Geographic Space and a Woman's Place: The Gendered Relationship between Geographic and Economic Mobility

Katariina Mueller-Gastell*and David S. Pedulla†

Abstract

Women are less likely to apply to jobs that require geographic relocation than men. This difference is starkest for partnered men and women. We use unique and novel application level job search data to show that the differences between partnered men and women persist even when controlling for the gendered occupational make-up of these groups. Prior research – without access to data on search behavior itself – has been constrained to look at completed job moves (into a target occupation) rather than job applications (from an originating occupation). We replicate that gender differences appear smaller when controlling for target occupation and theoretically explain why such conditioning misses existing substantively significant gendered difference in actual search behavior. Further, we examine other explanations for gendered differences: the relative earnings of spouses, performative gender roles, and parental status. We present evidence that gender norms continue to carry strong explanatory power.

Keywords: job change, household mobility, gender, occupational segregation, migration

*Department of Sociology, Stanford University. Please direct all correspondence to kata@stanford.edu
†Department of Sociology, Stanford University.

Women are less likely to geographically relocate for job purposes than men are. This broad association has been found in numerous previous studies, using a variety of data sources.[1] Qualitative studies have focused on the behavior of married or cohabiting (i.e. partnered) women and on household decision making. These studies have identified relative bargaining power, performative gender norms, and differential effects of parenthood as causal explanations for this gender gap in geographic mobility: women tend to earn less and their careers are less favored within the household (Hochshild, 1989), and their decisions are more influenced by childcare considerations (Bass, 2015).

Quantitative studies, on the other hand, have stressed the structural differences in the labor market positions that men and women inhabit. Most powerfully, Benson (2014) argues that the observed differences in geographic mobility can be fully explained by occupational sorting: women tend to work in occupations that are more geographically dispersed, while men are more often found in highly geographically concentrated occupations. Hence, a job change by a woman is simply less likely to *necessitate* a costly geographic relocation than a man's job change. Controlling for the geographic dispersion of the target occupation of job moves completely eliminates gender differences in the frequency of geographic relocations in the Current Population Survey (CPS).

Though structural features of the labor market and occupational differences between men and women are clearly an important factor in relocation decisions and frequencies, we wish to move the quantitative discussion back to some of the features identified in the qualitative literature. Using newly available data that overcomes limitations in existing survey data such as the CPS, we can show that gender differences in geographic mobility are largely driven by the lower propensity of married women to relocate, and that these differences cannot be explained solely by differences in the occupational make-up. Instead, we present evidence that gender norms, even above and beyond relative earnings, carry strong explanatory power.

Further, we advance the conceptual argument that concentrating on *target* occupations of a

[1] E.g., Shauman (2010) finds that women are less likely to transition geographically in the Panel Study of Income Dynamics (PSID) data and Current Population Survey (CPS) data; Boyle et al. (2001) also uses the CPS, as does Benson (2014).

job move (as has to be done due to data limitations in most existing survey data) may paint a misleading picture of gender differences in geographic mobility. Because we have access to data on individual-level job search behavior we can validate this theoretical consideration in our analysis by comparing target occupation controls to source occupation controls. The former – misleadingly, as we argue – accounts for a larger fraction of gendered mobility differences than the latter. We hence advocate for gathering more data, such as job search data, that allows to condition on *current* or source occupations, as this information more closely matches to what one theoretically wishes to control for when taking occupational sex segregation into account.

Why do men and women relocate for work at different rates?

Women do not relocate for work at equal rates to men, but instead are likely to change jobs locally when they do change them. At the same time, the relationship between geographic and economic (upward) mobility is strengthening. The ability to move for work is growing in importance in the United States. Internal migration has declined – Americans no longer move to greener pastures (Molloy et al., 2011, 2017) (though cf. Borjas et al. (1992)). There has been a significant literature tackling the possible causes of this decline and its effects on economic mobility (Hsieh and Moretti, 2019; Kaplan and Schulhofer-Wohl, 2017). The existing research provides a stylized fact: 'good' jobs are increasingly concentrated in specific geographic hubs, and the ability of individuals to access jobs in these hubs has declined. If women have a lower propensity to move for their careers, this may, depending on the source of this lower propensity, restrict their access to these good, geographically clustered jobs and may contribute to gender pay gaps.

To understand what the impacts of women's greater tendency to stay local could be, it is important to understand why women are not moving for economic mobility. First – and there is strong evidence that this is indeed the case (Benson, 2014, 2015; Shauman, 2010; Shauman and Xie, 1996) – women's decisions to work locally are structurally determined. Men and women face a different

labor market. They thus have different incentives to move and different opportunities, resulting in what appear to be different preferences to move. A second set of explanations turns to examine the relocation decision itself not as an individual decision, but as one made as a household. These arguments note that the gender disparity in job relocation is highest for partnered individuals, and they focus on explaining how this pattern emerges from the relative status or bargaining power of spouses within a household. Separating these strands of explanation is somewhat artificial – the structural constraints on women in large part arise from gender disparity within the household, and these disparities are in turn reproduced by unequal labor market opportunities and outcomes. Yet, while deeply endogenous, disentangling these sources of the disparity can help tease apart the precise mechanism that produces it.

We conduct analyses with our data which suggest that the structural determinant of women's relocation decisions – namely, the focus on occupational sex segregation – has been overstated in the existing literature. We are able to use application level data on job search pools to show that the observed differences in search behavior between partnered men and partnered women cannot be explained by differences in the present occupational make-up of job seekers. Our data replicate the attenuation of difference between men and women when conditioning on the occupation of the job to which the applicant is applying, and explain how conditioning on target occupation misses the existing stark difference in actual search behavior and its gendered component. To illustrate the importance of this subtle distinction for explaining gender difference, we test arguments about the relative bargaining power of spouses within a household using these application level job search behavior data. We find tentative support for what we call a performative gender hypothesis, which prior research – using *target* occupation measures – has largely dismissed.

The structural explanation: occupational sex segregation

Structural accounts dominate recent literature on why men and women differ in their propensity to move for work. According to these explanations, labor market sex segregation produces structural differences, e.g. men and women labor in different occupations. It is these positional differences

that then produce the observed difference in geographic mobility. Though most research has primarily highlighted the attenuating effect of conditioning on occupation, Benson (2014) provides a compelling account of the mechanisms for *how* gendered occupational composition can create this outcome. He shows that some of the lower observed frequency of job relocation of women can be explained by the geographic dispersion of female-dominated occupations. In his data, among *unmarried* women relocation rates are already smaller than among men, suggesting that women are not limited by their partners. Once one controls for the geographic clustering of occupations, however, the difference disappears. The interpretation is that women, be they partnered or not, are simply less likely to need to relocate, as they tend to have jobs which are available across the country.

Another mechanism through which occupation segregation could also operate is the lower average earnings of women-dominated occupations. If women are in lower wage occupations, they may have both fewer incentives and fewer resources to move (Garip, 2008). This mechanism is complementary to the account of the geographic dispersion of different occupations. It suggests an additional route through which occupational gender segregation can lead to differences in relocation rates, regardless of where those jobs are geographically available.[2]

In sum, existing research establishes that occupation has a clear mediating effect on the geographic mobility of women. What is less clear is whether occupational composition fully accounts for the disparity between men's and women's relocation rates. Occupational sex segregation remains stark in the United States, and composition undeniably explains gender differences – but in this paper, we suggest that other mechanisms continue to be at play. These other mechanisms may have been masked by the data and modelling strategies to which researchers have been limited previously.

[2]The choice of occupation is, of course, in some sense endogenous – anticipating future constraints of a spouse's job, women may choose more geographically dispersed jobs. Additionally, the relatively high geographic mobility of men may have aided the geographic clustering of male-dominated occupations. Similarly, clustered jobs may be more specialized and need to pay more to attract qualified workers.

An economic theory of household migration

Other accounts of the disparity of geographic mobility focus on explaining differences between married (or partnered) men and women. Mincer (1978) provides a seminal account of household migration decisions: A family moves when the benefits to the family as a whole outweigh the costs of doing so. This offers the simplest economic theory of migration. Women, who tend to be the secondary earner within most dual-earner households (Bertrand et al., 2015), could be unwilling or unable to search for jobs that would require relocation because it simply would not make financial sense to endanger a husband's more lucrative career.[3] Relocation for a spouse's job is costly (Boyle et al., 2001) and may not be a risk worth taking for the sake of the lower earner.

The **relative earnings hypothesis** that arises from this framework would then predict a gender difference – for partnered men and women, at least – beyond that accounted for by occupational sorting. Even within the same occupation, women tend to earn less (Goldin et al., 2017), and assortative mating results in women partnering with higher earning men (this remains the case even though the within-couple gap has declined over time (Schwartz, 2010)). On the other hand, this theory would predict that were the woman the primary earner, the household would then move for her work. Further, as the relative earnings share of a woman increases, so too should the likelihood that her job search is geographically broad.

Gender norms

Theories about the role of gender norms suggest that we ought not expect relative income to behave symmetrically for men and women. Instead, the decision to relocate for work is governed by beliefs or norms about gender-appropriate behavior or roles (Bielby and Bielby, 1992; Schneider, 2011; Sorenson and Dahl, 2016). Schneider (2011) finds a curvilinear association between earnings share and women's housework time, which provides strong evidence that women perform gender through increased engagement with gender-appropriate activities – housework – when the male-breadwinner

[3]Hence, evidence against the workings of a direct gender norm would come from the share of the woman's income being *positively* related to her probability to look for work further from home.

norm is violated (see also Fleche et al., 2018).

Other authors offer a complementary comparative advantage explanation, referred to as relative resource theory (Bielby and Bielby, 1992; Shauman, 2010). Beyond being primary earners, men have other resource advantages over women that can heighten their bargaining power within a household and consequently increase the household's propensity to move for his work. This spousal advantage can be operationalized as an upperhand in e.g. educational attainment or occupational prestige (e.g. Shauman, 2010). In a very different context, Udry (1996) finds evidence for the higher bargaining power of men within households in Burkina Faso, resulting in greater allocation of resources to male householders' plots of land even when this results in a poorer harvest for the entire family.

The **performative gender hypothesis** thus predicts that there will be residual difference between partnered men and women in the propensity to relocate for work. When gender norms are violated – such as by the woman being the primary earner in the household – this hypothesis suggests that the couple may compensate by increasing norm-conforming performance of gender in other respects. The couple may then be less willing to relocate for the woman's job to avoid further reinforcing her role as the household breadwinner.

Motherhood and fatherhood

Finally, parenthood might have a different effect on men and women. Parenthood could operate through many of the hypotheses suggested above. Women might sort into specific occupations in anticipation of a future greater responsibility for childcare compared to their male spouses (Bass, 2015). Relative earnings gaps might also be exacerbated, as women face a motherhood penalty (Correll et al., 2007), while men might experience a fatherhood premium (Killewald, 2013) in wages. Shauman and Xie (1996) argue that, for doctoral scientists, partnerships constrain women only once they become mothers. The presence of children may cement the household as unit and make long-distance relationships less desirable. Gender norms may become increasingly salient as parents specialize into their roles as a mother and as a father. For women responsible for childcare,

the support provided by existing social networks and local ties may feel increasingly crucial. The **parenthood hypothesis** predicts that having a child has a stronger effect on women, reducing their propensity to move for work. The other mechanisms predict that the effect of fatherhood would be of opposite sign for men than it is for women: increased household specialization, relative earnings, and more occupationally flexible spouses could mean fathers move for work *more* than do childless men.

Limitations of existing tests of these hypotheses

It has proven difficult to identify the cause of an economic relocation given data limitations. Previous research has, for example, identified the cause of the move by testing whether post-migration outcomes are predicted by measures of the man or woman's earnings potential and relevant local labor market conditions (Shauman and Noonan, 2007), or by asking recently relocated households about the impetus for the move (Benson, 2014) after-the-fact.

Further, the focus on successful household migrations muddies the distinction between supply-side processes – whether a household wants to move for work – and demand-side factors – whether an employer wants to hire a member of the household. Demand-side discrimination has a gendered component. Women might not be hired into positions that require a move, for fear that their partners will not follow (Rivera, 2017). There might be additional barriers in place that make it more difficult for women to enter those dispersed occupations that men dominate.

These demand-side processes can of course also constrain women's search for work. Even childless women make career decisions with the spectre of future family constraints at the back of their minds, in a way that men do not (Bass, 2015). Thus, women might sort into occupations that are more dispersed. Experiences of past discrimination can also change women's preferences.

Shifting our analysis to an earlier stage in the job matching process, when job seekers are *searching* for work, allows us to examine the opportunity set that households are faced with *before* they make a migration decision. Though the supply-side of household economic migration is informed by demand-side constraints, an analysis of job search behavior enables a better test of

how men and women behave at the first stage of the process.

Data and Methods

Data

Using a unique data-set of 2,060 job seekers, the National Longitudinal Study of Job Search (NLSJS),[4] we can study the supply-side of the job matching process. The NLSJS sampled non-institutionalized adult job seekers in the United States over an 18-month period. This data collection effort was conducted in collaboration with GfK, a survey research company. The respondents provided information on the most recent jobs they had applied to over the course of nine waves, including whether the job would require relocation. They also provided information on marital status, prior employment history, and other sources of household income. Once we restrict our sample to respondents for whom we have data on wages and occupation and exclude those who do not identify as heterosexual, our sample size is 1,826 job seekers who applied to 19,886 jobs. This data-set provides application level data on job search pools.

Modelling strategy

Our measure of geographic job mobility and the key dependent variable is whether a job applied to would require the respondent to move. For ease of interpretation, we employ linear probability models throughout but all of the reported results are robust to logit specifications.

As job seekers often apply to multiple jobs, we cluster all standard errors at the respondent level, and all models include survey wave fixed effects to account for time trends. All models also include controls for the number of applications sent by each applicant, logged to base ten, to ensure results are not driven by differences in sheer volume of applications in our data.

[4]These data have previously been described and used by Pedulla and Pager (2018).

Variable construction

Dependent variable: job requires move

For each reported application, respondents were asked to estimate a commute time or whether they would have to move in order to take the position. In the results presented here, this is treated as a binary outcome variable, where an application is either to a job that requires a move or does not. The analysis then is at the application level.

Alternatively, the dependent variable could also have been constructed as the proportion of applications that are to jobs that require a move. This analysis would occur at the applicant level. With this restriction, we would be unable to conduct the analyses of Table 3 which require application level variation.

Income

For job seekers who were employed at the time of the first survey wave, we use their present wage to construct the measure of their percentage of wage relative to their spouse. For unemployed respondents, we code their relative income as zero in the results presented here. Similarly, if a respondent reports that their spouse is presently unemployed, they are coded as providing zero percent relative to the respondent spouse. Here, we hence consider the realized relative wage, rather than the earning potential of the spouses. The results presented are, however, robust to instead including the most recent wage earned by unemployed respondents and spouses. The discussed coefficients remain statistically significant and similar in size.

Our measure of income is, unfortunately, relatively noisy. Respondents self-report their own most recent earnings and their spouses' income. When first admitted to the GfK panel, respondents also reported a household income bracket. For many respondents, this income bracket may not be up to date. However, the brackets are broad, so we do not think moves from one bracket to the other are likely. The exception to this is cases of unemployment of one or both (or more) household earners. Even in this case, however, we are comfortable using the previous earnings of the household when the earners were employed, because this probably better captures earning potential.

Occupation

Our data allow us to control for two different occupations: current (or most recent, for unemployed respondents) occupation, and the occupation of the job applied to, the *target occupation*. Respondents in the NLSJS were asked to provide the job title in an open text field. Trained coders from the University of Wisconsin Survey Center translated these in the 96 categories of the 2010 Standard Occupational Classification (SOC) system, with a 97 percent success rate (see Pager and Pedulla (2015)). In the analyses presented here, we use the 23 major SOC categories. We are concerned that the fine-grained nature of the 96 categories could mask some important patterns.

Parenthood

To test for the effect of parenthood, we use a measure of whether or not the respondent has a child under the age of 18 living with them at the time of the baseline survey.

Marital status

We collapse married and cohabiting respondents. Though others have argued for excluding cohabiting couples (e.g. Shauman and Noonan (2007)), we are concerned about selection effects into marriage compared to cohabitation (Sassler and Miller, 2017). If anything, we would expect that cohabiting women would be less constrained by their partners (if cohabitation is taken to signal lower commitment to the relationship), and thus including these respondents would downwardly bias our estimates of the relationship between gender and job relocation for married women and is thus the more conservative modeling decision given our hypothesis.

We keep formerly partnered respondents are their own category in our models. Separated, widowed, and divorced respondents are grouped together instead of being included as single respondents.

Demographic controls

We include a continuous measure of age at the time of the baseline survey. We also control for a categorical measure of self-reported race or ethnicity.

Table 1 describes the analytic sample and the distribution of key variables.

Table 1: Summary statistics of respondent sample

	Mean	Std.Dev.	Observations
Number of applications sent by applicant	10.89	8.45	1826
Age (Years)	40.77	13.58	1826
Woman	0.50	0.50	1826
Married	0.43	0.50	1826
Living with partner	0.11	0.31	1826
Never married	0.31	0.46	1826
Divorced	0.11	0.31	1826
Separated	0.03	0.16	1826
Widowed	0.01	0.12	1826
Parent	0.39	0.49	1826
% of applications requiring move	0.12	0.33	1826
Unemployed	0.35	0.48	1822
Employed full-time	0.52	0.50	1822
In liminal employment	0.13	0.34	1822

Notes: Four respondents did not provide their employment status in the baseline survey. Number of applications is across all nine survey waves. Liminal employment includes part-time and temporary employment.

Source: NLSJS

Results

Table 2: Linear probability models of moving on gender and marital status

	Dependent Var.: Job Would Require Move				
	(1)	(2)	(3)	(4)	(5)
Woman	-0.0593***	-0.0538***	-0.0116	-0.0155	-0.0838***
	(0.0127)	(0.0127)	(0.0222)	(0.0220)	(0.0175)
Married/cohabiting			0.0590**	0.0505*	
			(0.0228)	(0.0222)	
Separated/Divorced/Widowed			0.0920*	0.0990**	
			(0.0368)	(0.0351)	
Woman×Married/cohabiting			-0.0796**	-0.0625*	
			(0.0283)	(0.0274)	
Woman×Separated/Divorced/Widowed			-0.0689	-0.0735+	
			(0.0436)	(0.0417)	
Constant	0.160***	0.227***	0.142***	0.207***	0.408***
	(0.0390)	(0.0437)	(0.0395)	(0.0438)	(0.0921)
Wave fixed effects	Yes	Yes	Yes	Yes	Yes
Previous occupation	No	Yes	No	Yes	Yes
Household income	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes
Observations	19886	19886	19886	19886	10474
R^2	0.030	0.047	0.035	0.052	0.090

Standard errors in parentheses

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: In model 5, the sample is restricted to partnered respondents only. Previous occupation is the current or most recent occupation of the applicant, coded to the 23 major categories of the 2010 Standard Occupational Classification System. Household income is measured in brackets, and bottom-coded at 5,000 USD and top-coded at 175,000 USD. Demographic controls include respondent's age and self-reported race and ethnicity. Each model includes a control for the total number of applications reported by the applicant across all nine survey waves. All standard errors are clustered on the respondent.
Source: NLSJS

Partnered women have geographically narrow job searches

All models in Table 2 include controls for household income, survey wave fixed effects, demographic controls (age and race/ethnicity), and the number of applications sent by each applicant. Standard errors are clustered on applicant. Model 1 shows the association between gender and applying to

a position that would require a move, conditioning on all of the above. In model 2, we account for the current or most recent occupation of the applicant, with little effect on our coefficient of interest. In model 3, we repeat model 1 but now control for marital status and allow for a different slope for men and women by including an interaction. Here, we find that the coefficient on gender is greatly attenuated and no longer statistically significant. We find that the main effect of being partnered is positive, but the interaction with being a woman is negative. This suggests that partnered *men* are more likely to apply to positions that require relocation, but that the reverse is true for partnered women. The relationship between gender and applications to positions that require relocation appears then to be mediated by marital status. In model 4, we again control for the previous or current occupation of the applicant. The association between marital status and the outcome variable is attenuated, but remains statistically significant at the 5 percent level. In model 5, the sample is restricted to partnered respondents, which corroborates the gender disparity found in model 4. In sum, this table shows that partnered men and women apply for jobs that require geographic mobility at different rates, and that one's most recent occupation does not account for this variance. This finding is counter to others in the literature.

Occupational composition does not account for the difference in search behavior

In Table 3, we restrict our sample to applications sent by partnered respondents. Again in model 1, including all of the controls from model 1 in Table 2, we show a strong negative association between being a woman and applying to a position that requires a move for this subset of applicants. In model 2, we condition on most recent occupation, and find the main effect slightly attenuated but still strong and statistically significant at the 0.1 percent level. In model 3, we condition instead on the *target occupation*, i.e., the occupation of the job to which the application was sent. The main effect is further attenuated, but remains statistically significant at the 0.1 percent level. Combining both measures of occupation in table 4, we see the coefficient is not much attenuated compared to that of the preceding model.

Table 3: Linear probability models of moving on gender, restricted to partnered respondents

	Dependent Var.: Job Would Require Move			
	(1)	(2)	(3)	(4)
Woman	-0.0911***	-0.0838***	-0.0673***	-0.0670***
	(0.0171)	(0.0175)	(0.0169)	(0.0171)
Constant	0.350***	0.408***	0.443***	0.476***
	(0.0959)	(0.0921)	(0.0948)	(0.0924)
Wave fixed effects	Yes	Yes	Yes	Yes
Previous occupation	No	Yes	No	Yes
Target occupation	No	No	Yes	Yes
Household income	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes
Observations	10474	10474	10474	10474
R^2	0.066	0.090	0.099	0.112

Standard errors in parentheses

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Previous occupation is the current or most recent occupation of the applicant, coded to the 23 major categories of the 2010 Standard Occupational Classification System. Target occupation is the occupation of the job to which the applicant applied, similarly coded. Household income is measured in brackets, and bottom-coded at 5,000 USD and top-coded at 175,000 USD. Demographic controls include respondent's age and self-reported race and ethnicity. Each model includes a control for the total number of applications reported by the applicant across all nine survey waves. All standard errors are clustered on the respondent.
Source: NLSJS

Target and previous occupation – an example

Why does conditioning on target occupation attenuate the results further than previous occupation? To illustrate: Let us imagine we restrict our sample to nurses – a highly dispersed occupation. Some nurses apply for other positions as nurses (assume none of these require a move), and some apply to a new type of highly specialized nursing position (assume all of these require a move). Conditioning on *current* occupation (nurse) means we compare the rates at which male and female nurses apply to these two target occupations: if men are more likely to apply to the specialized position requiring a move, we will find a gender gap in mobility. For example, consider the case in which there are as many male as there are female nurses and that 80 percent of the applications

to specialized nursing are sent by men, and 20 percent by women. Assume that all remaining applications are to standard nursing positions that do not require a move. Thus, there is a large gender gap in the propensity to move when conditioning on current occupation. On the other hand, if we condition on target occupation, we see *no* gender effect: in this case, we compare men who send an application to the specialized nursing position to women who send an application to a specialized nursing position – each of which, in our example, requires a move, and thus there is no difference between men and women. Similarly, the applications by men sent to normal nursing positions never require a move, as do the applications sent by women, again resulting in no gender difference when conditioning on the target occupation. Thus, even though there is a four to one ratio of applications that require a move *amongst current* nurses, there is no measurable gender gap within the applications sent to nurses or within the applications sent to specialized nurses in this example.

When one conditions on the *current* occupation *after* a potential move would have occurred, as is necessitated by the type of data available in the CPS, one effectively conditions on the *target* occupation: in the example above, imagine applications sent by men and women are equally successful. Then, in the period after the applications are sent, the incumbents of the normal nursing position display no gender gap in the frequency of having experienced a move (none of them, man or woman, have, in the example above), and neither do the incumbents of the new specialized nursing positions (all of the men and all of the women have experienced a move).

While *statistically* correct that the target occupation explains away all gender differences in the propensity to move (male and female specialized nurses are equally likely to have moved, male and female normal nurses are equally likely to have moved), we argue that *sociologically*, this is not what we should analytically concentrate on when assessing gender (in)equality: the experienced opportunity set by current nurses considering their career advancement is clearly different for male and female nurses in our above example – men are four times more likely to seek advancement via a move than women. Only when conditioning on either *previous* occupation in data collected *after* potential moves, or using *current* occupation in data about *prospective* moves, are we able to

capture this experienced gender difference in the option set of career advancement. Controlling for target occupation, or the current occupation after a successful job move, is hence not an adequate test of the occupational hypothesis.[5]

Alternative mechanisms

We turn now to test the alternative mechanisms that existing research suggests could generate the difference in geographic mobility in the search for work between partnered men and women.

In Table 4, we control for previous occupation throughout. We hope we have made the case for doing so clear. Model 1 tests the **relative income hypothesis**. Here, relative income is coded as percentage of household income contributed by the respondent. E.g., if a respondent earns the same as her partner, this variable takes the value 50. Thus, for partnered men, a 10 percentage point increase in income share increases the probability of a position requiring a move by 0.7 percentage points. This association is significant but slight in size, yet what is interesting is that it is perfectly countered by the interaction between relative income and being a woman. This model hence tentatively suggests that if there is a relative earnings advantage for men, there is none for women. The effect of being a woman itself is no longer statistically significant (but remains sizable even in comparison to the previous models). This points towards a somewhat subtle distinction: the on average lower relative income of women in a relationship does explain a large amount of the overall observed gender difference in the demand for geographic mobility in job search; at the same time, however, we find that the relative income-mobility relationship remains strongly gendered, as women do not search for more geographically dispersed jobs as their relative income rises, in contrast to men.

In model 2, we test the **performative gender hypothesis**. For men, being the primary earner increases the probability of a job requiring a move by about five percentage points. The main effect for being a woman is marginally significant and negative, indicating that even among

[5]Note that this example only shows in what type of situation controlling for target occupation attenuates gender differences by more than controlling for current occupation. In theory, the inverse case is also possible, but these cases seem both *a priori* less likely and are not observed in any of the previous literature or in this paper.

Table 4: Linear probability models testing alternative mechanisms, restricted to partnered respondents

	Dependent Var.: Job Would Require Move			
	(1)	(2)	(3)	(4)
Woman	-0.0349	-0.0445*	-0.0395+	-0.0237
	(0.0235)	(0.0219)	(0.0222)	(0.0307)
Relative earnings 10%	0.00720*			
	(0.00282)			
Woman×Relative earnings 10%	-0.00779*			
	(0.00365)			
Primary earner		0.0528*		0.00234
		(0.0229)		(0.0257)
Woman×Primary earner		-0.0570+		-0.0476
		(0.0300)		(0.0383)
Parent			0.0634*	0.00102
			(0.0276)	(0.0305)
Woman×Parent			-0.0703*	-0.0378
			(0.0328)	(0.0389)
Primary earner×Parent				0.116*
				(0.0472)
Woman×Primary earner×Parent				-0.0351
				(0.0618)
Constant	0.124**	0.133**	0.144***	0.122**
	(0.0425)	(0.0408)	(0.0378)	(0.0410)
Wave fixed effects	Yes	Yes	Yes	Yes
Previous occupation	Yes	Yes	Yes	Yes
Observations	10474	10474	10474	10474
R^2	0.051	0.050	0.052	0.061

Standard errors in parentheses

+ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Previous occupation is the current or most recent occupation of the applicant, coded to the 23 major categories of the 2010 Standard Occupational Classification System. Each model includes a control for the total number of applications reported by the applicant across all nine survey waves. All standard errors are clustered on the respondent.
Source: NLSJS

secondary-earner men and women, women are less likely to apply for jobs that require relocation. The negative coefficient on the interaction between being a woman and a primary earner suggests women do not experience this boost when they are the primary earner, as already hinted at in the results to model 1. While our evidence does not amount to a full backlash effect, it again suggests

that if men do broaden their searches when they are the primary earner, this is not the case for women.

In model 3, we test the **parenthood hypothesis**. This hypothesis predicts that parenthood affects men and women differently. Women's local ties become stronger, men's breadwinner role is amplified, specialization within the household may occur. Mothers may face penalties in the labor market, while fathers experience premia. We are not able to disentangle between these different mediators through which parenthood may operate, but we do see a general association between parenthood and geographic breadth of search that is in line with the broader hypothesis. Note that even when controlling for children and the interaction of children with being a woman, the main effect on the gender dummy is still marginally significant and sizably negative.

In model 4, we test for a three-way interaction between gender, parenthood, and being a primary earner. Shauman and Xie (1996) find gender differences come into play once a family has children. This suggests that parenthood acts as a mediator: performative gender norms take effect when women become mothers and men fathers, but are not relevant for partnered childless couples. To test for this, we run the analysis of model 2 (performative gender norms hypothesis) allowing for an interaction between being a mother or father and being the primary earner. We do not find a strongly significant effect on this three-way interaction, but all signs remain the same as in previous models.

Discussion

Gendered geographic mobility remains a persistent fact. We show that partnered women are significantly less likely to apply for jobs that would require relocation than partnered men. In contrast to previous literature, we find that differences in the occupational make-up between genders cannot fully explain differences in mobility. While we cannot precisely pin down the exact mechanism through which this gendered inequality then emerges, we present strong evidence that performative gender norms play a crucial role: first, relative income does positively matter for mobility – but significantly less so for women than for men, suggesting that women who earn more have to compensate by reducing their further demands on the family and moderating their career aspirations; second, women who are the primary earner do not benefit from the boost to mobility that men who are the primary earner experience, further indicating that these women need to act more in accordance with gender norms; third, parenthood affects women's mobility more negatively than men's mobility, hinting at differences in how parenting tasks and concerns are distributed based on traditional norms.

Gender differences in the occupational make-up have been a key factor in the study of gender differences in geographic mobility. We argue that the analysis of what different types of occupation controls imply for the experienced mobility of men and women has received too little attention. We present a stylized theoretical example in which men are much more likely to relocate than women but controlling for the target occupation of the move hides this existing gap between men's and women's option sets. Contrariwise, the difference remains visible when controlling for current occupation in our example. In accordance with this theoretical consideration, our empirical results indicate that taking into account gender differences in the current (or source) occupation of job applicants does account for some of the total difference in mobility – but significantly less so than controlling for the target occupation of these job applications.

While it may at first appear ambiguous what the correct occupation control is, we strongly urge that current occupation is the more sociologically relevant control. Current occupations structure

the option sets that men and women face. In contrast, controlling for target occupation in a sense hides mobility differences by only considering the "winners" of the geographic mobility game.

Consider a highly clustered occupation that is heavily dominated by men. It is feasible that the proportion of women for who the current job in this occupation required a move is equal to the proportion of men for who this job required a move. Hence, when one controls for target occupation, and hence compares proportions of men and women in this occupation that had to move, one finds no gender difference. At the same time, it may be that women faced significantly higher obstacles to moving into this occupation from adjacent occupations than men in the same occupations. Hence, when comparing men and women in these adjacent occupations – i.e., when controlling for current occupation – one sees that the opportunity set of men and women is not equal.

Some of the hypotheses advanced in this paper cannot be tested completely adequately with the data to which we have access. For example, more fine grained income data, data on the application decisions for both spouses in a partnership, and a longer panel of job search behavior may shed greater light on questions of how income and gender norms interact. On an even more fundamental level, current parenthood is a poor control for how prospective, current, or past motherhood and fatherhood affect occupational and geographic sorting: women who anticipate higher child-bearing costs will seek different jobs and occupations than men (Bass, 2015). Our data set cannot differentiate between women who never intend to have children and hence can plan their careers without such concerns and women who strongly anticipate children and seek to accommodate this in their plans.

Our hope is that other researchers can test the mechanisms sketched in this paper and other more fine-grained mechanisms with even more detailed data – including target occupation.

References

Bass, Brooke Conroy. 2015. "Preparing for parenthood? Gender, aspirations, and the reproduction of labor market inequality." *Gender & Society* 29:362–385.

Benson, Alan. 2014. "Rethinking the Two-Body Problem: The Segregation of Women Into Geographically Dispersed Occupations." *Demography* 51:1619–1639.

Benson, Alan. 2015. "A Theory of Dual Job Search and Sex-Based Occupational Clustering." *Industrial Relations: A Journal of Economy and Society* 54:367–400.

Bertrand, Marianne, Emir Kamenica, and Jessica Pan. 2015. "Gender identity and relative income within households." *The Quarterly Journal of Economics* 130:571–614.

Bielby, William T and Denise D Bielby. 1992. "I will follow him: Family ties, gender-role beliefs, and reluctance to relocate for a better job." *American Journal of Sociology* 97:1241–1267.

Borjas, George J, Stephen G Bronars, and Stephen J Trejo. 1992. "Self-selection and internal migration in the United States." *Journal of Urban Economics* 32:159–185.

Boyle, Paul, Thomas J Cooke, Keith Halfacree, and Darren Smith. 2001. "A cross-national comparison of the impact of family migration on womens employment status." *Demography* 38:201–213.

Correll, Shelley, Stephen Benard, and In Paik. 2007. "Getting a job: Is there a motherhood penalty?" *American Journal of Sociology* 112:1297–1338.

Fleche, Sarah, Anthony Lepinteur, and Nattavudh Powdthavee. 2018. "Gender Norms and Relative Working Hours: Why Do Women Suffer More Than Men from Working Longer Hours Than Their Partners?" *AEA Papers and Proceedings* 108:163–168.

Garip, Filiz. 2008. "Social capital and migration: How do similar resources lead to divergent outcomes?" *Demography* 45:591–617.

Goldin, Claudia, Sari Pekkala Kerr, Claudia Olivetti, and Erling Barth. 2017. "The expanding gender earnings gap: evidence from the LEHD-2000 census." *American Economic Review* 107:110–14.

Hochshild, Arlie. 1989. *The second shift: Working parents and the revolution at home*. New York: Viking.

Hsieh, Chang-Tai and Enrico Moretti. 2019. "Housing Constraints and Spatial Misallocation." *American Economic Journal: Macroeconomics* .

Kaplan, Greg and Sam Schulhofer-Wohl. 2017. "Understanding the Long-Run Decline in Interstate Migration." *International Economic Review* 58:57–94.

Killewald, Alexandra. 2013. "A reconsideration of the fatherhood premium: Marriage, coresidence, biology, and fathers wages." *American Sociological Review* 78:96–116.

Mincer, Jacob. 1978. "Family migration decisions." *Journal of Political Economy* 86:749–773.

Molloy, Raven, Christopher L Smith, and Abigail Wozniak. 2011. "Internal migration in the United States." *Journal of Economic Perspectives* 25:173–96.

Molloy, Raven, Christopher L Smith, and Abigail Wozniak. 2017. "Job changing and the decline in long-distance migration in the United States." *Demography* 54:631–653.

Pager, Devah and David S. Pedulla. 2015. "Race, self-selection, and the job search process." *American Journal of Sociology* 120:1005–1054.

Pedulla, David S. and Devah Pager. 2018. "Race and Networks in the Job Search Process." *Working Paper* .

Rivera, Lauren A. 2017. "When two bodies are (not) a problem: gender and relationship status discrimination in academic hiring." *American Sociological Review* 82:1111–1138.

Sassler, Sharon and Amanda Miller. 2017. *Cohabitation nation: Gender, class, and the remaking of relationships*. Univ of California Press.

Schneider, Daniel. 2011. "Market earnings and household work: New tests of gender performance theory." *Journal of Marriage and Family* 73:845–860.

Schwartz, Christine R. 2010. "Earnings inequality and the changing association between spouses earnings." *American Journal of Sociology* 115:1524–1557.

Shauman, Kimberlee A. 2010. "Gender asymmetry in family migration: Occupational inequality or interspousal comparative advantage?" *Journal of Marriage and Family* 72:375–392.

Shauman, Kimberlee A and Mary C Noonan. 2007. "Family migration and labor force outcomes: sex differences in occupational context." *Social Forces* 85:1735–1764.

Shauman, Kimberlee A and Yu Xie. 1996. "Geographic mobility of scientists: Sex differences and family constraints." *Demography* 33:455–468.

Sorenson, Olav and Michael S. Dahl. 2016. "Geography, Joint Choices, and the Reproduction of Gender Inequality." *American Sociological Review* 81:900–920.

Udry, Christopher. 1996. "Gender, agricultural production, and the theory of the household." *Journal of Political Economy* 104:1010–1046.