ELSEVIER

Journal of Public Economics 70 (1998) 157–183

JOURNAL OF PUBLIC ECONOMICS

Relative income concerns and the rise in married women's employment

David Neumark[a], Andrew Postlewaite[b,*]

[a]*Department of Economics, Michigan State University, and NBER, East Lansing, MI 48824, USA*
[b]*Department of Economics, University of Pennsylvania, Philadelphia, PA 19104, USA*

Abstract

We ask whether women's decisions to be in the workforce may be affected by the decisions of other women in ways not captured by standard models. We develop a model that augments the simple neoclassical framework by introducing relative income concerns into women's (or families') utility functions. In this model, the entry of some women into paid employment can spur the entry of other women, independently of wage and income effects. We show that relative income concerns can help to explain why, over some periods, women's employment rose faster than can be accounted for by the simple neoclassical model. We test the model by asking whether women's decisions to seek paid employment depend on the employment or incomes of other women with whom relative income comparisons might be important, including sisters and sisters-in-law. The evidence is largely supportive of the relative income hypothesis.

Keywords: Relative income; Married women; Employment

1. Introduction

What are the important factors in the decision to work? The standard neoclassical model provides an explanation as to why such factors as the wage rate, education, age and past work experience, spouse's income, as well as other personal and household characteristics can affect the decision. This model has been extremely useful in understanding labor supply decisions, but – not

*Corresponding author. E-mail: apostlew@econ.sas.upenn.edu

0047-2727/98/$ – see front matter
PII: S0047-2727(98)00065-6

surprisingly – there are some cases in which it falls short of providing compelling explanations. We are particularly interested in the rise in the employment of married women in the U.S. in the twentieth century. Empirical estimates of the simple neoclassical model of married women's labor force participation suggest that positive substitution effects outweigh negative own and husbands' income effects, with the consequence that rising real wages draw women into the labor market (Mincer, 1962; Smith and Ward, 1985). But there is ample evidence that this simple model fails to fully explain the rise in the employment of married women in the U.S.

The theme of this paper is that there may be important variables in women's decisions to enter the work force that are omitted from standard neoclassical models. We specifically focus on the possibility that a woman's deciding whether to be in the workforce may be affected by the decisions of other women in ways not captured by standard models. Other women's decisions may affect a particular woman's decision in many ways. For example, other women's decisions will affect the ''quality'' of remaining out of the work force if there are positive externalities among women who remain at home. Perhaps most importantly (although this remains an open question), to the extent that people care about their relative income position, a given woman's employment decision can be influenced by other women's employment decisions.

In this paper we develop a model that augments the simple neoclassical framework by introducing relative income concerns into women's (or families') utility functions. In this model, the entry of some women into paid employment can spur the entry of other women, independently of wage and income effects. We show that relative income concerns can help to explain why, over some periods, women's employment rose faster than can be accounted for by the simple neoclassical model.[1]

We test the model by asking whether women's decisions to seek paid employment depend on the employment or incomes of other women with whom relative income comparisons might be important, independently of standard variables affecting the employment decision. A natural peer or reference group, which has the advantage of being exogenously given, is a woman's siblings. We first test whether women's employment is affected by the employment of their sisters-in-law; we use sisters-in-law rather than sisters because, with the former, unobserved heterogeneity is much less likely to bias the results in favor of the relative income hypothesis. We also test whether women's employment is affected by the income of their husbands relative to the income of their sisters' husbands, a test that is more directly motivated by the theoretical model of relative income concerns that we develop. Both tests support the predictions of the model.

[1]Goldin (1994) considers how another non-neoclassical factor may influence changes in married women's participation over time, specifically the social stigma against married women's employment in manufacturing.

2. The rise in married women's employment: evidence and explanations

The twentieth century was characterized by rapid increases in women's employment in the U.S., especially that of married women. There was a slow acceleration of women's employment prior to 1940, concentrated among younger women. From 1940 to 1960 employment exploded for women aged 35 and over, with participation rates doubling for women aged 45–54 (to about 47%) and women aged 55–64 (to about 35%), and increasing by about 50% for women aged 35–54 (Smith and Ward, 1985, Table 1). After 1960 the employment of younger women accelerated, while the employment of older women held steady or declined, with the participation rate of women aged 25–34 rising two percentage points a year in the 1970s.[2]

Mincer (1962) developed the basic neoclassical model of married women's labor supply to attempt to explain the increased employment of married women. In this model, women's employment is influenced by two potentially offsetting factors. Rising real wages act through the husband's income to exert a negative income effect on employment. But rising real wages also act through the wife's market wage, with opposing income and substitution effects. Mincer estimated a cross-section model of women's employment, using 1950 Census data on SMSA averages to focus on permanent components of earnings. He found that the compensated substitution effect via the wife's wage dominated the income effect, and the positive uncompensated substitution effect dominated the negative income effect through the husband's wage. Consequently, rising real wages over time predict rising employment of married women.

Mincer then used the cross-sectional estimates, along with time-series data on full-time earnings of men and women, to ask whether the neoclassical model could explain the time-series changes in women's employment. He found that the cross-section estimates overpredicted the increase in married women's employment for 1919–1929, accurately predicted the increase for 1929–1939, and underpredicted the increase for 1939–1959 (1962, Table 10). His cross-section model explained 77% of the increase for 1939–1949, and 68% of the increase for 1949–1959. Later work by Bowen and Finegan (1969) strengthened the conclusion that the neoclassical model cannot adequately explain the rise in married women's employment, finding, for example, that the neoclassical model can explain only 25% of this rise in the two decades following World War II.[3]

Of course no one would expect the simple neoclassical model to explain all of

[2]These increases were concentrated among married women, as single women had high employment rates at the beginning of the century. Participation rates among black women display similar trends to those of white women, but started at much higher levels, and declined during the Depression (Goldin, 1990).

[3]Smith and Ward (1985) reach a qualitatively similar conclusion from a more sophisticated empirical analysis of the neoclassical model.

the increase in married women's employment. Researchers have also identified rising education levels and the growth of the clerical sector as potentially partly exogenous factors spurring this increase. Fertility declines over the twentieth century (the baby boom excepted) may also have contributed, although fertility is probably best treated as jointly determined with employment. Researchers have also pointed to the effects of World War II in bringing women into the workforce, although Goldin (1991) presents evidence that the war had little permanent effect on women's employment, based on sample evidence that only about 20% of women working in 1950 had entered the labor force during the war, and about half of the wartime entrants left the labor force after the war (p. 755).[4]

An alternative argument is that changes in income and substitution effects led to a faster rise in women's employment than is predicted by a stable set of cross-sectional estimates. Mincer (1962) argued that substitutability between home-produced and market-bought goods has increased over time. With low substitutability, an increase in income goes into increased consumption of goods produced in the home (as well as leisure). With increased substitutability, the increase in income is more likely to go partly into increased consumption of goods purchased in the market, hence leading to less of a reduction in market hours of work. Similarly, the declining income effect would increase the uncompensated effect of the wage on women's employment. These changes in income and substitution effects would lead Mincer's (1962) time-series experiment with stable cross-section estimates to overpredict the growth of employment early in his sample period, and to underpredict it later, given rising income of both husbands and wives.

Goldin (1990) examines this hypothesis by assembling cross-section estimates at various points of time, also based on cross-city variation. Consistent with Mincer's conjecture, she finds that income elasticities appear to have declined consistently throughout the century. But uncompensated wage effects appear to have increased through 1950 (see also Goldin, 1994, for a discussion of increases in substitution effects over this period), and then fallen again, in contrast to Mincer's conjecture.[5]

To summarize, our reading of this literature suggests that the simple neoclassical model helps to explain the rise in married women's employment. However, the

[4]Bowen and Finegan (1969) also discount the importance of the war per se, finding that standard income, wage, and shadow wage variables are particularly powerful predictors of married women's employment for the 1940–1950 period. Cain (1966) reached a similar conclusion, focusing in particular on the decrease in family income owing to drafting of husbands, and the decline in the demand for women's household services.

[5]Goldin offers two reasons why the uncompensated wage effect may have fallen. First, for women work has become less solely a way to earn money, and more a means of seeking purpose and status in life. Thus, the decision to work has become less responsive to the wage. Second, with women gaining access to better jobs, work at one point in the life-cycle may increase earnings at other points in the life-cycle via, for example, training. This also would diminish the response of employment to wages.

model cannot fully explain the rise. In the next section, we develop a formal model that can explain the ''faster than expected'' growth in married women's employment, via peer- or reference-group effects driven by relative income comparisons. The incorporation of concerns over relative income has become popular in other models, especially in labor economics, to attempt to explain behavior that is difficult to fully reconcile or explain with the neoclassical model (Duesenberry, 1949; Akerlof, 1982; Summers, 1988; Frank, 1985). We are interested in exploring whether such concerns may also help to explain the rise in married women's employment.

3. A relative income model of women's employment

3.1. The neoclassical world

We begin with a neoclassical model that, on the one hand, replicates Mincer's (1962) finding that the positive uncompensated substitution effect of general wage increases on women's employment dominates the negative income effect through the husband's wage, and on the other hand, is easily extended to incorporate ranking concerns. There is a continuum of couples, each consisting of a man and a woman. Each woman is characterized by an ability level denoted by a, and each man is characterized by an ability level denoted by b. We assume that $a \in (0,A]$ and $b \in (0,B]$, $A, B < \infty$. A person who works earns $w(>0)$ per unit of ability. We assume all men work. We want to capture the idea that there is an opportunity cost to a woman working, which we interpret for now to represent lost household production. In theory this opportunity cost might depend on a number of variables including family income, the woman's ability, the number and ages of the children the couple might have, and so on. For our purposes, we will assume that the opportunity cost depends only on the husband's income, that is, there is a function $v(i)$ that specifies the value of the woman's home production.[6]

If the wage rate is w, a couple whose abilities are represented by (b, a) will have utility $(a+b)\cdot w$ if the woman works and $b\cdot w + v(b\cdot w)$ if she does not. We assume $v' > 0$, that is, that higher-income families put higher value on the home production of the woman. A consequence of this form for the opportunity cost is that of the women matched with men of the same ability level b, it is those women with the highest ability who will work, since women work if $a > v(b\cdot w)/w$. Fig. 1 illustrates the set of possible characteristics of couples, with the set divided into two components: those couples with (b, a) such that $a > v(b\cdot w)/w$, in which utility is higher if the woman works; and those for whom the reverse holds.

We are interested in the effect of an increase in the wage rate, w, on women's

[6] Section 3.4 discusses the possible implications of relaxing this specification of the opportunity cost.

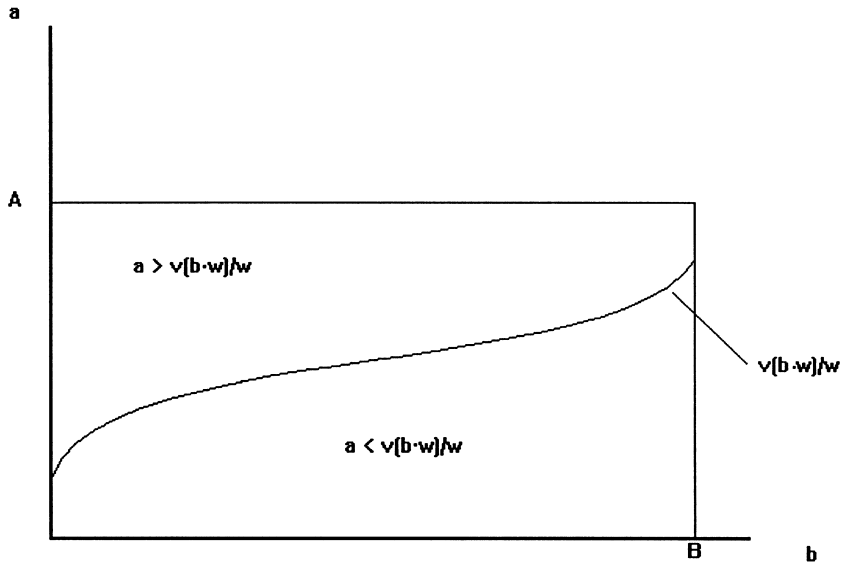

Fig. 1. The neo-classical employment decision of women.

employment.[7] If we fix the ability of the man at level b, a woman with ability a who is matched to this man, where a is such that $a \cdot w = v(b \cdot w)$, is indifferent between employment and non-employment, while any woman with higher ability married to a man of this ability will strictly prefer employment. In other words, the marginal woman matched with a man of ability b has ability $a = v(b \cdot w)/w$. If we differentiate this with respect to w, we find $\partial a/\partial w = (v' \cdot b - a)/w$. $\partial a/\partial w$ will be negative if $v'(b \cdot w) \cdot b < a$; that is, if v does not increase too quickly, increases in the wage rate will result in increased women's employment. We assume that this holds for all values of a, b, and w.

We can see from Fig. 1 what the effect of an increase in the wage would be on women's employment. As mentioned above, the woman in a couple with abilities (b, a) will work if $a \cdot w > v(b \cdot w)$. Suppose a woman does not work when the wage rate is w but enters the workforce when the wage increases to w^*. It must be the case then that $a \cdot w < v(b \cdot w)$ and $a \cdot w^* > v(b \cdot w^*)$. If we let $a' = (w^*/w) \cdot a$ and $b' = (w^*/w) \cdot b$, then $a \cdot w^* > v(b \cdot w^*)$ if and only if $a' \cdot w > v(b' \cdot w)$. That is, the increase in the wage rate from w to w^* will cause the woman to enter the workforce if and only if a proportional increase in their abilities that resulted in the same wages would cause the woman to enter the workforce. Alternatively, if one wished to predict whether the woman in the couple with abilities (b, a) would

[7]Because we are most interested in relative income concerns, we focus on employment, not participation, although we recognize that employment depends on a labor supply decision as well as a labor demand decision.

work when the wage increased by a specified amount, one could check whether a woman in a couple with higher abilities in proportion to the wage increase works at wage w. Thus, the assumption made at the end of the previous paragraph assures that in our model the positive uncompensated substitution effect of general wage increases on women's employment dominates the negative income effect through the husband's wage.

3.2. The effect of ranking concerns

The phenomenon that we are interested in is a particular externality between couples. The externality is one in which when the woman in one couple enters the workforce, this decreases the reservation wage for the woman in the comparison couple. There are a number of reasons that might underlie such a concern. For example, if in all the couples with whom one socializes, the wives begin to work, then dinner outings, cocktail parties, etc., may become more expensive. To the extent that we believe that a couple's utility may depend on the probability of retaining and socializing with their friends, a woman's reservation wage will naturally be affected by other women's decisions to work. One can think of this story as explaining how the indirect utility for money in a reduced form model might depend on the labor supply decisions of acquaintances. Alternatively, there are natural ways in which the opportunity cost associated with a woman's working might depend on other women's labor supply decisions. The reservation wage for a given woman is the utility a woman gets when at home rather than in the workforce. This may well be affected by whether a close friend is at home or working. Besides the pure recreational value of that person's company, it may be possible to share child care responsibilities or other work only when both are out of the workforce.

While there are many ways in which the optimal employment decision for a woman may depend on the decisions of other women, we will simply assume that couples are concerned with their relative income position. There is a broad array of specific forms such a concern might take, such as a concern with how far one is from the top of the income distribution, how far from the bottom, or one's percentile rank in the distribution. While we believe that it is perfectly plausible that there is a concern with relative position, we are not confident about the precise form. It is quite reasonable to believe that the specific form varies from couple to couple and often is some composite of a number of different aspects like those mentioned above. For our purposes, we will assume a simple form of the concern that has two advantages. First, it provides clear and unambiguous predictions about the effect of ranking concerns on women's employment decisions, and second, it is relatively easily translated into empirical tests. However, the qualitative characteristics of the equilibrium with which we are most interested are not driven by the particular form of the relative income concern; how the results of our model might change as we vary the form of the ranking concern is discussed in Section 3.4.

We assume that each couple is concerned with its relative position vis-a-vis a particular comparison couple. That is, the set of all couples is broken into pairs with members of each pair of couples comparing themselves with the other couple in the pair. We assume that the benefit of being the higher income family generates an increase in utility equal to $c > 0$. The increase c is independent of the size of the difference in incomes.[8] We will further assume for now that c is greater than the opportunity cost of the home production value of a woman's time, i.e., $c > v(b \cdot w)$ for all $b < B$. We assume that couples are matched randomly so that for any given couple, the distribution over the characteristics of couples with whom they compare themselves is the same as the unconditional distribution over couples' characteristics. As with the form of the concern, our assumptions about the size and determinants of the benefit and the couples with whom a given couple will compare themselves are made primarily for reasons of tractability. We will discuss below how alternative assumptions would affect the nature of our results.

To see how the concern with relative position affects women's employment decisions, we first note that for those couples (b, a) with $a \cdot w > v(b \cdot w)$, women will work. That is, if women were better off working than not when ignoring ranking concerns, the inclusion of the ranking consideration will not cause them to leave the workforce. Any effects of the ranking consideration will show up in the decisions of those women who in the absence of ranking concerns would decide not to work.

To see how the inclusion of ranking considerations might affect a woman's decision to enter the workforce, consider a couple (b, a) in which the woman would not be working in the absence of ranking considerations, depicted in Fig. 2. The set of couples who have the same combined ability as the given couple is shown by the straight line (with slope -1) through the point (b, a). If both the man and the woman in the couple whose characteristics are on this line work, they will have higher combined income than any couple whose characteristics are below the line. In addition, they will have higher income than any family with characteristics (b', a') with $b' < b$ such that the woman does not work. If the woman in the couple with characteristics (b, a) does not work, the couple will have higher income than those families with characteristics below the line with horizontal intercept b and slope -1, plus those couples in which the man's ability is less than b and the woman does not work.

We are most interested in those comparison couples whose incomes would be higher than (b, a)'s income if the woman in (b, a) does not work, but lower if she works, even if the woman in the comparison couple works. This is precisely the

[8]We should point out that there are naturally (at least) two distinct components of a concern for relative position. We discussed above why families might want their incomes to be similar to those of families with whom they socialize. This concern would properly be labeled ''conformism'' while our model might be thought of as closer to competition or emulation. For the problem we address, the distinction will not be important.

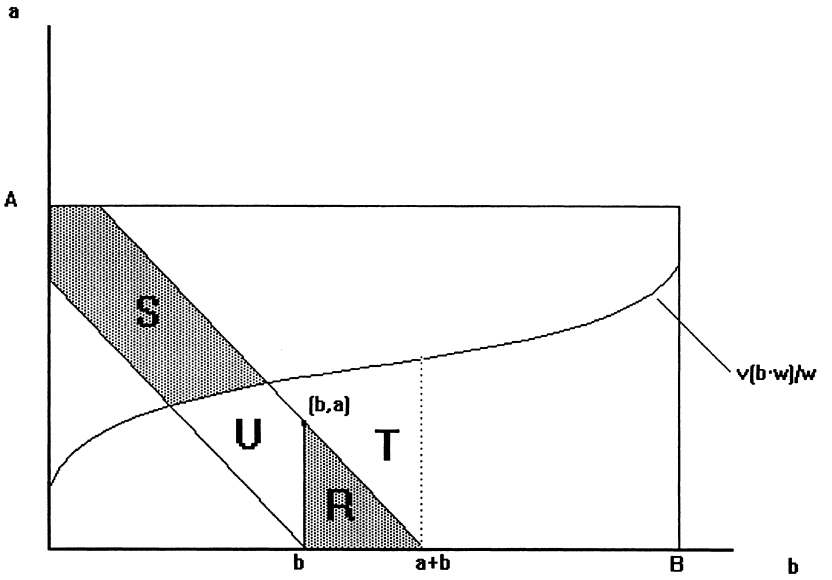

Fig. 2. The employment decision of women with relative income concerns.

set of couples who will be ''jumped over'' if the woman in (b, a) decides to work. The shaded regions in Fig. 2 represent these comparison couples. In the lower (right-hand) shaded region (R), the woman in the comparison couple is not working, but the comparison couple would have lower income even if she did work. Although the ability of the husband in the comparison couple, b', is greater than the ability of the husband in the given couple, this is more than offset by the lower ability of the woman in the comparison couple (a'). For any couple in the upper left-hand shaded region (S), the woman in the comparison couple is already working, and if the wife in couple (b, a) works, this couple earns higher income than the comparison couple. Thus, for the sets of couples in these shaded regions, we need not be concerned with possible reactions of the comparison couple in determining the incentives facing the woman in our given couple (b, a).

The issue of the comparison couple's reaction does arise, however, with respect to comparison couples that could revert to earning higher income if the wife responds by also going to work. Consider a couple (b', a') with $a' + b' > a + b > b'$ and $a' \cdot w < v(b' \cdot w)$. In Fig. 2, such a couple would be represented by a point in the region (T) formed by the vertical line through $a + b$ on the horizontal axis, the diagonal through (b, a), and the curve $v(b \cdot w)/w$. The last inequality implies that in the absence of ranking concerns, the woman in this couple (as the woman in our given couple) would not work. A consequence of the first inequality, however, is that starting from a position in which neither woman is working, the couple (b, a) can jump over the couple (b', a') if the wife goes to work, but if the

woman in (b', a') also enters the workforce, the ranking reverts to the original order.

Our exclusion of couples in region T from the set which would represent a reversal of the two couples' ranks warrants a brief discussion. We have not specified precisely the game form representing the strategic interaction of the couples, and will avoid doing so. It is straightforward to construct numerous such game forms for which the outcome we are focusing on – that in situations such as those above, neither woman enters the workforce – is an equilibrium. More importantly, however, for many such game forms the only equilibrium outcomes will be the ones on which we are focusing. Suppose the game is a sequential game in which there is no ''last mover,'' that is, whenever the woman in one couple enters the workforce the woman in the comparison couple can respond to that decision by entering herself. For such a game, a best response for a couple whose characteristics satisfy the inequalities in the above paragraph would be for the woman to enter the workforce if the woman in the couple (b, a) entered. Thus, in any Nash equilibrium, the woman in couple (b, a) entering the workforce would always be followed by the woman in the comparison couple entering as well. Thus, for equilibria of such game forms, it is appropriate to exclude those couples in region T.

Essentially the same issue arises with respect to a couple with characteristics satisfying $a + b > a' + b' > b$, $b' < b$, and $a' \cdot w < v(b' \cdot w)$. Couples (b', a') satisfying these inequalities are those in the unshaded trapezoidal region (U) in Fig. 2. For a couple represented by a point in this region, because of ranking concerns the woman in (b', a') would work if the woman in (b, a) did not, but the couple (b, a) could jump back ahead if the woman in this couple worked.

What then is the female labor supply with this ranking effect added? First, as mentioned above, all women in couples with characteristics above $v(b \cdot w)/w$ will work. In addition, some women in couples below that line will work, namely those for whom the comparison couple for that couple lies in the shaded regions. Thus, we see that for any characteristics below the line $v(b \cdot w)/w$, the proportion of the women in a couple with those characteristics who are in the workforce is equal to the proportion of women in the shaded regions. It is straightforward to see that the probability that a woman married to a man with ability b works is an increasing function of her ability.

3.3. ''Cross-section'' predictions of ''time-series'' increases in employment

To see that this model may explain why married women's employment rose faster than predicted by the simple neoclassical model, consider what predictions we might make about the effects of wage changes on employment if we fail to take account of ranking considerations. If we look at a couple (b, a) as making a decision independent of those made by other couples, an increase in the wage rate will have two effects: the extra income the woman can earn goes up because of her

own ability; and the opportunity cost goes up because of her husband's increased income. We can think about generating the analogue to the usual cross-section estimate of the effect of the wage on the proportion of women working from differences in the proportion of women working between couples with different levels of ability. For example, since the choice the couple (*b*, *a*) makes depends on the potential income of the woman and the actual income of the man, we can find existing couples whose potential and actual incomes today are precisely the same as those the couple (*b*, *a*) would be faced with given any hypothetical wage increase. Our prediction would then be that the proportion of women in couples (*b*, *a*) who will be in the workforce when the wage rate goes up by this amount is equal to the proportion of those higher ability couples for which women are working today. For the interesting case of women who would not work in the absence of ranking concerns, this proportion is represented by the two shaded regions in Fig. 3.[9] These are comparable to the shaded regions R and S in Fig. 2, but shifted to the right.

For our model, this would be an accurate estimate of the probability that the woman in the couple (*b*, *a*) would work if the wage increased for this couple only.

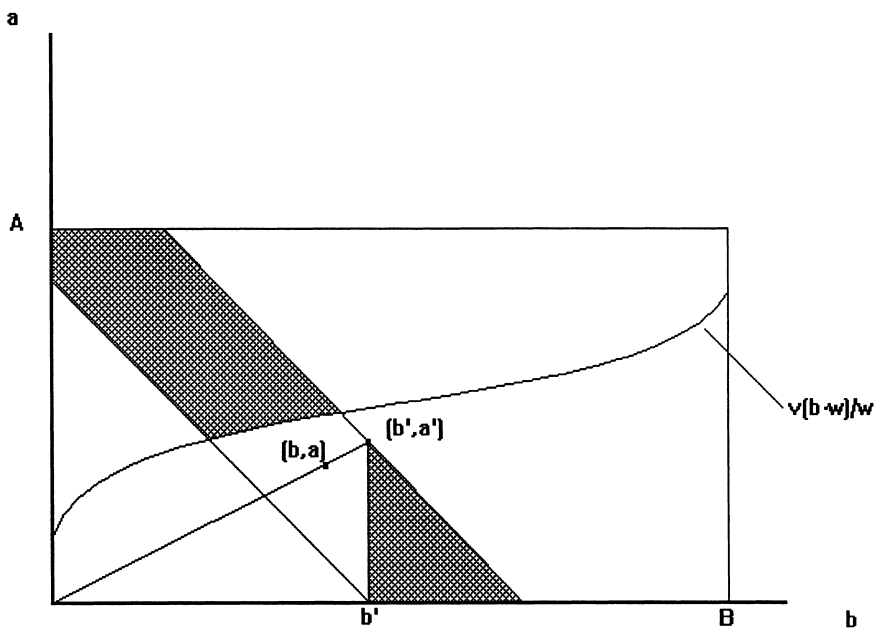

Fig. 3. The 'cross-sectional' effects of the wage on women's employment.

[9]Note that we are showing the proportion of women at (b', a') who work, not the difference between the proportion of those working at (b', a') vs. those working at (*b*, *a*).

If this is a general wage increase that affects all workers, however, this analysis will miss an important effect. We are assuming that the nature of the function v is such that when the wage rate goes up, women's employment goes up for purely economic reasons (that is, excluding ranking considerations). But this increases the set of couples for whom the ranking will be reversed if the woman in this couple enters the workforce. In Fig. 4, the increase in the set of comparison couples that would induce the woman in (b, a) to enter the workforce, after the wage increase to w', is shown by the lightly shaded region. Thus, the wage increase causes a ''purely economic'' increase in women's employment, which in turn amplifies the incentives for women to enter for ranking considerations.[10]

To sum up, our estimate from cross-section covariation between employment and wages of the proportion of women who will enter the workforce in response to a wage increase ignores the increased incentive for women to enter the workforce, which imparts a downward bias (toward zero) to the estimated effect of a general wage increase. Thus, the model with ranking concerns may help to explain why

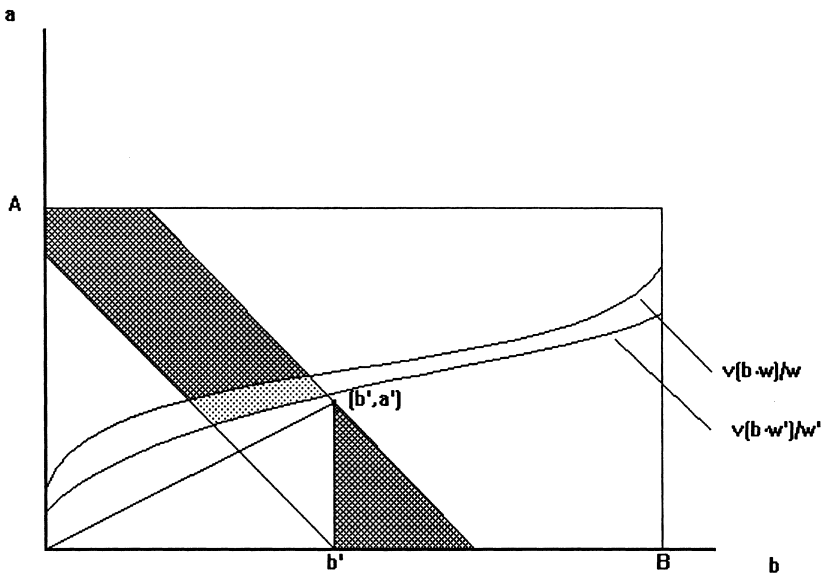

Fig. 4. The effect of general ('time-series') wage increases on women's employment.

[10]This result does not speak to the question raised in Section 2 regarding changes over time in cross-section estimates of the effects of changes in women's and men's wages. Nor does it speak to changes in relative income effects over time. Without further restrictions on functional forms, our model makes no predictions as to how income and substitution effects or relative income effects would vary as, for example, the proportion of women working rises. (See Vendrik, 1996, for an approach to such questions.)

cross-section estimates of women's employment equations fail to fully predict the time-series increase in women's employment.[11]

3.4. Discussion of model

Naturally our model incorporates a number of simplifying assumptions. While we take these assumptions to be plausible, we want to point out that they have been made for reasons of tractability and do not drive the main qualitative result in which we are interested, namely that the use of cross-section data to estimate the elasticity of female employment with respect to wages may result in estimates that are biased toward zero if ranking considerations are not taken into account. We now discuss briefly the nature of the various assumptions and how changes would affect our main conclusions.

We assumed that the opportunity cost of a woman's working was a function of the man's income only. This assumption allows us to focus on the effects of general wage increases on women's employment, and in particular how these increases cause a secondary effect due to the entrance of women into employment because the woman in the comparison couple entered the workforce. We would expect a similar effect even if the opportunity cost was a more complicated function that depended also on the number and ages of children, the woman's ability, etc. As long as wage increases lead to increased female employment excluding ranking considerations, including ranking considerations will amplify the employment increase. It would be more difficult, however, to make a simple argument that cross-section estimates of the effects of wages on employment are biased downward, because it would be significantly more difficult to trace through the model the effects of wage increases.

We also assumed that the wage rate per ability unit is the same for men and for women. If women earn some fixed percentage of what men earn, this could be incorporated into the abilities; that is, one could simply multiply women's abilities by this fixed percentage. What this assumption does not allow for is varying rates of wage increases for men and women. One could easily write down an extension

[11]We have attempted to distinguish between wage effects in which relative income concerns do not play a role (because the wage increase only occurs for one couple, paralleling the effect we would identify using micro-level data), and wage effects in which they do play a role (because the wage increase is general, as in time-series data). Mincer's (1962) estimates use cross-city variation, which may partially reflect general wage variation or at least wage variation that also affects the comparison group. However, as this section demonstrates, in that case the estimated ''neoclassical'' wage effects are likely biased upward, strengthening the conclusion that – based on his estimates – the neoclassical model fails to fully explain the rise in married women's employment.

Kapteyn and Woittiez (1990) use panel data to estimate a model with interdependent preferences in which the effects of variation grow over time because of interdependencies in labor supply decisions. Additional econometric work incorporating relative income concerns, along the lines of the research by Kapteyn and Woittiez which tries to embed such concerns in a fully-specified family labor supply model, would be valuable.

of our model to allow for such differences, and one would expect the same qualitative conclusion that estimates of wage elasticities from individual wage variation would yield downward biased estimates of the effects of general wage increases. This is for essentially the same reason as mentioned just above; the bias comes from the amplification of the employment effect due to ranking considerations.

We assumed that the couples were paired randomly with respect to the ranking comparison. This is likely the least plausible of our simplifying assumptions. We might believe that there is substantial homogeneity in the characteristics of couples in comparison groups. At first pass, we can say the same thing as in the previous two comments; this will not alter the fact that wage increases lead to increased female employment for purely economic reasons, and hence the ranking considerations will amplify the effect. Here, however, we will speculate a bit about the magnitude of the effects. We are only interested in the model if the size of the bias is not trivial. It could be that the amplification that we have identified is negligible. This would be the case if the probability that a given couple could jump over its comparison couple changed little when the wage rate increased. But if couples compare themselves to other couples with similar characteristics, this probability should be higher than in the case we considered. Thus, although one has to be careful about how one made the comparison relation more realistic, one would expect the effect we are focusing on to become more important.

We also assumed that the benefit of ''beating'' the comparison couple generated a fixed utility increment c that was greater than the opportunity cost of any woman's working. It is clear that allowing the benefit to differ across couples will have only a quantitative effect, not a qualitative one. Furthermore, given the discussion above, we might expect the comparison couples to be similar. But this implies that when a woman enters employment for purely economic reasons, the woman in the comparison couple may prefer not to work, but the decision will be close. Hence, a small increment c may send her to work.

Allowing the increment c to depend in a more sophisticated way on the ranking might raise one difficulty. If the increment depended on the amount by which a couple's income exceeded the comparison couple's, then in light of the previous suggestion that comparison couples have similar characteristics, a woman who would otherwise not have wanted to work will gain little by working since the amount by which incomes will differ if both women work will be small. On the other hand, if we allowed c also to be negatively related to the amount by which you ''lost'' the income race, the situation would be more or less as in our analysis.

Probably more significant than the detailed assumptions we made within our model are several that are embedded within the basic structure of the model. First, we have assumed that couples care whether they have higher income than their comparison couple. While we are convinced of the plausibility of such concerns, it is of interest to understand the basis of them. At the simplest level, it can be posited that these are the preferences of people (Frank, 1989). However, many

economists are more comfortable with explanations that can be derived from standard utility functions (e.g., Becker, 1976, Ch. 1), rather than alternative assumptions about tastes. We discussed at the beginning of this section how our concern for ranking might arise as a reduced form of models in which there was no direct concern for ranking, but rather that relative ranking affects some nonmarket decisions (e.g., the social opportunities mentioned above).[12]

A second defense of our approach is that it generates testable restrictions. The model does not automatically lead to the conclusion that cross-sectional estimates of the elasticity of female employment will be biased downward. For small enough gains from being the higher income couple or for large enough opportunity costs of entering the workforce, there will be no relative income effects on women's employment. Further, the model can be extended in simple ways so as to generate secondary and tertiary implications of the central ideas that can themselves be tested. For example, if we were to allow other shocks to a family's income such as shocks to the man's income or bequests from the man's family, such shocks should have an effect on the labor supply decisions of the woman in the comparison couple. The central idea of the model – that relative income matters – is specified in sufficient detail at the micro level that we need not be concerned that we cannot test the hypothesis.

We argued above that the important aspect of our model is that there be an externality associated with a woman's decision to work. One of the examples illustrating how an externality might arise centered on the decrease in utility a nonworking woman might suffer when a close friend or relative entered the workforce. The utility decrease might lower the reservation wage for such a woman. One can slightly alter our model to allow a utility increment c only if both the woman herself and the woman in the comparison couple are not working. The analysis and conclusions are essentially unchanged, with one exception. The sort of pure externality described above might lead to a utility loss to a woman who enters the workforce when the woman in the comparison couple is out of the workforce. This could lead to a situation with multiple equilibria; it could be optimal for each woman in a pair of matched couples to be out of the workforce if the other woman is, yet be optimal to be in the workforce if the other woman is.[13] A model that is similar to ours but based on such externalities rather than income comparisons could then have an ''inertial'' property that entry initially is slower than would occur in the absence of the externality. We say ''could'' because there would be delay in one of the multiple equilibria. Despite the possibility of the

[12]Cole et al. (1992) analyze a model in which no agents care directly about their relative income position, but in equilibrium relative income matters because nonmarket decisions (children's marriage prospects in that model) might be affected by relative income position. They discuss in more detail how reduced form utility functions may depend on relative wealth when direct utility functions do not.

[13]A model in a similar spirit is provided in Becker (1991), in which because of social interactions, individual demands for some goods depend on the aggregate market demand. In our case, a woman's labor supply depends to some extent on the labor supply of other women.

initial delay, once entry starts to occur, it would occur more rapidly than would be expected in the absence of the externality for much the same reason as in our model; one woman's entry into the workforce lowers the reservation wage for the woman in the comparison couple.[14]

Another feature of the model that should be discussed is the built-in assumption that the externality, or the comparison group for relative income concerns, is confined to two couples. There are several things to be said about extending our model to larger comparison groups. First, the basic logic of our model should carry over with little change. As more women become employed for purely economic reasons, this provides an additional incentive for other women not employed to become so. Second, extending the model does not present many technical difficulties. However, in discussing the effect of wage increases when ranking matters, we mentioned one difficulty that arose, namely that for a woman contemplating entering the workforce to know what the net benefit would be, she needed to predict the reaction of the woman in the comparison group. We made the simplifying assumption that the second woman would enter if doing so restored her to first place. This assumption simplified the analysis of the first woman's decision whether to work. Any extension to larger groups would exacerbate the problem of forecasting the decisions of other women in a comparison group.

4. Empirical analysis

4.1. The empirical approach

The model developed in the previous section suggests that women's employment decisions may depend not only on the individual-level variables included in standard employment equations, but also on the employment decisions or incomes of other women or couples with respect to whom relative income comparisons are made. This suggests that we might want to incorporate into a standard employment equation information on the employment status or income of such women or couples. In principle, surveys could elicit information on the employment status or incomes of women's peers, although we are not aware of any surveys that do so. A potential problem with such data, however, would be that individuals may subjectively define their peer groups so that their relative income is high compared to that of their peers. That is, the definition of one's peer group could be contaminated by relative income concerns. Thus, we instead focus on what we might think of as ''objectively-determined'' peers, namely family members.

First, we look at the relationship between women's employment and that of

[14]This type of phenomenon could, in principle, explain why the neoclassical model appears to fall short of explaining the rise in married women's labor force participation only in the latter part of this century.

their sisters-in-law, asking whether women's employment decisions depend on the employment decisions of their sisters-in-law, independently of standard variables suggested by the neoclassical model of the employment decision. We obtain data on sisters and sisters-in-law from the National Longitudinal Survey of Youth (NLSY).

Theoretical and empirical results on assortative mating suggest that common unobserved components may be much less important for sister/sister-in-law pairs than for sister/sister pairs. Thus, estimates for sister/sister-in-law pairs are much less likely to generate spurious evidence in favor of the relative income hypothesis. Becker (1981) derives theoretical conditions for positive or negative assortative mating on given traits. The general result is that we expect negative sorting on wages (as marriage markets generate a sex-based division of labor), and positive sorting on non-wage characteristics. Behrman et al. (1994) use data on twins and their spouses to study explicitly the relationship between earnings endowments (i.e., unobserved ability that affects earnings) of men and schooling of their wives. They find evidence of a negative relationship, which they interpret as evidence of negative sorting on wage-related characteristics, because schooling is positively related to the propensity to work. Assuming that brothers and sisters have positively correlated unobserved propensities to work (because of similarities in motivation or ability), then this negative sorting suggests that the unobserved propensities to work of sisters and their sisters-in-law should be negatively correlated, if anything. Based on this result, our evidence from sisters matched to sisters-in-law is not biased in favor of the relative income hypothesis. Below, we present additional findings using our data that confirms negative sorting on unobserved propensities to work of sisters and their sisters-in-law.

To begin, we posit a standard equation for the propensity to work

$$E_{it}^{*} = X_{it}\beta + \epsilon_{it}. \tag{4.1}$$

E_{it} is an observed dummy variable equal to one for woman i in period t if $E_{it}^{*} > 0$ and she is therefore employed, and X_{it} is a set of standard control variables. For the most part, the variables in X are assumed to be determinants either of the market wage or the value of home production. Thus, the argument made in the previous section regarding cross-section estimates of wage effects on employment carries over to the effects of most of the variables in X. We estimate Eq. (4.1) as a logit model.

Evidence on whether women's employment decisions are affected by the employment decisions of their sisters-in-law speaks to the general types of externalities we discussed in Section 3. But it is also directly related to relative income concerns, because the sister-in-law's employment decision affects the income of the comparison couple. To see this, note that in Fig. 2, a woman's employment decision should be unaffected by her sister-in-law's employment (denoted E') unless her sister-in-law and her brother fall in regions R or S. If they

fall in region R or S, the woman works even though she would not have worked in the absence of relative income concerns. Thus, a positive estimated relationship between E and E' arises from the fraction of women whose sisters-in-law (and their husbands) are in region S. On the other hand, the figure indicates that for women whose sisters-in-law are in region R, a negative relationship between E and E' is predicted. Thus, taking the figure literally, the model only predicts a positive relationship between E and E' when the fraction of sisters-in-law in S exceeds the fraction in R. Nonetheless, a positive estimate of the effect of E' on E is consistent with the relative income model, but not the standard neoclassical model.

We therefore add to the employment equation the observed employment status of the sister-in-law, E'

$$E_{it}^{*} = X_{it}\beta + E'_{it}\gamma + \epsilon_{it}.^{15} \tag{4.2}$$

To examine whether the results for Eq. (4.2) are influenced by heterogeneity bias – possibly, as the negative assortative mating argument above suggested, biasing the estimate of γ downward – we also estimate a specification including the woman's own lagged employment in the equation.[16]

After looking at the relationship between the employment of sisters-in-law, we turn to a test that is more directly motivated by the relative income model we developed in the previous section, and for which the model makes more precise predictions. The test is most easily explained by looking at Fig. 2, where the comparison is now made with the sister's family, since in this empirical exercise we will not be studying directly the employment of women and their sisters. First, if a woman's sister is not employed, the woman should be more likely to work if her sister's husband earns more than her own husband. To see this, note that if her sister is not employed, then her sister (and her husband) lie somewhere below the $v(b \cdot w)/w$ locus. The woman has no incentive to work if her sister's husband's ability is less than b (her own husband's ability), but does have an incentive to work if her sister's husband's ability is greater than b. This makes intuitive sense; if the sister's husband earns less than a woman's own husband, and the sister does not work, then the woman does not need to work in order for her family's income to be higher. However, if the sister's husband earns more than her own husband,

[15]Note that it is the dummy variable representing the sister-in-law's employment, rather than the latent variable underlying E'_{it}, that enters the equation. Heckman (1978) discusses models with mixed latent variable and dummy variable structures. He identifies cases in which sensible statistical models do not exist, as when the dummy variable for an event enters into the equation for the latent variable underlying that dummy variable. However, this is not the case here, since it is the *sister-in-law's* dummy variable that enters the equation for the woman's unobserved propensity to work.

[16]In addition, if the true value of γ is positive, there may be upward endogeneity bias in the estimate of γ because there is a parallel equation to Eq. (4.2) for the sister-in-law. However, endogeneity cannot alone explain a positive estimate of γ.

and her sister does not work, then by going to work the woman may be able to ''win'' the income race. She will be able to do so precisely when her sister lies in the region R. So a positive estimated relationship between the woman's employment and her sister's husband's relative income arises from the fraction of women with comparison couples in region R. By the same token, if her sister is employed, then the woman's employment should be *negatively* related to her sister's husband's relative income. In this case, the sister (and her husband) lies above the $v(b\cdot w)/w$ locus. The only case in which a woman's working changes the outcome – enabling her to ''win'' the income race – is if her sister (and her husband) lies in the region S, giving rise to a negative relationship between the woman's employment and her sister's husband's relative income.

Therefore, we also estimate the equation

$$E_{it}^{*} = X_{it}\beta + RI_{it}\gamma + \epsilon_{it}, \tag{4.3}$$

where *RI* is a dummy variable equal to one if the sister's husband's income exceeds the woman's own husband's income. The model predicts that the estimate of γ should be positive if the sister is non-employed, and negative if the sister is employed. We regard the evidence from this test as stronger for three reasons: first, the variable *RI* is more closely related to the specific model we have developed; second, it is likely to be less prone to remaining biases from common contemporaneous unobservables affecting sisters' employment; and third, for the test using relative income of husbands, the model makes more specific predictions, depending on the employment of the sister.

We do not interpret the equations we estimate as representing structural relationships obtained from the model. Specifying such structural relationships would require strong assumptions regarding functional form, the game facing women in their employment decisions, the nature of the rankings, etc., as well as data related to husbands' and wives' abilities or productive capacities. Rather, we view the empirical exercise primarily as testing the plausibility of a model that augments the neoclassical model with relative income concerns or externalities among women with respect to their employment decisions.

4.2. The data

The data are extracted from the NLSY. The survey contains many multiple-respondent households, consisting of individuals in the age range of the initial sample (14–22) in 1979 who were living in the same household. We excluded women in the NLSY's military subsample. We then matched up all sister–brother pairs, for the sister/sister-in-law equation, and all sister/sister pairs, for the relative income equation. When there was more than one possible match (because, for example, there were three or more sisters), we randomly paired up individuals,

until all possible members of the sibship were matched, without using data on any individual in more than one pair in a particular sample.[17] For each woman respondent in these samples, we identified the first observation after which she had left school, which means that she did not re-enroll in subsequent years of the survey (which extends through 1990 for the data used in this paper). We then extracted information on employment status in each year after leaving school, as well as standard control variables typically used in employment equations (education, number of children, marital status, husband's income, other non-labor income, etc.).

Because of the strong persistence in married women's participation and employment (Goldin, 1991), we view the relative income hypothesis primarily as an explanation of why some women choose to work, and others do not, rather than as an explanation of year-to-year transitions into and out of employment. Consequently, we use the earliest possible data on women, and do not utilize the full longitudinal structure of the NLSY. In particular, for the sister/sister pairs we selected data for the first two contiguous years on each sister in each pair for which both sisters were out of school and married. For the sister/sister-in-law pairs this requirement was imposed for the woman and her brother (although the woman was only required to be ever married). These restrictions led to 140 pairs of currently-married sisters, and 305 sister/sister-in-law pairs. Descriptive statistics are reported in columns (1) and (2) of Table 1.

4.3. Results

Logit estimates of the basic employment Eq. (4.1) are reported in column (3) of Table 1.[18] The specification includes controls for husband's income and other income, the woman's level of education, which is a proxy for her wage, the local unemployment rate, demographic controls, and a dummy variable for whether the husband was unemployed and collected unemployment compensation during the

[17]Siblings are identified in the NLSY sample if they were coresident in surveyed households in 1979. The issue of biases from this nonrandom selection of the sample has been raised in research using sister data to estimate the socioeconomic effects of teenage childbearing (Hoffman et al., 1993; Geronimus and Korenman, 1993). However, this concern is less important in the current context for two reasons. First, the issue originally arose for the NLS Young Women's sample, in which sisters had to be coresident between ages 14 and 24. There may be biases from this selection rule for 22–24 year-olds, among whom those remaining at home are likely to be nonrandom sample (Hoffman et al., 1993). In contrast, the NLSY age range was 14–21 in 1979. Second, the teenage childbearing literature is concerned with outcomes that are directly related to household formation, for which selection rules related to household structure seem most likely to create problems.

[18]Throughout, we report the implied partial derivatives of the probability of employment, evaluated at the sample means. The t statistic for the corresponding coefficient is reported in parentheses.

Table 1
Descriptive statistics and basic employment equation[a]

	Descriptive statics		Basic employment logit[c] (3)
	Sister/sister-in-law pairs (1)	Currently married sister/sister pairs[b] (2)	
Currently employed	0.820	0.678	...
Husband's income/$1000	12.761	16.902	−0.005
	(12.770)	(14.241)	(−2.396)
Other income/$1000	0.276	0.686	−0.005
	(0.992)	(3.257)	(−0.229)
Years of education	12.682	12.789	0.038
	(2.107)	(2.119)	(2.816)
Local unemployment rate	7.617	7.225	−0.016
	(3.585)	(3.270)	(−2.640)
Husband unemployed during year	0.059	0.079	0.049
			(0.594)
Urban	0.751	0.757	0.057
			(1.195)
Black	0.184	0.143	−0.039
			(−0.672)
Hispanic	0.170	0.132	0.040
			(0.533)
Age	24.954	25.389	0.016
	(3.127)	(3.042)	(1.782)
Number of children	1.082	1.050	−0.084
	(1.024)	(1.060)	(−3.427)
Number of children aged one year or less	0.236	0.250	0.079
	(0.448)	(0.442)	(1.731)
Divorced, widowed, separated, or spouse absent	0.161	...	−0.101
			(−1.684)
Number of individuals	305	280	305

[a]Standard deviations are reported in parentheses in columns (1) and (2), and asymptotic *t* statistics in column (3). Data from the second year on each woman are used. Husband's income and own income refer to labor income, and all income measures refer to the past calendar year. The income measures are top-coded. Because the top-codes vary across the years in the NLSY, sometimes rising and sometimes falling across the years, the smallest nominal top code for each income measure was chosen. The CPI was used to create a nominal top code for the other years of the survey, and the data were then treated as top-coded at those values. Finally, the income measures were deflated by the CPI. Husband's income and unemployment variables are set to zero for women not currently married with spouse present.
[b]The sister/sister data set is constructed to contain one record for each woman in the sample. Thus, there are two observations for each sibling pair, one with one sister as the unit of observation whose employment is to be explained (perhaps partly by her sister's employment), and another with the other sister as the unit of observation.
[c]Estimates are reported for women respondents in the sister/sister-in-law sample. Most of the estimates were similar for the women in the sister/sister sample. Partial derivatives of the probability of employment evaluated at the sample means are reported, with the asymptotic *t* statistic for the corresponding coefficient reported in parentheses.

past calendar year.[19] The estimated signs of the coefficients of the income variables (negative), education (positive), and the local unemployment rate (negative) are as expected. Children appear to have significant negative effects on women's employment, although the estimated effect of young children is positive, which is unexpected, although late child bearers (who have younger children, controlling for age) do have higher wages, possibly because of greater human capital investment (Blackburn et al., 1993).

The first test of the model entails adding to the employment equation the employment status of the sister-in-law, as in Eq. (4.2). Estimates are reported in column (1) of Table 2. The estimate of γ, the coefficient of the sister-in-law's employment, is positive and nearly statistically significant at the 10% level. The estimates imply that women with employed sisters-in-law are about seven percentage points more likely to be employed, when the probability expressions

Table 2
Logit estimates of sister/sister-in-law employment equations[a]

	(1)	(2)	(3)
Sister-in-law employed	0.067	0.072	0.059
	(1.612)	(2.001)	(1.522)
Own employment, lagged one year	···	0.259	0.247
		(6.473)	(5.668)
Husband's income/$1000	−0.005	−0.002	−0.003
	(−2.400)	(−1.611)	(−1.571)
Other income/$1000	−0.005	0.003	0.006
	(−0.232)	(0.185)	(0.227)
Years of education	0.035	0.019	0.014
	(2.573)	(1.627)	(1.169)
Currently married respondents only	No	No	Yes
Sample size	305	305	256

[a]Partial derivatives of the probability of employment evaluated at the sample means are reported, with the asymptotic t statistic for the corresponding coefficient reported in parentheses. The other control variables are the same as in column (3) of Table 1.

[19]Child support and alimony are other measures of exogenous income. However, data on these variables were not available for many years in the NLSY, and therefore could not be used.

We use education rather than an imputed wage to avoid having to impute wages for non-working women. Without knowing the form of the employment equation (which is the issue this paper addresses), we cannot confidently impute wages for non-working women.

The unemployment rate is for the standard metropolitan statistical area (SMSA), for those residing in an SMSA, and for the non-SMSA population of the state of residence otherwise. The only measure of husband's unemployment that is consistently available across the years in the NLSY is that based on receipt of unemployment compensation.

Controls such as marital status and number and age of children are potentially endogenous variables. The results reported below were qualitatively similar, although the statistical evidence was sometimes a bit weaker, when these variables were excluded.

are evaluated at the sample means. The other estimated coefficients – some of which are reported in the table – are little changed from those in column (3) of Table 1.

To examine whether the estimate of γ is influenced by unobservables, in column (2) we add a dummy variable for the woman's own lagged employment status, which is strongly positively related to her current employment status. The estimated effect of sister-in-law's employment rises, and becomes significant at the 5% level. The increase in the estimate of γ suggests that the residual in Eq. (4.1) is negatively correlated with the sister-in-law's employment status, which is consistent with the negative assortative mating on propensities to work that motivated the sister/sister-in-law test.[20]

The model described in Section 3 refers to married couples. The estimates in columns (1) and (2) are based on samples that include divorced, separated, or widowed women, and married women with absent spouses. This was done partly to boost sample sizes, partly because we are interested in feedback effects that stem from sources other than relative income concerns, such as other externalities, and partly because there is no reason to believe that relative income concerns are relevant only for currently-married women. To gauge the sensitivity of the results to the inclusion of previously-married women, column (3) reports estimates using only the subset of currently-married women (the sisters-in-law are always currently married, or there would not be any data on them, since their husbands are sampled). The estimate of γ falls only slightly, remaining positive and marginally significant.

Finally, we turn to evidence on the test that is more directly linked to the model described in Section 3, and report estimates of Eq. (4.3). This equation provides a more direct test of the model by isolating the impact of relative income concerns on women's employment decisions. However, this test entails smaller samples, as we must use the sample of currently-married sisters, split into subsamples with employed and non-employed sisters.

The results are reported in Table 3. Panel A reports results for the sample of women whose sisters are not employed. In columns (1) and (2), for logit estimates of the employment equation, the estimated effect of the husband's relative income variable is positive, as predicted; note that this arises even though the level of the husband's income is already included as a control variable. The estimated coefficient is significant at the 5% level when we exclude lagged own employment, and marginally significant when lagged own employment is included. The

[20]In contrast, in a preliminary version of this study (Neumark and Postlewaite, 1995), we found that for sister/sister pairs the estimate of γ falls toward zero when own lagged employment status is included, consistent with a positive correlation among sisters in unobserved propensities to work, and implying that studying the relationship between employment decisions of sisters is likely to lead to spurious evidence in favor of the relative income hypothesis if we cannot fully control for unobserved heterogeneity.

Table 3
Logit estimates of relative income employment equations[a]

	(1)	(2)	(3)
A. Sister Non-Employed (N=90)			
Sister's husbands's income	0.254	0.158	0.126
greater than own husband's income	(2.515)	(1.495)	(1.441)
B. Sister Employed (N=190)			
Sister's husband's income	−0.089	−0.085	−0.097
greater than own husband's income	(−1.490)	(−1.611)	(−1.772)
Other control variables included in Panels A and B:			
Lagged own employment	No	Yes	Yes
Age of sister	No	No	Yes

[a]Partial derivatives of the probability of employment evaluated at the sample means are reported, with the asymptotic *t* statistic for the corresponding coefficient reported in parentheses. The other control variables are the same as in column (3) of Table 1.

estimates indicate that – for women with non-employed sisters – when sisters' husbands' incomes are relatively higher, women's own employment probability is boosted by 16–25 percentage points. Panel B reports results for the sample of women whose sisters are employed. In columns (1) and (2) the estimated effect is negative, also as predicted, and the estimated coefficient is nearly significant at the 10% level when lagged own employment is included.[21] Overall, then, the evidence is consistent with the predictions of the relative income model that we developed in Section 3, although perhaps because of the small samples, the results are not strongly significant.

We also consider the possibility that the results in Table 3 stem from life-cycle developments that bias the results in the direction of the predicted effects. In particular, the sample in Panel A might tend to include women who are younger than their sisters. These women may still be employed, while their sisters are non-employed, in part because the latter's (generally older) husbands have attained higher earnings. Thus, there may be positive upward bias. In Panel B, the opposite might hold. The sample could be weighted toward older sisters who might be non-employed while the younger sister is employed, in part because the older

[21]If women's employment is negatively related to their husbands' income, either because husband's income responds endogenously to women's employment, or because of negative assortative mating on wage-related characteristics, there may be positive upward endogeneity bias in the estimated coefficient of the relative income variable (since husband's income appears in the denominator). This suggests that in Panel B there may be some bias against the hypothesis being tested, and in Panel A some bias in its favor. For this reason, and others discussed below, we regard the evidence in Panel B as more compelling.

sister's husband earns more, also creating positive bias, and hence strengthening the evidence of a negative effect in columns (1) and (2) of Panel B.[22] We explore this in column (3) by adding sister's age as a control variable. In Panel A, the effect of sister's husband's relative income falls slightly, consistent with this source of bias. Similarly, in Panel B the estimated effect strengthens, becoming significant at the 10% level. The estimates in both panels, and especially in Panel B, remain consistent with the relative income hypothesis.

5. Conclusions

The striking rise in the employment of married women can be partially but not fully explained by a neoclassical model including variables capturing individual women's market opportunities, home productivity, other sources of income, and taste shifters such as demographic controls. In this paper, we propose a ''relative income'' model of women's employment decisions which implies that, in addition to the variables identified by the neoclassical model, the employment decisions and incomes of other women may affect a woman's own employment decisions. We test the prediction of this model that a woman's employment decision depends on the employment or incomes of her peers in two ways.

First, we look at the interrelationship between the employment decisions of women and their sisters-in-law. Consistent with the model, we find that there is a positive effect of sister-in-laws' employment on women's own employment, after taking account of the explanatory variables suggested by the neoclassical model.

We also look at evidence on whether women's employment responds to the income of their husbands relative to that of their sisters' husbands, as the specific form of our model predicts. This evidence is largely consistent with the model's predictions. In particular, women with non-working sisters are more likely to be employed if their husbands earn less than their sisters' husbands; this result is consistent with women's employment decisions being partly driven by relative income concerns, because women with relatively low-earning husbands and non-working sisters may be able to attain higher relative family income if they work. In addition, women with working sisters are less likely to be employed if their husbands earn less than their sisters' husbands; this is also consistent with the relative income model, since such women are unlikely to be able to attain higher relative family income by working.

In our view, the weight of the evidence, combined with the theoretical results, suggests that relative income concerns can help to explain women's employment

[22] Although the inclusion of the level of husband's income as a control variable may partially control for these life-cycle effects, it may not fully do so because husband's income also influences the relative income variable.

decisions, and in particular why married women's employment in the U.S. rose faster than can be explained by the standard neoclassical model.

Finally, while this paper considers expanding the standard neoclassical model to include relative income comparisons in order to help illuminate a particular issue – the rapid rise in married women's employment – the general strategy of expanding the neoclassical model in this way opens up some interesting theoretical and empirical questions. First, if we take seriously the idea that individuals take into account relative position when making economic decisions, it seems clear that the comparison group could be quite wide. For example, in our particular context, a woman could be affected not only by the employment decisions of her sister-in-law, but also by the decisions of other friends and acquaintances. Our model incorporating relative comparisons is a particularly simple model that, as we pointed out above, gives a woman substantially less ability to alter the comparison set than possible alternative models. An interesting and more realistic model would endogenously determine the comparison relationships as well as the employment decisions. Both the theoretical development and the testing of such a model would be extremely difficult, however.

Second, there are other interesting implications of an individual's employment decision being affected by the employment status of a comparison group. One is that small exogenous shocks to the economy will be amplified. A new plant opening in a town might employ a relatively small number of people. But if those people are treated as comparisons by others, the marginal willingness to work for those others may be increased. If, as a consequence, some of those others decide to work, still others who compare themselves to this new group will have their willingness to work increased. The ripple effects of the direct employment increase will result in greater employment increases than a standard neoclassical model would suggest. There would be an analogous multiplier effect in the case of decreases in employment. A plant shutdown reduces the income for some group, and hence decreases the willingness to work for those who compare their incomes to that group. In general, the existence of the type of income externality we have presented would lead to larger responses to economic shocks than would occur in the absence of such externalities.

Acknowledgements

Postlewaite's research was partially supported by the National Science Foundation. We would like to thank Andrew Foster, George Mailath, Steve Matusz, Mark Rosenzweig, Robert Topel, Ken Wolpin, and seminar participants at Michigan State University, the University of Michigan, and the University of Pennsylvania for helpful comments and suggestions, and Sadiq Currimbhoy for research assistance. This is a substantially-revised version of NBER Working Paper No. 5044.

References

Akerlof, G.A., 1982. Labor contracts as partial gift exchange. Quarterly Journal of Economics 97, 543–569.

Becker, G.S., 1991. A note on restaurant pricing and other examples of social influences on price. Journal of Political Economy 99 (5), 1109–1116.

Becker, G.S., 1981. A Treatise on the Family. Harvard University Press, Cambridge.

Becker, G.S. 1976. The Economic Approach to Human Behavior. The University of Chicago Press, Chicago.

Behrman, J.R., Rosenzweig, M.R., Taubman, P., 1994. Endowments and the allocation of schooling in the family and in the marriage market: the twins experiment. Journal of Political Economy 102, 1131–1174.

Blackburn, M., Bloom, D.E., Neumark, D., 1993. Fertility timing, wages, and human capital. Journal of Population Economics 6, 1–30.

Bowen, W.G., Finegan, T.A., 1969. The Economics of Labor Force Participation. Princeton University Press, Princeton.

Cain, G. 1966. Married Women in the Labor Force. University of Chicago Press, Chicago.

Cole, H., Mailath, G., Postlewaite, A., 1992. Social norms, savings behavior and growth. Journal of Political Economy 100, 1092–1125.

Duesenberry, J.S., 1949. Income, Savings, and the Theory of Consumer Behavior. Harvard University Press, Cambridge MA.

Frank, R.H., 1989. Frames of reference and the quality of life. American Economic Review 79 (2), 80–85.

Frank, R.H., 1985. Choosing the Right Pond: Human Behavior and the Quest for Status. Oxford University Press, New York.

Geronimus, A.T., Korenman, S., 1993. The costs of teenage childbearing: evidence and interpretation. Demography 30 (2), 281–290.

Goldin, C.D., 1994. The u-shaped female labor force function in economic development and economic history. NBER Working Paper No. 4707.

Goldin, C.D., 1991. The role of World War II in the rise of women's employment. American Economic Review 81 (4), 741–756.

Goldin, C.D., 1990. Understanding the Gender Gap: An Economic History of American Women. Oxford University Press, New York.

Heckman, J.J., 1978. Dummy endogenous variables in a simultaneous equation system. Econometrica 46 (6), 931–959.

Hoffman, S.D., Foster, E.M., Furstenberg, F.F., 1993. Reevaluating the costs of teenage childbearing: response to Geronimus and Korenman. Demography 30 (2), 291–295.

Kapteyn, A., Woittiez., I., 1990. Preference interdependence and habit formation in family labor supply. In: Florens, J.P., Ivaldi, M., Laffont, J.J., Laisney, F. (Eds.), Microeconometrics: Surveys and Applications. Basil Blackwell, Oxford UK, pp. 231–63.

Mincer, J. 1962. Labor force participation of married women: a study of labor supply. In: Gregg, H., Lewis, H.G. (Eds.), Aspects of Labor Economics. Princeton University Press, Princeton NJ, pp. 63–97.

Neumark, D., Postlewaite, A., 1995. Relative income concerns and the rise in married women's employment. NBER Working Paper No. 5044.

Smith, J.P., Ward, M.P., 1985. Time-series growth in the female labor force. Journal of Labor Economics 3 (1, Supp.), S59–S90.

Summers, L.H., 1988. Relative wages, efficiency wages, and Keynesian unemployment. American Economic Review 78 (2), 383–388.

Vendrik, C.M.M., 1996. Bandwagon and Habit Effects on Female Labor Force Participation. Manuscript, Maastricht University.