

Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century*

Raquel Fernández
New York University

Revised July 2011[†]

Abstract

This paper investigates the role of changes in culture in generating the dramatic increase in married women's labor force participation over the last century. 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 endogenously via an intergenerational learning process. Women are assumed to learn about the long-term payoffs of working by observing (noisy) private and public signals. This process generically generates the S-shaped figure for female labor force participation found in the data. I calibrate the model to several key statistics and show that it does a good job in replicating the quantitative evolution of female LFP in the US over the last 120 years. I also examine the model's cross-sectional and intergenerational implications. The model highlights a new dynamic role for changes in wages via their effect on intergenerational learning. The calibration shows that this role was quantitatively important in several decades.

JEL Nos.: J16, J21, Z1, D19. Keywords: female labor force participation; cultural transmission; preference formation; learning; S shape; social norms; beliefs.

*Other affiliations: CEPR, NBER, IZA, and ESOP. I thank the anonymous referees for their many helpful remarks and Liz Potamites and Joyce Cheng Wong for excellent research assistance. Many people made helpful comments: I thank Christophe Chamley, John Knowles, Oriana Bandiera, Gianluca Violante, Elisabeth Schulte, and George Tridimas among others. I also wish to thank seminar audiences at the LAEF "Households, Gender and Fertility" conference, the NY/Philadelphia feds. "Quantitative Macroeconomics" Workshop, the NBER Summer Institute, the Silvaplan Political Economy workshop, the "Family Behavior and the Aggregate Economy", SITE workshop, LACEA, and Columbia U. for many helpful remarks. Lastly, I thank the NSF for financial support and the Russell Sage Foundation for its hospitality.

[†]An earlier version of the model and simulation in this paper were presented in my Marshall Lecture at the EEA, Vienna, August 2006. The slides for this presentation are available at <http://homepages.nyu.edu/~rf2/Research/EEAslidesFinal.pdf> (pp 48-52).

1 Introduction

A fundamental change over the last century has been the vast increase in female labor force participation, particularly for married women. The pace of change, however, has been markedly uneven. As shown in figure 1, white married women’s labor force participation was at around 2% in 1880 and increased very slowly to 1920, averaging 1 percentage point per decade. It grew somewhat more rapidly between 1920 and 1950 (on average 4.9 percentage points per decade), and then took off between 1950 and 1990, increasing on average 12.9 percentage points per decade. Since then, it has stayed relatively constant.¹

Many explanations have been proposed for this transformation. Depending on the particular time period under consideration, potential causal factors have included structural change in the economy, technological change in the workplace and in the household, medical advances, decreases in discrimination, institutional changes in divorce law, and the greater availability of childcare.² A striking fact that none of these theories have addressed is the accompanying revolution in social attitudes towards married women working. This social transformation can be seen everywhere, from changes in laws governing women’s work (e.g. the ”marriage bar” which made it difficult for schools to hire married women or to keep women as teachers once they married) to the depiction of married women in literature and the popular press.³

Quantitative evidence for the dramatic changes in social attitudes is provided by polls. Figure 1 also plots the evolution over time of the percentage of the sample that answered affirmatively to the question “Do you approve of a married woman earning money in business or industry if she has a husband capable of supporting her?”⁴ As shown by the stars (red), in 1945 fewer than 20% of individuals sampled agreed with the statement; in 1998 fewer than 20% of individuals disagreed with it.

To understand why social attitudes (i.e., culture) and female labor force participation moved in tandem requires a framework that is able to address both phenomena. The objective of this paper is to provide such a framework and to examine whether the suggested mechanism may be quantitatively significant. Taking inspiration from the fact that the path of female labor force participation follows an ”S-shape” over time and thus suggesting a

¹These LFP numbers were calculated by the author from the US Census for white, married women between the ages of 25-44, born in the US, in non-agricultural occupations and living in non-farm, non-institutional quarters.

²The 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), Greenwood, Seshadri, and Yorukoglu (2005), Gayle and Golan (2010), Albanesi and Olivetti (2009a, 2009b), Attanasio, Low, and Sanchez-Marcos (2008), and Knowles (2007).

³For example, Sarah Hale, editor of Godey’s from 1836-1877, lectured her readers about “the importance of sticking to home and hearth” so that they would be able to “influence, and ennoble, the entire world” and the popular novelist Grace Greenwood wrote that “true feminine genius is ever timid, doubtful and clingingly dependent, a perpetual childhood”). One way to follow this transformation is by examining the evolution of perceptions of women’s role by the influential Godey’s Lady’s Book, a popular periodical. The quotes are taken from Collins (2003), p. 86-87.

⁴The exact wording of this question varied a bit over time. See The Gallup Poll; public opinion, 1935-1971.

process of information diffusion, this paper develops a model in which cultural change is the result of a rational, intergenerational *learning* process in which individuals are endogenously learning about the married women’s long-run payoff from working.⁵ In this process, female labor force participation (LFP) and culture are co-determined, giving rise to an S-shaped process of aggregate labor supply and social attitudes.

In this paper I develop a simple model of a married woman’s work decision. Using a framework broadly similar to Vives (1993) and Chamley (1999), I assume that women possess private information about the long-run costs of working (e.g., about the severity of the consequences of working for a woman’s marriage, her children, etc.) and that they also observe a noisy public signal indicative of 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 following period, the next generation once again observes a noisy public signal generated by women’s decisions in the preceding generation and they in turn make their own work decision. Thus, beliefs evolve endogenously via a process of intergenerational learning.

The model has several attractive features. First, the model *generically* generates an S-shaped figure for female labor force participation and for social beliefs. Second, the model introduces a new role for changes in wages or technological change. Unlike traditional models in which changes in women’s wages affect female LFP solely by changing the payoff from working, in this model they also affect the *informativeness* of the public signal and hence the degree of intergenerational updating of beliefs. Thus, wages affect the pace of learning and consequently have dynamic effects on female LFP. Third, the model generates a path for culture, i.e., for the evolution of social beliefs.

To evaluate whether the proposed learning mechanism has the potential to be quantitatively significant, I calibrate the model to a few key statistics for the last decades of the sample (1980-2000). I find that the calibrated model does a fairly good job of replicating the dynamic path of married women’s LFP from 1880 to 2000. The calibrated model indicates that the paths of both beliefs and earnings played important roles in the transformation of women’s work. In the decades between 1880-1950 the growth in female LFP was small, and most of the change in LFP was the result of changes in wages. From 1950 -1970, both the dynamic and static effects of wage changes played a role in increasing female LFP, and from 1970 -1990 in what Goldin (2006) has called the “revolutionary” phase, the dynamic effect on beliefs of changes in earnings is critical in accounting for the large increase in the proportion of working married women. Thus, the model resurrects the importance of wage changes in explaining the dynamic path of female LFP with a novel mechanism.⁶ A welfare

⁵A curve that rises slowly at first, then rapidly, and then flattens is called an “s-shaped” diffusion shape even if it doesn’t, strictly speaking, look like an S. These curves are common in the technology and epidemiology literature. See Geroski (2000) for a review of this literature.

⁶It should be noted that Jones, Manuelli, and McGrattan (2003) show that changes in the gender gap can account for much of the changes in married women’s labor supply from 1950 to 1990 in a model with home production and leisure. The new channel for wages that I explore here does not require a home-production sector and in this way the papers can be seen as complementary.

analysis of the cost of imperfect information included in the online Appendix indicates, not surprisingly, that the latter was high.

A unique feature of the model is that it endogenously generates a path for culture or beliefs. A comparison of the predictions of the calibrated model with the poll data shows that the model does a good job in generating the basic features of the data. I also derive other testable implications of the model regarding elasticities, cross-sectional predictions across women of different education levels, and the dynamics of LFP and beliefs. In particular, cross-state variation in the importance of World War II as a shock that generated exogenous variation in female labor supply can be used to study whether the intergenerational implications of the model are consistent with the data.

The paper is organized as follows. Section 2 reviews recent literature on the role of culture and the evolution of married women’s LFP. Section 3 presents the learning model and Section 4 derives the main results. Section 5 presents the calibration of the model and decomposes the changes in LFP into a beliefs component, a static wage component, and a dynamic wage-belief component to assess their respective contributions to the evolution of LFP. Section 6 examines other testable implications of the model. Section 7 discusses the roles of various assumptions and concludes. An online Appendix presents details of the calibration procedure and data construction and conducts an analysis of the welfare cost of imperfect information.

2 The Literature on Married Women’s LFP, Culture, and Learning

There is a by now a large literature on married women’s labor supply.⁷ This literature has attempted to explain the change in married women’s labor supply primarily with changes in standard economic variables such as women’s education, earnings and the gender gap, fertility, and marriage/divorce prospects. It has sometimes included factors that are more difficult to measure such as changes in the cost of childbearing, culture or social norms, or technological change in household production. Overall, this literature has concluded, often by default, that this last category of difficult-to-measure variables plays an important quantitative role in accounting for the large changes seen in married women’s LFP (see, e.g., Eckstein and Lifshitz (2009), Lee and Wolpin (2010), and Heathcote, Storesletten, and Violante (2010)).

There is another strand of literature that provides micro-based evidence on the role of culture and married women’s labor supply using the “epidemiological” approach to studying culture.⁸ This approach attempts to separate the influences of culture from those of institutions and traditional economic variables by studying the descendants of immigrants in a given country. The basic argument is that whereas immigrants tend to transmit the

⁷This literature is too large to be reviewed here. See, e.g., the survey by Blundell and MaCurdy (2000).

⁸See the chapter in the Handbook of Social Economics (Fernández (2010a)) for a review of the literature using this approach.

preferences and beliefs of their country of origin to their offspring, the new economic and institutional environment in which these descendants make choices will be the same across the different immigrant groups (controlling for geographic variation within the country). One can then examine whether variables that reflect the beliefs in the country-of-ancestry are able to account for a significant portion of the variance in economic outcomes across the second-generation groups (see, e.g., Antecol (2000), Burda, Hamermesh, and Weil (2007), Fernández (2007), and Fernández and Fogli (2009)).

Neither the epidemiological approach to culture nor the dynamic quantitative models discussed above provide explicit mechanisms that explain why culture might vary, whether across countries or over time. At least in part this absence has been due to the paucity of models of cultural *change*. The few papers that model the dynamics of culture do so either in an evolutionary theory framework (see, e.g. Bowles (1998)) or in utility-maximizing models in which the proportions of people with different *fixed* beliefs (e.g., the proportion that is Catholic versus Jewish) evolve in the population over time (e.g., Bisin and Verdier (2000)).

An exception to this is Hazan and Maoz (2002) and Fernández, Fogli, and Olivetti (2004). Hazan and Maoz’s model is also inspired by the S-shape in the path of women’s LFP. In their model, working incurs the psychic cost of violating the social norm, which is assumed to decrease with the proportion of women who worked in the previous generation. For some parameters, the model can exhibit multiple locally-stable steady states – one with low participation and one with high participation. Which steady state is chosen depends on initial conditions, i.e. on the proportion of women working initially. The S-shape is obtained by assuming that there is a normal distribution in the disutility of work which implies that if few women work early on, LFP will evolve gradually giving rise to the shape of LFP found in the data.

In the model of Fernández et al (2004), working mothers are assumed to pass on to their male children a more positive view of working women, making these boys more amenable later on to having a working wife. This in turn renders it more attractive for girls invest in market-work human capital (as opposed to house-work human capital), as they know that men are more likely to be supportive of a working wife. Their model does not necessarily imply an S shape for married women’s LFP, however.

In contrast with the literature above, the present paper explicitly models the changes in culture as arising from a process of learning.⁹ In this process, the probability individuals assign to different views of the long-term consequences of married women working is updated in a Bayesian fashion as new information endogenously becomes available.¹⁰ Although there

⁹The idea that cultural change may be modelled as a learning process is already present in the seminal paper of Bikhchandani, Hirshleifer, and Welch (1992), though the focus is very different since they are interested in information cascades in which individuals stop learning.

¹⁰A recent paper by Fogli and Veldkamp (2007) independently (but subsequent to my presenting the learning model developed in this paper in my Marshall lecture at the EEA), develops a related idea. They study the labor force participation of women with children for a substantially shorter period (1940-2000) and assume that women learn about the ability cost to a child from having a working mother. Learning occurs through sampling the ability outcomes of children of a small number of other women. Whereas in my

is no direct evidence showing that learning is responsible for the change in social attitudes, as noted in the introduction the path followed by married women's LFP is the commonly denoted S-shape that arises in processes of technological diffusion.¹¹ Thus, the very shape of the LFP path may itself constitute a clue that a similar mechanism of information diffusion is also at play in this context, though on a very different time scale. Furthermore, using a learning mechanism to think about cultural change has the advantage, from an economist's perspective, that it lends itself to standard welfare analysis unlike mechanisms in which preferences themselves change.

For an approach that emphasizes learning to make sense, women need to have been uncertain about the consequences of working. Thus, one must come to grips with the question of what is it that women could have been learning about over the last 120 years. It is not an exaggeration to state that, throughout the last century, the consequences of women's (market) work have been a subject of great contention and uncertainty. In the US, as in other countries, a process of specialization accompanied industrialization and urbanization. Younger men and (unmarried) women were drawn into the paid workplace and away from sharing household chores, and the spheres of work and home became increasingly separate. This process left the wife in charge of the domestic realm and her husband responsible for supporting the family, and spawned a debate on the consequences of a working wife on her family and marriage as well as on her psyche and image (and on those of her husband's) that continues, in different guises, to this day.¹²

For example, as noted by Goldin (1990), at the turn of the 20th century most working women were employed as domestic servants or in manufacturing. In this environment, a married woman's employment signalled that her husband was unable to provide adequately for his family and, consequently, most women (around 80% prior to 1940) exited the workplace upon marriage.¹³ Over time, the debate shifted to the effect of a married woman working on family stability and to the general suitability of women for various types of work and careers.¹⁴ More recently, public anxiety regarding working women centers around the effect of a working mother on a child's intellectual achievements and emotional health. 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. Thus, throughout the last century the expected payoff to a married woman working has been the subject of an evolving process of discovery and debate.

model actions change because people modify their beliefs about the cost of working, in their model actions change only because of a reduction of uncertainty about the cost.

¹¹See, e.g., Griliches (1957). See also Foster and Rosenzweig (1995), Conley and Udry (2003), Munshi (2004), Munshi and Myaux (2006), and Bandiera and Rasul (2006) for models and evidence of learning about new technologies. See Chamley (2004) for a review of this literature.

¹²See Goldin (1990) for a very interesting account of this process of separation and specialization.

¹³Over 80% of married women, not employed in 1939 but who had worked at some point prior to marriage, exited the workplace at the precise time of marriage. These numbers are cited in Goldin (1990, p. 34) from the 1939 Retrospective Survey.

¹⁴For example, at the turn of the 20th century, work and marriage were not seen as very compatible for educated middle-class women. Jane Addams noted that men "did not want to marry women of the new type, and women could not fulfill the two functions of profession and home making until...public opinion tolerated the double role." Cited in Collins (2003), p.294.

3 A Simple Learning Model of Work and Culture

This section develops a simple model of a married woman's work decision in which the disutility from working includes an unknown *long-run* welfare consequence (i.e., the long-run consequences of working on a woman's identity, her marriage, or her children, as previously discussed). As these payoffs are revealed gradually over a long period of time, this uncertainty cannot be resolved by short-run experimentation. Thus, whether to work or not is modeled as a one-time decision.¹⁵ The first subsection describes the maximization problem faced by a woman in time t , given her prior beliefs and her private information; the second subsection derives how beliefs evolve across generations over time.

3.1 The Work Decision

A woman makes her work decision to maximize:

$$U(w_f, w_h, v_i) = \frac{c^{1-\gamma}}{1-\gamma} - \mathbf{1}(E_{it}v_i), \quad \gamma \geq 0 \quad (1)$$

where $\mathbf{1}$ is an indicator function that takes the value one if she works and zero otherwise. A woman's disutility of working is the sum of a known idiosyncratic component l_i (whose distribution in the population is $G(l)$ which we take to be $N(0, \sigma_l^2)$) and an unknown component B_i which is the individual-level realization of a random variable that is iid across women. The realization of this variable is revealed only if the woman works – it is the *long-run* disutility from working.¹⁶ Consumption is a household public good that is given by the sum of a woman's earnings, w_f , (which are positive only if she works) and her husband's earnings, w_h (husbands are assumed to always work). Thus,

$$c = w_h + \mathbf{1}w_f, \quad v_i = l_i + B_i, \quad B_i = \beta + u_i \quad \text{with} \quad E(u_i) = 0 \quad (2)$$

Note that since v_i enters linearly in utility, only the expected value of B_i , β , enters women's work decisions.¹⁷ For simplicity, I assume that β can take on only two values, high (H) and low (L), i.e., $\beta \in \{\beta_H, \beta_L\}$ where $\beta_H > \beta_L \geq 0$.

Prior to making her work decision, a woman inherits her mother's private signal s_i that yields (noisy) information about the true value of β .^{18,19} In particular, the value of the

¹⁵For simplicity I only consider the extensive margin as this is the one that has seen the largest changes over time.

¹⁶Note that the realization of B happens at the end of a woman's work period. Hence she is not able to exit the labor market if it turns out to be high. A multi-period model would allow B_i to be learned slowly over time. In that case, women would also exit the labor market if their expectations become more pessimistic. A multiperiod extension of the model would complicate the algebra considerably but not otherwise change the main conclusions of the model.

¹⁷A simple linear specification allows us to abstract from risk aversion (the only force at work in Fogli and Veldkamp (2007)) and focus purely on learning.

¹⁸It is equivalent from a modelling perspective to have the signal inherited from one's mother or drawn afresh (e.g., a draw from the scientific literature that exists at that moment regarding the welfare consequences of a woman working).

¹⁹This assumption guarantees that a mother's and daughter's attitudes towards women working will be positively correlated in accordance with the evidence of Farre and Vella (2009).

signal is given by:

$$s_i = \beta + \varepsilon_i \quad (3)$$

where $\varepsilon \sim N(0, \sigma_\varepsilon^2)$ with a cdf $F(\cdot; \sigma_\varepsilon)$ and pdf $f(\cdot; \sigma_\varepsilon)$.²⁰ The private signals are assumed to be i.i.d across women in any given generation.

After inheriting her private signal s_i , each woman updates her prior belief using Bayes' rule. Consider a woman i in period t who has a prior belief about β as summarized in the log likelihood ratio (LLR) $\lambda_t = \ln \frac{Pr(\beta=\beta_L)}{Pr(\beta=\beta_H)}$. By Bayes rule, her beliefs given her private signal s can be summarized in a new LLR, $\lambda_{it}(s)$, given by

$$\lambda_{it}(s) = \lambda_t + \ln \left(\frac{Pr(s|\beta = \beta_L)}{Pr(s|\beta = \beta_H)} \right) = \lambda_t - \left(\frac{\beta_H - \beta_L}{\sigma_\varepsilon^2} \right) (s - \bar{\beta}) \quad (4)$$

where $\bar{\beta} = (\beta_L + \beta_H)/2$.²¹ A few properties of $\lambda_{it}(s)$ are worth noting:

$$(i) \quad \frac{\partial \lambda_{it}(s)}{\partial s} < 0; \quad (ii) \quad \frac{\partial^2 \lambda_{it}(s)}{\partial s \partial \sigma_\varepsilon^2} > 0.$$

The first property follows from the fact that higher realizations of s increase the likelihood that $\beta = \beta_H$. The second property implies that the updating of λ is decreasing with the variance of the noise term, σ_ε^2 , since a greater variance lowers the informativeness of the signal. Note that since $s \sim N(\beta, \sigma_\varepsilon^2)$, it follows that $\lambda_{it}(\beta) \sim N\left(\lambda_t - \left(\frac{\beta_H - \beta_L}{\sigma_\varepsilon^2}\right)(\beta - \bar{\beta}), \frac{(\beta_H - \beta_L)^2}{\sigma_\varepsilon^2}\right)$.

Assume that women share a common prior in period t , λ_t .²² What proportion of women will choose to work that period? From (1) it follows that woman i will work iff $W(w_{ht}, w_{ft}) \equiv \frac{1}{1-\gamma} [(w_{ht} + w_{ft})^{1-\gamma} - w_{ht}^{1-\gamma}] - E_{it}(\beta) \geq l_i$. Note first that given $\{\beta_H, \beta_L\}$ and earnings (w_{ht}, w_{ft}) , irrespective of their beliefs and thus of the signal they receive, women with very low l ($l \leq \underline{l}(w_{ht}, w_{ft})$) will always choose to work whereas women with very high l ($l \geq \bar{l}(w_{ht}, w_{ft})$) will never choose to work, where

$$\underline{l}(w_{ht}, w_{ft}) \equiv W(w_{ht}, w_{ft}) - \beta_H \quad (5)$$

$$\bar{l}(w_{ht}, w_{ft}) \equiv W(w_{ht}, w_{ft}) - \beta_L \quad (6)$$

Next, for each women of type l , $\underline{l} < l < \bar{l}$, one can solve for the critical value of the private signal $s_i^*(\lambda)$ that she would need to inherit from her mother in order to be indifferent between working and not working. Let $p \equiv Pr(\beta = \beta_L)$ and let p_l^* be the critical probability such that a woman of type l is indifferent between working and not, i.e., $p_l^* \beta_L + (1 - p_l^*) \beta_H = W(w_{ht}, w_{ft}) - l$ and thus $p_l^*(w_{ht}, w_{ft}) = \frac{\beta_H + l - W}{\beta_H - \beta_L}$ or, using (5) and (6), $\ln \frac{p_l^*}{1 - p_l^*} = \ln \frac{l - \underline{l}}{\bar{l} - l}$. Hence the critical value $s_i^*(\lambda)$ is given by:

$$s_i^*(\lambda_t; w_{ht}, w_{ft}) = \bar{\beta} + \left(\frac{\sigma_\varepsilon^2}{\beta_H - \beta_L} \right) \left(\lambda_t + \ln \left(\frac{\bar{l}(w_{ht}, w_{ft}) - l}{l - \underline{l}(w_{ht}, w_{ft})} \right) \right) \equiv s_i^*(\lambda_t) \quad (7)$$

²⁰The results do not depend on ε being normally distributed. Rather, as will be made clear further on, they require a cdf that changes at first slowly and then more rapidly.

²¹To obtain (4) one uses the fact that $Pr(s|\beta)$ is equal to the probability of observing a signal s generated by a normal distribution $N(\beta, \sigma_\varepsilon^2)$.

²²The structure of the model will ensure that this is the case.

Invoking the law of large numbers, the proportion of women of type l , $\underline{l} < l < \bar{l}$, that will choose to work given the true value of β and a prior of λ_t , $L_{lt}(\beta; \lambda_t)$, is the proportion of them that receive a signal lower than $s_l^*(\lambda_t)$, i.e., $L_{lt}(\beta; \lambda_t) = F(s_l^*(\lambda_t) - \beta; \sigma_\varepsilon)$. Integrating over all l types, the overall proportion of women who will work in period t is given by:

$$L_t(\beta; \lambda_t) = G(\underline{l}) + \int_{\underline{l}}^{\bar{l}} F(s_l^*(\lambda_t) - \beta; \sigma_\varepsilon)g(l)dl \quad (8)$$

where $g(\cdot)$ is the pdf of $G(l)$. Note that L_t can in equilibrium only take on only two values – one for each possible value of β . How $L_t(\beta_L; \lambda_t) - L_t(\beta_H; \lambda_t)$ evolves as λ changes will play an important role in the results that follow.

Before specifying how beliefs are transmitted across generations, it is worth noting a few properties of $s_l^*(\lambda)$.²³

$$(i) \quad \frac{\partial s_l^*}{\partial \lambda} > 0; \quad (ii) \quad \frac{\partial s_l^*}{\partial w_f} > 0; \quad \frac{\partial s_l^*}{\partial w_h} < 0.$$

That is, *ceteris paribus*, a more optimistic prior implies that women are willing to work for higher values of s . Second, increases in own earnings make women more willing to work (and thus willing to do so at a higher signal) whereas the opposite holds for increases in husband's earnings. It follows directly that: $\frac{\partial L_t}{\partial w_f} > 0$ and $\frac{\partial L_t}{\partial w_h} < 0$. Thus the model satisfies the traditional (empirically desirable) comparative statics results with respect to wages, which is not surprising as these properties follow entirely from the specified preferences.

3.2 Intergenerational Transmission of Beliefs

The model thus far is purely static. To incorporate dynamics we need to specify how the state variable (beliefs) changes over time by being explicit as to the information passed on from generation t to generation $t + 1$.

I assume that generation $t + 1$ inherits the prior of generation t (its “culture”), λ_t , which each individual then updates with her own idiosyncratic private signal, s_i , inherited from her mother.²⁴ An equivalent assumption is that each woman transmits her belief to her daughter, i.e., a child inherits $\lambda_{it}(s)$ from her mother.²⁵ In either case, a woman in $t + 1$ ends up with the belief $\lambda_{it}(s)$. If solely this information were transmitted intergenerationally, then $\lambda_{it}(s) = \lambda_{it+1}(s)$, $\forall i, \forall t$ and then only wage changes would alter work behavior. There is, however, potentially an additional source of information available to women in $t + 1$ that was unavailable to women at time t : L_t .

If generation $t + 1$ observed L_t perfectly, they would be able to back out the true value of β as a result of the law of large numbers (i.e., using equation (8)). While assuming that information about how many women worked in the past is totally unavailable seems

²³Each property follows directly from taking the appropriate derivative of equation (7).

²⁴For simplicity, I assume that a daughter does not observe a working mother's realization of B . If she could observe B_i directly this would be another signal used to augment her information. This would make the model less analytically tractable since the population would bifurcate but with reasonable noise in B_i this would not alter the main results.

²⁵This can be thought of as the family's idiosyncratic culture since the signal is fixed along a dynasty but is updated with the common social belief of λ_t .

extreme,²⁶ the notion that this variable is completely informative seems equally implausible – it is merely an artifact of the simplicity of the model.²⁷ I employ, therefore, the conventional tactic in this literature and assume that women observe a noisy function of the aggregate proportion of women worked.²⁸

In particular, I assume that women observe a noisy signal of L_t , given by y_t , where

$$y_t(\beta; \lambda_t) = L_t(\beta; \lambda_t) + \eta_t \quad (9)$$

and where $\eta_t \sim N(0, \sigma_\eta^2)$ with a pdf denoted by $h(\cdot; \sigma_\eta)$. The assumption that η (or z in the example developed below) is distributed normally should be taken as an approximation made for analytical simplicity.²⁹

Alternatively, one could assume that individuals perfectly observe aggregate LFP but are uncertain about the value of a parameter that affects the distribution of the idiosyncratic disutility of work, $G(l)$. The realization of that parameter (e.g., the mean of the distribution of l which is now set at zero) would change randomly every period (for example, by depending on an unobservable aggregate shock in the economy).³⁰ One very simple formulation of this alternative is to assume that a fixed fraction $\omega > 0$ of the population is subject to extreme preference shocks such that a *random* proportion z_t of this fraction works whereas the remaining proportion $1 - z_t$ does not work, independently of wages. Thus the aggregate proportion of women who work, $\tilde{L}_t(\beta; \lambda_t, z_t)$, would be given by $\tilde{L}_t(\beta; \lambda_t, z_t) = (1 - \omega)L_t(\beta; \lambda_t) + z_t\omega$ where $L_t(\beta; \lambda_t)$ is given again by equation (8). Assuming that z is distributed normally then yields an expression equivalent to (9).

Returning to the original model, using Bayes law after observing y_t generates an updated common belief for generation $t + 1$ of:

$$\lambda_{t+1}(\lambda_t, y_t) = \lambda_t + \ln \frac{h(y_t|\beta = \beta_L)}{h(y_t|\beta = \beta_H)} = \lambda_t + \left(\frac{L_t(\beta_L; \lambda_t) - L_t(\beta_H; \lambda_t)}{\sigma_\eta^2} \right) (y_t - \bar{L}_t(\lambda_t)) \quad (10)$$

where $\bar{L}_t(\lambda_t) = \frac{L_t(\beta_L; \lambda_t) + L_t(\beta_H; \lambda_t)}{2}$.

Figure 2 summarizes the time line for the economy. Individuals start period t with a common prior, λ_t . Each woman updates the common prior with her inherited private signal

²⁶This is the assumption in Fogli and Veldkamp (2007).

²⁷In particular, with additional sources of heterogeneity in the model, backing out the true value of β would require agents to know the geographic distribution of male earnings and female (potential) earnings and how they were correlated within marriages, the distribution of preferences, the geographic distribution of shocks to technology and preferences, etc.

²⁸This is the strategy used in finance, for example, by introducing noise traders. An alternative assumption is that agents know the work behavior of a small number of other women in their social circle (as in Banerjee and Fudenberg (2004)) which is imperfectly correlated with that of other social circles. This yields similar results. It has the advantage, for the calibration, of not requiring a specification of an aggregate shock but the disadvantage of being sensitive to assumptions about the size of a woman's social group. Amador and Weill (2009) also obtain an S shape in the behavior of aggregate investment by assuming that agents observe a noisy private signal of other's *actions* as well as a noisy public signal of aggregate behavior. They are interested in the welfare properties of the two sources of information.

²⁹One can make alternative assumptions about this distribution (e.g., an appropriately truncated normal pdf) but this renders the analytical expressions and computations considerably more cumbersome.

³⁰See Chamley (1999).

and makes her work decision. This generates an aggregate L_t and a noisy public signal y_t . Generation $t + 1$ observes y_t and uses it to update the old common prior (λ_t), generating λ_{t+1} – the “culture” of generation $t + 1$.³¹ It should be noted that the assumption that women in $t + 1$ inherit λ_t (or λ_{it}) which they update with the information contained in y_t is mathematically equivalent to the assumption that each generation t is born with the same common prior of λ_0 and has access to the entire preceding history of y_τ , $\tau = 0, 1, 2, \dots, t$ which they use to update λ_0 . Both yield the same λ_{t+1} (or of λ_{it+1}).

4 The Dynamic Evolution of Culture and LFP

The learning model developed above has several important properties that will prove useful in generating LFP dynamics similar to those in figure 1. It also gives rise to a new role for wage changes and their effect on women’s labor supply. Below I discuss both in turn. It is useful to note first that beliefs in this model are unbounded. Hence, in the long-run beliefs must converge to the truth.³² Since female LFP has been increasing over time, this implies that it is more likely that $\beta = \beta_L$ and we shall henceforth assume that this is the case.

4.1 Generating an S-Shape in Culture and LFP

A key feature of this model is that it naturally generates an S-shaped LFP curve. To understand why this is so, we start by noting from (8) that, ceteris paribus, the size of change in λ is positively related to the size of the change in LFP. Thus, understanding the shape of the LFP curve requires us to know how λ varies over time.

To determine how λ tends to evolve, first note that given $\beta = \beta_L$, we can rewrite (10) as $\lambda_{t+1} = \lambda_t + \left(\frac{L_t(\beta_L; \lambda_t) - L_t(\beta_H; \lambda_t)}{\sigma_\eta^2} \right) \left(\eta_t + \frac{L_t(\beta_L; \lambda_t) - L_t(\beta_H; \lambda_t)}{2} \right)$. Taking its expected value, it is easy to see that $E_t(\lambda_{t+1}) - \lambda_t$ is increasing in the difference between the aggregate proportion of women who work when $\beta = \beta_L$ relative to those who work when $\beta = \beta_H$, i.e., in $\Delta L_t \equiv L_t(\beta_L; \lambda_t) - L_t(\beta_H; \lambda_t)$.

To understand the determinants of variation in ΔL_t , it is easiest to first abstract from the heterogeneity in l types. Note that for a given $l \in (\underline{l}, \bar{l})$

$$\Delta L_{lt} = F(s_l^*(\lambda_t) - \beta_L; \sigma_\varepsilon) - F(s_l^*(\lambda_t) - \beta_H; \sigma_\varepsilon). \quad (11)$$

Taking the derivative of (11) with respect to s_l^* yields the f.o.c.:

$$f(s_l^* - \beta_L) - f(s_l^* - \beta_H) = 0 \quad (12)$$

³¹As explained below, we can think of generation τ as having a shared culture given by λ_τ . Individual deviations around the median λ_τ (given by the normal distribution of $\lambda_{i\tau}(s)$) constitute the distribution of beliefs induced by different individual’s dynastic histories (i.e., by their inheritance of different s).

³²See, e.g., Smith and Sorensen (2001). Chamley (2004) gives an excellent explanation of the conditions required for cascades to occur.

Recalling that $f(s_l^* - \beta) = \frac{1}{\sqrt{2\pi}\sigma_\varepsilon} \exp - \left\{ \left(\frac{(s_l^* - \beta)^2}{2\sigma_\varepsilon^2} \right) \right\}$, this implies that (11) is minimized at $s_l^* = \pm\infty$ and maximized at $s_l^* = \bar{\beta}$.

It follows from the analysis above that if the critical signal s_l^* is far from $\bar{\beta}$ in absolute value, (11) will be small, the updating of beliefs will be small, and the change in work behavior next period will also be small.³³ Why does this happen? When women require an extreme signal in order to be willing to work, the difference in the proportion of woman who work across the two states L, H is small (see the small shaded area in figure 3). This renders the aggregate signal $y_t(\beta; \lambda)$ less informative as its variance across the two possible states will be swamped by the variance of the aggregate noise term η . Thus, the intergenerational updating of beliefs will be slight and the change in the proportion of women who work that period, ceteris paribus, will likewise be small. The opposite is true when s_l^* is close to $\bar{\beta}$ (see the larger shaded area in figure 3).³⁴

The intuition above can now be applied to the economy with heterogeneity in l types, for which a similar conclusion holds. Taking the derivative of ΔL_t with respect to λ and using (7), we obtain:

$$\frac{\partial}{\partial \lambda_t} (\Delta L_t) = \left(\frac{\sigma_\varepsilon^2}{\beta_H - \beta_L} \right) \int_{\underline{l}}^{\bar{l}} [f(s_l^*(\lambda_t) - \beta_L) - f(s_l^*(\lambda_t) - \beta_H)] g(l) dl \quad (13)$$

Thus, if the critical signal $s_l^*(\lambda_t)$ is far from $\bar{\beta}$ for a substantial proportion of individuals, (13) will be small in absolute value, intergenerational updating will be small, and the evolution of LFP over time will be slow. The opposite is true when the critical signal is close to $\bar{\beta}$ for a substantial proportion of individuals.³⁵

The S -shape follows from the logic above. If parameter values are such that initially a substantial proportion of women require an extreme signal (i.e. very large negative values of s) in order to be willing to work, then learning will be slow for several generations since little information will be revealed by the aggregate signal y . During this time, there will be little change in the aggregate proportion of women who work. This is the first (slowly rising) portion of the S curve for LFP. Once beliefs have become more moderate and a large proportion of women require less extreme values of s to change their work behavior, learning accelerates since the aggregate signal is very informative leading to large changes in aggregate LFP over time. This is the second (rapidly rising) portion of the S curve. Lastly, once beliefs are optimistic and women, on average, require extreme signals (i.e., very large positive values of s) in order to choose *not to work*, learning once again slows down since the aggregate signal becomes relatively uninformative. Changes in aggregate female

³³Note that the distance of s^* from $\bar{\beta}$ depends not only on the value of λ , but also on w_f and w_h as these affect a woman's willingness to work as well.

³⁴Note that a normal distribution is not critical for this result. Rather one requires a cdf that, in general rises slowly at the beginning and then more rapidly (either slowing down again or not towards the end). This is a common shape, e.g., binomial, Poisson, some log normal distributions, the chi distribution, etc.

³⁵Note that, at each λ , s_l^* ranges from $-\infty$ (for $l = \underline{l}$) to $+\infty$ ($l = \bar{l}$). Thus, what determines the general shape of the LFP curve is how the value of s_l^* evolves for the range of $l \in (\underline{l}, \bar{l})$ where is populated by a substantial mass of individuals. It should also be noted here that the assumption of a normal distribution for l does not play an important role.

LFP during this period are correspondingly small. This is the third and last (slowly rising) portion of the S curve. As the time horizon goes off to infinity, beliefs converge to the truth, so any further changes in LFP result solely from changes in wages.

4.2 A New Role for Wages

The learning model generates a novel role for wages. As in a standard labor supply model, an increase in women’s wages increases female LFP that period. Learning, however, introduces an additional *dynamic* effect. In particular, wage changes affect the *pace* of intergenerational learning, i.e., the magnitude of $E_t(\lambda_{t+1} - \lambda_t)$. This does not happen, as one might initially think, by greater numbers of working women providing more information about the welfare consequences of working.³⁶ Rather, the mechanism operates by having wage changes affect s^* and, through this, affect the *informativeness* of the aggregate work outcome. This is true not only for changes in female wages, but also for changes in male wages, for technological change that facilitates women’s market work (e.g., the washing machine in Greenwood et al (2005) or work-enhancing improvements in maternal health as in Albanesi and Olivetti (2009b)), or for policies that increase the attractiveness of work.

The above does not imply that increases in women’s wages will always increase the pace of learning. The opposite effect would be present if women were initially very optimistic about β and hence required, on average, very high values of s^* in order not to work. In that case, wage increases would tend to decrease the amount learned by the next generation.

To conclude, wage changes (or other changes that affect the pace of learning) have a dynamic externality in this model that is absent in more traditional settings. The model thus yields a very different perspective on how one should evaluate the effects of changes in wages, technology, and policy and one of the main objectives of the next section will be to evaluate the quantitative importance of this dynamic learning effect in the historical evolution of female LFP.

5 Quantitative Analysis

In this section I examine the quantitative potential of the simple learning model. Given a path of wages, the degree of heterogeneity in the taste for work/leisure (σ_l), an initial probability that $\beta = \beta_L$, a distribution of private information (σ_ε), and a variance for the noise in the observation of aggregate LFP (σ_η), the model delivers predictions for the time path of female LFP and beliefs. The model also delivers predictions with respect to the quantitative importance of wages and beliefs in changing female LFP over time. This will allow us to assess the significance of the new dynamic learning role of wage changes (discussed in section 4.2) over the last century.

It should be noted that a more realistic analysis would nest the learning mechanism within a richer multi-period dynamic model of female labor supply (e.g. Eckstein and Lifschitz (2009) or Attanasio, Low, and Sanchez-Marcos (2008)). This would require one

³⁶It would be easy though to incorporate this additional channel which is standard in the literature.

to restrict attention to a significantly shorter time period for which more detailed data is available (in most cases, after 1940). Wages though are a key variable that are included and several papers have found them to be an important determinant of the increase in married women’s LFP by several authors (see, e.g., Jones, Manuelli, and McGrattan (2003) or Heathcote, Storesletten, and Violante (2010)). As this is the first paper to explore the dynamic contribution of the evolution of cultural beliefs to women’s work, assessing its quantitative potential in a simple model that allows both the theory and the calibration to be fairly transparent is an important first step to subsequently developing more complicated quantitative dynamic models.

5.1 Calibration Strategy

In the model, married women decide whether to engage in market work taking their husbands’ earnings and their own potential earnings as given. Given the paucity of data prior to 1940, I use the median earnings of full-time white men and women for which some data is available as of 1890.³⁷ This choice exaggerates the earnings of working women in general, as some work less than full time (although part time work is not quantitatively important until after 1940). As shown in the online Appendix, however, the main conclusions are robust to reasonable alternatives. After 1940 I use the 1% IPUMS samples of the U.S. Census for yearly earnings.³⁸

The evolution of married women’s LFP from 1880 to 2000 is shown in Figure 1. These percentages are calculated from the US Census for married white women (with spouse present), born in the US, between the ages of 25 and 44, who report being in the labor force (non-farm occupations and non-group quarters). I calibrate the model to match female LFP in 1980, 1990 and 2000. The remainder of the LFP series is generated endogenously by the model. I also require the model to match the own and cross-wage elasticity in 2000, the cross-wage elasticity in 1990 and the relative probability of a woman working in 1980 (conditional on whether her mother worked). See table 1 for a list of the calibration targets.

I take the elasticity estimates from Blau and Kahn (2007) who use the March CPS 1989-1991 and 1999-2001 to estimate married women’s own-wage and husband’s-wage elasticities along the extensive margin. For the year 2000, Blau and Kahn estimate an own-wage elasticity of 0.30 and a cross-elasticity (husband’s wage) of -0.13. The cross elasticity in 1990 is -0.14.³⁹ The findings are robust to matching instead the 1990 own-wage elasticity (0.44). In general, the model is not very sensitive to the choice of calibration targets.

To calculate the probability that a woman worked in 1980 conditional on her mother’s work behavior, I use the response from the General Social Survey (GSS) question “Did your mother ever work for pay for as long as a year, after she was married?” (MAWORK) from several years in 1977-1983 to determine whether a woman’s mother worked. For each sample year, I calculated the ratio between the probability of a woman working given that

³⁷I restrict the sample to white women as black women have had a different LFP trajectory with much higher participation rates earlier on.

³⁸See the online Appendix for details.

³⁹See the online Appendix for a discussion of the exact specification chosen.

her mother worked to the probability of a woman working given that her mother didn't work (henceforth referred to as the work risk ratio). Averaging across the sample years yielded a risk ratio of 1.13., i.e., women whose mother worked are 13% more likely to work in 1980 than women whose mother didn't work. The GSS sample for this calculation was restricted to all white married women between the ages 25-45 who were born in the U.S. In the calibration each period is a decade and, for the purpose of computing the work risk ratio, daughters are assumed to make their work decisions two periods after their mothers (i.e., a separation of 20 years). The calibration targets are summarized in table 1.

5.2 Calibration Results

Before turning to the calibration of the full model, it is instructive to calibrate a simpler version of it in which β is perfectly known. Shutting down the learning channel implies that only wage changes can explain changes in labor supply. The unknown parameters are now three instead of seven (γ , β , and σ_l) allowing us to match a restricted set of targets. In particular, I choose to match female LFP, a woman's own-wage elasticity, and her cross-wage (husband's wage) elasticity, all in the year 2000. These are useful statistics as the ratio of the elasticities gives information about the curvature of the utility function and an elasticity and LFP value combined give information both about the magnitude of the common disutility of working, β , and about the necessary dispersion of l types in order to generate a given response to a change in wages. Furthermore, as shown in the online Appendix, the simplicity of the model allows one to solve for the parameter values analytically, yielding $\gamma = 0.503$, $\sigma_l = 2.29$ and $\beta = 0.321$. To interpret the magnitude of the common expected disutility of working, β , note that this is equivalent to 4.7% of the utility obtained from consumption in 2000 if a woman worked or 22.4% of the difference in the consumption utility between working and not working in that year. In 1880, these numbers are larger: 10.4% and 88.1% respectively.

As can be seen from the light line (with squares) with the caption "No Learning Model" in Figure 4, the model with no learning does a terrible job of matching the female LFP data (depicted by the unfilled circles), grossly overpredicting LFP in all decades other than 1990 and, by construction, the calibrated target in 2000. As discussed in the online Appendix, this failure of the model is robust to a wide range of values for the elasticities, to assuming that the woman obtains only some share of joint consumption, and to the exact choice of the earning series. It is also important to note that the failure of the model without learning does not follow from the structure of the model itself. If, instead of calibrating, one chooses parameter values for γ , σ_l , and β that minimize the sum of squared errors between the predicted LFP path and the data, this results in a curve that fits the data very well.⁴⁰

We now turn to calibrating the full model. As married women's LFP has been increasing throughout and, as learned from the calibration exercise of the model without learning,

⁴⁰The sum of squared errors in that hypothetical exercise is 0.021 which is less than half of what we obtain with the fully calibrated model below.

changes in wages alone cannot replicate this phenomenon, I assume that the true state of nature is $\beta = \beta_L$. In this case, learning over time about the true cost of working would, *ceteris paribus*, increase LFP.

There is an additional complication in calibrating this model that needs to be addressed—the presence of an aggregate observation 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. Each realization of η_t generates a corresponding different public belief λ_{t+1} in the following period, and consequently a different proportion of women who choose to work after receiving their private signals. Note that we cannot simply evaluate the model at the mean of the expected η shocks (i.e., at zero) since, although λ_{t+1} is linear in η , the work outcomes L_{t+1} are not.

I deal with the aggregate shock in the following way. For each period $t + 1$, given L_t , I calculate the proportion of women who would work, L_{t+1} , for each possible realization of the shock, η_t , i.e., for each induced belief $\lambda_{t+1}(\eta_t)$. Integrating over the shocks, I find the expected value of LFP for that period, $E_t(L_{t+1}(\lambda_{t+1}(\eta_t)))$, and then back out the particular public belief (or shock) that would lead to exactly that same proportion of women working, i.e., I solve for $\lambda_{t+1}^*(\eta_t^*)$ such that $\int_{\eta} L_{t+1}(\lambda_{t+1}(\eta); \lambda_t) h(\eta) d\eta = L_{t+1}(\lambda_{t+1}^*(\eta_t^*))$.⁴¹ Performing this exercise in each period determines the path of beliefs.

As shown in the Appendix, manipulating the ratio of elasticities in this model yields the same value of γ as absent the learning channel, i.e., $\gamma = 0.503$. As in the simplified model, the elasticities and female LFP levels yield information both about how bad women believe it is to work and how much heterogeneity there is across women are in their willingness to work at given wages. Unlike before, however, this dispersion is a function not only of the distribution of the l types, σ_l , but also of the distribution of private information, σ_ε . Furthermore, women’s beliefs about the value of β are evolving over time. The values of LFP from 1980-2000 yield information as well on how rapidly the commonly held portion of beliefs, λ_t , needs to evolve over these decades and hence on the variance, σ_η , of the noise η in the public signal.

Lastly, as mothers and daughters share the same private information, the conditional probability that a woman works as a function of her mother’s work behavior (the work risk ratio, R) also yields information on the evolution of λ and how different the values of β_H and β_L should be. I take a mother (M) and daughter (D) to be separated by twenty years. The work risk ratio is thus given by $R_t \equiv \frac{\Pr(DWork_t | MWork_{t-2})}{\Pr(DWork_t | MNotWork_{t-2})}$ and in the benchmark I assume that daughters inherit perfectly their mother’s private signal whereas their l type is a random draw from the normal distribution that is *iid* across generations.⁴² The details of the calculation are given in the online Appendix.

Table 1 below shows the calibration targets (column 1) and the values obtained in the

⁴¹For the computation, I take a large number of draws of entire histories for η (500 histories) in order to calculate the expected value of L . See the online Appendix for details.

⁴²Thus, this model yields a positive correlation between a mother and her daughter’s work “attitudes” ($E_{it}(\beta) + l_i$ and $E_{i',t+1}(\beta) + l_{i'}$ where i indexes the mother and i' the daughter). For evidence on this, see Farré-Olalla and Vella (2009).

<i>Calibration Targets</i>		Model with known β					Learning Model	
Own-Wage Elast.(2000)	0.30	0.30					0.29	
Cross-Wage Elast.(2000)	-0.13	-0.13					-0.13	
Female LFP(2000)	0.734	0.734					0.744	
Female LFP(1990)	0.725	0.725					0.716	
Cross-Wage Elast.(1990)	-0.14	-0.13					-0.14	
Female LFP(1980)	0.586	0.687					0.585	
Work Risk Ratio(1980)	1.13	1.00					1.13	

<i>Parameters</i>	γ	σ_L	β	β_H	β_L	$P_0(\beta = \beta_L)$	σ_ε	σ_η
Model with known β	0.503	2.293	0.321					
Learning Model	0.503	2.085		4.935	0.001	0.086	5.288	0.055

Table 1: *All elasticities are from Blau & Kahn (2007). The work risk ratio uses data from GSS (see text). The values **in bold** are the model's predicted values for its calibration targets.*

calibrated learning model (column 3). The second column reports the values obtained in the prior calibration exercise in the model with no learning.

5.3 The Calibrated Model: LFP

The LFP predictions from the calibrated model are shown in Figure 4. The (blue) solid line with circles shows the evolution of the expected value of female LFP and the (red) dashed line shows the evolution of the probability that the true state is β_L that is held by the median woman. As can be seen from the figure, the calibrated model on the whole does a fairly good job of replicating the historical path of married women's LFP.⁴³ It under-predicts LFP from 1930 to 1970, however, and slightly over-predicts it from 1880 to 1900. The undershooting during 1930-1970 may indicate that another factor, such as technological change in the household or less employment discrimination on the part of firms, was also responsible for the higher levels of LFP during this period. Recall that a characteristic of the learning model is that technological change that occurred in the 1930s and 1940s would have had repercussions in later decades through the dynamic impact of technological change on learning discussed earlier. World War II may also have played a role by increasing women's willingness to work during the war years and this, in turn, increasing the pace of intergenerational learning.

It is instructive to ask why the learning model yields such a different time path for female LFP than the model with no learning. As noted previously, the calibration implies that both models must have the same value of γ . Furthermore, the difference in the standard deviation of the normal distribution of types is relatively small: 2.29 versus 2.09. Lastly, the expected value of β in 2000 is not very different from the value of β obtained when learning is eliminated: the individual with median beliefs has an expected value of β of 0.26

⁴³The sum of squared errors (between actual and model predicted LFP) is 0.052.

whereas without learning $\beta = 0.32$.⁴⁴ Thus, it is the endogenous evolution of the expected value of β in the learning model that is responsible for the difference in the LFP behavior observed over time across the two models. Whereas by construction the expected value of β remains constant when learning is eliminated, the median expected value of β in the learning model is close to 4.51 in 1880 and evolves over time to 0.26 in 2000. This allows LFP to respond in dramatically different ways over time by as reflected in the evolution of wage elasticities over time.⁴⁵

5.4 The Calibrated Model: Beliefs

The model also generates predictions about the evolution of social attitudes. Individuals start out in 1880 with pessimistic beliefs about how costly it is to work: the median individual assigns around a 9% probability to the event $\beta = \beta_L$.⁴⁶ Beliefs evolve very slowly over the first seventy years (remaining below 20% for the median individual during this period). Then, as of 1960, the change in beliefs accelerates. The median individual jumps from assigning a probability of 26.0 to β_L in 1960 to 83.8% in 1980. By 2000, the median probability assigned to $\beta = \beta_L$ is 94.7%. Individual beliefs are highly dispersed as the private signal has a large variance. Figure 5 shows the equilibrium path of beliefs for the individuals with the median private signal and for those with private signals two standard deviations below and above this median.

In order to contrast the model’s prediction about social attitudes with those of the data I make use of the poll results discussed previously. The solid line in Figure 6 shows the percentage of white married women between the ages of 25-44 who answered that they approved when asked “Do you approve or disapprove of a married woman earning money in business or industry if she has a husband capable of supporting her?” To compare the model’s predictions and the poll data, one needs to map the distribution of probabilities at any point in time to the binary variable (“approve” versus “disapprove”) available in the data. To do this, I use the fact that $\lambda_t \sim N\left(\lambda_t - \left(\frac{\beta_H - \beta_L}{\sigma_\varepsilon^2}\right) (\beta_L - \bar{\beta}), \frac{(\beta_H - \beta_L)^2}{\sigma_\varepsilon^2}\right)$ to find the cutoff belief ($\tilde{\lambda}$) in the model’s predicted belief distribution in 1940 such that the proportion of women with more optimistic beliefs (i.e. those with $\lambda \geq \tilde{\lambda}$) is equal to the proportion of women who approved of women working in the poll data at that time, 0.27.⁴⁷ This allows me to identify, for every period, the proportion of women who would approve of married women working, i.e., those with beliefs above $\tilde{\lambda}$, and contrast it with the data. The dashed line in Figure 6 shows the model’s predictions for the proportion of women who approve of married women working. As can be seen from the figure, the shape of the

⁴⁴Note that the calibration does not require both models to have the same values of σ_l and β (for 2000) since the learning model has an additional source of heterogeneity (intra-generational heterogeneity in beliefs induced by private signals) which affects the elasticity.

⁴⁵See figure 13 in the online Appendix contrasting the own and cross wage elasticities of the model with and without learning.

⁴⁶It is worth noting that this pessimism is fully congruous with social views concerning married women working (see the quotes in the introduction and evidence in the poll data).

⁴⁷There is no poll data for 1940. Instead I use a linear interpolation between the 1938 and 1945 points to obtain it.

beliefs path is quite similar to the data path but the level is higher throughout by some 10 percentage points.

5.5 The Quantitative Contributions of Wages and Beliefs

To investigate the quantitative contributions of earnings and beliefs, we start by not allowing public beliefs to evolve (i.e., shutting down the public signal). First, we can freeze beliefs at the 1880 distribution and ask how labor force participation would have evolved had mothers continued to endow their girls with their private information but in the absence of any public signal (y_t). As show by the bottom line (with the caption “LFP if no public updating”) in Figure 7, female LFP would have barely exceeded 10% by the year 2000. Alternatively, we can assume that agents had full information about the true value of β , i.e., $\beta = \beta_L$. This scenario is given by the line with the caption “full information LFP”. It too predicts a very different trajectory, with LFP starting close to 63% in 1880 and slowly evolving to 80% by 2000. Thus, we can conclude that the actual dynamics of beliefs induced by learning is essential to producing the predicted path of female LFP.

To distinguish between the static and dynamic effects of wage changes on female LFP we can perform the following instructive decomposition. First, we can keep wages constant at their initial 1880 levels, henceforth denoted by w_0 , and let only beliefs change endogenously over time. This generates an LFP path denoted $L(\bar{\lambda}(w_0), w_0)$ and is depicted by the bottom (magenta) line in Figure 8.⁴⁸ Hence, the difference between $L(\bar{\lambda}(w_0), w_0)$ and the level of LFP in 1880 (the dotted horizontal line) measures the contribution of changes in beliefs to LFP absent any changes in wages.

Next we can disentangle the dynamic from the static effects of wages by confronting women with the actual historical earnings path, \bar{w} , but endowing them with the belief path obtained from the exercise above, $\bar{\lambda}(w_0)$. In this exercise, changes in wages affect the relative attractiveness of working for the usual reasons, but their dynamic effect on the intergenerational evolution of beliefs is eliminated by construction. We denote the LFP obtained this way by $L(\bar{\lambda}(w_0), \bar{w})$ and it is shown with (red) x’s in the figure. The difference between $L(\bar{\lambda}(w_0), w_0)$ and $L(\bar{\lambda}(w_0), \bar{w})$ measures the static contribution of wages to the evolution of LFP. Lastly, we allow wages to also influence intergenerational learning and thus beliefs; let the LFP path obtained this way be denoted by $L(\bar{\lambda}(\bar{w}), \bar{w})$.⁴⁹ It is the top (blue) curve shown in Figure 8 and the difference between $L(\bar{\lambda}(\bar{w}), \bar{w})$ and $L(\bar{\lambda}(w_0), \bar{w})$ measures the dynamic contribution of wages through its effect on beliefs.

As can be seen in Figure 8, for the first several decades the static effect of wages is mostly responsible for the (small) increase in LFP. Over time, both the dynamic effect of wages on beliefs and the evolution of beliefs independently of wage changes become increasingly important, with the dynamic effect of wages accounting for over 50% of the change in LFP between 1970 to 1990, which are the decades of largest LFP increases.⁵⁰ We conclude

⁴⁸In this section I use a bar to denote a vector of the time series of a variable that is not being kept constant.

⁴⁹Note that this LFP path is the one predicted by the model and depicted previously in Figure 4.

⁵⁰It should be noted that the decomposition of LFP is not unique. The alternative decomposition, as

from our decomposition that in some decades changes in beliefs induced by higher earnings were critical to the increases in female LFP. Overall, changes in beliefs – both those that would have occurred even had wages remained constant and those induced by changes in wages – played the largest quantitative role in generating the changes in female LFP over the last 120 years.

6 Some Testable Implications of the Model

The model can be used to derive levels of female LFP in out-of-sample years. For example, constructing median earnings for men and women from the 2008 CPS using the same sample selection criteria as before, the model predicts female LFP of 76.58% as compared to 75.57% in the CPS.⁵¹

A more demanding test for the model is to generate LFP paths for women of different education levels. As the Census does not provide wage or education data prior to 1940, I instead examine the implications of the calibrated model from 1940 onwards. To do this, for each decade I categorize women into “high school” if they did not study beyond high-school; women with at least some college education are categorized as “college”. These women are then matched with their husbands and I use the same procedure as before to calculate a median earnings series for each female education type and their respective representative husbands (i.e., the median earnings of the men married to “high school” women and those of the men married to “college” women).⁵²

More problematic is the absence of a prior belief for women in 1940. To deal with this, one can create a prior for 1940, by skill type, by running the model from 1880 onwards using the same initial prior for both types of women as in the calibrated model ($p_{1880} = 0.086$), but employing the appropriate wage series for each type. Given the absence of wage-education data before 1940, I assume that for each decade prior to that year, the ratio of the wages of each education type relative to the wages of the female population at large is at the 1940 level.⁵³ This implies a skill premium for women of 1.4, which seems a reasonable figure. I perform a similar procedure for these women’s husbands (i.e., I assume that the ratio of median wages for the husbands of each type of woman to the median wages of men was at the same ratio as in 1940). This calculation yields 1940 priors of 0.148 for college women and 0.089 for high-school women.

The results of this exercise are given in Figure 9. As can be seen from the figure, the model does a surprisingly good job at capturing the general shape of the increase in LFP for

shown in the online Appendix, yields a similar quantitative pattern.

⁵¹This is the last sample available prior to the start of the recession in December 2007 (the data, as usual, refer to work behavior in the preceding year). Not surprisingly, the recession’s impact renders the model’s predictions for 2010 significantly off: 77.3 as opposed to 72.8.

⁵²See Figure 16 in the online Appendix for a graph of these earning series.

⁵³This is probably not such a bad assumption. During the 1940-1970 period, the ratio of married college women’s median earnings to the median woman’s earnings was relatively constant as was the same ratio for married high school women. For college women the mean of this ratio was 1.21 with a standard deviation of 0.04; for high school women the mean of this ratio was 0.88 with a standard deviation of 0.04. This relatively tight band is also found for each type’s husbands.

both college and high school women. It also does a very good job at predicting LFP levels for college women for the four decades from 1970-2000 but it substantially underpredicts the LFP of high-school women in every decade prior to 2000. It should be noted that the fact that it consistently predicts that college women work more than high school women is a success as this is not guaranteed by the model since a woman's work decision depends on the wage levels of both wives and husbands as well as on beliefs.⁵⁴

The model also has implications concerning the intergenerational speed of learning. As explained in section 4.2, the model implies that wage changes, technological change in the household, or changes in policies that influence women's willingness to work, all have dynamic consequences as they affect the rate at which the next generation learns and consequently the path of LFP. Thus, if one could find exogenous variation in the magnitude of these changes one could use it to examine the dynamic implications of the learning model.

One historical episode that potentially created variation across US states in women's willingness to work is World War II (WWII). As shown by Acemoglu, Autor, and Lyle (2004), women worked more in 1950 (but not in 1940) in those states that had a greater mobilization rate of men during WWII, even after controlling for other differences across these states (e.g., differences in age, education, or racial composition, differences in the importance of farming or occupational structure, or Southern vs non-Southern state differences). They attribute the cause of the variation in female labor supply across states in 1950 to the greater participation of women during the war years, with more women staying in market work after the war ended in the states with higher mobilization rates. This allows them to interpret the mobilization rate as a source of exogenous variation in female LFP across states and to use it to analyze the effect of the latter on male and female wages.

For our purposes here we can make use of the same source of exogenous variation in female LFP to examine the intergenerational implications of the learning model. In particular, the model implies that if states with greater mobilization rates had higher female LFP during the war years either because aggregate demand for women's work was greater (i.e., women's wages were higher) or because these women needed to make up for the decline in household income (i.e. husbands' wages were lower) or in response to the campaign to get women to work as part of the war effort (recall, e.g., the iconic image of "Rosie the Riveter"), *ceteris paribus* this should result in the next generation of women in those states learning more and hence working more. Thus, states with greater mobilization rates during WWII should have greater female labor supply not only in 1950, but also in the next generation. For rather different purposes, this is what was shown in Fernández, Fogli, and Olivetti (2004). Specifically, their paper showed that states with higher mobilization rates tended to have higher labor supply of white married women within a narrow age range (those likely to have children) in 1950, even after controlling for other sources of inter-state variation as in Acemoglu et al (2004). They then found that the effect of the mobilization rate persisted

⁵⁴It should also be noted that we are making a number of data and modelling assumptions that may make it more difficult for the model's predictions to match the data, including the assumption that the unknown cost of working is the same for both types of women, that one learns only from one's type, and that there is no selection into education (e.g., by belief).

for white married women in the *next* generation. This is what they call an “echo” effect of the war since it affected women who would have been too young to feel the direct effects of the war but old enough to have witnessed and learned from the change in LFP of their mothers’ generation.⁵⁵

Fernández et al (2004) interpreted their finding as being consistent with the hypothesis that working mothers transmitted a better view of working women to their sons, making the latter more amenable to having a working wife. This in turn increased the attractiveness for the girls in that generation to prepare to work (e.g., by investing in market-specific, rather than home-specific, human capital) as they knew that men would tend to be more receptive to having a working wife. Their finding, however, is also consistent with the hypothesis that the next generation of women had learned more in those states with a greater mobilization rate. Note that both interpretations involve a change in culture: in one case through a change in men’s preferences (or beliefs) and in the other primarily through a change in women’s beliefs resulting from learning. In reality, both channels are likely to have played a role and the available evidence does not permit one to distinguish between the two.

Lastly, it may be of interest to ask whether other countries’ path of married women’s LFP also look S-shaped. If they did, this would help support the case that a process of cultural change driven by learning took place. Unfortunately there does not appear to be readily accessible data on most other countries’ evolution of female LFP over long periods of time and, for the few that exist, they tend not to distinguish between farm and non-farm workers, marital status, nor among different age groups.⁵⁶ I have, however, been able to plot married women’s LFP for two countries for which relatively lengthy data series exist in tables of published papers: France for 1921-1981 and Great Britain between 1911-1998. As shown in Figure 10, married women’s LFP in these countries seems to follow a general S-shaped curve.⁵⁷

7 Discussion and Conclusion

This paper modeled the joint dynamics of married women’s labor force participation and cultural change. In the model, married women compared the benefits of increased consumption from labor earnings with the expected utility cost of working. This cost was unknown and women’s beliefs about it evolved endogenously over time in a Bayesian fashion. I showed that a simple model with these features, calibrated to key statistics from the later part of the 20th century, generates a time trend of married women’s LFP that is roughly similar to the historical one in the US over the last 120 years. Furthermore, a mapping of the evolution of social beliefs in the model was shown to follow a path that is similar to the poll data.

⁵⁵For greater depth, the interested reader is referred to Fernández et al (2004), and in particular to tables III and IV in that paper.

⁵⁶The restrictions to non-agricultural, married, and within a narrow age group, as used to derive female LFP series in the US, are important as otherwise the shape of the path is also likely to be driven by the transition from an agricultural to an industrial economy or by demographic changes in the population.

⁵⁷See the notes in the online Appendix for details as to how these paths were constructed.

The model makes some simplifying assumptions, including an unchanged psychic cost of working over 120 years. It would not be difficult to incorporate changes in the cost structure, but without direct empirical evidence on the matter, an unchanged cost means fewer free parameters – an important disciplining device in a model with few data points. The model also ignored costs that are endogenous in nature. In particular, by modeling changes in culture as arising solely from a process of learning about exogenous costs, it neglected the endogenous, socially imposed, costs stemming from social (cultural) reactions to married women in the work force. Other assumptions in the model, such as the normal distributions of the noise terms, could easily be replaced with others (e.g., single-peaked distributions and relatively thin tails on the left hand side of the modal frequency) that would preserve the same qualitative features, particularly the S-shaped curve. Introducing risk aversion (with respect to the uncertainty about the long-run payoff from working) is straightforward and would create an additional reinforcing channel for learning.

The calibrated model finds that at the outset women were pessimistic about the true cost of working, consistent with the literature of that period. This lack of neutrality may indicate that particular social forces were at play in determining culture initially. Common economic interests for certain groups in industrial societies at that time (e.g., men or unions), may help explain why most countries shared the view that women working outside the home was harmful. Endogenizing this initial prior, however, is outside the model presented here and might require a political economy framework to explain why certain opinions become dominant.⁵⁸

In future work it would be interesting to investigate quantitatively both the informational role of different social networks and the contribution of social rewards and punishments to changing behavior over time relative to social learning. Some interesting work in this area has been done by Munshi and Myaux (2006) who incorporate strategic interactions in the context of a learning model with multiple equilibria in which individuals decide whether to adopt modern contraception. At a theoretical level, it would also be interesting to explore further the potential inefficiencies that arise due to learning externalities and to examine the possible role for policy. At the empirical level, it is important to depart from focusing exclusively on aggregate features of the data over a very long time horizon. In particular, sharper hypotheses about cultural change over a shorter time period would allow a greater use of microdata and permit one to learn more about the process of cultural diffusion.⁵⁹

⁵⁸As the economy changed, so may have the interests of firms (capitalists) and perhaps men in general with respect to having women in the work force (see, e.g., Fernández (2010b)).

⁵⁹See, e.g. Munshi and Myaux (2006). Bandiera and Rasul (2006) and Conley and Udry (2003) use self-reported data on social contacts to construct networks to test their models of learning about new technologies. Mira (2007) structurally estimates his model using Malaysian panel data.

References

- [1] Acemoglu, Daron, David Autor, and David Lyle, “Women, War, and Wages: The Effect of Female Labor Supply on the Wage Structure at Midcentury,” *Journal of Political Economy*, CXII (2004), 497–551.
- [2] Akerlof, G. and Kranton, R. E. (2000), “Economics and identity,” *Quarterly Journal of Economics* 115, 715–33.
- [3] Albanesi, Stefania and Claudia Olivetti (2009a), “Home Production, Market Production and the Gender Wage Gap: Incentives and Expectations,” *Review of Economic Dynamics* 12, 80-107.
- [4] Albanesi, Stefania and Claudia Olivetti (2009b), “Gender Roles and Medical Progress,” working paper.
- [5] Amador, Manuel and Pierre-Oliver Weill (2009), “Learning from Private and Public Observations of Others’ Actions,” working paper.
- [6] Antecol, H. (2000), “An Examination of Cross-Country Differences in the Gender Gap in Labor Force Participation Rates,” *Labour Economics* 7, 409-426.
- [7] Attanasio, Orazio, Hamish Low, and Virginia Sánchez-Marcos (2008), “Explaining Changes in Female Labor Supply in a Life-Cycle Model,” *American Economic Review*, 98, 1517-1552.
- [8] Bandiera, Oriana and Imran Rasul (2006), “Social Networks and Technology Adoption in Northern Mozambique,” *Economic Journal* 116, 869-902.
- [9] Banerjee, Abhijit and Drew Fudenberg (2004), “Word of Mouth Learning,” *Games and Economic Behavior*, Vol. 46 (1), pp. 1-22.
- [10] 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, 681-701.
- [11] Bikhchandani, S., D. Hirshleifer, and I. Welch (1992), “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades,” *Journal of Political Economy*, 100, 992-1026.
- [12] Bisin, A., and T. Verdier (2000), “Beyond the Melting Pot: Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits,” *Quarterly Journal of Economics* 115, 955-988.
- [13] Blau, Francine D. and Lawrence M. Kahn (2006), “Changes in the Labor Supply Behavior of Married Women: 1980-2000,” *Journal of Labor Economics* 25, 393-438.

- [14] Blundell, R.W. and T. MaCurdy, (2000), “Labor Supply: A Review of Alternative Approaches”, in Ashenfelter and Card (eds), *Handbook of Labor Economics*, Elsevier North-Holland.
- [15] Bowles, Sam (1998), “Endogenous preferences: the cultural consequences of markets and other economic institutions,” *Journal of Economic Literature* 36, 75–111.
- [16] Burda, Michael, Daniel Hamermesh, and Philippe Weil (2007), “Total Work, Gender and Social Norms,” NBER wp 13000.
- [17] Chamley, Christophe (1999), “Coordinating Regime Switches,” *Quarterly Journal of Economics* 114, 869-905.
- [18] Chamley, Christophe (2004), *Rational Herds*, Cambridge University Press.
- [19] Collins, Gail (2003), *America’s Women: 400 Years of Dolls, Drudges, Helpmates, and Heroines*, Harper Perennial.
- [20] Conley, T. and Chris Udry (2010), “Learning About a New Technology: Pineapple in Ghana,” *The American Economic Review* 100, 35-69.
- [21] 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.
- [22] Eckstein, Zvi and Osnat Lifshitz (2009) “Dynamic Female Labor Supply,” IZA discussion paper.
- [23] Farré-Olalla, Lidia and Francis Vella (2009), “The Intergenerational Transmission of Gender Role Attitudes and its Implications for Female Labor Force Participation,” working paper.
- [24] Fernández, Raquel (2007), “Women, Work, and Culture,” *Journal of the European Economic Association*, 5, 305-332.
- [25] Fernández, Raquel (2010a), “Does Culture Matter?,” in J. Benhabib, M. Jackson and A. Bisin (ed.), *Handbook of Social Economics*, edition 1, volume 1A, chapter 11, Elsevier.
- [26] Fernández, Raquel (2010b), “Women’s Rights and Development,” working paper.
- [27] Fernández, Raquel and Alessandra Fogli (2009), “Culture: An Empirical Investigation of Beliefs, Work, and Fertility,” *American Economic Journal: Macroeconomics* 1, 146-177.
- [28] 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.

- [29] Fogli, Alessandra and Laura Veldkamp (2007), “Nature or Nurture? Learning and Female Labor Force Participation,” working paper.
- [30] 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.
- [31] Gallup, George H. (1972), *The Gallup Poll; public opinion, 1935-1971* New York: Random House.
- [32] Galor, Oded and David Weil (1996), “The Gender Gap, Fertility, and Growth,” *American Economic Review* 86(3), 374-87.
- [33] Gayle, George-Levi and Limor Golan (2010), “Estimating a Dynamic Adverse Selection Model: Labor Force Experience and the Changing Gender Earnings Gap,” working paper.
- [34] Geroski, P. A. (2000), “Models of technology diffusion,” *Research Policy* 29, 603-625.
- [35] Goldin, Claudia (1990), *Understanding the Gender Gap: An Economic History of American Women*, New York: Oxford University Press.
- [36] 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.
- [37] Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu (2005), “Engines of Liberation,” *Review of Economic Studies* 72(1) 109-133.
- [38] Griliches, Zvi (1957), “Hybrid Corn: An Exploration in the Economics of Technological Change”, *Econometrica* 25, 501-22.
- [39] Hazan and Maoz (2002), “Women’s labor force participation and the dynamics of tradition,” *Economic Letters*, 75, 193-198.
- [40] Heathcote, Jonathan, Kjetil Storesletten and Giovanni L. Violante (2010), “The Macroeconomic Implications of Rising Wage Inequality in the United States,” *Journal of Political Economy* 118, 681-722.
- [41] Jones, Larry, Rodolfo Manuelli, and Ellen McGrattan (2003), “Why are Married Women Working so much,?” working paper.
- [42] Knowles, John (2007), “High-Powered Jobs: Can Contraception Technology Explain Trends in Women’s Occupational Choice?” working paper.
- [43] Lee, Donghoon and Ken Wolpin (2010), “Accounting for Wage and Employment Changes in the US from 1968–2000: A Dynamic Model of Labor Market Equilibrium,” *Journal of Econometrics*, 156 (1), 68-85.

- [44] Mira, Pedro (2007), “Uncertain infant mortality, learning and life-cycle fertility,” *International Economic Review* 48, 809-846.
- [45] Munshi, Kaivan (2004), “Social Learning in a Heterogeneous Population: Social Learning in the Indian Green Revolution,” *Journal of Development Economics*, 73, 185-213
- [46] Munshi, Kaivan and Myaux, J. (2006), “Social norms and the fertility transition,” *Journal of Development Economics*, 80, 1– 38
- [47] Riboud, Michelle (1985), “An Analysis of Women’s Labor Force Participation in France: Cross-Section Estimates and Time-Series Evidence,” *Journal of Labor Economics*, 3(1,part 2), S177-S200.
- [48] Smith, Lones and P. Sorensen (2001), “Pathological Outcomes of Observational Learning,” *Econometrica*, 68, 371-398.
- [49] Vives, Xavier (1993), “How Fast do Rational Agents Learn,” *The Review of Economic Studies*, 60 (2), 329-347.

8 Figures

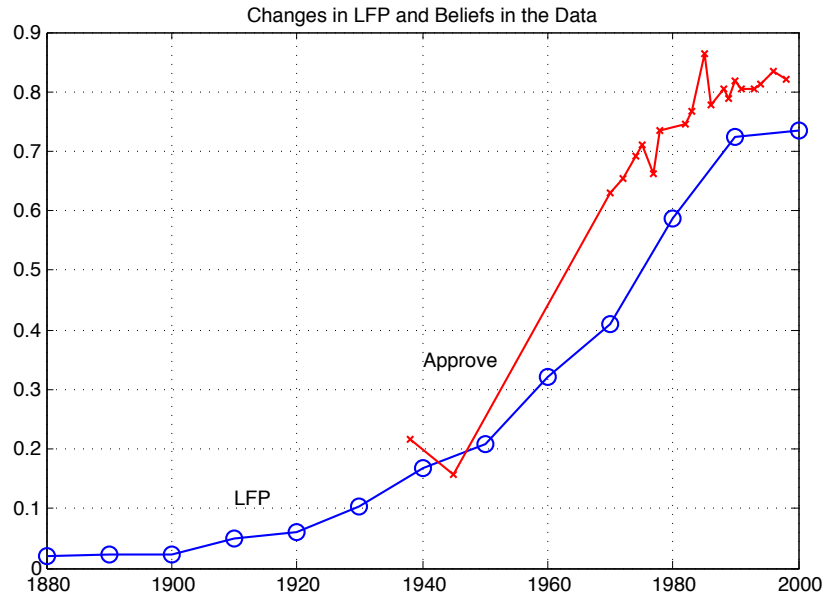


Figure 1: LFP of married women (circles) and percent who approve of woman working if husband can support her (stars). See the online Appendix for details of data construction.

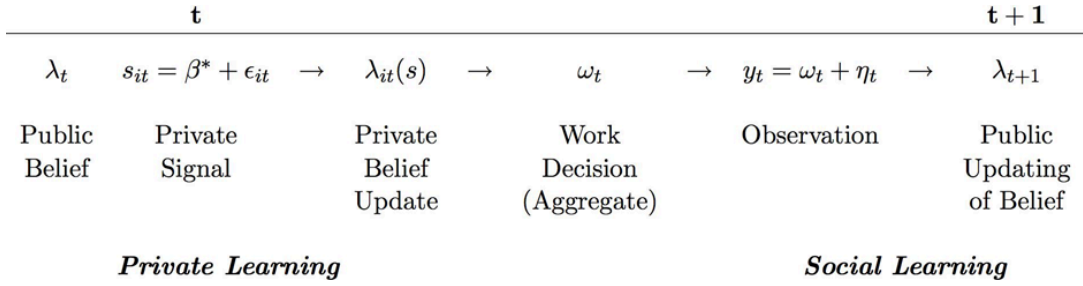


Figure 2: Timeline of Learning Model




Figure 3: Signal extraction

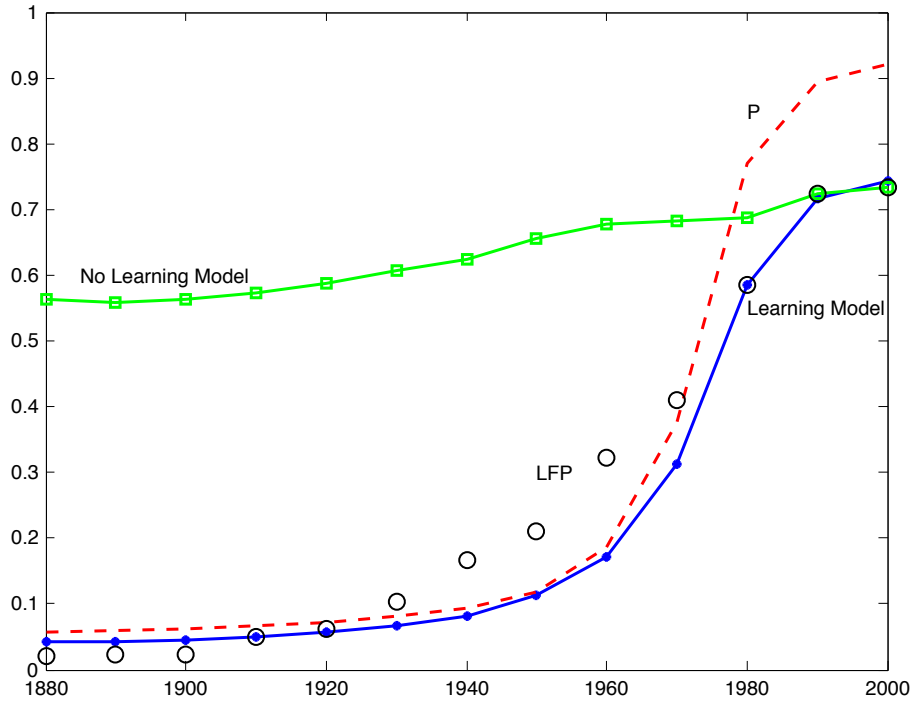


Figure 4: The dashed (red) line (P) is the belief path of the median individual. The sum of squared errors (distance of predicted LFP from actual LFP) is 0.052.

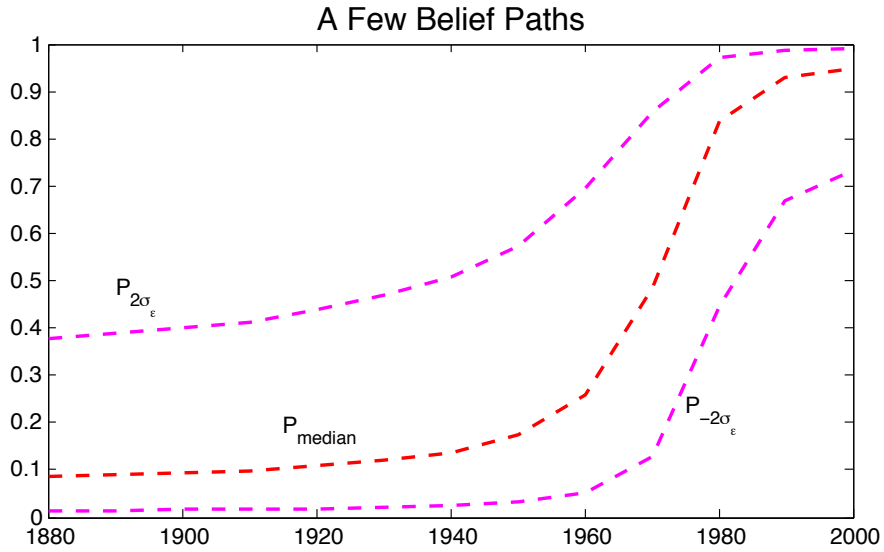


Figure 5: This shows $\Pr(\beta = \beta_L)$ for agents with $s = \beta_L$ and $s = \beta_L \pm 2\sigma_\epsilon$.

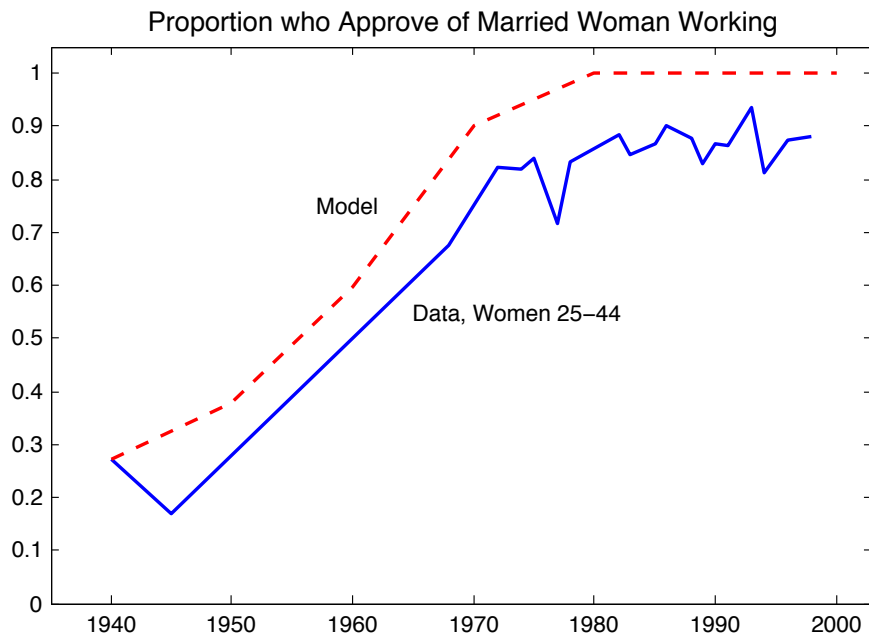


Figure 6: Proportion who approve of woman working if husband can support her, data vs model. Poll data as described in the online Appendix. The sample is restricted to white married women between the ages of 25 and 44.

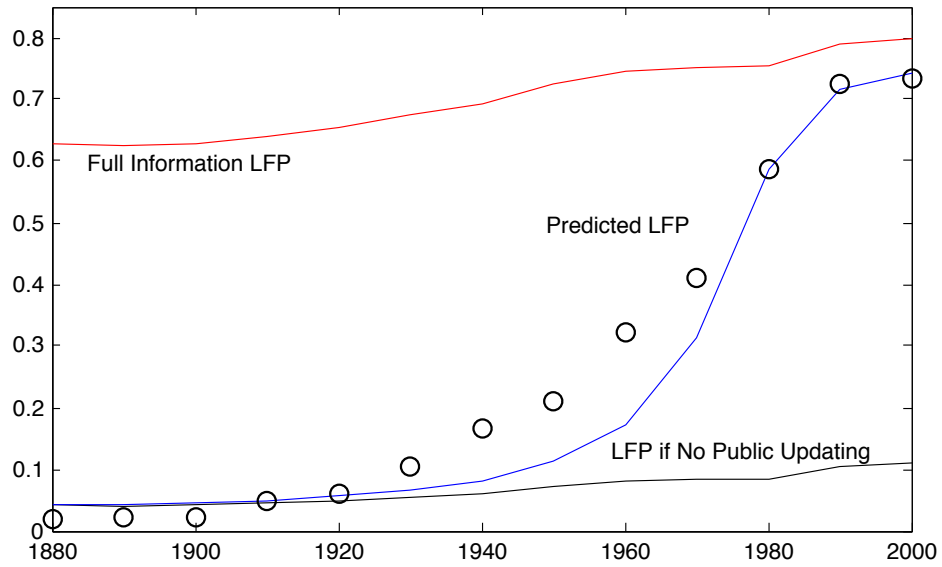


Figure 7: Uses the solution parameters from calibrated model but without public learning.




Figure 8: Decomposition of LFP. See the text for notation.

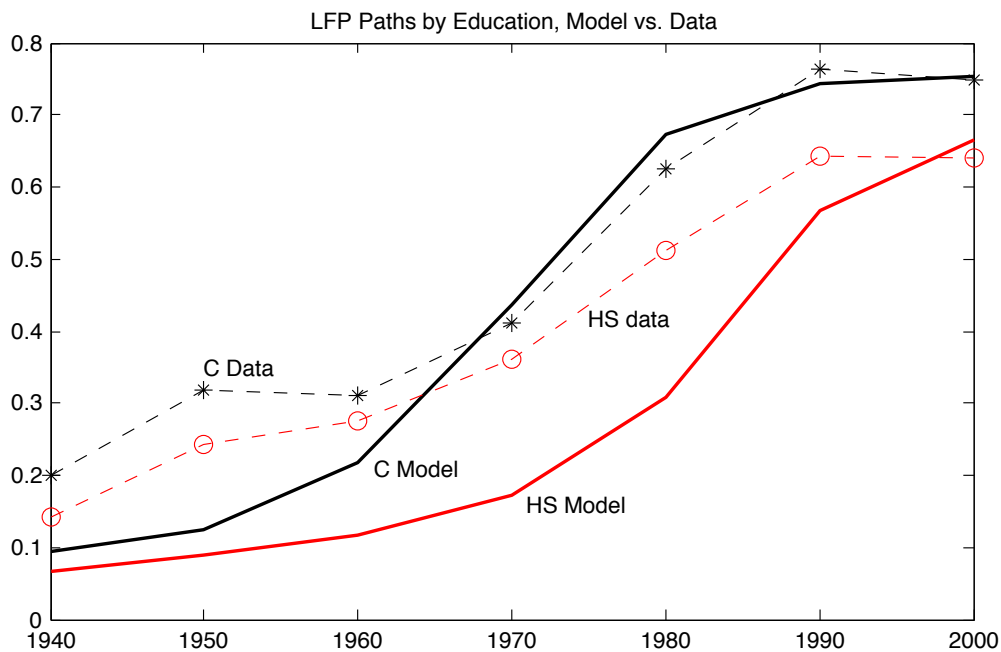


Figure 9: Predictions of the calibrated learning model for college (C) and high school (HS) types. See text for definitions of education categories. Bold = model, dashed = data.




Figure 10: Data for Great Britain is from Table 1 in Costa (2000). Data for France is compiled from Tables 1 and 2 in Reboud (1985). See the online Appendix for details on data.

A Online Appendix – For Online Publication Only

A.1 Data

Earnings

For earnings data prior to 1940, I rely on numbers provided in Goldin (1990) who 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 earnings for men and women.⁶⁰ As there is no data for earnings in 1880 and 1910, these points are constructed using a cubic approximation with the data from 1890 -1930 (inclusive).

To construct the earnings sample from 1940 onwards we used the 1% IPUMS samples of the U.S. Census for yearly earnings (incwage) and calculate the median earnings of white 25-44 years 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.⁶¹ As is commonly done, observations that report weekly earnings less than a cutoff are excluded.⁶² Prior to 1980, individuals report earnings from the previous year, weeks worked last year, and hours worked last week. From 1980 onwards, individuals are asked to report the "usual hours worked in a week last year." Hence for these years we require that people answer 35 or more hours to that question and we drop the restriction on hours worked last week. In 1960 and 1970, the weeks and hours worked information was reported in intervals. We take the midpoint of each interval for those years.

Sample weights (PERWT) were used as required in 1940, 1990, 2000. In 1950 sample line weights were used since earnings and weeks worked are sample line questions. The 1960-1980 samples are designed to be nationally representative without weights.

For the education categories, college is defined as having at least one years of schooling after high school. High School is defined as having at most a high school degree or no more than 13 years of education.

Figure 11 shows the evolution of female and male median earnings over the 120 year period 1880-2000 (with earnings expressed in 1967 dollars). In order to compare data sets, the figure plots both the numbers obtained from the calculations above as of 1940 (they are shown in (red) dots) as well as Goldin's numbers (which continue to 1980 and are shown in (blue) x's). The only significant difference is with male earnings in 1950 which are higher

⁶⁰See Goldin (1990) pages 64-65 and 129 for greater detail about the earnings construction for various years. I use the data for white men and women.

⁶¹The sample is limited to full-time year-round workers because hourly wages are not reported. The sample could have been restricted to include only married men and women, but I chose not to do this in order to be consistent with the data from the earlier time period.

⁶²The latter is calculated as half the nominal minimum wage times 35 hours a week and nominal weekly wages are calculated by dividing total wage and salary income last year by weeks worked last year. See, 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, I use the same cutoff rule as in Goldin and Margo (1992) as a way to eliminate unreasonably low wages. Note that by calculating median earnings, I do not have to concern myself with top-coding in the Census.

for Goldin.⁶³

Married Women's LFP: US and International

For the LFP numbers we used 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. Since the individual census data is missing for this 1890, we use the midpoint between 1880 and 1900. We restricted our sample to married white women (with spouse present), between the ages of 25-44, born in the US, in non-agricultural occupations and living in non-farm, non-institutional quarters.

For the education categories, college is defined as having at least one year of schooling after high school. High School is defined as having at most a high school degree or no more than 13 years of education.

The data for Great Britain (except 2010) is from Table 1 in Costa (2000) for married women between the ages either from 15-64 or 16-59, depending on the year. She obtains this data from various sources (see her footnote in Table 1). For 2010, the data is from the Office of National Statistics website (www.statistics.gov.uk).

For France the data for 1921-1975 is from Table 1 in Reboud (1985) for married women, in non-agricultural sectors, between the ages of 15-69. For 1978-1981, the data is from Table 2 in Reboud (1985), for married women, ages 15-64. For 1990-2008, the data is from the National Institute for Statistics and Economic Studies website (www.insee.fr), for married women, ages 15-64.

Elasticities

I take the elasticity estimates from Blau and Kahn (2006) who use the March CPS 1989-1991 and 1999-2001 to estimate married women's own-wage and husband's-wage elasticities along the extensive margin. They 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, along with the variables used to impute wages, both including and excluding education. The sample is restricted to married women 25-54 years old (with spouses in the same age range). I use the results obtained from the basic probit specification, which does not control for education, as this way the elasticity measure does not control for a measure of permanent income. This is preferable since an elasticity with respect to some measure of lifetime earnings is more appropriate for the model. I also chose the specification without children as a control variable as it is endogenous. Using the elasticities estimated from a specification with education controls does not affect the results as the elasticities are very similar (0.28 and

⁶³Goldin's 1950 number is from the Current Population Reports, series P-60 number 41 (January 1962). It is for all men over 14 which may explain the discrepancy since our census figure leaves out men older than 44 who would, on average, have higher earnings.

-0.12 for 2000 and -.15 in 1990). For the year 2000, Blau and Kahn estimate an own-wage elasticity of 0.30 and the cross-elasticity (husband's wage) of -0.13. The cross elasticity in 1990 is -0.14.

Poll Data

I use the Gallup Poll data for years 1938, 1945 and 1970. From 1972 to 1998, I use data from the General Social Survey (GSS). The exact question varied somewhat over time. It was "Do you approve or disapprove of a married woman holding a job in business or industry if her husband is capable to support her?" in 1945; and "Do you approve or disapprove of a married woman earning money in business or industry if she has a husband capable of supporting her?" for the remaining years. The possible answers included "Yes", "No" or "No Opinion" / "Don't know", depending on the year.

A.2 Calibration of the learning model

The Model with No Learning

Note that the wage elasticity ε (own, f , or cross, h) is given by:

$$\varepsilon_k = g(l^*) \frac{\partial l^*}{\partial w_k} \frac{w_k}{L} \quad (14)$$

$k = f, h$. Taking the ratio of the two elasticities and manipulating the expression yields a closed-form expression for γ , from which one can obtain a parameter value by using the earnings and elasticity numbers in 2000, i.e.,

$$\gamma = \frac{\log\left(1 - \frac{w_f \varepsilon_h}{w_h \varepsilon_f}\right)}{\log\left(1 + \frac{w_f}{w_h}\right)} = 0.503 \quad (15)$$

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

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

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

$$\sigma_l = \frac{A}{\exp\left(\frac{\Phi^{-1}(L)^2}{2}\right)} = 2.29 \quad (16)$$

where $A = \frac{w_f(w_f + w_h)^{-\gamma}}{\sqrt{2\pi\varepsilon_f L}}$. One can then solve for β directly from the definition of l^* , yielding $\beta = 0.321$.

This basic inability of the model absent learning to match the historical data is robust to a wide range of values for the elasticities (I explored with values ranging from twice to half of those in Blau and Kahn). It is also robust to alternative specifications of the share

of consumption that a woman obtains from her husband's earnings. In particular, one can modify the model so that the wife obtains only a share $0 < \alpha \leq 1$ of her husband's earnings as joint consumption. The results obtained from recalibrating the model using values of α that vary from 0.1 to 1 is shown in Figure 12. As is clear from the figure, this modification does little to remedy the basic problem. Furthermore, introducing any sensible time variation in this share would not help matters as it would require women to have obtained a much larger share of husband's earnings in the past in order to explain the much lower participation rates 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., greater bargaining power and hence a higher share of male earnings than in the past.⁶⁴

The failure of the model without learning is also robust to the exact choice of earnings series. For example, one might argue that, over time, the average hours worked by women has changed and this intensive margin is not incorporated into the model. In order to more fully account for this margin, rather than use the median earnings of full-time women, I constructed a series of the median annual earnings for all working women from 1940 to 2000. The sample consisted of 25-44 year old women who were born in the U.S., not living in group quarters, and working in a non-farm occupation. The adjustment to earnings was sizeable, ranging from 18% to 30% lower depending on the decade. This resulted in different parameter values ($\gamma = 0.49$, $\beta = .25$, $\sigma_l = 2.01$) but the predicted path of LFP generated was similar to the one obtained with the original series and hence still wildly overpredicted LFP.

The Model with Learning

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

$$\frac{\varepsilon_{w_f}}{\varepsilon_{w_h}} = \frac{\frac{\partial l}{\partial w_f} w_f}{\frac{\partial l}{\partial w_h} w_h}$$

Noting further that $\frac{\partial l}{\partial w_k} = \frac{\partial l^*}{\partial w_k}$, this implies that by performing the same manipulations as in the previous subsection one obtains (15), and thus the same value of γ as in the earnings only model, i.e., $\gamma = 0.503$.

In order to calculate a daughter's conditional probability of working (as a function of her mother's work behavior), one needs to specify, in addition to how private signals are inherited, how mothers and daughters are correlated in their l_j types. As a benchmark, I assume that the correlation is zero, i.e., the l type is a random draw from the normal distribution $G(\cdot)$ that is *iid* across generations.⁶⁵ Signals, on the other hand, are perfectly

⁶⁴Note that, in any case, to obtain the very low LFP numbers in 1880 would require women to fully share husband's earnings in that decade and to obtain a share of only 0.0001 of husband's earnings in the year 2000.

⁶⁵Thus, this model yields a positive correlation between a mother and her daughter's work "attitudes" ($E_{it}\beta + l_i$ and $E_{i',t+1}\beta + l_{i'}$ where i indexes the mother and i' the daughter).

inherited. Thus, given a signal s we can define l_s as the l type that is just indifferent between working and not at that signal value (i.e., $s_{l_s}^* = s$). Hence, the probability that a woman with signal s works is $G(l_s)$, i.e., it is the probability that her l type is smaller than l_s . Rearranging the expression for s_l^* in (7), we obtain

$$l_{st} = \frac{\underline{l}_t + \bar{l}_t \exp\left(\lambda_t - \left(\frac{\beta_H - \beta_L}{\sigma_\varepsilon^2}\right)(s - \bar{\beta})\right)}{1 + \exp\left(\lambda_t - \left(\frac{\beta_H - \beta_L}{\sigma_\varepsilon^2}\right)(s - \bar{\beta})\right)} \quad (17)$$

And, using Bayes rule and $\beta^* = \beta_L$, we can calculate the probability that a daughter works given that her mother worked as:

$$\begin{aligned} \Pr(DW_t | MW_{t-2}) &= \frac{\Pr(DW_t \text{ and } MW_{t-2})}{P(MW_{t-2})} \\ &= \frac{\int_{-\infty}^{\infty} \Pr(DW_t \text{ and } MW_{t-2} | s) f(s - \beta_L) ds}{L_{t-2}(\beta_L)} \\ &= \frac{\int_{-\infty}^{\infty} G(l_{st}) G(l_{s,t-2}) f(s - \beta_L) ds}{L_{t-2}(\beta_L)} \end{aligned} \quad (18)$$

where DW and MW stand for daughter works and mother worked, respectively. I use the predicted LFP from two periods earlier to calculate the probability that mothers worked (hence the $t - 2$ in expressions such as $G(l_{s,t-2})$). Note that in (18), the probability that both mother and daughter worked, $\Pr(DW_t \text{ and } MW_{t-2} | s)$, is multiplied by $f(s - \beta_L)$ as this is the proportion of daughters (or mothers) who have a private signal s in any time period.

A similar calculation to the one above yields

$$\Pr(DW_t | MNW_{t-2}) = \frac{\int_{-\infty}^{\infty} G(l_{st})(1 - G(l_{s,t-2})) f(s - \beta_L) ds}{1 - L_{t-2}(\beta_L)} \quad (19)$$

where MNW denotes a mother who did not work. The work risk ratio is thus given by

$$R_t = \frac{\Pr(DW_t | MW_{t-2})}{\Pr(DW_t | MNW_{t-2})} \quad (20)$$

In order to estimate $\lambda_0, \sigma_\varepsilon, \sigma_\eta, \beta_H, \beta_L$, and σ_l I minimized the sum of the squared errors between the predicted and actual values of our calibration targets (see table 1). All statistics were weighted equally.

The simplex algorithm was used to search for an optimal set of parameters. Multiple starting values throughout the parameter space were tried (specifically over 2,000 different starting values with λ_0 ranging between $[-10, -.01]$, σ_ε in $[0.1, 5]$, σ_η in $[0.01, 2]$, σ_l between $[0.5, 4]$, β_L in $[-.01, 1]$, and β_H to be between $[1, 10]$ 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.

300 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 8. 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 earnings or male earnings 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 was within $\pm .05$ of the true LFP that year). These elasticities were calculated individually for all histories meeting this criterion and were then averaged.

In order to approximate the integrals that are needed to compute $\Pr(DW_t|MW_{t-2})$ and $\Pr(DW_t|MNW_{t-2})$, 400 discrete signals from $\beta_L - 4\sigma_\varepsilon$ to $\beta_L + 4\sigma_\varepsilon$ were used.

Elasticity Paths

The path of elasticities predicted by the models is very different. See figure 13 below.

An Alternative Decomposition

It should be noted that there is not a unique way to decompose LFP in order to measure the quantitative importance of wages and beliefs. One could alternatively eliminate the $L(\bar{\lambda}(w_0), \bar{w})$ curve and replace it with the LFP path that would result if beliefs followed the path obtained from the historical earnings series, \bar{w} , but wages were kept constant at their 1880 levels. This curve is shown in Figure 14 as $L(\bar{\lambda}(\bar{w}), w_0)$. The dynamic effect of wages is now given by the difference between $L(\bar{\lambda}(\bar{w}), w_0)$ and $L(\bar{\lambda}(w_0), w_0)$. These paths are obtained using the same constant 1880 earnings, but in the first trajectory beliefs evolve as they would with the historical earnings profile, whereas in the second beliefs follow the path they would have taken had wages not changed over time. The static effect of earnings is now measured as the difference between $L(\bar{\lambda}(\bar{w}), w_0)$ and $L(\bar{\lambda}(\bar{w}), \bar{w})$, as beliefs evolve the same way for both series whereas earnings follow different paths.

The Welfare Costs of Imperfect Information

Another interesting exercise is to quantify the welfare costs of imperfect information. While this is not a policy-relevant calculation in that the government, for example, is not assumed to be better informed than any individual, it gives an idea of how costly mistaken beliefs were and how this cost evolved over time.

To quantify the losses from imperfect information we start by noting that if women were given the true value of β , only those individuals of type $l_j \in (l_t, \bar{l}_t)$ who as a result of their private information did not work in time t would change their decisions (and thus their

utility).⁶⁶ All women with $l_j \leq \underline{l}_t$ worked, and all those with $l_j \geq \bar{l}_t$ would choose not to work even if they knew the truth.⁶⁷

Thus, to quantify the loss of welfare, we can calculate at each moment in time, for each $l_j \in (\underline{l}_t, \bar{l}_t)$ type, the amount of consumption, z_{jt} , that a woman would have to be given in order to make her as well off as she would have been had she worked:

$$\frac{(w_{ht} + z_{jt})^{1-\gamma}}{1-\gamma} = \frac{(w_{ht} + w_{ft})^{1-\gamma}}{1-\gamma} - \beta_L - l_j \quad (21)$$

To interpret equation (21), note that the right-hand side is the utility enjoyed by a woman of type l_j who works. Thus, the left-hand side solves for the consumption equivalent of that utility.

The proportion of a given l_j type who made the wrong decision is given by those whose private signals lay above s_{jt}^* (as expressed in equation (7)), i.e., a fraction $1 - F(s_{jt}^* - \beta_L; \sigma_\epsilon)$. Integrating over the l_j types yields the aggregate welfare loss for these women at time t :

$$Z_t \equiv \int_{\underline{l}}^{\bar{l}} z_{jt}(1 - F(s_{jt}^* - \beta_L; \sigma_\epsilon))g(l_j)dl$$

In order to contrast this with the welfare enjoyed by working women, we need to translate the utility of an l_j type who worked into consumption units. This is given by finding x_{jt} such that:

$$\frac{(w_{ht} + w_{ft} + x_{jt})^{1-\gamma}}{1-\gamma} = \frac{(w_{ht} + w_{ft})^{1-\gamma}}{1-\gamma} - \beta_L - l_j \quad (22)$$

The proportion of a given l_j women who made the correct decision is given by $F(s_{jt}^* - \beta_L; \sigma_\epsilon)$ if $l_j \in (\underline{l}_t, \bar{l}_t)$ and equals one if $l_j \leq \underline{l}_t$. Thus, the aggregate welfare of working women, expressed in consumption units, is given by:

$$W_t \equiv \int_{\underline{l}}^{\bar{l}} (w_{ht} + w_{ft} + x_{jt})F(s_{jt}^* - \beta_L; \sigma_\epsilon)g(l_j)dl + \int_{-\infty}^{\underline{l}} (w_{ht} + w_{ft} + x_{jt})g(l_j)dl \quad (23)$$

Lastly, the total welfare of non-working women expressed in units of consumption, is given by:

$$N_t \equiv (1 - L_t) w_{ht}$$

Thus, the welfare lost as a result of imperfect information as a proportion of married women's welfare at time t , translated into consumption units, is given by \hat{c}_t :

$$\hat{c}_t \equiv \frac{Z_t}{W_t + N_t} \quad (24)$$

Note that \hat{c}_t gives the proportion of married women's average consumption lost as a result of imperfect information, when all utility is expressed in consumption units.

⁶⁶The top (red) line in Figure 7 shows what the evolution of female LFP would have been with full information.

⁶⁷Note that $\underline{l}_t, \bar{l}_t$ have time subscripts since their values depend on wages which are changing over time.

Figure 15 shows the evolution of \hat{c}_t over time. It starts out very high at some 39.29% of average consumption, decreases slowly until 1950, and then decreases dramatically to 0.19% by the year 2000. The very high numbers at the beginning are the consequence of the fact that the calibrated model implies that 63% rather than 2% of married women would have been working in 1880 had they possessed full information. A model in which the real cost of working evolved over time so that it was higher in 1880 than 2000 would imply smaller numbers as would a model in which the cost of working increased with the number of children.⁶⁸ Thus, these numbers should be taken as an upper bound to the costs of imperfect information.

Additional References for Online Appendix

References

- [1] 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.
- [2] Haines, Michael, "Fertility and Mortality in the United States". EH.Net Encyclopedia, edited by Robert Whaples. March 19, 2008. URL <http://eh.net/encyclopedia/article/haines.demography>
- [3] Katz, Lawrence and David Autor (1999), "Changes in the wage structure and earnings inequality," in O. Ashenfelter and D. Card (ed.), *Handbook of Labor Economics*, edition 1, vol. 3, chap. 26, 1463-1555, Elsevier.

⁶⁸The total fertility rate of white women in 1880, for example, is estimated to have been 4.24 as compared to 2.05 in 2000 (Haines (2008)). On the other hand, fertility is an endogenous variable and women might have chosen to have fewer children had they faced a higher known opportunity cost.

B Appendix Figures:




Figure 11: Crosses (blue) represent the yearly median earnings data from Goldin (1990), Table 5.1. Dots represent our calculations 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. All earnings are expressed in 1967 \$. See text for more detail.

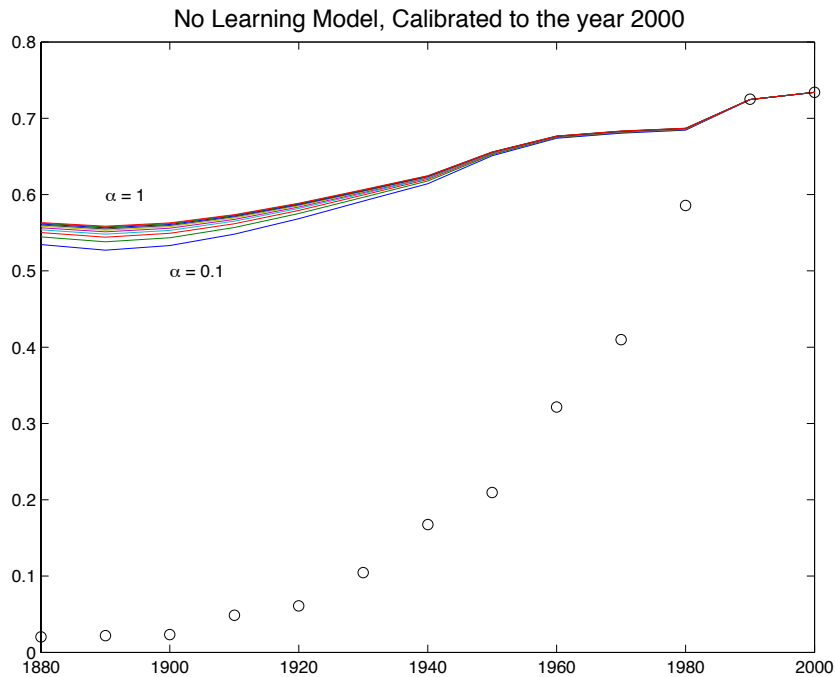


Figure 12: Parameters: $\gamma = 0.503$, $\beta = 0.321$, and $\sigma_L = 2.293$. α is the fraction of husband's earnings that enters a wife's utility via consumption.

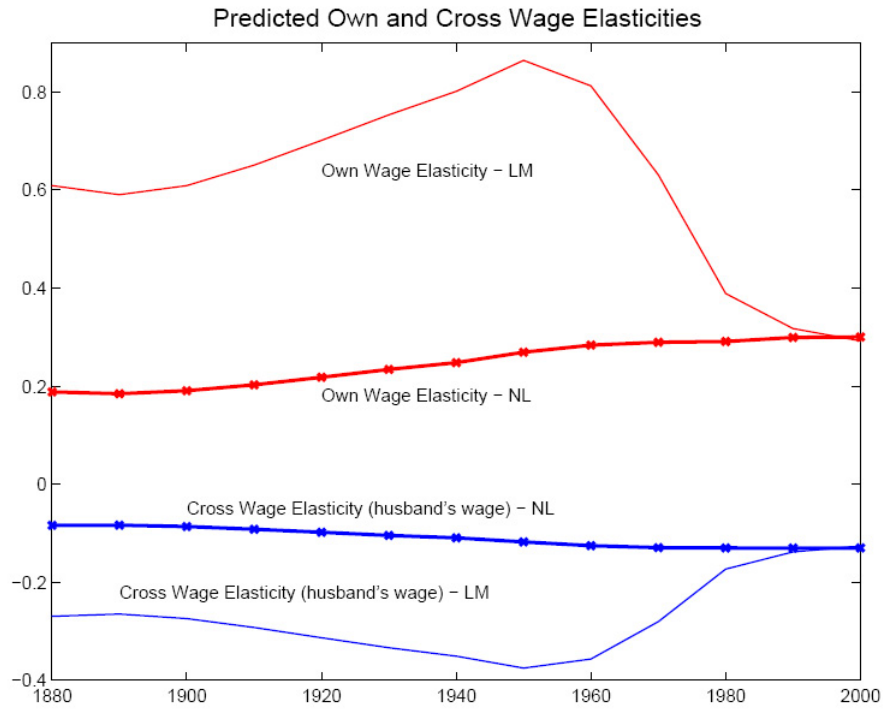


Figure 13: Parameter values from calibrated model. See the online Appendix for a description of how the elasticities were calculated. LM = Learning Model, NL = No Learning Model.

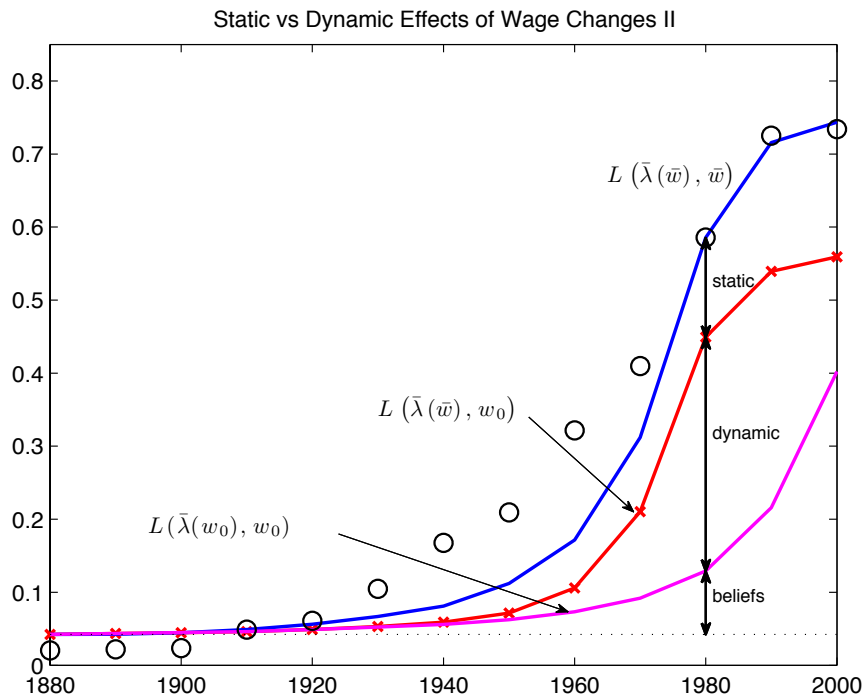


Figure 14: Alternative decomposition of LFP.

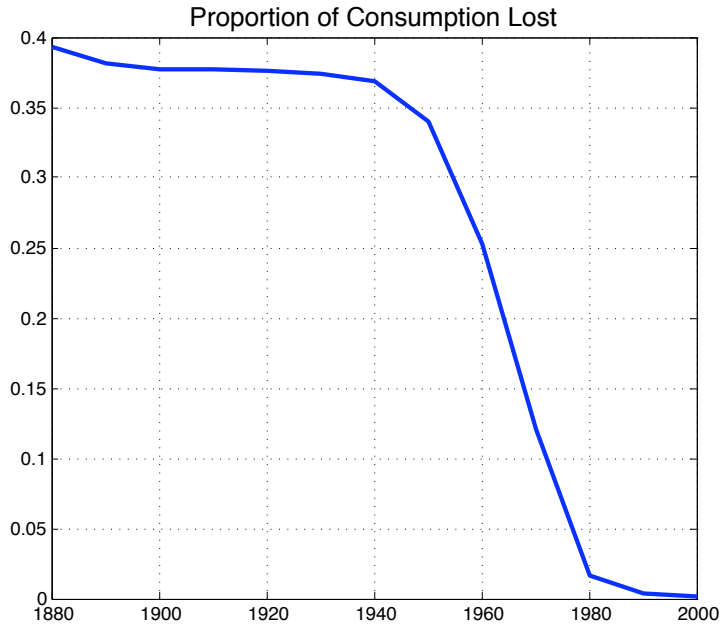


Figure 15: Percentage of welfare loss (in consumption terms) due to imperfect information about disutility from work. See text for details.

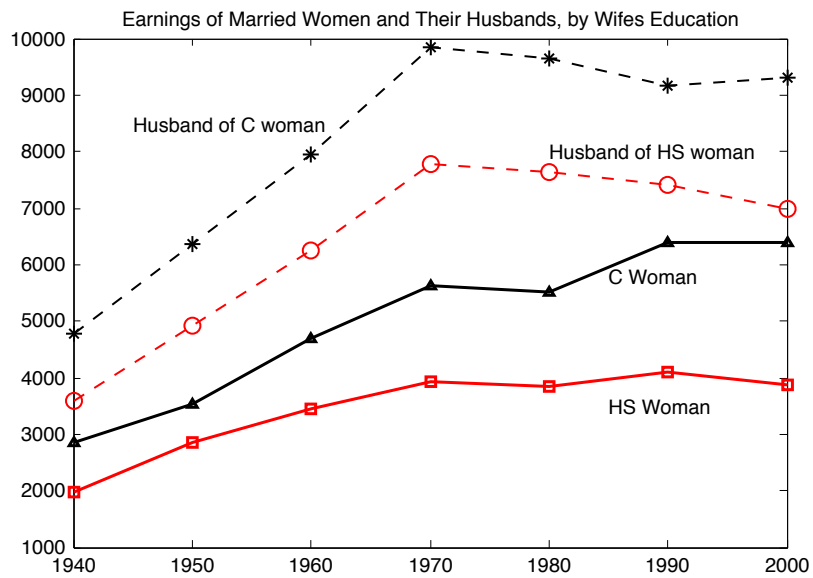


Figure 16: Median real earnings of married women by education type, college(C) and high school (HS), and median real earnings of each education type's respective husband. See text for definitions.