Estimating Marginal Treatment Effects of Transfer Programs on Labor Supply

Robert Moffitt*

Johns Hopkins University

October, 2014

Abstract

The standard static model of the effect of welfare program participation on labor supply is estimated allowing the effect of participation to be heterogeneous in the population and allowing the program participation decision to be endogenous and to exhibit incomplete takeup. Nonparametric methods are used to estimate the distribution of the marginal treatment effect over the range of participation rates generated by the instruments, which are measures of the non-financial costs of program participation. The results show that those with the greatest negative labor supply effects of program participation enter first and, as participation expands, individuals with smaller labor supply disincentives are drawn into the program.

*The author would like to thank Kai Liu and Marc Chan for research assistance. Comments on an earlier version of this paper from Joe Altonji, Mike Keane, and Thierry Magnac are appreciated. Research support from the National Institutes of health is gratefully acknowledge.

There is a long and extensive literature on estimating the effects of welfare programs on labor supply. Studies in the early 1970s estimating labor supply equations with cross-sectional data to forecast the effects of welfare programs (Cain and Watts (1973)) were followed by results from a series of negative income tax experiments (Burtless (1987),Moffitt and Kehrer (1981)). A literature using more sophisticated econometric methods to handle the selectivity of welfare participation developed in the 1980s (Burtless and Hausman (1978),Hausman (1981),Moffitt (1983)) which was followed by a literature of reduced-form estimates of the effects of specific welfare reforms in the 1980s and 1990s (Gueron and Pauly (1991),Grogger and Karoly (2005)). Reviews of these literatures have been published by Danziger et al. (1981), Moffitt (1992), and Blundell and Macurdy (1999). For the most part, the parts of this literature which aim to estimate the effects of welfare tax rates and guarantees on labor supply have found them to be in the expected direction, negative in both cases.

This work mostly predates the literature on identification and estimation of treatment effects developed in the last decade and a half. The feature of this newer literature which is the focus of this paper is the recognition of the potential importance of heterogeneity in response to treatment. Individual heterogeneity in response was permitted in the LATE model of Imbens and Angrist (1994) and is implicit in the work of Rubin (1974)). In the presence of heterogeneity, the average treatment effect need not equal the effect of the treatment on the treated or the marginal treatment effect, distinctions made mostly formally by Heckman and Vytlacil (2005) and Heckman et al. (2006). In the light of this framework, estimates of the effect of labor supply in the older literature have to be interpreted either as estimates under the assumption of homogeneous effects or else as some average of marginal treatment effects over the range of participation rates in the data.

This paper revisits the older literature allowing for individual heterogeneity in response. The paper uses the same static cross-sectional labor supply model used in the older literature to motivate the exercise, although the model used here is less structural because it makes labor supply a function of a dichotomous program participation indicator rather than a function of the budget constraint variables themselves. The more structural approach is left for future work. The response parameter representing the effect of participation on labor supply is assumed to be heterogeneous and to have a distribution to be nonparametrically estimated from the data. That parameter is set up as a random coefficient on the participation indicator, an approach introduced by Bjorklund and Moffitt (1987) and Heckman and Robb (1985). Unlike the former paper, however, which assumed heterogeneity to be normally distributed and imposed other distributional assumptions on the model, here the unobservables are allowed to be distribution-free. Their distributions are estimated with conventional series approximation methods.

In addition to allowing preference heterogeneity, the model allows incomplete welfare takeup, which means that welfare participation is a separate, endogenous choice to the individual. Moffitt (1983) first showed that not all eligible individuals participate in welfare programs and ascribed non-participation to stigma costs, although a more general interpretation allows time, money, and "hassle" costs. With a separate participation equation, exclusion restrictions are needed to identify the model and the natural set of exclusion restrictions are, in fact, measures of the cost of participation. The application here uses state-level measures of non-financial participation barriers to identify the model. These costs are continuous, albeit with limited support, and permit the estimation of the marginal treatment effect over a larger range of participation rates than would be the case with a dichotomous instrument, which permits the estimation of that effect only between two participation rates

Another feature of the model is that it imposes sufficient assumptions to allow extrapolation of the results to learn the effects of policy interventions which alter variables other than participation costs. A completely nonparametric approach would only permit the prediction of the labor supply effects of the instrument itself, but that instrument (fixed costs of participation) is not of particularly high policy interest. The model assumptions that welfare participation is a parametric function of the instrument as well as policy variables of greater interest (e.g., the guarantees and tax rates in the welfare benefit formula), which permits the prediction of the marginal labor supply effects of altering those variables as well.

Marginal treatment effects have been estimated for the effect of education on earnings by Carneiro et al. (2003) and Moffitt (2010). While the latter paper used estimation methods similar to those here, Heckman et al. used kernel methods. The main difference with the former paper, aside from the application, is the type of non-parametric estimation method used (kernel methods versus series methods).

The results show that there is significant heterogeneity in the response of labor supply to welfare participation. When costs are high and participation is low, the welfare caseload is disproportionately composed of those with the greatest labor supply disincentives (i.e., the most negative effects). As costs fall and participation rises, individuals with smaller disincentives enter the program and hence the average disincentive of those on the program falls. The coefficients on other variables in the model are shown to be consistent with the cross-sectional static model and how the size of the labor supply reduction varies with wage rates and nonlabor income. The extrapolated results from the estimated parametric cost-benefit relationship show that successive increases in the generosity of a welfare program have increase the labor supply disincentives of those initially on the program but bring in individuals with smaller labor supply responses, and that the net effect on the mean response of those on the program is approximately zero.

The first section of the paper formulates the standard static labor supply model but with individual heterogeneity and shows that changes in the cost of participation have ambiguous theoretical effects on the marginal labor supply disincentive, so that empirical examination is needed to answer the question. The next section of the paper lays out the econometric model, which is followed by the presentation of the data and results.

I Adding Heterogeneity to the Canonical Static Labor Supply Model of Transfers

The canonical static model of the labor supply response to transfers with variable (i.e., incomplete) takeup (Moffitt (1983)) posits utility to be

$$U(H_i, Y_i; \theta_i) - \phi_i P_i \tag{1}$$

where H_i is hours of work for individual i, Y_i is disposable income, P_i is a program participation indicator, θ_i is a vector of labor supply preference parameters, and ϕ_i is a scalar representing fixed costs of participation in utility units. The separability of P_i from the U function is for analytic convenience and is not required for any of the following results. Allowing for fixed costs of participation–in money, time, or utility, with the exact type unspecified–is required because many individuals who are eligible for transfer programs do not participate in them. If this were not the case, then all individuals would locate on the boundary of their budget sets and program participation would be automatically determined by the choice of H, meaning that there would be no separate participation decision.

The individual faces an hourly wage rate W_i and has available exogenous non-transfer nonlabor income N_i. The welfare benefit formula is $B_i = G - tW_iH_i - rN_i$ (assuming, for the moment, that the parameters G, t and r do not vary by i) and hence the budget constraint is

$$\begin{aligned} Y_i &= W_i(1-t)H_i + G + (1-r)N_i \;\; if \;\; P_i = 1 \\ Y_i &= W_iH_i + N_i \;\; if \;\; P_i = 0 \end{aligned} \tag{2}$$

The resulting labor supply model is represented by two functions, a labor supply function conditional on participation and a participation function:

$$H_i = H[W_i(1 - tP_i), N_i + P_i(G - rN_i); \theta_i] \tag{3}$$

$$P_i^* = V[W_i(1-t), G + N_i(1-r); \theta_i] - V[W_i, N_i; \theta_i] - \phi_i \tag{4}$$

$$P_i = 1(P_i^* \geq 0) \tag{5}$$

where V is the indirect utility function and 1(•) is the indicator function. Nonparticipants, those for whom P* is negative, are of two types: low-work individuals for whom a positive benefit is offered and a utility gain (in V) could be obtained but who do not participate because ϕ_i is too high, and high-work individuals for whom the utility gain (in V) is negative and who would not participate even if ϕ_i were zero (these individuals are above the eligibility point, or "above breakeven" in the terminology of the literature). Figure 1 is the familiar income-leisure diagram showing three different individuals who respond to the transfer program constraint by continuing to work above the breakeven point (III), below breakeven but off the program (II), and below breakeven and on the program (I'; I is the pre-program location for this individual).

The response to the program for individual i is

$$\triangle_i(\theta_i) = H[W_i(1-t), G + N_i(1-r); \theta_i] - H[W_i, N_i; \theta_i] \tag{6}$$

which is a heterogeneous response if θ_i varies with i. The response $\triangle_i$ includes both responses from below breakeven and above breakeven. Individual values of $\triangle_i$ will never be identified by the data, but the mean of those values over some populations or subpopulations can be. Letting

$$S_\phi = \sup port \;\; of \;\; \phi \tag{7}$$

$$S_{\theta(\phi)} = set \;\; of \;\; \theta \;\; s.t. \;\; P_i = 1 \;\; conditional \;\; on \;\; \phi \tag{8}$$

the mean effect of the transfer program over the entire population, participants and non-participants combined, conditional on the budget constraint, is

$$\widetilde{\triangle} = E(\triangle_i P_i \mid W_i, N_i, G, t, r) \tag{9}$$

$$= \int_{S_\phi} \int_{S_{\theta(\phi)}} \triangle_i \, (\theta_i \mid W_i, N_i, G, t, r) dG(\theta_i, \phi_i) \tag{10}$$

where $G(\theta_i, \phi_i)$ is the joint c.d.f. of the two heterogeneity components. Note that the two sets S are functions of the budget constraint parameters, which is not made explicit. Letting S_θ be the unconditional support of θ, the participation rate in the population is

$$P = E(P_i \mid W_i, N_i, G, t, r) \tag{11}$$

$$= \int_{S_\phi} \int_{S_\theta} 1\{V[W_i(1-t), G + N_i(1-r); \theta_i] - V[W_i, N_i; \theta_i] - \phi_i\} dG(\theta_i, \phi_i) \tag{12}$$

and the mean labor supply response among those who participate is

$$\widetilde{\triangle}_{P_i=1} = \widetilde{\triangle}/P \tag{13}$$

The marginal response to a change in program participation, which is often interpreted as the mean $\triangle$ of those who change participation, is $\partial\widetilde{\triangle}/\partial P$. These effects have been discussed extensively in the treatment effects literature and are defined within the econometric model in the next section.

The distribution of θ_i affects the mean response in the population in two ways: first, by affecting the distribution of $\triangle_i$ across the population–that is, the distribution of response if all individuals participate–and, second, by altering which of those individuals participate because θ_i appears in eqn(4). The distribution of ϕ_i affects mean response only through the latter mechanism, by altering the composition of the participant population; this feature will lead to an exclusion restriction in the econometric model below.

While $\triangle_i$, $\widetilde{\triangle}$, $\widetilde{\triangle}_{P_i=1}$, and $\partial\widetilde{\triangle}/\partial P$ must be negative according to theory, how they change as the participation rate changes is less clear and requires making a distinction between different sources of change in participation. How the effect varies with a change in participation induced by a change in ϕ_i, for example, is ambiguous in sign because the magnitude of $\triangle_i$ has no determinate relationship to the magnitude of the non-cost portion of the utility gain of going onto welfare, $dV = V[W_i(1-t), G + N_i(1-r); \theta_i] - V[W_i, N_i]$. For example, those with greater gains dV may be those with greater marginal utilities of consumption and hence those with smaller marginal utilities of leisure; it is the relative marginal utility of consumption and leisure that matters. An increase in participation induced by a reduction in ϕ will draw new individuals onto welfare whose values of dV are smaller than those of initial recipients (for any given value of ϕ, participation is positively selected on dV), but those smaller values of dV could be associated with either greater or smaller labor supply reductions. Thus, one central question of the analysis can only be determined empirically.

Participation rate expansions induced by changes in the budget constraint, on the other hand, have quite different effects because they induce changes in mean labor supply reductions for those initially on welfare as well. An expansion of the generosity of the program, for example, will increase participation and necessarily increase mean labor supply reductions. Thus $\triangle_i$, $\widetilde{\triangle}$, and $\widetilde{\triangle}_{P_i=1}$ will necessarily become more negative as participation rises. However, this gross marginal response, $\partial\widetilde{\triangle}/\partial P$–that is, not holding the budget constraint variables fixed–cannot be interpreted as the mean response of those brought into the program because it will include not only their responses but also the mean increase in labor supply reductions of those initially on the program. But the correlation between $\triangle_i$ and dV will still be at play in this case because it will determine whether the labor supply reductions of the new entrants are greater or smaller than those of the initial recipients after both face the same new budget constraint. Consequently, heterogeneity in response may make the increase in labor supply reductions arising from budget constraint expansions greater or smaller than would be predicted if responses had been assumed to be homogeneous and unchanging as the program expands. These effects will be separately identified in the econometric model in the next section.

II An Econometric Model

The object of the exercise is to estimate eqns(3)-(5). However, a choice model will not be imposed on the problem and we shall let H_i be a function of P_i and some additional covariates that proxy the budget constraint variables. The participation equation likewise will simply be allowed be a function of a set of variables including proxies for the budget constraint parameters and for costs of participation. A vector of other covariates will be added on the presumption that they affect the remaining unobservable portions of parameters θ and ϕ. Imposing a formal utility choice structure on each equation and on them jointly is left for future work.

To illustrate the structure of the model, we shall initially ignore all covariates and will focus on a model for H_i, P_i, and an observable proxy for participation cost, which we shall denote as Z_i. An unrestricted model with full individual heterogeneity can be written as follows [1]

$$H_i = \beta_i + \alpha_i P_i \tag{14}$$

$$P_i^* = m(Z_i, \delta_i) \tag{15}$$

$$P_i = 1(P_i^* \geq 0) \tag{16}$$

where β_i and α_i are scalar random parameters and δ_i is a vector of random parameters. All parameters are allowed to be individual-specific and to have some unrestricted joint distribution. A separate model of this type exists for each individual i. The function m

[1] This model is modified from a similar model constructed by Moffitt (2010).

can likewise be unrestricted and can be saturated if Z_i is assumed to have a multinomial distribution, although we shall discuss restrictions on δ_i below. The object of interest is the distribution of α_i. Selection in this model can occur either on the intercept (β_i) or the slope coefficient (α_i) because both may be related to δ_i and, in fact, the theoretical model implies that they must be because the participation equation contains the parameters of the labor supply function. Assuming Z_i is independent of the three parameters, we can condition both equations on it to determine what is identified and estimable:

$$E(H_i \mid Z_i = z) = E(\beta_i \mid Z_i = z) + E(\alpha_i \mid P_i = 1, Z_i = z)\,\Pr(P_i = 1 \mid Z_i = z) \tag{17}$$

$$E(P_i \mid Z_i = z) = \Pr[m(z, \delta_i) \geq 0] \tag{18}$$

What we wish to identify is $E(\alpha_i \mid P_i = 1, Z_i = z)$ (if we can identify that, we can also integrate over the support of Z_i to obtain the mean of α_i conditional only on participation). Identification requires that Z_i satisfy two mean independence requirements, one for the intercept and one for the slope coefficient:

$$A1.\;\; E(\beta_i \mid Z_i = z) = \beta \tag{19}$$

$$A2.\;\; E(\alpha_i \mid P_i = 1, Z_i = z) = g[E(P_i \mid Z_i = z)] \tag{20}$$

where g is the effect for those on the program (i.e., the effect of the treatment on the treated) conditional on Z_i, and depends on the shape of the distribution of α_i and how different fractions of participants are selected from different portions of that distribution. While the first assumption is familiar, the second may be less so. The usual assumption in the literature is that the two potential outcomes, β_i and $\beta_i + \alpha_i$, are fully independent of Z_i, which implies that α_i is as well. Eqn (20) is a slightly weaker condition which states that all that is required is that the effect of the treatment on the treated be dependent on Z_i only through the effect of the participation probability. If this were not so, different values of Z_i would lead to different conditional means of α_i through some other channel, which would rule it out as a valid exclusion restriction.

The "monotonicity" condition of Imbens and Angrist (1994) constitutes, in this model, a restriction on δ_i and can be expressed as

$$P_i(Z_i = z) - P(Z_i = z') \text{ is zero or the same sign for all } i \text{ for any distinct values } z \text{ and } z' \quad (21)$$

Inserting the two assumptions into the main model in eqns (17)-(18), and denoting the participation probability as $F(Z_i) = E(D_i \mid Z_i)$, we obtain two estimating equations

$$H_i = \beta + g[F(Z_i)]F(Z_i) + \epsilon_i \quad (22)$$

$$P_i = F(Z_i) + \upsilon_i \quad (23)$$

where ϵ_i and ν_i are mean zero and orthogonal to the RHS by construction. No other restriction on these error terms need be made, as this is a reduced form of the model. The first equation merely states that the population mean of H_i equals a constant plus the mean response of those in the program times the fraction who are in. The implication of the model is that preference heterogeneity is detectable by a nonlinearity in the response of the population mean of H_i (taken over participants and nonparticipants) to the participation probability. If responses are homogenous and hence the same for all members of the population, the function g reduces to a constant and therefore a shift in the fraction on the program has a linear effect on the population mean of H_i. If the responses of those on the marginal vary, however, the response of the population mean of H_i will depart from linearity. This feature of the heterogeneous-response treatment model has been noted by Heckman and Vytlacil (2005) and Heckman et al. (2006). However, here it will form the basis of the estimation of the model, as eqn(22) will be estimated directly.

.Nonparametric identification of the parameters of the model–β and the function g at every point F–is straightforward. F is identified at every point Z_i from the second equation. If there is a value of Z_i in the data for which $F(Z_i) = 0$, then β is identified from the mean of H_i at that point and hence g is identified pointwise at every other value of Z_i and hence F. If no such value is in the data, then g can only be identified subject to a normalization or multiple variables of g can be identified. For example, the LATE of Imbens and Angrist (1994) is identified by the discrete difference in H between two points z_i and z_j divided by the difference in F between those two points. A marginal treatment effect is a continuous version of this and requires some smoothing method across discrete values of Z, and is computed by $\partial H/\partial F = g'(F)F + g(F)$.

Exogenous covariates are now introduced and allowed to shift the parameters β and the functions g and F. Let X_i^β denote a vector which includes W_i , N_i, and sociodemographic characteristics (age, education, family composition, etc.), all of which affect labor supply when off welfare. Let X_i denote a vector which augments X_i^β with the welfare-program variables G and t, which will affect labor supply on welfare; X_i will shift the function g, the effect of welfare on labor supply. The vector X_i will affect the probability of participation as well, thus shifting F. While extensive interaction is in principle possible by estimating the model separately for every set of values of these covariates, a less ambitious and more conventional approach will be taken here, which is to introduce index functions of covariates and to allow these index functions to affect the means of β, g, and F, and to be additively separable with Z. With this formulation, the model specializes to

$$H_i = X_i^\beta \beta + [X_i\lambda + g(F(X_i\eta + \delta Z_i))]F(X_i\eta + \delta Z_i) + \epsilon_i \tag{24}$$

$$P_i = F(X_i\eta + \delta Z_i) + \nu_i \tag{25}$$

which leaves only the functions g and F unspecified. For g, we will estimate its shape with series methods, either splines or polynomials. We will assume normality for F and leave nonparametric estimation of that function for future work.[2] With these two functions specified, we will employ two-step estimation of the model, with a first-stage probit estimation of eqn(25) and second-stage estimation of eqn(24) using fitted values of F from the first stage. Robust standard errors correcting for estimation error in F and for the nonlinearity of F in eqn(24) are obtained by applying the asymptotic formulas in Newey and McFadden (1994).

The parametric assumptions on the participation function imply that the effect of changes in budget constraint variables in X on labor supply can always be calculated, even though they are not the source of the identification of the model because they appear in both equations. Since there is a mapping from Z_i to X_i in the participation equation, the effects of budget constraint variables on participation can be separated from their direct effects on labor supply conditional on participation (from the first equation). A completely nonparametric approach which allowed X_i and Z_i to have separate and unrelated effects on participation would not allow such a calculation.

III Data and Results

We study the labor supply effects of the well-known U.S. cash transfer program, the Aid to Families with Dependent Children (AFDC) program, using data from the Survey of Income and Program Participation (SIPP) in the early 1990s. The SIPP is a set of rolling, short (12 to 48 month) panels which are representative samples of the U.S. population. The first panel began in 1984 and subsequent panels for many years were begun annually, each with between 30,000 and 70,000 families. To increase sample sizes of the subpopulation we will examine (disadvantaged single mothers), we pool the SIPP panels having interviews between 1989 and 1991. We do not go farther than 1991 because a major restructuring of the program began shortly after that which introduced work requirements,

[2] Equations (25)-(26) are equivalent to the classic Lee (1979) two-regime switching regression model but with nonparametric assumptions on the unobservables (save for F).

time limits, and other features to the program which are not captured in our model. Prior to 1992, the program was close to a simple cash transfer program paying benefits according to a fixed schedule. To minimize seasonal variables, we draw our data from the Spring surveys of all SIPP panels interviewing families in the Spring of 1989, 1990, and 1991.[3]

Eligibility for AFDC in this period required sufficiently low assets and income and, for the most part, required that eligible families be single mothers with at least one child under 18. Our sample is therefore restricted to such families, similar to the practice in past AFDC research. To concentrate on the AFDC-eligible population, we restrict our sample to those with completed education of 12 years or less, nontransfer nonlabor income less than $1,000 per month, and between the ages of 20 and 55. The resulting data set has 5,722 observations.

The variables we use for estimation are average hours worked per week in the month prior to interview (H) (including zeroes), whether the mother was on AFDC at any time in the prior month (P), and we construct covariates for education, age, race, and family structure. The hourly wage is omitted because it is only available for workers and is assumed to be proxied by demographic characteristics, especially education. However, nontransfer nonlabor income is explicitly included among the covariates. The AFDC guarantee for a family of four in the individual's state of residence is also included. AFDC tax rates on earned and unearned income are not included because uniform levels were imposed on the states by the federal government over this period, both equal approximately to 100 percent.[4] The names, definitions, and means of the variables used in the estimation appear in Appendix Table A1. Thirty-one percent of the sample was on AFDC in the month prior to interview.

The exclusion restrictions (Z) are selected to proxy costs of participation in AFDC. Institutional descriptions of the program have revealed that non-financial barriers have

[3]We take data from the 1989, 1990, and 1991 panels and select all families interviewed in the February-May period in any year of the panel.

[4]The nominal tax rate was 100 percent but the effective tax rate differed somewhat from this because of various exemptions and allowances, including a four-month window when the tax rate was 67 percent.

always been present in the program and have hindered participation, perhaps intentionally on the part of the states to keep caseloads down. Data are available on a number of proxies for these barriers and information was collected from official documents on several of them and were pretested in OLS estimations of the welfare participation equation. From this exercise, three emerged as consistently significant and with the expected sign: the error rate made by the state resulting in incorrect denial of eligibility (collected by the federal government as part of its audit procedures of state records), the percent of applications denied because of a failure on the part of the applicant to comply with all procedure requirements (an indication of the amount of paperwork and bureaucracy imposed on prospective recipients), and administrative expenses per case in the state (interpreted as an indicator of the level of bureaucracy in the program). All three affect welfare participation negatively. The means and data sources of the three variables are given in Appendix A.

For the initial results we set $\lambda = 0$ (hence no interactions of X with participation) and estimate the hours equation by OLS, regressing hours on X and P. OLS gives a response estimate of -24.6 (s.e.=.46), which is only slightly smaller than the raw mean difference between participants and non-participants of -26.3, implying that conditioning on X has little effect. Next we estimate eqn(25) assuming a constant g, which is equivalent to the homogeneous-effect model and equivalent to 2SLS, though using probit for the first stage instead of the linear model. The estimate of g is shown in the first column of Table 1 (other parameter estimates are shown in Appendix Table A2). The estimate is -29.5 (s.e.=.3.5), a bit larger than the OLS estimate, suggesting that OLS is slightly biased downward.

The rest of the columns show the results of fitting splines and polynomials to the g function, of the form

$$g(F) = \gamma_0 + \sum_{j=1}^{J} \gamma_j Max(0, F - \pi_j) \tag{26}$$

$$g(F) = \gamma_0 + \sum_{j=1}^{J} \gamma_j F^j \tag{27}$$

where the π_j are preset spline knots. Column (2) shows the effect of allowing g to linearly decline with $F(J = 1, \pi_1 = 0)$. The coefficient on F is positive and significant, implying that the average labor supply disincentive of participants falls as participation expands. The marginal treatment effect is $-69.5 + 89.4F$, implying an even faster decline in the work disincentive with increases in F. Columns (3) and (4) show the effects of allowing splines at the median of the predicted F distribution and at its 25th and 75th percentile points. Column (3) shows that the significance of the g parameters declines considerably with the addition of a spline at the median, and column (4) shows a further decline in the stability of the fitted model, as many of the slopes are implausibly large and poorly determined. This indicates that the data do not have the ability to detect nonlinearities much higher than above or below the median, and perhaps only a linearly declining effect without any further nonlinearities.

Figure 2 plots the marginal treatment effect for the different models. For participation rates above about .30, the models uniformly predict declining work disincentives as participation expands, although the quadratic specification has a faster rate of increase, which is no doubt a result of poor extrapolation common to polynomials. Ninety-percent confidence interval bands (not shown) generally overlap in this region. However, for lower participation rates, some models show declining marginal treatment effects and others show increasing effects (confidence interval bands often do not overlap). The linear gamma specification, which shows declining marginal treatment effects, could be a result of poor extrapolation from higher values of F. In fact, a rising labor supply disincentive over the early range is theoretically possible because early entrants to welfare–those who participate even though costs are high–could be those with the lowest values of labor supply, and their labor supply reductions might be bounded below by $H = 0$ (and they could have $H = 0$ even if not on welfare).

However, a feature of the estimates which is not revealed by these methods is where the instruments have the most power, i.e., in what ranges the instruments Z move F the most. In fact, these instruments have very little power at low values of F. On the one hand, the fitted F distribution, which is shown in Figure 3, shows that the data have considerable density at low values of the participation rate (though very little above .60 or so; the estimates should not be reliable in this range). However, these densities partly arise from the X vector, and the question instead is how much the instruments move the participation probabilities in the different ranges. This is illustrated in Figure 4, showing that the instruments have their largest effect in the range (.30,.60), and have little or no impact at very low or very high values of the participation rate. Thus the estimates from the different models in the low ranges of F have little reliability.

Table 2 shows estimates of the g function and the parameters λ when the latter are not restricted to zero. This specification therefore allows the effect of welfare participation on labor supply to vary with observables. Not surprisingly, those γ parameters that were significant and hence reasonably well-determined continue to be so but the magnitudes of the parameter estimates are reduced. More interesting in this specification is that response does often vary significantly with observables, for the work disincentives of welfare participation are greater for younger women and those with more education and lower levels of nonlabor income. Those facing higher welfare guarantee levels have greater work disincentives, as predicted by the theory.

The magnitudes of the differences in work disincentives for those with different observable characteristics are considerable, as shown in Table 3. For example, white 25-year-old women with 12 years of education with \$100 of monthly income and facing a \$600 monthly welfare guarantee have work disincentives (-42.1 hours per week) about double those of black 35-year-old women with 8 years of education, \$150 of nonlabor income, and facing a \$400 monthly guarantee (-19.6 hours per week). The importance of interacing observables with participation is clearly demonstrated.

Although the budget constraint variables W, N, and G (recall that t is fixed at 100 percent for all women) are not entered structurally in the model, the latter two are represented explicitly in the X vector and education can be taken as a proxy for W. The simple static labor supply model of response to a welfare program with a 100 percent tax rate implies that the effects of W, N, and G on the magnitude of the labor supply response should be negative, positive, and negative, respectively, as shown in Figures 5-7. Women with a higher wage rate work more than those with a lower wage rate (assuming substitution effects dominate income effects, as past work has shown them to for single mothers) and hence reduce hours more when going to $H = 0$ (Figure 5). Women with greater levels of nonlabor income work less when off welfare, resulting in smaller reductions in hours when going onto welfare (Figure 6). Finally, and more obviously, higher levels of G result in greater H reductions (Figure 7). These theoretically predicted signs are indeed the signs estimated by the model.

A further confirmation of this interpretation can be obtained by examining the β and η coefficients for education, N, and G. According to Figures 5-7, W and N, should have positive and negative,effects on the level of H, off welfare, for example. These are the signs of the estimated coefficients in the β vector (see Appendix Table A2). The variables W, N, and G should have negative, negative, and positive effects on the probability of welfare participation, respectively. These are also the signs of the estimated coefficients in the η vector. The expected signs in the different vectors, all consistent with the data, are summarized in Table 4.

Finally, as noted previously, the parametric assumptions on the participation equation allow an estimate of the effect of an increase in G on the mean labor supply disincentive, taking into account the fact that new entrants will be brought into the program whose labor supply disincentives differ from that of those who are initially on the program. The total effect of a change in G on the mean labor supply disincentive (that is, on the mean effect on participants) is $\lambda_G + (\partial g/\partial F)(\partial F/\partial G)$. The first term is negative in sign (albeit insignificant) and the second term is positive, at least over the range for participation rates above about .30, as discussed above (the product of two positive partials). Thus the direct effect of an increase in monthly G on reducing labor supply is dampened by new entrants who have smaller labor supply reductions than those already on welfare. Taking the estimates for the model in column (1) of Table 2, $\lambda_G = -1.6$, $\partial g/\partial F = 61.7$, and $\partial F/\partial G = .028$ (evaluated at the mean of the normal density), and hence the total effect is 0.13 hours of work per week, actually positive but close to zero (this is for a \$100 change in G). Thus the selection effect about cancels out the direct effect, leaving no change in mean labor supply of participants.

IV Conclusion

In this paper we have modified the traditional static labor supply model of the effect of income transfers to allow for heterogeneous response, implying that changes in welfare generosity have not only direct effects on labor supply of participants but which also change the composition of the caseload and hence change mean work disincentives through compositional effects. Using a modified version of the conventional treatment effects model and estimating the distribution of the unobserved response heterogeneity with series approximation methods, the results show that the marginal treatment effect on hours of work is positive. This implies that an increase in participation brings into the program individuals who have smaller work disincentives than those initially on the program. This effect is shown by a variety of approximation methods in the range of the data where the instruments have power, which is about in the range of participation rates between 30 and 60 percent of the population. At lower and higher rates of participation, the data either have either thin distributions of participants (at the high end) or weak power of the instruments (at the low end); thus we have no strong evidence on the shape of heterogeneous responses in these ranges.

These results modify some of the findings in the large literature on estimating the effects of the response of labor supply to changes in welfare participation. Since that literature has generally assumed homogeneous responses, their estimates are best interpreted as some weighted average of responses over the ranges of participation in the data sets used in each of these studies. There is consequently no reason to expect that past estimates should generate the same response estimates since the data sets used have ranged over calendar time and across different states, where participation rates have no doubt differed.

The model used in the paper could use refinements in its method of series estimation of the heterogeneity distribution. Relaxation of the parametric assumptions on the observables could also be usefully conducted. More structural methods which incorporate budget constraint variables more formally is another direction of useful pursuit, as would be models of dynamic labor supply.

References

Bjorklund, A. and R. Moffitt (1987). The estimation of wage and welfare gains in self-selection models. *The Review of Economics and Statistics 69*(1), 42–49.

Blundell, R. and T. Macurdy (1999). Labor Supply: A Review of Alternative Approaches. In O. C. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Volume 3A.

Burtless, G. (1987). The Work Response to a Guaranteed Income: A Survey of Experimental Evidence. In A. Munnell (Ed.), *Lessons from the Income Maintenance Experiments*. Boston: Federal Reserve Bank of Boston and Brookings Institution.

Burtless, G. and J. Hausman (1978). The Effect of Taxation on Labor Supply: The Gary Income Maintenance Experiment. *Journal of Political Economy 86*, 1103–1130.

Cain, G. G. and H. W. Watts (1973). *Income Maintenance and Labor Supply*. Boston,MA: Rand McNally.

Carneiro, P., J. J. Heckman, and E. Vytlacil (2003). Understanding What Instrumental Variables Estimate: Estimating Marginal and Average Returns to Education.

Danziger, S., R. Haveman, and R. Plotnick (1981). How Income Transfers Affect Work, Savings, and the Income Distribution: A Critical Review. *Journal of Economic Literature 19*(3), 975–1028.

Grogger, J. and L. A. Karoly (2005). *Welfare Reform: A Decade of Change*. Cambridge: Harvard University Press.

Gueron, J. M. and E. Pauly (1991). *From Welfare to Work*. New York: Russell Sage Foundation.

Hausman, J. (1981). Labor Supply. In H. Aaron and J. Pechman (Eds.), *How Taxes Affect Economic Behavior*. Washington: Brookings Institution.

Heckman, J. J. and R. Robb (1985). Alternative Methods for Evaluating the Impact of Interventions. In J. J. Heckman and B. Singer (Eds.), *Longitudinal Analysis of Labor Market Data*. Cambridge: Cambridge University Press.

Heckman, J. J., S. Urzua, and E. Vytlacil (2006). Understanding instrumental variables in models with essential heterogeneity. *The Review of Economics and Statistics 88*(3), 389–432.

Heckman, J. J. and E. Vytlacil (2005, May). Structural equations, treatment effects, and econometric policy evaluation. *Econometrica 73*(3), 669–738.

Imbens, G. and J. Angrist (1994). Identification and estimation of local average treatment effects. *Econometrica 62*(2), 467–475.

Lee, L. (1979). Identification and Estimation in Binary Choice Model with Limited (Censored) Dependent Variables. *Econometrica 47*(4), 977–996.

Moffitt, R. (1983, December). An Economic Model of Welfare Stigma. *American Economic Review 73*(5), 1023–35.

Moffitt, R. (1992). Incentive Effects of the U.S. Welfare System: A Review. *Journal of Economic Literature 30*(1), 1–61.

Moffitt, R. (2010). Estimating Marginal Treatment Effects in Heterogeneous Populations. *Annales d'Economie et de Statistique*.

Moffitt, R. and K. Kehrer (1981). The Effect of Tax and Transfer Programs on Labor Supply: The Evidence from the Income Maintenance Experiments. In R. Ehrenberg (Ed.), *Research in Labor Economics*, Volume 4, pp. 103–50. Greenwich, Conn.: JAI Press.

Newey, W. and D. McFadden (1994). Large Sample Estimation and Hypothesis Testing. In

R. Engle and D. McFadden (Eds.), *Handbook of Econometrics*, Volume IV. Amsterdam: Elsevier.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *The American Economic Review 66*(5), 688–701.

Table 1

Gamma Parameter Estimates

	(1)	(2)	(3)	(4)	(5)
Constant	-29.5 (3.5)	-69.5 (6.7)	-44.3 (20.7)	-11.5 (42.6)	-40.4 (14.5)
F	--	44.7 (6.6)	-21.2 (52.9)	-135.0 (191.5)	-24.4 (32.2)
Max(0,F-F(.25))	--	--	--	101.7 (180.9)	--
Max(0,F-F(.50))	--	--	61.1 (49.0)	25.6 (53.9)	--
Max(0,F-F(.75))	--	--	--	48.1 (27.0)	--
F^2					52.2 (24.2)

Notes:

Standard errors in parentheses. Parameter estimates for β, δ, and η for column (1) are shown in Appendix Table A2. All models constrain λ=0. Percentile points for splines: F(.25)=.29, F(.50)=.17, F(.75)=.43

Table 2

Gamma and Lambda Parameter Estimates

	(1)	(2)	(3)
<u>Gamma</u>			
Constant	-42.1 (22.7)	-65.7 (38.6)	-29.7 (24.8)
F	61.7 (22.3)	116.1 (76.7)	27.8 (38.1)
Max(0,F-F(.50))	--	-44.2 (60.0)	--
F^2	--	--	31.5 (30.5)
<u>Lambda</u>			
Education	-2.6 (1.0)	-2.4 (1.1)	-2.4 (1.1)
Age	4.4 (1.9)	4.4 (1.9)	4.3 (1.9)
Black	0.4 (4.4)	-0.3 (1.4)	-0.3 (4.6)
No. Children Lt 6	-4.7 (3.0)	-5.5 (3.2)	-6.1 (3.4)
Family size	1.7 (1.2)	1.7 (1.3)	1.7 (1.2)
Nonlabor income	7.9 (1.8)	8.7 (2.1)	7.7 (1.8)
Welfare G	-1.6 (0.5)	-1.6 (0.5)	-1.8 (0.6)

Notes:
Standard errors in parentheses.

Table 3

Mean Treatment Effect on Treated at Different X

Age	Education	Race	N (nonlabor income)	G (guarantee)	g (s.e.)
35	8	Black	150	400	-19.6 (10.9)
35	8	Black	100	400	-23.5 (11.0)
35	8	White	100	400	-23.9 (9.0)
35	8	White	100	600	-27.0 (8.9)
25	8	White	100	600	-31.5 (9.1)
25	12	White	100	600	-42.1 (7.1)

Notes:

$g = X\lambda + \gamma_0 + \gamma_1 F$ evaluated at F=.33

At mean X, $X\lambda$=-10.9

Table 4

Expected Coefficient Signs Under Budget Constraint Interpretation

	λ	β	η
Wage (Education)	<0	>0	<0
Nonlabor income	>0	<0	<0
Welfare Guarantee	<0	<0	>0

Appendix A

Means and Data Sources

The means and standard deviations of the variables used in the analysis are shown in Table A-1. The sources of the state-level variables are as follows. The AFDC guarantee is the monthly maximum amount paid for a family of four in the state, and is obtained from unpublished data provided to the author by the U.S. Department of Health and Human Services for all three years 1989-1991. The state "negative case action error rate," the rate of error per applicant resulting in an incorrect denial of eligibility, is taken from the federal government's quality control program for AFDC and was obtained for 1991 from U.S. House of Representatives (1994, Table 10-39). The state percent of applicants denied for failure to comply with procedural requirements was obtained from the 1989, 1990, and 1991 issues of Quarterly Public Assistance Statistics published quarterly by the U.S. Department of Health and Human Services. The data on state administrative expenditures per case was obtained for 1989, 1990, and 1991 from http://www.acf.dhhs.gov/programs/opre/timetren/index.htm.

Appendix Table A1

Means and Standard Deviations of the Variables Used in the Analysis

Variable Name	Variable Definition	Total sample	P=1	P=0
Hours	Average hours of work per week in the month prior to survey	21.9 (19.4)	3.8 (.79)	30.1 (17.0)
P	Dummy variable equal to 1 if individual was on AFDC anytime in the month prior to survey	.31 (.46)	--	--
Age	Age in years at survey date divided by 10	3.2 (.88)	3.1 (.79)	3.3 (.91)
Education	Years of education at survey date	10.9 (2.0)	10.4 (2.1)	11.1 (1.9)
Family size	Number of individuals in the family at the survey date	3.3 (1.4)	3.4 (1.4)	3.2 (1.4)
No. Childress Lt 6	Number of children less than 6 in the family at the survey date	.72 (.89)	1.1 (1.0)	.55 (.76)
Black	Dummy variable equal to 1 if respondent is black	.33 (.47)	.44 (.50)	.28 (.45)

Appendix Table A1 (continued)

Variable Name	Variable Definition	Total sample	P=1	P=0
Nonlabor income	Nontransfer nonlabor income in the month prior to survey divided by 100	1.12 (2.07)	.40 (1.16)	1.45 (2.29)
Welfare G	State monthly AFDC guarantee for a family of four divided by 100	4.69 (1.97)	4.91 (2.01)	4.59 (1.94)
Admin	AFDC Administrative expenditures per case in the state averaged over 1989, 1990, and 1991, divided by 1000	.044 (.021)	.045 (.022)	.043 (.021)
Pctdenied	Fraction of applications denied for failure to meet procedure requirements in the state averaged over 1989, 1990, and 1991	.59 (.17)	.58 (.17)	.59 (.17)
Eligerror	Federally-audited percent error rate made by the state in 1991 in calculating eligibility	2.25 (2.26)	2.05 (1.87)	2.34 (2.40)
Sample size	--	5,722	1,783	3,939

Notes:

Standard deviations in parentheses
All dollar-valued variables are deflated by a 1990 price index using the GDP-based personal consumption expenditure deflator.

Appendix Table A2

Full Estimates for OLS and Basic 2SLS Specifications

	OLS	2SLS
Gamma	-24.6 (0.5)	-29.5 (3.5)
Beta		
Education	1.1 (0.1)	1.0 (0.2)
Age	2.0 (0.2)	2.0 (0.3)
Black	-1.4 (0.4)	-0.8 (0.6)
No. Children Lt 6	-0.9 (0.3)	-0.3 (0.5)
Family size	-0.6 (0.2)	-0.7 (0.2)
Nonlabor income	-0.5 (0.1)	-0.6 (0.2)
Constant	14.6 (1.6)	17.6 (2.9)
Nu	--	
Education	--	-0.10 (.01)
Age	--	-0.01 (0.02)
Black	--	0.44 (0.04)

Appendix Table A2 (continued)

	OLS	2SLS
No. Children Lt 6	--	0.39 (0.02)
Family size	--	-0.04 (0.01)
Nonlabor income	--	-0.19 (0.01)
Welfare G	--	0.08 (0.01)
Constant	--	0.47 (0.16)
Delta		
Admin	--	-2.70 (1.07)
Pctdenied	--	-0.31 (0.12)
Eligerror	--	-0.02 (0.01)

Notes:
Standard errors in parentheses
2SLS corresponds to Table 1, Column (1)

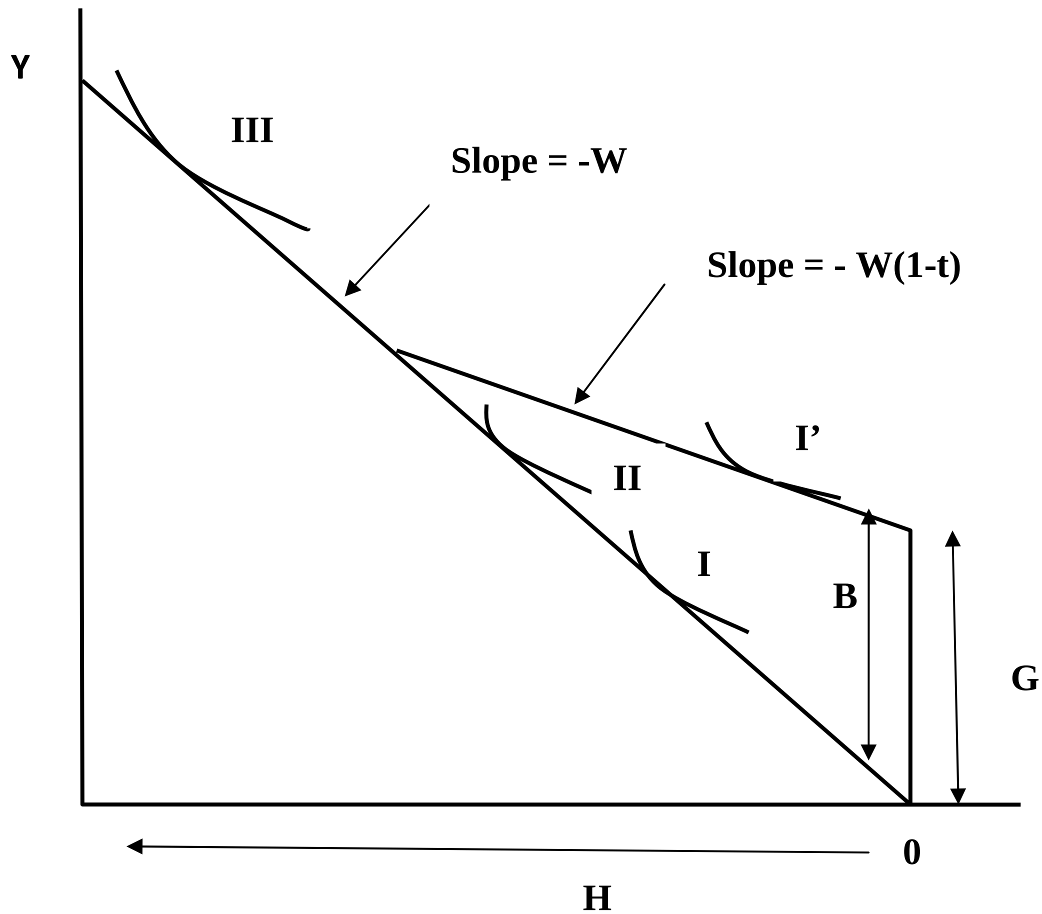

Figure 1

Figure 2: MTE for Different Models

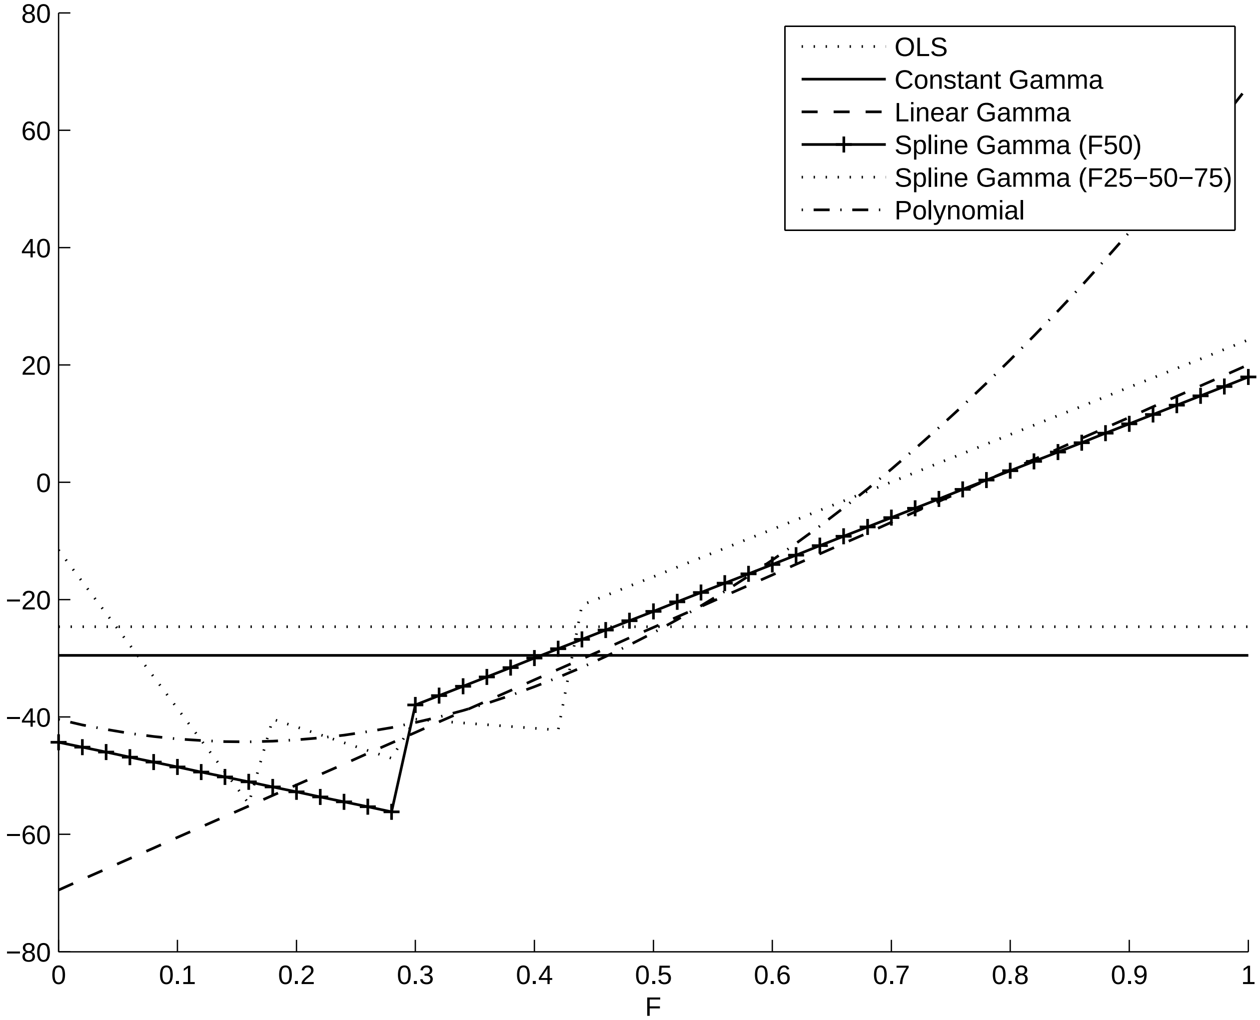

Figure 3: Histogram of Predicted Participation Rates

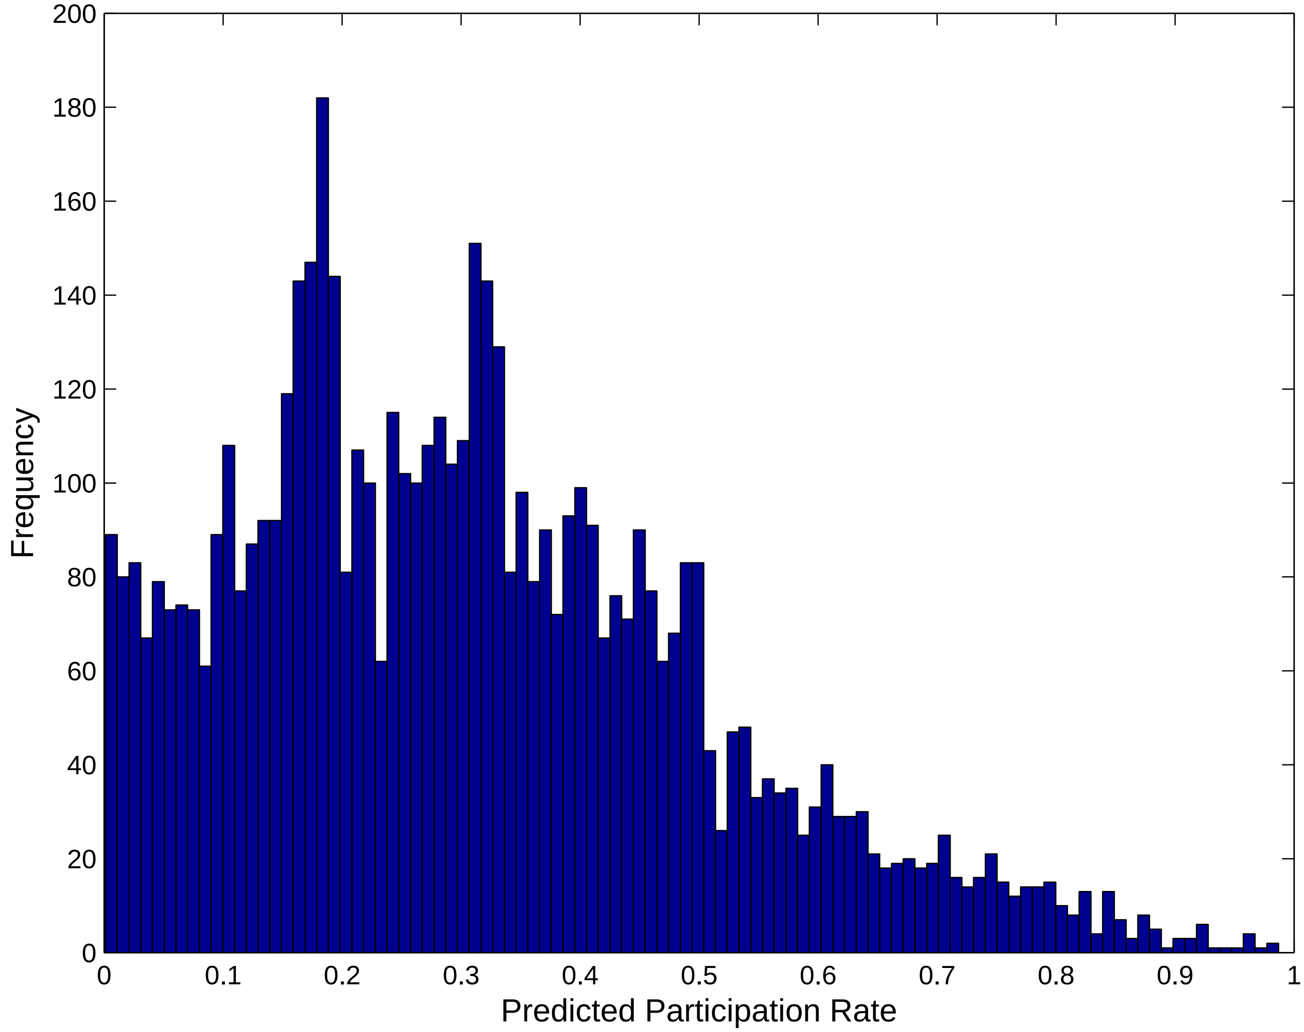

Figure 4: Baseline and Actual F Distribution
at Deciles of Baseline F

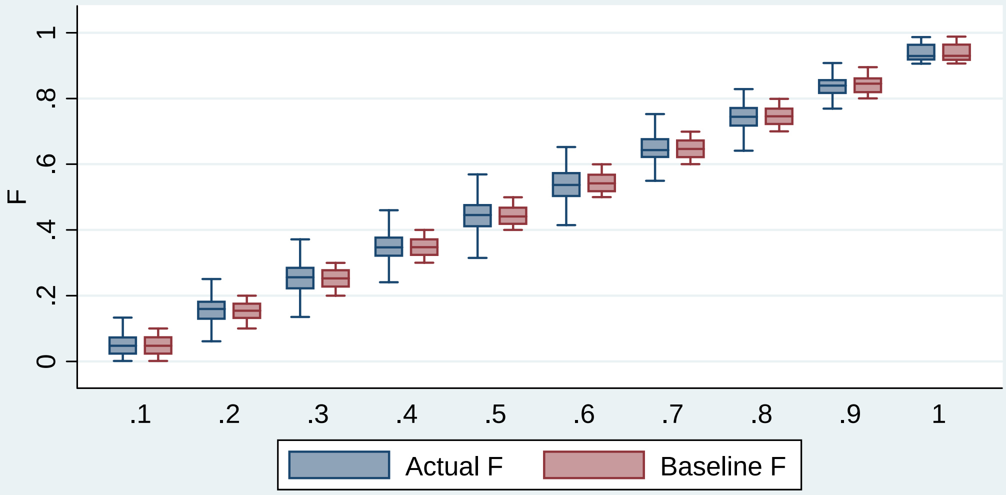

Baseline F is the predicted probability holding the Z vector at its mean.
Actual F is the predicted probability allow the Z vector to vary.
Horizontal axis represents decile ranges of Baseline F.
The upper and lower points of the rectangles are 75th and 25th percentile points of the distribution, respectively, and the horizontal lines inside the recentangles are medians.
Upper and lower tick marks above and below the rectangles are upper and lower ranges, respectively.

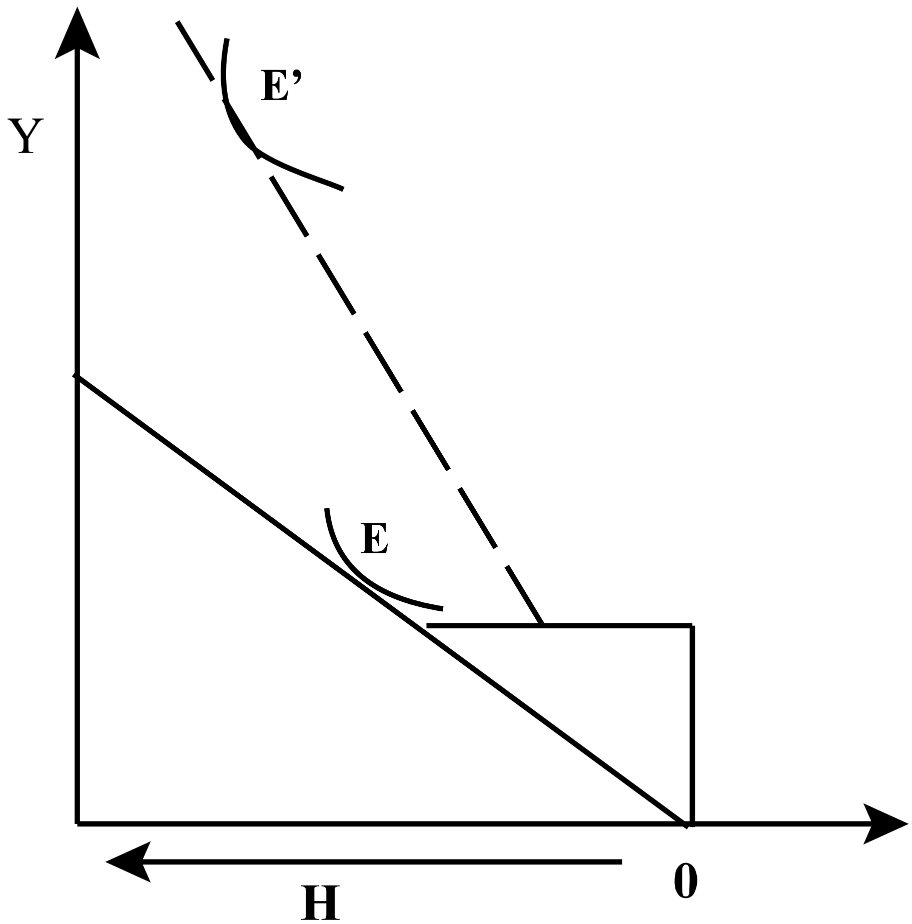

Figure 5. Effect of Increase in the Wage on Labor Supply assuming t=1

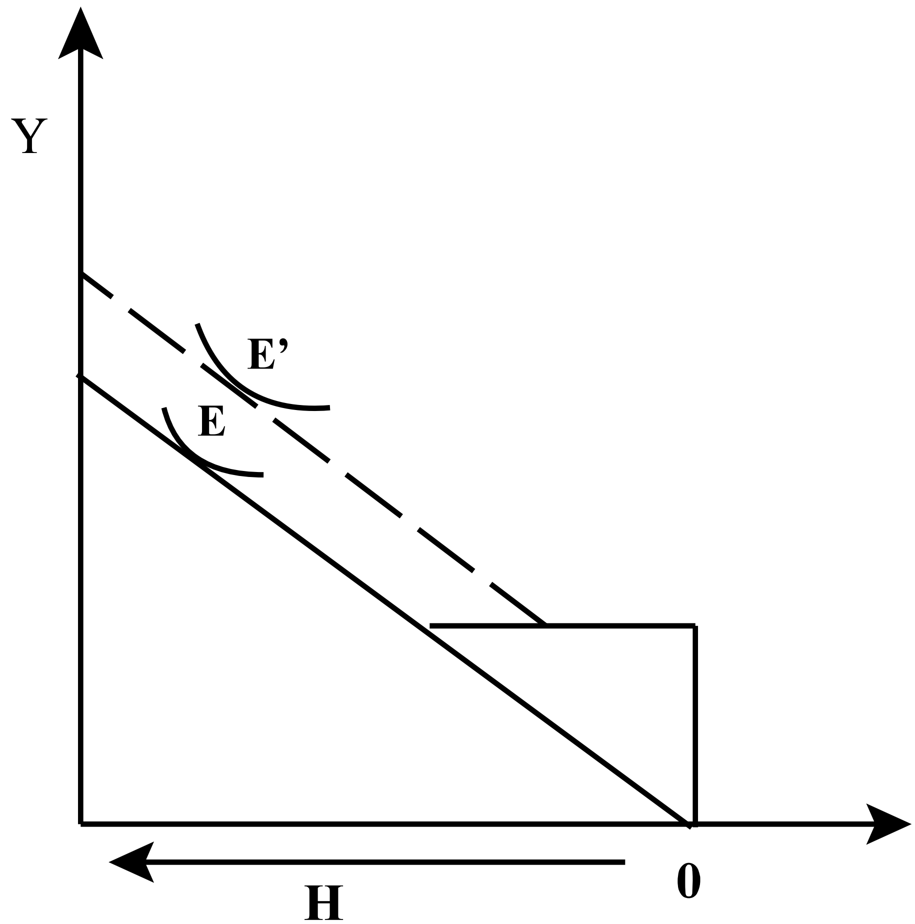

Figure 6. Effect of an Increase in Nonlabor Income on Labor Supply, assuming t=1

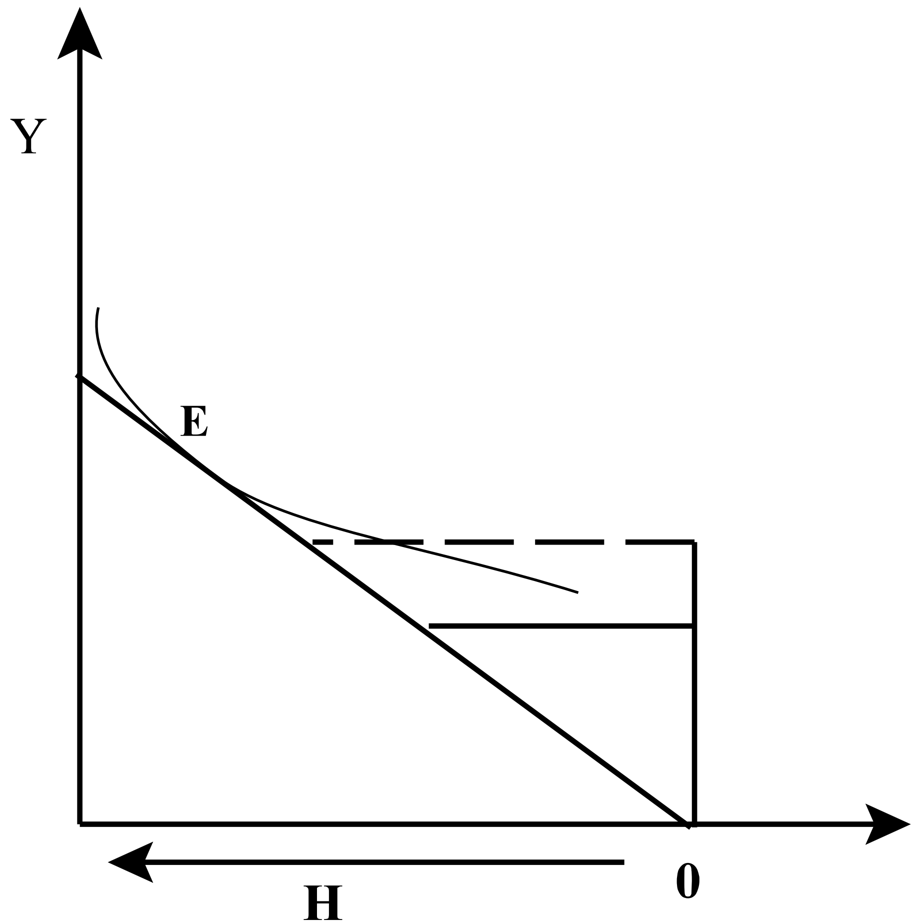

Figure 7. Effect of an Increase in G on Labor Supply, assuming t=1.