Appendix

A1 Data Sources

Our analysis is based on two publicly available federal government databases—the U.S. Department of Labor's Occupational Information Network (O*NET) and the Bureau of Labor Statistics' Current Population Survey (CPS).

A1.1 O*NET

In order to construct a worker's current skill portfolio and to compute the skill similarity between any two occupations, we rely on the U.S. Department of Labor's Occupational Information Network (O*NET) 25.2 Database. The O*NET database is one of the most comprehensive public data sources for occupational information, providing detailed information about 969 individual occupations categorized by a standard occupation code (SOC). The O*NET database is maintained and updated continually through its Data Collection program, a regularly conducted survey of workers and occupational analysts, assessing topics such as required education, abilities, experience, tasks, and skills.

For each occupation, the O*NET provides importance ratings for a set of 35 skills within seven broad categories: content, process, social, complex problem solving, technical, systems, and resource management. Each skill is rated on a scale from 1 to 5 based on how important that skill is in order to perform the tasks associated with a given occupation. For example, the skills rated most important for a customer service representative are active listening, speaking, and service orientation. In comparison, skills such as installation, equipment selection, and repairing are rated as particularly unimportant for success in this occupation.

Table A1.1. O*NET skill definitions.

Basic Skills	
Active Learning	Understanding the implications of new information for both current and future problem-solving and decision-making.
Active Listening	Giving full attention to what other people are saying, taking time to understand the points being made, asking questions as appropriate, and not interrupting at inappropriate times.

Critical Thinking	Using logic and reasoning to identify the strengths and weaknesses of alternative solutions, conclusions or approaches to problems.
Learning Strategies	Selecting and using training/instructional methods and procedures appropriate for the situation when learning or teaching new things.
Mathematics	Using mathematics to solve problems.
Monitoring	Monitoring/Assessing performance of yourself, other individuals, or organizations to make improvements or take corrective action.
Reading Comprehension	Understanding written sentences and paragraphs in work related documents.
Science	Using scientific rules and methods to solve problems.
Speaking	Talking to others to convey information effectively.
Writing	Communicating effectively in writing as appropriate for the needs of the audience.
Complex Problem Solving Skills	
Complex Problem Solving	Identifying complex problems and reviewing related information to develop and evaluate options and implement solutions.
Resource Management Skills	
Management of Financial Resources	Determining how money will be spent to get the work done, and accounting for these expenditures.
Management of Material Resources	Obtaining and seeing to the appropriate use of equipment, facilities, and materials needed to do certain work.
Management of Personnel Resources	Motivating, developing, and directing people as they work, identifying the best people for the job.
Time Management	Managing one's own time and the time of others.
Social Skills	
Coordination	Adjusting actions in relation to others' actions.
Instructing	Teaching others how to do something.
Negotiation	Bringing others together and trying to reconcile differences.
Persuasion	Persuading others to change their minds or behavior
Service Orientation	Actively looking for ways to help people.

Social Perceptiveness	Being aware of others' reactions and understanding why they react as they do.
Systems Skills	
Judgment and Decision Making	Considering the relative costs and benefits of potential actions to choose the most appropriate one.
Systems Analysis	Determining how a system should work and how changes in conditions, operations, and the environment will affect outcomes.
Systems Evaluation	Identifying measures or indicators of system performance and the actions needed to improve or correct performance, relative to the goals of the system.
Technical Skills	
Equipment Maintenance	Performing routine maintenance on equipment and determining when and what kind of maintenance is needed.
Equipment Selection	Determining the kind of tools and equipment needed to do a job.
Installation	Installing equipment, machines, wiring, or programs to meet specifications.
Operation and Control	Controlling operations of equipment or systems.
Operation Monitoring	Watching gauges, dials, or other indicators to make sure a machine is working properly.
Operations Analysis	Analyzing needs and product requirements to create a design.
Programming	Writing computer programs for various purposes.
Quality Control Analysis	Conducting tests and inspections of products, services, or processes to evaluate quality or performance.
Repairing	Repairing machines or systems using the needed tools.
Technology Design	Generating or adapting equipment and technology to serve user needs.
Troubleshooting	Determining causes of operating errors and deciding what to do about it.

A1.2 CPS

We use the Current Population Survey (CPS) to sample the worker population. Conducted by the U.S. Census Bureau for the Bureau of Labor Statistics (BLS), the CPS is a monthly household survey providing one of the most important sources for economic and social statistics about the labor force over an extended period of time. Each year, the Annual Social and Economic Supplement (ASEC) of the CPS provides a more in-depth snapshot of individuals' employment, earnings, and health insurance. The ASEC is particularly useful for us because workers are asked about both their current occupation and their occupation in the previous year, which allows us to identify cross-occupational transitions.

We access CPS data using the IPUMS-CPS dataset which integrates CPS data from 1962 to the present using harmonized variables which make cross-time comparisons more feasible [46].[1] Because the Census has made multiple updates to their occupation codes during this time frame, IPUMS-CPS provides a harmonized 2010 occupation coding, which condenses the larger set of occupational codes used over this time period to a comparable set of 441 occupations. For each worker in the CPS, we match the O*NET occupational skill requirements for that same occupation to the worker.

Because of the portion of younger workers attending college and making sizeable transitions from part-time work (often unrelated to their career choice) to full-time work after graduation, we adopt the norm in the literature and focus on members of the active labor force who are 25 years or older. We exclude unemployed workers and those in the CPS for whom income or occupational data is not available, primarily because we use a workers' current occupation to estimate her skill portfolio and the distance to alternative occupations. Finally, we restrict our analysis to workers with at least a high school diploma or its equivalent and to workers with no more than a bachelor's degree. While future work should consider the mobility frictions faced by workers with both higher and lower levels of education, we chose to focus our attention on the subset of workers on both sides of the four-year degree threshold.

To account for any differences in the average number hours worked per year by occupation and education level, we convert annual wage earnings to hourly wages using total earnings last year, the number of weeks earned last year, and the usual number of hours worked per week last year. For each occupation and education group, we then estimate median hourly wages weighted by the CPS weights. For data quality purposes, we only include workers who worked more than 13 weeks, worked more than 3 hours per week, and earned at least \$2/hr in the median hourly wage estimates. All hourly wages are

[1]The IPUMS-CPS dataset is publicly available at https://cps.ipums.org/cps/. In order to merge the O*NET and CPS data sets, we use two crosswalks from the Bureau of Labor Statistics to first move from O*NET's occupation codes to the Standard Occupation Classification (SOC) codes and then move from the SOC codes to the CPS occupation codes. Most of the O*NET and CPS occupations match at the 6-digit SOC level (78%), but we used 5-digit (11%) and 4-digit (11%) averages for the ones that did not match at the 6-digit level.

converted to real 2019 dollars using the BLS Consumer Price Index (R-CPS-U).

Throughout our work, we consider the data in rolling ten-year windows (e.g., 2010-2019, 2009-2018) in order to observe a reasonably large sample of cross-occupational transitions and to not miss less frequently occurring origin-destination pairs. While we ultimately extend our analysis back to 1976,[2] our analysis primarily uses the results from 2010 to 2019. In Table A1.2 we report summary statistics of the data. Overall, our sample consists of an average of 102.2 million workers in each year. This corresponds to roughly 62% of the active civilian labor force. The median age of the STARs and workers with bachelor's degrees in our sample are close at 45 and 43 years old, respectively, whereas the share of racial minorities and the share of men is higher among STARs. Workers with bachelor's degrees earn more than STARs and are also more likely to change jobs, however, conditional on making a cross-occupational shift in a given year, STARs are slightly more likely to transition to a job paying at least 10% higher wages.

Table A1.2. Summary statistics of labor force, average from 2010 to 2019

	Workers w/ Bachelor's Degrees	STARs
Total # of workers, annually	31,626,663	70,547,998
Median age	43	45
% Women	49.0	46.4
% Black	8.4	12.9
% Hispanic	8.4	15.4
Median hourly wage	$24.04	$16.35
% of workers who make cross-occupational transitions	9.2	10.2
% of cross-occupational transitions that lead to >10% higher wages	38.1	37.2

[2]This is the earliest that the CPS asked respondents about the total number of weeks worked per year which we use to calculate worker's hourly wages from their reported annual wages.

A2 Validating O*NET Skills

If we conceive of an occupation as a bundle of tasks, the skills rated by the O*NET represent the bundle of job-related competencies necessary for a worker to successfully perform those tasks [48]. While there is certainly variation in the skill sets of workers within an occupation, especially for recently hired workers, we assume that workers develop the full set of skills required to perform their current occupation through training and on-the-job experience. As a result, the 35-item vector of O*NET skill importance ratings skill vector for a worker's current occupation provides a good estimate of their current skill portfolio.

A2.1 Relationship between Skills and Wages

In order to demonstrate the value of our skills measure, we first implement a Mincer wage regression. Mincer's model of earnings [47] has been widely used to explain a worker's earnings as a function of years of schooling and potential work experience.[3] This function is most commonly formulated as:

$$\log(\omega_i) = \alpha + \rho s_i + \beta_1 x_i + \beta_2 x_i^2 \epsilon_i \tag{1}$$

where ω_i is the wages of worker i, α is the intercept which represents the level of earnings of an individual with no years of schooling and no work experience, s_i is the worker i's years of schooling, and x_i is years of potential work experience. Note that the parameters ρ, β_1, and β_2 can be interpreted as the returns to schooling and experience, respectively.[4] This model has been used to study racial and gender wage gaps – often using occupational fixed effects to control for selection effects which might lead female, Black, or Hispanic workers into a different subset of occupations than their male or white counterparts. Our interest in Mincer wage regressions centers around a comparison of the model with occupational fixed effects and an alternative model with controls for a worker's skill portfolio instead. Ultimately, we seek to show that an occupation can be thought of as a set of skills and we do not lose power in explaining the returns to formal schooling even as we validate work experience as a meaningful proxy for skill.

[3] Potential work experience is calculated as a person's age minus years of schooling minus six. Although this definition is accepted and used in the literature, it is worth noting that this definition ignores work experience that students might gain while attending school.

[4] The quadratic is included for years of potential work experience because of the diminishing marginal returns of additional experience.

A2.2 Validating Measure of Skill Distance

In Table A2.1, we present the results from a Mincer wage regression for the year 2019[5]. Column I represents the basic specification of the Mincer wage regression and an augmented regression with controls for race and gender is in Column II. Columns III and IV, which further augment the specification to include occupational fixed effects and the O*NET measures of skills, show that the O*NET measures of skills capture meaningful variation in wages, thereby validating their use as a measure of worker skill in our study.

First, we see that relative to the baseline model, in which the estimated returns to one additional year of school is 16.6% (Column II), the returns to school from the model with controls for O*NET skills is 9.3% (Column IV), which mirrors the estimated returns to school from the model with occupational fixed effects 8.8% (Column III) and the OLS estimates in the literature summarized in Card [49]. Moreover, the model with O*NET skills has an $R^2 = 0.28$, which is both larger than the R^2 in the baseline model ($R^2 = 0.17$) and closer to the explanatory power of the model with occupational fixed effects ($R^2 = 0.32$). Overall, the 35 skill measures in the O*NET database perform almost as well as the 439 occupational fixed effects in estimating the returns to education and in explaining the variance in wages across individual workers. For this reason, we are confident that the O*NET skills measures are good proxies of the skills that workers learn on the job.

[5]We used the 2019 CPS-ASEC because it was the most current available data at the time of analysis, but these results are robust over time

Table A2.1. 2019 Mincer Wage Regressions

	Dependent variable:			
	logwage			
	I	**II**	**III**	**IV**
Education	0.167***	0.166***	0.088***	0.093***
	(0.002)	(0.002)	(0.002)	(0.002)
Work Experience	0.027***	0.027***	0.022***	0.023***
	(0.001)	(0.001)	(0.001)	(0.001)
(Work Experience)^2	−0.0004***	−0.0004***	−0.0003***	−0.0003***
	(0.00001)	(0.00001)	(0.00001)	(0.00001)
Sex: Female		−0.180***	−0.101***	−0.115***
		(0.008)	(0.007)	(0.007)
Ethnicity: Black		−0.146***	−0.068***	−0.084***
		(0.007)	(0.007)	(0.007)
Ethnicity: Hispanic		−0.055***	−0.013	−0.023**
		(0.010)	(0.009)	(0.009)
Ethnicity: Other		−0.254***	−0.185***	−0.193***
		(0.005)	(0.006)	(0.006)
Constant	0.304***	0.750***	2.575***	0.276***
	(0.033)	(0.033)	(0.044)	(0.063)
Fixed Effects	None	None	Occ.	None.
Skill	False	False	False	True
Observations	53,443	53,443	53,443	53,443
R^2	0.119	0.170	0.317	0.281
Adjusted R^2	0.118	0.170	0.312	0.280
Note:				*p<0.1; **p<0.05; ***p<0.01

A3 Model

Let $S_{i,j}$ represent the total surplus generated by a worker starting in occupation 'i' matching to occupation 'j', further let us assume that this surplus is the following function:

$$S_{i,j} = (\alpha_0 + \xi_i) \times \mathbb{1}[j = 1] + \mathbb{1}[j \neq i)] \times \left(\theta \log(d_{i,j}) + \xi_j\right) + \epsilon_{i,j} \tag{2}$$

where ξ_i reflects the average match surplus from staying in occupation i, $d_{i,j}$ represents the skill distance between occupation i and j, ξ_j is a measure of the average match surplus in the destination occupation and $\epsilon_{i,j}$ is an idiosyncratic taste. To close the model, we assume that the $\epsilon_{i,j}$ follows a Gumbel distribution. In equilibrium, the log of the flow rate, i.e. the number of workers leaving origin 'i' for destination 'j' ($N_{i,j}$) divided by the number of workers originally in origin 'i' (N_i), is given by:

$$\log\left(\frac{N_{i,j}}{N_i}\right) = \begin{cases} \theta \log(d_{i,j}) + \xi_j - \log\left(e^{\alpha_0 + \xi_i} + \sum\limits_{k \neq i} e^{\theta \log(d_{i,k}) + \xi_k}\right), \text{ for } j \neq i \\ \frac{e^{\alpha_0 + \xi_i}}{e^{\alpha_0 + \xi_i} + \sum\limits_{k \neq i} e^{\theta \log(d_{i,k}) + \xi_k}}, \text{ for } j = i \end{cases}$$

$$\log\left(\frac{N_{i,j}}{N_i}\right) \approx \begin{cases} \theta \log(d_{i,j}) + \xi_j - \xi_i - \alpha_0, \; j \neq i \\ \frac{e^{\alpha_0 + \xi_i}}{e^{\alpha_0 + \xi_i} + \sum\limits_{k \neq i} e^{\theta \log(d_{i,k}) + \xi_k}}, \text{ for } j = i \end{cases}$$

The relationship between the flow rate, $\log\left(\frac{N_{i,j}}{N_i}\right)$ and the parameters of the model for the case $i \neq j$ follows from the assumption that the number of workers remaining in the origin occupation is close to 1. In practice the fraction of stayers in the origin occupation is higher than 90%, as shown in Table A1.2. Formally, if the fraction of those transitioning is close to 1, then:

$$\log\left(e^{\alpha_0 + \xi_i} + \sum_{k \neq i} e^{\theta \log(d_{i,k}) + \xi_k}\right) \approx \alpha_0 + \xi_i. \tag{3}$$

The intuition for the estimating equation relating the flow rate to the skill distance between origin and occupation jobs in the equation

$$\log\left(\frac{N_{i,j}}{N_i}\right) \approx \theta \log(d_{i,j}) + \xi_j - \alpha_0 - \xi_i, \text{ for } j \neq i \tag{4}$$

is the following: the better the origin occupation ξ_i, the less likely the worker is to leave occupation i, hence the lower flow rate. Conversely, the better the destination occupation,

ξ_j, the more likely the worker is to leave. The term θ captures the cost of making a transition. If $\theta < 0$, then further skill distance transitions $d_{i,j}$ reduce match quality between the firm and worker. If $\theta > 0$, then further skill distance transitions are associated with higher match quality between the firm and the worker. The descriptive results presented in the main paper suggest that skills learned on the job matter for all workers, however workers with bachelor's degrees face less friction than STARs when transitioning to higher-paying occupations. This analysis also shows in absolute terms that STARs' mobility as a function of skills is invariant to whether they are transitioning to higher or lower wages, whereas workers with bachelor's degrees experience less friction when moving upwards than moving downwards. While the descriptive findings are strongly suggestive, there are several threats to identification. The first is omitted variable bias. For example, it may be the case that unobservable attributes of the origin and destination occupations could impact the likelihood of transitioning from one to the other beyond what is captured by the skills distance. These attributes could take the form of non-wage employee benefits, job satisfaction, or occupation-specific degree requirements. To account for these unobservable attributes and move towards a potentially more causal interpretation of the findings, we now implement regression models with origin and destination occupation fixed effects. Including origin and destination fixed effects accounts for the differences in the way that STARs and workers with bachelor's degrees sort into origin jobs and transition to destination jobs for our estimates of the absolute skill mobility friction, purging them from some of the effect of these endogenous differences in worker sorting by type. A second concern is that there is a confounding variable that is correlated with both the skill distance and the flow rate between occupations. In our robustness section, we indirectly test this by using a permutation test in which we drop one of the skills from the calculation of the skill distance and recompute the absolute skill mobility friction. This exercise provides a bound on the extent to which an omitted skill may be impacting our results. In Table A3.1, we outline the four specifications that we use to estimate the relationship between the skill distance and the flow rate.

We use the first model specification with no heterogeneity to estimate the absolute skill mobility friction across all worker types and transitions ($\theta_{1,1}$). The second specification is a stacked regression of the flow rates for both upwardly mobile and downwardly mobile transitions from which we capture the average skill mobility friction for upwardly mobile transitions ($\theta_{1,2} + \theta_{2,2}$) and downwardly mobile transitions ($\theta_{1,2}$), averaged across all worker types. Because this regression is stacked by transition type, we test for differences between the skill mobility frictions for upwardly and downwardly mobile transitions ($\theta_{2,2}$). The third specification measures heterogeneity in the absolute skill mobility friction by worker type. The skill mobility friction for STARs is captured by $\theta_{1,3}$ and that for workers with bachelor's degrees is $\theta_{1,3} + \theta_{2,3}$. The difference between the two is $\theta_{2,3}$, or the relative skill mobility friction. In our fourth specification, we calculate flow rates by both worker type and transition type and then calculate a stacked regression of flow rates on skill distances, allowing for heterogeneity in the skill mobility friction by worker and transition type. The skill mobility friction for STARs is $\theta_{1,4}$ for downwardly mobile transitions and $\theta_{1,4} + \theta_{2,4}$ for upwardly mobile transitions. The skill mobility friction for workers with bachelor's degrees is $\theta_{1,4} + \theta_{3,4}$ for downward transitions and $\theta_{1,4} + \theta_{2,4} + \theta_{3,4} + \theta_{4,4}$

Table A3.1. Summary of model specifications

Interactions	Model Specification
None	$Y_{i,j} = \alpha_{0,1} + \theta_{1,1}\log(d_{i,j}) + \xi_{i,1} + \xi_{j,1} + \epsilon_{i,f,t,1}$
Transition (t)	$Y_{i,j,t} = \alpha_{0,2} + \alpha_{1,2} \times \mathbb{1}[\text{Up}] + (\theta_{1,2} + \theta_{2,2} \times \mathbb{1}[\text{Up}])\log(d_{i,j}) + \xi_{i,2} + \xi_{j,2} + \epsilon_{i,f,t,2}$
Worker (w)	$Y_{i,j,w} = \alpha_{0,3} + \alpha_{1,3} \times \mathbb{1}[\text{BD}] + (\theta_{1,3} + \theta_{2,3} \times \mathbb{1}[\text{BD}])\log(d_{i,j}) + \xi_{i,3} + \xi_{j,3} + \epsilon_{i,f,t,3}$
Both (t,w)	$Y_{i,j,t,w} = \alpha_{0,4} + \alpha_{1,4}\mathbb{1}[\text{Up}] + \alpha_{2,4}\mathbb{1}[\text{BD}] + \alpha_{3,4}(\mathbb{1}[\text{BD}] \times \mathbb{1}[\text{Up}]) +$ $(\theta_{1,4} + \theta_{2,4}\mathbb{1}[\text{Up}] + \theta_{3,4}\mathbb{1}[\text{BD}] + \theta_{4,4} \times \mathbb{1}[\text{Up}] \times [\text{BD}])\log(d_{i,j}) +$ $\xi_{i,4} + \xi_{j,4} + \epsilon_{i,j,t,m,4}$

Note: Up refers to upwardly mobile transitions. BD refers to transitions made by workers with bachelor's degrees.

for upward transitions. In our fourth specification we can calculate the difference in skill mobility friction within transition type and across worker type and the difference in the skill mobility friction across transition type and within worker type. These comparisons allow us to test whether the skill mobility friction is the same for a given worker type regardless of transition type or whether the skill mobility friction is different across worker type given the same transition.

A4 Results on Labor Market Tightness

It is well documented that firms in slack labor markets take advantage of surplus workers with bachelor's degrees. To test whether the differential treatment of workers with bachelor's degrees for upwardly mobile transitions is due to a market failure, we examine whether the gap in absolute skill mobility friction persists in tight labor markets or occurs primarily in loose labor markets. If we were to find that the gap still persists in tight labor markets, this would point to an advantage to workers with bachelor's degrees that cannot be competed away.

For this analysis, we focus on workers in the CPS who reported living in one of the 392 metropolitan statistical areas (MSAs). We use data from the BLS's Local Area Unemploy-

ment Statistics to calculate the average MSA-level annual unemployment rate, and then categorize workers as living in tight labor markets (unemployment rate is below 5%) or loose labor markets (unemployment rate is 5% or above).[6] A more direct measure of labor market tightness would be the number of effective job vacancies to the number of effective job searchers as in Abraham et. al. [50], however, the nature of the CPS does not let us reliably estimate this measure at the MSA-level.

Table A4.1. Absolute skill mobility friction by degree attainment and labor market tightness

	All Workers	STARs	BD	BD - STAR
All Labor Markets	-0.953***	-0.672***	-0.537***	0.135**
	(0.026)	(0.026)	(0.035)	(0.041)
Loose Labor Markets	-0.751***	-0.632***	-0.511***	0.120*
	(0.027)	(0.032)	(0.047)	(0.055)
Tight Labor Markets	-0.820***	-0.686***	-0.541***	0.146*
	(0.030)	(0.032)	(0.047)	(0.055)
Difference (Tight - Loose)	-0.068	-0.054	-0.029	0.025
	(0.038)	(0.048)	(0.068)	(0.081)
Note:				*p<0.05; **p<0.01; ***p<0.001

After classifying workers by the labor market tightness of their MSAs, we calculate the occupation-to-occupation flow rates separately for workers in tight and loose labor markets. We then run regressions similar to those shown in Table A3.1, only replacing the indicator for transition, i.e. $\mathbb{1}[\text{Up}]$), with an indicator for whether labor market is tight or loose (i.e. $\mathbb{1}[\text{Tight}]$). The results from this exercise are reported in Table A4.1.

In the aggregate, we find that flow rates decline more quickly as a function of skill distance in tight labor markets when compared to slack labor markets. Although the ASMF is generally higher in tight labor markets than in loose labor markets, these differences are small and not statistically significant. This mirrors the findings in the paper in which differences in the ASMF between STARs and workers with bachelor's degrees (see Fig. 3A-B) are small until we also distinguish between upwardly and downwardly mobile transitions (see Fig. 4).

We take the analysis one step further here and introduce heterogeneity by whether a job transition is upwardly mobile or downwardly mobile in addition to the heterogeneity

[6] Because of the generally monotonic decline in unemployment during the 10-year period from 2010 to 2019, most MSAs in the beginning of the decade are classified as loose labor markets while most MSAs at the end of the decade are classified as tight. As a result, this absolute measure of labor market may capture genuine variation in the mobility elasticities by labor market tightness, but it also might capture longitudinal changes in the way that transitions occur.

by whether the transition is occurring in a loose versus a tight labor market. First, we condition on the set of upward or downward transitions and then run our fully interacted model:

$$\begin{aligned} Y_{i,j,t,w} = \alpha_{0,5} + \alpha_{1,5}\mathbb{1}[\text{Tight}] + \alpha_{2,5}\mathbb{1}[\text{BD}] + \alpha_{3,5}(\mathbb{1}[\text{BD}] \times \mathbb{1}[\text{Tight}]) \\ + (\theta_{1,5} + \theta_{2,5}\mathbb{1}[\text{Tight}] + \theta_{3,5}\mathbb{1}[\text{BD}] + \theta_{4,5} \times \mathbb{1}[\text{Tight}] \times [\text{BD}])\log(d_{i,j}) \\ + \xi_{i,5} + \xi_{j,5} + \epsilon_{i,j,t,m,5} \end{aligned}$$

We use this specification to test whether labor market tightness matters when conditioned on the transition mobility type. When transitioning to jobs with higher earnings, we find that workers with bachelor's degrees face an ASMF that is between 50 and 60 log points lower than STARs in both tight and loose labor markets, as reported in Table A4.2. Likewise, for downwardly mobile transitions, workers with bachelor's degrees face an ASMF that is about 30 log points higher than STARs in both tight and loose labor markets.

Table A4.2. Relative skill mobility friction by transition type and labor market tightness

	Upwardly Mobile Transitions	**Downwardly Mobile Mobile Transitions**	**Difference (Up - Down)**
Loose Labor Markets	0.538***	-0.326***	0.864***
	(0.076)	(0.077)	(0.108)
Tight Labor Markets	0.558***	-0.327***	0.885***
	(0.088)	(0.092)	(0.127)
Difference (Tight - Loose)	0.020	-0.001	
	(0.117)	(0.120)	
Note:			*p<0.05; **p<0.01; ***p<0.001

What is clear from this analysis is that labor market conditions—whether the market is tight or loose—matter a lot less than the type of transition—whether it is an upward or downward transition—for explaining the differences in ASMF between STARs and workers with bachelor's degrees. Fundamentally, labor market inequality between workers with bachelor's degrees and STARs depends on whether they are moving up or down in the labor market. For workers with bachelor's degrees moving up is easier than moving down whereas for STARs moving down is easier than moving up: workers with bachelor's degrees face more friction when falling down in the labor market and less friction when climbing up, whereas STARs face less friction falling down and more when climbing up.

A5 Omitted Variable Bias

The main threat to identification in our context is omitted variable bias. In particular, suppose there is an unobserved or omitted skill from the O*NET classification, its importance affects the flow rate of workers, and its importance varies across occupations. This could bias our estimate of the skill mobility friction. We propose the following exercise to create potential bounds on this type of omitted variable bias. The key idea behind our approach is that we can synthetically generate an omitted variable in our context by excluding one of the 35 O*NET skills and directly measure how this affects our estimated mobility elasticities. Since we have 35 skills in total, we permute through each of the skills to generate 35 synthetic estimates of our mobility elasticities, each of which, by construction, suffers from an omitted variable problem.[7] To convert this into informative bounds, we first report a histogram of the percent bias that we find. Secondly, for each skill mobility friction, we plot the percent bias against the correlation between the included skill distance and the absolute value of the skill distance of the omitted component. The second exercise allows us to trace out the functional relationship between the correlation of the omitted variable and the size of the omitted variable bias.

A5.1 Simple Example of Bounding Exercise

To illustrate how we conduct this analysis, consider its application to our simplest model in which we estimate an ASMF that does not depend on worker type or transition type:

$$Y_{i,j} = \alpha_{0,1,-k} + \theta_{1,1,-k}\log(d_{i,j,-k}) + \xi_{i,1,-k} + \xi_{j,1,-k} + \epsilon_{i,f,-k}. \tag{5}$$

The omitted skill is the k-th skill. The term $d_{i,j,-k}$ is the Euclidean skill distance between occupation 'i' and 'j' where we consider all O*NET skills except k. Likewise, the $\theta_{1,1,-k}$, is the skill mobility friction obtained from excluding the k-th skill and $b_{1,1,-k} \equiv \frac{\theta_{1,1,-k}-\theta_{1,1}}{\theta_{1,1}}$ is the estimated omitted variable bias in percent terms. Permuting through all skills, we obtain a vector of 35 estimates of the ASMF from which we compute the percent bias in skill mobility friction $\{b_{1,1,-1}, b_{1,1,-2}, \ldots, b_{1,1,-k}, \ldots\ldots, b_{1,1,-35}\}$.

Next, for each omitted skill, we compute the correlation between the skill and the observed skill distance under this omission $\rho_{1,1,-k} = \text{Cor}(\log(\left|s_i^k - s_j^k\right|), \log(d_{i,j}^{-k}))$. We then plot $b_{1,1,-k}$ against $\rho_{1,1,-k}$ The range of the correlations tells us how much support we have for creating omitted variable bias bounds that are internally consistent. The slope of this line helps us to think about extrapolating outside of the sample if the omitted variable has a correlation with the observed distance that is larger than any of the correlations that we have computed.

[7] This is similar to the jack-knife procedure that is used to construct confidence intervals.

A5.2 Application of Bounding Exercise

We now apply our bounding exercise to the fully interacted model in which we allow the ASMF to vary by both worker type (STAR/bachelor's degree) and by transition type (upwardly mobile/downwardly mobile). In Fig. A5.1, we illustrate the density curves of the percent bias for each worker type. For STARs, we can bound the bias on the skill mobility friction by -3% on the low end and 3% on the high end. For workers with bachelor's degrees, we can bound our mobility elasticities on the low end by -6% and on the top end by 10%. The clear message here is that even the most extreme levels of omitted variable bias are unlikely to change our skill mobility friction estimate by more than 10%. In Fig. 6 of the paper, we extend our analysis to compute the percent bias in the difference in the mobility elasticities for workers with bachelor's degrees and STARs by transition.

In Fig. A5.2, we present the results from our plot of the percent bias against the correlation between the omitted and included skill distances. First, we find that there is a broad range in the correlation between the included and omitted skill, ranging from a correlation of 0.14 to a correlation of 0.46. Second, there is a positive approximately linear relationship between the percent bias and the correlation between the omitted and included skill distances: the more positive the correlation between the omitted and included skill distances, the larger the omitted variable bias.

To obtain an even more conservative set of bounds on the estimated elasticities than the bounds that we obtained from the distributions in Fig. A5.1, we extrapolate the linear relationship between the bias and the correlation to their predicted values for the maximum positive and minimum negative correlations possible, i.e., $\rho_{-k} = -1$ and $\rho_{-k} = 1$. Building on this relationship, we extrapolate to hypothetical extreme cases where the correlation between the omitted skill and the included skills is +1 or -1, illustrated in Fig. A5.2. These counterfactual bounds allow us to estimate the potential magnitude of bias even when the omitted factor is either perfectly aligned or entirely misaligned with the included measures. We construct bounds by computing the RSMF under the case that the omitted skills and included skills have a correlation of +1 and -1 respectively. This yields a bias interval of [-4.63, 9.29] for downwardly mobile transitions and [-11.98, 5.86] for upwardly mobile transitions. Note: because the RSMF includes a difference of RSMF, the RSMF calculated from this exercise is less biased that the underlying ASMFs because the RSMF differences out some of the bias in the ASMF, which are both positively correlated with the omitted skill.

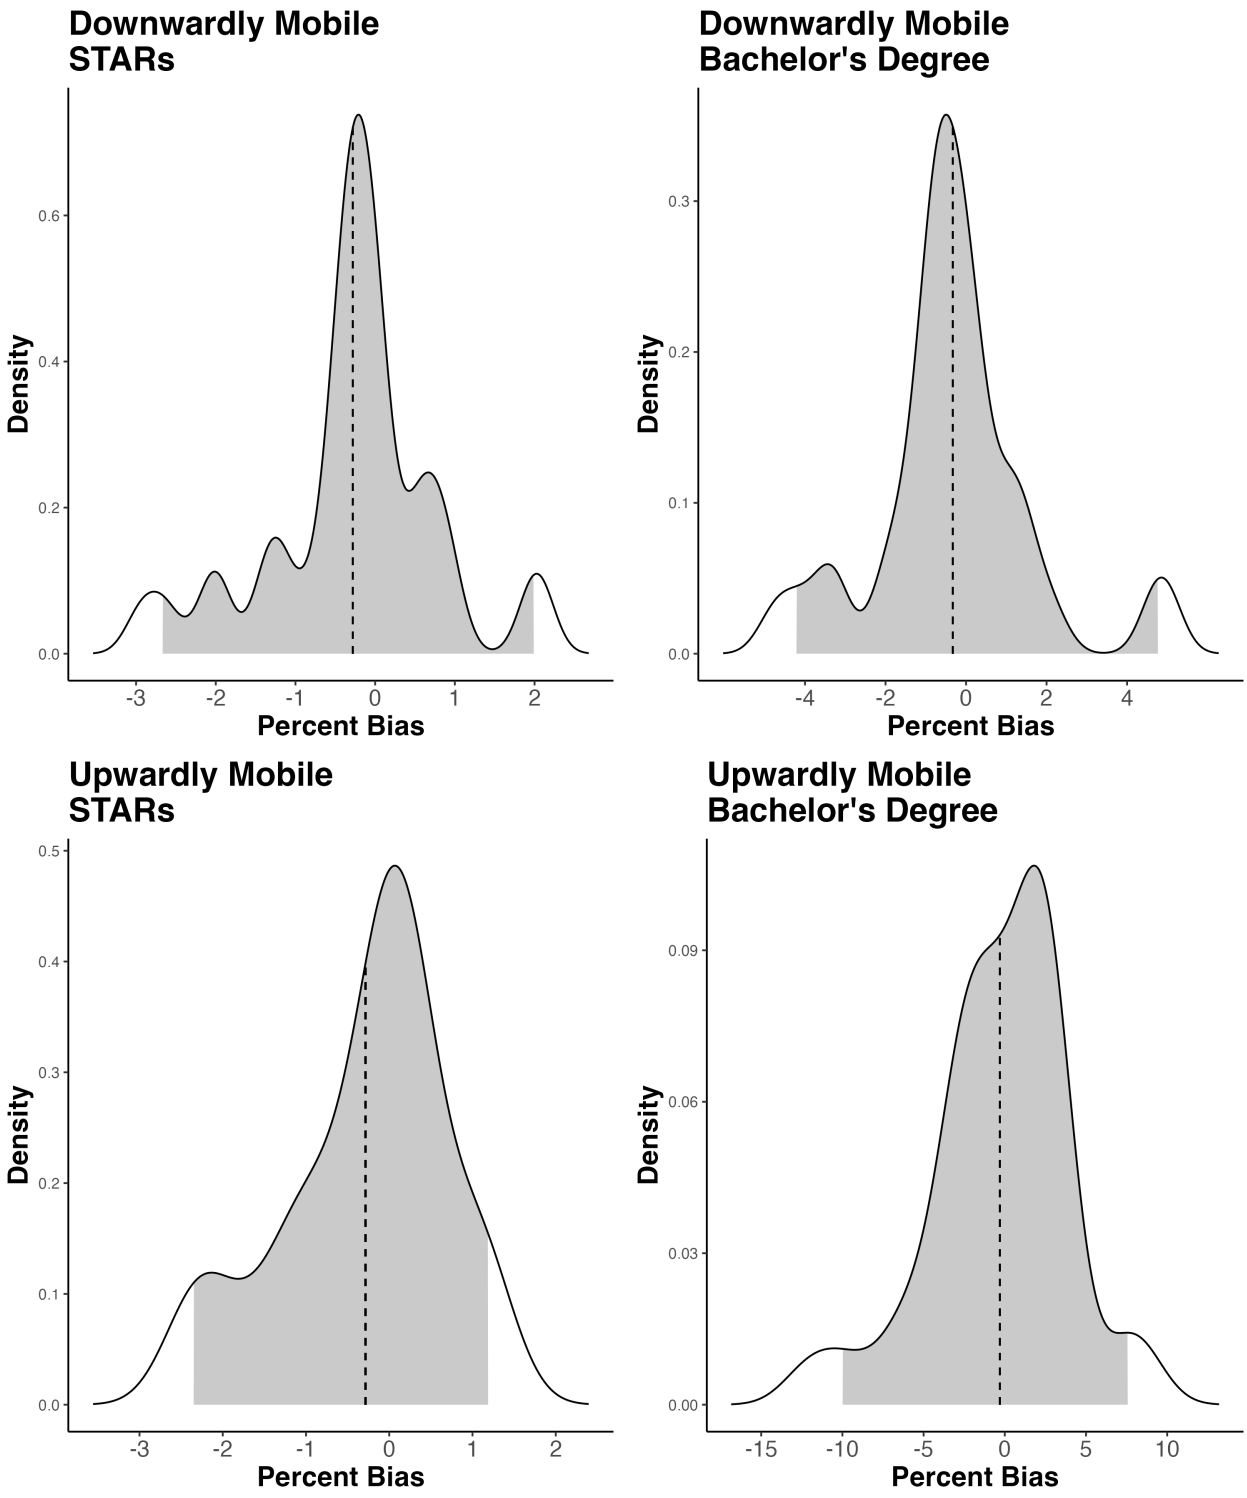

Fig. A5.1. Potential Omitted Variable Bias by Education and Transition Type

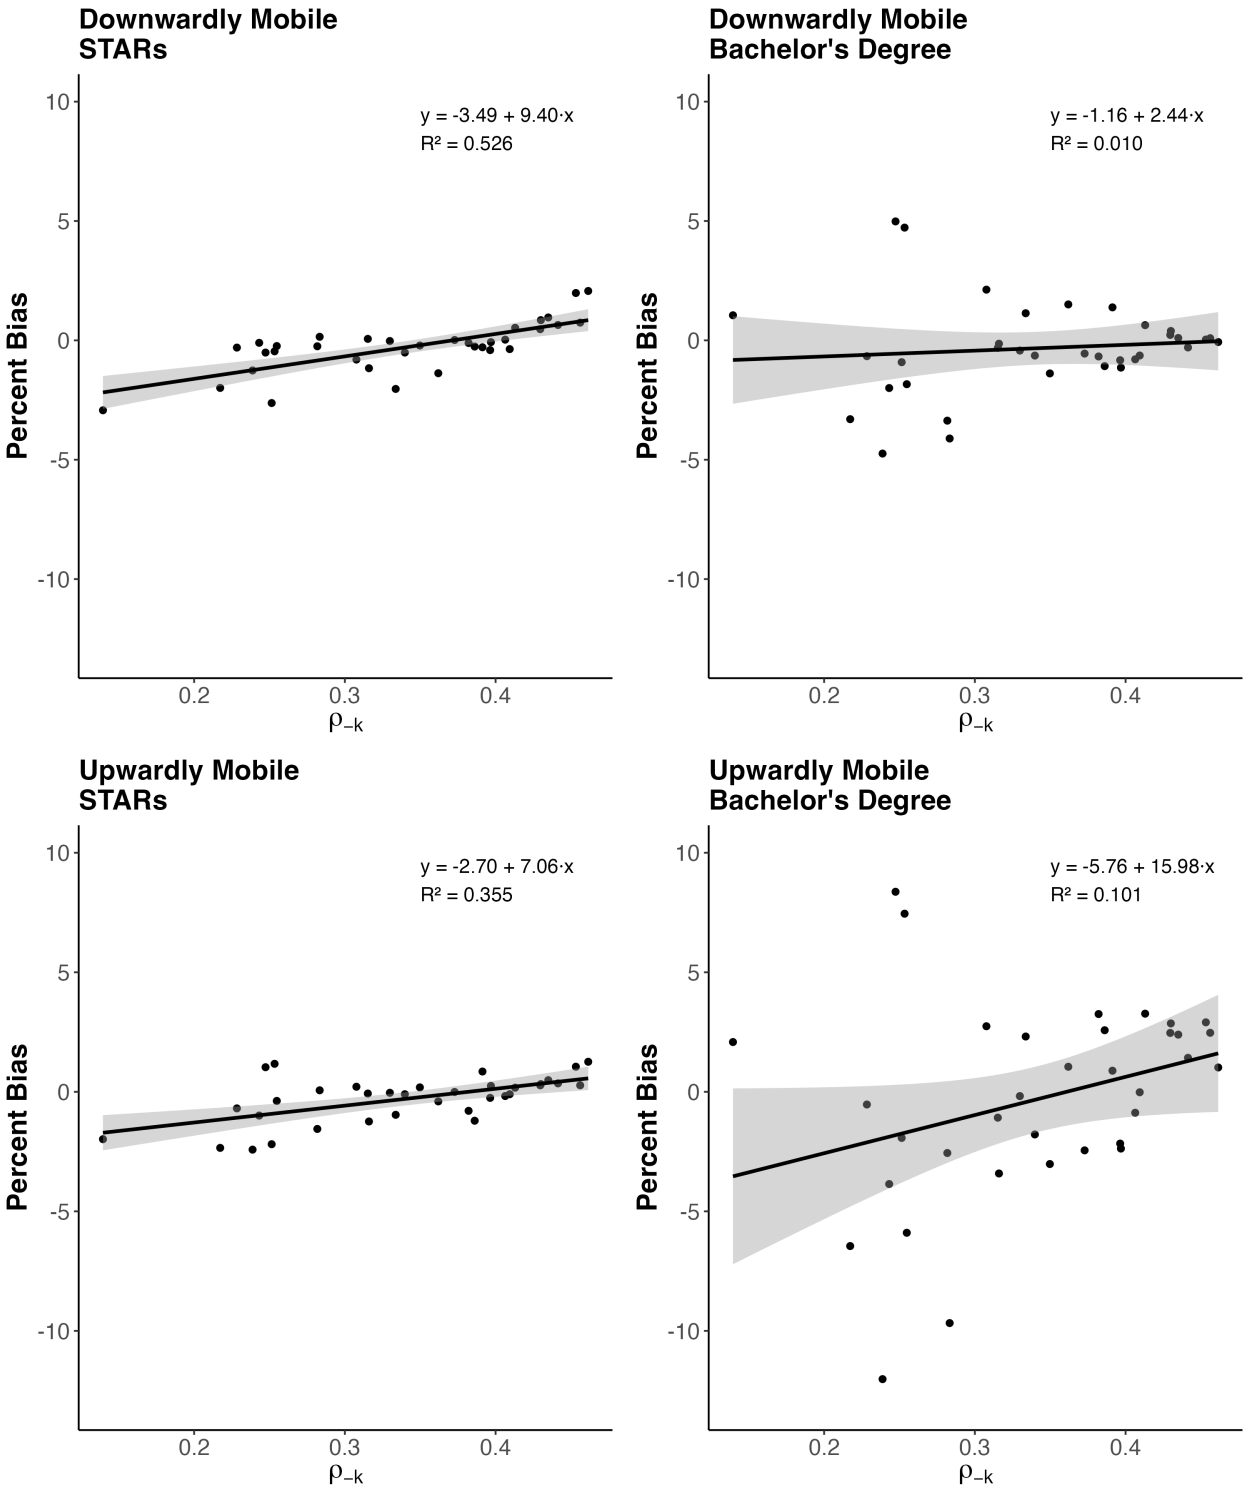

Fig. A5.2. Relationship Between Percent Bias and the Correlation Between Omitted and Included Skill Distances by Education and Transition Type

A6 Occupational Segregation and Selection Effects

A6.1 Occupational Composition

The differences in the absolute skill mobility frictions for workers with bachelor's degrees and STARs, particularly for transitions to higher wage occupations, may be explained in part by occupational segregation in the labor market by degree attainment. It is worth noting that this explanation is immediately hampered by the fact that these two groups of workers have very similar ASMFs when transitions require a lower skill distance or lead to lower wage occupations. However, as seen in Fig. A6.1, many occupations in the labor market are filled disproportionately by workers with a bachelor's degree or by workers without a bachelor's degree. As a result, STARs and workers with bachelor's degrees may be transitioning into or out of occupations that are exclusive to workers with a similar level of education.

Fig. A6.1. Distribution of occupations by the percent of workers with a bachelor's degree or higher, weighted by the number of workers in the occupation

In order to test this hypothesis directly, we recreated Fig. 4 in the paper after limiting the sample to origin and destination occupations with higher levels of educational diversity. Fig. A6.2 shows the relationship between the skill distance and flow rate for transitions for which the origin and destination occupations are made up of at least 10 percent of

workers with bachelors degrees and at least 10 percent of workers without bachelor's degrees. Fig. A6.3 limits the occupations of interest still further by requiring that at least 20 percent of workers have and do not have a bachelor's degree.

While reducing our data in this way increases the noise in these figures, the fundamental relationship between skill distance and flow rate remains the same across degree attainment and mobility type. STARs face greater skill mobility frictions when transitioning to higher wage occupations than when transitioning to lower wage occupations. In comparison, workers with bachelor's degrees experience less skill mobility friction when transitioning to higher wage work, especially when transitions have a particularly high skill distance. The findings in Fig. A6.2 and A6.3 rule out the hypothesis that workers with bachelor's degrees experience a positive ASMF for transitions with high skill distances because they are transitioning into occupations exclusive to worker with degrees (e.g. lawyers, doctors).

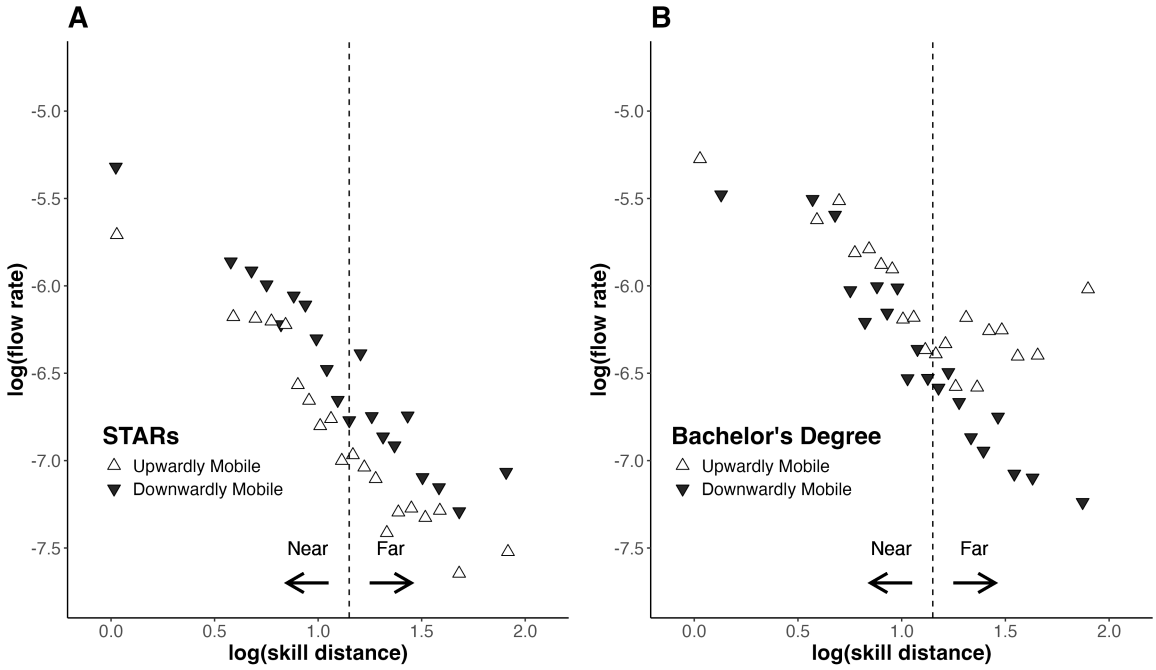

Fig. A6.2. Absolute skill mobility friction by degree attainment and mobility type for transitions into and out of occupations composed of at least 10% or more of both STARs and workers with bachelor's degrees

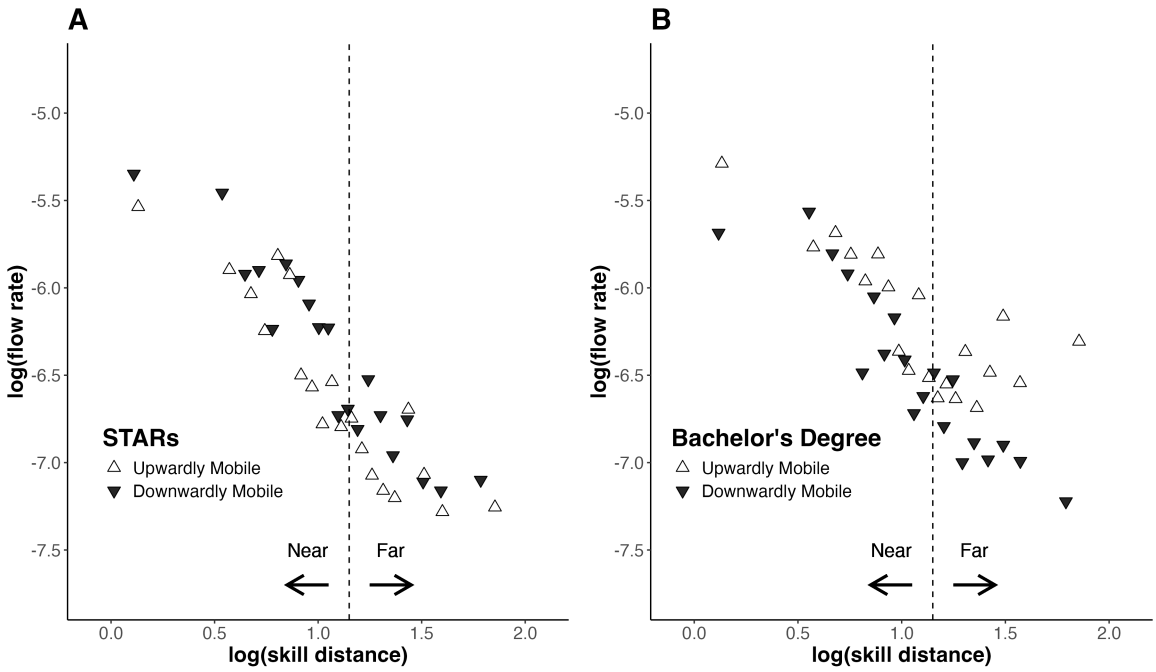

Fig. A6.3. Absolute skill mobility friction by degree attainment and mobility type for transitions into and out of occupations composed of at least 20% or more of both STARs and workers with bachelor's degrees

A6.2 Sample Size Constraints

Even with a pooled sample of 621,587 respondents for the years 2010 to 2019, only 60,382 respondents made a cross-occupational transition from year-to-year. Given that there are 422 occupations, the full set of possible permutations includes 177,662 cross-occupational transitions, though many of these are extremely unlikely. In practice, we observe 16,382 unique cross-occupational transitions which represent transitions made by over 100,000 workers. In a sample space so large, it is unsurprising that the modal cross-occupational transition is based on only one respondent, especially after disaggregating by education. However, as a result, it is not always possible to disentangle whether the resulting flow rate is low because a transition is unlikely or because the transition was not included in the sample

One potential remedy is to limit our sample to cross-occupational transitions with a larger number of unweighted respondents. We do so here, considering thresholds of 2 and 4 which represent the 61st and 84th percentile, respectively. Note that, as a result, the sample of cross-occupational transitions by education is reduced from 19,767 to 7,734 and then 3,219.

Even with underpowered samples, the findings are consistent with those presented in the main paper. The RSMF for downwardly mobile transitions is either negative or statistically indistinguishable from zero. In comparison, the RSMF for upwardly mobile transitions is consistently large and positive, suggesting that workers with bachelor's degrees

experience signficantly less friction than STARs when moving into higher wage occupations.

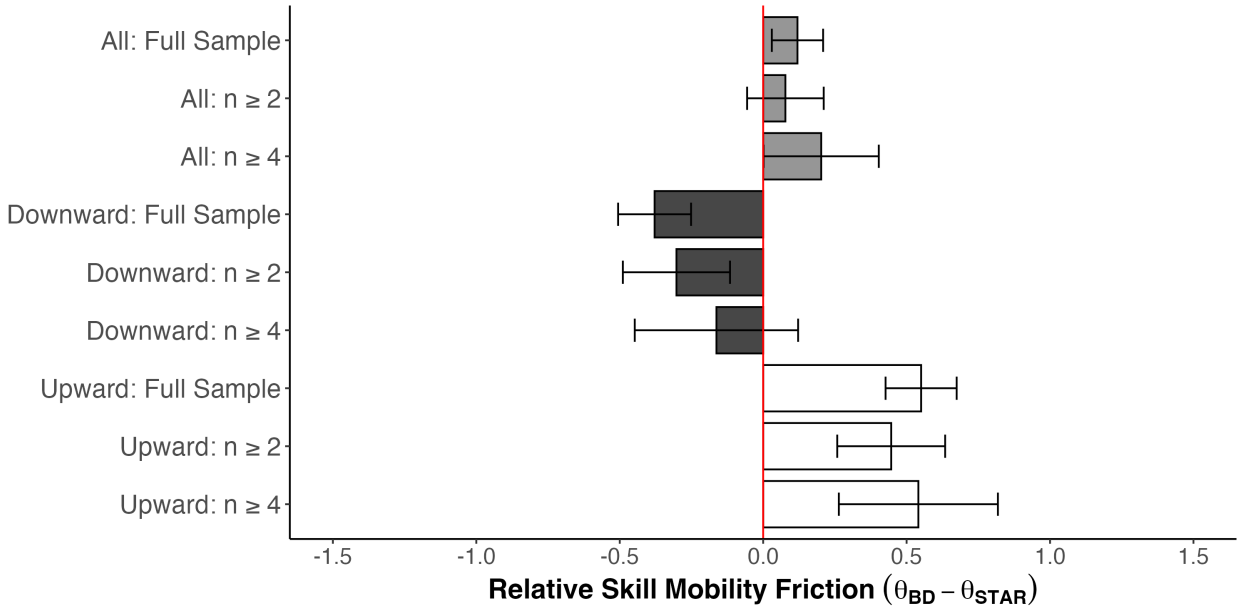

Fig. A6.4. Relative skill mobility friction by degree attainment and mobility type after limiting to transitions with larger numbers of respondents

A7 Skill Distance

A7.1 O*NET Skill Importance v. Skill Levels

In addition to the importance ratings that we use in this analysis, the O*NET provides level ratings for the 35 skills on a scale from 0 to 7 based on the degree to which a particular descriptor is required or needed to perform an occupation. Although it is theoretically possible for a skill to be given a high importance rating and a low level rating or vice versa, this does not occur in the data. The correlation coefficient between importance ratings and level ratings is 0.97. Further, the skill distances based on importance ratings and level ratings also have a correlation coefficient of 0.96. Fig. A7.1 shows the relationship between the importance and level ratings as well as the relationship between the Euclidean skill distances based on each skill rating.

Our results are robust to the use of either skill importance or skill level ratings. Fig. A7.2 shows the absolute skill mobility friction by degree attainment and mobility type using skill distances based on the skill level ratings. Fig. A7.3 compares the relative skill mobility friction estimates using the skill importance ratings and level ratings. Our conclusions are consistent regardless of this choice.

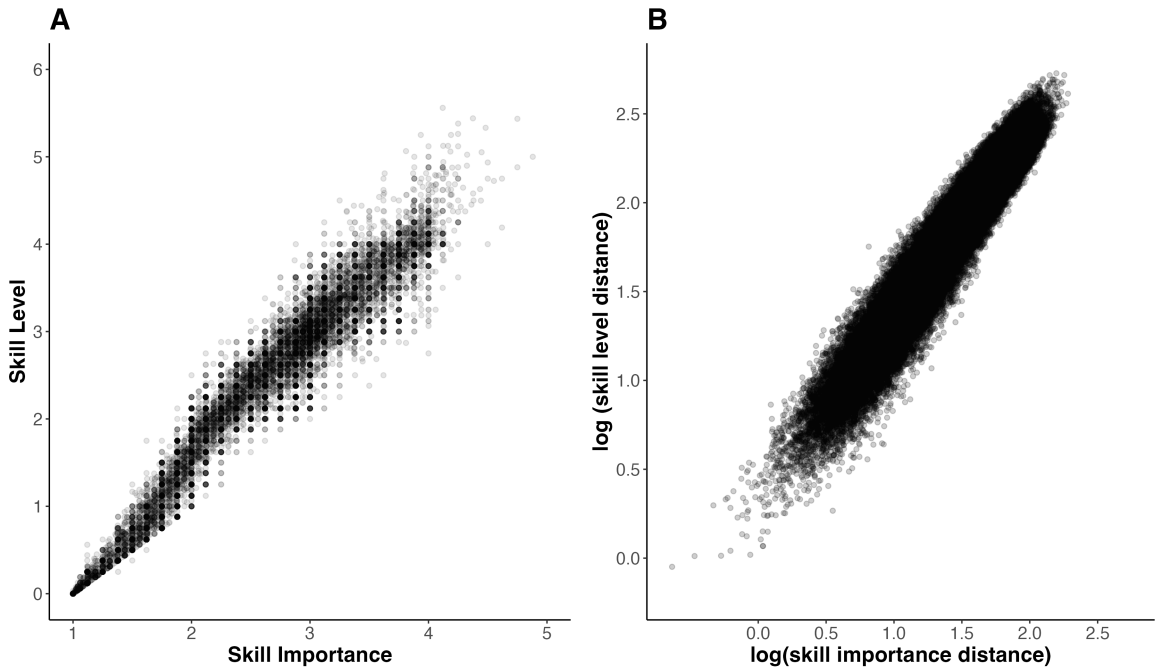

Fig. A7.1. Relationship between O*NET Skill Importance and Level Ratings

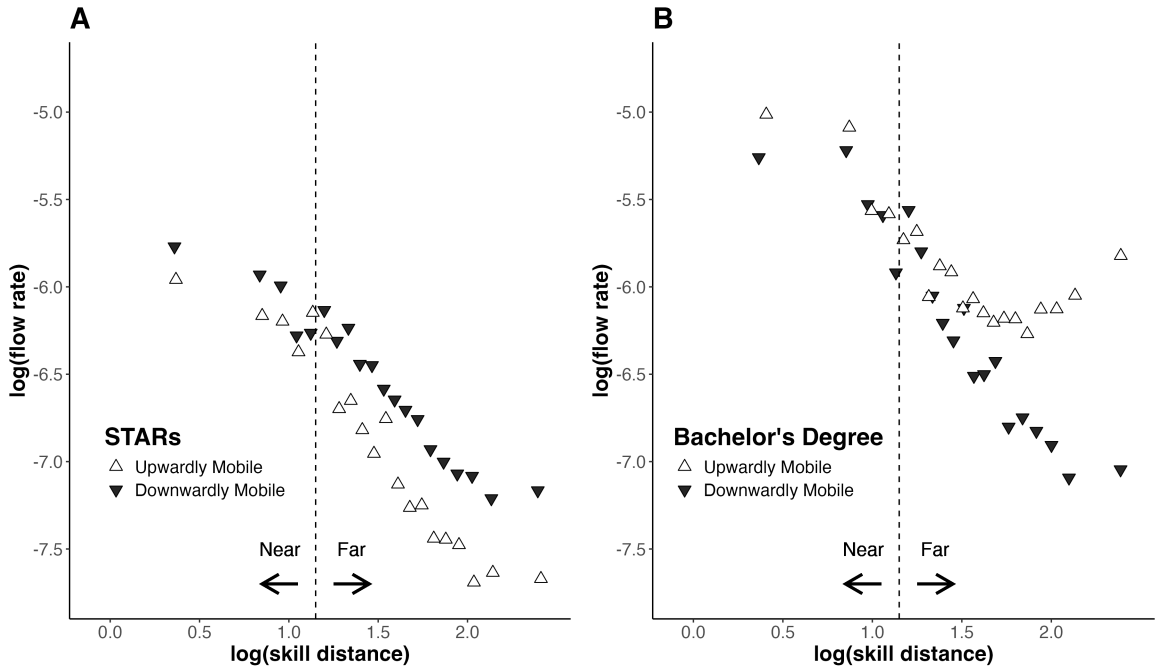

Fig. A7.2. Absolute skill mobility friction by degree attainment and mobility type using O*NET skill level ratings

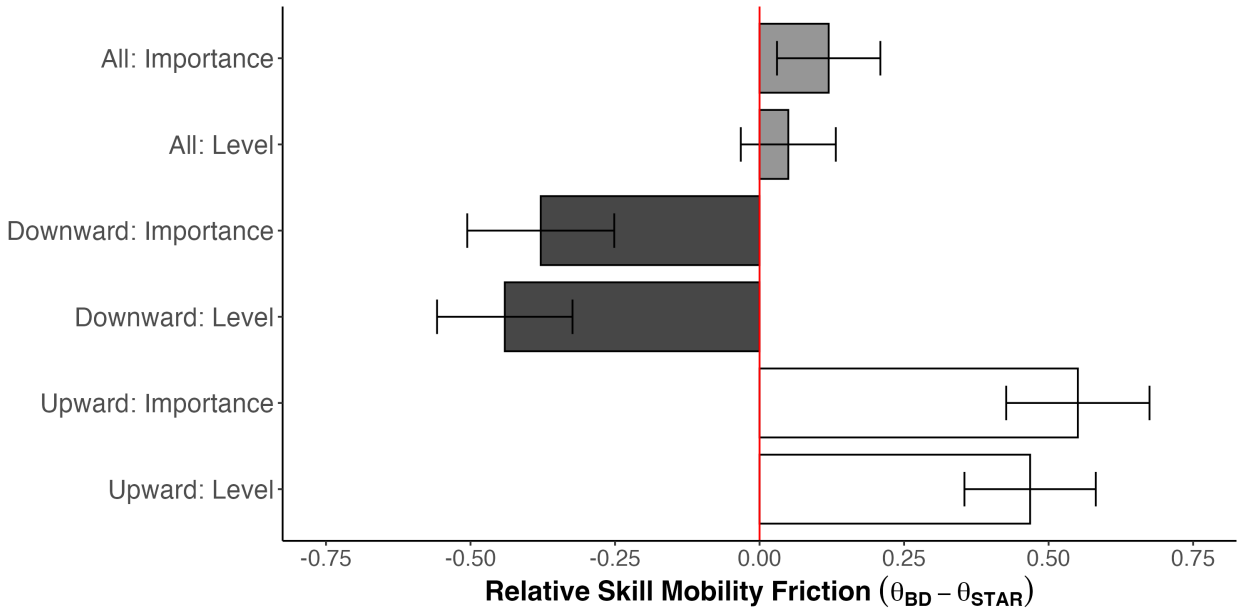

Fig. A7.3. Relative skill mobility friction by degree attainment and mobility type using skill importance and level ratings

A7.2 Alternative Measures of Skill Distance

Manhattan distance, or L_1 distance, has been used as an alternative to Euclidean distance, or L_2 distance [16]. As the number of dimensions increases, Euclidean distance can fail to discriminate between the furthest and nearest neighbor. For especially high-dimensional data, often in data mining and machine learning applications, the Manhattan skill distance can outperform Euclidean distance [51]. Manhattan skill distance can be calculated as follows:

$$d(\text{Occupation}_i, \text{Occupation}_j)_{\text{Manhattan}} = \sum_{k=1}^{35} \left| \text{Skill}_{k,\text{Occ.}_i} - \text{Skill}_{k,\text{Occ.}_j} \right| \tag{6}$$

An alternative way to measure the similarity of two objects with multiple attributes is to use cosine similarity. Unlike Euclidean or Manhattan distances, cosine similarity measures the degree to which two vectors point in the same direction and is insensitive to absolute differences in the magnitude of the two vectors. Cosine similarity can be calculated as follows:

$$s(\text{Occupation}_i, \text{Occupation}_j)_{\text{cosine}} = \frac{\sum_{k=1}^{35} \text{Skill}_{k,\text{Occ.}_i} \cdot \text{Skill}_{k,\text{Occ.}_j}}{\sqrt{\sum_{k=1}^{35} (\text{Skill}_{k,\text{Occ.}_i})^2} \sqrt{\sum_{k=1}^{35} (\text{Skill}_{k,\text{Occ.}_j})^2}} \tag{7}$$

For a measure that is directionally consistent with the distance-based measures, cosine

distance is often used as a complement of the cosine similarity such that larger values represent vectors that are more dissimilar. Cosine distance can be calculated as follows:

$$d(\text{Occupation}_i, \text{Occupation}_j)_{\text{cosine}} = 1 - s(\text{Occupation}_i, \text{Occupation}_j)_{\text{cosine}} \tag{8}$$

Although cosine similarity has been adopted in the past (DeMaria, Fee, and Wardrip 2020), we argue that an absolute measure of distance (i.e. Euclidean, Manhattan) is preferable to a relative measure of distance such as the cosine distance to understand the similarity of two skill portfolios. Consider a toy example with two workers who are rated from 1 to 5 for five dimensions of skill. Worker A has a rating of 2 for all five skills and worker B has a rating of 3 for all five skills. Both workers are applying for a job J in which employers prefer a rating of 4 for all five skills. Because the skill vectors of both workers are proportional to the skill vector of job J, worker A and worker B have a cosine skill distance of 0 from job J and would be equally likely to transition to job J. In comparison, the Euclidean and Manhattan distances both categorize worker A as being further than worker B from job J. When thinking about the likelihood that a worker's current skills allow them to move to a new occupation, we think this is a desirable property.

For most geographic applications, distance functions must meet a symmetry condition such that $d(A, B) = d(B, A)$. The Euclidean, Manhattan, and cosine distance functions each meet this condition. However, it is not obvious that a skill distance function should preserve this property. Consider another toy example with two occupations rated from 1 to 5 for five dimensions of skill. Occupation i has a rating of 2 for all five skills and occupation j has a rating of 4 for all five skills. A symmetric distance function would consider a transition from i to j to be the same distance as a transition from j to i. However, we might think that is easier for a worker in occupation j to transition to occupation i than the reverse. As a result, we may want to consider an asymmetric skill distance function that accounts for whether a worker is over- or under-skilled on any given skill dimension. To accomplish this, we consider an asymmetric Euclidean distance measure from occupation i to occupation j which is calculated as follows:

$$d(\text{Occupation}_i, \text{Occupation}_j)_{\text{asymmetric}} = \sqrt{\sum_{k=1}^{35} \left[\Delta\text{Skill}_{k,(i,j)}\right]^2} \tag{9}$$

where,

$$\Delta\text{Skill}_{k,(i,j)} \equiv= \begin{cases} \text{Skill}_{k,i} - \text{Skill}_{k,j}, & \text{if } \text{Skill}_{k,i} - \text{Skill}_{k,j} \leq 0 \\ 0.5 \cdot (\text{Skill}_{k,i} - \text{Skill}_{k,j}), & \text{if } \text{Skill}_{k,i} - \text{Skill}_{k,j} > 0 \end{cases} . \tag{10}$$

As written, this distance function gives less weight to the skill dimensions in which a worker has a higher level of skill than a destination occupation requires.

Alternatively, we could consider an asymmetric distance function that gives no weight to the skill dimensions in which a worker has a higher level of skill than a destination occupation requires. The theory behind this type of distance function is that workers should not be considered farther away from a job when they are overqualified on a given skill dimension. However, this type of measure overlooks the fact that, even when workers are overqualified on some skill dimensions, workers are more likely to transition to occupations that require the skills they have. Returning to the toy example, we think a worker in occupation j is more likely to move to occupation i than she is to move to an occupation with a rating of 1 for all five skill items.

Despite the differences in how each measure captures the similarity of multi-item vectors, all four skill distances are highly correlated with one another, as Fig. A7.4 shows.

Further, our analysis of the relationship between transition flow rates and skill distance is robust to the choice of distance measure. Fig. A7.5 is a replication of Fig. 4 in the paper and shows the absolute skill mobility friction by degree attainment and mobility type. In lieu of the Euclidean skill distance, Fig. S7.5A-B use the Manhattan skill distance, Fig. S7.5C-D use the cosine skill distance, and Fig. S7.5E-F use the asymmetric skill distance. Regardless of how we measure skill distance, the same pattern emerges: STARs experience more friction when moving to higher wage occupations and less friction when moving to lower wage occupations. Workers with bachelor's degrees experience the reverse, especially for transitions that have a high skill distance.

We also estimate the RSMF by transition type using all four measures of skill distance. Fig. A7.6 presents these results. The results using the Manhattan skill distance and asymmetric skill distance are highly consistent with the original results in the paper. Adopting the cosine skill distance reduces the magnitude of the RSMF for downwardly and upwardly mobile transitions, however, the estimates are directionally consistent and statistically significant.

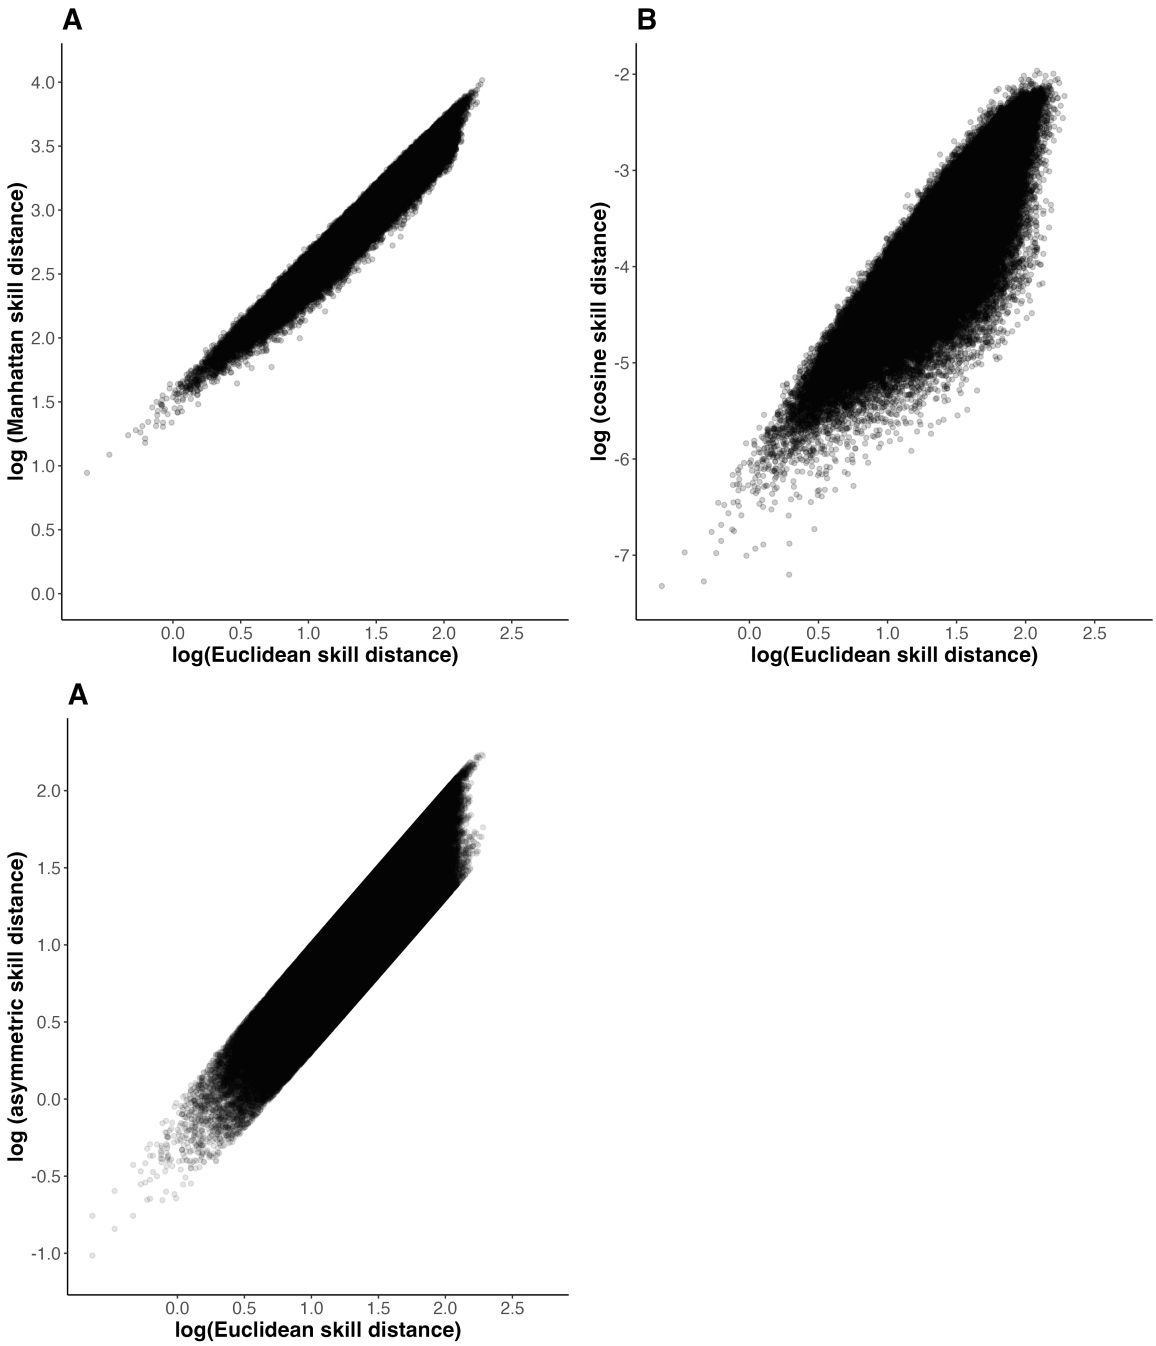

Fig. A7.4. Correlation between Euclidean, Manhattan, cosine, and asymmetric skill distances

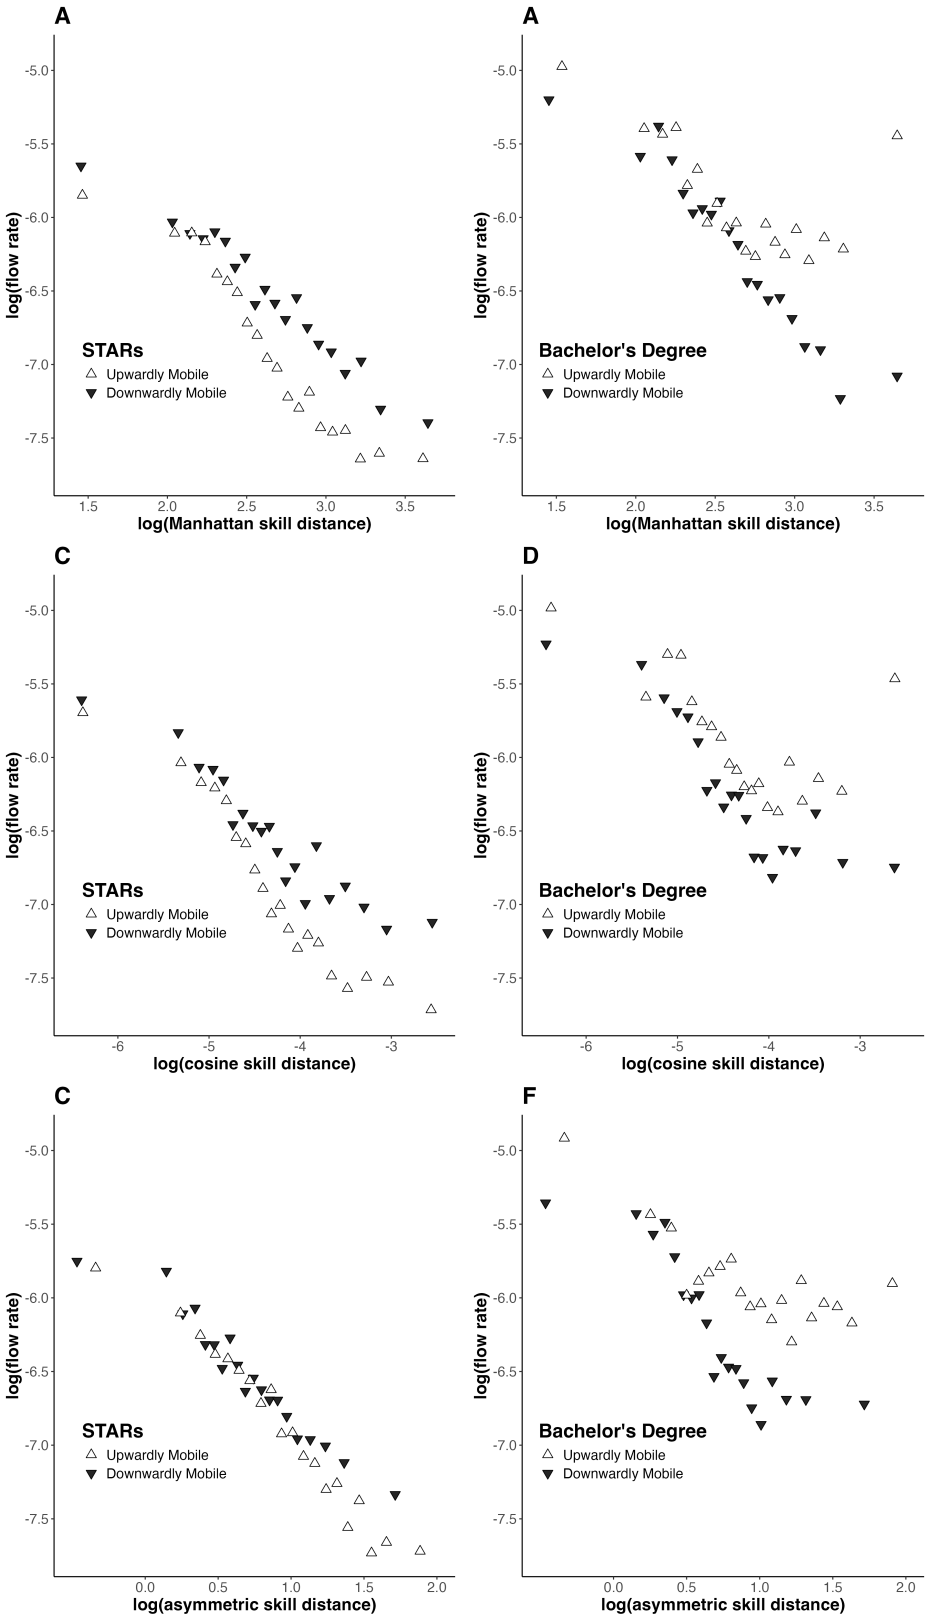

Fig. A7.5. Absolute skill mobility friction by degree attainment and mobility type using alternative measures of skill distance.

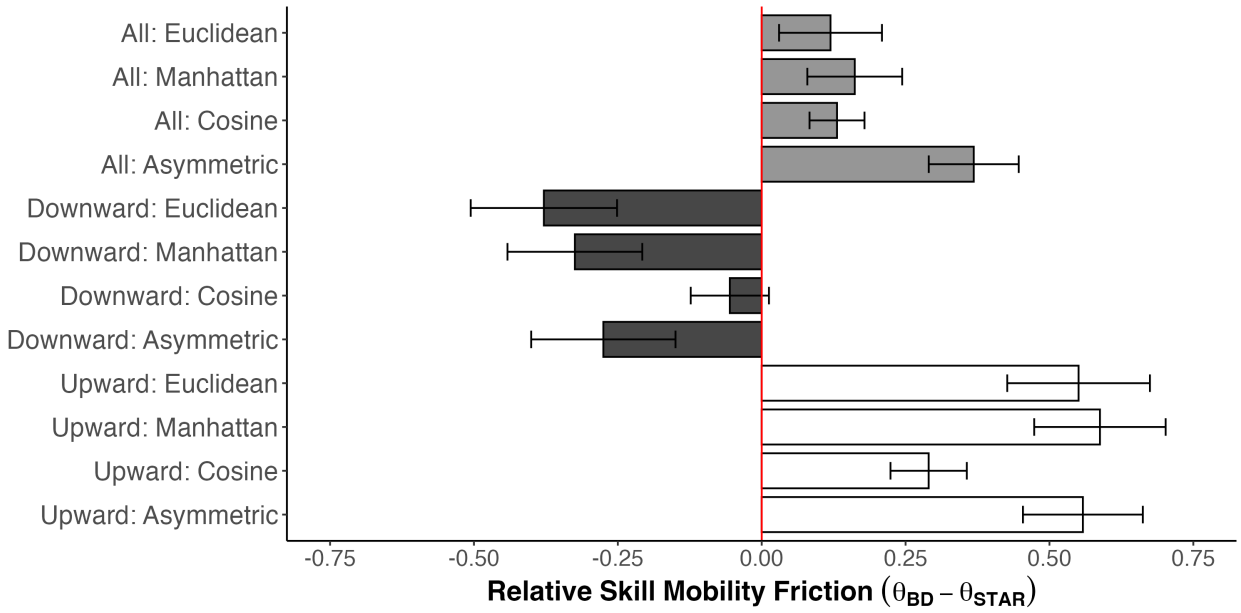

Fig. A7.6. Relative skill mobility friction by degree attainment and mobility type using alternative measures of skill distances

A7.3 Choosing the Skill Distance Threshold

Fig. 4 in the paper clearly shows that there is not a linear relationship between the log transformed skill distance and flow rate for workers with bachelor's degrees who are transitioning to higher wages. While this relationship is linear and negative for most other groups and types of transitions, this relationship has a positive slope for transitions with a large enough skill distance. In order to capture the difference in absolute skill mobility friction for STARs and workers with bachelor's degrees for transitions with higher skill distances, we categorize the transitions in our data as near transitions if the log skill distance was less than 1.15 and as far transitions if the log skill distance was greater than 1.15. We then estimate the RSMF separately for near and far transitions using the regressions in Table A3.1.

Fig. A7.7 shows the RSMF by transition type (upwardly or downwardly mobility) and skill distance (near or far) using a variety of skill distance cutpoints. Fig. S7.7A shows that our main finding is robust to our choice of threshold to distinguish between near and far transitions. Regardless of our choice, high skill distance transitions to higher wage jobs have a uniquely large and positive RSMF. We ultimately choose 1.15 as the cutpoint to use in our analysis because it minimizes the sum of the standard errors for the estimates of the RSMF for near and far transitions as can be seen in Fig. S7.7B.

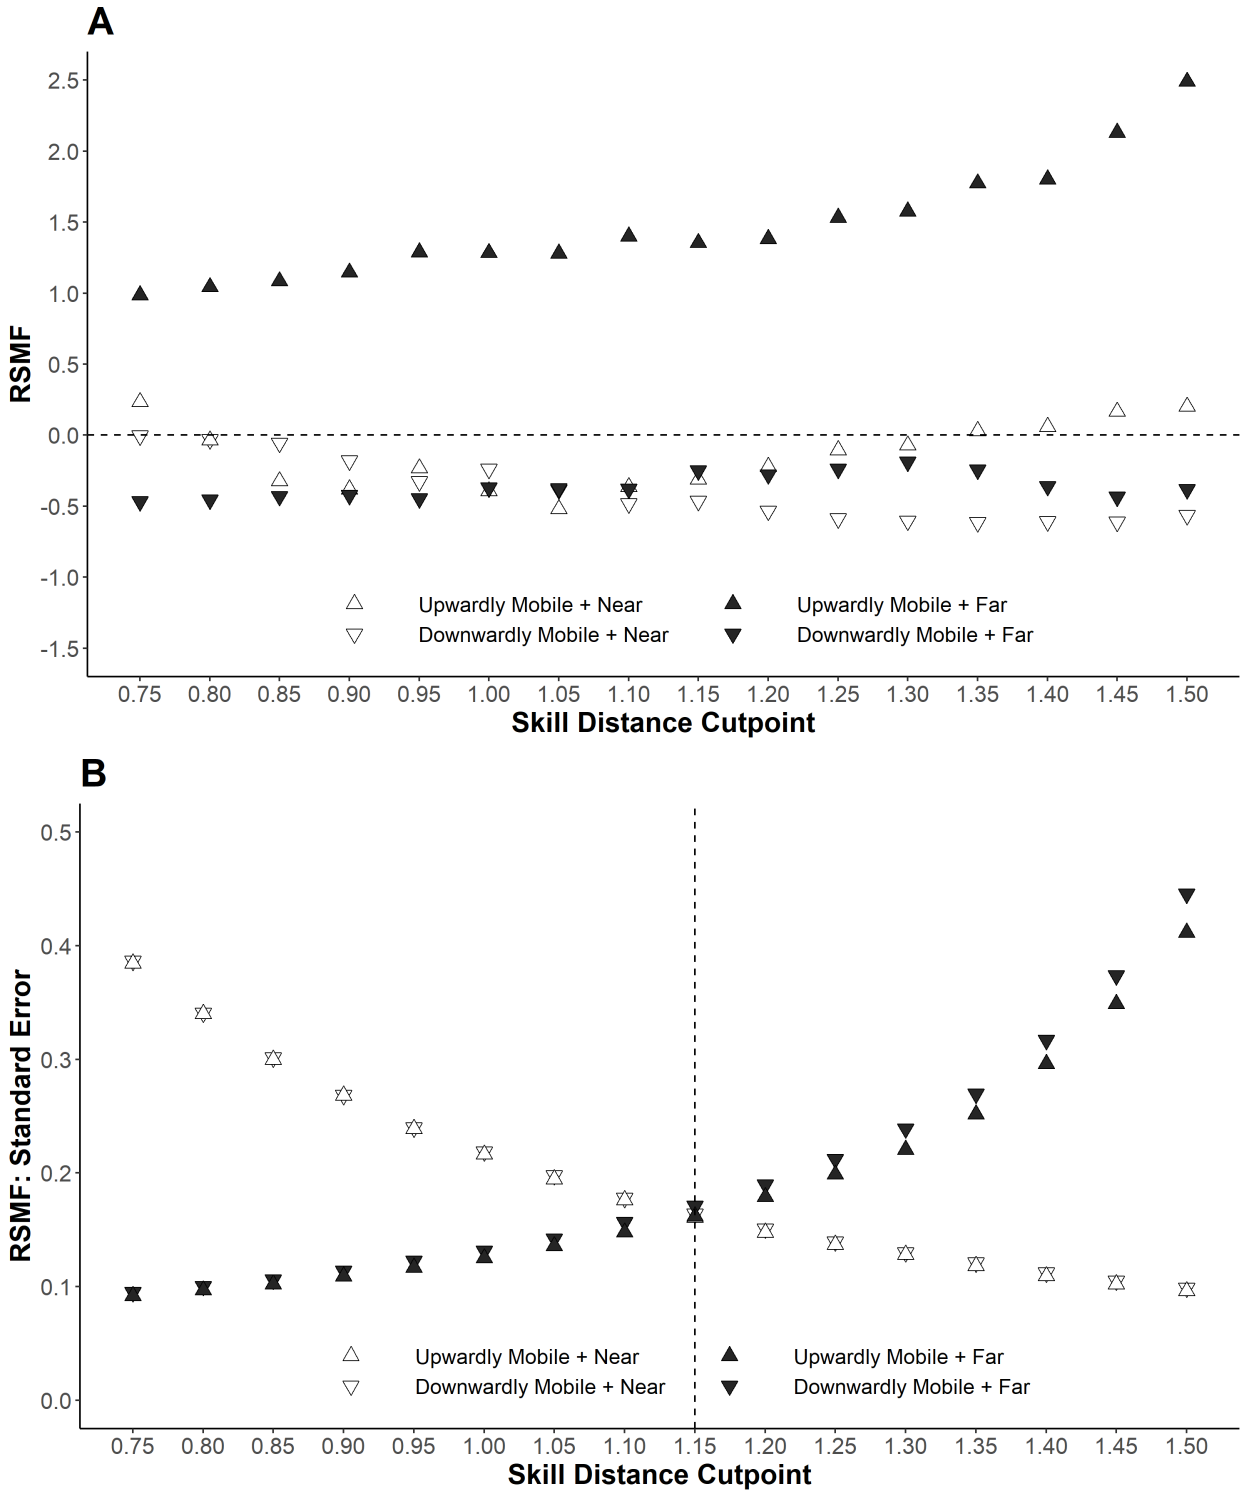

Fig. A7.7. Relative skill mobility friction by transition type and skill distance cutpoint to define near and far transitions.

A8 Tenure Considerations

The level of a worker's skills is likely to increase with additional tenure within an occupation or with an employer, which may confound observed differences in mobility frictions between STARs and workers with a bachelor's degree. Longer employer tenure can provide opportunities for skill development, internal promotions, and increased employer recognition of a worker's capabilities. As a result, workers with longer tenure could experience less skill mobility friction and move more easily into higher-paying positions that require a wider or deeper set of skills than their current role. Therefore, observed differences in opportunity between STARs and workers with bachelor's degrees may be influenced by differences in tenure.

To address the potential impact of variation in tenure by education on the likelihood of cross-occupational transitions, we use employer tenure reported in the CPS Employee Tenure and Occupational Mobility Supplement in combination with wage and occupational worker estimates obtained from the 2019 CPS 1-year ASEC. Because the Employee Tenure and Occupational Mobility Supplement is conducted every two years and on a different monthly cadence than the CPS ASEC, it is not possible to add tenure directly into our models of transition likelihood. In addition, the job tenure variable measures how long respondents have worked in their current job, not their current occupation. It is important to note that workers could have more years of experience in the same occupation at a previous employer.

Despite these limitations, we use the Employee Tenure and Occupational Mobility Supplement to explore potential differences in job tenure between STARs and workers with bachelor's degrees by occupation. First, we pool the supplements between 2000 and 2020 in order to generate a sufficient sample during the years of interest. Next, we calculate the median number of years the respondents remained in their current job by occupation and education.

Fig. A8.1 shows the relationship between the tenure premium and the wage premium among the 319 occupations for which tenure and wage data were available for both STARs and workers with a bachelor's degree. For each occupation, the tenure premium is calculated as the difference between the log of median job tenure for workers with a bachelor's degree and STARs. The wage premium is calculated as the difference between the log of the median wages. A linear regression line is plotted across the data as a visual representation of the overall trend.

No clear pattern or relationship emerges between these variables. First, in terms of tenure differences, the majority of occupations are clustered near zero because, on average, there are not significant differences in employer tenure between STARs and workers with a bachelor's degree. Second, linear regression does not reveal a linear relationship between the variables with a slope with a statistically insignificant coefficient of 0.001. The intercept of 0.171 is statistically significant, revealing a bachelor's degree wage premium that is consistent across occupations regardless of median job tenure by education. These

results are insensitive to alternative cuts, including limiting the occupations to those in which at least 20% of workers are STARs and 20% of workers have a bachelor's degree and limiting occupations to those in which median hourly wages are $25 or higher. Under these conditions, the slope remains statistically indistinguishable from zero and the bachelor's degree wage premium increases.

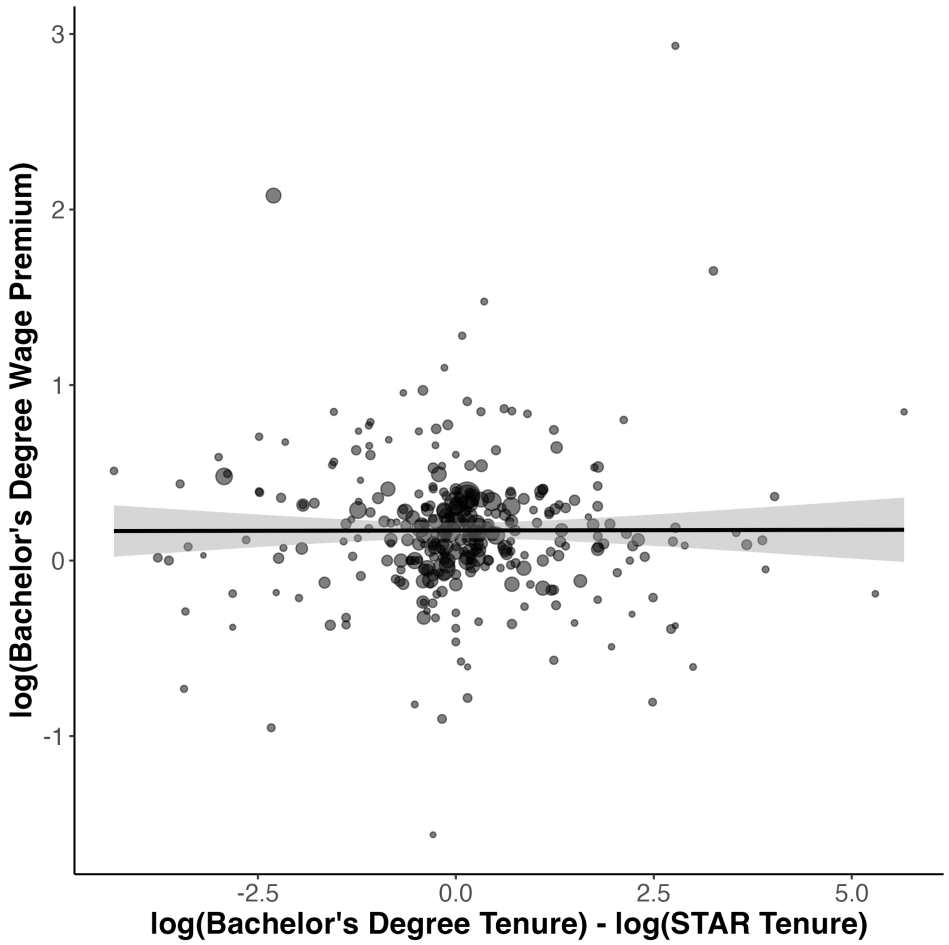

Fig. A8.1. Relationship between differences in job-level tenure and the bachelor's degree wage premium

A9 Routine Task Intensity

Given existing levels of occupational segregation by education, it is reasonable to wonder whether differences in task composition in the types of origin occupations that STARs and workers with bachelor's degrees start in might shape their likelihood of making upwardly mobile transitions. In particular, occupations with a high routine task intensity (RTI) might lead workers to develop a narrower and less adaptable skill set which may reduce the likelihood of upward mobility.

In order to test this hypothesis explicitly, we calculate an RTI measure using task data available from Autor and Dorn [45] (2013). This measure captures the degree to which an occupation involves routine, repetitive tasks that are susceptible to automation and offshoring. We categorize occupations into low RTI and high RTI groups based on whether their RTI scores are above or low the median, with high RTI occupations characterized by a greater share of routine task content.

We then calculate the RSMF separately for transitions out of occupations with a low RTI and high RTI to explore whether the differences in ASMF by education are potentially a result of a high proportion of STARs working in high RTI occupations. We present the findings in A9.1. Although there are marginal differences in the RSMF by RTI type, these differences are not statistically significant and they do not impact our core findings. There is no evidence to suggest that differences in the absolute skill mobility friction between STARs and workers with a bachelor's degree are the result of differences in the RTI of workers' starting occupations.

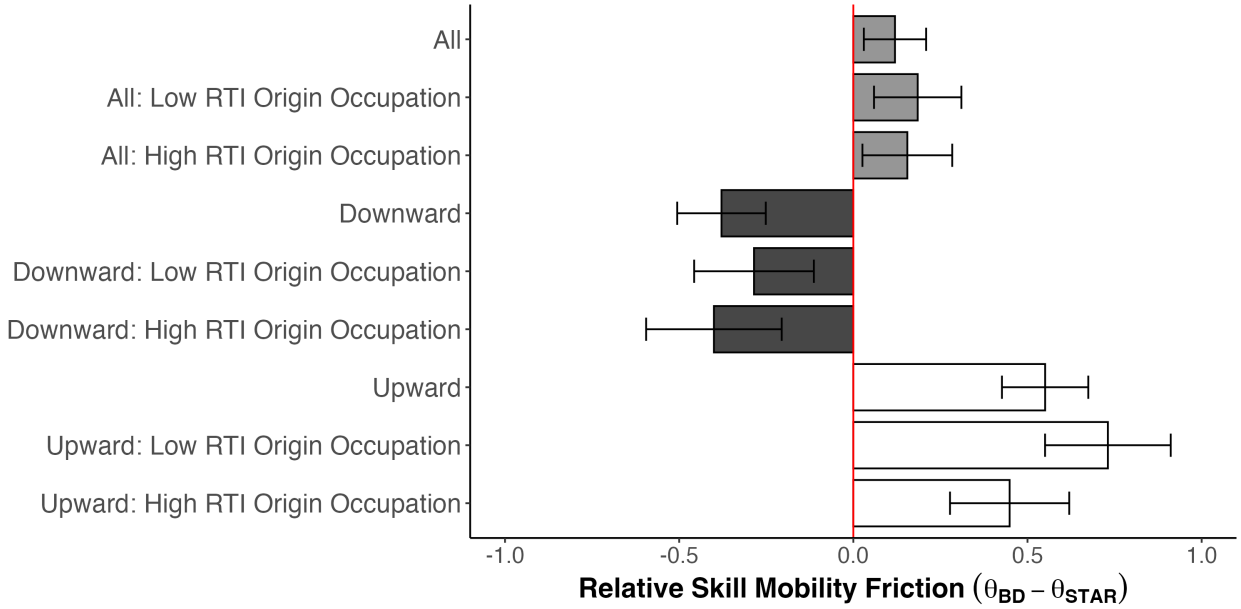

Fig. A9.1. Relative skill mobility friction by degree attainment and mobility type after limiting to transitions with larger numbers of respondents