Intrahousehold Bargaining Power and Leisure Externalities using the PSID 1968-2011

Harrison B. Wheeler
Advisor: Prof. Pierre-André Chiappori

April 20th, 2015

Abstract

Using a collective model of labor supply and the methodology of Vermeulen et al. (2006), I estimate intrahousehold bargaining power and leisure externalities from observed labor supply decisions and income in the Panel Study of Income Dynamics. This analysis provides several contributions to the collective model literature. First, I consider a large time period (1968-2011) – a scale rarely approached in this literature. Subsequently, I use this estimation to look at how intrahousehold bargaining power and leisure externalities have changed over time. I find an increasing and concave time profile to male intrahousehold bargaining power, which can be mapped to dramatic changes in the female labor market in the last 40 years. I also find that women's preference for shared leisure is higher than but converging on men's, with both preferences exhibiting decreasing and convex profiles over this same time period. Estimates of preferences over a leisure interaction are large in magnitude relative to preferences over private leisure, suggesting that models considering only private leisure are not sufficiently rich to capture the interdependence of labor choices in the household. Heterogeneity across households in intrahousehold bargaining power and preferences over shared leisure is also considered.[1]

[1] I would like to thank my advisor Pierre-André Chiappori for the countless hours of discussion. I am grateful for his help in making accessible such an important literature on intrahousehold allocations and labor supply. I would also like to thank Lena Edlund for her insights on various drafts, and my thesis class for sitting through presentations of varying quality as my thesis took shape. Finally, many thanks to Frederic Vermeulen for taking the time to discuss his 2006 paper – a work whose methodology forms the backbone of my thesis.

1 Introduction

Much of economics has viewed household decision-making through the lens of consumer choice theory, whereby the household acts as a single agent. This *unitary* household model can be seen in Samuelson's approach (1956). Here, households maximize one welfare function over the utilities of its members, subject to budget and household production constraints. The household, *in abstracto* or reinforced by a familial 'consensus,' has a common set of preferences that determine the nature of its welfare function. Also within the unitary framework, Becker's proposed household (1974) is headed by an altruistic patriarch whose utility directly depends on the consumption of all household members and who has dictatorial power over household decisions.

Common household preferences may be more of a moral principle than an economic one, however. Though the unitary model provides a clear path to estimate household behavior and preferences (hence, its continued use in the literature on household taxation), basic implications of the unitary framework have faced empirical criticism. One immediate prediction of the unitary model is that each family member's income should be pooled in the household problem. Consequently, changes in the intrahousehold distribution of income should only affect the household problem through their effect on total family income. This income-pooling hypothesis, as well as the prediction of a negative semi-definite Slutsky matrix, have been repeatedly rejected by the data.[2]

[2] The income-pooling hypothesis has been rejected in a number of papers, including Bourguignon et. al (1993), Fortin and Lacroix (1997), and Attanasio and Lechene (2002). In one particularly illustrative study, Lundberg, Pollak, and Wales (1997) saw that when a 1970s UK child benefit was changed from being given to the husband to being given to the wife (a change in the intrahousehold distribution of income, but not of total family income),

Moreover, the unitary model is unable to answer economic questions concerned with the distribution of resources within the household, such as the extent of individual-level poverty.

Chiappori (1988, 1992) expands this model of household behavior to a *collective* framework. For this model, households are composed of different members who are each endowed with their own utility function. We assume only that the household decision is Pareto efficient. This model has several implications that have fared considerably better in the literature than its unitary counterpart. Browning and Chiappori (1998) is an early example of empirical support for the collective model; in several recent works, Vermeulen (2005), Cherchye and Vermeulen (2008), and Cherchye et al. (2009), reject the symmetry and semi-definite negativity of the Slutsky matrix (as implied under a unitary framework) but fail to reject implications of the collective framework.

The collective model is a natural generalization of household behavior. It captures both the individuality of the household's members, as well as the fact that household members, being cognizant of each other's preferences, will exploit better alternatives. In this new setting, the household problem can be expressed as the maximization of the weighted sum of members' utilities, subject to budget and household production constraints. These Pareto weights correspond roughly to the relative decision power possessed by each member, as they dictate where along the Pareto frontier the household outcome will fall.[3] Consequently, these weights are functions of prices, relative wages, and

household consumption of women and children's clothing increased significantly. Browning and Chiappori (1998) and Kapan (2006) reject the slutsky matrix property for couples. See Chiappori and Mazzocco (forthcoming) for a review of empirical tests of the unitary model.

[3]The terminology 'relative decision power' is borrowed from Chiappori and Mazzocco

distribution factors. Distribution factors are variables that do not affect the household's budget constraint or preferences but affect the bargaining or decision process, as through changes in an individual's threat point. Two such examples are divorce legislation and the marriage market, where divorce is a possible threat point for a household member. It is important to note that if the Pareto weights were constant, we would have a unitary model with a welfare index over family members' utilities.

In the following, I reconstruct a proxy of household Pareto weights from observed consumption (as a Hicksian consumption aggregate) and labor supply decisions in the Panel Study of Income Dynamics (PSID). Borrowing the methodology of Vermeulen et al. (2006), I first estimate preferences over private consumption and private leisure on singles in the PSID. Making the identifying assumption that these parameters are the same for singles and for married individuals, I then proceed to construct the Pareto frontier for each household, and ultimately, find proxies for Pareto weights that correspond to observed household behavior. This analysis provides several contributions to the collective model literature. First, it adds robustness to Vermeulen et al.'s methodology by applying it to the PSID. The survey of U.S. households as well as its panel aspect provide two new features for the extension of Vermeulen et al.'s model. Second, Vermeulen et al.'s framework explicitly assumes a leisure externality between spouses. Estimating this externality reconsiders the private nature of leisure – an assumption made in some of the collective model and other economic literature. Finally, both intrahousehold bargaining power

(forthcoming).

and leisure externalities have rarely been considered over such a long time frame (more than 40 years). I find an increasing and concave time profile to male intrahousehold bargaining power, which can be mapped to dramatic changes in the female labor market in the last 40 years. I also estimate that women's preferences for shared leisure are higher than men's, with both preferences exhibiting decreasing and convex profiles over this same time period. My estimates of a leisure externality are large in magnitude relative to preferences over private leisure, suggesting that models using private leisure are not sufficiently rich to capture the interdependence of labor choices in the household.

2 Literature Review

To paint a more complete picture of the collective model literature, I introduce another notion used to estimate intrahousehold bargaining power: the sharing rule. The collective household decision is equivalent to a two-stage process, wherein income is first divided up among the family members according to a sharing rule, and then each member maximizes his or her own utility subject to their individual budget constraint. In the case of public goods, the conditional sharing rule is defined similarly with Lindahl prices for public goods in the budget constraint. The sharing rule offers the nice intuition of how much of the total household income is actually commanded by a given family member.[4]

One of the main concerns with looking at intrahousehold bargaining power, either through Pareto weights or the sharing rule, is identification. Several

[4] See Chiappori and Meghir (2014).

identification results exist in the literature on collective models. Within the restricted space of egoistic or Beckerian caring type preferences, the sharing rule is identifiable up to a constant (Chiappori, 1988). So while the researcher can consider changes in the sharing rule, he or she is unable to peg down a specific sharing rule level itself. Even so, egoistic and Beckerian caring type preferences impose the restriction that an individual's marginal rate of substitution between consumption and leisure is unaffected by other household member's leisure decisions. This restriction has faced scrutiny in recent literature on leisure externalities (discussed below).

Chiappori and Ekeland (2009) find that for more general preferences, if one can observe consumption such that for all household members there exists at least one good not consumed by a given individual, then the indirect utilities are identified. Furthermore, they show the same identification is possible if distribution factors and an *assignable* commodity (a good that is consumed solely by a given member of the household, like men's clothing for the husband in a two-person family) are available. In a setting with private and public goods, the sharing rule then is identified up to an additive function of prices. In order to identify the sharing rule, one needs either assignable goods (up to an additive function from the above result) or additional knowledge of individual preferences.

Given these identification difficulties, several strategies have emerged in the estimation of intrahousehold bargaining power.[5] The first is simply to only consider what variables affect the sharing rule (or Pareto weights), not

[5] This categorization of identification approaches is from Lewbel et al. (2012).

the actual sharing rule level itself.[6] Chiappori, Fortin, and Lacroix (2002) look at the effect of state sex ratio and divorce legislation on the sharing rule, and find that both significantly impact the intrahousehold allocation of resources. In particular, using a PSID sample of married couples, they estimate that husbands transfer on average \$2,163 to their wives upon a one-percentage-point increase in the male population and \$4,310 upon the passage of divorce legislation that is advantageous to women. Blundell, Chiappori, Meghir, and Magnac (2007) consider a collective model of labor supply including the issue of nonparticipation in the labor force. Specifically, they allow for a continuum of labor choices among women, but a discrete participation choice for men (not working versus full-time work), which coincides with the empirical distribution of U.K. labor hours. They assume that one's leisure is assignable which, in tandem with some functional form specifications and wage variation in the U.K. over this time, gives them the ability to identify the sharing rule up to a constant. Blundell et al. reject restrictions required by the unitary framework and fail to reject those required by a collective framework, while finding a significant effect of wages on the sharing rule for couples. They also estimate that in couples where the husband does not work, his consumption increases with improvements in the male labor market – an effect that could only occur in a collective setting.

Many important economic questions concerning income distributions and poverty lines, however, rely on the specific sharing rule level. In this vein, the second and most recent identification strategy aims at set-identification of the

[6] These are precisely the aforementioned distribution factors.

sharing rule through estimating bounds on its level. Cherchye, de Rock, and Vermeulen (2011) show under discrete prices and commodity bundles, and *ex ante* knowledge of whether goods are private or public, revealed preference theory can identify bounds on the sharing rule. Cherchye et al. (2012) extend this work to continuous demand functions while still relying on revealed preference theory to achieve lower and upper bounds for the sharing rule. Through these estimates, they are able to consider the sensitivity of poverty measures to household versus individual allocations. While 11% of couples in their sample have total family income below a two-person poverty line, estimation using their constructed sharing rule bounds shows that the number of individuals below the one-person poverty line is somewhere between 16% and 20%.

A third strategy makes restriction assumptions on sharing rules within households in order to pin down a specific level. Dunbar, Lewbel, and Pendakur (2013) estimate sharing rules for Malawi families by using consumption information on a private assignable good (clothing), restrictions on individual preferences, and, most importantly, the assumption that each member's share of total household consumption does not change with the level of total expenditure. The credibility of the last assumption has garnered recent empirical support (Cherchye et al. 2012; for children, Menon, Pendakur, and Perali 2012). Dunbar, Lewbel, and Pendakur find that as the number of children increases, the redistribution of resource shares falls disproportionately on the wife. Moreover, higher education attainment for the wife and more dispersion among the children's ages result in a higher share of resources for the wife. Ultimately, the authors use these sharing rule estimates to suggest that poverty measures conceal poverty among children.

The fourth strategy is to make identifying assumptions on preferences or sharing rules across households. Lise and Seitz (2011) use U.K. survey data to look at how interhousehold consumption inequality compares with intra-household consumption inequality. They perform this analysis by identifying the sharing rule for each family according to the assumption that the sharing rules are equal when the potential incomes are equal. Their estimation over U.K. data from 1968-2001 shows that while consumption inequality between households increased dramatically in the 1980s, this rise in inequality was compensated completely by a decrease in consumption inequality within households, so that consumption inequality at the individual level appeared unaffected.[7]

Vermeulen et al. (2006) identify proxies for Pareto weights through an assumption on household preferences. They assume explicitly that an individual's utility over consumption and leisure consists of a linear expenditure system (LES) part over consumption and leisure, and an interaction term between the individual's own leisure and the spouse's leisure. Under the assumption that the LES parameters do not change when one gets married, they proceed to estimate the parameters of the LES part on a sample of singles. Estimated household bargaining power and preferences over the leisure interaction are then used to consider labor supply responses to taxes under both a unitary and collective framework. Their approach, which I rely on heavily,

[7] Their work probably comes closest in the collective model literature to the time scale of my PSID sample.

will be described in more detail later.[8]

Vermeulen et al. note several precedents for this assumption in the literature. Barmby and Smith (2001) make the same assumption on preferences, while Manser and Brown (1980) make the stronger assumption that preferences are exactly the same before marriage and after. Browning, Chiappori and Lewbel (2013) make the identifying assumption that preferences over the same set of private good bundles are identical between single individuals and individuals in couples. In a recent paper, Michaud and Vermeulen (2011) assume that preferences are of a certain form upon a couple dissolving due to spousal death; namely, that preferences shift due to observable changes like mental health, happiness (as recorded in their Health and Retirement Study data set), and the widow(er)'s inability to share leisure with a spouse. They then simultaneously estimate preferences using labor and consumption decisions while in a couple and while a widow(er) through a conditional logit over discrete labor decisions. They find, among other results, that a spouse's leisure enters significantly in an individual's utility function.

The results of Michaud and Vermeulen suggest that leisure externalities are a very real part of the household's decision process. The approach of Vermeulen et al. (2006) has the flexibility to allow for such potential externalities in leisure through the inclusion of a leisure interaction term. This contrasts with a considerable part of the literature that has assumed leisure is private – in part, for convenience, and in part, for identification purposes. Exceptions include Browning and Chiappori (1998) who do not assume that individual

[8]Other examples of preference restrictions include Browning, Chiappori, and Lewbel (2013) as well as Lewbel and Pendakur (2008).

preferences are separable.

More generally though, sharing leisure time with one's spouse is a primary reason for marriage formation. Ruuskanen (2004) estimated by use of a Finnish time-use survey that spouses spent approximately 20%-25% of weekday leisure together and slightly more on weekends. In Hamermesh (2000), the observed distribution of a spouse's work timing suggests that the decision of when to work is not independent of the other spouse's decision; a result that would occur if couples coordinate market work so as to be able to share leisure together. Hallberg (2003) finds that this synchronization of leisure occurs if we consider time *physically* spent together (not just time where both spouses are not working). His estimates suggest that 2/3 of time that could be shared together by spouses was shared together. Fong and Zhang (2001) consider identification issues for estimating the component of leisure that is shared with one's spouse. In a more recent paper, Browning, Donni, and Gørtz (2012) consider a model of joint and private leisure and apply it to a Danish time-use survey. They find that wives value shared leisure more highly than their husbands, and that a couple's joint leisure increases in tandem with the wife's wage, but decreases with the husband's.

3 Methodology and Data

3.1 Model Setup

From Vermeulen et al. (2006), I assume utility functions take the following form:

$$u^i(c^i, l^i, l^j, \mathbf{z}) = \beta_c^i(\mathbf{z}) \ln(c^i) + \beta_l^i(\mathbf{z}) \ln(l^i) + \delta^i(\mathbf{z}) \ln(l^i) \ln(l^j) \tag{1}$$

In the above, i is male (m) or female (f), c^i is consumption, l^i is leisure, and $\mathbf{z}$ is a vector of household characteristics. Leisure is defined as total time endowment T net of labor hours L^i. For T, I simply take 50 possible weeks of work a year (to account for vacation time) with 17 possible daily hours of work, allowing for 7 hours of regenerative time a day (the latter specification was used in Browning, Donni, and Gørtz [2012], for instance). This is a slightly cruder measure of leisure than Vermeulen et al. (2006) use, as they consider leisure net of subsistence levels. Consumption is calculated as a Hicksian consumption aggregate $y + wL$, consisting of non labor income y, wages w, and labor L. Vermeulen et al. have consumption enter the utility function net of subsistence levels as well. The functional form used here is slightly more rigid, having individuals care about total consumption.

Including the leisure interaction term on top of the LES relaxes the assumption that leisure is private, and more generally, that preferences are separable. The parameter δ^i captures precisely how strong this externality is, and how greatly each spouse values joint leisure.[9] The interdependence of preferences

[9] I use the terms leisure interaction, shared leisure, and joint leisure interchangeably.

over leisure makes intuitive sense; as said before, we would expect that individuals who are married do so in part because they enjoy spending time together. *Ceteris paribus*, this functional form also necessarily gives more utility to married individuals. The utility function will be of the standard strictly increasing and quasi-concave form, so long as:

$$\beta_c^i(\mathbf{z}) > 0 \quad \text{and} \quad \delta^i(\mathbf{z}) > -\beta_l^i(\mathbf{z}) / \ln(l^j) \qquad\qquad i = m, f \text{ and } i \neq j$$

I consider $l^j > 1$ and positive δ^i in my estimation, and thus, the above inequality always holds. Less than 1% of my sample had an estimation of $\delta^i = 0$, suggesting that this simplification is not problematic.

The identification of the individual's preferences relies primarily on the assumption that preferences over the LES are the same before and after an individual gets married. In other words, an individual's utility function conditional on being married is the sum of that individual's utility function conditional on being single plus the leisure interaction term. This, in effect, allows the researcher to observe individual consumption before only observing aggregate household consumption. As such, so long as the utility function is sufficiently well-behaved (as the LES type preferences are), then the normal unitary model applies to the single individual, and identification is assured. The validity of this assumption on preferences should be subjected to future empirical tests, but for now, I only suggest that this is reasonable. Up to a chosen cardinalization of the utility function, this assumption pegs a specific Pareto weight

I should note that joint leisure and shared leisure are terms often used to denote leisure physically spent by spouses with each other. However, I use the terms more generally.

from the set of possible Pareto weights in a presumably meaningful way.

The household maximization, as outlined in the next section, will include discrete leisure bundles to allow for non-participation. Though there is considerable flexibility in this functional form and estimation – allowing for non-participation and leisure externalities – Vermeulen et al. note a few shortcomings of this model. Like much of the collective model literature, I am unable to incorporate household production, such as parental hours devoted to child care or housework. However, Donni (2004) suggests that this may not be too serious of an issue; the model specification being consistent with a collective model may be consistent with a richer model involving household production. Though the PSID does collect information on hours of housework, the data quality of these variables has been questioned (Achen and Stafford 2005). Moreover, modeling the endogenous choice of housework, though important, is outside of the scope of this paper. On a similar note, I am unable to distinguish between time not spent in the labor market and pure leisure.

3.2 Estimation of β_c and β_l on Singles

The estimation procedure consists of the following three stages: the estimation of preferences for singles, the estimation of the household Pareto frontier, and the estimation of the leisure interaction term. In the first stage, I estimate the LES parameters β_c^i and β_l^i over a subsample of singles constructed from the PSID. Missing wage information for the unemployed poses a serious problem. For this reason, and also to avoid "corner solution" issues arising from modeling the household's first order condition, some researchers have opted to consider households composed only of working individuals. Chiappori, Fortin, and Lacroix (2002) and Moreau and Donni (2002) are two such examples.

Considering the restrictive nature of this selection, I opt to impute wages for the non-working by the estimation of an in-sample Mincer equation. Log hourly wages are regressed separately for men and women on age, age squared, education, their interactions, year dummies, and some other demographic variables. The predicted wage values for the unemployed are then used as proxies for the wages they face in the labor market. A similar imputation is performed for couples with a non-working individual. Both estimations are included in the appendix, and their results are discussed in section 4.1.

While Vermeulen et al. employ a mixed multinomial logit model over discrete consumption and leisure bundles, I opt to use a Tobit model of labor supply. A histogram of hours worked shows significant frequency spikes at no work and full-time work, and a fairly well represented continuum of hours in between, justifying modeling labor supply for this sample as including both a discrete choice (to work or not) and a continuous choice (how many hours to work). This suggests that Tobit may be a better approach for this data set than the multinomial logit Vermeulen et al. employ. This graph is included in the appendix. Again, I estimate separately for men and women. I exploit the panel aspect of the PSID by allowing for unobserved heterogeneity through random effects. From solving individual j's labor problem, L^i is of the following form:

$$L^i_{jt} = \begin{cases} L^i_{jt^*} = \frac{\beta^i_c(\mathbf{z})T}{\beta^i_c(\mathbf{z})+\beta^i_l(\mathbf{z})} - \frac{\beta^i_l(\mathbf{z})}{\beta^i_c(\mathbf{z})+\beta^i_l(\mathbf{z})}\left(\frac{y_{jt}}{w_{jt}}\right) + \mu_j + \epsilon_{jt} & \text{if } L^i_{jt^*} > 0 \\ 0 & \text{if } L^i_{jt^*} \leq 0 \end{cases}$$

Issues with a Tobit model of labor supply include viewing employment purely as a choice (glossing over labor market frictions) and treating individuals who move between jobs as unemployed. However, the second shortcoming is somewhat mitigated by observing labor supply over an entire year. I estimate observed heterogeneity in the LES parameters through the form $\beta^i_c(\mathbf{z}) = \beta^i_{c0} + \sum_k \beta^i_{ck} z_k$ and $\beta^i_l(\mathbf{z}) = \beta^i_{l0} + \sum_k \beta^i_{lk} z_k$ for several demographic dummy variables z_k. Specifically, I estimate preferences for age-education cohorts by including five-year age group dummies, an education dummy for educational attainment equal to or below a high school degree, and their interactions. I also *ex-ante* normalize the utility function so that $\beta^i_c(\mathbf{z}) + \beta^i_l(\mathbf{z}) = 1$. This normalization allows dummies z_k to enter simply as interacted with T and y/w in my latent variable equation. Non-labor income y and wages w enter the above regression as a ratio, justifying the use of nominal values in this section. The results of this regression are included in the appendix and discussed in section 4.2.

3.3 Estimation of Intrahousehold Bargaining Power and the Leisure Interaction Term

Though nominal values sufficed in the estimation of singles, all variables that follow in this estimation are time-adjusted according to the CPI. In the next two steps, I use Vermeulen et al.'s exact specifications. I consider a sample of households consisting only of husbands and wives. First, discrete sets of possible leisure interaction terms δ^i and leisure bundles (l^m, l^f) are created. For each δ^i, the wife's dictatorial position on the Pareto frontier is constructed by giving the wife a maximal share of total consumption, and then choosing (l^m, l^f) so as to maximize her utility. As per Vermeulen et al., I let 90% denote a maximal share. The wife's and husband's dictatorial positions are solutions to the following household problem:

$$\max_{c^i, l^m, l^f} u^i(c^i, l^m, l^f, \mathbf{z}; \delta)$$

$$\text{subject to} \quad c^i \leq .9(y + L^m w^f + L^f w^f)$$

The solution to this problem gives: $u^i_{max} = u^i(c^{*^i}, l^{*^m}, l^{*f}, \mathbf{z}; \delta)$. The wife's minimum utility level is her utility at the husband's dictatorial position, and likewise for the husband's minimum utility level.

Between $u^f_{max}(\delta)$ and $u^f_{min}(\delta)$, I construct a discrete set $k \in \{0, K\}$ of female utility levels given by $\bar{u}^f_k(\delta) = u^f_{min}(\delta) + (k/K)[u^f_{max}(\delta) - u^f_{min}(\delta)]$. For each k of these utility levels $\bar{u}^f_k(\delta)$ chosen for each leisure interaction term δ, the husband solves the following problem:

$$\max_{c^m,c^f,l^m,l^f} u^m(c^m,l^m,l^f,\mathbf{z};\delta) \tag{2}$$

subject to $\qquad u^f(c^f,l^m,l^f,\mathbf{z};\delta) \geq \bar{u}^f_k(\delta) \qquad c^m+c^f \leq y+L^m w^m+L^f w^f$

The above maximization occurs over the set of possible leisure bundles (l^m, l^f). For now, I assume that $\delta = \delta^m(\mathbf{z}) = \delta^f(\mathbf{z})$ though this will be relaxed during the estimation of δ^i. Any Pareto efficient outcome (c^m, c^f, l^m, l^f) is a solution to the above problem.[10] Shifting $\bar{u}^f_k(\delta)$ from $u^f_{max}(\delta)$ to $u^f_{min}(\delta)$ completely maps the Pareto frontier.

For each δ, I choose the k_* corresponding to the bundle (l^m, l^f) that is closest to observed leisure behavior. This calculation is performed via the metric $d = [l^m(\delta,k) - l^m_0]^2 + \lambda[l^f(\delta,k) - l^f_0]^2$ where l^i_0 denotes observed leisure. I calibrate λ to be the ratio of the standard deviations of hours worked for men and women, accounting for a larger dispersion in hours worked among women. If two or more allocations minimize this metric, I take the allocation that minimizes consumption differences between the husband and wife. From this set of k_*, I again choose the δ_* that is closest to observed leisure behavior. I can then define

$$\mu_f = k_*(\delta_*)/K \tag{3}$$

as a measure of the wife's relative decision power. The same calculation can be performed with the roles reversed to get μ_m. Since μ_f and μ_m are measures of where along the Pareto frontier the household decision is, they are proxies for

[10] See Jehle and Reny, 2011.

the Pareto weights in the household maximization problem. Vermeulen et al. note that it may be more appropriate to consider the measure of bargaining power:

$$\omega^f = \mu_f^\alpha \qquad \text{where } \alpha \text{ is such that} \qquad \mu_f^\alpha + \mu_m^\alpha = 1.$$

This measure takes into account the curvature of the Pareto frontier, and by this construction, normalizes the sum of the bargaining powers to be 1. This ω^f can then be regressed on household characteristics $\mathbf{z}$ and a set of distribution factors $\mathbf{d}$ to get estimated $\omega^f(\mathbf{d}, \mathbf{z})$ and $\omega^m(\mathbf{d}, \mathbf{z})$. Vermeulen et al. suggest spousal age and education differences, as well as regional dummies, for distribution factors. The latter is included to capture cultural or marriage market differences (like the sex ratio, or divorce laws) that could affect intrahousehold bargaining power. The log of the ratio of spouses' wages will also be used as a regressor, picking up on the effect of differences in actual / potential earnings on intrahousehold bargaining power. This regression and its results are discussed in section 4.3.

In the third step, a similar algorithm produces estimates of δ^m and δ^f. As before, I first construct $u^f_{max}(\delta^m, \delta^f)$ and $u^f_{min}(\delta^m, \delta^f)$ from the discrete sets for δ^i. Again, I slice the interval $[u^f_{min}(\delta^m, \delta^f), u^f_{max}(\delta^m, \delta^f)]$ into K utility levels by the construction $u^f_k(\delta^m, \delta^f) = u^f_{min}(\delta^m, \delta^f) + (k/K)[u^f_{max}(\delta^m, \delta^f) - u^f_{min}(\delta^m, \delta^f)]$, for $k \in \{0, K\}$. Noting that $(\omega^f)^{1/\alpha}K = k_*(\delta_*)$, I pick k_*, the closest integer to $\omega^f(\mathbf{d}, \mathbf{z})^{1/\alpha}K$, and solve the program specified in (2) where the wife's utility level is $u^f_{k_*}(\delta^m, \delta^f)$. For each pair (δ^m, δ^f), the household chooses an optimal labor supply and consumption bundle. From these, I select the one that minimizes the distance criterion from observed behavior specified in the second stage of this estimation.[11] To account for potential asymmetries in the optimization process, I also run this optimization with the wife maximizing her utility; the average between the two estimates gives my final estimated values for δ^m and δ^f. The results and analysis of this estimation are considered in section 4.4.

3.4 Data and Sample Selection

The estimation is conducted on the PSID, a public data set collected by the University of Michigan.[12] Beginning in 1968, the PSID surveyed 5000 US households (of which 2000 were low-income, known as the SEO sample). In 1990, 2000 additional Latino households were added to the PSID to account for this rapidly growing demographic. However, my analysis will not consider either of the SEO or Latino subsamples. Each year (and every other year since 1997), the PSID continues to gather information on all surveyed households, their children's households, and so on. The PSID includes a large array of demographic information, as well as information concerning income, wealth, and labor supply.[13]

Labor supply is recorded as total hours worked in the year before the survey; other variables are considered similarly. However, the current year and age of the surveyed individual are not retrospective. Following Blundell, Pistaferri, and Preston (2008), I recode accordingly. Thus, the graphs here consider

[11] As per Vermeulen et al., and as was done in the estimation of the Pareto weight proxies, if there exist many values that minimize the distance criteria relative to observed leisure, then I pick the one that minimizes the difference between consumption values for the husband and wife. See Vermeuelen et al., (2006) for another description of the algorithms.

[12] Obtained via `psidonline.isr.umich.edu`

[13] For more information about the PSID, see Duffy (2013).

the time period 1967-2010. Following their work once more, age is recoded such that there are no gaps or jumps and the maximum grade achieved is an individual's education. Each individual survey-year data set is appended to one another to form one family-level data set from 1967-2010. Individuals with missing observations for education, region, or disability status, or who have top coded or negative income, are dropped. For simplicity, income and wage variables are observed pre-tax. I suspect that demographic and time trends for intrahousehold bargaining power will not be much affected by excluding taxes. In contrast, the related series of papers Vermeulen et al. (2006), Bargain et al. (2006), Beninger et al. (2006), and Myck. et al. (2006) focus on household responses to taxation (and thus, non-convex budget sets) in a unitary and collective setting. Similarly, Lise and Seitz (2011) incorporate joint taxation into their collective model.

An unfortunate reality of the PSID, however, is that information concerning the wife's race was only collected beginning in 1985. Race is an important variable both for the wage imputation for couples as well as the estimation of preferences for singles. Rather than discard 15 years of data, I have opted to discard non-white singles and non-white heads, noting that intermarriage for white men is extremely low; Rosenfeld (2005) estimates that more than 97% of white men marry white women. Moreover, one would expect this percentage to be even higher for the years when the PSID did not collect information on the wife's race. Though racial questions are also of great importance in looking at intrahousehold bargaining and leisure externalities, lack of information in this area is an unfortunate aspect of the PSID.

For the estimation of single's preferences, I use only those who report being not married (single, divorced, absent spouse, or widowed). I select individuals between the ages of 25 and 60 to avoid the liquidity-constrained young and the retirement-minded elderly. I consider singles who have no children or dependents. Though a richer model incorporating singles or families with children would be preferable, under this paper's identification strategy, I cannot include children's utility functions in the above maximization problem since I do not observe children consuming as singles.[14] One could also view children's consumption as a public good. Childcare expenditure in the PSID could be of service to this approach; however, its consideration is beyond the scope of this paper. Moreover, my basis for selection has precedents in the collective model literature (Blundell et al., 2007). This sample selection produces a population of singles that is slightly more educated, more employed, and more evenly split between the sexes (dropping single mothers); however, income, age, and other demographics remain relatively unchanged.[15] In general, this sample of singles contains more individuals from the Midwest and slightly fewer individuals from the South than the general population. Otherwise, the sample means do not suggest major deviations from American population characteristics, and capture general trends like higher educational attainment over time. However, once I select only white individuals, nominal income, employment, and education increase to levels above the general American population. Descriptive

[14]This approach can be taken of course under different identifying assumptions. For example, Dunbar, Lewbel, and Pendakur (2013) equip children with their own utility functions in their study of Malawi households.

[15]For couples, this sample selection produces similarly unaffected demographics, though I do observe an increase in household income and an increase in the husband and wife's ages. Also as expected, the wife work slightly more in households without children.

statistics of the singles sample are included in Table I of the appendix.

For the sample of couples, I consider households consisting of a married husband and wife with no children or dependents. The results of Hamermesh (2000) and Hallberg (2003) suggest that families without children have higher joint leisure, and thus, perhaps a higher leisure interaction term than families with children. Though this limits the application of my conclusions to all types of families, it provides a window into intrahousehold bargaining power when labor decisions are further intertwined, beyond their interdependence through the budget constraint. Given the manner household demographics are recorded in the PSID, more information concerning the spouse is collected when the house is "headed" by the husband. At the survey's inception, headship was assigned to husbands to comply with Census Bureau definitions and has arbitrarily remained that way since. Female heads exist when the husband is incapacitated, the wife or husband has insisted on it, or there is no husband. The first two represent exceptional cases and the third is already selected out of my sample of couples. For this reason, I keep only households headed by the husband, where the husband and wife are between ages 25 and 60. I also require no missing data for demographic information about the wife, such as age and education. Descriptive statistics of the couples sample are included in Table II of the appendix.

An obvious benefit of using the PSID is its panel aspect. The effect is two-fold. First, additional identification power is achieved through pulling my sample of singles and my sample of couples from the same data set. Given that singles marry and married couples divorce, some individuals overlap between the two samples. Second, I am able to use random effects regressions in my estimations of individual preferences, intrahousehold bargaining power, and leisure interaction terms, to account for believably significant unobserved heterogeneity.

As such, I briefly mention the way I have identified a coherent, consistent family unit across time. I apply a unique identifier to households that maintained the same head over time and have no change in family composition. The latter is not restrictive given that I only consider singles and two-person households. Moreover, I estimate separately on my singles sample and my couples sample, so changes in the household identification due to marriage do not affect my results. Individuals who are single, get married, and then re-enter the singles sample through a divorce are conservatively given a new identifier to allow for potential preference changes due to the divorce.

4 Results

4.1 Wage Imputation for Singles and Couples

To handle missing wage information for the roughly 10% of singles, 20% of married women, and 5% of married men not working, I estimate a Mincer equation for the working population, and from this, predict the wages the unemployed faced in the labor market. I regress separately for singles and couples, as well as for men and women. The regression will consider an individual's log hourly wages against their education level, age, age squared, age and age squared interactions with education, year (to capture wage inflation), region, and disability status.[16] Dummies are used for educational attainment of no high school (primarily found early in the data set), some high school, high school degree, some college, college degree, and more than college. Though this fine partition of educational attainment reduces the significance of individual coefficients, it increases the overall fit of the regression which is of greater importance in an imputation. The non-working excluded from this regression will have either reported zero hours worked or zero labor income.

An empirical concern of this methodology is the sample selection inherent in using wages of the employed to predict wages of the unemployed. This selection is generally thought to be a serious issue for women, but less so for men. To mitigate this issue, I employ a Heckman correction for the women's Mincer equation, estimated by full maximum likelihood. The selection equation is of the standard probit type, and contains education, age, region, and disability status. A similar selection equation is run for single men and married men to check for selection in this sample. As is required in a Heckman selection model, I need to observe variables that significantly affect the selection process but not the equation of interest. For both singles and couples, non-labor income serves as a good exclusion restriction. For couples, additional identification power is achieved by including in the selection equation a dummy for whether the spouse works.[17]

[16] Admittedly, using an individual's experience instead of age would be preferable, but the PSID only collects bracketed information about education for many survey-years. Even once the PSID began to ask how many years of education an individual completed, post-college education was bracketed off.

[17] Ideally, hourly wages would offer a better exclusion restriction; the idea being that the higher your spouse's wage, the less likely you are to work, but conditional on you working, your spouse's wage should not affect your own. However, this would introduce another selection as not all spouse's work.

The samples for singles and couples show similar results. The coefficient signs are generally in the direction expected; wages are increasing at a declining rate in age, the northeast has slightly higher wages, those who identify as disabled have slightly lower wages, and the coefficients on year dummies are increasing with time. The coefficient on age and age squared are similar to those estimated in Bargain, Orsini, and Peichl (2012). Time dummies show higher coefficients for women than for men, as would be expected since the wage gap between men and women has decreased over this time horizon. Education dummy variables are considered with respect to having an educational attainment beyond a bachelor's degree. Coefficients on education are generally in the expected direction; positive coefficients on lower educational attainment are largely offset by negative interactions with age or age squared. The lack of significance on some education and age variables is due to the large number of educational dummies and their interactions, used to increase the fit of the model. For couples, higher spousal educational attainment and spousal age significantly predict higher wages.

The adjusted R-squared for single men and single women (in the simple regression) are .34 and .51; for couples, these numbers are .46 and .46, supporting the reasonableness of this approach. For women, variables in the selection equation (most importantly, non-labor income) are significant. However, ρ is not significant for either single or married women, suggesting that selection is not too great a problem.[18] The selected sample of single and married women

[18] A selection model for single and married men was checked but it too produced nearly identical coefficients for the Mincer equation without selection. Though ρ is significant for married men, it is small and size and did not affect coefficients before the third non-zero digit.

without children may explain this observation. These regression results are included in the appendix.

4.2 Estimation of LES Parameters β_c and β_l

The PSID lends additional identification power to Vermeulen et al.'s methodology through its panel aspect. First, the model relies quite heavily on accurate predictions for preferences when single. With a panel, one can observe labor decisions not only when single, but also when married. Though I do not explicitly make use of this panel aspect in my estimation, the extension of estimated preferences of singles to married individuals gains further reliability since there is some overlap between my samples. Second, single men and women remain in the sample for an average of 4.4 and 5.0 years respectively. So, while I try to capture broad observable determinants of the individual's labor decision (namely, age and education), I can also account for believably strong idiosyncratic determinants: one's work ethic and productivity, etc. As such, I include random effects in my Tobit model. It is easy to see that for individuals with zero non-labor income, this random effect is simply a taste-shifter. This simple relationship between β_c, β_l, and μ_j breaks down for non-zero non-labor income, so μ_j should be considered more generally.

As described before, this regression includes a constant, the variable y/w, and interactions of these two with age cohorts, education, and age-education interactions to capture observed heterogeneity in preferences. Other determinants of observed heterogeneity were tried (including regional dummies) but age and education continued to be the most significant determinants. Both age and education were considered as dummies; age by 5-year cohorts and education as having a high school degree or less. The use of dummies instead of continuous variables is both functionally and conceptually meaningful. First, it allows for the normalization of β_c and β_l described in section 3.2. Second, age-education cohorts is an intuitive way of considering different preferences, rather than seeking additional preference variation in a continuous treatment of age or education. The estimation uses Gaussian random effects and is tabulated in the appendix.

The estimation on both women and men is significant at any reasonable level, with a corresponding Wald statistic of 430.46 and 389.52 respectively. The proportion of total variance attributable to panel-level variance, often labelled ρ, is large and significant. This test confirms that the panel estimator is more reasonable than the corresponding pooled Tobit. The lack of significance on some of the coefficients is a product of the fine partition of age-groups and the inclusion of many interactions, both done to increase fit. But for both estimations and in the majority of cases, either the coefficient on an age dummy or the coefficient on the age dummy interacted with education is significant.

As a product of both the inelastic nature of a single's labor supply, and the linear approximation to observed heterogeneity, some age-education cohorts have negative (but small, and insignificant) estimated values for β_l. For both women and men, college educated individuals aged 40-44 and high school educated individuals aged 45-49 exhibit such estimated preferences over leisure. Including a few other age-education cohorts, these individuals totaled to approximately 20% of singles. I floor these values at .001 (well below the next lowest estimated β_l). I then renormalize β_c and β_l accordingly. Note that though the lower bound on β_l homogenizes several age-education cohorts, I still observe variation in β_c and so, after renormalization, there is still variation among these cohorts in preferences over leisure. These cohorts represent the most labor inelastic, and as such, their labor choices when married will be driven primarily (and not unreasonably) by the value they place on joint leisure with their spouse.

In my estimation, β_l show averages decreasing in education for women, men, and the single's sample as a whole. β_l is decreasing on average in age until the prime working age group of 40-45, after which β_l increases in age; the latter trend corresponds to a decline in work hours as one gets older. As a brief robustness check to this estimation, I calculate naive estimates of own-wage labor elasticity for singles. First, I note that despite a few outliers driven by low wages or extreme non-labor income, the vast majority of my single sample's labor elasticities are within the range of 0 to .1. The similarity in labor decisions between single men and women was alluded to by the similar distribution over hours worked, and confirmed by negligible differences in labor elasticity. My estimates confirm the intuition that singles are quite labor inelastic. After eliminating outliers and those with zero non-labor income (who de-facto are perfectly inelastic in this model), I calculate labor elasticity averages of .03 for men and women.

Additional forces generating this observed labor inelasticity are the selection of only white individuals - a more educated and higher earning group of individuals - and a sample of individuals without children, who are also consequently younger. A recent structural analysis on Dutch data shows decreasing labor elasticities in education, quartile of income, and age for singles (Mastrogiacomo et al. 2013); their estimates being as low as .04 for singles with fourth quartile income. Bargain, Orsini, and Peichl (2012) confirm these trends more generally for Europe and the U.S., and note lower labor elasticities for individuals without children. It is not unreasonable then that all these forces working at once in my sample would generate these estimated labor elasticities.

4.3 Pareto Weight Proxies ω^i

4.3.1 Estimation

The time complexity of solving the household program is quite large – looping through each observation, each leisure interaction term δ^i, each utility level, and each bundle from the set of spouse's possible leisures.[19] As such and to simplify matters, I *ex ante* have limited the possible leisure allocations for the spouse to be near his/ her observed labor supply. After calculating the household's optimal allocation, I still *ex post* match up these optimal values with observed behavior. This simplification allows (in terms of time complexity) for a finer partition of utility levels, leisure interaction terms, and other spouse's leisure. It also reduces the distortionary effect that couples with both individuals working less than full time has on the household optimization, given the very high δ required to generate an optimal household allocation with a low number of hours worked. As evidence that the optimal allocations emerging from this program are reasonable, the average weighted squared distance

[19]This optimization is performed in MATLAB 2013.

(metric d, as defined in section 3.2) from observed labor behavior when the husband optimizes was 4337 while it was 5224 for women. If all of this error came from the husband, then taking the square root suggests that the estimated optimal allocation is, on average, different from observed behavior by approximately 1-2 labor hours on a weekly basis.

I should note that the precise level of μ^m and ω^m depends on the cardinality of the utility function. Here, this cardinality is fixed with $\beta_c + \beta_l = 1$ by the estimation on singles. Once a cardinality is chosen, changes in relative decision power μ^m and ω^m can be considered. Moreover, the results that follow will use ω^m as the relevant measure of intrahousehold bargaining power (given that it takes the Pareto frontier's convexity into account), though μ^m shows practically identical results. After eliminating a few outliers identified by extreme μ^m, μ^f, and poor fit of estimated optimal household labor supply to actual labor supply,[20] the following distribution of estimated ω^m appears (included on the next page).

The distribution of ω^m has an average of .49, and thus slightly lower than ω^f. Though the specific level depends on the chosen cardinality, the lower intrahousehold bargaining power for men is not unreasonable. The lower percentage of working wives relative to working husbands produces household outcomes with higher utility to women (relative to their minimum and maximum utility), and thus, higher bargaining power. The distribution is fairly concentrated around .5, with a standard deviation of .075, and is symmetri-

[20] As stated before, these couples were mainly those where both individuals worked well below full time. This occurred, for example, in couples that had a very high non-labor income; in all, these couples represented a small part of the sample.

Distribution of ω^m

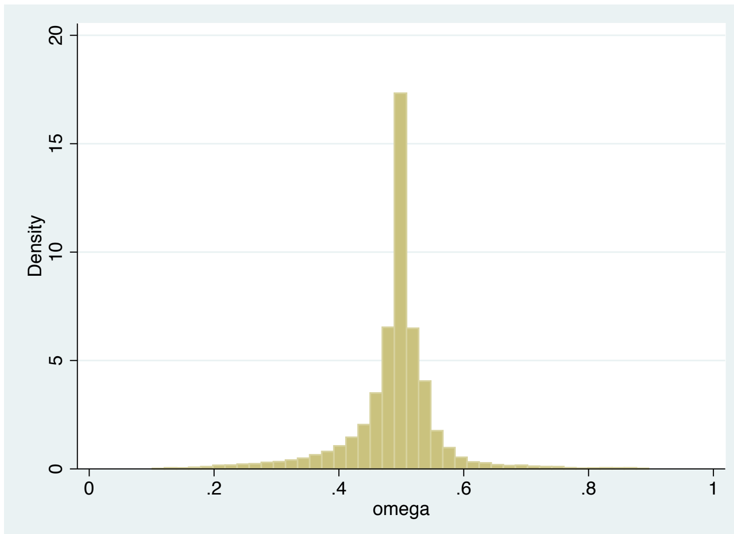

cally spread around the peak.

4.3.2 Trends and Analysis

A comparison of the mean of ω^m by year shows that male intrahousehold bargaining power has increased over time but at a decreasing rate (graph included on the next page).

This trend in intrahousehold bargaining power can be situated in a well-documented trend over this timeline: a decrease in the gender wage gap and an increase in the labor market participation rate of women. The female-to-male earnings ratio steadily increased from around .57 in 1973 to .75 by 2003 (Blau and Kahn, 2007). Decreasing gender differences in educational attainment over this period along with other factors played a role in this pay convergence. My sample of couples marks a clear increase in the wife's education, from 29% of wives in 1970 having more than a high school degree, to nearly 71% by

Mean of ω^m over Time

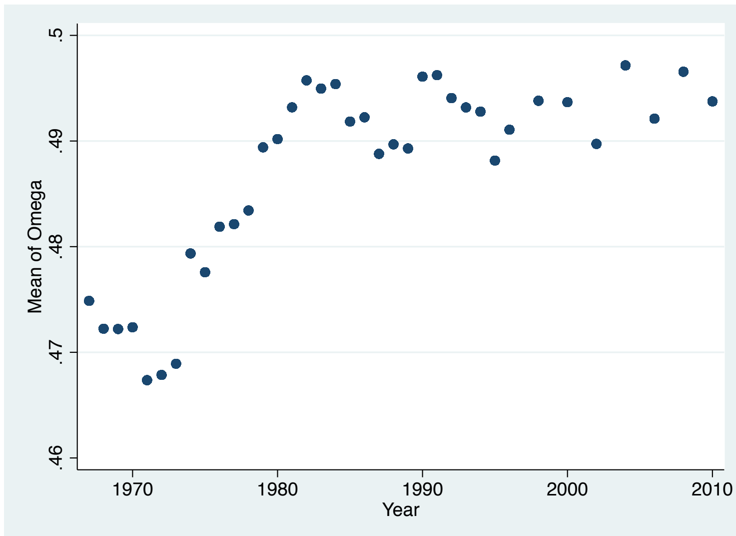

2010. The female labor force participation, especially among married women, has also increased dramatically over the last 50 years from around 40% in 1970 to nearly 75% by 2000 (Fernandez, 2013). Public access to birth control in the 1960s and the liberalization of abortion laws in the 1970s surely played a role in this meteoric rise. This trend is equally captured in my couples sample, with 58% of wives employed in 1970 increasing to 84% by 2010.

Together, these two stylized facts begin to explain the observed time trend in intrahousehold bargaining power ω^m. The substantial increase in married women labor participation began as early as 1920 and was climbing spectacularly by 1950; however, increases in the female-to-male earnings ratio did not begin to increase until the mid 1970s, and in fact, may even have decreased slightly from 1968-1973 (Blau and Kahn, 2007). Thus, the first decade or so of my sample features an increase in the average male intrahousehold bargaining power, corresponding to a decrease in the wife's leisure, and presumably, an increase in the husband's consumption from the additional family income. Considering the positive effect of relative earnings on own bargaining power, it is not surprising that the representative wife's additional consumption from working does not offset her lost utility from working more due to the low female-to-male earnings ratio. As wife participation *and* the female-to-male earnings ratio began to rise by the mid 1970s, I observe that the husband's bargaining power is still increasing (due to his consumption increasing by a fraction of the additional income brought to the family), but at a decreasing rate as the wife's relative earnings increase, and so, the wife's command of the additional income increases. The concavity could also partly capture the effect of reforms in divorce legislation. In the early 1970s, states began to pass no-fault divorce legislation; the ability of one spouse to unilaterally bring about separation surely would have benefited the threat point of non-moneyed wives. The significant effect of favorable divorce legislation on women's bargaining power has already been observed, albeit in the late 1980s, in Chiappori, Fortin, and Lacroix (2002).

Without a more rigorous attempt to identify the specific drivers of the intrahousehold bargaining power's time path, the above explanation is partly narrative. The observed trajectory of ω^m could also partly reflect the crude leisure measure used here. Massive changes in domestic technology over the early part of the sample time period meant that domestic work, largely performed by women in the 60s and 70s, became less time-consuming. Therefore, the decrease in women's leisure consumption may indicate a transfer from hours spent in domestic production to labor market. However, the time-saving benefits of domestic technology have faced scrutiny; Bittman, Rice, and Wajcman (2004) use Australian time use data to show that household appliances do not, in general, reduce women's hours spent on domestic work, and in some instances, can actually increase it. Moreover, the convincing match between ω^m and changes in the female labor market suggests that my model is capturing changes in household allocations that go beyond my coarse treatment of leisure.

Let $\mathbf{z}$ denote household characteristics of education, age, and time (measured in years since 1967) and $\mathbf{d}$ denote distribution factors of log ratio of husband's to wife's wages, age difference, education difference, and region. To quantify the relationship between intrahousehold bargaining power ω^m and time, demographics, and distribution factors, I perform the following regression:

$$\omega_{jt}^m = \frac{\exp(\beta \mathbf{z}_{jt} + \gamma \mathbf{d}_{jt} + \mu_j + \epsilon_{jt})}{1 + \exp(\beta \mathbf{z}_{jt} + \gamma \mathbf{d}_{jt} + \mu_j + \epsilon_{jt})}$$

This functional form necessitates that ω^m lies between 0 and 1 and has been employed in other literature as well (Michaud and Vermeulen 2011). Performing the transformation $g(\omega_{jt}^m) = \ln\left(\frac{\omega_{jt}^m}{1-\omega_{jt}^m}\right)$ allows for a linear regression. While both the husband and wife's age and education enter as demographics, I only include one in my regression to avoid multicollinearity with distribution factors. For education, I consider whether husbands and wives have a high school degree or less, or more than a high school degree. Dummies for the husband having a higher educational attainment and the wife having a higher educational attainment are included to allow for heterogeneity in the effect of educational differences by gender. I use random effects μ_j to capture unobserved heterogeneity in intrahousehold bargaining power. As performed throughout my work, I use random effects rather than fixed effects because I consider time-invariant variables like education. Cluster robust standard errors are also employed. The results of the random-effects GLS regression are tabulated below.[21]

$\mathbf{g}(\omega^m)$

	Coef.	S.E.
Log ratio of husband's wage to wife's	.137***	.007
Age diff.	.003***	.0008
Husband has higher education	.019	.015
Wife has higher education	-.019	.013
Time	.006***	.001
Time Sq.	-.00006*	.00002
Husband has HS degree or less	-.021*	.011
Age husb	-.009**	.003
Age Sq. husb	.0001*	.00004
Region Dummies	Yes	Not Significant
Cons	.068	.064

Wald = 430.7

Prob $> \chi^2$=0.00

$\rho = .23$ (fraction of variance attributable to μ_j)

First, both demographic variables and distribution factor age difference are significant. The high Wald indicates this significance as well. Regional dummies were not significant, suggesting that differences in cultural values

[21]* denotes significance at the 5% level, ** significance at the 1% level, and *** at the .1% level

between regions were not an important demographic determinant, while differences in sex ratio were not an important distribution factor. Though in a dynamic setting, Lise and Yamada (2014) find a similarly insignificant effect of sex ratio on Pareto weights. However, for the estimation here, the lack of significance may be the result of coarse geographic dummy variables that are unlikely able to capture sufficient variation in sex ratio. Interestingly, though higher educational attainment predicts more relative decision power ω^i, the effect is not significant. Lise and Yamada (2014), again in a dynamic model of Pareto weights, find that differences in relative education do not generate statistically different Pareto weights between the husband and wife. I include these remarks to suggest that the estimated insignificance is not unreasonable for regions or education differences.

However, the distribution factor age difference was highly significant. Moreover, the coefficient is in the direction expected, with a higher age difference (relative to the husband) predicting higher husband bargaining power. This is a rejection of stronger versions of the unitary model that require "distribution factor independence" (Browning, Chiappori, and Lechene 2006). Moreover, the unitary model, under a model specification where husbands and wives have different utility functions, requires price independence of the household utility function (Chiappori and Meghir 2014). Thus, the significance of the log ratio of wages offers another compelling rejection of the unitary model. A rejection of the unitary model was found on similar grounds in Michaud and Vermeulen (2011), with both results contributing to a now substantial and growing list of unitary model rejections.[22]

Moreover, under the same functional form for ω^m chosen here, Michaud and Vermeulen find that a unit increase in the husband's relative earning capacity (which, unlike the measure here, takes into account future earnings as well) increases his transformed ω^m by .130, close in size to the estimated .137. If all variables are evaluated at their means and the log of the ratio of spousal wages is evaluated at -1.61 and 1.61 (where the wife earns 5 times as much as the husband, and visa versa respectively), I predict a range for ω^m of .422 to .530. This is similar to Michaud and Vermeuelen's estimated range of .448 to .552. However, these numbers depend on the cardinalization of u^i, but give a frame of reference for predicted variation in intrahousehold bargaining power.

For demographics, I estimate that the husband's age has a significant negative (but at a decreasing rate with age) effect on ω^m. His education also impacts his relative decision power positively and significantly; however, the effect is small in comparison to the effect of the husband's age or the log ratio of wages. Time has a highly significant positive affect on male intrahousehold bargaining power, while time squared has a negative and significant effect on male intrahousehold bargaining power. These results quantify the relationship explored above. In 1967, the effect of an additional year on ω^m was equivalent in size to the effect of a 2 year age gap between husband and wife; this effect had been virtually halved by 2010.

Several other model specifications were considered. Time interactions with demographic and distribution factors, and years of marriage were not sig-

[22]Such rejections were mentioned briefly in the introduction.

nificant. Though the latter would be an attempt to measure the changing time profile of ω^m over the course of a marriage (and potentially, the differential effects of distribution factors over this period), my methodology is unable to account for dynamic aspects of intrahousehold bargaining power, as in a Limited-commitment Intertemporal Collective model (Chiappori and Mazzocco forthcoming). Such an empirical approach is taken in Lise and Yamada (2014).

4.4 Leisure Interaction Term δ^i

4.4.1 Estimation

The estimation of ω^m required the assumption that $\delta = \delta^m = \delta^f$. The average δ from the estimation of ω^m was approximately .1, on an order similar to the estimates of β_l. Recall that preferences over own leisure can be formalized as $\beta_l^i + \delta^i \ln l^j$. Considering that the preference term δ is magnified by the spouse's leisure, a picture emerges where labor decisions are driven strongly by preferences over shared leisure. I now use my estimates of ω^m to back out estimates of δ^m and δ^f through a similar household optimization and fit to observed labor supply outlined in section 3.3.

As before, I *ex ante* limit the search for optimal household allocations to those where the spouse's labor is close to his / her observed labor. This restriction should not be too onerous, given that I also *ex post* fit the estimated household allocations to observed labor supply; moreover, its inclusion provides significant time cost benefits. After estimation, I trim a small group of households whose observed labor supply significantly did not fit the estimated household optimizing allocation. To justify that this estimation is finding household allocations close to observed labor supply, I note that the mean weighted squared distance when the husband optimizes is 7365 and it is 6096 when the wife optimizes. If all the error came from the husband, this would be equivalent to a 1-2 hour difference in weekly labor supply. My estimation gives the following distributions for δ^m and δ^f.

Distributions of δ^m and δ^f

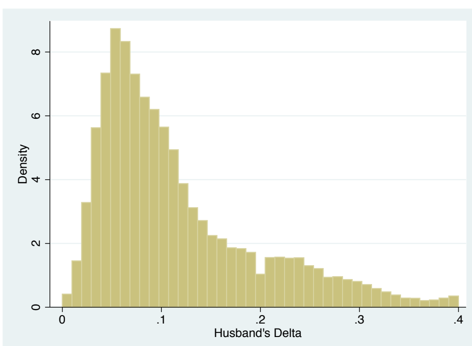

Histogram of δ^m

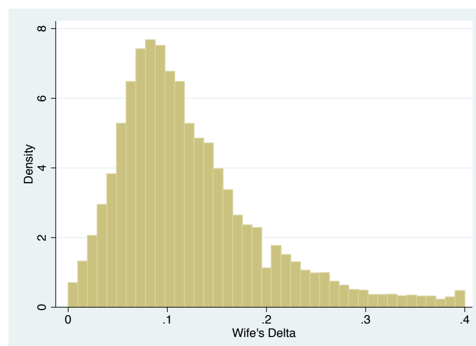

Histogram of δ^f

Both distributions show similar shapes, with a peak around .1 and a long right tail. The relatively few values at 0 and .4 suggests that these values for δ_{min} and δ_{max} were not too distortionary. It is clear, however, that in aggregate, δ^f is higher than δ^m, with their average being .123 and .115 respectively. Though the difference may seem small, the effect of δ^i (and consequently, the differences between δ^f and δ^m) on own leisure preferences is magnified by spousal leisure choice. In comparison, the average estimated β_l^m and β_l^f are .124 and .113 respectively. The observation that wives value joint leisure more highly than their husbands do, relative to preferences over private leisure, has been confirmed in Browning, Gørtz and Donni (2012). More generally, Michaud and Vermeulen (2011) find that the wife's leisure enters significantly into the husband's utility function. They consider an older sample, for whom coordination of retirement decisions may be a primary driver of complementarities in spouses' leisure. The work here suggests that the importance of leisure externalities is a more general phenomenon than can be explained by the role of joint retirement.

The elegance in Vermeulen et al.'s approach is that one deconstructs household preferences into those held when single and those held when married, delineating a specific way in which preferences change upon marriage (the addition of the leisure interaction term). The relative labor inelasticity of singles, as observed elsewhere and seen in my naive estimates as well, indicates that increased sensitivity of labor supply to wage changes must come from a mechanism other than preferences over private consumption and private leisure. Preferences over shared leisure provides such an outlet. As stated above, the mean of the estimated β_l^m and β_l^f are .124 and .113 respectively. Thus, δ^i is of a similar magnitude to preferences over private leisure. Moreover, consider the formalization of preferences over own leisure given by $\beta_l^i + \delta^i \ln l^j$. In light of the size of δ^i, preferences for own leisure appear *driven* by preferences over shared leisure (enlarged by the level of spousal leisure, $\ln l^j$) rather than preferences over private leisure β_l^i. This is in stark contrast with collective models that have assumed leisure is private. Blundell et al. (2007), for instance, require the assignability of leisure for identification of the sharing rule. Though the results should be considered with respect to my estimation methodology and sample selection, they are indicative of significant household leisure externalities that, in some of the literature on household decisions, have not been adequately accounted for.

4.4.2 Trends and Analysis

A comparison of means between δ^m and δ^f reveal the following time trends.

Mean of δ^m and δ^f over time

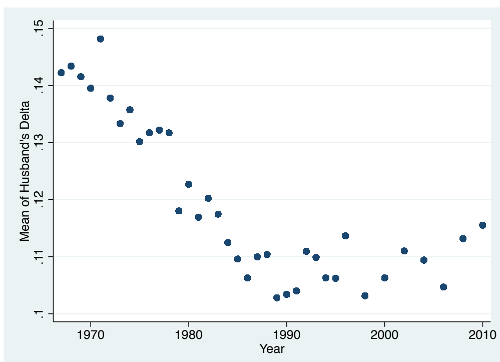

δ^m time trend

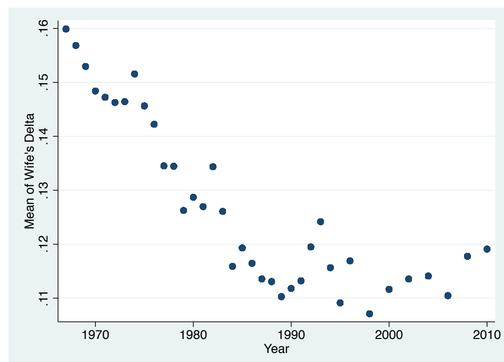

δ^f time trend

Several observations can be made immediately. First of all, both graphs show a clear decreasing, convex relationship over time. It could also be argued that both exhibit a dramatic decline from the late 60s to early 90s, after which there is virtually no trend. A second observation is that women possess a time-persistent higher average leisure interaction term. Women in the 60s have an estimated mean δ^f of around .16 while men possess an estimated mean δ^m of .14. After the mid 1980s, mean δ^m hovers in the range .1 to .11 while δ^f is more concentrated in the range .11 to .12. In fact, 1971 is the only year in which mean δ^m was higher than mean δ^f. However, while there is considerable variation in the difference between average δ^i, a third observation is that this gender difference in preferences for shared leisure is decreasing over time. Average δ^f begins .02 higher than average δ^m, but this gap decreases until 2010 when the difference is a more modest .005.

Again, this time path can be situated in larger labor market trends over the period. With female preferences over leisure formalized as $\beta_l^f + \delta^f \ln l^m$, everything else held constant, lower δ^f necessitates lower female valuation of own leisure. Thus, we would expect to see more hours worked over this time period, which is precisely what female labor participation trends have confirmed. The convergence of δ^m and δ^f also tracks well with decreasing asymmetries between female and male labor markets: particularly, in pay and participation. Labor elasticity with respect to tax rates is decreasing among married women and increasing among men over the last 30 years, lending further support to convergence in preferences over shared leisure (McCelland and Mok 2012).

As with ω^m, I now quantitatively estimate the relationship between δ^i, demographics, and time. To do so, I employ a GLS regression with random effects. Standard errors are cluster robust. As before, w^m denotes the husband's hourly wages, w^f the wife's hourly wages, and y non-labor family income.[23]

[23] As above, * denotes significance at the 5% level, ** significance at the 1% level, and *** at the .1% level. ρ denotes the fraction of variance attributable to μ_j.

Dependent Variable	δ^m		δ^f	
	Coef.	S.E.	Coef.	S.E.
w^m	.0014***	.0003	.0012***	.0003
w^f	.0004*	.0002	.0006***	.0002
y	-4.93×10^{-7}***	1.1×10^{-7}	-4.87×10^{-7}***	1.04×10^{-7}
HS deg. or less, husb.	-.0014	.0025	.0071***	.0022
HS deg. or less, wife	.0122***	.0028	.0095***	.0024
Time	-.0017***	.0003	-.0021***	.0003
Time sq.	.000032***	6.46×10^{-6}	.000035***	5.78×10^{-6}
Husb age	.0085*	.0010	.0002	.0002
Husb age sq.	-.00009***	.00001	–	–
Wife age	-.0036**	.0012	.0005*	.0002
Wife age sq.	.00005***	.00001	–	–
Midwest	-.0059*	.0027	-.0095***	.0023
Other regions	Yes	Not significant	Yes	Not significant
Constant	.0128	.0167	.1164***	.005
Wald = 406.73			Wald=411.64	
Prob > χ^2=0.00			Prob> $\chi^2 = 0.00$	
$\rho = .34$			$\rho = .30$	

In general, both regressions exhibit significant demographic effects on δ^i. The regressions for δ^m and δ^f both have high Wald statistics. Also, ρ is around 30% for both regressions justifying the use of random effects.

I have omitted the husband's and wife's squared ages from the δ^f regression due to insignificance. The wife's preferences over shared leisure display much less sensitivity to own and spousal age. Not only does the husband's age enter insignificantly in the δ^f regression, but the size of the coefficients on both the husband's age and the wife's age are dwarfed by their counterparts in the δ^m regression. In particular, husband's age has a significant, large and positive (but decreasing) effect on δ^m, while the wife's age has a significant, large, and negative (but increasing) effect on δ^m. The Midwest has a negative effect on both δ^m and δ^f and is significant, unlike the other regions.

While lower educational attainment for the wife has a significant positive effect on both δ^m and δ^f, lower educational attainment for the husband results in lower δ^m but higher δ^f (though the former is insignificant). Non-labor income has a similarly sized negative effect on δ^i between the two regressions. The husband's and wife's hourly wages were estimated to positively affect δ^m and δ^f, with the wife's wage rate affecting δ^f more than δ^m and the husband's wage rate affecting δ^m more than δ^f. Browning, Gørtz and Donni (2012) consider leisure actually spent with one's spouse (with the aid of time use data) and find that joint leisure is increasing with the wife's wage but decreasing with the husband's. Assuming that my measure of leisure interaction captures some of this time physically spent together, in contrast, I find that both spousal wage rates increase preferences for shared leisure. Finally, I also find significant time effects, with the δ^f regression producing stronger linear and squared effects. This quantifies both observations that δ^i are decreasing over time and that the wife's preferences over shared leisure are converging on the husband's.

5 Conclusion

To my knowledge, no collective model has been attempted on the time scale that my thesis has used. Moreover, Lise and Seitz (2011), whose period of consideration 1968-2001 is probably the closest to this work, analyze U.K. data, while I look at U.S. households. This work in and of itself suggests the viability of considering collective models and leisure externalities over long periods of time, to which the PSID can be invaluably useful. This work's exploration of intrahousehold bargaining power and leisure externalities over 40 years is made possible by Vermeulen et al.'s framework - one driven primarily by income and labor supply. In fact, though many estimation methods are available in the collective model literature, not all rely solely on information that has been collected consistently over such a time period.

The above results should of course be situated in the specific context of sample selection and model specification. While I am able to incorporate non-working individuals into the estimation (specifically, through a wage imputation, a Tobit estimation, and a discrete optimization that allows for non-participation), the sample only considers white singles and couples who are childless. These are significant selections that hinder the generalization of these results to a larger population. In terms of methodology, this model relies heavily on the assumption that preferences for private leisure and consumption remain unchanged whether the individual is single or married. Joint estimation under similar identifying assumptions (as used in Michaud and Vermeulen [2011]) not only would provide additional power but also offer scope for testing this preference restriction.

With these caveats in mind, these results offer a convincing depiction of intrahousehold bargaining power over time. They depict male intrahousehold bargaining power as increasing at a decreasing rate over the last 40 years, matching well with dramatic changes in the female labor market with respect to the gender wage gap and female labor participation. The increase in male intrahousehold bargaining power may seem counterintuitive in light of the "gender equality revolution" that occurred from the 70s to mid 90s. However, intrahousehold bargaining power as measured here is driven largely by labor choices. If more wives work, as the increasing labor participation rate for married women over this time period indicates, the wife's additional consumption from higher family income may not compensate for her lost leisure; this scenario seems likely given a persistent (albeit, decreasing) gender wage gap. In estimating the effect of demographics and time on intrahousehold bargaining power, the significance of the log ratio of spousal wages and spousal age difference offers another rejection of the unitary model. Moreover, I find women's preferences for joint leisure to be higher on average and through time than men's. Both female and male preferences for joint leisure exhibit a decreasing, convex time path and convergence over time; they are also large in magnitude. Heterogeneity in both intrahousehold bargaining power and leisure externalities were also estimated.

This work suggests several directions for future research. The construction of sharing rules over this time horizon based on the PSID would be useful. Then, one could also estimate the dollar value transfer between spouses over time. In general, while I have outlined how the trajectories for intrahousehold bargaining power and preferences over shared leisure match with changes in the female labor market, more work is needed to quantify this relationship. Framing these patterns in the context of broader historical and political transformations is another promising path to explore. More importantly, this work affirms what a budding literature has already proposed: namely, the large role of leisure externalities in the household decision process. Models of household decision making that treat leisure as purely private may be too simple to capture a richer interdependence of labor decisions between spouses. These results may have ramifications for how we approach public policy, including the effect of taxes on household labor supply.

References

[1] Achen, Alexandra and Frank Stafford. 2005. "Data Quality of Housework Hours in the Panel Study of Income Dynamics: Who Really Does the Dishes?" PSID Online. http://www.psidonline.isr.umich.edu/Publications/Papers/achenproxyreports04.pdf (accessed March 4, 2015).

[2] Attanasio, Orazio and Valeria Lechene. 2002. "Tests of Income Pooling in Household Decisions." *Review of Economic Dynamics* 5 (October): 720-748.

[3] Bargain, Olivier, Kristian Orsini and Andreas Peichl. 2012, "Comparing Labor Supply Elasticities in Europe and the US: New Results." IZA Discussion Papers 6735.

[4] Bargain, Olivier, Miriam Beblo, Denis Benninger, Richard Blundell, Racquel Carrasco, Maria-Concetta Chiuri, Francois Laisney, Valerie Lechene, Nicolas Moreau, Michal Myck, Javier Ruiz-Castillo, and Frederic Vermeulen. 2006. "Does the Representation of Household Behavior Matter for Welfare Analysis of Tax-benefit Policies? An Introduction." *Review of the Economics of the Household* 4: 99-111.

[5] Barmby, Timothy and Nina Smith. 2001. "Household Labour Supply in Britain and Denmark: Some Interpretation Using a Model of Pareto Optimal Behavior." *Applied Economics* 33(9): 1109-1116.

[6] Becker, Gary S. 1974. "A Theory of Social Interactions." *Journal of Political Economy* 82(6): 1063-93.

[7] Beninger, Denis, Olivier Bargain, Miriam Beblo, Richard Blundell, Raquel Car- rasco, Maria-Concetta Chiuri, Francois Laisney, Valérie Lechene, Ernesto Longob- ardi, Michal Myck, Nicolas Moreau, Javier Ruiz-Castillo, Frederic Vermeulen. 2006. "Evaluating the Move to a Linear Tax System in Germany and other European Countries." *Review of the Economics of the Household* 4: 159-180.

[8] Bittman, Michael, James M. Rice and Judy Wajcman. 2004. "Appliances and their Impact: the Ownership of Domestic Technology and Time Spent on Household Work." *The British Journal of Sociology* 55(3): 401-423.

[9] Blau, Francine D. and Lawrence M. Kahn. 2003. "The Gender Pay Gap: Have Women Gone as Far as They Can?" *Journal of Labor Economics* 21(January) : 106-144.

[10] Blundell, Richard, Luigi Pistaferri and Ian Preston, 2008, "Consumption Inequality and Partial Insurance." *American Economic Review* 98(5): 1887-1921.

[11] Blundell, Richard, Pierre-André Chiappori, Thierry Magnac and Costas Meghir. 2007. "Collective Labour Supply: Heterogeneity and Nonparticipation." *Review of Economic Studies* 74(2):417-445.

[12] Browning, Martin, Olivier Donni and Mette Gørtz. 2012. "Do You Have Time to Take a Walk Together? Private and Joint Leisure within the Household." Working Paper.

[13] Browning, Martin and Pierre-André Chiappori. 1998. "Efficient Intrahousehold Allocations: A General Characterization and Empirical Tests." *Econometrica* 66(6): 1241-1278.

[14] Browning, Martin, Pierre-André Chiappori and Arthur Lewbel. 2013. "Estimating Consumption Economies of Scales, Adult Equivalence Scales, and Household Bargaining Power." *Review of Economic Studies* 80(4): 1267-1303.

[15] Browning, Martin and Pierre-André Chiappori and Valérie Lechen. 2006. "Collective and Unitary Models: A Clarification." *Rev Econ Household* 4:5-15.

[16] Cherchye, Laurens, Bram De Rock, Arthur Lewbel and Frederic Vermeulen. 2012. "Sharing Rule Identification for General Collective Consumption Models." Center for Economic Studies - Discussion Papers ces12.05.

[17] Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen. 2009. "Opening the Black Box of Intra-household Decision-Making: Theory and Nonparametric Empirical Tests of General Collective Consumption Models." *Journal of Political Economy* 117(6): 1074-1104.

[18] Cherchye, Laurens, Bram De Rock and Frederic Vermeulen. 2011. "The Revealed Preference Approach to Collective Consumption Behavior: Testing and Sharing Rule Recovery." *Review of Economic Studies* 78(1): 176-198.

[19] Cherchye, Laurens, and Frederic Vermeulen. 2008. "Nonparametric Analysis of Household Labor Supply: Goodness of Fit and the Power of the Unitary and the Collective Model." *The Review of Economics and Statistics* 90(2): 267-274.

[20] Chiappori, Pierre-André. 1988. "Rational Household Labor Supply." *Econometrica* 56(1): 63-90.

[21] Chiappori, Pierre-André. 1992. "Collective Labor Supply and Welfare." *Journal of Political Economy* 100(3): 437-67.

[22] Chiappori, Pierre-André, Bernard Fortin and Guy Lacroix. 2002. "Marriage Market, Divorce Legislation and Household Labor Supply." *Journal of Political Economy* 110(1): 37-72.

[23] Chiappori, Pierre-André and Costas Meghir. 2014. "Intrahousehold Inequality." NBER Working Papers 20191.

[24] Chiappori, Pierre-André and Ivar Ekeland. 2009. "The Microeconomics of Efficient Group Behavior: Identification." *Econometrica* 77(3): 763-799.

[25] Chiappori, Pierre-Andre, and Maurizio Mazzocco. Forthcoming. "Static and Intertemporal Household Decisions."

[26] Donni, Olivier. 2004. "A Note on the Collective Model of Labor Supply with Domestic Prooduction." Mimeo, Cergy-Pontoise: Universite de Cergy-Pontoise.

[27] Duffy, Denise, Patricia Andreski, April Beaule, Mary Dascola, Eva Leissou, Katherine McGonagle, Jay Schlegel, and Robert Schoeni. 2013. "PSID Main Interview User Manual: Release 2013." Institute for Social Research, University of Michigan.

[28] Dunbar, Geoffrey R., Arthur Lewbel and Krishna Pendakur. 2013. "Children's Resources in Collective Households: Identification, Estimation, and an Application to Child Poverty in Malawi." *The American Economic Review*, 103(1): 438-471.

[29] Fernndez, Raquel. 2013. "Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century." *American Economic Review* 103(1): 472-500.

[30] Fong, Yuk-Fai, and Junsen Zhang. 2001. "The Identification of Unobserved Independent and Spousal Leisure." *Journal of Political Economy* 109(1): 191-202.

[31] Hallberg, Daniel. 2003. "Synchronous Leisure, Jointness and Household Labour Supply." *Labour Economics* 10(2) :185-2003.

[32] Hamermesh, Daniel S. 2000. "Togetherness: Spouses' Synchronous Leisure and the Impact of Children." NBER Working Paper 7455.

[33] Jehle, Geoffrey and Philip Reny. 2011, *Advanced Microeconomic Theory*. UK, Hampshire: Pearson Education Limited.

[34] Kapan, Tumer. 2006. "Collective Tests of Household Behavior: New Results." Mimeo, Columbia University.

[35] Lewbel, Arthur, and Krishna Pendakur. 2008. "Estimation of Collective Household Models with Engel Curves. *Journal of Econometrics* 147(2): 350358.

[36] Lewbel, Arthur, Krishna Pendakur, Geoffrey Dunbar, Martin Browning, Pierre-André Chiappori, Laurens Cherchye, Bram de Rock, and Frederic Vermeulen. 2012. "Identifying Sharing Rules in Collective Household Models: an Overview." Powerpoint presentation. Retrieved from `https://hceconomics.uchicago.edu/sites/default/files/file_uploads/lewbelslides_0.pdf`.

[37] Lise, Jeremy and Ken Yamada. 2014. "Household Sharing and Commitment: Evidence from Panel Data on Individual Expenditures and Time Use." IFS Working Papers W14/05.

[38] Lise, Jeremy and Shannon Seitz. 2011. "Consumption Inequality and Intra-household Allocations." *The Review of Economic Studies*, 78(1): 328-355.

[39] Lundberg, Shelly J., Robert A. Pollak, and Terence J. Wales. 1997. "Do Husbands and Wives Pool Their Resources? Evidence from The United Kingdom Child Benefit." *The Journal of Human Resources* 32(3): 463-480.

[40] Manser, Marilyn and Murray Brown. 1980. "Marriage and Household Decision Making: A Bargaining Analysis." *International Economic Review* 21(1): 31-44.

[41] Mastrogiacomo, Mauro, Nicole M. Bosch, Miriam D.A.C. Gielen and Egbert L.W. Jongen. 2013. "A Structural Analysis of Labour Supply Elasticities in the Netherlands." CPB Discussion Paper 235.

[42] McClelland, Robert and Shannon Mok. 2012. "A Review of Recent Research on Labor Supply Elasticities." Congressional Budget Office Working Paper.

[43] Menon, Martina, Krishna Pendakur and Federico Perali. 2012. "On the Expenditure-Dependence of Children's Resource Shares." *Economics Letters* 117(3): 739-742.

[44] Michaud, Pierre-Carl and Frederic Vermeulen. 2011. "A Collective Labor Supply Model with Complementarities in Leisure: Identification and Estimation by Means of Panel Data." *Labour Economics* 18(2): 159-167.

[45] Moreau, Nicolas, and Olivier Donni. 2002. "Estimation of a Collective Model of Labour Supply with Taxation." *Annals of Economics and Statistics* 65: 55-83.

[46] Myck, Michal, Olivier Bargain, Miriam Beblo, Denis Beninger, Richard Blundell, Raquel Carrasco, Maria-Concetta Chiuri, Franois Laisney, Valrie Lechene, Ernesto Longobardi, Nicolas Moreau, Javier Ruiz-Castillo, and Frederic Vermeulen. 2006. "The Working Families' Tax Credit and some European Tax Reforms in a Collective Setting." *Review of the Economics of the Household* 4: 129- 158.

[47] Rosenfeld, Michael, 2005, "A Critique of Exchange Theory in Mate Selection." *American Sociological Review* 70(4): 541-562.

[48] Ruuskanen, Olli-Pekka Ruuskanen. 2004. *An Econometric Analysis of Time Use in Finnish Households*. Ph.D.-thesis. Helsinki School of Economics. HeSE.

[49] Samuelson, Paul, 1956, "Social Indifference Curves." *The Quarterly Journal of Economics* 70(1): 1-22.

[50] Vermeulen, Frederic. 2005. "And the Winner is ... An Empirical Evaluation of Unitary and Collective Labour Supply Models." *Empirical Economics* 30(3): 711-734.

[51] Vermeulen, Frederic, Olivier Bargain, Miriam Beblo, Denis Beninger, Richard Blundell, Raquel Carrasco, Maria-Concetta Chiuri, Francois Laisney, Valerie Lechene, Nicolas Moreau, Michal Myck and Javier Ruiz-Castillo. 2006. "Collective Models of Labor Supply with Nonconvex Budget Sets and Nonparticipation: a Calibration Approach." *Review of Economics of the Household* 4(2): 113-127.

6 Appendix

6.1 Robustness Check

Vermeulen et al. note the possibility of using the estimates of δ^m and δ^f to re-estimate ω^m and ω^f. I perform this calculation as a robustness check, given that ω^m was originally estimated under the assumption $\delta = \delta^m = \delta^f$. Admittedly, using point estimates will reduce the flexibility of the model to shift preferences over shared leisure in order to increase fit with observed labor supply. In fact, this lack of flexibility doubled with heterogeneity in δ^i produced an estimated distribution of ω^m with significantly fatter tails. As such, I do not suggest that these estimates are a refinement on my previous ones; I only check that time trends in ω^m are preserved under this new estimation.

Mean of Re-estimated ω^m over Time

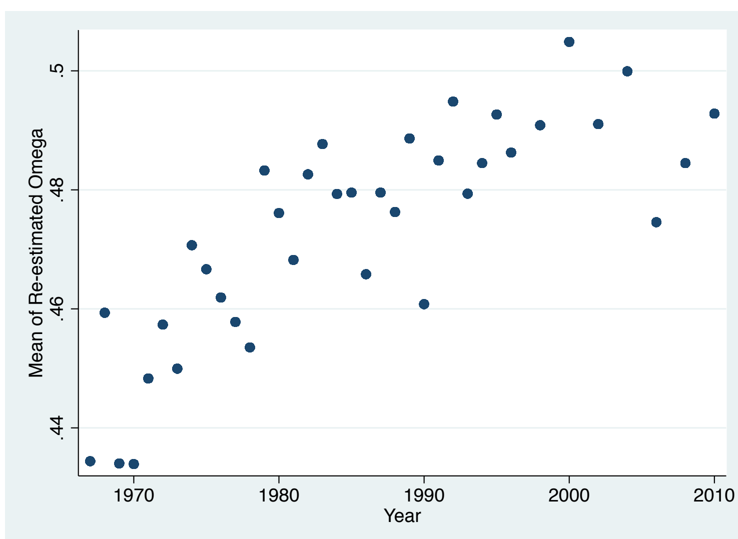

The range of mean ω^m and its variability are considerably larger than my estimates in section 4.3. However, there is still is a strong positive, concave time trend, suggesting that the main results are not overly sensitive to the optimization assumption $\delta = \delta^m = \delta^f$.

6.2 Graphs and Tables

Table I
Demographic Means in Single Sample

	Total	1970	1980	1990	2000	2010
Age	39.68	45.03	37.90	38.46	40.79	40.04
Male	0.52	0.38	0.49	0.52	0.56	0.54
No HS degree	0.12	0.32	0.12	0.10	0.08	0.09
HS degree	0.24	0.22	0.26	0.23	0.25	0.23
More than HS	0.64	0.46	0.61	0.66	0.67	0.69
Northeast	0.20	0.24	0.24	0.23	0.16	0.19
South	0.29	0.28	0.27	0.30	0.32	0.30
West	0.22	0.18	0.22	0.19	0.22	0.22
Midwest	0.29	0.30	0.27	0.27	0.30	0.29
Employed	0.91	0.91	0.94	0.92	0.93	0.86
Nominal Income	30,733	7,749	17,010	30,522	44,963	47,904

Total number of observations $= 13,387$

Table II
Demographic Means in Couple Sample

	Total	1970	1980	1990	2000	2010
Age	45.40	49.04	45.28	43.58	45.44	46.36
Age Wife	43.36	47.39	42.79	41.16	43.70	44.99
No HS degree Husb	0.14	0.34	0.20	0.11	0.09	0.07
HS degree Husb	0.28	0.34	0.28	0.27	0.28	0.25
More than HS Husb	0.58	0.32	0.53	0.62	0.63	0.68
No HS degree Wife	0.10	0.28	0.11	0.07	0.06	0.05
HS degree Wife	0.32	0.43	0.38	0.32	0.28	0.24
More than HS Wife	0.58	0.29	0.51	0.61	0.66	0.71
Northeast	0.19	0.18	0.20	0.20	0.16	0.17
South	0.31	0.31	0.26	0.31	0.32	0.31
West	0.20	0.18	0.20	0.21	0.21	0.22
Midwest	0.30	0.34	0.34	0.28	0.31	0.30
Employed Husb	0.94	0.96	0.94	0.95	0.93	0.90
Employed Wife	0.79	0.58	0.75	0.86	0.85	0.84
Nominal Income	64,951	14,010	35,797	61,524	88,819	119,577

Total number of observations $= 15,979$

Histograms for Hours Worked in Single Sample

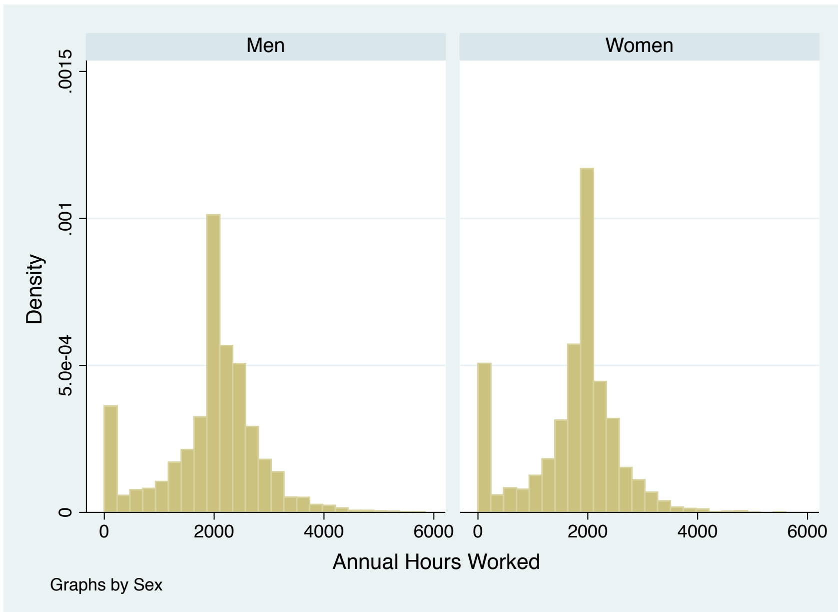

Histograms for Hours Worked in Couples Sample

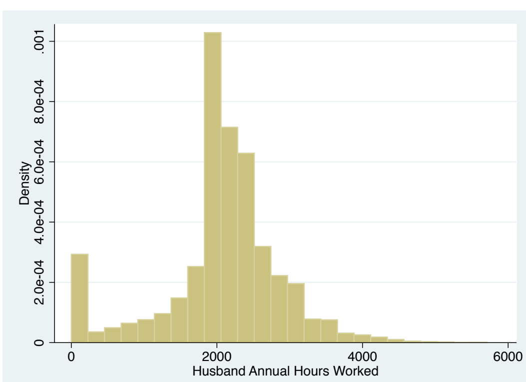

(a) Husband

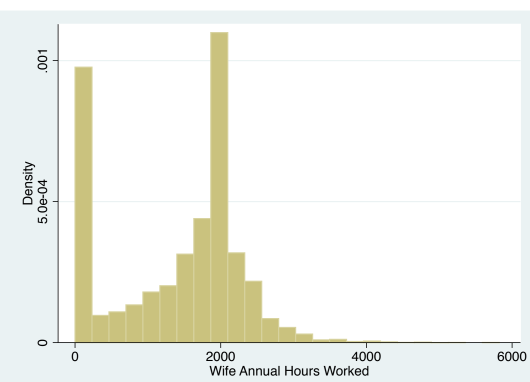

(b) Wife

Regression I: Wage Imputation for Single Women[24]

Women's Log Hourly Wages	Coef.	S.E.
No HS	2.56*	(1.19)
Some HS	1.47*	(.870)
HS	1.31**	(.405)
Some Coll	0.436	(0.396)
Coll	0.148	(0.417)
No HS x Age	-0.145**	(.053)
No HS x Age sq.	0.0015*	(.0006)
Some HS x Age	-0.110**	(.040)
Some HS x Age sq.	0.001**	(.0004)
HS x Age	-0.081***	(.021)
HS x Age sq.	0.001**	(.0002)
Some Coll x Age	-0.034	(.020)
Some Coll x Age sq.	0.0003	(.0002)
Coll x Age	0.0003	(0.022)
Coll x Age sq.	-0.0002	(.0003)
Age	0.0866***	(.015)
Age sq.	-0.001***	(.0002)
Northeast	0.084**	(.024)
Midwest	-0.086***	(.023)
South	-0.106***	(.023)
Disability	-0.190***	(.029)
Cons	-0.492	(.304)
Year Dummies	Yes	
Selection		
No HS	-0.853***	(.114)
Some HS	-0.763***	(.101)
HS	-0.349***	(.087)
Some Coll	-0.275**	(.089)
Coll	-0.109	(.107)
Age	-0.029***	(.003)
Northeast	-0.056	(.078)
Midwest	-0.0008	(.073)
South	0.078	(.074)
Disability	-1.20***	(.052)
Non labor income	-0.000021***	(1.57×10^{-6})
Cons	3.43***	(.147)
ρ	-.023	(.062)
Prob > χ^2		.00
N	6491 (non-working 717)	
Adj. R^2 (linear reg)=.51		

[24]Asterisks for given significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. χ^2 test given is for $\rho = 0$.

Regression II: Wage Imputation for Single Men[25]

Men's Log Hourly Wages

	Coef.	S.E.
No HS	-.569	(1.36)
Some HS	1.84**	(.655)
HS	1.419**	(.434)
Some Coll	0.260	(.440)
Coll	0.403	(.471)
No HS x Age	-0.029	(.062)
No HS x Age sq.	0.0004	(0.0007)
Some HS x Age	-0.122***	(.034)
Some HS x Age sq.	0.0015***	(.0004)
HS x Age	-0.084***	(.023)
HS x Age sq.	0.001**	(.0002)
Some Coll x Age	-0.019	(.023)
Some Coll x Age sq.	0.0002	(0.0003)
Coll x Age	-0.015	(.025)
Coll x Age sq.	0.00007	(0.0003)
Age	0.076***	(.016)
Age sq.	-0.0007***	(.0002)
Northeast	-0.001	(0.027)
Midwest	-0.105***	(.026)
South	-0.169***	(.025)
Disability	-0.162***	(.028)
Cons	-0.167	(.325)
Year Dummies	Yes	
N	6896 (non-working 634)	
Adj. $R^2 = .34$		

[25] Asterisks for given significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. ρ was small and insignificant when a Heckman correction was used.

Regression III: Wage Imputation for Married Women[26]

Wife's Log Hourly Wages

	Coef.	S.E.
No HS	3.85***	(1.14)
Some HS	-.038	(.708)
HS	1.14**	(.310)
Some Coll	1.68***	(.313)
Coll	1.03**	(.329)
No HS x Age	-.204***	(.052)
No HS x Age sq.	0.002***	(.0006)
Some HS x Age	-0.043	(.033)
Some HS x Age sq.	0.0006	(.0004)
HS x Age	-0.085***	(.016)
HS x Age sq.	0.001***	(.002)
Some Coll x Age	-0.105***	(.016)
Some Coll x Age sq.	0.001***	(.0002)
Coll x Age	-0.058**	(.017)
Coll x Age sq.	.001**	(.0002)
Age	0.106***	(.012)
Age sq.	-0.001***	(.0002)
Northeast	0.079***	(.019)
Midwest	-0.091***	(.017)
South	-0.076***	(.017)
Disability	-0.066***	(.018)
Cons	-0.816**	(.241)
Husband Dummies	Yes	
Year Dummies	Yes	
Selection		
No HS	-1.03***	(.081)
Some HS	-0.993***	(.061)
HS	-0.541***	(.047)
Some Coll	-0.353***	(.046)
Coll	-0.189***	(.050)
Age	-0.032***	(.003)
Northeast	.221***	(.041)
Midwest	.131***	(.035)
South	-.002	(.034)
Disability	.116**	(.035)
Non labor income	-7.66×10^{-7}*	(3.00×10^{-7})
Husband Works	.161***	(.042)
Cons	2.84***	(.090)
Husband Dummies	Yes	
ρ	.035	(.039)
Prob > χ^2		.407
N	15,979 (non-working 3352)	
Adj. R^2 (linear reg)=.46		

[26] Asterisks for given significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. χ^2 test given is for $\rho = 0$.

Regression IV: Wage Imputation for Married Men [27]

Husband's Log Hourly Wages

	Coef.	S.E.
No HS	.100	(1.10)
Some HS	1.11*	(.501)
HS	1.41***	(.317)
Some Coll	.063	(.328)
Coll	.920**	(.331)
No HS x Age	-0.027	(.046)
No HS x Age sq.	0.0002	(0.0005)
Some HS x Age	-0.081***	(.02)
Some HS x Age sq.	0.001***	(.0002)
HS x Age	-0.078***	(.016)
HS x Age sq.	0.001***	(.0002)
Some Coll x Age	-0.012	(.016)
Some Coll x Age sq.	0.0001	(0.0002)
Coll x Age	-0.051**	(.017)
Coll x Age sq.	0.0006**	(0.0002)
Age	0.097***	(.012)
Age sq.	-0.001***	(.0001)
Northeast	0.041*	(0.018)
Midwest	-0.066***	(.016)
South	-0.105***	(.016)
Disability	-0.162***	(.018)
Cons	-0.631**	(.240)
Wife Dummies	Yes	
Year Dummies	Yes	
N	15,979 (non-working 1,427)	
Adj. $R^2 = .46$		

[27] Asterisks for given significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. A Heckman correction was estimated and though ρ is significant at the 5% level for married men, it is small in size (.07) and did not affect coefficients before the third non-zero digit.

Regression V: Estimation of β_c and β_l for Women[28]

Women's Annual Labor Hours

	Coef.	S.E.
y/w	-0.045***	(.010)
y/w x HS	0.014	(.016)
y/w x age cohort 1	0.024	(.017)
y/w x age cohort 1 x HS	-0.015	(.026)
y/w x age cohort 2	0.062***	(.010)
y/w x age cohort 2 x HS	-0.043*	(.031)
y/w x age cohort 3	0.040**	(.013)
y/w x age cohort 3 x HS	-0.177**	(.062)
y/w x age cohort 4	0.045***	(.010)
y/w x age cohort 4 x HS	-0.015	(.017)
y/w x age cohort 5	-0.016	(.030)
y/w x age cohort 5 x HS	0.048	(.033)
y/w x age cohort 6	0.017	(.021)
y/w x age cohort 6 x HS	-0.065*	(.031)
age cohort 1	465.5***	(48.0)
age cohort 1 x HS	-.985	(88.7)
age cohort 2	517.4***	(49.9)
age cohort 2 x HS	45.2	(92.6)
age cohort 3	503.3***	(51.4)
age cohort 3 x HS	1.50	(95.1)
age cohort 4	411.5***	(51.1)
age cohort 4 x HS	67.3	(80.6)
age cohort 5	428.7***	(51.2)
age cohort 5 x HS	-145.5	(75.0)
age cohort 6	246.4***	(44.0)
age cohort 6 x HS	-21.88	(65.3)
HS	-321.4***	(66.8)
cons	1452.3***	(45.9)
σ_u	737.1***	(19.3)
σ_e	591.4***	(6.13)
ρ	.608	.014
N	6491	
Avg. Obs. Per Individual	5.0	
Wald $\chi^2(27)$=430.46		Prob> χ^2=.000

[28]HS denotes the dummy variable for having a high school degree or less. Age cohorts are dummy variables for 5-year groups between the ages 25 and 60.

Regression VI: Estimation of β_c and β_l for Men[29]

Men's Annual Labor Hours

	Coef.	S.E.
y/w	-0.061**	(.023)
y/w x HS	-0.096*	(.057)
y/w x age cohort 1	0.019	(.030)
y/w x age cohort 1 x HS	.018	(.062)
y/w x age cohort 2	.038	(.030)
y/w x age cohort 2 x HS	.093	(.057)
y/w x age cohort 3	.036	(.025)
y/w x age cohort 3 x HS	.144***	(.049)
y/w x age cohort 4	.094***	(.026)
y/w x age cohort 4 x HS	.055	(.055)
y/w x age cohort 5	.071*	(.034)
y/w x age cohort 5 x HS	.087	(.053)
y/w x age cohort 6	.057*	(.025)
y/w x age cohort 6 x HS	.107*	(.048)
age cohort 1	502.7***	(63.3)
age cohort 1 x HS	329.9**	(105.6)
age cohort 2	544.8***	(64.3)
age cohort 2 x HS	226.4*	(107.1)
age cohort 3	555.9***	(65.1)
age cohort 3 x HS	-12.3	(108.5)
age cohort 4	492.4***	(65.5)
age cohort 4 x HS	157.0	(108.6)
age cohort 5	482.7***	(66.6)
age cohort 5 x HS	-36.9	(109.2)
age cohort 6	207.8**	(64.2)
age cohort 6 x HS	78.9	(106.1)
HS	-353.8***	(98.8)
cons	1591.4***	(61.7)
σ_u	694.8***	(17.4)
σ_e	651.7***	(6.55)
ρ	.532	.014
N	6896	
Avg. Obs. Per Individual	4.4	
Wald $\chi^2(27)$=389.52		Prob> χ^2=.000

[29]HS denotes the dummy variable for having a high school degree or less. Age cohorts are dummy variables for 5-year groups between the ages 25 and 60.