

Culture as Learning: The Evolution of Female Labor Force Participation over a Century

Raquel Fernández
NYU, CEPR and NBER

April 2007[†]

Women's labor force participation has increased dramatically over the last century. Why this has occurred has been the subject of much debate. This paper investigates the role of culture as learning in this change. To do so, it develops a dynamic model of culture in which individuals hold heterogeneous beliefs regarding the relative long-run payoffs for women who work in the market versus the home. These beliefs evolve rationally via an intergenerational learning process. Women are assumed to learn about the long-term payoffs of working by observing (noisy) private and public signals. They then make a work decision. This process generically generates an S-shaped figure for female labor force participation, which is what is found in the data. The S shape results from the dynamics of learning. When either small or large proportions of women work, learning is very slow and the changes in female labor force participation are also small. When the proportion of women working is close to 50%, rapid learning and rapid changes in female LFP take place. I show that such a model does a very good job in replicating the evolution of female labor force participation in the US over the last 120 years.

[†]An earlier version of the model and simulation in this paper were presented in my Marshall Lecture at the EEA, Vienna, August 2006. I thank Liz Potamites for excellent research assistance and the NSF for financial support.

1 Introduction

The vast increase in female labor force participation over the last century fundamentally transformed society. This increase was especially pronounced for married women, whose labor force participation increased from under 5% in 1880 to over 60% in 2000.¹ The pace of this change has been very uneven. As shown in Figure 1, female labor force participation increased very slowly from 1880 to 1920, then a bit more rapidly between 1920 and 1950, it accelerated between 1950 and 1990, and has since stayed relatively constant.² Many different factors have been put forward as potentially explaining at least some portion of this lengthy and uneven process of transformation. Prominent candidates include technological change in the workplace and in the household, medical advances, the introduction and dissemination of the oral contraceptive, decreases in discrimination, changed preferences/skills transmitted from working mothers to their sons, institutional changes in divorce law, and the greater availability of childcare.³

A popular alternative explanation (though less popular with economists) has been that changes in culture or social norms have been important contributors to this evolution in women's role in the market place.⁴ And, from multiple sources of evidence, it certainly appears that opinions about the role of women in the workplace have changed radically over time. Figure 2, for example, shows the evolution of the percentage of the population that answered the question "Do you approve of a married woman earning money in business or industry if she has a husband capable of supporting her?" affirmatively.⁵ In 1936 fewer than 20% of individuals sampled agreed with the statement; in 1998 fewer than 20% of individuals sampled disagreed with it.⁶

Merely pointing to the fact that society has changed the way in which it regards women, however, is not particularly enlightening. It begs the question as to why culture changed over these 120 years, and why these changes affected work behavior in such a gradual and uneven fashion. One might be tempted, as surely some are, to dismiss these shifts in beliefs as mere accompanying changes in the superstructure that simply reflect the real changes

¹These numbers are for the LFP of married women over the age of 15 (16 in 2000) (Goldin (1990) and Costa (2000)).

²The LFP numbers were calculated for US-born, married white women between the ages of 25-44.

³The classic source for an economic history of female labor force participation is Goldin (1990). For various explanations for this change see, among others, Goldin (1990), Galor and Weil (1996), Costa (2000), Goldin and Katz (2002), Jones, Manuelli, and McGrattan (2003), Fernández, Fogli, and Olivetti (2004), Greenwood, Seshadri, and Yorukoglu (2005), and Albanesi and Olivetti (2007).

⁴The reluctance of economists to believe in cultural explanations stems, in large part, from the absence of empirical evidence that convincingly isolates cultural influences from their economic and institutional environment. There has been recent progress in this area, however (see Fernández (2007a) and Guiso, Sapienza, and Zingales (2006) for reviews of this literature). For example, Fernández and Fogli (2005) show that the variation in the work behavior of second-generation American women can be explained, in part, by variation in past values of female LFP in their parents' country of origin. Fernández (2007) shows that the attitudes towards women's work in the parental country of origin has important explanatory value for second-generation women's work behavior in the US.

⁵The exact wording of this question varied a bit over time. See The Gallup Poll; public opinion, 1935-1971.

⁶For additional evidence that individual attitudes and work behavior are correlated see, for example, Levine (1993), Vella (1994), and Fortin (2005).

in material conditions.⁷ Viewed from this perspective, as women’s work behavior changed (because of, e.g., technological change), beliefs simply marched right along in step and changed with them. An alternative view of culture often provided in economic theory—that of a selection mechanism among multiple equilibria—likewise does not provide a useful framework in which to think about these questions. Without a theory of why culture changes, one is left only with sunspots causing a switch among equilibria.

This paper puts forward the idea that in some contexts it may be useful to think about cultural change as the evolution of beliefs that occur over time as part of a rational intergenerational learning process. In particular, for the case of female labor force participation, the observation of the S-shaped curve shown in Figure 1 may be a clue indicating that a process of information diffusion reminiscent of those common in technology adoption, though on a different time scale, may also be at work in this context.⁸

Where might learning play a role in the transformation of women’s role in market work? It is not an exaggeration to state that, throughout the last century, the payoff to women’s work has been fraught with uncertainty. The effect of working on a woman’s marriage, on her image and psyche, and especially on a woman’s children has long been the subject of heated debate. Even today, studies debate the effect of a working mother on a child’s intellectual achievements and emotional health and the research evidence is far from conclusive though it commands a great deal of public attention.⁹ For example, a recent finding by Belsky et al. (2007) of a positive relationship between day care and subsequent behavioral problems became headline news all over the US, despite its finding only a small quantitative effect.

In this paper I develop a simple model of women’s work decisions in which beliefs about the (long-run) payoff to working evolve endogenously over time. Using a framework similar to that in Vives (1993) and Chamley (2003), I assume that women receive a private signal about how costly it is to work (e.g., how negative the outcome is for one’s marriage, children, etc.) and that they also observe a noisy public signal about past beliefs concerning this value. This signal is a simple linear function of the proportion of women who worked in the previous generation. Women use this information to update their prior beliefs and then make a decision whether to work. In the next period, the next generation once again observes a noisy public signal generated by the decisions of women in the preceding generation, obtains their individual private signals, and makes work decisions. Thus, beliefs evolve endogenously via a process of intergenerational learning.

The model described above generically generates an S-shaped figure for female labor force participation. The S shape results from the dynamics of learning. When very few

⁷See, e.g., Guner and Greenwood (2006) who argue that the change in sexual mores reflect changes in the efficacy of contraception. This is no doubt a partial explanation but does not explain, for example, why attitudes towards homosexuality have changed.

⁸See, e.g., Griliches (1957), Foster and Rosenzweig (1995), Conley and Udry (2003), Munshi (2004, 2006), and Bandiera and Rasul (2006).

⁹See, for example, Bernal (2007), Keane and Bernal (2005), Hill, Waldfogel, Brooks-Gun and Han (2005), and Ruhm (2006) for reviews and recent findings of this literature. The level of attention devoted to evidence in this area is tremendous. As an interesting indication of culture, note that the effect of having a working father is rarely investigated.

women participate in the labor market (as a result of initial priors that are very negative about the payoff from working), learning is very slow since the noisiness of the signal swamps the information content given by differences in the proportion of women who would work in different states of the world. As the proportion of women who work increases and beliefs about work become more positive, the information content in the signal improves. Once a large enough proportion of women work though, once again, the informational content in the public signal falls since the differences in the proportion of women who would work under different states of the world is swamped by the noise.

To quantitatively evaluate the potential ability of such a model to explain the evolution of female LFP, I first calibrate a version of the model without any evolution of beliefs to a few female LFP statistics for the year 2000. I show that such a model performs very badly and that it grossly overestimates the proportion of women who would have worked for basically every time period. I then introduce learning as discussed above and show that this simple modification greatly improves the capacity of the model to fit the data. The model indicates that both the dynamic paths of beliefs and earnings played an important role in the transformation of women's work.

2 A Simple Model of a Woman's Work Decision

We start with a very simple model of a woman's work decision. We include the main variables that should play a role in this decision, namely her consumption possibilities as a function of her work decision and her disutility from working. As we are interested in the difference in the long-run payoffs from working versus not working, we view the disutility from working as stemming not only from labor-leisure preferences, but also from what might happen to her marriage or her children as a result of decision. In this first model, we assume that the difference in disutility is known and constant. What is critical is that its expected value does not evolve endogenously over time.

A woman makes her work decision to maximize:

$$U_i(w_f, w_h, v_i) = \frac{c^{1-\gamma}}{1-\gamma} - \mathbf{1}.v_i \quad (1)$$

where $\gamma \geq 0$ and $\mathbf{1}$ is an indicator function that takes the value one if she works and zero otherwise. A woman's consumption is the sum of her earnings, w_f , (which are positive only if she works) and her husband's earnings, w_h . Husbands are assumed to always work, i.e.,

$$c = w_h + \mathbf{1}.w_f \quad (2)$$

The disutility of work, v_i , is assumed to consist of two parts,

$$v_i = \beta + l_i \quad (3)$$

where the first component β is common to all women and the second component is idio-

syncratic and normally distributed, $l \sim N(0, \sigma_l^2)$. We assume that there is a continuum of agents of mass one in each period.

Clearly, a woman will work iff

$$\frac{1}{1-\gamma}[(w_{ht} + w_{ft})^{1-\gamma} - w_{ht}^{1-\gamma}] - \beta \geq l_i \quad (4)$$

and the aggregate proportion of women who work at time t is given by

$$\omega_t = G(l_t^*; \sigma_l) \quad (5)$$

where $G(\cdot)$ is the cdf of the l distribution and l_t^* is the value of l such that (4) is a strict equality. As earnings evolve, so will l^* . Note that $\frac{\partial l^*}{\partial w_f} > 0$ whereas $\frac{\partial l^*}{\partial w_h} < 0$, i.e., the proportion of women working increases with their own earnings and decreases with those of their husband's.

3 The Simple Work Model with Learning

We next incorporate beliefs in the simple model above and modify (1) to reflect uncertainty in the payoff to working. In particular, women are assumed to be uncertain about the common value of the disutility of labor, β , e.g., they are unsure how bad working will be for their marriage, children, identity, etc. This is not something that can be learnt by entering the labor market for a short period of time or by experimenting but rather reveals its results over a lifetime. For simplicity we assume that β can take on only two values, high (H) and low (L), i.e., $\beta \in \{\beta_H, \beta_L\}$. Note that β_L is the "good" state of nature in which working is not so costly, i.e., $\beta_H > \beta_L > 0$. An individual woman now maximizes her expected utility, i.e.,

$$\frac{e^{1-\gamma}}{1-\gamma} - \mathbf{1} \cdot (E_{it} v_i) \quad (6)$$

Our model will incorporate two sources of learning. One is an individual source whereby a woman receives a noisy signal regarding the true value of β , β^* . The second is an intergenerational source whereby all women in generation t observe a noisy signal of the decisions taken by women in the preceding generation. The exact mechanics are made more precise below.

Consider a woman in period t who has a prior belief about β^* as summarized in the log likelihood ratio (LLR) $\lambda_t = \ln \frac{Pr(\beta^* = \beta_L)}{Pr(\beta^* = \beta_H)}$. Prior to making her work decision, she receives a private signal s_t regarding β^* . This signal can be thought of as arising from many sources, e.g., the scientific literature that existed at that time regarding the effect of mother's work on children or families. The private signal is assumed to have a normal distribution and is defined by

$$s_t = \beta^* + \epsilon_t \quad (7)$$

where $\epsilon_t \sim N(0, \sigma_\epsilon^2)$ and its cumulative distribution function is denoted by $F(\cdot; \sigma_\epsilon)$.

After receiving her private signal, s , each woman i updates her prior belief accordingly using Bayes' rule, resulting in a new LLR, $\lambda_{it}(s)$, given by

$$\begin{aligned}\lambda_{it}(s) &= \lambda_t + \ln \left(\frac{Pr(s|\beta^* = \beta_L)}{Pr(s|\beta^* = \beta_H)} \right) \\ &= \lambda_t + \frac{\beta_L - \beta_H}{\sigma_\epsilon^2} (s - \bar{\beta})\end{aligned}\tag{8}$$

where $\bar{\beta} = (\beta_L + \beta_H)/2$.¹⁰ Note that $\frac{\partial \lambda_{it}(s)}{\partial s} < 0$ since higher signals increase the likelihood that the true value of β is β_H .

Assume that women have a common prior in period t , λ_t (we shall assume that they start with a common prior in period zero and show that, thereafter, each generation continues to have a common prior). What proportion of women will choose to work that period? A woman will work in period t iff

$$\frac{1}{1-\gamma} [(w_{ht} + w_{ft})^{1-\gamma} - w_{ht}^{1-\gamma}] - l_i \geq E_{it}(\beta)\tag{9}$$

that is, the benefit of working must exceed the expected value of the disutility of work. For notational ease, we henceforth denote $\frac{1}{1-\gamma} [(w_{ht} + w_{ft})^{1-\gamma} - w_{ht}^{1-\gamma}]$ by $W(w_{ht}, w_{ft})$.

Note first that given $\{\beta_H, \beta_L\}$ and current earnings, irrespective of their beliefs, women with very low l 's ($l \leq \underline{l}$) will always work and women with very high l 's ($l \geq \bar{l}$) will never work, where

$$\underline{l}(w_{ht}, w_{ft}) \equiv W(w_{ht}, w_{ft}) - \beta_H\tag{10}$$

$$\bar{l}(w_{ht}, w_{ft}) \equiv W(w_{ht}, w_{ft}) - \beta_L\tag{11}$$

Next, for each women of type l_j , $\underline{l} < l_j < \bar{l}$, we can solve for the critical value of the private signal $s_j^*(\lambda)$ such that, for any $s \leq s_j^*$, given her prior belief λ , she would be willing to work. Let $p = Pr(\beta^* = \beta_L)$ and let p_j^* be the critical probability such that a woman of type l_j is indifferent between working and not, i.e.,

$$p_j^* \beta_L + (1 - p_j^*) \beta_H = W(w_{ht}, w_{ft}) - l_j\tag{12}$$

Using (10), we obtain

$$p_j^* = \frac{l_j - \underline{l}}{\beta_H - \beta_L}\tag{13}$$

and hence, $\frac{p_j^*}{1-p_j^*} = \frac{l_j - \underline{l}}{\beta_H - \beta_L + l - l_j} = \frac{l_j - \underline{l}}{\bar{l} - l_j}$.

Thus, the critical value of the private signal a woman of type j must receive in order to work, given a prior belief of λ_t , is given by

$$\lambda_t(s_j^*) = \lambda_t + \frac{\beta_L - \beta_H}{\sigma_\epsilon^2} (s_j^* - \bar{\beta}) = \ln \left(\frac{l_j - \underline{l}}{\bar{l} - l_j} \right)\tag{14}$$

¹⁰To obtain (8) one uses the fact that $Pr(s|\beta)$ is equal to the probability of observing a signal s from a normal distribution $N(\beta, \sigma_\epsilon^2)$.

and hence

$$s_j^*(\lambda_t; \beta_H, \beta_L, w_h, w_f, \sigma_\epsilon^2) = \bar{\beta} + \frac{\sigma_\epsilon^2}{\beta_H - \beta_L} \left(\lambda_t + \ln \left(\frac{\bar{l} - l_j}{l_j - \underline{l}} \right) \right) \equiv s_j^*(\lambda_t) \quad (15)$$

We can conclude from the above derivation that the proportion of women of type l_j , $\underline{l} < l_j < \bar{l}$, that will work in time t if the true state of nature is β , $\omega_{jt}(\beta)$, is given by the proportion of this type that receives signals lower than s^* , i.e.,

$$\omega_{jt}(\beta) = F(s_j^*(\lambda_t) - \beta; \sigma_\epsilon) \quad (16)$$

Thus, the total proportion of women that will work in period t if the state of nature is β , is given by:

$$\omega_t(\beta) = G(\underline{l}) + \int_{\underline{l}}^{\bar{l}} F(s_j^*(\lambda_t) - \beta; \sigma_\epsilon) g(l_j) dl_j \quad (17)$$

where $g(\cdot)$ is the pdf of the l distribution $G(\cdot)$.

3.1 Intergenerational Transmission

If the next generation ($t + 1$) of women were able to observe the aggregate proportion of women who worked in period t , they would be able to back out the true state of nature, β^* , as a result of the law of large numbers. While assuming that women do not know how many women worked in the previous generation seems extreme, the notion that this knowledge is completely informative seems equally implausible. We employ instead the conventional tactic in this literature (e.g. Vives (1993)) and assume that women are able to observe a noisy function of the aggregate proportion of women worked.¹¹ In particular, we assume that women observe a noisy signal, y_t , of ω_t , where

$$y_t(\beta) = \omega_t(\beta) + \eta_t \quad (18)$$

and where $\eta_t \sim N(0, \sigma_\eta^2)$ with a pdf denoted by $h(\cdot; \sigma_\eta)$.

Furthermore, we assume that women in generation $t + 1$ inherit the common prior of generation t , λ_t . This prior is updated with the information contained in y_t , which generates λ_{t+1} (i.e., the common prior that will be inherited next generation) and also individually with each individual's private signal (generating $\lambda_{it+1}(s)$). Alternatively, instead of assuming women inherit λ_t which they update with the information contained in y_t , we can assume that women observe the entire history of y_s , $s = 0, 1, 2, \dots, t$. Thus, in this setup, the common updated belief λ_τ can be thought of as the shared culture of generation

¹¹One way to think about this is that agents know the proportion of women that worked in their community (along with incomes of married men and women there), but are uncertain about the value of these variables across all other communities in the US. Thus, even if they had information about the aggregate proportion of female LFP, without the disaggregated information (at the community level) of incomes of wives and husbands as well as the community levels female LFP, they are unable to back out the true state of nature. An alternative assumption, pursued in Fernández and Potamites (2007), is that women know the work behavior of a small number of their social circle. This assumption of a limited (discrete) sample yields similar results.

τ , with the individual deviations around around λ_τ (given by the normal distribution of $\lambda_{i\tau}(s)$) constituting the distribution of beliefs induced by different individuals experiences (i.e., different realizations of s).

Thus, given a common belief λ_t , after observing last period's signal of aggregate female LFP, y_t , Bayes' law implies an updated common belief for generation $t + 1$ of:

$$\begin{aligned}\lambda_{t+1} &= \lambda_t + \ln \frac{h(y_t|\beta^* = \beta_L)}{h(y_t|\beta^* = \beta_H)} \\ &= \lambda_t + \frac{\omega_t(\beta_L) - \omega_t(\beta_H)}{\sigma_\eta^2} \left(y_t - \frac{\omega_t(\beta_L) + \omega_t(\beta_H)}{2} \right)\end{aligned}\tag{19}$$

Note that (19) is the law of motion of aggregate beliefs (culture) for the economy. Figure 3 summarizes the time line for the economy.

3.2 Some Properties of the Learning Model

The learning model has several important properties that will be prove useful when we try to match the data in Figure 1. Note first that beliefs in this model are unbounded. Hence, in the long run beliefs converge to the truth. As over time female LFP has been increasing, this implies that it is likely that $\beta^* = \beta_L$ and we shall henceforth assume that this is the case.

A key characteristic of this model is that it will naturally generate an S-shaped curve. To see why, note that given $\beta^* = \beta_L$, we can rewrite (19) as

$$\lambda_{t+1} = \lambda_t + \frac{\omega_t(\beta_L) - \omega_t(\beta_H)}{\sigma_\eta^2} \left(\eta_t + \frac{\omega_t(\beta_L) - \omega_t(\beta_H)}{2} \right)$$

Hence, the change in the LLR is increasing in the difference between the aggregate proportion of women who work when $\beta^* = \beta_L$ relative to when $\beta^* = \beta_H$. To see when this difference will be greater, we can start by examining, for a given l_j type, how the proportion that works varies over the two states of nature. Taking the derivative of $\omega_{jt}(\beta)$ with respect to β yields

$$\frac{\partial \omega_{jt}}{\partial \beta} = -\frac{1}{\sqrt{2\pi}\sigma_\epsilon} \exp - \left(\frac{(s_j^*(\lambda_t) - \beta)^2}{2\sigma_\epsilon^2} \right)\tag{20}$$

Thus, if the critical signal $s_j^*(\lambda_t)$ is far from β (either to the far right or far left of it), (20) will be small in absolute value. This means that the difference in the signal $y_t(\beta)$ across the two states will be swamped by the aggregate noise term η_t . Thus, the amount of intergenerational updating will be small and hence the change in the proportion of women who work over time, *ceteris paribus*, will likewise be small. This property of the normal distribution is illustrated in Figure 4. As can be seen in the figure, when s^* is far from extreme in value, the difference in proportion of women who work will be large (the shaded area). The opposite is true at $s^{*'}$.

Note that a similar conclusion holds once we aggregate over the l_j types. Taking the

derivative of (17) we obtain

$$\frac{\partial \omega_t}{\partial \beta} = -\frac{1}{\sqrt{2\pi}\sigma_\varepsilon} \int_{\underline{l}}^{\bar{l}} \exp\left(-\frac{(s_j^*(\lambda_t) - \beta)^2}{2\sigma_\varepsilon^2}\right) g(l_j) dl_j \quad (21)$$

Thus, if the critical signal $s_j^*(\lambda_t)$ is, for the average individual in (\underline{l}, \bar{l}) , far from β , (21) will be small in absolute value, intergenerational updating will be small, and the evolution over time will be slow.

Under what circumstances will the critical signal take an extreme value and lead to slow intergenerational learning? As can be seen from expression (15), for all l_j types, this occurs when λ_t is either very small or very large. To understand why this is the case, note that when an individual assigns a very low (high) probability to $\beta^* = \beta_L$, it takes a very low (high) realization of the private signal to update a woman's belief sufficiently to convince her to work (not work). The difference across states of nature, however, in the probability of obtaining such extreme values of s is very small. Once again, this is the result of the shape of the cdf of a normal distribution which is very flat for extreme (positive or negative) values of s .

It follows from the logic above that if parameter values are such that few women would choose to work if they assigned a low probability to $\beta^* = \beta_L$ whereas many women would choose to work if they assigned a high probability to this state, then the amount of intergenerational learning that will happen when female LFP is either very low or very high will be small as the aggregate noise term dominates in (19) and hence the period to period increase in female LFP will be likewise small. At these extremes, learning occurs, but it takes time. When, instead, the difference in the proportion of women who choose to work across states is large, i.e., when s^* is close to $\bar{\beta} \equiv \frac{\beta_H + \beta_L}{2}$, then observing the aggregate signal tends to be informative, intergenerational learning is rapid, and the period to period change in female LFP will be large. Putting these statements together, it is easy to see that this model will tend to generate an S-shaped curve with a slow evolution of female LFP at the beginning, followed by rapid increases over time, and then tapering off again to small increases in female LFP until there is no more learning. At that point, any further changes in female LFP result solely from changes in earnings.

4 Empirical Analysis

In this section we examine the degree to which incorporating intergenerational learning improves upon our simple model of labor supply in which beliefs did not evolve over time (we will call this the "earnings only" model). We first calibrate the three parameters of the earnings only model to key statistics of female LFP in the year 2000 to see how well it is able to replicate the aggregate dynamics of female LFP and then perform a similar exercise for the learning model allowing the remaining parameters to be chosen so as to best fit the historical LFP series.

4.1 The Data

Our model requires data on earnings for men and women as well as the historical data on female LFP. It is difficult to obtain earnings data prior to 1940. We rely on data provided in Goldin (1990). Goldin uses a variety of sources (Economic Report of the president (1986), Current Population Reports, P-60 series, and the U.S. Census among others) to calculate median earnings for men and women. We use the data for white men and women.¹² As Goldin does not provide data for earnings in 1880 and 1910, we construct these using a cubic approximation with the data from 1890 -1930 (inclusive).

After 1940 we use the 1% IPUMS samples of the U.S. Census for yearly earnings (incwage) and calculate the median earnings of white 25-44 year old men and women who were working full time (35 or more hours a week) and year round (40 or more weeks a year) and were in non-farm occupations and not in group quarters. As is commonly done, we exclude observations that report weekly earnings less than half the minimum wage. We use the half the nominal minimum wage times 35 hours a week as our cutoff for weekly wages and calculate nominal weekly wages by dividing total wage and salary income last year by weeks worked last year.¹³

Figure 5 shows the evolution of female and male earnings over the 120 year period 1880-2000 (with earnings expressed in 1967 dollars). In order to compare our data with Goldin's we also plot Goldin's figures (which continue to 1980). These numbers are show in blue. The only significant difference is with male earnings in 1950.¹⁴

The LFP numbers in 1880 and for the years 1900-2000 are for married white women with spouse present between the ages of 25 and 44 who report being in the labor force (non-farm occupations and non-group quarters). We use the 1% IPUMS samples for 1880, 1900-1920, 1940-1950, 1980-2000, and the 0.5% sample in 1930 and the 1970 1% Form 2 metro sample. For 1890, we use the midpoint between 1880 and 1900.¹⁵

We calibrate both models to match female LFP in the year 2000 as well as the own and cross-elasticity of labor force participation with respect to earnings in that same year. We use estimates of the latter computed by Blau and Kahn (2006). The authors use the March CPS 1999-2001 and compute married women's own-wage and husband's-wage elasticities along the extensive margin restricting their sample to married women of age 25-54 (with spouses in the same age range).¹⁶ We use the results obtained from the basic

¹²See Goldin (1990) pages 64-65 and 129 for greater detail about the earnings construction for various years. We look at white women as black women have had a different LFP trajectory.

¹³See, for example, Katz and Autor (1999). This procedure is somewhat more problematic for the decades 1940-1960, when the federal minimum wage did not apply to all workers (prior to the 1961 amendment, it only affected those involved in interstate commerce). Nonetheless, as in Goldin and Margo (1992), we use the same cutoff rule as a way to eliminate unreasonably low wages. Note that since we are calculating median earnings, we do not have to concern ourselves with the top-coding in the Census.

¹⁴Goldin's 1950 number is from the Current Population Reports, series P-60 number 41 (January 1962). It is for all men over 14 which probably explains the discrepancy since our census figure leaves out men older than 44 who would, on average, have higher earnings.

¹⁵The individual census data is missing for this year.

¹⁶They impute wages for non-working wives using a sample of women who worked less than 20 weeks per year, controlling for age, education, race and region, and a metropolitan area indicator (page 42). They run a probit on work (positive hours) including log hourly wages (own and husband's), non-wage income,

probit specification, which does not control for education, as this way the elasticity measure obtained does not control for a measure of permanent income. This is preferable since we are more interested in an elasticity with respect to some measure lifetime earnings. The specification we chose also did not control for children which we consider an endogenous variable. Blau and Kahn estimate an own-wage elasticity of 0.30 and the cross-elasticity (husband’s wage) of -0.13 for our preferred specification in the year 2000.

4.2 Calibrating the Model Without Learning

We start out by calibrating the model without learning. In that model, only changes in earnings (male and female) can explain why labor supply could have changed over time. The unknown parameters are γ, β , and σ_l .

Note that the wage elasticity (own or cross) is given by:

$$\varepsilon_k = g(l^*) \frac{\partial l^*}{\partial w_k} \frac{w_k}{\omega} \quad (22)$$

$k = f, h$. Taking the ratio of the two elasticities and manipulating the expression, yields a closed-form expression for γ :

$$\gamma = \frac{\log\left(1 - \frac{w_f \varepsilon_h}{w_h \varepsilon_f}\right)}{\log\left(1 + \frac{w_f}{w_h}\right)} = 0.503 \quad (23)$$

Next we can use the elasticity expression and the requirement that $G(l^*; \sigma_l) = \omega$ in 2000 to solve for β and σ_l . Note that since G is a normal distribution, we can write:

$$l^* = \sigma_l \Phi^{-1}(\omega)$$

where Φ^{-1} is the inverse of a standard normal distribution $N(0, 1)$. After some manipulation of (22), we obtain:

$$\sigma_l = \frac{A}{\exp\left(\frac{\Phi^{-1}(\omega)^2}{2}\right)} = 2.29 \quad (24)$$

where $A = \frac{w_f(w_f + w_h)^{-\gamma}}{\sqrt{2\pi\varepsilon_f\omega}}$. We can then solve for β directly from the definition of l^* , yielding $\beta = 0.321$. To interpret this value, note that this is 9.44% of the utility from working in 1880 or 46.8% of the difference in utility between working and not working in that year.

As can be seen in Figure 6, the calibrated model does a terrible job of matching the female LFP data (the data is shown in small circles). It grossly overestimates the amount of female LFP that should exist in all decades other than 1990.

This basic inability of the earnings only model to match the historical data is robust to a wide range of values for the elasticities (we explored with values twice and half that of Blau and Kahn). It is also robust to alternative specifications of the shares of consumption that

along with the variables used to impute wages, both including and excluding education.

women obtain from their husband’s earnings. In particular, we can modify the model so that the wife obtains only a share $0 < \alpha \leq 1$ of her husband’s earnings as joint consumption. Figure 7 shows the results obtained from recalibrating the model to values of α that vary from 0.1 to 1. As is clear from this figure, this does little to remedy the basic problem. Furthermore, introducing any sensible time variation in this share would also not help matters as it would require women to have obtained a much larger share of husband’s earnings in the past than in the present in order to explain why they worked so much less then. Since women’s earnings relative to men’s are higher now than in the past, most reasonable bargaining models would predict the opposite, i.e., a greater ability to obtain a higher share of male earnings now than in the past.¹⁷

4.3 Estimating the Learning Model

Next we calibrate the learning model to the same year 2000 set of statistics as the earnings only model. Of course, this model has many more parameters with which to explain the same data, so that it is not possible for it to do a worse job than the previous model. How much better it should do, however, is not clear *ex ante*. As we will show below, it greatly improves the ability of the model to match the data.

After some algebra and noting that $\frac{\partial \bar{l}}{\partial w_k} = \frac{\partial l}{\partial w_k}$, one can show that the ratio of the elasticities in this model can be written as

$$\frac{\varepsilon_{w_f}}{\varepsilon_{w_h}} = \frac{\frac{\partial l}{\partial w_f} w_f}{\frac{\partial l}{\partial w_h} w_h} \quad (25)$$

Noting further that $\frac{\partial l}{\partial w_k} = \frac{\partial l^*}{\partial w_k}$, this implies that following the same manipulations as in the previous section, the value of γ in the model with learning must be equal to that in the earnings only model, i.e., $\gamma = 0.503$. We choose the values of the remaining parameters so that, in addition to matching the elasticities and female LFP in 2000, they also minimize the sum of the squared errors between the expected value and actual LFP across the decades. We assume throughout that the true state of nature is given by $\beta^* = \beta_L$.

An additional complication in estimating this model, that makes it different from an otherwise straightforward exercise, is the presence of an aggregate shock in each period (i.e., individuals observe a noisy *public* signal of aggregate female LFP). This implies that the path taken by the economy depends on the realization of this shock. Furthermore, each realization of the public shock generates a corresponding different public belief in the following period, and consequently a different proportion of women who choose to work after receiving their private signals.

We used the following procedure to deal with the complication introduced by the aggregate shock. For each period t , given the labor force participation in the previous period ω_{t-1} , we calculated the proportion of women who would work, ω_t , for each realization

¹⁷Note that, in any case, to obtain the very LFP numbers in 1880 would require women to obtain a share of 1 of husband’s earnings in that decade and a share of 0.0001 in the year 2000.

of the shock, η_{t-1} , i.e., for each induced belief $\lambda_t(\eta_{t-1})$. Integrating over the shocks, we found the expected value of LFP for that period, $E\omega_t(\lambda_t(\eta_{t-1}))$, and then backed out the public belief, λ_t^* , that would lead to exactly that same proportion of women working, i.e., we found, $\lambda_t^*(\eta_{t-1}^*)$ such that:¹⁸

$$E\omega_t(\lambda_t(\eta_{t-1})) = \omega_t(\lambda_t^*(\eta_{t-1}^*)) \quad (26)$$

Performing this exercise in each period determines the path of beliefs.¹⁹

The results are shown in Figure 8 and Table 1 reports the parameter values. The blue line in Figure 8 shows the evolution of the expected value of female LFP and the green line shows the evolution of beliefs, i.e., the belief, p_t , that the true state is β_L in period t .

Individuals start out in 1880 with pessimistic beliefs about how costly it is to work. They assign around 11% probability to the event $\beta^* = \beta_L$ and this belief evolves very slowly over the first sixty years or so (remaining below 20% for this period). Then, especially as of 1950, the change in beliefs accelerate, jumping from 40.4% in 1960 to 62.4% in 1970. By 2000, the public belief assigns a probability of 93.4% to $\beta^* = \beta_L$.

The main qualitative difference between the two models is with respect to the behavior of the expected value of β . In the earnings only model this is constant at 0.32 whereas in the learning model, in 1880 the expected value of β was about 6.66 which then evolved over time to 0.49 by 2000.

We can ask two different questions about the quantitative role of beliefs in this model. First, we can freeze beliefs at the 1880 level (i.e., a prior of approximately 11% that $\beta^* = \beta_L$) and ask how labor force participation would have evolved in the absence of any updating of beliefs using the public signal. Thus, in each period women receive a private signal and decide how much to work but there is no intergenerational evolution of beliefs. As show by the bottom line (with the caption "LFP if no public updating") in Figure 9, female LFP would barely exceeded 10% by the year 2000. Alternatively, one can ask what female LFP would have been if, throughout the entire time period, agents had known the true value of β , i.e., $\beta^* = \beta_L$. This scenario is shown for the parameters of the estimated model in the red line (with the caption "full information LFP"). It predicts a very different trajectory than the one we estimated, with LFP starting a bit over 60% in 1880 and slowly evolving to around 80% by 2000. Thus, as can be seen from contemplating either of the two extremes regarding constant beliefs, the actual dynamics of beliefs is central to producing the path of female LFP shown in the figure. The model with dynamics induced solely by changes in male and female earnings along with unchanged beliefs grossly under or over estimates female labor supply over the entire time period.

One can also evaluate the importance of the evolution of male and female earnings. In Figure 10 we show how female LFP would have evolved in the learning model if earnings had

¹⁸We take a large number of draws of entire histories for η (1000 histories). See the Appendix for details on the calibration and estimation methods.

¹⁹Note that, although the expected value of the shock is 0 in each period and λ_t is a linear function of η_{t-1} , the proportion of women who work ω_t is not a linear function of λ_t .

been kept constant. We examine two scenarios. In the first the earnings are kept constant at their 1880 levels and in the second they are frozen at their 2000 levels throughout.

As can be seen in the figure, earnings also matter to the trajectory of female LFP. In one case, LFP increases slowly over time and is below 50% in the year 2000 (the bottom blue line), whereas at the higher earnings levels, LFP takes off much faster than it does in the data reaching around 70% by 1950 (the top blue line). The corresponding paths of beliefs also look very different. The figure shows the paths of beliefs (the top red for the 2000 earnings and the bottom red for the 1880 earnings) for the two artificial earnings scenario as well as the one obtained under the historical earnings series (the latter is shown in black). At the higher earning levels, beliefs evolve very quickly since many more women are working in the early decades and thus (endogenously) improving the informativeness of the public signal since the s_j^* required becomes less extreme for all types. The opposite is true when earnings remain at their 1880 level. In that case, the path of beliefs looks like the one obtained under the historical earnings series until 1950. At that point there is an important increase in women's historical earnings—they increase 32% from 1940 to 1950 (and 32% and 16% for the following two decades) and women find it more attractive to work under the historical series, thus leading to an increase in the divergence of beliefs. Note that since learning slows down once there are many women working, the final beliefs under the three possible paths are converging over time and by 2000 they assign to $\beta = \beta^*$ a probability of 93% in actual solution vs 96% for the constant 2000 earnings vs 84% for the constant 1880 earnings.

It is instructive to note from the figure as well that even if earnings had remained unchanged from their 1880 levels, the evolution of beliefs over time implies that by 2000 female LFP would have been almost 50%, i.e., a large increase over its initial value of 2%. This, once again, points to the important quantitative role played by beliefs.

We can also evaluate the role that heterogeneity in l types plays in the model. If we were to eliminate heterogeneity in preferences (and keep the original solution parameters) then labor force participation would remain under 10% for the entire time period (see the bottom curve in Figure 11) and the elasticities in 2000 are completely off.²⁰ Reestimating the model without heterogeneity, on the other hand, yields the labor force participation path depicted on the top path of Figure 11. As can be seen, the model still does well except for 1980 and 1990. Interestingly, the model now requires a very large value of β_H to best fit the data (213,993 versus a β_L of 0.63) and the initial probability estimate for $\beta^* = \beta_L$ is 0.99992. The small amount of uncertainty and subsequent learning, however, is crucial.²¹ If the belief is kept at its initial level, female LFP in the year 2000 would be under 35%.

As a last exercise, we can use the model to generate a prediction for future female LFP. Using median earnings for men and women in 2005 as our guess for 2010 earnings (\$7518 and \$5959 in 1967 dollars, calculated as described earlier), our model predicts that 76% of

²⁰The own wage elasticity is 1.59 and the cross wage elasticity is -0.68.

²¹In the year 2000, the model without heterogeneity assigns 0.99999 to $\beta^* = \beta_L$.

women would work in 2010.

5 Conclusion

This paper argues that in some contexts it may be useful to think about cultural change as a process of beliefs updating that occurs as part of a rational intergenerational learning process. In particular, we model the aggregate changes in married women's labor supply as the outcome of a Bayesian learning process in which women learn about the long-term implications of working on children's and marital welfare. We show that a simple model with these features is capable of generating the aggregate time trend of female labor force participation over the last 120 years.

This model naturally generates the S-shaped curve of female LFP found in the data, shown in Figure 1. This shape results from the dynamics of learning. When very few women participate in the labor market (as a result of initial priors that are very negative about the payoff from working), learning is very slow since the noisiness of the signal swamps the information content given by differences in the proportion of women who would work in different states of the world. As the proportion of women who work increases and beliefs about work become more positive, the information in the signal improves. Once a large enough proportion of women work though, once again, the informational content in the public signal falls since the differences in the proportion of women who would work under different states of the world is swamped by the noise.

To quantitatively evaluate the potential ability of such a model to explain the evolution of female LFP, I first calibrate a version of the model without any evolution of beliefs to a few key female LFP statistics for the year 2000, namely LFP, and the own and cross-wage elasticities of LFP. I show that such a model performs very badly and that it grossly overestimates the proportion of women who would have worked for basically every time period. Introducing learning in this simple model greatly improves the capacity of the model to fit the data. Analysis of the model indicates that both the dynamic paths of beliefs and earnings played an important role in the transformation of women's work.

In addition to exploring the role of social networks as in Fernández and Potamites (2007), in future work it would be of interest to incorporate the contribution that social rewards and punishments may play in changing behavior over time. Munshi and Myaux (2006) follow this approach in the context of a model with multiple equilibria. In their model, the payoff to an individual using birth control depends on whether she interacts (at random) with another woman who is doing likewise. If society starts in an equilibrium with no modern contraceptive use, whether it can transit to another equilibrium with contraceptive use will depend upon the proportion of individuals who are reformers (and by assumption use contraception in the initial period). If there exists uncertainty about this proportion, interaction with others will yield information and individuals may learn over time and change the long-run equilibrium.

The path followed by Munshi and Myaux seems a very nice way to model why it may

take time for culture to change in a model with multiple equilibria. Our model differs from theirs in the critical feature that our long-run equilibrium is unique. It therefore would be of interest to examine whether a model without multiplicity may still yield an important role for social punishments that depend on beliefs.

Our paper has concentrated mainly on aggregate features of the data over a very long time horizon. It would also be of interest to examine sharper hypothesis about cultural change over a shorter time horizon that would allow a greater use of microdata.²²

²²Munshi and Myaux test their hypothesis, for example, using microdata from a 10 year interval in Bangladeshi villages.

6 Appendix

In order to estimate $\lambda_0, \sigma_\epsilon, \sigma_\eta$ along with $\beta_H, \beta_L, \sigma_l$ we minimized the sum of the squared errors between predicted and actual LFP (12 obs) in every year before 2000.

The simplex algorithm was used to search for an optimal set of parameters. Multiple starting values throughout the parameter space were tried (specifically over 4,000 different places with λ_0 ranging between $[-10, -.01]$, σ_ϵ in $[.1, 2]$, σ_η in $[.1, 2]$, σ_l between $[.5, 4]$, β_L in $[.01, 5]$, and β_H to be between $[1, 20]$ units greater than β_L).

A period is 10 years. 500 different public shocks were generated for each period (these draws were held constant throughout the minimization process). For each shock, there is a corresponding public belief that subjects begin the next period with. For each belief, a different percentage of women will choose to work after they receive their private signals.

200 discrete types were assumed between $\underline{l}(w_h, w_f)$ and $\bar{l}(w_h, w_f)$ in each year to approximate the integral in equation 17. Then we average over the η shocks to determine the expected number of women working. We then back out the belief that would lead to exactly that many women working. This determines the path of beliefs.

The elasticities were calculated computationally by assuming either a 1% increase in female wages or male wages in the year 2000 and calculating the corresponding changes in LFP predicted by the model in those histories in which the (original) predicted LFP was close to the true LFP value (specifically those histories in which the predicted LFP in the year 2000 was in the range $.734 \pm .05$). These elasticities were calculated individually for all histories meeting this criterion and were then averaged.

References

- [1] Albanesi, Stefania and Claudia Olivetti (2007), "Gender Roles and Technological Progress," Working Paper.
- [2] Bandiera, Oriana and Imran Rasul (2006), "Social Networks and Technology Adoption in Northern Mozambique," *Economic Journal*, Forthcoming
- [3] Belsky, Jay, Deborah Lowe Vandell, Margaret Burchinal, K. Alison Clarke-Stewart, Kathleen McCartney, and Margaret Tresch Owen (2007), "Are There Long-Term Effects of Early Child Care?" *Child Development* 78 (2), 681-701.
- [4] Bernal, Raquel and Michael Keane (2006), "Child Care Choices and Children's Cognitive Achievement: The Case of Single Mothers," Working Paper.
- [5] Bernal, Raquel (2007), "The Effect of Maternal Employment and Child Care on Children's Cognitive Development," Working Paper.
- [6] Blau, Francine D., and Lawrence M. Kahn (2006), "Changes in the Labor Supply Behavior of Married Women: 1980-2000," IZA Discussion Paper.
- [7] Chamley, Christopher (2003), *Rational Herds*, Cambridge University Press.
- [8] Conley, T. and Chris Udry (2003), "Learning About a New Technology: Pineapple in Ghana," mimeo Yale University.
- [9] Costa, Dora (2000), "From Mill Town to Board Room: The Rise of Women's Paid Labor," *The Journal of Economic Perspectives* 14(4), 101-122.
- [10] Fernández, Raquel (2007a), "Culture and Economics," in the *New Palgrave Dictionary of Economics*, 2nd edition, edited by Steven N. Durlauf and Lawrence E. Blume, Palgrave Macmillan (Basingstoke and New York), forthcoming.
- [11] Fernández, Raquel (2007b), "Women, Work, and Culture," *Journal of the European Economic Association*, forthcoming.
- [12] Fernández, Raquel and Alessandra Fogli (2005), "Culture: An Empirical Investigation of Beliefs, Work, and Fertility," NBER Working Paper No. 11268.
- [13] Fernández, Raquel, Alessandra Fogli, and Claudia Olivetti (2004), "Mothers and Sons: Preference Formation and Female Labor Force Dynamics," *Quarterly Journal of Economics* 119(4), 1249-1299.
- [14] Fernández, Raquel and Elizabeth Potamites (2007) "Learning, Social Networks, and Changes in Female Labor Force Participation," mimeo, NYU.
- [15] Fortin, N. (2005), "Gender Role Attitudes and the Labour Market Outcomes of Women Across OECD Countries," *Oxford Review of Economic Policy*, 21, 416-438.

- [16] Foster, Andrew and Mark Rosenzweig (1995), "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture," *Journal of Political Economy* 103, 1176-1209.
- [17] Gallup, George H. (1972), *The Gallup Poll; public opinion, 1935-1971* New York: Random House.
- [18] Galor, Oded and David Weil (1996), "The Gender Gap, Fertility, and Growth," *American Economic Review* 86(3), 374-87.
- [19] Goldin, Claudia (1990), *Understanding the Gender Gap: An Economic History of American Women*, New York: Oxford University Press.
- [20] Goldin, Claudia and Lawrence Katz (2002), "The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions," *Journal of Political Economy* 110(4), 730-770.
- [21] Goldin, Claudia and Robert A. Margo (1992). "The Great Compression: The Wage Structure in the United States at Mid-century" *The Quarterly Journal of Economics*, 107(1), 1-34.
- [22] Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu (2005), "Engines of Liberation," *Review of Economic Studies* 72(1) 109-133.
- [23] Griliches, Zvi (1957), "Hybrid Corn: An Exploration in the Economics of Technological Change", *Econometrica* 25, 501-22.
- [24] Guiso, Luigi, Paola Sapienza and Luigi Zingales (2006), "Does Culture Affect Economic Outcomes," *Journal of Economic Perspectives* 20(2), 23-48.
- [25] Guner, Nezih and Jeremy Greenwood (2006), "Social Change," working paper.
- [26] Hill, Jennifer L., Jane Waldfogel, Jeanne Brooks-Gunn, and Wen-Jui Han (2005), "Maternal Employment and Child Development: A Fresh Look Using Newer Methods," *Developmental Psychology* 41(6), 833-850.
- [27] Jones, Larry, Rodolfo Manuelli, and Ellen McGrattan (2003), "Why are Married Women Working so much?" Working Paper.
- [28] Katz, Lawrence and David Autor (1999), "Changes in the wage structure and earnings inequality," in O. Ashenfelter D. Card (ed.), *Handbook of Labor Economics*, edition 1, volume 3, chapter 26, pages 1463-1555, Elsevier.
- [29] Levine, David (1993), "The Effect of Non-Traditional Attitudes on Married Women's Labor Supply," *Journal of Economic Psychology* 14, 665-679.
- [30] Munshi, Kaivan (2004), "Social Learning in a Heterogeneous Population: Social Learning in the Indian Green Revolution," *Journal of Development Economics*, 73, 185-213

- [31] Munshi, Kaivan (2006), “Social norms and the fertility transition,” *Journal of Development Economics*, 80, 1– 38
- [32] Ruhm, Christopher J. (2004), “Parental Employment and Child Cognitive Development,” *Journal of Human Resources* 39(1), 155-192.
- [33] Vella, Francis (1994), “Gender Roles and Human Capital Investment: The Relationship between Traditional Attitudes and Female Labour Force Performance,” *Economica* 61, 191-211.
- [34] Vives, Xavier (1993), “How Fast do Rational Agents Learn,” *The Review of Economic Studies*, 60 (2), 329-347.

Table 1: Calibration and Estimation Results

Parameter	Earnings	Belief
	Only Model	Learning Model
γ	0.503	0.503
σ_L	2.293	2.067
β	0.321	
β_H		7.481
β_L		.0004
$P_0(\beta = \beta_L)$		0.110
σ_ϵ		5.408
σ_η		0.157

Married Female Labor Force Participation in the U.S.




Figure 1: Source: U.S. Census data 1880-2000. White, married (spouse present) women born in the U.S. 25-44 years old who report being in the labor force.

Approve of Wife working if Husband can Support




Figure 2: Source: 1936-1938 and 1969 numbers are from the Gallup Poll (1972), 1945 is from Benjamin I. Page and Robert Y. Shapiro, *The Rational Public*, University of Chicago Press, 1992; pp. 101,403-404. 1972 onwards are from the General Social Survey.

Figure 3: Time Line of Learning Model

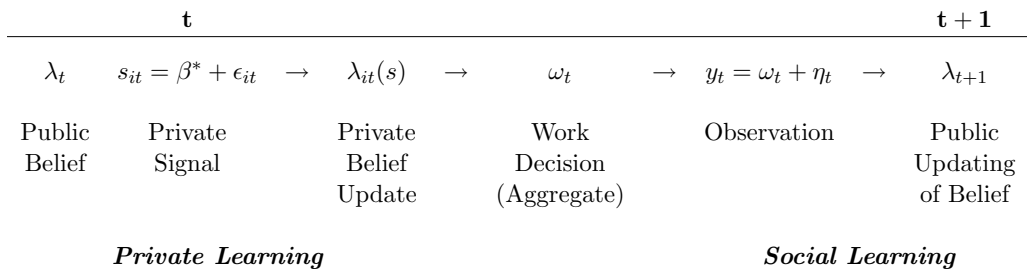



Figure 4:



Female and Male Earnings




Figure 5: Crosses (blue) represent the yearly earnings data (in 1967 \$) from Goldin table 5.1. Dots represent our calculation using U.S. census data (red). They are the median earnings of white men and women between the ages of 25-44, in non-farm occupations and not living in group quarters. See text for more detail.

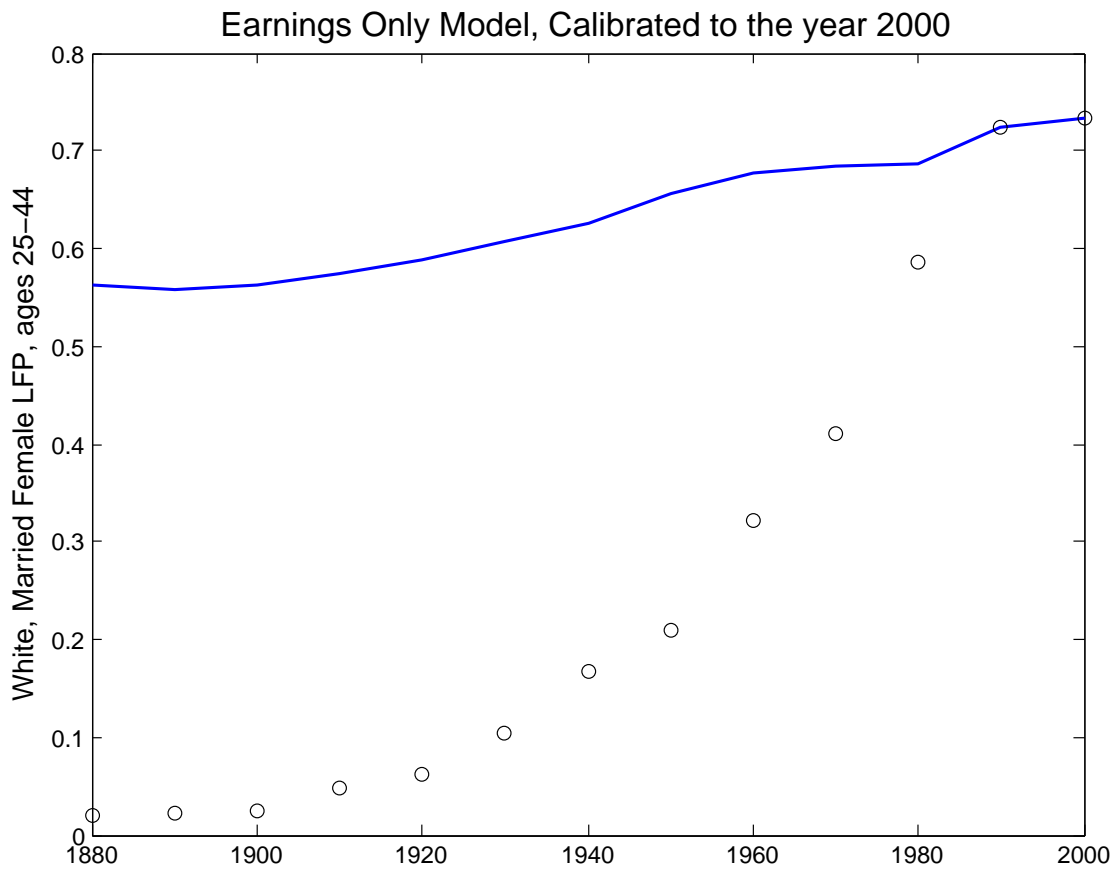


Figure 6: Parameters: $\gamma = 0.503$, $\beta = 0.321$, and $\sigma_L = 2.293$




Figure 7: α is the fraction of husband's earnings that enter a wife's utility via consumption.

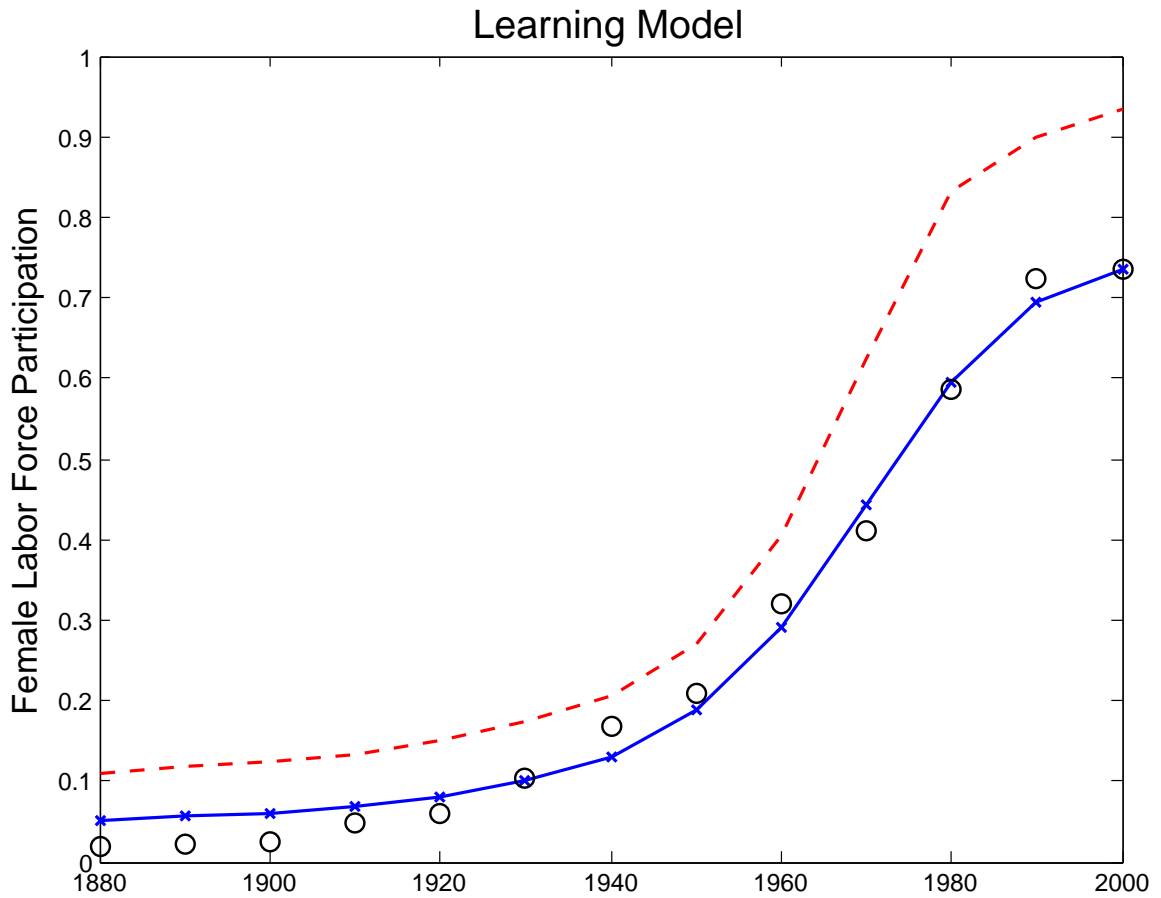


Figure 8: Dashed red line is belief path. Parameters: $\gamma = 0.503$, $\beta_H = 7.481$, $\beta_L = 0.0004$, $\sigma_L = 2.067$, $\sigma_\epsilon = 5.408$, $\sigma_\eta = 0.157$, and initial $Pr(\beta = \beta_L) = 0.110$. Sum of squared errors (distance of predicted LFP from actual LFP) is 0.009. Estimated elasticities (calculated computationally as described in text) are -0.13 with respect to husband's earnings and 0.30 with respect to own earnings.

The Role of Beliefs

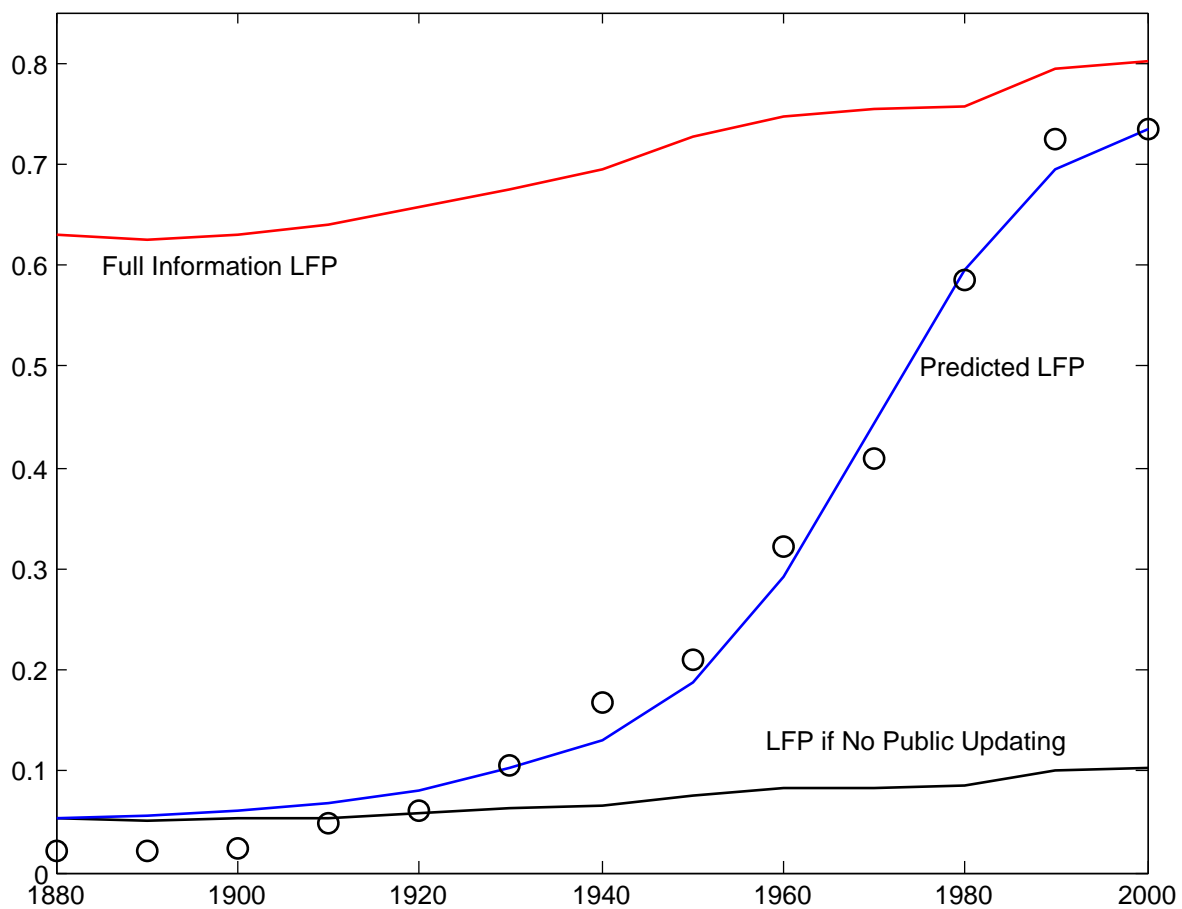


Figure 9: Uses same solution parameters as in figure 8.

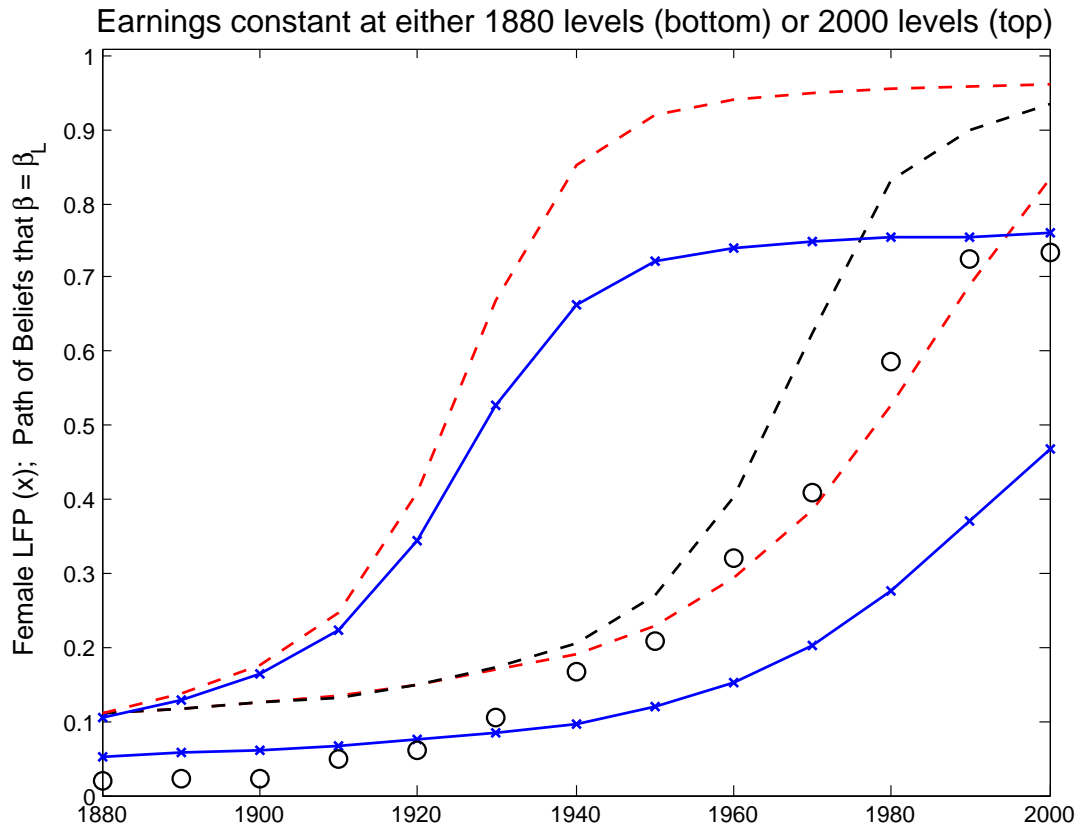


Figure 10: Uses same solution parameters as in figure 8. Dashed red lines are the belief paths if earnings were constant at 1880 levels (bottom) or at 2000 levels (top). The dashed black in middle is the trajectory of beliefs given actual earnings path.

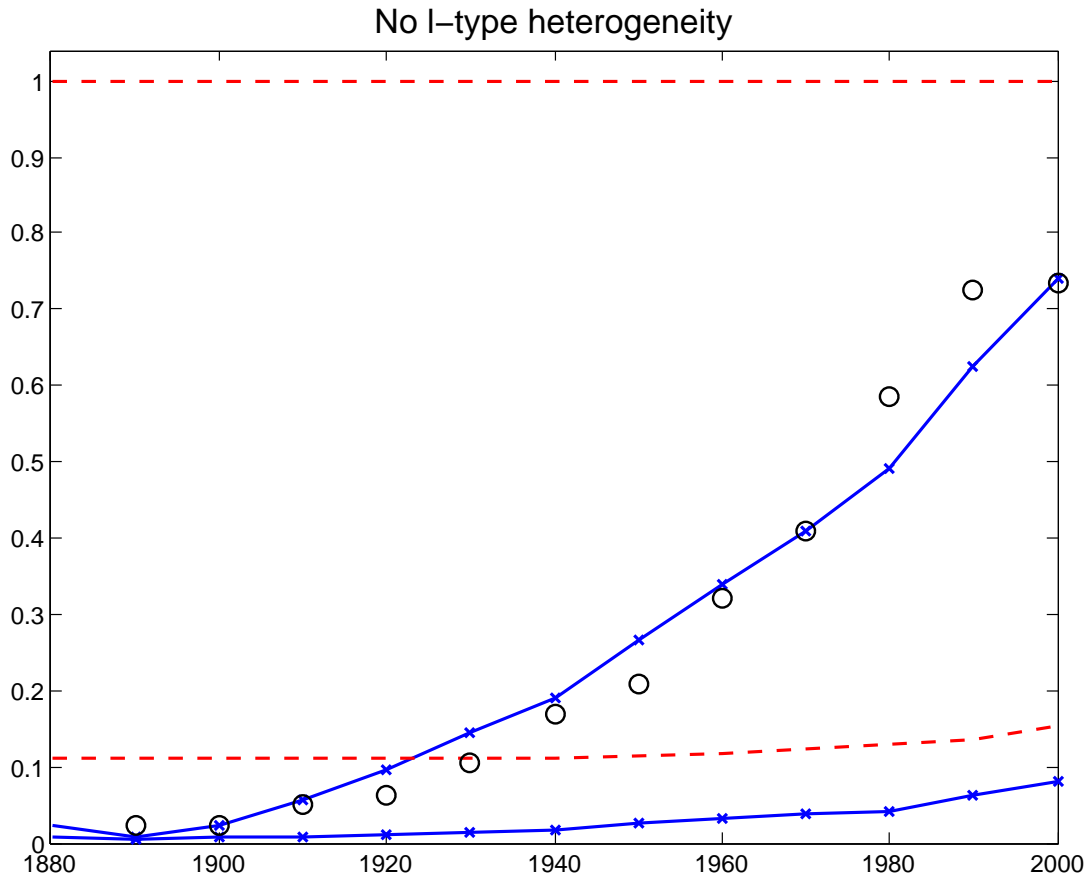


Figure 11: The two bottom lines are female LFP (x's) and beliefs using the same solution parameters as in figure 8 but with $\sigma_L = 0$. The two top lines are from the re-calibrated model with no l-type heterogeneity. Dashed red lines are beliefs in both cases. Parameters: $\gamma = 0.503$, $\beta_H = 213,993$, $\beta_L = 0.633$, $\sigma_L = 0$, $\sigma_\epsilon = 109,411$, $\sigma_\eta = 0.455$, and initial $Pr(\beta = \beta_L) = 0.99992$. Sum of squared errors (distance of predicted LFP from actual LFP) is 0.026.