(4) Median	(5) 75^{th} perc.	(6) 95^{th} perc.
Panel A: Opportunity characteristics						
Number of positions	7.68	1.00	1.00	3.00	10.00	27.00
Number of applicants	858.50	36.00	156.00	363.50	838.00	2492.00
Number of shortlisted applicants	29.47	10.00	18.00	20.00	24.00	75.00
Panel B: Applicant characteristics						
Average age	26.10	24.20	25.39	26.08	26.80	28.11
Fraction female	0.70	0.35	0.65	0.76	0.81	0.90
Fraction urban	0.89	0.43	0.89	0.98	1.00	1.00
Fraction any experience	0.60	0.25	0.44	0.60	0.75	0.93
English matric score	0.60	0.52	0.57	0.60	0.63	0.69
Math matric score	0.50	0.41	0.46	0.50	0.54	0.59
Soft skill score (0 to 1)	0.51	0.42	0.47	0.50	0.55	0.64

Thus, our unit of randomization, and the unit of analysis in our main hiring results, is at the job listing level. Our sample of jobseekers is a subset of the applicants to these job listings. In particular, we survey the jobseekers who are shortlisted for the job, as well as the "counterfactual" sample of jobseekers who would have been shortlisted if the job listing had been assigned to the alternative treatment group. We typically attempt to survey all shortlisted and counterfactual applicants, except in cases in which the shortlists are prohibitively long (e.g., because the employer specifically requested a very large shortlist),

in which we case we interview a semi-random subset in which higher-ranked applicants (those who are, all else equal, more likely to get the job) are more likely to be sampled. Job listings may cover multiple vacancies; for example, a job listing for a waiter may be looking to fill three waiter positions. Shortlists are then limited to either the top 20 applicants by rank or the top $n * 4$ applicants by rank (whichever is greater), where n is the target number of hires listed by the employer (e.g., n would be equal to three in the prior example), *except* in cases where employers explicitly ask for additional candidates (though this is rare).

Table 3.4: Balance

	(1) Control	(2) Treatment	(3) Δ	(4) p(Δ = 0)	(5) Obs.
Panel A: Opportunity characteristics					
Number of positions	7.30	8.05	0.75	0.51	985
Number of applicants	806.14	910.44	104.31	0.45	986
Number of shortlisted applicants	27.85	31.08	3.23	0.15	986
Panel B: Applicant characteristics					
Average age	26.14	26.06	-0.07	0.35	986
Fraction female	0.71	0.70	-0.01	0.56	986
Fraction urban	0.89	0.89	-0.01	0.68	986
Fraction any experience	0.60	0.60	0.00	0.74	986
English matric score	0.60	0.60	-0.00	0.75	986
Math matric score	0.50	0.50	0.01	0.07	986
Soft skill score (0 to 1)	0.51	0.51	-0.00	0.81	986

As shown in Table 3.3, the average job listing receives a far greater number of applicants (roughly 850 on average) than are included on the shortlist (roughly 30 on average). Thus, shortlisting is an extremely important driver of hiring outcomes. Applicants to the jobs in the sample tend to be young (about 26 years of age on average), located in urban areas, and skew female (70% of applicants are women on average). This reflects the demographics of the jobseekers on our partner's platform. As shown in Table 3.4, job listings and applicants are well-balanced between treatment and control groups. There are no significant differences in observable variables, other than a marginally significant difference in standardized math test scores, which is quite small in absolute terms.

3.5 Results

This section provides the results of our main analyses, which can be divided into two broad research questions. First, we seek to understand the impact of including soft skills in applicant ranking, compared to ranking using only traditional hiring criteria (education and formal work experience). Specifically, we examine hiring outcomes at the job posting level to assess how many and what types of workers are interviewed and hired, shedding light on whether our treatment is improving hiring process efficiency and match quality. Second, we use the random variation that our treatment introduces in which candidates are included on the shortlists sent to employers to investigate the impact of being shortlisted for a job on subsequent job search behavior, labor market beliefs, and labor market outcomes at the individual level. This second analysis assesses the potential spillover effects of our intervention, as well as other labor market interventions that impact the probability of shortlisting, interviewing, or hiring particular candidates relative to others.

Our main results are promising, largely suggesting that our treatment is effective and produces limited negative spillovers. Our primary hiring results show that treatment has a positive impact at all stages of the hiring funnel: employers interview, extend offers to, and ultimately employ more applicants from treated opportunities, on average. Meanwhile, our analysis of "counterfactual" applicants (those who would have been shortlisted under the alternative ranking algorithm to the one actually applied to each vacancy) does not find any evidence of lasting disadvantages for these jobseekers; after controlling for the direct employment effect of not being shortlisted for a specific opportunity in our sample, we do not observe a lower probability of employment, nor do we observe discouragement effects.

3.5.1 Primary Hiring Results

Our results suggest that treatment boosts hiring by helping employers find more suitable candidates. We estimate the effect of incorporating soft skills into the candidate ranking algorithm on hiring outcomes at the job posting level. For these outcomes, we estimate average treatment effects by comparing mean outcomes between postings assigned to treatment

and postings assigned to control, adjusting for randomization block fixed effects and covariates selected using Post Double Selection Lasso (Ahrens et al., 2019), using the following specification:

$$y_j = \alpha_0 + \alpha_1 T_j + \gamma X_j + \zeta_j + \epsilon_j \tag{3.1}$$

where y_j is the outcome for job posting j (e.g., number of hires), T_j captures the treatment assignment for j (1 if the job posting was assigned to the treatment group, 0 if control), X_j is a vector of covariates selected using Post Double Selection Lasso, and $zeta_j$ is randomization block fixed effects.

Note that these comparisons capture average treatment effects of a change in the *shortlisting criteria*. They do not directly capture average treatment effects of a change in the *shortlist composition*, in the *hiring criteria*, or (for post-hiring outcomes specifically) in the *composition of the hired applicants*.

We first confirm that the treatment alters the applicant pool considered by employers, and thus ultimately hiring outcomes. In line with this expectation, treatment increases the average soft skills of shortlisted applicants by about 0.45 standard deviations and of hired applicants by 0.27 standard deviations, as shown in Tables 3.5 and 3.6, respectively. The treatment also raises the share of less experienced workers, as shortlisted applicants in the treatment group are about 5.5% less likely to have any formal work experience, and hired applicants are about 6.4% less likely, than the control group average (though noting that the result for hired applicants is not significant at conventional levels).

Our main hiring results, as shown in Table 3.7, indicate that our treatment increases hiring at all stages of the recruiting pipeline. In particular, treated job openings yield significantly more interview invitations (about one additional interview on average compared to a control mean of about 4.5) and more actual employment (about 0.5 additional workers accepting and starting jobs compared to a control mean of about 1) than control group job openings.

Table 3.5: First Stage: Randomization Affects Skill Composition of Shortlisted Applicants

	Soft skill scores (z)				Traditional criteria		
	(1) Aggregate	(2) grit	(3) BA	(4) GM	(5) Any exp.	(6) English (z)	(7) Math (z)
Treatment	0.438***	0.357***	0.347***	0.324***	-0.055***	0.044**	0.019
	(0.019)	(0.016)	(0.015)	(0.016)	(0.009)	(0.018)	(0.019)
Control mean	-0.256	0.260	0.209	0.279	0.645	-0.060	0.028
Observations	984	984	984	984	984	984	984
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Randomization block FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 3.6: First Stage: Randomization Affects Skill Composition of Hired Applicants

	Soft skill scores (z)				Traditional criteria		
	(1) Aggregate	(2) grit	(3) BA	(4) GM	(5) Any exp.	(6) English (z)	(7) Math (z)
Treatment	0.265***	0.218***	0.068	0.386***	-0.064	0.068	-0.101
	(0.075)	(0.079)	(0.088)	(0.091)	(0.039)	(0.087)	(0.082)
Control mean	-0.209	0.258	0.378	0.223	0.620	0.041	0.049
Observations	468	468	468	468	469	469	469
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Randomization block FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 3.7: Hiring Results

	Interview		Job offer		
	(1) invite	(2) attended	(3) received	(4) accepted	(5) started
Treatment	0.750**	0.605**	0.383*	0.303	0.338**
	(0.310)	(0.267)	(0.220)	(0.197)	(0.162)
Control mean	4.008	3.071	1.758	1.387	0.894
Observations	984	984	984	984	984
Controls	Yes	Yes	Yes	Yes	Yes
Randomization block FE	Yes	Yes	Yes	Yes	Yes

Looking at the preliminary composition of hires to treated job postings in Table 3.8, we see that treatment particularly boosts hiring for less experienced, and to some extent younger, workers. Thus, incorporating soft skills not only increases the total number of hires but may also help less experienced applicants gain a foothold in the labor market.

Table 3.8: Demographics of Hired Applicants

	# hires					
	(1) female	(2) male	(3) exp.	(4) inexp.	(5) young	(6) old
Treatment	0.199	0.112*	0.131	0.169	0.182*	0.131
	(0.155)	(0.064)	(0.121)	(0.104)	(0.104)	(0.110)
Relative effect size	0.211	0.254*	0.158	0.302	0.250*	0.205
Control mean	0.945	0.440	0.827	0.560	0.729	0.640
Observations	984	984	984	984	984	984
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Randomization block FE	Yes	Yes	Yes	Yes	Yes	Yes

3.5.2 Counterfactual Results

We now explore potential spillover effects from our intervention through a "counterfactual" analysis. In particular, this analysis captures the effect of being shortlisted for a job posting, or not, on job search behavior and labor market beliefs and outcomes at the individual level. For these outcomes, we estimate average treatment effects by comparing mean outcomes between individuals who were shortlisted for a job posting and individuals who *would have been* shortlisted for the job posting if it had been assigned to a different treatment—whom we refer to as the "counterfactual" sample. To give a simplified example, think of comparing endline employment for a jobseeker $J1$ who's not shortlisted for vacancy T because it's treated, but would be (*counterfactually*) shortlisted if it were untreated, to employment for jobseeker $J2$, who's shortlisted for vacancy C because it's untreated, but would not be shortlisted if it were treated. Jobseeeker $J1$ represents the "counterfactual" jobseeker in this case. This can happen if jobseekers J1 and J2 both have work experience but lower soft skills, so they're recommended by the traditional algorithm that values experience but not by the new algorithm that values soft skills. We limit the sample to actually and counterfactually shortlisted applicants because shortlisting under both ranking schemes depends on variables that are likely to correlate with labor market outcomes, such as education and work experience.

We estimate these counterfactual results using the following specification:

$$y_{ij} = \beta_0 + \beta_1 S_{ij} + \gamma X_{ij} + \phi_{ij} + \zeta_j + \epsilon_{ij} \tag{3.2}$$

where y_{ij} is the outcome for applicant i for job listing j (e.g., employment status), S_{ij} captures shortlisting status (1 if *not* shortlisted, 0 if shortlisted), X_{ij} is a vector of covariates selected using Post Double Selection Lasso, ϕ_{ij} is the inverse odds ratio of being shortlisted for applicant i for job listing j, and $zeta_j$ contains randomization block fixed effects. We adjust for ϕ_{ij}—the probability that each applicant is shortlisted for a given job posting—because this probability is not constant across applicants, and is positively

correlated with skills.[4] Adjusting for this inverse odds ratio accounts for the variation in treatment assignment probabilities for the same reasons that adjusting for the probability of selection into a sample accounts for sample selection bias (J. Heckman, 1974).

Examining outcomes for counterfactual applicants at the individual level, we find that these applicants are less likely to be hired and have lower average wages than applicants who were actually shortlisted, as shown in Table 3.9. In particular, counterfactually shortlisted applicants are about 3 percentage points less likely to be employed at endline than actually shortlisted applicants, and this effect appears to be due completely to differences in wage employment (compared to negligible differences in self-employment levels). This is a sizable effect given low overall and wage employment probabilities of about 42% and 19%, respectively, in the control group. Further, counterfactual applicants are working fewer hours and earning less at endline, and report lower reservation wages and wage expectations.

Table 3.9: Counterfactual Analysis

	Work					Expectations	
	(1) Any	(2) Wage	(3) Self-employed	(4) Hours	(5) Earnings	(6) Res. wage	(7) Wage exp.
Not shortlisted	-0.031***	-0.034***	-0.003	-0.796**	-179.819**	-210.454***	-139.249**
	(0.010)	(0.008)	(0.007)	(0.373)	(91.683)	(76.276)	(69.592)
Control mean	0.429	0.194	0.148	9.407	2104.994	6327.021	6762.794
Observations	9988	9988	9988	9829	9829	9971	9965
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Opportunity FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

That said, these differences appear to reflect only *direct* effects of our treatment, rather than *indirect* spillover effects. Specifically, these results are driven almost completely by shortlisted applicants who were actually hired, as shown in Table **??**. After removing applicants who were actually hired to one of the randomized vacancies, we observe no statistically significant differences between counterfactually and actually shortlisted applicants in terms of employment probability, hours worked, or earnings at endline. In other words, while our intervention changes who is hired to treated jobs (as expected), it does not appear to

[4] Specifically, because we introduced a third treatment arm at some points in the study, some applicants have a 2/3 probability of appearing on a shortlist while others have only a 1/3 probability, for job postings randomized under our three-arm protocol.

have any impact on employment or wages in jobs outside the experiment. This implies that the direct displacement effect is approximately one-for-one, while overall hiring increases as shown in Section 3.5.1. Taken together, these results suggest that the net effect of our intervention is a positive expansion in employment.

Table 3.10: Counterfactual Analysis Excluding Hired Applicants

	Work					Expectations	
	(1) Any	(2) Wage	(3) Self-employed	(4) Hours	(5) Earnings	(6) Res. wage	(7) Wage exp.
Not shortlisted	-0.010	-0.005	-0.009	-0.164	-91.174	-159.129**	-109.077
	(0.011)	(0.008)	(0.008)	(0.372)	(95.343)	(78.242)	(70.901)
Control mean	0.403	0.158	0.154	8.730	2008.420	6282.655	6736.191
Observations	9551	9551	9551	9446	9413	9558	9552
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Opportunity FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

3.6 Conclusion

This paper presents novel evidence on the effects of including information about applicants' soft skills in hiring. Including a soft skills measure in an applicant recommendation algorithm on a large online job search and matching platform appears to boost hiring throughout the funnel, resulting in more interviews, job offers, and hires. Soft skills are included in addition to the more traditional criteria—such as educational credentials and formal work experience—used in the control algorithm, and there is some indication that this expansion results in more inclusive hiring, helping younger and less experienced workers break into the formal labor market. These results suggest that our intervention helps address the information frictions that are pervasive in developing-country labor markets.

Additionally, this paper provides, to our knowledge, the first evidence on spillover effects of an intervention that shifts hiring criteria in a developing-country setting. While our intervention mechanically lowers the probability of being hired into treated jobs for "counterfactually" shortlisted applicants, we do not observe indirect spillover effects for these jobseekers, either in terms of discouragement and subsequent job search behavior or in terms of endline employment probability. Thus, given that our intervention expands hiring overall, we anticipate that it would generate net positive employment gains in equilibrium.

Future work will explore the longer-term effects of our treatment using later jobseeker and firm follow-up surveys that are currently in progress. In particular, we hope to shed further light on the extent to which our intervention improves match quality by examining outcomes over a longer time horizon such as attrition, job performance and promotions, and firm performance and labor demand. We also plan to study the longer-term impacts on treated and counterfactual jobseekers' career trajectories, in order to better understand potential career ladder and general equilibrium effects.

4. The Relationship Between the "Norm to Work" and Labor Market Outcomes

4.1 Introduction

The past decade has seen growing interest among economists in the role of social norms in shaping labor market outcomes.[1] While much of this work has focused on gender norms, recent labor market upheavals in the wake of COVID-19 have drawn attention to broader social attitudes toward work. Many commentators have suggested that a pandemic-induced reassessment of work's value contributed to increased job quitting, higher wages, and a resurgence in worker power (e.g., Hirsch, 2021; Krugman, 2021; Sull et al., 2022). Although there is some evidence of shifting worker preferences and increased worker power in the form of unionization, this narrative has proven difficult to test empirically, in large part because social norms are so challenging to quantify (Faberman et al., 2022; Kinder and Stateler, 2022). The big-picture question remains: To what extent does the "norm to work"—i.e., the social norm that working-age adults should be employed—influence worker behavior and, ultimately, macro-level outcomes such as employment and wages?

To investigate this question, I construct novel measures of the NTW and related social norms using large-scale text data from social media. Specifically, I employ text analysis and machine learning methods to classify unstructured text from Twitter. By collecting a large dataset of tweets (roughly 20 million) related to work attitudes, labeling them through a supervised classification model, and aggregating results at the state-year and city-year level, I generate a geographically granular measure of work norms over time. This approach not only enables real-time tracking of work norms but also presents a scalable framework for measuring other social norms using similar methods.

To contextualize the mechanisms through which the NTW may influence labor market outcomes, I draw on identity economics and reference-dependent labor supply models

[1] Recent examples of papers in this area include, but are not limited to: Bertrand et al., 2016; Binzel and Carvalho, 2017; Breza et al., 2019; Bursztyn et al., 2017; Bursztyn, González, and Yanagizawa-Drott, 2020; Dube et al., 2022, Fortin, 2015; Jayachandran, 2021.

(Farber, 2008; Kranton, 2016; Pissarides, 2000). These frameworks suggest that workers' employment decisions are shaped by social comparisons and that stronger work norms—likely accompanied by greater unemployment stigma—should influence labor supply behavior much like unemployment benefit generosity (Jäger et al., 2020). Theoretical predictions suggest that labor force participation and employment should be positively associated with the NTW, while wages could be negatively affected due to increased labor supply.

Given the complexities introduced by the COVID-19 pandemic, this study focuses on the pre-pandemic period (2011-2019) to isolate underlying relationships between work norms and labor market outcomes. Future extensions of this work will explore the COVID pandemic as a potential exogenous shock to work norms, as well as the subsequent labor market effects of this shock.

This paper makes two primary contributions. First, it provides empirical evidence on the extent to which the NTW influences labor supply and employment outcomes, offering new insights into the broader labor market implications of work norms. Second, it introduces a novel methodology for constructing high-frequency, geographically specific measures of social norms using unstructured text data. By demonstrating the feasibility of this approach, the paper contributes to the growing literature on social norms in economics and presents a framework for future applications in labor market research.

Existing research has established that social norms are powerful drivers of behavior (e.g., Coleman, 2018; Elster, 1989; McAdams, 1997), particularly in the labor market. Prior work has shown that unemployment carries substantial psychological costs even after accounting for financial hardship (Blanchflower and Oswald, 2004; Brand, 2015; Jahoda, 1988). These costs are influenced by the prevailing social norms in one's community, with evidence suggesting that stronger work norms increase the psychological cost of unemployment and increase job search effort among the unemployed (Chadi, 2014; Clark, 2003; Eugster et al., 2017; Stutzer and Lalive, 2004).

While prior research has documented the individual-level effects of work norms, primarily for unemployed individuals, little is known about their aggregate labor market im-

plications. A key limitation has been the difficulty of measuring the NTW. Most prior studies have relied on survey-based measures that are infrequent and geographically coarse, or have proxied work norms using local unemployment rates. Notable exceptions include Stutzer and Lalive, 2004 and Eugster et al., 2017, who leverage natural experiments to infer employment norms. However, such approaches are highly context-specific. This paper addresses these gaps by constructing a scalable, real-time measure of work norms that can be applied across different contexts.

This study presents preliminary evidence that work norms shape labor supply decisions. The results suggest a positive association between the NTW and labor force participation and employment, consistent with theoretical expectations. While the causal direction remains an open question, future work will focus on distinguishing causality through instrumental variables, quasi-experimental designs, or plausibly exogenous shifts in work norms. These findings have significant implications for policy debates surrounding unemployment insurance, labor market participation incentives, and broader workforce engagement strategies.

The rest of the paper is organized as follows: Section 2 briefly presents motivating theoretical frameworks. Section 3 describes the data used in the measure construction workflow and labor market outcomes analysis, while section 4 discusses the measure construction approach in detail. Section 5 presents results of the main analysis exploring the relationship between labor market outcomes and social norms around work, while section 6 concludes.

4.2 Theoretical Frameworks

This section briefly discusses theoretical frameworks that are particularly helpful to fix thinking about the potential mechanisms through which the NTW may influence behavior and aggregate labor market outcomes. I also use these frameworks to generate simple predictions about the sign of expected effects. The resulting model sketches predict that labor supply will be increasing in the NTW, while wages will be decreasing.

Most notably, the framework of identity economics captures the idea that violations of

social norms associated with an agent's identity are costly (Akerlof and Kranton, 2010). This model rationalizes the psycho-social cost of unemployment, and predicts that this cost should vary with the strength of work norms in an agent's area, as existing empirical work has suggested.

Combining identity economics with a basic model of reference-dependent labor supply (Farber, 2008) yields the intuitive prediction that labor supply is weakly increasing in the strength of the NTW. In particular, we can formulate a simple model of worker utility as follows[2]:

$$U(c_t, h_t) = u(c_t, h_t) - I[h_t < H]\sigma$$

Where c_t is consumption, h_t is hours, H is the minimum number of hours an agent must work to be considered "employed," and σ is the strength of the NTW in the agent's community. We can see that the additive penalty for failing to work enough hours to reach the "employed" status is increasing in work norms, causing workers in areas with stronger work norms to supply more labor, all else equal. This in turn puts downward pressure on wages.

A simple wage bargaining framework (Pissarides, 2000), again combined with identity economics, also implies that wages are decreasing in the strength of the NTW. To the extent that being unemployed imposes an identity cost I, lowering a worker's outside option value (non-employment, or N) as follows:

$$\rho N = b + f\dot{(}E(w') - N) - I\sigma$$

Where the flow value (ρN) is equal to UI benefits (b), the potential "capital gain" into reemployment at wage w' ($E(w') - N$) at job-finding rate f, and the identity cost of unemployment ($I\sigma$), which is determined by the strength of the NTW (σ). A lower outside option value in turn lowers the Nash-bargained wage, which is a weighted average of productivity p and the outside option:

$$W = \sigma p + (1 - \sigma)\rho$$

[2] This is similar to the functional form used by Moffitt (1983) to model the effect of welfare stigma on worker behavior, though other specifications (without a discontinuity in the utility function) are possible in the event one would like to construct a fully estimable model.

The expected employment effect is ambiguous but could be estimated with a more complete model—a possible extension of the current project.

4.3 Data

4.3.1 Tweet Sample

The NTW measures in this paper are based on a sample of roughly 20 million tweets collected through Twitter's Academic API.[3] The sample was selected using keyword searches intended to filter for tweets relating to the norm to work and attitudes towards employment and unemployment[4]. The sample begins in 2011—the year when Twitter reached 100 million users—and ends in November 2022, when the tweets were collected (Meyer, 2019). After dropping tweets that were from outside the U.S. or missing geographic data (based on user-reported locations or geo-tags), the final sample consists of about 4.3 million tweets[5].

After collection, I pre-processed the text of each tweet according to standard text data analysis approaches (Grimmer et al., 2022). This included case normalization (converting all letters to lowercase), whitespace normalization (collapsing repeated spaces, tabs, and line-breaks into single spaces), punctuation and number handling (e.g., removing intra-word punctuation), and spelling correction (using the Python package Gingerit, which corrects spelling and grammar mistakes based on the context of complete sentences). The pre-processing also included normalization specific to the Twitter platform, such as stripping out links and replacing user tags with generic placeholders.

[3] The Academic API, along with all other affordable tiers of Twitter API access, was discontinued in 2022 following Elon Musk's acquisition of the platform.

[4] While some searches were based on single keywords (e.g., for "unemployment") many topic words (e.g., "work") were too vague to yield a reasonable proportion of relevant tweets. Therefore, I employed many multi-word searches intended to triangulate beliefs and attitudes towards work and employment (e.g., a keyphrase search for tweets containing ("work" OR "job") AND "lazy", meant to roughly capture sentiments along the lines of "people who don't work (or have a job) are lazy").

[5] The Twitter API does not permit filtering based on users' self-reported locations (the primary source of geographical data), so many of the collected tweets had to be dropped from the sample ex-post.

4.3.2 Labor Market Outcomes

Data for the individual-level labor market outcomes used in the analysis—labor force status, employment, unemployment duration, and wage data—come from the Census Bureau's Annual Social and Economic Supplement (ASEC) of the Current Population Survey for years 2011-2019. It should be noted that these data are comprised of repeated cross-sections, rather than a panel, so it is generally not possible to follow specific individuals over time. The ASEC contains data on respondent demographics, employment status, geography at the state and metro area level, and (importantly for this analysis) wage earnings over the past twelve months. The sample is restricted to civilian (non-military) adults aged 25-65. The sample used for studying wage earnings as an outcome is further restricted to individuals who are employed full-time over the measurement period.

4.4 Measuring the Norm to Work

The norms measures used in this paper are based on data from Twitter (as described in Section 4.3.1 above), labeled using a custom-trained text classification model. Figure 4.1 presents an outline of the measure construction workflow, which is described in more detail in the following sections.

This novel measure construction approach combines new access to large, unstructured text data sources with new technology for parsing large, unstructured text data. Using data from a large social media platform provides the necessary volume and density to measure norms with much more temporal and geographic specificity than is typically possible, as described in more detail in Section 4.4.1 below. At the same time, analyzing such a large volume of raw text data would be prohibitively time-consuming and expensive using most existing methods. In particular, even methods utilizing supervised machine learning for large-scale text analysis have historically required a large training dataset based on human labeling—a slow and expensive process. By completely automating the process of *both* generating training data (as described in Section 4.4.2) *as well as* generating discrete labels for each unit of text in the sample based on this training data (as described in Section

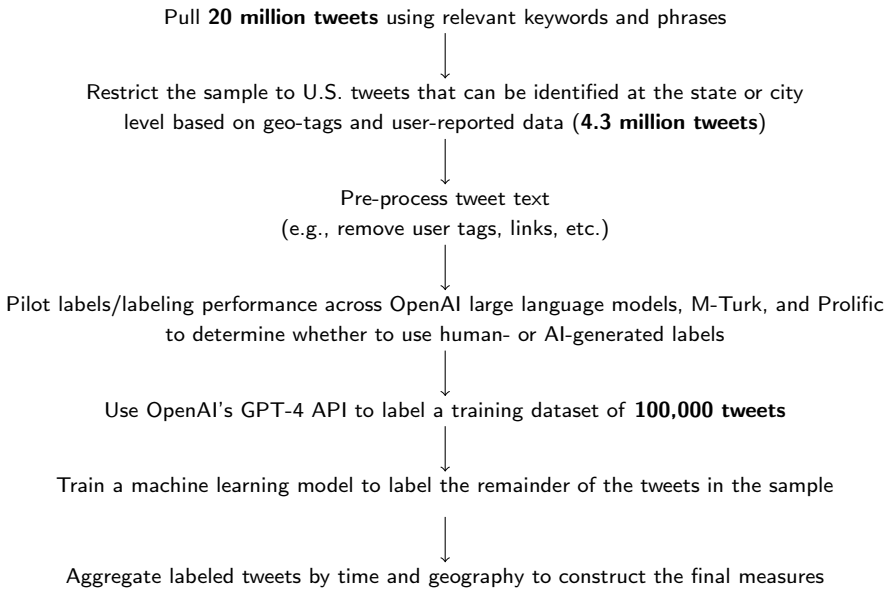

FIGURE 4.1: Outline of the Measure Construction Workflow

4.4.3), this approach presents a more feasible alternative.

4.4.1 Rationale for Using Twitter Data

Twitter data has a number of desirable features for this context, as compared to traditional media sources and other social networking platforms.

Relative to traditional media such as newspapers and TV news, there is simply a higher volume of Twitter data to exploit. Twitter users have been generating hundreds of millions of tweets per day since 2011, providing an extensive corpus of real-time thoughts and attitudes. (Twitter, 2011). This democratization of public speech also means that Twitter offers a much more direct and comprehensive picture of public opinion than traditional mass media, which is produced and gate-kept by a small cadre of journalists, editors, and executives. Twitter posts are unfiltered and opinion-based, in contrast to the formal and (theoretically) objective content of major news outlets, making them ideal for studying something as inherently subjective as social norms.

Twitter also confers a number of practical advantages compared to other social media platforms. First, Twitter has historically provided free API access to academic researchers,

making its data more accessible than any other major social media platform[6]. Second, Twitter data is largely text-based, and therefore easier to parse and analyze than data from image- or video-oriented platforms (e.g., Instagram, YouTube). Twitter posts are also short by design[7], further simplifying search and analysis. Finally, Twitter has geographic information for users in the form of geo-tagging and self-reported profile locations, enabling the construction of a geographically granular measure of norms.

4.4.2 Labeling Pilot: Comparison of Human and Large Language Model Labeling Performance

A key innovation of this measure construction approach is the automation of the "ground truth" training data. This data is used for training the machine learning model used to classify the bulk of the text in the sample. Existing text analysis workflows using supervised machine learning models have largely relied on human-generated responses to provide the "true" answers (e.g., in this case, the correct label classifications) that the model is then trained to predict. However, machine learning models perform best with large training datasets, and obtaining sufficiently large samples of text analysis from human workers can be time-consuming or expensive. I instead automate this process by using a large language model ("LLM")—specifically, OpenAI's GPT-4 model[8]—to label an initial dataset of tweets with which to train a custom machine learning model.

To verify that this approach results in training data of comparable quality to that generated by human labelers, I first conducted a pilot study comparing the two. Specifically, for a pilot sample of 700 randomly-selected tweets, I compared labels from two OpenAI LLMs (GPT-3.5 and GPT-4) to human-generated labels sourced from Amazon MTurk and Prolific (two digital gig work platforms commonly used in social science research). The labels and labeling prompt used in this pilot study are the same as are used in the

[6] Following Elon Musk's takeover of the platform, Twitter announced that it will be discontinuing all free tiers of API access in early 2023.

[7] Twitter posts were initially limited to 140 characters; the limit was doubled in 2017 and remained at 280 characters through 2022.

[8] This model was chosen as it was widely considered the best large language model available at the time of the analysis, and significantly outperformed OpenAI's lower-tier GPT-3.5 model in the pilot.

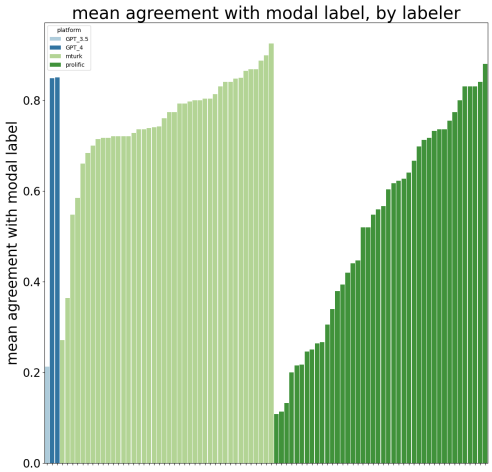

FIGURE 4.2: Comparison of Human and Large Language Model Labels from Pilot

Note: Each tweet in the pilot sample was labeled by nine separate labelers: GPT-3.5, GPT-4 at two different "temperature" settings, and three different Prolific and M-Turk workers, respectively. The figure shows the extent to which each labeler's labels agreed with the modal label for each tweet. Distributions for M-Turk and Prolific reflect different labelers as captured by worker ID. As shown by the dark blue bars, the GPT-4 labels aligned with the modal label over 80% of the time—higher than most human labelers from either platform.

full measure construction workflow, and are described in Section 4.4.3 below. I collected nine labels for each tweet in the sample: one from GPT-3.5, two from GPT-4 (each at a different "temperature" setting), three from three distinct MTurk workers, and three from three distinct Prolific workers. Based on these nine labels, the majority-chosen or modal label was designated as the "correct" label for each tweet, and used to benchmark the performance of each labeling source. Figure 4.2 shows accuracy (defined as percent of total labels that were correct) for each labeler in the pilot sample. The distributions for M-Turk and Prolific reflect accuracy for each unique worker in the sample (allowing for the possibility that some workers on the platform are higher-quality than others, which appears to be the case[9]). The GPT-4 labels aligned with the correct (modal) label over 80% of the time—higher than most human labelers from either platform. These results suggest that text classification from GPT-4 is of comparable quality to analysis by human workers.

[9] There was significant variance in worker quality, as measured by label accuracy, among both MTurk and Prolific workers, in spite of limiting the MTurk sample to "Master" workers (as designated by the platform) and limiting the sample on both platforms to workers located in the U.S.

4.4.3 Tweet Labeling

I employ a supervised machine learning approach to perform multi-class label classification on each tweet in the sample. As there is little concrete guidance regarding the exact language that best characterizes the NTW, I include several different (exclusive) category labels capturing different facets of work norms, as follows:

1. Anti-work
2. Norm to work
3. Pro-workers' rights
4. Unemployment stigma
5. Not applicable

The final label captures relevance, as many of the tweets in the sample are simply off-topic. The labels were chosen to distinguish between positive versus negative framings of the social norm to work, as well as positive versus negative framings of the opposing (anti-work) attitudes. In practice, the labels appear to capture meaningfully different concepts. Upon inspection of the labels chosen by workers on MTurk or Prolific, "norm to work" and "anti-work" seem to be interpreted more broadly or abstractly, while "pro-workers' rights" or "unemployment stigma" tend to be chosen for tweets containing more concrete and specific language around, e.g., advocating for better working conditions or criticizing unemployment benefits. See Table 1 for an example of tweets classified into each label in the training data (as discussed in the following paragraphs).

Supervised machine learning requires a pre-existing set of correctly-labeled data to train a model—typically a dataset that has been manually labeled by humans. After conducting a pilot study comparing labels for 700 tweets generated by workers on M-Turk and Prolific to labels chosen by OpenAI's GPT-3.5 and GPT-4 models, I instead opted to use GPT-4 to label the training dataset, as the data quality appeared to be similar across these sources (see Figure 4.2 for detail). I used GPT-4 to label a training dataset of 100,000 tweets, with each tweet classified into one of the categories listed above. I then used this set of 100,000

Table 4.1: Example Labeled Tweets

Tweet text	**Label***
thrive global: why we need to talk about elizabeth holmes's destructive work ethic. (link) via @User	anti-work
people get burned out because they're working so hard not chasing their dream @User	anti-work
looking forward to getting back to work at the ballpark early in the morning and being a productive designer. #earlymorningdesign #natitude	norm to work
being on time means 15 minutes early. employers will notice. followed this for 35 years	norm to work
@User i wonder how many people commenting pro wwe sentiment think it's smart and clever when their boss rips them off.	pro-workers' rights
happy work/life balance being addressed, but y anyone wld ever consider it a "women's issue" is beyond me: (link) @User	pro-workers' rights
@User what the heck is wrong with you? get a job!	unemployment stigma
yeah, people who are lazy and don't work for their money while taking handouts from the government usually are "chill and laid back"	unemployment stigma
another hard working @User family safeguarding their cash and time by not overpaying ridiculously high list fees (link)	not applicable
framingham jobs: server: amc dine-in theatres - framingham, ma - purpose deliver superior service. (link) #jobs #framingham	not applicable

*Reflects labels generated by OpenAI's GPT-4 model and used to train the current version of the machine learning model.

labeled tweets to train a machine learning model[10] for this multi-class labeling task, which generated predicted labels for the remainder of the tweets in the sample.

The details and performance of the preliminary tweet classification model are summarized in Tables 2 and 3 below. The overall accuracy of the multi-class classification on hold-out data is $\sim$ 62%, though accuracy considering each label individually (binary classification) is higher. This level of accuracy is in line with the accuracy of text classification

[10] The machine learning model used for labeling the tweets in this sample was built using the Keras and Tensorflow packages in Python. It is a convolutional neural network model trained to conduct multi-class text classification. Convolutional neural networks are generally considered the gold standard for classification tasks with high-dimensional, unstructured data (such as tweets), as their multi-layer and flexible structure allows for highly complex data representations (Chollet, 2021). Specifically, the models used in this paper are recurrent neural network (RNN) models that include long short-term memory ("LSTM") layers. RNNs are ideal for sequential data (such as text data), as they allow the model to have "memory" of data from earlier in a sequence of elements (essentially, state-dependence).

models used in similar contexts[11], but introduces a significant amount of noise into the final NTW measures.

This process resulted in a sample of about 4.3 million tweets, each classified into one of the categories listed above.

Table 4.2: Model Summary

Layer (type)	**Output shape**	**Parameters**
embedding 1 (Embedding)	(None, 67, 200)	10000000
lstm 1 (LSTM)	(None, 64)	67840
dense 1 (Dense)	(None, 64)	4160
dense 2 (Dense)	(None, 5)	325

Table 4.3: Model Performance on Hold-out Data

Model	**Accuracy**	**Precision**	**Recall**
Multi-class classification (all labels simultaneously) (used for preliminary measure)	62%	71%	51%
Binary label classification: "not applicable"	78%	78%	67%
Binary label classification: "anti-work"	85%	60%	38%
Binary label classification: "norm to work"	80%	70%	59%
Binary label classification: "pro workers' rights"	93%	63%	33%
Binary label classification: "unemployment stigma"	95%	63%	42%

[11] E.g., text classification models achieved accuracy of about 70% in Adams-Prassl et al., 2022 (identifying training offers in job ad text), 74%-85% in Lassébie et al., 2021 (identifying skills in job ad text), and 54%-97.5% in Osnabrügge et al., 2023 (identifying topics in parliamentary speeches).

4.4.4 Resulting NTW Measures

The labeled tweets are aggregated by time and geography to create four NTW measures (corresponding to categories 1-4 listed in the previous section). Each measure is adjusted for the volume of *relevant* tweets in a given area at time t. The measure for each label is constructed as follows:

$$NTW_{irt} = \frac{\sum_{\text{tweet} \in \text{tweets}_{rt}} \text{label i}}{\sum_{\text{tweet} \in \text{tweets}_{rt}} \text{relevance label}}$$

Where r denotes region (with measures constructed at both the U.S. state and metro-area level), t denotes time (year), and i denotes the label being used (ranging from 1-4). Measures were constructed for all 50 states and the top 50 largest cities by population. Recall that all labels are binary, so this is equivalent to dividing the count of positive (label=1) tweets by the count of relevant (i.e., not categorized as "Not applicable") tweets for each label.

While I examined results using alternative denominators (population and gross tweet volume), adjusting for relevant tweet volume is theoretically the most robust approach. Area population does not account for variation in the amount people tweet in different areas at different times if this is not directly proportional to population. Adjusting for gross tweet volume solves this issue, but is sensitive to changes in popular interest in work-related discussions. For example, if people become more interested in discussing unemployment policy from one year to the next, but gross tweet volume remains unchanged, the NTW could look artificially higher simply because there are more relevant tweets in the latter sample (both supporting and opposing the NTW). Indeed, the gross volume-adjusted and relevant volume-adjusted measures are highly correlated (typically >95%) within years, but produce different time trends. For the reasons stated above, this paper focuses on results using the relevant volume-adjusted measures.

As shown in Figure 4.3, the share of tweets expressing a "norm to work" attitude has been slowly declining since 2011, while the share of tweets expressing "pro-workers' rights" has grown over the sample period. "Anti-work" attitudes have gained moderate share since

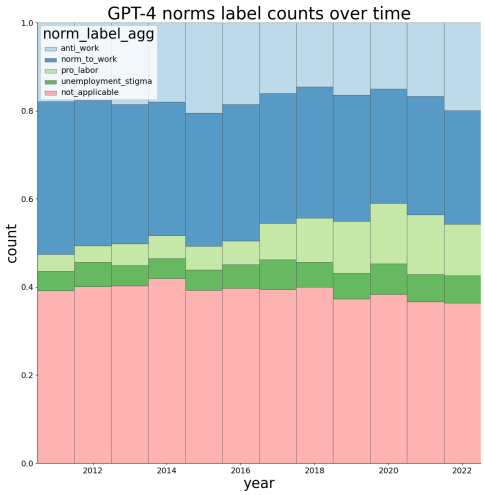

FIGURE 4.3: Label Shares by Year

the beginning of the sample period, with fluctuations throughout. Tweets communicating "unemployment stigma" have remained the smallest share of the sample throughout, perhaps as a result of the specificity of this category. Finally, the share of non-relevant tweets is large—close to 40% of the sample, reflecting the difficulty of precisely targeting keyword searches when collecting tweets—but has fallen gradually over time, consistent with the idea that norms and attitudes around work have been gaining increased media and public attention. Interestingly, there was not a dramatic shift in label shares post-COVID, suggesting that worker dissatisfaction may have been steadily building leading up to the pandemic, with COVID serving as a catalyst for discussion and action.

Levels of each preliminary norm measure by state are shown in Figure 4.4, aggregated over all of the sample years. The level of "anti-work" sentiment by state appears somewhat noisy, but is generally consistent with what one might expect based on media narratives (e.g., stronger anti-work sentiment on the coasts, weaker in the South). The level of "norm to work" sentiment is roughly inverse to the level of "anti-work" attitudes (e.g., stronger in the South and midwest), as we should expect conceptually. Note that this is not by construction, as each tweet could be sorted into one of four possible relevant categories (in

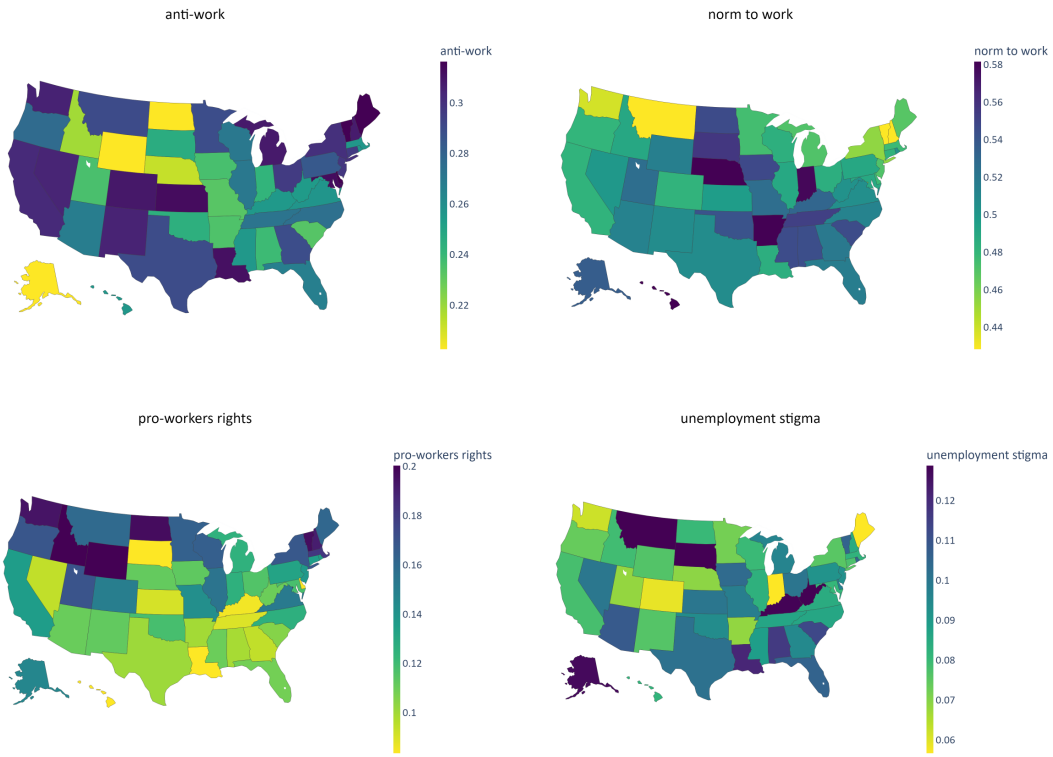

FIGURE 4.4: Preliminary Norm Levels by State (2011-2022 Average)

other words, while the exclusivity of the labels means each label necessarily takes shares from all the others, the correlation between levels of any two particular labels could be null or even positive). Northeastern and upper midwestern states have the highest levels of "pro-workers' rights" attitudes, consistent with historical union strongholds, while "unemployment stigma" levels approximately invert this pattern.

4.5 Results

Using the norms measures described above, I examine the empirical relationship between work norms and labor supply (as measured by labor force participation and employment) and wages. Data for the individual-level labor market outcomes used in the analysis—labor force status, employment, and wage data—come from the Census Bureau's Annual Social and Economic Supplement (ASEC) of the Current Population Survey for years 2011-2019 (see Section 4.3.2 for details). There appears to be a meaningful relationship between work

norms and labor supply that is directionally consistent with theory. In particular, higher levels of "norm to work" ("anti-work") attitudes are associated with increased (decreased) labor supply. The relationship between norms and wages is unclear, perhaps due to noise or opposing direct and indirect effects of norms. In all cases, the results should be interpreted as descriptive rather than causal, as the current empirical approach cannot distinguish the extent to which norms are driving labor market outcomes or vice versa.

These results are based on the time period from 2011-2019. While the COVID-19 pandemic presents a potentially interesting shock to work norms, it also impacted the labor market in a number of other ways, both directly and indirectly, making any interpretation of post-COVID results inherently challenging. I hope to explicitly address the impact of COVID-19 on the NTW, and possible implications for labor market dynamics post-COVID, in future work.

All analyses employ two-way fixed effects to control for unobserved heterogeneity due to year and time-invariant geographical characteristics. However, this specification is not robust to simultaneity, an obvious concern in this setting. Therefore, the current results should be interpreted with caution—in particular, it is not currently possible to say anything definitive about the direction of causality. Teasing out causality, for example through policy or labor events (e.g., large strikes) that could offer plausibly exogenous variation in labor supply or work norms, is a focus of ongoing work.

The estimating equation is given by:

$$Y_{irt} = \alpha + \beta_1 norm_{rt} + \mathbf{X}_i\beta_X + \mathbf{Z}_{rt}\beta_Z + \phi_r + \tau_t + \epsilon_{irt} \tag{4.1}$$

where β_1 is the coefficient of interest, Y_{irt} is individual i's binary employment status [unemployment status, wage] at time t in region r, $norm_{rt}$ is the norm measure at time t in region r, $\mathbf{X}_i$ is a vector of individual-level controls (e.g., demographics, union membership), $\mathbf{Z}_{rt}$ is a vector of region-by-time controls (e.g., lagged regional unemployment, unemployment benefit generosity), and ϕ_r and τ_t are region and time fixed effects. The norms measures

are standardized standardized throughout.

4.5.1 Labor Supply

Using the framework of reference-dependent labor supply from Farber, 2008, I predict that stronger work norms (in terms of "norm to work" and "unemployment stigma") will be associated with higher labor force participation and employment. In particular, to the extent that workers share a common understanding of the minimum number of hours they need to work to be considered employed, thus complying with the norm to work, we should see a positive relationship between the NTW and binary labor force participation and employment, as well as hours up to the reference point. Preliminary results are consistent with these predictions.

The observed relationships between norms and labor supply are broadly consistent with theoretical predictions, noting however that not all results are significant at conventional levels. In particular, as shown in Tables 4.4 and 4.5, there is a strong and significant negative relationship between "anti-work" attitudes and labor force participation and employment, and a strong positive relationship between NTW and these outcomes. An increase of one standard deviation in anti-work attitudes in a given state-year is associated with a roughly 0.6 percentage point decrease in the probability of both labor force participation and employment, representing a 7-8% decrease relative to the sample means of these outcomes. Conversely, and increase of one standard deviation in the NTW is associated with a roughly 0.3 percentage point increase in the probability of both labor force participation and employment. These results control for U.S. macroeconomic conditions (through year fixed effects) as well as state-specific labor market conditions (through lagged state unemployment rates and the state's unemployment insurance replacement ratio, a proxy for benefit generosity), offering some suggestion that variation in norms that is orthogonal to labor market conditions may be driving these results.

Labor supply results for the "pro-workers' rights" and "unemployment stigma" measures are imprecisely estimated and thus more difficult to interpret. As shown in Tables 4.4 and

Table 4.4: Effect of State-level Norms on Labor Force Participation

	(1)	(2)	(3)	(4)
Anti-Work	-0.0057***			
	(0.0019)			
Norm to Work		0.0028**		
		(0.0011)		
Pro-Workers Rights			0.0040*	
			(0.0022)	
Unemp. Stigma				-0.0005
				(0.0013)
Controls:				
UI Replacement Ratio	0.0009	-0.0017	-0.0073	-0.0079
	(0.0315)	(0.0317)	(0.0313)	(0.0314)
Lagged Unemp. Rate	0.0858	0.0698	0.1043	0.0899
	(0.1008)	(0.1021)	(0.1005)	(0.1004)
Population	-0.0073***	-0.0069***	-0.0067***	-0.0066***
	(0.0017)	(0.0016)	(0.0016)	(0.0015)
Tweet Count	-0.0002	-0.0005	-0.0002	-0.0004
	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Demographics	✓	✓	✓	✓
Year and region FEs	✓	✓	✓	✓
Constant	0.3194***	0.3191***	0.3118***	0.3133***
	(0.0305)	(0.0299)	(0.0294)	(0.0292)
Mean LFP	0.782	0.782	0.782	0.782
R2	0.113	0.113	0.113	0.113
N	785432	785432	785432	785432

Standard errors clustered by state-year in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Reflects individual labor force participation (binary) regressed on state-level norms measures and various controls, including respondent demographics and state and year fixed effects. Additional controls include (at the state-year level) tweet count, lagged unemployment, population, and state Unemployment Insurance replacement ratio. Norms measures are standardized.

4.5, "pro-workers' rights" attitudes are generally positively associated with labor supply, though these relationships are only marginally statistically significant. This could reflect the idea that caring enough about work and working conditions to want to change them for the better indicates a greater degree of labor market attachment and investment in work. The estimated coefficients for "unemployment stigma," on the other hand, are technically negative but indistinguishable from zero. It is not surprising that the results for these

Table 4.5: Effect of State-level Norms on Employment

	(1)	(2)	(3)	(4)
Anti-Work	-0.0060***			
	(0.0017)			
Norm to Work		0.0033***		
		(0.0011)		
Pro-Workers Rights			0.0031	
			(0.0024)	
Unemp. Stigma				-0.0012
				(0.0013)
Controls:				
UI Replacement Ratio	-0.0123	-0.0143	-0.0212	-0.0212
	(0.0321)	(0.0322)	(0.0329)	(0.0327)
Lagged Unemp. Rate	-0.4899***	-0.5096***	-0.4736***	-0.4888***
	(0.1011)	(0.1023)	(0.1030)	(0.1022)
Population	-0.0079***	-0.0076***	-0.0073***	-0.0070***
	(0.0017)	(0.0017)	(0.0017)	(0.0017)
Tweet Count	0.0001	-0.0002	0.0001	-0.0001
	(0.0005)	(0.0005)	(0.0006)	(0.0005)
Demographics	✓	✓	✓	✓
Year and region FEs	✓	✓	✓	✓
Constant	0.2544***	0.2547***	0.2472***	0.2466***
	(0.0337)	(0.0330)	(0.0330)	(0.0329)
Mean Employment	0.742	0.742	0.742	0.742
R2	0.108	0.108	0.108	0.108
N	785432	785432	785432	785432

Standard errors clustered by state-year in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Reflects individual employment (binary) regressed on state-level norms measures and various controls, including respondent demographics and state and year fixed effects. Additional controls include (at the state-year level) tweet count, lagged unemployment, population, and state Unemployment Insurance replacement ratio. Norms measures are standardized.

norms categories are unclear as these labels were relatively rare in the training data, and ML models tend to struggle with low-incidence labels, meaning that these measures are likely noisier.

4.5.2 Wages

I next examine the relationship between wages and work norms. These results, shown in Table 4.6, use the log of total income, reported for the last twelve months as of the survey response, as the outcome variable. Wage results are imprecisely estimated (indistinguishable from zero) for all norms measures. This may reflect noise in both the outcome data (income is notoriously difficult to measure accurately in surveys, and existing research has documented specific data quality issues in the ASEC (Weinberg, 2004)) as well as in the norms measures as previously discussed. It should also be noted that this analysis controls for union coverage, the minimum wage, and the unemployment insurance replacement ratio (a proxy for benefit generosity) in each state-year, which are all potentially determined by work norms to some extent. Unsurprisingly, income is positively associated with minimum wage and unemployment insurance levels, suggesting a possible indirect route through which norms may influence wages.

Another potential explanation is simultaneity, as theoretical frameworks predict opposing effects of norms on total income. On the one hand, stronger NTW attitudes would presumably cause workers to view work more favorably and work more or more productively, thus earning more income. On the other hand, stronger NTW attitudes could indirectly depress wages by increasing labor supply. Exploring additional causal inference methods to decompose these effects will be a focus of future work.

Table 4.6: Effect of State-level Norms on Wages

	(1)	(2)	(3)	(4)
Anti-Work	0.00137			
	(0.00256)			
Norm to Work		-0.000602		
		(0.00186)		
Pro-Workers Rights			-0.00365	
			(0.00418)	
Unemp. Stigma				0.00318
				(0.00236)
Controls				
UI Replacement Ratio	0.149**	0.149**	0.151**	0.150**
	(0.0610)	(0.0613)	(0.0614)	(0.0612)
Lagged Unemp. Rate	-0.185	-0.181	-0.199	-0.174
	(0.223)	(0.224)	(0.222)	(0.224)
Population	4.02e-09	3.93e-09	3.81e-09	2.99e-09
	(4.34e-09)	(4.34e-09)	(4.28e-09)	(4.41e-09)
Tweet Count	-2.03e-06*	-1.97e-06*	-2.13e-06*	-1.76e-06
	(1.16e-06)	(1.16e-06)	(1.18e-06)	(1.17e-06)
Union Coverage	-0.104	-0.104	-0.102	-0.0976
	(0.132)	(0.132)	(0.132)	(0.132)
Min. Wage	0.00697***	0.00697***	0.00685***	0.00661***
	(0.00220)	(0.00220)	(0.00220)	(0.00218)
Demographics	✓	✓	✓	✓
Year and region FEs	✓	✓	✓	✓
Constant	9.121***	9.121***	9.126***	9.134***
	(0.0756)	(0.0758)	(0.0749)	(0.0755)
Mean Log Wage	10.66	10.66	10.66	10.66
Observations	503,621	503,621	503,621	503,621
R-squared	0.378	0.378	0.378	0.378

Standard errors clustered by state-year in parentheses *** p<0.01, ** p<0.05, * p<0.1

This table shows the relationship between state-level norms and wages, which appears to be small and imprecisely estimated. Columns show the relationship between log wages by state and year with the level of different norms in the same region and year, estimated according to equation (1). Controls include those shown above as well as individual-level demographic controls (e.g., age, gender) and state and year fixed effects. Norms measures are standardized.

4.5.3 Heterogeneity in Labor Supply Results by Gender

While there are large and statistically significant relationships between anti-work and NTW attitudes and both LFP and employment for the full sample, these results are larger and more robust for men. For example, as shown in Tables 4.7 and 4.8, an increase of one standard deviation in anti-work attitudes is associated with an decrease of nearly 0.75 percentage points in the probability a man is in the labor force, and an decrease of almost 0.9 percentage points in the probability a man is employed. By comparison, adult male LFP dropped by about 2.2 percentage points over the course of the 2008 financial crisis (from January 2008 to its December 2009 low) (Bureau of of Labor Statistics, 2023). Additionally, pro-workers' rights attitudes are significantly predictive of labor force participation and labor supply for men, but not for women.

These results are consistent with the idea that men are on average more influenced by social norms around work—e.g., the "breadwinner" norm—than women in the U.S. (Potuchek, 1997; Thébaud and Pedulla, 2016). Importantly, Twitter users are roughly evenly split between men and women, reducing concern that these results are due to measure composition effects (Wojcik and Hughes, 2019). This in turn provides additional suggestive evidence that the social norms measured here are, in fact, driving labor supply decisions to some extent.

It's also interesting to briefly consider these results in the context of recent literature positing supply-side explanations for declining LFP among (particularly younger) working-age men (ranging from the evergreen "generalized moral decay" argument found in Eberstadt, 2016 to the dystopian "video games are crowding out reality" narrative in Aguiar et al., 2021). On the one hand, these results could be interpreted as supporting some version of the argument that young men are opting out of the labor force due to a cultural shift, to the extent that they imply this subgroup is in fact responsive to norms. An alternative explanation is that causality runs in the opposite direction: young men struggling to find

jobs in the context of weakening demand for low- and middle-skill workers[12]. adjust their attitudes to align with their realities (a motivated-reasoning framework similar to, e.g., Masera and Rosenberg, 2020).

Table 4.7: Effect of State-level Norms on Labor Force Participation by Gender

	Men				Women			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Anti-Work	-0.0074***				-0.0041**			
	(0.0022)				(0.0019)			
Norm to Work		0.0030**				0.0025**		
		(0.0013)				(0.0013)		
Pro-Workers Rights			0.0077***				0.0001	
			(0.0028)				(0.0028)	
Unemp. Stigma				-0.0011				0.0001
				(0.0016)				(0.0016)
Controls:								
UI Replacement Ratio	-0.0444	-0.0493	-0.0547	-0.0558	0.0418	0.0413	0.0355	0.0354
	(0.0366)	(0.0373)	(0.0400)	(0.0391)	(0.0450)	(0.0455)	(0.0438)	(0.0437)
Lagged Unemp. Rate	-0.0192	-0.0352	0.0134	-0.0158	0.1955	0.1797	0.2006	0.2009
	(0.1155)	(0.1182)	(0.1195)	(0.1189)	(0.1433)	(0.1437)	(0.1410)	(0.1415)
Population	-0.0107***	-0.0101***	-0.0099***	-0.0096***	-0.0045**	-0.0043**	-0.0040**	-0.0040**
	(0.0026)	(0.0025)	(0.0025)	(0.0024)	(0.0019)	(0.0019)	(0.0019)	(0.0019)
Tweet Count	-0.0002	-0.0006	-0.0001	-0.0005	-0.0002	-0.0004	-0.0003	-0.0003
	(0.0005)	(0.0006)	(0.0006)	(0.0006)	(0.0007)	(0.0007)	(0.0007)	(0.0007)
Demographics	✓	✓	✓	✓	✓	✓	✓	✓
Year and region FEs	✓	✓	✓	✓	✓	✓	✓	✓
Constant	0.3755***	0.3743***	0.3638***	0.3667***	-0.0910**	-0.0903**	-0.0945***	-0.0942***
	(0.0409)	(0.0399)	(0.0399)	(0.0390)	(0.0357)	(0.0354)	(0.0352)	(0.0355)
Mean LFP	0.854	0.854	0.854	0.854	0.716	0.716	0.716	0.716
R2	0.111	0.111	0.111	0.111	0.083	0.083	0.083	0.083
N	373936	373936	373936	373936	411496	411496	411496	411496

Standard errors clustered by state-year in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Reflects individual labor force participation (binary) regressed on state-level norms measures and various controls, including respondent demographics and state and year fixed effects. Additional controls include (at the state-year level) tweet count, lagged unemployment, population, and state Unemployment Insurance replacement ratio. Norms measures are standardized.

[12] E.g.: Autor and Dorn, 2013; Charles et al., 2019.

Table 4.8: Effect of State-level Norms on Employment by Gender

Table: Employment

	Men				Women			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Anti-Work	-0.0087***				-0.0034**			
	(0.0024)				(0.0017)			
Norm to Work		0.0039***				0.0027**		
		(0.0014)				(0.0012)		
Pro-Workers Rights			0.0075**				-0.0014	
			(0.0033)				(0.0027)	
Unemp. Stigma				-0.0012				-0.0011
				(0.0018)				(0.0015)
UI Replacement Ratio	-0.0132	-0.0182	-0.0258	-0.0269	-0.0143	-0.0134	-0.0200	-0.0189
	(0.0410)	(0.0415)	(0.0453)	(0.0443)	(0.0420)	(0.0425)	(0.0413)	(0.0415)
Lagged Unemp. Rate	-0.7603***	-0.7821***	-0.7266***	-0.7554***	-0.2212	-0.2389*	-0.2212	-0.2231
	(0.1381)	(0.1406)	(0.1442)	(0.1418)	(0.1390)	(0.1393)	(0.1377)	(0.1392)
Population	-0.0087***	-0.0081***	-0.0078***	-0.0075***	-0.0075***	-0.0074***	-0.0072***	-0.0069***
	(0.0027)	(0.0027)	(0.0027)	(0.0026)	(0.0022)	(0.0022)	(0.0022)	(0.0022)
Tweet Count	-0.0001	-0.0006	-0.0001	-0.0004	0.0003	0.0001	0.0002	0.0002
	(0.0007)	(0.0007)	(0.0008)	(0.0008)	(0.0007)	(0.0007)	(0.0007)	(0.0006)
Constant	0.2663***	0.2655***	0.2538***	0.2565***	-0.0685*	-0.0670*	-0.0703*	-0.0739*
	(0.0461)	(0.0451)	(0.0456)	(0.0450)	(0.0395)	(0.0392)	(0.0392)	(0.0395)
Controls:								
Demographics	✓	✓	✓	✓	✓	✓	✓	✓
Year and region FEs	✓	✓	✓	✓	✓	✓	✓	✓
Mean emp	0.808	0.808	0.808	0.808	0.681	0.681	0.681	0.681
R2	0.111	0.111	0.111	0.111	0.082	0.082	0.082	0.082
N	373936	373936	373936	373936	411496	411496	411496	411496

Standard errors clustered by state-year in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Reflects individual employment (binary) regressed on state-level norms measures and various controls, including respondent demographics and state and year fixed effects. Additional controls include (at the state-year level) tweet count, lagged unemployment, population, and state Unemployment Insurance replacement ratio. Norms measures are standardized.

4.6 Conclusion and Directions for Future Work

This paper provides novel empirical evidence on the relationship between the social norm to work (NTW) and aggregate labor market outcomes. Using a unique measure of work norms constructed from social media data, I document a positive correlation between the NTW and labor supply—stronger work norms are associated with higher labor force participation and employment rates. Conversely, higher levels of anti-work attitudes are linked to reduced labor market attachment. While theoretical frameworks suggest an inverse relationship between work norms and wages via labor supply effects, empirical results in this dimension remain inconclusive, likely due to noise in both wage measurement and the constructed NTW measures. The findings contribute to a growing literature on the role of social norms in labor markets and highlight the potential for real-time, high-resolution measurement of such norms using machine learning and unstructured text data.

Important questions remain regarding both the mechanisms underlying the observed relationships and the direction of causality. Future work on this project will seek to address these questions through two main extensions.

First, the COVID-19 pandemic presents a natural experiment in shifting work norms, offering an opportunity to examine how an exogenous shock influences both work attitudes and labor market behavior. While this study deliberately restricts its sample to pre-pandemic years to avoid confounding effects, subsequent work should explore how the pandemic altered work norms and whether these changes had persistent labor market consequences.

Second, while the present analysis documents correlations, establishing causality remains a crucial challenge. Future work should employ more sophisticated econometric tools to disentangle the causal relationship between work norms and labor market outcomes. One potential approach involves leveraging plausibly exogenous variation in work norms—such as policy changes, large-scale labor events (e.g., strikes or unionization drives), or sudden shifts in unemployment benefit generosity—that may influence norms independently of la-

bor market conditions. Additionally, instrumental variable strategies or quasi-experimental methods such as difference-in-differences designs could provide stronger causal inference.

Ultimately, understanding the role of work norms in shaping labor market outcomes has important policy implications, particularly for the design of unemployment benefits, workforce participation incentives, and labor market interventions aimed at marginalized populations. This work also sheds important light on the potential impacts of how unemployment is framed in political discourse and mass media more broadly. Further exploration of these dimensions will help refine our understanding of how deeply embedded social attitudes interact with economic behavior and policy.

By integrating a novel approach to measuring social norms that leverages newly available text data sources and text analysis tools, this research provides a foundation for future work on the evolving role of work norms in the labor market. In an era of rapid labor market transformation, driven by technological change, demographic shifts, and exogenous shocks such as COVID-19, further investigation into these questions will be critical for designing policies that align with both economic realities and prevailing social attitudes toward work.

5. Conclusions

This dissertation explores the role of non-standard variables such as risk aversion, soft skills, and social norms in shaping these labor market outcomes.

Chapter 2 examined job transition decisions and their wage growth implications over workers' careers, demonstrating that larger career transitions can boost wage growth, but carry greater risk. Given that workers are risk averse (to varying degrees), this risk plays an important role in shaping job mobility decisions. These results motivate a model, incorporating both heterogeneous skills and risk preferences, which suggests that active labor market policies focusing on reducing the risk of job transitions (e.g., by addressing information frictions) may be more effective than policies focusing on skills (e.g., through training) in some cases.

Chapter 3 presented results from a field experiment testing the inclusion of soft skills in the hiring criteria of a large online job search and matching platform. Results suggest that including measures of soft skills in applicant ranking—in addition to more traditional criteria such as education and formal work experience—leads to more efficient hiring and better firm-worker matches.

Chapter 4 developed a novel measure of social norms toward work and used this measure to study the relationship between these norms and labor market outcomes in the U.S. over roughly the past decade. Results suggest that there is a meaningful relationship between social norms and labor supply, particularly for men.

Appendix A. Skill Distance Measures

A.1 Distributions of Skill Distance with Random Job Transitions

The distributions of skill distances with moves to randomly drawn occupations are shown in green, while the actual distributions (based on actual observed job transitions in the data) are shown in blue.

(a) Distances between worker skills and destination job skills

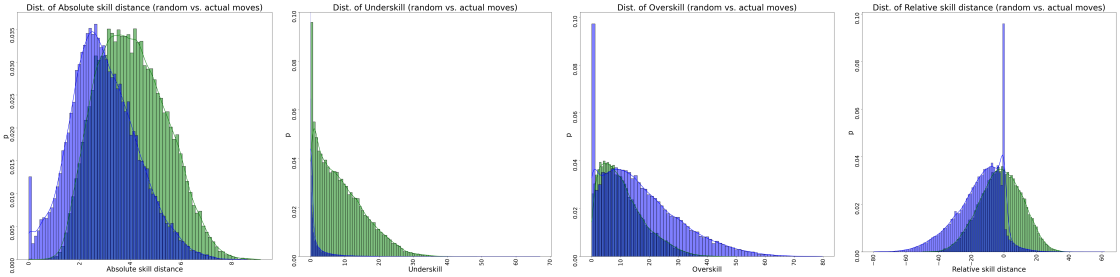

FIGURE A.1: Skill distances of job transitions (random vs. actual moves)

A.2 Distributions of Skill Distance by Transition Type (Voluntary vs. Involuntary)

(a) Distances between worker skills and destination job skills

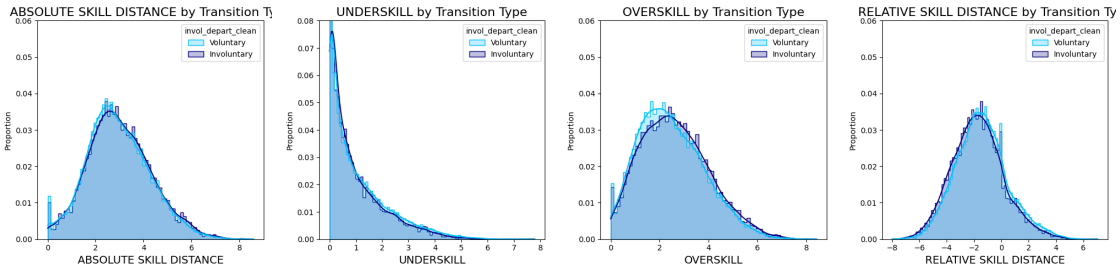

(b) Distances between source and destination job skills

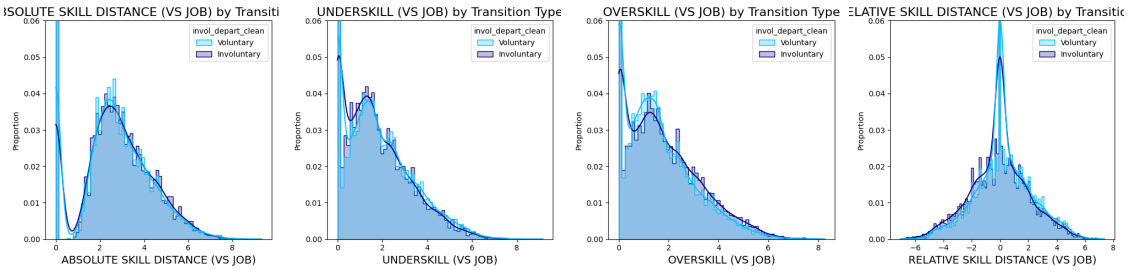

FIGURE A.2: Skill distances of job transitions

Appendix B. Main Results with Various Control Sets

The main results presented above are generally robust to the set of control variables used.

B.1 Training

Table B.1: Any employer-provided training in first year of new job

	(1)	(2)	(3)	(4)
Absolute skill distance	0.003***	0.004***	0.004***	0.003*
	(0.001)	(0.001)	(0.001)	(0.001)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.013	0.018	0.042	0.267
N	33508.000	32707.000	32707.000	31616.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table B.2: Any employer-provided training in first year of new job

	(1)	(2)	(3)	(4)
Underskill	0.012***	0.011***	0.009***	0.007***
	(0.001)	(0.001)	(0.001)	(0.001)
Overskill	-0.000	0.000	-0.000	-0.000
	(0.001)	(0.001)	(0.001)	(0.002)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.016	0.020	0.043	0.267
N	33508.000	32707.000	32707.000	31616.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table B.3: Any employer-provided training in first year of new job

	(1)	(2)	(3)	(4)
Relative skill distance	0.009***	0.008***	0.007***	0.006***
	(0.001)	(0.001)	(0.001)	(0.002)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.015	0.019	0.042	0.267
N	33508.000	32707.000	32707.000	31616.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

B.2 Wage Level vs. Occupation Mean

Table B.4: Wage level in first year of new job vs. occupation mean

	(1)	(2)	(3)	(4)
Absolute skill distance	-0.041***	-0.026***	-0.025***	-0.028***
	(0.002)	(0.002)	(0.002)	(0.002)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.290	0.455	0.554	0.687
N	32685.000	32031.000	32031.000	30933.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table B.5: Wage level in first year of new job vs. occupation mean

	(1)	(2)	(3)	(4)
Underskill	0.024***	0.031***	0.000	-0.004*
	(0.002)	(0.002)	(0.002)	(0.002)
Overskill	-0.057***	-0.042***	-0.031***	-0.033***
	(0.002)	(0.002)	(0.002)	(0.003)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.300	0.465	0.555	0.687
N	32685.000	32031.000	32031.000	30933.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table B.6: Wage level in first year of new job vs. occupation mean

	(1)	(2)	(3)	(4)
Relative skill distance	0.064***	0.056***	0.026***	0.020***
	(0.002)	(0.002)	(0.002)	(0.003)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.299	0.465	0.554	0.686
N	32685.000	32031.000	32031.000	30933.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

B.3 Wage Growth from Job Transitions

Table B.7: Wage Growth from job transitions

	(1)	(2)	(3)	(4)
Absolute skill distance	-0.013***	-0.026***	-0.025***	-0.028***
	(0.002)	(0.002)	(0.002)	(0.002)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.007	0.255	0.391	0.582
N	32031.000	32031.000	32031.000	30933.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table B.8: Wage Growth from job transitions

	(1)	(2)	(3)	(4)
Underskill	0.043***	0.031***	0.000	-0.004*
	(0.002)	(0.002)	(0.002)	(0.002)
Overskill	-0.032***	-0.042***	-0.031***	-0.033***
	(0.002)	(0.002)	(0.002)	(0.003)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.024	0.269	0.392	0.583
N	32031.000	32031.000	32031.000	30933.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table B.9: Wage Growth from job transitions

	(1)	(2)	(3)	(4)
Relative skill distance	0.057***	0.056***	0.026***	0.020***
	(0.002)	(0.002)	(0.002)	(0.003)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.023	0.269	0.390	0.581
N	32031.000	32031.000	32031.000	30933.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

B.4 Relationship Between Risk Aversion and Skill Distances

Table B.10: Absolute J2J skill distance and risk tolerance

	(1)	(2a)	(2b)	(3)	(4)
Risk tolerance (general)	0.004**	0.005**	0.005**	0.005**	0.005***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Controls:					
Demographics	✓	✓	✓	✓	✓
Job & transition vars.		✓	✓	✓	✓
Year & industry/occ cat. FEs				✓	✓
R2	0.0291	0.0349	0.0352	0.0643	0.1249
N	31214.0	30488.0	30488.0	30488.0	30461.0

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table B.11: J2J Underskill distance and risk tolerance

	(1)	(2a)	(2b)	(3)	(4)
Risk tolerance (general)	-0.003	-0.002	-0.002	-0.003	-0.002
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Controls:					
Demographics	✓	✓	✓	✓	✓
Job & transition vars.		✓	✓	✓	✓
Year & industry/occ cat. FEs				✓	✓
R2	0.0327	0.0395	0.0405	0.0559	0.0831
N	31214.0	30488.0	30488.0	30488.0	30461.0

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table B.12: J2J Overskill distance and risk tolerance

	(1)	(2a)	(2b)	(3)	(4)
Risk tolerance (general)	0.006***	0.006***	0.006***	0.007***	0.007***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Controls:					
Demographics	✓	✓	✓	✓	✓
Job & transition vars.		✓	✓	✓	✓
Year & industry/occ cat. FEs				✓	✓
R2	0.0440	0.0488	0.0490	0.0870	0.1605
N	31214.0	30488.0	30488.0	30488.0	30461.0

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table B.13: Relative J2J skill distance and risk tolerance

	(1)	(2a)	(2b)	(3)	(4)
Risk tolerance (general)	-0.006***	-0.006***	-0.006***	-0.006***	-0.006***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Controls:					
Demographics	✓	✓	✓	✓	✓
Job & transition vars.		✓	✓	✓	✓
Year & industry/occ cat. FEs				✓	✓
R2	0.0501	0.0552	0.0558	0.0856	0.1407
N	31214.0	30488.0	30488.0	30488.0	30461.0

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Appendix C. Main Results Using Skill Distance vs. Source Job

The results presented in the main analysis above are directionally similar when using the skill distance between jobs ('J2J') instead of the distance between the destination job skills and the worker's current skill bundle based on their full labor market history.

C.1 Training

Table C.1: Any employer-provided training in first year of new job

	(1)	(2)	(3)	(4)
J2J Abs. distance	0.005***	0.006***	0.005***	0.003**
	(0.001)	(0.001)	(0.001)	(0.001)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.013	0.018	0.042	0.269
N	33320.000	32550.000	32550.000	31461.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table C.2: Any employer-provided training in first year of new job

	(1)	(2)	(3)	(4)
J2J Underskill	0.011***	0.010***	0.009***	0.008***
	(0.001)	(0.001)	(0.002)	(0.002)
J2J Overskill	-0.002*	-0.001	-0.001	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.015	0.020	0.043	0.268
N	33508.000	32707.000	32707.000	31616.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table C.3: Any employer-provided training in first year of new job

	(1)	(2)	(3)	(4)
J2J Relative distance	0.009***	0.008***	0.008***	0.006***
	(0.001)	(0.001)	(0.002)	(0.002)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.015	0.019	0.042	0.267
N	33508.000	32707.000	32707.000	31616.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

C.2 Wage Level vs. Occupation Mean

Table C.4: Wage level in first year of new job vs. occupation mean

	(1)	(2)	(3)	(4)
J2J Abs. distance	-0.022***	-0.009***	-0.023***	-0.018***
	(0.002)	(0.002)	(0.002)	(0.002)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.285	0.453	0.554	0.686
N	32506.000	31878.000	31878.000	30786.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table C.5: Wage level in first year of new job vs. occupation mean

	(1)	(2)	(3)	(4)
J2J Underskill	0.014***	0.032***	0.008***	0.003
	(0.002)	(0.002)	(0.002)	(0.002)
J2J Overskill	-0.045***	-0.045***	-0.039***	-0.029***
	(0.002)	(0.002)	(0.002)	(0.002)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.292	0.464	0.556	0.687
N	32685.000	32031.000	32031.000	30933.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table C.6: Wage level in first year of new job vs. occupation mean

	(1)	(2)	(3)	(4)
J2J Relative distance	0.040***	0.052***	0.031***	0.021***
	(0.002)	(0.002)	(0.002)	(0.002)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.290	0.464	0.554	0.686
N	32685.000	32031.000	32031.000	30933.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

C.3 Wage Growth from Job Transitions

Table C.7: Wage Growth from job transitions

	(1)	(2)	(3)	(4)
J2J Abs. distance	0.001	-0.009***	-0.023***	-0.018***
	(0.002)	(0.002)	(0.002)	(0.002)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.006	0.251	0.390	0.581
N	31878.000	31878.000	31878.000	30786.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table C.8: Wage Growth from job transitions

	(1)	(2)	(3)	(4)
J2J Underskill	0.052***	0.032***	0.008***	0.003
	(0.002)	(0.002)	(0.002)	(0.002)
J2J Overskill	-0.052***	-0.045***	-0.039***	-0.029***
	(0.002)	(0.002)	(0.002)	(0.002)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.034	0.268	0.394	0.583
N	32031.000	32031.000	32031.000	30933.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table C.9: Wage Growth from job transitions

	(1)	(2)	(3)	(4)
J2J Relative distance	0.071***	0.052***	0.031***	0.021***
	(0.002)	(0.002)	(0.002)	(0.002)
Controls:				
Demographics	✓	✓	✓	
Job & transition vars.		✓	✓	✓
Year & industry/occ cat. FEs			✓	✓
Individual FEs				✓
R2	0.034	0.267	0.391	0.581
N	32031.000	32031.000	32031.000	30933.000

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

C.4 Relationship Between Risk Aversion and Job-to-job ('J2J') Skill Distances

Table C.10: Absolute J2J skill distance and risk tolerance

	(1)	(2a)	(2b)	(3)	(4)
Risk tolerance (general)	0.006***	0.006***	0.006***	0.006***	0.007***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Controls:					
Demographics	✓	✓	✓	✓	✓
Job & transition vars.		✓	✓	✓	✓
Year & industry/occ cat. FEs				✓	✓
R2	0.0170	0.0233	0.0237	0.0629	0.1640
N	31044.0	30347.0	30347.0	30347.0	30322.0

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table C.11: J2J Underskill distance and risk tolerance

	(1)	(2a)	(2b)	(3)	(4)
Risk tolerance (general)	0.002	0.002	0.002	0.003	0.003
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Controls:					
Demographics	✓	✓	✓	✓	✓
Job & transition vars.		✓	✓	✓	✓
Year & industry/occ cat. FEs				✓	✓
R2	0.0120	0.0261	0.0271	0.0938	0.2001
N	31214.0	30488.0	30488.0	30488.0	30461.0

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table C.12: J2J Overskill distance and risk tolerance

	(1)	(2a)	(2b)	(3)	(4)
Risk tolerance (general)	0.004*	0.004	0.004*	0.004*	0.004**
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Controls:					
Demographics	✓	✓	✓	✓	✓
Job & transition vars.		✓	✓	✓	✓
Year & industry/occ cat. FEs				✓	✓
R2	0.0081	0.0125	0.0128	0.1393	0.3460
N	31214.0	30488.0	30488.0	30488.0	30461.0

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Table C.13: Relative J2J skill distance and risk tolerance

	(1)	(2a)	(2b)	(3)	(4)
Risk tolerance (general)	-0.002	-0.001	-0.001	-0.001	-0.001
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Controls:					
Demographics	✓	✓	✓	✓	✓
Job & transition vars.		✓	✓	✓	✓
Year & industry/occ cat. FEs				✓	✓
R2	0.0014	0.0148	0.0157	0.1621	0.3899
N	31214.0	30488.0	30488.0	30488.0	30461.0

Robust standard errors in parentheses.
* p<0.1, ** p<0.05, ***p<0.01

Appendix D. Validation of Soft Skills Measure

D.1 Soft Skills Predict Planned Search and Work Outcomes in the Control Group

As shown in Table D.1, the soft skills measure used in our study predicts both planned job search and employment outcomes for untreated jobseekers, as captured in our follow-up survey. Also, based on data from a pilot survey as well as data from our partner's job search platform, our soft skills measures predict having a small business, business earnings, on-platform job search activity, and the length of on-platform "career goals" text.

Table D.1: Soft Skills and Planned Search and Work Outcomes in the Control Group

	Planned search			Work		
	(1) Hours	(2) Expend	(3) Apps	(4) Any	(5) Hours	(6) Earnings
Soft skill score (z)	3.393***	23.671***	3.699***	0.014***	0.625***	194.811***
	(0.510)	(6.649)	(0.295)	(0.005)	(0.179)	(45.740)
Any work experience	-1.850*	-25.007*	1.134*	0.009	-0.371	14.117
	(1.014)	(13.580)	(0.586)	(0.010)	(0.361)	(88.524)
Math score (z)	-0.374	-23.492***	-0.413	-0.009*	-0.381**	-66.164
	(0.496)	(6.473)	(0.294)	(0.005)	(0.176)	(45.250)
English score (z)	0.694	0.109	1.905***	0.011**	0.530***	94.881*
	(0.510)	(6.647)	(0.275)	(0.005)	(0.189)	(51.090)
Control mean	46.576	572.365	31.143	0.395	9.133	2016.698
Observations	9832	9816	9810	10017	10006	9981

D.2 No Evidence of Jobseekers Giving Deliberately Untruthful Answers

There may be some concern that jobseekers, rather than answering the soft skills questions honestly, are deliberately manipulating their responses in order to appear more attractive to employers and boost their probability of getting a job. To get some sense of whether this seems to be happening, we conduct a pilot study ($N = 611$) in which we vary the framing of the soft skills elicitation:

- *High stakes*: Please answer this survey as truthfully as possible. #SAYouth may use your answers to match you to work and learning opportunities.

- *Low stakes*: Please answer this survey as truthfully as possible. #SAYouth want to learn more about you so they can best support you in your journey.

The "high stakes" framing is intended to highlight the potential importance of responses for employment opportunities, thus boosting the incentive to intentionally misrepresent oneself.

Our pilot yields no evidence of such deliberate misreporting. In the high stakes treatment group, soft skill scores are only 0.05 standard deviations higher, and this result is insignificant at conventional levels ($p = 0.52$). There is no evidence of bunching at the very top; differences are largest in the left tail of the distribution. Further, internal consistency of subscales (Cronbach's alpha) does not vary across treatments.

Bibliography

Abebe, G., Caria, A. S., Fafchamps, M., Falco, P., Franklin, S., & Quinn, S. (2021). Anonymity or distance? job search and labour market exclusion in a growing african city. *The Review of Economic Studies*, *88*(3), 1279–1310.

Abebe, G., Caria, A. S., & Ortiz-Ospina, E. (2021). The selection of talent: Experimental and structural evidence from ethiopia. *American Economic Review*, *111*(6), 1757–1806.

Abebe, G., Caria, S. A., Fafchamps, M., Falco, P., Franklin, S., Quinn, S., & Shilpi, F. J. (2023). *Matching frictions and distorted beliefs: Evidence from a job fair experiment* (tech. rep.). working paper.

Abel, M., Burger, R., & Piraino, P. (2020). The value of reference letters: Experimental evidence from south africa. *American Economic Journal: Applied Economics*, *12*(3), 40–71.

Acemoglu, D., & Pischke, J.-S. (1998). Why do firms train? theory and evidence. *The Quarterly journal of economics*, *113*(1), 79–119.

Adams-Prassl, A., Barbanchon, T. L., & Marcato, A. (2022, May). *On the job training and labor competition* [Presented at the 5th IDSC of IZA Workshop: Matching Workers and Jobs Online - New Developments and Opportunities for Social Science and Practice].

Adda, J., & Dustmann, C. (2023). Sources of wage growth. *Journal of Political Economy*, *131*(2), 456–503.

Adhvaryu, A., Kala, N., & Nyshadham, A. (2023). Returns to on-the-job soft skills training. *Journal of Political Economy*, *131*(8), 2165–2208.

Africa, S. S. (2023). *Quarterly labour force survey*.

Afridi, F., Dhillon, A., Roy, S., & Sangwan, N. (2023). *Social networks, gender norms and labor supply: Experimental evidence using a job search platform* (tech. rep.). Competitive Advantage in the Global Economy (CAGE).

Aguiar, M., Bils, M., Charles, K. K., & Hurst, E. (2021). Leisure luxuries and the labor supply of young men. *Journal of Political Economy*, *129*(2), 337–382.

Ahn, H., & Powell, J. L. (1993). Semiparametric estimation of censored selection models with a nonparametric selection mechanism. *Journal of Econometrics*, *58*(1-2), 3–29.

Ahrens, A., Hansen, C., & Schaffer, M. (2019). Pdslasso: Stata module for post-selection and post-regularization ols or iv estimation and inference.

Akerlof, G. A. (1980). A theory of social custom, of which unemployment may be one consequence. *The quarterly journal of economics*, *94*(4), 749–775.

Akerlof, G. A., & Kranton, R. E. (2010). Identity economics. In *Identity economics*. Princeton University Press.

Allison, P. D., Williams, R., & Moral-Benito, E. (2017). Maximum likelihood for cross-lagged panel models with fixed effects. *Socius*, *3*, 2378023117710578.

Almlund, M., Duckworth, A. L., Heckman, J., & Kautz, T. (2011). Personality psychology and economics. In *Handbook of the economics of education* (pp. 1–181, Vol. 4). Elsevier.

Anderson, A., & Smith, L. (2010). Dynamic matching and evolving reputations. *The Review of Economic Studies*, *77*(1), 3–29.

Antonovics, K., & Golan, L. (2012). Experimentation and job choice. *Journal of Labor Economics*, *30*(2), 333–366.

Apesteguia, J., & Palacios-Huerta, I. (2010). Psychological pressure in competitive environments: Evidence from a randomized natural experiment. *American Economic Review*, *100*(5), 2548–64. https://doi.org/10.1257/aer.100.5.2548

Archsmith, J. E., Heyes, A., Neidell, M. J., & Sampat, B. N. (2021). *The dynamics of inattention in the (baseball) field* (tech. rep.). National Bureau of Economic Research.

Arcidiacono, P., Beauchamp, A., Hull, M., & Sanders, S. (2015). Exploring the racial divide in education and the labor market through evidence from interracial families. *Journal of human capital*, *9*(2), 198–238.

Argaw, B., Maier, M. F., & Skriabikova, O. J. (2017). Risk attitudes, job mobility and subsequent wage growth during the early career. *ZEW-Centre for European Economic Research Discussion Paper*, (17-023).

Autor, D. H., & Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American economic review*, *103*(5), 1553–1597.

Autor, D. H., & Handel, M. J. (2013). Putting tasks to the test: Human capital, job tasks, and wages. *Journal of labor Economics*, *31*(S1), S59–S96.

Autor, D. H., & Scarborough, D. (2008). Does job testing harm minority workers? evidence from retail establishments. *The Quarterly Journal of Economics*, *123*(1), 219–277.

Bagger, J., Fontaine, F., Postel-Vinay, F., & Robin, J.-M. (2014). Tenure, experience, human capital, and wages: A tractable equilibrium search model of wage dynamics. *American Economic Review*, *104*(6), 1551–1596.

Bailes, L. P., & Guthery, S. (2020). Held down and held back: Systematically delayed principal promotions by race and gender. *AERA Open*, *6*(2), n2.

Banerjee, A., & Sequeira, S. (2023). Learning by searching: Spatial mismatches and imperfect information in southern labor markets. *Journal of Development Economics*, *164*, 103111.

Barbulescu, R., & Bonet, R. (2024). Looking for greener grass? prior status and exploration-exploitation decisions in job search. *Organization Science*, *35*(1), 159–176.

Barlevy, G. (2001). Why are the wages of job changers so procyclical? *Journal of labor economics*, *19*(4), 837–878.

Barron, J. M., Berger, M. C., & Black, D. A. (1999). Do workers pay for on-the-job training? *Journal of Human Resources*, 235–252.

Barron, J. M., Black, D. A., & Loewenstein, M. A. (1989). Job matching and on-the-job training. *Journal of labor Economics*, *7*(1), 1–19.

Bassi, V., & Nansamba, A. (2022). Screening and signalling non-cognitive skills: Experimental evidence from uganda. *The Economic Journal*, *132*(642), 471–511.

Becker, J., Grun, G., Seemann, R., Mantz, H., Jacobs, K., Mecke, K., & Blossey, R. (2003). Complex dewetting scenarios captured by thin-film models. *Nat. Mater.*, *2*(1), 59–63.

Behaghel, L., Gurgand, M., Dromundo, S., Hazard, Y., & Zuber, T. (2024). The potential of recommender systems for directing job search: A large scale experiment. *Manuscript, Paris School of Economics.*

Belot, M., Kircher, P., & Muller, P. (2019). Providing advice to jobseekers at low cost: An experimental study on online advice. *The review of economic studies*, *86*(4), 1411–1447.

Ben Dhia, A., Crépon, B., Mbih, E., Paul-Delvaux, L., Picard, B., & Pons, V. (2022). *Can a website bring unemployment down? experimental evidence from france* (tech. rep.). National Bureau of Economic Research.

Benjamini, Y., Krieger, A. M., & Yekutieli, D. (2006). Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*, *93*(3), 491–507.

Benner, C., Lewis, C., & Omar, R. (2007). The south african call centre industry: A study of strategy, human resource practices and performance. *Human Resource Practices and Performance (June 22, 2007)*.

Ben-Porath, Y. (1967). The production of human capital and the life cycle of earnings. *Journal of political economy*, *75*(4, Part 1), 352–365.

Berger, J., & Pope, D. (2011). Can losing lead to winning? *Management Science*, *57*(5), 817–827. https://doi.org/10.1287/mnsc.1110.1328

Berman, E., Bound, J., & Machin, S. (1998). Implications of skill-biased technological change: International evidence. *The quarterly journal of economics*, *113*(4), 1245–1279.

Bertoni, M., & Brunello, G. (2022). Training during recessions: Recent european evidence. *IZA Journal of Labor Policy*, *12*(1).

Bertrand, M., Cortés, P., Olivetti, C., & Pan, J. (2016). *Social norms, labor market opportunities, and the marriage gap for skilled women* (tech. rep.). National Bureau of Economic Research.

Bhattacharya, V., & Howard, G. (2021). Rational inattention in the infield. *AEJ: Microeconomics (Forthcoming)*.

Binzel, C., & Carvalho, J.-P. (2017). Education, social mobility and religious movements: The islamic revival in egypt. *The Economic Journal*, *127*(607), 2553–2580.

Black and white: A matter of capitalization. (2020, June). http://cmosshoptalk.com/2020/06/22/black-and-white-a-matter-of-capitalization/

Blair, P. Q., Debroy, P., & Heck, J. (2021). *Skills, degrees and labor market inequality* (tech. rep.). National Bureau of Economic Research.

Blanchflower, D. G., & Oswald, A. J. (2004). Well-being over time in britain and the usa. *Journal of public economics*, *88*(7-8), 1359–1386.

Blattman, C., & Dercon, S. (2018). The impacts of industrial and entrepreneurial work on income and health: Experimental evidence from ethiopia. *American Economic Journal: Applied Economics*, *10*(3), 1–38.

Blommaert, L., Coenders, M., & Van Tubergen, F. (2014). Discrimination of arabic-named applicants in the netherlands: An internet-based field experiment examining different phases in online recruitment procedures. *Social forces*, *92*(3), 957–982.

Böckerman, P., & Ilmakunnas, P. (2012). The job satisfaction-productivity nexus: A study using matched survey and register data. *Ilr Review*, *65*(2), 244–262.

Bodvarsson, Ö. B., & Brastow, R. T. (1999). A test of employer discrimination in the nba. *Contemporary Economic Policy*, *17*(2), 243–255.

Bonn, D., Eggers, J., Indekeu, J., Meunier, J., & Rolley, E. (2007). *Wetting and spreading* [Submitted to *Rev. Mod. Phys.*]. http://itf.fys.kuleuven.be/~joi/papers/Wetting%20and%20spreading.pdf

Brand, J. E. (2015). The far-reaching impact of job loss and unemployment. *Annual review of sociology*, *41*, 359–375.

Breunig, C., & Burauel, P. (2021). Testability of reverse causality without exogeneous variation. *arXiv preprint arXiv:2107.05936*.

Breza, E., Kaur, S., & Krishnaswamy, N. (2019). *Scabs: The social suppression of labor supply*. National Bureau of Economic Research.

Brunello, G., & De Paola, M. (2004). Market failures and the under-provision of training. *Available at SSRN 608944*.

Brunello, G., & Wruuck, P. (2021). Skill shortages and skill mismatch: A review of the literature. *Journal of Economic Surveys*, *35*(4), 1145–1167.

Burdett, K. (1978). A theory of employee job search and quit rates. *The American Economic Review*, *68*(1), 212–220.

Burdett, K., & Mortensen, D. T. (1998). Wage differentials, employer size, and unemployment. *International Economic Review*, 257–273.

Bursztyn, L., Egorov, G., & Fiorin, S. (2020). From extreme to mainstream: The erosion of social norms. *American economic review*, *110*(11), 3522–3548.

Bursztyn, L., Fujiwara, T., & Pallais, A. (2017). 'acting wife': Marriage market incentives and labor market investments. *American Economic Review*, *107*(11), 3288–3319.

Bursztyn, L., González, A. L., & Yanagizawa-Drott, D. (2020). Misperceived social norms: Women working outside the home in saudi arabia. *American economic review*, *110*(10), 2997–3029.

Busso, M., Montaño, S., & Munoz-Morales, J. S. (2023). *Signaling specific skills and the labor market of college graduates* (tech. rep.). IZA Discussion Papers.

Cappelli, P. H. (2015). Skill gaps, skill shortages, and skill mismatches: Evidence and arguments for the united states. *ILR review*, *68*(2), 251–290.

Card, D., & DiNardo, J. E. (2002). Skill-biased technological change and rising wage inequality: Some problems and puzzles. *Journal of labor economics*, *20*(4), 733–783.

Card, D., Kluve, J., & Weber, A. (2018). What works? a meta analysis of recent active labor market program evaluations. *Journal of the European Economic Association*, *16*(3), 894–931.

Carranza, E., Garlick, R., Orkin, K., & Rankin, N. (2022). Job search and hiring with limited information about workseekers' skills. *American Economic Review*, *112*(11), 3547–3583.

Chade, H., & Lindenlaub, I. (2022). Risky matching. *The Review of Economic Studies*, *89*(2), 626–665.

Chadi, A. (2014). Regional unemployment and norm-induced effects on life satisfaction. *Empirical Economics*, *46*(3), 1111–1141.

Chakravorty, B., Bhatiya, A. Y., Imbert, C., Lohnert, M., Panda, P., & Rathelot, R. (2023). Impact of the covid-19 crisis on india's rural youth: Evidence from a panel survey and an experiment. *World Development*, *168*, 106242.

Chang, H.-C., & Demekhin, E. (2002). Complex wave dynamics on thin films. Elsevier.

Charles, K. K., Hurst, E., & Notowidigdo, M. J. (2019). Housing booms, manufacturing decline and labour market outcomes. *The Economic Journal*, *129*(617), 209–248.

Chekmasova, S. (2020). Smart but scared: The risk-ability tradeoff in entrepreneurship. *Available at SSRN 3742229*.

Chen, D. L., Moskowitz, T. J., & Shue, K. (2016). Decision making under the gambler's fallacy: Evidence from asylum judges, loan officers, and baseball umpires. *The Quarterly Journal of Economics*, *131*(3), 1181–1242.

Chen, J., & Roth, J. (2023). Logs with zeros? some problems and solutions. *The Quarterly Journal of Economics*, qjad054.

Cheung, M., Egebark, J., Forslund, A., Laun, L., Rödin, M., & Vikström, J. (2019). *Does job search assistance reduce unemployment? experimental evidence on displacement effects and mechanisms* (tech. rep.). Working Paper.

Cheung, M., Egebark, J., Forslund, A., Laun, L., Rödin, M., & Vikström, J. (2025). Does job search assistance reduce unemployment? evidence on displacement effects and mechanisms. *Journal of Labor Economics*, *43*(1), 47–81.

Chollet, F. (2021). *Deep learning with python*. Simon; Schuster.

Clark, A. E. (2003). Unemployment as a social norm: Psychological evidence from panel data. *Journal of labor economics*, *21*(2), 323–351.

Colbét, S. (1999). Name with proper names Hello Kitty. *Imag. Res.*, *1*, 123–127.

Coleman, J. S. (2018). The emergence of norms. In *Social institutions* (pp. 35–60). Routledge.

Coles, P., Cawley, J., Levine, P. B., Niederle, M., Roth, A. E., & Siegfried, J. J. (2010). The job market for new economists: A market design perspective. *Journal of Economic Perspectives*, *24*(4), 187–206.

Conlon, J. J., Pilossoph, L., Wiswall, M., & Zafar, B. (2018). *Labor market search with imperfect information and learning* (tech. rep.). National Bureau of Economic Research.

Cozzi, M. (2012). Risk aversion heterogeneity, risky jobs and wealth inequality. *Working paper*.

Crépon, B., Duflo, E., Gurgand, M., Rathelot, R., & Zamora, P. (2013). Do labor market policies have displacement effects? evidence from a clustered randomized experiment. *The quarterly journal of economics*, *128*(2), 531–580.

de Gennes, P.-G., Brochard-Wyart, F., & Quere, D. (2003). *Capillarity and wetting phenomena: Drops, bubbles, pearls, waves*. Springer Verlag.

De Chaisemartin, C., & d'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, *110*(9), 2964–96.

Deming, D., & Kahn, L. B. (2018). Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics*, *36*(S1), S337–S369.

Deming, D. J., & Noray, K. (2020). Earnings dynamics, changing job skills, and stem careers. *The Quarterly Journal of Economics*, *135*(4), 1965–2005.

Dietz, D., & Zwick, T. (2021). The retention effect of training: Portability, visibility, and credibility1. *The International Journal of Human Resource Management*, *33*(4), 710–741.

Dillon, E. W. (2018). Risk and return trade-offs in lifetime earnings. *Journal of labor economics*, *36*(4), 981–1021.

Dinerstein, M., Megalokonomou, R., & Yannelis, C. (2022). Human capital depreciation and returns to experience. *American Economic Review*, *112*(11), 3725–3762.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., & Wagner, G. G. (2011). Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the european economic association*, *9*(3), 522–550.

Donovan, K., Lu, W. J., & Schoellman, T. (2023). Labor market dynamics and development. *The Quarterly Journal of Economics*, *138*(4), 2287–2325.

Dube, A., Naidu, S., & Reich, A. D. (2022). *Power and dignity in the low-wage labor market: Theory and evidence from wal-mart workers* (tech. rep.). National Bureau of Economic Research.

Duckworth, A. L., & Quinn, P. D. (2009). Development and validation of the short grit scale (grit–s). *Journal of personality assessment*, *91*(2), 166–174.

Dweck, C. S. (2006). *Mindset: The new psychology of success*. Random house.

Dyreng, S. D., Mayew, W. J., & Williams, C. D. (2012). Religious social norms and corporate financial reporting. *Journal of Business Finance & Accounting*, *39*(7-8), 845–875.

Eales, M. (1989). Shame among unemployed men. *Social Science & Medicine*, *28*(8), 783–789.

Eberstadt, N. (2016). *Men without work: America's invisible crisis*. Templeton Foundation Press.

Eckardt, D. (2019). Are chemists good bankers? returns to the match between training and occupation.

Elster, J. (1989). Social norms and economic theory. *Journal of economic perspectives*, *3*(4), 99–117.

Elvira, M., & Town, R. (2001). The effects of race and worker productivity on performance evaluations. *Industrial Relations: A Journal of Economy and Society*, *40*(4), 571–590.

Engbom, N. (2022). *Labor market fluidity and human capital accumulation* (tech. rep.). National Bureau of Economic Research.

Eugster, B., Lalive, R., Steinhauer, A., & Zweimüller, J. (2017). Culture, work attitudes, and job search: Evidence from the swiss language border. *Journal of the European Economic Association*, *15*(5), 1056–1100.

Faberman, R. J., Mueller, A. I., & Şahin, A. (2022). Has the willingness to work fallen during the covid pandemic? *Labour Economics*, *79*, 102275.

Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D. B., & Sunde, U. (2015). *The nature and predictive power of preferences: Global evidence* (tech. rep.). IZA Discussion Papers.

Farber, H. S. (2008). Reference-dependent preferences and labor supply: The case of new york city taxi drivers. *American Economic Review*, *98*(3), 1069–82.

Farber, H. S., & Gibbons, R. (1996). Learning and wage dynamics. *The Quarterly Journal of Economics*, *111*(4), 1007–1047.

Fernando, A. N., Singh, N., & Tourek, G. (2023). Hiring frictions and the promise of online job portals: Evidence from india. *American Economic Review: Insights*, *5*(4), 546–562.

Ferracci, M., Jolivet, G., & van den Berg, G. J. (2014). Training the unemployed in france: How does it affect employment and unemployment duration? *Labour Economics*, *27*, 157–169.

Field, E., Jayachandran, S., Pande, R., Rigol, N., & Schaner, S. (2023). Can aspirations and information improve job search? experimental evidence from urban pakistan [Available at SSRN: https://ssrn.com/abstract=4311994]. *Working Paper*.

Fortin, N. M. (2015). Gender role attitudes and women's labor market participation: Opting-out, aids, and the persistent appeal of housewifery. *Annals of Economics and Statistics/Annales d'Économie et de Statistique*, (117/118), 379–401.

Frazis, H., Loewenstein, M. A., et al. (2007). On-the-job-training. *Foundations and Trends® in Microeconomics*, *2*(5), 363–440.

Gable, S. L., Reis, H. T., & Elliot, A. J. (2000). Behavioral activation and inhibition in everyday life. *Journal of personality and social psychology*, *78*(6), 1135.

Garlick, R., & Hyman, J. (2022). Quasi-experimental evaluation of alternative sample selection corrections. *Journal of Business & Economic Statistics*, *40*(3), 950–964.

Gathmann, C., & Schönberg, U. (2010). How general is human capital? a task-based approach. *Journal of Labor Economics*, *28*(1), 1–49.

Gauriot, R., Page, L., et al. (2014). *Does success breed success? a quasi-experiment on strategic momentum in dynamic contests* (tech. rep.). QUT Business School.

Gautier, P., Muller, P., Van der Klaauw, B., Rosholm, M., & Svarer, M. (2018). Estimating equilibrium effects of job search assistance. *Journal of Labor Economics*, *36*(4), 1073–1125.

Genakos, C., Pagliero, M., & Garbi, E. (2015). When pressure sinks performance: Evidence from diving competitions. *Economics Letters*, *132*, 5–8. https://doi.org/https://doi.org/10.1016/j.econlet.2015.04.005

Gibbons, R., Katz, L. F., Lemieux, T., & Parent, D. (2005). Comparative advantage, learning, and sectoral wage determination. *Journal of labor economics*, *23*(4), 681–724.

Gibbons, R., & Waldman, M. (2004). Task-specific human capital. *American Economic Review*, *94*(2), 203–207.

Glenn, E. N. (2009). *Shades of difference: Why skin color matters*. Stanford University Press.

Goldsmith, A. H., & Veum, J. R. (2002). Wages and the composition of experience. *Southern Economic Journal*, *69*(2), 429–443.

González-Díaz, J., Gossner, O., & Rogers, B. W. (2012). Performing best when it matters most: Evidence from professional tennis. *Journal of Economic Behavior & Organization*, *84*(3), 767–781. https://doi.org/https://doi.org/10.1016/j.jebo.2012.09.021

González-Díaz, J., & Palacios-Huerta, I. (2016). Cognitive performance in competitive environments: Evidence from a natural experiment. *Journal of Public Economics*, *139*, 40–52. https://doi.org/https://doi.org/10.1016/j.jpubeco.2016.05.001

Grimmer, J., Roberts, M. E., & Stewart, B. M. (2022). *Text as data: A new framework for machine learning and the social sciences*. Princeton University Press.

Groh, M., McKenzie, D., Shammout, N., & Vishwanath, T. (2015). Testing the importance of search frictions and matching through a randomized experiment in jordan. *IZA Journal of Labor Economics*, *4*, 1–20.

Groothuis, P. A., & Hill, J. R. (2013). Pay discrimination, exit discrimination or both? another look at an old issue using nba data. *Journal of Sports Economics*, *14*(2), 171–185.

Guvenen, F., Kuruscu, B., Tanaka, S., & Wiczer, D. (2020). Multidimensional skill mismatch. *American Economic Journal: Macroeconomics*, *12*(1), 210–244.

Hallsworth, M., List, J. A., Metcalfe, R. D., & Vlaev, I. (2017). The behavioralist as tax collector: Using natural field experiments to enhance tax compliance. *Journal of public economics*, *148*, 14–31.

Hamilton, B. H. (1997). Racial discrimination and professional basketball salaries in the 1990s. *Applied Economics*, *29*(3), 287–296.

Hanushek, E. A., Schwerdt, G., Woessmann, L., & Zhang, L. (2017). General education, vocational education, and labor-market outcomes over the lifecycle. *Journal of human resources*, *52*(1), 48–87.

Hardy, M., & McCasland, J. (2023). Are small firms labor constrained? experimental evidence from ghana. *American Economic Journal: Applied Economics*, *15*(2), 253–284.

Hashimoto, M. (1981). Firm-specific human capital as a shared investment. *The American Economic Review*, *71*(3), 475–482.

Heckman, J. (1974). Shadow prices, market wages, and labor supply. *Econometrica: journal of the econometric society*, 679–694.

Heckman, J. J., & Robb Jr, R. (1985). Alternative methods for evaluating the impact of interventions: An overview. *Journal of econometrics*, *30*(1-2), 239–267.

Heckman, J. J., & Sedlacek, G. (1985). Heterogeneity, aggregation, and market wage functions: An empirical model of self-selection in the labor market. *Journal of political Economy*, *93*(6), 1077–1125.

Heckman, J. J., Stixrud, J., & Urzua, S. (2006). The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior. *Journal of Labor economics*, *24*(3), 411–482.

Heller, S. B., & Kessler, J. B. (2024). Information frictions and skill signaling in the youth labor market. *American Economic Journal: Economic Policy*, *16*(4), 1–33.

Hensel, L., Tekleselassie, T. G., & Witte, M. (2024). Formalized Employee Search and Labor Demand. *Working Paper*.

Hershbein, B., Macaluso, C., & Yeh, C. (2018). Labor market concentration and the demand for skills. *IDSC of IZA Workshop: Matching Workers and Jobs Online*, 2–6.

Hirsch, P. B. (2021). The great discontent. *Journal of Business Strategy*.

Hoffman, M., Kahn, L. B., & Li, D. (2018). Discretion in hiring. *The Quarterly Journal of Economics*, *133*(2), 765–800.

Hoffmann, F., Laptiev, A., & Shi, S. (2016). Long-run sectoral reallocation, job to job transitions, and earnings inequality: An empirical investigation. *Job to Job Transitions, and Earnings Inequality: An Empirical Investigation (August 9, 2016)*.

Horton, J. J., & Johari, R. (2018). Buyer signaling improves matching: Evidence from a field experiment. *Available at SSRN 3245704*.

Ingram, B. F., & Neumann, G. R. (2006). The returns to skill. *Labour economics*, *13*(1), 35–59.

Iqbal, H., & Krumer, A. (2019). Discouragement effect and intermediate prizes in multistage contests: Evidence from davis cup. *European Economic Review*, *118*, 364–381.

Jäger, S., Schoefer, B., Young, S., & Zweimüller, J. (2020). Wages and the value of nonemployment. *The Quarterly Journal of Economics*, *135*(4), 1905–1963.

Jahoda, M. (1988). Economic recession and mental health: Some conceptual issues. *Journal of social Issues*, *44*(4), 13–23.

James, V. K., & VANDER HILL, C. W. (1988). Is there discrimination in the "black man's game"? *Social Science Quarterly*, *69*(1), 83.

Jayachandran, S. (2021). Social norms as a barrier to women's employment in developing countries. *IMF Economic Review*, *69*(3), 576–595.

Jetter, M., & Walker, J. K. (2015). Game, set, and match: Do women and men perform differently in competitive situations? *Journal of Economic Behavior & Organization*, *119*, 96–108.

Jetter, M., & Walker, J. K. (2017). Good girl, bad boy? evidence consistent with collusion in professional tennis. *Southern Economic Journal*, *84*(1), 155–180. https://doi.org/https://doi.org/10.1002/soej.12213

Jones, S., & Sen, K. (2022). Labour market effects of digital matching platforms: Experimental evidence from sub-saharan africa.

Kahn, L. M. (2006). Race, performance, pay, and retention among national basketball association head coaches. *Journal of Sports Economics*, *7*(2), 119–149.

Kahn, L. M., & Sherer, P. D. (1988). Racial differences in professional basketball players' compensation. *Journal of Labor Economics*, *6*(1), 40–61.

Kahn, L. B. (2010). The long-term labor market consequences of graduating from college in a bad economy. *Labour economics*, *17*(2), 303–316.

Kambourov, G., & Manovskii, I. (2009). Occupational specificity of human capital. *International Economic Review*, *50*(1), 63–115.

Karabarbounis, L., & Neiman, B. (2014). The global decline of the labor share. *The Quarterly journal of economics*, *129*(1), 61–103.

Kaymak, B. (2014). Postschooling training investment and employer learning. *Journal of Human Capital*, *8*(3), 318–349.

Keane, M. P., & Wolpin, K. I. (1997). The career decisions of young men. *Journal of political Economy*, *105*(3), 473–522.

Kelley, E. M., Ksoll, C., & Magruder, J. (2024). How do digital platforms affect employment and job search? evidence from india. *Journal of Development Economics*, *166*, 103176.

Kim, J. W., & King, B. G. (2014). Seeing stars: Matthew effects and status bias in major league baseball umpiring. *Management Science*, *60*(11), 2619–2644.

Kinder, M., & Stateler, L. (2022). Frontline workers were excluded from companies' pandemic windfalls. no wonder so many are forming unions. *The Avenue*. https://www.brookings.edu/blog/the-avenue/2022/05/04/frontline-workers-were-excluded-from-companies-pandemic-windfalls-no-wonder-so-many-are-forming-unions/

Koning, R. H. (2011). Home advantage in professional tennis [PMID: 21104519]. *Journal of Sports Sciences*, *29*(1), 19–27. https://doi.org/10.1080/02640414.2010.516762

Koursaros, D. (2017). Labor market dynamics when (un) employment is a social norm. *Journal of Economic Behavior & Organization*, *134*, 96–116.

Kranton, R. E. (2016). Identity economics 2016: Where do social distinctions and norms come from? *American Economic Review*, *106*(5), 405–09.

Kropp, B. (2021). Trends that will shape work in 2021 and beyond. *Harvard Business Review*.

Krug, G., Drasch, K., & Jungbauer-Gans, M. (2019). The social stigma of unemployment: Consequences of stigma consciousness on job search attitudes, behaviour and success. *Journal for Labour Market Research*, *53*(1), 1–27.

Krugman, P. (2021). Working out: Is the great resignation a great rethink? *New York Times*, *29*.

Laajaj, R., & Macours, K. (2021). Measuring skills in developing countries. *Labour Economics*, *70*, 101993.

Laschever, R., & Weinstein, R. (2021). Preference signaling and worker-firm matching: Evidence from interview auctions.

Lassébie, J., Marcolin, L., Vandeweyer, M., & Vignal, B. (2021). Speaking the same language: A machine learning approach to classify skills in burning glass technologies data.

Latif, E. (2010). Crisis, unemployment and psychological wellbeing in canada. *Journal of Policy Modeling*, *32*(4), 520–530.

Lazear, E. P. (2009). Firm-specific human capital: A skill-weights approach. *Journal of political economy*, *117*(5), 914–940.

Lee, D. S. (2005). Training, wages, and sample selection: Estimating sharp bounds on treatment effects.

Leszczensky, L., & Wolbring, T. (2022). How to deal with reverse causality using panel data? recommendations for researchers based on a simulation study. *Sociological Methods & Research*, *51*(2), 837–865.

Light, A. (2005). Job mobility and wage growth: Evidence from the nlsy79. *Monthly Lab. Rev.*, *128*, 33.

Light, A., & Ahn, T. (2010). Divorce as risky behavior. *Demography*, *47*(4), 895–921.

Light, A., & McGarry, K. (1998). Job change patterns and the wages of young men. *Review of Economics and Statistics*, *80*(2), 276–286.

Lindenlaub, I., & Postel-Vinay, F. (2023). Multidimensional sorting under random search. *Journal of political Economy*, *131*(12), 3497–3539.

Lindqvist, E., & Vestman, R. (2011). The labor market returns to cognitive and noncognitive ability: Evidence from the swedish enlistment. *American Economic Journal: Applied Economics*, *3*(1), 101–128.

Lippens, L., Baert, S., Ghekiere, A., Verhaeghe, P.-P., & Derous, E. (2020). Is labour market discrimination against ethnic minorities better explained by taste or statistics? a systematic review of the empirical evidence.

Lise, J., & Postel-Vinay, F. (2020). Multidimensional skills, sorting, and human capital accumulation. *American Economic Review*, *110*(8), 2328–2376.

Macaluso, C. (2023). Skill remoteness and post-layoff labor market outcomes. *CESifo Working Paper*.

Macaluso, C., & Borovičková, K. (2024). Heterogeneous job ladders [In preparation for the Carnegie-Rochester 2024 Special Issue of the Journal of Monetary Economics]. *working paper*.

Mack, K., & Palfrey, J. (2020). Capitalizing black and white: Grammatical justice and equity. *MacArther Foundation*, *26*.

Madden, J. F., & Vekker, A. (2017). Output-based performance pay, performance-support bias, and the racial pay gap within a large retail stock brokerage. *Industrial Relations: A Journal of Economy and Society*, *56*(4), 662–687.

Malueg, D. A., & Yates, A. J. (2010). Testing contest theory: Evidence from best-of-three tennis matches. *The Review of Economics and Statistics*, *92*(3), 689–692.

Marira, T. D., & Mitra, P. (2013). Colorism: Ubiquitous yet understudied. *Industrial and Organizational Psychology*, *6*(1), 103–107.

Masera, F., & Rosenberg, M. (2020). *Tell me what you grow and i'll tell you what you think: Westward expansion and the politics of slavery in the us south* (tech. rep.). Working Paper. https://www. dropbox. com/s/pkac59dax0wy5ne/Slavery% 20Paper. pdf.

McAdams, R. H. (1997). The origin, development, and regulation of norms. *Mich. L. Rev.*, *96*, 338.

Meyer, J. (2019). *History of twitter: Jack dorsey and the social media giant*. TheStreet. Retrieved March 28, 2023, from https://www.thestreet.com/technology/history-of-twitter-facts-what-s-happening-in-2019-14995056

Mitts, T. (2022). Countering violent extremism and radical rhetoric. *International Organization*, *76*(1), 251–272.

Moffitt, R. (1983). An economic model of welfare stigma. *The American Economic Review*, *73*(5), 1023–1035.

Mohanty, M. (2019). Effects of job satisfaction on the worker's wage and weekly hours: A simultaneous equations approach. *Journal of Behavioral and Experimental Economics*, *79*, 27–42.

Mooi-Reci, I., & Ganzeboom, H. B. (2015). Unemployment scarring by gender: Human capital depreciation or stigmatization? longitudinal evidence from the netherlands, 1980–2000. *Social science research*, *52*, 642–658.

Mortensen, D. T., & Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *The review of economic studies*, *61*(3), 397–415.

Muehlemann, S., & Leiser, M. S. (2018). Hiring costs and labor market tightness. *Labour Economics*, *52*, 122–131.

Mueller, A. I., & Spinnewijn, J. (2023). Expectations data, labor market, and job search. *Handbook of Economic Expectations*, 677–713.

Munasinghe, L., Reif, T., & Henriques, A. (2008). Gender gap in wage returns to job tenure and experience. *Labour economics*, *15*(6), 1296–1316.

Neal, D. (1995). Industry-specific human capital: Evidence from displaced workers. *Journal of labor Economics*, *13*(4), 653–677.

Neal, D. (1999). The complexity of job mobility among young men. *Journal of labor Economics*, *17*(2), 237–261.

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

Newey, W. K. (2009). Two-step series estimation of sample selection models. *The Econometrics Journal*, *12*(suppl_1), S217–S229.

Nguyên, A. T., & Pendleton, M. (2020). Recognizing race in language: Why we capitalize "black" and "white." *Center for the Study of Social Policy*, *23*.

of Labor Statistics, B. (2023). *Civilian labor force participation rate, men 20 years and older, seasonally adjusted* [[accessed on: March 30, 2023]]. https://www.bls.gov/data/

Osnabrügge, M., Ash, E., & Morelli, M. (2023). Cross-domain topic classification for political texts. *Political Analysis*, *31*(1), 59–80.

Palacios-Huerta, I. (2003). Professionals play minimax. *The Review of Economic Studies*, *70*(2), 395–415.

Pallais, A. (2014). Inefficient hiring in entry-level labor markets. *American Economic Review*, *104*(11), 3565–3599.

Parent, D. (2000). Industry-specific capital and the wage profile: Evidence from the national longitudinal survey of youth and the panel study of income dynamics. *Journal of Labor Economics*, *18*(2), 306–323.

Parsons, C. A., Sulaeman, J., Yates, M. C., & Hamermesh, D. S. (2011). Strike three: Discrimination, incentives, and evaluation. *American Economic Review*, *101*(4), 1410–35.

Paserman, M. D. (2007). *Gender differences in performance in competitive environments: Evidence from professional tennis players* (tech. rep.). IZA. IZA Discussion Paper.

Passel, J. S., & Taylor, P. (2009). *Who's hispanic?* Pew Hispanic Center Washington, DC.

Peeters, T. (2018). Testing the wisdom of crowds in the field: Transfermarkt valuations and international soccer results. *International Journal of Forecasting*, *34*(1), 17–29.

Pissarides, C. A. (1974). Risk, job search, and income distribution. *Journal of Political Economy*, *82*(6), 1255–1267.

Pissarides, C. A. (2000). *Equilibrium unemployment theory*. MIT press.

Pope, D. G., & Schweitzer, M. E. (2011). Is tiger woods loss averse? persistent bias in the face of experience, competition, and high stakes. *American Economic Review*, *101*(1), 129–57. https://doi.org/10.1257/aer.101.1.129

Potuchek, J. L. (1997). *Who supports the family?: Gender and breadwinning in dual-earner marriages*. Stanford University Press.

Powdthavee, N. (2007). Are there geographical variations in the psychological cost of unemployment in south africa? *Social indicators research*, 629–652.

Powell, J. L. (1987). *Semiparametric estimation of bivariate latent variable models* (tech. rep.).

Price, J., & Wolfers, J. (2010). Racial discrimination among nba referees. *The Quarterly journal of economics*, *125*(4), 1859–1887.

Rae, M. (2022). The great resignation strains us labor force. In *Sage business cases*. SAGE Publications: SAGE Business Cases Originals.

Rambachan, A., Kleinberg, J., Ludwig, J., & Mullainathan, S. (2020). An economic perspective on algorithmic fairness. *AEA Papers and Proceedings*, *110*, 91–95.

Riach, P. A., & Rich, J. (2002). Field experiments of discrimination in the market place. *The economic journal*, *112*(483), F480–F518.

Rich, J. (2014). What do field experiments of discrimination in markets tell us? a meta analysis of studies conducted since 2000.

Rubinstein, Y., & Weiss, Y. (2006). Post schooling wage growth: Investment, search and learning. *Handbook of the Economics of Education*, *1*, 1–67.

Rumbaut, R. G. (2011). Pigments of our imagination: On the racialization and racial identities of 'hispanics' and 'latinos'. *How the US racializes Latinos: White hegemony and its consequences*, 15–36.

Sackmann, J. (2021). The match charting project [data provided by Jeff Sackmann; a subset is publicly available here: https://github.com/JeffSackmann/tennis_MatchChartingProject].

Sanders, C., & Taber, C. (2012). Life-cycle wage growth and heterogeneous human capital. *Annu. Rev. Econ.*, *4*(1), 399–425.

Schönberg, U. (2007). Testing for asymmetric employer learning. *Journal of Labor Economics*, *25*(4), 651–691.

Serengil, S. I. (2021). Deepface [documentation here: https://github.com/serengil/deepface].

Shuster, N. (2022). Big-box retail workers, the great resignation, and the anti-work movement.

Stavrova, O., Schlösser, T., & Fetchenhauer, D. (2011). Are the unemployed equally unhappy all around the world? the role of the social norms to work and welfare state provision in 28 oecd countries. *Journal of Economic Psychology*, *32*(1), 159–171.

Stutzer, A., & Lalive, R. (2004). The role of social work norms in job searching and subjective well-being. *Journal of the European Economic Association*, *2*(4), 696–719.

Sull, D., Sull, C., & Zweig, B. (2022). Toxic culture is driving the great resignation. *MIT Sloan Management Review*, *63*(2), 1–9.

Sullivan, P. (2010). Empirical evidence on occupation and industry specific human capital. *Labour economics*, *17*(3), 567–580.

Telles, E., & Paschel, T. (2014). Who is black, white, or mixed race? how skin color, status, and nation shape racial classification in latin america. *American Journal of Sociology*, *120*(3), 864–907.

Thébaud, S., & Pedulla, D. S. (2016). Masculinity and the stalled revolution: How gender ideologies and norms shape young men's responses to work–family policies. *Gender & society*, *30*(4), 590–617.

Topel, R. H., & Ward, M. P. (1992). Job mobility and the careers of young men. *The Quarterly Journal of Economics*, *107*(2), 439–479.

Tsacoumis, S., & Willison, S. (2010). O* net analyst occupational skill ratings: Procedures. *Alexandria, VA: Human Resources Research Organization*.

Twitter. (2011, June). *200 million tweets per day*. Twitter blog post. https://blog.twitter.com/official/en_us/a/2011/200-million-tweets-per-day.html

van Huizen, T., & Alessie, R. (2019). Risk aversion and job mobility. *Journal of Economic Behavior & Organization*, *164*, 91–106.

Van der Meer, P. H. (2014). Gender, unemployment and subjective well-being: Why being unemployed is worse for men than for women. *Social Indicators Research*, *115*(1), 23–44.

Webb, M. A., & Tangney, J. P. (2022). Too good to be true: Bots and bad data from mechanical turk. *Perspectives on Psychological Science*, 17456916221120027.

Weinberg, D. H. (2004). Income data quality issues in the annual social and economic supplement to the current population survey. *Washington: US Census Bureau*, *126*.

Wheeler, C. H. (2008). Local market scale and the pattern of job changes among young men. *Regional Science and Urban Economics*, *38*(2), 101–118.

Wheeler, L., Garlick, R., Johnson, E., Shaw, P., & Gargano, M. (2022). Linkedin (to) job opportunities: Experimental evidence from job readiness training. *American Economic Journal: Applied Economics*, *14*(2), 101–125.

Winkelmann, L., & Winkelmann, R. (1998). Why are the unemployed so unhappy? evidence from panel data. *Economica*, *65*(257), 1–15.

Wojcik, S., & Hughes, A. (2019). Sizing up twitter users.

Wolthoff, R. (2018). Applications and interviews: Firms' recruiting decisions in a frictional labour market. *The Review of Economic Studies*, *85*(2), 1314–1351.

Ziegler, L., et al. (2020). *Skill demand and posted wages: Evidence from online job ads in austria*. Department of Economics, University of Vienna Vienna.

Zizzamia, R., & Ranchhod, V. (2019). Measuring employment volatility in south africa using nids: 2008-2017.

Zschirnt, E., & Ruedin, D. (2016). Ethnic discrimination in hiring decisions: A meta-analysis of correspondence tests 1990–2015. *Journal of Ethnic and Migration Studies*, *42*(7), 1115–1134.