WHY DO SO FEW WOMEN WORK IN NEW YORK (AND SO MANY IN MINNEAPOLIS)? LABOR SUPPLY OF MARRIED WOMEN ACROSS U.S. CITIES

DAN BLACK, NATALIA KOLESNIKOVA, AND LOWELL J. TAYLOR

Abstract. This paper documents a little-noticed features of U.S. labor markets—that there is very large variation in the labor market participation rates and annual work hours of married women across cities. We focus on cross-city differences in commuting times as a potential explanation for this variation in women's labor supply. Our starting point is the analysis of labor supply in a model in which commute times introduce non-convexities into the budget set. Empirical evidence appears consistent with the model's predictions: In the cross section, labor force participation rates of married women are negatively correlated with the metropolitan area commuting time. Our analysis also indicates that metropolitan areas which experienced relatively large increases in average commuting time from 1980 through 2000 had slower growth in the labor force participation of married women.

JEL: J21, J22, R23, R41.

Keywords: female labor supply, local labor markets, commuting time, non-convex budget sets

Introduction

Women's labor supply has, for good reason, been the object of extensive empirical study. After all, the dramatic rise in female labor force participation that occurred over the past 60 years in the U.S. (and in many other countries) has been the most visible and important shift in the labor market. Also, women's labor supply is often the margin of adjustment in households' responses to policy shifts, e.g., changes in the taxation of household income or welfare entitlement programs, and thus holds the key to proper policy evaluation.

Although many empirical studies of female labor supply have been conducted, it appears that an interesting, potentially important feature of the U.S. markets have gone largely unnoticed: There is wide variation in female labor supply across metropolitan areas in the United States. Consider, for example, one large group of women: married non-Hispanic white women aged 25 to 55 with a high school level education. In Minneapolis 79 percent of such women are employed, while in New York the proportion is only 52 percent. This cross-city variation in female labor supply within the U.S. is as large as the well-known and

Date: October 2008. Black is affiliated with University of Chicago and NORC, Kolesnikova with the Federal Reserve Bank of St. Louis, and Taylor with Carnegie Mellon University. The views expressed are those of the authors and do not necessarily reflect the official position of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors.

widely-studied variation across OECD countries in female employment rates.[1] In an effort to make sense of international comparisons, analysts typically focus on policy differences across countries (in paid parental leave, marginal taxes, employment protection, welfare benefits, etc.[2]), but of course such policy differences are much smaller across locations in the U.S. than OECD countries; the cross-city variation in female labor supply in the U.S. is apparently generated by characteristics of the local markets themselves.

Furthermore, while the labor supply of women has over the past 60 years increased substantially in all cities in the U.S., there has been big differences in these cities in the timing and magnitude of the increase. Figure 1 illustrates, for 1940 through 2000, the well-known large increases in the labor supply of married non-Hispanic white women generally, and shows also how different the paths are for two particular urban locations, New York and Minneapolis. In 1940 labor supply among women was lower in Minneapolis than New York, but the subsequent growth in female labor supply was much more rapid in Minneapolis than in New York, leading to the large disparities observed in 2000. These results are especially interesting in light of the on-going discussion about the possibility that the U.S. labor market has now achieved a "natural rate" of female labor force participation.[3]

The goals of this paper are to carefully document the cross-city variation in married women's labor supply across U.S. labor markets, to explore potential economic explanations for observed cross-city variation in married women's labor supply, and to examine the implications for the study of female labor supply generally. We believe that many factors are at play in producing the large observed local variation in female labor supply across the U.S., but, we argue, one explanation stands out: married women are very sensitive to commuting times when making labor force participation decisions.

[1]For the sake of rough comparison, one might look at employment rates among women with "upper secondary education" from selected OECD countries: United Kingdom, 80 percent; Sweden, 78 percent; Netherlands, 74 percent; France, 71 percent; Canada, 69 percent; U.S., 66 percent; Italy, 64 percent; Japan, 59 percent; and Germany, 52 percent. (These statistics, for women aged 25 to 64, are from OECD, 2007.)

[2]Ruhm (1998), for example, focuses on the impact paid parental leave policy on female labor supply in nine European countries. More generally, a large literature compares labor policy differences in U.S., Canadian, and European to explain differences in labor market outcomes. Nickell (1997), Card et al. (1999), Freeman and Schettkat (2001), Alesina et al. (2005) are just a few examples.

[3]Many authors have documented the fact that female labor force participation slowed considerably in the mid-1990's, and levelled off in the 2000s, e.g., Blau et al. (2002), Blau and Kahn (2000), and Juhn and Potter (2006). Goldin (2006) points to the importance of considering different age groups separately (rather than simply looking at aggregate measures of female labor supply), noting that for some groups of women "a plateau ... was reached a decade and a half ago." Looking at variation across local labor markets brings an additional dimension of complexity. Should we expect that cities with low rates of labor force participation will continue to experience an increase in female labor supply until they reach the national average, or are there reasons to expect that some markets have a lower "equilibrium" participation rate than others?

In building our argument about the importance of commuting cost, we start with the theory of labor supply when there is a fixed cost of participation (i.e., commuting time). The introduction of a fixed cost of participation introduces non-convexities into the budget constraint. This complication is easily handled in a one-period case for a one-person household: Assuming leisure and consumption are normal, and assuming also that initially the individual is at an "interior solution," an increase in the fixed cost reduces both leisure and labor supplied, up to a threshold at which the individual moves to a "corner solution" of supplying zero labor. Matters are far more interesting in a multiple-period model, or in a model in which there are two individuals in the household. When a person can work more than one period, an increase in commute time can *reduce* the number of periods worked (i.e., lead to more periods of non-employment), while *increasing* hours worked during periods when the individual does work. Similarly, in a two-person household, increases in the commute time can induce one partner to move out of the labor force (perhaps the wife) while inducing the other partner (perhaps the husband) to work longer hours.

As mentioned above, there are many studies of women's labor supply; Blundell et al. (2007) and Blundell and MaCurdy (1999) provide valuable discussions of key issues in this literature, and Killingsworth and Heckman (1986) overview earlier results. Most studies use national data, with results are aggregated at the national level, and no attention is given to the possibility of meaningful local variation. A small body of work in economic geography does provide some evidence about cross-location variation in labor supply (e.g., Odland and Ellis (1998) and Ward and Dale (1992)), but this work does not seek to provide an explanation for the observed variation. In particular, we know of no work that posits the importance of fixed commuting costs for explaining local labor supply, and then evaluates predictions empirically. We carry out such an analysis in five additional sections:

Section 1 provides the basic facts about the city-specific employment rates of non-Hispanic white married women in 50 large U.S. metropolitan areas from 1940 through 2000 using Public Use Samples of the U.S. Census.[4] We document significant variation across cities in current levels of women's employment and average annual work hours, and also substantial variation across cities in the magnitude and timing of the increase in female labor supplied over the past 60 years.

[4]For comparison, some evidence is also presented for non-Hispanic white married men. Issues concerning local variation in the labor supply of minority men and women are also very interesting, but are not the focus here.

Section 2 is a discussion of economic forces that might serve as potential explanations for the observed cross-cities variation in women's labor supply. We argue that the variation in observed employment rates stems are unlikely to be due primarily to differences across cities in labor *demand*.

Section 3 contains the primary economic contribution of this study. We develop an argument about the effect of cross-city differences in commuting times (owing, for example, to differences in congestion across cities) for labor force participation. Our model allows us to examine the effect of commuting time on individuals' and households' labor supply.

Section 4 has empirical evidence concerning the predictions of the model. The cross-sectional evidence indicates that in cities with longer average commuting times, female labor force participation rates are lower. Women with young children are particularly sensitive to longer commute. What is more, an examination of cross-city changes in female labor force participation from 1980 to 2000 indicates that increases in commuting times were associated with relatively slower growth in the participation rates of married women. These results are all consistent with the theory presented in Section 3.

Finally, Section 5 provides a conclusion and discussion of directions for future research.

1. Differences in Labor Supply Across Labor Markets

This study focuses on the labor supply of married non-Hispanic white women who live in the 50 largest Metropolitan Statistical Areas (MSAs) in the United States. The focus on married women is motivated by the fact that these women are most responsible for the large changes in female labor supply that have been observed over the last several decades (see, e.g., Juhn and Potter (2006)). Women in racial and ethnic minorities are excluded to avoid the complications of dealing with additional dimensions to the analysis, and because sample sizes for these groups are much smaller. The sample is restricted to individuals aged 25 to 55.

For much of our analysis we rely on the 2000 Census 5 percent Public Use Micro Sample (PUMS).[5] We also exploit comparable data from 1940 through 1990 (except 1960, owing to the lack of location identifiers for that year).

The PUMS data provide information on employment status. The three main categories are employed, unemployed and not in the labor force. For the most part, the analysis below looks at the "employment rate" as the measure of labor force participation in a local labor

[5]Individuals with imputed values are excluded from the analysis. The resulting sample is quite large, 423,300 observations.

market, thereby including women who are reported as unemployed with those who are out of labor force.[6] Included in the sample are women in the armed forces; they constitute 0.1 percent of the sample in 2000.

We begin by estimating participation rates for each of the 50 MSAs, for selected years, 1940 through 2000. Results are given for women in the largest educational categories—women with a high school diploma. The results are sorted by participation rates in 2000, from lowest to highest. The variation evident in the statistics is striking. In 2000 the participation rates of high school educated women vary from just 52 percent in New York City to 79 percent in Minneapolis. Similarly wide variation is evident in other years as well; for instance, in 1970 MSA-specific participation rates varied from 30 percent to 59 percent. There are also substantial differences in the growth of married women's labor supply over time. For example, from 1940 through 2000, the participation rate increased by 36 percentage points in New York, but by 64 percentage points in Minneapolis (as we have seen in Figure 1).

We examined the extent to which cross-city differences in the age distribution of women account for the observed labor force participation rates, and found that this matters very little.[7] We also constructed tables similar to Table 1 separately for women without children, women with older children, and women with young children. The summary is reported in Table 2. In each case, we found big differences across cities in labor force participation. For example, in 2000, among high school graduate women with children younger than 5, participation rates varied from 29 percent (New York) to 68 percent (Minneapolis).

The data indicate comparably wide variation in an alternative measure of labor supply—annual work hours, which can be computed as a product of individuals' reported "weeks worked last year" and "usual weekly hours." For instance, in 2000 this measure of labor supply varied from 862 hours worked per year (New York) to 1,456 hours (Minneapolis).

Similarly, we repeated the analysis for a second large educational group, women with college degrees. Participation rates of college educated women are generally higher than for women with high school diplomas, but we found significant cross-MSA differences in labor

[6]The PUMS defines unemployed persons as those without a job and looking for a job. It is difficult to know the extent to which a married woman might indicate that she is looking for a job if she intends to return to work at some point, but is not actively currently seeking a job. In any event, none of the conclusions below are altered if unemployed women are instead included in the labor force. (Unemployed women are only 1.5 percent of the sample; the average unemployment rate is only 2 percent.)

[7]In particular we try "standardizing" using the national age distribution. First the age distribution $f(a)$ is calculated for all the women in a sample. Then in each MSA j, for each age a the employment rate $E_j(a)$ is found. Finally, a standardized participation is obtained for MSA j rate: $E_j = \sum_{a=25}^{55} E_j(a)f(a)$. This makes virtually no difference for Table 1.

force participation for these women as well. For example, participation rates of college educated women with children range from 65 percent (Honolulu) to 86 percent (Albany).

2. Possible Explanations

It is likely that many factors are at play in producing the large variation across cities in observed married women's labor force participation. However, none of the most obvious explanations seems to be key. In particular, we focus initially on three factors that intuitively might influence participation: local housing prices, local wages, and the local unemployment rate (which might be an indicator of local *demand* for labor).

To set our discussion, we examine MSA-level regressions, separately for high school- and college-educated married women, in which local labor force participation is the dependant variable. As independent variables we have a measure of local housing costs (based on a quality-adjusted housing price index developed by Chen and Rosenthal (2008)), local wages (which we calculate using wages of *single* women, i.e., women who are not in the analyses), and the local unemployment rate. Results are reported in Table 3.[8]

One might expect that married women are more likely to work in relatively expensive cities, if only because their income is crucial to pay for high housing costs.[9] Our regressions do not indicate that this is the case. Moreover, we see in the first column in Table 1 that expensive cities such as New York, Los Angeles and San Francisco are actually close to the top of the list as cities with *low* participation; there is a negative correlation between female participation rate and our housing price index (-0.42).

We might also expect that women's labor decisions are influenced by local wages. In fact, in our regression the local wage rate for women—as indicated by the local wage or log-wage of single women (who generally have high labor force participation)—does not appear to have a large impact on local labor force participation.[10]

As we have emphasized, the cross-MSA variation is very large. Differences across cities are also quite persistent across decades. For example, the correlation between the 1990 MSA-level participation rates and 2000 MSA-level participation rates is 0.82, and comparable

[8]For all regressions, we also used the ACCRA city cost of living index as an alternative to the Chen-Rosenthal housing price index, using the 46 MSAs for which the ACCRA index is available. Key results were virtually unchanged.

[9]For instance, Fortin (1995), using data from the 1986 Canadian Family Expenditures Survey, shows that labor supply of some married women is affected by mortgage commitments.

[10]If instead we simply look at the relationship between observed wages for *married* women (among women who work) and the participation rate, we find a correlation coefficient of -0.46 (significant at the 0.01 level). That is, higher wages are associated with *lower* participation rates. Of course, given the selection process of women into the labor market, interpretation is difficult. For instance, this negative correlation would be expected if women who work are disproportionately drawn from the high end of the talent pool.

correlation coefficients are 0.82 for 1980–1990, and 0.81 for 1970–1980. These large and persistent differences are likely not due to local demand shocks. Having said that, transitory local shocks surely do matter, and in our regressions we include the local unemployment rate for men as a way of examining this possibility.[11] As expected, married women's labor force participation is lower when male unemployment is relatively high. It is helpful to keep this feature of the markets in mind when conducting explorations of other potential explanations.

Yet another contributing factor to variation in local female labor supply is child-care costs.[12] We make no attempts here to directly evaluate cross-city differences in child-care costs on the labor supply of mothers with young children. But there are good reasons to believe that child-care costs are at best a small part of the story. First, participation rates and average hours worked vary greatly across cities even among married women *without* children—women who presumably are little affected by differences in child-care costs. Second, across the 50 cities in the study, labor supply measures for women *with* children and women *without* children are positively correlated; the correlation is 0.80 (significant at the 0.01 level) for women with high school dimploma. It appears that the same city-specific factors affect both married women with children and without children.

A study of "power couples" by Costa and Kahn (2000) potentially offers some clues about cross-MSA variation in married women's labor supply. In that paper, college educated couples are shown to concentrate heavily in large metropolitan areas. If college educated women disproportionately locate in large cities, the average participation rate in these cities would be higher than in smaller cities. This argument, however, does not help explain the large differences in female labor supply that exist even among large cities. Nor is the argument helpful in understanding why labor force participation varies so widely across cities among women with a high school level of education.

"Social norms" might be important for understanding cross-MSA differences in female labor supply. Fernández et al. (2004), for instance, show that wives of men whose mothers

[11] Women's unemployment rates are mechanically related to participation rates as defined in this study, so as an alternative we use unemployment rates constructed for men as a measure of local labor demand. In particular, we compute unemployment rates of white men with high school diploma or above aged 25 to 55 years old in each of the 50 MSAs (using the usual definition, i.e., the proportion of people in the labor force who are unemployed). To account for differences in demographic distributions across cities, the unemployment rates are standardized using the national distribution of age and education. These standardized unemployment rates for men in the 50 MSAs vary between 1.2 percent and 4.4 percent in 2000.

[12] The evidence on the magnitude of labor supply elasticity with respect to child care costs is mixed and the range of elasticity estimates is wide. (See Blau (2003) for a review of this literature.) Baker et al. (2005) analyze the impact of Quebec's subsidized child care program on labor supply of mothers, finding a significant but relatively small impact.

worked are more likely to work. Thus, local history can matter. We do not immediately see how these considerations would lead to the sort of MSA-level variation in female labor force participation that we observe here, and more generally we do not pursue explanations based on social norms (or historical artifacts) here.

In this study we turn to a simple economic explanation that differs from initially appealing ideas mentioned above; we look at the role of commuting costs that vary across cities.

3. A Model of Labor Supply with Commuting Time

Among the important systematic ways in which cities differ are traffic patterns, congestion, and the resulting length of the commute to work. Commuting time can be viewed as a fixed time cost and/or monetary cost of going to work. Oi (1976), in a classic paper, introduces the idea that commuting time might play an important role in family labor supply decisions, including the joint decision about where to live. Cogan (1981) presents an analysis of labor supply with fixed money and time costs of labor market entry, arguing that fixed costs "are of prime importance in determining the labor supply behavior of married women." The general idea is developed in the urban economics literature, motivating a small empirical literature. For example, Gordon et al. (1989) show that women have shorter commute times than men regardless of income, occupation, or marital status.

Existing theoretical work on the role of fixed costs in labor supply does not treat the issue in a dynamic setting, nor does it focus on issues related to joint household labor supply decisions. As we show below, these issue may be crucial for understanding the role of fixed costs in labor force participation and hours worked (among individuals who do participate). Also, we know of no work that is using cross-city variation in commuting times as a way of learning about the potential importance of fixed costs in labor supply decisions.

As a starting point to motivate our empirical work, we present in this section a simple theory of fixed cost and labor supply. The analysis proceeds in two parts. First, we consider a one-person one-period case, i.e., a "static model." We then turn to a two-period formulation. The logic in this latter case can be used to understand how commuting affects *lifetime* labor supply (and indeed the results generalize easily to an n-period model). Alternatively, the logic can help sort out effects of fixed cost on labor supplied by a two-person household in a "collective model" of intrafamilial decisions (along the lines of Chiappori, 1988 and 1992).

3.1. **A Static Model.**

Consider the standard labor supply model with the addition of a commuting cost, which we assume here is a time cost. Let $c > 0$ denote the time cost of commuting incurred whenever hours of work are positive. In a one-period one-person model, we can write the budget constraint and time constraint, respectively, as

$$F = N + wT = pX + wL + wIc, \qquad \text{and} \tag{1}$$

$$T \geq L + Ic, \tag{2}$$

where T is the endowment of time, N is non-labor income, p is the price of the consumption good X, w, the wage, is the implicit price of leisure L or commuting, I is an indicator function equal to one when $L < T$ and equal to zero when $L = T$, and then F is full income.

For the specification of preferences, we will assume that the agent has a twice differentiable, strictly concave utility function in which both the consumption good and the leisure are normal. (The assumption of strict concavity provides the agent with an incentive to smooth her consumption over time when we move to a multiple period model.) In such a world, the objective of the agent is to maximize utility, $u(X, L)$, subject to the constraints that leisure may never be larger than the time endowment and the budget constraint, which, from (1) and (2), can be written

$$N + w(T - L - Ic) = pX. \tag{3}$$

From a technical standpoint, the only difficulty is that the indicator function I make the budget constraint nonconvex.

To solve this problem, for any prices and nonlabor income, we can simply solve the problem under two regimes: (i) the agent pays commuting costs and is free to work or not and (ii) the agent does not pay the commuting cost and cannot work. The agent then selects the regime that provides the higher utility. Each of the two regimes provides us with a standard convex budget set, which in turn allows us to apply the theorems of concave optimization.

The problem facing the agent when she does not commute (and therefore cannot work) is simple: she merely sets consumption equal to $X = N/p$. In solving the problem when the agent *does* commute, necessary conditions are

$$u_L(X, L) \geq \lambda w, \tag{4}$$

$$u_X(X, L) = \lambda p, \qquad \text{and} \tag{5}$$

$$N + w(T - L - Ic) = pX. \tag{6}$$

As we have said, the agent solves each of these problems and chooses the outcome that yields higher utility.

We depict the indirect utility function that results from each of the two optimizations problems in Figure 2, which shows the relationship between the indirect utility functions and full income F. Because leisure is a normal good, we know that the *no commute/no work* indirect utility, V_{nc}, must cut the *commute/work* indirect utility, V_c from below. Let F^* depict full income such that the agent is indifferent between working and not.[13] To the right of F^*, the agent chooses not to work, and to the left of F^* the agent works. To develop some intuition for the agent's choices, consider behavior at $F = F^*$. If the agent works, she consumes more, but pays both the cost of commuting and the cost of foregone leisure due to work. If she chooses not to work, she accrues additional leisure and does not need to pay the commuting costs, but must reduce her consumption.

Comparative statics are easy in this set-up. Consider, for example, an increase in commuting costs c. The impact on the indirect functions is depicted in Figure 3, which shows the comparative statics for an increase in commuting cost to $c' > c$. If the agent is initially not commuting (i.e., if her income initially exceeds F^*), the increase of course has no impact on her utility. If she initially is working, the utility of working must decline. The critical value of full income is now smaller. An increase in commuting cost of course never induces a non-working agent to join the labor force, but it can induce her to withdraw from the labor force.

Our model applies to any individual, but there is some reason to think that it may be of greater relevance to women than to men if women assume greater responsibilities at home. Suppose we extend the model by allowing workers to allocate time between work, commuting, leisure, and also home production. Kolesnikova (2007) follows this route, showing that individuals who have a fixed time commitment to home production (e.g., women who are committed to the care of young children at home) will be especially likely to have labor supply that is sensitive to the fixed commuting cost.

We have one final observation in the static model. Conditional on staying employed, individuals facing an increased commuting cost will work fewer hours. Figure 4 illustrates the logic: When commuting cost is c_0, the choice is $(L^*(c_0), X^*(c_0))$. After an increase in c, the new budget constraint is BC_1. Since leisure and consumption are both normal, the new optimal point must lie on a segment of BC_1 between points y and z. Thus the decrease

[13]We are looking for a non-trivial solution, so we do not consider a case when $V_{nc} > V_c$ for all positive values of F.

in leisure L is less than $x - y = \Delta c$. Since the decrease in leisure is less than the increase in commuting time, work hours decline.

3.2. **A Two-Period Problem.** The static model provides insight, but it misses a key feature of the intertemporal problem. In a multiple-period setting, an agent can move financial resources between periods to smooth her consumption, and this in turn results in some subtle differences in behavior. We discuss here a simple two-period case; the idea generalizes quite easily to more periods. To focus on the intertemporal aspects of the problem, we assume that the prices (w and p) are the same in each period, and to conserve on notation, we assume the interest rate and discount rate are zero (though nothing in our analysis changes if we assume they are positive constants). To keep from having to characterize all permutations, we assume that if the agent works at all, she always work early in her career, and if she withdraws from the labor market at all, she does so later in her career.

With two periods, total income becomes

$$\sum_{i=1}^{2}(wT + N_i) = F_1 + F_2.$$

where F_i is full income in period i. The non-convex budget set is now

$$F_1 + F_2 = \sum_{i=1}^{2}(wT + N_i) = \sum_{i=1}^{2}(pX_i + wL_i + wI_ic), \tag{7}$$

and the objective function is

$$U = \sum_{i=1}^{2} u(X_i, L_i). \tag{8}$$

Our approach to the utility maximization problem is to exploit the intertemporal separability of the utility function and make this a two-stage budgeting problem. Let y_i be a (possibly negative) transfer in period i from the other period.[14] In the first stage, conditional on the price, the wage, and a given allocation of income, the consumer chooses the optimal consumption bundle in the same manner as she did in the single period problem. Let $V_c(F_i, w, p)$ denote the indirect utility function if she pays the commuting cost and $V_{nc}(F_i, p)$ denote the indirect utility function if she does not pay the commuting cost, where $F_i \equiv y_i + N_i + wT$. Now define

$$V^*(F, w, p) = \max\big\{V_c(F, w, p), V_{nc}(F, p)\big\}. \tag{9}$$

[14] Of course $y_1 = -y_2$.

The function $V^*(F, w, p)$ is continuous in its variables, but of course is not differentiable at the crossing point.

For the second-stage, the agent's problem is:

$$\max_{y_i} \left\{ \sum_{i=1}^{2} V^*(F_i, w, p) \right\} \tag{10}$$

subject to the budget constraint

$$\sum_{i=1}^{2} F_i = \sum_{i=1}^{2} (wT + N_i + y_i) = \sum_{i=1}^{2} (wT + N_i).$$

The utility function is continuous in F_i and the budget set is convex so the function has an optimum. (As we noted, there may not be a unique optimum because the agent is indifferent over the timing of employment. We will follow our convention of assuming any withdrawals from the labor market start with a second period withdrawal.)

While the explicit derivation of the optimal conditions is tedious, armed with our insights from the static model this is a simple problem to solve. If her nonlabor income is sufficiently high, the woman chooses not to work in either period. Conversely, if her nonlabor income is sufficiently low, then she will choose to work both periods. In each of these cases, the analysis parallels the static case.

More interesting, however, is the case where the agent finds it optimal to work in one period and not in the other. To see why, consider Figure 5. In Figure 5 we see that the two indirect functions create a nonconcave objective function, and this in turn has an important impact on the intertemporal decision. Consider an individual with full income F^* such that $V_c(F^*, w, p) = V_{nc}(F^*, p)$. In a one-period model, this person would be indifferent between working or not working. This same person, though, given full income $2F^*$ in the *two-period* model, would work one period, allocating $F^* - \delta$ to consumption and leisure in that period, and *not* work the next period, allocating $F^* + \delta$ for consumption and leisure in the second period. The average utility of the two periods exceeds what the agent would get from consuming full income of F^* in each perod. The intertemporal allocation essentially "convexifies" the objective function. Note that our argument holds for values of full income that differ modestly from $2F^*$; an individual with lifetime full income sufficiently close to $2F^*$ will optimally choose to work one period and withdraw from the labor force the second period.

Comparative statics in the two-period model are not particularly difficult. As in the static model, an increase in c can induce individuals to withdraw from the labor market—in this case cause a person who would have worked two periods to now work only one period, or cause a person who would have worked one period to withdraw completely from the labor market.

In contrast to the one-period model, in the two-period model an increase in the commuting cost (or, for that matter, an increase in income) can result in the worker supplying *more* hours in an a period in which she works. To see how this occurs, consider Figure 6. In Figure 6, the worker initially has just enough full income F^* to be indifferent between working both periods or working one period only.[15] Now if there is a small increase in the commute cost, the worker unambiguously prefers working in one period only. Relative to her behavior in when working in both period, when she moves to working only one period she has much lower full income F_i. Given that leisure is normal, she will consume less leisure, i.e., work longer hours in the period in which she works. (Since the consumption good is also normal, she has lower consumption as well). As Heim and Meyer (2004) note, the presence of nonconvexities in the budget constraint may result in the agent exhibiting discontinuous behaviors similar to an agent with nonconvex preferences (but conventional convex budget sets).

We can illustrate our point quite easily with an example in which preferences are Cobb-Douglas:

$$U = ln(X_1^{\alpha} L_1^{1-\alpha}) + ln(X_2^{\alpha} L_2^{1-\alpha}),$$

where the subscripts index the time period. The budget set is given by

$$N + w(1 - cI_1) + w(1 - cI_2) = pX_1 + pX_2 + wL_1 + wL_2, \tag{11}$$

where $N = N_1 + N_2$ is lifetime non-labor income, which we can assign to period 1 without loss of generality, and T is set to 1. We solve the problem via two stage budgeting. In particular, we solve the first period problem,

$$\max\{X_1^{\alpha} L_1^{1-\alpha}\} \quad \text{s.t.} \quad N - y + w(1 - I_1 c) = pX_1 + wL_1, \tag{12}$$

where we let y be the (possibly negative) transfer from period 1 to period 2. Next we solve the second period problem,

$$\max\{X_2^{\alpha} L_2^{1-\alpha}\} \quad \text{s.t.} \quad y + w(1 - I_2 c) = pX_2 + wL_2. \tag{13}$$

[15] She can work both periods, receiving $V_c(F^*, w, p))$ in each period. Or she can work one period, allocating full income $F^* - \delta$ to that period, and then withdraw during the other period, allocating $F^* + \delta$ full to that latter period.

Let $V(1,1)$ be lifetime indirect utility for an agent who works both periods. Then, in the case in which the agent works both periods, maximization problems (12) and (13) yield

$$V(1,1) = k_1 + \ln(N - y + w(1-c)) + k_2 + \ln(y + w(1-c)), \tag{14}$$

where k_j are constants. Solving the second-stage budgeting problem, we find $y^* = N/2$. So the full solution when the individuals works both periods entails

$$V(1,1) = 2k_1 + 2\ln(N/2 + w(1-c)). \tag{15}$$

Now let $V(1,0)$ be indirect utility when the agent works in the first period but not the second:

$$V(1,0) = k_1 + \ln(N - y + w(1-c)) + k_2 + \alpha \ln(y). \tag{16}$$

The second-stage necessary condition then yields

$$\frac{dV(1,0)}{dy} = -\frac{1}{(N - y^* + w(1-c))} + \frac{\alpha}{y} = 0.$$

Rearranging we get

$$y^* = \frac{\alpha}{1+\alpha}.$$

Thus, the indirect utility is

$$V(1,0) = k_1 + \ln\left(\frac{N + w(1-c)}{(1+\alpha)}\right) + k_2 + \alpha \ln\left(\frac{\alpha}{(1+\alpha)}(N + w(1-c))\right). \tag{17}$$

Our focus is on the situation in which the worker is indifferent about whether to work one period or two periods, i.e., in which $V(1,1) = V(1,0)$. Our goal is to demonstrate that if the agent works one period rather than two, she will work longer hours in the period in which she does work. To see this, note first of all that when working both periods, leisure demanded is

$$L_1(1,1) = \frac{(1-\alpha)}{w}(N/2 + w(1-c)). \tag{18}$$

In contrast, the leisure demand function is

$$L_1(1,0) = \frac{(1-\alpha)}{w}\frac{(N + w(1-c))}{(1+\alpha)} \tag{19}$$

for a worker who works in the first period only. Thus period-one leisure demand is lower when working one period rather than when working both periods if

$$(N/2 + w(1-c)) > \frac{(N + w(1-c))}{(1+\alpha)}, \tag{20}$$

which simplifies to

$$\frac{N}{2} < \frac{\alpha}{1-\alpha} w(1-c). \tag{21}$$

It is easily verified that this last condition must hold. Suppose to the contrary, that the condition fails, e.g., holds with equality, $\frac{N}{2} = \frac{\alpha}{1-\alpha} w(1-c)$. Then leisure demand when working both periods is

$$L_1(1,1) = \frac{(1-\alpha)}{w}(N/2 + w(1-c)) = \frac{(1-\alpha)}{w}\left(\frac{\alpha}{1-\alpha} w(1-c) + w(1-c)\right) = (1-c).$$

But this means the worker is receiving no wage income despite paying the commuting cost, i.e., is commuting but not working at all! Clearly this regime cannot provide the same utility as the regime in which the worker does not commute in period two.

In short, the Cobb-Douglas example nicely demonstrates the point we made above: If a worker is indifferent between working one or two periods, and if she opts to work one period only, her optimal choice must entail working relatively *longer* hours in that one period.

3.3. **Observations from our Models.** There are three key points that we wish to emphasize from our analysis:

(1) In both the static and two-period models we have an unambiguous prediction about the impact of commuting cost c on labor force participation. An increase in c can never cause an agent to work when she otherwise would not have worked, but *can* induce withdrawal from the labor market.

In a cross-section of cities in which commuting times vary, we would expect, all else equal, to see lower labor force participation rates in cities with high commuting times. Also, as mentioned above, we might expect this pattern to particularly pronounced for women with children.

(2) In a multiple-period model, an increase in c can cause an agent to work fewer periods, and in such a case can cause her to work *longer* hours in a period in which she does work. This might well be the typical lifetime response to an increase in c. Consider, for instance, a person who is making a lifetime decision over say 600 months of potential work (e.g., 50 years $\times$ 12 months). If the individual initially intends to supply labor in, say, 480 months, she is presumably roughly indifferent between that decision and working 479 months. Then if a small increase in c does induces her to work 479 months, it will also induce her to work longer hours in the months in which she does work.

While the focus of this paper is on labor decisions of women, our logic should apply also to men. We predict that work hours will be longer in cities with high commuting times, and in these same cities men will work fewer periods. Of course, in the U.S. most healthy "prime age" men participate in the labor force and work "full time." But "full time" might nonetheless entail longer hours for men in cities with relatively long commuting times. Also, if periods of leisure are concentrated at the end of the potential work life cycle, our theory predicts that the age of retirement will be lower in cities with high commute times.

(3) Some readers may have noticed that our two-period model is formally the same as a two-person "collective model" of household behavior (along the lines of Chiappori (1988) and Chiappori (1992)), with the objective function (8) now having a person index instead of a period index. Under this interpretation, an increase in commuting times can cause one person to withdraw from the labor force while inducing the other person, who remains in the labor force, to work longer hours. If the wife is generally the person withdrawing from the labor force in a husband-wive household, this reinforces the general observation that an increase in commuting time can be associated with *longer* work hours among married men.

With these ideas in mind, we turn now to an empirical investigation of the relationship between commuting times across U.S. cities and labor supplied.

4. An Empirical Analysis of Labor Supply and Commuting Time

4.1. **MSA-Level Labor Force Participation and Commuting Time.** To begin the empirical analysis, average commuting time is computed for each of the 50 MSAs. The Census asks respondents about how long it takes them to get to work; this number is multiplied by two to obtain the daily commute measure. For each MSA, average daily commuting time is estimated for working married women, women with children and married men. These results are summarized in Table 4. Women on average have shorter commute than men, and women with children commute even less. For men, commuting to work and back takes on average at least 43.5 minutes a day (in Dayton) and can be as much as 76.1 minutes a day (in New York). For women, the daily commute varies between 38.6 and 63.4 minutes a day on average, depending on the MSA. In short, the cross-city differences in average daily commute are very substantial. The cities with the longest commute times are New York, Washington, D.C., and Chicago. Dayton, Oklahoma City, and Buffalo have the shortest average commute times in the sample of 50 MSAs.

Clearly, commuting time is important. A married man in the median MSA works approximately 9 hours per day, and typical total daily commutes easily tops one hour in many cities.

To examine an effect of commute times on participation decision of women it is necessary to have a measure of average commuting time cost in an MSA. Since many women do not work, there is a selection bias in our estimates of their average commute. Thus, women's average commute time cannot be used as such a measure. On the other hand, most white married men do work. Their average commuting time can serve as an indicator of the difficulty of the commute in a city.[16] In what follows this measure is referred to as an MSA daily commuting time.

Our theory above gives an unambiguous prediction about the effects of commuting time on labor force participation, so we focus initially on the cross-MSA relationship between labor force participation by married women and commuting times.

Table 5 reports the results from an MSA-level linear regression in which the dependent variable is the MSA labor force participation rates, and average MSA commute time (for men) is the key explanatory variable. The analysis is performed separately for high school educated women and for women with a college degree. Results indicate that higher levels of the MSA daily commuting time are generally associated with lower levels of labor force participation rates among married women. The association is stronger for the women with a high school level of education. A 1 minute increase in the MSA daily commute time is seen to be associated with a 0.3 percentage point reduction in the labor force participation rate of high school educated women. We have included some of the other control variables that we discussed when we presented Table 3.[17] Among these variables, only the unemployment rate of non-Hispanic white men is found to be statistically significant at conventional levels. Failure to include this latter variable makes virtually no difference to the estimates of the impact of commuting time; there is a near zero correlation between MSA daily commuting time and the MSA unemployment rate.[18]

[16]We also used city commute data provided by Census. The numbers are very close to the estimates reported in Table 4, and produced similar results in our analysis.

[17]As with the analysis reported in Table 3, all regressions were also conducted using the ACCRA city cost of living index instead of the Chen-Rosenthal housing price index. Results were unchanged.

[18]As we discuss above, we do not have measures of child-care costs, which might also be relevant. Cortes and Tessada (2007) consider the possibility that in a city with a high concentration of low-skilled immigrants, who often work as maids and nannies, the cost of day care might be low, which in turn could increase the labor supply of high-skilled American women. They find no such effects except for an increase in work hours of women with graduate degrees (a group whom we do not study). Nonetheless, to satisfy curiosity, we did include their measure of low-skilled immigrants in our city-level regressions (for the 25 cities for which Cortes and Tessada measure is available). We similarly find no evidence that low-skill immigrants increase women's labor supply. More importantly, our key results about the effect of commuting time were virtually unchanged.

4.2. **Individual Level Analysis of Labor Force Participation.** We next look at individual level data, examining the relationship between the average MSA commute and women's labor force participation decisions. We estimated both a probit model and a linear probability model. Results were nearly identical, so we focus on the linear model, which is easier to interpret. In our analysis the dependent indicator variable, which equals 1 if the woman participates in the labor force, is multiplied by 100 (so that probability can be discussed in percentage points).

The usual problem when estimating labor supply is that wages are observed only for those who work. The approach here is to include in each regression only women with exactly the same level of education, i.e., women who are likely to have similar levels of market productivity.[19] As we discuss above, the relative value of women's time spent at home is plausibly higher among women with young children than those without children, which in turn can make these women more sensitive to commute time than other women. To account for this possibility, the sample is also divided into three separate groups: women with children younger than 5 years old, women with children none of whom are younger than 5, and women with no children.

Panel A of Table 6 presents the first piece of individual-level empirical evidence concerning the effect of the average MSA commuting time on work force participation decisions by married women. It indicates that an increase in MSA commuting time is associated with a decrease in the probability of being in the labor force for all the groups. The effect is the largest for high school graduates with young children. For them, a 1 minute increase in the average MSA commuting time is associated with a decrease in the probability of labor force participation of 0.59 percentage points. For high school educated women with children but no young children, the associated decrease in probability of labor force participation is 0.38 percentage points, while for high school educated women with no children the estimated effect is 0.22 percentage points. For women with a college degree the effect of longer average commuting time on labor force participation decisions is somewhat smaller, as one would expect if the opportunity cost of their time at work is higher than women with a high school level of education.

[19] While there is substantial debate on the validity of education as an instrument in this context, in fact education is frequently used as an instrument in estimating labor supply. See, for example, Ham and Reilly (2002), MaCurdy (1981), Altonji (1986), Ham (1986), and Reilly (1994).

In Panel B of Table 6, the model specification includes other factors that affect labor supply decisions, such as the number of children, non-labor income,[20] and controls for 5 years age intervals. In addition, an MSA unemployment rate for white men is included to control for local labor market conditions. When these factors are included, the estimated effects of the average MSA commuting time on labor supply decisions are slightly smaller (in absolute value) than the ones reported for the simple model in Panel A. However, the overall pattern is similar.

The labor supply function in the specification presented in Panel B of Table 6 includes non-labor income (primarily husband's income). Because the data are cross-sectional, this income includes a transitory component. As is discussed in the classic work on female labor supply by Mincer (1962), the husband's *permanent* income may be the more relevant construct for affecting a married woman's labor supply.[21] We thus substitute husband's education for non-labor income to create a "reduced form" variant of our regression, and report results in Panel C of Table 6. Estimates of the relationship between MSA commuting time and women's labor supply are little changed.[22]

Our measure of the MSA average commuting time relies on individual commuting times reported by Census respondents, which are likely to be quite noisy. The National Transportation Statistics (2005) of the Bureau of Transportation Statistics provides some alternative measures of MSA traffic congestion, such as annual roadway congestion index, annual person-hours of highway traffic delay, and annual highway congestion cost.[23] These measures are available for 48 of the 50 MSAs used in this study. When any of these measures are used (instead of average MSA commuting time) in the regressions reported above, the estimated relationship between women's labor force participation and the MSA traffic congestion is negative and significant. Alternatively, we can use these later variables as instruments. The estimated coefficients are quite similar to OLS results.

4.3. **Changes in MSA Labor Force Participation, 1980 to 2000.** In this section we return to the MSA-level analysis, but now examine data from the 1980 and 1990 Census

[20]Non-labor income is calculated as a difference between family income and woman's income from wages. Using husband's income (i.e., the difference between family income and woman's total own income) gives essentially the same results, as the two measures are highly correlated.

[21]In fact, our empirical exercise is similar in spirit to Mincer (1962), who also uses cross-city variation to examine female labor supply. His work of course does not consider commuting times, and in our work we are not trying to estimate labor supply elasticities.

[22]In all analyses standard errors are obtained using a Huber-White sandwich estimator of the variance using clustering on MSA level.

[23]The reported source is Schrank and Lomax (2005) from the Texas Transportation Institute.

PUMS. We compute labor force participation rates of white married women in 50 large MSAs in 1980 and 1990. The average MSA daily commuting time is also calculated using 1980 and 1990 Census PUMS data, and is used as an explanatory variable.[24] In analyzing 2000 data, neither wage nor housing prices seem to matter much for married women's participation at the MSA level (see Table 5) so we do not include them in our regressions here.

As it was the case in 2000, there is a significant variation across cities in labor force participation rates of white married women in 1980 and 1990. For instance, the participation rates of high school educated women vary from 41 percent in Pittsburgh to 67 percent in Greensboro in 1980 and from 51 percent in New York City to 78 percent in Minneapolis in 1990. As in 2000, this variation is "explained" in part by commuting times: longer average commuting times within an MSA are associated with lower levels of married women's labor force participation.[25]

We next ask if *changes* in commuting time over the decades under study are correlated with MSA-specific trends in women's labor force participation.

In order to conduct a differences-in-differences analysis, we first calculate the differences in female participation rates between the Censuses dates for each city. When we undertake this first step we find that the labor force participation rose about 16 percentage points on average between 1980 and 2000 for high school educated married women and about 14 percentage points for college educated women, with most of the increase occurring between 1980 and 1990. However, the rise of participation is not uniform across MSAs. From 1980 through 2000 high school educated women's employment rose by only 5 percentage points in San Francisco, while increasing by 24 percentage points in Buffalo. During the 1990s, some MSAs experienced moderate increases in participation rates of married women while in others participation declined (e.g., it fell by 7 percentage points in Honolulu).

Our second step is to look is to look at MSA-specific changes in the average commuting time. We find that between 1980 and 2000 the daily MSA commute increased by about 5 minutes on average.

We then examine a relationship between changes in the MSA-level commuting time and changes in women's participation rates over the same period of time. To control for changes in labor market demand conditions over the years, changes in unemployment rates of white

[24]Unfortunately, we cannot examine earlier decades because we do not have commuting time data.

[25]As was true of the 2000 data, including the MSA unemployment rate of white men in the regression to control for labor demand conditions does not change the effect of a longer commute on the employment rates of white women. Compare columns (1) and (2) for each group in Table 7.

men are also included (in specification (2) for each regression). Results are presented in Table 8. Panel A shows the results of the linear regression estimation for white married women with a high school degree. Panel B present the results for women with a college degree. We find a negative relationship between the change in commuting time in an MSA and the change in the employment rate married women in that MSA; cities in which commuting time increased most rapidly generally also experienced slower growth in female labor force participation.

4.4. **Commute Times and the Labor Supply of Men.** The focus of our paper is the labor supply of married women, but as we note above, our analysis also has predictions about work hours for married men. In particular, suppose most prime age married men participate in the labor market.[26] Then our model predicts that work hours will be *longer* for men in cities that have longer commute times. We also expect that they will work fewer periods over the course of a lifetime. If our assumption is correct that the extra periods of leisure are most likely to be taken at the end of the work life cycle, this translates into a prediction that the age of retirement will be younger in cities that have long commute times.

Table 9 provides an examination of the first of our predictions. We find that indeed men work longer hours in cities that have longer average commuting times. An extra minute of daily commute is associated with a 2 hour increase in annual hours worked. This effect is not negligible. If city A has a commute time that is 20 minutes longer than city B, we predict that men will work 40 hours (almost a full week) longer per year in city A.

Table 10 reports results relevant to the second of our predictions, about the age of retirement. We perform a crude but useful calculation for men ages 55 to 65. We form an indicator that equals 100 if the man is retired (0 otherwise), and then regress that variable against some control variables—age (entered as dummy variables), race, education, and the log of wage in the MSA where he lived 5 years previously—and also the commute time in the city in which the individual lived 5 years previously. We expect longer commute times to be associated with a higher probability that the individual is retired, i.e., that he is consuming periods of leisure late in his potential work life cycle. We find that this is indeed the case. Again, the association is not negligible. If city A has a commute time that is 20

[26]This indeed is the case. Participation of married men aged 25 to 55 is well above 90 percent in every city in our analysis, and the correlation between commuting time and participation is very close to zero.

minutes longer than city B, we predict that probability of being retired (for a man aged 55 to 65) is 1.34 percentage points higher in city A.

5. Concluding Remarks

Using Census Public Use Micro Sample (PUMS) data for 1950 through 2000, we find wide variation in labor market participation rates of white married women in 50 large U.S. metropolitan areas. This wide variation is found in all years, and appears for women with different levels of education as well as for women with and without children. Among a number of explanations for observed cross-city differences in female labor supply one emerges as particularly important: Married women's labor force participation decisions appear to be related to MSAs' commuting times.

While the possibility that fixed employment costs might matter for female labor supply was raised in the literature many years ago (e.g., Oi, 1976, and Cogan, 1981), we know of no work that has explored implications for the cross-city variation in labor supply. With this in mind we have undertaken an analysis that, first, sets out a simple theory of labor supply. In our theoretical exploration, we emphasized the twin predictions: An increase in commuting cost c will, all else equal, (1) reduce the number of periods an individual will want to work, and (2) increase the number of hours worked in periods when the individual does work.

Cross-section evidence is consistent with the theory. We find a negative association between commuting time and women's labor force participation rates in the three decades for which we have commuting time data (1980, 1990, and 2000). Similarly, the negative correlation between commuting times and participation appears also in a differences-in-differences analysis; metropolitan areas which experienced relatively large increases in average commuting times between 1980 and 2000 experienced slower growth of labor force participation of married women.

While we do not focus on the labor supply of men, we would expect the logic of our theory to pertain for them as well. Evidence bears out this expectation: We find that men tend to work longer hours in cities with relatively long commute times. Furthermore, men tend to consume more periods of leisure—concentrated at the end of the potential work lifecycle—in cities with relatively long commute times. That is, retirement appears to occur at younger ages for men in cities where commuting is more onerous.

For women, the effects of commute times are quite large. Consider, for example, high school educated women who have older children at home. From Table 6 we see that a 10 minute increase in an MSA's commute time is associated with about a 3 percentage point decline in participation. It is worth noting that the 10 minute increase in MSA commute time will generally be associated with an increase in the time required for many household activities, e.g., travel to the grocery store, school, or the little league baseball. In short, traffic congestion increases the cost of many activities, and these costs can in turn indirectly affect the labor supply decision.

In addition, of course, our findings are based on *equilibrium* correlations, so we must in the end be cautious about make statements concerning causality. The evidence is remarkably consistent with a theoretical story in which those factors that affect commute times (local topography, the quality of roads, weather, levels of traffic congestion, etc.) in turn have important causal effects on lifetime labor supply decisions in households. But the logic of our model also suggests that there will be *selection* of couples into cities that best suit their own preferences for commuting and working. For instance, households in which the woman chooses not to work are less likely to locate into (presumably expensive) locations that have good work opportunities *and* short commuting times. Such behavior, if it occurs, represents additional evidence that fixed costs are an important part of household location and labor supply decisions, but such behavior also makes it harder to isolate the casual impact of commute times on labor supply.

Naturally we would like to find appropriate instruments that affect commuting times but otherwise have no effect on labor supply behavior, as this would in principle allow the opportunity to estimate causal effects. However, this is not an easy task.[27] Finding such instruments, perhaps provided by natural experiments, is a useful goal for future research. For the moment we proceed instead to consider several interesting implications, given that fixed costs do seem to play some important role in labor supply decisions.

One possible implication concerns the century-long increase in the female labor force participation—the increase in married women's participation from only about 7 percent in 1900, to current rates (near 70 percent). There are doubtless many factors contributing to this trend, many of which have received careful examination in the literature. Little attention has been given, though, to the possibility that part of this trend is due to the reduction in commuting costs, owing to improvements in transportation technology (the expansion of modern public transportation, the introduction and continued improvement in

[27]Consider, for example, the construction of a rail transit extension to a particular neighborhood. This decreases commuting times but also makes the neighborhood more attractive to commuting people. Baum-Snow and Kahn (2000) show that while rail transit improvement increases ridership, the increase is mostly due to new migrants.

automotive technology, improvements in roads, etc.) and because of changes in residential patterns.

Our findings about the wide cross-city variation in the labor force participation of married women also introduce a new dimension to the current discussion about trends in the female labor supply. In particular these findings complicate discussions about women having reached a "natural rate" of labor force participation. The issue is how close to 1 we can expect this participation rate to be. Goldin (2006), Juhn and Potter (2006) and others show that labor force participation rates depend on a combination of demographic factors such as age, presence of children, education, and race. The "natural rate" of participation is expected to be different for different groups. Our research suggests that the maximum achievable rate of labor force participation for each group would also vary across cities (and also across countries) because of differences in commuting time.

Of course, commuting times in local communities also depend on population density, the resources devoted to transportation, and local planning (e.g., zoning laws that may sometimes serve to isolate residential communities from job locations). Thus, from a public policy perspective, it may be that targeted actions that reduce commuting times, would thereby increase labor force participation by women.[28]

Yet another open policy issue concerns the importance of variation in labor supply across cities for tax and welfare policy. In our model it is easy to see that because the time cost of commuting increases the reservation wage, a decrease in the income tax, for instance, would induce larger increases in work participation by women in an MSA with shorter average commutes than in an MSA with longer commutes. It would be interesting to empirically test this prediction.

Finally, we note that most empirical research in labor economics is conducted at the national level, with little attention given to the possibility that local labor markets differ in important ways. Our research points to the value of work that allows for the possibility of differences across locations in labor supply responses. More generally, there is surely a rich set of interesting issues yet to be examined around the interactions of urban characteristics and labor market outcomes.

[28]In this sense the work here is related to the "spatial mismatch hypothesis" literature that typically looks at job accessibility as a determinant of the generally poor employment prospects of minority workers. This idea was first introduced by Kain (1968).

References

Alesina, A., E. Glaeser, and B. Sacerdote (2005). Work and Leisure in the US and Europe: Why So Different? *NBER Working Paper 11278*.

Altonji, J. (1986). Intertemporal Substitution in Labor Supply: Evidence from Micro Data. *The Journal of Political Economy 94*(3), 176–215.

Baker, M., J. Gruber, and K. Milligan (2005). Universal Childcare, Maternal Labor Supply, and Family Well-Being. *NBER Working Paper 11832*.

Baum-Snow, N. and M. Kahn (2000). The Effects of New Public Projects to Expand Urban Rail Transit. *Journal of Public Economics 77*(2), 241–263.

Blau, D. (2003). Childcare Subsidy Programs. *Means-Tested Transfer Programs in the United States*, Robert A. Moffitt, ed. Chicago: University of Chicago Press.

Blau, F., M. Ferber, and A. Winkler (2002). *The Economics of Women, Men, and Work.* NJ: Prentice Hall.

Blau, F. and L. Kahn (2000). Gender Differences in Pay. *Journal of Economic Perspectives 14*(4-Fall), 75–99.

Blundell, R. and T. MaCurdy (1999). Labor Supply: A Review of Alternative Approaches. *Handbook of Labor Economics 3*, 1559–1695.

Blundell, R., T. MaCurdy, and C. Meghir (2007). Labor Supply Models: Unobserved Heterogeneity, Nonparticipation and Dynamics. *Handbook of Econometrics 6A*, 4667–4775.

Card, D., F. Kramarz, and T. Lemieux (1999). Changes in the Relative Structure of Wages and Employment: A Comparison of the United States, Canada, and France. *The Canadian Journal of Economics 32*(4), 843–877.

Chen, Y. and S. Rosenthal (2008). Local Amenities and Life Cycle Migration: Do People Move for Jobs or Fun? *Journal of Urban Economics 64*(3), 519–537.

Chiappori, P. (1988). Rational Household Labor Supply. *Econometrica 56*(1), 63–90.

Chiappori, P. (1992). Collective Labor Supply and Welfare. *Journal of Political Economy 100*(3), 437.

Cogan, J. (1981). Fixed Costs and Labor Supply. *Econometrica 49*(4), 945–963.

Cortes, P. and J. Tessada (2007). Cheap Maids and Nannies: How Low-skilled Immigration is Changing the Labor Supply of High-skilled American Women. *Working Paper*.

Costa, D. and M. Kahn (2000). Power Couples: Changes in the Locational Choice of the College Educated, 1940-1990. *The Quarterly Journal of Economics 115*(4), 1287–1315.

Fernández, R., A. Fogli, and C. Olivetti (2004). Mothers and Sons: Preference Formation and Female Labor Force Dynamics. *Quarterly Journal of Economics 119*(4), 1249–1299.

Fortin, N. (1995). Allocation Inflexibilities, Female Labor Supply, and Housing Assets Accumulation: Are Women Working to Pay the Mortgage? *Journal of Labor Economics 13*(3), 524–557.

Freeman, R. and R. Schettkat (2001). Marketization of production and the US-Europe employment gap. *Oxford Bulletin of Economics and Statistics 63*(s 1), 647–670.

Goldin, C. (2006). The Quiet Revolution that Transformed Women's Employment, Education, and Family. *American Economic Review 96*(2), 1–21.

Gordon, P., A. Kumar, and H. Richardson (1989). Gender Differences in Metropolitan Travel Behaviour. *Regional Studies 23*(6), 499–510.

Ham, J. (1986). Testing Whether Unemployment Represents Intertemporal Labour Supply Behaviour. *The Review of Economic Studies 53*(4), 559–578.

Ham, J. and K. Reilly (2002). Testing Intertemporal Substitution, Implicit Contracts, and Hours Restriction Models of the Labor Market Using Micro Data. *The American Economic Review 92*(4), 905–927.

Heim, B. and B. Meyer (2004). Work costs and nonconvex preferences in the estimation of labor supply models. *Journal of Public Economics 88*(11), 2323–2338.

Juhn, C. and S. Potter (2006). Changes in Labor Force Participation in the United States. *Journal of Economic Perspectives 20*(3), 27–46.

Kain, J. (1968). Housing Segregation, Negro Employment, and Metropolitan Decentralization. *The Quarterly Journal of Economics 82*(2), 175–197.

Killingsworth, M. and J. Heckman (1986). Female Labor Supply: A Survey. *Handbook of Labor Economics 1*, 103–204.

Kolesnikova, N. (2007). Three Essays on Local Labor Markets. Doctoral Dissertation, Carnegie Mellon University.

MaCurdy, T. (1981). An Empirical Model of Labor Supply in a Life-Cycle Setting. *The Journal of Political Economy 89*(6), 1059–1085.

Mincer, J. (1962). Labor Force Participation of Married Women: A Study of Labor Supply. *Aspects of Labor Economics*, 63–97.

National Transportation Statistics, . (2005). *National Transportation Statistics (NTS) 2005.* Research and Innovative Technology Administration/Bureau of Transportation Statistics.

Nickell, S. (1997). Unemployment and Labor Market Rigidities: Europe versus North America. *The Journal of Economic Perspectives 11*(3), 55–74.

Odland, J. and M. Ellis (1998). Variations in the Labour Force Experience of Women Across Large Metropolitan Areas in the United States. *Regional Studies 32*(4), 333–347.

Oi, W. (1976). Residential Location and Labor Supply. *The Journal of Political Economy 84*(4), 221–238.

Reilly, K. (1994). Annual Hours and Weeks in a Life-Cycle Labor Supply Model: Canadian Evidence on Male Behavior. *Journal of Labor Economics 12*(3), 460–477.

Ruhm, C. (1998). The Economic Consequences of Parental Leave Mandates: Lessons from Europe. *The Quarterly Journal of Economics 113*(1), 285–317.

Schrank, D. and T. Lomax (2005). *The 2005 Urban Mobility Report.* Texas Transportation Institute.

Ward, C. and A. Dale (1992). Geographical Variation in Female Labour Force Participation: An Application of Multilevel Modelling. *Regional Studies 26*(3), 243–255.

Figure 1. Married Women Labor Force Participation Trends

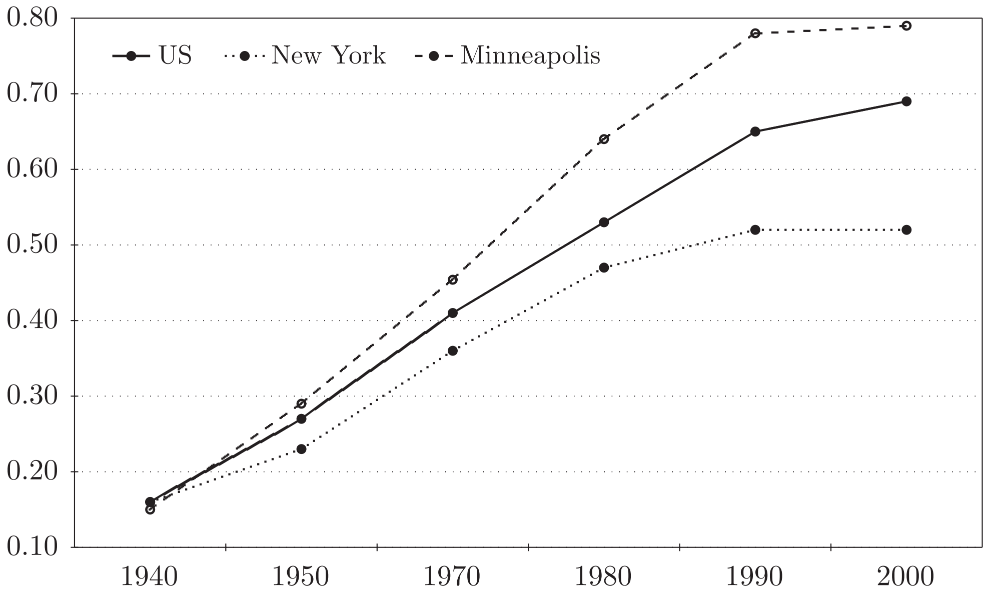

FIGURE 2. Indirect Utility as a Function of F—Commuting and Not Commuting

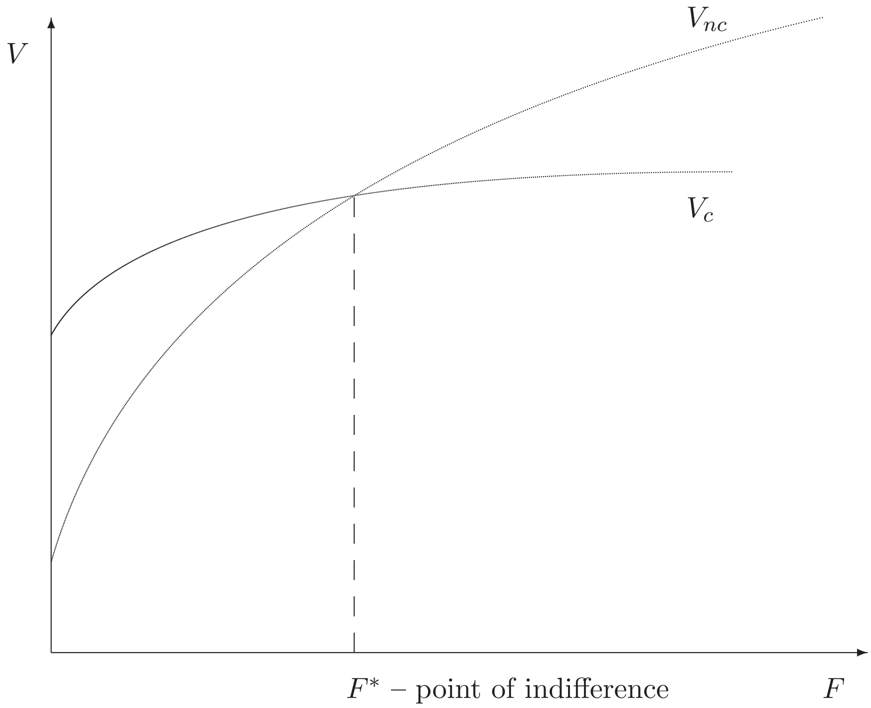

Figure 3. Consequence of an Increase in Commuting Time from c to c'

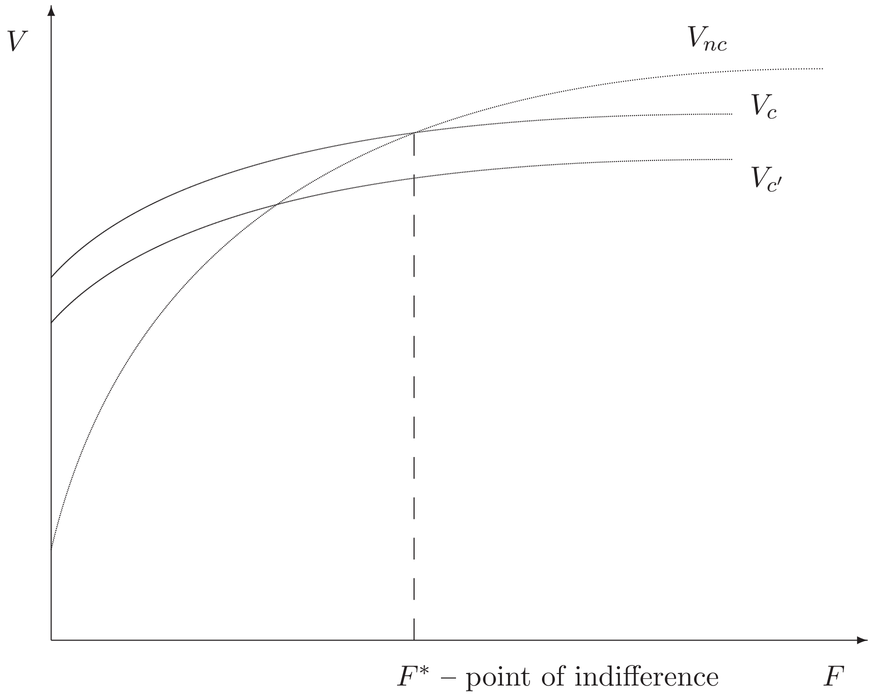

FIGURE 4. Effect of an Increase in c on Consumption and Leisure in the Static Model

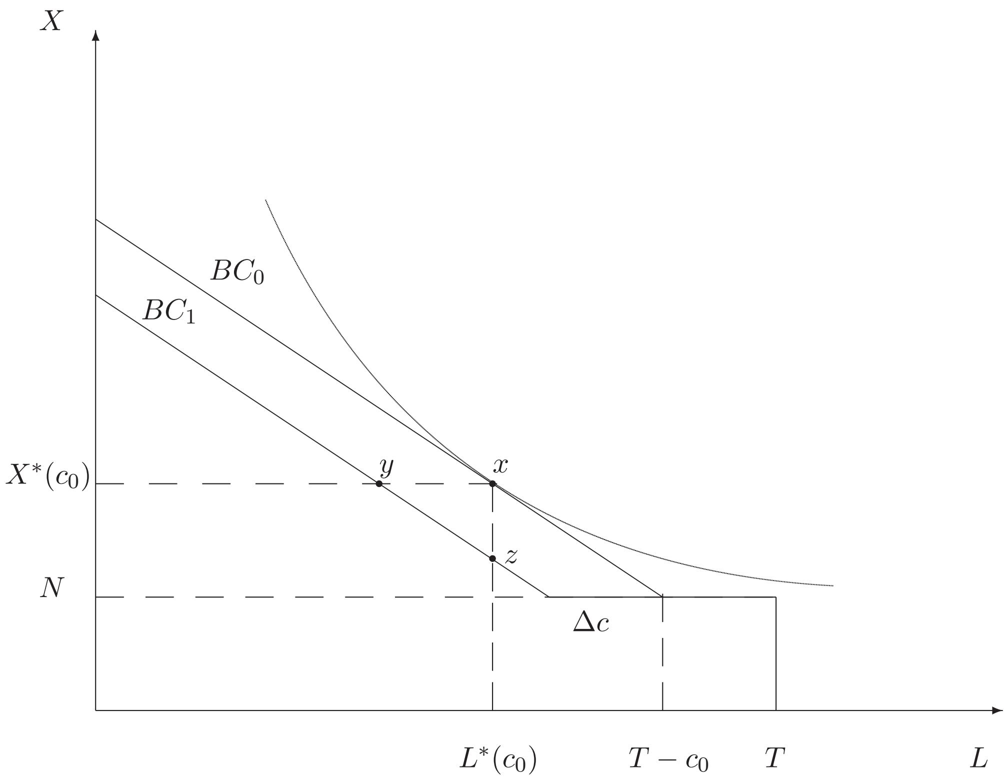

FIGURE 5. "Convexifying" the Objective Function with an Intertemporal Transfer

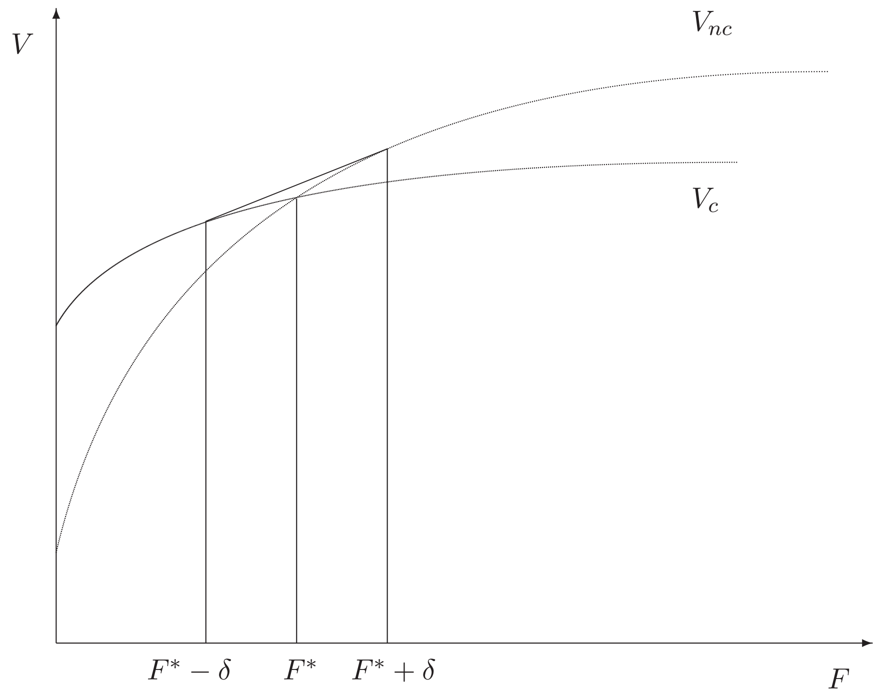

Figure 6. Indirect Utility Functions for a Person who is Indifferent between Working One Period and Two Periods

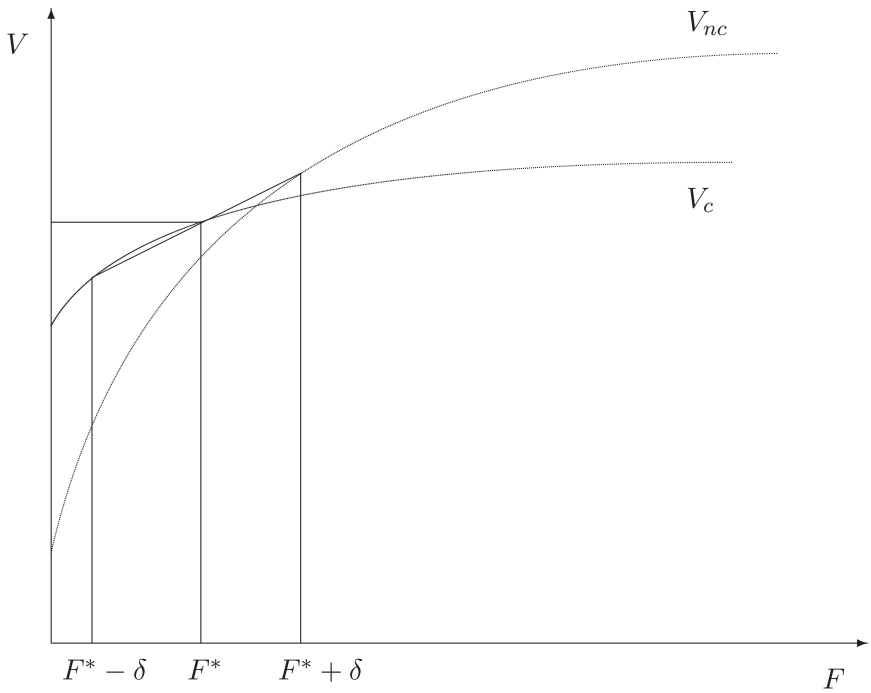

Table 1. Participation Rates of Non-Hispanic White Married Women with a High School Degree, Census PUMS

MSA	**Census Year**					
	2000	**1990**	**1980**	**1970**	**1950**	**1940**
United States	**0.69**	**0.65**	**0.53**	**0.41**	**0.27**	**0.16**
New York	0.52	0.52	0.47	0.36	0.23	0.16
Honolulu	0.55	0.62	0.46	0.34	—	—
Los Angeles	0.59	0.59	0.54	0.41	0.31	0.21
Miami	0.61	0.65	0.55	0.45	0.29	0.23
Houston	0.61	0.63	0.56	0.37	0.34	0.16
San Francisco	0.62	0.66	0.57	0.42	0.34	0.20
San Diego	0.62	0.62	0.51	0.39	0.29	0.14
San Antonio	0.63	0.60	0.53	0.37	0.37	0.18
Detroit	0.63	0.59	0.45	0.31	0.26	0.17
New Orleans	0.64	0.62	0.44	0.36	0.15	0.20
Sacramento	0.64	0.64	0.54	0.41	0.50	0.20
Birmingham	0.65	0.63	0.53	0.43	0.29	0.16
Pittsburgh	0.65	0.58	0.41	0.30	0.14	0.08
West Palm Beach	0.65	0.65	0.53	0.40	—	—
Portland	0.66	0.66	0.54	0.41	0.31	0.20
Norfolk	0.67	0.61	—	0.40	0.16	0.11
Oklahoma City	0.67	0.68	0.57	0.49	0.51	0.26
Chicago	0.67	0.64	0.53	0.40	0.28	0.17
Las Vegas	0.68	0.63	0.56	0.45	—	—
Phoenix	0.68	0.64	0.54	0.49	0.22	0.20
Seattle	0.68	0.67	0.55	0.41	0.25	0.16
Atlanta	0.68	0.69	0.59	0.44	0.34	0.26
Dallas-Fort Worth	0.69	0.69	0.60	0.50	0.35	0.24
Memphis	0.69	0.66	0.55	0.41	0.40	0.18
Philadelphia	0.69	0.64	0.49	0.35	0.23	0.15
Austin	0.69	0.71	0.65	0.47	0.37	0.23
Orlando	0.69	0.69	0.56	0.38	0.35	—
Tampa	0.69	0.69	0.53	0.43	0.26	0.22
Dayton	0.69	0.66	0.51	0.42	0.41	0.19
Cleveland	0.69	0.65	0.49	0.36	0.25	0.17
Charlotte	0.70	0.73	0.65	0.54	0.25	0.29
Nashville	0.70	0.71	0.60	0.44	0.34	0.34
Salt Lake City	0.70	0.67	0.53	0.44	0.22	0.12
Boston	0.71	0.65	0.54	0.40	0.24	0.11
Buffalo	0.71	0.65	0.47	0.37	0.23	0.12
Baltimore	0.72	0.69	0.54	0.41	0.29	0.18
Indianapolis	0.72	0.70	0.57	0.46	0.30	0.22
St. Louis	0.72	0.67	0.53	0.38	0.25	0.17
Cincinnati	0.72	0.68	0.52	0.35	0.21	0.11
Louisville	0.73	0.68	0.54	0.40	0.31	0.19
Washington	0.73	0.72	0.58	0.45	0.35	0.26
Richmond	0.73	0.73	0.60	0.39	0.35	0.25
Kansas City	0.74	0.69	0.57	0.49	0.31	0.23
Denver	0.74	0.72	0.58	0.47	0.38	0.14
Albany	0.74	0.70	0.56	0.43	0.29	0.15
Columbus	0.74	0.70	0.57	0.39	0.40	0.20
Rochester	0.75	0.68	0.54	0.43	0.31	0.23
Greensboro	0.77	0.76	0.68	0.59	0.37	0.30
Milwaukee	0.78	0.74	0.60	0.46	0.33	0.17
Minneapolis	0.79	0.78	0.64	0.44	0.29	0.15

Table 2. Summary of Married Women's Labor Force Participation by Education and Presence of Children in 50 Large MSAs

	With Children Under 5		With Children, None Under 5		No Children	
	High School	College	High School	College	High School	College
lowest MSA	0.29	0.41	0.56	0.65	0.54	0.78
10th percentile	0.44	0.54	0.65	0.7	0.65	0.82
25th percentile	0.48	0.56	0.67	0.72	0.69	0.84
median MSA	0.53	0.6	0.72	0.74	0.73	0.86
75th percentile	0.59	0.65	0.74	0.78	0.75	0.88
90th percentile	0.62	0.67	0.77	0.79	0.78	0.89
highest MSA	0.68	0.71	0.83	0.86	0.82	0.93

Note: Data are from the 2000 Census PUMS. The sample consists of non-Hispanic white married women aged 25 to 55 years with non-imputed data. The unit of observation is the MSA. There are 50 MSAs. The MSA-specific averages are standardized using the sample age distribution.

Table 3. MSA-Level Regression: Labor Force Participation of Married Women

N=50	**High School**	**Bachelor's Degree**
Quality Adjusted	−4.90	−4.09
Housing Index $\times 10^6$	(3.75)	(2.77)
Single Women	−0.002	−0.001
Wage	(0.003)	(0.002)
Unemployment Rate	−3.01**	−0.343
	(1.02)	(0.767)
Adjusted R^2	0.266	0.062

Significance Levels, One-Tailed Tests: *5%, **1%

Note: Data are from 2000 Census PUMS. The unit of observation is an MSA. There are 50 MSAs included in the analysis.

Table 4. Summary of Daily Commuting Time in 50 Large MSAs

	Married Men	Married Women	Married Women with Children
lowest MSA	43.5	38.6	37.4
10th percentile	47.0	41.1	40.5
25th percentile	50.4	44.0	43.0
median MSA	54.3	47.4	46.5
75th percentile	58.1	50.6	49.2
90th percentile	63.1	53.4	51.8
highest MSA	76.1	63.4	61.2

Note: Data are from the 2000 Census PUMS. The sample consists of non-Hispanic white married men and women aged 25 to 55 years with non-imputed data. The unit of observation is the MSA. There are 50 MSAs. The MSA-specific averages are standardized by age and education.

TABLE 5. MSA-Level Regression: Labor Force Participation of Married Women (with Commute Time as an Explanatory Variable)

N=50	**High School**	**Bachelor's Degree**
Quality Adjusted	−1.56	−3.08
Housing Index $\times 10^6$	(3.58)	(2.79)
Single Women	0.001	0.001
Wage $\times 10^2$	(0.003)	(0.002)
Unemployment Rate	−3.516**	−0.553
	(0.940)	(0.763)
MSA Average Daily	−0.003**	−0.002*
Commute	(0.001)	(0.001)
Adjusted R^2	0.390	0.097

Significance Levels, One-tailed Test: *5%, **1%

Note: Data are from 2000 Census PUMS. The unit of observation is an MSA. There are 50 MSAs included in the analysis.

Table 6. Individual-Level Regression: Labor Force Participation of Married Women, by Presence of Children and Woman's Education

	With Children Under 5		With Children, None Under 5		No Children	
	H. School	College	H. School	College	H. School	College
Panel A. Labor Force Participation, OLS Regression						
MSA Commute	−0.59**	−0.39**	−0.38**	−0.20**	−0.22*	−0.09†
	(0.231)	(0.075)	(0.142)	(0.071)	(0.110)	(0.068)
Intercept	83.9**	81.1**	90.6**	85.2**	83.0**	91.8**
	(12.7)	(4.5)	(7.8)	(4.0)	(6.2)	(3.8)
R^2	0.0087	0.0037	0.0040	0.0011	0.0012	0.0004
Panel B. Labor Force Participation, OLS Regression (with Covariates)						
MSA Commute	−0.44**	−0.21**	−0.24**	−0.04	−0.16†	−0.01
	(0.153)	(0.075)	(0.105)	(0.060)	(0.103)	(0.054)
Number of Children	−6.71**	−9.39**	−3.88**	−4.91**		
	(0.4)	(0.5)	(0.3)	(0.3)		
Non-Labor Income	−0.16**	−0.15**	−0.12**	−0.15**	−0.10**	−0.09**
	(0.016)	(0.009)	(0.009)	(0.005)	(0.007)	(0.006)
MSA Unemp. Rate	−5.01**	0.02	−4.49**	−1.31*	−2.39**	−1.47**
	(1.35)	(0.76)	(0.76)	(0.51)	(0.80)	(0.37)
Intercept	113.3**	109.5**	104.5**	106.3**	100.1**	100.3**
	(14.1)	(6.1)	(6.6)	(4.0)	(6.1)	(3.1)
Controls						
Woman's Age	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.0560	0.1002	0.0344	0.1099	0.0329	0.0772
Panel C. Labor Force Participation, Reduced Form using Husband's Education						
MSA Commute	−0.49**	−0.35**	−0.29**	−0.13*	−0.21*	−0.07
	(0.140)	(0.082)	(0.098)	(0.071)	(0.099)	(0.059)
Number of Children	−6.96**	−10.12**	−4.28**	−5.48**		
	(0.4)	(0.5)	(0.3)	(0.3)		
MSA Unemp. Rate	−5.26**	−1.01	−4.48**	−1.61**	−2.41**	−1.69**
	(1.21)	(0.78)	(0.75)	(0.60)	(0.78)	(0.43)
Intercept	100.5**	127.8**	103.9**	109.9**	100.2**	103.9**
	(15.9)	(7.2)	(6.0)	(5.3)	(5.9)	(4.4)
Controls						
Woman's Age	Yes	Yes	Yes	Yes	Yes	Yes
Husband's Education	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.0434	0.0716	0.0218	0.0512	0.0236	0.0513
N	14,572	28,254	44,602	34,478	23,956	29,286

Significance Levels, One-Tailed Test: †10% *5% **1%

Note: Standard errors are clustered at MSA level. Data are from the 2000 Census PUMS. The dependent variable is a labor force participation dummy multiplied by a 100 (so interpretation is in percentage points). The sample, drawn from the 2000 PUMS, consists of non-Hispanic white married women aged 25 to 55 years with non-imputed data, living in 50 largest MSAs. The MSA unemployment rate is for white men and is measured in percentage points. Controls for the woman's age are a set of dummy variables for each of the 5-year intervals between 25 and 55 years old. Controls for husband's education are a set of dummy variables for each education category (e.g. less than high school, high school diploma etc.)

TABLE 7. MSA-Level Regressions: Labor Force Participation of Married Women, 1990 and 1980

	High School Educated		College Educated	
	(1)	(2)	(1)	(2)

Panel A. MSA Labor Force Participation Rates in 1990

MSA Commute	-0.314^{**}	-0.279^{**}	-0.112^{*}	-0.110^{*}
	(0.097)	(0.086)	(0.056)	(0.060)
MSA Unemployment		-2.64^{**}		-0.15
		(0.674)		(0.449)
Intercept	82.38^{**}	88.46^{**}	79.14^{**}	79.49^{**}
	(4.93)	(4.60)	(2.86)	(3.06)

Panel B. MSA Labor Force Participation Rates in 1980

	High School Educated		College Educated	
	(1)	(2)	(1)	(2)
MSA Commute	-0.231^{**}	-0.226^{**}	-0.011	-0.010
	(0.104)	(0.085)	(0.082)	(0.082)
MSA Unemployment		-27.3^{**}		-0.639
		(0.561)		(0.540)
Intercept	66.04^{**}	73.57^{**}	58.61^{**}	60.37^{**}
	(5.21)	(4.56)	(4.15)	(4.39)

Significance levels : $*$: 5% $**$: 1%

Note: Data are from 1990 and 1980 Census PUMS. The sample size is 50. Standard Errors are in parentheses. The dependent variable is MSA-specific labor force participation of non-Hispanic white married women. Unemployment rate is for white men. Both unemployment and labor force participation are measured in percentage points.

Table 8. Differences-in-Differences Estimation

Panel A. Δ Labor Force Participation, High School Graduates

	1980-2000		1980-1990		1990-2000	
	(1)	(2)	(1)	(2)	(1)	(2)
Δ Commute	-0.51^{**}	-0.31^{*}	-0.28^{*}	$-0.25^{\dagger}$	-0.62^{**}	-0.65^{**}
	(0.188)	(0.171)	(0.167)	(0.158)	(0.196)	(0.177)
Δ Unemp. Rate		-2.61^{**}		-1.11^{*}		-1.69^{**}
		(0.649)		(0.425)		(0.484)

Panel B. Δ Labor Force Participation, College Graduates

	1980-2000		1980-1990		1990-2000	
	(1)	(2)	(1)	(2)	(1)	(2)
Δ Commute	-0.63^{**}	-0.48^{**}	-0.12	-0.08	-0.33	$-0.37^{\dagger}$
	(0.139)	(0.128)	(0.177)	(0.167)	(0.224)	(0.201)
Δ Unemp. Rate		-1.91^{**}		-1.20^{**}		-1.94^{**}
		(0.485)		(0.448)		(0.552)

Significance levels : †10% ∗5% ∗∗1%

Note: Data are from 1980, 1990 and 2000 Census PUMS. The unit of observation is an MSA. There are 50 MSAs included in the analysis. In each of the regressions, the dependent variable is change of LFP rates of non-Hispanic white married women in each of the 50 MSAs in the indicated time period. It is measured in percentage points. The independent variable is the change in average MSA daily commuting time of non-Hispanic white married men over the same period. The sample size is thus 50 in each regression.

TABLE 9. Individual-Level Regressions: Annual Work Hours of Married Men

	Full Sample	Men with Non-Working Wives	Men with Working Wives
MSA Commute	2.11**	1.38**	2.36**
	(0.16)	(0.32)	(0.18)
Log Wage	−76.29**	−76.34	−92.09**
	(2.71)	(4.85)	(3.31)
Number of Children	28.04**	27.06**	17.86**
	(1.03)	(2.06)	(1.20)
MSA Unemp. Rate	−43.35**	−46.47**	−42.53**
	(1.46)	(2.93)	(1.67)
Intercept	2312**	2271**	2406**
	(13)	(25)	(17)
Controls			
Age	Yes	Yes	Yes
Education	Yes	Yes	Yes
R^2	0.0329	0.0483	0.0286
N	323,468	82,509	240,959

Significance Level, One-Tailed Test: **1%

Note: Data are from the 2000 Census PUMS. The dependent variable is annual work hours. The sample consists of non-Hispanic white married men aged 25 to 55 years with non-imputed data. The MSA unemployment rate is for white men and is measured in percentage points. Controls for the man's age are a set of dummy variables for each of the 5-year intervals between 25 and 55 years old. Controls for education are a set of dummy variables for each education category (e.g. less than high school, high school diploma etc.)

Table 10. Individual-Level Regressions: Probability of Being Retired for Men Aged 55 to 65

	Full Sample	Men who Moved	Men who Did Not Move
Commute in MSA, Five Years Ago	0.067** (0.0082)	0.234** (0.0308)	0.030** (0.0085)
Controls			
MSA wage	Yes	Yes	Yes
Age	Yes	Yes	Yes
Education	Yes	Yes	Yes

Significance Level, One-Tailed Test: **1%

Note: Data are from the 2000 Census PUMS. The sample consists of men aged 55 to 64 years with non-imputed data. The dependent variable is binary variable (0 or 100) indicating whether or not a person is in the labor force. MSA wage is calculated for men 45 to 54 years old who are employed full time.)