Are Children "Normal"?

Preliminary Draft

Dan Black
Natalia Kolesnikova
Seth G. Sanders
Lowell J. Taylor

Current version: October 2007

Key words: economics of fertility, location choice, Appalachian fertility

Author affiliations, respectively: University of Chicago, Chicago, IL 60637; Federal Reserve Bank of St. Louis, St. Louis, MO 63166; University of Maryland, College Park, MD 20742; and Carnegie Mellon University, Pittsburgh, PA 15213.

Abstract

In his classic work on the economics of fertility, Gary Becker (1960) presents arguments about the likely "normality" of children. We examine Becker's contention. Our starting point is an empirical regularity about the cross section of married couples in the U.S.: as long as we restrict comparisons to households living in broadly similar locations (e.g., households located in expensive urban areas, or in rural areas), completed fertility is positively correlated with husband's income. Alas, two models rationalize the data—one in which children are "normal" and a second in which the observed pattern emerges solely as a consequence of rational sorting by households. In an effort to sort out causal effects, we undertake a rather specialized empirical exercise, an analysis of the localized impact on fertility of the mid-1970s increase in world energy prices—an exogenous shock that substantially increased men's incomes in the Appalachian coal-mining region. The evidence indicates that children are indeed "normal."

1 Introduction

In his influential 1960 paper Gary Becker sets out the basic neoclassical theory of fertility—a framework in which children are recognized as providing utility to parents in much the same way as other goods. In consumer theory in general, goods that have no ready substitutes are "normal;" increases in income induce greater consumption of these goods. Because children seemingly have few substitutes we therefore expect desired fertility to be increasing in income *ceteris paribus*.

The idea that an increase in income induces parents to have more children at first blush seems counterfactual. After all, the transition to lower fertility that occurred across the developed world since the mid 19th century has coincided with massive increases in household income. Also, in many societies—both developed and developing—high-income families have fewer children than their lower-income counterparts. For example, below we report such a relationship for married couples in the U.S. in 1990. Empirical evidence seems to suggest, in short, that higher income is associated with *lower* fertility.

Of course, in his original work on the economics of fertility, and in important subsequent work, Becker demonstrates that the neoclassical model itself provides a sensible explanation for the observed inverse relationship between income and fertility. The key insight is that the underlying economic forces that lead to increasing household prosperity often have two offsetting effects for parents: First, there is an income effect, which in isolation is expected to induce higher fertility. Second, though, increased earnings opportunities increase the value of work provided in the labor market. Because rasing children is intensive in time, particularly the mother's time, the "shadow price" of children increases relative to other relevant prices; a substitution effect reduces the desired level of child-bearing.

As Joseph Hotz, Jacob Klerman, and Robert Willis (1997) suggest, there are considerable challenges to testing the central propositions of the neoclassical theory of fertility, i.e., the predicted income and substitution effects.[1] Nonethe-

[1]The authors note that this same observation can be made about identifying predictions in

less, over the last four decades or so, numerous studies have provided relevant empirical evidence.

Many of the empirical papers on fertility focus on the proposition that increases in the value of women's time reduce desired fertility. T. Paul Schultz (1985) provides a well-crafted example of this latter empirical enterprize with his analysis of fertility in Sweden in the late 19th and early 20th century. The idea is to exploit exogenous variation in the relative world prices of basic commodities. Schultz argues that the production of dairy products was "women's work," while the production of grains and forestry products was undertaken primarily by men. It appears that a marked increase in the price of dairy products (relative to other prices) in the 1880s increased the relative market value of women's time, which in turn led to a statistically significant and economically important decline in fertility.

If children are "normal" and if men's time is relatively unimportant to the household production of children, we might have expected to see that a positive relationship between men's wages and fertility. The evidence from Sweden is not clear-cut on this issue, though. And, more generally, our reading of the literature suggests that there is little systematic evidence concerning the central proposition from the basic economic model of fertility the children are "normal goods."[2]

It is possible of course, theoretical reasoning notwithstanding, that children simply are not normal goods; more precisely, fertility might not be increasing in the men's earnings. For instance, an increase in male wages might have a substitution effect of its own (if the father's time is indeed important for raising children). Alternatively, it may be that "child services" are normal, but that these "services" have both "quantity" and "quality" components (see, e.g., Becker and Lewis, 1973, and Willis, 1973), and "quality" might be increasing in income while "quantity" is not.

---

many applied domains in economics.

[2]In his 1960 paper, Becker suggested that desired fertility might be normal even if completed fertility is not, owing to income-related differences in knowledge about contraceptives. He presents evidence from an Indianapolis survey consistent with this conjecture. This concern is less relevant now in the U.S. than it was in 1960, because of the widespread availability of oral contraceptives.

Before jumping to any conclusions, though, it appears to us that there is value in an empirical investigation of the simple proposition that children are normal goods. We focus, as a practical (and theoretically defensible) matter, on the relationship between fertility and husbands' earnings. Our empirical work examines U.S. fertility in the late 20th century.

Our starting point is a careful look at the cross-sectional relationship between income and the fertility of married non-Hispanic white women in the U.S. in 1990.[3] It is true that for our study sample there is a negative correlation between completed fertility and husbands' income. Matters are different, though, when we take a bit of care in constructing economically sensible comparison groups. Theory tells us, in particular, that in looking for an income effect we must hold constant the opportunity cost of women's time (e.g., the wage that women could earn if they spent additional time in the labor market). Thus in constructing comparison groups we take a step in that direction by looking at women with similar levels of education. Also, in theory we should be comparing households have face similar relevant prices, which would surely include housing prices. Thus, for example, one of our investigations of income and marital fertility restricts attention to college-educated women aged 40 to 50 who reside in urban areas with relatively high housing prices (New York, Boston, San Francisco, etc.). Among these women, those who have husbands with income in the lowest decile have on average 1.48 children while those with husbands' income in the top decile average 2.08 children.

In general we find a robust *positive* relationship between husbands' income and fertility in the U.S. Having established this basic stylized pattern we turn next to possible theoretical explanations. The first model we consider is a standard urban model along the lines of Roback (1982). In our set-up, local amenities differ across locations and, therefore, so too do equilibrium housing prices. We assume that housing is an input to the production of children (children need somewhere to live!). In our model children are normal; taking location as given, high-income households have more children than low-income households. But households make *joint* decisions over location and family size, and all else equal a couple that lives in an expensive location will have fewer children than a couple

[3]1990 is last year for which the Census collects "Children Ever Born" (CEB), which proves to be a valuable piece of information for our analysis.

that chooses an inexpensive location. In the resulting equilibrium, the correlation between income and fertility will be positive within each location, though it need not be positive the across the population as a whole.

Alas, there is a sensible alternative model, based on sorting, that rationalizes the basic patterns we observe for fertility, income, and location choice. In this model children need not be a normal good. Suppose, in particular, that both income and fertility is randomly assigned to households, and that the two are not related. Households then can be expected to sort into expensive high-amenity locations and inexpensive low-amenity locations as follows: Large families would generally sort into expensive cities only when they are unusually wealthy, while small families would generally sort into inexpensive rural locations only when they are unusually poor. In this sorting equilibrium, the correlation between income and fertility will be positive in each location. Such a pattern cannot, obviously, be taken as evidence that children are normal.

How can one distinguish between the two theories? We suggest that two types of empirical analysis can in principle be helpful. The first possibility, following in the intellectual footsteps of Schultz (1985), is to seek out a large and sustained exogenous shock to men's income (i.e., a shock that makes a meaningful difference to lifetime income) in a particular location, and then evaluate the impact on fertility. The second possibility is to identify households whose ability or inclination to undertake sorting by location is limited, for reasons that are external to the model, and then evaluate the relationship income and fertility for these households specifically. We present an extended example of the first approach, and then provide some thoughts about the second approach (along with some suggestive evidence).

Our empirical study focuses on fertility in the Appalachian coal-mining region from 1950 through 1990. Our focus in particular is on the chain of events that followed the high world prices of coal that were sustained over the 1970s. Our interest in the coal boom is spurred by the observation that high coal prices resulted in robust economic activity in the Appalachian coal-producing region (certain rural counties of West Virginia, Kentucky, Ohio, and Pennsylvania), while having no such effect in other counties in these same states. Specifically, the coal boom greatly improved the long-run economic prospects for men in coal counties, which provides us the opportunity to examine the causal effect on fertility of an increase in men's income. We read the evidence from the coal boom as being supportive of the proposition increases in husbands' income leads to higher fertility.

We then also provide some additional evidence about the relationship between income and fertility for groups of individuals who have strong, arguably exogenous, location ties. *Authors: One idea is to focus on the fertility of Mormons in the U.S. Mormons make location decisions that differ dramatically from other Americans. In particular, the vast majority of Mormons locate in Utah and regions of Idaho and Nevada that border Utah; there are many counties in Utah in which more than 80 percent of the population is Mormon. Thus it may be that the "effective" location choice set for many Mormon families is far smaller than for most other American households. Logically then we should expect this population to be a good study group for testing the proposition that children are normal. Similar observations might be possible for other ethnic or religious groups.*

2 Fertility and Income in U.S.: Evidence from the Cross-Section

We begin by looking at the cross-sectional relationship between income and fertility among non-Hispanic white women in the U.S. in 1990. We examine data from 1990 because it is last year for which the Census collected the variable "Children Ever Born" (CEB). We look at this variable for married women aged 40 to 50, i.e., among women born in the years 1940 through 1950. In general we observe very few births among women over age 50. Further analysis of the women aged 48 to 50 indicates that only 2.7% of these women had a child after age 40 (with most of this fertility occurring at ages 41 and 42). Thus our CEB variable is quite close to "completed fertility" for most of these women.

Table 1 presents some basic facts about fertility. Our first observation is that fertility is moderately lower among high-income families than among low-income families. It is evidence of this sort that, upon casual inspection, might lead one to believe that children are not normal. We observe other regularities: First, women with high levels of education have fewer children than women with lower levels of education. As discussed above, this relationship has been the subject of extensive research, both in the U.S. and in other countries.[4] Second, there is substantial variation across location; women who live in rural areas have more children than women who live in large cities. Third, there is variation across cohorts in completed fertility rates, with women in the younger cohorts group generally having fewer children than the older women.

With these facts in mind, we turn to regression-based depictions of the basic correlates of fertility. The goal in particular is to look at the relationship between fertility and husbands' income, conditional on other important determinants.

As far as income is concerned, we would like some measure of permanent income exclusive of income earned by the woman.[5] We have no ideal candidate for our income variable; as a viable measure we use the log of the husband's current annual income. Most of the men in our sample are in their prime earnings years, i.e., in their 40s or 50s, so current income might be quite highly correlated with lifetime income (for example, few men in the age range are in school or are retired). To the extent that current income is a noisy measure of lifetime income, the coefficient in our regression will be biased toward zero.

In setting up our regressions we want to restrict attention to women with similar levels of human capital (i.e., similar time opportunity costs). The best we can do on this account is to examine among women who have the same levels of education. We choose, in particular, two large groups for further study: those with a high school degree, and those with a college degree.

Finally, as we discuss in some detail in Section 3 below, we are concerned about the possibility that couples in the U.S. face rather different housing prices depending on where they live. Thus we also divide our sample in a rather crude fashion to account for this issue. Results are presented in Table 2.

[4]See, e.g., Schultz (1985), as discussed above, and many other papers cited in Hotz, Klerman, and Willis (1997).

[5]The mother's income is clearly endogenous, if for no other reason than the fact that the presence of additional children has the causal effect of reducing women's labor supply, as Angrist and Evans (1998) demonstrate.

The first column provides estimates for a regression that has CEB as the dependent variable and log of husband's income as a regressor, along with dummy variables for cohorts (one for each cohort of women in birth cohorts 1940 through 1950). Evidence from Table 1 indicates that fertility and *household* income are inversely related for non-Hispanic white married women aged 40 to 50. The first regression reported in Table 2, indicates that the same is true if we use the *husband's* income (in a regression that also includes cohort fixed effects). In the remaining columns, we report results for samples in which all women have the same level of education—a high school degree or a college degree—and we also break the samples apart by the household's location.

In columns (2) through (4) we consider households located in the 50 largest MSAs in the U.S. There is massive variation in the price of housing across these cities. In one variant of our regression we therefore include city fixed effects. As an alternative, we include in our regression a measure of the MSA housing costs, using Chen and Rosenthal's (2005) quality-adjusted housing values.[6] As we could expect, CEB is lower in relatively expensive cities. In both specifications log of husband's income is positively correlated with children ever born.

Figure 1 provides an alternative non-parametric presentation of the basic relationship between CEB and husband's income, in this instance dividing our sample of 50 large MSAs roughly into thirds, and presenting outcomes for the least-expensive cities and most-expensive cities. We notice that fertility is generally lower for college-educated women than for women with a high school education. Fertility is higher for women in cities with relatively lower housing costs. Most important, for our purposes, CEB is generally increasing in husband's income

We unfortunately do not have housing cost indices for rural areas comparable to those that are available for the 50 largest MSAs. We nonetheless present regressions for women who live in rural areas, in columns (6) and (7) of Table 2, using state fixed effects (which, we acknowledge, is probably not a particularly effective way of controlling for location-specific differences in housing costs). We notice that in this regression as well husband's income is positively correlated

[6]Specifically, our measure is 1000 times their index. We are grateful to Stuart Rosenthal for providing us with these data.

with CEB.

3 Theoretical Issues Concerning Fertility and Location Choice

From our brief look at U.S. data we conclude that fertility and husbands' income is negatively correlated across the population generally, but is positively correlated within similar locations. Is this persuasive evidence that children are normal? We investigate this question by briefly presenting two competing theoretical perspectives. In the first model of fertility and location, children are indeed normal. In the second model the observed pattern arises instead as a result of sorting.

3.1 A Model of Fertility and Location Choice

We begin with a very simple equilibrium model in which households jointly choose fertility and location. As will be readily apparent, our model is really just an example, intended to set the basic idea: even when children are normal, we needn't observe a positive correlation between income and fertility across the population.

In our example newly married couples act as a unitary household whose utility is determined by three factors: (1) where the family lives—in an urban location that has a valued local amenity level $A = A_U$, or a rural area that has a lower level of the amenity, $A = A_R < A_U$, (2) the number of children (c) they have, and (3) the consumption of some other good (x).

Each household i has the same tastes. Households do differ, though, by having varying levels of an exogenously given endowment of human capital, $\theta_i > 0$. For simplicity we let workers be equally productive in the two locations; they supply θ_i units of effective labor at wage w. Thus lifetime income, $y_i = w\theta_i$, is independent of location.

The price of x be the same in the two locations, and equal to 1. Thus, as is typical in urban location models (e.g., Jennifer Roback, 1984), in equilibrium housing prices must be higher in the high-amenity location; if p_U is the price of a "unit" of housing in the urban area, and and p_R is the price in the rural area, we expect to show that $p_U > p_R$. (Otherwise everyone would simply live in the high-amenity urban location.)

In our model, couples do not derive utility directly from housing. They buy housing solely for the purpose of accommodating their family. In particular, we suppose that for each extra child a couple chooses to have, they must purchase an additional unit of housing.

We let utility be given by a simple Stone-Geary form. In particular, utility is

$$A_U(c-\alpha_c)^{\gamma}(x-\alpha_x)^{(1-\gamma)} \qquad \text{in the urban location, and}$$
$$A_R(c-\alpha_c)^{\gamma}(x-\alpha_x)^{(1-\gamma)} \qquad \text{in the rural location,}$$

with $\gamma \in (0,1)$. Utility is maximized subject to a budget constraint given by $(p_U\ c)+x=y_i$ if the households is in the urban location, and by $(p_R\ c)+x=y_i$ in the rural location. Indirect utility for household i then is:

$$V^{iR} = \frac{\Gamma A_U}{p_U^{\gamma}}\left[w\theta_i - \alpha_c\left(p_U\right) - \alpha_x\right] \qquad \text{in the urban location,} \tag{1}$$
$$V^{iU} = \frac{\Gamma A_R}{p_R^{\gamma}}\left[w\theta_i - \alpha_c\left(p_R\right) - \alpha_x\right] \qquad \text{in the rural location,} \tag{2}$$

where Γ is the constant $\gamma^{\gamma}(1-\gamma)^{(1-\gamma)}$. Household i's location decision boils down to choosing the location that provides the higher lifetime utility, i.e., determining which of (1) or (2) is larger.[7] In general, that decision depends on the household's endowment θ_i.

As it turns out, the nature of our model's equilibrium is quite simple: First, as noted above, the price of housing is clearly higher in the urban location than in the rural location.[8] Second, we can show that in equilibrium there will be some

[7] We assume that income is high enough that the terms in brackets in (1) and (2) are positive.

[8] V is increasing in A and decreasing in p. Given that $A^U > A^R$, V^{iU} would exceed V^{iR} for all households i if p_R were not lower than p_U.

cut-off value, say $\widehat{\theta}$, such that relatively wealthy families ($\theta_i > \widehat{\theta}$) will live in the urban location, while poorer families ($\theta_i < \widehat{\theta}$) will live in the urban location.[9] This outcome is pictured in Figure 2.

Now consider the demand for children, which is given by

$$c(\theta_i) = (1-\gamma)\alpha_c + \frac{\gamma\,(w\theta_i - \alpha_x)}{p_U} \tag{3}$$

for the wealthier urban families ($\theta_i > \widehat{\theta}$), and by

$$c(\theta_i) = (1-\gamma)\alpha_c + \frac{\gamma\,(w\theta_i - \alpha_x)}{p_R} \tag{4}$$

for the poorer rural families ($\theta_i < \widehat{\theta}$). The demand for children, as a function of income, is shown in Figure 3.

There are two features of child demand that merit attention. First, as a consequence of our assumption that preferences are Stone-Geary, children are "normal;" so long as we restrict attention to changes in income that do not induce the family to change location, an increase in income increases desired family size. Second, for a household that is indifferent between the rural and urban locations, i.e., a family with $\theta_i = \widehat{\theta}$, fertility is lower if the couple chooses the urban location. The economic forces that drive this result are clear enough: the urban location has the advantage of having the higher level of local amenities, but this is offset by higher housing prices. Higher housing costs in turn induce the family to choose a smaller family size.

As for the overall relationship between income and fertility, we note that a *negative* correlation might well pertain. This is true even though there is a positive relationship between income and fertility in each of the two locations. If we are hoping to learn about the relationship between income and fertility

[9] To see this, first note that for the "marginal" family, (1) is equal to (2), from which one can show $\left(\frac{A_U}{p_U^\gamma} - \frac{A_R}{p_R^\gamma}\right)\left(w\widehat{\theta} - \alpha_x\right) = \alpha_c\left(A_U p_U^{(1-\gamma)} - A_R p_R^{(1-\gamma)}\right)$. Now the right-hand side of this latter expression is clearly positive, as is the second term in parentheses on the left-hand side. So $\frac{A_U}{p_U^\gamma} > \frac{A_R}{p_R^\gamma}$, which in turn shows that $\frac{\partial V^U}{\partial\theta} > \frac{\partial V^R}{\partial\theta}$ (see (1) and (2)), as depicted in Figure 2.

by looking at cross-sectional evidence, it is necessary to be careful to restrict comparisons to families that live in locations are reasonably similar in housing prices.[10]

3.2 An Alternative Model of Family Size and Sorting by Location

There is a simple alternative model—a model in which children need *not* be normal—that also rationalizes the stylized facts presented in Section 2. Consider, in particular, a situation in which family size is completely independent of income. Perhaps, for example, fertility is simply "assigned" to couples at random. Duesenberry (1960) claimed: "Economics is all about how people make choices. Sociology is all about why they don't have any choices to make." If so, then we might say that fertility is determined by "sociological factors" in this second model.

Having been assigned an expected family size, couples might be expected to choose where to live based on economic considerations. We continue to consider the case in which families choose to live in one of two locations—an expensive high-amenity urban location or a relatively less expensive low-amenity rural location. It is reasonable to expect that the relatively high cost of housing in the urban location is a greater deterrent to large families than to small families, and of course a greater deterrent to low-income families than high-income families.

Figure 4 illustrates. When families sort optimally we observe that in the cross section: (1) the high-amenity urban location disproportionately attracts high income families, (2) average family size is smaller in the urban location than in the rural location, and (3) there is a positive correlation between income and family size in both locations. These are essentially the same predictions for

[10] As we discuss shortly, there are very large differences in housing prices across locations in the U.S. Of course, in a more general model, this same logic would lead us to worry also about locational differences in other prices, e.g., the price of education and wage rates. For now we restrict attention to the very large price differences in housing, while acknowledging that a more thorough investigation might also consider price variation along other dimensions.

the cross-section we derive for our first model—the model in which fertility is a normal good.

3.3 Implications for Empirical Evaluation

How can we determine that causal impact of income for fertility? Based on the analysis above we suggest two possible ways to proceed.

The first research strategy is to find an exogenous shock to income. In particular, suppose there is a substantial increase in the market wage w that pertains in the rural location. Now if the first of our models pertains (i.e., the model in Section 3.1), fertility will increase among affected households. It might be reasonable to assume that the supply of housing in rural areas is perfectly elastic (land in rural areas might be close to free). Then for the rural location, the effect of the observed change in income, $y_i = w\theta_i$, on fertility comes from taking the derivative of (3):

$$\frac{\partial c(\theta_i)}{\partial y_i} = \frac{\gamma}{p_R}. \tag{5}$$

In our simple set-up this latter expression is in principle the same as the relationship between c and y that one would observe in a cross-sectional sample of rural households.[11]

On the other hand, if children are exogenously assigned (or in any event the demand for children is unrelated to income), as assumed in the second model outline above, the income shock will have no effect on fertility of affected households.

The second research strategy is to look for a class of households for which location sorting is impossible, or at least relatively expensive, for reasons that lie outside the model. Consider, for example, a identifiable population of households

[11] For an urban location with scarce land, it may be less plausible that the supply of housing is perfectly elastic; an increase in income shifts the demand for housing outward, raising the equilibrium price p_U. We would still expect fertility to increase if there is a positive income shock, but the equilibrium impact would be dampened by the increase in housing costs.

that is effectively constrained to live in the rural area, but that in other ways is similar to unconstrained households.

Now if our first model of fertility and location choice pertains (i.e., the model in Section 3.1), then the demand for children will be generally increasing in income for households in the constrained class of households. As illustrated in Figure 5, wealthy households in this group who cannot locate in the high-amenity urban area instead will have fertility higher than those households that can locate in the urban area.

Suppose instead that our second model pertains (i.e., the model in Section 3.2). In this case, there is no causal relationship between income and fertility. Because constrained households now simply live in the rural area, none of the sorting issues are relevant. Thus for the class of constrained households will expect to observe no correlation between fertility and income.

In the two sections that follow we provide examples of each of the approaches. In Section 4 we look at fertility in the Appalachian coal-mining region during the 1970s coal boom and subsequent bust. In Section 5 we consider fertility among groups that are arguably constrained in their location choices.

4 Fertility in Appalachia: The Effect of the Coal Boom and Bust 1970-1990

Our major empirical contribution centers on the chain of events caused by the sharp run-up in the price of coal that occurred in early 1970s, and the subsequent collapse in coal prices that followed in the early 1980s. Our argument is that this exogenous shift caused a meaningful increase in lifetime income for men in the Appalachian coal-mining region, which allows us to identify the income effect on fertility.

4.1 The Economic Environment

During the 1970s a series of events, notably changes in U.S. regulatory policy and then the 1974 oil embargo by the Organization of Petroleum Exporting Countries (OPEC), caused a massive increase in the real price of coal. In particular, after holding steady through most of the 1960s, the real price of coal (i.e., the ratio of the producer price index for bituminous coal to the consumer price index) rose significantly between 1969 and 1970 and then spiked dramatically, increasing by 44%, between 1973 and 1974. Prices stabilized for the remainder of the decade, at levels that were more than twice the 1960s level. Then the relative price of coal declined throughout the 1980s, with rapid declines occurring after 1983. By 1990 the price of coal was approximately the same as it had been in 1970.

As illustrated in Figure 6, in certain counties of Kentucky, Ohio, Pennsylvania, and West Virginia there are massive coal reserves. In many of these counties coal production plays an important role in the local economy. For instance, in 32 counties in this four-state region, more than 10% of total earnings derived directly from the coal industry in 1969. In these counties the median fraction of earnings from the coal industry is 25.3% (the mean is 30.4%). Not surprisingly, in these 32 counties, there were dramatic effects on industry employment and wages owing to the massive price swings. Table 3 provides a summary for the 32 counties (relatively to other counties). We can divide the 1970s and '80s into three periods: 1970 through 1977 was a boom period, with both employment and earnings per worker increasing rapidly in the coal-mining industry. Employment and earnings in mining were robust, and reasonably stable, during the peak period 1978 through 1982. Then employment and earnings declined substantially during the 1983 through 1989 bust period. Figure 7 shows trends in coal prices and documents the massive spike in employment in coal mining.

It appears that the labor market impact of the coal boom and bust extended beyond the mining industry in these counties; there were modest spillovers to other sectors. Dan Black, Terra McKinnish, and Seth Sanders (2005) estimate that in the 32 counties, during the coal boom two jobs were created in construction, retail and services for every 10 jobs created in the mining industry. Conversely, during the bust approximately 3.5 jobs were lost in construction, retail and services for every 10 jobs lost in mining.[12]

In short, the changes in world energy prices created a remarkable sustained period of economic prosperity in the Appalachian coal-mining region. Because mining is a heavily male occupation, this prosperity is that it had an especially large impact on men.

4.2 Fertility over the Economic Cycle: Evidence from Natality Files

Our starting point is to ask a simple question for the four-state region under examination: Are observed year-to-year changes in fertility correlated with the economic activity?

In our examination of the temporal association between fertility and the business cycle we are following a long tradition in statistical demography. Undy Yule (1906), for example, shows a pro-cyclical fertility pattern for nineteenth century England and Wales, a relationship that was driven in part by a pro-cyclical relationship between favorable economic circumstances and rates of marriage, and then subsequent fertility. As a second example, Virginia Galbraith and Dorothy Thomas (1941) find a positive association between fertility and economic activity in the early twentieth century U.S. They show that this relationship holds not only for first births, as one would expect if the relationship were being driven by changes in the marriage rate, but also by higher-order births (e.g., third and fourth births).[13]

Yoram Ben-Porath (1973) provides evidence showing that fertility was procyclical in Israel over the 1950–1970 period, and gives a nice discussion of the economic forces that would lead to such a cyclical response. As Ben-Porath argues, the role of aggregate business activity on fertility, while "somewhat fuzzy,"

[12] The boom and subsequent bust appears to have had little impact on employment in manufacturing; there was no evidence that the coal boom crowded out "export sector" jobs in the local economies.

[13] Morris Silver (1965) gives additional evidence for the U.S., using data over a longer time span.

likely includes at least three elements: First, an economic boom or bust will typically be of uncertain duration and severity, and may in some cases be taken by families as having a meaningful effect on lifetime income. As Ben-Porath notes, an economic dip can have permanent effects because "children constitute an irreversible commitment of expenditure over a fairly long period; if the depression lingers, causing a revision in the desired number of children, there will be no way back" (p. 186). Second, if households are liquidity constrained, in general there will be a stronger temporal relationship between income and current consumption, and this tie could extend to the timing of births. Third, through, to the extent that couples believe that they can time their births in relation to economic fluctuations, couples might want to have babies *during* recessions, when the opportunity cost of the mother's time is relatively low.

We are particularly interested in the fluctuations in economic activity that are the consequence of the coal boom, as over the period we study these changes were large and sustained. In coal-rich counties these swings in economic activity would be more likely than other fluctuations to be seen as having an effect on permanent income. Also, as we have argued, economic activity associated with the coal boom had particularly large effects on men's income, not women's income, and therefore likely had a relatively small impact on the opportunity cost of women's time.

Table 4 presents regression results concerning county-level fertility in West Virginia, Kentucky, and Pennsylvania, using data from 1969 through 1988.[14] Using county-level measures of fertility, based on the Detail Natality Files, our basic regression has changes in county-level birth rates as a dependent variable and has as independent variables the lagged value of difference in county-level log earnings as well as a full complement of state-year fixed effects intended to pick up secular trends in birthrates.[15]

Consider first the upper panel of Table 4. Column (1) shows that fertility is mildly pro-cyclical, as in previous literature. A 10 percent increase in earn-

[14] Ohio is excluded here because of a lack of consistent data on marital status over the period of study. We also looked at total fertility, and included Ohio for that analysis. Results are very similar to those presented in Table 4.

[15] The practice of using lagged values for economic activity stems back to Yule and Thomas, who typically use lags of two to three years on economic variables when examining fertility.

ings is associated with a 1 percent increase in births in the following year. As Ben-Porath suggests, it is not completely clear how one should interpret these findings.

Column (2) is intended to capture the contemporaneous effects of economic activity on fertility associated specifically with the coal economy. In particular, we calculate for every county each year the value of that county's coal reserves. We then form an instrument that is equal to the change in this variable, lagged by one year.[16] The idea is to look at the effects of large, exogenous, reasonably permanent changes in earnings that are focused primarily toward men. We find the contemporaneous impact of coal-related fluctuations in income to be much larger than the effects of economic activity generally; a 10 percent increase in income associated with the coal boom results in a 7 percent increase in birth rates.

The lower panel of Table 4 repeats our analysis but focuses on higher-order births. The effect of economic activity is, not surprisingly, somewhat lower than on birth rates generally, but the general pattern is the same. A 10 percent increase in income due to the coal boom results in a 5 percent increase in birth rates.

In sum, the coal boom appears to have had a large positive impact on fertility; changes in income associated with the coal boom—income increases that are large, sustained, and focused primarily toward men—lead to substantial temporal increases in fertility. We turn next to analysis with Census data to ask how these fluctuations impact cohort-level timing of fertility patterns, and to see if there are discernable impacts for completed fertility.

4.3 Evidence on the Timing of Births for Four Cohorts

As has been well documented (e.g., by Hotz, *et al.,* 1997, and many other scholars), there have been large sustained swings in U.S. fertility over the past sixty

[16]The additional lag here reflects the fact that it takes a year or longer for mining firms to expand production in response to changes in the market.

years. Our goal here is to look for an effect of the coal boom independent of these secular trends. Our basic idea is to construct patterns of age-specific fertility rates (ASFRs) for several cohorts of women—specifically women who would likely have been affected by the coal boom and comparable women who were less likely to have been so affected. We then conduct a differences-in-differences analysis, using women in non-coal counties as the base against which we look for effects of the coal boom and bust. Our data sources are the 1960, 1970, 1980, and 1990 waves of the U.S. Decenniel Census. We use the complete long-form data, which allows us to assemble samples large enough for our analysis.

Unfortunately, Census data do not allow us to construct the age-specific fertility rates directly, but for some cohorts we can make a reasonable approximation as follows. Consider a woman age 34 in a given year. Suppose the following assumptions pertain: (a) she did not give birth before she reached age 17, (b) none of her children died or left home before they reached age 17, and (c) she has no children after age 34. Then we can construct her age-specific pattern of fertility by looking at the ages of the children in the household. Of course these assumptions won't hold for all women, but this method does provide a consistent way of looking at fertility across different cohorts. We undertake this exercise for two groups of women: women who live in counties with high concentrations of coal mining, and women who live in counties that have relatively little coal. For this analysis we omit counties that have a moderate level of coal mining.

Figure 8 illustrates our first piece of evidence about the cohort effect of fertility for affected women. The first panel shows ASFRs as of age 34 for women born in 1936 (using 1970 Census data), separately for women in coal counties and non-coal counties in the four-state region. The ASFRs up to age 34 are very similar for these two groups of women. The second panel shows ASFRs up to age 34 for women born in 1946 (using 1980 data). Again, these are calculated separately for women in coal counties and non-coal counties. Women in the 1946 cohort were age 24 in 1970, at the beginning of the coal boom, and were age 26 in 1972 when the effects would be first seen in fertility rates (assuming a two-year lag). Notice that at precisely that age there appears a divergence in the ASFR patterns for these two groups of women. At every age, 27 through 34, women in coal counties have higher cohort-specific fertility rates than women in the non-coal counties.

Table 5 reports the results of a comparable differences-in-differences calculation. Column (2) of that table is formed as follows: First we calculate the differences between ASFRs in the coal counties and non-coal counties for the 1936 cohort (i.e., we just use the differences that appear between the ASFRs in the first panel of Figure 8). Second we calculate these same differences for the 1946 cohort (i.e., use the differences that appear between the ASFRs in the second panel of Figure 8). Finally we take the differences in these differences, and report the results in column (2). The boldface type highlights the years when the coal economy was booming: from ages 26 through 34, differences-in-differences are all positive. The accumulated effect is quite large; over these ages accumulated decline in fertility was on average 0.181 greater in *non-coal* counties. Before the boom, coal counties actually had a slight decline in fertility relative to non-coal counties. Thus the difference-in-difference for the total fertility rate (accumulated over the 17–34 age range) is 0.121.

Column (1) of Table 5 provides a comparable exercise for women born in 1926 and 1936. These are cohorts that would not have been much affected by the coal boom. (Women born in 1936 would have been in their mid- to late-thirties before there was any substantial economic effect of the coal boom.) This exercise suggest that fertility was in fact declining in coal counties, relative to non-coal counties, prior to the coal boom.[17]

Together, the results of columns (1) and (2) indicate that the 1970s coal boom had a substantial positive impact on completed fertility. This inference is supported by the observed timing and accumulated effects on fertility in the coal counties.

Finally, we repeat our exercise for women born in 1946 and 1956, and report the results in column (3). Again we present in boldface statistics that pertain for women that would have potentially affected by the coal boom. Women in the 1956 birth cohort, who were aged 34 in 1990, would have been subject to

[17] More precisely, the decline in the total fertility rate (TFR), which we define over the 17–34 age span, was larger in coal counties than in non-coal counties (by an average of 0.162 children) for the 1926 and 1936 birth cohorts of women. Notice from column (2), the absolute value of the decline in accumulated fertility (for ages 17 through 25) was greater in *coal* counties by on average 0.060 children per woman. In this respect they seem to be continuing a trend documented in column (1).

the impact of the 1970s coal boom and peak over ages 17 through 28. In fact, our differences-in-differences statistics show little observed effect on fertility rates for women in these ages, and during the bust years that followed, fertility declines continued to be slightly larger in coal counties than in non-coal counties. Thus overall TFRs declined by slightly more in coal counties than in non-coal counties.

We expect that further analysis is merited to explore why the coal boom apparently have a positive impact on the fertility of women in their late 20s and early 30s, but no impact on women in their teens and early 20s.

4.4 Evidence using County-Level Data on Children Ever Born

Authors: A disadvantage of the analysis presented in the previous section is that we exploit data from four cohorts only. Also, because of data limitations we can look at ASFRs only up to age 34, i.e., we're really not quite getting to completed fertility. We can undertake an additional analysis, focusing on cohort-level CEB for all women in birth cohorts 1935 through 1960. The first problem for this empirical work is constructing a CEB measure for each cohort. Given that very little fertility occurs after age 40, we can get CEB for the 1935–1950 cohorts by simply using the CEB measure in the 1990 Census. Since CEB was not collected in the 2000 Census, we can use a different approach. For each birth cohort of women 1951 through 1960, we can make "synthetic cohort" estimates, calculating the average CEB as of the 1990 Census and then using the 2000 Census calculate the number of children in the household aged 0 through 10. Adding these two numbers for each cohort should give a reasonably good cohort-specific measures of CEB. We can undertake this exercise for each of the 36 cohorts under study for each coal county and each non-coal counties in our four-state region.

This analysis must be done at a Census Center.

5 Fertility for Location-Bound Populations

Authors: Our idea here is to look at the relationship between income and fertility for individuals who are constrained in their choice of location for some reason that is independent of our model. For example, it may be that many Mormons are drawn to Utah as their family location. Preliminary analysis shows a strongly positive correlation between fertility and husband's income for families that live in heavily-Mormon areas of Utah. We can also analyze other ethnic or religious groups.

6 Concluding Remarks

Authors: To be written.

References

Angrist, Joshua D., and William N. Evans. 1998. "Children and Their Parents' Labor Supply: Evidence from Exogenous Variation in Family Size," *American Economic Review,* 88(3), 450–477

Becker, Gary. 1960. "An Economic Analysis of Fertility," in *Demographic and Economic Change in Developed Countries.* Universities–National Bureau of Economic Research Conference Series 11. Princeton, NJ: NBER.

Becker, Gary, and H. G. Lewis. 1973. "On the Interaction Between Quantity and Quality of Children," *Journal of Political Economy,* 81, S279–S288.

Becker, Gary. 1991. *A Treatise on the Family,* enlarged edition. Cambridge: Harvard University Press.

Ben-Porath, Yoram. 1973. "Short-Term Fluctuations in Fertility and Economic Activity in Isreal," *Demography,* 10(2), 185–204.

Black, Dan, Gary Gates, Seth Sanders, and Lowell Taylor. 2002. "Why Do Gay Men Live in San Francisco?" *Journal of Urban Economics,* 51(1), 54–76.

Black, Dan, Terra McKinnish, and Seth Sanders. 2005. "The Economic Impact of the Coal Boom and Bust," *Economic Journal,* 115 (503), 449–476.

Chen, Yong and Stuart Rosenthal, 2005. "Local Amenities and Life Cycle Migration: Do People Move for Jobs or Fun?" Unpublished paper, Syracuse University.

Coale, Ansley J. 1973. "The Demographic Transition Reconsidered," in *International Population Conference,* vol. 1, Leige, Belgium: International Union Scientific Study of Population.

Dehejia, Rajeev, and Adriana Lleras-Muney. 2004. "Booms, Busts, and Babies' Health," *Quarterly Journal of Economics,* 119(3), 1091–1130.

Galbraith, Virginia L., and Dorothy S. Thomas. 1941. "Birth Rates and The Interwar Business Cycles," *Journal of the American Statistical Association,* XXXVI, 465–476.

Jacobsen, Joyce P., James W. Pearce III, and Joshua L. Rosenbloom. 1999. "The Effects of Childbearing on Married Women's Labor Supply and Earnings: Using Twins as a Natural Experiment," *Journal of Human Resources,* 34(3), 449–474.

Hotz, V. Joseph, Jacob Klerman, and Robert Willis. 1997. "The Economics of Fertility in Developed Countries," in Mark Rosenzweig and Oded Stark, eds., *Handbook of Population and Family Economics.* Elsevier Science.

Roback, Jennifer. 1982. "Wages, Rents, and the Quality of Life," *Journal of Political Economy*, 90(6), 1257-1278.

Schultz, T. Paul. 1985. "Changing World Prices, Women's Wages and the Fertility Transition: Sweden, 1960–1910," *Journal of Political Economy,* 93, 1126–1154.

Silver, Morris. 1965. "Births, Marriages, and Buiness Cycles in the United States," *Journal of Political Economy,* 73(3), 237–255.

Willis, Robert. 1973. "A New Approach to the Economic Theory of Fertility Behavior," *Journal of Political Economy,* 81, S14–S64.

Yule, G. Undy. 1906. "On the Changes of Marriage and Birth Rates in England and Wales During the Past Half Century: With an Inquiry as to Their Probable Causes," *Journal of the Royal Statistical Society," LXIX, 88–147.*

Figure 1: Children Ever Born for Couples in Expensive and Inexpensive MSAs, by Husband's Income

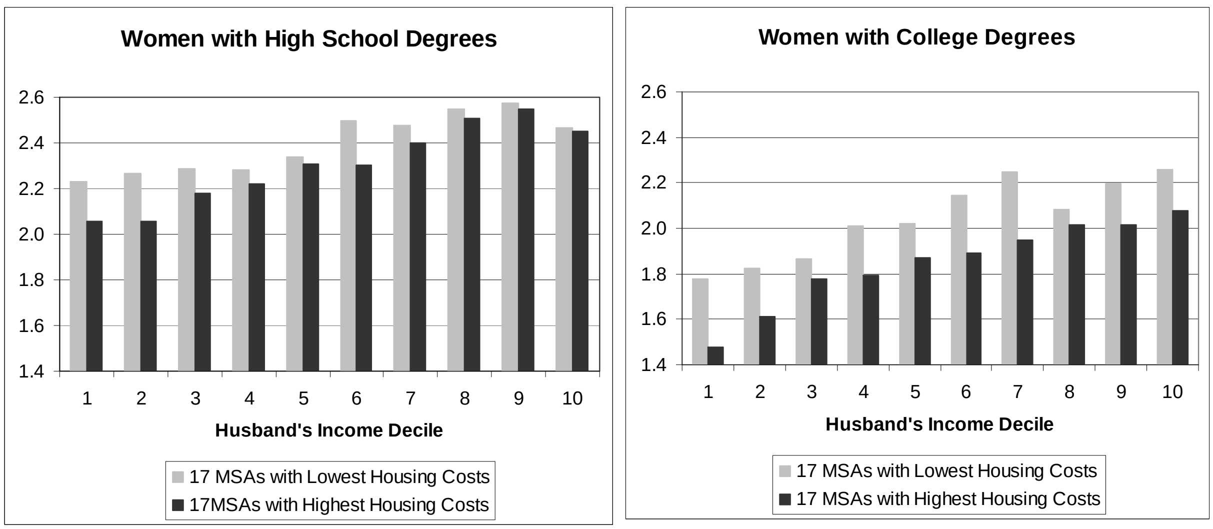

Figure 2: Indirect Utility in Urban and Rural Locations

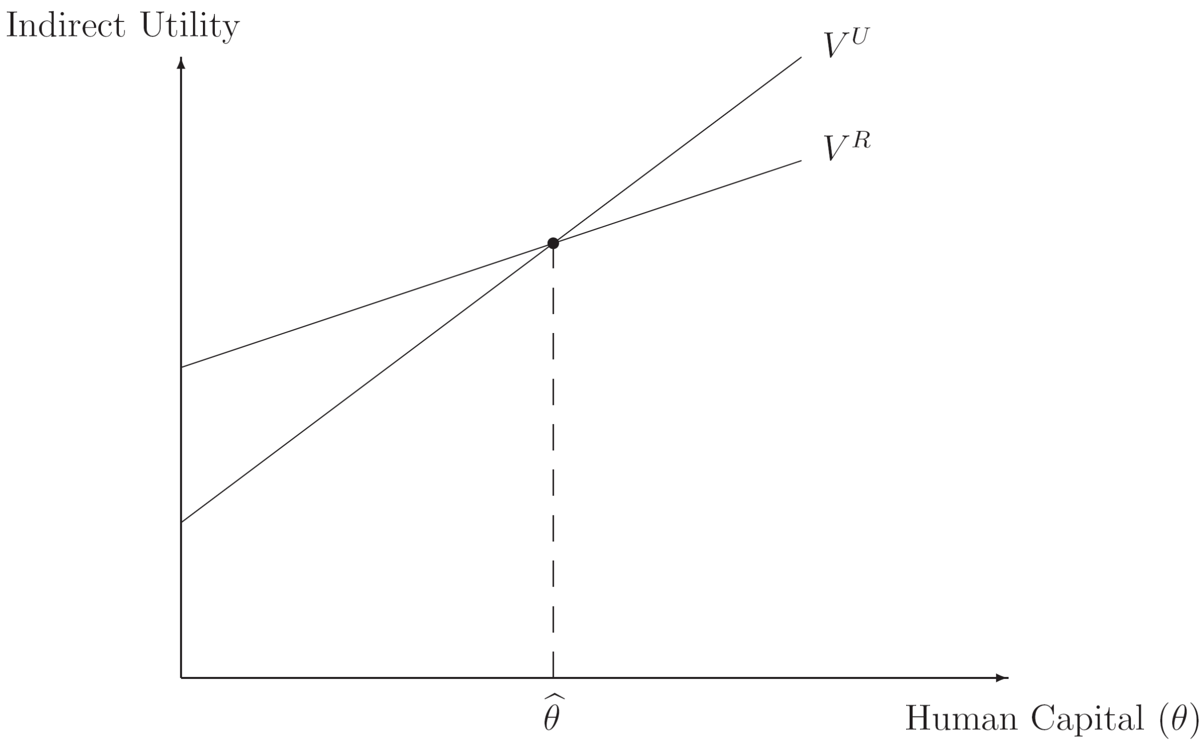

Figure 3: Income and the Demand for Children

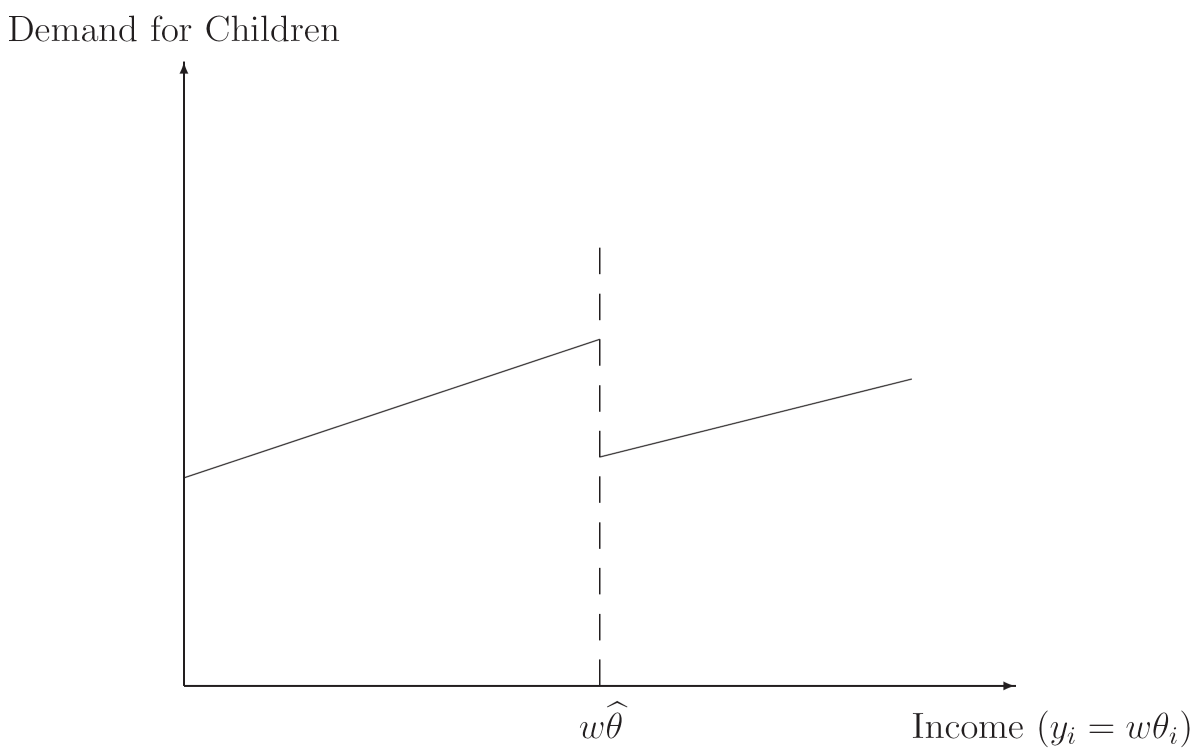

Figure 4: Sorting by Income and Family Size

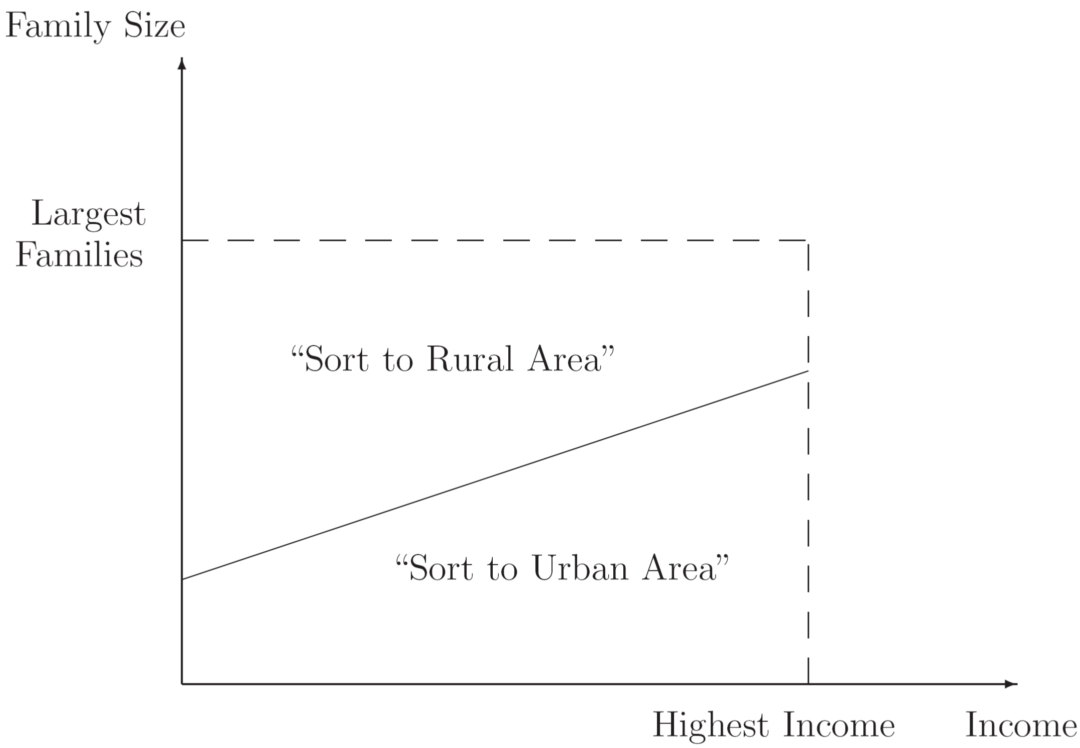

Figure 5: Fertility when Households Locate in the Rural Area Only

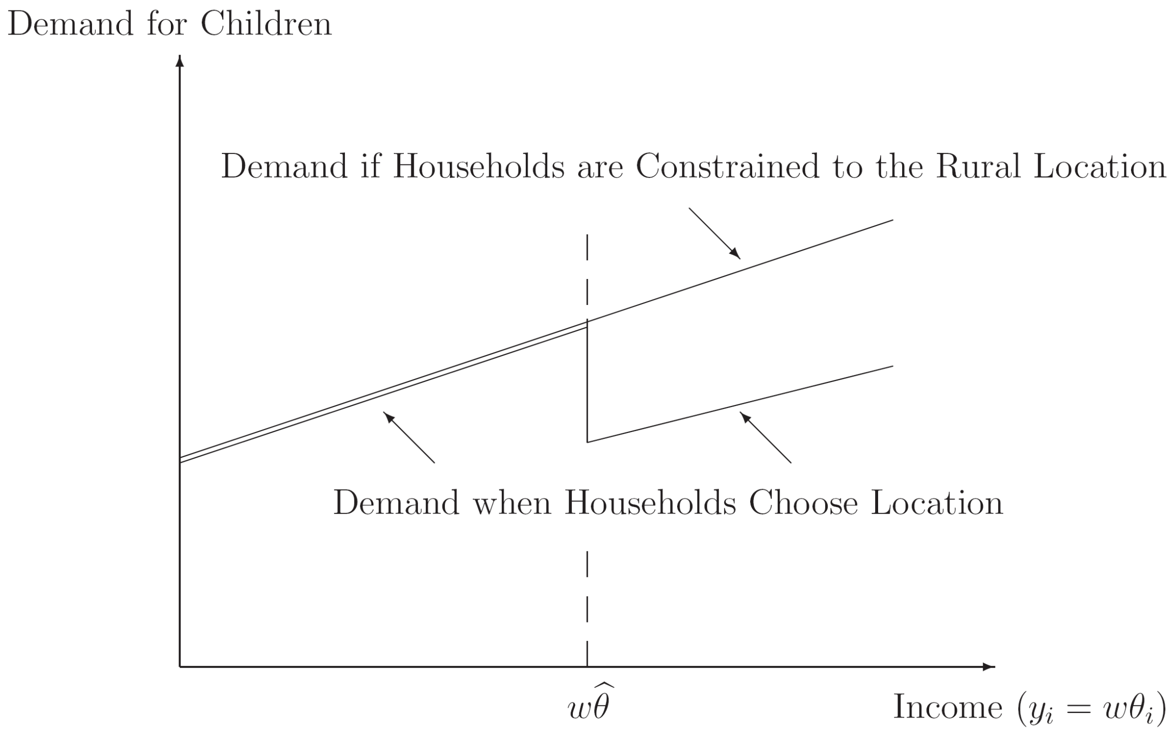

Figure 6: Coal Reserves in Ohio, Pennsylvania, Kentucky, and W. Virginia

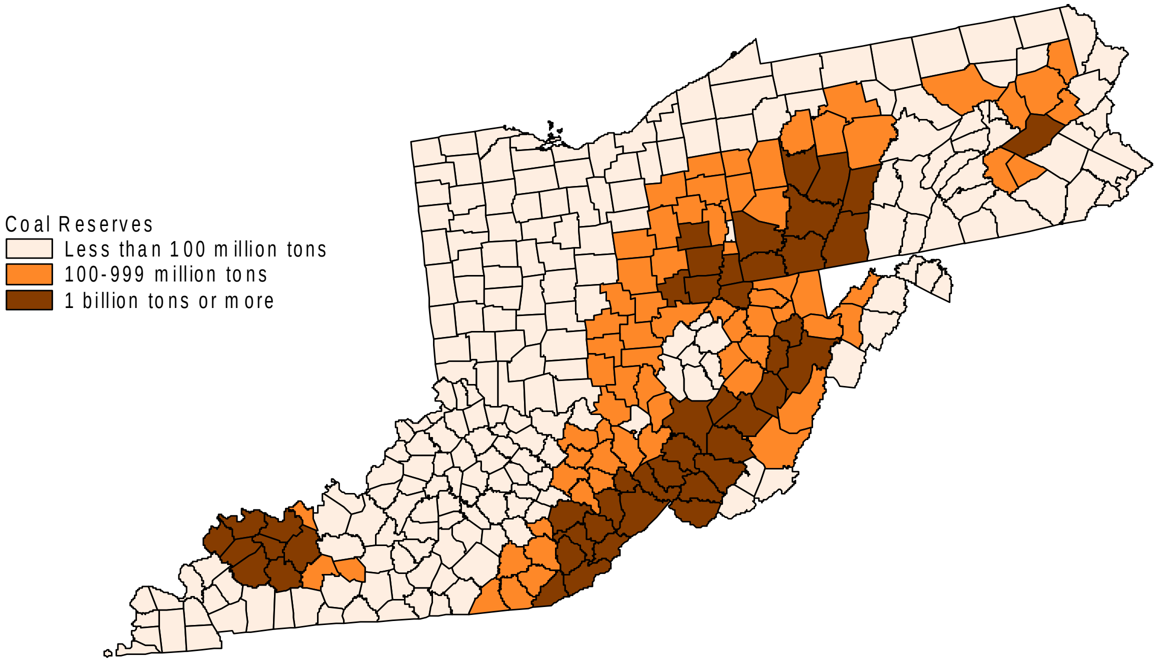

Figure 7: Price of Coal and Employment from Mining

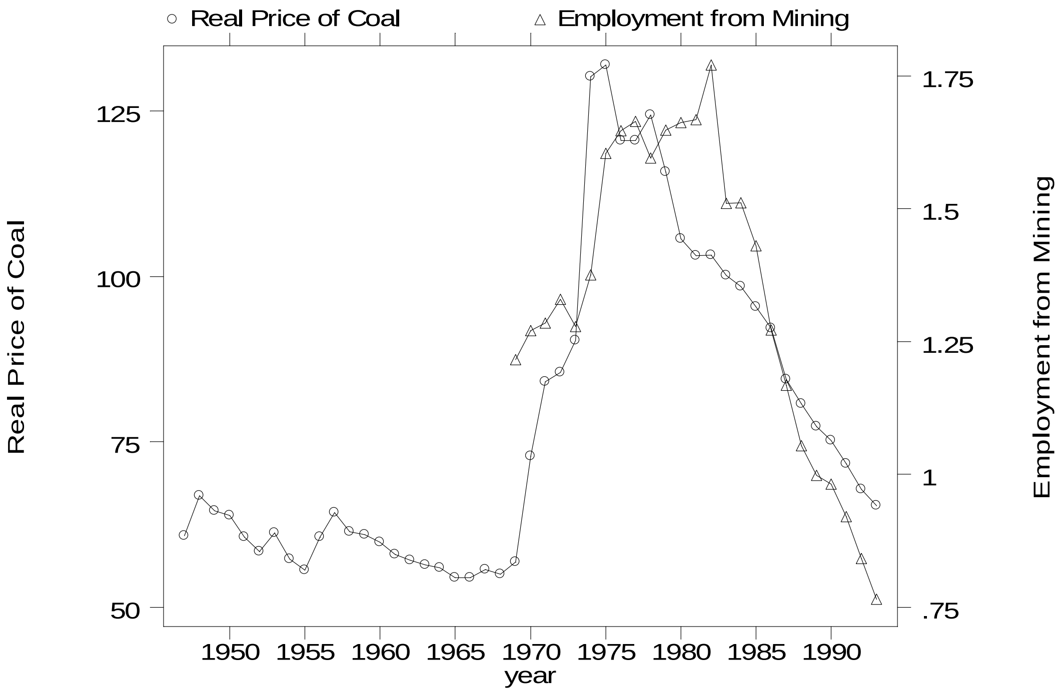

Figure 8: Age-Specific Fertility Rates for Women Born in 1936 and 1946

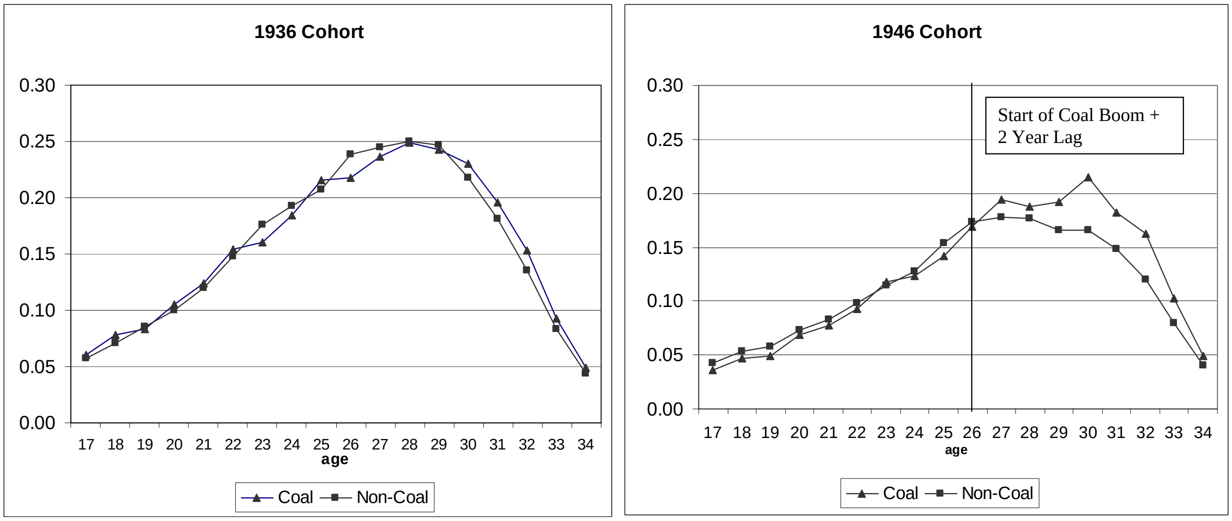

Table 1: Descriptive Statistics, Children Ever Born, Women Aged 40-50 in 1990

By Household Income:		
	Lowest Quartile	2.50
	Second Quartile	2.32
	Third Quartile	2.25
	Highest Quartile	2.16
By Husband's Income:		
	Lowest Quartile	2.32
	Second Quartile	2.26
	Third Quartile	2.29
	Highest Quartile	2.33
By Mother's Education:		
	High School Degree or Less	2.53
	Some College or College Degree	2.17
	Advanced Degree	1.75
By Location:		
	Rural	2.47
	Urban, other than 50 Largest MSAs	2.29
	Largest 50 MSAs	2.19
By Age Group:		
	Born 1945-50 (now aged 40-45)	2.15
	Born 1940-45 (now aged 46-50)	2.54

Table 2: Children Ever Born to Women Aged 40 to 50 in 1990, Regression Results

	U.S., All Women	50 Large MSAs, Women with HS Degree		50 Large MSAs, Women with College Deg.		Rest of MSAs, Women with HS	Rest of MSAs, Women with Col.	Rural Areas, Women with HS	Rural Areas, Women with Col.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
ln(Hus. Inc.)	-0.028	0.159	0.154	0.193	0.194	0.141	0.196	0.063	0.119
	(0.003)	(0.009)	(0.009)	(0.010)	(0.010)	(0.009)	(0.011)	(0.007)	(0.012)
Housing Cost		-0.026		-0.046					
		(0.003)		(0.004)					
Fixed Effects?									
Cohort	yes	yes	yes	yes	yes	yes	yes	yes	yes
MSA	—	no	yes	no	yes	yes	yes	—	—
State	—	—	—	—	—	—	—	yes	yes
N	422,427	44,925	44,925	23,671	23,671	42,710	18,398	72,734	17,834
R^2	0.027	0.032	0.03	0.034	0.028	0.028	0.027	0.057	0.057

Table 3: Growth in Employment, Earnings and Earnings per Worker: Difference between Treatment and Comparison Counties, 1970–89

Average Annual Growth in:	Difference between Treatment and Comparison Counties (Std. Error)
Total Employment	
Boom, 1970–77	0.020 (0.004)
Peak, 1978–82	-0.001 (0.004)
Bust, 1983–89	-0.027 (0.004)
Total Earnings	
Boom, 1970–77	0.050 (0.007)
Peak, 1978–82	0.005 (0.007)
Bust, 1983–89	-0.055 (0.006)
Earnings per Worker	
Boom, 1970–77	0.030 (0.004)
Peak, 1978–82	0.005 (0.004)
Bust, 1983–89	-0.028 (0.004)

Table 4: Changes in County-Level Marital Fertility Rates

Births	**OLS**	**IV**
	(1)	(2)
State-year fixed effects	Yes	Yes
Instruments:		
Value of coal reserves lagged one and two periods	No	Yes
The lagged value of differences in the log of earnings in county	0.120	0.745
	(0.0281)	(0.1249)
First-stage F-stat for instruments	—	59
N	4,356	4,356

Higher Ordered Births	**OLS**	**IV**
State-year fixed effects	Yes	Yes
Instruments:		
Value of coal reserves lagged one and two periods	No	Yes
The lagged value of differences in the log of earnings in county	0.068	0.464
	(0.0373)	(0.2166)
First-stage F-stat for instruments	—	59
N	4,356	4,356

Notes: Ohio is excluded

Table 5: Changes in Cohort ASFRs: Differences between Treatment and Comparison Counties (Boom/Peak Years in **Boldface**)

Age of Fertility	1926 and 1936 Cohorts	1936 and 1946 Cohorts	1946 and 1956 Cohorts
	(1)	(2)	(3)
17	0.002	-0.010	**-0.010**
18	0.005	-0.014	**-0.010**
19	0.004	-0.006	**-0.001**
20	0.004	-0.010	**-0.016**
21	0.000	-0.010	**0.004**
22	0.005	-0.012	**-0.007**
23	-0.024	0.018	**-0.011**
24	-0.024	0.003	**0.009**
25	0.003	-0.019	**-0.002**
26	-0.030	**0.016**	**0.007**
27	-0.026	**0.025**	**0.000**
28	-0.016	**0.011**	**0.013**
29	-0.026	**0.031**	0.008
30	-0.019	**0.037**	-0.017
31	-0.008	**0.020**	-0.010
32	0.000	**0.025**	-0.014
33	-0.005	**0.013**	-0.001
34	-0.005	**0.003**	-0.002
Diff-in-Diff in TFR (Over Ages 17 to 34)	-0.162	0.121	-0.062