A Study of U.S. Employment Rates with Emphasis on Gender Considerations

Carl D. Lantz and Pierre-Daniel G. Sarte

Policymakers who are concerned with understanding the cyclical state of the economy will be interested in identifying the sustainable rate of employment growth. If productivity growth is taken to be exogenous, the rate at which jobs are created will determine growth in output. Assuming that employment is measured reasonably well, a policymaker armed with a measure of trend job growth will be able to glean information about the economy's status from the monthly employment report.

When attempting to determine this trend rate, one might posit as a first approximation that employment growth should be equal to the growth rate of the population. At the very least, this assumes that the fraction of the population employed will remain constant over time. However, a cursory look at U.S. data reveals that this fraction has been rising persistently over the postwar period. A robust determination of the sustainable rate of job creation should therefore take account of this trend. Upon further inspection of the available data, it becomes apparent that the driving forces behind this process have been demographic trends, i.e., the "baby boom," and, to a larger extent, the increased participation of women in the labor market. In this article we focus our attention on the latter determinant.

Female employment, expressed as a fraction of the working age population, has been increasing steadily over the past 50 years. While most work in this area has focused on the socioeconomic factors underlying this phe-

The opinions expressed herein are the authors' and do not necessarily represent those of the Federal Reserve Bank of Richmond or the Federal Reserve System. We wish to thank Michael Dotsey, Wenli Li, and Yash Mehra for a number of useful suggestions. Any remaining error is, of course, our own.

nomenon, we explore the phenomenon's effects on the larger economy within the framework of a stochastic dynamic general equilibrium model. Specifically, we use a model that allows households to make labor supply decisions along both the extensive margin—the decision to work outside the home or not—and the intensive margin, along which workers adjust the number of hours supplied to the outside market.

In our first experiment, we model a reduction in participation costs for female workers, that is, the implicit utility cost to the representative household incurred when a female member devotes time to outside work. Several factors may have contributed to such a phenomenon. For instance, advances in household technology such as those embedded in modern appliances have automated many of the chores previously accomplished entirely through physical labor. A general decline in birthrates over much of the period in question may have reduced the cost of childcare for the average working family. Furthermore, part of the utility cost to females choosing to work outside the home can be interpreted as the unfavorable social stigma historically associated with working women. As the severity of this stigma has declined in recent decades, the cost of participating in the outside labor market has fallen for women. We find that this decline in participation costs leads not only to an increase in female participation, but also to a decline in the fraction of males who are market employed. The latter result follows naturally as women replace men in the workplace, but it is also reinforced by a wealth effect that leads men optimally to choose greater leisure as overall consumption increases. We also find that the real rate of interest in our model economy is temporarily affected by this process. What is left unexplained in this first experiment is the well-documented earnings gap that has existed historically between male and female workers.

In our second experiment, we attempt to capture the effects of demand-side, or employer-based, discrimination against female workers in order to partially explain the behavior of the gender earnings gap. We do not attempt to identify which of the potential causes of discrimination is most relevant, nor do we take a stand on its microeconomic underpinnings. Instead, we model discrimination as a reduced form constraint on the amount of female labor that firms are willing to employ. We then allow for this exogenous constraint to be relaxed over time. In this scenario, we find that reduced employment opportunities for women generate both higher male employment and a larger gender earnings gap than can be explained by human capital differences alone. As opportunities for women expand in the model, female employment naturally increases and the earnings gap narrows. As in the first experiment, male employment falls partly as the result of optimal household behavior with respect to leisure. Interestingly, our model predicts that the gender earnings gap will be countercyclical when changes in total factor productivity are the main sources of shocks to the economy. These results are all consistent with postwar U.S. labor market experience.

The article is divided into three parts. In the first section, we analyze some labor market data and attempt to estimate some reasonable approximations of trend employment growth. The fact that employment rates have been far from constant over time complicates this estimation. Two of the main forces driving the secular change in employment rates have been changes in demographics and the gender composition of employment. We leave the analysis of demographics to further research, but understanding the role of gender motivates the work in sections two and three.

1. SOME DATA ANALYSIS OF U.S. LABOR MARKETS

To start, let us think of the working population (i.e., age 16 and over), P, in terms of three distinct pools: employed workers, denoted E; the unemployed, U; and those outside the labor force who are not actively looking for work, Z. We denote the labor force by N. Thus, we have that

$$P = \underbrace{E + U}_{N} + Z. \tag{1}$$

Given equation (1), perhaps the simplest way in which to think about the evolution of aggregate employment is to consider a long-run equilibrium where each pool on the right-hand side of the equation grows at a constant rate. Letting γ_X denote the constant growth rate of pool X over a given period, it must then be the case that $\gamma_P = \gamma_E = \gamma_U = \gamma_Z$. In this scenario, we can think of the level of job creation, ΔE, over the period used to measure growth as

$$\Delta E = \gamma_P E. \tag{2}$$

With employment currently at about 130 million and population growth hovering around 1 percent annually, it turns out that ΔE equals roughly 110,000 at a monthly rate.

Equation (2) is often used to approximate a sustainable level of job creation, that is, a level of job creation that does not generate undue strain in labor markets. Accordingly, any job creation in excess of 110,000 might be considered above trend. However, the simple calculation we have just carried out is subject to an important problem. If the growth rate of each pool is constant and equal to population growth in some long-run equilibrium, the number of employed workers, unemployed workers, and people outside the labor force as a fraction of the total population should also be constant. Looking at Figure 1, we can see that this has never in fact been the case.

While the unemployment rate, U/P, is often the focus of attention in both academic research and the popular press, Figure 1 clearly shows that other considerations also deserve attention. Specifically, in Figure 1, panel b, we

Figure 1 U.S. Employment Trends

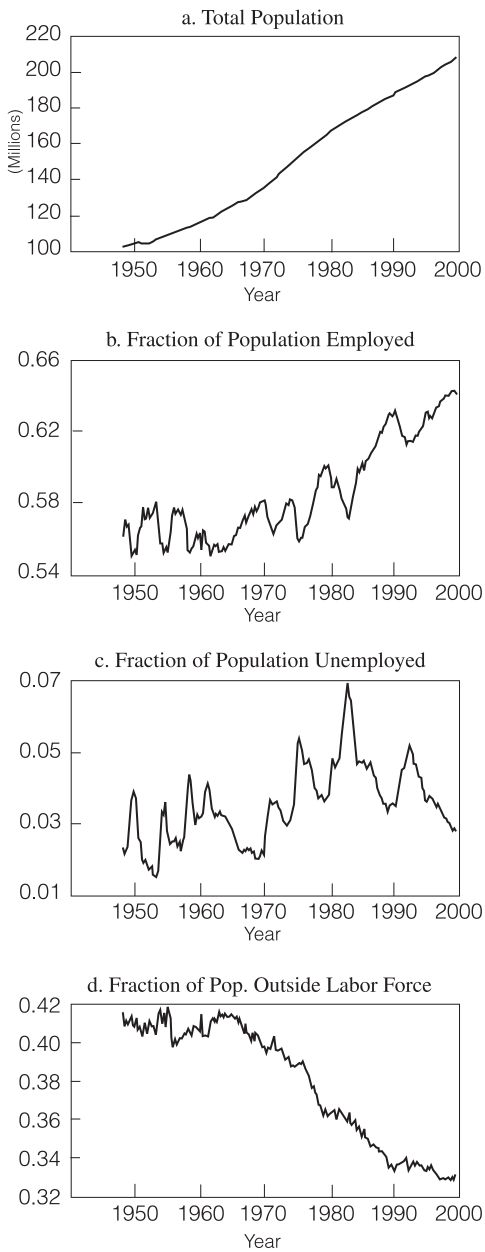

see that the employment rate, E/P, has shown a steady but important change over the years. It was as low as 0.55 in 1950 and has gradually increased to a high of 0.64 in 1999. Similarly, the fraction of population outside the labor force, Z/P, has steadily fallen from approximately 0.41 in the early 1950s to 0.33 today. In contrast, Figure 1, panel c, suggests that the unemployment rate has remained roughly constant throughout the years at around 0.045. Thus, it seems that to acquire a grasp of U.S. labor markets, it is as important to understand movements in and out of employment that possibly stem from changes in the population outside the labor force as it is to understand variations in unemployment. Moreover, because changes in E/P and Z/P appear to have been very gradual over time, the source of these changes is likely to have been structural in nature.

As we have just argued, it is not entirely clear in Figure 1, panel b, that the fraction of employed workers has ever been in a steady state. More importantly, at first glance it seems anyone's guess where this fraction might settle, if ever. One way to approach this issue would be simply to make a statistical guess as to the behavior of E/P that allows for non-constant growth over time. We would then be able to infer sustainable levels of job creation without assuming that the share of each labor market pool in population is constant. Thus, define ρ_E as E/P and let

$$\rho_E(t) = \frac{a_0}{1 - a_1 e^{-a_2(t-1948)^2}}, \tag{3}$$

where $t > 0$ denotes time. We choose the functional form depicted in (3) because it will generally give rise to S-shaped curves as suggested by Figure 1, panel b. Furthermore, this functional form is convenient since the function starts at $a_0/(1 - a_1)$ in 1948 and eventually asymptotes to a_0. The parameters a_0 and a_1, therefore, determine the bounds of the function while a_2 controls its degree of curvature. These features will allow us to keep $\rho_E(t)$ bounded between zero and one. Under different assumptions about what $\rho_E(t)$ might be in the long run (i.e., different values of a_0), we can then use the data available on E/P to estimate a_1 and a_2 by Non-Linear Least Squares. The results from this estimation are presented in Figure 2.

Under the assumption that the fraction of employed workers should eventually settle at 0.65, close to its current level, it turns out that a sustainable level of job creation for the year 2000 would be around 110,000 jobs at a monthly rate. However, this number may be as high as 130,000 jobs for 2000 if we expect that 3/4 of the working population should eventually be employed in the long run. It is interesting to note that in all reasonable scenarios, current employment to population ratios are above trend. In fact, for current employment to population ratios to be considered on trend, we would have to expect all of the working population to be employed in the long run, which is perhaps a somewhat unrealistic scenario.

Figure 2 Non-Linear Estimation of Trend Employment

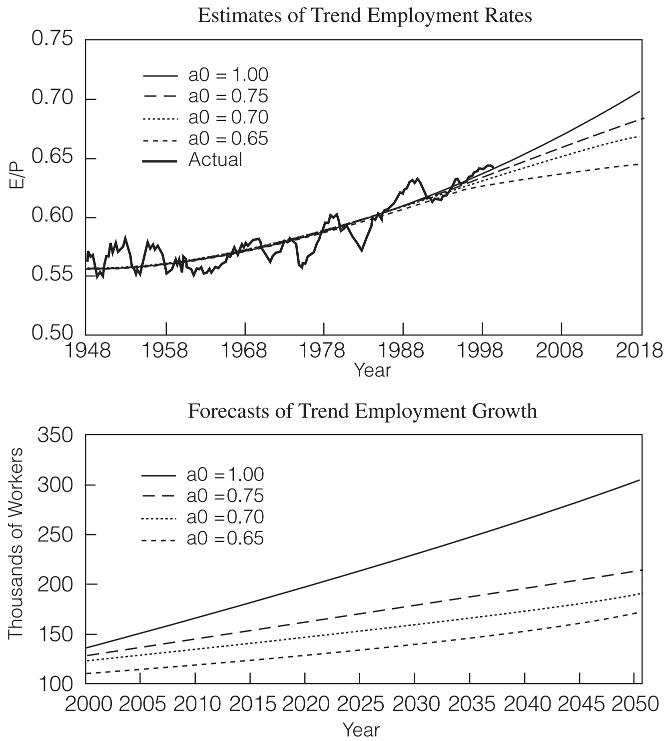

Forecasts of Average Monthly Employment Growth by Year

a_0	2000	2001	2002	2003	2004	2005	2010	2020	2030	2040	2050
0.65	109,827	110,621	111,410	112,195	112,978	113,761	117,753	119,847	122,041	124,358	126,814
0.70	122,033	123,417	124,786	126,138	127,475	128,797	135,197	138,285	141,318	144,312	147,289
0.75	128,446	130,264	132,076	133,881	135,679	137,469	146,283	150,595	154,840	159,020	163,136
1.00	138,573	141,253	143,970	146,724	149,513	152,338	166,985	174,622	182,457	190,479	198,676

Once it has been established that the fraction of employed workers has been far from constant over the last 50 years, a natural question to ask is: What structural changes have driven the evolution of E/P during this period? We explore the two most important sources of structural change in this section: the continually changing role of gender in labor markets and the impact of demographic considerations.

The Role of Gender

Figure 3 illustrates employment rates by gender and the fraction of females and males outside the labor force relative to their working population. The most striking feature of Figure 3 is that the employed males to male population ratio and the males outside the labor force to male population ratio have steadily moved in a direction opposite to that suggested by their aggregate counterparts. The fraction of employed males began to increase in 1990 but currently falls well short of its historic high in the early 1950s. Similarly, the fraction of males not looking for employment (relative to the male population) is now twice what it was in the late 1940s. In sharp contrast, the employment rate of women has gradually risen from a low of 0.30 in the late 1940s to a high of 0.58 today; that is, only 30 percent of the female working population was employed at the end of World War II. The fraction of females outside the labor force has also steadily fallen over the postwar period. Hence, it appears that aggregate employment rates are largely driven by the increasing participation of women in the labor force. Furthermore, assuming that there are no fundamental differences between men and women in their preferences towards work, we might expect that E/P and Z/P should converge to the same values for these two groups.

Figure 4 shows employment rates by gender as well as cohort and further reinforces the ideas we have just presented. Consider, for instance, the cohort of women who are 20 years old in 1948–1949. The employment rate for women actually falls between ages 20 and 30, increases up to age 50, and drops off as they retire. The fall in the employment rate between ages 20 and 30 can presumably be attributed to childbirth or the rearing of young children. In contrast, consider the cohort of women 20 years of age in 1968–1969. Not only is their employment rate higher across the board, but also the dip that occurs at age 30 does not stand out. It is noteworthy that the male employment rate only shows a small decrease by cohort. Overall, therefore, increases in the employment rate at all ages are mostly driven by female labor behavior.

The continuously increasing female employment rate just documented is generally thought to be the result of both demand- and supply-side factors. On the demand side, Jacobsen (1994) argues that part of the increase can simply be attributed to a general rise in labor demand stemming from technological advances in production. In addition, the noticeable rise in women's education over the past 50 years has led to increased demand for female workers in a world that is becoming more and more service oriented. Of course, a question immediately arises as to what factors prompted the growth in women's education in the first place. According to Jacobsen (1994, p. 128), these factors include "a relaxation of social restrictions on appropriate levels and types of education for women, and greater resources on the part of families who might previously have had to ration education among their children." Finally, a decline in labor market discrimination against women may also have contributed to a rise in female labor demand.

Explanations of demand-side discrimination usually fall into two broad categories. On one hand, neoclassical models propose that labor market out-

Figure 3 U.S. Employment Rates by Gender

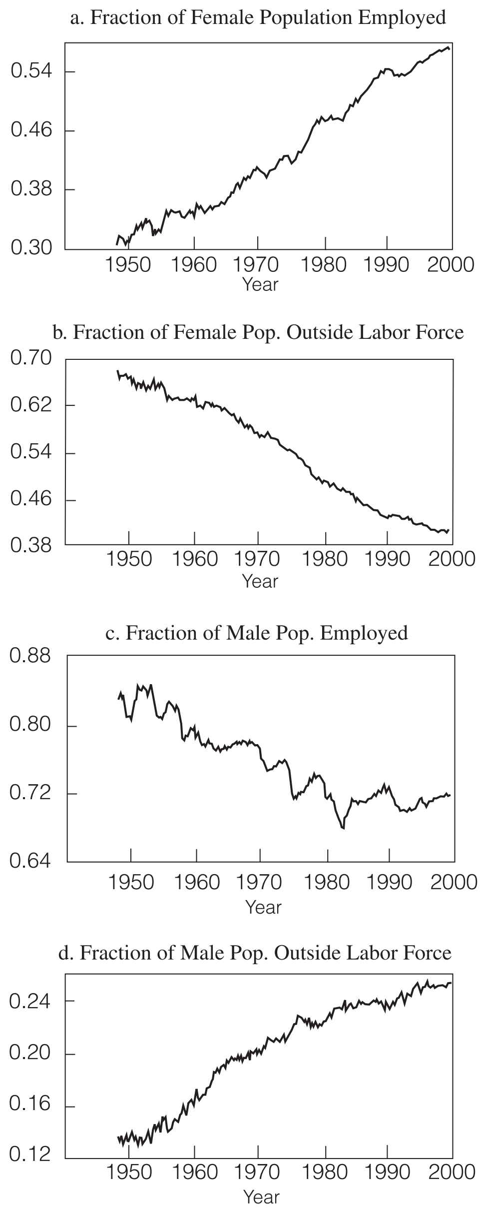

Figure 4 Employment Rates by Cohort

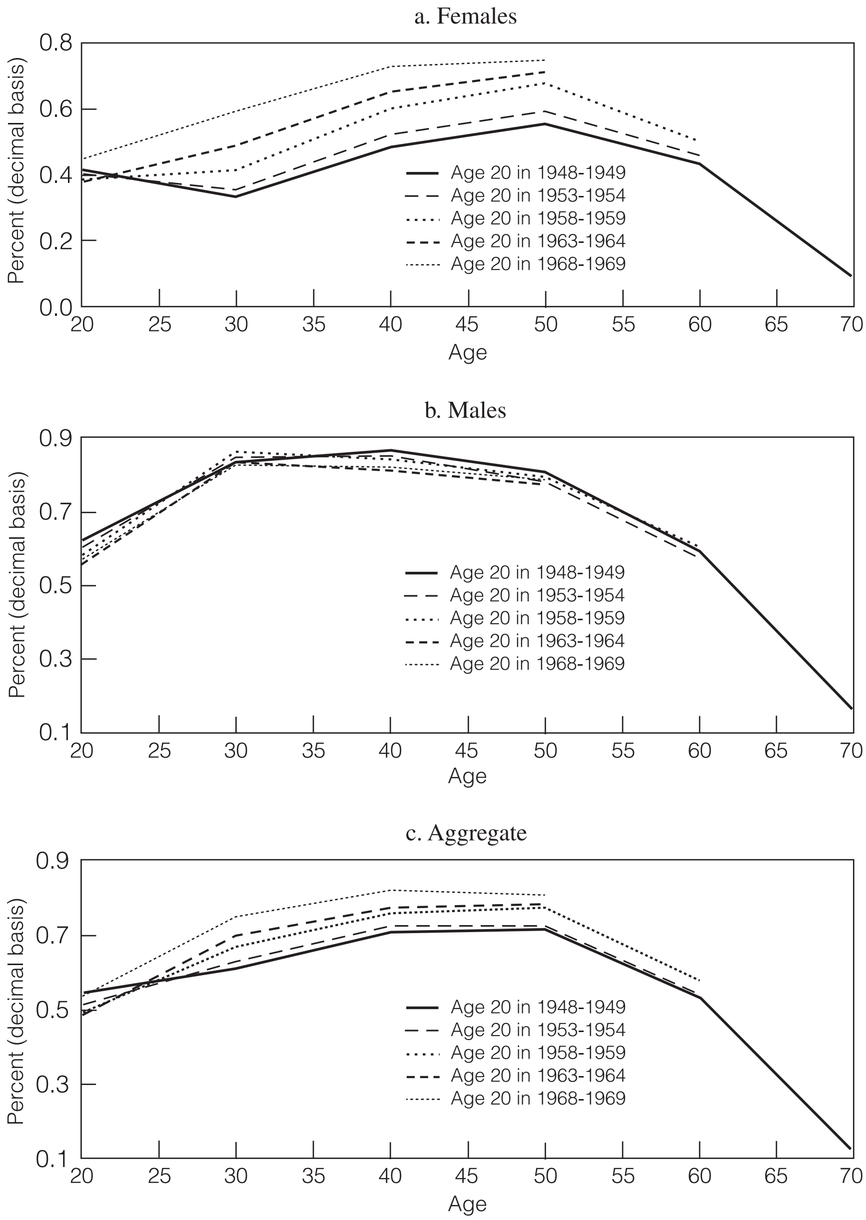

comes are in part a function of agents' personal prejudices or tastes against associating with particular demographic groups, as in Becker (1957). In the case of gender discrimination, these tastes may be a function of employers' or customers' perception of appropriate roles for women and men. On the other hand, models of statistical discrimination assume that employers must make hiring decisions in the face of incomplete information or uncertainty, as in Phelps (1972). Since it is impossible to assess the exact level of productivity associated with a particular job candidate, employers make inferences based on observed or perceived correlation between productivity and various employee characteristics. To the degree that women are perceived to be less productive or dependable than their male counterparts, they will face reduced employment opportunities and wages. One might expect such perceptions to be biased by long-standing attitudes toward gender differences, which may adjust only slowly over time, even in the face of accumulated evidence to the contrary.

Antidiscrimination legislative efforts may have played an important role in reducing the constraints on female employment opportunities during the postwar period. The Equal Pay Act of 1963 required employers to pay the same wage to women and men who do substantially equal work. The notion that this act was even necessary is suggestive that some degree of discrimination against women was taking place in the years that followed World War II. Title VII of the Civil Rights Act of 1964 prohibited employment discrimination on the basis of race, religion, national origin, or sex. The Equal Opportunity Act of 1972 strengthened the 1964 legislation by expanding its coverage to state and local governments as well as educational institutions. Furthermore, it granted the Equal Employment Opportunity Commission (EEOC) the power to sue private sector respondents. Beller (1982) finds that enforcement of Title VII, as measured by the number of completed EEOC investigations and successful settlements, was significant in reducing the gender gap both in wages and in the probability of being employed in male dominated occupations for the period 1967 to 1974. Furthermore, these studies indicate that Title VII was more effective following the passage of the 1972 amendment.

On the supply side, economic considerations that are likely to have induced increased female employment rates include rising wages for women, changes in family composition, and especially changes in non-market production technology. Jacobsen (1994, p. 129) writes that during the "twentieth century, technology has been widely adopted that has enabled families to produce non-market output at lower cost. In particular, we have seen the spread of market goods and services that serve as critical inputs into non-market production. In 1920, one-third of homes had electricity; . . . by 1960, practically all homes were electrified. In 1940, 17 percent of farm homes had indoor running water, . . . by 1970, 93 percent of rural homes had running water." Of course, for advances in home production technology to cause a rise in female employment rates in the postwar period, we must assume that the responsibility for housework (e.g., meal preparation and cleanup, clothing maintenance, housecleaning, etc.) has disproportionately fallen on women. A 1965 study found that women spent an average of 37.8 hours per week engaged in unpaid household chores and childcare. Men, on average, spent only 10.0 hours per week involved in these tasks.[1] A later study, conducted in 1986, revealed that women were spending 31.9 hours per week on these duties, compared to 18.1 hours for men.[2] These data clearly suggest that female household members have more likely been responsible for household work over the years.

It is interesting that as the female employment rate steadily increases in Figure 3, panel a, the male employment rate progressively falls in panel c. This observation is often interpreted to mean that as the male employment rate fell, women were hired to fill the newly created vacancies. However, this reasoning sidesteps the question of why men were gradually less willing to work at particular jobs. We shall argue in this article that, in fact, the direction of causality may well run the other way. If important changes in the economic environment made it less costly for women to work in the marketplace, the resulting increase in family income may have led to a wealth effect that reduced male labor supply. Note that in the latter scenario, males are not displaced by females in jobs but would optimally choose to work less.

2. GENDER AND LABOR MARKETS: AN APPLICATION OF THE CHO-COOLEY MODEL

Although both demographic considerations and the changing role of gender have been crucial determinants of the U.S. aggregate employment rate, for simplicity we shall confine our theoretical analysis to the role of gender within a general equilibrium dynamic framework. In exploring how to think about gender within a neoclassical model, we shall address the various implications of advances in home production technology, the different factors that might underlie differences in gender earnings, and the role of discrimination in general equilibrium.

A Basic Framework without Discrimination

As we saw in the previous section, both the female and male employment rates have displayed considerable variation over the postwar period. To address this fact within the context of an artificial economy, we shall need a model

[1] Source: Multinational Comparative Time-budget Research Project—National Survey, 1965 (Institute for Social Research, University of Michigan).

[2] Source: Study of Americans' Use of Time, 1986 (Survey Research Center, University of Maryland).

that allows workers to adjust their labor supply both along the intensive and extensive margins as in Cho and Cooley (1994).

Consider a closed economy populated by a large number of households that comprise a continuum of members uniformly distributed on $[0, 2]$. Households are composed of men and women in equal proportion, working members of the current generation. Each household cares about the welfare and resources of its present as well as future descendants. As in Barro and Sala-i-Martin (1995, chapter 2), this intergenerational consideration may be modeled by assuming that the current generation maximizes utility subject to a budget constraint over an infinite horizon. Thus, we shall in effect analyze the decisions of an immortal extended family.[3] For the sake of transparency, we assume that the size of each extended family is constant over time.

Let $0 \leq h_f \leq 1$ represent the normalized number of hours supplied by a female worker to the market on any given day. (Throughout the remainder of the article, the subscripts f and m will stand for females and males respectively.) Further, we define $a_f h_f^{\gamma+1}/(\gamma+1)$ as the disutility that a female household member experiences when she provides h_f hours of work. If $0 \leq e_f \leq 1$ stands for the measure of female workers within the household, then $\left[a_f h_f^{\gamma+1}/(\gamma+1)\right] e_f$ designates total household disutility derived from female labor. Similarly, we let $\left[a_m h_m^{\gamma+1}/(\gamma+1)\right] e_m$ denote total family disutility derived from male labor. Observe that if $a_f = a_m$, then men and women have identical preferences with respect to the number of hours spent working and, consequently, should have identical labor supply schedules along the intensive margin.

Turning our attention to the extensive margin, we follow Cho and Cooley (1994) and assume that there exists a cost associated with each time a household member chooses to work. The idea is that participating in the labor force requires that real resources be spent in the replacement of household production. For example, we can think of these resources as the cost of replacing services such as child care or domestic maintenance while agents are away at work.[4] Cho and Cooley (1994) express the utility costs associated with the replacement of household services as an increasing function of the number of working household members, $b_i e_i^{\tau}/(\tau+1)$, $i = m, f$. In the context of the model presented here, the assumption that the burden of housework fell more heavily upon women for a considerable fraction of the past 50 years may be interpreted as $b_f > b_m$. To a degree, this inequality in the cost of market work can also be interpreted as the unfavorable social stigma that working women may have carried earlier in this century. Observe that from the vantage point

[3] See Barro (1974) for a formal derivation of this modeling assumption with parental altruism.

[4] Cho and Cooley (1994) show explicitly how to map household production into agents' preferences.

of the household as a whole, the total costs associated with having female members work is $b_f e_f^{\tau+1}/(\tau+1)$ and is similar for male members.

In each period, female and male workers receive wages $w_{f,t}$ and $w_{m,t}$ respectively in exchange for their labor services. Households also own capital, K_t, from which they earn interest income, $r_t K_t$, and discount the future at the rate $0 < \beta < 1$. Income is either saved in the form of capital accumulation or used to purchase consumption goods, C_t. Given the features of the model we have just described, the representative household maximizes its expected utility into the infinite future,

$$\begin{aligned} \max \mathcal{U} &= E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{\sigma}-1}{\sigma} - \sum_{i=f,m} \frac{a_i h_{i,t}^{\gamma+1} e_{i,t}}{\gamma+1} - \sum_{i=f,m} \frac{b_i e_{i,t}^{\tau+1}}{\tau+1} \right], \\ 0 &< \sigma < 1, \end{aligned} \tag{4}$$

subject to the following budget constraint,

$$C_t + K_{t+1} - (1-\delta)K_t = \sum_{i=f,m} w_{i,t} h_{i,t} e_{i,t} + r_t K_t, \tag{5}$$

$$C_t \geq 0,\ K_t \geq 0,\ K_0 > 0 \text{ given},$$

where $0 < \delta < 1$ is the depreciation rate on capital. Note that in this framework, $e_{i,t}$ carries the interpretation of the employment rate of gender i in period t. The solution to this problem yields the following first order conditions:

$$-a_i h_{i,t}^{\gamma} + C_t^{\sigma-1} w_{i,t} = 0, \qquad i = f, m, \tag{6}$$

$$-a_i \frac{h_{i,t}^{\gamma+1}}{\gamma+1} - b_i e_{i,t}^{\tau} + C_t^{\sigma-1} w_{i,t} h_{i,t} = 0, \qquad i = f, m, \tag{7}$$

$$-C_t^{\sigma-1} + \beta E_t \left\{ C_{t+1}^{\sigma-1} \left[r_{t+1} + 1 - \delta \right] \right\} = 0. \tag{8}$$

Equation (6) can be thought of as a labor supply schedule for each gender along the intensive margin. As noted earlier, when $a_f = a_m$, these labor supply schedules are identical. In fact, the number of weekly hours spent at work has been slightly lower for women than for men over the last two decades. However, the ratio of female hours to male hours has been creeping up somewhat during that period. In addition, differences between genders in weekly hours at work are nowhere near as large as differences in employment rates. In our model economy, employment rates for each gender are determined by equation (7). The first two terms of this equation denote the marginal costs of having an additional household member work in the marketplace while the third term captures its marginal benefit in utility terms. Finally, equation (8) equates the marginal benefit and the marginal cost of saving an additional unit of the consumption good.

Figure 5 Trend and Cyclical Components of the Gender Earnings Ratio

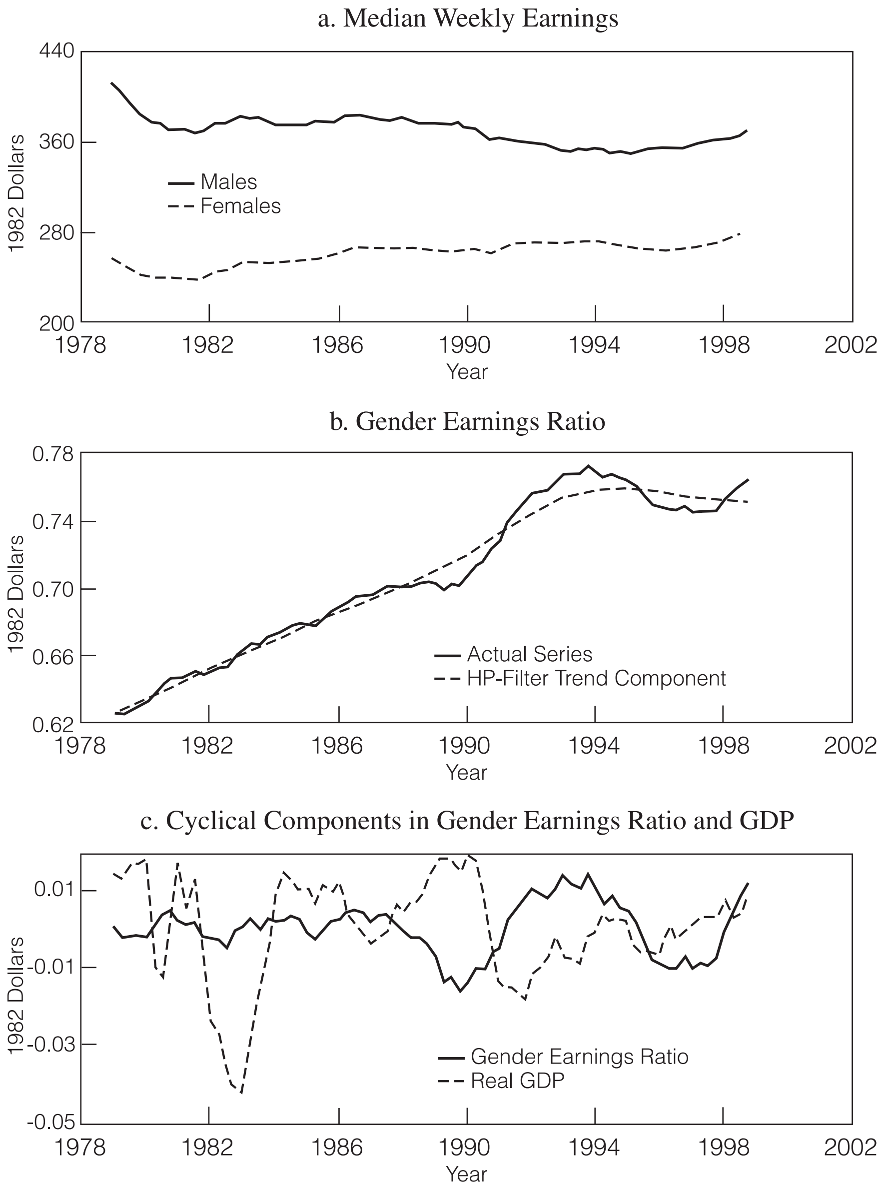

Firms in the economy produce goods, pay wages to female and male labor L_f and L_m respectively, and make rental payments on capital. Each firm has access to an identical constant-returns-to-scale production technology,

$$Y_t = A_t K_t^{1-\alpha} \left[z_f L_{f,t}^{\rho} + z_m L_{m,t}^{\rho} \right]^{\frac{\alpha}{\rho}}, \; 0 < \alpha < 1, \text{and } \rho \leq 1, \tag{9}$$

where A_t embodies shifts in total factor productivity. We interpret z_f and z_m as factors that might influence the productivity of genders differently in the market. For example, we noted earlier the substantial increase in women's education over the postwar period. In principle, therefore, both z_f and z_m should be endogenously determined and time varying. However, in the case of education, the gradual narrowing of differences in education between genders was partly due to a relaxation of social restrictions on women. It will be simplest, therefore, to take z_f and z_m as reduced form parameters. The parameter ρ captures the elasticity of substitution between male and female labor in production, which in this case is given by $1/(\rho - 1)$. Hence, male and female labor are perfect substitutes when $\rho = 1$, and one might think of this scenario as a suitable benchmark.

At each point in time, firms maximize profits and solve,

$$\max \; \Pi_t = Y_t - r_t K_t - w_{f,t} L_{f,t} - w_{m,t} L_{m,t}, \tag{10}$$

which gives the following first order conditions,

$$(1-\alpha) A_t K_t^{-\alpha} \left[z_f L_{f,t}^{\rho} + z_m L_{m,t}^{\rho} \right]^{\frac{\alpha}{\rho}} - r_t = 0, \tag{11}$$

$$\alpha A_t K_t^{1-\alpha} \left[z_f L_{f,t}^{\rho} + z_m L_{m,t}^{\rho} \right]^{\frac{\alpha-\rho}{\rho}} z_f L_{f,t}^{\rho-1} - w_{f,t} = 0, \tag{12}$$

$$\alpha A_t K_t^{1-\alpha} \left[z_f L_{f,t}^{\rho} + z_m L_{m,t}^{\rho} \right]^{\frac{\alpha-\rho}{\rho}} z_m L_{m,t}^{\rho-1} - w_{m,t} = 0. \tag{13}$$

Combining equations (12) and (13) immediately yields an expression for the gender earnings ratio,

$$\frac{w_{f,t}}{w_{m,t}} = \frac{z_f}{z_m} \left(\frac{L_{f,t}}{L_{m,t}} \right)^{\rho-1}. \tag{14}$$

Thus, when male and female workers are perfect substitutes into production, only factors that affect differences in gender productivity affect the gender earnings ratio. Of course, at this stage, equation (14) abstracts from discrimination. Figure 5, panel a, shows that over the last 20 years, male median weekly earnings have consistently exceeded those of females (i.e., $w_{f,t}/w_{m,t} < 1$). However, panel b shows an increasing trend in the gender earnings ratio during that period. Interestingly, this reduction in the male/female earnings gap appears to have slowed down as of 1994.

An equilibrium for this economy consists of households' optimality conditions (6) through (8), firms' optimality conditions (11) through (13), and goods and labor market clearing conditions,

$$C_t + K_{t+1} - (1-\delta)K_t = Y_t \tag{15}$$

and

$$L_{i,t} = e_{i,t}h_{i,t} \qquad i = f, m. \tag{16}$$

3. NUMERICAL EXAMPLES

In this section, we study several numerical examples in order to gain insight into the dynamic general equilibrium effects of different changes in the economic environment. We investigate the effects of a reduction in the female cost of market work. We also examine how a loosening of discriminatory hiring practices against women affects the overall economy. The idea of women facing reduced employment opportunities in the workplace may seem somewhat outdated by today's standards. Nevertheless, it remains that this notion may have substantially contributed to the gender wage gap over the last 50 years.

To explore these issues in greater detail, we must first assign values to the exogenous parameters of the model we have just presented. The parameters a_f and a_m are set to 19 and 16 respectively, implying that $h_f = 0.313$ while $h_m = 0.374$. Assuming that agents can work a maximum of 16 hours a day, 7 days a week, these values translate into 35 weekly hours spent working for females and 42 hours for males. We set τ_f and τ_m to 1.31 and 1.10 respectively to generate employment rates of 0.58 for women and 0.72 for men. This calibration assumes that the present employment rates depicted in Figure 3 are approximately at their steady state. We normalize the Total Factor Productivity parameter to 1 and assume that male and female workers are perfect substitutes in production, $\rho = 1$. We normalize z_f to 1 and set z_m to 1.20 so as to obtain a gender earnings ratio of 0.83. While Figure 5, panel b, suggests that the gender earnings ratio is currently 0.76, Blau and Kahn (1997) find that 43 percent of the gap cannot be explained by human capital differences with the implication that much of the unexplained portion results from discriminatory practices.[5] On the other hand, Kim and Polacheck (1994) have conducted empirical research suggesting that the "unexplained" portion can be reduced 50 percent when estimates allow for unobservable individual-specific effects, which the authors think of as individual differences in motivation. We choose an intermediate level and model 30 percent of the gender earnings ratio as

[5] Blau and Kahn consider 1979 and 1988 PSID data. They find that the earnings ratio increased from 0.62 to 0.72 over that period. Adjusting for human capital variables, the ratios are 0.72 and 0.80, respectively. Adjusting for human capital and industry, occupation and collective bargaining status, the ratios are 0.78 and 0.88. The unexplained portion of the pay gap (in 1988) is therefore given by $\frac{1-0.88}{1-.72} = .43$.

Table 1a Benchmark Steady Rate

Aggregate Characteristics		
Consumption	0.469	
Rate of Interest	0.070	
Investment:Output Ratio	0.255	
Gender Earnings Ratio	0.833	
Labor Force Gender Characteristics	**Male**	**Female**
Weekly Hours	42	35
Employment Rate	0.720	0.577

Table 1b Steady State with Reduced Female Employment Opportunities

Aggregate Characteristics		
Consumption	0.452	
Rate of Interest	0.070	
Investment:Output Ratio	0.255	
Gender Earnings Ratio	0.760	
Labor Force Gender Characteristics	**Male**	**Female**
Weekly Hours	42	33
Employment Rate	0.73	0.523

resulting from discrimination. In other words, the women's to men's wage ratio is 0.76 and 70 percent of that gap is attributable to factors other than discrimination. This means that absent discrimination, this ratio would be 0.83, and that is how the no-discrimination model is calibrated. We shall think of a period as a year in the numerical examples so that β is set to 0.98. All other parameters are chosen symmetrically for men and women and set to the values in Cho and Cooley (1994). In particular, we have $b_m = b_f = 0.8$, $\sigma = 2$, $\gamma = 0.8$, and $\alpha = 0.64$. Table 1a summarizes key aspects of the model steady state that arise from the calibration presented here.

The first experiment we carry out considers the effects of a protracted fall in the female labor market participation cost. As advances in household production—as well as changing attitudes—have made it progressively easier for women to join the labor force, we wish to analyze their general equilibrium implications for other variables. In the second experiment, we introduce reduced employment opportunities for women as a way of modeling discrimination. We explain how such reduced opportunities can generate a gender wage gap and analyze the effects of a permanent change in total factor productivity in that environment. As in the data, changes in total factor productivity will generate a countercyclical gender earnings ratio. Finally, we shall analyze the effects of a gradual reduction in discrimination against female labor.

The Dynamic Effects of a Reduction in Female Labor Market Participation Costs

Figure 6 illustrates the effects of a gradual fall in the cost of market work for women. In Figure 6, panel a, we show that this is modeled as a permanent 10 percent fall in the value of b_f over 16 years. This gradual reduction in the dynamics of $b_{f,t}$ can be modeled as

$$\ln b_{f,t} = \sum_{j=1}^{16} \omega_j u_{t-j}, \tag{17}$$

where

$$u_t = u_{t-1} + \varepsilon_t \text{ and } \sum_{j=1}^{16} \omega_j = 1. \tag{18}$$

Since $b_{f,t}$ falls gradually, we should expect that both male and female employment rates should also respond incrementally over time as suggested in Figures 1 and 3. Observe in Figure 6, panel d, that the reduction in women's labor market costs directly implies a 4 percent permanent rise in the female employment rate. As we had conjectured, therefore, historical changes in the female employment rate are consistent with technological advances in home production. In addition, some interesting general equilibrium effects emerge. Because the utility cost of market work for women is permanently lower, aggregate consumption eventually rises to a higher steady state in Figure 6, panel b. This rise in aggregate consumption translates into a wealth effect that actually causes a fall in the male employment rate (and male work hours) in Figure 6, panel c. Note that equations (6) and (7), which characterize labor supply for each gender, both depend on aggregate consumption. Of course, wages also adjust downward in this experiment, leading to a substitution effect, which reinforces the reduction in male employment. In other words, some degree of crowding out does take place as women enter the labor market. On the whole, in a manner consistent with Figure 3, panels a and c, advances in home production technology lead not only to a rise in the female employment rate but also a fall in its male counterpart. However, in the final steady state, the magnitude of the change in the male employment rate is relatively small compared to that in the female employment rate.

It is also worth noting that a reduction in the labor participation cost of women examined here implies a temporary rise in the rate of interest. As is typical of neoclassical frameworks, changes in the rate of interest mimic changes in the growth rate of consumption (see equation (8)). Finally, observe

Figure 6 The Effects of a Reduction in the Female Labor Participation Cost

that this particular example is also useful in illustrating the forward-looking behavior of household members. Since advances in home production technology are gradual, reductions in women's cost of market work today signal further reductions in the future. Anticipating these future reductions, male workers cut back their labor supply contemporaneously despite the fact that the initial decline in market costs is quite small. Given that the capital stock is predetermined when the shock occurs, production also falls contemporaneously. It follows that consumption falls on impact as shown in Figure 6, panel b.

The General Equilibrium Impact of Reduced Employment Opportunities for Women

In the numerical experiment we have just carried out, the gender earnings ratio remained unaffected and constant as suggested by equation (14). However, this equation only captures one notion of the earnings gap based on differing gender productivity. We argued earlier, for instance, that the human capital embodied in female workers had risen substantially over the past five decades because of increased education. We now show how demand-side discrimination in the form of reduced employment opportunities for women can contribute to lowering female earnings.

Consider the model presented in the previous section. With the household side unchanged, male and female labor supply continue to obey equations (6) and (7). However, suppose that firms are unwilling to manage more than $\widetilde{L}_f > 0$ units of female labor. As in Phelps (1972), we can imagine that this labor demand constraint stems from the perception that female labor is less dependable or productive than male labor. According to Goldin (1990), firms have often viewed gender as a sign of shorter expected job tenure, leading to job segregation and limited opportunities for women. Firms then maximize profits in (10) subject to

$$L_{f,t} \leq \widetilde{L}_f. \tag{19}$$

Letting $\widetilde{w}_{f,t}$ denote the new female wage that emerges from this constrained maximization problem, we obtain

$$\frac{\widetilde{w}_{f,t}}{w_{m,t}} = \frac{z_f}{z_m} - \frac{\phi_t}{w_{m,t}} \leq \frac{w_{f,t}}{w_{m,t}}, \tag{20}$$

where $\phi_t \geq 0$ is the Lagrange multiplier associated with the constraint (19). Note that in the previous section, the market equilibrium value of female labor was increasing in female productivity, z^f. Thus, the more productive female labor, the more likely equation (19) is to bind. When this is the case, $\phi_t > 0$ and a lower gender earnings ratio emerges.

Table 1b describes the steady state that obtains when $\widetilde{L}_f$ is calibrated to generate a gender earnings ratio of 0.76 as suggested by the most recent ratios

Table 2a Gender Earnings Ratio and GDP

U.S. Economy

$\text{Corr}(Y_t, w_{f,\, t+k}/w_{m,\, t+k})$ for $k =$

−3	−2	−1	0	1	2	3
−0.050	−0.107	−0.169	−0.234	−0.310	−0.330	−0.332

Table 2b Gender Earnings Ratio and GDP

Model Economy

$\text{Corr}(Y_t, w_{f,\, t+k}/w_{m,\, t+k})$ for $k =$

−3	−2	−1	0	1	2	3
−0.354	−0.550	−0.762	−0.985	−0.661	−0.395	−0.165
[0.14]	[0.11]	[0.07]	[0.00]	[0.07]	[0.11]	[0.13]

in Figure 5, panel b. Aside from the fall in both the female employment rate and female work hours that naturally follows, we wish to stress the following two points. First, the male employment rate rises somewhat to compensate for the drop in female labor input; however, this increase is less than one-for-one since $z_m > z_f$. Second, and more important, steady state household consumption falls by 3.62 percent. This decline in consumption reflects the efficiency loss that emerges when 30 percent of the gender wage gap is due to discriminatory hiring practices against women.

Figure 5, panel c, suggests a negative correlation between the cyclical components of the gender earnings ratio and real GDP. Table 2a confirms this negative correlation at different leads and lags. An argument often cited to explain this negative correlation is that in periods of boom, women are more likely to be new entrants or re-entrants in the job market since the male participation rate is already high. Relative to women who are already part of the labor force, however, the marginal female worker is likely to be less productive. Hence, one expects the median female wage to decrease in times of economic expansion. An implicit assumption here is that there exists some heterogeneity among female workers.

Aside from the compositional effect we have just described, there may be other factors that contribute to a countercyclical gender earnings ratio. In fact, in the framework with reduced female employment opportunities presented in this section, the gender earnings ratio will move in a direction opposite to that of output when the economy is mainly driven by shifts in total factor productivity. Figure 7 shows the impulse responses of output and the gender

Figure 7 The Effects of an Increase in Total Factor Productivity

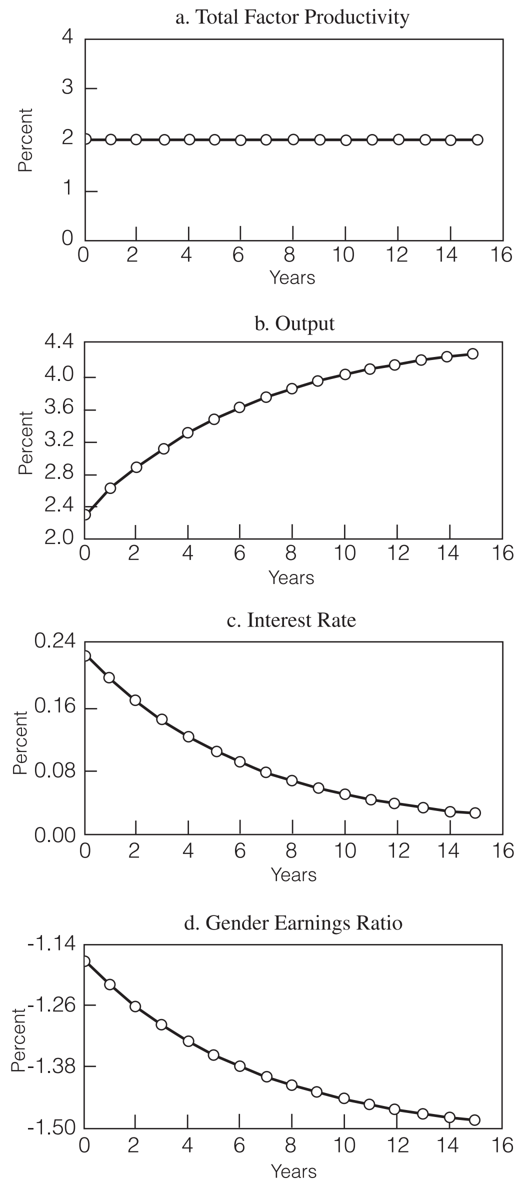

earnings ratio to a 2 percent permanent increase in total factor productivity. As expected, output rises on impact and increases monotonically to a higher steady state level. In contrast, the gender wage ratio falls both contemporaneously and in the long run. Because firms are unwilling to hire more than a given level of female labor input, the female wage exhibits inertia despite the rise in productivity. The male wage, on the other hand, unambiguously increases as the demand for male labor shifts out in response to the productivity shock. Consequently, the gender earnings ratio falls in periods of boom.

In Table 2b, we present the cross-correlations of output and the gender earnings ratio obtained at different leads and lags when the productivity process is calibrated as

$$\ln A_t = \rho_A \ln A_{t-1} + \varepsilon_{A,t}, \tag{21}$$

where $\rho_A = 0.95$ and $\varepsilon_{A,t}$ is a random variable with mean zero and standard deviation of 0.01.[6] The model statistics are the mean values calculated from 200 simulations of samples with 80 observations each, the size of the sample for the data in Figure 3. In square brackets are the standard deviations of the sample statistics. As suggested by the impulse responses in Figure 7, output and the gender earnings ratio are negatively correlated at all leads and lags, much more so in fact than in the data. While our model misses some important dimensions, such as heterogeneity among workers that would give rise to the type of compositional effect discussed above, our result also suggests that there may be other key sources of shocks aside from shifts in total factor productivity.

Finally, Figure 8 shows the dynamic effects of a gradual 2 percent improvement in female employment opportunities (i.e., $\Delta\widetilde{L}_f > 0$). As the constraint on female employment becomes less binding, the economy becomes more efficient and aggregate consumption rises in Figure 8, panel a. At the same time, the female wage rises closer to its unconstrained equilibrium so that the gender earnings ratio increases in panel b. Note in Figure 8 that as the female employment rate rises in response to a looser employment constraint, the male employment rate correspondingly falls. This is not only because the rise in female employment is making work opportunities scarcer for male labor; in this case, the fall in male employment is partly a reflection of the wealth effect induced by the increase in aggregate consumption. Thus, as the economy becomes more efficient, men choose to reduce how much they work. Further, this exercise suggests that the rise in the female employment rate in Figure 3, panel a, and the simultaneous decline in the male employment rate in panel c, are consistent with a continuing relaxation of discriminatory hiring practices against female labor. A loosening of the constraint on female employment

[6] This particular process for total factor productivity is standard in the real business cycle literature.

Figure 8 The Effects of Relaxing the Female Labor Constraints

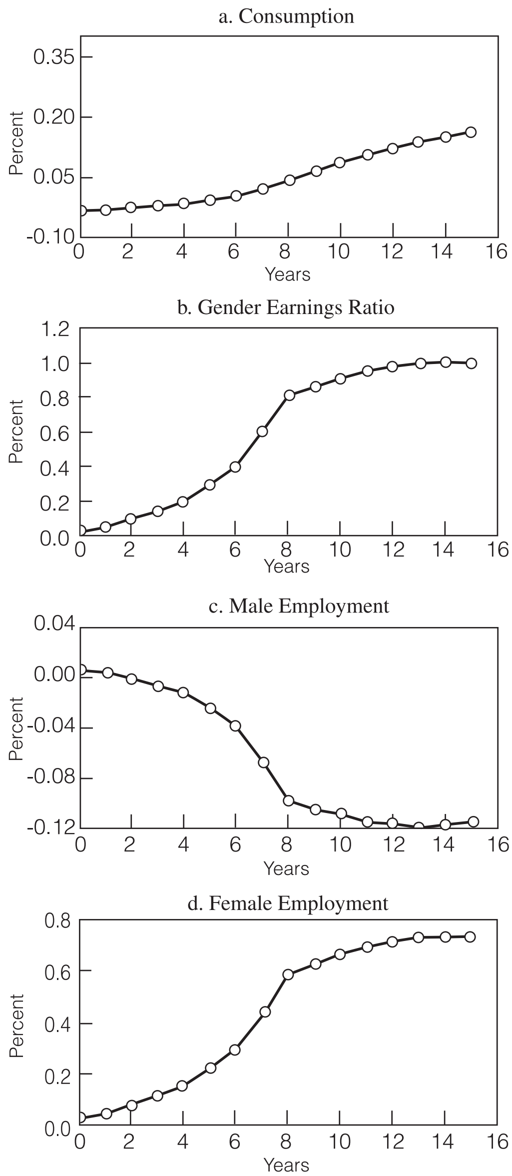

opportunities has the same effects on gender employment rates as a decrease in women's costs of market work. Both of these changes in the economic environment are consistent with the behavior of employment rates by gender since the end of World War II.

4. CONCLUDING REMARKS

We have documented that U.S. employment rates have changed considerably over the postwar period. The data suggest that both demographic and gender specific factors have been important forces underlying the evolution of these rates. We have focused specifically on the role of gender heterogeneity in determining aggregate employment rates, and we have also developed the implications of gender differences for the overall economy.

Within a stochastic general equilibrium framework, we have modeled the effects of a reduction in female employment participation costs, as well as a reduction in employer-based gender discrimination. Reduced employment costs, while giving rise to higher levels of female employment, also generate lower rates of male employment through both a wealth and a substitution effect at the household level. Reduced participation costs do not, however, explain the historical behavior of the gender earnings gap. Employer based discrimination is then introduced as a constraint on the measure of female labor that firms are willing to employ. In the presence of such discrimination, a gender earnings gap emerges in excess of what can be explained by relative differences in productivity. This model-generated pay gap is countercyclical in nature, rising in periods of economic contraction and diminishing when the economy booms, thus approximating what is seen in U.S. data. When the constraint on female labor is relaxed over time, the model predicts gradually rising levels of female employment, which endogenously lead to lower rates of male employment and a narrowing gender wage gap. The latter predictions have also been stylized features of U.S. labor market experience.

REFERENCES

Barro, Robert J., and Xavier Sala-i-Martin. *Economic Growth*. New York: McGraw-Hill, Inc., 1995.

Becker, Gary S. *The Economics of Discrimination*. Chicago: University of Chicago Press, 1957.

Beller, Andrea H. "The Impact of Equal Opportunity Policy on Sex Differentials in Earnings and Occupations," *The American Economic Review*, vol. 72 (May 1982), pp. 171–75.

Blau, Francine D., and Lawrence M. Kahn. "Swimming Upstream: Trends in the Gender Wage Differential in the 1980s," *Journal of Labor Economics*, vol. 15 (January 1997), pp. 1–42.

Cho, Jang-Ok, and Thomas F. Cooley. "Employment and Hours Over the Business Cycle," *Journal of Economic Dynamics and Control,* vol. 18 (March 1994), pp. 411–32.

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

Jacobsen, Joyce P. *The Economics of Gender*. Cambridge: Blackwell Publishers, 1994.

Kim, Moon-Kak, and Solomon W. Polachek. "Panel Estimates of the Gender Earnings Gap," *Journal of Econometrics*, vol. 61 (March 1994), pp. 23–42.

Phelps, Edmund S. "The Statistical Theory of Racism and Sexism," *The American Economic Review*, vol. 62 (September 1972), pp. 659–61.