

Online Appendix

Marrying Your Job: Matching and Mobility with Geographic Heterogeneity

Davide Alonzo*

May 19, 2025

Abstract

This Online Appendix provides information and analysis supporting the main text.

*Department of Economics, Université de Montréal, 3150 Rue Jean-Brillant, Montréal, QC H3T 1N8, Canada. E-mail:davidealonzo03@gmail.com.

C Additional Figures

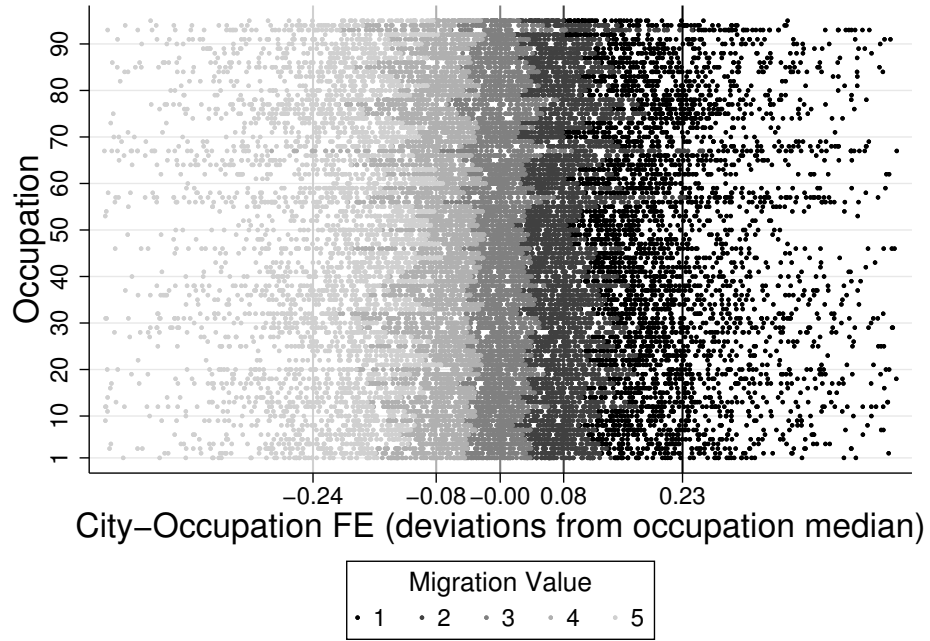


Figure 9: Estimated city-occupation fixed effects; deviations from the occupation median. Different shades of gray represent the five levels of migration value. The vertical lines indicate the averages for each group.




Figure 10: Graphic representation of occupation-specific city rankings. On the y-axis, cities are ordered based on average ranking across all occupations. Each point in the graph corresponds to a particular city-occupation pair and it is placed in correspondence to its ranking on the x-axis. The figure shows that there is a clear correlation by which some cities tend to pay higher wages to many occupations, but city effects do not explain all the variation.




Figure 11: Estimated probability of having married at least once as a function of age for men with the lowest and highest levels of migration value. The left panel shows the estimated probabilities while the right panel shows the difference between the two. For each observation i , I predict the average probability of marriage at different ages \hat{p}_i^a . The probability of marrying exactly at age \bar{a} is given by $\tilde{p}_i^{\bar{a}} = [\prod_{a < \bar{a}} (1 - \hat{p}_i^a)] \hat{p}_i^{\bar{a}}$. The estimated probability of marrying before age a is computed as $p_i^a = \sum_{\bar{a} \leq a} \tilde{p}_i^{\bar{a}}$. The values reported in this figure are obtained as the averages of p_i^a for the specified level of migration value.




Figure 12: Marital sorting as a function of the spouses' value of migration for all couples. The figure focuses on the diagonal values and reports the baseline measure with its confidence interval for reference.




Figure 13: This graphs shows a measure of marital sorting that is conditional on the spouses having the same education level (left panel: college; right panel: non-college). This measure compares the observed distribution of marriages across two dimensions, migration value and education, to the one that would prevail if marriage was random with respect to migration value but not with respect to education. Formally, I compute

$$s_{h,w} = \frac{P(h, w, H, W)}{\sum_c P(h|H, c)P(w|W, c)P(H, W|c)P(c)}$$

where H and W represent the education of husbands and wives, respectively. The figure shows assortative mating for educationally homogeneous couples and focuses on the diagonal values ($h = w$).

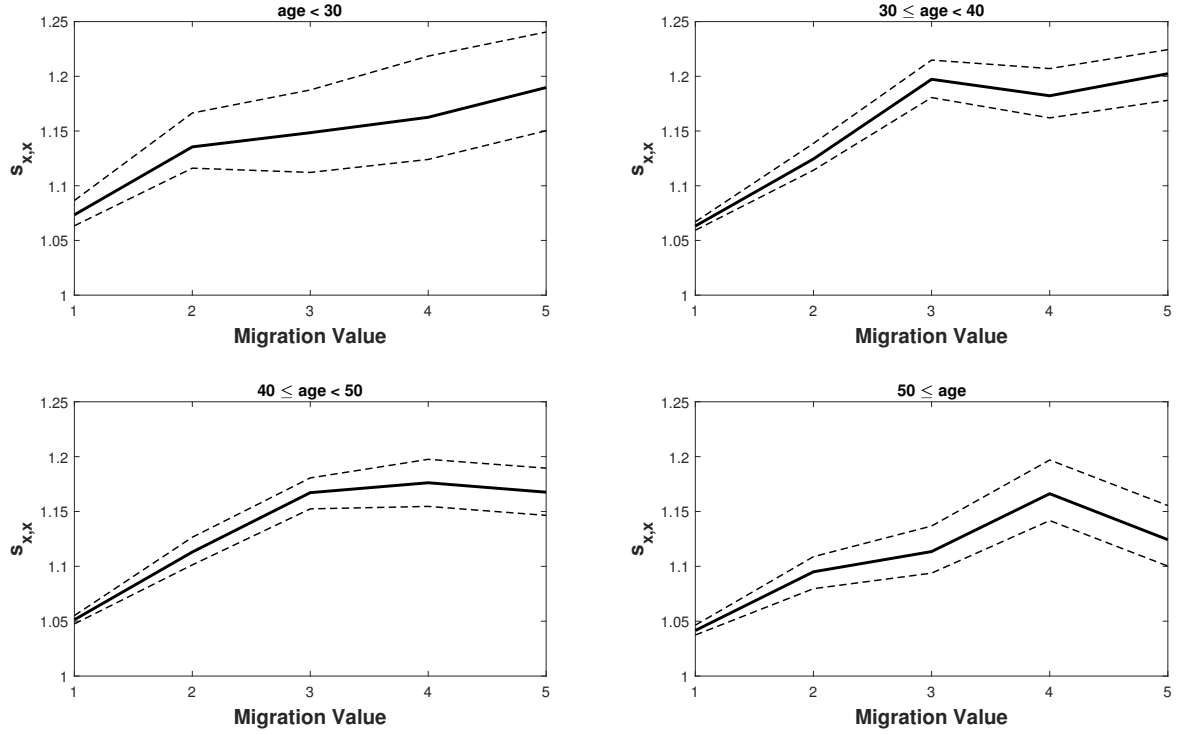


Figure 14: This graphs shows a measure of marital sorting that is conditional on the spouses belonging to the same age group. This measure compares the observed distribution of marriages across two dimensions, migration value and age, to the one that would prevail if marriage was random with respect to migration value but not with respect to age. Formally, I compute

$$s_{h,w} = \frac{P(h, w, H, W)}{\sum_c P(h|H, c)P(w|W, c)P(H, W|c)P(c)}$$

where H and W represent the age group of husbands and wives, respectively. The figure shows assortative mating for couples of the same age group and focuses on the diagonal values ($h = w$).




Figure 15: Average (negative) change in migration value upon migration for recent divorcees for men (left) and (women). The solid line shows what this change would look like if the city of destination was chosen at random. The main takeaway is that recent divorcees who migrate are more likely to relocate to cities paying higher wages to their occupation which implies that their migration choices respond to pecuniary incentives just as those of singles who have never married.




Figure 16: Comparison of occupation shares from the ACS sample to the stationary distribution implied by the transition probabilities from the CPS sample. The latter coincides by construction with the occupational shares in the model.




Figure 17: Model fit: non-targeted moments relative to migration and marriage patterns of women as a function of migration value.




Figure 18: The effects of a tax incentive to married households. Left: fraction of the total population residing in each city. Right: change in the geographic concentration (HH index) of employment; each bubble is an occupation and the size of the bubble is proportional to the total employment in such occupation.

D Additional Tables

		Women			Total
		White	Black	Other	
Men	White	51,502	476	3,585	55,563
		71%	0%	5%	76%
	Black	1,143	4,789	461	6,393
		2%	7%	1%	10%
	Other	2,489	164	7,372	10,025
		4%	0%	10%	14%
Total		55,134	5,429	11,418	71,981
		77%	7%	16%	100%

Table 9: Racial composition of the 71,981 marriages observed in the ACS sample.

Yearly Probability of Marriage (Odds-ratios) - Men				
Migration Value 1	1 (0)	1 (0)	1 (0)	1 (0)
Migration Value 2	0.928*** (0.0158)	0.953*** (0.0168)	0.939*** (0.0176)	0.926*** (0.0154)
Migration Value 3	0.915*** (0.0173)	0.933*** (0.0187)	0.923*** (0.0197)	0.916*** (0.0170)
Migration Value 4	0.906*** (0.0205)	0.968 (0.0235)	0.914*** (0.0229)	0.903*** (0.0200)
Migration Value 5	0.826*** (0.0240)	0.908*** (0.0283)	0.827*** (0.0262)	0.832*** (0.0236)
Observations	313,745	313,369	313,745	329,826
Occ. FE	NO	YES	NO	NO
State FE	NO	NO	YES	NO

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 10: Estimates of the marriage equation for men. The table reports the odds-ratios relative to an individual who has a low migration value (level 1). Controls include: a quartic in age; the logarithm of city size; the logarithm of the city-level sex ratio; the logarithm of own wage; the standard deviation of the cross-city distribution of wage fixed effects for own occupation; a dummy for the presence of children in the household; a college dummy; year fixed effects. The second column adds occupation fixed effects while the third controls for state fixed effects. The fourth column includes men who are out of the labor force for which an occupation is observed in the data.

Yearly Probability of Marriage (Odds-ratios) - Women				
Migration Value 1	1 (0)	1 (0)	1 (0)	1 (0)
Migration Value 2	0.985 (0.0173)	1.002 (0.0185)	1.000 (0.0200)	0.987 (0.0168)
Migration Value 3	1.014 (0.0197)	1.001 (0.0208)	1.024 (0.0236)	1.010 (0.0191)
Migration Value 4	0.988 (0.0240)	0.954* (0.0249)	0.993 (0.0274)	0.983 (0.0231)
Migration Value 5	0.957 (0.0313)	0.880*** (0.0305)	0.969 (0.0347)	0.940* (0.0297)
Observations	278,232	278,055	278,232	291,622
Occ. FE	NO	YES	NO	NO
State FE	NO	NO	YES	NO
Robust standard errors in parentheses				
*** p<0.01, ** p<0.05, * p<0.1				

Table 11: Estimates of the marriage equation for women. The table reports the odds-ratios relative to an individual who has a low migration value (level 1). Controls include: a quartic in age; the logarithm of city size; the logarithm of the city-level sex ratio; the logarithm of own wage; the standard deviation of the cross-city distribution of wage fixed effects for own occupation; a dummy for the presence of children in the household; a college dummy; year fixed effects. The second column adds occupation fixed effects while the third controls for state fixed effects. The fourth column includes women who are out of the labor force for which an occupation is observed in the data.

Probability of Marriage Within Occupation (Odds-ratios) - Men

Migration Value 1	1 (0)	1 (0)	1 (0)	1 (0)
Migration Value 2	1.062 (0.0649)	1.053 (0.0657)	1.091 (0.0764)	1.051 (0.0638)
Migration Value 3	1.289*** (0.0852)	1.228*** (0.0834)	1.330*** (0.102)	1.302*** (0.0848)
Migration Value 4	1.392*** (0.113)	1.240** (0.105)	1.420*** (0.131)	1.400*** (0.112)
Migration Value 5	1.492*** (0.156)	1.221* (0.134)	1.515*** (0.175)	1.492*** (0.154)
Observations	23,653	23,370	23,635	24,448
Occ. FE	NO	YES	NO	NO
State FE	NO	NO	YES	NO

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 12: Estimates of the equation for the probability of marrying within occupation for men. The table reports the odds-ratios relative to an individual who has a low migration value (level 1). Controls include: the logarithm of the city-level fraction of women working in the same occupation; a college dummy for each of the spouses and their interaction; the logarithm of the wage of both spouses and their interaction; year fixed effects. The second column adds occupation fixed effects while the third controls for state fixed effects. The fourth column includes men who are out of the labor force for which an occupation is observed in the data.

Probability of Marriage Within Occupation (Odds-ratios) - Women				
Migration Value 1	1 (0)	1 (0)	1 (0)	1 (0)
Migration Value 2	1.108 (0.0699)	1.105 (0.0715)	1.089 (0.0818)	1.114* (0.0685)
Migration Value 3	1.318*** (0.0887)	1.165** (0.0801)	1.271*** (0.107)	1.324*** (0.0866)
Migration Value 4	1.558*** (0.133)	1.361*** (0.121)	1.532*** (0.157)	1.596*** (0.132)
Migration Value 5	1.962*** (0.222)	1.559*** (0.183)	1.845*** (0.235)	1.960*** (0.214)
Observations	19,631	19,091	19,626	20,832
Occ. FE	NO	YES	NO	NO
State FE	NO	NO	YES	NO

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

Table 13: Estimates of the equation for the probability of marrying within occupation for women. The table reports the odds-ratios relative to an individual who has a low migration value (level 1). Controls include: the logarithm of the city-level fraction of men working in the same occupation; a college dummy for each of the spouses and their interaction; the logarithm of the wage of both spouses and their interaction; year fixed effects. The second column adds occupation fixed effects while the third controls for state fixed effects. The fourth column includes women who are out of the labor force for which an occupation is observed in the data.

Yearly Probability of Divorce (Odds-ratios) - Men				
Migration Value 1	1 (0)	1 (0)	1 (0)	1 (0)
Migration Value 2	1.143*** (0.0234)	1.131*** (0.0233)	1.057** (0.0242)	1.152*** (0.0231)
Migration Value 3	1.164*** (0.0255)	1.134*** (0.0251)	1.063** (0.0266)	1.164*** (0.0249)
Migration Value 4	1.200*** (0.0297)	1.136*** (0.0287)	1.089*** (0.0306)	1.196*** (0.0290)
Migration Value 5	1.147*** (0.0347)	1.085*** (0.0333)	1.039 (0.0348)	1.141*** (0.0337)
Observations	1,370,554	1,310,532	1,370,554	1,413,718
Occ. FE	NO	YES	NO	NO
State FE	NO	NO	YES	NO
Robust standard errors in parentheses				
*** p<0.01, ** p<0.05, * p<0.1				

Table 14: Estimates of the divorce equation for men. The table reports the odds-ratios relative to an individual who has a low migration value (level 1). Controls include: a quartic in age; the logarithm of own wage; the standard deviation of the cross-city distribution of wage fixed effects for own occupation; a dummy for the presence of children in the household; a college dummy; year fixed effects. The second column adds occupation fixed effects while the third controls for state fixed effects. The fourth column includes men who are out of the labor force for which an occupation is observed in the data.

Yearly Probability of Divorce (Odds-ratios) - Women				
Migration Value 1	1 (0)	1 (0)	1 (0)	1 (0)
Migration Value 2	1.080*** (0.0199)	1.059*** (0.0196)	0.993 (0.0210)	1.082*** (0.0194)
Migration Value 3	1.115*** (0.0217)	1.109*** (0.0218)	0.981 (0.0226)	1.122*** (0.0213)
Migration Value 4	1.119*** (0.0251)	1.121*** (0.0253)	0.990 (0.0259)	1.128*** (0.0246)
Migration Value 5	1.102*** (0.0302)	1.138*** (0.0314)	0.958 (0.0295)	1.099*** (0.0293)
Observations	1,073,317	1,060,785	1,073,317	1,146,909
Occ. FE	NO	YES	NO	NO
State FE	NO	NO	YES	NO
Robust standard errors in parentheses				
*** p<0.01, ** p<0.05, * p<0.1				

Table 15: Estimates of the divorce equation for women. The table reports the odds-ratios relative to an individual who has a low migration value (level 1). Controls include: a quartic in age; the logarithm of own wage; the standard deviation of the cross-city distribution of wage fixed effects for own occupation; a dummy for the presence of children in the household; a college dummy; year fixed effects. The second column adds occupation fixed effects while the third controls for state fixed effects. The fourth column includes women who are out of the labor force for which an occupation is observed in the data.

	Recent Divorcees vs. Singles	All Divorcees vs. Singles	Recent Divorcees vs. Married	All Divorcees vs. Married
Men	1.95	1.37	3.46	2.38
Women	2.19	1.41	4.18	3.28

Table 16: Regression-based odds-ratios for the probability of migration comparing different groups. Each column compares the migration probabilities of recent divorcees (less than one year) and all divorcees to that of single or married households. I estimate a logit model for the probability of migration on the relevant sub-sample including a dummy for divorcees. The results clearly show that the yearly probability of migration is higher for divorcees when compared to both single and married households. Moreover, the difference is stronger when considering only recent divorcees. The estimated logit model contain controls for age, wages, geographic dispersion of occupation-specific wages, presence of children, and education. All the estimated odds-ratios are significant at 0.1% or less.

Moment	Model	Data
Average probability of migration (single women)	0.95%	0.71%
Average probability of marriage (women)	9.29%	10.22%
Average probability of divorce (women)	1.61%	1.69%

Table 17: Additional non-targeted moments.

Fraction of occupationally homogeneous marriages					
	Baseline (%)	SR (%)	% Change	LR (%)	% Change
All					
	11.0	10.8	-2.0	10.8	-2.4
Migration Value					
1	12.0	11.8	-1.4	11.6	-3.2
2	10.0	9.7	-2.6	9.7	-2.7
3	7.8	7.6	-3.5	7.8	-0.1
4	9.6	9.4	-2.4	10.1	+5.5
5	9.2	8.6	-6.2	10.3	+11.3
Men					
Migration Value					
1	11.4	11.2	-1.4	10.9	-4.2
2	10.6	10.4	-2.6	10.3	-2.8
3	9.6	9.2	-3.9	9.8	+1.6
4	10.6	10.3	-2.2	11.4	+8.1
5	9.8	9.1	-7.8	11.2	+14.3
Women					
Migration Value					
1	11.4	11.2	-1.4	10.9	-4.2
2	10.6	10.4	-2.6	10.3	-2.8
3	9.6	9.2	-3.9	9.8	+1.6
4	10.6	10.3	-2.2	11.4	+8.1
5	9.8	9.1	-7.8	11.2	+14.3

Table 18: Yearly fraction of marriages involving spouses employed in the same occupation in the baseline and counterfactual scenario where migration is not allowed. In aggregate, the absence of migration reduces the incentive to enter same-occupation marriages both in the short and in the long-run. In the short-run, the biggest drop occurs among high-value workers, which suggests that many of these marriages in the baseline model are sustained by the prospect of migration.

E Validating the measure of the value of migration

In this appendix, I validate my measure of migration value as a proxy for the option value of migration by showing that high levels of this measure are associated with a higher probability of migration.

The ACS surveys contain information regarding the city of residence of respondents in the year preceding the survey. Using this information, I construct a dummy variable that equals one if the household migrated (if the current city of residence is different from the previous one) and zero otherwise. I estimate a logit model for the probability of migration to analyze how the latter changes with migration value.²⁹ In the regression, I control for the presence of children, a quadratic function of age, and a set of dummies for education and year fixed effects. The analysis is carried out separately for men and women.

The left panels of Figure 19 show the average probabilities of migration obtained from the estimated logit model as a function of the five levels of migration value. The probability of migrating is monotonically increasing in migration value. For singles, the yearly migration rate is about 0.8% for men and 0.7% for women, but high-value workers are about twice more likely to migrate than the low-value. To get a sense of the direction of migration, the right panels of Figure 19 display the negative average change of migration value upon migration. The negative sign allows for the interpretation of a positive value as a movement towards higher-ranked cities. The solid line represents the counterfactual pattern that one would observe if migration was random. The fact that the estimated changes are always above the random case is evidence that migration is on average directed toward cities that pay higher wages conditional on one's own occupation.

For more details and robustness, Tables 19 and 20 show all the regression results with occupation and state fixed effects and robustness using an extended sample including all observations for which an occupation is observed (leftmost column).³⁰ Finally in Table 21 the same analysis is carried out jointly for men and women. The larger sample size allows for the inclusion of city fixed effects as additional controls (fourth column).

Finally, I also estimate a version of the logit model for married households that flexibly estimates the migration probabilities as a function of the migration value of both husbands and wives. Figure 20 shows the predicted average probability of migration for each combination of migration values of husbands and wives. Overall, migration probabilities are increasing in the migration value of both spouses.

²⁹This measure does not take into account the possibility of occupational changes during migration. In Online Appendix G, I argue that this is not a major concern of the analysis.

³⁰In the ACS respondents are asked to report either their current occupation or the most recent in the previous 5 years.

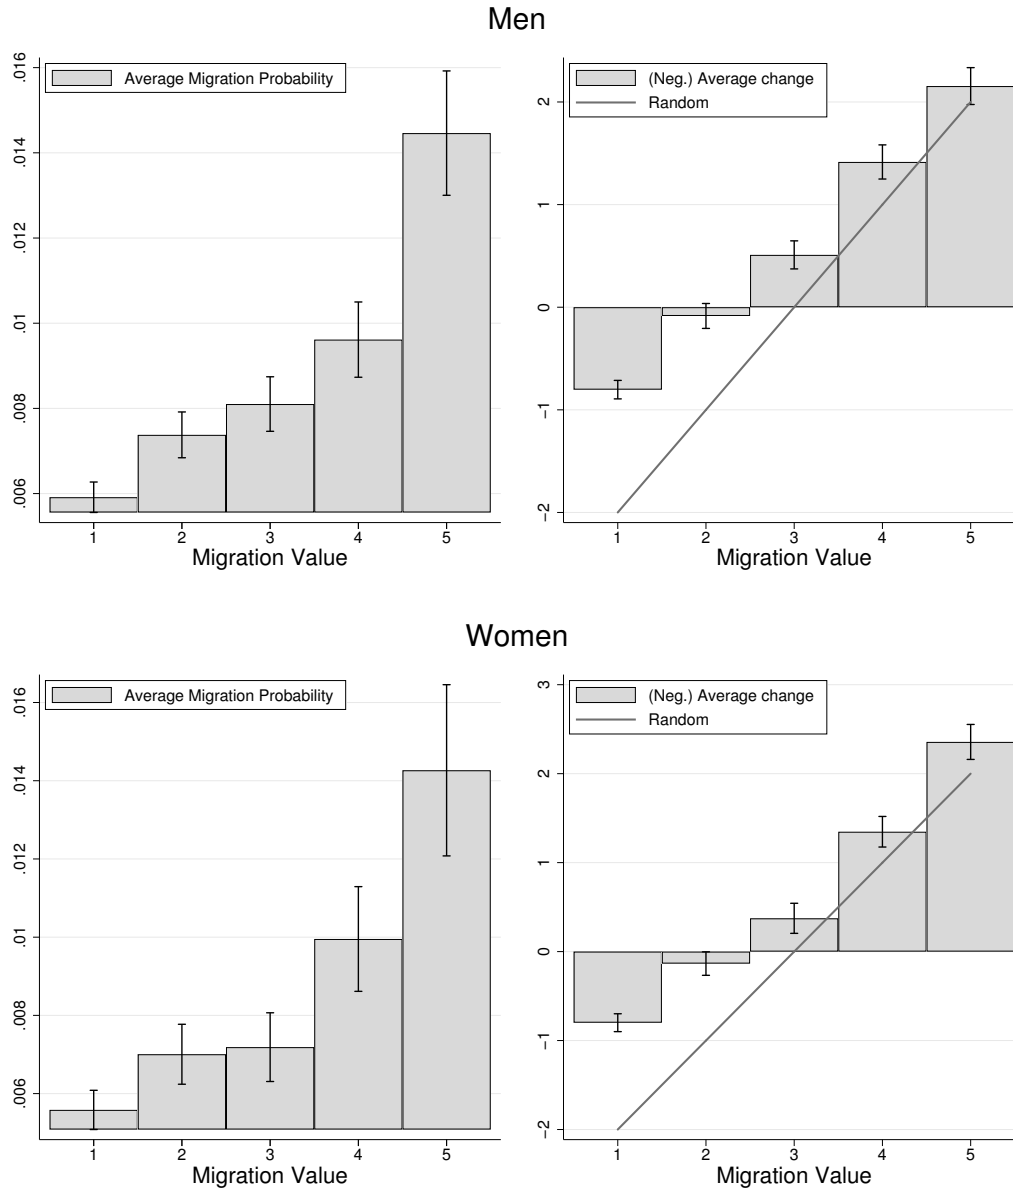


Figure 19: The left panels show the average migration rates of single men and women as a function of the value of migration in the origin city. For both men and women, migration propensities increase with the value of the migration option. The right panels show the average (negative) change in the value of future migration upon migrating. A positive value indicates that workers tend to move to cities paying higher wages to their occupation (since they move to better cities, the pecuniary value of future migration opportunities falls). The solid line represents what this change would look like if migration was random. The fact that the actual values are above the solid line suggests that, on average, workers move to better paying cities.




Figure 20: Model-predicted average migration rates of married households as a function of the migration value of both spouses (smoothed). The logit regression includes as controls a quadratic function of age for both spouses, the logarithm of wages of each both spouses, the standard deviation of the cross-city distribution of wage fixed effects for the occupation of each spouse, a dummy for the presence of children in the household, a college dummy for each spouse and year fixed effects.

F Cohabitation

In the main text, I argue that cohabiting couples behave differently in comparison to legally married couples because of the less stringent nature of cohabitation versus marriage. In this appendix, I present evidence in support of this claim, and show that cohabiting couples display migration rates that are similar to those of singles. The ACS sample allows to distinguish between simple housemates from cohabiting, non-married couples thanks to a variable (“related”) that defines the relationship of household members with the head. One of the possible values for this variable is “unmarried partner.” Among all non-married households, about 13.9% report the head being in an unmarried relationship. I use this information to compare the migration rates of singles to the subset of unmarried couples. The comparison is carried out estimating a logit model for the probability of migration of the following form:

$$Pr(Y_{i,t} = 1) = \varphi(\beta_0 + \beta_1 Cohabiting_{i,t} + \beta_2 \mathbf{X}_{i,t})$$

where $Cohabiting_{i,t}$ is a dummy that indicates whether individual i has a partner cohabiting with them and $\mathbf{X}_{i,t}$ contains controls for age, education, and the presence of children. Table

Yearly Probability of Migration (Odds-ratios) - Men				
Migration Value 1	1 (0)	1 (0)	1 (0)	1 (0)
Migration Value 2	1.239*** (0.0822)	1.246*** (0.0839)	1.427*** (0.110)	1.253*** (0.0801)
Migration Value 3	1.435*** (0.0991)	1.422*** (0.102)	1.983*** (0.170)	1.392*** (0.0940)
Migration Value 4	1.508*** (0.120)	1.529*** (0.128)	2.035*** (0.197)	1.516*** (0.117)
Migration Value 5	2.382*** (0.203)	2.474*** (0.221)	3.464*** (0.369)	2.314*** (0.191)
Observations	232,568	232,372	222,391	244,917
Occ. FE	NO	YES	NO	NO
State FE	NO	NO	YES	NO

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

Table 19: Estimates of the migration equation for men. The table reports the odds-ratios relative to an individual who has a low migration value (level 1). Controls include: a quadratic function of age; the logarithm of own wage; the standard deviation of the cross-city distribution of wage fixed effects for own occupation; a dummy for the presence of children in the household; a college dummy; year fixed effects. The second column adds occupation fixed effects while the third controls for state fixed effects. The fourth column includes men who are out of the labor force for which an occupation is observed in the data.

22 reports the estimated coefficient (in odds-ratio form) on the *Cohabiting_{i,t}* dummy. It clearly shows that there is no statistically significant difference between the migration rates of singles and unmarried couples.

G Occupational and Geographic Mobility

This appendix discusses the challenges that occupational mobility poses to the reliability of the measure of the value of migration. Ideally, one would like to be able to observe occupational mobility to directly account for it. Yet, the ACS sample does not contain any information on the occupational history of workers. I argue here that this is not detrimental to the validity of the results.

First, I show that that over a one-year horizon (the frequency at which the analysis is performed) occupational change is not a pervasive phenomenon. Table 23, constructed using CPS-ASEC data (Flood et al., 2023) from 2008 to 2017 applying the same sample restrictions

Yearly Probability of Migration (Odds-ratios) - Women				
Migration Value 1	1 (0)	1 (0)	1 (0)	1 (0)
Migration Value 2	1.258*** (0.0928)	1.198** (0.0912)	1.761*** (0.153)	1.243*** (0.0890)
Migration Value 3	1.291*** (0.102)	1.207** (0.0998)	2.088*** (0.202)	1.256*** (0.0968)
Migration Value 4	1.797*** (0.154)	1.635*** (0.149)	2.766*** (0.292)	1.725*** (0.145)
Migration Value 5	2.596*** (0.248)	2.440*** (0.248)	4.188*** (0.495)	2.540*** (0.235)
Observations	204,639	203,539	196,454	213,840
Occ. FE	NO	YES	NO	NO
State FE	NO	NO	YES	NO

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 20: Estimates of the migration equation for women. The table reports the odds-ratios relative to an individual who has a low migration value (level 1). Controls include: a quadratic function of age; the logarithm of own wage; the standard deviation of the cross-city distribution of wage fixed effects for own occupation; a dummy for the presence of children in the household; a college dummy; year fixed effects. The second column adds occupation fixed effects while the third controls for state fixed effects. The fourth column includes women who are out of the labor force for which an occupation is observed in the data.

as in the main ACS sample, reports the fraction of workers who have changed occupation in the year preceding the survey conditional on their migration status. It shows that the vast majority of moves are not accompanied by occupation changes, with the overall fraction of occupational switches averaging 5.8%.

Secondly, the measure of migration value computed in the main text, which is based on post-migration occupations, is still pertinent in a setting in which workers are forward-looking. A forward-looking worker who wants to change occupation and receives an offer another location, will make a migration decision based on the associated costs and benefits of the new occupation, rather than those of the previous one. This suggests that the relevant occupation for measuring the option value of migration is the post-migration occupation, rather than the occupation held before migration.

Finally, one can advance an equilibrium argument. Occupational mobility imposes costs, and the cost of transitioning between two occupations is contingent on their similarities.³¹

³¹See Poletaev and Robinson (2008), Lazear (2009), Yamaguchi (2012), Gathmann and Schönberg (2010),

Yearly Probability of Migration (Odds-ratios) - Men and Women					
Migration Value 1	1 (0)	1 (0)	1 (0)	1 (0)	1 (0)
Migration Value 2	1.251*** (0.0617)	1.223*** (0.0616)	1.559*** (0.0895)	1.509*** (0.0908)	1.251*** (0.0597)
Migration Value 3	1.375*** (0.0714)	1.325*** (0.0717)	2.029*** (0.130)	1.982*** (0.136)	1.335*** (0.0677)
Migration Value 4	1.636*** (0.0954)	1.572*** (0.0970)	2.309*** (0.164)	2.292*** (0.178)	1.610*** (0.0911)
Migration Value 5	2.483*** (0.157)	2.458*** (0.165)	3.716*** (0.294)	3.753*** (0.318)	2.418*** (0.148)
Observations	437,207	437,207	418,845	418,845	458,757
Occ. FE	NO	YES	NO	NO	NO
State FE	NO	NO	YES	NO	NO
MSA FE	NO	NO	NO	YES	NO

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

Table 21: Estimates of the migration equation on the sample including both men and women. The table reports the odds-ratios relative to an individual who has a low migration value (level 1). Controls include: a quadratic function of age; the logarithm of own wage; the standard deviation of the cross-city distribution of wage fixed effects for own occupation; a dummy for the presence of children in the household; a dummy for gender; a college dummy; year fixed effects. The second column adds occupation fixed effects while the third controls for state fixed effects. With the larger sample size, MSA fixed effects can be added as controls in the fourth column. In column five, I include workers who are out of the labor force but for which an occupation is observed in the data.

According to a simple equilibrium argument, returns between similar occupations should be equalized, up to some compensating differentials for non-wage characteristics, within a given labor market. This would result in similar city rankings. In this case, my empirical analysis is less likely to be hindered by occupational transitions, given that most of shifts will occur between occupations with similar rankings.

H Existence of the Equilibrium

In Online Appendix I, I construct an *update function* Ψ that has a fixed point if and only if there exists a stationary equilibrium. This function takes as argument an element from the set of feasible distributions of singles \mathcal{D} , computes all the endogenous objects in the

and Cortes and Gallipoli (2018).

Yearly Probability of Migration (Odds-ratios)						
	Men			Women		
Cohabiting	1.072 (0.0514)	1.026 (0.0507)	1.027 (0.0542)	1.510*** (0.0721)	1.124** (0.0549)	1.089 (0.0574)
Observations	329,214	329,214	278,924	277,617	277,617	240,311
Demographic Controls	NO	YES	YES	NO	YES	YES
Robust standard errors in parentheses						
*** p<0.01, ** p<0.05, * p<0.1						

Table 22: Regression-based odds-ratios for the probability of migration of cohabiting partners vs singles. Demographic controls include: a quadratic function of age; a dummy for the presence of children; a dummy for education; year fixed effects. In the third and sixth column, I restrict the sample to men and women who report being employed and earn a positive income.

Occupation switchers conditional on mobility					
	Same house	Moved within county	Moved within state	Moved between states	Total
Same occupation	94.4%	91.8%	87.1%	80.8%	94.2%
Different occupation	5.6%	8.2%	12.9%	19.2%	5.8%
Total	100%	100%	100%	100%	100%

Table 23: Migration and job switching (2008-2017 CPS-ASEC).

model, and returns an updated distribution of singles from the set of feasible distributions. To prove the existence of an equilibrium, it is sufficient to show that Ψ and \mathcal{D} satisfy the conditions stated in Brouwer's theorem, namely that \mathcal{D} is compact and convex and that Ψ is a continuous mapping from \mathcal{D} to itself.

\mathcal{D} is compact and convex. The stated properties of the set \mathcal{D} follow trivially from the definition of the set itself. Let $\boldsymbol{\mu}$ be a vector of values μ_{g,x_g} for all g and x_g describing a distribution of single males and females over the state space. Then, the set of feasible distributions, \mathcal{D} , is defined as

$$\mathcal{D} = \left\{ \boldsymbol{\mu} \mid \sum_{\ell} \sum_{x_f} \mu_{f,x_f} \leq F; \sum_{\ell} \sum_{x_m} \mu_{m,x_m} \leq M \right\}. \quad (30)$$

Compactness and convexity of this set are trivially satisfied.

Ψ is a continuous mapping from \mathcal{D} to itself. The mapping Ψ is described in detail in Online Appendix I. For each $\mu \in \mathcal{D}$, the mapping gives $\mu' = \Psi(\mu) \in \mathcal{D}$ simply because the mass of males M and females F is exogenous to the model and kept constant by construction. The continuity of Ψ follows from the continuity and smoothness (no probability masses) of the bliss (ζ) and preference (ξ) shocks, and from the continuity of the matching function. Intuitively, the continuity of the shocks, ensures that there are no discrete “jumps” in the flows of workers across states, as the value functions associated with each of the discrete choices that agents make change continuously. Moreover, these value functions are continuous in the matching probabilities that are continuous in the distribution μ because of the continuity of the matching function. It follows that all the value functions in the model are continuous in μ . Summing up, continuous changes in μ cause continuous changes in the value functions, that imply continuous changes in the flows of workers across states and, thus, continuous changes in μ' .

I Model Solution

This appendix describes the algorithm employed for the numerical solution of the model.

Step 1: Guess initial feasible³² distributions of singles, μ_{m,x_m} and μ_{f,x_f} .

Step 2: Given the guessed distributions, compute the matching probabilities using equation (14).

Step 3: Using the probabilities computed above, solve the fixed-point problem described by the value functions. The solution to this problem is computed through value function iteration. The policy functions for marriage, divorce and migration are obtained as a byproduct of this process.

Step 4: Using the policy functions, obtain a new guess for the distribution of singles using equations(24) through (29).

Step 5: Evaluate the distance between the initial guesses and the new guesses. If it is less than a specified tolerance level then terminate. Otherwise, repeat from step 2 using the updated distributions as the initial guesses.

³²The distributions have to satisfy $\int_{x_f} d\mu_{f,x_f} = F$ and $\int_{x_m} d\mu_{m,x_m} = M$.

Notice that the algorithm outlined above describes a mapping $\Psi : \mathcal{D} \rightarrow \mathcal{D}$ from the set of feasible distributions of singles \mathcal{D} to itself. The stationary equilibrium then can be defined as the solution to the fixed point problem described by $\Psi(\mathcal{D}) = \mathcal{D}$.

J Dealing with the curse of dimensionality

This appendix describes the process used to construct the clusters of cities fed to the model. Clustering is necessary to reduce the computational power needed to numerically compute the equilibrium distributions. Since the model dynamics are driven by the wage heterogeneity across cities and occupations, the algorithm is aimed at grouping together cities offering similar wages across occupations.

To perform the clustering, I employ a process that combines two unsupervised machine learning techniques. First, I perform a permanent component analysis (PCA) on the city-occupation wage premiums. Then, I apply k-means clustering on a subset of the principal components. Performing k-means clustering on a set of principal components is not a novelty in the scientific literature (Ibes, 2015). Although there is no evidence that this two-step procedure improves the quality of the resulting clustering over a direct application of the k-means algorithm (Yeung and Ruzzo, 2001), the dimensionality reduction achieved with PCA is helpful to improve the workings of the k-means algorithm. The k-means algorithm is, in fact, initialized by randomly drawing initial centroids for the clusters. The algorithm recursively updates these centroids as the average of the closest elements to the initial centroid (I use the euclidean distance to measure the distance from centroids). Issues arise with high dimensionality, i.e., if the vector of characteristics on which the clustering is performed is too big. In this case, the algorithm tends to produce results that are highly dependent on the starting point. One way to deal with this is to try with different initial points and then pick the best clustering. Nevertheless, with high dimensionality, this is almost akin to a random search of the best clustering. An additional advantage of using a subset of permanent components is that it reduces the impact of the estimation noise coming from the wage regression.

To perform PCA, I consider each city as one observation and the set of occupation fixed effects (the estimates of $\alpha_{\ell,j}$ from equation 21) as the covariates. Figure 21 shows the share of variation explained by the first 30 principal components.

In the second step, I perform k-means on the first 5 components. This choice is rather arbitrary as there is no objective criterion nor any rule of thumb. The 5 components together explain 49% of the total variation.

One of the shortcomings of k-means is that the algorithm does not identify the optimal number of clusters. The desired number of clusters has to be provided by the user. A sim-




Figure 21: Explanatory power of the first 30 principal components.

ple and popular solution consists of inspecting the dendrogram produced using hierarchical clustering to see if it suggests a particular number of clusters. Nevertheless, this approach is also subjective. An alternative popular way is the elbow method that consists of a graphical analysis of the total intra-cluster variation as a function of the number of clusters.³³ The optimal number of clusters is identified by an “elbow” in the graph, i.e. the point in which the gains from increasing the number of clusters in terms of the reduction of the total intra-cluster variation i becomes small.

To choose the number of clusters I adapt the idea behind the elbow method to my particular problem. As discussed in the main text, the most important aspect of the wage distribution to preserve after the clustering is the covariance structure of the occupation-city premiums. With this goal in mind, I perform k-means varying the number of clusters from 5 to 150 and, for each clustering, I estimate the wage regression (21) and compute the wage covariance structure.³⁴ Next, I compare the estimated covariance structure to the original one (i.e. the one obtained from real cities). In practice, I compute the Euclidean distance between the covariance vectors from the cluster and the original data. The computed distances as a function of the number of clusters are shown in Figure 22.

As expected, the function is downward sloping as adding more clusters makes it easier to match the actual wage covariance structure. Nevertheless, the gains from adding more clusters are drastically reduced after 20/30 clusters. Given this, I deemed 25 clusters to be a fair compromise between accuracy and the need for speed. A table that shows the composition of the 25 clusters is available on request.

³³As the number of clusters increases the intra-cluster variation is mechanically reduced.

³⁴For each call of k-means, I run the clustering algorithm 10000 times for different initial points and select the clustering that returns the lowest total intra-cluster variation.




Figure 22: Euclidean distance between the actual and clustered covariance vectors as a function of the number of clusters.

K The gender wage gap

The exogenous gender gap used in the model has to be intended as an ex-ante income gap reflecting factors that are orthogonal to marriage and divorce (e.g., discrimination). This value is, in fact, obtained from the estimation of equation (21) which controls, among other factors, for marital status and city fixed effects.

The interaction between marriage and divorce induces an additional income gap that is generated from the different migration patterns of men and women. To assess the performance of the model in reproducing the ex-post gap, I compute the log difference between the average wage of men and the average wage of women from both the model and the data. For the latter, I compute the gap using the residuals from a regression of log wages over dummies for education and year fixed effects, and a quadratic function of potential work experience. This adjustment is necessary to make the quantities from the data comparable to that of the model since the latter does not present life-cycle effects nor heterogeneity in education. Reassuringly, the empirical and modeled gender gaps are very close, the former being equal to 26.0% and the latter to 26.3%.

L Computing bounds for the change in total earnings

This appendix explains in details the calculations performed to compute the boundaries for the change in aggregate labor earnings reported in Table 5. Both bounds are obtained computing counterfactual wages under different assumptions.

The lower bound

To determine the lower bound, I use the demand elasticities obtained from the estimated production function from Alonzo and Gallipoli (2025). This production function displays decreasing marginal productivity of labor, such that an increase in labor supply to one occupation translates into lower wages. In that paper, the labor demand side of the economy is characterized by a continuum of firms operating in monopolistic competition that demand single occupations to produce intermediate goods that are subsequently combined by a representative firm to produce the consumption good. This production structure delivers the following inverse demand function for occupation j' , that is

$$w_{ij} = \rho A \alpha_j \beta_j \left[\sum_{j'} \alpha_{j'} \left(\sum_i \beta_{ij'} L_{ij'} \right)^\rho \right]^{\frac{1-\rho}{\rho}} \left(\sum_{i'} \beta_{i'j} L_{i'j} \right)^{\rho-1}. \quad (31)$$

where i identifies demographic groups defined on gender, education, and age. The corresponding elasticity of wages to labor is

$$\begin{aligned} \epsilon_{ij} &= \frac{\partial w_{ij} / w_{ij}}{\partial L_{ij} / L_{ij}} \\ &= A \alpha_j \beta_j^2 \rho (1 - \rho) \left(\sum_{i'} \beta_{i'j} L_{i'j} \right)^{\rho-2} \left[\sum_{j'} \alpha_{j'} \left(\sum_{i'} \beta_{ij'} L_{ij'} \right)^\rho \right]^{\frac{1-\rho}{\rho}} \left[\frac{\alpha_j (\sum_{i'} \beta_{i'j} L_{i'j})^\rho}{\sum_{j'} \alpha_{j'} (\sum_{i'} \beta_{ij'} L_{ij'})^\rho} - 1 \right] \frac{L_{ij}}{w_{ij}} \\ &= \beta_{ij} \rho (1 - \rho) \left(\sum_{i'} \beta_{i'j} L_{i'j} \right)^{-1} \left[\frac{\alpha_j (\sum_{i'} \beta_{i'j} L_{i'j})^\rho}{\sum_{j'} \alpha_{j'} (\sum_{i'} \beta_{ij'} L_{ij'})^\rho} - 1 \right] L_{ij} \\ &= \beta_{ij} \rho (1 - \rho) \frac{L_{ij}}{(\sum_{i'} \beta_{i'j} L_{i'j})} \left[\frac{\alpha_j (\sum_{i'} \beta_{i'j} L_{i'j})^\rho}{\sum_{j'} \alpha_{j'} (\sum_{i'} \beta_{ij'} L_{ij'})^\rho} - 1 \right] \end{aligned} \quad (32)$$

These elasticities cannot be directly applied to the framework of this paper for two reasons: (i) my occupational system is more disaggregated than in Alonzo and Gallipoli (2025) and (ii) the demographic definition is more aggregated. To overcome the latter issue, I compute an employment-weighted average of these elasticities for each occupation,

$$\epsilon_j = \sum_i \mu_{ij} \epsilon_{ij} \quad (33)$$

where μ_{ij} is the share of workers of demographic group i in occupation j . To deal with the first issue, I simply assume that the wage elasticities are constant for all the occupations composing each aggregated occupation used in Alonzo and Gallipoli (2025). In other words, let Ξ_j be the set of occupations j (from the occupational system used in this paper) that are contained in the definition of occupation \hat{j} (from the occupational system used by Alonzo

and Gallipoli, 2025). I assume that $\epsilon_j = \epsilon_{\hat{j}}$ for all $j \in \Xi_{\hat{j}}$.

The lower bound is obtained by using these elasticities to compute counterfactual wages based on the changes in employment in the counterfactual scenarios.

The upper bound

To compute the upper bound, I rely on the literature on agglomeration economies. In a meta-analysis of 729 elasticities of wages to city size, Melo et al. (2009) identify an average elasticity of 3.2% with a standard deviation of 7.6%.³⁵ To compute the counterfactual wages, I use a conservative value of 18.4% (the mean plus two standard deviations). As for the lower bound, this elasticity is used to compute counterfactual wages based on the changes in city sizes in the counterfactual scenarios.

³⁵See Table 2 in Melo et al. (2009).

References

- Alonzo, D. and Gallipoli, G. (2025). The Changing Value of Employment. Working Paper.
- Cortes, G. M. and Gallipoli, G. (2018). The costs of occupational mobility: An aggregate analysis. *Journal of the European Economic Association*, 16(2):275–315.
- Flood, S., King, M., Rodgers, R., Ruggles, S., Warren, J. R., Backman, D., Chen, A., Cooper, G., Richards, S., Schouweiler, M., and Westberry, M. (2023). Integrated Public Use Microdata Series, Current Population Survey: Version 11.0 [dataset]. *Minneapolis, MN: IPUMS*.
- Gathmann, C. and Schönberg, U. (2010). How General Is Human Capital? A Task-Based Approach. *Journal of Labor Economics*, 28(1):1–49.
- Ibes, D. C. (2015). A multi-dimensional classification and equity analysis of an urban park system: A novel methodology and case study application. *Landscape and Urban Planning*, 137:122–137.
- Lazear, E. P. (2009). Firm-Specific Human Capital: A Skill-Weights Approach. *Journal of Political Economy*, 117(5):914–940.
- Melo, P. C., Graham, D. J., and Noland, R. B. (2009). A meta-analysis of estimates of urban agglomeration economies. *Regional Science and Urban Economics*, 39(3):332–342.
- Poletaev, M. and Robinson, C. (2008). Human Capital Specificity: Evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984-2000. *Journal of Labor Economics*, 26(3):387–420.
- Yamaguchi, S. (2012). Tasks and Heterogeneous Human Capital. *Journal of Labor Economics*, 30(1):1–53.
- Yeung, K. Y. and Ruzzo, W. L. (2001). An empirical study on principal component analysis for clustering gene expression data. *Bioinformatics*, 17:763–774.