Home Production, Market Production and the Gender Wage Gap: Incentives and Expectations*

Stefania Albanesi[†], Claudia Olivetti[‡]

December 2005 (First draft: April 2005)

Abstract

This paper explores the hypothesis that gender wage differentials arise endogenously from the interaction between the intra-household allocation of labor and the contractual relation between firms and workers in the presence of private information on workers' labor market attachment.

In our model, households efficiently choose the contribution of each spouse to home production. Workers with high home hours are less attached to market work. Individual home hours and effort applied to market work are private information. Firms offer incentive compatible labor contracts that imply an inverse relation between earnings and home hours. Absent ex ante gender differences in productivities, if firms believe women to be less attached to market work, they will offer them contracts with lower earnings and lower hours than men. If firms believe that labor market attachment is the same across genders, they will offer the same contract to male and female workers. Spouses' optimal allocation of home hours will respond to firms' beliefs, thus generating the potential for statistical discrimination by gender. If women have a comparative advantage in home production, then the interaction between the labor market and household decisions amplifies the resulting gender differences in earnings.

The large variation in gender earnings differentials across industries and occupations observed in the data motivates the central role of incentive problems in the determination of gender differences in labor market outcomes. We document this variation using Census and PSID data and we relate it to the severity of incentive problem across industries and occupations.

*We wish to thank Mark Bills, Peter Gottschalk, Luigi Guiso, Kevin Lang, Marjorie McElroy, Bentley MacLeod, Stephanie Schmitt-Grohe, B. Ravikumar, Bernard Salanie, Pedro Teles, Gianluca Violante, and seminar participants at Brown University, Columbia University, Duke University, the Federal Reserve Bank of Minneapolis, NBER SI, SITE, ESSIM, Northwestern University, the University of Pittsburgh, University of Rochester, Queen's University, University of Southern California, and SUNY Albany for very helpful comments.

[†]Columbia University. Corresponding author. Address: 1022 International Affairs Building, 420 West 118th Street, New York, NY 10027. Tel. 212–854-4478. E-mail: sa2310@columbia.edu.

[‡]Boston University

1 Introduction

The purpose of this paper is to study the emergence and persistence of gender differences in wages and in the division of labor within the household. We explore the hypothesis that gender wage differentials arise from the interaction between the intra-household allocation of labor and the contractual relation between firms and workers in the presence of private information on workers' labor market attachment.[1]

Our theoretical analysis is based on a model in which the intra-household allocation is efficient. We model the households according to Chiappori's (1988, 1997) "collective labor supply" model.[2] In particular, households are assumed to optimally choose the contribution of each spouse to home production. Workers with high home hours are less attached to market work. Individual home hours and effort applied to market work are private information. Firms, in an extension of the framework developed by Holmstrom and Milgrom (1991), face both adverse selection and moral hazard in contracting with workers. They offer incentive compatible labor contracts that are constrained-efficient. Optimal contracts imply earnings that are inversely related to home hours.

If firms believe women to be less attached to market work than men, they will offer women labor contracts with lower earnings and lower hours even in the absence of gender differences in productivity. This implies that it is efficient for wives to allocate more time to home production. Hence, women will be less attached to market work and firms' expectations will be confirmed. If firms believe that labor market attachment is the same across genders, they will offer the same labor contracts to male and female workers. As a consequence, spouses will face the same earning opportunities, and the efficient intra-household allocation of labor will not be related to gender. Hence, absent ex-ante differences in productivity for male and female workers, there are two types of equilibria. One in which there is a systematic gender differential in earnings, accompanied by a gendered intra-household division of labor; and one in which earnings and home hours are not systematically related to gender. It is statistical discrimination that determines gender differentials in the first type of equilibria. Given that firms use gender as a screening device in this type of equilibrium, discrimination actually reduces the incentive problem for firms, eliminating adverse selection. This property of the model implies that equilibria with gender discrimination will be hard to break, which is consistent with the persistence of the gender earnings gap and the intra-household division of labor observed in the data.[3]

If, instead, we allow for higher relative productivity of women in home production, female

[1] There is a vast empirical literature that studies the gender wage gap and its evolution over time. Altonji and Blank (1999) provide an extensive review of this literature for the post-war period. Goldin (1990) presents an extensive historical analysis for the United States.

[2] This paradigm does not focus on a particular model of spousal interaction, rather it merely restricts household decisions to be Pareto efficient. This framework is consistent with a variety of "household bargaining" models, as in McElroy and Horney (1981) and Manser and Brown (1980).

[3] O'Neill (2003) shows that there is still a 10% differential in female and male wages in the U.S. in 2000 that remains unexplained by gender differences in schooling, actual experience and job characteristics. Moreover, PSID data for the period 1976-2001 show that husbands' home hours are roughly one third of wives' and that this difference is stable over time.

workers will be less attached to market work if the relative productivity difference is large enough. In this case, the interaction between the optimal intra-household allocation process and the incentive problem in the labor market serves as an amplification mechanism, so that the gender differential in earnings is larger than the difference in relative productivity.

Our model bridges three literatures: the literature on the sexual division of labor in the Beckerian tradition; the one on statistical discrimination, as in Coate and Loury (1993) and Lundberg and Starz (1983); and finally the literature on incentive contracts and job design, as in Holmstrom and Milgrom (1991). As argued by Becker (1985), what distinguishes gender and racial discrimination is that the feedback on the optimal intra-household division of labor generates a larger impact on earnings. The centerpiece of our model is to identify the source of statistical discrimination with the incentive problem in the interaction between firms and workers, given that labor market attachment, which is inversely related to hours devoted to home production, is private information. The central role of informational problems for gender differences in labor market outcomes is motivated by the large variation in gender earnings differentials across industries and occupations that we find in the data.

We use a variety of data sources to support several important predictions of our model. First, our model predicts that gender earning differentials should be higher for married workers in industries/occupations in which the incentive problem is more severe. We use Census data for year 2000 to study gender earnings differentials across industries for four broad occupational categories: high management, low management, sales and production. We argue that incentive problems are most stringent in these occupations. High level managers have a wide range of responsibilities, hence, the uncertainty associated with their performance, given their effort should be greater than for low level management. Similarly, sales volumes depend to a large degree on variables that are not directly related to sales personnel's effort. These considerations are less important for production workers. We find that, for the sample of married workers, gender differentials in earnings are greatest in high management and sales occupations. By contrast, for single workers, there is no clear pattern and the gender earnings ratio is highest in high management occupations in most industries. Second, according to our model we should observe a negative relation between the male/female difference in the fraction of incentive pay and the female/male earnings ratio. Since the Census does not include information on the structure of earnings, we use PSID data from the late 1990s to document this fact. We find a negative and significant correlation between the two ratios. Moreover, we also find that there is a strong negative correlation between the fraction of incentive pay and the female/male earnings ratio across industries/occupations. This confirms our Census findings, since incentive pay is used more in those industries/occupations where the incentive problem is more severe, as discussed in MacLeod and Parent (2003).

Our model also predicts a negative correlation between the female/male ratio of home hours and the female/male earnings ratio for married workers. We document this correlation across industries/occupations using the PSID and find that is negative and significant for married workers, much smaller for single workers. Finally, we show that there is no relation between the female/male earnings ratio or the fraction of incentive pay and the percentage of female workers across industries and occupations. This suggests that occupational sorting by gender based on comparative advantage or differences in preferences, such a risk aversion, cannot fully account for the variation in gender earnings differentials across industries and occupations.

Our model highlights the importance of incentives and differences in the pay structure in determining gender differences in earnings. This prediction also resonates with current debates on gender discrimination in personnel policy. For example, in June 2004 a federal judge ruled in favor of class-action status for the Dukes vs Wal-Mart gender discrimination lawsuit. The ruling was based on extensive evidence presented by the plaintiffs, Drogin (2003), showing that women working at Wal-Mart stores face pay disparities in most job categories, and take longer to enter management positions.[4] Finally, it is also interesting to note how expectations of a gender wage gap characterize both male and female workers. As documented by Babcock and Laschever (2003): "Women report salary expectations between 3 and 32 percent lower than those of men for the same jobs; men expect to earn 13 percent more than women during their first year of full-time work and 32 percent more at their career peaks."

Our paper is organized as follows. Section 2 presents the model and discusses the results of numerical simulations. Section 3 reports evidence supporting the model's predictions. Finally, Section 4 concludes.

2 The Model

The economy is populated by a continuum of adult agents, ex ante identical except for gender, and a continuum of identical firms. The agents are equally divided by gender, they are all married and belong to a household. Households are made up of two agents of different gender. There are two types of goods in this economy- a private market good and public home goods. Households combine market goods and home hours of each spouse to produce the public home good, which is household specific. Each agent is employed by a firm. Firms produce the market good using the agents' labor as the only input. Hence, there are two components of our model economy, the labor market and households. On the labor market, firms and individual agents negotiate the terms of labor contracts. Households efficiently choose the allocation of home hours across spouses. We assume that individual utility is increasing in consumption of the market good and decreasing in the number of hours worked at home and on the labor market, and in the effort applied to market work. Following Becker (1985), we posit that an agent's utility cost of market hours and effort is *increasing* in home hours. We also assume that agents' home hours and effort are not observed by firms. Then, firms face *adverse selection* and *moral hazard* when contracting with workers. Firms will offer incentive compatible labor contracts that maximize the surplus from the employment relationship subject to incentive compatibility constraints stemming from the private information. Individual agents' labor market outcomes will depend on their home hours, which are chosen at the household level. On the other hand, an household's efficient choice of home hours will depend on the spouses relative earnings, which are determined on the labor market. Hence, there is a feedback from

[4] Discrimination lawsuits based on analogous complaints where filed by a team of women brokers at Merrill Lynch and by women researchers working at Rand corporation during the summer of 2004. See The New York Times, August 22, 2004 and The New York Times, September 5, 2004, respectively.

household decisions to labor market outcomes.

We now describe the labor contracts and the household decision problem in detail, and present our definition of equilibrium.

2.1 Labor Contracts

On the labor market, each firm hires agents to produce output. The output of one agent is related to her market hours and effort, according to:

$$y = f(n, e) + \omega,$$

The function $f(n, e)$ denotes expected output, where f is strictly increasing in both hours, n, and effort, e, twice continuously differentiable and weakly concave. The random variable ω is distributed normally with zero mean and variance $\Sigma^2 > 0$.

Each agent has a utility function:

$$U(c, h, n, e) = -\exp(-\sigma [c - v(h, n, e)]), \tag{1}$$

where c is individual consumption of the market good, h denotes home hours and n denotes market hours, and e denotes effort applied to market work. We adopt a CARA specification, where the coefficient of relative risk aversion, σ, is strictly greater than zero, and $v(\cdot)$ denotes the disutility of market and home work. We assume that v is increasing in all its arguments, twice continuously differentiable and that:

$$v_{hn} \geq 0, \;\; v_{he} > 0. \tag{2}$$

Hence, the marginal cost of market hours and effort is increasing in home hours.

Firms choose labor contracts to maximize the surplus from the employment relationship. We assume that effort, e, and home hours, h, are *not observed* by firms, while output, y, and market hours, n, are *observable*. Hence, labor contracts will be constrained-efficient, since firms will be subject to incentive compatibility constraints. Since home hours do not influence agents' output directly, home hours correspond to an agent's *type* from the standpoint of firms. Labor contracts will specify an earnings function $w(y)$, market hours and effort to be implemented for type of agent, h, in the population. The dependence of earnings on output is required to implement effort, given that it is private information. Moreover, since home hours are also unobserved, the optimal menu of contracts will depend on the firms' belief over the distribution of home hours, and can be represented by the mapping, $\mathcal{C}(\pi(h)) = \{w(y), n, e\}(h)$, where $\pi(h)$ represents such distribution.

We assume:

$$f_{ne} \geq 0, \; v_{ne} \leq 0. \tag{3}$$

This assumption, jointly with (2) ensures that the analogue of a single crossing condition holds for this model. This ensures that, given that contracts are incentive compatible, agents with home hours h will self-select into the appropriate contract in the menu implied by $\mathcal{C}(\pi(h))$.

It is important to note that gender is observable, so firms can offer different contracts to female and male workers. However, since the contract space is *unrestricted*, firms will find it optimal to do so *if and only if* they believe that the distribution of home hours differs across genders.

To elucidate the role of our informational assumptions in the determination of labor market outcomes, we derive the properties of constrained-efficient labor contracts when home hours are observable first, and then consider the case in which home hours are also private information. For simplicity, we assume that effort and home hours can only take on two values, $e \in \{0,1\}$ and $h \in \{h_L, h_H\}$ with $h_L < h_H$, respectively. Firms take the values h_L and h_H as given, but they will determined in equilibrium from the optimal behavior of households.

If firms observe home hours but effort is not observable, firms face a moral hazard problem. Firms will choose labor contracts to solve:

$$\max_{\{w(y),n,e\},n\in N(h),e\in\{0,1\}} S\left(n,e;h\right) \qquad \text{(Problem 1)}$$

subject to

$$E\left[u\left(n,1\right)\right] \geq E\left[u\left(n,0\right)\right], \tag{4}$$

where the objective function is the expected surplus from the employment relationship, $N(h)$ is a feasible range for market hours, and (4) is the incentive compatibility constraint. As shown in Holmstrom and Milgrom (1991), CARA utility implies that, without loss of generality, we can restrict attention to earnings functions of the form: $w(y) = \bar{w} + \tilde{w}y$. We refer to $\bar{w}$ and $\tilde{w}y$ as salary and incentive pay, respectively. This implies that under CARA, the expected surplus from the employment relationship can be written as:

$$S\left(n,e;h\right) = f\left(n,e\right) - v\left(h,n,e\right) - \sigma\Sigma^2\left(\tilde{w}\right)^2/2. \tag{5}$$

The first term is expected output, the second term is the utility cost of working, given home hours h. The last term stems from the need to provide incentives by making earnings depend on output, y. This implies that earnings are stochastic and reduces the surplus from the employment relationship, since workers are risk averse. Given the CARA assumption on preferences, the incentive compatibility constraint simplifies to:

$$\tilde{w}\left[f\left(n,1\right) - f\left(n,0\right)\right] - v\left(h,n,1\right) + v\left(h,n,0\right) \geq 0. \tag{6}$$

To implement $e = 1$, firms must set $\tilde{w} > 0$, which implies that the incentive compatibility constraint will be binding since $\Sigma^2\sigma > 0$.[5] Then:

$$\tilde{w} = \frac{v\left(h,n,1\right) - v\left(h,n,0\right)}{f\left(n,1\right) - f\left(n,0\right)} > 0. \tag{7}$$

We impose a zero profit condition on firms, which implies $\bar{w} = y\left(1-\tilde{w}\right)$ and $w = y$.

[5] If n and e are perfect substitutes for the worker, the firm cannot influence and agent's choice of e. We exclude this case.

The first order necessary conditions for Problem 1, which are derived in the appendix, imply that the optimal value of n is decreasing in h for given e, from (2). The optimal choice of the effort depends on the cost of implementing high effort for given h. The multiplier on the incentive compatibility constraint, μ, can be interpreted as a measure of this cost. As shown in the appendix:

$$\mu = \frac{\sigma \Sigma \tilde{w}}{f(n,1) - f(n,0)}. \tag{8}$$

This multiplier is increasing in risk aversion, σ, and in the parameter Σ, which represents the variance of an agent's output for given market hours and effort. High values of Σ make it harder for firms to provide incentives for high effort. The value of μ is also increasing in incentive pay $\tilde{w}$. As shown in Proposition 3 in the appendix, under assumption (3), $\tilde{w}$ is increasing in h. This means that workers with higher home hours are less responsive to incentives, so that the optimal level of effort is decreasing in home hours. Hence, total earnings, w, will be decreasing in home hours.

If both home hours and effort are unobserved, this introduces additional constraints on the optimal contracts, which we refer to as adverse selection incentive compatibility constraints. These constraints imply that workers will self-select the contract on the menu appropriate to their level of home hours. Adverse selection implies that workers can extract an informational rent T_j, $j = L, H$, which reduces the surplus generated from the employment relation and may reduce the level of effort that can be implemented.

The pattern of binding adverse selection incentive compatibility constraints depends on the properties of the functions f and v. We assume that $f(n,1) - \sigma\Sigma > 0$ for all $n \in N(h_j)$, $j = L, H$ and we impose $\tilde{w}_j \leq 1$[6]. This ensures that a worker's utility is increasing in $\tilde{w}_j$ for each type of worker. Under (3), $\tilde{w}$ is increasing in h by proposition 3. Thus, under these conditions, the adverse selection incentive compatibility constraint for workers with high home hours will not be binding. We focus on this case here.

Constrained-efficient contracts solve the following problem:

$$\max_{\{\bar{w}_j, \tilde{w}_j, n_j \in N(h_j), e_j \in \{0,1\}\}_{j=L,H}} \sum_{j=L,H} \pi(h_j) \left[f(n_j, e_j) - v(h_j, n_j, e_j) - \sigma \Sigma^2 \frac{(\tilde{w}_j)^2}{2} - T_j \right] \tag{Problem 2}$$

s.t.

$$\tilde{w}_j \left[f(n_j, 1) - f(n_j, 0) \right] - v(h_j, n_j, 1) + v(h_j, n_j, 0) \geq 0, \text{ for } j = L, H, \tag{9}$$

$$T_L + \tilde{w}_L f(n_L, e_L) - v(h_L, n_L, e_L) - \sigma \Sigma^2 \frac{(\tilde{w}_L)^2}{2} \geq \max_{e=0,1} \{\tilde{w}_H f(n_H, e) - v(h_L, n_H, e)\} - \sigma \Sigma^2 \frac{(\tilde{w}_H)^2}{2}. \tag{10}$$

[6]The condition $\tilde{w}_j < 1$ for $j = H, L$ does not ensure that the firms make non-negative profits ex ante. However, it does imply that the contract is renegotiation proof. That is the firm does not have an incentive to not pay the worker when output is realized, since ex post profits are guaranteed to be positive.

We do not impose this constraint directly on the problem. We solve a relaxed contracting problem without this constraint and verify ex post that it is not violated.

Here, (10) is the adverse selection incentive compatibility constraint for workers with low home hours, hence $T_H = 0$. The moral hazard incentive compatibility constraints (9) are the same is in Problem 1.

The additional loss of efficiency determined by adverse selection implies that, even if high effort is implemented for workers with high and low home hours when home hours are known, high effort might not be optimal for either or both types of workers in Problem 2.

The labor contracting environment described above elegantly embeds elements of job design and of optimal compensation policy for a wide class of occupations. The assumption that workers' output depends on both effort and market hours captures the fact that most jobs comprise a variety of tasks, that might vary in their degree of observability from the standpoint of an employer, as discussed in Holmstrom and Milgrom (1991). The optimal combination of tasks into jobs will depend on this relative observability, as well as on the relative disutility associated with the tasks for the worker. Hence, a menu of contracts in which one specifies high effort and one specifies low effort can be interpreted as two different jobs or positions within a firm.

The incentive pay component in the optimal earnings schedule is consistent with a variety of widely used compensation schemes. For example, for production workers, y can be interpreted literally as units of output produced, where the incentive pay component, $\tilde{w}$, depends on number of units produced. For sales workers, y corresponds to volume of sales, and $\tilde{w}$ represents the optimal commission rate. For management position, y may stand for profits corresponding to a unit or division under a manager's supervision. Then, $\tilde{w}$ captures the dependence of the manager's total earnings on this observable measure of performance. As discussed in Milgrom and Roberts (1992), bonuses received by workers in addition to their basic salary are most often implicitly or explicitly linked to observable performance. Hence, $\tilde{w}y$ can be interpreted as a bonus, the size of which, depends on output.

2.2 Households

We model households according to the "collective labor supply" approach developed by Chiappori (1988, 1997)[7]. Three ingredients of this paradigm are crucial from our standpoint. Each spouse *individually* chooses consumption of the market good, market hours and effort. Spouses *jointly* choose home hours, the level of production of the home public good and a sharing rule for household wealth. Individual and joint decisions occur simultaneously.

Each household is endowed with wealth a. We denote the amount of household wealth attributed to each spouse with s_i, for $i = f, m$, where f, m stand for female and male, respectively. The production function for the home public good is

$$G = g(h_f, h_m, k), \tag{11}$$

where k is the amount of market good used in home production. We assume that g is increasing in each argument and concave, and restrict attention to specifications in which h_f and h_m are

[7] This way of modelling the household is consistent with a broad class of efficient bargaining models. See Bergstrom (1997) for a review.

substitutes in the production of the public home good.

Households and individual agents take as given the price of the market good and the labor contracts offered by firms, $\mathcal{C}_i(\pi(h|i))$, where $i = f, m$.. If $\mathcal{C}_i(\pi(h|i))$ satisfies Problem 2 in section 2.1, incentive compatibility and the CARA specification of preferences imply that individual optimality of market consumption, market hours and effort for given home hours is satisfied for each spouse for given h_i, independently from s_i . The households' problem is to choose G, k h_i and s_i to maximize:

$$\sum_{i=f,m} \lambda_i V_i(s_i, h_i; \mathcal{C}) + \theta \log(G), \tag{12}$$

subject to (11), $h_f, h_m \in \mathcal{H}$, and $\sum_i s_i + k = a + \Pi$. Here, $\mathcal{H}$ denotes the set of feasible values of home hours and V_i for $i = f, m$ is the value function for the spouses' individual problem. The parameters, λ_i, for $i = f, m$, represent the weight of each spouse in household decisions.

2.2.1 Choice of Home Hours

We describe in detail the optimal allocation of home hours for a given k and G. This amounts to a cost minimization problem and is independent of the weights λ_i. The opportunity cost of home hours for a spouse is her labor earning potential, which depends on labor contracts. The substitutability of spousal hours in the production of the public home good implies that marginal differences in market earnings will give rise to an asymmetric intra-household allocation of home hours, with the spouse with lower earning potential in market work devoting more time to home production. We interpret the intra-household allocation of home hours as a long term arrangement of the spouses, that may be costly to reverse in the short run. Hence, the choice of home hours determines long-run spousal roles within the household.

For simplicity, we adopt the following functional form for g:

$$g(h_f, h_m, k) = H(h_f, h_m)^{\delta} k^{1-\delta}, \tag{13}$$

$$H(h_f, h_m) = \left[h_m^{\zeta} + h_f^{\zeta}\right]^{1/\zeta}, \tag{14}$$

with $\delta, \zeta \in (0, 1)$ and $h_f, h_m \in \mathcal{H}$. The function $H(\cdot)$ aggregates the contribution of spousal home hours. The parameter δ denotes the contribution of market goods to the production of the public home good, while ζ determines the substitutability of spousal home hours in home production.

The optimal choice of h_f and h_m solves the following cost minimization problem: $C\left(\bar{H}; \mathcal{C}\right) = \min_{h_f, h_m \in \mathcal{H}} w(h_f) + w(h_m)$ subject to $\left[h_m^{\zeta} + h_f^{\zeta}\right]^{1/\zeta} \geq \bar{H}$ for given $\bar{H} > 0$ and given $\mathcal{C}_j(h)$ for $j = f, m$. The first order condition for this problem is:

$$\left(\frac{h_f}{h_m}\right)^{1-\zeta} = \frac{E\left[w_m'(h_m)\right]}{E\left[w_f'(h_f)\right]}, \tag{15}$$

where $w'(h)$ denotes the derivative of total earnings with respect to home hours, which corresponds to the opportunity cost of home hours, and the expectation is taken with respect to ω.

Equation (15) clarifies that the adjusted opportunity cost of home hours for each spouse depends on labor contracts and determines the optimal allocation of home hours. The substitutability of spousal hours in the production of the public home good give rise to specialization, with the spouse with lower opportunity cost devoting more time to home production. The difference in spousal home hours for given labor contracts depends on the elasticity of substitution in H. If $w_f(h) = w_m(h)$ for all $h \in \mathcal{H}$ households are indifferent over the allocation of home hours across spouses and they will randomize.

By substituting (13) and (15) into the objective function, the household optimization problem amounts to the choice of h_f, k , and s_i subject to the household budget constraint. The solution to this problem gives rise to the policy functions: $s_i(a, \mathcal{C})$, $h_{i=f,m}(a, \mathcal{C})$, $k(a, \mathcal{C})$ and $G(a, \mathcal{C})$.

2.3 Equilibrium

We now provide a definition of equilibrium for our economy. We focus on symmetric equilibria, in which all firms have the same beliefs.

Definition 1 *An equilibrium is given by beliefs $\pi(h|i)$ for $i = f, m$, labor contracts $\mathrm{C}(\pi(h|i)) = \{w_i(y), n_i, e_i\}(h)$ for $i = f, m$, and policy functions for the household $\{G, k, h_f, h_m, s_f, s_m\}(a, \mathrm{C})$, such that:*

i) Labor contracts solve Problem 2, given beliefs;

ii) Household policy functions solve the household problem, given labor contracts;

iii) The resulting distribution of home hours in the population is consistent with firms' beliefs.

If we assume that all households are homogeneous with respect to θ and a. This implies that all households will have the same allocation of home hours $\{h_f, h_m\}$ in equilibrium. Given that individuals of different genders are ex ante identical, the equilibrium distribution of home hours across genders depends on firms' self-fulfilling beliefs about this distribution[8]. If $\{h_L, h_H\}$ is the support of the distribution of home hours perceived by firms then in equilibrium $h_j \in \{h_L, h_H\}$, where the values h_L and h_H are endogenously determined. We say that there is *gender discrimination* when firms believe that the distribution of home hours is different in the population of female and male workers.

We can show the following.

[8] Francois (1998) presents a model in which equilibria with gender wage differentials are self-fulfilling. His result relies on exogenously given job heterogeneity. One class of jobs operate under an efficiency wage setting while a second class of jobs operate under piece rate wage setting. Earnings are higher in the efficiency wage jobs. In an equilibrium with female wage discrimination, the first class of jobs is assigned to men, the second to women. The female wage differential stems from job segregation. If all workers were to operate under the same job, the gender wage gap would be reversed. Hence, this model cannot account for gender differentials within the same occupational categories.

Proposition 1 *If all households are homogeneous with respect to θ and a, there are two classes of equilibria:*

i) Equilibria with gender discrimination, with $\pi(h_H|i) = 1$ and $\pi(h_L|j) = 1$ for $i, j = f, m$ and $i \neq j$;

ii) Equilibria without gender discrimination, with $\pi(h|f) = \pi(h|m)$.

In an *equilibrium with gender discrimination* $\pi(h|f) \neq \pi(h|m)$ for $h \in \{h_L, h_H\}$. Then, either $\pi(h_H|f) = 1$ and $\pi(h_L|m) = 1$, or $\pi(h_H|m) = 1$ and $\pi(h_L|f) = 1$. To see that these can be equilibria, note that the distribution of home hours for each gender is degenerate, so that labor contracts for female and male workers solve Problem 1. Then, $w_j(h_j) = f(n_j, e_j) + \omega$ for $j = f, m$ and by concavity of f, $w_j(\cdot)$ is concave in h, so that (15) holds for either $h_f = h_H$ and $h_m = h_L$ or $h_f = h_L$ and $h_m = h_H$.

In an *equilibrium without gender discrimination*, firms believe that the distribution of home hours is the same for male and female workers, so that $\pi(h|f) = \pi(h|m)$ for all possible values of h. The same selection of labor contracts will be offered to female and male workers. Households will be indifferent over which spouse should be assigned high home hours and they will randomize. Lemma 4 in the appendix shows that in an equilibrium without gender discrimination with two possible values of home hours in the population, $\{h_L, h_H\}$, the only equilibrium distribution of home hours is $\pi(h_j|f) = \pi(h_j|m) = 0.5$ for $j = L, H$. Hence, labor contracts will satisfy problem 2 with distribution of home hours $\pi(h_H|m) = \pi(h_H|f) = 0.5$, where h_L and h_H is determined in equilibrium from (15). While there will be heterogeneity in home hours in the population, the distribution of home hours will be the same for female and male workers, thus confirming firms' expectations. Labor market outcomes will be ungendered. These equilibria have the feature that adverse selection is present in the labor market. This is not true in the equilibria with gender discrimination, since the distribution of home hours is degenerate conditional on gender. Hence, in the equilibria with gender discrimination, the incentive problem is less severe.

Equilibria without gender discrimination with distribution of home hours $\pi(h_H|m) = \pi(h_H|f) = 0.5$ are not stable. Households are indifferent over the allocation of home hours across spouses, hence, they could choose any randomization strategy. If one household deviates from $\pi(h_H|m) = \pi(h_H|f) = 0.5$, then, in equilibrium $\pi(h|f) \neq \pi(h|m)$, which is a contradiction. There are also equilibria without gender discrimination in which the distribution of home hours in the population is degenerate and given by one point. In this case, (15) implies $h_f = h_m = \bar{h}$, where $\bar{h}$ is determined in equilibrium from the household problem and labor contracts solve Problem 1. These equilibria are stable and do not feature any adverse selection.

Equilibria with female or male discrimination and equilibria without gender discrimination are equally likely, given that there are no ex ante differences across genders. However, the prevailing gender role distinction in most societies is one in which men specialize in market production and women in home production. Typically, gender differences in labor market outcomes, such as the earnings gap, have been ascribed to this division of labor, which is seen as the result of biological differences, in particular, women's ability to bear children. In the next section, we explore this argument in the context of our model.

2.3.1 Equilibrium with Ex-ante Differences Across Genders

We assume that female and male workers are equally productive in market work, but female workers are more productive in home work. Specifically, we posit that:

$$H\left(h_f, h_m\right) = \left[h_m^{\zeta} + (1+\varepsilon)\, h_f^{\zeta}\right]^{1/\zeta},$$

where $\varepsilon > 0$. A strictly positive value of ε corresponds to higher relative productivity in market vs. home work of male workers with respect to female workers. Women's higher relative productivity in home production is related to their ability to bear children. The parameter ε can be interpreted as tied to the decreased relative market productivity of women during and after pregnancy. Alternatively, if children are viewed as a component of the public home good, ε captures women's higher relative contribution to the nourishment of children via breast feeding. Hence, lower values of ε can be interpreted as corresponding to technological advances, such as medical improvements reducing the physical stress associated with pregnancy and the introduction of baby formula.

The following result holds.

Proposition 2 *In the economy with homogeneous households, if θ and a are such that $h_f \neq h_m$, there exists $\bar{\varepsilon}$, such that:*

i) for $\varepsilon < \bar{\varepsilon}$, there are three possible equilibrium distribution of beliefs $\pi\left(h_H|f\right) = 1$ and $\pi\left(h_L|m\right) = 1, \pi\left(h_H|m\right) = 1$ and $\pi\left(h_L|f\right) = 1$, and $\pi\left(h_H|f\right) = 0.5$ and $\pi\left(h_L|m\right) = 0.5$;

ii) for $\varepsilon \geq \bar{\varepsilon}$, there exists one equilibrium distribution of beliefs, with $\pi\left(h_H|f\right) = 1$, $\pi\left(h_L|m\right) = 1$.

Proof. TBA ■

If ε is greater than $\bar{\varepsilon}$, no self-fulfilling equilibria are possible. In this case, the incentive problem in the labor market provides an amplification mechanism for the relative productivity differences across genders. Hence, the equilibrium gender differences in earnings will be greater than the relative productivity difference. The threshold $\bar{\varepsilon}$ depends on the home production technology, specifically, on the parameter ζ. Technological changes that impact the complementarity between spouses' hours in the production of the public home good change this threshold and the region of self-fulfilling equilibria.

Based on this result, we can interpret the prevailing pattern of gender specialization in the context of our model in the following way. Initially, high values of ε due to poor medical technology imply that the only possible equilibrium is one in which women are mostly devoted to home production and men specialize in market work. Subsequent improvements in medical technology reduce the value of ε, thus making ungendered equilibria possible. However, the self-fulfilling nature of equilibria with gender discrimination for low ε, coupled with the gendered initial conditions, may have contributed to the persistence in gender differences in labor market outcomes and household roles, despite the lack of significant differences in relative productivities. We explore the effect of evolution of medical technologies on labor market outcomes and home hours in Albanesi and Olivetti (2005).

2.4 Numerical Experiments

We illustrate the properties of efficient labor contracts and the resulting implications for the gender earnings differentials numerically for equilibria with and without gender discrimination. We restrict attention to the following functional forms. The disutility from labor is:

$$v(h,n,e) = -\gamma \log(T - h - \eta(n,e)), \tag{16}$$

where η represents the time cost of market work. We assume:

$$\eta(n,e) = \left[\nu (n+1)^{\rho} + (1-\nu)(e+1)^{\rho}\right]^{1/\rho} - 1, \tag{17}$$

with $\rho, \nu \in (0,1)$, which implies $\eta_{ne} > 0$ and $v_{ne} > 0$. The production function is:

$$f(n,e) = \left[(1-\phi)(1+\tau e)^{z} + \phi(1+n)^{z}\right]^{1/z} - 1, \tag{18}$$

with $\tau > 0$ and $\phi, z \in (0,1)$. The functional forms for η and f are extremely flexible and allow us to explore the properties of optimal contracts as a function of the degree of substitutability between effort and hours in each function.

Our benchmark parameterization is based on empirical evidence where possible and is reported in Table 2. The value of δ is set following Benhabib, Rogerson and Wright (1991). The parameter ζ is chosen to be consistent with the ratio of home hours for married women and men, which ranges from 4.3 in 1976 to 2.8 in 1993 based on PSID data. We target $h_f/h_m = 3$, the average value of this statistic for the 1990's, which implies $\zeta = 0.87$. The values of γ and T are set so that the ratio of market to home hours matches data from the PSID. The average of this ratio for the 1990's is equal to 1.25 for married female workers and to 6.04 for married male workers. The parameters ρ, ν, z, ϕ, τ and Σ, which pertain to the utility cost of working and the firm's production technology, cannot be calibrated based on aggregate data. We restrict attention to the range $\phi, \nu > 0.5$ and $z, \rho > 0.5$. We set τ so that it is optimal to implement high effort for both types of workers, given the other parameters. The value assigned to the parameter Σ implies that the standard deviation of output conditional on effort and hours is approximately equal to 9% for males and 10% for females.

Table 2: Benchmark Parameters

σ	γ	T	ν	ρ	ϕ	τ	Σ	z	δ	ζ	θ
1.1	0.3	1	0.68	0.8	0.68	2.5	1	0.75	0.68	0.87	1

We study the basic properties of our model in equilibria with and without female discrimination[9]. Table 3 reports summary information for the equilibrium with female discrimination:

[9] Equilibria with male discrimination, which are also possible, are the same as equilibria with female discrimination, with all outcomes relabelled by gender.

Table 3: Equilibrium with Gender Discrimination								
	w	$\widetilde{w}$	$\overline{w}$	n	e	Σ/y	h	n/h
Male	3.32	0.89	0.33	0.66	1	0.09	0.1	6.6
Female	2.88	0.91	0.24	0.37	1	0.10	0.3	1.23

The female/male earnings ratio is equal to 0.89. We use a simple Oaxaca-Blinder decomposition of female earnings, at the male parameters, to evaluate the component of the earnings gap not explained by differences in market hours. The decomposition is based on the earnings schedule, hours and output generated by the model. We find that 13% of the earnings gap is not explained by differences in hours worked on the market across genders. These values are consistent with the corresponding data statistics, reported by O'Neill (2000). Using CPS data, she shows that the gender earnings ratio is equal to 0.80, and that there a 10% unexplained gender differential in earnings after controlling for observable differences in human capital and job characteristics.

The properties of the equilibrium without gender discrimination for our benchmark parameterization are displayed in table 4:

Table 4: Equilibrium without Gender Discrimination					
	w	$\widetilde{w}$	n	e	n/h
h_L	2.14	0	0.80	0	8
h_H	1.74	0	0.55	0	1.83

,

In this case, low effort is implemented for both types of workers, despite the fact that high effort is optimal for both males and females in the equilibrium with gender discrimination. Total earnings are lower for both types of workers relative to the equilibrium with discrimination. The ratio of earnings for workers with high home hours relative to workers with low home hours is equal to 0.81- lower than in the equilibrium with gender discrimination. These findings illustrate the additional inefficiencies associated with unobserved home hours, in equilibria without gender discrimination.

We now focus on the equilibrium with female discrimination in which female workers have high home hours and the only incentive problem is moral hazard due to hidden effort. A firm may decide to implement low or high effort for a worker depending on the cost of implementing high effort. Such a cost, which is larger for female workers, depends on the degree of substitutability of n and e in utility and in production and on the workers' relative cost of effort, as discussed in section 2.1.

To investigate this, we vary each relevant parameter in isolation, maintaining the other parameters at their benchmark values. We also fix home hours for both genders. Figure 1 displays how the fraction of incentive pay in earnings for females and males and the earnings ratio vary with ρ, which corresponds to the degree of complementarity between effort and hours in the utility cost of working. The dashed and dash-dotted lines represent the fraction of incentive pay in earnings, which is equal to $\tilde{w}$, for females and males, respectively. The

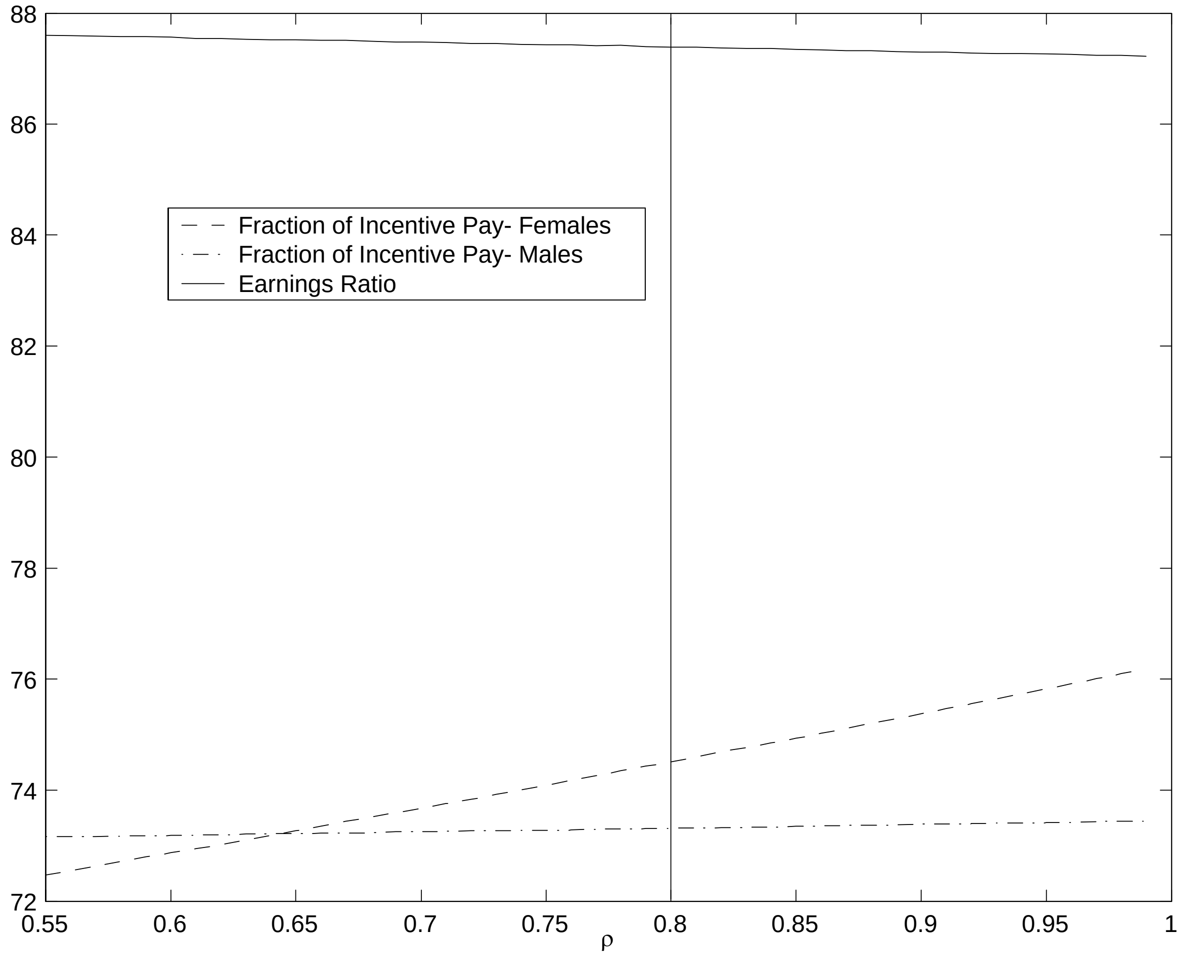

Figure 1: Female and male earnings and the complementarity between effort and hours in utility

solid line represents the earnings ratio. All numbers are expressed in percentages. The vertical line corresponds to our benchmark parameterization. As ρ increases, the fraction of incentive pay increases for both male and female workers, but more so for females. Higher values of ρ correspond to higher substitutability between effort and hours in utility. This exacerbates the incentive problem, more so for women who have a higher marginal cost of effort. For high enough values of ρ, the fraction of incentive pay is higher for women than men.

Figure 2 plots the fraction of incentive pay in earnings for females and males and the earnings ratio as a function of z. Lower values of z correspond to lower substitutability between effort and hours in production. This reduces the marginal product of effort for firms at lower values of hours. It follows that lower values of z increase the cost of implementing high effort more for female workers, since they work fewer hours. As long as high effort is being implemented for male and female workers, the fraction of incentive pay in earnings is very

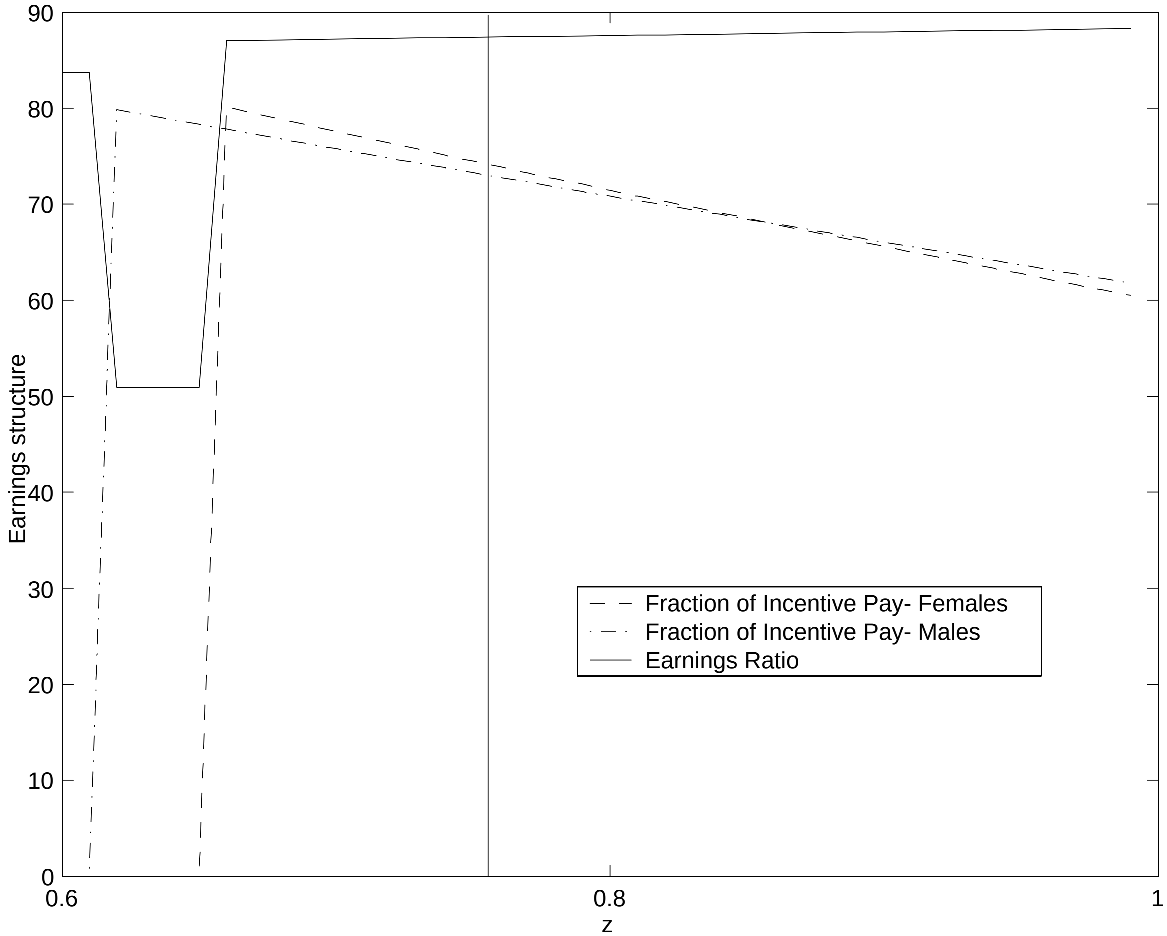

Figure 2: Female and male earnings and the substitutability between effort and hours in production

similar for both genders, though higher for women. For sufficiently low values of z, however, firms will find it optimal to implement low effort for female workers, while high effort is still implemented for males. This determines a large drop in the earnings ratio.

Figure 3 plots the fraction of incentive pay in earnings for females and males and the earnings ratio as a function of ν,which determines the relative cost of effort for workers. As ν declines, the relative cost of effort increases, and this exacerbates the incentive problem. The interpretation is similar to the one for Figures 1 and 2.

Our findings for ν are consistent with historical evidence on the gender wage gap and on the structure of earnings for female and male workers in manufacturing and in the clerical sector in Goldin (1986, 1990). Goldin argues that clerical sector work was perceived as less onerous, involving less physical fatigue and discomfort than manufacturing work. She documents that in manufacturing, piece-rate compensation was more prevalent for women. Specifically, 47% of female operatives and 13% of males in the same positions are paid by the piece in manufac-

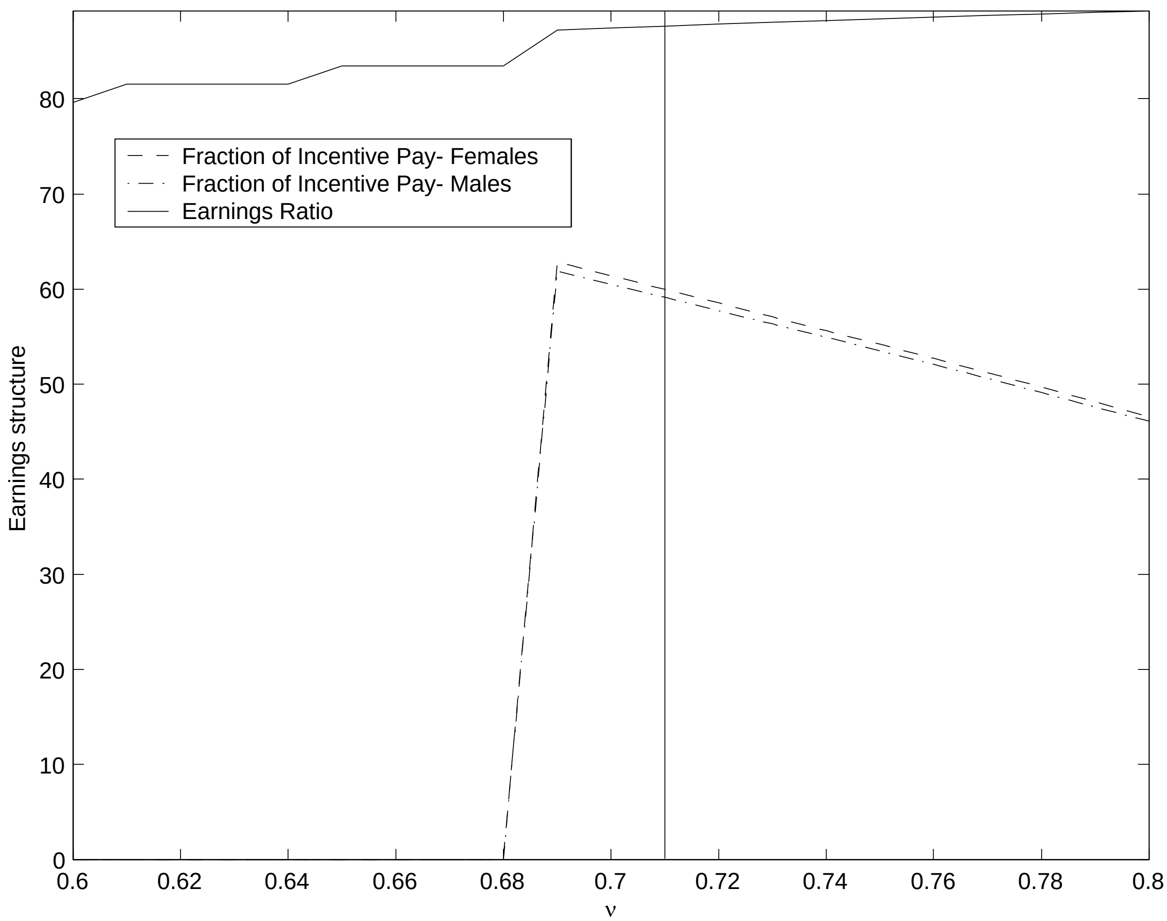

Figure 3: Female and male earnings and the relative cost of effort

turing in 1890. This accords with equilibria in our model in which male and female workers both exert high effort. Goldin argues that in the clerical sector, "career tracks" emerged as a standard motivational device, replacing the piece-rate compensation schemes prevalent in manufacturing. She maintains that in the clerical sector: "Firms often used sex as a signal of shorter expected job tenure. ... By segregating workers by sex into job ladders (and some dead-end positions), firms may have been better able to use the effort-inducing and ability-revealing mechanisms of the wage structure." If we adopt a job ladder interpretation, labor contracts implementing low effort can be thought of as "dead end jobs", while labor contracts that implement high effort can be thought of as positions that allow for career growth. Then, our results on the effect of a higher relative cost of effort on the earnings structure are consistent with the differences in earnings structure by gender across these two sectors. Moreover, these findings are in line with the transition to the clerical sector in the period 1930/1940, when, as discussed in Goldin (1990), there was also a significant rise in the gender wage gap.

Finally, we explore the properties of the model as a function of the parameter Σ, which corresponds to the variance of output for given hours and effort. An increase in this parameter makes it harder to infer effort from observed output. Figure 4 displays the fraction of incentive pay in earnings for females and males and the earnings ratio for different values of Σ. We vary Σ from values approximately equal to 10% of output to values approximately equal to 60% of output for males. As Σ increases, the surplus from the employment relationship declines, given that workers are risk averse. This can be seen from the multiplier on the incentive compatibility constraint, derived in section 2.1, which is proportional to Σ. It follows that the returns from implementing high effort for firms decline. This effect is greater for women, who have a higher marginal cost of effort. Then, there will be a critical value of Σ, for which it becomes too costly for firms to implement high effort for females, but not for males. As discussed for Figure 1-3, as long as high effort is being implemented for both, the fraction of incentive pay in earnings is similar across genders, though higher for women and the earnings ratio roughly constant. The earnings ratio drops significantly, when low effort is implemented for females but not for males. Extrapolating these findings to a version of the model with continuous effort, they translate into the following predictions: the F/M earnings ratio should be lower when the incentive problem is more severe and the difference in the fraction of incentive pay for male and female workers should be negatively related to the female/male earnings ratio. We argue that industries/occupations may differ in the severity of the incentive problem as captured by Σ. We explore the implications for gender earnings differentials across industries and occupations in the next section.

3 Connecting the Model with the Evidence

Our model delivers several predictions about the relation between gender earnings differentials, incentive problems in the labor market and the household division of labor. First, the gender earnings differentials should be higher in industries/occupations where the incentive problem is more severe. Moreover, this pattern should be stronger for married workers than for single workers. Second, our model predicts a negative correlation between the difference in the

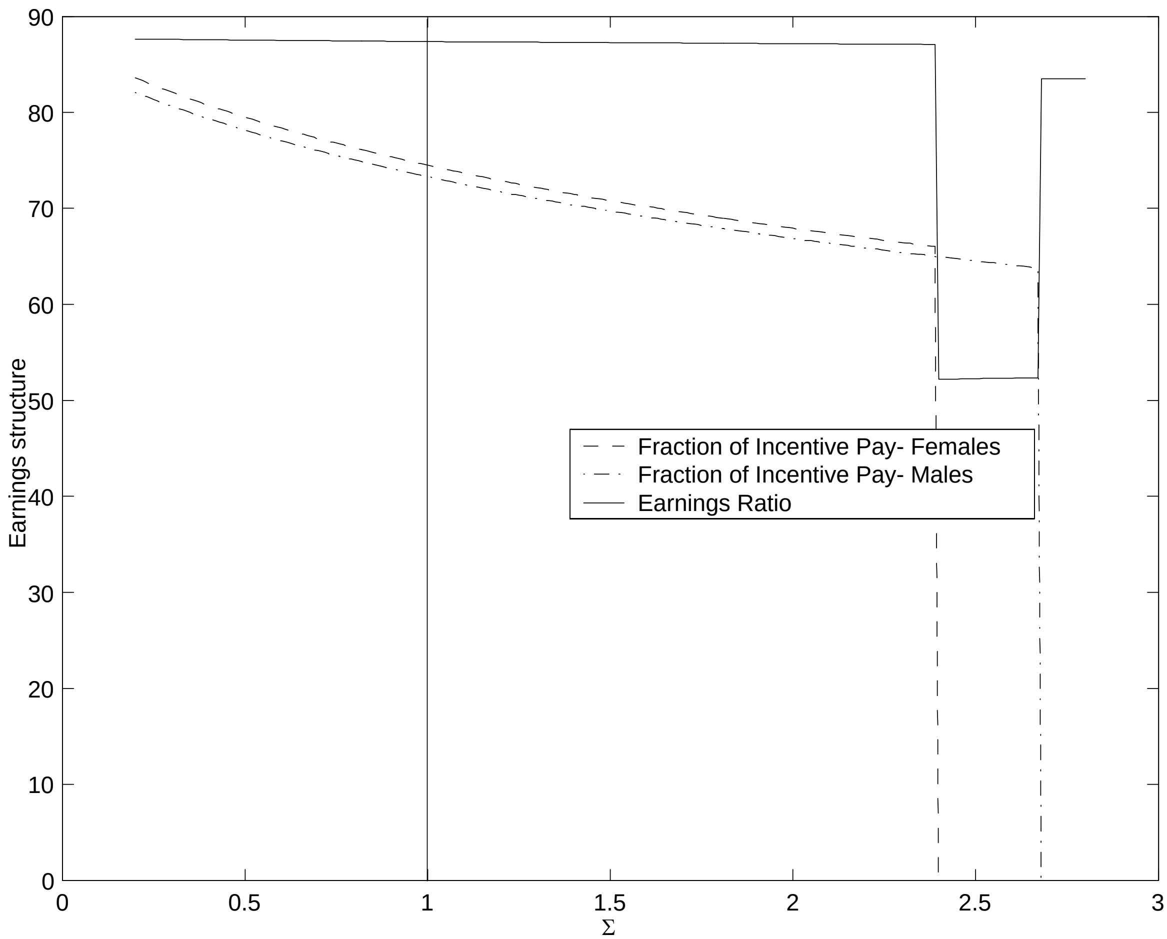

Figure 4: Female and male earnings and the severity of incentive problems

fraction of incentive pay for male and female workers and the female/male earnings ratio.

We use Census 2000 data to study gender earnings differentials by marital status across industries and across four broad occupational categories. We find that for married individuals gender differentials are greatest in high management and sales occupations across all industries. Evidence suggests that these occupations are likely to be the ones in which incentive problems are most stringent, as discussed below. For the sample of single workers, we do not observe the same pattern of gender differentials across industries and occupations. Since the Census does not include information on the fraction of incentive pay, we use PSID data from the late 1990s to document the second prediction. We find that there is a strong negative correlation between the male/female difference in the fraction of incentive pay and the female/male earnings ratio.

According to our model, we should also observe a negative correlation between the female/male ratio of home hours and the female/male earnings ratio. In keeping with the previous analysis, we study this correlation across industries and occupations using the PSID. We find that this correlation is negative and significant for married workers. It is much smaller for single workers.

3.1 Earnings Gaps and Incentive Problems

We use Census data to document differences in gender earnings differentials across all industries for four broad occupational categories: higher level management, lower level management, sales and production.[10] We consider this occupational classification based on the notion that incentive problems are more stringent in sales and high management relative to production and low management. In our model, the severity of the incentive problem is linked to the degree of uncertainty over the workers' effort for given observable measures of performance, which correspond, respectively, to the parameter Σ and output[11]. As shown in MacLeod and Parent (2003), high management and sales occupations are characterized by greater workers' autonomy and larger variety of tasks, which increase the severity of the incentive problem. More generally, for management occupations, the uncertainty associated with managers' effort given observable performance measures should be related to the complexity of the job, which is greater for high level managers. For sales occupations, sales volumes are typically used as a benchmark measure of performance. Yet, these depend to a large degree on variables that are not directly related to a sales personnel's effort and hours of work and may be uncertain[12]. These considerations are less important for production workers.

[10] We use Census 2000. The industry variable, INDNAICS, reports the type of establishment in which a person worked in terms of the good or service produced. Industries are coded according to the North American Industrial Classification System developed in 1997. We use the variable OCCSOC for occupation. OCCSOC classifies occupations according to the 1998 Standard Occupational Classification (SOC) system. The Census also provide an aggregation of all the occupations in 23 broader categories that include the four categories considered in the analysis. The definition of production occupations also includes construction and extraction workers.

[11] See sections 2.1 and 2.4.

[12] For example, sales workers are typically assigned to specific territories or products. Hence, sales volumes will fluctuate with shocks to local demand. See Catalyst (1995) for a description of the sales occupation, especially in relation to gender.

Following the Census classification we consider 19 industries: Agriculture; Forestry Fishing and Hunting, Mining, Utilities, Construction, Manufacturing, Wholesale Trade, Retail Trade, Transportation and Warehousing, Information, Finance and Insurance, Real Estate and Rental/Leasing, Professional, Scientific and Technical Services, Administrative and Support and Waste Management and Remediation Services, Educational Services, Health Care and Social Assistance, Arts, Entertainment and Recreation, Accommodation and Food Services, Other Services (except Public Administration), and Public Administration. Our sample includes all individuals between 15 and 64 years of age, who are not in school, do not reside on a farm or live in group quarters. We also exclude the armed forces and restrict attention to those individuals who work at least 50 weeks in the previous year and who usually work at least 30 hours per week.

For each industry and for each of the four broad occupational categories we compute the gender gap in earnings by running median regressions that control for a gender dummy as well as for human capital variables - age and its square term and education. Our measure of the median gender gap in earnings is the coefficient of the gender dummy in the regression. For the education variable, we group individuals according to four broad educational categories: less than high school, high school completed, some college and college completed. We construct four education dummies based on this categorization. The first dummy is equal to one if an individual has completed less than twelve years of schooling and is equal to zero otherwise. The second dummy variable is equal to one if he or she has completed twelve years of schooling, and is equal to zero otherwise. The third dummy variable equals one if the individual has completed between twelve and fifteen years of schooling and it is equal to zero otherwise. Finally, the fourth dummy variable is equal to one if an individual has completed at least sixteen years of education and it equals zero otherwise. The omitted dummy variable corresponds to individuals who completed less than twelve years of schooling. The dependent variable is the log of annual earnings. We use total labor earnings in our analysis because this is the data counterpart of the measure of total labor compensation that we consider in our model. However, one could argue that this is not the appropriate measure of labor compensation when making gender comparisons, since women tend to work less hours on the labor market than men do. Hence, our analysis could be confounding gender differences in market hours and in hourly compensation. This concern is attenuated by the fact that we only consider individuals that usually work at least 30 hours per week and who were employed for at least 50 weeks in the previous year. This sample selection criterion reduces considerably the variation in the number of market hours within and between gender groups.[13]

Table 5 reports the results of this analysis. The first column in the table reports the female/male ratio of median earnings for full-time year-round workers across all industries for the four broad occupational categories. The second and third column report, respectively, the frequency with which the gender wage ratio is lowest (second column) and highest (third

[13] We also conducted our analysis using the log of hourly wages as a dependent variable. Our findings in this case are consistent with the ones reported in this section, and results are available upon request. However, we do not report them here, since the hourly wage does not have a model counterpart, given that in our framework, labor contracts are non-linear in hours.

column) across the four categories. In each column we report the results obtained both for the sample of married individuals and for the sample of individuals who were never married.

We find that for married workers there is a systematic pattern in the ranking of gender earnings gaps over the four occupational categories across all the industries. In particular, we find that the female/male median earnings ratio is lowest for sales occupations in 13 industries - 68% of the cases. In the 6 remaining industries (32% of the cases) the occupation that display the lowest ratio of female to male median earnings is high management. On the other end, when we look at the frequency with which the gender wage ratio is highest across all industries we find the opposite pattern. Workers in production occupations display the largest ratio in 13 industries - 68% of the cases. In the 6 remaining industries the ratio is highest in low level management occupations. As shown in column 1, for the sample of married workers the average female/male ratio of median earnings in all industries varies substantially across the four occupation categories, from a minimum of 0.57 in sales to a maximum of 0.79 in production.

Table 5: Gender differences in earnings across industries
for four broad occupational categories. (Full time, year round workers.)

	% female/male earnings ratio		% times lowest ratio		% times highest ratio	
	married	*single*	*married*	*single*	*married*	*single*
higher level mgmt	65	94	32	5	0	53
lower level mgmt	74	86	0	16	32	16
sales	57	86	68	21	0	16
production	79	83	0	47	68	5

For the sample of never married workers the gender gap is much lower and varies less across industries and occupations. Single women earn at least 83% of single men's median earnings across all industries. In contrast to our findings for married workers, we find that the gender earnings ratio for single workers is largest in high management occupations in 10 of the 17 industries - 53% of the cases. For the 7 remaining industries we find that the earning ratio is highest in sales (3 times), low management (3 times) and production (1 time). The ratio is lowest in production occupations for 9 industries - 47% of the cases. For the 8 remaining industries we find that the earning ratio is lowest in sales (4 times), low management (3 times) and high management (1 time). Note that for this sample we can only consider 17 industries since in two industries, Agriculture and Forestry and Mining, there are two few observations in sales occupations.

Table 6: Variation in female/male median earnings ratios by industries, occupation, and marital status

	Married			Never Married		
	Lowest	Highest	All	Lowest	Highest	All
Accommodation and Food	0.38	0.78	0.61	0.82	1	0.86
	(sales)	(prod.)		(sales)	(high mgt)	
Administrative and Support	0.64	0.79	0.70	0.85	0.91	0.90
	(sales)	(prod.)		(prod.)	(sales)	
Agriculture, forestry	0.20	0.98	0.75			
	(sales)	(low mgt.)				
Arts, Ent. & Recreation	0.63	0.82	0.73	0.81	1.12	0.92
	(sales)	(prod.)		(high mgt)	(low mgt)	
Construction	0.51	0.82	0.59	0.67	0.94	0.89
	(high mgt.)	(prod.)		(low mgt)	(high mgt)	
Educational Services	0.60	0.83	0.72	0.88	0.95	0.91
	(sales)	(low mgmt.)		(prod.)	(high mgt)	
Finance and Insurance	0.55	0.82	0.54	0.83	0.92	0.85
	(sales)	(prod.)		(sales)	(high mgt)	
Health Care & Soc. Assistance	0.54	0.74	0.65	0.62	0.97	0.91
	(sales)	(prod.)		(sales)	(high mgt)	
Information	0.68	0.82	0.64	0.87	0.91	0.90
	(high mgt.)	(prod.)		(prod.)	(high mgt)	
Manufacturing	0.60	0.74	0.59	0.61	0.95	0.83
	(sales)	(low mgt.)		(prod.)	(high mgt)	
Mining	0.56	1.01	0.59			
	(high mgmt)	(low mgt.)				
Other Services (no Public. Adm.)	0.44	0.73	0.57	0.72	1.02	0.85
	(sales)	(low mgt.)		(prod.)	(low mgt)	
Profess,Scientific&Tech. Services	0.61	0.81	0.61	0.86	0.90	0.86
	(sales)	(prod.)		(prod.)	(high mgt)	
Public Administration	0.64	0.82	0.69	0.88	1.23	0.87
	(sales)	(prod.)		(prod.)	(sales)	
Real Estate and Rental/Leasing	0.58	0.83	0.70	0.88	1.02	0.95
	(sales)	(prod.)		(prod.)	(high mgt)	
Retail Trade	0.44	0.68	0.51	0.71	0.91	0.82
	(sales)	(prod.)		(prod.)	(high mgt)	
Transportation and Warehousing	0.55	0.82	0.72	0.85	1	0.91
	(sales)	(prod.)		(low mgt)	(high mgt)	
Utilities	0.64	0.83	0.64	0.55	0.88	0.80
	(high mgt.)	(prod.)		(sales)	(prod.)	
Wholesale Trade	0.61	0.85	0.62	0.70	0.97	0.88
	(high mgt.)	(low mgt.)		(low mgt)	(high mgt)	

Our findings suggest that there is a considerable variation in the female/male ratio of median earnings across industries, and across the four occupational categories within an industry, even after controlling for human capital variables. We report these ratios by industry, occupation and marital status in Table 6. For each industry, we include the value for the occupation with the lowest and highest ratio, as well as for all occupations in the Census. The variation of the median earnings ratio across industries and occupational categories is striking. For married workers it ranges from 0.20 for sales occupation in the Agriculture and Forestry industry to 1.012 for low management occupation in Mining. For never married workers the range is smaller, from a minimum of .55 for sales occupation in the Utilities industry to a maximum of 1.23 for sales occupation in the Public Administration.

3.2 Fraction of incentive pay, earnings gaps and home hours

Our model predicts that the difference in the fraction of incentive pay for male and female workers should be negatively related to the female/male earnings ratio, and that we should observe a negative correlation between the female/male ratio of home hours and the female/male earnings ratio across industries/occupations. We use PSID data for the late 1990s to document these facts. As we did with the Census data set, we select our sample to include all individuals between 15 and 64 years of age who are not in school and who are not in the armed forces. In our analysis, we study gender earnings ratios at the occupation/industry level. This level of disaggregation requires a larger sample size than the one available in each wave of the PSID. Hence, we do not exploit the panel dimension of the data set but simply pool together all the individuals in the 1994 to 2001 waves of the PSID. The resulting statistics can be interpreted as medium run averages of the relevant variables.[14]

To analyze the relation between the fraction of incentive pay and gender earnings differentials, we use the information on bonuses and commissions. We calculate the fraction of incentive pay as the ratio of bonuses and commissions to labor income, defined as wages and salaries, plus bonuses and commissions.[15] The PSID coding of occupations differs from the one available from the Census 2000. In our analysis, we construct occupation categories that are similar to the ones used in our Census analysis. We consider the following categories: management positions in administration, management positions in banking, finance and in the clerical sector, lower level management occupations, professional occupations (engineers,architects, lawyers, and medical doctors), technical occupations (in the health sector, engineering, and social sciences), occupations in community/social services, social scientists and university professors, teachers other than college professors, occupations in arts and entertainment, design, sports and the media, sales occupations, clerical occupations, craftsmen, operatives, physical laborers, laborers in services excluding private households, and laborers in private households.

[14] In the PSID, data on hours worked, total labor earnings, bonuses and commision income, are reported for the previous calendar year. Hence, our data covers the time period 1993-2000.

[15] Disaggregated data on bonuses are available from 1984 to 1992. However, data on commissions were made available only for 1993 (in the Income Plus PSID files). Hence, we restrict attention to the more recent waves. The majority of workers report either bonus or commission pay. We exclude workers with real weekly earnings below $67 in 1982 dollars from the sample. We deflate nominal variables using the CPI with base year 2000.

We find a strong negative correlation between the female/male earnings ratio and the male/female difference in the fraction of incentive pay. The correlation coefficient is -0.82 and it is significant at the one percent level. Figure 5 displays a scatter plot of these two variables. Consistent with our Census findings, sales and high management occupations are characterized by the lowest female/male earnings ratio and the highest male/female difference in the fraction of incentive pay.[16] According to our model, this relation should be stronger for married workers. However, the PSID only reports information on bonuses and commissions for household heads, that are mostly married males or single women. Since observations on married women are extremely limited,[17] we do not condition on marital status for this analysis.

We use the PSID data on bonuses and commissions to corroborate our findings on the severity of the incentive problem and gender earnings differentials discussed in section 3.1. Figure 6 reports a scatter plot of the aggregate fraction of incentive pay and the female/male earnings ratio across occupations. The correlation between these two variables is -0.62, significant at the one percent level. Consistent with MacLeod and Parent (2003), the occupations where the incentive problem is more severe exhibit a higher fraction of incentive pay. These same occupations also have low female/male earnings ratios.

There is a large variation in the percentage of female workers across the occupations considered in this analysis. This raises the question of whether occupational sorting can in fact explain the gender differential in earnings and fraction of incentive pay across occupations. One could think of two reasons why occupational choice might depend on gender- comparative advantage or differences in preferences.[18] Occupational sorting based on comparative advantage would predict that, if workers are paid their marginal product, the female/male earnings ratio should be higher in occupations in which women have a comparative advantage, thus giving rise to a positive correlation between percentage of female workers and the female/male earnings ratio across occupations. This prediction is strongly rejected by the data. As shown in Figure 7, there is no clear relation between these two variables and their correlation is not significantly different from zero.[19]

Survey and experimental evidence seems to suggest that women are more risk averse than men[20]. Then, this could provide the basis for occupational sorting of workers by gender in the contest of our model. If risk aversion is not observed and all workers are homogenous in home hours, firms will offer incentive compatible contracts that induce workers to self-select according to their risk-aversion. Hence, female workers would choose occupations with less variable earnings- with a lower fraction of incentive pay.[21] This prediction is also strongly

[16] To account for the role of differences in hours worked in determining gender earnings differentials, we also conduct this analysis for hourly wages. We find that the correlation between the female/male difference in log hourly wages and the male/female difference in the fraction of incentive pay is -0.83, significant at the one percent level.

[17] Married women with information on bonuses and commission are 206 out of approximately 10,000 women in the sample.

[18] In our theoretical analysis, workers are identical in market productivity and talents. Hence, we do not derive any predictions on occupational sorting.

[19] The same is true for the difference in log hourly wages.

[20] See for example Niederle and Vesterlund (2005).

[21] Jullien, Salanie and Salanie (2003) analyze a model with hidden effort and private risk aversion. They show

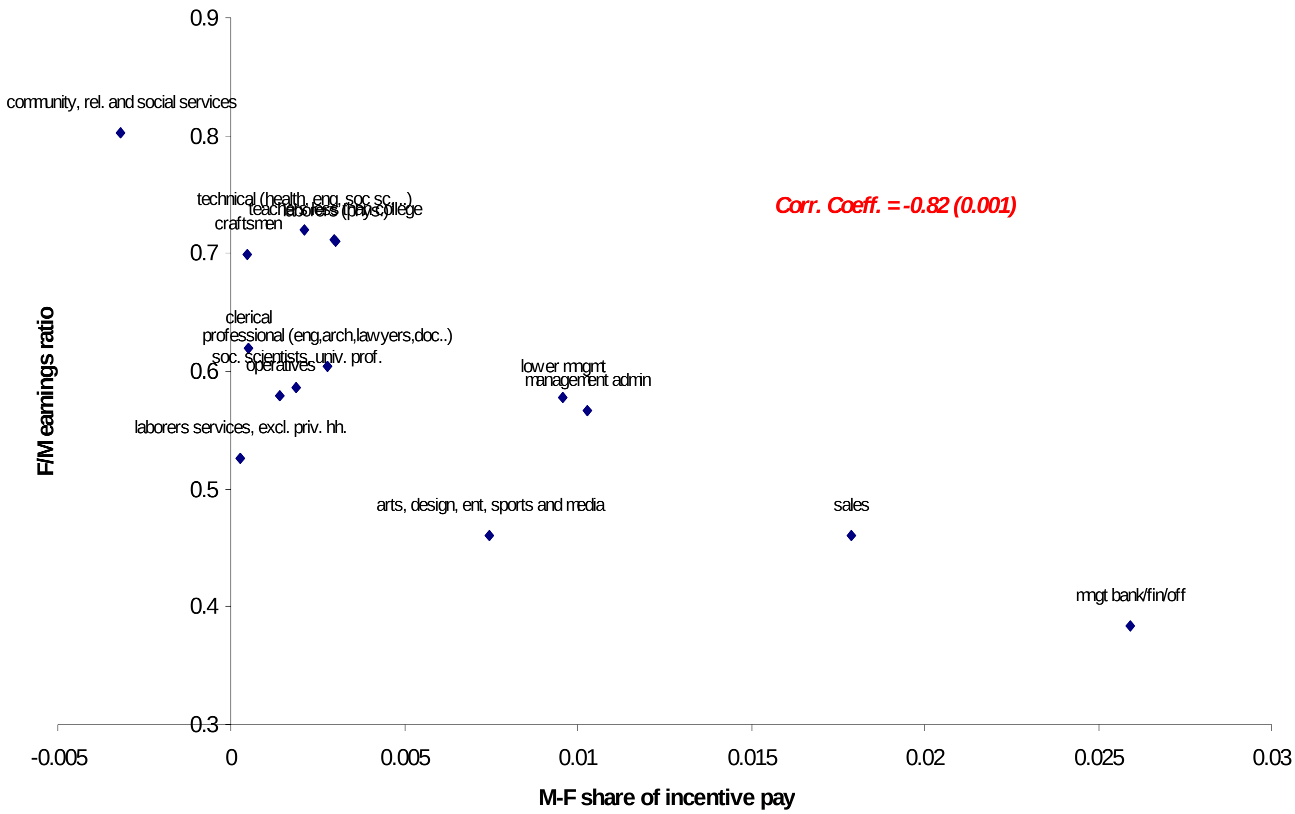

Figure 5: Correlation between the M-F difference in the fraction of incentive pay and the F/M earnings ratio.

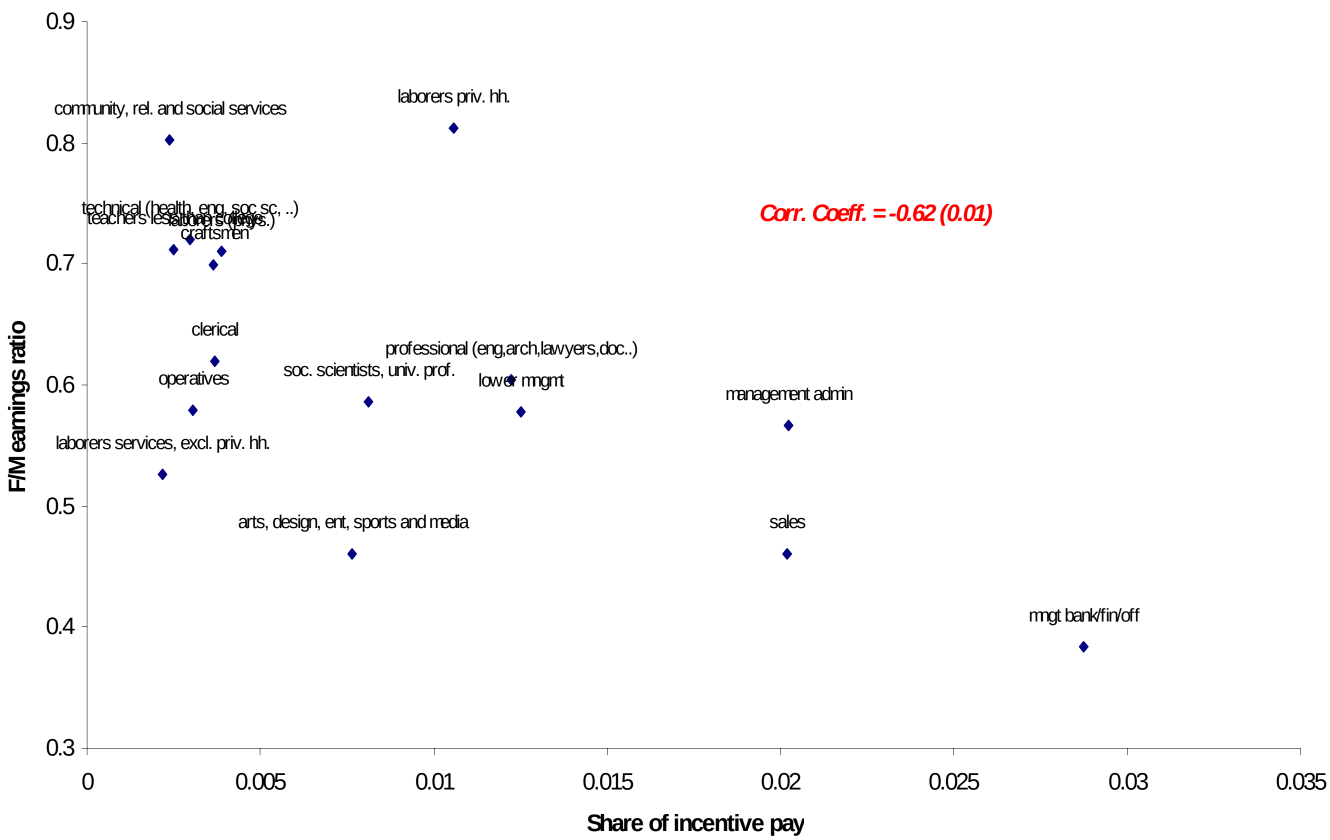

Figure 6: Correlation between the F/M earnings ratio and the aggregate fraction of incentive pay.

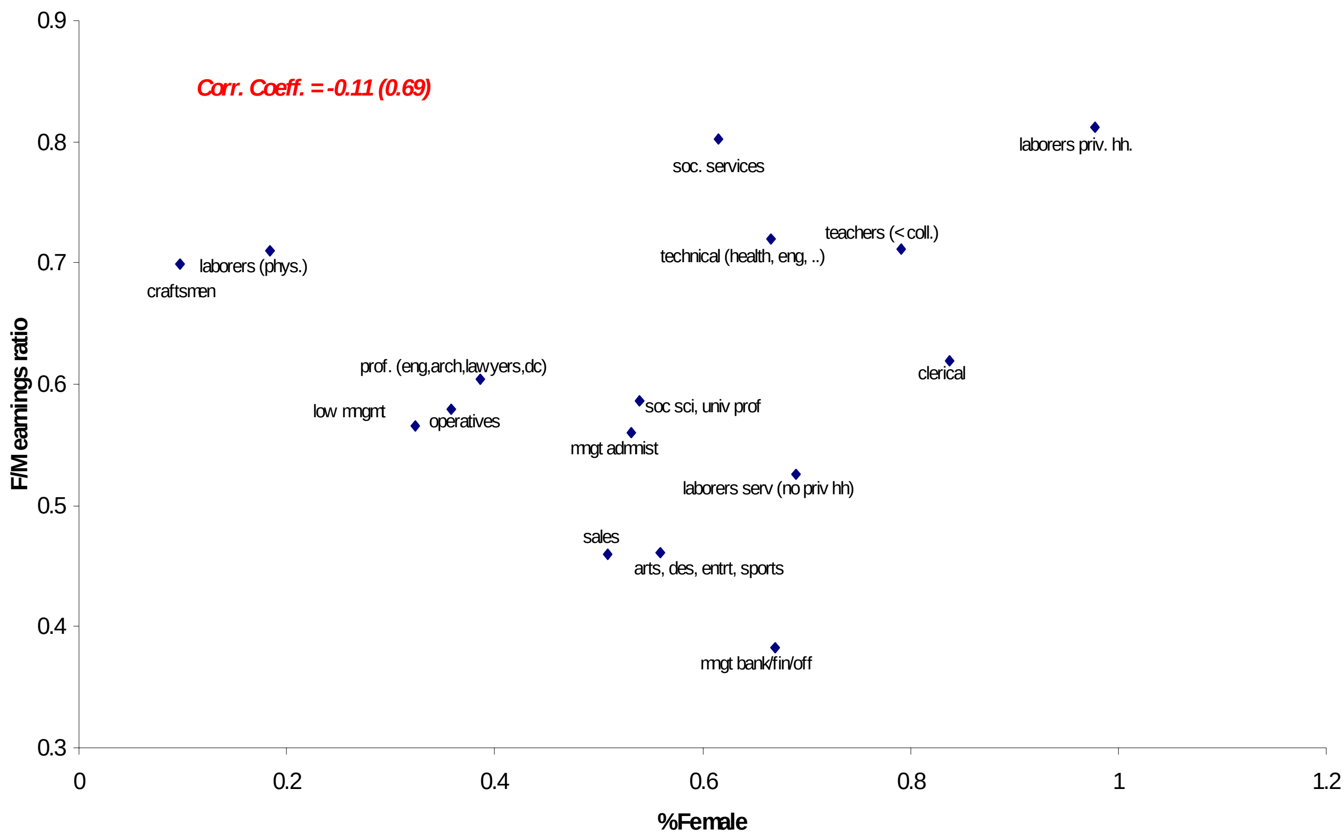

Figure 7: Correlation between the percentage of female workers and the F/M earnings ratio.

rejected by the data. In Figure 8, we plot the aggregate fraction of incentive pay against the percentage of female workers. There is no clear relation between these two variables and their correlation is not significantly different from zero.

To analyze the relation between the gender differentials in home hours and in earnings, we use PSID data on weekly home hours[22] and annual labor earnings. Our sample includes 49,301 observations of which 33,379 are married and only 7,728 are single. Figure 9 plots the female/male ratio of average home hours against the female/male ratio of median earnings across occupations. Lower female/male earnings ratios are strongly associated with higher female/male ratios of home hours for workers in management occupations in banking and finance and in the clerical sector, and for workers in sales occupations. However, we find that for workers in operative occupations and for laborers higher gender earnings ratios are associated with higher female/male ratios of home hours. We think that wealth effects might explain this observation. The correlation between the two ratios is twice as large for married workers than for never married workers. For married workers, the correlation is -0.31, and significant at the 1 per cent level. For single workers the correlations is equal to -0.15.

Finally, our model is based on the assumption that households make efficient decisions on the allocation of home hours, which implies that the spouse with higher earnings will contribute fewer hours to the production of the home public good. This prediction is supported by the PSID data. We find that 86.3% of married couples in the PSID sample display a negative relation between the husband/wife ratio of earnings and the husband/wife ratio of home hours.[23] The fraction of married couples for which the relation between these two ratios is negative ranges from 73% in the bottom quartile to approximately 95% in the top quartile.

4 Concluding Remarks

The data available through the Census and the PSID enables us to provide suggestive evidence in support of the mechanisms for wage determination embedded in our model, as well as for some of its implications for the interaction with household variables. However, neither data set is rich enough to *directly* evaluate the predictions regarding the interaction between the labor market and the household or the predictions on the structure of earnings by gender. Specifically, the Census sample does not include information on home hours or on the fraction of incentive pay. This information is included in the PSID, although this data set does not include information on the structure of earnings for *both* husbands and wives. Our model predicts a negative correlation between the difference in the fraction of incentive pay for male

that in a separating equilibrium, agents with higher risk aversion will have lower fraction of incentive pay, under the analogue of a single-crossing condition. However, this does not imply that they exert lower effort, since a given level of earnings variability has a stronger incentive effects for more risk-averse agents.

[22] The information on home hours in the PSID is not obtained from time diaries, rather the survey respondent is asked to provide a measure of weekly hours worked by him- or herself and by his or her spouse (if married.) As a consequence, the home hours variable is potentially subject to serious measurement problems.

[23] In 3.7% of the remaining households the husband is both the main earner and the main provider of home hours. This leaves a 10% of households where wives have both higher earnings and higher home hours than their husbands.

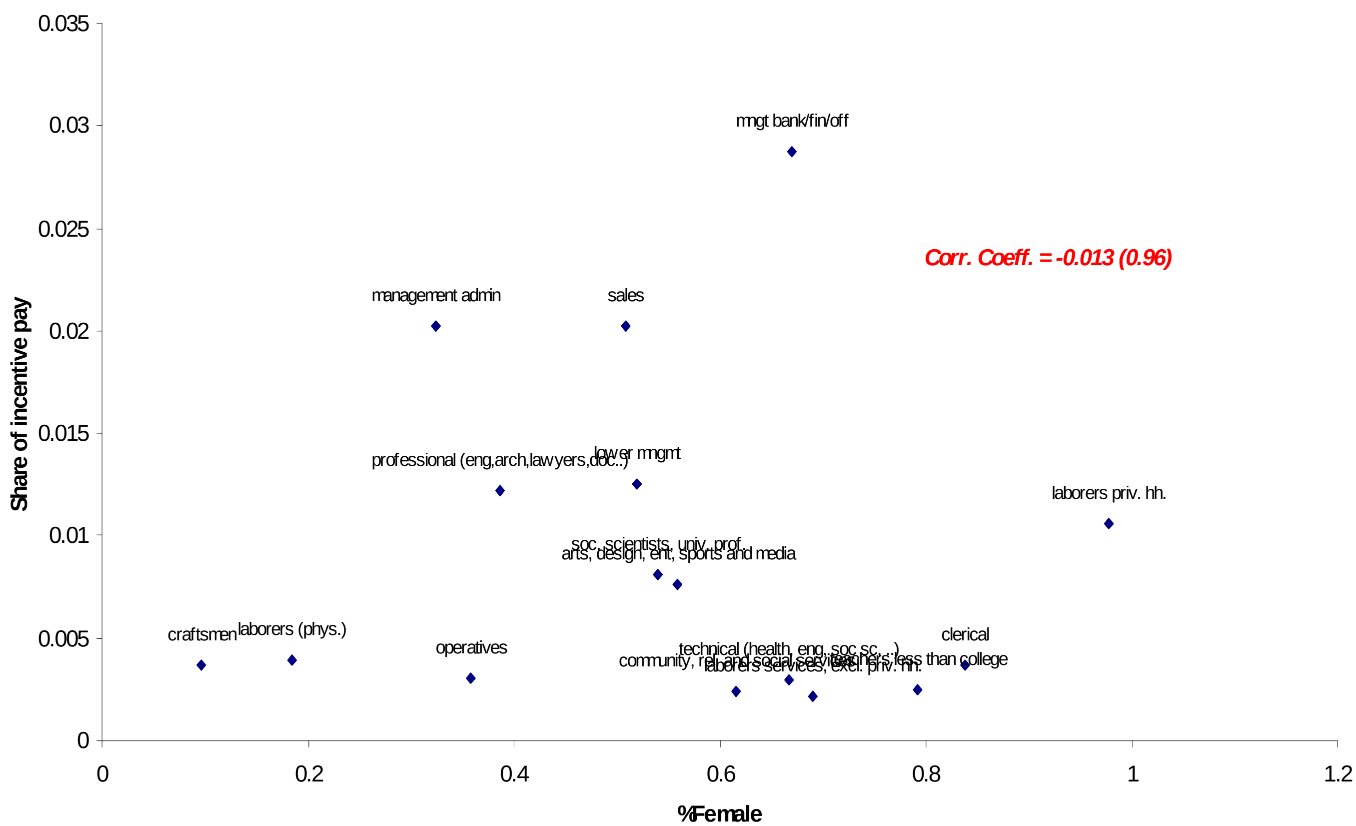

Figure 8: Correlation between the aggregate fraction of incentive pay and the percentage of female workers.

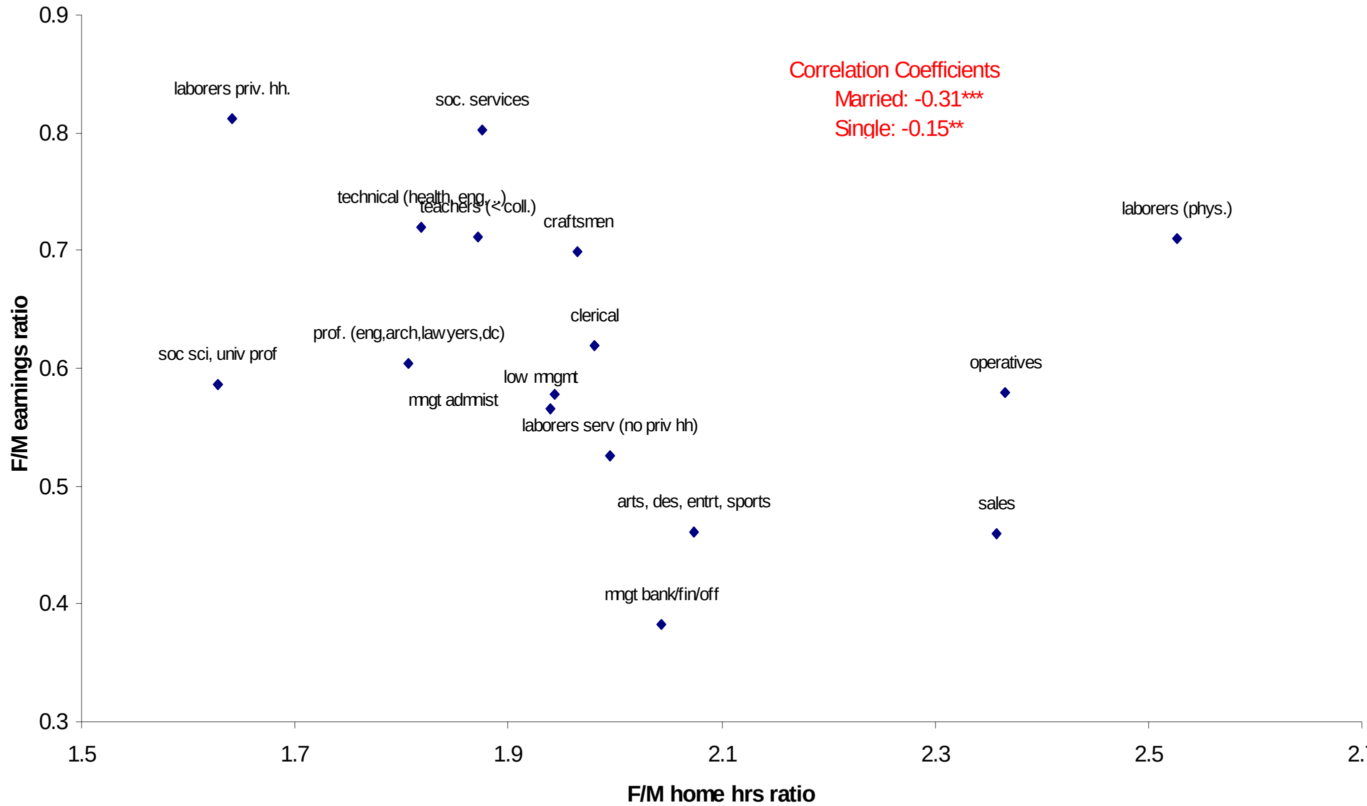

Figure 9: Correlation between F/M earnings ratio and F/M weekly home hours

and female workers and the female/male earnings ratio across industries and occupations. A corollary of this prediction is that husbands should on average receive a larger fraction of incentive pay relative to wives, conditional on their labor market characteristics. Direct evidence on these important predictions of our model would be important to discriminate our mechanism against alternative theories of gender differentials in labor market outcomes and to confirm the importance of feedback effects between household decisions and these outcomes. To explore these predictions, one would need a data set that includes information on the structure of earnings for a broad class of sectors and jobs, as well as, detailed household level information. To our knowledge, this information is not available for the U.S.[24]

Our model can also be used to think about how different sources of technological progress might impact gender differences in labor market outcomes and the household division of labor. For example, Greenwood, Seshadri and Yorukoglu (2005) have emphasized that new consumer durables introduced in the twentieth century acted as "engines of liberation" for women's time. The consumer durable revolution obviously also influenced the home production technology, in particular the degree of complementarity in spousal home hours. Similarly, improvements in medical technology that reduced the physical stress associated with pregnancy, as well as the introduction of breast milk substitutes, would determine a reduction in women's comparative advantage in home production. In Albanesi and Olivetti (2005), we use a version of our model to simulate the effect of these developments on women's employment and earnings and on households' division of labor.

Finally, in our framework, there are no efficiency losses associated with gender discrimination, given that all workers are equally productive in market work. We plan to extend the model to allow for a non-degenerate distribution of individual productivities, symmetric across genders. In an equilibrium with female discrimination, female workers with high productivity may be offered contracts in which they exert low effort despite their high productivity. This would generate misallocation costs associated with gender discrimination. This extension would allow an assessment of the welfare implications of gender discrimination.

References

[1] Albanesi, Stefania and Claudia Olivetti. 2005. Gender Roles and Technological Progress. Manuscript, Columbia University.

[2] Joseph G. Altonji, and Rebecca M. Blank. 1999. Race and Gender in the Labor Market. Handbook of Labor Economics, Volume 3, Elsevier.

[3] Benhabib, Jess, Richard Rogerson, and Randall Wright. 1991. Homework in Macroeconomics: Household Production and Aggregate Fluctuations. *The Journal of Political Economy*, Vol. 99, No. 6, 1166-1187.

[24] Suitable data is available for Norway.

[4] Babcock, Linda and Sarah Laschever. 2003. Women Don't Ask: Negotiation and the Gender Divide. Princeton University Press.

[5] Becker, Gary S. 1985. Human Capital, Effort and the Sexual Division of Labor. *Journal of Labor Economics*, Vol. 3, No.1, 33-58.

[6] Becker, Gary S. and Kevin Murphy. 2000. Social Economics: Market Behavior in a Social Environment. Harvard University Press.

[7] Bergstrom, Theodore. 1997. A Survey of Theory of the Family. Handbook of Family and Population Economics, Elsevier.

[8] Catalyst. 1995. Knowing the Territory: Women in Sales. A Catalyst Performance Report. Catalyst Publications, New York.

[9] Chiappori, Pierre-Andre'. 1988. Introducing Household Production in Collective Models of Labor Supply. *Econometrica*, Vol. 56, No. 1, 63-90.

[10] Chiappori, Pierre-Andre'. 1997. Introducing Household Production in Collective Models of Labor Supply. *The Journal of Political Economy*, Vol. 105, No. 1, 191-209.

[11] Coate, Stephen and Glenn Loury.1993. Will Affirmative Action Policies Eliminate Negative Stereotypes? *The American Economic Review*, Vol. 83, No.5, 1220-1240.

[12] Drogin, Richard. 2003. Statistical Analysis of Gender Patterns in Wal-Mart Workforce. Plaintiff's Expert Statistical Report, Dukes vs Wal-Mart, http://www.walmartclass.com.

[13] Francois, Patrick. 1998. Gender Discrimination without Gender Difference: Theory and Policy Responses. *Journal of Public Economics*, Vol. 68, No.1, 1-32.

[14] Goldin, Claudia. 1986. Monitoring Costs and Occupational Segregation by Sex: A Historical Analysis. *Journal of Labor Economics*, Vol. 4, 1-7.

[15] Goldin, Claudia. 1990. Understanding the Gender Wage Gap: An Economic History of American Women. Oxford University Press.

[16] Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu. 2005. Engines of Liberation. *The Review of Economic Studies* 72, 109-133.

[17] Holmstrom, Bengt and Paul Milgrom, 1991. Multitask Principal Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design. *Journal of Law, Economics and Organization*, Volume 7.

[18] Jullien, Bruno, Bernard Salanie and Francois Salanie. 2003. Screening Risk-Averse Agents under Moral Hazard: Single-Crossing and the CARA case. Manuscript, Columbia University.

[19] Laffont, Jean-Jack and David Martimort. 2002. The Theory of Incentives. The Principal Agent Model. Princeton University Press.

[20] Light, Audrey and Manuelita Ureta. 1992. Panel Estimates of Male and Female Job Turnover Behavior: Can Female Non-Quitters Be Identified. *The Journal of Labor Economics*, Vol. 10, No. 2, 156-181

[21] Lundberg, Shelly and Richard Starz. 1983. Private Discrimination and Social Intervention in Competitive labor Market. *The American Economic Review*, Vol. 73, No. 3, 340-347.

[22] MacLeod, W. Bentely and James M. Malcomson. 1988. Reputation and Hierarchy in Dynamic Models of Employment. *The Journal of Political Economy*, Vol. 96, No. 4, 832-854.

[23] MacLeod, Bentley and Daniel Parent. 2003. Job Characteristics and the Form of Compensation. *Research in Labor Economics*, Volume 18, pages 177-242.

[24] McElroy, Marjorie and Mary Jean Horney. 1981. Nash Bargained Household Decisions: Toward a Generalization of the Theory of Demand. *International Economic Review*, 22, 333-349.

[25] Manser, Marilyn and Murrey Brown, 1980, "Marriage and Household Decision Making: A Bargaining Analysis", *International Economic Review*, 21, 31-44.

[26] Milgrom, Paul and John Roberts. 1992. Economics, Organization and Management. Prentice Hall.

[27] Niederle, Muriel, and Lise Vesterlund. 2005. Do Women Shy Away from Competition? Do Men Compete too Much?. Manuscript, Stanford University.

[28] O'Neill, June. 2003. The Gender Wage Gap in Wages, circa 2000. *The American Economic Review, Papers and Proceedings.*

5 Appendix

5.1 Labor Contracts

The first order necessary conditions from Problem 1 are:

$$0 = f_n(n,1) - v_n(h,n,1) + \mu\left\{\tilde{w}\left[f_n(n,1) - f_n(n,0)\right] - v_n(h,n,1) + v_n(h,n,0)\right\}, \qquad (19)$$

$$\mu\left[f(n,1) - f(n,0)\right] - \sigma\Sigma\tilde{w} \leq 0, \qquad (20)$$

with equality at $\tilde{w} > 0$,

$$\left[\tilde{w}\left[f(n,1) - f(n,0)\right] - v(h,n,1) + v(h,n,0)\right]\mu = 0,\ \ \mu \geq 0,$$

μ is the multiplier on (6).

Proposition 3 *If the optimal contract with observable home hours implements $e = 1$ and (3) holds, $\tilde{w}$ is increasing in h.*

Proof. Conditions (7) and w=y(1-w) imply that the fraction of earnings given by incentive pay is $\tilde{w}$. We are interested in:

$$\frac{d\tilde{w}}{dh} = \frac{\partial \tilde{w}}{\partial h} + \frac{\partial \tilde{w}}{\partial n}\frac{dn}{dh}.$$

Condition (7) implies:

$$\frac{\partial \tilde{w}}{\partial h} = \tilde{w}\frac{v_h(h,n,1) - v_h(h,n,0)}{v(h,n,1) - v(h,n,0)} < 0,$$

by (2). In addition:

$$\frac{\partial \tilde{w}}{\partial n} = \tilde{w}\left[\frac{v_n(h,n,1) - v_n(h,n,0)}{v(h,n,1) - v(h,n,0)} - \frac{f_n(n,1) - f_n(n,0)}{f(n,1) - f(n,0)}\right].$$

Hence, under (3), $\frac{\partial \tilde{w}}{\partial n} \leq 0$ and $\frac{d\tilde{w}}{dh} \geq 0$. ■

5.2 Equilibrium

Lemma 4 *If $\pi(h|f) = \pi(h|m)$ for $h \in \{h_L, h_H\}$, it must be that $\pi(h_H|f) = 0.5$ and $\pi(h_H|m) = 0.5$ if $h_L < h_H$.*

Proof. Suppose not. Since households are homogeneous, they will all choose the same randomization strategy for the allocation of home hours. If this randomization strategy does not imply $\pi(h_H|m) = \pi(h_H|f) = 0.5$, then, the distribution of home hours is not the same for female and male workers. Hence, in equilibrium firms will not be offering the same menu of contracts to female and male workers. It follows that the equilibrium will be gendered. Contradiction. ■