NBER WORKING PAPER SERIES

HOUSEHOLD VS. PERSONAL ACCOUNTS OF THE U.S. LABOR MARKET, 1965-2000

Casey B. Mulligan
Yona Rubinstein

Working Paper 10320
http://www.nber.org/papers/w10320

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2004

We appreciate the comments of George Akerlof, Gary Becker, Claudia Goldin, Bob Michael, Robert Pollak, seminar participants at the University of Chicago, and the financial support of the University of Chicago's Alfred P. Sloan Center for the Study of Working Families. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

Household vs. Personal Accounts of the U.S. Labor Market, 1965-2000
Casey B. Mulligan and Yona Rubinstein
NBER Working Paper No. 10320
February 2004
JEL No. J22, J12

ABSTRACT

The empirical labor supply literature includes some simple aggregate studies, and some individual-level studies explicitly accounting for heterogeneity and the discrete choice, but sometimes leaving open the ultimately aggregate questions that motivated the study. As a middle ground, we construct household-based measures of labor supply by within-household aggregating answers to the usual weeks and hours worked questionnaire items. Household (H) measures are substantially different than the more familiar person (P) measures: H employment rates are relatively higher, with little trend, and relatively little fluctuations. From the H point of view, essentially all aggregate hours trends and fluctuations can be attributed to changes on the "intensive" margin and not the "extensive" margin – a characterization that is opposite of that derived from P measures. The cross-H distribution of hours is richer, and less spiked, than the cross-P distribution. Labor supply is more wage elastic from an H point of view.

Casey B Mulligan
University of Chicago
Department of Economics
1126 East 59th Street, #506
Chicago, IL 60637
and NBER
c-mulligan@uchicago.edu

Yona Rubinstein
Tel Aviv University
yonar@post.tau.ac.il

Table of Contents

While labor supply is often modeled as a family decision (e.g., Becker 1965, Ashenfelter and Heckman 1974, Smith 1977, Becker 1985, Costa and Kahn 2000), and so many indicators of economic activity are measured at the family level, both microeconomic and macroeconomic labor supply continue to be measured at the person level. For example, the unemployment rate is the fraction of *persons* who are not employed during the week, the year, or some other time interval, and are actively seeking work (or on layoff) during that interval. It is the fraction of *persons* working, not the fraction of families working, that is one of the primary indicators used by the National Bureau of Economic Research to declare a recession. Another example: the gender wage gap is calculated as the average wage of working female persons relative to the average wage of working male persons, rather than the wage of a wife relative to that of her husband.

These standard personal labor market indicators are incomplete and potentially misleading because they do not put labor market activity in a family context. First of all, the majority of working-aged adults live with a spouse and apparently share their income. Family structures have changed substantially over time, and differ substantially across races. Important examples include the marked decline in marriage rates among working-aged adults (although those rates are still above 50%), and the rapid growth in dual-earner couples.[1] Second, family situations may have a big impact on the labor market.[2] It has been argued in the literature (eg., Sweet 1973, pp. 67-8, or Bowen and Finegan 1969, pp. 96f), and is a part of common knowledge, that the presence and age of children affects married female labor supply. More recently, Johnson and Skinner (1986) argue that an important fraction of increased

---

[1]These trends have been documented in various publications. See for example, Waite's (1995, Figure 1) PAA Presidential Address and Hayghe (1990, Chart 1) for some documentation of U.S. marriage and dual-earner trends, respectively.

[2]It has also been argued (eg., Winkler 1998) that family, social, and financial implications of female earnings are best seen in a family context. For example, she (pp. 42-3) suggests that a variety of expenditure decisions might be made differently when the wife earns more than her husband. Perhaps the family expenditure literature surveyed by Lundberg and Pollak (1996) is consistent with this view. On the other hand, Tenn (2001) suggests than husbands dominate family migration decisions as much in 1990 as they did decades earlier.

female labor supply and the increase in dual-earner couples is a consequence of increased divorce rates. Daniel (1993) and others have even argued that marriage increases a man's wage. The purpose of our paper is to help complete the usual set of labor market indicators by constructing household (H) labor supply accounts, comparing them with the well-known personal (P) accounts, and considering some of the implications of those differences.

In theory, H accounts of the labor market could be very similar to, or very different from, P accounts. We find some of each, namely:

- the employment rate is higher in the H account than in the P account
- H nonemployment is age-related, and few H nonemployed are actively seeking work
- the H employment rate trends slightly down, while the P employment rate trends up
- neither the level, trend, nor changes in the H employment rate can be predicted from the P employment rate and its changes
- the H employment rate is well modeled by gender-partitioned P employment rates, if the model presumes very little sorting (as opposed to positive or negative sorting) of spouses on hours
- husband-wife hours correlations are much weaker than other husband-wife correlations (eg., age, years of school, race, religion, etc.)
- when aggregate hours are decomposed into the H employment rate and hours per employed household, the latter accounts for essentially all of the trend and fluctuations.
- labor supply is more wage elastic in the H accounts than in the P accounts
- the gender wage gap closes substantially more in the H accounts than in the P accounts
- P and H accounts of labor supply are more different than are P and H accounts of earnings[3]

Some of these findings, such as the first, were to be expected given that households are aggregates of persons. Indeed, in theory, the entire H accounts might be constructed from P accounts although, without seeing the H accounts, it is impossible to predict exactly which theory should be used for the construction. For example, should we expect positive, negative, or zero "sorting" of household members on hours? Our H accounts show that the H employment rate, and the cross-H distribution of hours, is similar to what we would find if households were formed by randomly pairing a man with a

[3]P and H distributions of earnings have been compared by Mincer (1974), Smith (1979), Cancian, Danzinger, and Gottschalk (1993) and Blackburn and Bloom (1995).

woman, and holding both person's hours fixed at what they are in the observed pairings.

The easy availability of micro-data and computing power during the last three decades, plus the development of important microeconometric techniques, has led to a proliferation of individual-level studies of labor supply, as well as other behaviors. Our paper is part of a more recent trend toward considering more aggregate measures of behavior, even when less aggregate measures are available, when the aggregate measures correspond more closely with the key concepts in the economic models. For example, public finance economists now study "taxable income" and its relation with the tax law, rather than decomposing it into its many distinct behaviors (tax cheating, labor supply on various margins, occupational choice, demand for fringes, etc.) because taxable income is a key concept in the theory and the various sub-behaviors are of interest mainly through their contribution to taxable income.[4] We argue that individual-level analyses of labor supply may be too disaggregated because individuals are sharing substantial resources, and personal characteristics with their family members, so that individual behavioral changes may be misleading unless they are observed in the family context.

The H-P account differences can have important implications for calculating the cost of business cycles, understanding female labor supply, and modeling co-movements of aggregate wages and hours. For example, fluctuations in work time and their cross-sectional incidence are part of the cost of business cycles, and the H accounts give a different picture of that incidence than do the P accounts. Labor supply is more wage elastic in the H accounts than in the P accounts. Our Sections I-III therefore calculate the H-P differences, and Section IV offers some interpretations of them, drawing heavily on the important theoretical results of Ashenfelter, Becker, Heckman, and others.

I. Aggregate H Employment Rates

I.A House Structures in Our Data

To begin the analysis we adopt the CPS definition of household people living and eating together.[5] Figure 1's dashed line shows that most prime-aged adults (aged 25-54) are married and living

[4]Namely, taxable income is one of the main variables in the government budget constraint. See Slemrod (1998) for a survey of studies of taxable income.

[5]"A household consists of all the persons who occupy a house, an apartment, or other group of rooms, or a room, which constitutes a housing unit. A group of rooms or a single room is regarded as a housing unit when it is occupied as separate living quarters; that is, when the occupants do not live and eat with any other person in the structure, and when there is direct access from the outside or through a common hall. The count of households excludes persons living in group quarters, such as

with a spouse aged 16-65, even though that fraction has declined over time. The level and trend of this series may be relevant for studying the labor market, because it suggests that a significant fraction of the working-aged population shares resources with other working-aged people, or can share them.

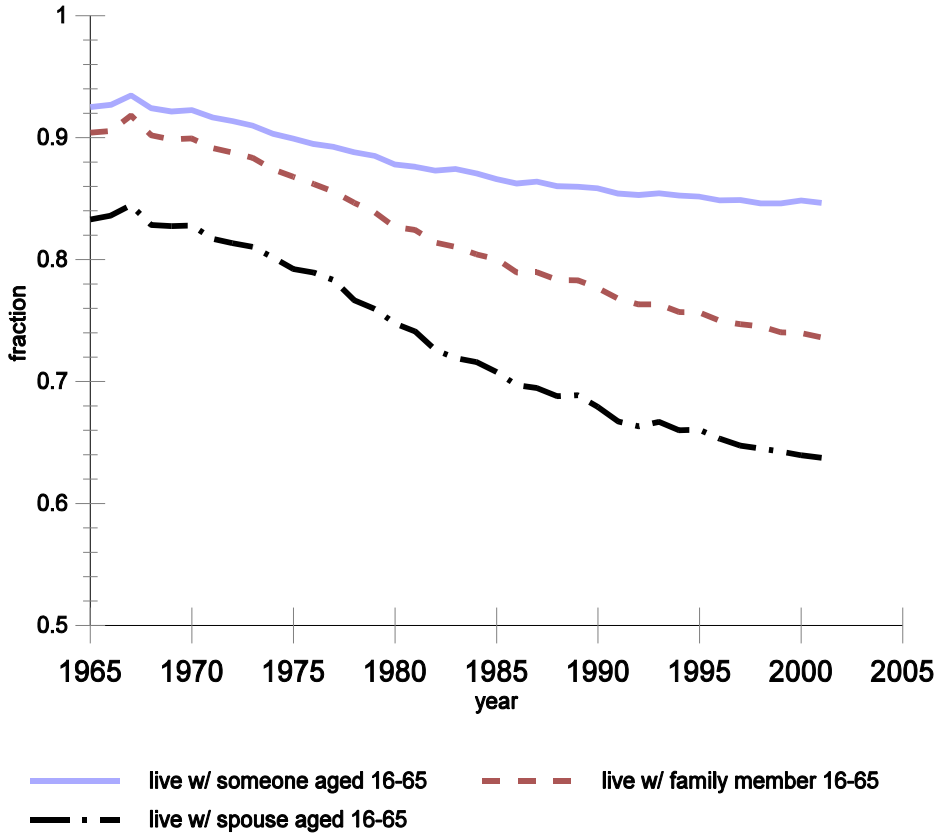

Figure 1 Household Structures for Persons aged 25-54 (CPS)

Significant resource sharing may occur perhaps not only among husbands and wives, but also among other adults living together. Figure 1's solid line shows that almost 90% of adults aged 25-54 live with another adult aged 16-65, although the other adult is sometimes not a spouse or family member.[6] Perhaps Figure 1 reveals only the obvious – that few working-aged adults live alone – but this obvious

rooming houses, military barracks, and institutions. Inmates of institutions (mental hospitals, rest homes, correctional institutions, etc.) are not included in the survey" (Census Bureau, http://www.bls.census.gov/cps/ads/1995/sglosary.htm).

[6]See also Juhn, Murphy, and Topel (1991, Table 7). A "family" is a group of two or more persons related by marriage, birth or adoption (Census Bureau, http://www.bls.census.gov/cps/ads/1995/sglosary.htm). For example, a husband and wife living with their son and his wife would be considered two families in one household.

fact is the starting point of our argument, and its departure from much of the literature, that labor market activity looks much different from the household perspective.

II.B. Household and Person Measures Compared

Let the H employment rate for year *t* be the fraction of persons aged 25-54 living in households with at least one of its members aged 16-65 employed during year *t*, and the P employment rate the fraction of persons aged 25-54 who were themselves employed. Figure 2 displays the 1967-97 H and P employment rates as solid and dash-dot lines, respectively. We see that the H and P employment rates have fluctuations that are qualitatively similar, but the P rates have larger fluctuations and a different trend.

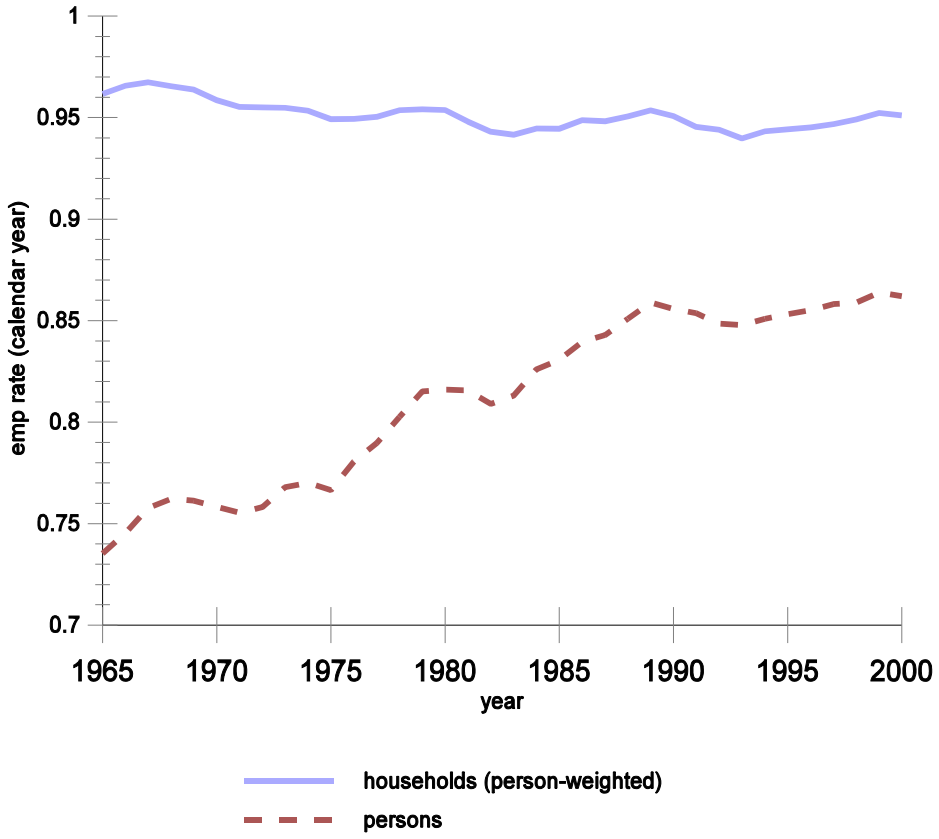

Figure 2 H and P employment rates (ages 25-54, including unmarried persons; CPS)

As we show below, it is analytically convenient to look at husbands and wives only. Let the MH employment rate for year *t* be the fraction of married couples, with husband aged 25-54, who supply positive hours during year *t*. The MH rate differs from the H rate because: (a) it is calculated in a sample

of married households only, and (b) the labor supply of adults other than head or wife is not included in the calculation. Whether the MH rate is weighted by households or persons is irrelevant because, according to the CPS definitions, there is exactly one husband for every wife and vice versa. Of course, limiting the analysis to husbands and wives has the potential disadvantage that husbands and wives are not representative of the entire sample of men and women. The purpose of Figure 3 is to study how restrictive a labor supply data set with only husbands and wives might be. Figure 3 graphs the MH and H employment rates. The two employment rates have similar level, trend, and changes.

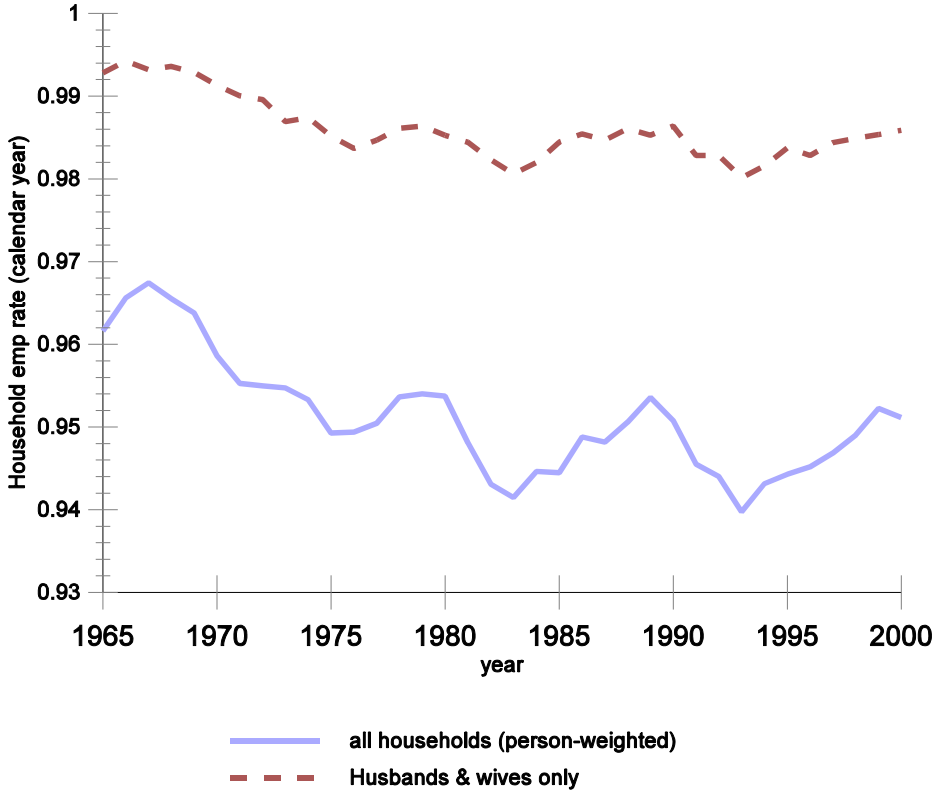

Figure 3 H and MH employment rates (CPS)

Figure 4 graphs P employment rates calculated separately for husbands and wives (since only married people are sampled, we reference the series as "MP"). We see how the husband rate is trending down, the wife rate up, and, since the wife trend is stronger, the MP rate for husbands and wives pooled trends up. Concurrently with these various trends, the MH employment rate trends slightly down (see also the additional detail in Figure 3). Also notice that the P and MP employment rates and their changes are fairly similar whether all adults aged 25-54 are included as in Figure 2, or whether the sample is restricted to husbands and wives as in Figure 4.

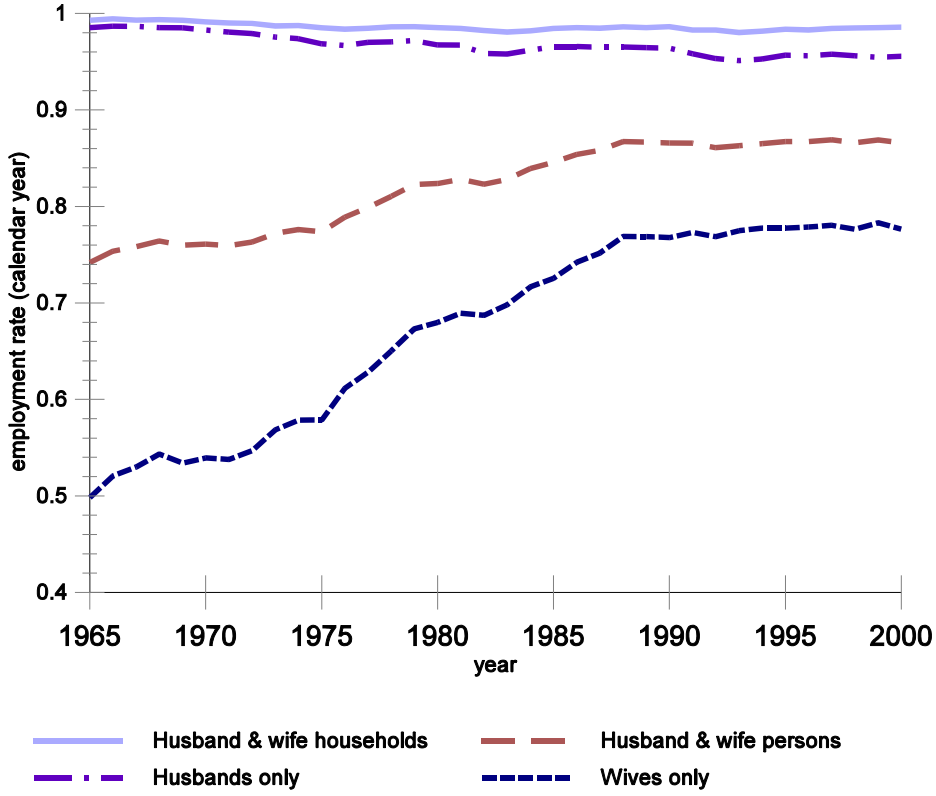

Figure 4 MH vs. MP employment rates (CPS)

The P employment rates fluctuate more than the H employment rates. Perhaps this suggests that the business cycle is more of a person phenomenon than a household phenomenon, or that there is an "added-worker effect." We consider this possibility below, and compare our calculations to those in the added-worker literature, but first we report more on the characteristics of nonemployed households.

I.C. Who Lives in a Prime-Aged NonEmployed Household?

Figure 3 shows how, in the prime-aged population, the household supplying zero labor during

a calendar year is a rarity. We believe that a sizeable majority of the few nonemployed households are retired or have permanently left the labor force because of an age-related disability. Figure 5 shows how most (65%) of the nonemployed households we observe in the PSID 1968-93 have head aged 40+, and how the household's rate transition out of nonemployment is quite low among the relatively old households. In other words, a lot of prime-aged household nonemployment is age-related and permanent labor force exit.[7]

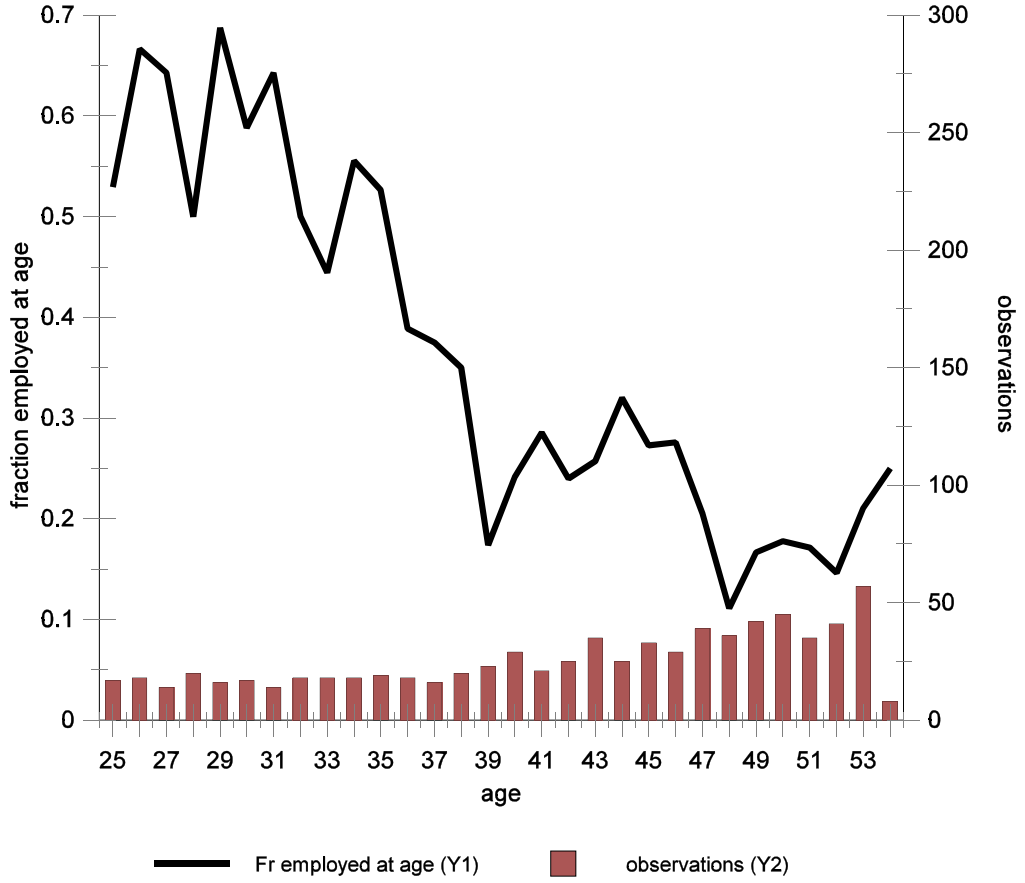

Figure 5 Transitions of Married PSID Households Not Employed at (Age-1), by Age

The CPS has larger samples in any given year, permitting a more detailed decomposition of nonemployed households, but it is harder to analyze transitions over time. We use the CPS Annual Demographic Surveys 1996-2001 to create a matched sample of prime-aged husband-wife pairs supplying zero labor in a calendar year 1995-2000. First of all, we find the MH nonemployment rate in

[7]Another simple calculation illustrates this point: the nonemployment rate of households with husband aged 25-40 is 0.7%, significantly less than the nonemployment rate of households with husband aged 41-54.

the CPS to be essentially the same as in the PSID: 98%.[8] Most (57%) of the men cite illness or disability as their reason for not working during a calendar year. This reason is cited more for men aged 40+ (61%) than for men aged less than 40 (44%).[9] Most of the women (61%) cite homemaking as their reason for not working, although illness/disability is an important reason for women aged 40+ (33% of those aged 40+). "Retirement" is the second most common reason (15%) for men, and third most common (11%) for women. Only 8% of men and 2% of women in these households report inability to find work as their reason for not working during the calendar year. Only 8% of these households, and therefore only one-tenth of one percent of all married households, supplied zero labor because both husband and wife reported inability to find work.

We also notice that 86% of nonemployed households have positive income during the year of nonemployment, and this income includes retirement income only for a very small minority. The median (mean) annual income of married households supplying zero labor during the calendar year is $13000 ($19900). Husbands are high school dropouts in most of the nonemployed households, so it is interesting to compare these nonemployed family incomes to the median (average) family income of $35000 ($42000) for married households headed by *employed* high school dropout men.

The fact that MH nonemployment is closely associated with age and disability suggests that even the nonemployed households have a work history. This is a very different characterization from that of the P accounts, where it appears that many of the prime-aged nonemployed (which are typically women) are not employed for their lifetime.[10]

I.D Person-Based Statistical Models of the Household Employment Rate

We might expect that the MH employment rate could be approximated as a function of these two MP employment rates, but which function? In order to explore this question, consider the following identity, where we denote household, personal, husband, and wife employment rates as E^{MH}, E^{P}, E^{M}, and

[8]88% of households are not employed at the time of the interview, which is in March of the year following the year of nonemployment.

[9]Since public and private disability benefits often depend on work history, the nonwork option may not become attractive for many disabled until age 40 or so. These benefits have also been most generous in the last decade or two, so our 1995-2000 CPS calculations are probably not representative of the corresponding calculations (yet to be made) for the 1970's (we thank Mark Duggan for these points).

[10]Heckman and Willis (1979) present some evidence on this point, although Mincer and Ofek (1979) take the contrary view that a lot of nonemployment by women is temporary.

E^F, respectively:

$$E_t^{MH} = E_t^M + E_t^F - E_t^M E_t^F - (1 - E_t^M) E_t^M \delta_t$$

where δ_t is the gap between the female year *t* employment rate among wives with working husbands and that among wives with nonworking husbands.[11] This identity relates *MH* employment rates to *MP* employment rates, and becomes a model of the H account E^{MH} if we specify a model for δ_t. One prototypical H account model is "random" sorting of men and women – that is, $\delta_t = 0$. A second model is perfect positive sorting – that is, $\delta_t = E_t^F/E_t^M$.[12] Without seeing the H accounts, a third and perhaps most plausible *personal* model of household employment would suppose that sorting on hours is the same as sorting on other personal characteristics. For example, δ_t is 0.34 in our sample if we take the personal characteristic to be "Having a college diploma."[13] Figure 6 displays measured E^{MH} (solid line) together with the E^{MH} predicted by these three models.

[11] δ_t is therefore the coefficient from the date *t* cross-family regression of wife employment on husband employment. Our identity could equivalently written in terms of husband-wife correlations, or husband wife-covariances, rather than the regression coefficient δ (δ is just the correlation times $\sqrt{\frac{(1 - E_t^F)E_t^F}{(1 - E_t^M)E_t^M}}$). We find δ to be a convenient means of thinking about husband-wife sorting, because it is most easily compared across characteristics – even when some of those characteristics are discrete (eg., employment) and others are continuous (eg., wages). Remember that, with discrete variables, sorting can be perfect (eg., all working women have husbands working) but correlations of the discrete variables imperfect.

[12] The formula is derived from the assumption that $E_t^F \leq E_t^M$. If $E_t^F > E_t^M$, then the appropriate formula for perfect positive sorting uses $\delta_t = (1-E_t^F)/(1-E_t^M)$.

[13] In words, the female college degree rate is 34 percentage points higher among women whose husbands have college degrees than among with whose husbands do not have the degree.

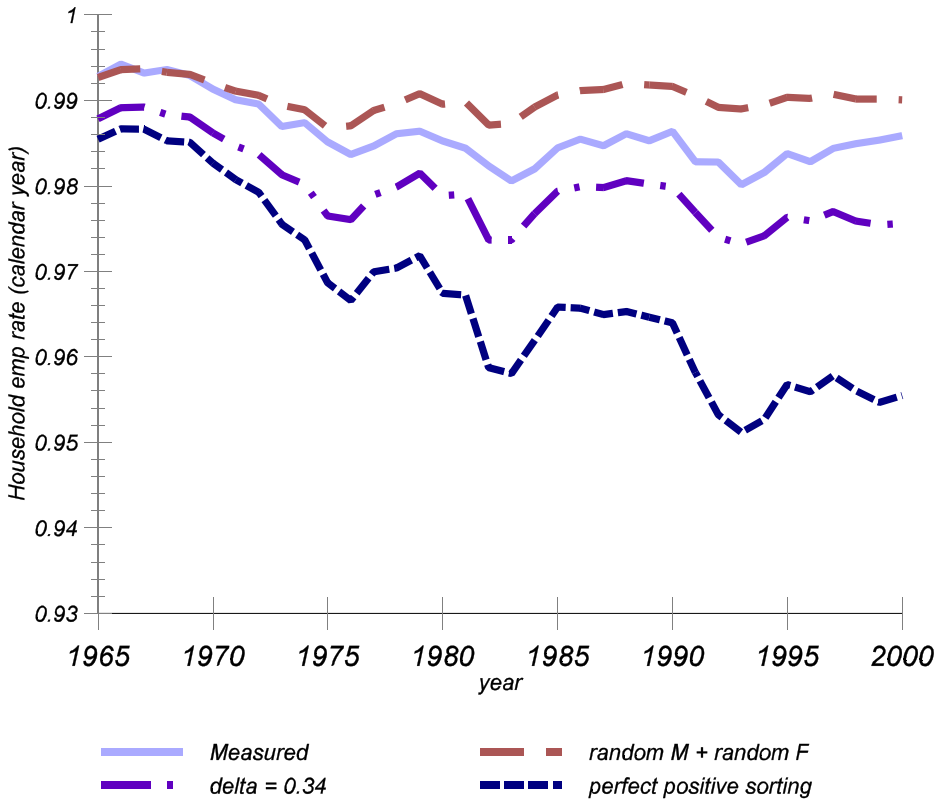

Figure 6 P-based Models of the MH Employment Rate (CPS)

Husbands and wives sort pretty strongly on so many characteristics – including age, schooling, race, wages, and age. Hence, it is notable that the *random* sorting model of family employment (the long dashed line in Figure 6) predicts actual family employment (the solid line) as well or better as the imperfect positive sorting model we calibrated from measured sorting on college graduation ($\delta_t = 0.34$, dash-dot line). Or, to put it another way, $\delta_t = 0.15$ fits the measured MH employment rate very well, which means that sorting on employment is much weaker than sorting on college graduation. This finding is not surprising from the point of view of household economics since there are at least three sets of economic influences relevant for a person's labor supply in such models, and all of them suggesting that husband-wife employment or hours correlations should be no more than correlations for other characteristics. First are the personal determinants of a person's labor supply – such as age, schooling, proximity to the job market, etc. If personal characteristics were the whole story, then employment or hours would be correlated like other characteristics. But a second economic influence is that the hours of one spouse to has a negative "income" effect on the hours of the other spouse (under the usual

assumption than nonwork time is a "normal good"), which by itself would lead to a negative husband-wife work hours correlation. Third, the second effect may be reinforced in many contexts where male and female nonwork time are substitutable. Hence, employment and hours correlations are predicted to be less than correlations of other characteristics. We return to husband-wife correlations in the hours section of our paper, and in our section on applications of the H accounts of the labor market.

II. Aggregate MH Employment and Hours

Aggregate hours are the same in the H and P accounts, and the same in the MH and MP accounts, but their decomposition into "employment" and "hours per employed" are different. Coleman's (1984) calculations for the P accounts are well cited. In his words, "in the annual data, hours vary significantly less than the number of *employees*." (p. 13, italics added). Coleman's statement can be seen in the CPS, as displayed in our Figures 7a and 7b (axes are scaled in proportion with the levels, so that a fluctuation of an inch in the figure is roughly the same proportional fluctuation for either series), although less dramatically than in his results since we use a household survey rather than an establishment survey. Still, our figures make it clear that all of the trend, and more than half of the fluctuations, in aggregate hours can be accounted for by hours per employed in the P and MP accounts.[14]

[14]Although it has fewer observations, we use the PSID for hours-employment calculations because the PSID questionnaire is designed to calculate hours worked in the year prior to the interview (whereas the CPS asks only about usual weekly hours and weeks paid in the prior year). Employment rates shown in Figures 7 and 8 (PSID) are therefore slightly different from those shown in Figures 2-4, 6 (CPS).

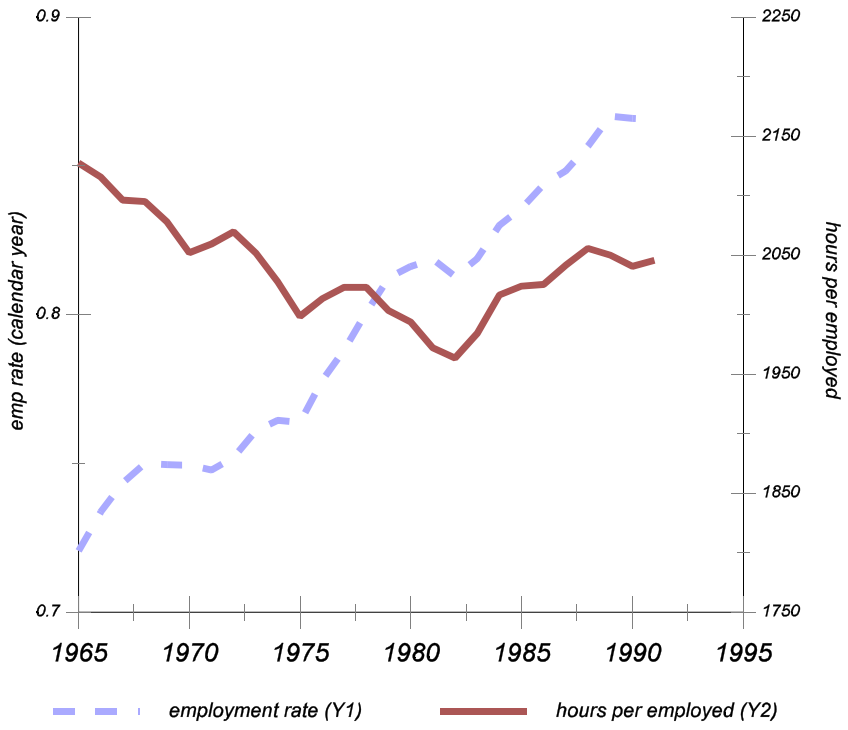

Figure 7a Emp. Rates and Hours/Employed: P-Accounts (CPS)

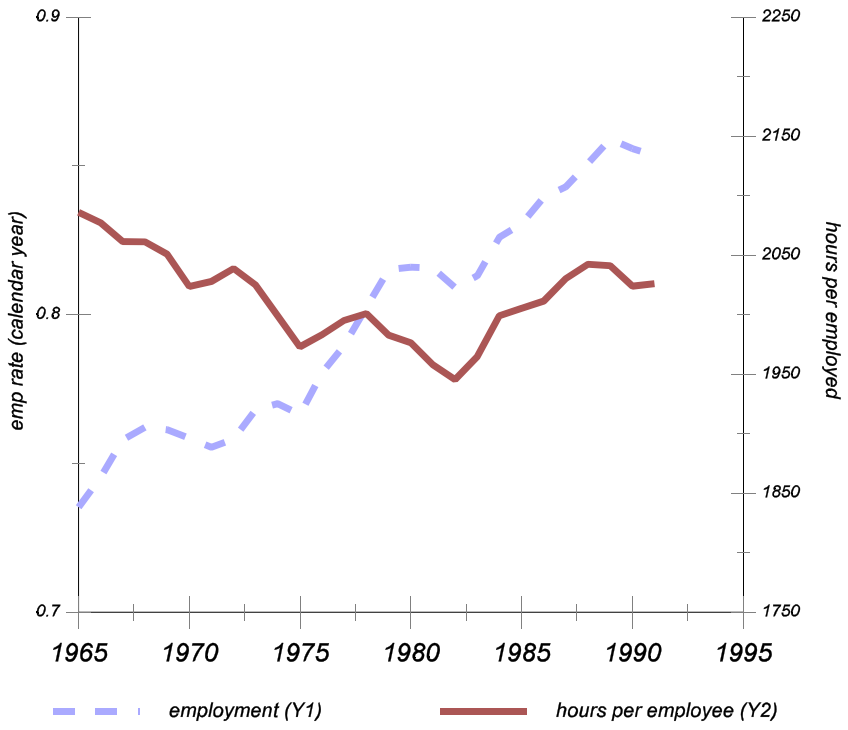

Figure 7b Emp. Rates and Hours/Employed: MP-Accounts (CPS)

The picture is very different in the H accounts. Figures 8a and 8b show how essentially *none* of the fluctuations in aggregate hours can be accounted for by the aggregate H or MH employment rates. In other words, a lot of the hours fluctuations can be attributed to couples shifting to and from the "traditional" and "dual-earner" categories – women supplying or not supplying hours in households where the man works. From a H point of view, most labor supply changes are in the number of hours supplied by the household, and very few are changes in whether or not the household supplies any hours at all.

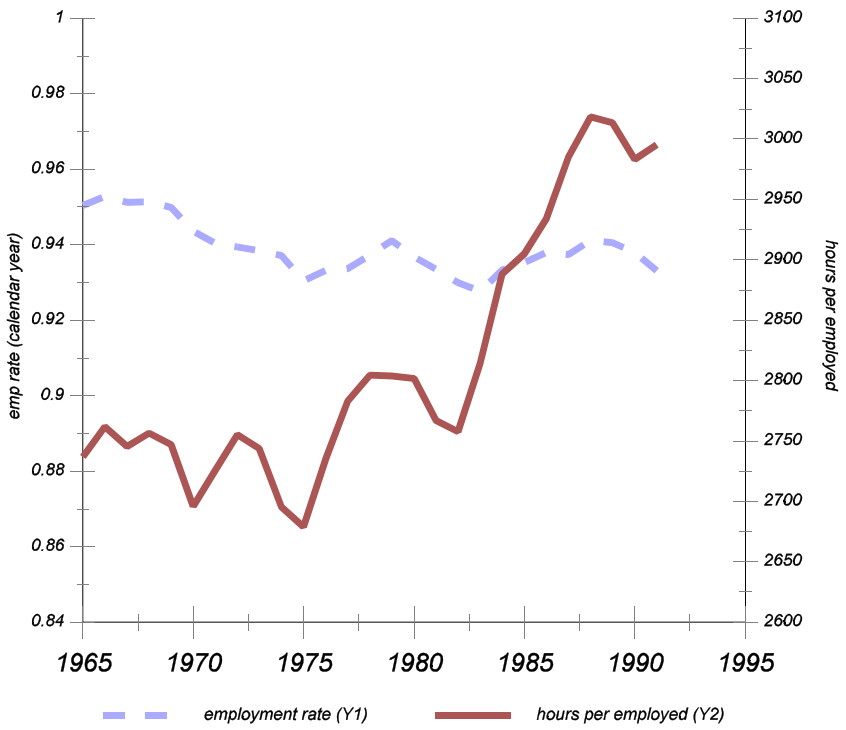

Figure 8a Emp. Rates and Hours/Employed: H-Accounts (CPS)

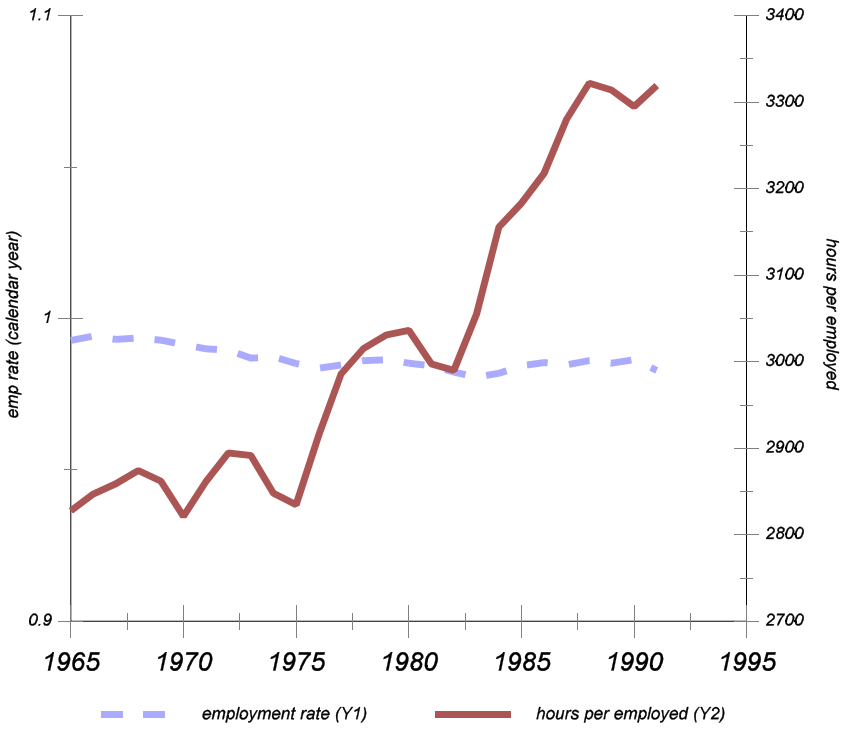

Figure 8b Emp. Rates and Hours/Employed: MH-Accounts (CPS)

III. Cross-sectional Hours Distributions in the MH accounts

Figure 9 has three panels, each with a histogram for a cross-section of hours measures. The first is husband's hours, the second wife's hours, and the third household hours. The first two show the familiar result that annual hours spike at 2000.[15] We also see the familiar result of a spike at zero, mainly for women. But the MH accounts give a very different picture (notice that the scale on the third panel is 5 times smaller than the scale on the other two panels). There are spikes at 0, 2000, and 4000, but they are less pronounced. The distribution of hours between 2000 and 4000 is fairly uniform.

Can the MH histogram be modeled from the MP histograms? Two prototypical models to consider are perfect positive sorting of spouses on hours, and random sorting on hours. In the first model, average hours for the *p*th centile of the husband distribution is summed with average hours for the *p*th centile of the wife distribution in order to calculate average hours for the *p*th centile of the MH distribution. In order to compute a histogram for the random model, we randomly pick a man from the husbands, and woman from the wives, and create a "family" observation by adding their hours. These histograms for these two models are graphed in the first two panels of Figure 10. The last panel graphs the histogram for actual families. We see that the histogram from the data is almost identical to that from the random model. One exception is the frequency of zero, a difference which we have already shown in our comparison of MH and MP employment rates (compare the solid and long dashed lines in Figure 6). A smaller exception is the higher frequency of observed families supplying 4000 hours.

IV. Conclusion: A Little Aggregation Goes a Long Way

Our paper seeks a middle ground between two empirical approaches to studying labor supply. The first is a macro approach, or a traditional labor economics approach, which considers large aggregates of, say, women by city (e.g., Mincer 1962), or even all persons in the entire country (Lucas and Rapping 1969, Prescott 2002), and thereby misses some of the important differences across persons in tastes, etc., that may be correlated with wages, income, and other variables of interest. The second approach is from modern labor economics (Heckman 1993, Blundell and MaCurdy 1999), using micro data (without aggregating), explicitly accounting for heterogeneity and the discrete choice that *seems*

[15] In our data, there are nearby spikes at 1920, 1960, 2040, and 2080, which are not discernable in the figure.

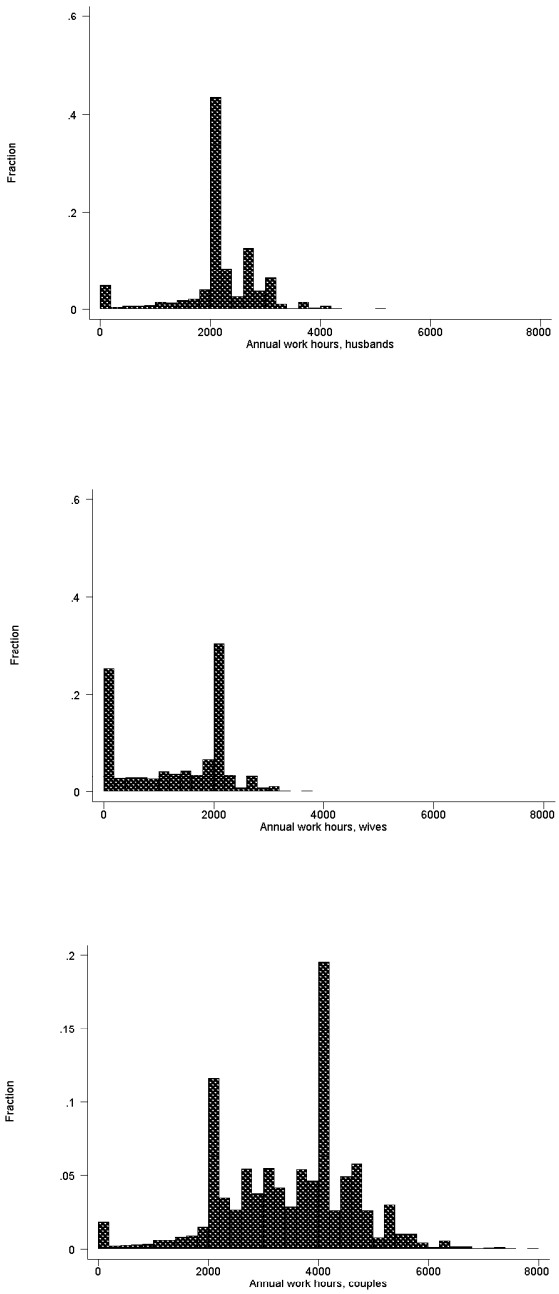

Figure 9: MP and MH Annual Hours Histograms

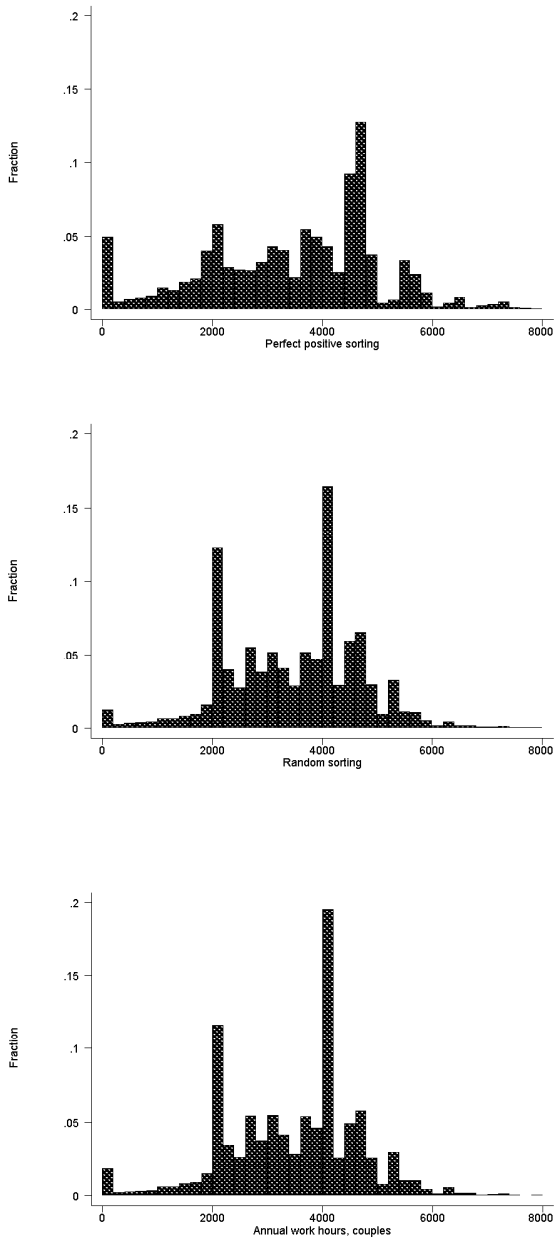

Figure 10. Simulated and Measured MH Hours Histograms

important from the microeconomic perspective, but often leaving open some of the ultimately aggregate questions that motivate labor supply study. Our proposal is to aggregate just a little: namely to average the behavior within the household. Aggregating by household has some theoretical justification, because household members are obviously sharing resources. Yet, it still permits analysis of a lot of the heterogeneity (namely, the cross-household heterogeneity) that has been justly emphasized in modern labor economics.

Calculating the H and P accounts is a step in this direction, and the main purpose of our paper. We use this concluding section to offer three illustrations of how the accounts can influence modeling the labor market. As our first illustration, we begin with Heckman's (1993, p. 116) emphasis on the important role of the "participation-hours dichotomy" in the legacy of labor supply research. In particular, it is said that hours choices are more amenable to neoclassical marginal calculations, but explain only a minority of aggregate trends and fluctuations. But our Figures 7 and 8 show that, from the household point of view, essentially all trends and fluctuations are on the intensive margin. If we view the primary household labor supply decision to be the number of hours supplied by the household, and the identity of the person or persons supplying the hours as secondary, then we can indeed model most trends and fluctuations using neoclassical marginal calculations.

Mulligan's (2001) study suggests one practical consequence of our view. Namely, while person-based analyses of employment like Ben-Porath's (1973, p. 702) predict that *personal* employment rates must in theory increase with wage rates, a household analysis *can* predict that personal employment rates decline with wages.[16] Consider, for example, a decline in wages for all blacks. This has income and substitution effects in opposite directions, but Ben-Porath's analysis predicts that, holding constant the other determinants of labor supply such as tastes and nonlabor income, the black *personal* employment rate must decline because there is no income effect on persons not working. If households are typically segregated by race, a household analysis predicts that the black *household* employment rate must decline even though the income and substitution effects are in opposite directions, because the income effect does not apply to the unemployed households. But the *personal* employment rate might increase, and thereby contradict Ben-Porath's prediction, because the nonworking persons are typically living with working persons, and thereby feel some effect of the wage decline. For example, it might be that black wage reductions lead some black men to exit employment and some black women to begin employment with, from a chauvinistic personal point of view, the latter response deriving from the adverse income

[16]Of course, a logic like Ben-Porath's implies that *household* employment rates must increase with wages.

effect of the wage reduction on black households. To put it another way, because so many adults live with another adult, and household decisions are interrelated, the comparative static analyzed by Ben Porath is of limited applicability because he assumes that the *personal distribution* of nonlabor income – where nonlabor income includes by definition money transfers from other household members – is held fixed in response to a wage change. The more interesting case from the household point of view has the personal distribution of nonlabor income responding to wage changes with, for example, workers transferring less to nonworkers in response to a wage reduction.[17]

Another practical consequence of modeling labor supply on the intensive margin by households, rather than on the extensive margin by individuals, may be a different wage elasticity of labor aggregate supply. Consider, for example, a pooled cross-section (1996-2001) regression of log *family* hours worked on log average *family* hourly earnings (namely family earnings divided by family hours), instrumenting log hourly earnings with family productivity characteristics.[18] Table 1 displays the estimated (uncompensated) wage elasticity in the first column: 0.267 (OLS s.e.= 0.005). Compare 0.267 with the wage elasticities estimated from individual level data, namely from a regression of log *individual* hours worked on log *individual* hourly earnings, instrumenting with individual productivity characteristics. The estimated wage elasticities are 0.182 for husbands, and 0.087 for wives (see specifications (4) and (6) which, unlike specifications (1)-(3), necessarily exclude zero-hours individuals). For the purposes of building models of labor market aggregates, it may well be better to use the elasticities generated by the H accounts than by the P accounts.

[17]This is essentially a cross-person version of Heckman's (1978, p. 205) cross-time-period (ie, life cycle) discussion of female labor force participation.

[18]A "family characteristic" is the average of the husband's characteristic with the wife's. Characteristics measured are years of schooling, experience, experience squared, race, region of residence, and city size (the last two are necessarily identical for husband and wife).

Table 1: The Uncompensated Wage Elasticity of Labor Supply
(pooled CPS cross-sections 1996-2001)

	husband+wife annual hours, log			individual annual hours, log			
				husband		wife	
regressor	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log family hourly wage	0.267 (0.005)	0.155 (0.005)			0.184 (0.005)		-0.008 (0.011)
log husband wage			-0.038 (0.008)	0.182 (0.004)			
log wife wage			0.315 (0.008)			0.087 (0.010)	
dual earner dummy		0.805 (0.018)					
observations	97,507	97,507	97,507	94,718	94,718	77,997	77,997

Notes: (1) standard errors in parenthesis. All regressions include year dummies.
(2) Husband hourly wages, wife hourly wages, and dual earner dummy are all predicted based on husband and wife school years completed, experience, experience squared, race and region of residence. Family wage is the average of the predicted wages for each spouse.
(3) Sample is 1996-2001 CPS husbands and wives from couples where both spouses are present and aged 25-54, at least one spouse has earnings, and both spouses report demographic characteristics.

We do not intend to say that 0.267 is the best conceivable wage elasticity estimate, because there are many possible improvements on Table 1, such as using panel data, better instruments, better hours measures, etc. Our point is that, among the many specification choices, an important one is whether to aggregate individuals into households. Aggregating seems to increase the wage elasticity estimate from roughly 0.15 to 0.27. Is this from aggregating (within family) wages on the RHS? or from aggregating hours on the LHS? It comes from the aggregation of family hours, in particular the relation between family wages and the propensity for dual earnership. We see this in specification (2) where we control for dual earnership – so that the elasticity of 0.267 can be understood as an 0.155 hours effect conditional on dual earnership plus an effect of family wage on the propensity for dual earnership.[19] In other words,

[19]As we see from the table, dual earnership is associated with 0.805 log points more hours. A probit of dual earnership on log family wage and dummies for number of children shows that one log

family labor supply is an aggregate of two distinct behaviors, each getting roughly equal weight, so it may be beneficial to study the aggregate directly rather than first decomposing the distinct behaviors for separate study (and later reaggregating the results in order to predict aggregates).

The nature of the prime-aged employment "decision" is also different from the household point of view. Heckman (1993), and many of those in the literature he surveys, characterize this decision discretely, in terms of a person's working or not working. But disability, and permanent labor force exit after a history of work, seem to be the important factors for the majority of those very few households not supplying any labor to the market. Rather than being a discrete decision, this may be a continuous decision in terms of the fraction of lifetime to be employed, or perhaps not a labor supply decision at all.[20] To the extent that the middle case is the relevant one, the H accounts thereby revive Mincer's (1962) view – a view thought to be discredited by analysis of the *personal* accounts by Heckman and Willis (1979) and others[21] – that employment (namely, H employment) can be usefully modeled as a continuous decision in terms of the fraction of lifetime to be employed. A little aggregation goes a long way.

V. References

Ashenfelter, Orley and James Heckman. "The Estimation of Income and Substitution Effects in a Model of Family Labor Supply." *Econometrica*. 42(1), January 1974: 73-86.

Becker, Gary S. "A Theory of the Allocation of Time." *Economic Journal*. 75(299), September 1965: 493-508.

Becker, Gary S. "Human Capital, Effort, and the Sexual Division of Labor." *Journal of Labor*

point higher family wage is associated with 0.163 more probability of dual earnnership. 0.805 times 0.163 is 0.131, which is about the difference between the coefficients in specifications (1) and (2).

To see the limited effect of aggregating the LHS, consider specification (3) which enters husband and wife wages separately.

[20]ie, becoming disabled may just be bad luck.

[21]Heckman and Willis (1979) estimate that a significant fraction of married *women* do not work during their married life, while we estimate that "virtually all" married *households* supply a significant number of hours to the labor market during their married life.

Economics. 3(1), Part 2, January 1985: S33-S58.

Ben-Porath, Yoram. "Labor-Force Participation Rates and the Supply of Labor." *Journal of Political Economy*. 81(3), May-June 1973: 697-704.

Blackburn, McKinley L. and David E. Bloom. "Changes in the Structure of Family Income Inequality in the United States and Other Industrial Nations during the 1980s." in Solomon W. Polachek, ed. *Research in Labor Economics*. 14 (1995): 141-70.

Blundell, Richard W and MaCurdy, T.E. "Labor Supply: A Review of Alternative Approaches." in Orley Ashenfelter and David Card, eds., Handbook of Labor Economics. Volume 3A.. New York: Elsevier Science, 1999: 1559-1695.

Bowen, William G. and T. Aldrich Finegan. *The Economics of Labor Force Participation*. Princeton, NJ: Princeton University Press, 1969.

Cancian, Maria, Sheldon Danzinger, and Peter Gottschalk. "Working Wives and Family Income Inequality among Married Couples." in Sheldon Danzinger and Peter Gottschalk, eds. *Uneven Tides: Rising Inequality in America*. New York: Russell Sage Foundation, 1993: 195-222.

Coleman, Thomas Sedgwick. "Essays on Aggregate Labor Market Business Cycle Fluctuations." PhD dissertation, University of Chicago, December 1984.

Costa, Dora L. and Matthew E. Kahn. "Power Couples: Changes in the Locational Choice of the College Educated, 1940-1990." *Quarterly Journal of Economics*. 115(4), November 2000: 1287-1315.

Daniel, Kermit. *Does Marriage Make Workers More Productive?* PhD Dissertation, University of Chicago, August 1993.

Eissa, Nada and Jeffrey B. Liebman. "Labor Supply Response to the Earned Income Tax Credit." *Quarterly Journal of Economics*. 111(2), May 1996: 605-37.

Goldin, Claudia. *Understanding the Gender Gap*. New York: Oxford University Press, 1990.

Gruber, Jonathan and Julie Cullen. "Spousal Labor Supply As Insurance: Does Unemployment Insurance Crowd Out the Added Worker Effect?" NBER wp 5608, June 1996.

Hall, Robert E. and Alvin Rabushka. *The Flat Tax*. 2nd edition. Stanford, CA: Hoover Institution Press, 1995.

Hayghe, Howard V. "Family Members in the Work Force." *Monthly Labor Review*. 113(3), March 1990; 14-19.

Heckman, James J. "A Partial Survey of Recent Research on the Labor Supply of Women." *American Economic Review*. 68(2), May 1978: 200-7.

Heckman, James J. "What has been Learned about Labor Supply in the Past Twenty Years?" *American*

Economic Review. 83(2), May 1993: 116-21.

Heckman, James J., and Thomas E. MaCurdy. "A Life Cycle Model of Female Labour Supply." *Review of Economic Studies*. January 1980, 47-74. also "Corrigendum on a Life Cycle Model of Female Labour Supply." *Review of Economic Studies*, October 1982, 659-60.

Heckman, James J. and Robert J. Willis. "The Distribution of Lifetime Labor Force Participation of Married Women: Reply to Mincer and Ofek." *Journal of Political Economy*. 87(1), February 1979: 203-11.

Hoynes, Hilary Williamson. "Welfare Transfers in Two-Parent Families: Labor Supply and Welfare Participation under AFDC-UP." *Econometrica*. 64(2), March 1996: 295-332.

Johnson, William R. and Jonathan Skinner. "Labor Supply and Marital Separation." *American Economic Review*. 76(3), June 1986: 455-69.

Juhn, Chinhui, and Kevin M. Murphy. "Wage Inequality and Family Labor Supply." *Journal of Labor Economics*. 15(1), Part 1 January 1997: 72-97.

Juhn, Chinhui, Kevin M. Murphy, and Robert H. Topel. "Why Has the Natural Rate Increased over Time?" *Brookings Papers on Economic Activity*. 2, 1991: 75-126.

Killingsworth, Mark R. *Labor Supply*. Cambridge: Cambridge University Press, 1983.

Kim, Sookon. "Cross-Substitution between Husband and Wife as One of the Factors Determining Married Woman's Labor Supply." *Journal of Economic Development*. 2(1), July 1977: 131-45.

Lucas, Robert E., Jr. and Leonard A. Rapping. "Real Wages, Employment, and Inflation." *Journal of Political Economy*. 77(5), September 1969.

Lundberg, Shelly and Robert A. Pollak. "Bargaining and Distribution in Marriage." *Journal of Economic Perspectives*. 10(4), Fall 1996: 139-58.

Mincer, Jacob. "Labor Force Participation of Married Women: A Study of Labor Supply." in H.G. Lewis, ed. *Aspects of Labor Economics*. Princeton: Princeton University Press, 1962.

Mincer, Jacob. "Labor Force Participation and Unemployment." in A. Gordon and M. Gordon, eds. *Prosperity and Unemployment*. New York: Wiley, 1966: 73-112.

Mincer, Jacob. *Schooling, Experience, and Earnings*. New York: Columbia University Press, 1974.

Mincer, Jacob and Haim Ofek. "The Distribution of Lifetime Labor Force Participation of Married Women: Comment." *Journal of Political Economy*. 87(1), February 1979: 197-201.

Mulligan, Casey B. "Aggregate Implications of Indivisible Labor." *Advances in Macroeconomics*, 1(1), Article 4, 2001.

Owen, John D. *Working Hours*. Lexington, MA: DC Heath and Company, 1979.

Prescott, Edward C. "Prosperity and Depression." *American Economic Review*. 92(2), May 2002: 1-15.

Slemrod, Joel. "Methodological Issues in Measuring and Interpreting Taxable Income Elasticities." *National Tax Journal*. 51(4), December 1998: 773-88.

Smith, James P. "Family Labor Supply over the Life Cycle." *Explorations in Economic Research*. 4(2), Spring 1977: 205-76.

Smith, James P. "The Distribution of Family Earnings." *Journal of Political Economy*. 87, October 1979: S162-192.

Smith, James P. and Michael P. Ward. "Time-Series Growth in the Female Labor Force." *Journal of Labor Economics*. 3(1), Part 2, January 1985: S59-S90.

Smith, James P. and Michael P. Ward. "Women in the Labor Market and in the Family." *Journal of Economic Perspectives*. 3(1), Winter 1989: 9-23.

Sweet, James A. *Women in the Labor Force*. New York: Seminar Press, 1973.

Tenn, Steven. "Three Essays on the Relationship Between Migration and Occupational Licensing." PhD Dissertation, University of Chicago, August 2001.

United States, Department of Commerce, Bureau of the Census. "Married-Couple Families with Wives' Earnings Greater Than Husbands' Earnings: 1981 to 1999." Historical Income Tables – Families, Table F-22, published October 2000 (http://www.census.gov/hhes/income/histinc/f22.html, accessed July 5, 2001).

Waite, Linda J. "Does Marriage Matter?" *Demography*. 32(4), November 1995: 483-507.

Winkler, Anne E. "Earnings of Husbands and Wives in Dual-Earner Families." *Monthly Labor Review*. 121(4), April 1998: 42-48.