Accounting for Changes in Labor Force Participation of Married Women: The Case of the U.S. since 1959

Bar, Michael and Leukhina, Oksana

University of North Carolina at Chapel Hill

2005

Accounting for Changes in Labor Force Participation of Married Women: The Case of the U.S. since 1959.

Michael Bar and Oksana Leukhina*

June, 2008

Using a model of family decision-making with home production and individual heterogeneity, we quantitatively investigate the role of changes in several aspects of the joint earnings distribution of husbands and wives (gender earnings gap, gender-specific inequality and assortativeness of matching) and the decline in prices of home appliances in accounting for the dramatic rise in labor force participation of married women since 1959. The implications of the factors examined are tested against changes in participation for disaggregated groups of couples and leisure trends of married individuals, documented from the U.S. population census and time-use survey data.

Keywords: labor force participation, married couples, family time allocation, gender earnings gap, home production

JEL Codes: J2, J3, E0

*Michael Bar: Department of Economics, San Francisco State University, San Francisco, CA; Oksana Leukhina (*CORRESPONDING AUTHOR)*: Department of Economics, 107 Gardner Hall, CB 3305, University of North Carolina, Chapel Hill, NC 27599-3305. *tel.* (919) 259-4428, *fax* (919) 966-4986. *ACKNOWLEDGEMENTS:* We are grateful to Larry Jones, V.V. Chari and Michele Boldrin for their advice and support. We also thank the participants of the IZA-EUI Workshop on secular trends in labor markets (September, 2006) and the Stanford Institute for Theoretical Economics Workshop on family behavior and aggregate economy (August 2007) for their useful feedback.

I. Introduction

The main observation investigated in this paper is the dramatic increase in the proportion of two-earner households,[1] from 33% to 76%, during the period 1959-1999 (Figure 1).[2] This change was due mainly to a large number of married females joining the workforce. In fact, married females' labor force participation (LFP) increased by approximately 130%, while married males' participation remained roughly constant. Figure 2 reveals that the average number of annual hours worked, conditional on working a positive number of hours, has remained roughly constant for males and only slightly increased for females. Hence, the overall trend in the labor supply of married couples was driven by the rise in the labor supply of married women, which transpired primarily at the *extensive* (participation) rather than the *intensive* (hours conditional on participation) margin.

Goldin (1990) and Costa (2000) provide a comprehensive documentation of historical trends in female labor supply and proposed explanations. Most commonly cited among the economic explanations of the rise in female labor supply are the home production revolution and factors that tend to close the gender wage gap. The former refers to the widespread diffusion of electrical appliances, such as washing machines, dishwashers and vacuum cleaners, often assumed to be a result of falling prices of home appliances. Jones et al. (2003) find that a small reduction in the gender wage differential, modeled as a discrimination tax on female income, can account for the entire observed increase in the labor supply of married females during the period 1950-2000, while the decline in prices of home appliances is much less quantitatively important. By contrast, Greenwood et al. (2005) focus on the period 1900-1990, finding the decline in the relative price level of home appliances to be the main driving force underlying the rise in female LFP.

One objective of this paper is to extend the test of these competing explanations to several important empirical observations documented in this paper, in particular, changes in participation for disaggregated groups of married women and gender-specific trends in leisure, i.e., time not spent on market or home production. Moreover, instead of focusing on the gender earnings gap as the only aspect of the joint earnings distribution of husbands and wives, we break down the earnings distribution into several components[3] (gender earnings gap, gender-specific inequality and assortativeness of matching) and investigate each in isolation. The idea that within-gender inequality may be an important determinant

[1] We regard a person as a labor force participant (interchangeably used with "an earner" and "a worker") if he/she works a positive number of hours in a given year. We emphasize that the use of the term "labor force participation" is not in line with the standard use of this term in macroeconomics.

[2] All datawork is original. We use the U.S. population census data available through IPUMS (2004). Although censuses were conducted in 1960, 1970, ..., 2000, the income and worktime questions therein referred to the previous year. Hence, the observations used in this paper are for 1959, 1969, ..., 1999. Only nonfarm married couples with each spouse between the ages of 25 and 64 are considered. (See the appendix for further discussion.) We focus on this particular time period because the U.S. census provides income information for both spouses (and not just the main respondent) only beginning with the 1960 census.

[3] More precisely, the mean vector and covariance matrix of any bivariate distribution, $\mathbf{m} = \begin{bmatrix} m_1 \\ m_2 \end{bmatrix}$ and $\mathbf{S} = \begin{bmatrix} s_1^2 & s_{12} \\ s_{12} & s_2^2 \end{bmatrix}$, can be uniquely represented as $\mathbf{m} = \begin{bmatrix} m_1 \\ m_1 GG \end{bmatrix}$ and $\mathbf{S} = \begin{bmatrix} (m_1 CV_1)^2 & \rho\,(m_1 CV_1)\,((m_1 GG)\,CV_2) \\ \rho\,(m_1 CV_1)\,((m_1 GG)\,CV_2) & ((m_1 GG)\,CV_2)^2 \end{bmatrix}$, where $GG = m_2/m_1$ represents the gender gap (a measure of inequality between genders), $CV_1 \equiv s_1/m_1$ and $CV_2 \equiv s_2/m_2$ are coefficients of variation (a measure of inequality within gender i) and $\rho = s_{12}/\left(s_1 s_2\right)$ is the correlation coefficient (a reflection of the degree of assortativeness of matching). We should point out that the value of the correlation of spouses' potential earnings could reflect factors other than the assortativeness of matching (e.g., a network of business connections acquired by a couple post marriage, a startup of a family business).

of aggregate participation is somewhat novel. In connection to this, Mulligan and Rubinstein (2002) present a reduced-form model of household specialization in which, for some parameter values, increasing within-gender inequality has a greater impact on female labor supply than the closing of the gender gap.

To the extent that we study the aggregate and disaggregated responses of households to changes in their potential earnings, without inquiring into the underlying causes of these changes, our paper is close in spirit to Juhn and Murphy (1997). However, because we employ a family time allocation model and work with the 1960-2000 U.S. census data, our methods and dataset are different. Their main focus is on documenting features of time allocation for disaggregated groups of couples. Qualitatively, the same features carry over to our dataset.

We construct a model of heterogenous couples, in which the potential market earnings of husbands and wives are jointly log-normally distributed. The assumption of log-normality enables us to correct for the selection bias when calibrating the model. Spouses jointly decide on their time allocation between market work, home production, and leisure. Because the change in married female labor supply occurred predominantly at the *extensive* rather than *intensive* margin, we focus on the participation decision, assuming that the market hours of work are fixed, although at different levels for men and women, to allow for a better fit with the data.[4] The home consumption good is produced by combining home appliances, purchased in the market, with perfectly substitutable male or female time.

We calibrate the parameters of the model to ensure a match with several important moments that we document using the 2000 U.S. census data, such as the fraction of two-earner couples, several moments of the observed group-specific earnings and hours worked conditional on participation. We then conduct five counterfactual experiments within the calibrated model; the first four isolate the impact of different aspects of the earnings distribution, and the last one investigates the fall in the relative price of home appliances. When estimating the parameters of the earnings distribution needed to conduct the experiments, we face the selection bias problem, as the potential earnings of non-workers are not observed. In order to correct for this selection bias and to take advantage of the observed characteristics of individuals in our sample, we estimate parameters of the earnings distribution by applying a censored regression model in conjunction with the participation (censoring) rule implied by the calibrated model. Finally, the data on relative prices of home appliances, needed for conducting the last experiment, are taken from the Bureau of Economic Analysis.

The mechanism used in this paper embodies earnings heterogeneity across individuals and provides clear predictions with respect to leisure time. This information helps in the evaluation of the relative importance of each force examined. Thus, in addition to investigating the ability of each counterfactual experiment to account for the observed aggregate trend in the fraction of two-earner couples, we also test its predictions against the cross-sectional features of the rise in LFP of married females and trends in leisure time experienced by married people. To briefly highlight these empirical facts, we document that although female participation increased for all groups differentiated by husband's real earnings, the increase was greater for females married to men with labor earnings in the upper range of the distribution.

[4] In fact, the data reveal that the hours worked by those who participate are concentrated around full-time work hours, indicating that a decision to participate is a lumpy time investment. Moreover, a person's working hours are not entirely his/her choice, as they are often determined by the employer. The assumption of a discrete work choice is common in the literature (e.g., Greenwood, 2003, Attanasio et al. 2006).

Put differently, the strong negative correlation of female participation and the husband's income in 1959 became much weaker by 1999. With regard to empirical leisure trends, we highlight that since 1965, both working and stay-home wives experienced gains in leisure time, with a greater increase enjoyed by stay-home wives. The average leisure time of married females, however, declined, due to the compositional change, as more women joined the labor force.

The main contributions of this paper are summarized below.

i. We build a model of family decision making capable of shedding light on how different aspects of the joint potential earnings distribution of husbands and wives (the gender earnings gap, within-gender inequality and assortativeness of matching) affect aggregate time allocation.

ii. We quantitatively assess the impacts of changes in several aspects of the joint potential earnings distribution and the decline in prices of home appliances on family time allocation since 1959. The model's structure and embodied heterogeneity allow us to subject the factors under consideration to empirical tests, previously unexplored in this context. In particular, we investigate the implications of the factors examined for (1) participation among groups of women disaggregated according to the husband's earnings and (2) gender-specific leisure trends among two-earner couples, one-earner couples, and on aggregate.

iii. We find that changes in the potential earnings distribution of husbands and wives account for nearly 90% of the observed increase in LFP of married females. The decomposition of this overall impact into the impact generated by each of the aspects reveals that the closing of the gender earnings gap drives most of the aggregate increase in participation (accounting for over 70% of the observed rise) and in a manner consistent with the cross-sectional pattern of female participation and changes in female leisure among two-earner couples, one-earner couples and women on aggregate. The decline in the relative price of home appliances accounts for only 5% of the rise in female LFP, while implying counterfactually strong increases in gender-specific leisure time.

iv. Our work establishes the economic significance of the rise in purchasing power that results from the closing of the gender gap, as we find that it generates widespread diffusion of home appliances,[5] whereas the home production revolution is commonly regarded as an outcome of falling prices of home appliances alone. It is through its effect on purchases of home appliances that the closing of the gender gap generates gains in leisure for both stay-home and working wives, with stay-home wives enjoying most gains. Intuitively, as many of the one-earner male families are close to the case of full specialization of stay-home wives in home production, stay-home wives do not share gains in leisure time from the diffusion of home appliances with their spouses. As the closing of the gender gap generates a rise in the fraction of dual-earner couples, the average female time spent on leisure falls, thus the composition effect dominating group-specific effects, just as seen in the data.

v. In addition, our results shed light on reasons behind different findings of Jones et al. (2003) and Greenwood et al. (2005). Discussion of this is given in Section VIII.

It is important to emphasize that we employ a family decision making model that treats couples' potential earnings and prices of home appliances as exogenously given. We do not inquire about the underlying reasons for their change. Thus, our results only speak to the *direct* effects of the distribution

[5] Greenwood et al. (2005) provides a comprehensive summary of the home production revolution that took place in the U.S. in the 20th century.

aspects and appliance prices on households' time allocation choices.

Our results are important because they suggest that in order to understand the dramatic rise in LFP of married women, it is essential to understand the determinants of the joint earnings distribution of husbands and wives, and in particular, those leading to the closing of the gender gap. These can potentially include the introduction of anti-discriminatory laws (Jones et al., 2003), factors affecting selection into marriage (Caucutt et al. 2002), changes in production technology, such as women-biased technical change (Galor and Weil, 1996), changes in returns to experience (Olivetti, 2006), and finally, factors affecting women's decisions regarding gains in education and experience, e.g., diffusion of the contraceptive pill (Goldin and Katz, 2002), changes in cultural norms (Fernández et al. 2004), reduction in the cost of child care (Attanasio et al. 2006), or improvement in home production technology.[6] More work is needed to disentangle the effects of these various factors. Our work further suggests that decision making at the *household* level should be incorporated in these attempts. Attanasio et al. (2006), which uses a lifecycle model to investigate the relative importance of changes in returns to experience, cost of child care and rate of depreciation of human capital when out of the labor force, is an example of such a work. Similar in spirit is a work of Gayle and Golin (2006), which investigates the role of labor-market attachment, on-the-job human-capital accumulation, occupational sorting and discrimination in the closing of the gender gap and the increase in female experience.

The rest of this paper is organized as follows. Section II describes the empirical trends in labor force participation and leisure of married couples. The model is presented in Section III, and its calibration is described in Section IV. Estimation of the earnings distribution parameters is given in Section V. The main quantitative findings are reported in Section VI. Extensive sensitivity analysis is reported in Section VII. In Section VIII, we compare our results to Jones et al. (2003) and Greenwood et al. (2005). Section IX concludes.

II. Empirical Trends in the Time Allocation of Married Couples

We use the U.S. census data on married couples, restricting our attention to couples for which each of the spouses is between the ages of 25 and 64.[7] (See the appendix for more details on the sample.) All the couples in the sample can be identified as either two-earner couples, male-earner couples, female-earner couples, or no-earner couples. Since nearly the entire increase in the fraction of two-earner couples is due to the decline in the fraction of male-earner couples, we choose to focus on only these two types of couples, and eliminate female-earner and no-earner couples from the original sample. Note that in doing so, we ignore only a small fraction of the married population (6%). Moreover, by focusing on the sample in which only females suffer from selection bias, we significantly simplify the procedure of estimating the joint earnings distribution. Note that in the remaining sample, two-earner couples are equivalent to couples with working females.

[6] Note that the finding of this paper that the fall in appliance prices accounts for a very small part of the rise in female LFP, does not rule out its indirect effects. For example, a part of the effect of falling prices of home appliances could be manifested in the closing of the gender gap through its impact on education and experience gains.

[7] Since we do not model human capital accumulation, we consider only those individuals who are sufficiently old that they can be regarded as having completed their education.

Note that the sample we are working with is different from samples used in labor studies (e.g., Mulligan and Rubinstein, 2008). Labor studies usually aim to estimate the returns to observable or unobservable characteristics, the degree of gender discrimination, etc., and work with samples of white individuals (not married couples) that are employed full time and full year. These studies refer to the gender gap as the difference in gender-specific hourly wage not accounted for by differences in characteristics. By contrast, we aim to study how household-level time allocation choice of all married couples of working age is affected by potential earnings and prices of home appliances. Hence, when estimating the distribution of potential earnings (in Section V), we work with the same sample of married households for which we report changes in participation, rather than full time full year working individuals. Also, what we refer to as the gender gap is the difference in annual potential earnings of men and women, the measure not purged of differences in gender characteristics, discrimination, or annual hours of work.

The main observation we investigate is the increase in the fraction of two-earner couples, from 0.33 to 0.76 during the period 1959-1999 (Figure 1). This increase in the aggregate female labor supply was driven primarily by more women joining the workforce rather than working women extending their hours of work (130% increase in participation rate vs. 17% increase in hours conditional on working). Since the hours of work conditional on working changed relatively little (Figure 2), we focus solely on the extensive margin. (Table 8 in the appendix reports descriptive statistics of the U.S. census sample used here.)

Notably, the increase in female LFP occurred across groups of couples differentiated according to the husband's earnings. Since all male labor earnings[8] are observed in the sample, we were able to split the couples into groups, characterized by the husband's real income. Specifically, we split the sample into ten groups, corresponding to the following arbitrarily chosen ranges of the husband's labor income, measured in 1999 dollars:[9] $(0,\ 12,000]$, $(12,000,\ 24,000]$, $(24,000,\ 36,000], \ldots, (108,000,\ +)$. Figure 3 plots female participation as a function of the husband's real income, revealing that female participation increased for all groups of couples, but the increase was greater for females with husbands in the upper range of the income distribution. Indeed, in 1959 less than 10% of females with husbands earning over 108,000 per year participated in the workforce, while in 1999, this number was over 60%. Thus, in some sense, female LFP became less tied to the husband's earnings.[10] Table 1 reports the percent change in female LFP between 1959 and 1999 for couples differentiated according to the husband's labor earnings.[11]

Table 1: Percent Change from 1959 to 1999 in Female LFP by the Interval of the Husband's Earnings

Husband's Income (in 000)	[0,12]	(12,24]	(24, .]	(36, .]	(48, .]	(60, .]	(72, .]	(84, .]	(96, .]	(108,+)
Data: % $\triangle$ in Female LFP	52.7	56.4	75.7	96.7	116.9	124.9	134.9	137.6	140.9	150.3

Figure 4 highlights the rotation in the schedule of female LFP, by taking out the trend. Precisely, each value reported in Figure 3 was divided by the average female participation in that year.

[8] Labor earnings are defined as the sum of wage and business income, the latter representing labor income for the self-employed.

[9] We use CPI to compute real incomes.

[10] This cross-sectional feature of the aggregate increase in participation was also pointed out in Juhn and Murphy (1997).

[11] Percent change computations throughout the paper are based on a midpoint formula. This table essentially decomposes the aggregate change ($\frac{76\%-33\%}{.5(76\%+33\%)} \approx 80\%$) across subgroups of the married population.

We also document several patterns in leisure. We use the dataset compiled from several time-use surveys by Aguiar and Hurst (2006).[12] The only modification here is that we retain the variable, contained in the source files, that provides information regarding the spousal participation status. We use the same sample characteristics guidelines (e.g., age group, marital status, working husband, etc.) as we applied to the U.S. census data.

Time-use surveys are not representative of the married population in the U.S. In fact, 26.6% and 47.8% of females declared that they work a positive number of hours in the 1965 and 2003 surveys, respectively, whereas the corresponding numbers inferred from the U.S. census data are 39.7% and 76.2%.[13] Thus, while we use time-use surveys to obtain gender-specific trends in leisure for two-earner and male-earner couples, to document the aggregate gender-specific trends, we combine group-specific findings for leisure time with the composition numbers obtained from the census data.

We use the same activity variables as defined by Aguiar and Hurst (2006). However, to be consistent with our conceptual framework, we compute leisure as a fraction of productive time not spent on market or home work. First, we compute weekly productive time for each respondent as 24×7 less the time spent on sleeping, eating and personal care. Weekly leisure hours are defined as the productive time less the time spent working (including work-related travel) less the time spent on home production (including basic child care, own medical care and care for others). To obtain leisure time as a fraction of productive time, we divide weekly leisure hours by weekly productive time.[14] Table 2 reports gender-specific time allocation patterns for two-earner couples and couples with only the husband working. Aggregate gender-specific time allocation patterns are computed using the composition of married couples from the census data. (The fraction of two-earner couples in 1965 and 2003 are 39.7% and 76.2%.) We also report the averages obtained from the time-use surveys.

The main qualitative patterns are summarized as follows.

(1) Working wives and stay-home wives both experienced an increase in leisure time.

(2) Stay-home wives enjoyed a greater increase in leisure time than working wives.

(3) Stay-home wives have always enjoyed greater leisure time than working wives.

(4) The average leisure time of married women declined due to the compositional change.

(5) The relative male-to-female leisure time declined among both male-earner and two-earner couples, the larger decline experienced by two-earner couples.

(6) Married men experienced a slight reduction in leisure time on average. The main factor underlying this change was the decline in leisure time (and an increase in home work time) of men among 2E couples to a level comparable to that enjoyed by their wives.

We emphasize that even though both working and stay-home wives experienced an increase in leisure time, the average leisure time of married females in our sample declined, due to the compositional change, as more women joined the labor force, and working women enjoy less leisure.

[12] The dataset is available at http://troi.cc.rochester.edu/~maguiar/timeuse_data/datapage.html.

[13] For 2003, we use the information on the composition from the 2000 census (76% of married women working) and for 1965, we use the linear interpolation of the 1960 and the 1970 census composition (39.7% of married women working).

[14] With the Stata variable names defined by Aguiar and Hurst (2006), our measure of leisure is [24×7- esp - work - (home production + child_care_basic + own_medical_care + care_others)] / (24×7 - esp).

Table 2. Time Allocation Trends: 1965-2003

	Two-earner Couples (working wives)		Male-earner Couples (stay-home wives)		Average Computed with Census Composition		Average in the Time-Use Surveys	
	1965	2003	1965	2003	1965	2003	1965	2003
WOMEN								
Work	0.51075	0.52679	0	0	0.20315	0.40256	0.13607	0.25442
Home Work	0.22967	0.18246	0.46925	0.38593	0.37406	0.23087	0.40555	0.28859
Leisure	**0.2596**	**0.2908**	**0.5304**	**0.6095**	**0.4228**	**0.3666**	0.45838	0.45699
MEN								
Work	0.63704	0.60717	0.65746	0.61842	0.64935	0.60985	0.65095	0.61062
Home Work	0.05082	0.09418	0.04398	0.08218	0.0467	0.09133	0.04616	0.0905
Leisure	0.31214	0.29865	0.29856	0.29939	0.30395	0.29883	0.30289	0.29888
$\text{Leisure}_m/\text{Leisure}_f$	**1.2024**	**1.0270**	**0.5629**	**0.4912**	0.7189	0.8151	0.6608	0.6540

Note that the time allocated to work for working individuals (Table 2) appears much higher than that reported in the census (Table 8); the difference arises from the differences in the precise definitions of market work in the two surveys. Hence, we refrain from comparing these quantities to those generated by the model.[15] Here, we only wish to draw attention to the trends. The quantities in the last two columns are presented only for curious readers, as the model's predictions will not be compared to these.

III. Model

There is a continuum of measure 1 of heterogeneous households. Each household consists of two people, a male and a female. Individuals are heterogeneous with respect to their earning ability. In particular, couples i's potential earnings are drawn from a bivariate log-normal distribution, $\left(w_m^i, w_f^i\right) \sim LN\left(\mathbf{m}, \mathbf{S}\right)$, where $\mathbf{m}$ and $\mathbf{S}$ refer to the mean vector and covariance matrix of the log-normal distribution.[16] This distribution reflects both observable (e.g., education, age, experience, number of young children) and non-observable (e.g., innate ability, ambition, leadership skills) characteristics of married individuals, as well as the state of production technology, market conditions and factors affecting selection into marriage. Exploring the underlying reasons for changes in this distribution is outside the scope of this paper. We only aim to investigate family time allocation decision for a given distribution.

All agents are endowed with 1 unit of productive time, which is allocated between market work $\left(l^1\right)$, home work $\left(l^2\right)$ and leisure $\left(1-l^1-l^2\right)$. Agents have identical preferences over consumption of the market good $\left(c^1\right)$, consumption of the home good $\left(c^2\right)$ and leisure, represented by $u\left(c^1, c^2, 1-l^1-l^2\right) = \mu \log\left(c^1\right) + \nu \log\left(c^2\right) + (1-\mu-\nu)\log\left(1-l^1-l^2\right)$. Motivated by the discussion in Section II, we

[15] There are other reasons for not making quantitative comparisons. One is that all the information reported is based on *weekly* time allocation, whereas the model is based on a yearly time allocation. In addition, the time-use surveys only extend back to 1965, while the model is employed to investigate the period 1959-1999.

[16] The convention is to write $\left(w_m^i, w_f^i\right) \sim LN\left(\boldsymbol{\mu}, \boldsymbol{\Sigma}\right)$, where the parameters are the mean vector and covariance matrix of the underlying normal distribution. Our notation, however, is more convenient for the purpose at hand. Note that there is a one-to-one mapping between $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and $(\mathbf{m}, \mathbf{S})$ (see the appendix).

assume that all men participate in market production. Furthermore, market hours are indivisible; i.e., the male works $\bar{l}_m^1$, while the female chooses her market worktime from the set $\{0, \bar{l}_f^1\}$. Market hours differ for working men and working women to allow for a better mapping of observables into the model. Throughout the paper, individual variables are subscripted by the individual's gender.

The home good is produced according to the constant elasticity of substitution (CES) production, $F(k,l) = [\theta k^\rho + (1-\theta) l^\rho]^{1/\rho}$, where k denotes the stock of home appliances and l denotes family time allocated to home production. Appliances are purchased in the market at the price q, measured in terms of units of the market good.

Assuming that the bargaining problem within households is solved efficiently, each household's problem can be written as a social planning problem, with λ denoting the relative weight of the male's utility. This weight represents the relative bargaining power of the husband. We assume it to be constant across households and over time. (We return to this assumption later when discussing the results.[17]) Depending on the maximum value associated with each time allocation choice, a household chooses to be either a two-earner household (2E) or a male-earner household (1M). Formally, given its draw (w_m, w_f), a couple chooses $\max\{V_{2E}(w_m, w_f), V_{1M}(w_m, w_f)\}$, where

$$
\begin{aligned}
V_{2E}(w_m, w_f) = \max_{c_m^1, c_m^2, c_f^1, c_f^2, k, l_m^2, l_f^2} & \lambda \left[\mu \log\left(c_m^1\right) + \nu \log\left(c_m^2\right) + (1-\mu-\nu)\log\left(1 - \bar{l}_m^1 - l_m^2\right)\right] \\
& + (1-\lambda)\left[\mu \log\left(c_f^1\right) + \nu \log\left(c_f^2\right) + (1-\mu-\nu)\log\left(1 - \bar{l}_f^1 - l_f^2\right)\right] \\
\text{s.t.} \quad & c_m^1 + c_f^1 + qk \le w_m + w_f, \\
& c_m^2 + c_f^2 \le F\left(k, l_m^2 + l_f^2\right), \\
& 0 \le l_j^2 \le 1 - \bar{l}_j^1, \quad j \in \{m, f\},
\end{aligned}
\tag{1}
$$

and $V_{1M}(w_m, w_f)$ is identical to $V_{2E}(w_m, w_f)$ with $w_f = \bar{l}_f^1 = 0$. Note from (1) that female and male time inputs are perfect substitutes in home production.

After substituting for the optimal consumption of the market and home good, derived analytically,[18] we simplify the maximum value functions associated with each of the two time allocation choices as follows:

$$
\begin{aligned}
V_{2E}(w_m, w_f) = \max_{k, l_m^2, l_f^2} & \mu \log(w_m + w_f - qk) + \nu \log\left(F\left(k, l_m^2 + l_f^2\right)\right) \\
& + (1-\mu-\nu)\left[\lambda \log\left(1 - \bar{l}_m^1 - l_m^2\right) + (1-\lambda)\log\left(1 - \bar{l}_f^1 - l_f^2\right)\right] + \kappa \\
\text{s.t.} \quad & 0 \le k \le (w_m + w_f)/q, \\
& 0 \le l_j^2 \le 1 - \bar{l}_j^1, \quad j \in \{m, f\}.
\end{aligned}
\tag{2}
$$

Then, $V_{1M}(w_m, w_f)$ is the special case of the above $V_{2E}(w_m, w_f)$ with $w_f = \bar{l}_f^1 = 0$. The constant $\kappa = (\mu + \nu)(\lambda \log \lambda + (1-\lambda)\nu \log(1-\lambda))$ is irrelevant for the household's optimization problem.

[17] Alternatively, one could model the bargaining game explicitly as, for example, in Knowles (2005), or one could model the relative weight as an explicit function of certain important factors, for example, relative earnings (Browning and Gortz, 2006).

[18] See Bar and Leukhina (2007) for details. Derivations of sufficient equilibrium conditions are included.

The model implies a partition of the earnings space into two regions: 2E and 1M. We define the decision rule threshold separating the two regions as a function $L(w_m)$ that solves $V_{1M}(w_m, 0) = V_{2E}(w_m, L(w_m))$. The solution can be found only numerically. Figure 5 illustrates the mechanism of the model pertaining to the aggregate labor force participation. Any given point in the earnings space represents a possible realization, (w_m, w_f). The contour plots of the bivariate log-normal density indicate how the couples are distributed over the space. Couples with potential earnings realizations above (below) the threshold choose to be 2E (1M) households. (Couples with realizations on the decision threshold, $(w_m, w_f) = (w_m, L(w_m))$, are indifferent.) In other words, the wife chooses to participate in market production if and only if her potential earnings are large enough relative to those of her husband. Note that the parameters of the earnings distribution determine where the couples are located in the earnings space, while the rest of the parameters (λ, μ, ν, $\bar{l}_m^1$, $\bar{l}_f^1$, θ, ρ and q) determine the shape and location of the decision rule threshold. As shown in the next proposition, two special cases of our model give rise to a linear decision rule: one is the case with no home production ($\nu = 0$),[19] and one is the case with a Cobb-Douglas home production function. With a general CES home production function, the decision rule threshold is non-linear.

Proposition 1 *In our model with (i) $\nu = 0$ or (ii) a Cobb-Douglas home production function, the decision rule threshold is a linear function with intercept zero. (Proof is given in the appendix.)*

Note that both a downward shift of $L(w_m)$ and a shift in the mass of the distribution towards the 2E region cause an increase in the proportion of two-earner couples.

The result summarized in Proposition 2 allows us to interpret the drop in q as a capital-augmenting technological change in home production.

Proposition 2 *In our model, capital-augmenting technological progress in home production is equivalent to a decline in the relative price of home appliances. (Proof is given in the appendix.)*

IV. Calibration

For computational accuracy, it is convenient to work with the logs of the earnings rather than the earnings themselves. Let $X = \log(w_m)$ and $Y = \log(w_f)$, so that $(X, Y) \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.

Because the selection bias problem is least severe in the latest census (as only 24% of women in the 2000 census sample do not work), we choose to calibrate the model to 1999. Assuming 5000 hours of annual productive time, the fixed work hours (conditional on working), $\bar{l}_m^1$ and $\bar{l}_f^1$, are set to 0.44 and 0.34 respectively, to match their 1999 data counterparts.

Note that technological improvements in the home production sector affect the relative returns to work at home and in the market, but the direction of this effect depends on the substitutability of inputs in the home production. Consider a home good, say home-cooked meals, produced by combining home appliances and labor. According to Proposition 2, the decline in the relative price of home appliances is equivalent to the capital-augmenting technological change. If the inputs are complements, capital-augmenting technological progress would cause households to allocate more labor to home production.

[19] Since preferences exclude utility from the home good, no home production will take place.

If the inputs are substitutes, then capital-augmenting technological progress would have the opposite effect.[20] We borrow parameters of the home production function, $\theta = 0.206$ and $\rho = 0.189$, from the estimation given in McGrattan et al. (1997), thus maintaining the assumption of substitutibility and not ruling out the story of falling prices of home appliances.

We set the preference parameters as $\lambda = 0.5$, $\mu = 1/3$, $\nu = 1/3$, lacking criteria for making more meaningful choices (see Section VII on sensitivity). The remaining parameters are those of the earnings distribution and the relative price of home appliances, $\Theta \equiv (\mu_X, \mu_Y, \sigma_X, \sigma_Y, \sigma_{XY}, q)$.

We calibrate Θ to match several important 1999 data moments, summarized in Table 3. The last moment helps us capture the cross-sectional feature of participation, namely, that in 1999, female participation was not closely tied to the husband's earnings.

Table 3. Summary of Moments Used for Calibration, 1999

0.2379	Proportion of 1M couples
10.537	Mean of log of male earnings
0.7543	Variance of log of male earnings
9.802	Mean of log of the working (observed) female earnings
1.1266	Variance of log of the working (observed) female earnings
0.2285	Proportion of 1M couples among couples with $\log(w_m)$ below the mean

Denoting these data moments by $\mathcal{M}$ and the corresponding moments implied by the model by $M(\Theta)$, we calibrate the remaining parameters, $\Theta \equiv (\mu_X, \mu_Y, \sigma_X, \sigma_Y, \sigma_{XY}, q)$, by solving the following minimization problem:

$$\min_{\Theta} \sum_{i=1}^{6} \left(\frac{\mathcal{M}_i - M_i(\Theta)}{0.5\left(\mathcal{M}_i + M_i(\Theta)\right)} \right)^2 . \tag{3}$$

Bar and Leukhina (2007) present derivations of $M(\Theta)$ and discuss the numerical integration methods adopted in its computation.[21] Table 4 summarizes the calibrated parameters.[22]

Table 4. Summary of Calibrated Parameters

Category	Value
Preferences	$\lambda = 0.5,\ \mu = 1/3,\ \nu = 1/3,$
Market hours	$\bar{l}_m^1 = 0.44,\ \bar{l}_f^1 = 0.34,$
Home production	$\theta = 0.206,\ \rho = 0.189,\ q = 0.99586,$
Earnings distribution	$\mu_X = 10.374,\ \mu_Y = 9.5586,$
	$\sigma_X = 0.86894,\ \sigma_Y = 1.2346,\ \sigma_{XY} = 0.66739.$

[20] In fact, Aguiar and Hurst (2006) use the substitutability of inputs (labor and capital in our case) to classify an activity as home production.

[21] Difficulty arises when computing the moments in the model, because the limit of the integration, i.e., a point on the decision rule threshold, must be found for every point at which we evaluate the integrand. This threshold must be computed numerically by equating the value functions of the 2E and 1M problems, which may have corner solutions.

[22] As a part of the robustness check, we repeated the entire analysis performed in this paper under a slightly different calibration procedure. Instead of setting μ, ν and λ as we do here, we allowed these to vary along with Θ in the minimization problem (3), with the initial values for the minimizers taken from Table 4. Although the match improved significantly when this was done, the overall results obtained were very similar to those presented here.

V. Estimation of the Earnings Distribution Parameters

Although the calibration procedure yields the parameters of the potential earnings distribution for 1999, it does not do so for the rest of the years under consideration. Since one of our goals is to study households' responses to changes in all aspects of the joint earnings distribution over the period 1959-1999, we must first estimate these changes.

Because of the non-random selection of women into workforce, we cannot simply set the earnings distribution parameters to their data counterparts. To overcome the selection bias, we predict the unobserved earnings by employing the year-specific censored regression model in conjunction with the participation (censoring) rule implied by the calibrated model. Once the unobserved earnings are predicted, we infer $(\mathbf{m}_t, \mathbf{S}_t)$ from their sample counterparts for $t = 1959, 1969, ..., 1999$. These are then used to conduct the counterfactual experiments in the context of the calibrated model.

Formally, the censored regression model is given by

$$y_i^* = \mathbf{x}_i\beta + u_i, \quad u_i \sim N\left(0, \sigma^2\right), \tag{4}$$

$$y_i = \begin{cases} y_i^* & \text{if} \quad y_i^* \geq p\left(z_i, \Omega\right) \\ 0 & \text{otherwise} \end{cases}, \tag{5}$$

where y_i^* denotes the log of the potential earnings of married female i, and $\mathbf{x}_i$ denotes her personal attributes, which determine her potential earnings, such as years of education, experience, race (see the appendix). Use of the Mincer equation (4) allows us to extract information from individuals' characteristics observed in our sample when predicting the missing potential earnings. If female i chooses to work, her potential earnings are revealed; i.e. the log of her observed earnings, y_i, is equal to y_i^*. The participation rule, $y_i^* \geq p\left(z_i, \Omega\right)$, is given by the time allocation decision rule implied by our calibrated model. Note the explicit dependence on the log of the husband's income, z_i, and the parameters of the model, $\Omega = (\lambda, \mu, \nu, \bar{l}_m^1, \bar{l}_f^1, \theta, \rho, q)$. Thus, our procedure is similar to applying a standard Heckman selection model,[23] but with the selection rule implied by a micro-founded model.

Since we apply this censored regression model to each year separately, the participation rule $y_i^* \geq p\left(z_i, \Omega\right)$ is also year-specific, as we adjust q to reflect the fall in the relative price of home appliances over time, while the rest of the parameters remain fixed at their calibrated values.[24, 25]

Note that we do not aim to develop a novel procedure for the earnings estimates. The purpose of

[23] This is a standard model of selection used in labor literature. See Heckman (1979).

[24] First, if q were fixed at its calibrated value, our main result, that the closing of the gender earnings gap was the main driving force behind the rise in the female LFP, would be reinforced. Consider, for example, estimating potential earnings for non-working females in 1959. In case of the low level of q implied by the 1999 calibration, the participation rule would have a relatively small slope (as home appliances and labor are substitutes in home production), resulting in lower values of predicted missing earnings, and hence a larger estimated gender earnings gap in 1959. A more dramatic closing of the gender gap would then reinforce its quantitative power. Second, even though the results would be reinforced, the change would be quantitatively small, because the decision rule threshold is not very sensitive to changes in q.

[25] The fact that we keep the relative bargaining power, reflected in λ, fixed throughout the entire period under consideration may appear to be problematic. Let us consider the implications of allowing women to gain more bargaining power over time (a decrease in λ). Given our calibration, this would result in a threshold with a relatively small slope in 1959. Hence, the predicted values for the missing earnings in 1959 would be lower, and our estimates would imply a greater closing in the gender gap over the period 1959-1999. Thus, allowing λ to decline over time would reinforce the quantitative importance of the closing gender gap.

applying the year-specific censored regression model along with our model used to provide participation (censoring) rules is simply to obtain estimates of the earnings distribution parameters employed in the counterfactual experiments. In our model, individuals are heterogeneous with respect to their earnings potential, but the model is silent regarding factors that underly this heterogenity. Rather, it only provides results concerning family time allocation across different activities. However, the dataset contains a great deal of information about personal characteristics useful for assessing individual earnings potential. Thus, employing the Mincer regression given in (4) within the context of our model, used to correct for the selection bias, appears to be a natural procedure for estimating $\{\mathbf{m}_t, \mathbf{S}_t\}_{t=1959,\ 1969,...}$. The derivation of the log-likelihood function is given in the appendix. In Section VII on sensitivity, we also study how our results change if the Heckman selection model was applied to arrive at these estimates.

The estimated censored regression model is used to predict the log of potential earnings of stay-home women, thus filling in the missing values. We then convert logs into annual earnings. Finally, we record $\{\mathbf{m}_t, \mathbf{S}_t\}_{t=1959,\ 1969,...}$ from the completed sample. Table 5 highlights our findings by reporting the estimated means, coefficients of variation, correlation coefficient, and the gender earnings gap.

Table 5. Aspects of the Estimated Joint Earnings Distribution

Parameter \ year, t	1959	1969	1979	1989	1999
m_1	36712	47597	47382	48078	52888
m_2	6779	11434	13702	18958	23739
$CV_1 = s_1/m_1$	0.781	0.744	0.769	0.849	1.035
$CV_2 = s_2/m_2$	1.212	0.997	0.977	0.956	1.075
$\rho = s_{12}/(s_1 s_2)$	0.004	0.042	0.052	0.143	0.159
$GG = m_2/m_1$	0.185	0.240	0.289	0.394	0.449

We reiterate that these estimates are not directly comparable to those derived in studies of wage distributions. First, we consider all married individuals of working age and regard a person to be a worker if he/she works a positive number of hours, while most related studies consider full-time full-year working-age employees. Second, while other studies consider either hourly wages or annual labor earnings of full-time full-year individuals, we consider annual labor incomes without controlling for gender differences in hours worked, or any other characteristic. In fact, women work much less on average in our sample, which is reflected in their annual earnings. These two facts are responsible for the large size in our measure of the gender discrepancy in annual earnings.[26]

We find that the gender earnings gap, corrected for the selection bias, narrowed monotonically (m_2/m_1 increased by 143% over the 40 years), in contrast to the pattern in the observed gender earnings gap, which remained roughly unchanged until the 1980s. To a large extent, the closing of the gap is due to gains in observable characteristics. Finally, we find that within-gender inequality increased for men and slightly decreased for women. Also, note that the estimated correlation of husbands' and wives' earnings increased over time, although remaining at a very low value even today.

[26] Note that even the observed gap for 1999 (see Table 8) is very large: 27303/50097=0.545.

VI. Quantitative results

A. Counterfactual Experiments

To recap, we calibrated the model to match several important moments in 1999. Year-specific censored regressions with the censoring rule, generated by the calibrated model and adjusted for 1959, 1969, ..., 1989 to reflect the fall in the relative price of home appliances over time, were applied to the dataset and used to predict the unobserved earnings of stay-home females. The mean vectors and covariance matrices of the year-specific completed samples, $\{\mathbf{m}_t, \mathbf{S}_t\}_{t=1959,1969...}$, provided the estimates of the moments of year-specific earnings distribution faced by households. Table 5 reports changes in several aspects of the distribution imputed from these estimates.

One of our goals is to assess the impact of changes in the entire earnings distribution and isolate the impact of its various aspects. Note that the mean vector and covariance matrix, $\mathbf{m}$ and $\mathbf{S}$, of any bivariate distribution can be represented by either of the following two forms:

$$(6)\qquad \begin{bmatrix} m_1 \\ m_2 \end{bmatrix}, \begin{bmatrix} (m_1CV_1)^2 & \rho\,(m_1CV_1)\,(m_2CV_2) \\ \rho\,(m_1CV_1)\,(m_2CV_2) & (m_2CV_2)^2 \end{bmatrix},$$

$$(7)\qquad \begin{bmatrix} \frac{m_1+m_2}{1+GG} \\ \frac{GG(m_1+m_2)}{1+GG} \end{bmatrix}, \begin{bmatrix} \left(\frac{CV_1(m_1+m_2)}{1+GG}\right)^2 & \rho\left(\frac{CV_1(m_1+m_2)}{1+GG}\right)\left(\frac{CV_2(GG)(m_1+m_2)}{1+GG}\right) \\ \rho\left(\frac{CV_1(m_1+m_2)}{1+GG}\right)\left(\frac{CV_2(GG)(m_1+m_2)}{1+GG}\right) & \left(\frac{CV_2(GG)(m_1+m_2)}{1+GG}\right)^2 \end{bmatrix},$$

where $GG = m_2/m_1$ is the gender earnings gap (a measure of inequality between genders), $CV_1 = s_1/m_1$ and $CV_2 = s_2/m_2$ are gender-specific coefficients of variation (a measure of inequality within gender i), $m_1 + m_2$ represents the potential purchasing power, and $\rho = s_{12}/(s_1s_2)$ is the correlation of spousal potential earnings. Note that the first representation expresses the distribution parameters in terms of the mean vector, gender-specific inequality and correlation alone. Representation (7) provides a more detailed look into the aspects of the distribution, as it further breaks down the mean vector into the relative female to male earnings and purchasing power of the potential household income. It is obtained from representation (6) by setting $m_1 = (m_1+m_2)/(1+GG)$ and $m_2 = GG\,(m_1+m_2)/(1+GG)$. Conducting the counterfactual experiments whose goal is to isolate the impact of each of the aspects of the distribution amounts to computing the appropriate experimental distribution parameters $\{\mathbf{m}_t^{Exp\ j}, \mathbf{S}_t^{Exp\ j}\}_{t=1959,1969...}$ and then feeding them into the calibrated model.

Experiment 1: *Changes in the entire earnings distribution.* We allow all of the distribution parameters to vary across time according to our estimates; that is, the experimental distribution parameters are set to the distribution estimates, $\{\mathbf{m}_t^{Exp\ 1}, \mathbf{S}_t^{Exp\ 1}\}_{t=1959,1969...} = \{\mathbf{m}_t, \mathbf{S}_t\}_{t=1959,1969...}$. The remaining parameters of the model, including the relative price of home appliances, are kept at their calibrated values. We solve the model for each year. We then interpret the model's predictions regarding time allocation, say for 1989, as revealing what the time allocation patterns would have been in 1989 if the environment in 1989 were identical to that of 1999, with the only difference being the parameters of the potential earnings distribution.

Experiment 2: *Changes in the mean vector: closing of the gender earnings gap and the rise in purchasing power.* We let the mean values of husbands' and wives' potential earnings vary across time according to the distribution estimates, $\{\mathbf{m}_t^{Exp\ 2}\}_{t=1959,1969...} = \{\mathbf{m}_t\}_{t=1959,1969...}$. In order to keep the correlation and coefficients of variation fixed at the 1999 level (thus isolating the effect of the mean vector alone), we compute $\{\mathbf{S}_t^{Exp\ 2}\}_{t=1959,\ 1969,...}$ according to (6), in which we use the year-specific mean values but keep CV_1, CV_2 and ρ at their 1999 levels from Table 5.

To complement Experiment 2, we conduct two additional experiments. This allows us to decompose the impact from changing the mean vector into the impact due to the increase in the purchasing power associated with it, $m_1 + m_2$ (Experiment 2.1), and the impact due to the change in the relative earnings, m_2/m_1 (Experiment 2.2). We obtain $\{\mathbf{m}_t^{Exp\ 2.1}, \mathbf{S}_t^{Exp\ 2.1}\}_{t=1959,1969...}$ from (7), in which we use the year-specific m_1 and m_2, while keeping GG, CV_1, CV_2 and ρ at their 1999 levels given in Table 5. Similarly, $\{\mathbf{m}_t^{Exp\ 2.2}, \mathbf{S}_t^{Exp\ 2.2}\}_{t=1959,1969...}$ are computed from (7), using the year-specific GG, while keeping m_1+m_2, CV_1, CV_2 and ρ at their 1999 levels.[27]

Experiment 3: *Changes in within-gender inequality.* This experiment aims to isolate the impact of changing within-gender inequality. We compute $\{\mathbf{m}_t^{Exp\ 3}, \mathbf{S}_t^{Exp\ 3}\}_{t=1959,1969...}$ from (7) by letting gender-specific coefficients of variations, CV_1 and CV_2, vary across time (Table 5), while keeping other aspects of the distribution, $m_1 + m_2$, GG and ρ, at their estimated values for 1999.

Experiment 4: *Changes in assortativeness of matching.* We isolate the impact of the growing correlation between husbands' and wives' earnings by computing $\{\mathbf{m}_t^{Exp\ 4}, \mathbf{S}_t^{Exp\ 4}\}_{t=1959,1969...}$ from (6), in which we vary ρ across time according to its estimates from Table 5, while keeping m_1, m_2, GG, CV_1, CV_2 at their estimated values for 1999.

Experiment 5: *Changes in the relative price level of home appliances.* In addition to investigating the impact of different aspects of the earnings distribution, we also study the impact of the fall in the relative price level of home appliances. National Income and Product Accounts Table 2.5.4 provides detailed price indices for numerous components of personal consumption expenditures. According to those, the relative price of durable consumption to non-durable consumption halved in the period under consideration, while the relative price of housing operation, a category including electricity, gas, telephone and water, did not change over the time period under consideration. The change in a more narrow category, titled "kitchen and other household appliances" (item 30),[28] was most dramatic. Table 6 reports the price index of this category relative to the price index of personal consumption expenditures. This series is normalized to 1 in 1999. We set q_{1999} to the calibrated value and vary it across time in accordance with the price index of appliances, given in Table 6. The rest of the parameters, including the parameters of the earnings distribution, are kept at their calibrated values.

[27] The distribution parameters corresponding to Experiment 2.2 are such that the mean value of male earnings declines over time (as GG increases), in order to eliminate the income effect associated with both male and female income growth. It may, however, be more natural to eliminate only the income effect associated with male income growth. After all, the "catching up" of female earnings inevitably incorporates the increase in purchasing power. Hence, we also perform an experiment (Experiment 2.3) isolating the effect of the closing of the gender gap compiled with the purchasing power due to the female income growth. We compute the experimental distribution from the representation of $\mathbf{m}$ and $\mathbf{S}$ given in footnote 3, in which we use the year-specific GG, but keep m_1, CV_1, CV_2 and ρ at their 1999 levels. Because m_1 stagnated over the years (except in the 1960s), most gains in purchasing power $m_1 + m_2$ are due to m_2, so the results generated by Experiment 2.3 are very close to those of Experiment 2 (which incorporates the closing of the gender gap and the rise of purchasing power due to both male and female incomes). In the 1960s, the quantitative results lie between those generated by Experiment 2 and those generated by Experiment 2.2. We do not focus on the results of Experiment 2.3, because the main message they provide is the same as that provided by Experiment 2.

[28] This category includes refrigerators, freezers, cooking ranges, dishwashers, laundry equipment, stoves, room air conditioners, sewing machines, vacuum cleaners and other appliances.

Table 6. Index of the Relative Price of Home Appliances

1959	1969	1979	1989	1999
3.112	2.312	1.848	1.347	1

B. Aggregate Implications for Female LFP

Figure 6 summarizes the impact of the main counterfactual experiments performed in the framework of the calibrated model on the evolution of the fraction of 2E couples. The bold line represents the empirical trend.[29] To assess the quantitative significance of the channels explored in this paper, we compare changes generated by our experiments with the corresponding changes in the data.

The overall change in the fraction of two-earner couples in the data during the period 1959-1999 was 79% (from 0.33 to 0.762). We find that changing the entire earnings distribution (Experiment 1) within the framework of the calibrated model generates a smaller rise, of about 69.6% (from 0.36 to 0.747), thus accounting for 88% of the observed increase in the fraction of two-earner couples (69.6% out of 79%).

Experiment 2, which isolates the impact due to changing gender-specific means, generates a rise in the fraction of 2E couples from 0.416 to 0.747, accounting for 72% of the observed change and hence appearing to be the aspect of the distribution with the most important quantitative implication for married female LFP. Also note that in the 1970s, the closing of the gap appears to be less important, although still more important than other channels considered in this paper.

Decomposing the impact from Experiment 2 into an impact arising from the increase in the purchasing power and an impact arising from the change in the relative earnings reveals that nearly the entire effect on the aggregate female participation is due to the latter change. Indeed, Experiment 2.1 generates a very small rise in female participation (from 0.734 to 0.747), while Experiment 2.2 generates a substantial rise (from 0.432 to 0.747), nearly identical to the rise implied by Experiment 2.

Other aspects of the earnings distribution have a much smaller impact on women's participation, with the change in within-gender inequality (Experiment 3) accounting for 10% of the empirical trend, and the assortative matching (Experiment 4) for 4%.

Finally, the decline in the relative price of home appliances (Experiment 5), accounts for a very small part (5%) of the observed rise in the fraction of two-earner couples.

C. Cross-sectional Implications

Note that we used neither the cross-sectional participation patterns nor leisure patterns in the calibration of the model. However, each experiment generates clear predictions with respect to these quantities, hence providing an additional test for evaluating the relative importance of different channels. We describe the cross-sectional implications of our experiments for participation and leisure, contrasting them with the corresponding empirical trends.

[29]It should be noted that because the selection bias is relatively small in 1999 and we use the selection rule predicted by the model, the obtained parameters of the earnings distribution for 1999 are in close accordance with the calibrated parameters for the earnings distribution reported in Table 4. There is, however, still a small mismatch in the fraction of two-earner couples at the initial point between the data (0.762) and the model (0.747).

Recall that Figure 3 illustrates female participation as a function of the husband's labor income in 1999 dollars. Participation increased during each decade and for nearly all intervals of the husband's income, except for the second interval (12,000 - 24,000], in the 1990s.[30] The magnitude of the increase was larger for women with husbands in the upper intervals.

The analogs to Figures 3 and 4 generated by Experiment 2 are given by Figures 7 and 8. We do not report the cross-sectional implications associated with changing the entire earning distribution (Experiment 1), as they are largely driven by changes in the mean vector (Experiment 2). As a result of Experiment 2, female participation increases across all intervals of the husband's income during the period 1959-1999, but the increase is more significant for females with husbands in the upper range of the distribution (Table 7). Comparing the results reported in Table 7 to their data counterparts reported in Table 1 confirms that the mean vector of the earnings distribution drives not only the overall increase in female LFP but also the cross-sectional schedule of female LFP as a function of the husband's income.[31]

Table 7. Experiment 2: Percent $\triangle$ in Female LFP by the Interval of the Husband's Earnings, [1959-1999]

Husband's Income (in 000)	[0,12]	(12,24]	(24, .]	(36, .]	(48, .]	(60, .]	(72, .]	(84, .]	(96, .]	(108,+)
Model: %$\triangle$ in female LFP	25	54.9	80.8	99	112.4	122.6	130.7	137.2	142.6	153.2

The asymmetric effect experienced by couples across different income intervals is driven by both the increase in the purchasing power and the change in the relative earnings, associated with changing the mean vector. As a result of Experiment 2.1 (change in $m_1 + m_2$), couples in the lowest interval increase participation by 3.8%, while couples in the top interval do by 53%. Due to the relatively low correlation of spousal earning ability, for most couples in the category of low husband income, the gender gap is already narrow in 1959. Moreover, females with husbands at the lower end of the distribution have lower reservation earnings. Hence, these women already experienced high participation rates in 1959 and had less to gain, despite being more strongly affected by the diffusion of home appliances. Note that there is no inconsistency of the result that couples across all income intervals are positively affected by the rise in purchasing power and the result, reported above, that the aggregate impact on participation generated by Experiment 2.1 is negligible: due to the downward sloping schedule of female participation as a function of the husband's income, the effect of changing the composition of couples across the intervals obtained from this experiment ($m_{1,1959}^{Exp\ 2.1} < m_{1,1999}^{Exp\ 2.1}$) works to substantially reduce the aggregate impact. As a result of Experiment 2.2 (change in GG), female participation among couples in the lowest interval increases by 11.3%, while couples in the top interval experience a much larger increase of 105.6%.[32] The intuition obtained from this experiment is similar to the one obtained from Experiment 2.1, except that diffusion

[30] As suggested in Bar and Leukhina (2007), female participation among many couples in this income range was negatively affected by the major expansion of the Earned Income Tax Credit in the 1990s.

[31] The model, however, generates participation rates in 1999 that are too high for low income intervals and too low for high income intervals. The reason is that our estimate of the correlation of spousal labor incomes is positive but low, resulting in high relative earnings of females married to husbands at the bottom of the earnings distribution. In reality, husbands with low labor income may have other sources of income, not incorporated into our model (e.g., welfare or dividend income, income from informal activities, and family transfers).

[32] Compare these percent changes to the corresponding 25% and 153.2% generated by Experiment 2 (Table 7). It may at first seem puzzling that the aggregate increase in the fraction of 2E couples generated by Experiment 2.2 is nearly identical to that generated by Experiment 2, yet, when decomposing the aggregate increase into income intervals, the contribution obtained from Experiment 2.2 appears somewhat smaller. There is no inconsistency of the two results: due to the downward

of home appliances is even across all categories of husbands' income. Closing the gender gap does not have as strong an effect on females with husbands in lower income intervals, as they already experienced high levels of participation.

Recall that changing within-gender inequality (Experiment 3) did little to explain trends in aggregate participation; we find that it is also unimportant for the cross-sectional schedule, which remains roughly unchanged.

Figure 9 demonstrates the cross-sectional implications of the increase in the assortativeness of matching (Experiment 4). Because this experiment generates a very small rise in aggregate female participation (3.33%), we report the participation pattern relative to the year's average, in order to only highlight the relative impact experienced by couples across different income intervals (Figure 9). Increasing correlation tends to close the average gender gap for women with husbands in the upper range of the earnings distribution (18% rise in $E(m_2)/E(m_1)$ for the top interval), encouraging their participation, and to widen it for women with husbands in the lower range (35% fall in $E(m_2)/E(m_1)$ for the lowest interval), thus inducing a fall in their participation. Thus, increasing assortativeness of matching appears to partly contribute to the rotation of the female participation schedule as a function of the husband's income. It also appears that the seemingly small aggregate impact of Experiment 4 fails to reflect a greater impact experienced by disaggregated groups of couples. In particular, the experiment generates an increase in female LFP of as much as 18.2%, 22.4% and 38.2% for couples in the top three income intervals, respectively.

The fall in the relative price of appliances generates a modest increase in wives' participation across all levels of the husbands' labor income, with a greater impact experienced by couples in the top intervals (Figure 10). The impact ranges from less than 2% increase in participation for the lower two intervals to increases of 9% and 11% for the top two intervals.

D. *Implications for the Diffusion of Home Appliances and Leisure Trends*

Up to this point, we have demonstrated that changes in the mean vector of the earnings distribution experienced by married couples account for a large part of the aggregate rise in female LFP as well as participation trends for groups of females disaggregated according to the husband's income. Decomposing this impact reveals that it is the change in the relative female-to-male earnings that predominantly drives the aggregate and disaggregated participation trends. The remaining discussion will reveal that the increase in the purchasing power associated with the change in the mean vector (primarily due to female incomes catching up, as male incomes stagnated in the 1970s and 1980s) also has quantitatively important implications, as it generates a widespread diffusion of home appliances and group-specific leisure trends, in close match with their empirical counterparts.

In what follows, we will focus on implications of Experiment 2 (and its decomposition into Experiments 2.1 and 2.2) and Experiment 5 regarding purchases of home appliances and group-specific leisure trends.[33]

sloping schedule of female participation as a function of the husband's income, the change in the composition of couples across the intervals implied by Experiment 2.2 works to boost the aggregate impact ($m_{1,1959}^{Exp\ 2.2} > m_{1,1999}^{Exp\ 2.2}$), while the opposite is true for Experiment 2 ($m_{1,1959}^{Exp\ 2} < m_{1,1999}^{Exp\ 2}$).

[33] We do not discuss Experiments 3 and 4 because we find them to be quantitatively unimportant in driving the diffusion of home appliances and leisure trends. Experiment 1 is left out of the discussion because it generates results that are nearly

Finally, we will show that even though Experiment 5 produces a greater diffusion of home appliances than that implied by Experiment 2, its implication for female leisure trends is in contradiction with the data.

As a result of Experiment 2, the average expenditure on home appliances (k) increases by 60%. This impact is mainly due to the increase in purchasing power, as Experiment 2.1 yields nearly the same increase (58%). In Experiment 5, a three-fold decline in the price of home appliances leads to a larger (112%) increase in k.[34] As home appliances diffuse through different subsets of the married population, they influence time allocated to leisure.

Figures 11-16 summarize leisure trends generated by Experiments 2 and 5. The change in leisure trends generated by Experiment 2 is also mainly due to the increase in purchasing power. First, note from Figures 11 and 12 that men with stay-home wives enjoy less leisure time on average than their spouses. Also note that for each year, the average leisure time of men with stay-home wives is very close to $1 - \bar{l}_m$, indicating that nearly all 1M couples are in the case of complete specialization, in which the optimal production of the home good is reached with only the wife's time. Hence, men with stay-home wives experience a nearly zero increase in leisure due to both Experiments 2 and 5, as very few of them devote any time to home production and thus have something to gain from the diffusion of home appliances. Their wives, on the other hand, enjoy a much greater increase in leisure (3.63% (3.25%) due to Experiment 2 (Experiment 2.1) and 6.8% due to Experiment 5), because they do not have to share gains from the greater use of home appliances with their spouses. This finding provides economic insight into why the relative male-to-female leisure declined among 1M couples (a 3.5% (3.2%) drop as a result of Experiment 2 (Experiment 2.1) and a 6.3% drop as a result of Experiment 5, see Figure 13). All of these implications are qualitatively in line with their empirical counterparts reported in Table 2.

All 2E couples in the calibrated model reach an interior optimal solution. It is then optimal to allocate the same amount of leisure time to both spouses because $\lambda = 0.5$. The wife works fewer hours in the market ($\bar{l}_m^1 > \bar{l}_f^1$) but works more at home. Since both spouses work at home, they both enjoy an increase in leisure associated with the greater use of appliances (Figure 14). The 4.4% increase implied by Experiment 5 is again greater than the 2.8% (2.1%) rise due to Experiment 2 (Experiment 2.1). The prediction that both spouses experienced equal gains in leisure, shared by both experiments, is counterfactual. In fact, Table 2 reveals that the relative male-to-female leisure among two-earner couples dropped significantly, from 1.2 to 1. This suggests that the bargaining process has drastically changed among these couples and provides further motivation for works explicitly modeling spousal bargaining (e.g., Knowles, 2005).

Comparison of Figures 12 and 14 reveals that on average, stay-home women enjoy more leisure than working women, which is in close agreement with the data. Also in line with our empirical findings, Experiments 2 and 5 predict that stay-home wives enjoyed a higher increase in leisure than working wives, again because stay-home wives devote more time to home production and thus have more to gain from the diffusion of home appliances.

identical to those of Experiment 2.

[34] The changes in the ratio k/c_2, which captures the evolution of capital intensity in household production, convey a similar message. Experiment 2 generates a 24% increase of this ratio, most of the increase arising from the change in the purchasing power, as Experiment 2.1 generates a comparable (20.7%) increase. Experiment 5 generates a larger (45%) increase in k/c_2.

Figure 15 reports the leisure trend averaged over all women, thus incorporating the composition effect. Consistently with our empirical findings, Experiment 2 (and Experiment 2.1) predicts that although both working and stay-home females increase their leisure, the average female leisure declines due to the composition effect, as more women work in the market today, and working women enjoy less leisure. Recall that Experiment 5 generates stronger group-specific leisure gains; it also predicts a very weak composition effect, as it generates very little change in the fraction of 2E couples. Thus, Experiment 5 implies that leisure time averaged over all married women increases (Figure 15), in contradiction with the empirical trend (Table 2).

Finally, Figure 16 illustrates that the qualitative predictions for the average leisure of men are similar to those for women, with the negative composition effect dominating for Experiment 2 and the positive group-specific effect dominating for Experiment 5. Although Experiment 2 generates a fall in male leisure, as seen in the data, this fall is due to the composition effect, while in the data, it is mainly due to the drop in male leisure among 2E couples, which the model fails to generate, as discussed above.

It can be concluded that the increase in purchasing power associated with changes in the mean vector of the earnings distribution generates a widespread use of home appliances and leisure trends that are qualitatively consistent with empirical leisure trends experienced by working women, stay-home women and women on average. It also accounts for the fall in the relative male-to-female leisure among 1M couples. It appears that to account for the fall in relative male-to-female leisure among 2E couples, one would need to extend the model to allow for a change in bargaining power. As discussed in Section V, allowing for λ to drop over time would only reinforce our finding that the narrowing of the gender gap was a more important force than the decline in prices of home appliances.

VII. Sensitivity Analysis

Earnings definition adjusted for discrepancy in annual hours

Clearly, a low income observation in our sample does not necessarily imply the person has a low earning ability, as he/she may be working only part-time. We repeated the entire analysis of this paper under an alternative definition of labor income, aimed to capture each person's full time earnings based on his or her part-time earning ability. Precisely, we created an artificial income variable, which consists of the observed labor income adjusted by the discrepancy between his/her annual hours worked and the gender-specific and year-specific average hours in the sample. For example, if a person worked 1000 hours in 1999, earning \$20,000 in that year, while the average male hours were 2215.4, we create an adjusted income for that person of \$20,000(2215.4/1000). We recalibrated the model, and following the procedure in Section V, we estimated the parameters of the adjusted earnings distribution for each decade. We found that the closing of the gender gap is monotonic from 0.2 to 0.46. Further, all the results are quantitatively very close to those obtained using annual earnings.

A larger drop in the price of home appliances

We also performed an experiment of a 25-fold drop in q (as opposed to the 3-fold drop investigated by Experiment 5). This change is equivalent to an 8.3% annual drop in the relative price of home appliances, the quantity used in Greenwood et al. (2005). Although this experiment generates a larger rise in the fraction of two-earner couples, accounting for 14.5% of the rise found in the data, it still appears to be a less important channel than the closing of the gender gap.[35] Moreover, it implies even larger leisure increases among working and stay-home wives (11.7% and 15.34%), which were already counterfactually large in the original Experiment 5. By contrast, Table 2 reveals that stay-home wives, although experiencing the largest gains in leisure relative to other groups, enjoyed a much smaller rise of 6.9%. Recall that in the data, the average female leisure declined due to the composition effect. This experiment, however, generates a 10.8% increase in leisure for women on average.

Estimation of the Earnings Distribution using the Heckman selection model

It is also possible to use the Heckman selection model to estimate the parameters of the earnings distribution used in our experiments, instead of microfounding the selection rule with our calibrated model. The Heckman model is given by

$$y_i^* = \mathbf{x}_i\beta + u_i, \tag{8}$$

$$y_i = \begin{cases} y_i^* & \text{if} \quad z_i^* > 0 \\ 0 & \text{otherwise} \end{cases}, \tag{9}$$

where $z_i^* = \mathbf{w}_i\gamma + v_i$, and the error terms are jointly normally distributed:

$$\begin{bmatrix} u_i \\ v_i \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma^2 & \rho\sigma \\ \rho\sigma & 1 \end{bmatrix}\right).$$

Again, y_i^* denotes the log of the potential earnings of female i, while y_i denotes the observed log of her potential earnings. We use years of education, potential years of experience and potential years of experience squared as the components of $\mathbf{x}_i$. In addition to these, $\mathbf{w}_i$ includes the log of the husband's annual earnings. Using the maximum likelihood method, we estimate the model for each census year separately, finding that the marginal effect of the log of the husband's earnings on probability of participation is always significantly negative, and it becomes less negative over time (from -0.16 in 1959 to -0.05 in 1999). Using the predicted values to complete the sample of logs of earnings and exponentiating, we record $\{\mathbf{m}_t, \mathbf{S}_t\}_{t=1959,\ 1969,\dots}$, again obtaining a monotonic closing of the gender gap since 1959,[36] although the closing is less drastic than that obtained in Section V (62% vs. 83%).[37] It is possible to

[35] In fact, closing the gender gap by only half the amount investigated in Experiment 2.2 still accounts for a larger part of the aggregate increase in female participation (21%) than a 25-fold drop in q.

[36] $\{m_2/m_1\}_{1959,\dots,1999} = \{0.2528, 0.3064, 0.3429, 0.4268, 0.4803\}$.

[37] Using a similar Heckman selection model, Mulligan and Rubinstein (2008) (hereafter MR) argue that the observed closing of the gender gap between the late 1970s and the late 1990s was due to changing composition of women in the workforce from those with relatively low to those with relatively high unmeasured characteristics, and not due to changes in the degree of discrimination or measured characteristics. MR work with Current Population Survey sample of both single and married individuals. They regard a female as a participant if she works full time full year and use a two-step estimator of the Heckman model. To facilitate comparison of our results to those of MR, we performed a two-step estimation of the Heckman model, including the number of small children in addition to characteristics used in $\mathbf{w}_i$, including men in the second step and a female dummy in addition to characteristics used in $\mathbf{x}_i$, and using our measure of adjusted income. The inverse Mills ratio is set to zero for men, as all men in our sample work. The observed gender gap, $E(y_i^*|\text{working woman}) - E(y_i^*|\text{man})$, can be decomposed into gender difference in the valuation of observed characteristics, $E(\mathbf{x}_i\beta|\text{working woman}) - E(\mathbf{x}_i\beta|\text{man})$, valuation of the unobserved characteristics, $E(u_i|\text{working woman}) - 0$, and discrimination (coefficient on a female dummy). Comparing 1979 to 1999, the observed gap changed from -0.95 to -0.63. In agreement with the MR's findings, only a small part of the closing is explained by the change in the observed characteristics (from -0.024 to 0.0007), and the composition of females in the workforce evolves mostly due to changes in the unobserved characteristics (from 0.007 to 0.098), although the latter change is not from a negative to positive selection bias as found by MR, but rather from a nearly zero to positive

use these distribution parameters to repeat experiments 1-4 in the calibrated framework from Section IV; however, because of the different selection rules implied by that calibration and by the Heckman model, several important moments implied by the model, when the normal counterpart of $(\mathbf{m}_{1999}, \mathbf{S}_{1999})$ is used in place of the model's $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, will be different from the corresponding 1999 empirical moments. We thus recalibrate the model, i.e., search for the parameters that make the selection rule predicted by the model more consistent with that predicted by the Heckman model. We set the parameters of the earnings distribution to their empirical counterparts given by the Heckman model, normalize $q = 1$, and choose the remaining parameters $\theta, \rho, \lambda, \mu, v$ (starting with the initial guess from the calibrated values from Section IV) to match $P(2E)$, $E(X|2E)$, and $E(Y|2E)$. There is no overidentification problem, because the match is imperfect. This gives $\theta = 0.215$, $\rho = 0.181$, $\lambda = 0.49$, $\mu = 0.36$, $v = 0.35$. We find that changes in the mean vector alone account for 40% of the increase in the proportion of dual earner couples, and in a manner consistent with the cross-sectional participation[38] and leisure patterns. Again, most of the aggregate and group-specific increases in participation are driven by the closing of the gender gap, while the increase in purchasing power is the dominant force behind the leisure patterns. We find that the drop in the price of home appliances accounts for even less (3.8%) of the increase in participation.

Sensitivity with respect to λ, μ, v

Rather than setting λ, μ, v to particular values, we check for the robustness of our results with respect to the following 9 sets of calibrations. In the first three, we vary the value of λ among $\{0.3, 0.5, 0.7\}$, while allowing for the remaining parameters (except ρ and θ, which are set to their values from Table 4) to vary to match the moments from Table 3. The second set of three calibrations is similar, except that we now vary the value of μ among $\{0.2, 0.4, 0.5\}$, while allowing λ to change. In the last set, we vary the value of v among $\{0.2, 0.4, 0.5\}$. Because the moments of interest remain the same across these calibrations, the 1999 decision rule threshold varies little, which makes it unnecessary to reestimate the earnings distribution parameters with each recalibration. Overall, these calibrations assign a somewhat smaller role to changes in the entire distribution (accounting for 70-73% of the observed rise in participation). The change in the mean vector is still the main driving force, accounting for 51-56% of the rise in participation. The role of home appliances is again much smaller, accounting for 2-5% of the rise in participation. Implications for cross-sectional and leisure trends are similar to what was found in the main body of the paper.

VIII. Comparison to Jones et al. (2003) and Greenwood et al. (2005)

How do we reconcile the difference in the results found in Greenwood et al. (2005) and Jones et al. (2003), both closely related to our work? Before answering this question, we note that as we find all aspects of the earnings distribution other than the closing of the gender gap quantitatively less important

bias. Finally, in contrast to MR's findings, a large part of the closing of the gender gap is driven by declining discrimination (coefficient on a female dummy rises from -0.93 to -0.73). Recall that, in contrast to MR, we regard a person as a worker if he/she works a positive number of annual hours, and we work with the Census data on married couples with working husbands. Although outside the scope of this paper, further inquiry is desired to pinpoint the reasons behind the discrepancy between our results. We did find, however, that excluding the log of the husband's earnings from the selection equation (as is done in MR) generates results that are closer to MR, in particular, the reversal of the selection bias from negative in 1979 to positive in 1999.

[38] In comparison to Table 7, Experiment 2 here implies increase in female participation ranging from 117% in the lowest interval to 95% in the highest interval.

for female LFP, our results reinforce the importance of focusing on the factors investigated by Greenwood et al. (2005) and Jones et al. (2003).

Our results regarding the relative importance of the closing of the gender gap and the decline in prices of home appliances are more in line with those of Jones et al. (2003). We estimate changes in the earnings distribution from the census data, using our model to correct for the selection bias, and then examine the quantitative response of households to the estimated change. By contrast, Jones et al. (2003) ask whether a small exogenous change in the relative female-to-male wages (due to anti-discriminatory law) through its direct effect and its indirect effect on human capital accumulation, can generate trends in participation and the *observed* gender wage gap in a manner consistent with their empirical counterparts. In fact, because we focus on households' response to the closing of the gender gap (direct effect) without inquiring into the reasons underlying this change, and because we find that the response is quantitatively important, our work suggests that in order to further understand the rise in female LFP, it is important to disentangle the relative contributions of factors to the closing of the gender earnings gap (including human capital accumulation, changes in antidiscriminatory laws, etc., as discussed in Section I). Jones et al. (2003) is a step in this direction.

While we focus on 1959-1999, Greenwood et al. (2005) investigate 1900-1990.[39] It is difficult to interpret the results of their counterfactual experiments for the period 1960-1990 in isolation, because their baseline model, which incorporates both the closing of the gender gap and the decline in prices of home appliances, actually generates a *fall* in female participation for this period.[40] Still we attempt to provide some insight into why our findings regarding the relative importance of the fall in prices of home appliances and the narrowing of the gender gap are different. Our basic message here is that the discrepancy is not due to differences in the modeling choice but rather due to differences in the time series representing exogenous changes.

Importantly, when performing the experiment of the closing of the gender wage gap, Greenwood et al. (2005) purposefully shut down the effect of the increasing purchasing power associated with the catching up of female wages, thus slightly decreasing the importance of the closing of the gender gap. A more important factor driving the descrepancy, however, is that the gender gap time series used in the analysis of Greenwood et al. (2003) is taken from Goldin (1990). This time series reflects the relative *observed* earnings of full time full year workers. Correcting for the selection bias, as we show, results in a greater closing of the gender gap and hence a greater impact on participation, at least for the time period that we study.

Furthermore, the prices of home appliances experienced a much more dramatic fall in the first half of the 20th century, as opposed to the period considered in this paper and in Jones et al. (2003). For this reason, it is indeed likely that the decline in prices of home appliances was a more important force behind changes in female LFP in the first half of the 20th century. However, the 8.3% annual fall, used in Greenwood et al. (2003), still seems to overstate the relevant fall. First, the quality adjustment is made only for home appliances, although other consumption goods also experienced large quality improvements. Second, the price index time series is based on appliances that include TVs and VCRs,

[39] As stated above, we cannot extend our analysis to an earlier period, because the U.S. census provides income information for both spouses (not just the main respondent) beginning with the 1960 census.

[40] This fall is an artifact of transitional dynamics.

which would appear to complement leisure time, not substitute for time spent on home production. In fact, Vandenbroucke (2006) hypothesizes that the fall in prices of leisure goods (such as TVs and VCRs) should *reduce* market labor supply.

Are models in Jones et al. (2003) and Greenwood et al. (2005) suitable for evaluating implications of different factors against group-specific leisure trends and married female participation across groups differentiated by the husband's income? The answer is no. Jones et al. (2003) uses a representative agent model. Greenwood et al. (2005) assumes perfectly assortative matching (i.e., spouses have identical earning ability). In addition, it is assumed that the return to husbands' time allocated to home production is zero. Our results show that both the low correlation of spousal earnings ability and husbands' ability to contribute to home production are crucial for generating disaggregated participation and leisure trends consistent with the data.

IX. Conclusions

In this paper, we have quantitatively investigated the roles played by falling prices of home appliances and changes in *different aspects* of the earnings distribution in driving the dramatic increase in the fraction of two-earner couples observed in the data since 1959.

We built a model of heterogeneous households, capable of shedding light on the question of how different aspects of the earnings distribution (gender earnings gap, purchasing power, within-gender inequality, assortativeness of matching) may affect aggregate time allocation patterns. The rich heterogeneity of the model and its clear predictions regarding individual and aggregate leisure allow us to subject the factors under consideration to a number of empirical tests, previously unexplored in this context.

The main finding is that the closing of the gender earnings gap is the main force underlying the rise in the fraction of two-earner couples, accounting for over 70% of the observed rise. In close agreement with the data, we find that the closing of the gender earnings gap also induces a higher response from females with husbands in the upper range of the earnings distribution. Other aspects of the earnings distribution yield quantitatively small changes in female participation, although an increasing correlation appears to partly account for the greater rise in participation experienced by women with husbands at the top of the earnings distribution.

We also find that the rise in purchasing power associated with "the catching up" of female earnings generates a widespread diffusion of home appliances. This contrasts with the common belief that the home production revolution is a result of falling prices of home appliances alone. Through such an effect on home appliances, the narrowing of the gender gap also accounts for the rise in leisure of working and stay-home women, and the fall in the leisure of women on average, due to the change in the composition of working and stay-home females in the married population. It also accounts for the fall in the relative male-to-female leisure among 1M couples.

Finally, we find that the three-fold decline in the relative price of home appliances accounts for only 5% of the rise in female LFP, while implying counterfactually strong increases in average leisure trends.

Appendix

Relationship between the moments of the log-normal and the underlying normal distributions

The relationship between the moments of the log-normal distribution, $\mathbf{m} = \begin{bmatrix} m_1 \\ m_2 \end{bmatrix}$ and $\mathbf{S} = \begin{bmatrix} s_1^2 & s_{12} \\ s_{12} & s_2^2 \end{bmatrix}$, and the underlying normal distribution, $\boldsymbol{\mu} = \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix}$ and $\boldsymbol{\Sigma} = \begin{bmatrix} \sigma_X^2 & \sigma_{XY} \\ \sigma_{XY} & \sigma_Y^2 \end{bmatrix}$, is given by

$$\mu_X = \log\left(\frac{m_1^2}{\sqrt{m_1^2 + s_1^2}}\right),\ \mu_2 = \log\left(\frac{m_2^2}{\sqrt{m_2^2 + s_2^2}}\right), \tag{10}$$

$$\sigma_X^2 = \log\left(1 + \frac{s_1^2}{m_1^2}\right),\ \sigma_Y^2 = \log\left(1 + \frac{s_2^2}{m_2^2}\right),\ \sigma_{XY} = \log\left(1 + \frac{s_{12}}{m_1 m_2}\right). \tag{11}$$

U.S. Census Data

See Bar and Leukhina (2007) for a more detailed record of the data analysis. We work with the U.S. Census data available through IPUMs (2004). Although the census was conducted in 1960, 1970, ..., 2000, income and worktime questions refer to the previous year. Hence, the observations used in this paper are for 1959, 1969, ..., 1999. We match spouses using the household serial number and create a household-level dataset. Only non-farm married couples with each spouse between the ages of 25 and 64 are considered. We refer to an individual as an earner if he/she works a positive number of hours. No-earner couples and couples with a female as the only earner (around 6% of the original sample) are excluded.

We correct for the topcoding of all income types in 1959, 1969, 1979 only, because the topcoded observations in 1989 and 1999 are already replaced by the state mean or median. We infer labor income as the sum of wage and business incomes and convert all income observations into 1999 dollars using the 12 months averages of the seasonally adjusted consumer price index for the five census years: 29.17, 36.68, 72.58, 123.94, 166.58.

Worktime variables titled "actual weeks worked last year" and "usual weekly hours worked (last year)" are available since the 1980 census only. The 1960 and 1970 censuses, however, report information regarding the interval of weeks worked and the interval of weekly hours worked in the previous year. We must choose the appropriate midpoints for each of the intervals. To do so, using the information provided in the 1980 census on actual and intervalled hours and weeks, we compute gender-specific averages for each interval.

Table 8 reports a few descriptive moments of the married population sample that we work with.

Table 8. Selected Descriptive Moments

	1959	1969	1979	1989	1999
Fraction of two-earner couples	0.3297	0.4424	0.6145	0.7429	0.7621
E(weekly hours | 2E, male)	42.726	42.653	43.252	44.42	45.026
E(weeks | 2E, male)	48.73	49.954	48.699	48.525	48.925
E(annual hours | 2E, male)	2088.4	2134.9	2118	2169.4	2217.3
E(weekly hours | 1M, male)	43.209	43.072	43.639	44.722	45.609
E(weeks | 1M, male)	49.313	50.172	48.662	47.812	48.174
E(annual hours | 1M, male)	2136.9	2165.1	2135.8	2157.7	2215.4
E(weekly hours | male)	43.049	42.886	43.401	44.497	45.165
E(weeks | male)	49.121	50.076	48.685	48.341	48.746
E(annual hours | male)	2121	2151.8	2124.8	2166.4	2216.9
E(weekly hours | 2E, female)	33.967	33.465	34.388	35.939	36.967
E(weeks | 2E, female)	41.562	43.12	41.05	43.333	45.101
E(annual hours | 2E, female)	1449.2	1475.9	1459.8	1600.8	1708.1
E(labor earnings | 2E, male)	30304	41691	44044	45910	50097
E(labor earnings | 1M, male)	39864	52282	52702	54344	61827
E(labor earnings | male)	36712	47597	47382	48078	52888
E(labor earnings | 2E, female)	14577	18756	18213	22221	27303
CV(labor earnings | 2E, male)	0.6118	0.6212	0.7073	0.7979	0.9724
CV(labor earnings | 1M, male)	0.8042	0.7803	0.8179	0.9292	1.132
CV(labor earnings | male)	0.7812	0.7435	0.7698	0.849	1.0354
CV(labor earnings | 2E, female)	0.7229	0.7379	0.8414	0.8934	1.0295
E(labor earnings | 2E, female) / E(labor earnings | male)	0.3971	0.3941	0.3844	0.4622	0.5162

Proof of Proposition 1

Proof. (i) Suppose $\nu = 0$. Then no home production takes place, $k = l_f^2 = l_m^2 = 0$. The decision threshold can be derived analytically as

$$\begin{aligned} V_{1M}(w_m, L(w_m)) &= V_{2E}(w_m, L(w_m)), \\ \mu \log(w_m) &= \mu \log(w_m + L(w_m)) + (1-\mu)(1-\lambda)\log\left(1-\bar{l}_f^1\right), \\ L(w_m) &= w_m (A-1), \end{aligned}$$

where $A = \exp\left(\frac{-(1-\mu)(1-\lambda)\log\left(1-\bar{l}_f^1\right)}{\mu}\right)$.

(ii) Suppose $F(k,l) = k^\theta l^{1-\theta}$. Consider the maximization problem of a 2E household, given in (2). The objective function becomes

$$\mu \log(w_m + w_f - qk) + \nu\left[\theta \log k + (1-\theta)\log\left(l_m^2 + l_f^2\right)\right]$$
$$+(1-\mu-\nu)\left[\lambda \log\left(1-\bar{l}_m^1 - l_m^2\right) + (1-\lambda)\log\left(1-\bar{l}_f^1 - l_f^2\right)\right] + \kappa.$$

Because the marginal utility from k approaches infinity as k approaches zero, the optimal k is interior. It can be obtained by solving the first-order condition,

$$\frac{\mu q}{w_m + w_f - qk} = \frac{\nu\theta}{k},$$

which yields $k = \frac{\nu\theta}{\mu+\nu\theta}\left(\frac{w_m+w_f}{q}\right)$. Finally, because the optimal time inputs into home production are independent of w_m and w_f, the maximum value function associated with being a 2E household is of the form

$$V_{2E}(w_m, w_f) = \mu \log\left(w_m + w_f - \frac{\nu\theta(w_m + w_f)}{\mu + \nu\theta}\right) + \nu\theta \log\left(\frac{\nu\theta}{\mu+\nu\theta}\left(\frac{w_m + w_f}{q}\right)\right) + \kappa_{2E}, \text{ i.e.,}$$

$$V_{2E}(w_m, w_f) = (\mu + \nu\theta)\log(w_m + w_f) + \tilde{\kappa}_{2E}, \tag{12}$$

where κ_{2E} and $\tilde{\kappa}_{2E}$ are constants.

Similarly, the maximum value function associated with being a $1M$ household becomes

$$V_{1M}(w_m, w_f) = (\mu + \nu\theta)\log(w_m) + \tilde{\kappa}_{1M}. \tag{13}$$

Using the derivations in (12) and (13), we obtain the decision rule threshold $L(w_m)$,

$$\begin{aligned} V_{1M}(w_m, L(w_m)) &= V_{2E}(w_m, L(w_m)), \\ (\mu+\nu\theta)\log(w_m) + \tilde{\kappa}_{1M} &= (\mu+\nu\theta)\log(w_m + L(w_m)) + \tilde{\kappa}_{2M}, \\ L(w_m) &= w_m(B-1), \end{aligned}$$

where $B = \frac{\tilde{\kappa}_{1M} - \tilde{\kappa}_{2M}}{\mu+\nu\theta}$. ■

Proof of Proposition 2

Proof. In this version of the model, $F(k,l) = [\theta(Ak)^\rho + (1-\theta)l^\rho]^{1/\rho}$. We want to show that an increase in A is equivalent to a proportional decrease in q. Define the new variable $\tilde{k} \equiv qk$ and rewrite the problem of the two-earner household as follows:

$$\begin{aligned} V_{2E}(w_m, w_f) = \max_{c_m^1, c_m^2, c_f^1, c_f^2, \tilde{k}, l_m^2, l_f^2} & \lambda\left[\mu \log\left(c_m^1\right) + \nu \log\left(c_m^2\right) + (1-\mu-\nu)\log\left(1 - \bar{l}_m^1 - l_m^2\right)\right] \\ &+ (1-\lambda)\left[\mu\log\left(c_f^1\right) + \nu\log\left(c_f^2\right) + (1-\mu-\nu)\log\left(1-\bar{l}_f^1 - l_f^2\right)\right] \\ &\text{s.t. } c_m^1 + c_f^1 + \tilde{k} = w_m + w_f, \\ &c_m^2 + c_f^2 \leq \left[\theta\left(A\frac{\tilde{k}}{q}\right)^\rho + (1-\theta)\left(l_m^2 + l_f^2\right)^\rho\right]^{1/\rho}, \\ &0 \leq l_j^2 \leq 1 - \bar{l}_j^1, \quad j \in \{m, f\}. \end{aligned}$$

It is clear that A and q appear in this problem as a ratio. Hence, increasing A by a factor of ζ is equivalent to decreasing q by the same factor. The same proof applies to $V_{1M}(w_m, w_f)$. ■

Estimating $\{\mathbf{m}_t, \mathbf{S}_t\}_{t=1959,\ 1969,\ldots}$ *for Conducting Several Counterfactual Experiments*

We need to estimate the parameters of the potential earnings distribution to be used for conducting several of the experiments. We use a censored regression model consisting of a Mincer earnings equation and the participation rule implied by the calibrated model and the appropriate price of home appliances (see the main text).

It is instructive to derive the log-likelihood function. The contribution to the log-likelihood function made by observations with $y_i = 0$ is given by

$$\begin{aligned}\Pr(y_i = 0) &= \Pr(x_i\beta + u_i \leq p(z_i, \Omega)) \\ &= \Pr\left(\frac{u_i}{\sigma} \leq \frac{p(z_i,\Omega) - x_i\beta}{\sigma}\right) \\ &= \Phi\left(\frac{p(z_i,\Omega) - x_i\beta}{\sigma}\right),\end{aligned}$$

where Φ is the cumulative standard normal distribution function.

Conditional on $y_i > 0$, the density of y_i is $f(y_i|y_i > 0) = f(y_i)/\Pr(y_i > 0) = (1/\sigma)\,\phi((y_i^* - x_i\beta)/\sigma)/\Pr(y_i > 0)$. Thus, we obtain the log-likelihood function,

$$\begin{aligned}\log L &= \sum_{y_i>0} \log\left(\frac{1}{\sigma}\phi\left(\frac{y_i^* - x_i\beta}{\sigma}\right)\right) + \sum_{y_i=0}\log\left(\Phi\left(\frac{p(z_i,\Omega) - x_i\beta}{\sigma}\right)\right) \\ &= \sum_{y_i>0} -\frac{1}{2}\left[\log(2\pi) + \log\sigma^2 + \frac{y_i - x_i\beta}{\sigma^2}\right] + \sum_{y_i=0}\log\left(\Phi\left(\frac{p(z_i,\Omega) - x_i\beta}{\sigma}\right)\right).\end{aligned}$$

The independent variable in the Mincer equation is the natural log of female annual earnings. The dependent variables are the number of years of schooling, her experience, her experience squared, a dummy variable for the white race and a dummy variable for the black race. Next, we explain how "years of schooling" and the race dummies are constructed.

The U.S. Census distinguishes between 9 education codes. With each code, we associate a certain number of years of schooling. Our assumptions are summarized in the table below.

Table 9. Educational Record

	Education code (defined in the U.S. census)	Years of schooling
None or preschool	1	0
Grade 1, 2, 3, or 4	2	3
Grade 5, 6, 7, or 8	3	7
Grade 9	4	9
Grade 10	5	10
Grade 11	6	11
Grade 12	7	12
1 to 3 years of college	8	14
4+ years of college	9	18

The U.S. census defines 9 different race codes: (1) White, (2) Black, (3) American Indian, (4) Chinese,

(5) Japanese, (6) Other Asian or Pacific Islander, (7) Other race, n.e.c., (8) Two major races, (9) Three or more major races. We define a dummy variable "black_id" to equal 1 whenever code 2 is observed, and a dummy variable "white_id" equal 1 whenever codes 1, 4, 5, 6 are observed.

References

Aguiar, Mark and Erik Hurst (2006), "Measuring Trends in Leisure: Evidence from Five Decades of Time Use Surveys," *The Quarterly Journal of Economics*, forthcoming.

Attanasio, Orazio, Hamish W. Low and Virginia Sanchez-Marcos (2006), "Explaining Changes in Female Labour Supply in a Life-cycle Model," Working Manuscript, University of Cambridge

Browning, Martin and Metter Gortz (2006), "Spending time and money within the household," Working Manuscript, University of Oxford, Department of Economics.

Bar, Michael and Oksana Leukhina (2007), "To Work or Not to Work: Did Tax Reforms Affect Labor Force Participation of Secondary Earners?" University of North Carolina, Working Manuscript.

Bar, Michael and Oksana Leukhina (2007), Technical Appendix to "Accounting for Changes in Labor Force Participation of Married Women: The Case of the U.S. since 1959," online supplement, www.unc.edu/~oksana/Appendix2.pdf.

Caucutt, Elizabeth, Nezih Guner and John Knowles (2002), "The Timing of Births: A Marriage Market Analysis," *Review of Economic Dynamics*, 5: 815-855.

Costa L. Dora (2000), "From Mill Town to Board Room: The Rise of Women's Paid Labor", *The Journal of Economic Perspectives,* Vol. 14, No. 4., pp. 101-122.

Fernández, Raquel, Alessandra Fogli and Claudia Olivetti (2004), "Mothers and Sons: Preference Formation and Female Labor Force Dynamics," *Quarterly Journal of Economics*, 119:4, pp.1249-1299.

Galor, Oded and David N. Weil (1996), "The Gender Gap, Fertility, and Growth," *The American Economic Review*, Vol. 86, No. 3., pp. 374-387.

Gayle, George-Levi and Limor Golan (2006), "Estimating a Dynamic Adverse Selection Model: Labor Force Experience and the Changing Gender Earnings Gap," Working Manuscript, Tepper School of Business, Carnegie Mellon University.

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

Goldin, Claudia and Lawrence F. Katz, "The Power Of The Pill: Contraceptives And Women's Career And Marriage Decisions," *Journal of Political Economy*, 2002, v110, pp. 730-770.

Greenwood, Jeremy, Ananth Seshadri and Mehmet Yorukoglu (2005), "Engines of Liberation," *Review of Economic Studies,* Vol. 72, No. 1, pp. 109-133, 01.

Heckman, James J. (1979), "Sample Selection Bias as a Specification Error." Econometrica, 47(1), pp.153-61.

Steven Ruggles, Matthew Sobek, Trent Alexander, Catherine A. Fitch, Ronald Goeken, Patricia Kelly Hall, Miriam King, and Chad Ronnander. Integrated Public Use Microdata Series: Version 3.0

[Machine-readable database]. Minneapolis, MN: Minnesota Population Center [producer and distributor], 2004, http://www.ipums.org.

Jones, Larry. E., Rodolfo E. Manuelli and Ellen R. McGrattan (2003), "Why are Married Women Working So Much?" Federal Reserve Bank of Minneapolis, Research Department Staff Report 317.

Juhn, Chinhui and Kevin M. Murphy (1997), "Wage Inequality and Family Labor Supply," *Journal of Labor Economics*, Vol. 15, No. 1, pp. 72-97.

Knowles, John (2005), "Why are Married Men Working So Much?" University of Pennsylvania, Working Manuscript.

McGrattan, Ellen R., Richard Rogerson, and Randall Wright (1997), "An Equilibrium Model of the Business Cycle with Household Production and Fiscal Policy", *International Economic Review,* Vol 38, No.2, May, pp. 267-90.

Mulligan, B. Casey and Yona Rubinstein (2002), "Specialization, Inequality, and the Labor Market for Married Women," University of Chicago, Working Manuscript.

Mulligan, B. Casey and Yona Rubinstein (2008), "Selection, Investment, and Women's Relative Wages Over Time," forthcoming in *Quarterly Journal of Economics.*

Olivetti, Claudia (2006), "Changes in Women's Hours of Market Work: The Role of Returns to Experience," *Review of Economic Dynamics*, vol. 9(4), pages 557-587, October.

Vandenbroucke, Guillaume (2006), "A Model of the Trends in Hours," University of Southern California, Working Manuscript.

Figure 1

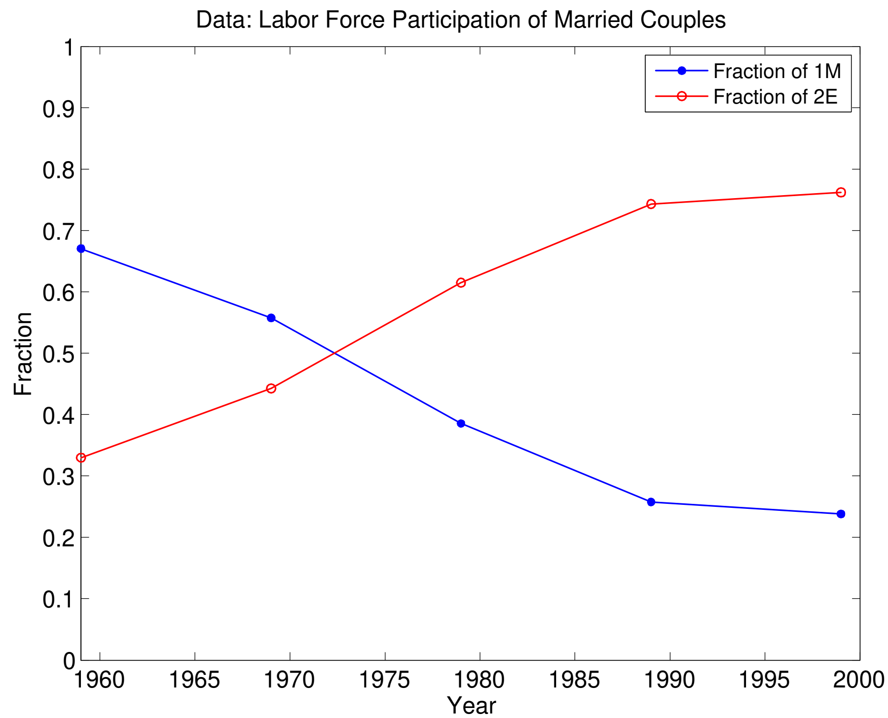

Figure 2

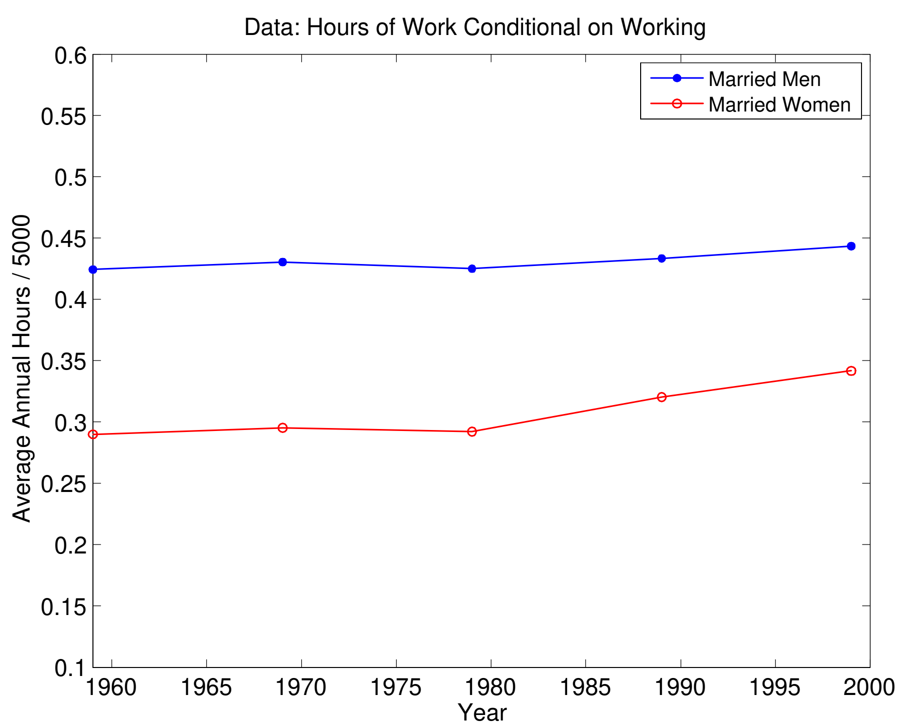

FIGURE 3

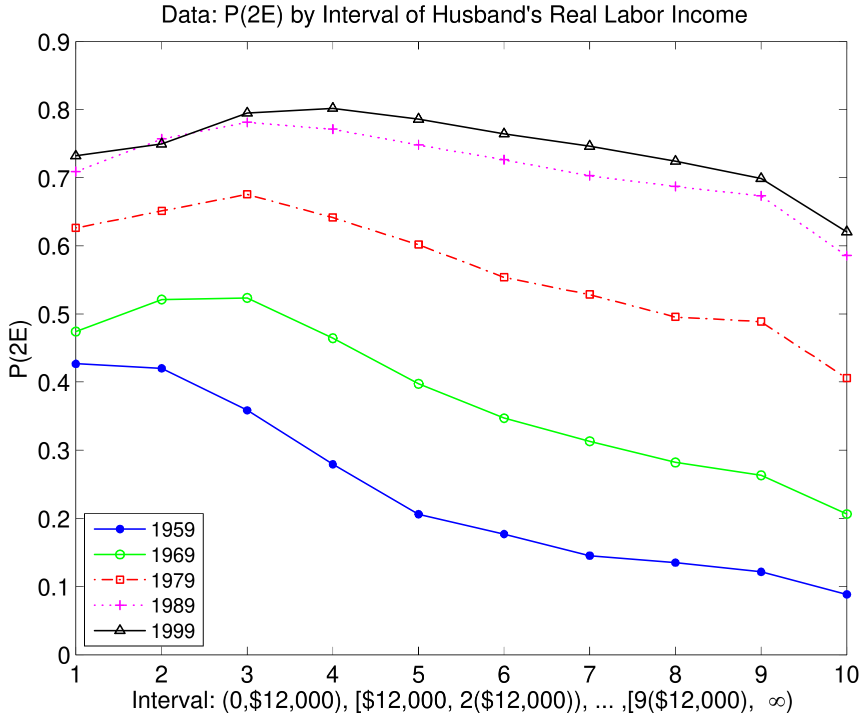

FIGURE 4

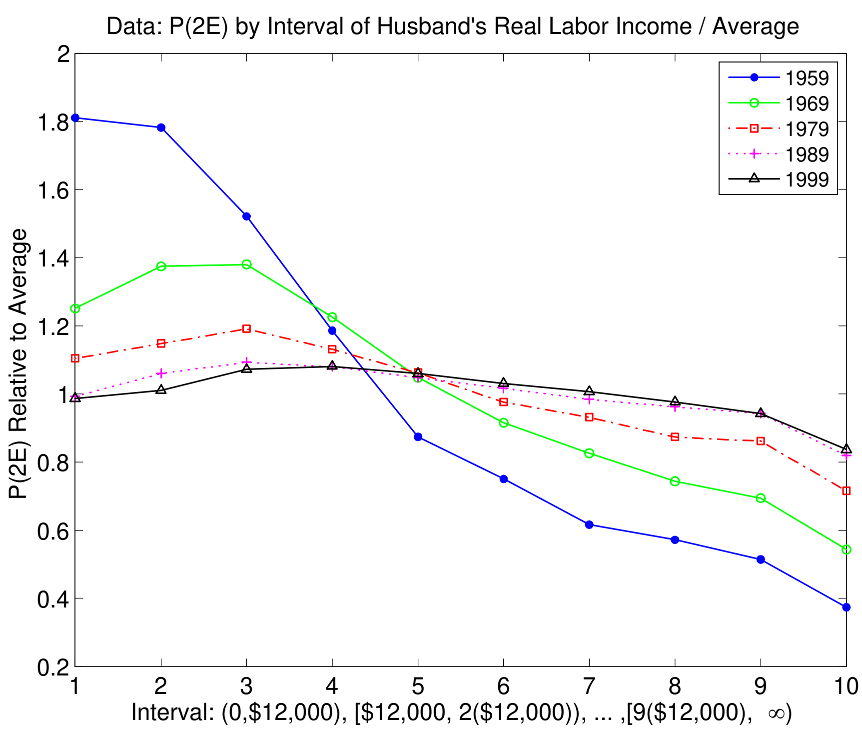

FIGURE 5

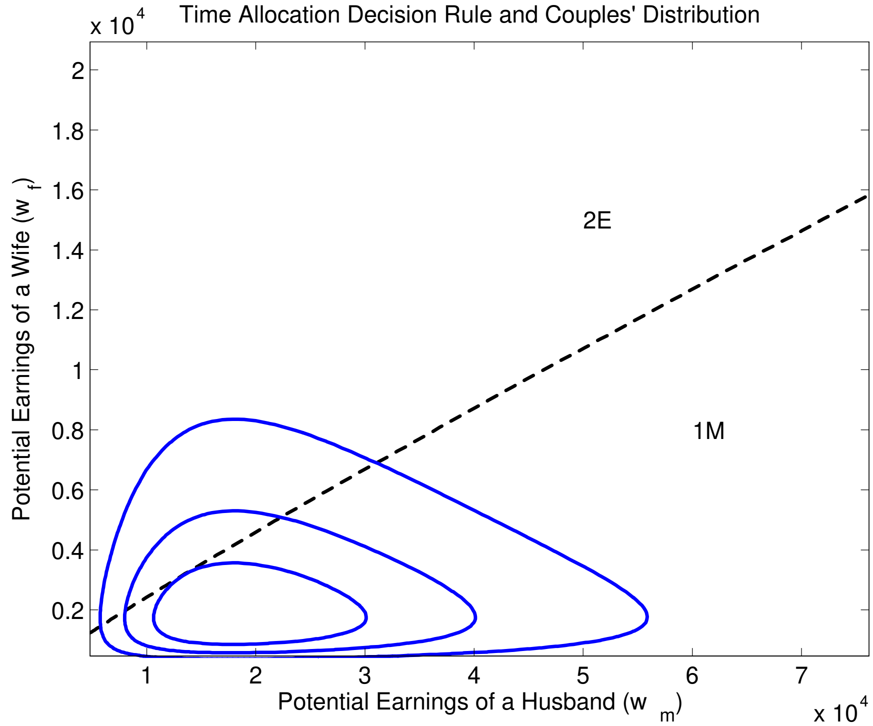

FIGURE 6

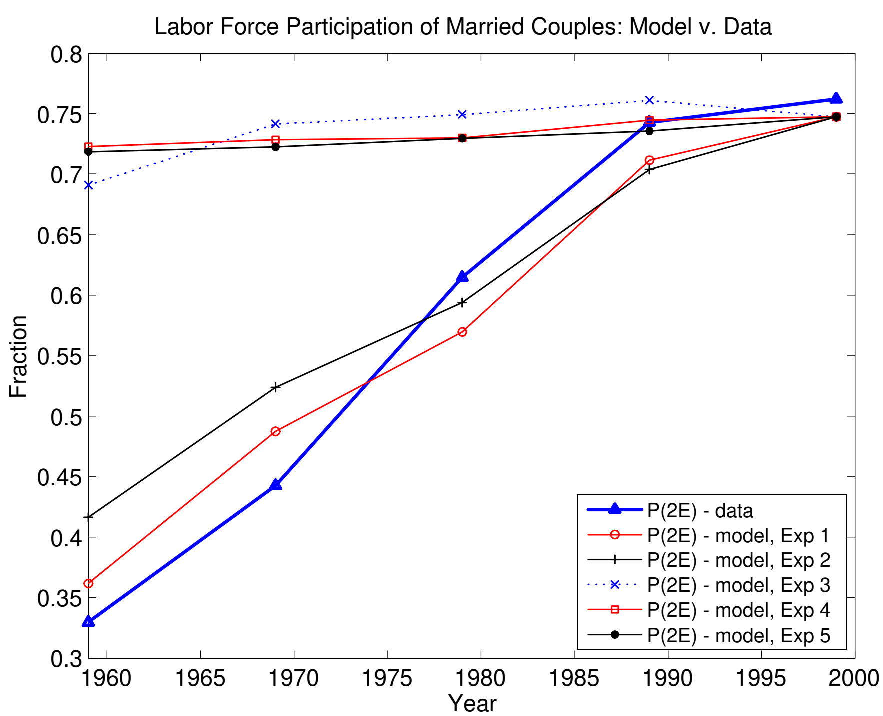

FIGURE 7

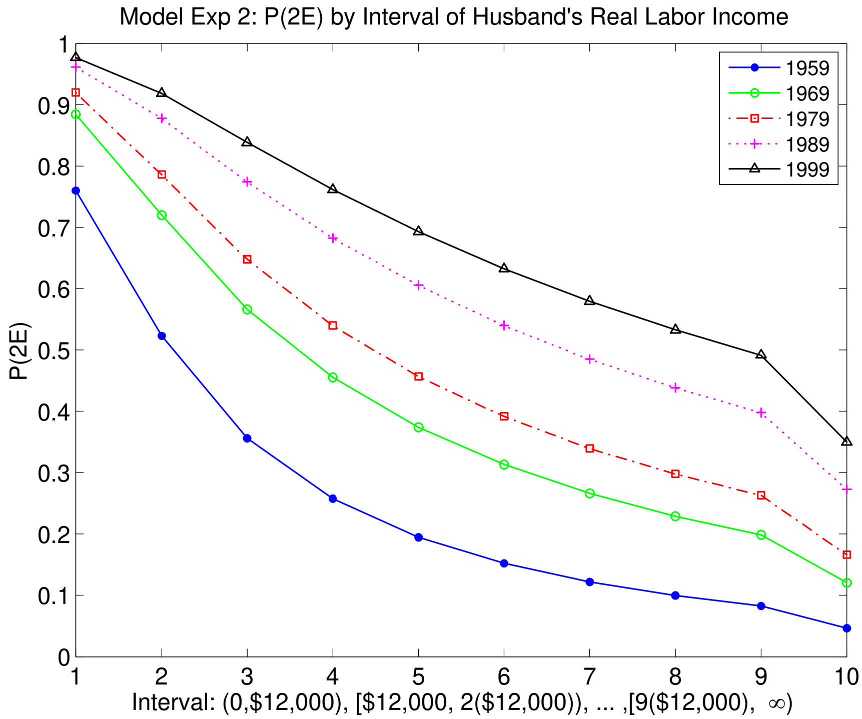

FIGURE 8

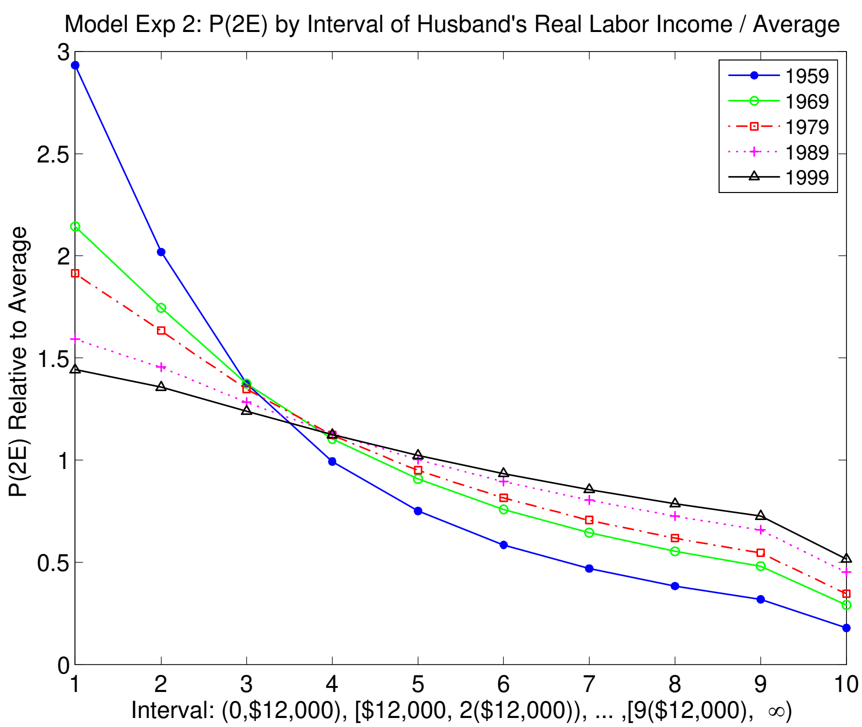

Figure 9

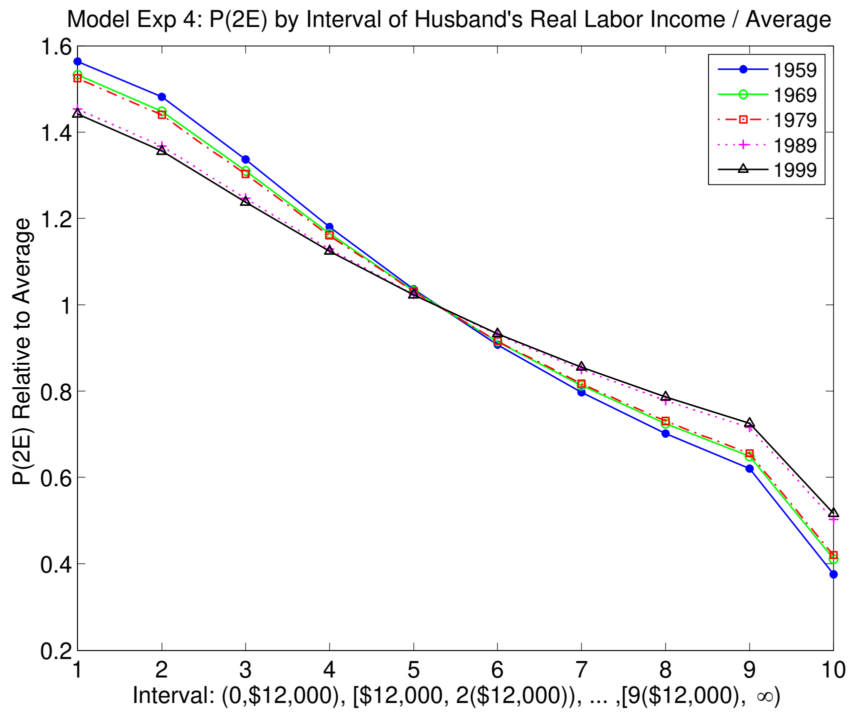

Figure 10

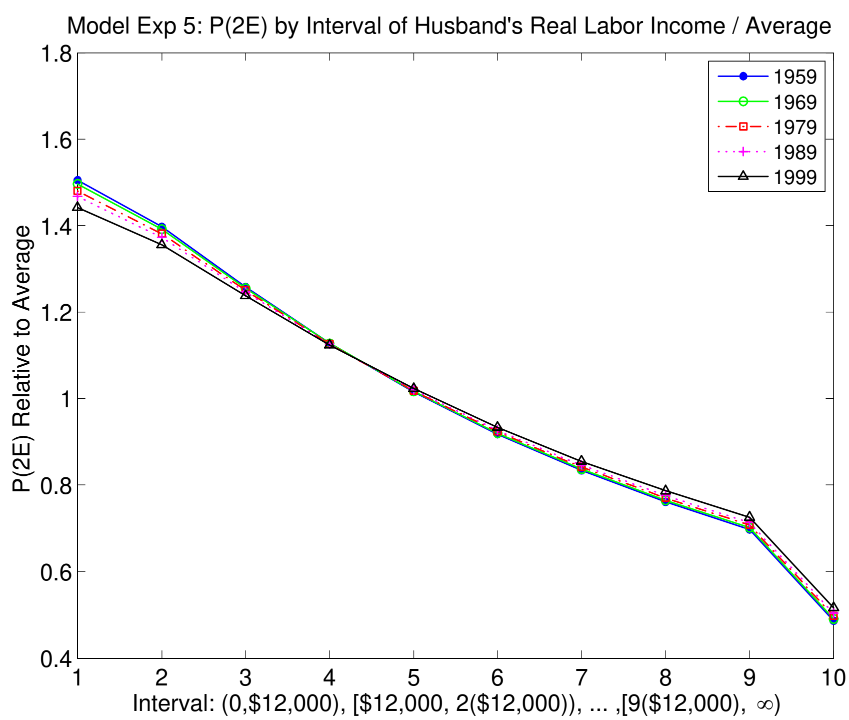

FIGURE 11

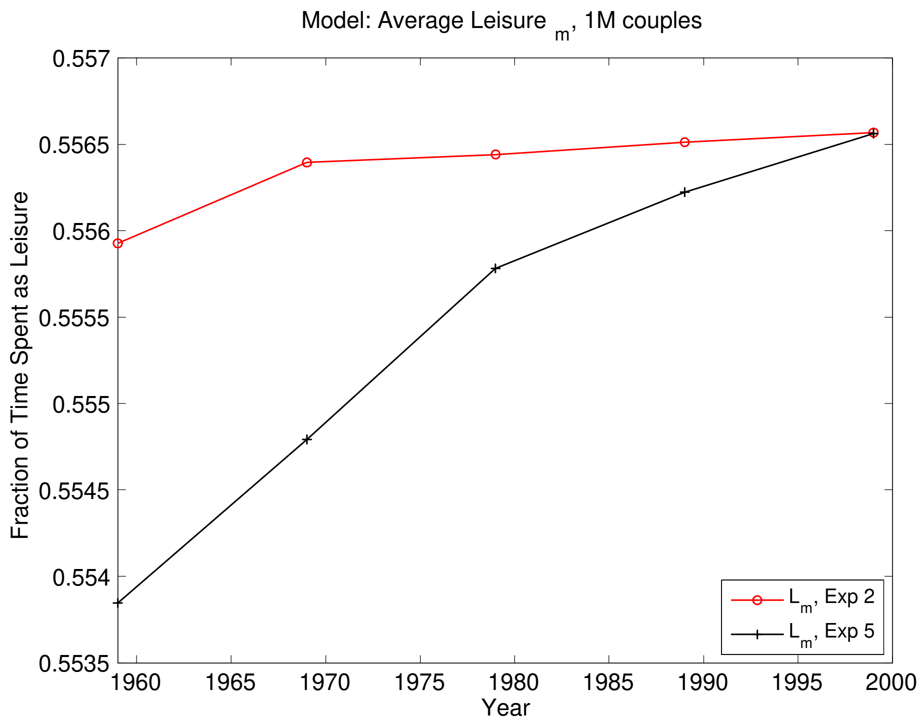

FIGURE 12

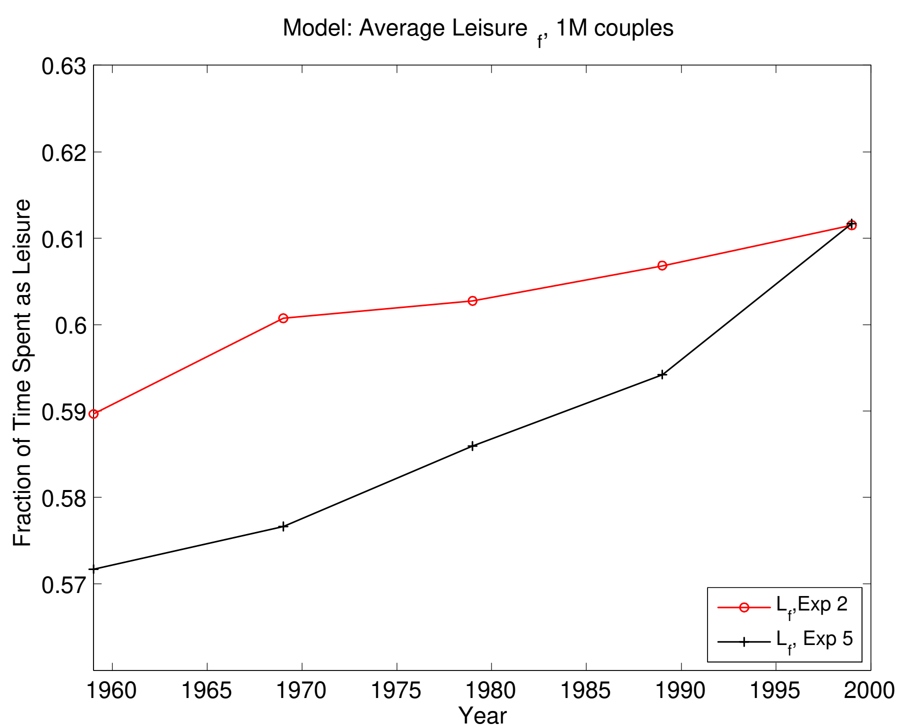

FIGURE 13

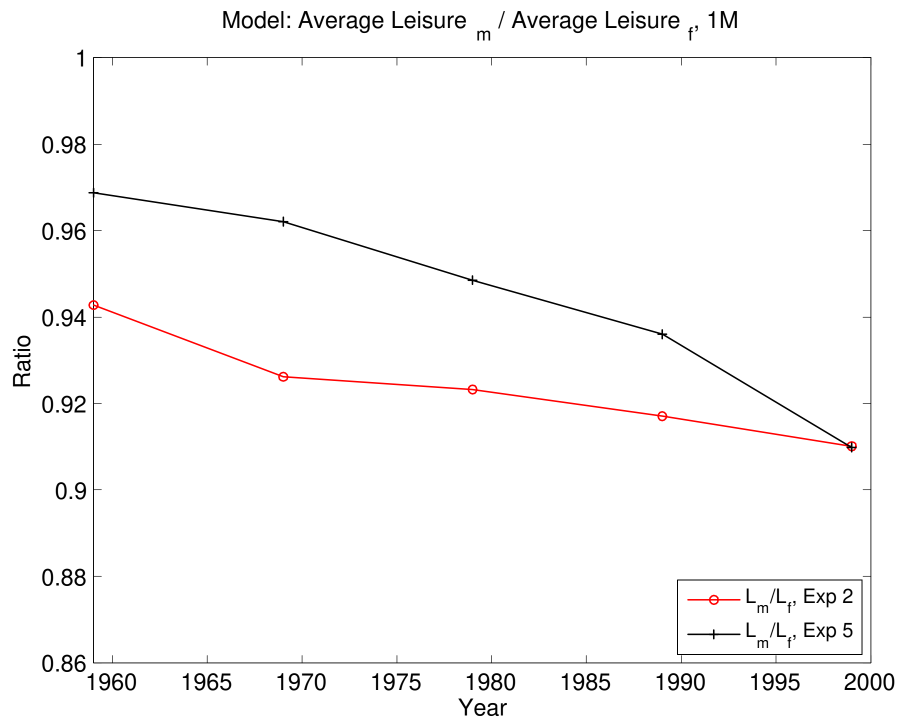

FIGURE 14

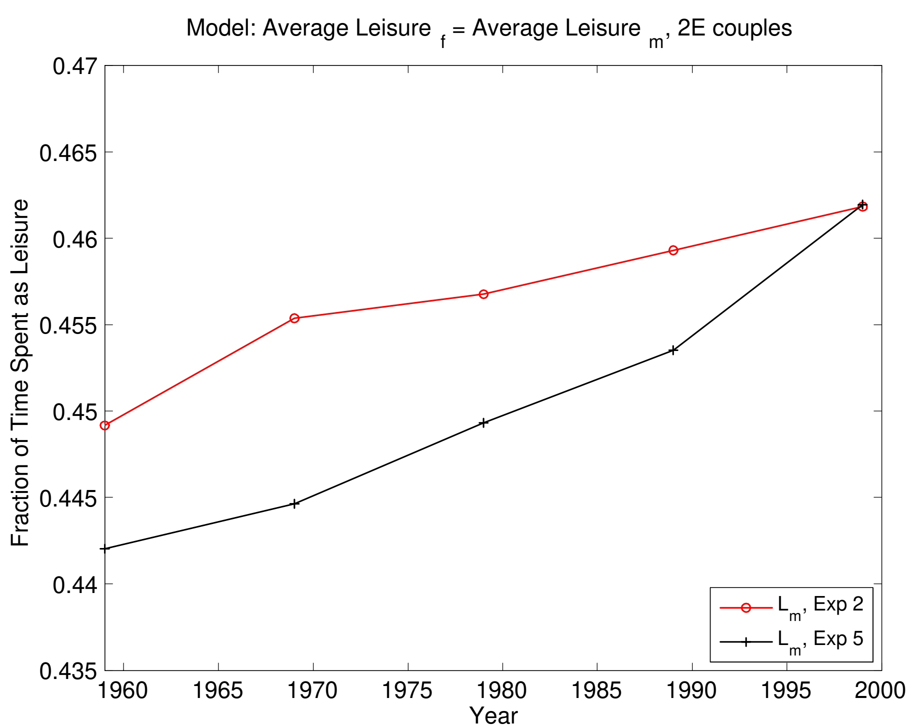

FIGURE 15

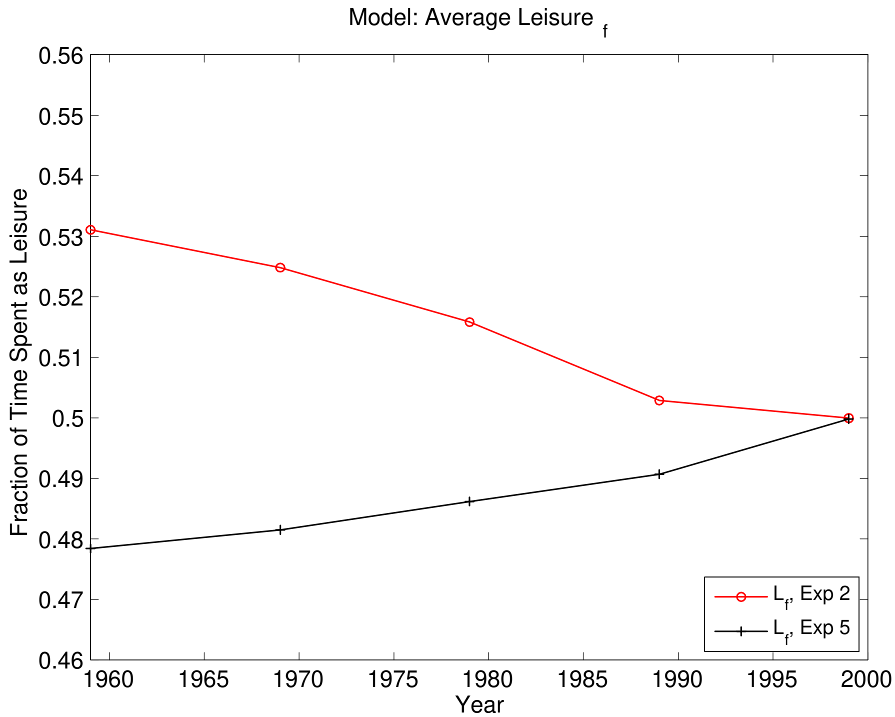

FIGURE 16