

Federal Reserve Bank of Minneapolis
Research Department Staff Report 386

February 2007

Nature or Nurture? Learning and Female Labor Force Dynamics

Alessandra Fogli*

Federal Reserve Bank of Minneapolis,
New York University,
and CEPR

Laura Veldkamp*

New York University

ABSTRACT

In the last century, the evolution of female labor force participation has been S-shaped: It rose slowly at first, then quickly, and has leveled off recently. Central to this dramatic rise has been entry of women with young children. We argue that this S-shaped dynamic came from generations of women learning about the relative importance of nature (endowed ability) and nurture (time spent child-rearing) for children's outcomes. Each generation updates their parents' beliefs by observing the children of employed women. When few women participate in the labor force, most observations are uninformative and participation rises slowly. As information accumulates and the effects of labor force participation become less uncertain, more women participate, learning accelerates and labor force participation rises faster. As beliefs converge to the truth, participation flattens out. Survey data, wage data and participation data support our mechanism and distinguish it from alternative explanations.

*Comments welcome. We thank Stefania Marcassa for excellent research assistance. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

In the last century, one of the most dramatic economic changes has been the rise in female labor force participation. Central to this rise has been the entry of married women with young children. While only 6% of the mothers of preschool age children worked in the labor market in 1940, 60% of these mothers are employed today. This phenomenon took place at the same time of a radical change in reported beliefs about the effects of mothers' employment on children's outcomes. While almost 70% of women reported in the early 1970s that a preschool child is likely to suffer if her mother works, this percentage is down to 42% today. A large empirical literature has argued that such attitudes and beliefs, broadly defined as culture, are an important determinant of labor force participation.¹ However, less is known about where these beliefs come from, how they evolve and how they interact with participation rates.

We argue that beliefs are formed and evolve over time as women learn about the relative importance of nature and nurture in determining children's outcomes. A crucial factor in a woman's choice to work is the effect of her employment on her children. The extent to which labor force participation trades off with children's future utility is unknown. Women pass down beliefs about the importance of nurture to their children. Each generation updates those beliefs, using a set of observations on other children's outcomes. However, observations are only informative about the cost of labor force participation if women in the previous generation work. Initially, very few women participate in the labor market; information about the role of nurture diffuses slowly and beliefs are nearly constant. As information accumulates and the effects of labor force participation become less uncertain, more women participate, learning accelerates and labor force participation rises more quickly. As beliefs converge to the truth, learning slows down and participation flattens out. This interaction between beliefs and employment generates a simultaneous S-shaped evolution of women's beliefs and of labor force participation that mimics the data from the past century.

The literature that has studied the dramatic rise of women's labor force participation over the last century has mostly focused on technological explanations, the less physical nature of most jobs

¹See Fernández and Fogli (2005), Fernández, Fogli, and Olivetti (2002), Fortin (2005), Alesina and Giuliano (2007). A related literature shows beliefs or culture affect a variety of other economic outcomes: Giuliano (2007), Guiso, Sapienza, and Zingales (2004) and Barro and McCleary (2006). For theoretical work, see Bisin and Verdier (2000).

and increases in women's wages.² This paper shares the idea that the main difference between men and women in their work behavior is linked to the different amount of time devoted to household work, in particular to children's nurturing. However, we argue that focussing on the availability of technology or actual labor conditions misses the fact that women's *beliefs* are an important determinant of their participation decision. A new technology may enable a woman to complete housework and pursue a career. But, even with such technology, a woman may give up a career if she believes that working harms her family.

As a theoretical contribution, the model fills a gap between the literature on S-shaped learning dynamics and that on endogenous information. The S-shaped learning dynamic is similar to Amador and Weill's (2006) model where agents, arranged on a lattice, learn what their neighbors know. It also is reminiscent of the S-shaped value of information in Keppo, Moscarini, and Smith (2005) and Radner and Stiglitz (1984), where agents choose how much information to acquire, before taking a risky action. But rather than endowing everyone with signals or an information production technology, our information comes from observing labor market outcomes. The fact that information and actions are mutually interdependent delivers additional testable predictions. Not only should beliefs predict actions, as in exogenous information models, but actions should also predict changes in beliefs. This idea that information is a by-product of economic activity appears in Van Nieuwerburgh and Veldkamp (2006). But there was no S-shaped learning dynamic. Because the variable being learned about was constantly changing, beliefs did not converge. Combining these two sets of ideas generates predictions that are unique to this model and supported by the data.

Over the last century, female labor force participation and beliefs about the welfare of children with working mothers have risen concurrently, following an S-shaped time path. Section 1 establishes this fact and presents evidence on the two key ingredients of our story: The change in the labor supply of married women with small children has been a major factor in the rise of aggregate female labor force; and there is great uncertainty about the effect of maternal employment on

²See Greenwood, Seshadri, and Yorukoglu (2001), Goldin and Katz (2002), and Albanesi and Olivetti (2006) on technologies and Goldin (1990), Jones, Manuelli, and McGrattan (2003) on nature of jobs.

children.

Section 2 develops an overlapping generation model where a woman faces a trade-off between enjoying more consumption if she enters the labor force and earns a wage, and staying out of the labor force to nurture her children and ensure they have high future utility. Women use Bayesian updating to learn about the importance of maternal employment in determining children's outcomes. The results (section 3) show how two competing forces speed up, then slow down learning, creating the S-shaped participation dynamic. Learning is slow at first because few women work, making data about labor force participation scarce. As women learn, participation increases and speeds up learning. The offsetting force is that as beliefs become more informed, new information affects them less. We also show that, as learning converges, labor force participation becomes less responsive to wages, a feature of recent data (Blau and Kahn 2005).

To evaluate the quantitative predictions of the model, section 4 uses moments of the labor force participation data and the distribution of wages from the census to calibrate and simulate a dynamic learning model. The simulated time path of labor force participation looks strikingly similar to the data. Learning speeds up, then slows down, at a rate only slightly slower than the observed rate of learning. Furthermore, the pattern of the mean belief in the population matches the change in the survey response over time.

Section 5 evaluates other testable implications of our learning model for belief dispersion, wage elasticity, history dependence of participation decisions, and the change in survey responses over the life-cycle. We use these predictions to distinguish our learning theory from other competing explanations (section 6). Section 7 concludes.

1 Motivating Evidence: Beliefs, Labor Force Participation, and its Toll on Children

Before proceeding with our theory, we document the S-shaped learning dynamic we are building the theory to explain. We also establish support for two critical assumptions of our model: First, the key change in labor force participation has been among women with small children. Second,

the effects of maternal employment on these children is uncertain.

1.1 The S-shape in beliefs and participation

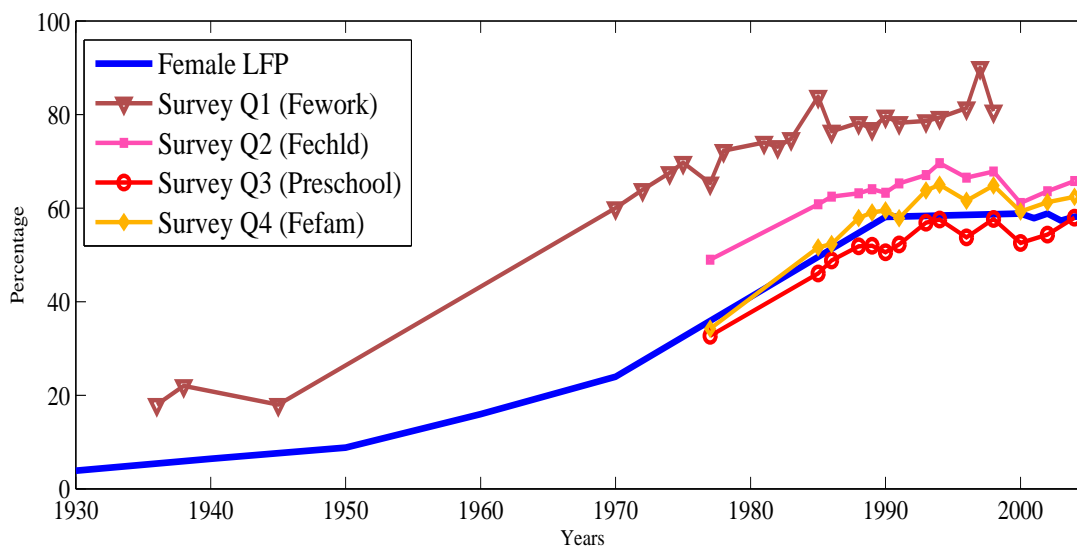


Figure 1: LABOR FORCE PARTICIPATION AND AVERAGE SURVEY RESPONSES.

Labor force participation (LFP) is the S-shaped line with squares. Survey questions are about the effect of women's labor force participation on children. A higher level indicates a more favorable attitude toward participation.

One of the reasons to think that the main force is changes in beliefs about the effect of market work on children is survey evidence revealing that such beliefs have changed. Figure 1 plots labor force participation, along with the answers to four questions about the effect of women's labor force participation on children, from 1930-2005. The survey questions are about whether a married woman should work (fework), whether she compromises her relationship with her children (fechld), whether pre-school age children, in particular, suffer (preschool), and whether her family suffers (fefam). (Data details in appendix B.) The fact that the survey responses are increasing in their average level indicates that, over time, people believe that women's labor force participation is less harmful to their families. What is striking about this graph is that the labor force participation tracks the survey responses so closely and that both display an S shape.

1.2 Employment increase of women with young children

The bulk of the increase in labor force participation came from married women with children. Single mothers and women without children started with much higher participation rates, which increased, but not dramatically. In contrast, the participation of married women with children increased by more than three-fold. Even among this subgroup, the largest increase came from women with children under 5 years of age. The participation rate for this group was just under 4% in 1930 and rose more than 10-fold to 59% in 2004. For this reason, we focus our theory on the concerns and trade-offs that women with children face. The rest of our data analysis will pertain to the group of women whose labor force participation showed the most change over the last century: married women, with children under 5 years of age. Details of the data underlying these facts and figure 1 can be found in appendix B.




Figure 2: LABOR FORCE PARTICIPATION AMONG SUB-GROUPS OF WOMEN.

1.3 Psychological evidence on the true value of nurture

Our theory is based on the premise that the effect of mothers' employment on children is uncertain. This is a reasonable assumption because even academic psychologists have not reached a consensus. Harvey (1999) summarizes studies that started in the early 60s and flourished in the 1980s when the

children of the women interviewed in the National Longitudinal Survey of Youth (NLSY) reached adulthood. Six different published studies have used the NLSY to examine the longitudinal effects of early maternal employment on children’s development, controlling for various family background variables. Harvey compares these studies and concludes that “their results have been surprisingly mixed considering they used the same data set. . . . Rather than shedding light on the issue, these conflicting results have added further confusion.” Her analysis of the data indicates that working more hours is associated with slightly lower cognitive development through age 9 and slightly lower academic achievement scores before age 7.

A more recent study by Hill, Waldfogel, Brooks-Gunn, and Han (2005) finds small but significant negative effects of maternal employment on children’s cognitive outcomes for full-time employment in the first year post-birth as compared with employment postponed until after the 1st year. Bernal (2006) concurs: Having a full-time working mother who uses child care during one of the first five years after childbirth is associated with a 1.8 percent reduction in the child’s test score. Belsky (1988) also finds that maternal employment is detrimental: Infants who were in non maternal care for more than 20 hours per week were at elevated risk for being insecurely attached at age 1 and were more aggressive and disobedient between ages 3 and 8. In sum, while there is little evidence of negative effects of maternal employment when children are older, there is substantial debate about the size of the effects of maternal employment in early childhood.

2 The Model

We assume a discrete infinite horizon, $t = 1, 2, \dots$, and we consider an overlapping generation economy made up of a large finite number of agents living for two periods. Each agent is nurtured in the first period and consumes and has one child in the second period of her life. Preferences of an individual in family i born at time $t - 1$ depend on their consumption in period t , c_{it} and the potential wage of their child $w_{i,t+1}$.

$$U = \frac{c_{it}^{1-\gamma}}{1-\gamma} + \beta \frac{w_{i,t+1}^{1-\gamma}}{1-\gamma} \quad \gamma > 1 \quad (1)$$

The budget constraint of the individual from family i born at time $t - 1$ is

$$c_{it} = n_{it}w_{it} + \omega_{it} \tag{2}$$

where ω_{it} is an endowment which could represent a spouse's income and $n_{it} \in \{0, 1\}$ is the discrete labor force participation choice. If the agent works in the labor force, $n_{it} = 1$.

The key feature of the model is that an individual's earning potential is determined by a combination of endowed ability and nurturing. Endowed ability is $a_{i,t} \sim N(\mu_a, \sigma_a^2)$. If a mother stays home with her child, the child's full natural ability is achieved. If the mother joins the labor force, some unknown amount θ of the child's ability will be lost. Wages depend exponentially on ability:

$$w_{i,t+1} = \exp(a_{i,t+1} - n_{i,t}\theta) \tag{3}$$

The constant θ determines the importance of nurture and is not known when making labor supply decisions. Agents learn about it over generations. Young agents inherit their prior beliefs about θ from their parents' beliefs. In the first generation, initial beliefs are $\theta_{i,0} \sim N(\mu_0, \sigma_0^2)$.

Updating beliefs At the end of each period, potential earnings are realized and agents observe both potential earnings and parental employment decisions for J of their peers, randomly and independently chosen from the population. The set of family indices for the outcomes observed by agent i is \mathbb{J}_i . Agents use this information to update their prior, according to Bayes' law.

Bayesian updating with J signals is equivalent to running a regression of children's potential wages on parents' labor choices and then forming a linear combination of the estimated weight on labor choices $\hat{\theta}$ and the prior belief μ_t .

At the end of each period t , the regression agents run to form their signal is

$$W - \mu_a = N\theta + \varepsilon_i$$

where W and N are $J \times 1$ vectors $\{\log w_{j,t+1}\}_{j \in \mathbb{J}_i}$ and $\{n_{i,t}\}_{j \in \mathbb{J}_i}$. Let $\bar{n}_{i,t}$ be the sum of the labor

decisions for the set of families that (i, t) observes: $\bar{n}_{i,t} = \sum_{j \in \mathbb{J}_i} n_{j,t}$. The resulting estimated coefficient $\hat{\theta}$ is normally distributed with mean $\hat{\mu}_{i,t} = \sum_{j \in \mathbb{J}_i} (\log w_{j,t+1} - \mu_a) n_{j,t} / \bar{n}_{i,t}$ and variance $\hat{\sigma}_{i,t}^2 = \sigma_a^2 / \bar{n}_{i,t}$.

Posterior beliefs about the value of nurturing are normally distributed $\theta \sim N(\mu_{i,t+1}, \sigma_{i,t+1}^2)$. The posterior mean is a linear combination of the estimated coefficient and the prior beliefs, where each component's weight is its relative precision:

$$\mu_{i,t+1} = \frac{\hat{\sigma}_{i,t}^2}{\sigma_{i,t}^2 + \hat{\sigma}_{i,t}^2} \mu_{i,t} + \frac{\sigma_{i,t}^2}{\sigma_{i,t}^2 + \hat{\sigma}_{i,t}^2} \hat{\mu}_{i,t} \quad (4)$$

The posterior precision (inverse of the variance) is the sum of the prior precision and the signal precision. Thus posterior variance is

$$\sigma_{i,t+1}^2 = (\sigma_{i,t}^{-2} + \hat{\sigma}_{i,t}^{-2})^{-1}. \quad (5)$$

Definition of equilibrium An equilibrium is a sequence of distributions that characterize beliefs about θ for each individual i in each generation t , work and consumption decisions and wages such that:

1. Taking beliefs and wages as given, consumption and labor decisions maximize expected utility (1) subject to the budget constraint (2).
2. Wages of agents born in period t are consistent with the labor choice of the parents, as in (3).
3. An agent i born at date t chooses consumption and labor at date $t + 1$. That optimization is conditioned on beliefs $\mu_{i,t+1}, \sigma_{i,t+1}$ that are equal to the posterior beliefs of the parent, born at $t - 1$. Those posterior beliefs come from the parent's prior beliefs $\mu_{i,t}$, which are updated using observed wage outcomes $\mathbb{J}_{i,t}$, according to Bayes' law (4).

3 Results

Substituting the budget constraint (2) and the law of motion for wages (3) into expected utility (1) produces the following optimization problem for agent i born at date $t - 1$:

$$\max_{n_{it} \in \{0,1\}} \frac{(n_{it}w_{it} + \omega_{it})^{1-\gamma}}{1-\gamma} + \beta E_{a_{i,t+1},\theta} \left[\frac{\exp((a_{i,t+1} - n_{i,t}\theta)(1-\gamma))}{1-\gamma} \right]. \quad (6)$$

Taking the expectation over the unknown ability a and the importance of nurture θ reveals that the optimal policy is to join the labor force when the following inequality is satisfied:

$$\begin{aligned} \mathcal{N}_{i,t} \equiv & \frac{(n_{it}w_{it} + \omega_{it})^{1-\gamma}}{1-\gamma} + \frac{\beta}{1-\gamma} \exp \left((\mu_a - n_{i,t}\mu_{i,t})(1-\gamma) + \frac{1}{2}(\sigma_a^2 + n_{i,t}^2\sigma_{i,t}^2)(1-\gamma)^2 \right) \\ & - \frac{(\omega_{it})^{1-\gamma}}{1-\gamma} - \frac{\beta}{1-\gamma} \exp \left(\mu_a(1-\gamma) + \frac{1}{2}\sigma_a^2(1-\gamma)^2 \right) > 0. \end{aligned} \quad (7)$$

3.1 The Role of Beliefs in Labor Force Participation

We begin by establishing some intuitive properties of the labor force participation decision rule. Women who think nurture is more important and those who are more uncertain about the importance of nurture are less likely to work.

Proposition 1 *A higher expected value of nurture reduces the probability that a woman will participate in the labor force.*

Proof: in appendix A.1. The logic of this result can be seen in condition 7. Increasing the expected value of nurture increases the expected utility of labor force participation: $\partial \mathcal{N}_{i,t} / \partial \mu_{i,t} = -\beta$. Since $-\beta < 0$, a higher $\mu_{i,t}$ reduces the utility gain from labor force participation and therefore reduces the probability that a woman will participate.

Proposition 2 *Greater uncertainty about the value of nurture reduces the probability that a woman will participate in the labor force.*

Proof: in appendix A.2. Again, the logic of the result is in the fact that the benefit to working is falling in uncertainty: $\partial \mathcal{N}_{i,t} / \partial \sigma_{i,t} = (1 - \gamma)\beta$. Since $\gamma > 1$, then $(1 - \gamma)\beta < 0$, meaning that more uncertainty reduces the probability a woman will participate.

3.2 S-shaped Learning Dynamics

The model generates an S-shaped pattern of labor force participation because beliefs about the cost of participation move slowly at first, then faster, and then slow down. There are two competing forces that generate this dynamic. One force is that as the number of women in the labor force increases, the average sized revision in beliefs increases. The second, competing force is that as beliefs converge to the truth, revisions in beliefs become smaller.

The information gleaned from observing others' labor market outcomes can be described as a signal with mean $\hat{\mu}_{i,t} = \sum_{j \in \mathbb{J}_i} (\log w_{j,t+1} - \mu_a) n_{j,t} / \bar{n}_{i,t}$ and variance $\hat{\sigma}_{i,t}^2 = \sigma_a^2 / \bar{n}_{i,t}$. Let ρ be the fraction of women who participate in the labor force. Then, the expected precision of this signal is $E[\hat{\sigma}_{i,t}^{-2}] = \rho N \sigma_a^2$.

A higher signal precision increases the expected size of squared changes in beliefs. This conditional variance of $t+1$ beliefs, conditional on period t beliefs is the difference between prior variance and posterior variance: $var(\mu_{i,t+1} | \mu_{i,t}) = \sigma_{i,t}^2 - \sigma_{t,t+1}^2$. Substituting in the formula for posterior variance (5), we find

$$var(\mu_{i,t+1} | \mu_{i,t}) = \sigma_{i,t}^2 - \frac{1}{\sigma_{i,t}^{-2} + \hat{\sigma}_{i,t}^{-2}}. \quad (8)$$

Since $\partial var / \partial \hat{\sigma}_{i,t}^{-2} > 0$, the expected size of revisions is increasing in the precision of the observed signals. For the average agent, this precision is increasing in the fraction of women who work. This is the first force as work that increases the rate of learning over time and causes the labor force participation rate to increase at an increasing rate, early in the century.

The second force is the convergence of beliefs to the truth. Over time, the variance of beliefs about θ declines. This effect can be seen in the posterior variance formula (5). If there is no information in any of the observed labor market outcomes, then $\sigma_{i,t+1}^2 = \sigma_{i,t}^2$. Given any new information, $\hat{\sigma}_{i,t}^{-2} > 0$, posterior variance is lower than in the previous period: $\sigma_{i,t+1}^2 < \sigma_{i,t}^2$. As σ^2

falls, the conditional variance of changes in beliefs falls as well: $\partial var/\partial \sigma_{i,t}^2 > 0$.

The increase in signal quality early on in the century was the dominant force, causing learning to speed up. This effect diminished later on because the convergence of beliefs to the truth reduces the effect of higher signal precision. This effect is about the cross-partial derivative: $\partial^2 var/\partial \hat{\sigma}_{i,t}^{-2} \partial \sigma_{i,t}^{-2} = -2/(\sigma_{i,t}^{-2} + \hat{\sigma}_{i,t}^{-2})^3 < 0$. The convergence of beliefs is the dominant force, slowing down learning later in the century both because higher prior precision makes belief revisions smaller and because higher prior precision reduces the effect of more informative signals.

3.3 The Effect of Wealth

An alternative explanation for the flattening out of women's labor force participation is that women participate less when their wealth level is higher. This force is present in the model.

Proposition 3 *Greater initial wealth $\omega_{i,t}$ reduces the probability that a woman will participate in the labor force.*

Proof in appendix A.3. The effect of wealth on the value of participating in the labor force is $\partial \mathcal{N}_{i,t}/\partial \omega_{i,t} = (w_{i,t} + \omega_{i,t})^{-\gamma} - (\omega_{i,t})^{-\gamma}$. This is negative because $(w_{i,t} + \omega_{i,t}) > (\omega_{i,t})$ holds with probability one. This means that more wealth reduces the utility gain from labor force participation.

3.4 Learning and Labor Supply Elasticity

The wage elasticity of the labor supply is the marginal change in the log probability that a woman participates due to a log change in the average wage. Since wages make working more attractive, this elasticity is always positive. However, it declines as women learn more about the effects of labor force participation on their children.

Proposition 4 *The wage elasticity of labor force participation is increasing in the level of uncertainty about the value of nurture: $\partial^2 \ln(\text{Prob}(\mathcal{N}_{i,t} > 0))/\partial \ln(\text{mean}(w_{i,t})) \partial \sigma_{i,t}^2 > 0$.*

Proof in appendix A.4. There are two pieces to this result, a direct effect on the benefit of working and an indirect effect on the expected utility cost to children. The direct effect is that

when uncertainty about the value of nurture is higher, fewer women work and changes in the wage generate larger percentage changes in the number of working women. When uncertainty declines, most women are already in the labor force and most have wages far above the cutoff wage for working or not working. The density of women near this cutoff is small. At this stage, proportional changes in the wage have a small effect on the number of women who work.

The indirect effect is that an increase in the average wage in the next generation makes mothers expect a higher income for their children. If their expected level of income is higher, the marginal utility cost of reducing children’s future realized income level by working is mitigated. This also causes more women to work when wages increase. When the value of nurture is very uncertain, the increase in the children’s expected utility from an increase in the expected wage is greater than if the value of nurture is well-known. This is a Jensen’s inequality effect that arises with CRRA utility. This second effect also pushes the elasticity of labor force participation down over time: As agents learn, the value of nurture becomes more certain. This diminishes the effect of increasing the average wage on labor force participation. Figure 3 illustrates this effect. When θ is known, utility is a concave function of ability, which determines the wage. When θ is uncertain, expected utility is an average of potential utility outcomes. Increasing the average wage moves expected utility by shifting the set of possible outcomes to the right. Because the realized utility curve (A) is less steep than the expected utility curve (B), the marginal utility effect on the uncertain agent is larger.

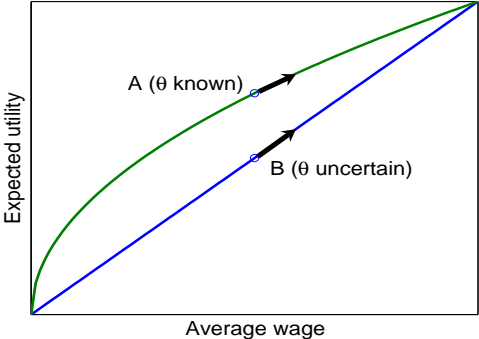


Figure 3: UNCERTAINTY AMPLIFIES THE MARGINAL UTILITY EFFECT OF A WAGE INCREASE. A is the utility of the agent who is certain about the value of nurture. B is the expected utility of the uncertain agent. The slope at A is smaller than the slope at B.

4 Simulation Results

We simulate a model with 2000 agents for 10 periods. We want to explain the rise in labor force participation of women with small children, so we focus on women that are 25-35 years old. Agents are born at age 15 with beliefs inherited from their parents. At the end of the first 10-year period outcomes are observed and agents update their beliefs. We assume that each individual observes the potential wage and the employment history of 2 other individuals of the same generation, that are randomly drawn from the distribution. Given the updated beliefs, at age 25 they enter adulthood, have a child and decide whether to participate in the labor market or not.

4.1 Calibration

We need to choose parameters for earnings of husbands and wives, initial beliefs about the value of nurture, a true value of nurture and a number of outcomes to observe each period. This is a total of 8 parameters, summarized in table 1.

outcomes observed	J	2	average family size
mean log endowment	μ_ω	-0.3	men's wage distribution
std log endowment	σ_ω	0.7	men's wage distribution
mean log ability	μ_a	-1.3	women's wage distribution
std log ability	σ_a	0.9	women's wage distribution
prior mean θ	μ_0	2.6	beginning LFP
prior std θ	σ_0	2.0	beginning LFP
true value of nurture	θ	0.2	2004 LFP
intertemporal substitution	γ	1.2	commonly used

Table 1: Parameter values for the simulated model and the features of the data each parameter matches.

Our strategy is to choose parameters of the earnings distributions in our model to match the empirical distribution in 1940, the earliest year for which we have the wage distribution data. (See appendix B.3 for more detail.) We set the previous period labor supply equal to 3 percent, that is the labor force participation of married women with children under age 5 in 1930 from Census data. Since we interpret women's endowment ω as being the earnings of their husbands,

we use a log-normal distribution, $\ln(\omega) \sim N(\mu_\omega, \sigma_\omega^2)$, which is frequently used to describe earnings. We normalize the average endowment (not in logs) to 1 and use σ_ω to match the dispersion of 1940 annual wage income of husbands with children under 5 at home. For women, the mean and standard deviation of ability, μ_a and σ_a , match the earnings gap between working women and their husbands and its log standard deviation to data in 1940. Initial beliefs μ_0 and σ_0 are the same across individuals and chosen to match women’s 1940 labor force participation. The true value of nurture θ causes the model to converge to the 2005 labor supply of our subset of women. Since not everybody is working in the data, and the only cost from working in the model is associated to the cost on the child, θ is small but positive (0.1). Interestingly, this value finds empirical support in recent micro literature. Finally, the elasticity parameter in preferences is a commonly used value.

4.2 S-shaped dynamics of participation and beliefs

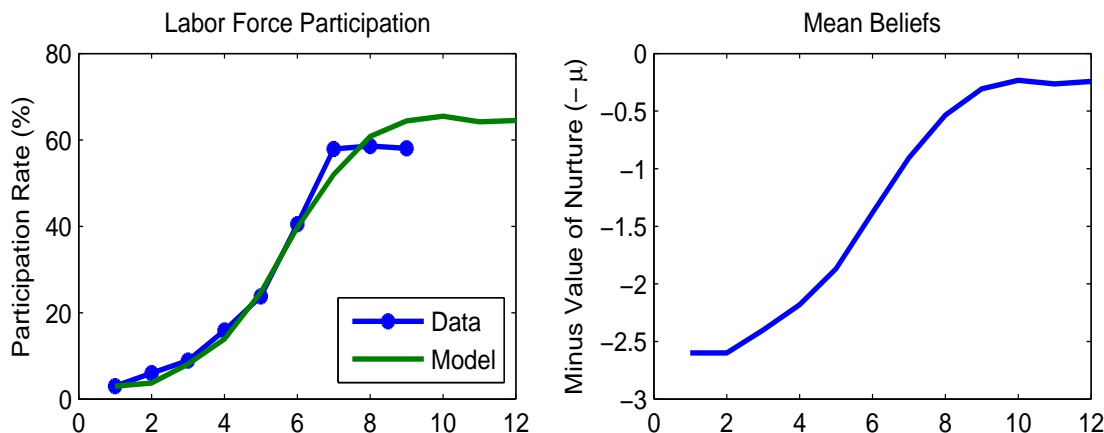


Figure 4: LABOR FORCE PARTICIPATION AND AVERAGE BELIEFS IN THE SIMULATED MODEL. LFP is the probability a woman participates in the labor force. The rise in mean beliefs shows that women believe labor force participation is less harmful to their children’s future outcomes over time.

Figure 4 illustrates the two main predictions of the model. The first (left panel) is the S-shaped rise in labor force participation in the simulated model. If we interpret 1 period as 10 years, then our model matches the realized labor force participation rates reasonably well. Using the model to project forward in time, it forecasts that there will be a 61% participation rate in 2010, after which the rate will change very little because learning has mostly converged.

The second prediction of the model is that beliefs about the value of nurture fall, in an S-shaped pattern. Figure 4 (right panel) plots the simulated value of $-\mu$, which rises. We reverse the sign of beliefs to make them comparable to our beliefs data. More positive answers to our survey questions means more support for a woman working, which implies a lower believed value of θ . The S-shaped increase is consistent with the survey evidence in figure 1.

Taken together, these two graphs illustrate the model’s main mechanism: When labor force participation is low, information about the cost of maternal employment on children is scarce, and beliefs change slowly. The slow change in beliefs means a slow rise in the participation rate. As that rate rises, more information is generated by working women whose children’s outcomes are observed. Beliefs and participation change at a faster rate. Once information becomes sufficiently abundant, beliefs and labor force participation converge, in unison, to their new steady-state values.

4.3 Decline in uncertainty

One of the predictions that distinguishes the learning model from other explanations is that learning entails a reduction in uncertainty. A common proxy for uncertainty is the dispersion in beliefs. The idea is that when there is lots of uncertainty, there is also lots of disagreement about what the right answer is.




Figure 5: UNCERTAINTY AND BELIEF DISPERSION IN THE SIMULATED MODEL. Belief dispersion is the cross-sectional standard deviation of the beliefs: $\text{std}(\mu_{i,t})$. Uncertainty is the average posterior standard deviation: $\text{mean}(\sigma_{i,t})$.

In the model, uncertainty declines monotonically (figure 5, right panel), but belief dispersion increases, then decreases (left panel). Belief dispersion starts out low because we assume that agents have common prior beliefs. Common beliefs, by definition, means no dispersion. As women begin working, some agents observe labor market outcomes, while others do not. Furthermore, the inferences that agents make from observing different labor market outcomes vary because the unobserved innate ability of the workers they observe is different. As new information arrives, beliefs diverge. Recall that what slows learning down at the end of the S-curve is that beliefs converge to the truth. Since the true importance of nurture is the same for all agents, their beliefs converge to each other and dispersion falls. The testable prediction here is that, in periods where the increase in labor force participation is slowing down, dispersion in beliefs should be falling. We test this prediction in the next section.

4.4 Wages and wealth of working women

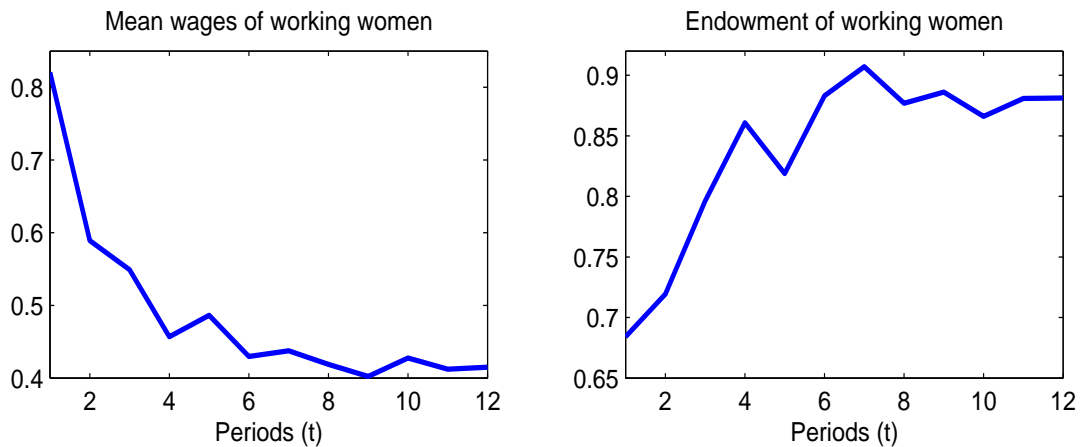


Figure 6: AVERAGE WAGE ($w_{i,t}$) AND ENDOWED WEALTH ($\omega_{i,t}$) FOR A WORKING WOMAN IN THE SIMULATED MODEL.

Figure 6 shows the mean endowment and the mean wage of the subset of women who choose to participate in the simulated model. We assumed that the unconditional distribution of endowments and abilities is constant. What is changing is the nurture inputs and the selection of women who work. In other words, this is primarily a selection effect.

In the data, women working at the beginning of the century are married to poorer husbands than the women that are in the labor force today. The same is true in the model: Early on, women perceive of maternal employment as costly and often choose to work because their husband is poor, making their marginal utility of income high. As the perceived cost of maternal employment falls, women with richer husbands join the labor force.

One counter-factual prediction of the model is that women's wages decline (figure 6, right panel). This is a selection effect. Early on, only very high-ability women work. As they learn that labor force participation is less detrimental to children, less-able women join the workforce, pushing wages down. This suggests that whatever caused women's wages to rise (more education, less discrimination, etc.) is an additional effect, outside our model. Obviously, changes in women's wages should also affect their willingness to participate. Future work will use the model and the wage data to disentangle what part of the increase in participation came from wages and what part came from learning. If we re-interpret the model as a decision between entering a high-wage, highly time-consuming occupation or a low-wage, less demanding job, the model could predict an increase in wages. Over time, women would perceive a lower and less uncertain cost of working longer hours and would choose careers that pay higher wages.

5 Empirical Support for Indirect Model Predictions

Next, we examine empirical evidence in support of the testable predictions of the model. The next section uses these facts to distinguish between our learning model and alternative theories of labor force participation.

5.1 Uncertainty Declined (1972-1996)

If the learning hypothesis is correct, then uncertainty about the effects of labor force participation on children should decline over time. (See section 3.2.) As beliefs become less uncertain, they are also converging to the truth. This convergence results in lower dispersion of beliefs in the population, over time.

This prediction of falling cross-sectional dispersion in beliefs is supported by our survey data. Using the data that are available through the 70's, 80's and 90's, we find that the dispersion in survey answers declined steadily throughout the last three decades.

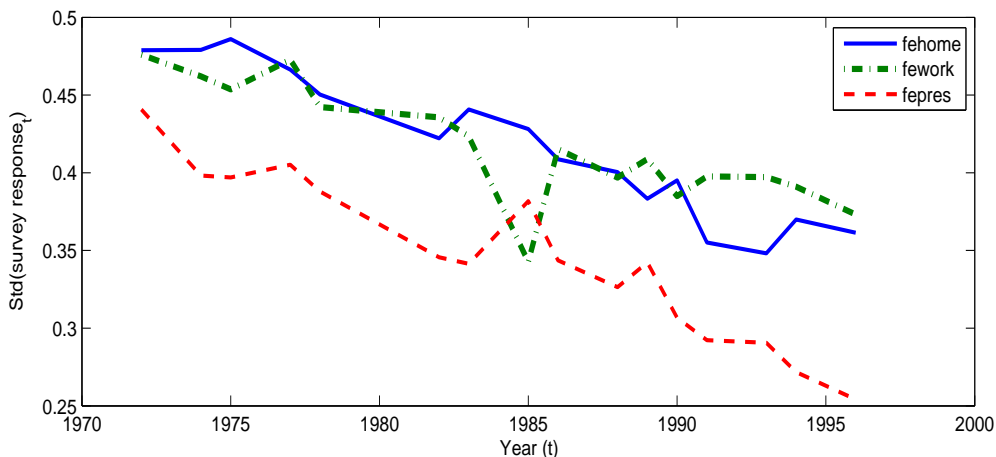


Figure 7: DISPERSION IN SURVEY RESPONSES DECLINES OVER TIME.

For the three variables that are observed in 1977 and then from 1985-2004, dispersion clearly decreases over the sample period. The only exception is *fefam*. That variable has no clear trend in its standard deviation. For the one variable that does not exhibit a fall in its dispersion, other measures of uncertainty show a decrease. The responses to the *fefam* question have both an agree or disagree dimension, but also a strength of agreement and disagreement and a possibility of answering “I don’t know.” Two features of this data point to a decline in uncertainty. First, the fraction of respondents who reply “I don’t know” fell from 1.5% in 1977 to 1.2% in 2004. But, this is a small number of replies. If we include also the responses agree or disagree, but exclude the strongly agree and strongly disagree responses, the fraction of less certain replies falls from 76% in 1977 to 71% in 2002.

5.2 Wage elasticity of labor declined

Blau and Kahn (2005) document that the elasticity of labor force participation to wages has decreased significantly over the last 20 years. The decline in wage elasticity is predicted by the model (proposition 4). A related fact shows up in our data: In the late 1980's and throughout the 1990's,




Figure 8: WAGES (SOLID LINE) AND FEMALE LABOR FORCE PARTICIPATION (DASHED LINE), 1930-2005.

women’s wages continued to increase and yet labor force participation stagnated (figure 8).

5.3 Female employment attitudes depend on children

The predictions of the model for beliefs depend crucially on the presence of children. A woman without a child is like one who knows that $\theta = 0$ for certain. Her actions cannot possibly harm her offspring. The data show clearly that the presence of a small child is crucial in determining individuals’ attitudes toward labor force participation. We have 3 survey measures about whether women, in various stages of their lifecycle, should work. Since we have these for only 3 years each (1988, 1994, 2002), we pool the three years. When asked if a woman should work full time outside the home after marrying and before there are children, 70% say yes. But when that woman has children who are in school, only 34% say yes. When a woman has a baby, the fraction of respondents who believe she should participate full time in the labor force falls to 10%. (See appendix B for data details.) These data tell us that the presence of children has an enormous influence on labor force participation decisions.

5.4 Ancestry affects labor choices

Using data on second generation American women, Fernández and Fogli (2005) show that cultural heritage, measured by country of origin, matters for female labor decisions.³ While it could be that childhood or background affect individual preferences, this evidence is also supportive of an information-based theory in which differences in cultural heritage or upbringing expose individuals to different information. By observing different outcomes of previous generations, women from different cultures form different beliefs about the costs of maternal employment.

Likewise, Fernández, Fogli, and Olivetti (2002) show that men who are born and raised by working mothers are more likely to marry a working woman. In our model only women's information affects their labor force participation decision, but it is logical to think that the husband's information is also aggregated when the labor force participation decision is jointly made.

6 Competing Explanations for Labor Force Participation

While the learning explanation appears to explain the qualitative and quantitative features of the labor force participation data, there are other logical explanations that could potentially do so as well. This section compares the predictions of our model with those of alternative explanations for the increase in women's labor force participation.

6.1 Women entered because their wages were rising

Wage-based explanations have a hard time reconciling the evolution of female labor force and wages over the last decade, particularly in the last 20 years while wages were rising and labor force participation stagnating (figure 8). Our theory (proposition 4) predicts this. Labor elasticity fell because uncertainty about the effects of labor market participation fell. Lower uncertainty made expected utility less sensitive to wage increases. Also, if women participate in the labor force simply because wages rise, there should be no systematic change in their uncertainty about

³Similar results, showing that ancestry affects labor choices can be found in Antecol (2000), Fortin (2005) and Alesina and Giuliano (2007).

the desirability of labor force participation. Wages have not become more observable over time. If anything, the increase in the dispersion of wages in recent years (Attanasio and Davis (1996), Krueger and Perri (2006)) should make the wage-based benefit of working more uncertain, not less. Finally, the wage-based theory misses cross-section differences in participation rates documented in section 1. Increases in wages are obviously an important component of the change in female labor force participation. But our theory can modify the predictions of wage based theories to allow them to match more nuanced features of the data.

6.2 Availability of child care and new technologies

Goods and services to assist working mothers encouraged labor force participation (Greenwood and Guner 2005). One might argue that as more women entered the labor force, technology reduced the cost of participation and its negative effects on children, accounting for the survey responses. We can think of this technological change as being like a change in the true value of nurture θ in our setting.

One shortcoming of this explanation, without incorporating a role for learning, is that it does not explain the observed decrease in uncertainty. If the availability of child care is known, then the average response about the effect of working mothers on families might change, but the dispersion of those responses should not. The second shortcoming is that it does not explain the fall in the labor supply elasticity. If anything, technology should increase elasticity. An increase in wages makes paying for child care or purchasing a washing machine more affordable, if that option is available. Therefore, if the availability of such goods and services is increasing over time, a second effect on the labor supply elasticity emerges: As wages rise, technology and child care become more affordable. Being able to afford them makes working trade off less with a nurturing home environment. Inside the model, this is like being able to sacrifice some labor earnings to reduce θ . That means that not only does an increase in wage increase the benefit to working, it also reduces its cost. This second effect makes labor supply more sensitive to the wage. Finally, the availability of child care or technologies would not explain some of the cross-sectional differences that learning can capture. For example, adding learning explains why women in the same place, at the same

time, make different decisions based on their parents' labor force decisions.

We do not deny that the arrival of new technologies and services influenced women's labor force participation. Rather, both mechanisms could be operating simultaneously. Perhaps technological innovations and child care became available when there were enough working mothers demanding them, which reinforced the trend to higher participation. The increased demand could come from a change in beliefs about the costs of labor force participation.

6.3 Preferences changed

The hardest alternative explanation to distinguish is a preference change. Perhaps women prefer staying at home when lots of other women stay at home, or preferences change adaptively as a consequence of previous history (Fernández, Fogli, and Olivetti 2002) or optimally in response to changing economic conditions (Doepke and Zilibotti 2007).

Three features of the data help us argue that learning is going on. They don't rule out simultaneous changes in preferences, but they support a role for our theory. First, the questions about beliefs in figure 1 questions are not about preferences. They are about the effect of work on children. Second, a change in preferences does not entail a fall in uncertainty. Third, and most importantly, changes in preferences would not cause the labor supply elasticity to fall over time. Preference changes don't have Jensen's inequality-driven effects. It is uncertainty, combined with the concavity of the utility function, that can generate the decreased sensitivity of participation to wages that we observe in the data.

Finally, this story and a preference story are not necessarily competitors. Our model can be interpreted as a theory of where preference changes come from. Such a theory is a useful contribution because it links observable determinants of information to observable consequences of learning and changing preferences; it isolates which preferences matter most for labor force participation, and it offers a testable, systematic way of thinking about why women's preferences have changed.

7 Conclusion

Female labor force participation is highly correlated with reported beliefs about the toll mothers' employment takes on their children. Not only is there a similar trend in these two variables, but they also share a non-linear S-shaped dynamic that other potential explanatory variables do not exhibit. We document these facts and argue that the S-shaped pattern in beliefs comes from Bayesian learning. Women learn about the effect of female employment on children by observing others' upbringing and outcomes. The feedback effect of learning on labor force participation and participation on the set of observable outcomes from which agents can learn generates the S-shaped dynamic. Our learning model delivers indirect testable implications for dispersion in beliefs and labor supply elasticity that are consistent with the data and distinguish our hypothesis from others.

More broadly, the theory provides some insight about how social changes may arise through learning from endogenous information. All kinds of cultural norms are adopted because they are thought to be best practice at the time. If there is no experimentation with alternatives, no new information is learned and the cultural norms stay fixed. Eventually, a few people with extreme preferences or abilities deviate from these norms. In doing so, they provide information that others can observe. This experimentation slowly begins to change beliefs, reducing uncertainty about the desirability of the alternative practice, which encourages others to engage in it. As more people deviate from the cultural norm, learning speeds up, social change quickens and a social revolution takes hold.

A Technical Appendix

A.1 Proof of proposition 1

Step 1: Define a cutoff wage \bar{w} such that all women who observe $w_{i,t} > \bar{w}$ choose to join the labor force.

A woman joins the labor force when the inequality in (7) is satisfied. Note that $\partial \mathcal{N}_{i,t} / \partial w_{i,t} = (n_{i,t} w_{i,t} + \omega_{i,t})^{-\gamma} > 0$. Since $\mathcal{N}_{i,t}$ is monotonically increasing in the wage w , there is a unique \bar{w} for each set of parameters, such that at $w = \bar{w}$, $\mathcal{N}_{i,t} = 0$.

Step 2: Describe the probability of labor force participation.

Let Φ denote the cumulative density function for the unconditional distribution of wages in the population. This is a log-normal c.d.f. Since the lognormal is unbounded and has positive probability on every outcomes, its c.d.f. is therefore strictly increasing in its argument. Then, the probability that a woman participates is $1 - \Phi(\bar{w})$, which is then strictly decreasing in \bar{w} .

Step 3: The effect of beliefs on labor force participation

Taking the partial derivative of the net utility gain from labor force participation yields $\partial \mathcal{N}_{i,t} / \partial \mu_{i,t} = -\beta$. By the implicit function theorem, $\partial \bar{w} / \partial \mu_{i,t} > 0$. Thus, $\partial(1 - \Phi(\bar{w})) / \partial \mu_{i,t} = (\partial(1 - \Phi(\bar{w})) / \partial \bar{w})(\partial \bar{w} / \partial \mu_{i,t}) < 0$, which completes the proof.

A.2 Proof of proposition 2

Steps 1 and 2 of the proof are as in appendix A.1.

The benefit to participating is falling in uncertainty: $\partial \mathcal{N}_{i,t} / \partial \sigma_{i,t} = (1 - \gamma)\beta$. Since $\gamma > 1$ and $\beta > 0$, by assumption, this means that $\partial \mathcal{N}_{i,t} / \partial \sigma_{i,t} < 0$. As before, the implicit function theorem tells us that $\partial \bar{w} / \partial \sigma_{i,t} > 0$. Thus, $\partial(1 - \Phi(\bar{w})) / \partial \sigma_{i,t} = (\partial(1 - \Phi(\bar{w})) / \partial \bar{w})(\partial \bar{w} / \partial \sigma_{i,t}) < 0$, which completes the proof.

A.3 Proof of proposition 3

Steps 1 and 2 of the proof are as in appendix A.1.

The benefit to participating is falling in wealth: $\partial \mathcal{N}_{i,t} / \partial \omega_{i,t} = (w_{i,t} + \omega_{i,t})^{-\gamma} - (\omega_{i,t})^{-\gamma}$. This is negative because $(w_{i,t} + \omega_{i,t}) > (\omega_{i,t})$ holds as long as $w_{i,t} > 0$. Since $w_{i,t}$ has a log-normal distribution, it is greater than zero with probability one. As before, the implicit function theorem tells us that $\partial \bar{w} / \partial \omega_{i,t} > 0$. Thus, $\partial(1 - \Phi(\bar{w})) / \partial \omega_{i,t} = (\partial(1 - \Phi(\bar{w})) / \partial \bar{w})(\partial \bar{w} / \partial \omega_{i,t}) < 0$, which completes the proof.

A.4 Proof of proposition 4

There are two effects of increasing the average wage. The first is the direct effect on increasing the number of women whose wages are above the cutoff that determines whether they join the labor force or not. The second effect is how the increase in the expected future wages of children affects the expected cost of joining the labor force.

Step 1: Direct effect. The probability that a woman participates is $1 - \Phi(\bar{w})$, which is then strictly decreasing in \bar{w} , as defined in appendix A.1. Rewrite the cumulative density function as the standard normal c.d.f. $\tilde{\Phi}$ of the log of \bar{w} , adjusted by the mean and standard deviation of log wages, μ_w and σ_w . The probability of labor force participation is then $1 - \tilde{\Phi}(\ln(\bar{w}) - \mu_w / \sigma_w)$. The marginal effect of increasing the average log wage is the standard normal probability density $\tilde{\phi}(\ln(\bar{w}) - \mu_w / \sigma_w) > 0$. The elasticity of labor force participation is $\partial \ln(1 - \tilde{\Phi}(\cdot)) / \partial \mu_w$, since μ_w is the log average wage. The elasticity is therefore $\tilde{\phi}(\ln(\bar{w}) - \mu_w / \sigma_w) / (1 - \tilde{\Phi}(\cdot))$, which is a standard normal hazard function.

Next, the question is how does the decreasing uncertainty about the effect of nurture change this elasticity over time? Uncertainty affects $(1 - \tilde{\Phi}(\ln(\bar{w}) - \mu_w / \sigma_w))$ through its effect on the cutoff wage \bar{w} . The proof of proposition A.2 tells us that $\partial \bar{w} / \partial \sigma_{i,t} > 0$. Therefore, the effect of an increase in uncertainty is given by the cross-partial derivative $\partial^2 \ln(1 - \tilde{\Phi}(\cdot)) / \partial \mu_w \partial \bar{w} > 0$. This is positive because a standard normal hazard function is always increasing in its argument.

Step 2: Indirect effect on children. This effect is the increase in the expected value of $(\mu_a - n_{i,t} \mu_{i,t}) = \mu_w$ in the second term of (6). This indirect effect works through the effect of μ_w on \bar{w} . Putting the two effects together gives us the total elasticity of labor force participation to a change in average wage:

$$\text{Elasticity} = H(\ln(\bar{w}) - \mu_w / \sigma_w) + \frac{\partial \ln(1 - \tilde{\Phi}(\cdot))}{\partial \bar{w}} \frac{\partial \bar{w}}{\partial \mu_w}$$

$$= H(\ln(\bar{w}) - \mu_w/\sigma_w)(1 - \frac{\partial \bar{w}}{\partial \mu_w})$$

Lastly, we need to compute the derivative of the total elasticity with respect to uncertainty $\sigma_{i,t}^2$.

$$\frac{\partial \text{Elasticity}}{\partial \sigma_{i,t}^2} = H'(\ln(\bar{w}) - \mu_w/\sigma_w) \left(\frac{\partial \bar{w}}{\partial \sigma_{i,t}^2} \right) \left(1 - \frac{\partial \bar{w}}{\partial \mu_w} \right) - H(\ln(\bar{w}) - \mu_w/\sigma_w) \frac{\partial^2 \bar{w}}{\partial \mu_w \partial \sigma_{i,t}^2}.$$

Step 1 established that the first term is positive. Since hazard functions are always positive, the last remaining step is to show that $\partial^2 \bar{w} / \partial \mu_w \partial \sigma_{i,t}^2 < 0$.

Step 3: Show that $\partial^2 \bar{w} / \partial \mu_w \partial \sigma_{i,t}^2 < 0$. Note that

$$\partial \mathcal{N}_{i,t} / \partial \mu_w = \beta \exp \left(\mu_a(1 - \gamma) + \frac{1}{2} \sigma_a^2(1 - \gamma)^2 \right) \left(\exp \left((\gamma - 1)n_{i,t}\mu_{i,t} + \frac{1}{2} n_{i,t}^2 \sigma_{i,t}^2 (1 - \gamma)^2 \right) - 1 \right).$$

This is positive because $(\gamma - 1)n_{i,t}\mu_{i,t} + \frac{1}{2} n_{i,t}^2 \sigma_{i,t}^2 (1 - \gamma)^2 > 0$. This tells us that when expected wages are higher, $\mathcal{N}_{i,t}$, the net benefit of working is higher. This arises because the marginal utility cost of a reduction in expected wages for children is lower if the expected level of those wages is higher. Differentiating this expression again with respect to uncertainty yields:

$$\frac{\partial^2 \mathcal{N}_{i,t}}{\partial \mu_w \partial \sigma_{i,t}^2} = \beta \exp \left(\mu_a(1 - \gamma) + \frac{1}{2} \sigma_a^2(1 - \gamma)^2 \right) \exp \left((\gamma - 1)n_{i,t}\mu_{i,t} + \frac{1}{2} n_{i,t}^2 \sigma_{i,t}^2 (1 - \gamma)^2 \right) \frac{1}{2} n_{i,t}^2 (1 - \gamma)^2 > 0.$$

Note also that $\partial^2 \mathcal{N}_{i,t} / \partial \mu_w \partial w = 0$ and $\partial^2 \mathcal{N}_{i,t} / \partial \sigma_{i,t}^2 \partial w = 0$. Therefore, applying the implicit function rule and differentiating a second time tells us that

$$\frac{\partial^2 \bar{w}}{\partial \mu_w \partial \sigma_{i,t}^2} = - \frac{\partial^2 \mathcal{N}_{i,t} / \partial \mu_w \partial \sigma_{i,t}^2}{\partial \mathcal{N}_{i,t} / \partial w} + 0.$$

Since the numerator is positive and the denominator is positive, the cross-partial is negative.

B Data Description

B.1 Survey data about beliefs

Data come from <http://webapp.icpsr.umich.edu/GSS/>.

Attitudes We have 6 measures for individuals' attitudes toward women working that have a time series dimension.

fehome Women should take care of running their homes and leave running the country up to men. (AGREE=1, DISAGREE=2, NOT SURE=8, NO ANSWER=9, NA=0). We generate a dummy variable fehome=1 if FEHOME=1 and =0 if FEHOME=2. exclude observations for which FEHOME= 0 or FEHOME > 2.

The data span 1974-1998, with nine missing years. There are between 890 and 1884 responses per year, with 16 years of observations, for a total 22,538 data points. The fraction of respondents who agree ranges from 38% in 1977 to 14% in 1995.

fework Do you approve or disapprove of a married woman earning money in business or industry if she has a husband capable of supporting her? (AGREE=1, DISAGREE=2, NOT SURE=8, NO ANSWER=9, NA=0). We generate dummy variable fework same as for fehome.

The data span 1972-1998, with ten missing years, for a total of 17 years with observations. There are between 902 and 1,933 responses per year, and 24,401 total observations. The fraction of respondents who agree ranges from 48% in 1972 to 37% in 1996.

fePRES If your party nominated a woman for President, would you vote for her if she were qualified for the job? (YES=1, NO=2, WOULD NOT VOTE=5, DON'T KNOW=8, NO ANSWER=9, NA=0). The dummy variable is fePRES=1 if FEPRES=1 and =0 if FEPRES=2. We exclude observations for which FEPRES > 2 or FEPRES= 0.

The data span 1972-1998, with ten missing years. This leaves 17 years in which we have observations. There are between 902 and 1,933 responses per year, for a total of 24,072 observations. The fraction of respondents who agree ranges from 44% in 1972 to 25% in 1998.

fechld A working mother can establish just as warm and secure a relationship with her children as a mother who does not work. (STRONGLY AGREE=1, AGREE=2, DISAGREE=3, STRONGLY DISAGREE=4, DON'T KNOW=8, NO ANSWER=9, NA=0). We use this variable as is, except observations with FECHLD= 0 and FECHLD> 4 which we treat as missing.

There is one set of observations in 1977, and then from 1995-2004, at least every two years, for a total of 14 years in which data are available. There are between 897 and 1,948 responses per year, totalling 19,270 observations. The average reply ranges from 2.5 in 1977 to 2.1 in 1994.

preschool A preschool child is likely to suffer if his or her mother works. (STRONGLY AGREE=1, AGREE=2, DISAGREE=3, STRONGLY DISAGREE=4, DON'T KNOW=8, NO ANSWER=9, NA=0). The only modification we make to this variable is to treat observations with preschool= 0 and preschool> 4 as missing.

There is one set of observations in 1977, and then from 1995-2004, at least every two years, for a total of 14 years in which data are available. There are between 890 and 2,344 responses per year, totalling 19,005 observations. The average reply ranges from 2.2 in 1977 to 2.6 in 2004.

fefam It is much better for everyone involved if the man is the achiever outside the home and the woman takes care of the home and family. (STRONGLY AGREE=1, AGREE=2, DISAGREE=3, STRONGLY DISAGREE=4, DON'T KNOW=8, NO ANSWER=9, NA=0). The only modification we make to this variable is to treat observations with FEFAM= 0 and FEFAM> 4 as missing.

There is one set of observations in 1977, and then from 1995-2004, at least every two years, for a total of 14 years in which data are available. There are between 894 and 2,353 responses per year, totalling 19,024 observations. The average reply ranges from 2.2 in 1977 to 2.6 in 2004.

Acceptance We have 3 measures of acceptance of women working over life cycle. We do not have much of times series dimension (only start in 1988), but they are interesting to show how the presence of a small child is crucial in determining individuals' attitudes.

Attitudes toward women working over life cycle: 3 stages of life cycle corresponding to 3 variables. We have these 3 variables for only 3 years (1988, 1994, 2002). Not much variation over 14 years, but can be used to make cross-section point. First 9 tables show each variable in each year. Last 3 show the three variables when all observations are pooled across years.

wrknokid Should a woman work outside the home after marrying and before there are children? (1=FULL, 2=PART, 3=STAY HOME, 8=CAN'T CHOOSE, 9=NO ANSWER, 0=NA). Generate variable wrknokid=WRKNOKID but recode WRKNOKID=0 and WRKNOKID=9 as missing.

wrkbaby Should a woman work outside the home when there is a child under school age? (1=FULL, 2=PART, 3=STAY HOME, 8=CAN'T CHOOSE, 9=NO ANSWER, 0=NA). Generate variable wrkbaby=WRKBABY but recode WRKBABY=0 and WRKBABY=9 as missing.

wrksch Should a woman work outside the home after the youngest child starts school? (1=FULL, 2=PART, 3=STAY HOME, 8=CAN'T CHOOSE, 9=NO ANSWER, 0=NA). Generate variable wrksch=WRKSCH but recode WRKSCH=0 and WRKSCH=9 as missing.

Pre-1972 Data The attitudes data from GSS begin only in 1972. However, the increasing speed of female entry in the labor force (start of the S) precedes that date. To establish the contemporaneous S-shaped evolution of beliefs, it is vital to have more historical data.

We have one measure of beliefs that is collected infrequently, since the 1930's. That is FEWORK. This data are from IPOLL databank, maintained by the Roper Center for Public Opinion Research. Unfortunately, the phrasing of the questions differs slightly over time. We describe below the questions and the replies.

August 1936 The Gallup Poll asked: "Should a married woman earn money if she has a husband capable of supporting her?" 18% said yes, 82% no. No uncertain or no response entries were allowed.

October 1938 The Gallup Poll asked: "Do you approve of a married woman earning money in business or industry if she has a husband capable of supporting her?" 22% approve, 78% disapprove.

November 1945 The Gallup Poll (AIPO) asked: "Do you approve or disapprove of a married woman holding a job in business and industry if her husband is able to support her?" 62% disapprove, 18% approve. The rest of the replies are miscellaneous open answers (e.g., if she has a good job, if she has no children, etc.).

June 1970 The Gallup Poll asked: "Do you approve of a married woman earning money in business or industry if she has a husband capable of supporting her?" 60% approve, 36% disapprove, 4% do not know.

B.2 Labor force participation and wage data

The wages data come from IPUMS (Integrated Public Use Microdata Series). It covers 1920-2005. It is every 10 years until 2000, and yearly thereafter. The data are a 1-in-100 national random sample of the population until 2000. After 2000, it is the American Community Survey 2001 sample: a 1-in-232 (approximately) national random sample of the population. We used only data on white women who are not agricultural workers and live with family members in their household: (race=1, gq=1, farm=1, sex=2, excluding occ=810, 820, 830, 840, 123, 100). Residents of Hawaii (statefip=15) and Alaska (statefip=2) are excluded.

Labor Force Observations are dropped if the occupation is in agriculture or the respondent lives on a farm or in an institution. Ages range from 25-54. They are white, married women with spouse present and at least one child under age 5 living in the household. Data are weighted using perwt (sample weights from census). There are no missing labor force observations.

Wages Data start in 1940. (Inctot and incwage start in 1940.) We deflated wages using the consumer price index (Series Id: CUUR0000SA0). The base year is 1982 – 84 = 100. Note that since the survey always asks about the previous year’s wages, the 1940 wage data are really about 1939 wages.

We make three modifications to these data. First, we recode the labor force variable and the empstat (employment status) variable as missing value if labforce=0 or empstat=0, meaning that the answer was N/A. Second, to account for the top-coding of the income variable, we multiply the top-coded values of INCWAGE (wage income) by a factor of 1.4. Finally, we take the log of each woman’s wage, then average across women.

Hours and weeks worked The weeks worked variable (WKSWORK2) is an indicator variable that describes which interval the respondent’s answer fell in. We replace each indicator with the middle point of the reported interval. We do the same for hours worked (hrswork2). The reporting procedure changes in 1980. From 1980 on, we used the reported average hours per week (UHRSWORK).

Full time This is a dummy variable that indicates if two conditions are satisfied: 1) hrswork2 recoded ≥ 40 , meaning that the worker worked at least 40 hours per week; and 2) WKSWORK2 recoded ≥ 50 , meaning that the woman reports being employed for at least 50 weeks.

B.3 Calibration details

Endowment distribution Data come from the census. We use wages in 1940 (first available year). We use the full sample of individuals not living in institution, not living in farm, not working in agriculture and white. From this, we construct two pools of matched data: One is only married women; the other is their husbands. We look at women 25-54, with their own child younger than 5 living in the household.

The log endowment is normal. For these two sets of wage data, we take the log of wages over previous year. For husbands, $\text{mean}(\log_incwage_husb) = 7.043089$ and $\text{std}(\log_incwage_husb) = .7348059$. Therefore, we set $\sigma_\omega = 0.73$. We choose the mean log endowment $\mu_\omega = -(\text{sig}_\omega^2)/2$ such that mean endowment is normalized to 1.

Initial labor force participation We need a period 0 participation rate that determines the period 1 wage distribution of women and a period 1 participation rate to start the simulation. Period zero FLFP=3%, period one FLFP=6%. These are given from Census 1930 and 1940 for women, married, white, not living in farms, not living in institutions, age 25-45 with a child younger than 5 living in the household. Finally, the true θ matches 1990-2000 participation rate of 60%.

Abilities $\sigma_a = .9$ is the standard deviation of log ability and $\mu_a = \ln(\text{earnings gap}) - (\text{sig}_a^2)/2$ is the mean of log ability. These parameters match the initial ratio between average earnings of working women to average earnings of all husbands (0.4 in the data) and to match the standard deviation of log earnings of women in the data (1.08 in the data).

References

ALBANESI, S., AND C. OLIVETTI (2006): “Home Production, Market Production and the Gender Wage Gap: Incentives and Expectations,” Working Paper.

- ALESINA, A., AND P. GIULIANO (2007): “The Power of the Family,” Harvard University Working Paper.
- AMADOR, M., AND P.-O. WEILL (2006): “Learning from Private and Public Observations of Others’ Actions,” Working Paper.
- ANTECOL, H. (2000): “An Examination of Cross-Country Differences in the Gender Gap in Labor Force Participation Rates,” *Labour Economics*, 7, 409–426.
- ATTANASIO, O. P., AND S. J. DAVIS (1996): “Relative Wage Movements and the Distribution of Consumption,” *Journal of Political Economy*, 104(6), 1127–1262.
- BARRO, R., AND MCCLEARY (2006): “Religion and Economy,” *Journal of Economic Perspectives*, Spring.
- BELSKY, J. (1988): “The Effects of Infant Day Care Reconsidered,” *Early Childhood Research Quarterly*, 3, 235–272.
- BERNAL, R. (2006): “The Effect of Maternal Employment and Child Care on Children’s Cognitive Development,” Northwestern University, Working Paper.
- 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(3), 955–988.
- BLAU, F., AND L. KAHN (2005): “Changes in the Labor Supply Behavior of Married Women: 1980-2000,” NBER Working Paper 11230.
- DOEPKE, M., AND F. ZILIBOTTI (2007): “Occupational Choice and the Spirit of Capitalism,” Working Paper.
- FERNÁNDEZ, R., AND A. FOGLI (2005): “An Empirical Investigation of Beliefs, Work and Fertility,” NBER Working Paper 11268.
- FERNÁNDEZ, R., A. FOGLI, AND C. OLIVETTI (2002): “Marrying Your Mom: Preference Transmission and Women’s Labor and Education Choices,” NBER Working Paper 9234.
- FORTIN, N. (2005): “Gender Role Attitudes and the Labor Market Outcomes of Women Across OECD Countries,” Working Paper.
- GIULIANO, P. (2007): “Living Arrangements in Western Europe: Does Cultural Origin Matter?,” *Journal of the European Economic Association*, forthcoming.
- GOLDIN, C. (1990): *Understanding the Gender Gap*. Oxford University Press.
- GOLDIN, C., AND L. KATZ (2002): “The Power of the Pill: Oral Contraceptives and Women’s Career and Marriage Decisions,” *Journal of Political Economy*, 100, 730–770.
- GREENWOOD, J., AND N. GUNER (2005): “Social Change,” *Economie d’avant gard*, research Report 9, University of Rochester.

- GREENWOOD, J., A. SESHADRI, AND M. YORUKOGLU (2001): “Engines of Liberation,” *Economic d’avant gard*, research Report 2, University of Rochester.
- GUISSO, L., P. SAPIENZA, AND L. ZINGALES (2004): “Cultural Biases in Economic Exchange,” NBER Working Paper 11005.
- HARVEY, E. (1999): “Short-Term and Long-Term Effects of Early Parental Employment on Children of the National Longitudinal Survey of Youth,” *Developmental Psychology*, 35(2), 445–459.
- HILL, J., J. WALDFOGEL, J. BROOKS-GUNN, AND W. HAN (2005): “Maternal Employment and Child Development: A Fresh Look Using Newer Methods,” *Developmental Psychology*, 41(6), 833–850.
- JONES, L., R. MANUELLI, AND E. MCGRATTAN (2003): “Why Are Married Women Working So Much?,” Research Department Staff Report 317, Federal Reserve Bank of Minneapolis.
- KEPPO, J., G. MOSCARINI, AND L. SMITH (2005): “The Demand for Information: More Heat Than Light,” NBER Working Paper.
- KRUEGER, D., AND F. PERRI (2006): “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory,” *Review of Economic Studies*, 73(1).
- RADNER, R., AND J. STIGLITZ (1984): “A Nonconcavity in the Value of Information,” in *Bayesian Models in Economic Theory*, ed. by M. Boyer and R. Kihlstrom. Elsevier Science Publishers B.V.
- VAN NIEUWERBURGH, S., AND L. VELDKAMP (2006): “Learning Asymmetries in Real Business Cycles,” *Journal of Monetary Economics*, 53(4), 753–772.