# Essays in Labor Economics

Jason A. Gates

April 25, 2025

*Submitted in partial fulfillment of the requirements for*

*the degree of Doctor of Philosophy in Economics*

Tepper School of Business

Carnegie Mellon University

## Abstract

In the first chapter, "Family Ties: The Reciprocal Influence of Parent and Adult Children Location Choices", I examine the relationship between the location decisions of adult children and those of their parents. It establishes the mutual influence of location decisions between agents and the role that family ties play in keeping young adults in less productive locations is investigated. Using early parental death as an instrument, the migration-altering effect of parental ties on young adults is estimated, finding parental ties result in substantial reductions in young adults' mobility.

In the second chapter, "Together or Apart: A Structural Model of Intergenerational Location Decision-Making", I estimate a dynamic discrete choice model of the co-location decision of adult children and their parents. Both the moving costs and the heterogeneous utility of parent-child proximity are estimated. Using these estimates, I conduct a counterfactual analysis in which I remove the mutual influence that parents and children have on each other's location preferences while maintaining the current regime of moving costs. In this scenario, I find substantial increases in overall migration rates and a significant reallocation of adult children across labor markets.

The third chapter, "Depressing Payment: Hospital Mergers and the Wage-Benefit Tradeoff" examines how hospital mergers affect local labor markets through their impact on the cost of employer sponsored health insurance. I estimate the effect of a hospital merger occurring within a commuting zone on the wages, employment, hours worked, part-time work, and full-time work of non-health care employees. I find evidence that mergers reduce wages by approximately 1-2 percent in local commuting zones where they occur. I also find evidence that they reduce average hours worked and that this can be explained in part by substitution from full-time workers to part-time workers, possibly to avoid the then costlier benefit provision.

# Chapter 1

## Family Ties: The Reciprocal Influence of Parent and Adult Children Location Choices

### Introduction

High levels of internal migration have long been a distinguishing feature of the U.S. economy compared to the rest of the world (see Long 1992, Esipova et al. 2013). At first glance, this high level of domestic mobility of labor should allow for a more efficient allocation of labor and reduced geographic disparities in the price of labor across labor markets; those in low-wage, low-productivity markets could reallocate their labor to higher-wage, higher-productivity ones. However, geographic wage inequality in the US is persistently high and increasing, causing concern among public officials (U.S. Department of Commerce 2023). Partly as a result of this, there has been a wave of research on the economic causes and consequences of domestic migration in the US. Key questions have been raised as to whether the disparity in wages is attributable to factors preventing migration or is explicable in terms of overall price level differences between local markets, which would result in little incentive for workers to move to higher wage markets. Recent work has pointed to a variety of push and pull factors that prevent or encourage migration to high productivity areas (see Hsieh and Moretti 2019; and Diamond 2016 as recent examples).

While high moving costs and an attachment to a home location have been noted features of US domestic migration (see Kennan and Walker 2011) and there has been work on the role family ties in the form of marriage play in inhibiting migration (see Gemici 2013) the direct role that parental ties and the desire to be near parents have on the migration decisions of working age adults is less studied. This study attempts to illuminate several patterns in the location-choice behavior of young adults and their parents that are of particular relevance to the internal migration literature. I document the overall patterns in the rate that adult children and their parents migrate and co-locate, indicating parent mobility is a factor in need of consideration. I find evidence that home location is not as important for child migration decisions as the current location of the parent. Finally, by analyzing adjusted means of migration rates between children whose parents are deceased and those who are not, I find evidence that having parental ties reduces overall young adult migration rates. Collectively, this suggests parental ties and location preferences are a key determinant of young adults migration choices and that they serve an inhibiting role in the reallocation of labor at a national level.

### Related Literature

This work fits into a large body of literature recently looking at domestic migration rates within the US, of which Jia et al. (2023) provides an extensive overview. In particular Molloy et al. (2011) and Molloy et al. (2014) have documented national trends in migration rates in recent decades. Molloy et al. (2011) examines migration rates across a variety of demographic characteristics, documenting an overall decline in domestic migration rates in recent years. Molloy et al. (2014) extends this work and finds that individual moving choices tend to be associated with changes in labor market circumstances rather than other life events. They also find that the declining migration rate tracks declines in employer and occupation switching, indicating that changes in the US labor market more broadly may be an important factor in determining overall migration rates. A related literature examines the role international migrants play in local labor markets. Cadena and Kovak (2016) study the role international Mexican migrants play in helping local labor markets adjust given this recent decrease in domestic migration among native-born workers.

Another relevant body of literature documents patterns in parent and child co-location and cohabitation. Choi et al. (2020) provide an overview of national estimates of the spatial distance between adult children and their parents. They find around three-quarters of these pairs live within 30 miles of each other. They did however find large variations in proximity across demographic characteristics indicating these decisions are in part related to socioeconomic characteristics. Compton and Pollak (2015) used the National Survey of Families and Households in conjunction with Census Data to document predictors of parent and child proximity. Their findings suggest college graduates and older children are less likely to live close to their parents. They find having grandchildren is not correlated with close proximity and that mothers having a disability is correlated with living in the same house but not with overall proximity. This suggests proximity is not necessarily driven by children assisting parents or parents assisting with grandchildren is, but may have benefits outside these scenarios. On the other hand, Choi et al. (2014) use the Health and Retirement Study to examine older adults developing at least one activity of daily living limitation. Using a multinomial logit model, they found residential proximity of spouses and children prevented nursing home enrollment. This indicates the proximity of parents and children potentially provide substantial benefits to older adults experiencing adverse health events.

Thirdly, there is a group of reduced form literature that looks at migration of children and its relationship to parental proximity. Spring et al. (2017) look specifically at local movers within metropolitan areas using a discrete choice model. They find parent age is a strong predictor of an adult child moving close and that living close to a parent reduces the likelihood that a young adults moves. Reyes and Shang (2024) use the HRS to assess changes in parent and child proximity when parents experience health shocks in the form of the onset of a functional limitation or cognitive impairment. They find evidence that parents and children remain closer or move closer in response to these shocks and find significant gendered interactions, with mother-daughter combinations more likely to move closer. Chan et al. (2025) look at return migration among young adults aged 25 to

29 in the form of returning to the same household as their parents. They document patterns in socioeconomic characteristics of these movers and find that those who return home tend to return to weaker labor markets.

## Data

The data used for this analysis comes from the Panel Study of Income Dynamic (PSID). The PSID is a longitudinal study following families that was started in 1968 and collects information on employment, wages, expenditures, education, health, and, critically, location. The use of the PSID for this study lies in the fact that when the children in a family unit move out of a household, the survey establishes a separate family unit for them and continues to follow them. This setup provides a richly detailed panel dataset with multiple generations of families. Here I use data from the years 2001-2019 limiting the sample to parent-child pairs where the children are between the ages of 22 and 43. At this age the children of the adult children begin to leave their family units in large numbers to start their own, and many of the children become the parents of adult children themselves. During this period the survey is conducted bienially. Though the survey inquires about between-wave moves, I only use information from the years the survey is conducted, in order to have consistent data on characteristics only collected point-in-time, such as parental health. Therefore each time-period is a two year period and any rates of change should not be interpreted as annual, but bi-annual. In order to avoid double counting a single parent or child I limit each individual to one appearance in a pair, preferencing those who participate in more waves and then randomly selecting among ties.

I gather data on parent and child education, work status, marital status, retirement status and parental health in the form of the parent having a reported Activities of Daily Living (ADL) limitation. Using restricted geocoded data, I also have information on parent and child location at a granular level. For this analysis I consider parent and child location to be which Core Based Statistical Area (CBSA) they live in. CBSAs consist of Metropolitan and Micropolitan Statistical Areas as determined by the United States Census Bureau. I also use the 9 Census Divisions as alternative geographic definitions for comparison, as these are frequently used in the migration choice literature (see Chapter 2). Individual's education status is considered fixed over time at the highest level ever attained, as is parent's marital status since there is little variation over time in this measure. Home location when used is considered the individual's location at age 17 or 18, depending on which age they were in the year the survey was conducted. The location and demographic-specific wage data is gathered from the Current Population Survey for the years 2001-2019.

### Results

In Table 1 I document migration and co-location patterns among parents and their adult children by a variety of demographic characteristics. Across all demographic slices, a majority of adult children and their parents live in the same CBSA. This could be driven by a strong preference to locate near one another, or simply reflect the fact children begin their lives with their parents are moving away is costly. Migration rates for children are higher than for parents across all demographic characteristics. However the difference is not of an order of magnitude. and parents are fairly mobile, with over 4 percent moving CBSAs on average over any given two-year period. This is similar to the rate that adult children move Divisions, which have frequently been used in empirical economic models of migration and are considered economically relevant.

Table 1. Migration and Co-location Patterns among Adult Children and Their Parents

| | By Child Characteristic | | | By Parent Characteristic | | |
|---|---|---|---|---|---|---|
| | Moves CBSA | Moves Division | Same CBSA | Moves CBSA | Moves Division | Same CBSA |
| All | 0.117 | 0.049 | 0.717 | 0.043 | 0.016 | 0.717 |
| Married | 0.112 | 0.049 | 0.597 | 0.041 | 0.015 | 0.697 |
| Unmarried | 0.120 | 0.049 | 0.798 | 0.048 | 0.015 | 0.697 |
| College | 0.157 | 0.074 | 0.599 | 0.053 | 0.022 | 0.613 |
| Non-College | 0.090 | 0.033 | 0.796 | 0.039 | 0.014 | 0.755 |
| Retired | - | - | - | 0.044 | 0.017 | 0.690 |
| Non-Retired | - | - | - | 0.042 | 0.016 | 0.738 |
| Grandchild | 0.094 | 0.036 | 0.698 | - | - | - |
| No Grandchild | 0.139 | 0.063 | 0.736 | - | - | - |

Among both parents and children, migration rates at both the CBSA and Division level are higher for the college-educated. Rates are also higher for adult children without children. Retirement status and marital status do not have clear differences in migration rates. It is important to note these are just subsample means and are not adjusted for other demographic characteristics. Across the board, the migration rates are substantially higher at the CBSA level than at the Division level, which is expected given many CBSA moves would still be within the same Division. The differences seem to be somewhat more pronounced among the parents, indicating they may be less likely to make farther cross-Division moves relative to the rate they move CBSAs.

The rate at which parent-child pairs co-locate can vary substantially by demographic, particularly with the characteristics of the children. College-educated children and parents are less likely to live in the same CBSA, as are retired parents and adult children who are married and or have children of their own. The relationship between child marital status and co-location is not surprising given that these couples would have another set of parents and may need to decide between living with one or the other. This may drive the difference among pairs with grandchildren and those without, since the grandchild would have another set of grandparents to possibly be close to and provision care.

Table 2. Average Hourly Income in Next Period Location

| Kids | Mean | SE |
|---|---|---|
| Non-movers | 24.61 | 0.07 |
| Movers | 27.86 | 0.26 |
| | | |
| Parents | | |
| Non-Movers | 25.21 | 0.07 |
| Movers | 25.97 | 0.46 |

Note: Average income is estimated from the Current Population Survey (CPS) and is adjusted for age, sex, education, and location.

Table 2 presents the inflation adjusted average hourly income in the CBSAs of movers and non-movers of parents and adult children, adjusting for age, sex, and education. It is widely observed than young individuals tend to move to higher wage locations and that pattern holds for this adult children sample. The next-period location of moving children have, on average, wages three dollars per hour higher than the locations of those who choose to remain. This pattern does not hold among parent movers, with no significant difference in average wages of these destination choices. This suggests wages may play a significant role in young adults migration decisions, but a less important role for parents.

Table 3. Child Location Decisions when Not Home or With Parent

| | Est. | SE |
|---|---|---|
| Move Home, Parent Home | 0.140 | 0.010 |
| Move Home, Parent not Home | 0.006 | 0.002 |
| Move to Parent, Parent not Home | 0.055 | 0.005 |

Note: Locations defined as CBSA, apart defined as not residing in the Same CBSA. Home location defined as adult child's location at 17 or 18.

Table 3 presents the results of my descriptive investigation into the comparative relevance of adult children's home location and parent location. These estimates are for parent-child pairs not currently living in the same location, defined as CBSA. Here we see different 'return rates' at which children move home or to their parent. By row, these are the rate at which children move (1) to their home location and parent location when they are the same, (2) to their home location location when their parent is not there, and (3) to their parent location when the parent is not in the child's home location. The return rates to the parent locations are both significant, with the rate when they coincide being the highest. However, when the home location does not align with the parent location, very few adult children move to their home location. The rate is nearly an order of magnitude smaller than the other rates. This strongly suggests that home location is substantially less important than parent location in driving return migration rates for young adults.

Table 4. Adult Children and their Parents' Migration Rates when Apart

| | Est. | SE |
|---|---|---|
| Child Moves to Parent | 0.087 | 0.005 |
| Share of All Child Moves | 0.421 | 0.018 |
| Parent Moves to Child | 0.031 | 0.003 |
| Share of All Parent Moves | 0.427 | 0.033 |

Note: Estimates are for adult children and their parents who are not currently residing in the same CBSA.

Child 'move-to' rates as well as parents' are compared in Table 4. The subsample investigated here includes parent-child combinations that are not currently residing in the CBSA. We the rate at which children move to parents is higher than the rate at which parents move to children. As a share of all moves however, we see that parent move-to rates are about the same share of moves as the children's are. The difference in the move-to rates seems entirely driven by the overall differences in migration rates between parents and children. This suggests that parent and child decisions are mutually influencing and that the relationship between their location decisions is not a one-way influence of parent location on child location decisions and that when examining the dynamics of both location choices it is important to consider the influence of children's choices on their parents as well.

Table 5. Migration Rates Among Adults
Without and With Deceased Parents

| Est. | SE | Est. | SE |
|---|---|---|---|
| 0.067 | 0.002 | 0.064 | 0.009 |
| N=15,349 | | N=781 | |

The rest of the results in this chapter attempt to tease out the influence of parents on their children's migration choices by comparing young adults with parents living to those whose parents have died. Therefore the sample expands to include the individuals 22-43 who report all of their parents as dead. Table 5 shows the baseline migration rates for these two groups. Here we see the rates do not differ significantly between the two groups. However, we expect the socio-demographic characteristics of those whose parents died while they were young to be substantially different from those who did not.

Table 6. Sample Characteristics

| | Living Parents | Deceased Parents |
|---|---|---|
| Female | 0.515 | 0.638 |
| Married | 0.399 | 0.376 |
| Kid College | 0.403 | 0.230 |
| Parent College | 0.270 | 0.051 |
| Grandchild | 0.497 | 0.657 |
| Parents Married | 0.766 | 0.332 |
| | | |
| Sample Size | 15,349 | 781 |

In Table 6 we see this is born out by the data. Those with living parents tend to be much more likely to attend college and have (had) a parent who attended college, as well as to have (had) parents who were married. They are also less likely to have a child of their own. Therefore I run a least-squares regression adjusting the mean in migration rate by these characteristics. These results are presented in Table 7. Here we see that adjusted for the same characteristics, migration rates among those with deceased parents are around 3 percentage points higher than those without.

While not causal, these results are suggestive that parental ties play a significant role in child migration decisions and resulting rates.

Table 7. Relationship Between Parental Death and Migration Rate

| | Est. | p-value |
|---|---|---|
| All Parents Deceased | 0.034 | 0.001 |
| Female | -0.006 | 0.123 |
| Married | 0.009 | 0.043 |
| College | 0.041 | $<0.001$ |
| Grandchild | -0.018 | $<0.001$ |
| Parent College | 0.033 | $<0.001$ |
| Parents Married | 0.004 | 0.341 |
| | | |
| Age Fixed Effects | Yes | |
| Location Fixed Effects | Yes | |
| | | |
| R-Squared | 0.0271 | |

## Conclusion

The results presented here suggest there are several important aspects of parent and child migration decisions to consider in any analysis of them. Firstly, the rate of co-location is significant and needs to be explained. This could be driven by starting locations and moving costs, or by a strong preference to locate near each other, or both. Secondly, parents do move, while at a lower rate than their children. Parent migration rates between CBSAs are similar to the rates of children at the Census Division level. Given the extent to which Divisions are used as locations within the migration literature, this suggests parent migration is significant enough to warrant attention on its own. Thirdly, children seem to respond more to their parent's location rather than their home location per se. This suggests it is more important and relevant to track parent location when analyzing young adult migration than it is to track child home location. I also find that child locations seem to have around the same relevance for parent choices as vice versa. While overall migration rates are lower for parents, the share of moves that are to their children are about the same as the share of children moves that are to their parent. Lastly, by looking at the young adults who have lost their parents, I see elevated levels of mobility, suggesting parental ties play an inhibiting role on migration for young adults.

# Chapter 2

# Together or Apart: A Structural Model of Intergenerational Location Decision-Making

## Introduction

Do family ties between adult children and their parents inhibit internal migration in the United States and if so, how? The evidence presented above suggests several things that need to be considered if we hope to answer this question. (1) Co-location is an important aspect of parent and child decisions and needs to be explained. (2) Parent movement seems to be significant enough and influential enough to warrant consideration in young adults' location decisions. (3) Home locations of adult children are not nearly as important for their migration decisions as the location of their parents. (4) The influence of child locations on where parents move is as important as the influence of parent locations on where children move. Therefore analysis of the problem needs to grapple with the fact that the location decision process of adult children and their parents is a necessarily dynamic problem with interlacing benefits and costs to both parties. A structural model with built-in assumptions around decision-making process could help come to terms with this nest of facts. In this chapter I build and a model capable of disentangling the mutually-influencing and simultaneous nature to the joint-decision problem in order to capture how much (im)mobility of labor in the United States is due to this type of family tie.

The purpose of this chapter is to isolate the utility benefit parent-child proximity from other factors that determine whether individuals migrate or not and why. The model incorporates both parent and adult children location decisions, is capable of capturing the differences in utilities between co-locating and not, and captures moving costs in a way that allows me to isolate the effect of these separate from an unwillingness of parents and children to separate. I can separately identify the moving cost structure from the utility benefits of proximity through differentials in initial migration and return rates. With the model estimated I can perform counterfactual analyses that reveal the location choices of parents and adult children with this type of family tie removed. While there have been several papers that have incorporated the important influence of family-ties on migration decisions, to the author's knowledge no previous work has sought to isolate the effect that parent-child ties have on both agents' mobility and quantify their importance relative to other factors.

The model is a dynamic discrete choice model based on the migration model of Kennan and Walker (2011), but with both adult children and parents choosing their combination of locations together. This allows both agents to coordinate across physical space and time in a manner that reflects the inter-related migration documented here. As in Kennan and Walker (2011), the value of a particular combination choice is dependent upon possible location as well as the previous period location. Unlike in Kennan and Walker (2011), given the irrelevance of home locations without the parent residing there, I do not need to track home locations. The choice set I construct for the agents is flexible enough to accurately capture parent and child proximity without expanding the level of geography to the point that estimation using full-solution methods is infeasible.

The model is structurally estimated using the Panel Study of Income Dynamics (PSID), the generational structure of which allows me to track adult children after they leave their parent's household. Model results and counterfactual analyses suggest that the utility of parent-child proximity is extremely influential over the location decisions of adult children, of a similar magnitude to the influence of wages on the decision process. Specifically, I find that eliminating the utility of the parent-child proximity results in larger flows of young adults to higher wage locations than in the baseline model. The effect on parent location choice is comparatively muted, suggesting parent-child ties are not the cause of parents not relocating to potentially preferred locations. I also find that factors often hypothesized to be important contibutors to the benefit of parent-child proximity, such as negative parental health shocks or the presence of a grandchild, are not hugely influential in the decision to live close to each other compared to other factors, which seems in line with the descriptive findings of Compton and Pollak (2015).

### Related Literature

This chapter is part of a large and growing literature of structural models of location choice. In their pioneering work, Kennan and Walker (2011) used the National Longitudinal Survey of Young (NLSY79) to estimate a dynamic discrete choice model of location choice using US states as location options, finding that income prospects play a substantial role in migration decisions. They also find that young adults' home locations are very relevant and result in a lot of return migration home even after many periods. Bishop (2007) also uses the NLSY79 to estimate a dynamic discrete choice model of location choice using metropolitan area as the relevant geographic choice. Using a two-step estimator and estimating the choice probabilities in the first step, the author estimates a variety of amenity values and moving costs. Ishimaru (2024) estimates a three stage model with college choice, labor market choice, and subsequent career being the three stages. The first two stages are discrete choice while the third is a spatial job search model. He finds that spatial differences in college and local labor market opportunities play a more important role in adult outcomes than variation in neighborhood quality.

Several studies utilize dynamic models that incorporate marital ties into the decision-making process. Gemici (2013) estimates a dynamic discrete choice model of married couples using an intra-household bargaining framework using the PSID. The results show that this form of family ties reduces both mobility and wages. Lessem (2018) estimates a dynamic discrete choice model of immigration and location choice that incorporates spousal ties through independent utility functions that vary with spouse location choice. The results suggest changes to border enforcement that alter wives migration choices, in turn lead to higher return migration among husbands. Venator (2024) estimates a unitary structural model of married couples' migration decisions, particularly looking at the effect of unemployment insurance eligibility for trailing spouses. Each member of the household shares a joint utility function and preference shocks are at the joint husband-wife level. Divorce or non-cooperation among spouses is shut down as a possibility.

While not dealing with migration, Voena (2015) studies the joint decision process of married couples. The study examines divorce laws in the United States using a dynamic model of married couples' decisions over savings, consumption and divorce. The agent's here make decisions independently by using a model of risk-sharing with limited commitment with the additional feature of a taste for marriage.

Several papers model parent child dynamics while not directly dealing with migration, but are still potentially relevant to the decision to colocate or not. Skira (2015) estimates a dynamic discrete choice model of adult daughters decisions around providing care for their parent and work. The results suggest that women are unlikely to return to work after a caregiving spell. This is potentially important for this study if a caregiving spell induces daughters to return to their parent and if any subsequent migration decisions would be for economic reasons. McMurry (2021) estimates a two period dynamic discrete choice problem of mothers with newborn children. While the study is primarily concerned with child skill outcomes, the findings also suggest a significant relationship between access to informal care and labor force participation. Mommaerts (2025) estimates a dynamic discrete choice model of parent and child decisions over long-term care using a similar framework with bargaining rather than a non-cooperative model. This is motivated by high levels of coresidence observed among adult children and parents, suggesting a non-cooperative game would be poor model choice. Here the results suggest that parent-child proximity plays a signficant role in informal care for parents and the demand for long-term care insurance.

The final three studies are the most relevant and closely related to the work presented here. Anstreicher and Venator (2022) estimate a dynamic discrete choice model of labor force participation and migration among women of child bearing age. The model incorporates grandparent contributions to childcare when proximate. They find that transfers from helpful-type grandparents are capable of covering all the childcare needs of unmarried women. Model estimates suggest a negative utility of being located close to a parent independent of the time-transfers from parents for childcare. This indicates the presences of grandchildren may play a substantial role in the willingness to locate close to a parent. Anstreicher (2024) estimates a model of intergenerational human capital investment, migration and child-rearing using a four period model with two potential moving decisions. This work is focused on the role migration plays in income mobility in the United States, finding that economic mobility is strongly influenced by migration for children from low-wage areas. The model does not incorporate adult migration. This model does incorporate future parent moves, but only when these agents are being modeled as children themselves. Therefore no simultaneous migration among parents and children is possible within the model as there is here.

Coate (2013) estimates a model of young adult location decisions, particularly examining parental proximity and occupation on children's wages. The model allows for stochastic parent moves based on a transition process. The choices the agents face are to move to stay, move to a home location, move to their parent, or make what is referred to as a 'national move' which is any other location in the United States. This national move is a random assignment to another metropolitan area following a transition process. The study is focusing mostly on proximity alone rather than the actual migration choices, but this setup would not allow me to conduct the migration counterfactuals I estimate in this work.

It is worth noting that among the papers cited above, Kennan and Walker (2011), Bishop (2007), Ishimaru (2024), Gemici (2013), Lessem (2018), Coate (2013), and Venator (2024) all include some form home location preference based on residence at an early age (varying by study). The final three papers reviewed, Anstreicher and Venator (2022), Anstreicher (2024), and Coate (2013), all explicitly incorporate the importance of the actual parent location in the decision process, with Coate (2013) also incorporating a separate home location preference as mentioned. One contribution of this chapter and the previous is to establish that, while home location may be a good proxy for parent location when this is not available to the researcher, it appears to be unnecessary to include in the child decision process when parent location is available.

### The Model

The model is a dynamic discrete choice model of joint parent-child location decisions. The model begins when children turn 22. At this point parents and children have a given initial location. In each period, the parent-child combinations know their marital, grandchild, parental health, and retirement status. As in the previous chapter, education levels and parental marital status are fixed at the maximum level they are observed to report in the same manner as Coate (2013). Transitions over child marital status, grandchildren, health, and parent retirement are stochastic. The parent-child combination also receives a payoff shock for each possible location combination. Given this state they must decide jointly where each will locate in the next period. The objective is to maximize joint utility over a 20 year horizon. Utility is affected by each parent and child's location and current state space, as well as whether the parent and child live in the same location. In the model utility is unitary within parent-child combinations. This is similar to how utility is handled in Venator (2024) for spousal pairs. While it may be less realistic to assume as much in this setting, the evidence of significant coordination and mutual-assistance found in previous studies suggests this is not a farfetched modeling choice given other options.

Let $v_t(l_{c,t}, l_{p,t}, \Omega_t)$ be the value of a parent-child joint choice with $l_{c,t}$ being the child's location and $l_{p,t}$ the parent's. Here $\Omega_t$ is the state space including previous locations $l_{c,t-1}$ and $l_{p,t-1}$. Before making a choice the parent-child duo receives a vector of payoff shocks $\xi_t$ across each possible location combination in $L = l_p \times l_c$, distributed Type 1 Extreme Value. They then jointly choose the locations to reside in to maximize

$$V_t(\Omega, \xi_t) = \max_{l \in L} v_t(l_{c,t}, l_{p,t}, \Omega_t) + \xi_{it} \tag{1}$$

Flow payoff includes both the one-period utility of the choice as well as the discounted expected value of $V_{t+1}(\Omega_{t+1})$:

$$v_t(l_{c,t}, l_{p,t}, \Omega_t) = u_t((l_{c,t}, l_{p,t}, \Omega_t) + \beta \mathbb{E}[V_{t+1}(\Omega_{t+1}|\Omega_t)] \tag{2}$$

The flow uility function is defined by,

$$\begin{aligned}\tilde{u_i}(l_{c,t}, l_{p,t}, tog_t, \Omega) = {} & \alpha_c \cdot wage(l_{c,t}, \Omega_c) + \alpha_p \cdot wage(l_{p,t}, \Omega_p)(1 - \omega_{ret}) \\ & + \phi_c \cdot sun(l_{c,t}) + \phi_p \cdot sun(l_{p,t}) + \theta \cdot tog_t + \theta_{gk} \cdot tog_t \cdot \omega_{gk} + \theta_{hlth} \cdot tog_t \cdot \omega_{hlth} \\ & + \theta_{pmar} \cdot tog_t \cdot \omega_{pmar} + \theta_{kmar} \cdot tog_t \cdot \omega_{kmar} + \theta_{sex} \cdot tog_t \cdot \omega_{sex} + \theta_{ret} \cdot tog_t \cdot \omega_{ret}\end{aligned} \tag{3}$$

Here $\alpha_c$ and $\alpha_p$ are the utility parameters associated with expected wages $wage(l_{c,t}, \Omega_c)$ and $wage(l_{p,t}, \Omega_p)(1 - \omega_{ret})$. Expected wages are estimated using least-squares regression, controlling for age, sex, and education status. Using expected wages by location follows after the work of Kennan and Walker (2011) and Lessem (2018). $\phi$ parameters reflect the utility benefit of average annual sunshine in each location $sun(l_t)$. Each $\theta$ parameter represents, in some form, a part of the utility benefit of co-locating in a location, with $tog_t$ being an indicator function whether the combination choice is one where they are together and each $\omega$ is an indicator for whether a the state space has that particular value. Specifically these are over grandchild presence, parental health, marital statuses, sex, and current retirement status.

Moving costs allow for heterogeneity over age, location distances, and location populations. See Appendix 1 for the full parameterization.

Given payoff shocks are distributed type 1 extreme value, the choice probabilities have logit form. Therefore the probability of selecting joint location combination $(l_{c,t}, l_{p,t})$ is

$$P_t(l_{c,t}, l_{p,t}|\Omega_t) = \frac{e^{u_t(l_{c,t}, l_{p,t}, \Omega_t) + \beta \mathbb{E}[V_{t+1}(\Omega_{t+1}|\Omega_t)]}}{\sum_{i \in L} e^{u_t(i, \Omega_t) + \beta \mathbb{E}[V_{t+1}(\Omega_{t+1}|\Omega_t)]}} \tag{4}$$

The model is estimated using maximum likelihood estimation. The value functions are solved via backward induction over the 20 year period, with the final period being when the children reach the starting age of the parents.

## Data and Geographies

The data used to estimate the model come from the PSID over the years 2001-2019. The sample is defined in the same manner as Chapter 1, with each unit being a parent-child combination. All adult children are ages 22-43 and each time period is two years given the bienniel nature of the PSID.

A perenniel issue for the migration literature that needs to be addressed here is the relevant unit of geographic analysis. Too fine a measure may not capture genuine parent and child proximity or may overstate the importance of nearby moves. Too broad a measure could overstate the rate of parent and child proximity while undercounting economically signficant moves. One possibility would be to follow the path of Kennan and Walker (2011) and use the 50 states as possible location choices. However, given the dual nature of the problem this becomes $50 \times 50 = 2500$ choices with a potentially very large state space when multiplied by the other state space elements. Ways around very large choice sets such as those employed by Bishop (2007) and Coate (2013) would not allow me to conduct the counterfactual analysis I conduct here.

One possible alternative is to use the US Census Bureau's 9 Census Divisions. Census Divisions have been used as relevant geographies in a variety of papers examining location choices including Gemici (2013), Diamond (2016), and Anstreicher and Venator (2022). However there are several potential issues with relying solely on Census Divisions. First, living in the same Census Division as a corresponding parent or child does not seem relevant for studying the co-location phenomenon I am interested in here. Living in a combination of Pittsburgh and Albany could hardly be justified as living in close proximity in the way I am trying to capture here. Another issue is that is misses out on a critical aspect of migration seen in the data, which is that young adults tend to migrate towards large metropolitan areas away from smaller ones and rural areas.

To overcome these issues I alter the previously used Census Division option in two ways. For the first issue of accurately capturing 'togetherness', I give each parent-child combination the choice to live together within a location or apart within the same location, along with the choices of living apart in each combination of differing locations. Here I define together as choosing to live within the same CBSA within a Census Division. Therefore if the parent and child choose to live in the Pittsburgh MSA, their choice tuple would be (Mid-Atlantic Division, Mid-Atlantic Division, Together). If the child lived in the Albany MSA and the parent in the Pittsburgh MSA, this would count as a choice tuple of (Mid-Atlantic Division, Mid-Atlantic Division, Apart). In addition, they would have choices of each differing Division combination and Apart. This allows me to capture with greater accuracy whether they are parents and children are living near one another, without vastly expanding the choice set to be every CBSA combination. For example in the basic 9 Division model, on top of the 81 combinations there would be 9 more choices.

To overcome the second issue of trying to realistically capture the migration patterns of individuals to and from large MSAs, I divide each Census Division in two, with the counties located in large metro areas being one part and the remainder being the rest. Here I define large as the 15 largest MSAs in the 2000 Census. This has the effect of adding 7 more locations to the choice set. In this way we can think of example locations as 'Large Metropolitan Pacific Division' or 'Non-Large Metropolitan Pacific Division'. Figure 1 shows the resulting locations with each color representing one of the 16 possible locations. The together and apart option also exists for the 'Metropolitan' options that consist of more than one MSA. So, for example, Philadelphia MSA and New York City MSA constitute one location, but a combination of parent and child residing in each separately would still be considered a choice tuple of (Large Metro Mid-Atlantic, Large Metro Mid-Atlantic, Apart). This change also increase the number of between CBSA moves captured by the model. Cross-Division moves alone capture 40 percent of all moves while this split captures 60 percent. With both of these changes the choice set expands to 270 options, but remains computationally tractable.

Figure 1. Map of Location Options

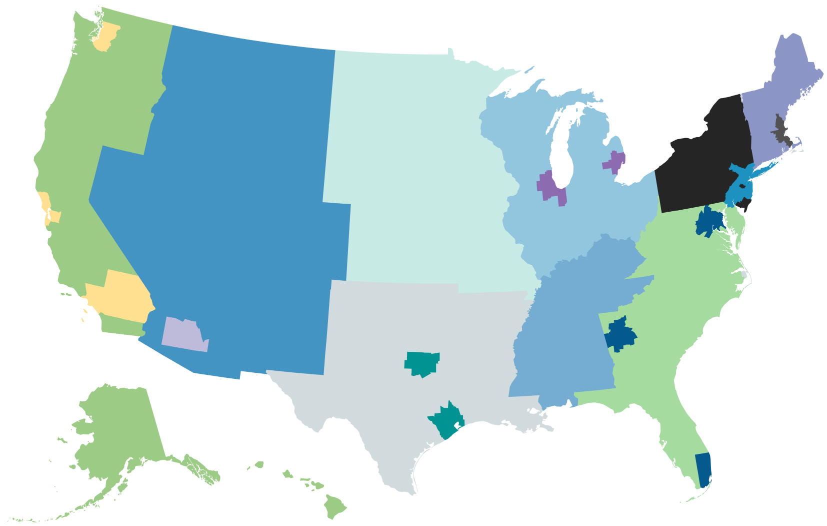

## Results

The parameter estimates and standard errors are presented in Table 8. The first section of parameter estimates are the utility benefits, while the bottom half reflect the estimates of moving costs. Here the utility benefit of hourly expected hourly wages is positive for both parents and children, although higher for children than parents. This difference may be explained by children being earlier in their careers and having more to gain over their remaining working career by relocating. The benefits of living in a sunnier place do not differ between adult children and their parnets. The results of the utility benefits of residing together and how it varies across demographic groups are suggestive of some drivers of this decision process. While there is substantial benefits to residing in the same CBSA, of the same utility benefit as a $10 increase in average hourly wages for children.

The subsequent terms are the utility benefit of residing together given specific aspects of the state space. The benefit of the duo residing in the same CBSA as a parent with an ADL limitation is positive, suggesting parent-child pairs prefer to reside near each other when negative health shocks to the parent occur. The utility of residing together when the child is married is negative. This likely captures the fact that married children also face another set of parents they may choose to reside near that is not observable to me. The utility of co-residing when there is a grandchild present is negative, which is surprising given previous findings in the literature such as Anstreicher and Venator (2022). It may be the fact that the model is picking up the effect of another set of grandparents when the adult child is unmarried to the other parent. But in either case we should not expect the effect to be so negative if the utility benefit of living together was on average the same for the grandparents we observe as those we do not.

Table 8. Parameter Estimates

| | Estimate | SE |
|---|---|---|
| *Utility* | | |
| Wage Child (per $10) | 0.641 | 0.027 |
| Wage Parent (per $10) | 0.249 | 0.029 |
| Same CBSA | 0.671 | 0.027 |
| Same CBSA with grandkid | -0.088 | 0.026 |
| Same CBSA with parent ADL Limitation | 0.200 | 0.053 |
| Same CBSA with kid married | -0.061 | 0.029 |
| Same CBSA with parent married | 0.053 | 0.019 |
| Same CBSA Female | -0.093 | 0.017 |
| Same CBSA parent retired | -0.321 | 0.025 |
| Child hundred hours of sunshine | 0.172 | 0.027 |
| Parent hundred hours of sunshine | 0.166 | 0.024 |
| *Moving Costs* | | |
| Kid age | 0.197 | 0.012 |
| Parent age | 0.872 | 0.030 |
| Population | 0.505 | 0.020 |
| Distance (hundreds of mi) | 0.362 | 0.007 |
| Switching to/from same CBSA | 1.470 | 0.044 |
| Marital moving cost | 0.527 | 0.056 |
| College moving cost | 0.898 | 0.048 |
| Grandkid moving cost | 0.459 | 0.050 |
| To/from large metro move | 0.629 | 0.044 |

Of the moving costs two terms may need some greater explication. The term for switching to or from the same CBSA captures the cost of switching from a parent-child combination choosing the same two locations twice, but switching from Together to Apart or vice versa. The cost of moving to and from a large metro area reflects the fact that distances between the large metro and non-large metro parts of the divisions are not really calculable. The distance between Seattle-San Francisco-LA and the rest of the Pacific states is not a comprehensible measure given the incontiguous nature of the large metro locations. Since this can't be captured by the distance moving cost these types of moves a captured by single switching cost parameter.

Table 9. Model Fit

| | |
|---|---|
| Top 3 Accuracy | 0.919 |
| Brier Score | 0.266 |
| Child Actual Migration Rate | 0.067 |
| Child Predicted Migration Rate | 0.086 |
| Parent Actual Migration Rate | 0.024 |
| Parent Predicted Migration Rate | 0.066 |
| Actual Together Rate | 0.691 |
| Predicted Together Rate | 0.726 |

Model fit diagnostics are presented in Table 9. The Top 3 Accuracy is 92 percent, meaning among the actual observed choices, they were among the top 3 of the 270 possible choices 92 percent of the time. The migration rates for both parent and child are somewhat higher in the model than in the observed data. It is important to note again that these rates are over a two year period so for a rough calculation of annual rates we'd see actual verses predicted child migration rates of 0.033 vs 0.043 and actual verses predicted parent migration rates of 0.012 and 0.033. The rate at which parents and adult children choose to reside in the same CBSA is very close to the model prediction.

### Counterfactual Analysis

The primary counterfactual of interest here is the effect of eliminating any utility differences between adult children and their parents residing in the same location and not. This allows me to see the alternative migration and location decisions where the moving cost structure remains the same, but any binding effect of parent-child ties is gone. In practice this means eliminating all of the $\theta$ terms from the utility function. Table 10 presents the comparison between the baseline model and the counterfactual under this regime. In this scenario migration rates incresae overall for parents and children, while at the same time the number of moves that are to the parents location decrease. The share of parents and adult children choosing to live in the same CBSA decreases from 0.715 to 0.399. This reduction simulatenously suggests that the utility benefits of co-locating are substantial, while also indicating that moving costs also play a signficant role in keeping children close to their parent. If the utility benefit were the only factor at play, we would expect choices to be much more spread out among the 900+ CBSAs and co-location rates much lower than 40 percent.

Table 10. Counterfactual Comparison

| | Baseline | Counterfactual |
|---|---|---|
| Child Migration Rate | 0.086 | 0.138 |
| Child Move to Parent Rate | 0.044 | 0.032 |
| Parent Migration Rate | 0.066 | 0.086 |
| Together Share | 0.715 | 0.399 |
| Child Wage | \$53,672 | \$54,204 |
| Parent Wage | \$56,483 | \$56,389 |

I also observe the expected child and parent wages under the baseline and counterfactual regime. This measure essentially takes the expected hourly wages and assumes they work full-time 40 hour weeks, unless retired in which case they are excluded from the calculation. Among adult children we see just over a \$500 or about 1 percent increase in annual wages. For parents, there is a negligible decline. This suggests adults children are being kept from moving to higher wage locations by these family ties, while the ties do not have the same effect on parents.

Beyond top-line rates of migration, part of the purpose of this analysis is to see how location choices are altered for the agents when these ties no longer hold. Figure 1 presents these results for both parents and children. In the first panel we see location choices among parents remain roughly stable. This indicates the distribution of parent choices are not heavily influenced by the preference to locate near their children. The right panel in comparison shows fairly substantial swings in location choices among adult children between the baseline and the counterfactual. These represent all location choices of all time periods for all agents in a kind of grand-census of locations over the twenty year period.

Figure 1. Location Choices among Adult Children and Their Parents

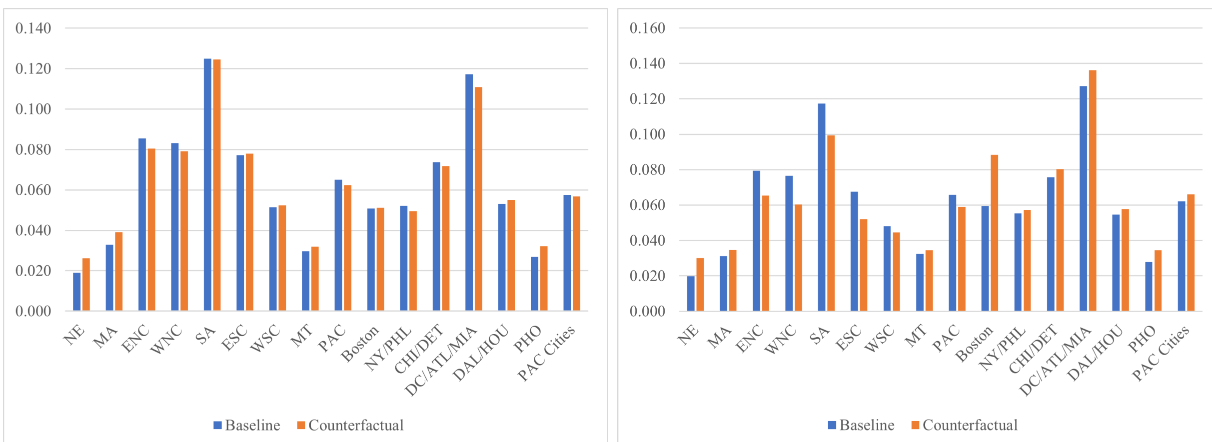


(a) Share of Parent Choices by Location

(b) Share of Child Choices by Location

Figure 2 maps the differences between the counterfactual and the baseline scenario. Here again we see parent location swings are of a much smaller magnitude than the children's. The children's location swings are large. We see each of the large metropolitan areas sees an increase in children locating in them under the counterfactual scenario with the largest losses among the the non-lare metropolitan location choices, particularly the East North Central, West North Central, South Atlantic, and East South Central Divisions. Among non-metro Divisions that saw increases, the Mountain, Mid-Atlantic and New England all saw increases in adult children location shares.

Figure 2. Percentage Point Changes in Location choice among Adult Children and their Parents

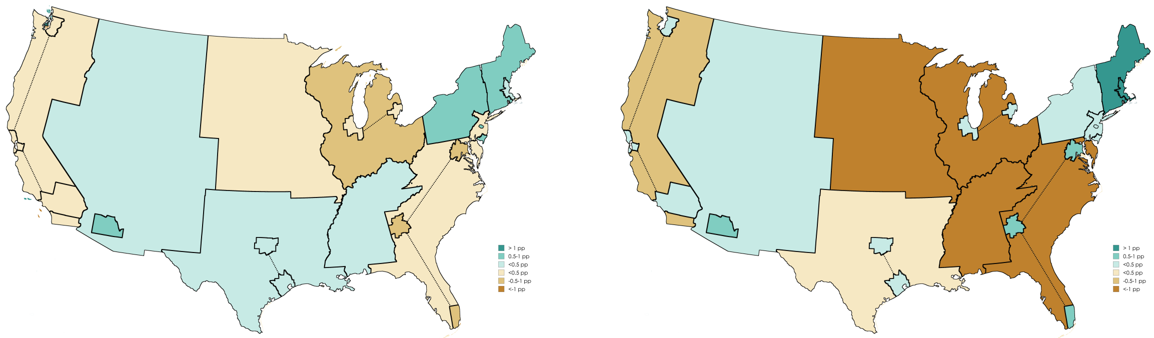

(a) Parent Percentage Point Change (b) Child Percentage Point Change

While these patterns are interesting they raise the question of what is driving the changes in migration rates and location choice distribution we are seeing. This will be further investigated in two steps. First, I am interested in how each of the aspects of the location utility affect overall results. Second, I am interested in what in the model is driving particular locations to see increases and decreases in Figure 2.

Recall, there are 7 different ways for utility to differ when parents and children reside in the same CBSA, the baseline utility increase of proximity and the subsequent variations depending on the particular realizations of the state space parent-child combinations experience. Table 7 presents the changes relative to baseline by eliminating each of the together-utility parameters individually.

Table 11. Counterfactuals Comparison Eliminating Parameters Individually

| | Child Migration | Move-to-Parent | Parent Migration | Together Share |
|---|---|---|---|---|
| Baseline | 0.086 | 0.044 | 0.066 | 0.715 |
| Full Counterfactual | 0.138 | 0.032 | 0.086 | 0.399 |
| No Base Together | 0.159 | 0.028 | 0.094 | 0.280 |
| No Grandchild x Together | 0.084 | 0.045 | 0.065 | 0.739 |
| No Health x Together | 0.089 | 0.043 | 0.066 | 0.698 |
| No Child Married x Together | 0.085 | 0.045 | 0.066 | 0.730 |
| No Parent Married x Together | 0.090 | 0.043 | 0.066 | 0.692 |
| No Sex x Together | 0.084 | 0.045 | 0.065 | 0.739 |
| No Retired x Together | 0.077 | 0.048 | 0.064 | 0.788 |

The first two rows of Table 11 represent the values of child and parent migration, together share, and the rate at which children move to their parent for the baseline and counterfactual model from Table 10. Each subsequent row contains the estimates when the individual $\theta$ parameters are removed from the baseline model. Here we see that the vast majority of variation is caused by the baseline preference to locate together or not, which closely aligns with the full counterfactual values. Among the other parameters the most substantial change from baseline is the retirement interaction, where it seems without the disutility of being together and the parent being retired, an additional 7 percent of pairs would live together. This would suggest that even given the role and benefits in terms of parental care or grandchild care, the general preference to be close to one another independent of those is the most important factor driving parent-child proximity apart from moving costs.

It is important to not think of Figure 2 as representing where individual parents and children would move if they were not stuck in their current location. Instead it is important to note that even in the baseline scenario there is substantial inflows and outflows of young adults in particular. Figure 3 shows the inflows and outflows by location of the adult children. Here we see over the 20 year time period examined, substantial flows into the large metropolitan locations in particular, and the northeastern regions more generally compared to other parts of the country. For outflows in the baseline model we see higher outflow rates among the non-large metropolitan locations, particularly in the eastern half of the US. Here it is worth noting that the Mid Atlantic and New England seem to have moderate inflow and outflow in the baseline model indicating that their overall mobility seems to be higher in northeast of the country relative to other parts.

Figure 3. Baseline Inflow and Outflow for Adult Children

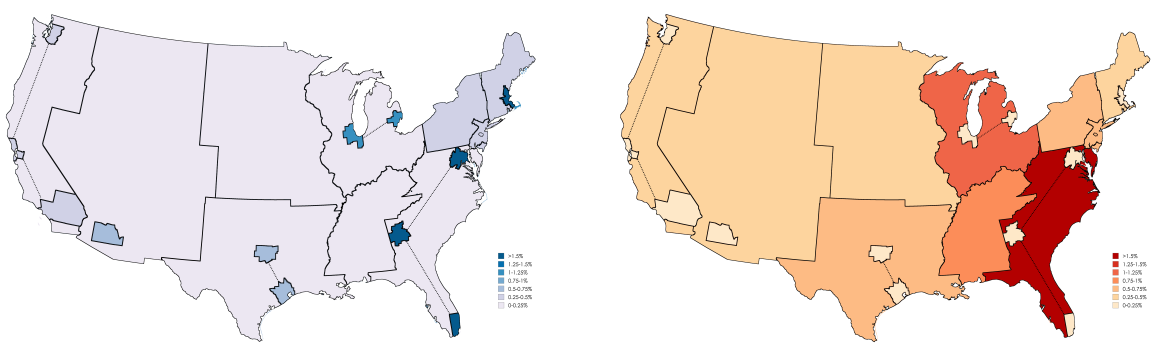

(a) Baseline Inflow

(b) Baseline Outflow

Figure 4 plots the net flow of the adult children in the baseline model. Here it makes it clear we tend to see large net inflows into the large metropolitan areas and net outflow of the non-large metrolitan locations. This suggests the model accurately captures the tendency of young adults to move to large urban areas. The exception is the large metro Mid-Atlantic region consisting of New York and Phildelphia, which had moderately high inflows and outflows, but the net was a slight outflow in the model.

Figure 4. Net Flow in Baseline

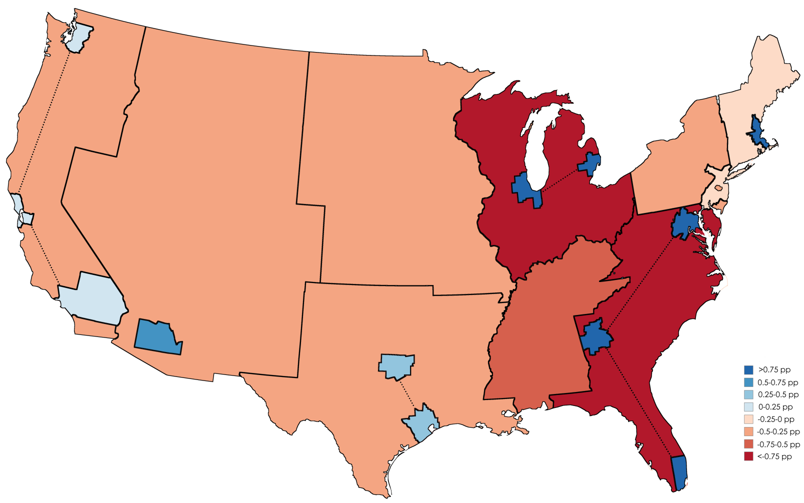


Now turning to the counterfactual with no family tie utility, we see a similar and elevated pattern of inflows and outflows in Figure 5. Inflows increase across the board when family ties are eliminated, with more moves to the large metro areas as well as more moves to the non metro geographies. Similarly outflows were up across the board as well. This highlights that overall mobility is increasing when family ties are eliminated. While young adults tied to their parents in Kentucky may now be able to move to Atlanta, there are young adults tied to their parents in Atlanta who will now move to Kentucky as well.

Figure 5. Counterfactual Inflow and Outflow for Adult Children

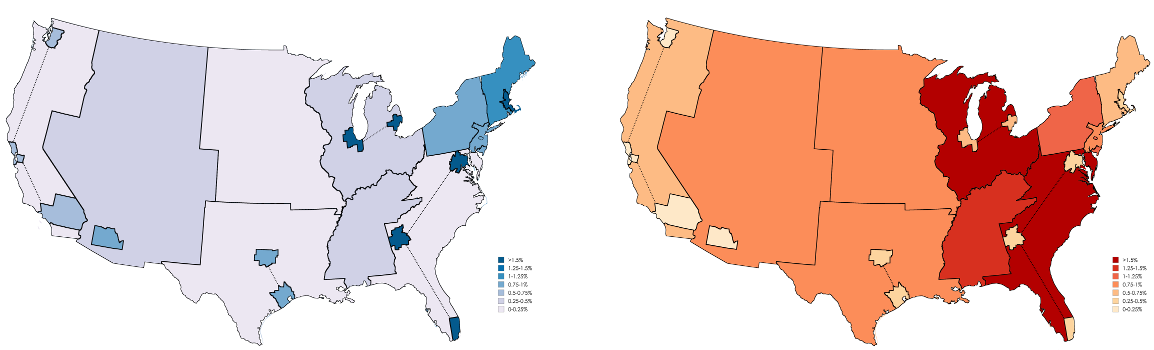

(a) Baseline Inflow

(b) Baseline Outflow

Figure 6 shows the net flow in the counterfactual scenario. Here we see broadly the same pattern as in the baseline scenario. However several net ouflow areas have their net outflows exacerbated or attenuated in this scenario, while every net inflow region sees their net inflows increased. The net outflow New England saw in the baseline flips to a net inflow in the counterfactual scenario.

Figure 6. Net Flow in Counterfactual

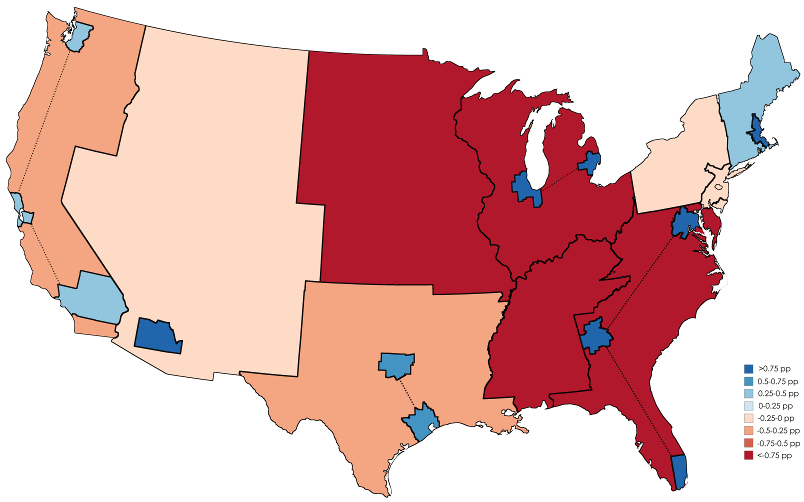


Figure 7 maps the differences between Figures 4 and 6. This shows the changes in netflow from the baseline to the counterfactual scenario. Here we see that every lare metropolitan geography sees a positive change in their net flow rate. The non large metro New England, Mid-Atlantic and Mountain Division also see incresaes. In practice this figure very closely matches Figure 2b, with the signs, but not order of changes matching between the two. This is not representing the same change though, since Figure 7 only examines flows while Figure 2 is total location share choices, even among those who never move. For instance in Figure 7, an individual who moves from rural Texas to Houston in period 1 and stays for the remaining periods counts the same as someone who lives in rural Texas for all periods except the last, in which they move to Houston. Whereas in Figure 2, these individuals would have different contributions to the total location choices in all years being counted.

Figure 7. Change in Net Flow from Baseline to Counterfactual

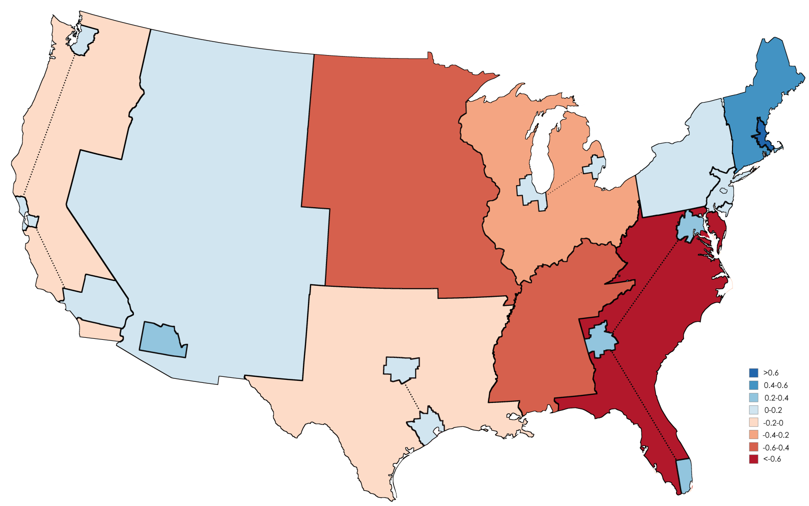


This excercise further elucidates what is happening in the model and counterfactual, with large differences in net flows between regions in both the baseline and counterfactual, but with the net flows made larger or smaller depending. The question on what exactly is driving these changes in the net flow remains a question. Within the model the largest determinants of change are the together utility, the wage utility, and the distance moving cost parameter. Appendix Figure A.1.1 contains plots of the resulting changes when variation in each is eliminated. What we care about is how the wage utility and the distance moving cost contribute to the reallocations seen in Figure 2. To further explore this I re-estimate the counterfactual while also eliminating variation in wages and in distances and see how each affects the counterfactual.

Here it is important to note that eliminating wage variation is relatively straightforward, however eliminating variation in distance moving costs is not. If all distance moving costs were eliminated altogether, migration rates would increase substantially and most importantly, the role that random shocks play in the decision process would increase in importance. However, if distance moving costs were set too high, say by just taking the average distance between all locations, only shocks great enough to justify cross-country moves would overcome them, effectively just keeping individuals in their home locations for the most part. To overcome this I find the uniform distance between locations that results in the same migration rate as in the baseline scenario. In this case it is 730 miles.

Figure 7. Change in Net Flow from Baseline to Counterfactual

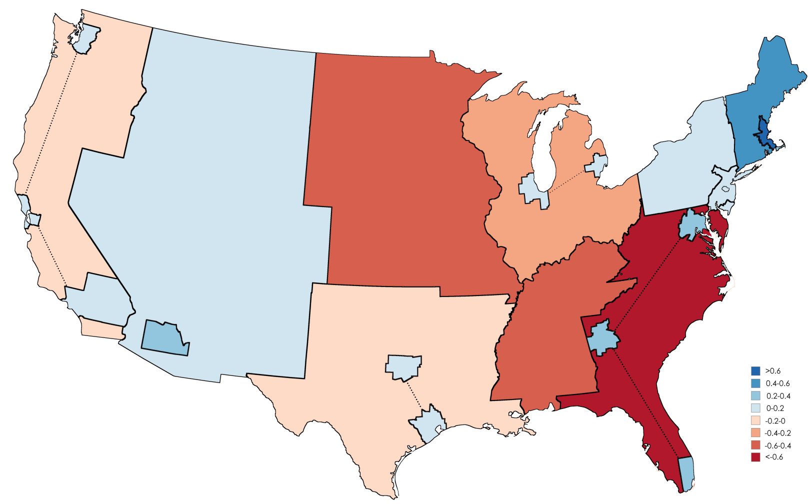


Figure 8 plots the change from baseline of the three scenarios, with the first bar demonstrating the effect of eliminating family ties, the second of eliminating family ties and elminating wage differences, and the third of eliminating family ties and eliminating distance variation. Here we see that the effect of eliminating distance tends to increase the magnitude of the effect of eliminating family ties. What this suggests is that the distance moving costs are actually working in the opposite direction as the counterfactual, and keeping the counterfactual closer to the baseline scenario than it otherwise would be. Eliminating wage variation on the otherhand, not only tends to attenuate the effect of the counterfactual on location choices, but in most cases flips the sign of the change. This suggests that without wages the effect of the counterfactual on location share changes would be the opposite of the effect. From this it seems clear that the wage variation is playing the determining role on the effect that eliminating family ties has on location choice distribution among young adults.

## Conclusion

This paper analyzes and quantifies the relative importance of parent-child ties in the joint location decisions of adult children and their parents. I find that these ties significantly hinder the mobility of the adult children, resulting in them remaining in lower wage areas than they would prefer without these ties. There is a near-halving of parent-child pairs residing within the same metropolitan area, with those remaining doing so because of moving costs and not any attachment to the physical location, which seems to be irrelevant without parents locating there. The results suggest these ties are mostly due to the utility benefit of proximity itself and not mainly driven by other observable features like grandchild presence or parental health concerns. Without these ties there would be greater migration and movement across all parts of the country, with areas that have higher wages being net recipients of young adults. Further work is necessary to ensure that these moves are in fact due to wages and not other features correlated with wages that are not captured in the model.

# Chapter 3

# Depressing Payment:

# Hospital Mergers and the Wage-Benefit Tradeoff

## Introduction

In recent decades the United States has experienced a significant rise in the concentration of health care provider markets (Gaynor and Town 2011). In particular, there has been rapid consolidation within the hospital industry as a result of a wave of hospital mergers. At the same time, health insurance premiums paid by employers and employees have been rising quickly, with the cost of a family premium for employer-based insurance reaching $\$19,616$ in 2018, a 55 percent increase over the previous 10 years (Kaiser Family Foundation 2019). Seeking to identify a relationship between these trends, a substantial literature has developed measuring the impact of hospital concentration on health care costs, with nearly all finding evidence that hospital concentration increases the prices hospitals charge (e.g. Moriya et al. 2010; Gowrisankaran et al. 2015; Cooper et al. 2019). In light of this, and given that 60 percent of non-elderly adults receive health insurance from an employer as part of their compensation (Kaiser Family Foundation 2018), these mergers may also be affecting the labor market through their impact on health care costs. For instance, while the median worker cost an employer $18.73 per hour in wages in 2019, they also cost $3.06 in health insurance premiums, making them a substantial portion of employers' compensation expenses (Bureau of Labor Statistics 2019). Any change to the cost of providing insurance will therefore directly affect how employers determine their compensation offerings. Thus, in order to fully account for the effects of these mergers, we need to understand how they affect the local labor market.

Yet to the author's knowledge no study has examined the full impact of hospital concentration on local labor markets. One previous study, Prager and Schmitt (2019), examined the impact of hospital mergers on health care workers, a market where hospitals are a major employer. This work is part of a larger recent literature studying the effect of firm concentration on local labor markets through monopsony power (Azar et al. 2017; Benmelech et al. 2018). This string of research has led to discussions of whether regulators should consider the monopsonistic effects of mergers on labor markets (Federal Trade Commission 2018). In this paper I extend this argument a step further for hospital mergers. Due to the unique way health care is financed in the United States, concentration in the health care industry reverberates not just to workers in this industry, but to workers in all industries via increased health care costs.

In this paper, I estimate the impact of hospital mergers on non-health care workers' labor outcomes including annual wages, employment status, and number of hours worked. Using both a difference-in-differences and event-study design, I compare the outcomes of workers in local labor markets exposed to a hospital merger with those in ones that were not using both the Quarterly Census of Employment and Wages (QCEW) and the American Community Survey (ACS). I find strong evidence that hospital mergers resulted in reductions in annual wages and that these effects grew over time. The point estimates for average wage suggests that hospital mergers in the period studied decreased local annual average wages by 0.8 to 1.5 percent. I also find evidence that these mergers reduced the number of hours worked and that this may be driven by shifts away from full-time employment towards part-time employment.

Altogether the results here suggest that hospital mergers do have significant impacts on local labor markets including suppressing the annual wages of workers outside of the health care industry. Beyond its relevance to merger analyses, this study also has implications for discussions surrounding wage stagnation and inequality that have occurred in recent decades. If hospital mergers are detectably suppressing wages through their ability to increase insurance premiums, they may be playing a role in the recent stagnation of real wages. I note that the work presented here is limited in that it only estimates relatively short run effects and does not incorporate any general equilibrium effects. For example, locally reduced demand as a result of lower disposable income may have macroeconomic effects as well. Data limitations also constrain me in determining how much of the wage effects are due to changes in employment-level and hours worked rather than changes in hourly wages. Further work using data better suited to determining these relationships would allow for a more accurate accounting of the dynamics caused by these mergers at a micro-level. The remainder of the paper proceeds as follows. I first review the theoretical framework of how mergers may affect premiums and how employers may respond to these changes. I then review previous findings of how employers respond to premium cost changes. The empirical model and data are then overviewed and the last section concludes.

## Conceptual Framework

### Insurers

In order for hospital mergers to have an effect on workers' wages and employment through health care costs, it would need to be established that they actually have a meaningful impact on such costs. In their study of 366 mergers occurring between 2007 and 2011, Cooper et al. (2019) find that the prices merging hospitals charge to insurance companies increases by approximately 6 percent when the merging hospitals are within 5 miles of each other. This effect dissipates to just over 2 percent for those within 25 miles of each other, with no statistically detectable effect beyond that radius. They also find some evidence that these effects increase over time, with mergers up to 30 miles away increasing prices by 6 percent after 2 years.

However these are the costs the hospitals charge to insurance companies, raising the issue of how much the incidence of these costs ultimately fall on employers and how quickly. Here it is important to note that 61 percent of workers who receive insurance through their employer are enrolled in a self-insured health plan (Kaiser Family Foundation 2019). Under self-insured plans, the employers pay medical costs directly and only contract insurance companies to design and administer these plans. Therefore any increases in prices that hospitals can effect on payers would be directly borne by self-insured employers.

For the 39 percent of employees in fully-insured plans, whose health care costs are paid by insurers while employers pay a premium to them for this coverage, the impact is less easily established. Potential insight into insurer behavior in this arrangement can be gleaned from recent work in the hospital merger literature. In their structural model of the impact of a potential hospital merger on an HMO in northern Virginia, Gowrisankaran et al. (2015) estimate the merger would increase the prices that the insurer was charged by 7.2 percent and that the insurer would then raise premiums charged to employers by 3.4 percent. While this result may not generalize to all mergers and fully-insured plans, it suggests that at least in some instances price increases faced by insurers are passed onto employers in fully-insured plans as well.

**The Model**

There remains the question of how employers will handle increases in health care costs that result from these mergers. The standard economic model of employer-provided benefits, and the one most frequently employed in this literature, is the compensating differential framework first posited by Summers (1989) and formalized by Gruber and Krueger (1991). Suppose labor demand is given by:

$$L_d = f_d(W + C) \tag{5}$$

Where $W$ is wages and $C$ is insurance costs. Furthermore suppose supply is given by

$$L_s = f_s(W + \alpha C) \tag{6}$$

Where $\alpha C$ is the monetary value that employees place on health insurance. Then

$$\frac{dW^*}{dC} = -\frac{\eta^d - \alpha\eta^s}{\eta^d - \eta^s} \tag{7}$$

where $\eta^d$ and $\eta^s$ are the elasticities of demand and supply for labor. From equation (3) we can see that if $\alpha = 1$ and employees fully value the benefit, wages fall by the full cost of the benefit. If workers do not value the benefit at all and $\alpha = 0$, then the result is identical to the incidence of a payroll tax. The impact on employment will be

$$\frac{dL}{L} = \frac{W_0 - W_2 - dC}{W_0} \cdot \eta^d \tag{8}$$

Where $W_0$ and $W_2$ are the initial and final wage levels.

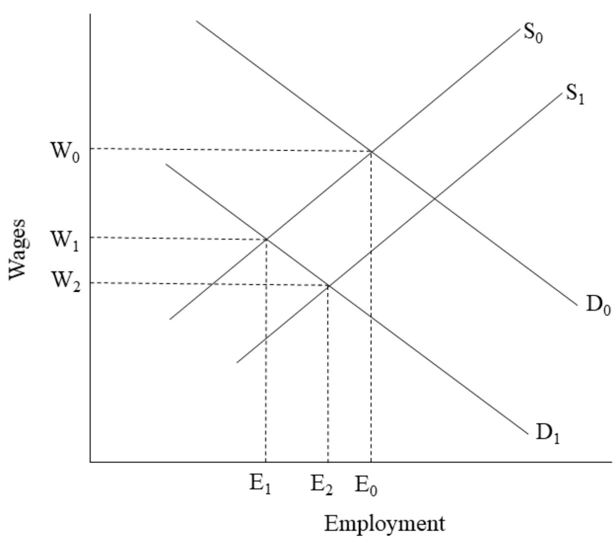


Figure (5) Labor Market Effects of Employer Sponsored Insurance

In this framework the effect of rising health insurance premiums on wages is predicted to be negative, while its effect on employment is a function of $\alpha$. If workers value the insurance they receive, that is $\alpha > 0$, the effect on employment is ambiguous. The basic workings of the model are depicted in Figure 1. The increased cost to employers shifts labor demand from $D_0$ to $D_1$, moving wages downward from $W_0$ to $W_1$ and employment from $E_0$ to $E_1$. To the extent that these benefits are positively valued by workers, supply also shifts out from $S_0$ to $S_1$, further reducing wages to $W_2$ and moving employment to $E_2$.

According to the model, the effects on employment are highly dependent on worker valuation of the benefits, $\alpha$. If $\alpha < 1$, and workers value change in insurance at less than its cost, employment will fall. If $\alpha = 1$ and employees value it at cost, there will be no impact on employment. If $\alpha > 1$, that is workers value insurance at greater than its monetary cost, employment will rise. One might think in this empirical setting we would expect $\alpha$ to be zero, since the change in the cost of insurance I am discussing is coming from an increase in market concentration by hospitals and not from value-adding innovations. However, even in this circumstance it is likely that employees value this marginal change in insurance. If insurance costs are going up for all private payers and not just employers, the cost of purchasing insurance individually or of being uninsured and paying out of pocket has increased by more than the employer premium cost. This is because, as noted by Currie and Madrian (1999), employers are capable of purchasing insurance much more cheaply than individuals and, unlike wages, benefits are not taxed. So a worker receiving the compensating differential in the form of wages and seeking individually purchased insurance would be forced to pay a higher premium than their employer is charged (or bear all the risk of being without insurance) and this money would be taxed first. This leaves us with theoretically ambiguous impacts on employment.

### Caveats

While a useful analytical tool, there are several ways in which employers and workers may behave contra the predictions of this model. For one, employment broadly defined incorporates both having a job and number of hours worked. As Cutler and Madrian (1998) note, employers may respond to increases in premiums with an increase the number of hours worked among workers with health insurance, since this is a "fixed" cost of employment. This may lead to total incomes rising for these workers. Furthermore, they may decide to replace full-time workers who receive employer sponsored insurance (ESI) with part-time workers they can exclude, potentially reducing average hours worked.

It is also possible that employers might respond by altering health insurance generosity and structure. For one, they could simply stop offering ESI or, for new firms, choose not to offer it from the start. Alternatively, they could switch to lower quality insurance with higher deductibles and copays, which could keep the total premium at the same level in the face of higher health care costs. They could also increase employee contributions to premiums. This might not make sense at first glance if the incidence of insurance is ultimately borne by the worker anyway; the only difference between an employer and employee premium is that the latter are often taxed before they are paid. However, as noted by Gruber and McKnight (2003), increased employee premiums may reduce take-up rates of insurance among employees, reducing the number of employees whose health care costs the firm is responsible for. That some individuals would not take up insurance while others would indicate heterogeneity in $\alpha$ across employees within a firm. From the point of view of this study, I am unable to see if firms absorbed the costs through increases in employee premiums since I cannot observe employee contributions.

Interestingly, in theory, increased employee contributions may increase per-enrollee premiums even if they reduced overall firm costs. Sicker workers more likely to have higher health care costs are more willing to pay the employee premium than healthier workers - the adverse selection problem. Therefore reliance on passing on increases in health care costs through increases in employee premiums may lead to worse adverse selection, resulting in a sicker and more expensive remaining pool, increasing health insurance premiums further and so on, potentially resulting in a death spiral. This would suggest firms would be reluctant to overuse this mechanism. While employee contributions have grown rapidly, they have remained a relatively constant share of overall premiums in the past 20 years (Kaiser Family Foundation 2019), potentially lending credence to the limits of this mechanism.

With this theoretical ambiguity and potential limitations of the basic compensating-differentials model, it remains an empirical question of how labor markets respond to health care cost pressures arising from hospital mergers.

## Literature Review

### Wages

The wage literature has been mixed on its findings of whether employers pass through increases in the cost of health insurance to workers in the form of reduced wages as the standard incidence model predicts. A major difficulty in studying the issue is that most surveys commonly used to study labor markets such as the Current Population Survey (CPS) or Survey of Income and Program Participation (SIPP) do not report the generosity or cost of health insurance premiums workers receive. One that does is the National Compensation Survey (NCS), which collects job-level wage and benefit information. Two key studies of how wages are altered by increasing premiums are Buchmueller and Lettau (1997) and Anand (2017), both of which use the NCS. Using a first-differences specification and regressing job-level changes in wages on job-level changes in employer premiums, Buchmueller and Lettau (1997) find no relationship between changes in the cost of insurance and wages. However, the study does not account for changes in the health insurance plan design or generosity between years, making the changes in the cost of insurance for that job potentially endogenous to the changes in wages they observe.

In an update and extension of Buchmueller and Lettau (1997), Anand (2017) uses the NCS from 2003 to 2010 to again check for a wage-benefit tradeoff. Like the earlier study, she finds no evidence of changes in wages resulting from increases in employer premiums. Unlike the earlier analysis, she restricts the data only to plans that were offered multiple years in a row in order to ensure that the price change is the result of premium changes and not changes in plan characteristics. An unfortunate drawback of this method is that 31 percent of the occupation-establishments in the sample were dropped. It is possible that these plans that were offered in one year and not offered in the next were dropped because their costs grew faster than those that were not dropped. It is also possible that plan-dropping is correlated with other employer characteristics and dynamics. Whether the third of the sample that was dropped differs in ways observable in the data is not pursued.

Gruber and Hanratty (1995) study the issue from the opposite direction, studying province- and industry-level variation in the implementation of Canada's national health insurance (NHI), which replaced employer-sponsored plans. They find that while provincial-level implementation was associated with increases in wages, workers in industries that had the highest levels of private insurance coverage prior to NHI saw wage increases lower than those with less coverage. This runs contrary to the prediction of the compensating differential model that wages would rise by the amount employer-premiums decrease (or are eliminated in this case).

Albeit not directly examining health insurance premiums, Gruber and Krueger (1991) uses the workers' compensation program, another type of employer provided insurance, to gain insight into whether health insurance premium changes would affect wages. They find that between 56 and 86 percent of increases in the cost of workers' compensation are borne by the employees in the form of lower wages. Clemens and Cutler (2014) study the impact of increasing employer premiums on public school districts. The authors find that for every \$1 increase in the cost of benefits, wages fall by \$0.15 and this point estimate was not statistically different from 0. Lubotsky and Olson (2015) also examined teachers, finding no evidence of reductions in wages as a result of premium increases.

While the above studies seem to contradict the incidence model predictions, there have also been numerous finding evidence of a tradeoff. Bhattacharya and Bundorf (2009) find that, controlling for observables, insured obese workers have lower wages relative to insured non-obese workers, while no such wage gap exists for workers without health insurance, suggesting workers pay the incidence through wages. The magnitude of their estimates suggests employees bear nearly all of the incidence of the cost of their coverage. Gruber (1994) finds evidence of a full shifting of the cost of mandated maternity benefits to the wages of those who benefit from it. Qin and Chernew (2014) look at public sector workers' wages and find a wage offset of 15 to roughly 50 percent depending on their specification.

Baicker and Chandra (2006) use state-level changes in medical malpractice payments as an instrument for health insurance premiums. They find a dollar-for-dollar offset of increases in premiums coming from the wages of workers with ESI, with no decline in wages for those without. Kolstad and Kowalski (2016) studied the impact of Massachusetts' health care reform on the labor market. Part of Massachusetts' reform was a mandate for certain employers to provide ESI, allowing the authors to check for a compensating differential. They find a compensating differential close to the average cost of ESI to the employer.

The conflicting findings within the literature on whether employees pay for changes in benefit costs via reductions in wages led Sommers (2005) to hypothesize that nominal wage constraints may prevent some employers from passing on the full incidence of premium changes to employers, at least in the short term. The idea is that nominal wage cuts are costly to firms due to perceptions of fairness on the part of workers. Therefore only through erosion by inflation can wages be cut, essentially putting a lower limit on how much real wages can change. He tests his hypothesis using the CPS by exploiting variation in inflation rates by region and finds that employers are constrained in passing premium costs onto workers by the inflation rate. However, the author still does find evidence of a partial wage offset even if constrained.

**Employment and Hours Worked**

Many of the studies that examined wages also checked for impacts on employment, particularly hours worked and binary employment measures. Similar to the findings on wages, the results on employment also vary between studies. In their study of workers' compensation premiums, Gruber and Krueger (1991) found no statistically significant impact on employment from increases in premiums. Qin and Chernew (2014) found no impact of increased premiums on hours worked among public employees. Lubotsky and Olson (2015) also found no evidence that school districts reduced the number of teachers in response to higher health insurance costs.

However, other studies have found impacts on employment and hours worked. Gruber (1994) finds evidence of both a rise in the number of hours worked and a fall of employment among the treated group of married women as a result of mandated maternity benefits. Using the SIPP and CPS, Cutler and Madrian (1998) find that rising insurance costs increased the number of hours worked of insured workers by 3 percent, suggesting firms may substitute hours per worker for the number of workers employed in response to this increase in fixed cost per worker. Montgomery and Cosgrove (1993) and Buchmueller (1999) both found an association between higher health premiums for full-time workers and the use of part-time workers within a firm.

Baicker and Chandra (2006) find a decrease in hours worked as a result of rising premiums, part of which is attributed to increases in the probability of unemployment along with increases in the probability of part-time work. In their study of the Massachusetts' health reform, Kolstad and Kowalski (2016) found a small but statistically significant decrease in the number of hours worked as a result of the employer mandate for certain firms to provide coverage to their employees. Sommers (2005) found the hazard rate of unemployment was higher and statistically significant for those with health insurance when employers faced high premium growth.

**Benefit Generosity**

As previously mentioned, while wage and employment data are relatively easy to access, data on plan design, generosity and employee contributions is quite limited. Several studies however have attempted to estimate an impact of premium growth on these. Perhaps the simplest way for employers to change benefits in response to higher premium costs is to drop the benefit altogether or change those eligible for them. Baicker and Chandra (2006) found increases in premiums decreased the likelihood of having ESI but were unable to see how much was attributable to declines in full-time employment as opposed to decreases in the offering of ESI.

Anand (2017) found that a $1 increase in health insurance costs for a worker resulted in just over a $0.50 increase in employee premium contributions. Lubotsky and Olson (2015) find that a $1 increase in premiums for individual health insurance increase the employee contribution to this premium by $0.17, while the number for family coverage was $0.46. This suggests that rather than reducing wages, employers may increase the employee contribution, thereby reducing the effective pay of those who use the insurance, while not affecting the pay of those who do not take it up. This would seem to contradict the idea presented above that employers are constrained in their ability to raise employee contributions while avoiding serious adverse selection problems.

Vistnes and Selden (2011) offer perhaps the most informative study on the non-wage effects of premium growth. Using the Medical Expenditure Panel Survey – Insurance Component (MEPS-IC), a survey of employers, they examine the impact of metropolitan area variation in insurance costs on employer offerings of insurance, employee eligibility for this insurance, employee premium contributions, and deductibles. They find that higher insurance costs reduce the likelihood that employers offer coverage to any of their employees and, among those who continue to offer coverage, to reduce the number of workers eligible for it. Like Anand (2017) and Lubotsky and Olson (2015), they find that employers increase employee premiums in response to health care price increases. They also find that firms increase deductibles in response to cost pressure, suggesting cost changes may cause firms to change plan generosity and design. If this is indeed how employers respond, there may be no observable effect on wages or employment levels in the data used for this study.

## Data

As previously mentioned, a serious issue with investigating the relationship between health care costs and wages is data availability. Ideally it would be possible to observe the price and type of service each hospital charged for each individual treated, the plan that individual was enrolled in, the cost of the premium of that plan, their employer, and a full breakdown of their compensation in terms of wages, benefits and hours worked and the same breakdown for every worker in their local labor market. It would then be possible to estimate the impact the merger had on hospital prices, the extent to which these price changes increased employer premiums, how employers dealt with these premium increases, and whether there were any spillovers into the labor market of these workers (e.g. if depressed demand for high-insurance workers at affected firms have impacts on workers at unaffected firms due to depressed demand in the market).

Even short of this it would be useful to have employer-sponsored premiums available on the surveys usually used to analyze the labor market such as the CPS or SIPP. Although even if these surveys asked respondents, employees are likely not aware how much it costs their employer to insure them. Without being able to observe these costs directly, estimating the impact that mergers have on premiums along with $dW^*/dC$ or $dL/L$ from equations (3) and (4) respectively is impossible. In light of this shortcoming, I therefore adopt a method from the "treatment effect" literature and estimate the impact of mergers by comparing the dynamics of local labor markets that were "treated" with a hospital merger to those that were not using a difference-in-differences, two-way fixed effect methodology. The next section discusses the empirical methodology in depth.

In order to estimate this effect of hospital mergers, I rely on two sources of local labor market data, each with relative advantages. I use both the Quarterly Census of Employment and Wages (QCEW) from the Bureau of Labor Statistics (BLS) and the American Community Survey (ACS) from the Census Bureau. For both datasets, I define local labor markets using Commuting Zones (CZs). CZs were developed by the Economic Research Service of the U.S. Department of Agriculture to more accurately delineate local labor markets than earlier methods.

The main advantage of the QCEW is that it is a near-census of worker earnings and employment at the county and industry level. It is based on ES-202 filings that every establishment is required to submit in order to calculate payroll taxes for unemployment insurance, and covers 98 percent of all workers. The industry level information allows for workers within the health care and insurance industries to be excluded from the sample. This is important for this analysis since I am interested in how these mergers affect labor markets through their role in the health care system rather than their role as employers in the health care labor market. For the part of the analysis using the QCEW, industries are broken out by their two-digit NAICS codes. Since it is aggregated to the county level and CZs are defined as small groups of counties, the QCEW also allows for easy aggregation to the CZ level. One drawback of the QCEW is that it only reports total payroll and total employment. Therefore, while it gives an accurate measure of average earnings of workers in a county and industry, there is no way to see the hours worked among these employees. So it will not be possible to tell whether findings of lower average annual wages are accompanied by (and potentially driven by) decreases in the number of hours worked or not.

In comparison, the ACS is an individual-level dataset with a rich set of demographic information as well as several employment outcomes including whether an individual was employed, their annual wages, and their usual hours worked per week. Furthermore, by the standard of surveys, the ACS is quite large, with over 3,000,000 observations in 2017. Responses are also legally obligated, reducing non-response issues relative to other surveys. However, to avoid issues of individual identifiability, information on low population geographies is limited. Therefore not every county is separately identifiable in the ACS. As an alternative, the ACS releases geographic information in the form of Public Use Microdata Areas (PUMAs), which are geographies that contain at least 100,000 people, are constructed from Census tracts, and are nested within states. Therefore, while not entirely commensurate with CZs, it is possible in most cases to map PUMAs into CZs, with the exceptions being particularly rural geographies. To address this limitation of the data I assign individuals in a certain PUMA to the CZ only if at least 80 percent of that PUMA's population is within that CZ according the Census.

Unfortunately, PUMAs are redrawn with every decennial Census, making the geographic identifiers from 2011 and earlier incompatible with identifiers from 2012 and later. Therefore, for the ACS portion of the analysis I restrict the years from 2005 (the first year the ACS was fielded) to 2011. While restrictive, this set of years roughly corresponds to the years examined by Cooper et al. (2019) in their study of merger's effects on prices. The outcomes I examine using the ACS are logged annual wages, logged hours worked per week, whether an individual was employed or not, and whether they worked over or under 35 hours per week (in order to see if workers are shifted from full to part-time work). Unfortunately, while I have information on usual hours worked per week in the past year and income over that time, the question asking the number of weeks worked in the past year was discontinued in 2007, leaving me unable to calculate average hourly wages. I restrict the sample to adults aged 22 to 64 since these are the workers most likely to receive insurance through their employer rather than as a dependent (those under 22) or from Medicare (those 65 and older).

Data on hospital mergers are provided by Cooper et al. (2019) and are a census of all ownership changes of hospitals registered with the American Hospital Association (AHA) from 2001 to 2014. The data includes the latitude and longitude of each hospital, the hospital that was targeted in the transaction and the hospitals that were members of the acquiring system. For more information on their data cleaning process see the Online Appendix of Cooper et al. (2019).

For this analysis I consider a CZ to have experienced a merger if a hospital with a latitude and longitude within that CZ is acquired by a hospital or system with a hospital whose latitude and longitude also place it in that CZ. It is important to note here that due to the frequency of hospital mergers in recent decades, many CZs experience multiple mergers over our time period examined here. In such cases, there is not a clearly defined the pre- and post- period to conduct a difference-in-differences analysis. Therefore, for the main analysis I restrict the treatment group to CZs that experienced one merger over the period analyzed in a manner similar to both Prager and Schmitt (2019) and Cooper et al. (2019). As a sensitivity test of this restriction, as in Cooper et al. (2019), I also estimate a model with a cumulative merger measure without dropping any CZs and find the main results remain.

Table 1. Summary Statistics

| | Treated | Control |
|---|---|---|
| Quarterly Census of Employment and Wages 2001-2014 | | |
| Average Wages | 30,981 | 28,412 |
| Employment | 96,011 | 33,666 |
| Observations | 1,946 | 5,502 |
| | Treated | Control |
| American Community Survey 2005-2011 | | |
| Annual Wages (Dollars) | 29,386 | 28,409 |
| Usual hours worked per week | 31.67 | 31.63 |
| Employed | 70.6% | 70.6% |
| Part-Time | 14.2% | 14.2% |
| Full-Time | 64.4% | 64.4% |
| Parent | 44.8% | 45.2% |
| Age | 42.24 | 42.42 |
| Female | 45.6% | 45.5% |
| Less than HS | 13.4% | 12.1% |
| College | 27.0% | 25.4% |
| Foreign Born | 14.2% | 10.5% |
| Noncitizen | 8.6% | 5.7% |
| Black | 8.8% | 11.4% |
| Hispanic | 13.0% | 8.4% |
| Observations | 1,177,867 | 1,144,141 |

Notes: Data for first panel is from the 2001-2014 Quarterly Census of Employment and Wages (QCEW). Sample includes all Commuting Zones (CZs) that experienced zero or one merger between 2001 and 2014. Data for second panel is from the 2005-2011 American Community Survey (ACS). Sample includes all Commuting Zones that experienced zero or one merger between 2005 and 2011.

Table 1 presents the summary statistics for the treated and untreated CZs in both datasets, with the top panel including the two measures available on the QCEW and the bottom panel including the outcomes and covariates used in the ACS analysis. The samples are broadly similar across a variety of characteristics including employment rates, prevalence of part-time work, demographic characteristics, and education levels. One key difference between the two samples is the average size of the labor market in each. In the QCEW sample, CZs that experienced a merger had 96,011 workers on average while those that did not had 33,666 on average. This issue is mitigated by the facts that (1) this analysis uses relative changes in employment as an outcome rather than absolute changes and (2) it uses the difference-in-difference study design will difference out any fixed characteristics of the labor markets. However this disparity in size could be an issue to the extent that labor market outcomes in smaller markets have different trends relative to larger markets over the period studied. To address this concern, along with the main difference-in-differences analysis, I test for differential trends between treated and untreated CZs leading up to merger events.

### Empirical Strategy

As mentioned above, for this study I adopt a difference-in-differences research design, considering individuals in labor markets that experienced a merger as "treated" and those in labor markets that were not as "controls". I will be estimating the following two-way fixed effects regression equation for individual $i$, CZ $j$, and year $t$ using ordinary least squares

$$\text{Outcome}_{ijt} = \alpha + \beta \text{ MERGE}_{jt} + \mathbf{\Gamma X_{it}} + \eta_j + \delta_t + \varepsilon_{ijt}. \tag{9}$$

Here Outcome is the labor market outcome of interest, either at the individual or aggregated level (in which case the subscript $i$ is dropped). $\text{MERGE}_{jt}$ is a binary indicator of whether the CZ experienced a hospital merger up to that point in time. For the individual-level regression, $\mathbf{X}$ is a vector of individual level covariates. Here, $\eta_j$ is a CZ fixed effect, controlling for any fixed differences between commuting zones' outcomes, and $\delta_t$ is a year fixed effect, which controls for any nationwide yearly shocks experienced across all CZs. As in all difference-in-differences designs, $\beta$ is identified by the assumption that absent a merger, changes in labor market outcomes in treated CZs would have been the same as changes in labor markets in control CZs had they not experienced a merger. While this assumption cannot be tested directly (as a CZ is not observable in two different states at the same time), I can test whether the outcomes being studied were evolving along the same trend prior to the merger event in treated and control CZs, which I do below.

The parameter of interest, $\beta$, is the average "treatment effect" of a merger on a local labor market. Due to the limitations listed above, this unfortunately cannot be interpreted as a parameter defined in the above incidence model. Given the data available, it is impossible to measure the impact that a given merger has on premium costs in that labor market or the impact that the subsequent cost change has on wages, $dW/dC$, or the other outcomes examined. Furthermore, I cannot differentiate between mergers involving larger or smaller shares of the hospital market, due to a lack of data on hospital size. Unfortunately this limits the economic interpretability of this point estimate and its applicability to particular future mergers under review. However, under the identifying assumption stated above, it can inform us retrospectively of the average overall effect that hospital mergers of any size had on each labor outcome in the study period.

Along with the basic difference-in-differences analysis, I also estimate an event-study model in order to test for both pre-trends and dynamic treatment effects. In an event-study an observation's indicator for being in the treatment group is interacted with a dummy for its year relative to its treatment. This, in effect, incorporates leads and lags into the regression. I estimate the following event-study model for individual $i$, CZ $j$, and year $t$ using ordinary least squares

$$\text{Outcome}_{ijt} = \alpha + \sum_{\substack{k=-r \\ k\neq 0}}^{s} \sigma^{k}\text{EVERMERGE}_{j}^{k} + \mathbf{\Gamma X_{it}} + \eta_j + \delta_t + \varepsilon_{ijt}. \tag{10}$$

Here $k$ is year relative to the merger and $k = 0$ is the year of merger, where $r$ is the number of leads and $s$ is the number of lags. Here the dummy for year of merger is excluded so that the point estimates of $\sigma^{-r}$ to $\sigma^{s}$ are relative to it. EVERMERGE$^{k}$ is an indicator for being in the treatment group and in year $k$.

## Results

Table 2 presents the difference-in-difference results for both log wage and log employment using the QCEW from 2001 to 2014. The results suggest that a hospital merger reduces the average wage in an exposed industry-commuting zone by 1.5 percent and this result is significant at the .01 level. The result for employment suggests mergers had no statistically significant impact on employment in exposed industry-commuting zones.

Table 2. Effect of Mergers on Labor Market Outcomes

| | Est. | P-Value | N |
|---|---|---|---|
| Log(Average Wage) | -0.015 | 0.000 | 94102 |
| Log(Employment) | 0.001 | 0.954 | 94102 |

Note: Data from the 2001-2014 Quarterly Census of Employment and Wages (QCEW). Observation is the industry-commuting zone level.

Table 3 presents the difference-in-differences results using the 2005-2011 ACS. In the ACS analysis I control for a host of individual-level demographic characteristics including age, sex, race/ethnicity, education, citizenship status, and parental status. The results suggest that mergers resulted in a 0.8 percent reduction in wage and salary income on average, with a p-value of 0.019. The results for hours worked, employment, and part-time status were all statistically insignificant. In the appendix, I also estimate a cumulative merger effect without dropping CZs with multiple mergers over the time period, finding similar results.

Table 3. Effect of Mergers on Labor Market Outcomes for ACS

| | Est. | P-Value | N |
|---|---|---|---|
| Log(Wage and Salary Income) | -0.008 | 0.019 | 1696943 |
| Log(Hours) | 0.000 | 0.823 | 1812034 |
| Employed | -0.002 | 0.158 | 2322008 |
| Part Time | 0.001 | 0.348 | 2322008 |
| Full Time | -0.001 | 0.337 | 2322008 |

Note: Data from the 2005-2011 American Community Survey (ACS). Sample is limited to those ages 22-64. Full-time work defined as working at least 35 hours per week. Covariates include year and Commuting Zone fixed effects, marital status, sex, age, education, 2-digit industry number, citizenship status, race/ethnicity, and parental status. Regression are weighted by ACS weights.

In conjunction with this basic difference-in-differences design, I also estimate event-study specifications using the same data. Event-study designs have two primary benefits. First, they allow for testing whether treated and control labor markets had differential trends in the run-up to mergers that could bias the difference-in-difference result. Secondly, they allow testing whether there were dynamic effects as a result of treatment, including whether any effects fade or increase over time.

The event-study results for the 2001-2014 QCEW are presented in Figure 1. Here the coefficients for year relative to merger interacted with treatment status are plotted, along with their 95 percent confidence intervals. The year the merger occurred is the excluded year which the point estimates are relative to. I find no evidence of pre-trends in the years leading up to mergers in the QCEW data for either wages or for employment. Again, I find that mergers reduced wages and here it appears this effect grows over time. As in the difference-in-differences analysis, no effect on employment is detectable.

Figure 1. QCEW Event Study Estimates

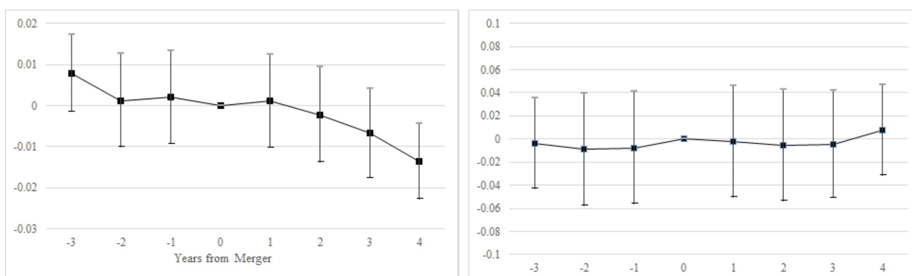


(a) Log(Average Wage) (b) Log(Employment)

Note:Data from the 2001-2014 Quarterly Census of Employment and Wages (QCEW). Point estimates and 95% Confidence Intervals are depicted for the three years prior to and four years subsequent to the merger. Observation is the industry-commuting zone level.

The event-study results for the 2005-2011 ACS are presented in Figure 2. Here again I find no evidence of pre-trends for wages, hours, or the work-status measures. Like the QCEW result for wages, the ACS results also suggest that the impacts of mergers on wages grew over time. While the difference-in-differences results found no overall average effect of mergers on non-wage outcomes in post-period years, the event-study results suggest some evidence that these measures were affected by mergers. In panel (b) of Figure 2, while coefficients on hours worked for one and two years post-merger are statistically insignificant, the coefficient for three years post-merger is marginally significant and negative, suggesting that hours were affected by mergers gradually. Similarly the event study results suggest mergers decreased full-time work and increased part-time work among adults in the sample, with the estimated effect on both measures statistically different from the base year after three years. These impacts on part-time and full-time work are consistent with the findings of Montgomery and Cosgrove (1993), Buchmueller (1999), and Baicker and Chandra (2006). It may be easier for firms to adjust the composition of their workforce between full- and part-time workers than deal with the costs through other potential channels.

Figure 2. ACS Event Study Estimates

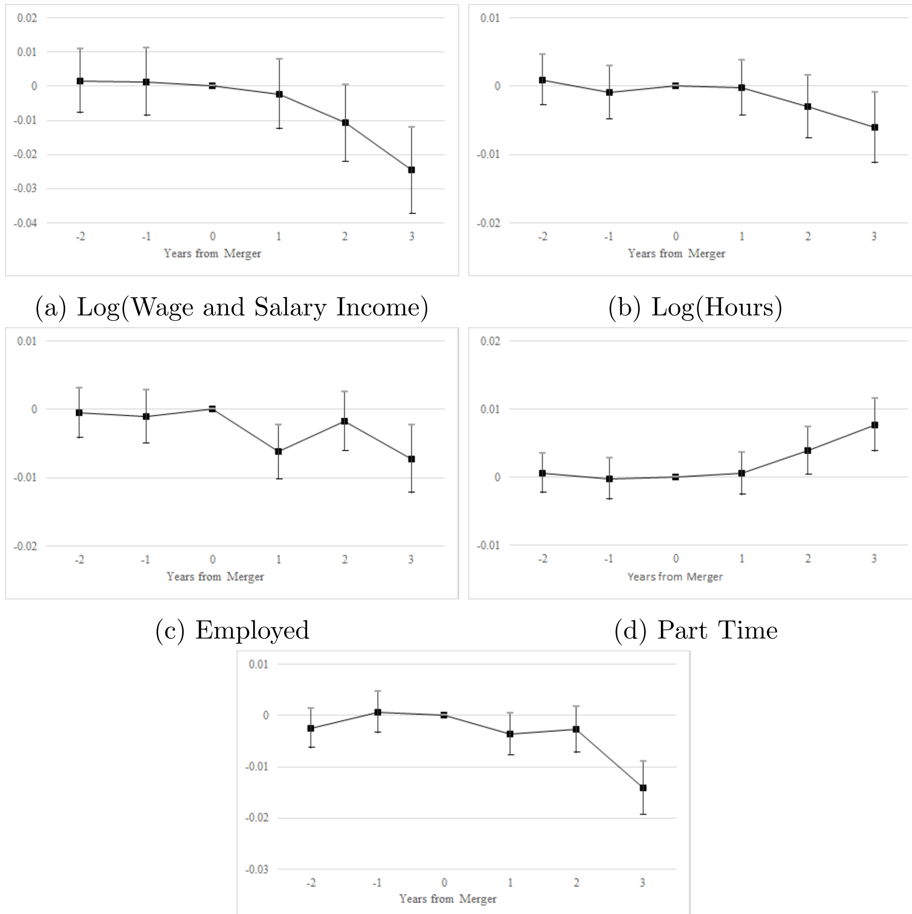


(a) Log(Wage and Salary Income)

(b) Log(Hours)

(c) Employed

(d) Part Time

(e) Full Time

Notes: Data from the 2005-2011 American Community Survey (ACS). Point estimates and 95% Confidence Intervals are depicted for the two years prior to and three years subsequent to the merger. Sample is limited to those ages 22-64. Full-time work defined as working at least 35 hours per week. Covariates include year and commuting zone fixed effects, marital status, sex, age, education, 2-digit industry number, citizenship status, race/ethnicity, and parental status. Regressions are weighted by ACS person weights.

The delay in the impacts of mergers on labor markets for both the QCEW and ACS is not surprising in light of the findings of Cooper et al. (2019). The authors found that the effect of mergers on prices charged to insurers appear to grow over time, with results greatest 2+ years after mergers occurred. Therefore, the growth in effects over time is in line with the idea that prices charged by hospitals are driving these labor market changes.

These effects on hours and full-time status present in the event-study analysis are important for interpreting the impact in both the QCEW and ACS analysis on annual wages. If hours worked, full-time status, and part-time status were unaffected by mergers, it would be likely that the decline in annual income was being driven by a decline in hourly wages, in line with a standard incidence story of the Summers (1989) model presented above. However with hours and work status changing as well, it could be either changes in hourly wages or the work arrangements themselves that are driving the income changes.

There may be some concern that the Great Recession is influencing the results presented here, say if it affected treated labor markets differently than control labor markets. In the appendix, I reestimate the QCEW event-study results on the years 2001 to 2007 and find similar results to the main analysis. Unfortunately the sensitivity of the ACS estimates cannot be similarly tested for since the Great Recession occurs in the middle of that sample period.

## Conclusion

This paper examines the impact of hospital mergers on local labor markets using both micro-level data from the American Community Survey and aggregate payroll data from the Quarterly Census of Employment and Wages. The results suggest that consolidation in the hospital industry affects local labor markets in several ways. I find strong evidence that mergers have negative effects on annual wages. I also find evidence that mergers resulted in reductions in full-time employment and increases in part-time employment in commuting zones where they occurred. Due to data limitations I am not able to determine whether these reductions are from changes in hourly wages or whether they are solely being driven by the changes in hours worked. Thus whether individual workers bear the incidence of changes in the cost of their insurance through hourly wage reductions cannot be determined. It may be the case they simply substitute away from now-costlier full-time workers who are eligible for insurance coverage towards part-time workers usually ineligible for such firm coverage. The findings of Cooper et al. (2019) suggest that the impact that hospital mergers have on prices increases over time, and here I find that the impact of the mergers on labor market outcomes also appears to grow over time.

These findings have important implications for policy-makers and regulators. Recent work on employer monopsony power such as Azar et al. (2017) and Benmelech et al. (2018) has suggested that the welfare implications of mergers extend beyond consumers to the employees in the industries being affected and that these impacts should be accounted for in merger reviews (see Prager and Schmitt (2019) for an example in the hospital industry). The findings in this paper suggest that, due to the unique relationship between health care costs and worker compensation in the United States, the implications of hospital mergers extend to workers outside of the directly relevant industry as well.

While beyond the direct scope of this paper, the findings in this study may also be relevant for understanding broader trends in slow wage growth experienced by the United States in recent years. As noted by Cooper et al. (2019), there were over 700 hospital mergers from 2001 to 2011 and the average Herfindahl-Hirschman Index (HHI) in the hospital industry increased by 19 percent over that period. If mergers do suppress wages through their impact on premiums, this rapid consolidation could be one factor contributing to slow wage growth seen in recent decades.

In light of this paper's findings, I believe there are several areas where future research would be best focused. This paper was only able to estimate the short run effects of these mergers and was not able to estimate an overall effect this concentration has had on the US economy. For example, if hospital concentration reduces local demand due to reductions in disposable income, there may be macroeconomic effects that cannot be explored in this partial equilibrium framework. Perhaps even more pressing is the need to incorporate the premium costs that employers face when studying labor markets. As repeatedly noted here, data limitations are a serious issue for understanding how employers choose their compensation packages. Given that the surveys frequently used to study labor markets including the ACS, CPS and SIPP do not have this information, studies examining labor costs or workers' compensation using them are missing a large piece of the picture. Better data can potentially shed light on what is currently a black box. With private health insurance expenditures surpassing 6 percent of GDP in 2017 (Centers for Medicare & Medicaid Services 2017), the matter could not be more pressing.

## References


Anand, P. (2017). Health insurance costs and employee compensation: Evidence from the national compensation survey. *Health Economics 26*(12), 1601–1616.

Anstreicher, G. (2024). Spatial influences in upward mobility. *Journal of Political Economy 132*(1).

Anstreicher, G. and J. Venator (2022). To grandmother's house we go: childcare time transfers and female labor mobility. https://www.bc.edu/content/dam/bc1/schools/mcas/economics/pdf/working-papers/1000/wp1051.pdf. Accessed: 2023-03-09.

Azar, J., I. E. Marinescu, and M. Steinbaum (2017). Labor market concentration. *NBER Working Paper 24147*.

Baicker, K. and A. Chandra (2006). The labor market effects of rising health insurance premiums. *Journal of Labor Economics 24*(3), 609–634.

Benmelech, E., N. Bergman, and H. Kim (2018). Strong employers and weak employees: How does employer concentration affect wages? *NBER Working Paper 24307*.

Bhattacharya, J. and M. K. Bundorf (2009). The incidence of the healthcare costs of obesity. *Journal of Health Economics 28*(3), 649–658.

Bishop, K. C. (2007). A dynamic model of location choice and hedonic value. https://are.berkeley.edu/ ligon/ARESeminar/Papers/bishop07.pdf. Accessed: 2023-03-09.

Buchmueller, T. (1999). Fringe benefits and the demand for part-time workers. *Applied Economics 31*(5), 551–563.

Buchmueller, T. and M. Lettau (1997). Estimating the wage-health insurance tradeoff: More data problems?

Bureau of Labor Statistics (2019). *Employer Costs for Employee Compensation-March 2019*. Bureau of Labor Statistics.

Cadena, B. C. and B. K. Kovak (2016). Immigrants equilibrate local labor markets: evidence from the great recession. *American Economic Journal: Applied Economics 8*(1).

Centers for Medicare & Medicaid Services (2017). National health expenditure fact sheet. https://www.cms.gov/research-statistics-data-and-systems/statistics-trends-and-reports/nationalhealthexpenddata/nhe-fact-sheet.html. Accessed: 2019-09-01.

Chan, S., H.-L. Liao, and K. O'Regan (2025). Moving back: spatial and demographic differences in boomeranging to parents. *Journal of Urban Affairs*.

Choi, H., R. F. Schoeni, K. M. Langa, and M. M. Heisler (2014). Spouse and child availability for newly disabled older adults: socioeconomic differences and potential role of residential proximity. *Journals of Gerontology, Series B: Psychological Sciences and Social Sciences 70*(3), 462–469.

Choi, H., R. F. Schoeni, E. E. Wiemers, V. J. Hotz, and J. A. Seltzer (2020). Spatial distance between parents and adult children in the united states. *Journal of Marriage and Family 82*(2), 822–840.

Clemens, J. and D. M. Cutler (2014). Who pays for public employee health costs? *Journal of Health Economics 38*, 65–76.

Coate, P. (2013). Parental influence on labor market outcomes and location decisions of young workers. https://citeseerx.ist.psu.edu/document?repid=rep1type=pdfdoi=ce6544f86819e70919a3a7a8df380b3c86582a Accessed: 2023-03-09.

Compton, J. and R. A. Pollak (2015). Proximity and co-residence of adult children and their parents in the united states: descriptions and correlates. *Annals of Economics and Statistics* (117/118).

Cooper, Z., C. C. Stuart, M. Gaynor, and J. V. Reenan (2019). The price ain't right? hospital prices and health. *The Quarterly Journal of Economics 134*(1), 51–107.

Currie, J. and B. Madrian (1999). Health, health insurance and the labor market. *Handbook of Labor Economics 3*, 3309–3416.

Cutler, D. M. and B. Madrian (1998). Labor market response to rising health insurance costs: Evidence on hours worked. *RAND Journal of Economics 29*(3), 509–530.

Diamond, R. (2016). The determinants and welfare implications of us workers' divergin location choices by skill: 1980-2000. *American Economic Review 106*(3), 479–524.

Esipova, N., A. Pugliese, and J. Ray (2013). The demographics of global internal migration. *Migration Policy Practice 3*(2), 3–5.

Federal Trade Commission (2018). Hearings on competition and consumer protection in the 21st century. https://www.ftc.gov/news-events/events-calendar/2018/10/ftc-hearing-3-competition-consumer-protection-21st-century. Accessed: 2019-07-01.

Gaynor, M. and R. Town (2011). Competition in health care markets. *Handbook of Health Economics 2*, 499–637.

Gemici, A. (2013). Family migration and labor market outcomes. http://sites.googlegroups.com/site/ahugemici/gemici$_{p}aper3.pdf. Accessed : 2021-23-11.$

Gowrisankaran, G., A. Nevo, and R. Town (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review 105*(1), 172–203.

Gruber, J. (1994). The incidence of mandated maternity benefits. *American Economic Review 84*(3), 622–641.

Gruber, J. and M. Hanratty (1995). The labor market effects of introducing national health insurance: Evidence from canada. *Journal of Business Economic Statistics 13*(2), 163–173.

Gruber, J. and A. B. Krueger (1991). The incidence of mandated employer-provided insurance: Lessons from workers' compensation insurance. In *Tax Policy and the Economy*, Volume 5.

Gruber, J. and R. McKnight (2003). Why did employee health insurance contributions rise? *Journal of Health Economics 22*(6), 1085–1104.

Hsieh, C.-T. and E. Moretti (2019). Housing constraints and spatial misallocation. *American Economic Journal: Macroecnomics 11*(2), 1–39.

Ishimaru, S. (2024). Geographic mobility of youth and spatial gaps in local college and local labor market opportunities. *Journal of Labor Economics Forthcoming.*

Jia, N., R. Molloy, C. Smith, and A. Wozniak (2023). The economics of internal migration: advances and policy questions. *Journal of Economic Literature 61*(1), 144–180.

Kaiser Family Foundation (2018). State health facts: Health insurance coverage of adults 19-64. https://www.kff.org/other/state-indicator/adults-19-64. Accessed: 2019-07-01.

Kaiser Family Foundation (2019). Employer health benefits: 2018 annual survey.

Kennan, J. and J. R. Walker (2011). The effect of expected income on individual migration decisions. *Econometrica 79*(1), 211–251.

Kolstad, J. T. and A. E. Kowalski (2016). Mandate-based health reform and the labor market: Evidence from the massachusetts reform. *Journal of Health Economics 47*, 81–106.

Lessem, R. (2018). Mexico-u.s. immigration: effects of wages and border enforcement. *Review of Economic Studies 85*, 2353–2388.

Long, L. (1992). Residential mobility differences among developed countries. *International Regional Science Review 14*(2), 133–147.

Lubotsky, D. and C. A. Olson (2015). Premium copayments and the trade-off between wages and employer-provided health insurance. *Journal of Health Economics 44*, 63–79.

McMurry, J. (2021). Child care policy and informal care. https://joelmcmurry.github.io/files/research/draft$_CCPIC.pdf.Accessed: 2023-08-12$.

Molloy, R., C. L. Smith, and A. Wozniak (2011). Internal migration in the united states. *Journal of Economic Perspectives 25*(3).

Molloy, R., C. L. Smith, and A. Wozniak (2014). Declining migration within the us: the role of the labor market. *NBER Working Paper 20065*.

Mommaerts, C. (2025). Long-term care insurance and the family. *Journal of Political Economy 133*(1).

Montgomery, M. and J. Cosgrove (1993). The effect of employee benefits on the demand for part-time workers. *ILR Review 47*(1), 87–98.

Moriya, A. S., W. B. Vogt, and M. Gaynor (2010). Hospital prices and market stucture in the hospital and insurance industries. *Health Economics, Policy, and Law 5*(4), 459–479.

Prager, E. and M. Schmitt (2019). Employer consolidation and wages: Evidence from hospitals. *Washington Center for Equitable Growth Working Paper*.

Qin, P. and M. Chernew (2014). Compensating wage differentials and the impact of health insurance in the public sector on wages and hours. *Journal of Health Economics 38*, 77–87.

Reyes, A. M. and Y. Shang (2024). Geographic relocation in response to parents' health shocks: who moves and how close? *Journal of Marriage and Family 86*(1), 49–71.

Skira, M. M. (2015). Dynamic wage and employment effects of elder parent care. *International Economic Review 56*(1).

Sommers, B. D. (2005). Who really pays for health insurance? the incidence of employer-provided health insurance with sticky nominal wages. *International Journal of Health Care Finance and Economics 5*(1), 89–118.

Spring, A., E. Acker, K. Crowder, and S. J. South (2017). Influence of proximity to kin on residential mobility and destination choice: examining local movers in metropolitan areas. *Demography 54*, 1277–1304.

Summers, L. H. (1989). Some simple economics of mandated benefits. *American Economic Review 79*(2), 177–183.

U.S. Department of Commerce (2023). Geographic inequality on the rise in the us. https://www.commerce.gov/news/blog/2023/06/geographic-inequality-rise-us. Accessed: 2025-03-05.

Venator, J. (2024). Dual earner migration, earnings, and unemployment insurance. https://joannavenator.com/wp-content/uploads/2024/09/venator$_{u}ifortrailingspouses_{8}2024.pdf.Accessed:$ $2025-01-02$.

Vistnes, J. and T. Selden (2011). Premium growth and its effect on employer-sponsored insurance. *International Journal of Health Care Finance and Economics 11*(1), 55–81.

Voena, A. (2015). Yours, mine and ours: do divorce laws affect the intertemporal behavior of married couples? *American Economic Review 105*(8), 2295–2332.

## Appendix 1

Figure A.1.1 Counterfactual Change from Baseline

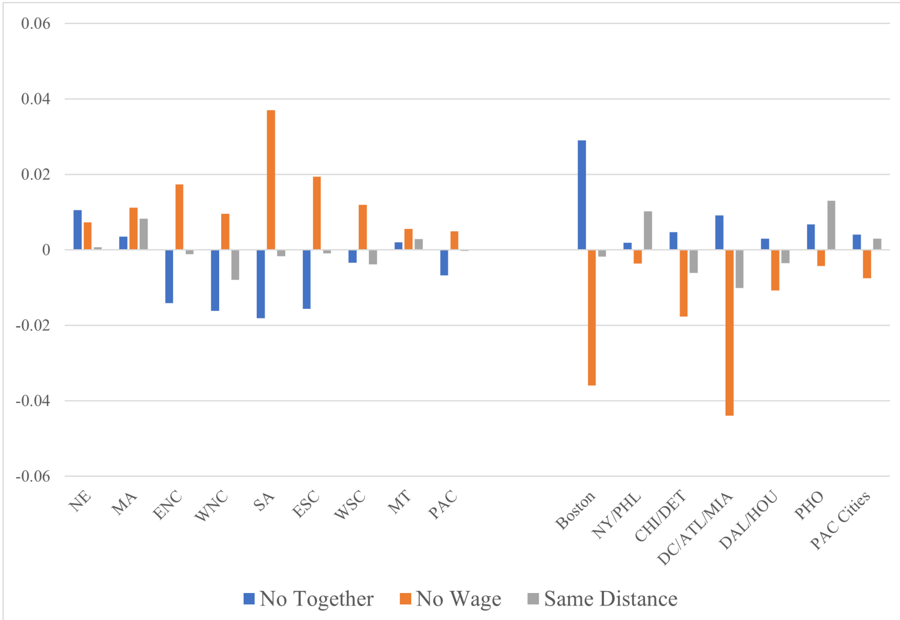


Figure A.1.2 Effect of Eliminating Wage and Distance Variation in Baseline and Main Counterfactual

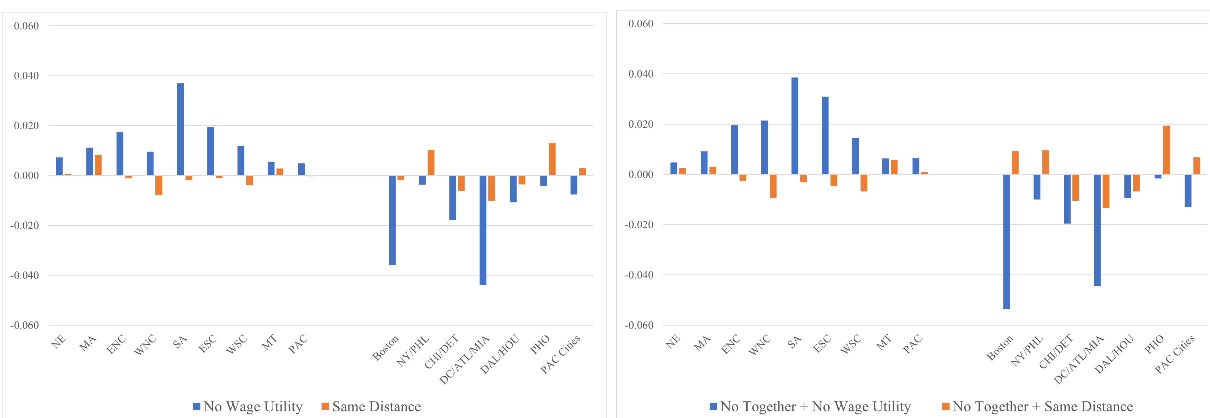


(a) Change from Baseline

(b) Change from Counterfactual

Figure A.1.3 Effect of Eliminating Wage Variation in Baseline and Main Counterfactual


Figure A.1.4 Effect of Eliminating Distance Variation in Baseline and Main Counterfactual

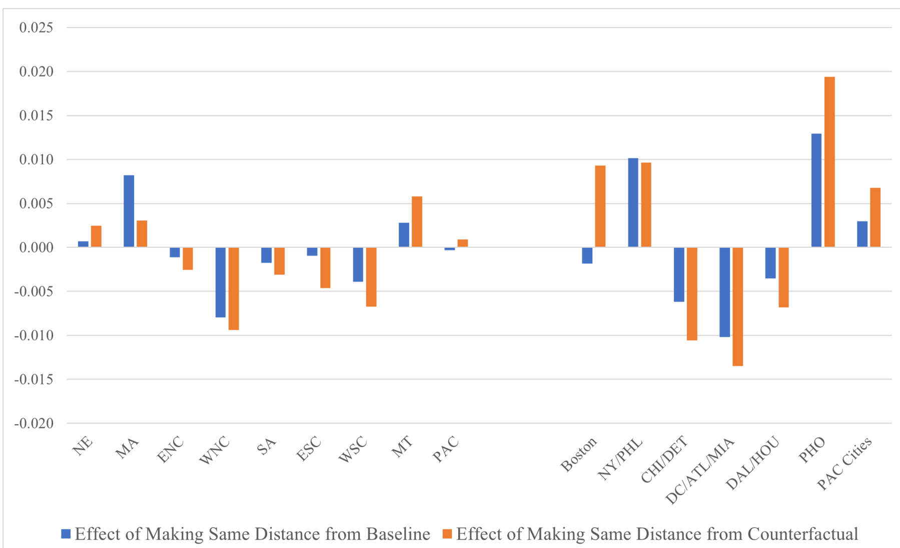

The moving cost is parameterized,

$$
\begin{aligned}
MC(l_{c,t}, l_{p,t}, l_{c,t-1}, l_{p,t-1}, tog_t, tog_{t-1}, \Omega) = {} & \lambda_{dist} \cdot dist(l_{c,t}, l_{c,t-1}) + \lambda_{dist} \cdot dist(l_{p,t}, l_{p,t-1}) \\
& + 1(l_{p,t} \neq l_{p,t-1})\big(\lambda_{parage} \cdot \omega_{parage} \\
& + \lambda_{pop} \cdot pop(l_{p,t}) + \lambda_{sd} samediv(l_{p,t}, l_{p,t-1})\big) \\
& + 1(l_{c,t} \neq l_{c,t-1})\big(\lambda_{kidage} \cdot \omega_{kidage} \\
& + \lambda_{pop} \cdot pop(l_{c,t}) + \lambda_{col} \cdot \omega_{gk} + \lambda_{mar} \cdot \omega_{mar} \\
& + \lambda_{coll} \cdot \omega_{coll} + \lambda_{sd} samediv(l_{c,t}, l_{c,t-1})\big) \\
& + \lambda_{togmove} \cdot 1(l_{p,t} = l_{p,t-1}) \cdot 1(l_{c,t} = l_{c,t-1}) \\
& \cdot 1(tog_t \neq tog_{t-1})
\end{aligned}
$$

## Appendix 2

### Cumulative Merger Effect

Here I present an alternative specification to the difference-in-differences analysis presented in the main text. Rather than restricting the sample only to commuting zones that experienced either no or one merger to isolate the impact of a single merger, I regress the outcome on a cumulative merger measure. The dependent variable of interest here is a count variable that increases by one for every year during the sample period that a merger occurs in a commuting zone. In this way no commuting zones are dropped from the sample. An issue with this method is that increases in the measure associated later mergers would capture any dynamic effects of earlier mergers in that same commuting zone, potentially biasing the estimates. This is particularly an issue for this analysis since I find strong evidence of dynamic post-merger effects.

Table A1 presents the results for the QCEW. Here the effect on wages remains negative and statistically significant and the effect on employment remains statistically insignificant, just as in the the main difference-in-differences analysis. Table A2 presents the results for the ACS. Here the impact of a merger on wage and salary income remains negative and statistically significant. Here the point estimates on employment and full-time employment become statistically significant. While this differs from the basic difference-in-differences analysis, these were findings that were present in the event-study model presented later in the paper.

Table A1 presents the results for the Quarterly Census of Employment and Wages from 2001 to 2014.

Table A1. Cumulative Effect of Mergers on Labor Market Outcomes

| | Est. | P-Value | N |
|---|---|---|---|
| Log(Average Wage) | -0.004 | 0.000 | 119302 |
| Log(Employment) | -0.003 | 0.483 | 119302 |

Note: Data from the 2001-2014 Quarterly Census of Employment and Wages (QCEW). Observation is the industry-commuting zone level.

Table A2. Cummulative Effect of Mergers on Labor Market Outcomes for ACS

| | Est. | P-Value | N |
|---|---|---|---|
| Log(Wage and Salary Income) | -0.012 | 0.000 | 2261765 |
| Log(Hours) | 0.000 | 0.659 | 2413717 |
| Employed | -0.002 | 0.004 | 3101002 |
| Part Time | 0.000 | 0.546 | 3101002 |
| Full Time | -0.003 | 0.000 | 3101002 |

Note: Data from the 2005-2011 American Community Survey (ACS). Sample is limited to those ages 22-64. Full-time work defined as working at least 35 hours per week. Covariates include year and Commuting Zone fixed effects, marital status, sex, age, education, 2-digit industry number, citizenship status, race/ethnicity, and parental status. Regression are weighted by ACS weights.

**Excluding the Great Recession**

Here I present QCEW event-study results, restricting the data to 2001 to 2007. It may be possible that the results above are driven by a relationship between which commuting zones experienced a merger and the impact the Great Recession had at a local level. Again we find no evidence of pre-trends prior to the merger events occurring. We also find no evidence of an impact of the mergers on employment. There is evidence here again that mergers reduced average annual wages. The point estimate of -0.012 percent after 3 years is similar in magnitude to the earlier results and statistically significant at the 0.1 level.

Table A3. QCEW Event Study Estimates 2001-2007

| | Est. | P-Value |
|---|---|---|
| Log(Average Annual Wage) | | |
| Two years prior to merger | 0.000 | 0.932 |
| One year prior to merger | -0.001 | 0.858 |
| One year after merger | -0.001 | 0.837 |
| Two years after merger | -0.004 | 0.598 |
| 3+ years after merger | -0.012 | 0.094 |
| Log(Employment) | | |
| Two years prior to merger | 0.001 | 0.962 |
| One year prior to merger | -0.003 | 0.925 |
| One year after merger | 0.003 | 0.912 |
| Two years after merger | 0.003 | 0.922 |
| 3+ years after merger | 0.006 | 0.868 |
| Observations | 56623 | |

Note: Data from the 2001-2007 Quarterly Census of Employment and Wages (QCEW). Observation is the industry-commuting zone level.