

SORTING AND LONG-RUN INEQUALITY*

RAQUEL FERNÁNDEZ AND RICHARD ROGERSON

Many social commentators have raised concerns over the possibility that increased sorting in society may lead to greater inequality. To investigate this, we construct a dynamic model of intergenerational education acquisition, fertility, and marital sorting and parameterize the steady state to match several basic empirical findings. We find that increased sorting will significantly increase income inequality. Four factors are important to our findings: a negative correlation between fertility and education, a decreasing marginal effect of parental education on children's years of education, wages that are sensitive to the relative supply of skilled workers, and borrowing constraints that affect educational attainment for some low-income households.

I. INTRODUCTION

Many social commentators claim that American society is becoming more stratified, in the sense that individuals are tending to interact more with others who are similar to themselves, and less with others who are different. These interactions can include with whom one works, with whom one forms a household, with whom one goes to school, and whom one has as neighbors. Income, aptitude, skills, education, tastes, race, and ethnicity are all dimensions along which this sorting may occur.¹

It has been argued in many contexts that patterns of social and economic interactions may have important effects on society. Increased sorting, for example, may affect the degree and popularity of redistribution, the extent and geographic dispersion of crime or disease, the formation of social capital or trust, or the

* We thank Alberto Alesina and two anonymous referees for helpful suggestions. We also thank Michael Kremer, Nezh Guner, John Knowles, Victor Rios-Rull, and seminar participants at numerous conferences and universities. We thank Dan Levy for his research assistance. The authors gratefully acknowledge support from the National Science Foundation, and the first author also thanks the C. V. Starr Center for financial support.

1. Whether stratification or what we will generally call "sorting" has actually increased is a separate question and one that we will not investigate here. Kremer and Maskin [1996] present some evidence in support of the argument that sorting by skill level in the workplace has increased. Evidence on marital/household sorting appears mixed. Based on years of schooling, the evidence leads to the conclusion that sorting has not increased; however, the decreased probability that certain educational barriers will be crossed (e.g., high school graduate married to college graduate), suggests greater sorting [Mare 1991]. Sorting at the neighborhood level by income (and controlling for racial and ethnic group) appears to have increased across all groups in U. S. metropolitan areas, but particularly for Blacks and Hispanics [Jargowsky 1996], whereas anecdotal evidence of increased tracking in schools and the proliferation of magnet schools suggests that sorting in schools may be increasing.

© 2001 by the President and Fellows of Harvard College and the Massachusetts Institute of Technology.

The Quarterly Journal of Economics, November 2001

costs and benefits associated with peer effects in the classroom. Furthermore, some observers, (e.g., Wilson [1987] and Reich [1991]) have argued that over time increases in sorting may have significant consequences for the degree of inequality in society.

Will increased sorting increase inequality? What are the mechanisms by which this happens? How significant are these channels quantitatively? These are important questions that to date have received relatively little theoretical or empirical analysis. The objective of this paper is to provide some answers to these questions in the area of marital or household sorting.² In particular, we wish to examine the consequences of increased household sorting on inequality in income and educational attainment.

We develop a simple model of intergenerational education acquisition and marital sorting that allows us to highlight the key interactions between parental income, fertility, and education. Individuals are assumed to be either skilled (college educated) or unskilled (high school educated). They meet, match, and have children. Unable to borrow against future human capital, families decide how many of their children to send to college based on their family income, their children's abilities, and the expected wage differential for skilled relative to unskilled labor. The distribution of education determines wages and the distribution of income.

We solve for the steady states of the dynamic model. A steady state is a constant ratio of skilled to unskilled workers (and thus constant wages and inequality). Different initial conditions may lead to different steady states as the former determines the number of individuals who are borrowing constrained. We parameterize the model to match several basic empirical findings and use it to examine the quantitative consequences of a change in the degree of sorting.

The degree of sorting in this model is captured by the fraction of the population that gets perfectly (as opposed to randomly) matched with a partner. An increase in the degree of sorting can, in theory, either increase or decrease the skilled fraction of the population, depending on a number of factors that we discuss in our analysis. In our calibrated model we find that if marital sorting increases, then a smaller fraction of children will become skilled. This drives down wages for unskilled workers and in-

2. Throughout we use the term marriage to mean household matching independently of whether the couple is officially married.

creases those of skilled workers and also increases the degree of wage inequality. If, as a result of lower wages, borrowing constraints become tighter for low-income families, the effect on wage inequality is further magnified. Our main finding (robust to a number of alternative specifications) is that increased sorting can significantly increase income inequality.

There is a relatively small literature in the area of sorting and distribution.³ The work most directly related to ours is Kremer's [1997] recent and provocative paper that argues that even very large increases in sorting—whether marital or residential—are likely to have negligible effects on the distribution of income and education. As will be made clear in Section VI which contrasts our findings with that of Kremer's, we find three factors (all absent in Kremer's analysis and, we argue, present in the data), to be central to our results. In particular, a negative correlation between fertility and education, a decreasing marginal effect of parental education on children's years of education, and a process of wage determination that is sensitive to the relative supply of skilled to unskilled workers all contribute to our qualitative and quantitative conclusions.

In addition to the paper by Kremer, our work is related to several others in the literature. Bénabou [1993], Caucutt [1997], Cooper [1998], Durlauf [1995], Epple and Romano [1996], and Fernández and Rogerson [1996, 1997] examine the effects of neighborhood and school sorting generated either endogenously by education policies or exogenously via increased neighborhood stratification. Alesina and La Ferrara [2000] examine the effect of sorting on social capital. Banerjee and Newman [1993], Bénabou [1996], Fernández and Rogerson [1998], Galor and Zeira [1993], Loury [1981], and Ljungqvist [1993] examine the effects of the existence of borrowing constraints on the dynamic evolution of the economy and income inequality. The effects of endogenous fertility on income distribution (and vice versa) have recently been the subject of analysis of Dahan and Tsiddon [1998], Fernández, Guner, and Knowles [2001], Greenwood, Guner, and Knowles [1999], and Kremer and Chen [1999] among others.

The outline of the paper follows. In the next section we describe the model and its steady states. In Section III we analyze the effects of changes in sorting. In Section IV we use data to parameterize the model, and in Section V we use our parameter-

3. See Fernández [2001a] for a review of this literature.

ized model to assess the effects of a large increase in sorting. Section VI reviews Kremer's analysis and contrasts it with our own. Section VII examines the robustness of our results to alternative parameterizations, and Section VIII concludes.

II. THE MODEL

To examine the effects of marital sorting on the process of intergenerational education transmission and income inequality requires a dynamic model that incorporates marriage, fertility, education, and the determination of income. The interaction of these factors easily yields a nontractable model (see Greenwood, Guner, and Knowles [1999] for a computational approach to this problem) so, wherever possible, we choose to model these decisions in as simple a way as possible, keeping many elements exogenous (in particular, fertility and marriage decisions) in order to highlight the interactions that are central to our analysis.⁴

The story our model tells is a simple one. In each period the adult population is characterized by a distribution of education or skill levels. We assume that individuals are either skilled or unskilled and that a competitive labor market determines the relative wages of these workers. These individuals meet and match with their household partner via an exogenous matching process that exhibits positive assortative matching. Couples have children and, based on the number of children, their aptitudes, family income, and expected wages, they decide the education levels of their children. This generates the next generation's distribution of education (skill levels). A more formal description follows.

A. Marriages

Consider a population at time t whose number is given by N_t and some division of that population into skilled workers, N_{st} , and unskilled workers, N_{ut} , where

$$(1) \quad N_t = N_{st} + N_{ut}.$$

For our purposes, skill levels will be synonymous with an educational attainment. All college-educated workers are skilled (s); all others are unskilled (u).⁵

4. As we argue in the conclusion, we believe that most plausible ways of endogenizing fertility and marital decisions will reinforce our conclusions.

5. Our model abstracts from any possible differences between women and

Each individual is matched with another, resulting in a "marriage" according to the following mechanical process. In order to capture the degree of sorting in the economy, we allow some fraction of marriages, say θ , to be perfectly matched; i.e., a skilled worker matches with another skilled worker or an unskilled worker matches with another unskilled worker. The remaining fraction of the population is matched in a random fashion resulting in homogamous and nonhomogamous households. Thus, marriages will belong to one of three categories: skilled matches with skilled (denoted by h for high type), skilled matches with unskilled (denoted by m for mixed or middle type), and unskilled matches with unskilled (denoted by l for low type). These categories will also correspond to the relative position of couples in the income distribution.

Given the degree of assortative matching θ and the distribution of the population at time t into skilled and unskilled, the fraction of all matches at time t that are high type, λ_{ht} , is given by

$$(2) \quad \lambda_{ht} = \beta_t[\theta + (1 - \theta)\beta_t],$$

where β is the fraction of skilled workers in the population; i.e.,

$$(3) \quad \beta_t = N_{st}/N_t.$$

To see why, note that a person is skilled with probability β_t . Conditional on being skilled, with probability θ she will be matched for sure with another skilled individual; and with probability $(1 - \theta)$ she will be randomly matched, in which case there is probability β_t of being matched with another skilled agent. Similarly, the fraction of all matches at time t that are middle type is given by

$$(4) \quad \lambda_{mt} = 2(1 - \theta)\beta_t(1 - \beta_t);$$

whereas the fraction that are low type is given by

$$(5) \quad \lambda_{lt} = \theta(1 - \beta_t) + (1 - \theta)(1 - \beta_t)^2.$$

Of course, $\lambda_{ht} + \lambda_{mt} + \lambda_{lt} = 1$.

men such as preferences, earnings, or differences in how families educate girls and boys.

B. Fertility

Fertility probably depends on parental education, income, culture, and technology among other things. We simplify matters by assuming that the probability distribution over fertility is determined entirely by the educational backgrounds of the parents. Thus, the probability of a family having a certain number of children, n , is a function only of marriage type and can be denoted by ϕ_{nj} , where $n = \{0, 1, 2, \dots, \bar{n}\}$ and $j = h, m$, or l . We denote the average fertility for families of type j by $f_j = \sum_{n=0}^{\bar{n}} n \phi_{nj}$.

A child can be of two "aptitude" types which we denote by either high or low, the significance of which will be made clear shortly. The probability that a given child is of high aptitude, $\gamma_j, j = h, m, l$ is allowed to differ across family types but not across families within the same category.⁶ Realizations are independent across children. The probability, therefore, that a family with a total number of children n has $a \leq n$ children of high aptitude is $\gamma_j^a (1 - \gamma_j)^{n-a} \binom{n}{a}$, where $\binom{n}{a}$ is the binomial coefficient (equal to the number of combinations of n things taken a at a time).⁷

C. Education

A family's decision to send a child on to college is determined by the child's aptitude, family income, and expected wages. If a child with high aptitude obtains a college education, we assume she receives one unit of skilled human capital, whereas a low-aptitude child who goes on to college is assumed to obtain zero units of skilled human capital.⁸ The quantity of unskilled human capital that a child obtains is assumed to be independent of her aptitude level; i.e., all individuals who obtain only a high school education have the same level of human capital. The aptitude (and education) of a child is assumed to be perfectly observable to all.

6. In this sense perhaps the term aptitude is a misnomer since, strictly speaking it is not genetically determined (otherwise we would have to keep track of whether a couple included 0, 1, or 2 high-aptitude individuals). It is best thought of as a high or a low ability to obtain marketable skills from college. This ability is assumed to depend on parental education and hence differs across family types.

7. Our specification assumes that children's aptitude depends only on average parental education, and not on how this is distributed across the mother and the father. In developing countries there is some evidence to suggest that the mother's education is more important. However, in the United States, our assumption is consistent with the findings reported in Kremer [1997]. He found no evidence for differential effects between mother's and father's education.

8. We could easily assume that a low-aptitude child ends up with $\varphi < 1$ units of skilled human capital. This would multiply the number of potential steady states we have to examine but not add any new factor of interest to our analysis.

We assume that the cost of sending a child to high school is zero; whereas a positive (constant) cost v must be incurred before obtaining a higher education. To render the decision of whether to send a child to college as simple as possible, we assume that, subject to obtaining a minimum per family member consumption level of \bar{c} , a family would always desire to send a high-aptitude child to college if the net return from doing so, $w_s - v$, exceeded the return from high school, w_u . More formally, if a family with n children of which a are of high aptitude sends $r \leq a$ of them on to college, and has per capita consumption equal to c , we assume they receive utility,

$$(6) \quad U = \begin{cases} (c - \bar{c}) & \text{for } c < \bar{c} \\ (c - \bar{c}) + \frac{r}{(2+n)}w_s + \left(\frac{n-r}{2+n}\right)w_u, & \text{otherwise,} \end{cases}$$

where w_s and w_u are next period's wages for skilled and unskilled workers, respectively.⁹

We assume a constant returns to scale aggregate production function given by

$$(7) \quad F(N_s, N_u) = N_u F\left(\frac{N_s}{N_u}, 1\right) \equiv N_u F\left(\frac{\beta}{1-\beta}, 1\right) \equiv N_u f(\beta) \\ f' > 0, \quad f'' < 0.$$

Assuming a competitive labor market, it follows that wages are determined only by the value of β :

$$(8) \quad w_s(\beta) = (1 - \beta)^2 f'(\beta) \quad \text{and} \quad w_u(\beta) = f(\beta) - \beta(1 - \beta)f'(\beta),$$

where the assumptions in (7) imply that skilled wages are decreasing in the ratio of skilled to unskilled workers and the opposite for unskilled wages. Note that no family would want to send their child to college if the fraction of skilled workers exceeds $\bar{\beta}$, where $\bar{\beta}$ is defined by

$$(9) \quad w_s(\bar{\beta}) = w_u(\bar{\beta}) + v.$$

We assume henceforth that $\bar{\beta}$ is strictly positive. Note, furthermore, that $\bar{\beta}$ would be the fraction of the population that would

9. We use this linear utility function for simplicity only; we could specify a concave utility function. It would be equally simple to incorporate discounting of children's future income or differential weights on family members' consumptions.

attend college if there were no borrowing constraints and, in aggregate, the fraction of high-aptitude children exceeded $\bar{\beta}$.

D. Budget Constraints

The utility maximization problem of a family of type j with n children, a of whom are high aptitude is given by the maximization of U as specified in (6) subject to a budget constraint. Note that in the absence of any impediments to borrowing against future income, all high-aptitude children would attend college as long as $\beta \leq \bar{\beta}$ in the subsequent period. With borrowing constraints, however, household income is an important determinant of the number of children that a family can afford to send to college.

In what follows, we assume that families are unable to access credit or insurance markets.¹⁰ For interpretational purposes, however, we think that it is important to note that these borrowing constraints need not be thought of as constraining directly the capacity of a family to send a child to college (which is debatable as some colleges are close to free).¹¹ Instead, in a richer model the inability to borrow against a child's future income could serve to constrain a family's residential choice and consequently the quality of the high school their children can attend.¹² This would then affect both the amount of human capital obtained from high school attendance and the probability that the child attends college.

Thus, the utility maximization problem of equation (6) is subject to a household-income budget constraint:

$$(10) \quad \begin{aligned} (2 + n)c + rv &\leq I_j(\beta) \\ 0 &\leq r \leq a, \end{aligned}$$

10. It is, of course, not necessary to shut down capital markets altogether in order to obtain the result we desire—that the maximum number of children a family can afford to send to college is a function of family income. It is simple to write down microfoundations (e.g., moral hazard or imperfect enforcement technology) for this or less extreme assumptions (see, for example, Ljungqvist [1993], Banerjee and Newman [1993], or Galor and Zeira [1993]). Note also that families would want to pool risk since the number of high-aptitude children each has is stochastic.

11. Although, of course, there are subsistence costs to be met, etc. Indeed Behrman, Pollak, and Taubman [1989] argue that unequal access to financing for college can help explain differences in educational attainment.

12. For a model that examines the consequences of local provision of education in which the cost of housing in a wealthy community prevents lower-income individuals from accessing high-quality primary and secondary education, see Fernández and Rogerson [1998].

where

$$(11) \quad I_j(\beta) = \begin{cases} 2w_s(\beta) & \text{for } j = h \\ w_s(\beta) + w_u(\beta) & \text{for } j = m \\ 2w_u(\beta) & \text{for } j = l. \end{cases}$$

Note that a higher fraction of skilled workers implies lower wages for skilled workers and higher ones for unskilled workers. Hence, an increase in β implies tighter budget constraints for high-type families and looser ones for low-type families. Whether the budget constraint for middle-type families is loosened or tightened depends on whether the increase in the wage of unskilled workers is greater than the accompanying decrease in skilled wages; i.e., on whether $N_s - N_u$ is positive.

E. Steady States

It is straightforward to show that if β_t is the fraction of the population that is skilled in period t , then next period's value of β is uniquely determined. The dynamic evolution of this economy will of course depend on the fertility of each family type, the fraction of children of each type that are of high aptitude, wages, minimum required consumption, and the cost of college.

Although the economy will follow a unique path starting from any initial condition, in general this economy may have multiple steady states. To see why this is the case, note that the fraction of skilled workers in the economy determines the income level for each marriage type, which in turn determines who can afford to attend college. A higher fraction of skilled workers implies a higher wage for unskilled workers and a lower one for skilled workers. This tightens constraints for high-type families, loosens constraints for low-type families, and loosens (tightens) them for middle-type families if $\beta > (<)$ 0.5. Thus, a low initial proportion of skilled workers can be reinforcing if as a consequence of low unskilled wages a large fraction of families find themselves constrained. Similarly, a high initial proportion of skilled workers can be reinforcing if as a consequence of high unskilled wages a small fraction of families find themselves constrained. This positive feedback effect can give rise to multiple steady states.

Suppose that in equilibrium a family of type j can afford to

send z_{nj} of their n children of college (and finds it desirable to do so). To solve for the fraction of children that type j families will send to college in aggregate, $\Gamma_j(z_j)$, $z_j = (z_{1j}, \dots, z_{nj})$ requires finding the distribution of high-aptitude children over type j families and evaluating which of these are constrained. In particular, we have

$$(12) \quad \Gamma_j(z_j) \equiv \frac{1}{f_j} \sum_{n=1}^{\bar{n}} \phi_{nj} \left[\sum_{a=0}^{z_{nj}} \binom{n}{a} \gamma_j^a (1 - \gamma_j)^{n-a} a + \sum_{a=z_{nj}+1}^n \binom{n}{a} \gamma_j^a (1 - \gamma_j)^{n-a} z_{nj} \right].$$

Recall that f_j is average fertility, ϕ_{nj} is the probability of having n children, and γ_j is the probability that a child is of high aptitude, in each case for a match of type j . Hence, the first summation term within the square brackets is the number of children that attend college from unconstrained families of type j with n children (i.e., those whose number of high-aptitude kids is fewer than z_{nj}) and the second summation is over the number of children that attend college from constrained families of type j with n children.¹³

The steady states of the economy are the fixed points of the dynamic system below:

$$(13) \quad \beta_{t+1}(\theta) = \frac{N_{st+1}}{N_{t+1}} = \frac{\sum_j \Gamma_j(z_j(\beta_t)) f_j \lambda_{jt}(\beta_t; \theta)}{\sum_j f_j \lambda_{jt}(\beta_t; \theta)};$$

i.e., a $\hat{\beta}$ such that $\beta_{t+1} = \beta_t = \hat{\beta}$.

Note that each z_j must be consistent with the equilibrium family budget constraint (i.e., $\forall n$ subject to $\phi_{nj} > 0, \forall j, (2 + n)\bar{c} + z_{nj}v \leq I_j(\hat{\beta})$ and either $(2 + n)\bar{c} + (z_{nj} + 1)v > I_j(\hat{\beta})$ or $z_{nj} = n$) and, of course, parents must wish to send their children to college; i.e., $\hat{\beta} \leq \bar{\beta}$. We will restrict our attention to locally stable steady states, and thus impose $\partial\beta_{t+1}/\partial\beta_t |_{\beta_t=\hat{\beta}} < 1$ as an additional constraint.

III. CHANGES IN SORTING

How will a change in the degree of sorting (i.e., in the level of θ) affect the steady-state level of β ? In answering this question, it

13. If a family of type j with n children is not constrained, we simply indicate this by $z_{nj} = n$.

is useful to distinguish two cases: one in which the change in sorting does not affect the maximum number of children any family type can afford to send to college, and the other in which it does. In the first case what we will call the “bindingness” of borrowing constraints is not affected; in the second case it is.¹⁴

Assume initially that the bindingness of constraints is not affected (i.e., assuming that the z_j 's do not change and hence that the Γ_j 's are constant). Using the implicit function rule on (13) yields

$$(14) \quad \frac{d\hat{\beta}}{d\theta} = \frac{\sum_j f_j(\partial\lambda_j(\hat{\beta};\theta)/\partial\theta)(\Gamma_j - \hat{\beta})}{\sum_j f_j\lambda_j(\hat{\beta};\theta) + \sum_j f_j(\partial\lambda_j(\hat{\beta})/\partial\hat{\beta})(\hat{\beta} - \Gamma_j)}.$$

Taking the derivatives of the λ_j 's (given by (2), (4), and (5)), evaluating at $\beta = \hat{\beta}$, and substituting into the expression above yields

$$(15) \quad \begin{aligned} \frac{d\hat{\beta}}{d\theta} &= \frac{\hat{\beta}(1 - \hat{\beta})[f_h(\Gamma_h - \hat{\beta}) - 2f_m(\Gamma_m - \hat{\beta}) + f_l(\Gamma_l - \hat{\beta})]}{D} \\ &= \frac{\hat{\beta}(1 - \hat{\beta})[(f_h\Gamma_h - 2f_m\Gamma_m + f_l\Gamma_l) - \hat{\beta}(f_h - 2f_m + f_l)]}{D}, \end{aligned}$$

where $D = \sum_j f_j\lambda_j(\hat{\beta};\theta) + \sum_j f_j(\partial\lambda_j(\hat{\beta};\theta)/\partial\hat{\beta})(\hat{\beta} - \Gamma_j)$. It is easy to show that local stability requires that

$$(16) \quad \frac{\sum_j f_j(\partial\lambda_j(\hat{\beta};\theta)/\partial\hat{\beta})(\Gamma_j - \hat{\beta})}{\sum_j f_j\lambda_j(\hat{\beta};\theta)} < 1$$

implying that D is positive.

Note that one way to think about what an increase in sorting does is that for every two middle-type marriages it destroys, it creates one high- and one low-type marriage. With this in mind, note that an interpretation of (15) is that increased sorting increases the steady-state fraction of the population that attends college if the result of substituting two middle types by one high and one low type on net increases the number of children who attend college by more than what would result from that same

14. In a model with a continuous income distribution, there would always be a change in the bindingness of borrowing constraints for some families (as long as some of them were constrained in the initial equilibrium). Thus, this second channel would always be present. In our discrete model, whether constraints become more binding depends on the cost of college relative to family income. Thus, small changes in sorting may not affect the extent to which families are constrained.

substitution and all three types sending a fraction $\hat{\beta}$ of their children to college.

It is easy to evaluate (14) or (15) in a few special instances. Consider first the case where the Γ_j 's are constant across family types; i.e., $\Gamma_j = \Gamma, \forall j$. Then, by (13), $\hat{\beta} = \Gamma$, and a change in the degree of sorting has no effect on the fraction of the population that attends college and hence on wages and on the personal income distribution (e.g., if all family types send 10 percent of their children to college, the steady-state fraction of the population that attends college is 10 percent, irrespective of the degree of sorting).¹⁵

Next consider the case where the Γ_j 's are not identical but where average fertility is constant across family types; i.e., $f_j = f$. In such case the sign of (14) is given by the sign of $\Gamma_h + \Gamma_l - 2\Gamma_m$. The intuition behind this is simple given the earlier observation: since average fertility is the same across family types, the effect of increased sorting depends on whether the fraction of children sent to college on average by two middle-type marriages ($2\Gamma_m$) is smaller than the combined fraction of children that go to college on average in one high- and one low-type family ($\Gamma_h + \Gamma_l$). Thus, if the relationship between parents' education and children's education is linear, changes in sorting will have no effect on $\hat{\beta}$; if it is concave, increased sorting will decrease $\hat{\beta}$, and the reverse if the relationship is convex. More generally, what matters is whether the production of skilled children as a function of two variables (mother's and father's years of education) has a positive or negative cross partial in these variables.

Another case for which it is relatively easy to derive an expression is if $\Gamma_h + \Gamma_l - 2\Gamma_m = 0$ and $f_h + f_l - 2f_m = 0$. In this case, after manipulating (15), it is easy to see that the sign of the effect of an increase in θ is given by the sign of $(\Gamma_h - \Gamma_l)(f_h - f_l)$. This is an interesting case since it implies that if both fertility and the probability of attending college are linear in parents' average years of education, the effect of increased sorting is to decrease the fraction of the population that attends college if children of high-type parents have a greater probability of attending college and if the fertility of low-type parents is greater than that of high types. This points out that although every individual relationship can be linear, what matters is a nonlinear combination of the two relationships, which can give rise to nonlinearities.

15. The household income distribution, though, will be affected.

Last, it is useful to note from (15) that a sufficient condition for increased sorting to impact negatively on $\hat{\beta}$ is for $f_h\Gamma_h - 2f_m\Gamma_m + f_l\Gamma_l \leq 0$ and $f_h + f_l - 2f_m \geq 0$ (with at least one inequality strict). The first expression captures whether the number of children that on average attend college is increased or decreased by substituting two m couples by an h and an l . Thus, it indicates by how much the population that attends college would increase given this substitution. The second expression captures the amount by which the population as a whole is increased or decreased by substituting two m couples by an h and an l . Obviously, a decrease in the population attending college will, *ceteris paribus*, serve to reduce $\hat{\beta}$, as will an increase in the overall population (since it dilutes further the gain/loss of the first term). As we shall see farther on, our parameterization implies that both inequalities hold strictly and hence that increases in sorting decrease the fraction of the population that goes to college.

As mentioned previously, the degree of sorting can also affect the steady-state level of $\hat{\beta}$ via its effect on the tightness of borrowing constraints. To see this, suppose that keeping all z_j 's constant as before, an increase in θ decreases $\hat{\beta}$. This smaller proportion of skilled workers is associated with lower unskilled wages and higher skilled wages. The change in wages will increase family income for high types and decrease it for low types, and thus may lead to less binding constraints for the first group and tighter ones for the second. If this should happen, the original equilibrium values of the z_j 's and hence of the Γ_j 's would no longer be feasible and $\hat{\beta}$ would change as a result. That is, a change in the degree of sorting can affect the feasibility (in steady-state equilibrium) of different values of z_j 's.

IV. PARAMETERIZING THE MODEL

In this section we parameterize our model. We choose parameters so that the cross-section data generated in a steady state of the model are consistent with cross-section relationships in actual U. S. data. This ensures that the reduced-form relationships implied by this steady state of the model are "reasonable."

Many of our relationships will be based on a sample of individuals from the PSID. We construct a sample that matches parents and children by selecting all individuals over 25 in the 1993 PSID whose parents were in the PSID in 1968, and for

whom we have data on educational attainment. The resulting sample size is 645 parental units and 1385 children. We split this sample into skill categories by counting all individuals with high school or below as unskilled, and all individuals with some college or above as skilled.

Recall that in the model there are three types of marriages—high, middle, and low—which differ in both the average education and the average income of the couple. Each type of marriage j is further characterized by two statistics: ϕ_{nj} , the fraction of families of that type that have n children, and γ_j , the fraction of children (on average) from that marriage type that have the aptitude to benefit from skill acquisition. These two profiles are central to our analysis, so much of our discussion will focus on them.

It is empirically well established that fertility rates are negatively correlated with both income and education. To calibrate our fertility profiles, we use the fertility rates for the parent portion of our PSID sample. We use the parents since the children do not have completed fertility profiles. For our sample the average number of children from high-, middle-, and low-type families is equal to 1.84, 1.90, and 2.26, respectively. In each case we assume that families of a given type have realized fertility corresponding to the two integers that bracket the actual fertility rate, with the probability of obtaining a given integer specified so as to match the average fertility rate for that family type. This gives $\phi_{1h} = .16$, $\phi_{2h} = .84$, $\phi_{1m} = .10$, $\phi_{2m} = .90$, $\phi_{2l} = .74$, $\phi_{3l} = .26$, and all other ϕ_{nj} 's are equal to 0.¹⁶

There is also information on the relationship between women's education and fertility.¹⁷ Although these data do not provide information about fertility and *family* education, and hence are not directly relevant to our calibration, it is still instructive to observe the magnitudes of fertility differences they indicate.¹⁸ For the period from 1960 on, Mare [1997] finds that the gross reproductive rate for white women with high school education or less varies between 1.18 and 1.35 times larger than the gross

16. We could choose the ϕ_{nj} 's to match the distributions of fertility within each marriage type, but since this does not affect the results, we have chosen the simpler procedure.

17. Mare [1997], for example, tracks this relationship over the period 1920–1990 for both blacks and whites.

18. Average fertility rates have varied considerably over this time period. In our model, proportional shifts in fertility are not important, so we focus on relative fertility rates.

TABLE I
APTITUDE PROFILES UNDER VARIOUS SCENARIOS $z_{nm} = n, z_{nb} = n$

	$z_{nl} = n$	$z_{nl} = 2$	$z_{2l} = 2, z_{3l} = 1$	$z_{nl} = 1$
γ_h	.81	.81	.81	.81
γ_m	.63	.63	.63	.63
γ_l	.30	.303	.334	.401

reproductive rate for white women who have at least some college. For black women the corresponding range is between 1.41 and 1.61. Our data from the PSID suggest differences between high- and low-type marriages that are roughly in accordance with the range that Mare found for white women. We will investigate the importance of different fertility profiles in our sensitivity analysis.¹⁹

Next we consider the choice of the γ_j 's. We have no direct measure of the fraction of children from marriages of different types that have an aptitude for skill acquisition. From our PSID sample, however, we have data on the relationship between the educational attainments of parents and their children. We find that the fraction of children from high-type families that acquire education beyond high school is .81, whereas the values for middle- and low-type families are .63 and .30, respectively. In the steady state of our model, we require that the Γ_j 's match these values.

As is evident from equation (12), Γ_j is a function of γ_j and z_j . Thus, the probability that a child from a particular marriage type is of high aptitude (i.e., the γ_j) can be deduced from the value of Γ_j in conjunction with an assumption about the maximum number of children that the distribution of families of each family type can afford to send to college (i.e., the z_{nj}) in the steady state. Table I illustrates this mapping by showing the values of the γ_j 's implied by various assumptions regarding the tightness of borrowing constraints subject to the requirement that each (γ_j, z_j) pair yield the aggregate Γ_j found in the data.

19. One factor in favor of considering relatively larger fertility differences is that our model does not allow for the fact that lower income families have their first child some five years before richer families. This would increase the relative size of the poorer group in steady state by more than what would be predicted based solely on differences in the number of children. See Knowles [1999] for details.

The first column of Table I corresponds to a case in which no one is constrained—all high-aptitude children become skilled. In this case the values of Γ_j 's and γ_j 's must coincide. In the remaining columns we consider various scenarios under which low-type families face some credit constraints but the middle- and high-type families do not. The second column assumes that all low-type families can afford to send at most two children to college. The third column assumes that low-type families with three children suffer different constraints than do low-type families with two children. Whereas the latter can afford to send two children to college, the former can afford to send only one child to college. Last, the fourth column assumes that all low-type families can afford to send only one child to college.

Our analysis, for the most part, is independent of which of these scenarios we take to represent the steady state. If the maximum number of children that different family types can afford to send to college remains unchanged when the degree of sorting increases, then as equation (15) indicates, the effect of sorting depends only on the Γ_j 's which are given by the data; the mix of γ_j 's and z_{nj} 's used to generate them is irrelevant. It is only when we allow the steady-state equilibrium value of the z_{nj} 's to be affected by the increased degree of sorting that the exact specification might matter. But even in that case all that matters to our results, as will be seen in the next section, is the change in the set of people affected by a tightening of the borrowing constraints, not the number who are initially affected.

We choose the second column of the table for our benchmark specification; i.e., we assume that borrowing constraints do not affect middle and high marriage types, but that low types with three children are able to send at most two of their children to college. This is actually a very mild constraint: since only families who have three children all of whom are high aptitude are constrained, less than 1 percent of low-type families are affected. Of course, as outlined in the previous section, it is necessary to check that our assumptions on credit constraints are consistent with wages, consumption requirements, and the cost of skill acquisition. We leave this for later in the analysis.

Next we assign a value to θ , the fraction of marriages that are matched perfectly as opposed to at random. Note that this matching procedure implies that θ is equivalent mathematically to the correlation between the education levels of spouses. We use our

sample from the PSID to obtain an estimate of this correlation for the United States, yielding $\theta = .6$.

Given the values assigned thus far, we can solve for β (the fraction of population that goes to college in the steady state). Doing so, we obtain $\beta = .60$. This turns out to be slightly higher than the corresponding number found in current data, as, according to the 1996 CPS, roughly 55 percent of individuals aged 25–34 have at least some college. Since average educational attainment in the United States continues to increase and our calculation is for the steady state, there is no real inconsistency here, especially as our model abstracts from immigration and childbirth outside of two-parent households.

It remains to specify the production function, for which we choose a constant elasticity of substitution production function:

$$y = A[bN_s^p + (1 - b)N_u^p]^{1/\rho}.$$

Note that the ratio of skilled to unskilled workers can be written as $\beta/(1 - \beta)$, and that the relative wage of skilled to unskilled workers is given by $w_s/w_u = b/(1 - b)(\beta/(1 - \beta))^{\rho-1}$. As is well-known, the ratio of skilled to unskilled wages has varied considerably over the last 30 years in the United States.²⁰ Recall that our two skill groups are those with at least some college and those with high school or less. Based on the data in Katz and Murphy [1992], we match a ratio of 1.9 for our benchmark case. This value is at the upper end of what has been observed in the United States, so in our robustness check we redo our analysis assuming a ratio of 1.4 and find that it has no impact on our results.

There is a literature that attempts to estimate the degree of substitutability between skilled and unskilled labor that we can use to obtain an estimate for ρ . The survey by Katz and Autor [1999] suggests that a reasonable range of values for the elasticity is between 1 and 2.5. We match an elasticity of substitution of 1.5 for our benchmark case, which implies that $\rho = .33$. As we will see shortly, the elasticity of substitution is a key parameter for our analysis—if we use a value that is substantially larger, the model generates smaller effects from changes in sorting. We explore the range of values suggested by Katz and Autor in our sensitivity analysis. Our chosen value of ρ and the above-men-

20. See, e.g., Katz and Murphy [1992].

tioned value for the skill premium implies that $b = .7135$. Last, for ease of interpretation of our results, we choose a value of A to scale steady-state unskilled wages to some “reasonable” value, which we set to be 30,000. This is purely an issue of normalization.²¹

Having assigned parameter values, we can solve for the steady state in which educational attainment is dictated by the observed values of the Γ_j 's as discussed previously. We now report some additional properties of this steady state. The model produces a distribution for individual income, with mass at two points, corresponding to the skilled and unskilled wage rates. The standard deviation of log income in the steady state equals .315. Distributions of annual income in the United States typically imply a value of around .6 for this figure. Alternatively, the lifetime income distribution generated by Fullerton and Rogers [1993] using PSID data yields a value around .4. Since we are relying entirely on the skill premium to generate our variation in income, it is not surprising that we produce less variation than is found in the data.

Last, we can also compute the standard deviation and mean of the educational attainment distribution. We assume that a high school education corresponds to $e = 11.3$ and a college education corresponds to $e = 15.0$, our choice of numbers given by the average educational attainments of children with high school or less and those with college or more in our PSID sample. The resulting standard deviation and mean of the steady-state educational attainment distribution are equal to 1.81 and 13.5. For our sample of children from the PSID, the corresponding values are 2.56 and 12.9. Given our restriction to two levels of education, it is not surprising that we generate less variation than the data. The fact that our mean is somewhat higher is related to the fact that it is the steady-state value.

V. THE EFFECTS OF INCREASED SORTING

We now use the parameterized model to assess the effects associated with an exogenous increase in the degree of sorting in

21. It is not clear what the “best” normalization is, since in our model these are lifetime earnings. Rescaling of this variable, of course, implies that the parameters \bar{c} and v need to be scaled accordingly as well.

TABLE II
EFFECTS OF INCREASED SORTING ON STEADY STATE

	$\theta = .6$	$\theta = .7$		$\theta = .8$	
	$\Gamma_l = .30$	$\Gamma_l = .30$	$\Gamma_l = .27$	$\Gamma_l = .30$	$\Gamma_l = .27$
mean(e)	13.52	13.48	13.40	13.44	13.36
std(e)	1.81	1.82	1.83	1.83	1.84
cv(e)	.134	.135	.137	.136	.138
β	.600	.589	.568	.579	.556
N_s/N_u	1.50	1.43	1.31	1.38	1.25
w_s/w_u	1.900	1.95	2.07	2.01	2.14
w_u	30,000	29,375	28,125	28,774	27,453
std(log y)	.315	.330	.361	.345	.379

the marriage market. Our objective is to examine whether the concern that some writers have expressed—namely, that increased sorting will lead to increased inequality—has any significant quantitative support. To this end, we analyze the consequences of increases in sorting corresponding to increases in θ to .7 and .8.

Table II displays the results. The first column gives the values for the original steady state (i.e., $\theta = .6$). The second column shows what the steady state would be if θ were to increase to .7 and the tightness of borrowing constraints were unchanged (i.e., all high-aptitude children from middle- and high-type families could afford to attend college but among low-type families at most two children per household could be sent to college). The third column reports the new steady-state values ensuing from the θ change, but assumes that the wage change associated with this increase tightens constraints for low-type families with three children to the point that they can afford to send at most one of their children to college.²² The fourth and fifth columns are analogous to the second and third columns except that they correspond to a θ increase to .8. Below we discuss each case in turn, first examining those cases in which the borrowing constraints are assumed to be unaffected by the change in sorting (i.e., columns 2 and 4).

22. Later in this section we show that this outcome is consistent with choices for \bar{c} and v .

A. Case 1: Set of Borrowing Constrained Individuals Remains Constant

We begin by comparing the first column with the second column. This amounts to examining the effects of an increase in sorting from .6 and .7 holding the pattern of college attendance fixed. The first three rows report the mean, standard deviation, and coefficient of variation for the steady-state distribution of educational attainment. The effect of the increase in sorting is to cause a small decrease in both the mean and the standard deviation of the distribution (less than 1 percent), and a small increase in the coefficient of variation (slightly more than 1 percent).²³

The decrease in mean educational attainment results from the decrease in the fraction of the population that goes to college; the fourth row shows that the fraction of the population that becomes skilled falls to 59 percent in the new steady state. Although this is a seemingly small decrease, it implies a fall in the ratio of skilled to unskilled workers of about 4 percent, as shown in the fifth row. As the next row indicates, this change in relative labor supply induces an increase in the skill premium of roughly 2.5 percent.²⁴ The next to last row shows that the standard deviation of the log income distribution increases by almost 5 percent. This is a sizable increase. By way of comparison, the much publicized increase in wage inequality that took place in the 1980s resulted in roughly a 10 percent increase in the standard deviation of male log weekly wages.

Thus far, we have reported properties of the individual income distribution. It is also of interest to examine the changes in the distribution of family income. This distribution changes both because household partners are more correlated and because of the general equilibrium effect this has on wages and on the fraction of the population that becomes skilled. Thus, to decompose the two effects, note that keeping wages and β constant at their original steady state values, the increase in θ from .6 to .7

23. Note that since our educational attainment distribution is over two levels, its variance is maximized when the population is evenly distributed across them. Hence, whether a change in β results in an increase or decrease in the standard deviation of education depends entirely upon whether the starting value was above or below .5. Having said this, we think that what is most relevant to notice is that the change in the standard deviation is very small, rather than the direction in which it changes.

24. Recall that the elasticity of substitution in the production function is 1.5, which implies that the percent change in N_s/N_u will be 1.5 times as large as the percent change in w_s/w_u for small changes. For large changes this expression continues to hold exactly in logs, but only approximately in ratios.

itself would increase the standard deviation of log family income by 3.2 percent. Allowing the skill premium and β to adjust, the increased correlation of household partners yields an increase in the standard deviation of log family income of 8.3 percent. We conclude that both components are important.

The results for the case in which θ increases to .8 follow a pattern similar to the one above. The change in the mean and standard deviation of the education distribution are small, as before. The increase in the skill premium is now greater—approximately 5 percent—and the increase in the standard deviation of log income is now almost 10 percent relative to the $\theta = .6$ case. This increase is roughly double that found in the case of $\theta = .7$, suggesting that the effect of changes in θ on inequality is close to linear. The increase in the standard deviation of log family income is 16.5 percent, again roughly double the increase found in the previous case.

B. Case 2: Set of Borrowing Constrained Individuals Changes

The preceding analysis assumed that borrowing constraints did not change as a result of the increase in the degree of sorting. As shown in row 7, however, associated with the increase in θ to .7 is a decrease of a bit over \$600 in the wages of unskilled workers and hence a decrease in low-type family income of more than \$1200. This wage decrease makes it possible that some low-type families will be able to send fewer children to college than previously and hence that the steady-state equilibrium values of z_l used in column 2 are no longer feasible.

In the third column we assume that as a result of the θ increase, in the new steady state low-type families who have three children can afford to send a maximum of one child to college, rather than two (i.e., $z_{3l} = 1$); low-type families with two children are assumed to remain unconstrained. This constraint implies that the fraction of children from low-type marriages that go to college drops from .30 to .2745 (as indicated by the reported values of Γ_l). The reason that this drop is relatively small is that only 24 percent of low-type families have three children and of these, only some 22 percent have at least two children of high aptitude.

As the table shows, the tightening of borrowing constraints has a sizable effect on how the θ increase affects the income distribution. In particular, although the change in the mean level of education is still relatively small (roughly 1 percent) and the

change in the steady-state equilibrium seemingly not large (β now equals .568), this implies a drop of almost 15 percent in the ratio of skilled to unskilled workers, relative to the $\theta = .6$ case. The skill premium (w_s/w_u) also increases by more than 9 percent and the standard deviation of the distribution of log income increases by almost 15 percent. The standard deviation of log family income increases by almost 19 percent.

The final column of Table II shows results for the case in which θ increases from .6 to .8 and the increase in the degree to which credit constraints bind is again assumed to reduce Γ_l from .30 to .27.²⁵ The resulting change in the standard deviation of log income is more than 20 percent. Once again, we note that the change in the mean and standard deviation of the education distribution are still small. For example, the change in mean education is roughly 1 percent relative to the $\theta = .6$ case.

We next verify that the structural change in college attendance decisions is a feasible equilibrium outcome. We do so only for the case of the change in θ to .7; the change to .8 is similar. In what follows, let $w_u(\beta_i)$ be the unskilled wage rate when the equilibrium ratio of skilled to unskilled workers, is given in column i , $i = 1, \dots, 5$, of Table II.

For column 1 to represent an equilibrium steady state, it must be that type l families with three kids can send two but not three children to college. This requires (i) $2v + 5\bar{c} < 2w_u(\beta_1) < 3v + 5\bar{c}$.²⁶ For the allocations in column 2 to be infeasible because at those wages l -type families with three kids cannot afford to send two of them to college, requires (ii) $2v + 5\bar{c} > 2w_u(\beta_2)$. Last, to ensure that the outcome in column 3 is an equilibrium requires checking that it allows type l families with three kids to send one child and type l families with two kids to send two to college. That is, it requires (iii) $v + 5\bar{c} < 2w_u(\beta_3)$ and (iv) $2v + 4\bar{c} < 2w_u(\beta_3)$. Note that if type l families with two kids can afford to send two kids to college then so can type m and h families. Inequalities (ii) and (iii) imply that $v > 2(w_u(\beta_2) - w_u(\beta_3))$, so that $v > 1250$ given the numbers in Table II. There are many combinations of v and \bar{c} that satisfy these inequalities. For example, $v = 10,000$, and $\bar{c} = 7800$.

We do not attach too much significance to the magnitudes of

25. Although we have kept the decrease in Γ_l constant, it may be of interest to consider even larger decreases since the drop in unskilled wages is now larger.

26. Of course, if low-type families can afford to send two children to college, so can higher-type families.

TABLE III
 TRANSITION FROM $\theta = .6$ STEADY STATE TO $\theta = .7$ STEADY STATE

t	β	std(log y)
0	.5996	.3145
1	.5949	.3218
2	.5922	.3256
3	.5908	.3276
4	.5901	.3287
5	.5897	.3292
6	.5895	.3295
10	.5892	.3298

v and \bar{c} . The simple choices that we made about utility functions and the fact that we abstract from life-cycle income dynamics and the timing of college attendance make us reluctant to do so as does our unwillingness to interpret the borrowing constraints literally as the ability to afford college. The main point of the above paragraph is to establish the logical consistency of our argument that the change in sorting can lead to a change in the extent to which credit constraints bind. It is perhaps not surprising that this can be done, given that we have not imposed any discipline on our choices of \bar{c} and v .²⁷

Our analysis thus far has focused on steady states. Given that our model predicts large changes in income inequality, it may also be of interest to ask how long it may take for these changes to occur. To pursue this, we have solved for the transition path between steady states. Our basic finding is that the value of β moves roughly halfway to its new steady state value each period. In view of this we conclude that the changes in income inequality that we are finding are large not only in the steady state but also at small horizons as well. We illustrate this in Table III, which shows the time series for β and the standard

27. Having said this, however, we do offer one check of "reasonableness" for the value of v . Specifically, we can compute the annual rate of return to spending on education. This obviously depends on how many years one assumes there are between the expenditure and the return in the form of higher wages, since the expenditure takes place in the first period of life but yields higher wages in the second period. If one interprets a period to be a generation, then twenty years may be reasonable. On the other hand, if one wants to look at the time between college expenditures and the midpoint of a typical working life, then a slightly smaller period length may be appropriate. In any case, for the steady state in column 1, the annual rate of return lies between 5 and 10 percent as we vary the number of years between ten and twenty.

deviation of log income as the economy moves from the $\theta = .6$ steady state to the $\theta = .7$ steady state, assuming no change in the degree to which credit constraints bind. In the table we denote period 0 to be the initial period of the change in the degree of sorting. Since β is a state variable, it does not respond until the following period.

VI. DISCUSSION

Our results support a conclusion very different from that reached by Kremer [1997]. Whereas he concluded that a large increase in sorting would have little effect on steady-state inequality given a reasonable parameterization, we have concluded that even moderate changes in the degree of sorting can have sizable effects on inequality. In this section we analyze what lies behind this difference. Having identified the factors that generate such different conclusions, we then examine the robustness of our results to different specifications. We first turn to a brief review of Kremer's analysis.

Kremer posits an intergenerational model of marriage, fertility, and educational attainment in which a child's educational attainment e can be written as a linear function of parental and neighborhood average education. For expositional purposes we consider the argument in the simplest context, and hence abstract from neighborhood effects.

The model assumes that all individuals marry and have two kids. A child's educational attainment is determined by the following linear relationship:

$$e_{i,t+1} = \kappa + \alpha[(e_{i,t} + e_{i',t})/2] + \varepsilon_i,$$

where $e_{i,t+1}$ is the educational level for the child, $e_{i,t}$ and $e_{i',t}$ are the education levels of the two parents, and ε is a normally distributed random shock that is *iid* across families, with mean 0 and standard deviation equal to σ_ε . An exogenous (assortative) matching of individuals takes place such that ρ_m is the correlation between the education levels of parents. Assuming that parameter values are constant over time, the distribution of education will converge to a normal distribution with mean and standard deviation given by $\mu_\infty = \kappa/(1 - \alpha)$ and $\sigma_\infty = \sigma_\varepsilon/([1 - \alpha^2(1 + \rho_m)/2]^{.5})$, respectively.

Kremer's objective was to determine how changes in sorting

among marriage partners (i.e., ρ_m) would affect the level of inequality in the steady state. His main measure of inequality was the standard deviation of educational attainment, and he argued that since there is a linear relationship between educational attainment and log of income in the cross section, this measure of inequality would probably be a good proxy for inequality in log of income as well. We shall return to this point later.

The effect of an increase in ρ_m on the steady-state distribution of education can be read off the limiting distribution equations. Because of the assumption of linearity, there is no effect of ρ_m on the mean of the distribution of education, but its standard deviation is increasing in ρ_m . Obviously, this model is at least qualitatively consistent with the view that increased sorting leads to increased inequality.

Kremer's main contribution, however, was to show that while the model supported this view qualitatively, there was little support for the view that this effect was important quantitatively. As the equation showing the limiting standard deviation makes clear, the percentage change in the standard deviation of income due to a change in the sorting parameter ρ_m is determined solely by the magnitude of the parameter α . Using data from the PSID (the same source that we used to parameterize our model), he obtained an estimate of α of about .4 and $\rho_m = .6$.²⁸ In this case, an increase in ρ_m from .6 to .8 would result in only a 1 percent increase in the standard deviation of education.

These results, as Kremer showed, are fairly insensitive to the exact value of α used in this vicinity. The easiest way to see this is by asking how large α would need to be in order that an increase in ρ_m from .6 to .8 results in a 10 percent increase in the standard deviation of log income. It is easy to show that this requires $\alpha = .852$.

Kremer's paper is mainly about the effect of increased neighborhood and marital sorting on the distribution of education. However, if one takes the view (as Kremer does in his introduction) that log earnings are approximately linear in years of education, and that the coefficients in this relationship are invariant to changes in the distribution of education, then the same conclusion applies to inequality in income; a large increase in sorting

28. When neighborhood effects were included, the sum of coefficients on parental and neighborhood education was about .55. This does not change his conclusions.

will not significantly affect income inequality in the United States; an increase in marital sorting from .6 to .8 will increase the standard deviation of log income by only 1 percent.

It should be clear that the results we report in Table II do not contradict Kremer's finding that sorting has a small impact on the level of inequality in the *skill* distribution, especially if the increase in sorting does not affect the bindingness of borrowing constraints. To further demonstrate that there is no inconsistency between our results and his, we perform his analysis on data generated from our model. Specifically, using data generated by the steady state of our calibrated model (i.e., column 1 in Table II), we take a random sample of 1200 families and run a regression of a child's educational attainment ($e_{i,t+1}$) on a constant and the average educational attainment of its parents ($\bar{e}_{i,t}$).²⁹ As noted previously, we assume that a high school education corresponds to $e = 11.3$ and a college education corresponds to $e = 15.0$.³⁰ We do this 100 times and average across the trials. The result of this exercise is³¹

$$e_{i,t+1} = 6.69 + .51\bar{e}_{i,t}.$$

It follows that if Kremer had performed his exercise using data generated from our model he would still have reached the same conclusion; i.e., he would have concluded that the coefficient on average parental education is too small to matter.

What gives rise to our very different conclusion about income inequality is the interaction between changes in the skill distribution and the price of skill in our general equilibrium model. This interaction is governed by three elements that are absent in Kremer's analysis but that are central to generating this effect on the price of skill: (i) the existence of a nonlinear relationship between parental years of education and those of their children; (ii) a negative correlation between fertility and parental education; and (iii) wage rates that are sensitive to changes in the skill

29. Note that since our model is not linear, it does not lead one to run this regression. We run this regression simply to illustrate how Kremer's analysis would look in our setup.

30. The specific values chosen here affect the constant term in the regression but have very little effect on the coefficient on parent's education.

31. Running this regression on our sample from the PSID yields a coefficient of .37. This discrepancy is accounted for by the fact that in our model we compress the education distribution to two levels, thereby increasing the correlation between the education levels of parents and their children. We have verified this via simulation.

distribution. As we shall see, it turns out that if we had only incorporated any one of these three elements, we would have reached the same conclusion as Kremer.³² But, allowing for the interaction of all three factors (especially (i) and (iii)) leads to a very different conclusion.

We begin with a discussion of the third factor. The distribution of labor earnings can be thought of as depending on the interaction of two factors. One is the distribution of skill (in our model, education) across individuals, and the second is the price of skill (i.e., the skill premium). As stated in our discussion of Table II, the impact of sorting on the level of inequality in the skill distribution is small. In fact, if wages were not responsive to the distribution of skills, the change in the standard deviation of log income would have been around one-half of 1 percent. What drives our results is that a change in sorting produces a significant change in the skill premium, even if it seemingly does not produce "large" effects on mean educational attainment. As can be seen from a comparison of columns 1 and 2 in Table II, a less than 1 percent decrease in the mean of the education distribution is associated with an almost 4 percent decrease in the relative supply of skilled labor. This translates to a 2.5 percent increase in the wage premium, leading to a significant change in the distribution of income.

To better understand how various elements interact to yield the increase in the skill premium, note first that in our model the impact of a change in θ on w_s/w_u can be decomposed into two distinct effects. The first concerns how a given change in θ affects β , and the second with how a given change in β affects w_s/w_u . This decomposition is useful because college attendance and fertility profiles are only relevant for the first effect, whereas the elasticity of substitution in the production function is only relevant for the second.³³

32. In fact, Kremer considers a Markov model in Section IV of his paper, and finds little effect of sorting on the standard deviation of education. This is obviously consistent with our findings. We have also rewritten Kremer's model to account for differential fertility and used numerical techniques to compute the steady-state distribution of education. Once again, changes in sorting have little effect on the standard deviation of this distribution.

33. In particular, the change in $\log w_s/w_u$ equals the inverse of this elasticity times the change in $\log \beta/(1 - \beta)$. Moreover, how this change in w_s/w_u is split between changes in each of the two wages is entirely determined by the elasticity and the initial value of β . If we had assumed that $\rho = 1$, i.e., a linear production function, there would be no effect of sorting on wages and, as discussed previously, we would have found very small effects from increased sorting on inequality.

Consider now the roles of fertility differences and of the function relating parental education to children's education in generating the change in $\hat{\beta}$. Recall from the discussion in Section III that a sufficient condition for increased sorting to impact negatively on $\hat{\beta}$ is for $f_h\Gamma_h - 2f_m\Gamma_m + f_l\Gamma_l \leq 0$ and $f_h - 2f_m + f_l \geq 0$ (with at least one strict inequality). Our parameter values strictly satisfy both inequalities, guaranteeing that increased sorting will decrease the fraction of the population that attends college. The magnitude of the respective contributions of our fertility profile and the concavity of the intergenerational education transmission function will be discussed in the next section on robustness.

One can ask under what conditions our model would give rise to the conclusion that changes in sorting do not have significant effects on the income distribution (without shutting down the effect of changes in the skill distribution on wages). A simple condition is given by the combination of a linear relationship between parents' and children's education (i.e., $2\Gamma_m = \Gamma_h + \Gamma_l$) and no fertility differentials (i.e., $f_j = f$ for all j). But these are precisely the assumptions made by Kremer in his paper—all parents have two kids and the child's years of education are linear in average parental years of education. Thus, if we had adopted Kremer's assumptions, our model would not have generated any effect from increased sorting on the steady-state value of β , and hence no effect on wage rates or inequality either. Moreover, the fact that wage rates would not have changed would necessarily imply that the bindingness of borrowing constraints would be unaffected and consequently there would be no scope for any change in college attendance decisions via this channel either.

Last, our analysis also suggests that one should exercise caution in interpreting regressions of child's educational attainment on parental educational attainment. In our discussion of Kremer's work, this regression coefficient was denoted α and was treated as a structural parameter that would not be affected by changes in sorting. However, as should be clear from our model, the degree to which education is heritable may differ across family types for a variety of reasons including the presence of borrowing constraints. The degree of sorting, as evidenced in columns 3 and 5 of Table II, affects the bindingness of borrowing

constraints and hence the degree to which parents' education is passed on to their children.³⁴

VII. ROBUSTNESS

In this section we report how the results from our benchmark model are affected by changes in our parameterization. Our finding of a quantitatively important increase in income inequality arising from changes in the degree of sorting is robust to reasonable variation in the model's parameterization.

We begin by considering how alternative profiles for fertility affect our results. Specifically, we consider two alternative fertility profiles. In the first, we assume that all families have exactly two children. In the second, we explore the consequences of having greater fertility differences between low-type marriages and the other marriages, and hence consider a 25 percent increase in f_l to 2.83, keeping $f_h = 1.84$ and $f_m = 1.90$. In each case we recalibrate our model to match the same statistics as before. The comparative statics exercises are the same as in columns 2 and 3 in Table III. That is, we examine the effect of an increase in θ from .6 to .7 first assuming that the bindingness of constraints is unchanged, and subsequently assuming that they are tightened so as to cause a decrease in Γ_l from .30 to .27.

The results are reported in Table IV. The basic message is the same for both of the alternative fertility profiles. Even with no fertility differences the increase in income inequality is still substantial, albeit somewhat less than in Table II (3.5 percent versus 4.8 percent, with no change in credit constraints and 11.5 percent versus 14.9 percent if Γ_l decreases to .27). For the case in which low-type families have 2.83 kids, the increase is 5.3 percent assuming no change in constraints and 16.9 percent if the constraints are tightened.³⁵

Next we examine how our findings are affected by changes in the profile of Γ_j 's used in the calibration. Table V shows the effect

34. Running a linear regression for the steady state in column 3 of Table II (using the same procedure described earlier), we obtain .53 rather than the .51 obtained for the scenarios in columns 1 and 2, although the true "heritability" of education is unchanged as reflected in the γ_j 's.

35. We also considered the case in which the fertility profile is linear in average parental education, a result that would obtain if, for example, fertility were solely determined by the mother's education. To check this, we increased f_m to 2.05. We found this to have minimal effect relative to our benchmark case and so do not report the results.

TABLE IV
EFFECT OF ALTERNATIVE FERTILITY PROFILES
 $\Gamma_h = .81, \Gamma_m = .63$

	$f_h = 2, f_m = 2, f_l = 2$			$f_h = 1.84, f_m = 1.9, f_l = 2.83$		
	$\theta = .6$	$\theta = .7$	$\theta = .7$	$\theta = .6$	$\theta = .7$	$\theta = .7$
	$\Gamma_l = .30$	$\Gamma_l = .30$	$\Gamma_l = .27$	$\Gamma_l = .30$	$\Gamma_l = .30$	$\Gamma_l = .27$
mean(e)	13.7	13.64	13.59	13.33	13.28	13.18
std(e)	1.78	1.78	1.80	1.84	1.85	1.85
cv(e)	.130	.131	.132	.138	.139	.140
$\hat{\beta}$.640	.634	.619	.547	.535	.508
N_s/N_u	1.78	1.73	1.62	1.21	1.15	1.03
w_s/w_u	1.90	1.94	2.03	1.90	1.96	2.11
w_u	30,000	29,543	28,539	30,000	29,312	27,873
std(log y)	.308	.319	.343	.320	.337	.374

of varying Γ_m from its value of .63 in our benchmark model. First, we examine the consequences of decreasing the degree of concavity in the relationship between parental and children's education to the point where it is linear ($\Gamma_m = .555$), and then we explore the consequences of increasing the degree of concavity. In each case the production function parameters are recalibrated to match the same statistics as before.

Qualitatively the results are not surprising. As we move closer to the linear case, the increase in the standard deviation of log income caused by an increase in sorting becomes smaller. What is most interesting is the quantitative impact of the in-

TABLE V
EFFECT OF ALTERNATIVE Γ_m 's
 $f_h = 1.84, f_m = 1.90, f_l = 2.26, \Gamma_h = .81$

	$\Gamma_m = .58$			$\Gamma_m = .555$			$\Gamma_m = .68$		
	$\theta = .6$	$\theta = .7$	$\theta = .7$	$\theta = .6$	$\theta = .7$	$\theta = .7$	$\theta = .6$	$\theta = .7$	$\theta = .7$
	$\Gamma_l = .3$	$\Gamma_l = .3$	$\Gamma_l = .27$	$\Gamma_l = .3$	$\Gamma_l = .3$	$\Gamma_l = .27$	$\Gamma_l = .3$	$\Gamma_l = .3$	$\Gamma_l = .27$
mean(e)	13.44	13.42	13.34	13.41	13.40	13.32	13.59	13.53	13.46
std(e)	1.83	1.83	1.84	1.83	1.83	1.84	1.80	1.81	1.82
cv(e)	.136	.136	.138	.137	.137	.138	.132	.134	.136
$\hat{\beta}$.580	.575	.553	.570	.567	.545	.619	.604	.583
N_s/N_u	1.38	1.35	1.23	1.33	1.31	1.19	1.62	1.52	1.39
w_s/w_u	1.90	1.93	2.05	1.90	1.91	2.04	1.90	1.98	2.10
w_u (000' s)	30.00	29.68	28.41	30.00	29.83	28.55	30.00	29.07	27.86
std(log y)	.317	.325	.357	.318	.322	.354	.312	.335	.365

TABLE VI
 CHILDREN'S EDUCATION AS A FUNCTION OF PARENT'S EDUCATION
 Dependent variable is years of education for the child.
 (Standard errors are in parentheses.)

	(1)	(2)	(3)
<i>constant</i>	8.71	12.30	16.51
	(.247)	(.608)	(1.285)
\bar{e}	.378	-.347	-1.798
	(0.021)	(0.114)	(0.407)
\bar{e}^2	—	0.034	0.184
		(0.005)	(0.041)
\bar{e}^3	—	—	-0.005
			(0.001)
N	1385	1385	1385
R^2	0.185	0.208	0.216

crease in sorting when the Γ profile is linear. In this case we still have an increase in the standard deviation of log income that exceeds 1 percent, assuming no change in credit constraints. This is roughly one-quarter of the impact obtained in our benchmark case. Assuming that constraints are tightened on low types (same exercise as in previous tables), then, as reported in the next column, the standard deviation of log income is again more than 10 percent higher. If Γ_m were increased to .68, the increase in the standard deviation of log income is higher by about one and a half percentage points assuming no change in the degree to which credit constraints bind. If this is not true and the effective value of Γ_l decreases to .27 when θ increases to .7, then the increase in inequality is 17 percent.

Given the importance in our analysis of a nonlinear relationship between children's and parents' schooling, we think it is of interest to document these beyond the Markov transition probabilities reported earlier. Table VI presents several regression results that incorporate higher-order terms in Kremer's original regression.³⁶ These regressions are based on our sample of parents and children from the PSID and are run on the reported

36. Although Kremer runs a regression that includes the square of parental average education, he also includes the square of average neighborhood education and an interactive term between parental and neighborhood effects. In that regression all variables are statistically insignificant, including average parental education.

years of education rather than the two categories we used for our model.³⁷

Column (1) in this table is the equivalent of column (5) in Table II in Kremer, with basically identical results. What columns (2) and (3) show, however, is that there is strong support for the notion that this relationship is nonlinear. In every specification, all terms are significant at the 1 percent level. Note that in the cubic specification the second derivative changes from positive to negative at $\bar{e} = 12.26$. Hence, up to this point there are "increasing returns" to parental education in terms of "producing" child's education, but beyond this point there are "decreasing returns" to parental education.³⁸ The fact that there are increasing returns in the lower part of the distribution suggests that increased sorting within this part of the distribution may actually increase mean educational attainment within this group. Our analysis abstracts from this issue since it is concerned with the degree of sorting between the top and bottom parts of the income distribution rather than the within-group sorting. There we find a concave relationship between children and average parental years of education.

We have also investigated this issue in another exercise. Consider the population segmented by four educational attainments: less than high school, high school, some college, and college and above. There are now six types of marriages. We investigate concavity through pairwise comparisons; e.g., we consider two couples, each of whom is perfectly matched, although the two marriages have different educational attainments. We then consider what happens to the average educational attainment of children if we turn these two marriages into two mixed marriages, holding the number of kids constant. The only violations of concavity that we find are for the cases involving the less than high school group with either the high school or some college group.³⁹ Hence, these findings tell the same story as the above

37. As in Kremer [1997], individuals with more than sixteen years of schooling are treated as having seventeen years of schooling.

38. We have also run regressions by splitting the sample into two groups: parents with average education less than or equal to twelve years, and parents with average education greater than or equal to twelve years. These results confirmed the above finding concerning the switch in returns to scale.

39. The same finding emerged if instead of years of educational attainment we considered the probability of obtaining at least some college or the probability of finishing college. These results are of interest since they are not affected by the fact that years of schooling are effectively bounded above.

TABLE VII
EFFECTS OF INCREASED SORTING FOR ALTERNATIVE VALUES OF ρ
 $f_h = 1.84, f_m = 1.9, f_l = 2.24, \Gamma_h = .81, \Gamma_m = .63$

	$\rho = .33$		$\rho = .6$		$\rho = 0$	
	$\Gamma_l = .30$	$\Gamma_l = .27$	$\Gamma_l = .30$	$\Gamma_l = .27$	$\Gamma_l = .30$	$\Gamma_l = .27$
$\% \Delta w_s / w_u$	2.63	9.12	1.72	5.35	4.35	13.91
$\% \Delta w_u$	-2.08	-7.46	-1.25	-3.74	-3.10	-9.19
$\% \Delta \text{std}(\log y)$	4.76	14.85	3.06	9.32	7.06	21.62

regression analysis—we find evidence for concavity except at the lower end of the distribution.⁴⁰

Our conclusion is also not sensitive to the choice of the value for the wage premium. Although the extent of income inequality in the steady state is affected by this ratio, using values for the wage premium anywhere in the range of 1.4 to 1.9 has virtually no impact on the extent to which the increase in sorting increases the steady-state standard deviation of log income.

Last, we consider how our results are affected by considering alternative values for the elasticity of substitution between skilled and unskilled workers. In our benchmark model we assumed a value for this elasticity equal to 1.5 (i.e., $\rho = .33$). Here we report how our conclusions are affected by assuming values of 1.0 ($\rho = 0$) and 2.5 ($\rho = .6$), since this is the range of estimates suggested by Katz and Autor [1999]. Table VII contains the results, with the first column repeating the findings from Table II in order to facilitate comparisons. As the change in $\hat{\beta}$ (and hence all changes in the distribution of education) is not affected by the value of this elasticity, we only include information on wages and inequality. As expected, the change in the standard deviation of log income is decreasing in this elasticity, but even for $\rho = .6$ the resulting change is still substantial—more than 3 percent with no change in the extent to which credit constraints bind and almost 10 percent if low-type families become more constrained.

40. Other researchers have also found evidence for nonlinearities in the relationship between the earnings of parents and children. See, for example, Cooper, Durlauf, and Johnson [1993] and Corak and Heisz [1999].

VIII. CONCLUSION

This paper investigated the effects of increased assortative matching in marriage. We constructed a dynamic model of education acquisition and parameterized it to U. S. data. Using this calibrated model, we conclude that an increase in sorting is likely to have quantitatively significant effects on the degree of income inequality. Although our conclusion holds even in the absence of imperfect borrowing markets, if borrowing constraints exist and are tightened as a result of the increase in sorting, this will magnify the increase in inequality.

The factors that contribute to our obtaining this conclusion are a negative correlation between fertility and education, a decreasing marginal effect of parental education on children's years of education, a process of wage determination that is sensitive to the relative supply of skilled to unskilled workers, and the potential tightening in borrowing constraints.

Our model interpreted borrowing constraints as high-aptitude individuals unable to borrow to cover the cost of obtaining a college education. We do not necessarily take this interpretation literally. An alternative formulation would be to assume that a child's aptitude is determined jointly by parental educational attainment and the resources that they devote to the child's development (for example, the quality of K-12 education the child obtains). If parents are unable to borrow against their child's future income to provide them with greater schooling resources, parental income is again a factor determining investment in a child's future education. This alternative interpretation does not require borrowing constraints to be operative at the time a person decides whether to attend college. Children who grow up in poor families will be less likely to attend college, not because they cannot obtain a loan to finance their college education, but because they have had lower quality K-12 educations and are less able to benefit from a college education.

One important question we have ignored in our analysis is how sorting is determined.⁴¹ How and why does sorting differ

41. Becker [1973] is the classic static model of marriage. See, for example, Burdett and Coles [1997] and Cole, Mailath, and Postlewaite [1992] for models that endogenize the degree of marital sorting and Fernández and Galf [1999] for a model that incorporates borrowing constraints into the matching process. Recent work by Greenwood, Guner, and Knowles [1999] provides a computational analysis that endogenizes fertility, marriage, divorce, and the bargaining power of males and females.

across time and across countries? Household sorting is undoubtedly a complex process that depends upon variables such as income differentials, male relative to female education and income distributions, the degree of sorting in other spheres of society such as residence or schools, on the functioning of networks, and on beliefs (e.g., racism or the desire to transmit a particular culture).⁴² Furthermore, the process may well exhibit multiple equilibria since the desire to match with a certain type may depend on the sorting behavior of others (see, e.g., Burdett and Coles [1997]).

In recent work, Fernández, Guner, and Knowles [2001] and Fernández and Pissarides [2000] endogenize sorting and examine the links with inequality. Suppose that matches are characterized both by a match-specific quality (e.g., love) and by the level of household income. If, *ceteris paribus*, greater income differentials induce high-skill individuals to search longer or harder for another high skill (high-income) partner (i.e., they become pickier about the quality of match with a low-income individual), then any process that increases wage differentials (e.g., skill-biased technological change) or reduces search costs for partners (e.g., internet dating) could well lead to greater sorting and hence greater inequality. Endogenizing fertility so that lower family income generates greater fertility (as in Fernández, Guner, and Knowles) would further reinforce the negative consequences of any increase in wage differentials and also serve to increase sorting.

Our work has ignored education and wage differentials between males and females and bargaining problems within the family. How might greater female education and labor force participation affect sorting and inequality? The most natural possibility would be for greater inequality in the female earning distribution to lead to greater sorting and hence even greater inequality. The findings of Juhn and Murphy [1997] provide empirical support for this effect. They conclude that the increase in women's participation and earnings in the United States has been associated with an increase in the correlation of incomes among spouses. This process would be reinforced if fertility differentials widened. Incorporating other elements, however, such

42. See, for example, Alesina and La Ferrara [2000] for a discussion of how racism leads to sorting and Bisin and Verdier [2000] for an analysis of the dynamic evolution of cultural beliefs.

as the greater ease of divorce, on the other hand, may differentially affect females and males and lead to sorting on a different set of characteristics. By modeling aptitude types as high and low only, we have also ignored the possibility that increases in the skill premium would have the effect of increasing the relative supply of skilled workers by encouraging those individuals with higher effort costs or lower aptitude to become skilled. Fernández [2001b] explores this possibility in the context of the United Kingdom. She finds that results are very sensitive to the assumed elasticity of this supply response. In general, this is an area in which much work remains to be done.

NEW YORK UNIVERSITY, CEPR, AND NBER
ARIZONA STATE UNIVERSITY AND NBER

REFERENCES

- Alesina, Alberto, and Eliana La Ferrara, "Participation in Heterogeneous Communities," *Quarterly Journal of Economics*, CXV (2000), 847–904.
- Banerjee, Abhijit, and Andrew Newman, "Occupational Choice and the Process of Development," *Journal of Political Economy*, CI (1993), 274–298.
- Becker, Gary, "A Theory of Marriage," *Journal of Political Economy*, LXXXI (1973), 813–846.
- Behrman, Jere, Robert Pollak, and Paul Taubman, "Family Resources, Family Size, and Access to Financing for College Education," *Journal of Political Economy*, XCVII (1989), 398–419.
- Benabou, Roland, "Workings of a City: Location, Education, and Production," *Quarterly Journal of Economics*, CVIII (1993), 619–652.
- , "Heterogeneity, Stratification and Growth," *American Economic Review*, LXXXVI (1996), 584–609.
- Bisin, Alberto, and Thierry Verdier, "Beyond the Melting Pot: Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits," *Quarterly Journal of Economics* CXV (2000), 955–987.
- Burdett, Ken, and Melvyn Coles, "Marriage and Class," *Quarterly Journal of Economics*, CXII (1997), 141–168.
- Caucutt, Elizabeth, "Education Policy When There Are Peer Group Effects," University of Rochester, mimeo, 1997.
- Cole, Harold, George Mailath, and Andrew Postlewaite, "Social Norms, Savings Behavior and Growth," *Journal of Political Economy*, C (1992), 1092–1125.
- Cooper, Suzanne, "A Positive Theory of Income Redistribution," *Journal of Economic Growth*, III (1998), 171–195.
- Cooper, Suzanne, Steven Durlauf, and Paul Johnson, "On the Evolution of Economic Status across Generations," *Proceedings of the American Statistical Association*, (1993), 50–58.
- Corak, Miles, and Andrew Heisz, "The Intergenerational Earnings and Income Mobility of Canadian Men," *Journal of Human Resources*, XXXIV (1999), 504–533.
- Dahan, Momi, and Daniel Tsiddon, "Demographic Transition, Income Distribution and Economic Growth," *Journal of Economic Growth*, III (1998), 29–52.
- Durlauf, Steven, "A Theory of Persistent Income Inequality," *Journal of Economic Growth*, I (1995), 75–93.
- Epplé, Dennis, and Richard Romano, "Ends against the Middle: Determining Public Service Provision When There Are Private Alternatives," *Journal of Public Economics* LXII (1996), 297–325.
- Fernández, Raquel, "Sorting Education and Inequality," in *Advances in Economic*

- Theory*, M. Dewatripont, L. Hansen, and S. Turnovsky, eds. (Cambridge University Press, 2001a), forthcoming.
- , "Education, Segregation and Marital Sorting: Theory and an Application to U. K. Data," NBER Working Paper No. 8377, 2001b.
- Fernandez, Raquel, and Jordi Galí, "To Each According to . . . ? Tournaments, Markets and the Matching Problem with Borrowing Constraints," *Review of Economic Studies*, LXVI (1999), 799–824.
- Fernandez, Raquel, Nezh Guner, and John Knowles, "Love and Money: A Theoretical and Empirical Analysis of Household Sorting and Inequality," 2000.
- Fernandez, Raquel, and C. Pissarides, "Matching and Inequality," work in progress, 2000.
- Fernandez, Raquel, and Richard Rogerson, "Income Distribution, Communities, and the Quality of Public Education," *Quarterly Journal of Economics*, CXI (1996), 135–164.
- Fernandez, Raquel, and Richard Rogerson, "Education Finance Reform: A Dynamic Perspective," *Journal of Policy Analysis and Management*, XVI (1997), 67–84.
- Fernandez, Raquel, and Richard Rogerson, "Public Education and Income Distribution: A Dynamic Quantitative Evaluation of Education Finance Reform," *American Economic Review*, LXXXVIII (1998), 813–833.
- Fullerton, Don, and Diane Rogers, *Who Bears the Burden of the Lifetime Tax Burden?* (Washington, DC: Brookings Institution, 1993).
- Galor, Oded, and Josef Zeira, "Income Distribution and Macroeconomics," *Review of Economic Studies*, LX (1993), 35–52.
- Greenwood, Jeremy, Nezh Guner, and John Knowles, "More on Marriage, Fertility, and the Distribution of Income," University of Rochester, mimeo, 1999.
- Jargowsky, Paul, "Take the Money and Run: Economic Segregation in U. S. Metropolitan Areas," *American Sociological Review*, LXI (1996), 984–998.
- Juhn, Chinhui, and Kevin M. Murphy, "Wage Inequality and Family Labor Supply," *Journal of Labor Economics*, XV (1997), 72–97.
- Katz, Lawrence, and David Autor, "Changes in the Wage Structure and Earnings Inequality," in *Handbook in Labor Economics*, volume 3A, Orley Ashenfelter and David Card, eds. (Amsterdam: North-Holland, 1999).
- Katz, Lawrence, and Kevin Murphy, "Changes in Relative Wages, 1963–1987: Supply and Demand Factors," *Quarterly Journal of Economics*, CVII (1992), 35–78.
- Knowles, John, "Marriage, Fertility and Income Distribution," University of Pennsylvania, mimeo, 1999.
- Kremer, Michael, "How Much Does Sorting Increase Inequality?" *Quarterly Journal of Economics*, CXII (1997), 115–139.
- Kremer, Michael, and Daniel Chen, "Income Distribution Dynamics with Endogenous Fertility," *American Economic Review Papers and Proceedings*, LXXXIX (1999), 155–160.
- Kremer, Michael, and Eric Maskin, "Wage Inequality and Segregation by Skill," Harvard University, mimeo, 1996.
- Ljungqvist, Lars, "Economic Underdevelopment: The Case of a Missing Market for Human Capital," *Journal of Development Economics*, XL (1993), 219–239.
- Loury, Glen, "Intergenerational Transfers and the Distribution of Earnings," *Econometrica*, XLIX (1981), 843–867.
- Mare, Robert, "Five Decades of Educational Assortative Mating," *American Sociological Review*, LVI (1991), 15–32.
- , "Differential Fertility, Intergenerational Educational Mobility and Racial Inequality," *Social Science Research*, XXVI (1997), 263–291.
- Reich, Robert, "Secession of the Successful," *New York Times Magazine*, January 20, 1991.
- Wilson, William J., *The Truly Disadvantaged: The Inner City, the Underclass, and Public Policy* (Chicago, IL, University of Chicago Press, 1987).