Economic Research Initiative on the Uninsured
CONFERENCE DRAFT

HOUSEHOLD SEARCH AND HEALTH INSURANCE COVERAGE

Draft: Please do not cite or quote without permission.

Matthew S. Dey
BLS-DOL
2 Massachusetts Avenue, NE
Washington, DC 20012
dey.mathew@bls.gov

Christopher Flinn
New York University
269 Mercer Street
New York, NY 10003
christopher.flinn@nyu.edu

Economic Research Initiative on the Uninsured
University of Michigan
555 South Forest Street, 3rd Floor
Ann Arbor, MI 49104-2531

July, 2005

This research is supported in part by grants from the Economic Research Initiative on the Uninsured at the University of Michigan and the C.V. Starr Center for Applied Economics at New York University. James Mabli provided extremely helpful comments. We are responsible for all errors, omissions, and interpretations.

1 Introduction

Health insurance in the United States is typically acquired through an employer-sponsored program. Even though health insurance can be purchased through private markets, the cost is considered prohibitive in comparison with the effective cost of purchasing health insurance though an employer. There are many possible reasons for this difference, such as tax subsidies to firms who offer such insurance to their employees, risk-pooling among a large group of relatively healthy individuals (i.e., individuals employed at a given firm), or sharing of a cost (health insurance) that improves the quality of the employment match to both sides of the contract (e.g., Dey and Flinn, 2005).

Another empirical regularity regarding health insurance purchase and coverage is that in households in which both husbands and wives work health insurance is often only purchased (through their employer) by one of the spouses. Apparently this reflects the fact that health insurance is largely a public (household) good in that most employers who offer health insurance to their employees also include the option to cover spouses and dependent children. In this research our goal is to investigate the implications of the "publicness" of health insurance coverage for the labor market careers of spouses and the cross-sectional distribution of wages and health coverage statuses of spouses. We use a relatively innovative household search framework to address this question.

A large empirical literature exists on the relationship between health insurance coverage and wage and employment outcomes, though most of it is formulated at the individual level; reasonably comprehensive surveys can be found in Gruber and Madrian (2001) and Currie and Madrian (1999). The research objective in these studies is almost invariably the estimation of a distribution of marginal willingness to pay (MWP) parameters characterizing the population, and the framework is that of compensating differentials. When a formal modeling framework is developed, it is a variant of a static labor supply model, with reference made to household rather than individual choice on rare occasions. This is a questionable choice given the great deal of concern in this literature with assessing the impact of employer-provided health care coverage on job mobility. Dey (2001) and Dey and Flinn (2005) take the position that to analyze mobility behavior requires a dynamic model with labor market frictions, which led them to employ a search framework with both unemployed and on-the-job search. Estimates from the equilibrium matching-bargaining model in Dey and Flinn (2005) led them to conclude that the productive inefficiencies resulting from the employer-provided health insurance system were not large.

All of these studies suffer from their focus on individual rather than household behavior. A few attempts have been made to look at the impact of health insurance coverage of a spouse on the other's employment probability. For example, Wellington and Cobb-Clark (2000) estimate that having an employed husband with a job covered by health insurance reduces a wife's probability of employment by 20 percent. However, their econometric model does not allow for simultaneity in these decisions, labor market frictions, and does not even condition on the husband's wage rate. To understand the distribution of health insurance and wages across spouses and households it is necessary to formulate a more appropriate framework for the analysis.

To simplify the modeling and estimation problem, and to promote comparability with previous analyses, we adopt a very simple specification of household behavior. We assume the existence of a (instantaneous) household utility function in which the valuations of consumption and health insurance coverage are independent. The subutility function associated with consumption is a quasiconcave function of (instantaneous) household income, and the instantaneous payoff if at least one of the spouses has employer-provided health insurance is ξ. This "taste" for health insurance is what most studies attempting to estimate the MWP set as their goal.

Two extremely valuable papers make clear the perils of attempting to infer tastes from cross-sectional relationships generated by dynamic choices among jobs offering different combinations of utility-yielding characteristics. Hwang et al. (1998) make the point using the equilibrium search framework of Burdett and Mortensen (1998), and Gronberg and Reed (1994) provide an empirical example by estimating a MWP parameter within a compensating differentials model using job duration data from the National Longitudinal Survey of Youth 1979. The point of both of these studies is to illustrate how the cross-sectional relationship between wages and job characteristics is determined by the primitive parameters characterizing the search equilibrium. The cross-sectional "trade off" between wages and health insurance coverage, for example, is an extremely complicated function of ξ and the parameters characterizing the labor market environments of the spouses. The only way to consistently estimate ξ is to simultaneously estimate all model parameters, a path we follow in this paper.

The contributions of this paper with respect to those mentioned in the previous paragraph are (1) the extension to a multiple agent setting in which job attributes have a public goods aspect and (2) estimation of the behavioral model. We provide a lengthy discussion regarding the challenges of estimating a multiple agent model in continuous time given the discreteness of the data to which we have access. We use the method of simulated moments (MSM) in conjunction with data from the Survey of Income and Program Participation (SIPP) to estimate the model parameters. We find evidence that utility is a concave function of instantaneous income and that there is a positive valuation of health insurance coverage by the household. We show that this estimate is sensitive to the specification of the instantaneous utility function, as is to be expected. The estimate of the MWP for the preferred specification of the model is significantly less than that estimated by researchers using static linear regression methodologies. As we argue throughout the paper, those estimates should be viewed with suspicion for a variety of reasons.

Our model is very much partial equilibrium in nature, so that the types of policy experiments that can legitimately be performed with our estimates are limited. We focus on two, both involving changes in the way health insurance is supplied to the household. We first consider moves by employers to curtail the offering of family coverage, and assume that the offer of health insurance, if it is made at all, is extended only to the employee. In this case health insurance becomes a private good and alters properties of the decision rules and steady state equilibrium distribution of employment outcomes. In the second experiment, we attempt to mimic a situation in which coverage is provided universally, i.e., not through employers. We do this by setting the valuation of health insurance to zero (which indicates no willingness to pay with reduced wages). The results of both experiments are to be viewed with caution since we don't allow any adjustment in the wage-health insurance offer distributions we estimated. In equilibrium it is likely that these would adjust to the new institutional environments.[1]

The plan of the paper is as follows. In Section 2 we develop the model of household search using a household utility function approach. In the following Section we analyze the implications of the model under various specifications of the household utility function. Section 4 includes a discussion of the data source and presents some descriptive statistics. In Section 5 we develop the econometric model and discuss why cross-sectional, regression-based estimates of the MWP bear little relation to the true value of that function. Section 6 contains the estimates of model parameters, and Section 7 carries out the policy experiments described above. A brief conclusion is provided in Section 8.

[1] The framework of Hwang et al (1998), based on the equilibriurm search model of Burdett and Mortensen (1998) would be more appropriate to use. The downside of employing that model are some of its counterfactual empirical implications.

2 The Modeling Framework

In this section we develop our modeling framework and point out the innovations. Due to data limitations and for reasons of tractability, we assume that household preferences can be represented by a utility function with reasonably standard properties. In particular, we assume that the utility flow to the household is given by

$$U(I, d; Z, \gamma, \xi) = g(I; Z, \gamma) + \xi d,$$

where Z is a vector of household characteristics, assumed to be time invariant, I is instantaneous income of the household, d is an indicator variable that assumes the value 1 when *anyone* in the household purchases health insurance through their employer, g is a differentiable, concave function of I, γ is an unknown parameter vector, and ξ is a non-negative random variable the distribution of which can depend on Z. We assume that all household consumption is public in the sense that

$$I = w_1 + w_2 + Y_1 + Y_2,$$

where w_i is the instantaneous wage rate of spouse i and Y_i is the instantaneous receipt of nonlabor income of spouse i. As in most search-theoretic models, we ignore the capital market and assume that all income is consumed the same moment it is received.[2]

The labor market is structured as follows. When not employed spouse i receives offers of employment at a rate λ_i^N and while employed they receive offers at rate λ_i^E. When employed, spouse i is subject to "involuntary" dismissals at a rate η_i. Job opportunities are characterized by the pair (w, h), where w is the wage offer and h is an indicator variable that assumes the value 1 when the job offers health insurance. We do not assume that spouses draw from the same distributions; we denote the job offer distribution faced by spouse i as $F_i(w, h)$. Spouse i receives independently and identically distributed (i.i.d.) draws from F_i, and the wage draws of the two spouses are independently distributed conditional on observable and/or unobservable spouse-specific characteristics.

We denote the vector-valued state variable characterizing the household's decision problem by S, which includes $(w_1, h_1, w_2, h_2)'$;[3] the steady state

[2] This is in contrast to the model developed in Garcia-Perez and Rendon (2004). In their model of household search, which is set in discrete time, households are allowed to make savings decisions, though they are not allowed to borrow against future uncertain income.

[3] For ease of notation we omit the time invariant household characteristics Y and Z from the list of state variables.

value associated with the state vector S is given by $V(S)$. When spouse i is employed $w_i > 0$ and when not employed $w_i = 0$. While it is possible to write down one generic value function summarizing the problem for all possible values of S, doing so obscures some of the more interesting implications of the model regarding the relationship between the labor market decisions of the spouses. Thus we prefer to outline the features of each of the three qualitatively distinct decision problems faced by the household, corresponding to the cases in which zero, one, or two members are currently employed.

We begin with the more straightforward situation in which neither member is currently working. Consider a small decision period of length ε, during which at most one event can occur to the household (which in this case means that at most one of the unemployed spouses can receive a job offer). The value of the household's problem in this case is

$$\begin{aligned} V(0,0,0,0) &= (1+\rho\varepsilon)^{-1}\{g(Y)\varepsilon + \lambda_1^N \varepsilon \int \max[V(\tilde{w}_1, \tilde{h}_1, 0, 0), V(0,0,0,0)] dF_1(\tilde{w}_1, \tilde{h}_1) \\ &\quad + \lambda_2^N \varepsilon \int \max[V(0,0,\tilde{w}_2,\tilde{h}_2), V(0,0,0,0)] dF_2(\tilde{w}_2, \tilde{h}_2) \\ &\quad + (1 - \lambda_1^N \varepsilon - \lambda_2^N \varepsilon) V(0,0,0,0) + o(\varepsilon)\}. \end{aligned}$$

where for simplicity we have omitted the arguments Z and γ from the function g and where $o(\varepsilon)$ is a function with the property that $\lim_{\varepsilon \to 0} o(\varepsilon)/\varepsilon = 0$. Because $g(I)$ is a monotone increasing function of I, because w_1 and w_2 are perfect substitutes in consumption, and because $d = \max[h_1, h_2]$ in the case we consider here), it is straightforward to show that the decision problem faced by the household has a critical value property, that is, that there exists a function $w_i^*(h_i)$ that gives the minimal acceptable wage offer to spouse i given that the job has health insurance state h_i and given that both spouses are currently unemployed. Using this result, rearranging terms, and taking the limit of the function as $\varepsilon \to 0$ we can write

$$\begin{aligned} \rho V(0,0,0,0) &= g(Y) + \lambda_1^N \sum_{\tilde{h}=0}^{1} \int_{w_1^*(\tilde{h})} [V(\tilde{w}, \tilde{h}, 0, 0) - V(0,0,0,0)] dF_{w_1|h_1}(\tilde{w}|\tilde{h}) p_1(\tilde{h}) \\ &\quad \lambda_2^N \sum_{\tilde{h}=0}^{1} \int_{w_2^*(\tilde{h})} [V(0,0,\tilde{w},\tilde{h}) - V(0,0,0,0)] dF_{w_2|h_2}(\tilde{w}|\tilde{h}) p_2(\tilde{h}), \end{aligned}$$

where $F_{w_i|h_i}$ is the conditional distribution of wage offers given health insurance status for spouse i and p_i is the marginal distribution of health insurance statuses of job offers to spouse i. Our only comment concerning this particular value function is that it indicates that the job acceptance decisions of unemployed spouse i are a function of whether or not health insurance is offered. Moreover, the reservation wage rates for spouse i depend not only on the characteristics of spouse i's labor market environment but also on the labor market environment faced by the other spouse (who is also unemployed in this case). We also note that because we are assuming that g is concave, the critical values of each spouse depend on the level of nonlabor income. This would not be the case in the standard search framework in which g is assumed to be linear.

Next consider the situation in which one spouse is currently employed; let us assume that it is individual 1, at a job characterized by (w_1, h_1). Performing similar operations to what we did above, we can write the steady state value of this case as

$$(\rho + \eta_1)V(w_1, h_1, 0, 0) = g(Y + w_1) + \xi h_1$$
$$+\lambda_1^E \sum_{\tilde{h}=0}^{1} \int_{\hat{w}_1(\tilde{h}; w_1, h_1, 0, 0)} [V(\tilde{w}, \tilde{h}, 0, 0) - V(w_1, h_1, 0, 0)] dF_{w_1|h_1}(\tilde{w}|\tilde{h}) p_1(\tilde{h})$$
$$+\lambda_2^N \sum_{\tilde{h}=0}^{1} \int_{\hat{w}_2(\tilde{h}; w_1, h_1, 0, 0)} [\max[V(w_1, h_1, \tilde{w}, \tilde{h}), V(0, 0, \tilde{w}, \tilde{h})] - V(w_1, h_1, 0, 0)] dF_{w_2|h_2}(\tilde{w}|\tilde{h}) p_2(\tilde{h})$$
$$+\eta_1 V(0, 0, 0, 0).$$

The functions $\hat{w}_i(\tilde{h}; w_1, h_1, w_2, h_2)$ $i = 1, 2$, denote the critical value for job acceptance regarding a wage offer to spouse i associated with a health insurance status $\tilde{h}$ given a current job status of (w_1, h_1) for spouse 1 and (w_2, h_2) for spouse 2. (Thus $w_i^*(\tilde{h}) \equiv \hat{w}_i(\tilde{h}; 0, 0, 0, 0)$.) We note the following important points concerning this value function and the decision rules associated with it.

1. In this case, when spouse 1 is employed, the receipt of an offer by spouse 2 can result in three outcomes. Firstly, the offer can be rejected and the status quo maintained. Secondly, the offer can be accepted and spouse 1 can remain employed at his job, resulting in an outcome with value $V(w_1, h_1, \tilde{w}, \tilde{h})$. Thirdly, the offer could be accepted and spouse 1 could "quit" into unemployment, resulting in a value of $V(0, 0, \tilde{w}, \tilde{h})$. When a job offer is accepted by spouse 2, which of the last two possibilities occurs is determined by comparing the values associated with each of them. For example, a quit into unemployment by spouse 1 will

be relatively more likely when he is working at a low wage job without health insurance. To get any quits into unemployment, it must be the case that the rate of arrival of offers in that state be less than it is when employed. The estimates of primitive parameters we obtain confirm that $\lambda_i^N \gg \lambda_i^E,\ i = 1, 2$.

2. The critical values for spouse 1 have the following properties. When the health insurance status of the current job and the potential job are the same, then the critical wage rate is simply the current wage (since there are no mobility costs), or $\hat{w}_1(h_1; w_1, h_1, 0, 0) = w_1$ for $h_1 = 0, 1$. When the current job offers health insurance and the potential job doesn't, then $\hat{w}_1(0; w_1, 1, 0, 0) \geq w_1$, where the nonnegative "wedge" between the wages is a form of "dynamic" compensating differential. Conversely, we have $\hat{w}_1(1; w_1, 0, 0, 0) \leq w_1$ due to the value of gaining health insurance for household welfare.

3. Perhaps the most interesting feature of this case is the form of the unemployed spouse's decision rule. Say that the employed spouse's job offers health insurance so that his employment is characterized by $(w_1, 1)$. Even though the family has insurance coverage at this moment in time, it is not the case that the critical wage value for the unemployed spouse is independent of the health status of a job offered to her. In particular,

$$\hat{w}_2(1; w_1, 1, 0, 0) \neq \hat{w}_2(0; w_1, 1, 0, 0).$$

These values are not the same, in general, because having access to a job with health insurance has an "option value" even when the other spouse's current job also offers health insurance. This is due to the fact that the spouse may lose his job, either involuntarily (at rate η_1) or may have the opportunity to move to a high-paying job that does not offer health insurance. The only situation in which the inequality above will be an equality is when $\eta_1 = 0$ and $\lambda_1^E = 0$; in this case the first spouse will keep his current job forever and the family will have perpetual health insurance coverage. When this is not the case, we will have

$$\hat{w}_{2,1}(0; w_1, h_1, 0, 0) > \hat{w}_2(1; w_1, h_1, 0, 0),\ h_1 = 0, 1.$$

This is an important result since a number of empirical studies attempt to impute the implicit price of health insurance in terms of foregone wages by looking at average wage rates of spouses (possibly conditional on other covariates as well) given the health insurance status of their spouse (see, e.g., Olson (2001)). Even though the value of health insurance is less to a woman who has an employed husband paid w_1 when he has a job covered by health insurance than when he doesn't, she is still willing to pay for health insurance with a reduced wage rate. Thus the difference in average wages of these two groups of women does not represent a pure valuation of health insurance to the family. Furthermore, the average wage earned by a woman as a function of the current health insurance status of her husband depends on when she took her job (e.g., before the husband had accepted a job with health insurance, at a point when both held jobs in the past, at a point when her husband was unemployed, etc.). Thus labor market dynamics must be accounted for in assessing the valuation of health insurance to the household and its impact on labor market outcomes.

The last case to consider is when both spouses are currently employed. The steady state value of this case can be written as

$$(\rho + \eta_1 + \eta_2)V(w_1, h_1, w_2, h_2) = g(w_1 + w_2 + Y) + \xi \max[h_1, h_2]$$
$$+\lambda_1^E \sum_{\tilde{h}=0}^{1} \int_{\hat{w}_1(\tilde{h}; w_1, h_1, w_2, h_2)} [\max[V(\tilde{w}, \tilde{h}, 0, 0), V(\tilde{w}, \tilde{h}, w_2, h_2)] - V(w_1, h_1, w_2, h_2)] dF_{w_1|h_1}(\tilde{w}|\tilde{h}) p_1(\tilde{h})$$
$$+\lambda_2^E \sum_{\tilde{h}=0}^{1} \int_{\hat{w}_2(\tilde{h}; w_1, h_1, w_2, h_2)} [\max[V(0, 0, \tilde{w}, \tilde{h}), V(w_1, h_1, \tilde{w}, \tilde{h})] - V(w_1, h_1, w_2, h_2)] dF_{w_2|h_2}(\tilde{w}|\tilde{h}) p_2(\tilde{h})$$
$$+\eta_1 V(0, 0, w_2, h_2) + \eta_2 V(w_1, h_1, 0, 0).$$

Given the assumptions we have made regarding the household utility function, it is not difficult to establish the existence of these critical value functions and the following properties of these functions and $V(w_1, h_1, w_2, h_2)$.

1. If the health insurance status of a job offered to spouse i is the same as that of their current job then the critical wage is equal to the current wage, or
$$\hat{w}_1(\tilde{h}; w_1, \tilde{h}, w_2, h_2) = w_1, \ \tilde{h} = 0, 1.$$

2. Even when the other spouse is employed at a job with health insurance, the individual is willing to pay a "premium" for a job that includes health insurance. For example, say that spouse 2 is working at a job with health insurance and spouse 1 is not. Then
$$\hat{w}_1(1, w_1, 0, w_2, 1) < w_1,$$

and if both spouses currently have jobs that provide health insurance spouse 1 will have to be "compensated" for accepting a new job that does not include it,

$$\hat{w}_1(0, w_1, 1, w_2, 1) > w_1.$$

3. Isomorphic to these properties of the critical value function is the ordering of the value functions:

$V(w_1, 1, w_2, 1) > \max[V(w_1, 1, w_2, 0), V(w_1, 0, w_2, 1)] > V(w_1, 0, w_2, 0).$[4]

3 Analysis of the Model

Our modeling framework is of interest not only for the analysis of health insurance and labor market transition issues, but also can be thought of as a critique of single agent models of labor market mobility. Our claim is that previous single-agent models of labor market decisions will be misleading representations of the mobility process unless certain conditions hold. One of the goals of the empirical work reported below is to assess how misleading the single-agent models are likely to be.

Our instantaneous household utility function has the form

$$U(I, d; Z, \gamma, \xi) = g(I; Z, \gamma) + \xi d.$$

We consider the following special cases.

3.1 No valuation of health insurance and linear g.

This is the standard partial-partial equilibrium model of search; the only novelty in this case is the fact that there are two agents involved in the problem. But now we have

$$\begin{aligned} U(I, d; Z, \gamma, 0) &= \gamma(Z)I \\ &= \gamma(Z)(w_1 + w_2 + Y). \end{aligned}$$

Since nonlabor income is received by the household in any state of the world, and the marginal utility of income is constant, we have

$$V(w_1, w_2, Y) = \tilde{V}(w_1, w_2) + \frac{Y}{\rho}.$$

[4] The strict inequalities hold as long as the possibility of mobility (voluntary or involuntary) is positive for both spouses.

Given the constant marginal utility of income, no decision of spouse i can depend on the wage of spouse i'. Given this separability, we can write

$$\tilde{V}(w_1, w_2) = \tilde{V}_1(w_1) + \tilde{V}_2(w_2).$$

The value functions are indexed by spouse number since the search environments for the two are not constrained to be equal, where the search environments are characterized by $(\lambda_i^N, \lambda_i^E, \eta_i, F_i)$. An individual is better off within a household strictly due to income pooling. As a single agent, the value of spouse $i's$ problem is

$$\tilde{V}_i(w_i) + \frac{Y_i}{\rho},$$

where Y_i is that spouse's nonlabor income. The surplus i gets from being a member of the household is

$$\tilde{V}_{i'}(w_{i'}) + \frac{Y_i}{\rho}.$$

The others spouse's wage is simply another form of nonlabor income (albeit transitory in nature), and agent i's welfare is maximized by having the spouse act so as to maximize the expectation of the present value of their wage stream.

3.2 Valuation of health insurance and linear g

Health insurance is a very particular type of good. In our model consumption within the household is considered to be a public good, and health insurance is thought of in this way as well. This makes it fundamentally different than other components of a compensation package, such as the characteristics of one's office, the personalities of one's colleagues, etc., that yield a payoff which is primarily accrued to the individual employee. This, plus the fact that health insurance is such an important component of compensation in dollar value, makes it and pension benefits the preeminent parts of remuneration after wages and salary.

The instantaneous payoff function in the present case is given by

$$\gamma(Z)(w_1 + w_2 + Y) + \xi d.$$

At first glance it might seem that the arguments applied to the previous case applied here as well, i.e., that household would maximize welfare by having the spouses act in a totally "decentralized" manner. This would be true if the payoff function was given by

$$\gamma(Z)(w_1 + w_2 + Y) + \xi(h_1 + h_2).$$

In this case we would have

$$V(w_1, h_1, w_2, h_2, Y) = \tilde{V}_1(w_1, h_1) + \tilde{V}_2(w_2, h_2) + \frac{Y}{\rho}.$$

However, we have assumed that health insurance benefits are perfect substitutes, so that

$$d = \max(h_1, h_2).$$

In this case the decisions cannot be uncoupled in the sense that spouse $i's$ decision of whether to accept a job offer of (w_i, h_i) will depend on the health insurance status of the spouse, $h_{i'}$, as well as their own current wage and health insurance status. As was true above in the previous case, labor market decisions will be independent of nonlabor income given the constant marginal utility of income.

3.3 No Valuation of Health Insurance and Concave g

We now consider the form of the decision rules when the household does not value health insurance and when the marginal utility of income is decreasing. Since the payoff function is not separable in the arguments (w_1, w_2, Y), the value function can not be expressed as a sum of individual value functions either.

For spouse i currently employed at a job paying a wage of w_i, any offer greater than $w_i^*(w_i, w_{i'}, Y)$ will be accepted, where

$$w_i^*(w_i, w_{i'}, Y) = \begin{cases} w_i & \text{if} \quad w_i > 0 \\ w_i^*(0, w_{i'}, Y) & \text{if} \quad w_i = 0. \end{cases}$$

In other words, for a currently unemployed individual, the value of the spouse's current wage $(w_{i'})$ and nonlabor income Y are both arguments of the critical value function. A high value of Y permanently reduces the marginal utility of income at any pair of wages (w_1, w_2), which means that

$$\frac{\partial w^*(0, w_{i'}, Y)}{\partial Y} > 0$$

for all $(w_{i'}, Y)$. Note that the critical value in this case depends on the separate arguments $w_{i'}$ and Y rather than simply their sum, $w_{i'} + Y$. Even though instantaneous consumption of the household when i is not working is given by the sum, Y is permanent and $w_{i'}$ is transitory instead. By the structure of the model then

$$\frac{\partial w^*(0,x,x)}{\partial Y} \geq \frac{\partial w^*(0,x,x)}{\partial w_{i'}} > 0, \ \forall x > 0, \tag{1}$$

which is due to the differential degree of "permanence" attached to $w_{i'}$ and Y. We would have equality between the first and second terms in (1) when $\lambda_{i'}^E = \eta_{i'} = 0$, for in this case the wage of spouse i' would be as permanent as Y. The decision rule when i is already employed is, as is commonly the case, to accept any offer greater than the current one since household welfare is monotone increasing in the wages of both spouses.

3.4 Valuation of Health Insurance and Concave g

This is the most general case we consider, and is the focus of our empirical analysis. By extension of the previous arguments, particularly those related to the cases of linear g with positive ξ and concave g with $\xi = 0$, the critical value for job change is given by

$$\hat{w}_i(\tilde{h}; w_1, h_1, w_2, h_2, Y),$$

as defined previously (where we had omitted the argument Y since it is time invariant). In general, all arguments appear individually in the function and are necessary to characterize the turnover decision (where by turnover we also implicitly include the change from the unemployment to the employment state). In other words, the vector (w_1, h_1, w_2, h_2, Y) is a *minimal sufficient statistic* for the job acceptance decisions of household members.

3.5 Discussion

The differences in the properties of the objective functions in the four cases we have considered produce critical value functions that differ in terms of their arguments, as well as qualitative differences in household labor market histories. We present a summary of some of these differences in the following table.

Objective Function	**Arguments of w_i^***	**Simultaneous Change**	**Lower Wage**
$\alpha + \beta I$	w_i	*No*	*No*
$\alpha + \beta I + \xi d$	$w_i, \tilde{h}, h_1, h_2$	*Yes*	*Yes*
$g(I; Z, \gamma)$, g concave	$w_i, w_{i'}, Y$	*Yes*	*No*
$g(I; Z, \gamma) + \xi d$, g concave	$w_i, w_{i'}, \tilde{h}, h_1, h_2, Y$	*Yes*	*Yes*

By varying the assumptions regarding the objective function of the household, we will be able to trace out their impact on observed labor market behavior. We will be particularly interested in seeing how the standard assumption made in the literature of linear g and $\xi = 0$ compares against the others, all of which involve some jointness in the labor market decisions of the spouses.

The jointness of household decision-making is best illustrated through examining the reservation value functions, which we now do graphically. In this set of examples and the empirical work that follows we restrict the form of g. In particular, we assume that g has the Constant Relative Risk Aversion form, or

$$g(I; Z, \gamma, \delta) = \gamma(Z)\frac{I^\delta}{\delta},$$

where $\gamma(Z) > 0$, $\forall Z$ and $\delta \in [0, 1]$. As is well-known, in this case

$$\begin{aligned}\lim_{\delta \to 1} g(I; Z, \gamma, \delta) &= \gamma(Z) I \\ \lim_{\delta \to 0} g(I; Z, \gamma, \delta) &= \gamma(Z) \ln(I).\end{aligned}$$

This functional form allows us to nest the standard expected wealth maximization model as a special case.

In Figures 1.a-1.d we plot the reservation wage function for the wife when she is unemployed as a function of her employed husband's wage, an indicator variable for whether his job provides health insurance, and an indicator variable for whether her offered job provides health insurance. The four figures correspond to the four cases considered above. In plotting these functions we have used model estimates wherever possible.[5]

Figure 1.a contains the graph of the wife's (conditional) reservation wage function for the simplest case examined, the one with a constant instantaneous marginal utility from wages and no valuation of health insurance. The independence of the wife's decision rule from the husband's wage is reflected in the constant reservation wage function. This is the reservation wage she would set in a single agent model facing the labor market environment she faces. There is no dependence of this function on the health insurance status

[5] To date we have only estimated the most general version of the model, the one with an unrestrcted δ and a positive ξ.Thus for the general case (Figure 1.d), the graph corresponds to the decision rules generated by the model estimates. For the other cases, they do not. For example, in Figure 1.b we present the decision rules for the case of linear g but a positive valuation of health insurance. The valuation of health insurance we use is almost assuredly too low, since it corresponds to what is estimated for the concave g case. As a result, only qualitiative features of these graphs can legitimately be compared.

of her offer or her husband's job since the household does not value this job attribute in this case.

In Figure 1.b things get more interesting. The household is still assumed not to value employer-provided health insurance, but the (instantaneous) marginal utility of wages is decreasing in household income. As a result, there is one conditional reservation wage function for the unemployed wife, but it is no longer constant. Beginning at the reservation wage for an unemployed husband with an unemployed wife (since this is the lowest wage he would ever accept), we see a rapid increase in the function until when the husband's wage is approximately 10.30. After reaching that point, the function is still increasing, but at a slower, approximately constant, rate.

The reason for the differences in the properties of the function over these two intervals is the husband's response to the wife's accepting employment at the reservation value. At low wages, the husband quits his job and begins a spell of unemployed search. At higher wages, when the wife accepts the reservation value the husband continues employment in his relatively high paying job. Consider the case in which the husband is employed at the lowest acceptable wage, which is approximately 8.50. The wife's reservation wage at this point is approximately 5.90. If she receives a wage offer slightly greater than this value, the household will continue with one employed member, after substituting the wife for the husband. The utility level in the household changes markedly in this case, from $(8.50)^{.75}/.75$ to $(5.90)^{.75}/.75$, but the household is willing to make the trade off because the search environment of the husband dominates that of the wife on virtually every dimension. In the extreme case in which they both faced identical search environments, on this part of the function the reservation wage of the wife would be identical to the current wage of the husband, since neither would have any comparative advantage in search. In such an instance, the household should opt to have employed that partner with the highest current wage offer.

The final two figures examine the cases in which the household values employer-provided health insurance. As a result, there are four separate reservation wage functions in each figure, one for each combination of h_1 (the health insurance status of the husband's job) and $\tilde{h}$ (the health insurance status of the job offered to the wife). Figure 1.c plots these four functions for the case of a constant marginal utility of consumption. The only reason for the dependence of her decision rules on his job is through the public good aspect of employer-provided health insurance.

First consider the cases in which the husband's current job does not provide health insurance, $(0, 0)$ and $(0, 1)$. The value of a job offer to the wife that provides health insurance is quite high in this case, which is reflected in the difference in the reservation wage functions.[6] The other two conditional reservation wage functions correspond to the case in which the husband's job provides health insurance. Comparing these two functions reveals an interesting difference. Though the having a job with health insurance provides no gain in household welfare at the moment it is accepted (since the household is already covered by the husband's policy), the wife is willing to accept a lower wage for such a job than for one that doesn't offer her health insurance. The reason is the "option value" associated with having both spouses covered by health insurance; if the husband should lose his job or quit into one without health insurance (the likelihood of which depends on whether the wife's job is covered by health insurance), the household will still have coverage. This option value only exists with forward-looking household members.

Figure 1.d plots the four conditional reservation wage functions for the most general case. The qualitative properties of these functions have been discussed in presenting the other cases, so we won't belabor these issues any further.

3.6 Single versus Multiple Agent Search

The model we propose in this paper is novel along a few dimensions. First, it is one of the few search models that consider the case in which remuneration varies along more than one dimension (references). Second, and most importantly, it is the only one (besides Garcia-Perez and Rendon (2004)) to consider the case of simultaneous search by more than one agent.

Let us begin by considering the difference between the two cases when health insurance has no intrinsic value to the household, or $\xi = 0$. The state variables for the household are simply (w_1, w_2) under this assumption, and the value of the household problem is given by $V(w_1, w_2)$. The exogenous labor market processes (i.e., arrival rates, dismissal rates, and wage offer distributions) each spouse faces are invariant whether or not we consider the decision-makers in isolation or jointly. Without getting into unnecessary technicalities, we will posit that the equilibrium labor market process of an agent is the same in the two cases only when the decision rules are identical.[7]

[6] This difference would be even more appreciable if we had used a more appropriate value (i.e., higher) for ξ in computing these functions for the linear utilty case.

[7] Technicalities relate to requirements that the decision rules differ on sets of measurable sets of labor market states. For example, say that the rate of offers to unemployed married women was equal to 0. Then women would never enter the employoment state, and the reservation wages of married men would be identical in the separable and nonseparable cases since the wage of women would be identically equal to 0 at all points in time.

In our discussion of Case A, we noted that

$$V(w_1, w_2; Y) = \tilde{V}_1(w_1) + \tilde{V}_2(w_2) + \frac{Y}{\rho}. \quad (2)$$

Thus the value of search for the household is independent of the level of nonlabor income. Moreover, $\tilde{V}_1(w_1)$ defines w_1^* uniquely and $\tilde{V}_2(w_2)$ defines w_2^* uniquely. Since the labor market process of spouse i is defined by $(\lambda_i^N, \lambda_i^E, F_i, w_i^*)$, the labor market process for individual i is the same whether we use $V(w_1, w_2; Y)$ or $\tilde{V}_i(w_i; Y_i)$. Given our functional form assumption for U, (2) holds if and only if $\delta = 1$.

When $\delta \neq 1$, separability does not hold and we have to be clear about what is meant by the single agent search model. The general value of the problem is given by $V(w_1, w_2; Y_1 + Y_2)$. Define the single agent problem value (for spouse i) by Q. There are a number of ways in which this might be specified. Obvious ones are:

1. $Q_i(w_i; Y_i)$
2. $Q_i(w_i; Y_1 + Y_2)$
3. $Q_i(w_i; Y_1 + Y_2 + w_{i'})$.

Under (1) the other agent is ignored altogether, both their nonlabor income and their labor market earnings, if they are employed. Then the correct household payoff function,

$$\gamma(Z)\frac{(w_1 + w_2 + Y_1 + Y_2)^\delta}{\delta},$$

is replaced with the incorrect payoff function

$$\gamma(Z)\frac{(w_i + Y_i)^\delta}{\delta}.$$

The objective function is monotone increasing in w_i in either case, so that agent i will always accept a higher offer job given that they are currently employed, no matter what the wage of the other spouse or household nonlabor income level. Thus this feature of the decision rule does not change. What does change is the reservation wage required to terminate unemployed search. We know that the household utility maximizing value is given by the function $w_i^*(0, w_{i'}; Y_1 + Y_2)$, and that

$$\begin{aligned} w_i^*(0, w_{i'}; Y_1 + Y_2) &= w_i^*(0, 0; Y_1) \Leftrightarrow \delta = 1 \\ &\Rightarrow w_i^*(0, 0, Y_i) = w_i^*. \end{aligned}$$

A constant reservation wage rule is inconsistent with maximization of expected household welfare, and incidentally, yields a different labor market process for individual i than the one associated with the optimal reservation wage rule.

In situation (2) the situation is only marginally changed. Once again our attention focuses on the reservation wage rule. While the problem is now evaluated at the correct level of nonlabor income for the household, the wage process of the other spouse is not considered in setting the reservation wage. Once again, the reservation wage function is independent of the other spouse's wage, yielding a suboptimal policy (from the perspective of the correct objective function) and different labor market processes for the spouses than that implied by the optimal rule.

The last situation (3) is a bit more subtle. In this case, the current period payoff function is correctly evaluated at the value $w_1 + w_2 + Y_1 + Y_2$.As has been true throughout, any offer greater than the current one is accepted by individual i when employed, and we only must consider the reservation value used by the unemployed agent. For simplicity, and due to limitations imposed by the data, we have not been overly concerned with the specification of the nonlabor income process, and have assumed that the flow value is constant. If we include the earnings of the spouse as a form of nonlabor income and treat it as constant over time, we have misspecified the nonlabor income process and therefore will define suboptimal behavioral rules whenever w_i^* is not independent of nonlabor income.

Conversely, what if we treat the "other income" process of the household - that is, all income outside of the individual's labor earnings - as being a stochastic process? The correct decision rules imply a conditional earnings process for the household that is a function of the current wages of both spouses and total nonlabor income. From this we can form a marginal conditional "other income" process for spouse i given by

$$J_i(w_i, w_{i'}; Y).$$

For notational simplicity consider the single agent choice problem for spouse

1 when he is *unemployed* and set $\gamma(Z) = 1$. Using $J_1(w_1, w_2; Y)$, write

$$\rho\tilde{Q}_1(0; J_1(0, w_2; Y)) = \frac{(w_2 + Y)^\delta}{\delta}$$
$$+ \lambda_1^N \int_{w_1^*} [\tilde{Q}_1(w_1; J(w_1, w_2; Y)) - \tilde{Q}_1(0; J_1(0, w_2; Y))] dF_1(w_1)$$
$$+ \tau_1(0, w_2; Y) \int [\tilde{Q}_1(0; J_1(0, x; Y)) - \tilde{Q}_1(0; J_1(0, w_2; Y))] dR_2(x|0, w_2; Y)$$
$$+ \tau_2(0, w_2; Y)[\tilde{Q}_1(0; J_1(0, 0; Y)) - \tilde{Q}_1(0; J_1(0, w_2; Y))].$$

We have chosen to represent the "other income" process facing spouse 1, $J_1(0, w_2; Y)$, by the hazard rate functions $\tau_1(0, w_2; Y)$ associated with changes in other income resulting from changes in w_2 that do not result in unemployment for spouse 2; $\tau_2(0, w_2; Y)$ that gives the rate at which a nonzero value of w_2 changes to 0 (by convention, say $\tau_2(0, 0; Y) = 0$); and the spouse 2 conditional new wage distribution $R_2(\cdot|0, w_2; Y)$. Then equation (**??**) defines a value of w_1^* for a given J_1 process. Now if

$$\begin{aligned} R_2(x|0, w_2; Y) &= \chi[x > w_2]\chi[w_2 > 0]F_2(x) \\ &+\chi[x > w_2^*(w_1 = 0; Y)]\chi[w_2 = 0]F_2(x) \\ \tau_1(0, w_2; Y) &= \chi[w_2 = 0]\lambda_2^N + \chi[w_2 > 0]\lambda_2^E \\ \tau_2(0, w_2; Y) &= \chi[w_2 > 0]\eta_2, \end{aligned} \quad (3)$$

then the decisions of spouse 1 acting in isolation are the same as those obtained by solving the full household maximization problem. But obviously this is all sleight of hand; the other income process given in (3) is defined with reference to the optimal rules associated with the joint maximization problem. Any other time-varying income process (even one state dependent on current wage draws in the household) will lead to single-agent decisions not consistent with their behavior given household maximization.

We now briefly consider the situation when we bring health insurance back into the picture by allowing $\xi > 0$. Since we have assumed that the instantaneous payoff from health insurance

$$\xi \max(h_1, h_2),$$

we have made the utility function a nonlinear function of number of health insurance plans held by household members and induced a nonseparability in the household value function even when $\delta = 1$ (constant marginal utility of consumption). If in this case we had redefined the payoff from health insurance as $\xi(h_1+h_2)$, we would be able to define separable value functions just for g linear and $\xi = 0$, so that

$$V(w_1, h_1, w_2, h_2; Y) = \tilde{V}_1(w_1, h_1) + \tilde{V}_2(w_2, h_2) + \frac{Y}{\rho},$$

and the labor market processes for individual i implied by the decision rules associated with $\tilde{V}_i(w_i, h_i)$ would be the same as those implied by the joint optimization problem. By imposing the nonlinearity in this part of the payoff function, we have ruled out the possibility of separability, even when $\delta = 1$ is linear. The most general case considered, with concave g and $\xi > 0$, does not lead to a separable value function even with linear payoffs from household health insurance. Therefore the decision rules in this case will also differ in the single and joint maximization cases, as will the resulting labor market processes for the spouses.

4 Data and Descriptive Statistics

Data from the 1996 panel of the Survey of Income and Program Participation (SIPP) will be used to estimate the models. The SIPP interviews households every four months for up to twelve times, so that at a maximum a household will have been interviewed relatively frequently over a four year period. The SIPP collects detailed monthly information regarding individual household members' demographic characteristics and labor force activity, including earnings, number of weeks worked, average hours worked, as well as whether the individual changed jobs during the month. In addition, at each interview date the SIPP gathers data on a variety of health insurance variables including whether an individual's private health insurance is employer-provided and covers other household or non-household members.

The main advantage of using SIPP data is the richness of the data across individuals in the household and the relative ease of creating these links. This allows a detailed investigation of the relationship between spouse's labor market decisions that no other dataset allows. The main disadvantage of using the SIPP when investigating the effects of health insurance at the household level involves the inability to distinguish between the lack of coverage and the decision not to takeup coverage. For example, if the husband receives health insurance coverage through his employer that also covers his wife, we will not necessarily observe whether she has the option of purchasing health insurance through her employer.[8]

[8] While this is the case for the majority of waves (four month data collection intervals)

The sample used in the empirical work that follows is selected from original sample households that contain only one family and a married couple. Since we are using transition information in our empirical work, we select only those households (or families) that remain intact from the original interview to the eighth interview. In addition, we select households in which both spouses meet certain standard requirements at each point over the interview period. In particular, both spouses are aged between 20 and 54, not enrolled in school, not in the Armed Forces, not self-employed, not retired, not disabled, not a contingent worker, and not receiving welfare benefits. The application of these selection conditions limits our sample to 1,826 married couples.

Table 1 contains some descriptive statistics from the sample of households used in our empirical analysis. As we discuss in detail below, while the behavioral model sets out the relationship between the employment decisions of both husbands and wives, the parameters of the model can be estimated using data from the spouses separately. Therefore, Table 1 shows key labor market outcomes of husbands and wives in our sample unconditionally on the labor market status of their respective spouses. We should note that the health insurance coverage rate is the percent of *employed* husbands or wives who are covered by insurance provided by their *own* employers. The transition probabilities are simply the proportion of individuals who occupy the original state who exit that state at some point over a one-year period.

Three features of the data deserve further comment. First, husbands are much more likely to be employed than wives. Almost all the husbands in our sample are employed at the initial observation period, while only 76 percent of wives are employed. Second, conditional on being employed, husbands are much more likely to be covered by health insurance provided by their employers than wives are. Nearly 80 percent of husbands are covered by their own employer-provided health insurance coverage, while slightly under half of wives are covered by their own employer-provided health insurance coverage. Third, for both husbands and wives, the wages in insured jobs are significantly higher than the wages associated with uninsured jobs.

in the SIPP, there are periodic (twice over the four year duration of the survey) topical modules that would, in principle, allow us to more fully characterize the status of employer-provided health insurance for both spouses. In the core data, when an employed wife reports that she has health insurance through her husband's employer, we simply do not know whether health insurance is not available through her employer or she chooses not to purchase coverage that is available. Using the data from the topical module, we would be able to determine the wive's insurance status and takeup decision.

5 Econometric Issues

The model is parsimoniously characterized in terms of the parameter vector

$$\theta = (\lambda_1^N, \lambda_2^N, \lambda_1^E, \lambda_2^E, \eta_1, \eta_2, F_{w_1|h_1}, F_{w_2|h_2}, p_1, p_2, \gamma, \rho, \xi)',$$

where all the parameters have been previously defined. In this section we discuss issues connected with the estimation of this household model.

In previous research (Dey and Flinn, 2005), we estimated a single-agent equilibrium version of this model using a simulated maximum likelihood estimator. It is difficult to follow the same strategy in the two-agent case when using a continuous-time framework. As noted when describing the behavioral model, certain shocks will lead to simultaneous changes in the labor market status of both members of the household. For example, a wife at a low wage job (with or without health insurance) whose unemployed husband receives a sufficiently high wage offer (with health insurance, say), will quit her job at the same instant the husband accepts the offer. While the continuous time framework is a fiction, of course, there is no nonarbitrary way to "fix" this problem.[9]

There are several alternatives one can pursue in this situation. One obvious choice is to abandon the continuous time framework altogether in favor of a discrete time setting. This approach is not without its pitfalls, however, there being at least two serious problems. The first is the arbitrariness of the choice of decision unit. Given the characteristics of the SIPP data, the most obvious choice would be a monthly unit of analysis. Changes in labor market state of both spouses could then be considered simultaneous, creating making the filtered data more coherent with the theory. But once a time period is chosen, we have no model of multiple changes of state within a decision period. While two or more changes in labor market status of an individual within a given month may be rare, an even more serious time problem exists. That problem is the arbitrariness of the boundaries of the decision period. Say that we choose the first day of a calendar month as the beginning of a decision period and the last day of that month as the end of that period. Then if the wife accepts a new job on February 21 and the husband quits his job on February 28, those two changes in state would be considered simultaneous. However, if she accepted her new job on February 28 and he quits on March 1, those would be considered two independent

[9] For example, one might assume that any time spouses changed states with a period of length Δ, that the moves were simultaneous. In this case, estimated parameters will critically depend on the choice of Δ. Moreover, for large enough Δ we will observe changes of state of the same individual, which by definition cannot be made coincidentally.

events under this definition of decision periods. The general point is that the time aggregation scheme adopted is arbitrary, with impacts on estimates and inferences that are difficult to assess.

The other argument against "time aggregation" of the type required to map a continuous time process into a discrete time one is the impact on equilibrium outcomes. In a continuous time point process type model at most one event occurs at any given moment in time, and both agents respond to this same event. This allows us to avoid the multiple equilibria-type problems we encounter in the context of simultaneous move games. In a discrete time model, both agents may receive offers in a period (an event that has positive probability, in general). In our household utility case, in which we can think of there being one decision maker, there is no problem in defining a single optimal choice, in general. However, as we extend the model to look at household behavior when the spouses have distinct preference maps, we can easily produce examples of multiple equilibria in the simultaneous move context.[10]

We have chosen to estimate the model off of moments of the stationary distribution of labor market outcomes and the steady state transition function. By not using the "fine detail" of the individual event histories, we do not have to directly confront the lack of simultaneity issue that is apparent only in the individual level event history data.

The algorithm used in obtaining the estimates is as follows. Consider some particular sample path of one simulated history. A simulated history is a mapping from some fixed (over iterations) draws using a pseudo-random number generator and a value of the parameter vector θ_k into an event history for a household. We denote the r^{th} vector of pseudo-random draws by ψ_r, where the dimension of ψ_r is $L \times 1$ and L is large. Then the event history associated with the r^{th} replication when using parameter vector θ is

$$\Im_r(\theta) = J(\psi_r, \theta).$$

We then define outcomes as functions of the event history, and these outcomes ultimately are used to compute the simulated moments upon which the estimator is based. In particular, a *data mapping* is a function that maps characteristics of event history $\Im$ into point-sampled or transitional "data" x, and is given by $x = B(\Im)$. By plugging a number of independently generated event histories into B we create an artificial data set $\{x\}$. From this set of "observations" the simulated moments are calculated.

[10] For example, say two unemployed spouses receive offers of x and y, respectively. It is easy to find cases where two Nash equilibria exist in which the first accepts x and the second declines y or the first declines x and the second accepts y.

To give a concrete example, one of the moments used throughout is the proportion of married women who are employed in the steady state. To compute this moment we will need to measure whether a wife is employed at an arbitrarily selected point in the event history that is sufficiently far away from the initialization of the process. Without loss of generality, all simulated household histories begin at time 0 with both spouses unemployed. After fixing a value of the primitive parameters, θ_k, say, we generate events, such as offer arrivals or dismissals, and their impact on the state variable describing household labor market characteristics is determined by passing the events through the expected welfare maximizing decision rules. At a point $T \gg 0$, we look at the household's state, $(w_1(T), h_1(T), w_2(T), h_2(T))$. If the first element of x is the wife's labor market status at time T, then

$$x_1 = \begin{cases} 1 & \Leftrightarrow \quad w_2(T) > 0 \\ 0 & \Leftrightarrow \quad w_2(T) = 0 \end{cases} .$$

If we compute a total of R simulation histories evaluated at the parameter θ_k, then

$$E(x_1|\theta_k) = \operatorname*{plim}_{R\to\infty} R^{-1} \sum_{r=1}^{R} x_1(r),$$

where $x_1(r)$ is the value of x_1 in replication r.

In order to estimate the model we must attempt to match at least as many characteristics of the stationary distribution and transition function as there are primitive parameters. Let the dimension of θ be K. The moments and transition parameters generated from $\{x\}$ are a mapping given by Γ, or

$$\begin{aligned} Q(\theta_k) &= \Gamma(\{x(r;\theta_k)\}) \\ &= \Gamma(\{B(\Im_r(\theta_k))\}). \end{aligned}$$

Let the corresponding sample moments and transition parameters be given by q_S. Then we define the method of simulated moments estimator by

$$\hat{\theta}_{SMM} = \arg\min_{\theta} (Q(\theta) - q_s)' W (Q(\theta) - q_s),$$

where $\dim(Q(\theta)) = \dim(q_s) = L \geq K$, and W is a symmetric, positive definite weighting matrix that is $L \times L$.[11] Now the q_s are computed from a sample of size N. Given identification of the elements of θ, we have

$$\operatorname*{plim}_{N\to\infty, R\to\infty} \hat{\theta}_{MSM} = \theta.$$

[11] In the initial work, the weighting matrix is simply chosen to keep the scale of the moments roughly similar. Our weighting matrix is a diagonal matrix where the diagonal

Consistency requires that both the sample size and the number of simulation histories become indefinitely large due to the nonlinearity of the model. We believe that our sample size is large enough to satisfy the first requirement, and that the large number of simulated histories ($R = 5000$) satisfies the second.

In computing standard errors we do not rely upon asymptotic approximations. Instead we compute bootstrap estimates of the standard errors by varying the simulation set of draws. By not redrawing samples from the actual data we are underestimating the amount of variability in our estimates. While this is straightforward to do in principle, varying both simulation samples and data samples is extremely time intensive. Given the large N, we felt that ignoring sampling error in the moments and transition parameters from the data would not produce seriously misleading estimates of precision.

It is notoriously difficult to determine analytically whether a rather complicated nonlinear model such as this one is identified. From Flinn and Heckman (1982) we know that the c.d.f.s $F_{w_i|h_i}$, $i = 1, 2$, are not identified nonparametrically. We assume that they both are (conditional) lognormal distributions, which means that we must estimate 8 parameters (2 lognormal parameters for 2 spouses for 2 health insurance states). The marginal health insurance offer functions p_1 and p_2 are each characterized by 3 parameters. While estimation of ρ is in principle possible given some set of assumptions on g, we will not attempt to do so and will instead fix it using the prevailing interest rate. At present we assume that the population is homogeneous in the sense that all face the same set of primitive parameters describing the search environment. The households only differ in terms of the observed state variable y, household nonlabor income.

5.1 Estimation of the MWP for Health Insurance

Almost the entire empirical literature on the relationship between health insurance status and wages is based on cross-sectional analysis. Most papers use a linear regression approach in which a function of an individual's wage is regressed on whether or not the job provides health insurance coverage and a vector of conditioning variables designed to capture the value of the

elements serve to appropriately scale the moments. For example, the mean wage of husbands in insured jobs is 18.10 and the employment rate of wives is 0.499. In order to give these moments roughly the same weight in the objective function we divide the mean wage by 2 (scaled moment is then 9.05) and multiply the coverage rate by 10 (scaled moment is then 4.99).

individual to his or her employer. A few analyses have proposed instrumental variable estimators to potentially deal with the lack of independence between the disturbance in such a regression and the health insurance status of the job. None of these approaches are likely to lead to credible estimates of the MWP, most obviously due to the fact that the regression framework is an inappropriate one to use with an inherently dynamic phenomenon.

Gronberg and Reed (1994) and Hwang et al. (1998) provide instructive examples and analysis of the problem of inferring the MWP using cross-sectional regression methods when the individuals make job acceptance decisions in a job search environment. We will discuss the problem in the context of our specific application (in which the nonpecuniary characteristic is binary) and where the offer distribution is taken as fixed. We will begin with the simpler case of individual search. Let an individual have a linear payoff function,

$$u(w, h) = w + \xi h, \tag{4}$$

and assume (for now) that all labor market participants share a common value of ξ. In this setup, ξ measures the willingness to pay for health insurance, for an individual will be indifferent between any two jobs, with and without health insurance, such that

$$w + \xi = w',$$

where the first job (with wage w) is the one that includes health insurance. In our partial equilibrium search model, the searcher faces an exogenously-given (w, h) offer distribution, $F(w, h)$. But given (4), the payoff from a job (w, h) is given by the scalar random variable

$$\nu = w + \xi h.$$

Thus from the point of view of labor market decisions and the resulting labor market process, only the distribution of ν is relevant. The distribution of ν is given by M, and is a function of F and ξ. In particular, there exists a reservation value of ν, ν^* say, such that any job offer with an associated value of ν at least as great as ν^* will be accepted (by unemployed searchers) while any $\nu < \nu^*$ will be rejected. Then in this simple case (with binary h), we have

$$\nu^* = w^*(1) + \xi = w^*(0),$$

where $w^*(1)$ is the reservation wage associated with a jobs offering health insurance and $w^*(0)$ the reservation wage for a job without health insurance. In a model with homogeneous agents then and no measurement error in wages and health insurance status, the following consistent estimator of ξ is suggested by the analysis of Flinn and Heckman (1982). Define

$$\begin{aligned} \underline{w}(1) &= \min_{S_1}\{w_k\}_{k=1}^{N_1} \\ \underline{w}(0) &= \min_{S_0}\{w_k\}_{k=1}^{N_0}, \end{aligned}$$

where S_1 is the set of wage offers associated with jobs offering health insurance that were accepted by unemployed individuals in the sample, N_1 is the cardinality of that set, and S_0 and N_0 are similarly defined for the accepted jobs not offering health insurance. Then Flinn and Heckman (1982) show that

$$\operatorname*{plim}_{N_j \to \infty} \underline{w}(j) = w^*(j),\ j = 0, 1,$$

so that

$$\hat{\xi} = \underline{w}(0) - \underline{w}(1)$$

is a consistent estimator of ξ.

Cross-sectional regression-type estimators, which amount to differences in means in our application with a single binary nonwage characteristic, are not generally expressible simply in terms of reservation wages. Instead, we will think of the cross-sectional relationship between mean wages in the two subpopulations of jobs defined by health insurance provision as generated by the steady state equilibrium distribution of jobs. It is well known that with OTJ search in a stationary environment, the steady state distribution of ν is given by

$$R(\nu) = \frac{M(\nu)}{1 + \kappa\tilde{M}(\nu)},$$

where

$$\kappa = \frac{\lambda^E}{\eta}.$$

The relationship between the steady state density of ν and the offer distribution of ν is given by

$$r(\nu) = \frac{1+\kappa}{[1 + \kappa\tilde{M}(\nu)]^2} m(\nu).$$

It is reasonably immediate to go from the density of $r(\nu)$ to the steady state wage distributions associated with the two health insurance states. Since all the matters in terms of welfare is ν, there is no difference between the proportion of jobs providing health insurance given ν in the steady state and the proportion providing health insurance at ν from the offer distribution. Thus the probability of health insurance given a value of ν is determined as follows. If a firm offers health insurance given ν, then the wage offer is $\nu - \xi$. If the firm does not offer health insurance given ν the wage offer is ν. The likelihood of health insurance and a wage offer of $\nu - \xi$ is given by $f(\nu - \xi, 1)$, while the likelihood of no health insurance and a wage offer of ν is $f(\nu, 0)$.Then the probability of receiving health insurance given ν is

$$p(h = 1|\nu) = \frac{f(\nu - \xi, 1)}{f(\nu - \xi, 1) + f(\nu, 0)}.$$

The marginal probability of health insurance in the steady state is

$$p(h = 1) = \int p(h = 1|\nu) r(\nu) d\nu.$$

Then the conditional steady state distribution of ν given h is

$$r(\nu|h) = \frac{p(h|\nu) r(\nu)}{p(h)}, \quad h = 0, 1.$$

The mean of the steady state distribution of wage offers given h is

$$E_{SS}(w|h) = \int [\nu - h \cdot \xi] r(v|h] d\nu. \tag{5}$$

Using (5) we can look at the issue of bias in estimates of the willingness to pay using differences in the mean wages. The difference in means in the steady state (taken to represent the cross-section) is

$$\begin{aligned} E_{SS}(w|h &= 0) - E_{SS}(w|h = 1) \\ &= \int \nu r(\nu|h = 0) d\nu - \int \nu r(\nu|h = 1) d\nu + \xi. \end{aligned}$$

Then the difference in cross-sectional mean wages is a consistent estimator of the willingness to pay if and only if

$$\int \nu r(\nu|h = 0) d\nu = \int \nu r(\nu|h = 1) d\nu.$$

There is nothing in the construction of the model that suggests this condition should be satisfied, though it is possible to construct examples in which it is. Given estimates of the primitive parameters of the model, we can compute this expression and determine how badly biased the cross-sectional estimator of the MWP would be. We conclude this section with an example to fix ideas.

Example 1 *Let $\xi = 1$. To keep things simple, suppose the wage-health insurance offer distribution is discrete and assumes four values (with equal probability) and let $\kappa = 2$. The characteristics of the distributions of interest appear below.*

(w,h)	ν	$M(\nu)$	$R(\nu)$	$p(h\|\nu)$	$R(\nu\|h=1)$	$R(\nu\|h=0)$
$(2,1)$	3	.25	.10	1	.286	0
$(4,0)$	4	.50	.25	0	.286	.231
$(4,1)$	5	.75	.50	1	1.00	.231
$(6,0)$	6	1.00	1.00	0	1.00	1.00

Among the population in the health insurance state, there are only two observable wages, 2 and 4, with $p_{SS}(w = 2|h = 1) = .286$ and $p_{SS}(w = 4|h = 1) = .714$, so that the mean wage is 3.428 among those with health insurance. In the population without health insurance we have $p_{SS}(w = 4|h = 0) = .231$ and $p_{SS}(w = 6|h = 0) = .769$. Then the mean wage in this group is 5.538. The difference in means in this example is

$$\begin{aligned} E_{SS}(w|h &= 0) - E_{SS}(w|h = 1) \\ &= 2.11 \\ &> 1 = MWP. \end{aligned}$$

Thus this difference severely overestimates the individual's marginal valuation of health insurance provision.

The example serves to illustrate the point that there is little relationship between the cross-sectional differences in means and the MWP defined in terms of the utility function. By altering the offer distribution probabilities (*i.e.*, M), we could get the steady state differences to equal MWP, or be less than it, etc. The fundamental indeterminacy we are illustrating is *not* due to the fact that we are not utilizing an equilibrium search framework (in which the offer distribution is endogenous). In fact, Hwang et al. (1998) develop a parallel argument using the equilibrium search model framework of Burdett and Mortensen (1998).

The case of household search is a bit more involved, clearly. The main lesson we have learned from the model analysis presented above is that, even when health insurance coverage at the spouses' jobs are perfect substitutes in an instantaneous sense, they are not in a dynamic one. Moreover, we have shown that the estimation of a dynamic, *single-agent* model of labor market decisions will lead to inconsistent estimates of the parameters of that agent's labor market environment and preferences.

6 Empirical Results

(NOT AVAILABLE IN THIS DRAFT)

7 Experiments

Although we are not working within an equilibrium search framework, some suggestive exercises can be performed using our estimates if the contextual changes considered do not have "first order" effects on the parameters estimated, in particular the offer distributions.

We consider two stylized experiments, and evaluate their impact using the steady state labor market distributions which they produce given our estimates of primitive parameters. Since the focus of our analysis is health insurance provision and household search, both experiments involve a change in the way employer-provided health insurance is offered or perceived. In conducting both experiments, we use estimates from the most general of the model specifications we consider, in which the household payoff function is

$$U(w_1, h_1, w_2, h_2) = \frac{(w_1 + w_2)^\delta}{\delta} + \xi \max\{h_1, h\}.$$

We first ask what the wage distribution and employment rate would be by an individual and a household if there was no explicit valuation of having health insurance. We think about this situation arising due to population coverage due to the existence of a national health insurance system, for example. Thus individuals still value having health insurance, but it plays no role in employment decisions given the additive separability we have assumed. Clearly, labor income taxes associated with financing such a system may be expected to alter the wage offer distribution in unspecified ways, and we ignore these effects. Thus all results have to be interpreted with a great degree of caution.

In generating the steady state distribution for this experiment, we assume that no firms offer employer-provided health insurance since it is costly and redundant. Consequently, we derive decision rules given $\xi = 0$ and use the estimated gender-specific wage offer distributions that were associated with no health insurance offer in the baseline model.

The second experiment envisions changes in the health insurance offers made by employers. Instead of offering family coverage, which makes employer-provided health insurance a public good, we will imagine a case in which it becomes a purely private good. To incorporate this, we will assume that the objective function is given by

$$U(w_1, h_1, w_2, h_2) = \frac{(w_1 + w_2 + Y)^\delta}{\delta} + \frac{\xi}{2}(h_1 + h_2).$$

Thus only when *both* spouses have health insurance is the instantaneous payoff the same in the two cases. If $\delta = 1$, then the decisions of the spouses are independent, as was previously discussed. However, given that the estimated $\delta < 1$, the decisions of the spouses are still interrelated.

(RESULTS OF EXPERIMENT TO BE DISCUSSED)

8 Conclusion

(NOT AVAILABLE IN THIS DRAFT)

References

[1] Burdett, K. and D. Mortensen. "Wage Differentials, Employer Size, and Unemployment." *International Economic Review* 39 (1998): 257-273.

[2] Currie, J. and B. Madrian. "Health, Health Insurance, and the Labor Market," in *Handbook of Labor Economics, Volume 3C*, 1999: 3309-3416. Amsterdam: North Holland.

[3] Dey, M. "Welfare and Mobility Effects of Employer-Provided Health Insurance," Mimeo, U.S. Bureau of Labor Statistics, 2001

[4] Dey, M. and C. Flinn. "An Equilibrium Model of Health Insurance Provision and Wage Determination." *Econometrica* 73 (2005): 571-627.

[5] Flinn, C. and J. Heckman. "New Methods for Analyzing Structural Models of Labor Force Dynamics." *Journal of Econometrics* 18 (1982): 115-168.

[6] Gronberg, T. and W.R. Reed. "Estimating Worker's Marginal Willingness to Pay for Job Attributes Using Duration Data." *Journal of Human Resources* 29 (1994): 911-931.

[7] Gruber, Jonathan and Brigitte C. Madrian. (2001) "Health Insurance, Labor Supply, and Job Mobility: A Critical Review of the Literature," University of Michigan, ERIU Working Paper, 2001.

[8] Hwang, H., D. Mortensen, and W.R. Reed. "Hedonic Wages and Labor Market Search." *Journal of Labor Economics* 16 (1998): 815-847.

[9] Garcia-Perez, J.I., and S. Rendon. "Family Job Search and Consumption: The Added Worker Effect Revisited." Mimeo, University of Western Ontario, May 2004.

[10] Olson, C. "A Comparison of Parametric and Semiparametric Estimates of the Effect of Spousal Health Insurance Coverage on Weekly Hours Worked by Wives." Journal of Applied Econometrics 13 (1998): 543-565.

[11] Rendon, S. "Job Search and Asset Accumulation under Borrowing Constraints." Forthcoming, *International Economic Review*, 2005.

[12] Royalty, A. and J. Abraham. "Health Insurance and Labor Market Outcomes: Joint Decision-Making within Households." University of Michigan, ERIU Working Paper, January 2005.

[13] Wellington, A. and D. Cobb-Clark. "The Labor-Supply Effects of Universal Health Coverage: What Can We Learn From Individuals with Spousal Coverage?," in S. Polachek, ed., *Worker Well-Being: Research in Labor Economics, Volume 19*, 2000. Amsterdam: Elsevier.

Figure 1a: delta = 1, csi = 0

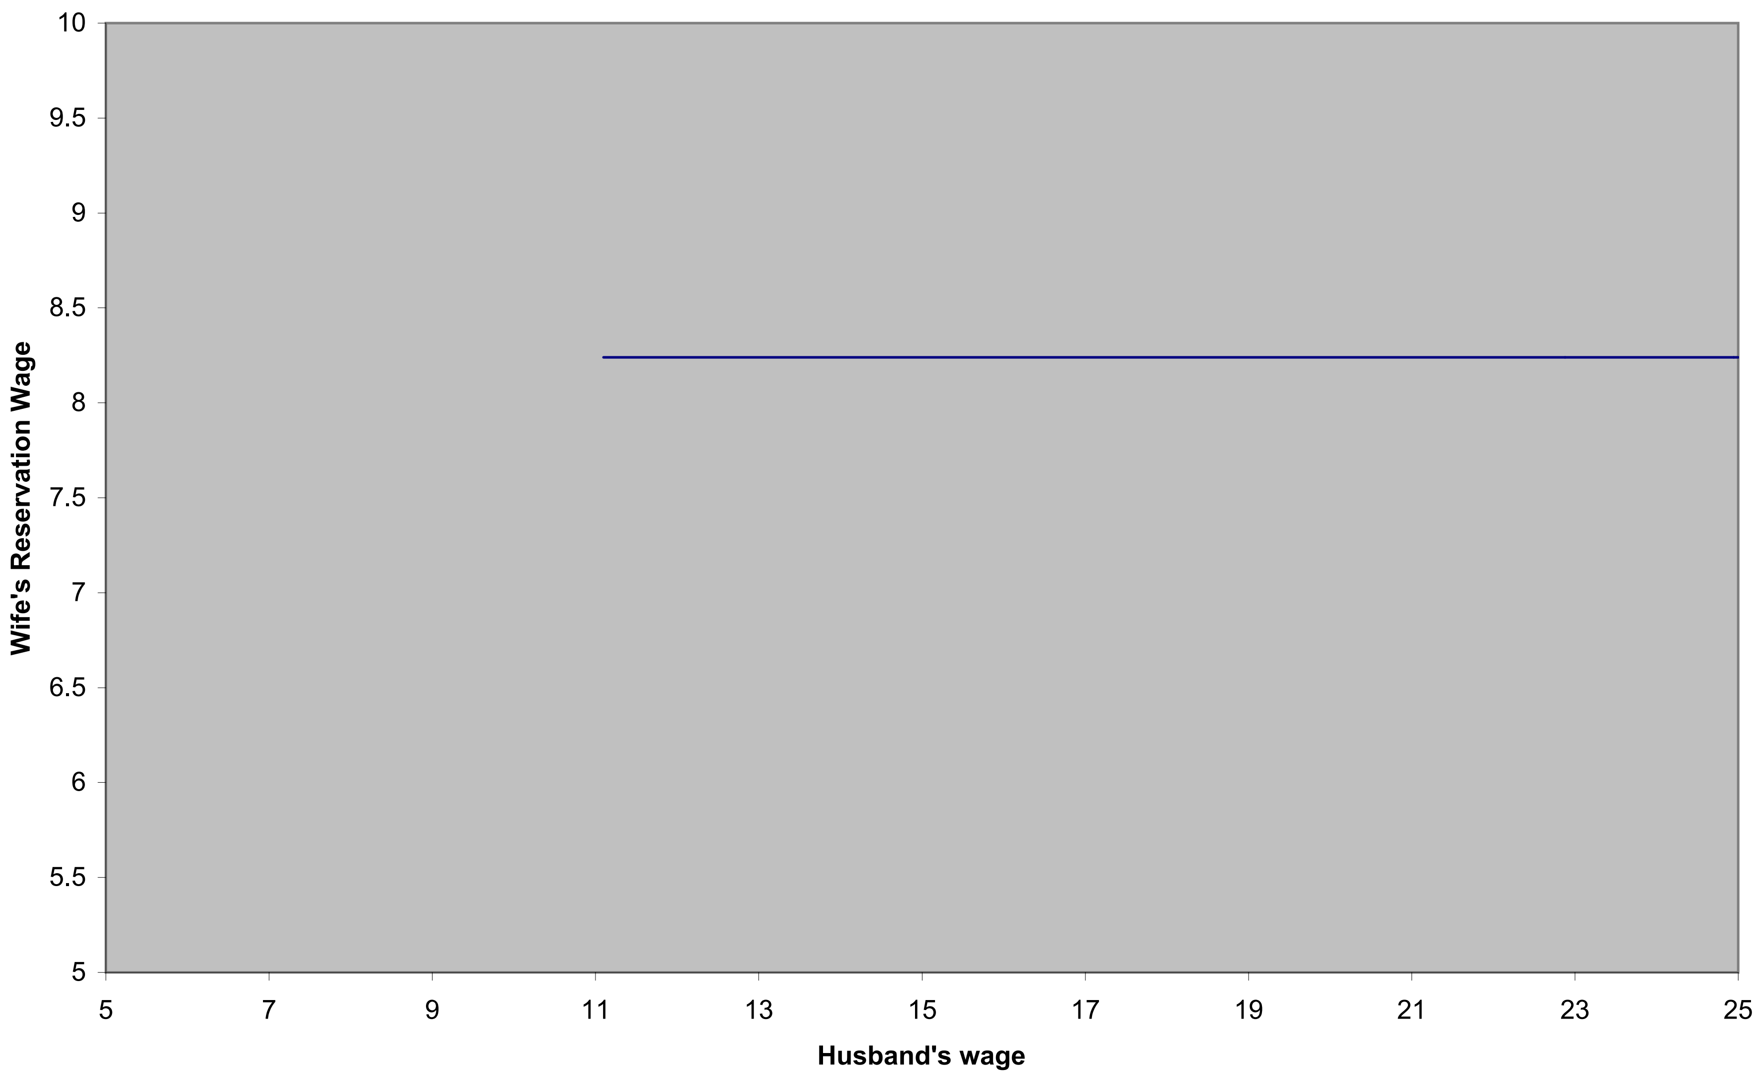

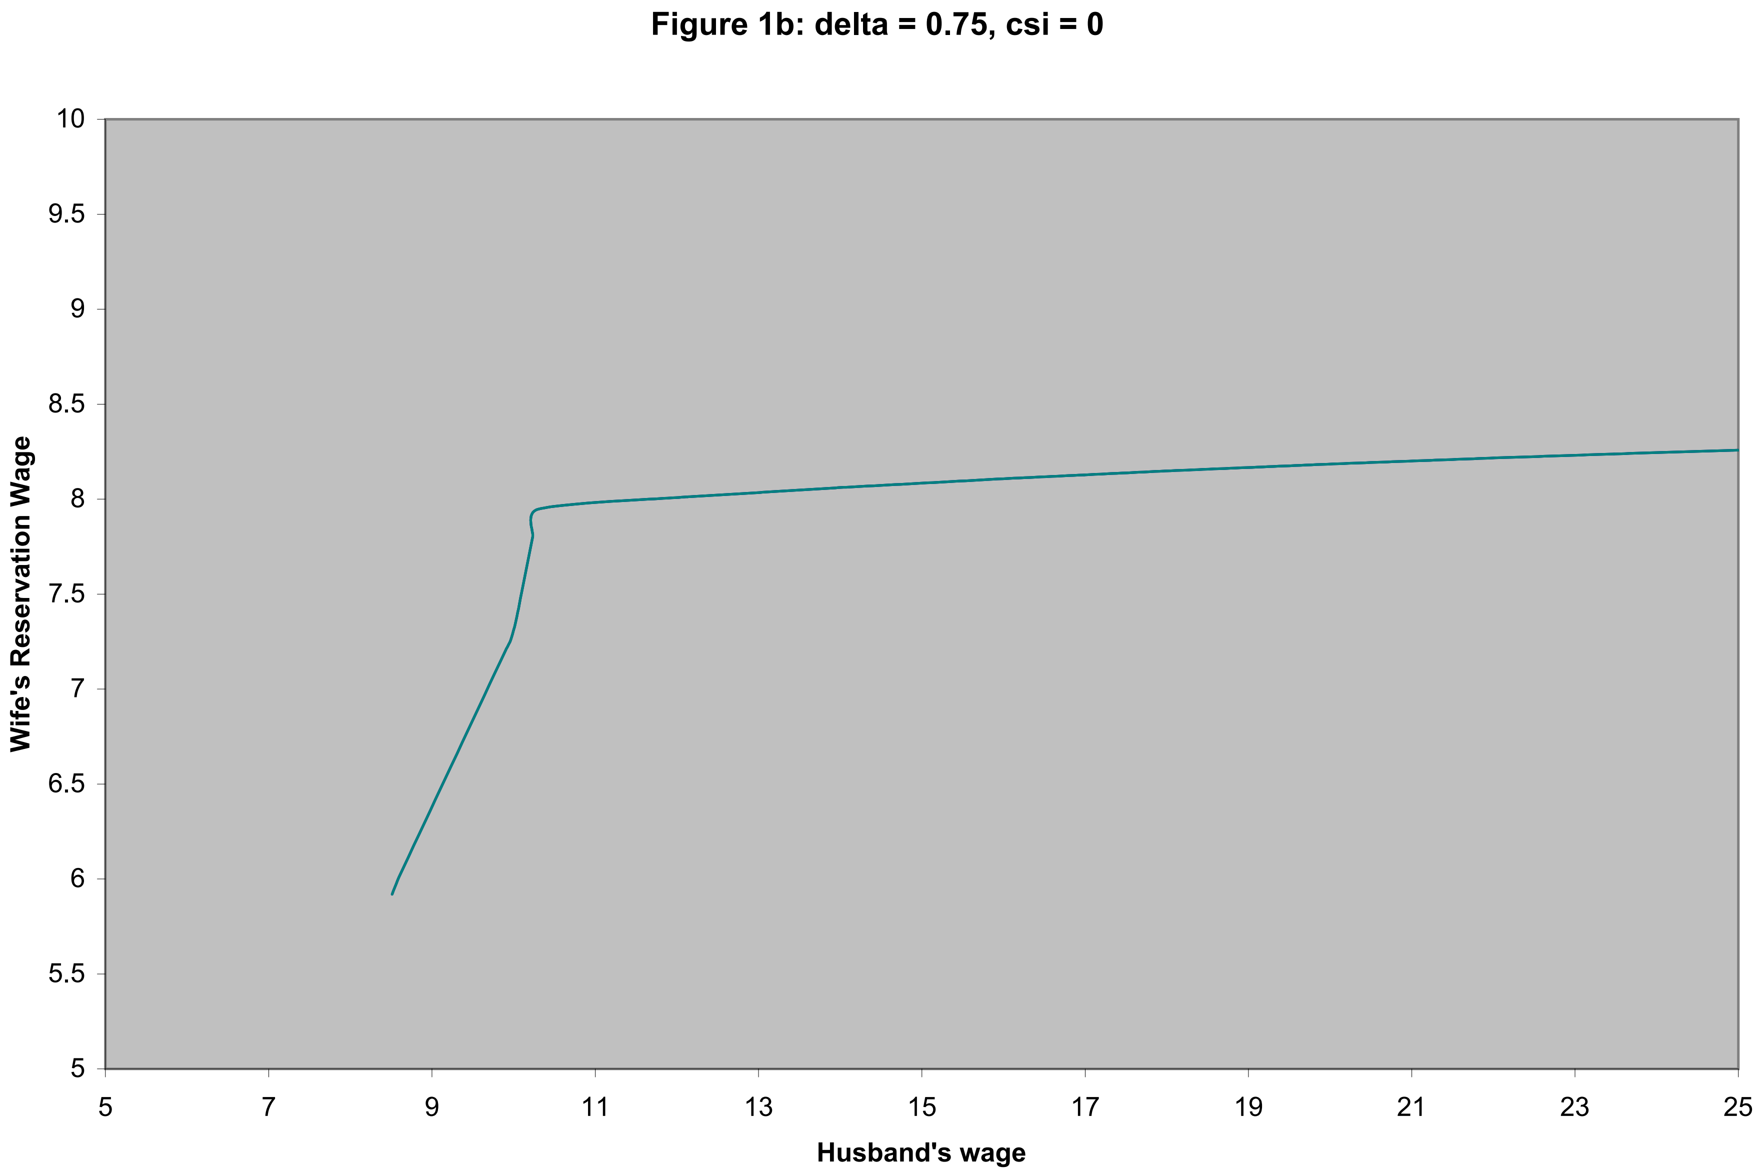

Figure 1b: delta = 0.75, csi = 0
Wife's Reservation Wage
Husband's wage
10
9.5
9
8.5
8
7.5
7
6.5
6
5.5
5
5
7
9
11
13
15
17
19
21
23
25

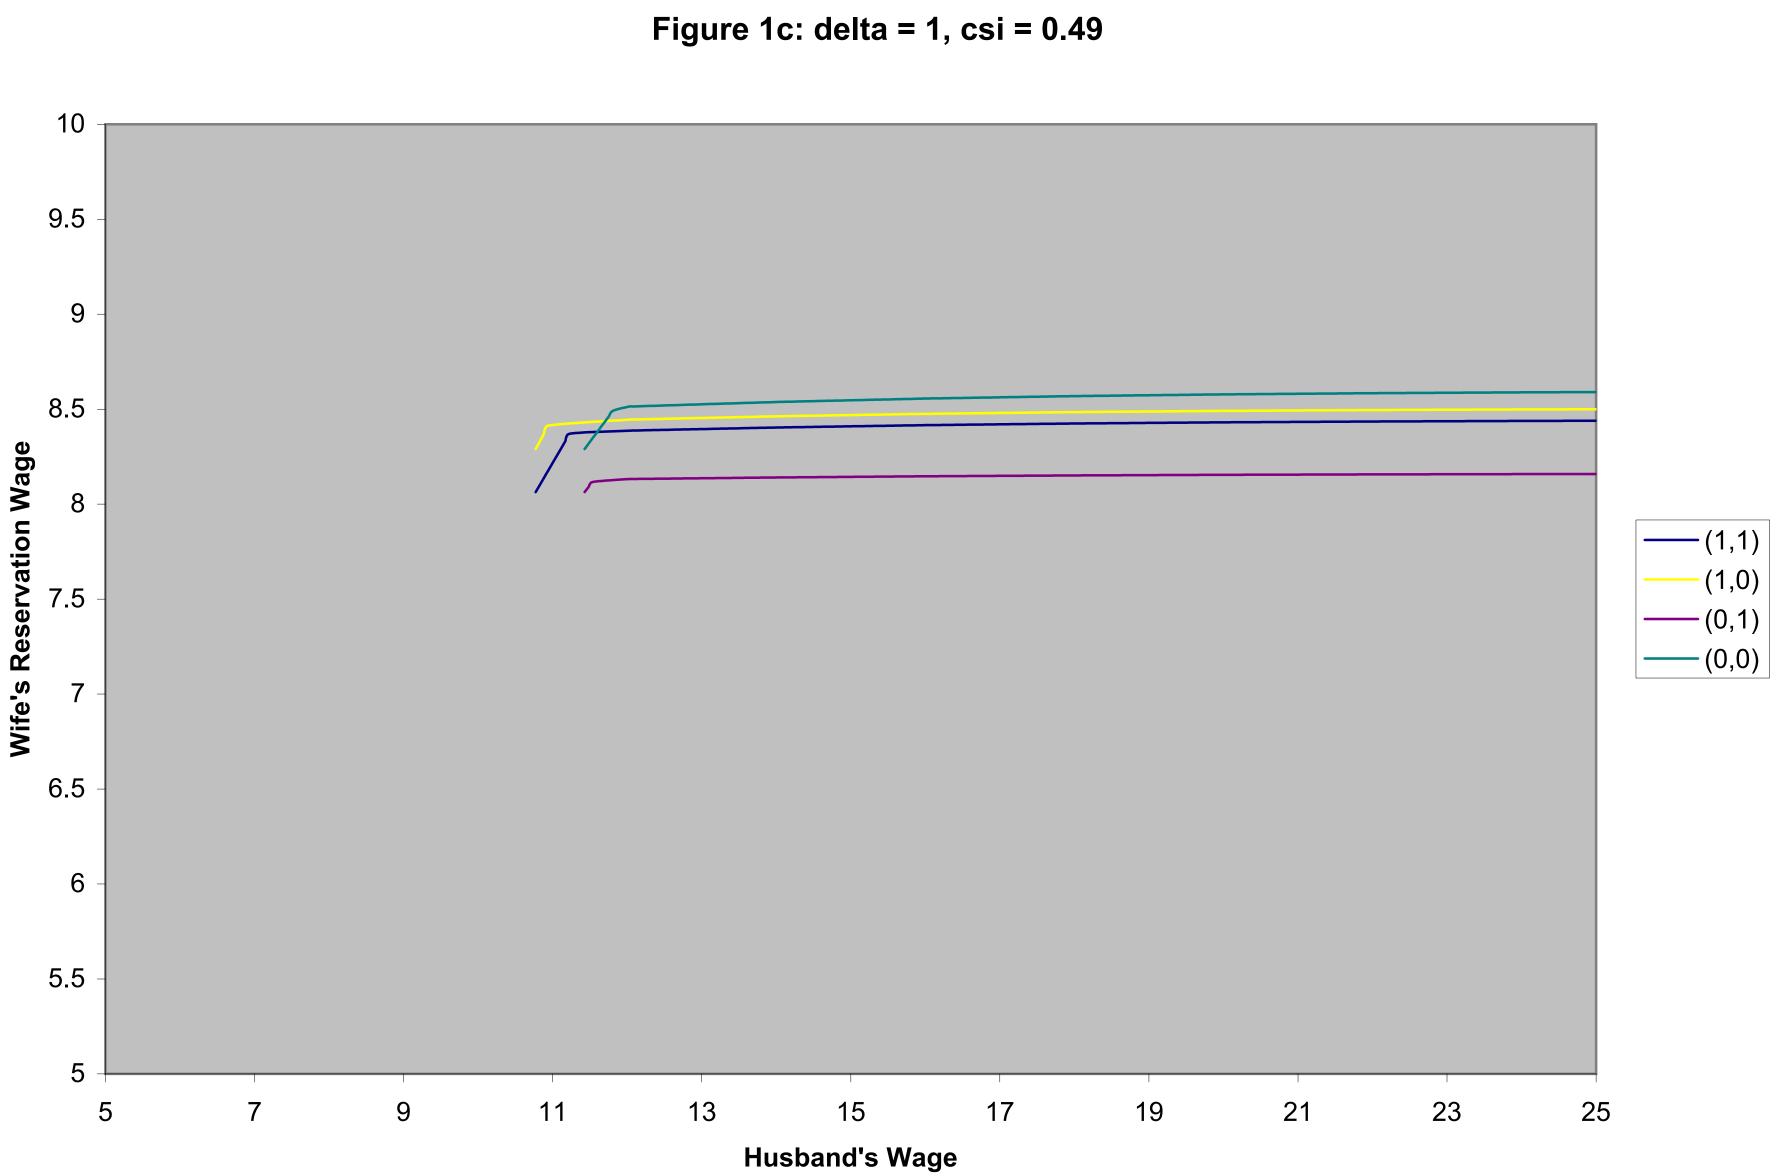
Figure 1c: delta = 1, csi = 0.49
Wife's Reservation Wage
Husband's Wage
10
9.5
9
8.5
8
7.5
7
6.5
6
5.5
5
5
7
9
11
13
15
17
19
21
23
25
(1,1)
(1,0)
(0,1)
(0,0)

Figure 1d: delta = 0.75, csi = 0.49

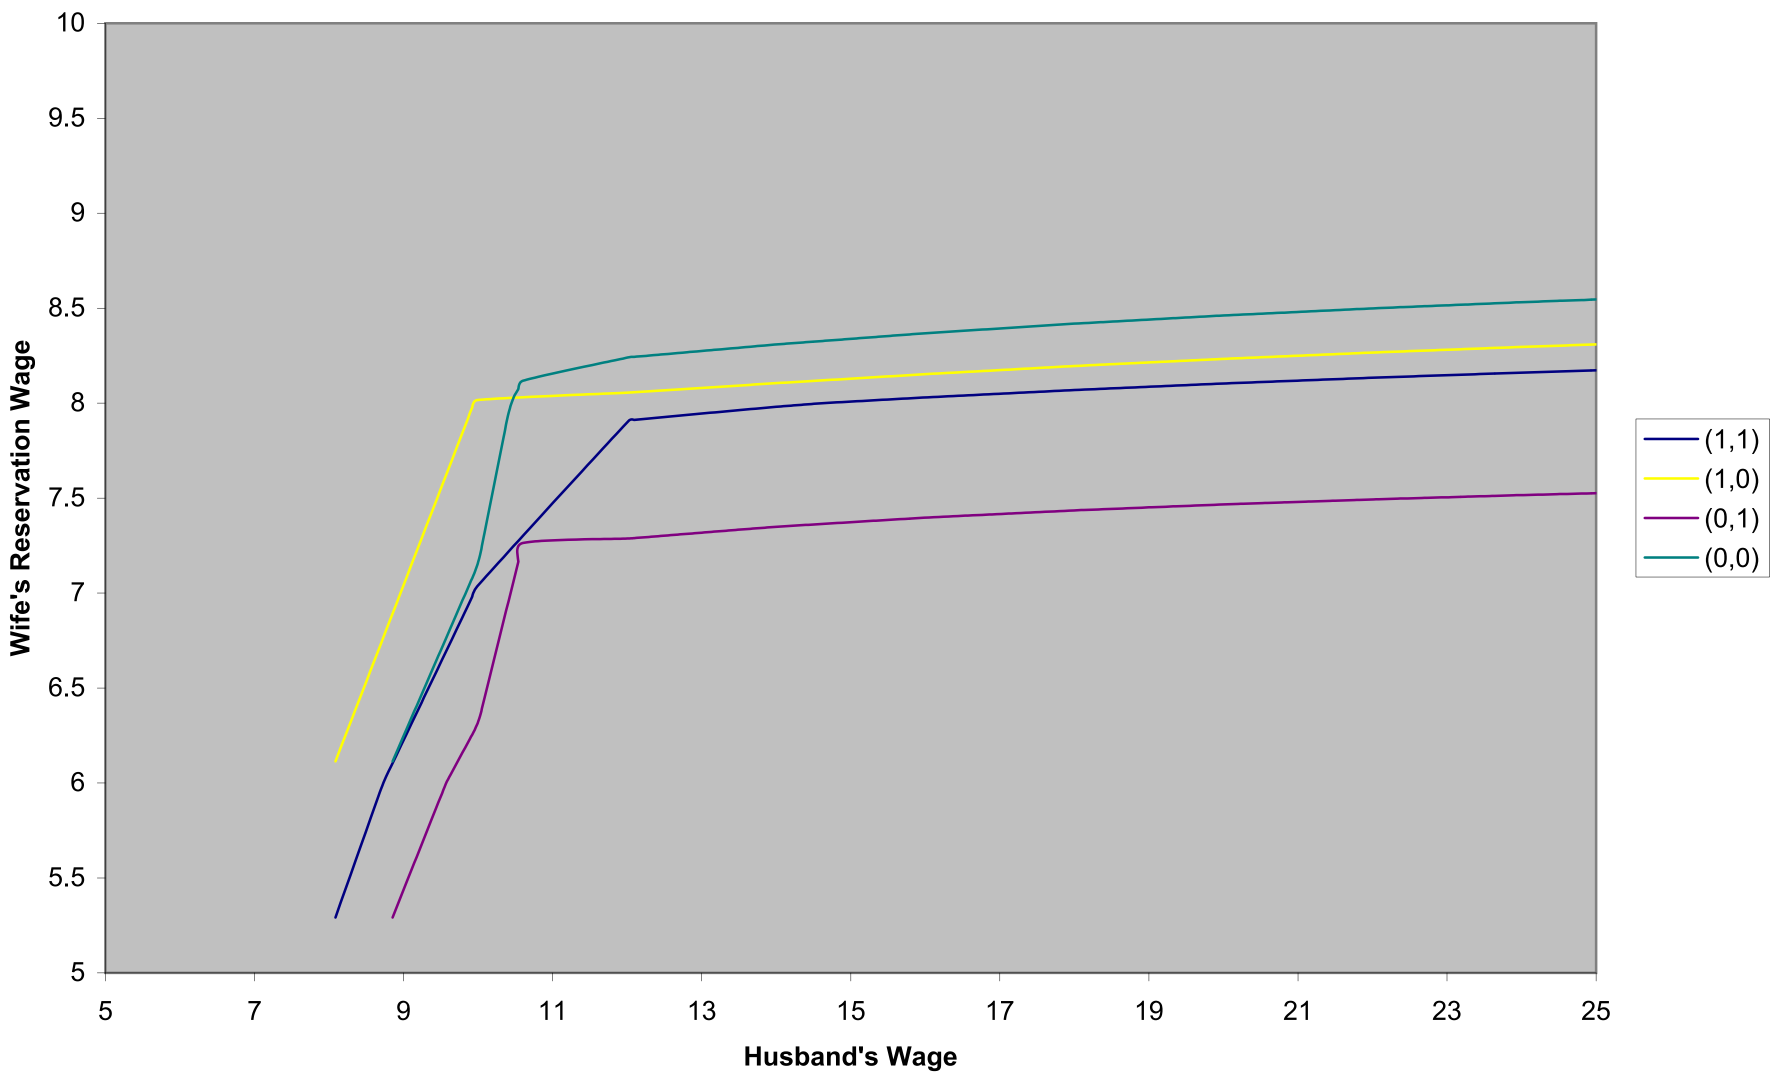

Table 1: Summary Statistics

Statistic	Husbands	Wives
Employment Rate	98.58	76.51
Health Insurance Coverage Rate	80.17	49.89
Mean Wage in Insured Jobs	18.10	14.45
	(7.73)	(6.97)
Mean Wage in Uninsured Jobs	15.06	12.09
	(6.90)	(5.94)
Probability of Transition out Unemployment	66.67	16.06
Probability of Transition out of Uninsured Job	40.67	30.65
Probability of Transition out of Insured Job	14.25	18.92

Note: Based on the 1996 panel of the Survey of Income and Program Participation (SIPP). The sample includes 1,826 married couples that meet certain selection criteria. Standard deviation of the various conditional wages are in parentheses.

Table 2: Parameter Estimates

Parameter	Husbands	Wives
Job Offer Arrival Rates in Unemployment	0.076	0.054
Job Offer Arrival Rates in Employment	0.012	0.008
Dismissal Rates	0.005	0.008
Probability Job Offer includes Health Insurance	0.713	0.467
Insured Jobs Wage Distribution Location Parameters	2.146	2.089
Uninsured Jobs Wage Distribution Location Parameters	2.066	1.983
Wage Distribution Shape Parameter	0.550	
Value of Health Insurance Coverage	0.489	
Utility Function Shape Parameter	0.75	
Reservation Wages in Insured Jobs	8.09	5.29
Reservation Wages in Uninsured Jobs	8.85	6.11

Note: Parameter estimates based on 1996 panel of the Survey of Income and Program Participation. The model is estimated using the simulated method of moments estimator described in the text. Non-labor income is set to 0 and the discount rate is set to 8 percent annually.

Table 3: Predicted Summary Statistics

Statistic	Husbands	Wives
Employment Rate	82.84	77.08
Health Insurance Coverage Rate	76.77	51.22
Mean Wage in Insured Jobs	17.42	14.28
	(6.69)	(5.99)
Mean Wage in Uninsured Jobs	17.73	13.49
	(7.17)	(5.20)
Probability of Transition out Unemployment	61.85	62.01
Probability of Transition out of Uninsured Job	29.22	38.80
Probability of Transition out of Insured Job	29.93	39.41

Note: Based on the parameter estimates in Table 2.

Table 4: Predicted Summary Statistics - Private Coverage

Statistic	Husbands	Wives
Employment Rate	84.96	75.88
Health Insurance Coverage Rate	77.87	54.40
Mean Wage in Insured Jobs	17.39	14.15
	(6.73)	(6.07)
Mean Wage in Uninsured Jobs	17.50	13.61
	(6.35)	(5.91)
Probability of Transition out Unemployment	67.46	64.05
Probability of Transition out of Uninsured Job	29.44	39.27
Probability of Transition out of Insured Job	27.71	38.95

Note: Based on the parameter estimates presented in Table 2.

Table 5: Predicted Summary Statistics - "Universal" Coverage

Statistic	Husbands	Wives
Employment Rate	83.08	75.20
Mean Wage	17.46	14.05
	(6.38)	(5.93)
Probability of Transition out Unemployment	67.32	61.20
Probability of Transition out of Employment	31.62	40.35

Note: Based on the parameter estimates presented in Table 2.