CHAPTER FOUR

Intergenerational Mobility and Income Inequality

Jörg L. Spenkuch

My remarks today are based on recent work with the late Gary Becker, Kevin Murphy, and Scott Kominers (see Becker et al. 2014). It bears emphasizing that Gary has already written the two seminal papers in the literature on intergenerational mobility and income equality (Becker and Tomes 1979, 1986). Yet, he felt it was worth going back and giving these issues more thought. Although the economics of our analysis are very simple—in fact, we rely only on basic price theory—the conclusions that we derive differ radically from those in the literature.

Before delving into the analysis, let us start out with some simple facts that others have documented (see Corak 2013a, Corak and Heisz 1999, Mazumder 2005, among others). These facts cover much of what is currently known about intergenerational mobility in the United States and elsewhere. After presenting the facts, I will lay out a simple theory in order to make sense of the data and to link intergenerational persistence of economic status with cross-sectional inequality. Lastly, I will discuss the role of government spending in reducing both. In particular, I will be asking how governments can achieve these goals, given the economics of the problem. Is there a role for government spending and, if so, what are "good" government interventions?

Looking at the cross-country data in figure 4.1, we see a strong positive relationship between inequality as measured by the Gini coefficient and intergenerational persistence in income—more

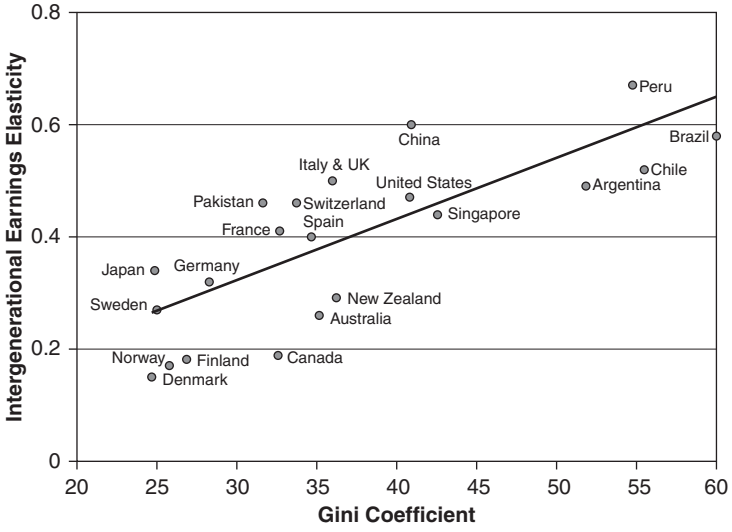

FIGURE 4.1. The Great Gatsby Curve

Notes: Figure shows the relationship between inequality, as measured by the Gini coefficient on the x-axis, and intergenerational persistence, as measured by the intergenerational earnings elasticity on the y-axis. Higher values indicate more inequality and more persistence, respectively.

Source: Miles Corak, "Inequality from Generation to Generation: The United States in Comparison," in *The Economics of Inequality, Poverty, and Discrimination in the 21st Century, vol. 1,* edited by Robert Rycroft (Santa Barbara, CA: Praeger, 2013)

popularly known as the "Great Gatsby Curve" (Krueger 2012). The US, Italy, and the United Kingdom, for instance, are countries with relatively high levels of inequality and low intergenerational mobility. By contrast, the Scandinavian countries feature high levels of mobility and much less inequality.

As figure 4.2 shows, there is also a strong positive correlation between the college earnings premium and intergenerational persistence in incomes. Compared with other countries, the US has a very large college earnings premium, but it also has a lot more persistence in incomes across generations.

Intriguingly, this graph would look very similar if, instead of the college premium, we were to plot government spending on

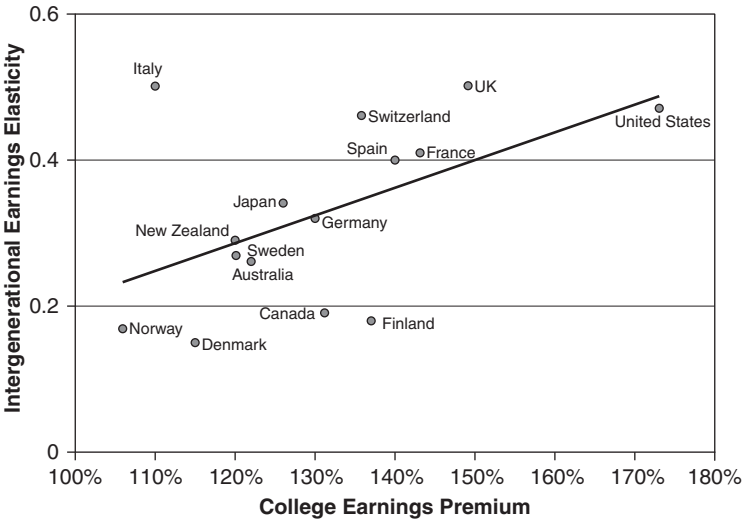

FIGURE 4.2. Intergenerational persistence and the college earnings premium

Notes: Figure shows the relationship between the college earnings premium (*x*-axis), as measured by the ratio of the average employment income of 25–34-year-old men with a college degree over that of their high-school educated counterparts, and intergenerational persistence (*y*-axis), as measured by the intergenerational earnings elasticity. Higher values indicate larger earnings premiums and more persistence, respectively.

Sources: Miles Corak, "Income Inequality, Equality of Opportunity, and Intergenerational Mobility," *Journal of Economic Perspectives* 27 (3); Miles Corak, "Inequality from Generation to Generation: The United States in Comparison," in *The Economics of Inequality, Poverty, and Discrimination in the 21st Century, vol. 1,* edited by Robert Rycroft (Santa Barbara, CA: Praeger, 2013); and OECD (2011)

tertiary education on the *x*-axis—i.e., government spending on higher education, elite universities, etc. The relationship would reverse, however, if one were to consider all government spending on education instead. This suggests that there is something about government spending on higher education (as opposed to government spending on primary education) that is correlated with intergenerational persistence.

Instead of looking across countries, let us look at intergenerational mobility within countries next. Figure 4.3 plots a son's probability of falling into a given earnings decile given that his father

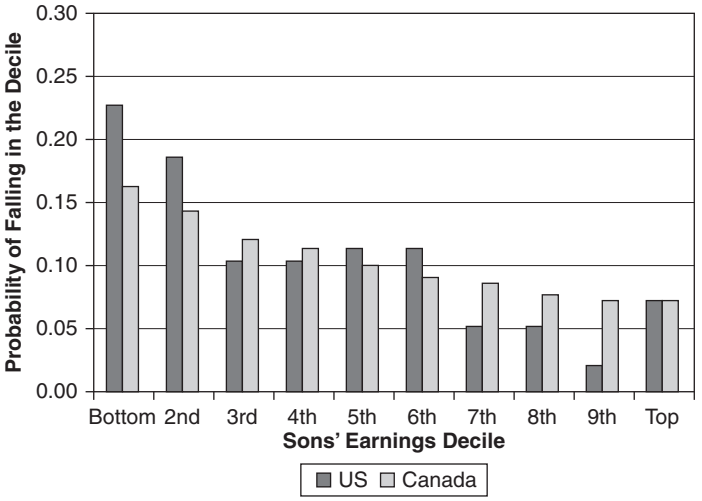

FIGURE 4.3. Earnings deciles of sons born to bottom-decile fathers

Notes: Figure shows the probability that a son whose father was in the first decile of the income distribution ends up in any given decile of his generation's distribution.

Sources: Miles Corak, "Income Inequality, Equality of Opportunity, and Intergenerational Mobility," *Journal of Economic Perspectives* 27 (3); Miles Corak and Andrew Heisz, "The Intergenerational Earnings and Income Mobility of Canadian Men: Evidence from Longitudinal Income Tax Data." *Journal of Human Resources* 34 (3) (1999); and Bhashkar Mazumder, "The Apple Falls Even Closer to the Tree than We Thought: New and Revised Estimates of the Intergenerational Inheritance of Earnings," in *Unequal Chances: Family Background and Economic Success,* ed. Samuel Bowles, Herbert Gintis, and Melissa Osborne Groves (New York: Russell Sage Foundation)

was in the bottom decile of the income distribution. It does so for both the US and Canada. For instance, if a US father was in the bottom decile of the income distribution of his generation, there is approximately a 22 percent probability that his son ends up in the bottom decile of the earnings distribution of the next generation. Strikingly, there is only about a 7 percent probability of the son falling in the top decile of the distribution. The same basic pattern holds in the US and Canada, although it is less pronounced in the latter. Loosely speaking, the data show that there is a lot of persistence at the bottom of the distribution. Simply put, children of poor parents are much more likely to end up being poor than rich.

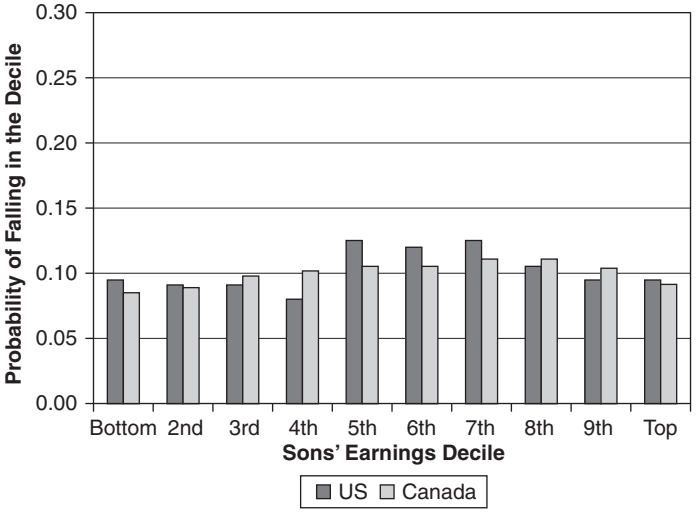

FIGURE 4.4. Earnings deciles of sons born to middle-two-decile fathers

Notes: Figure shows the probability that a son whose father was in either the fifth or the sixth decile of the income distribution ends up in any given decile of his generation's distribution.

Sources: Miles Corak and Andrew Heisz, "The Intergenerational Earnings and Income Mobility of Canadian Men: Evidence from Longitudinal Income Tax Data." *Journal of Human Resources* 34 (3) (1999); and Bhashkar Mazumder, "The Apple Falls Even Closer to the Tree than We Thought: New and Revised Estimates of the Intergenerational Inheritance of Earnings," in *Unequal Chances: Family Background and Economic Success,* ed. Samuel Bowles, Herbert Gintis, and Melissa Osborne Groves (New York: Russell Sage Foundation)

A very different picture emerges in figure 4.4, where, instead of conditioning on the father being in the bottom decile of the income distribution, we condition on him being in one of the two middle deciles. Although there is a weak tendency for the sons of these fathers to remain in the middle of the distribution, the probability of making it to the top—or the very bottom—is almost as high.

The last figure in this series, i.e., figure 4.5, deals with fathers in the top decile. Again, we see a lot of persistence, especially in the US. For instance, conditional on the father being in the top decile of the distribution, there is a 27 percent probability that the son will end up in the top decile as well.

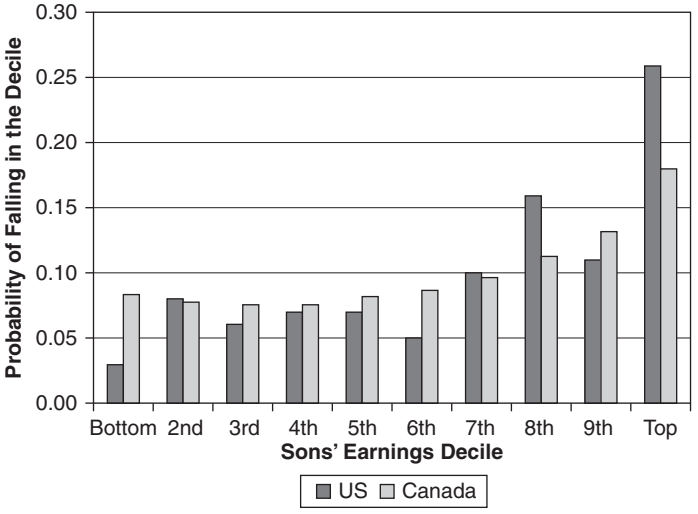

FIGURE 4.5. Earnings deciles of sons born to top-decile fathers

Notes: Figure shows the probability that a son whose father was in the tenth decile of the income distribution ends up in any given decile of his generation's distribution.

Sources: Miles Corak, "Income Inequality, Equality of Opportunity, and Intergenerational Mobility," *Journal of Economic Perspectives* 27 (3); Miles Corak and Andrew Heisz, "The Intergenerational Earnings and Income Mobility of Canadian Men: Evidence from Longitudinal Income Tax Data," *Journal of Human Resources* 34 (3) (1999); and Bhashkar Mazumder, "The Apple Falls Even Closer to the Tree than We Thought: New and Revised Estimates of the Intergenerational Inheritance of Earnings," in *Unequal Chances: Family Background and Economic Success,* ed. Samuel Bowles, Herbert Gintis, and Melissa Osborne Groves (New York: Russell Sage Foundation)

Why is there so much persistence at the top of the income distribution but not in the middle? Are there economic forces that imply high persistence in one part of the distribution but not in the other? One potential explanation is that rich parents invest a lot more in their children than poorer ones.

Figure 4.6 shows enrichment expenditures per child in the US. The data differentiate between parents with different incomes and are available for the period from 1972 to 2006. Even as early as the 1970s, there has always been a nontrivial difference in how much parents at opposite ends of the income distribution invest in their

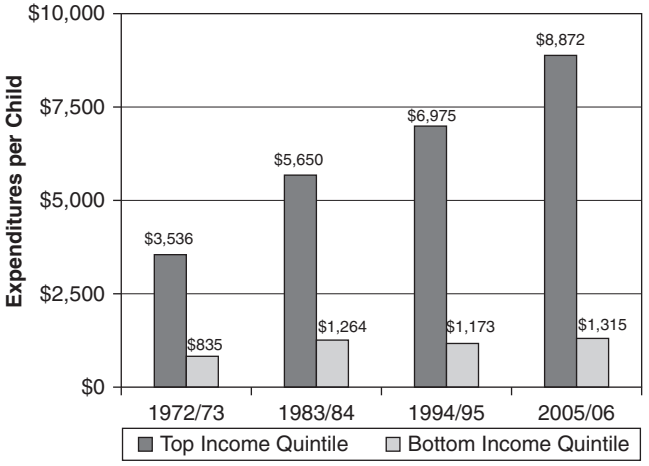

FIGURE 4.6. Enrichment expenditures per child in the US, by parental income

Notes: Numbers are in 2008 dollars and refer to parental spending on books, computers, child care, summer camps, private school, and other expenditures to promote children's abilities.

Sources: Miles Corak, "Income Inequality, Equality of Opportunity, and Intergenerational Mobility," *Journal of Economic Perspectives* 27 (3) (2013); and Greg Duncan and Richard Murnane, "Introduction: The American Dream, Then and Now," in *Whither Opportunity? Rising Inequality, Schools, and Children's Life Chances*, ed. Greg Duncan and Richard Murnane (New York: Russell Sage Foundation, 2011)

children. Perhaps more importantly, the gap has widened dramatically over the last three to four decades, and the same basic pattern holds not only for monetary expenditures, but also for parental time spent with children (see Guryan et al. 2008, Ramey and Ramey 2010). Guided by these facts, we develop a theory of the intergenerational transmission of resources. Our main goal is to understand why persistence among wealthy families is so much higher than in the middle class.

Our approach is simple yet powerful. Since labor income constitutes the vast majority of individuals' earnings—even among the "top 1 percent" (see Piketty and Saez 2003)—we believe that we have to understand inequality in human capital if we want to understand inequality within and across generations. In an

attempt to do so, we model earnings (E) by assuming that each unit of human capital (H) yields a given rate of return (r), as in $E = rH + e$, where the random term e refers to all other determinants of earnings, such as market luck or macroeconomic conditions beyond individuals' control. Consistent with the empirical record, in our setup cross-sectional inequality increases directly with the return to education, i.e., r. Over the last three decades or so, the returns to education have risen dramatically, especially in the US (Juhn et al. 1993, Murphy and Katz 1992). At the same time, inequality has increased, as well.

Since education is the main determinant of earnings, intergenerational mobility in income depends critically on the persistence of human capital across generations. That is, how much of my own human capital is determined by the human capital or the earnings of my parents? All else equal, intergenerational mobility will be low whenever parents exert a great influence over their offspring's human capital.

Gary Becker's earlier work had already emphasized the importance of human capital for linking cross-sectional inequality and intergenerational mobility. Progress since, however, has been largely confined to empirical work (see Solon 1999 and Black and Devereux 2011 for useful reviews). We build on the early theoretical literature, but allow for complementarities in the "production" of human capital. That is, we allow more educated parents to be "better" at investing in the human capital of their children. This seemingly minor departure turns out to be very important. In fact, allowing for complementarities radically changes some of the existing literature's conclusions.

A simple formulation for the production of children's human capital is $H_c = F(y,H_p,G)$, where H_c denotes children's human capital, H_p is that of parents, y gives parental investments in children, and G is government spending. In words, we assume that children's human capital depends on the investments they receive

from their parents, how educated their parents are, and how much the government spends on schools, etc. Naturally, all of these factors should have a positive effect on children's human capital.

The crucial assumption is that parental human capital and spending on children are complements. Technically, we posit that $\partial^2 H_c / \partial y \partial H_p > o$. Essentially, this means that educated parents are more productive or more efficient at investing in their children. In this respect our analysis departs from previous work. Existing research has simply assumed that the productivity of parents' investments is independent of their own human capital.

Taking their human capital, i.e., H_p, and government spending, G, as given, parents choose how much to invest in their children. For analytic simplicity, I focus on the case with perfect capital markets, meaning that parents are not credit-constrained. When capital markets are perfect, even poor parents can invest as much as they want in the education of their children, say, by borrowing against the income of future generations. This does not imply that parents invest infinitely much. On the contrary, in the case of perfect capital markets, the return on investing in children has to be, at the margin, exactly as high as the return parents would get from putting their money or time to the next best use.

Of course, it is not clear whether the assumption of perfect capital markets is, in fact, satisfied. Much current research wrestles with precisely this question (see Lochner and Monge-Naranjo 2012 and the references therein). What makes this assumption attractive for our purposes is that it allows us to considerably simplify the analysis. Moreover, it brings out more starkly the difference between our results and those of previous work. Existing research generally concludes that perfect capital markets break the connection between parents' human capital and that of their children. The reason is that perfect capital markets allow *all* parents to invest the optimal amount in their children. Hence, equally able children will receive equal investments, independent of their parents' income,

which, in turn, leads to similar levels of intergenerational mobility among children of the middle class and those of the rich. We have already seen that this is not true in the data, and our model explains why.

As a side note, we obtain broadly similar conclusions when we allow for lower-income families to be credit-constrained. The most important difference between the cases of perfect and imperfect capital markets is that credit constraints impose tight limits on how much poor parents can borrow and, therefore, invest in the human capital of their children. Thus, credit constraints directly lower the degree of intergenerational mobility at the bottom of the income distribution, as in figure 4.3. Since this is hardly surprising and since we are primarily interested in explaining why there is so much more persistence at the top of the distribution than among members of the middle class, we abstract from credit constraints—though we do pay careful attention to them in Becker et al. (2014).

If the return on savings is given by the economy-wide return on capital, R_k, then, at the optimum, the marginal return on investing in children's human capital must equal R_k. Or in symbols, $R_y \equiv rF_y = R_k$. This condition is, of course, familiar from Gary's seminal works on the economics of the family. The novel implication of our analysis is that high human capital parents invest more in children than their low human capital counterparts. The reason is simple: due to the assumed complementarities, investments of high human capital parents are more productive. Thus, highly educated parents invest more than their less educated counterparts.

After allowing for complementarities in the production function, even equally able children will receive different levels of investments from their parents, depending on whether the parents had high or low human capital. Not only does this prediction match the evidence in figure 4.6, it directly implies that children

born into different families will end up in different parts of the income distribution.

By contrast, without complementarities in the production function, i.e., if parents' human capital had no effect on that of their offspring, equally able children would receive the same investments and, therefore, fall in the same part of the income distribution. In a world where the productivity of parental investments is independent of parental characteristics and other environmental factors, there would be perfect intergenerational mobility in income, at least conditional on ability. However, once we allow for complementarities—which we think are important in the real world—this conclusion breaks down.

Let us consider a closed-form example. Suppose that the production function for children's human capital is given by $H_c = \mu + \kappa y + \varphi y^2 + \theta y H_p + v_c$, where $\kappa, \theta > o$ and $\varphi < o$. In words, children's human capital depends positively on the investments they receive from their parents (y) as well as other factors that are beyond agents' control (v_c). Investments in children are subject to diminishing marginal returns but complementary to parental human capital (because $\theta > o$).

Given these simplifying assumptions, it is straightforward to derive the optimal level of parental investments, i.e., $y^* = (R_k/r - \kappa - \theta y H_p)/(2\varphi)$. We use this expression to obtain a reduced-form relationship between the human capital of children and that of their parents: $H_c = \mu^* + \delta^* H_p + \gamma^* H_p^2 + v_c$, where $\mu^* \equiv \mu - (\kappa^2 r^2 - R_k^2)/(4\ \varphi r^2)$, $\delta^* \equiv -(\kappa\theta)/(2\varphi)$, $\gamma^* \equiv -(\kappa\theta)/(4\varphi)$. The fact that the relationship between children's and parent's human capital is quadratic turns out to be important.

Why? A convex relationship implies that the influence of parents' human capital on that of their children increases as parents become more educated. Mathematically, $\partial^2 H_c/\partial H_p^2 > o$. That is, not only will children of more educated parents be more educated

themselves, but the marginal impact of additional parental human capital rises. Hence, children born to parents in the upper part of the income distribution will be more similar to their parents (in terms of human capital) than children born to caregivers in the lower part of the distribution, which implies greater intergenerational persistence at the top.

While the example above may seem quite stylized, our conclusions hold much more generally. Unless there are strong diminishing marginal returns to parental human capital, complementarity in the production function implies more persistence at the top. Particular to the example is that the degree of persistence in human capital, and therefore the degree of intergenerational mobility, does not depend on the market return of human capital, i.e., r. In more general formulations, the effect of a rise in the return to education can go either way. That is, it can either lower or increase the degree of intergenerational persistence.

The media often convey the impression that increases in the market return to human capital lower mobility. Our analysis shows that need not be the case. Although rising returns to human capital aggravate the consequences of the birth lottery (because the income distribution spreads out), they do not necessarily change the intergenerational transmission of human capital itself. Measuring intergenerational mobility with respect to individuals' position in their generation's income distribution, our model predicts almost no change over the last few decades—despite the large increase in the college premium. Perhaps surprisingly, this prediction is supported by the best (newly) available empirical evidence (see Chetty et al. 2014).

To think about steady states and the long-run distribution of human capital, let us go back to the example. Figure 4.7 depicts the relationship between children's human capital (on the y-axis) and that of their parents (on the x-axis). In panel A on the top left, we analyze the usual case, in which there is only one steady state. In

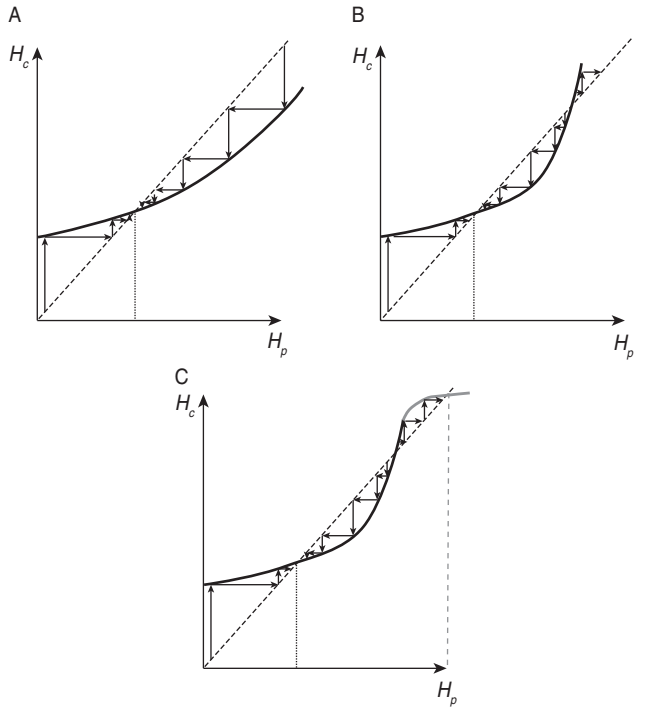

FIGURE 4.7. Steady state analysis
Source: Author's illustration

this panel, the transmission function crosses the forty-five-degree line from above. Thus, the steady state is stable; and over the long run, there is strong regression to the population mean.

In our model, however, this is not the only possible case. Panel B on the top right depicts a scenario in which complementarities in the production of children's human capital are so strong that the transmission function crosses the forty-five-degree line twice—first from above and then from below. Economically, this means that children of very highly educated parents end up having even more human capital than the previous generation. As a consequence, there is bifurcation. One set of families regresses toward the mean, whereas an "elite" keeps regressing away from the mean.

Naturally, unbounded regression away from the mean is not necessarily realistic. It seems reasonable that, at some point, even the elite's human capital would level out. To incorporate this feature we can add a small cubic term to the transmission function in Panel C. The result is a second stable steady state. There would thus be a highly persistent elite atop the distribution (i.e., around the stable steady state featuring high human capital) as well as a persistent "underclass" at the bottom (i.e., around the stable steady state featuring low human capital). Among the middle class, however, there would be a lot of churning—especially as the point where the transmission function crosses the forty-five-degree line from below constitutes an unstable steady state. Our model can, therefore, rationalize why there seems to be a lot of persistence at both ends of the income distribution. Importantly, there need not be any credit constraints for this to be true.

Before concluding, I will briefly discuss the role of government spending. To incorporate government spending into our analysis, we let G enter the production function for children's human capital. We assume that, all else equal, government spending weakly raises the human capital of children—simply because well-intentioned governments are probably not doing any direct harm. That does not mean, however, that there are no unintended side effects. The reason is that other factors are not necessarily going to be equal. When parental inputs adjust in response to government interventions, the key question becomes: how do G, y, and H_p interact?

We analyze two cases, which together cover a wide range of real-world government programs. In the first case, government spending (G) and parental investments (y) are perfect substitutes. For instance, if the government provides school books for children, parents are unlikely to incur the same expense. A similar argument may apply to, say, computers, free lunches, and a number of other educational supplies. Of course, if government spending is a

perfect substitute for spending by parents, then private spending is going to be crowded out one for one. Total investments in children's human capital will only increase if the government spends more than parents would.

Given that parental investments were optimal to begin with (due to the assumption of perfect capital markets), this gives rise to an important equity-efficiency trade-off. If crowding out is quantitatively important, then government interventions that increase the human-capital acquisition of children have to be very large, which raises concerns about the deadweight loss of taxation. Moreover, the marginal return on this spending is (weakly) lower than the return on capital. For this kind of intervention to be justified, a social planner would have to put a lot of weight on equity.

At least in part, the welfare effect of government interventions will depend on whether families are credit-constrained. In the presence of credit constraints, human-capital investments among poor families are not optimal. Thus, interventions that raise total spending on the children of these families may actually be welfare-improving. In particular, if the deadweight loss of taxation is not too high, there may not be an equity-efficiency trade-off at all. By contrast, if capital markets are perfect, that is, when parental investments are (close to) optimal, then there is little role for government spending.

As a side note, there is no consensus (yet) as to how important credit constraints are (for opposing results, see, for instance, Lovenheim 2011, Carneiro and Heckman 2002, or Belley and Lochner 2007). One strand of the literature finds that, conditional on a host of observables, parental resources and children's educational attainment are almost uncorrelated. These studies argue that credit constraints cannot be quantitatively important. Another set of studies does find a positive correlation, especially in recent years, which is then interpreted as evidence of credit constraints. Our analysis implies that such tests are theoretically ill-founded.

As we have seen earlier, even in a world with perfect capital markets, if there are complementarities in production, then children's human capital will generally be correlated with that of their parents and, therefore, with parental income.

Next, we analyze the (perhaps more interesting) case in which government spending and private investments are complements. Under this assumption, additional government spending increases the productivity of parental investments, which in turn induces families to invest more. Examples where such multiplier effects are plausible might include government spending on better teachers or elite universities.

Even when government spending does not crowd out private investments, there is still an equity-efficiency trade-off, though of a different kind. To clearly bring out this trade-off, consider a government that spends an equal amount of money on children in rich and poor families. Although such a government program may seem neutral, its effects are not. Government interventions of this kind actually *increase* inequality.

Why? The reason is that when government spending and private investments are complements, then additional government funds will have a bigger impact on the human capital of children in rich families. After all, their parents have more human capital, making the government spending more productive. Again, this result is a direct consequence of complementarities in the production function. Moreover, higher human capital families increase their own investments in children because these investments are being made more productive by the government intervention. To be just "neutral" in the sense of raising the human capital of all children equally, government spending would have to be biased toward poor families.

This, however, raises efficiency concerns. When capital markets are perfect, poor families are the ones in which additional government spending is the *least* productive. The efficient use of funds

in this case would be to invest in the children of rich families. Yet, from the perspective of a social planner who is also concerned about equity, that may be undesirable.

Again, these conclusions are greatly altered if credit constraints limit poor families' ability to optimally invest in their children. If credit constraints are sufficiently important, then biasing government spending toward children in poor families may be equity-*and* efficiency-enhancing.

To conclude, our analysis builds on earlier theoretical work that links human capital to intergenerational mobility and cross-sectional inequality. We depart from the existing literature by emphasizing complementarities in the production of children's human capital. If these complementarities are sufficiently strong, our theory predicts lower intergenerational mobility at the top of the income distribution than in the middle—even in the absence of credit constraints.

Moreover, our analysis implies that well-intentioned government policies may end up having perverse consequences, depending on whether a particular intervention is a substitute or complement to private spending. In reality, many government programs likely have elements of both, which makes predicting their effect even harder.

Beyond what I have already mentioned, our model is useful for thinking about the likely consequences of many recent changes in the marketplace. For instance, not only has the US experienced a dramatic increase in the returns to education, but assortative mating has become more important as well. The probability that a college-educated man marries a college-educated woman is much higher today than it was forty years ago. Suitably extended, our theory implies that such changes result in less regression to the mean in ability and, therefore, in less intergenerational mobility. A number of other extensions and implications are explained in our working paper (see Becker et al. 2014).

Question and Answer Session

QUESTION: *If we take to heart your interpretation of the facts, it seems to me it becomes quite important for researchers to unpack the sources of complementarity in the production of children's human capital. Another way to put the issue is: what is it about higher levels of human capital that shifts up the schedule that describes the return of parental investment in children's human capital? Is it something as simple as reading a bedtime story to a four-year-old, which is a skill that might be easily transmitted to many people? Or is it something that requires the kind of vocabulary and thinking process that comes along with a college education, which is obviously a much more costly thing to transmit?*

And the answer to that question, it seems to me, also has profound implications for the proper role of government, if government sets it as a goal to reduce inequality, about how to do it. It may be that if there are simple methods of teaching parents with lower levels of human capital, what are the skills, what are the habits, what are the practices that increase the returns on their investments in their children, that maybe government policy ought to be directed in that way.

SPENKUCH: I agree with you, and I wish that I knew the exact source of these complementarities. As you said, there are certain tasks, like reading a bedtime story, that parents can do about equally well—at least parents with a minimal level of human capital. But once we think about tasks that require just a little more education, complementarities may start to become important. For instance, it is easy for you to help your son with his algebra homework, but if his parents were high school dropouts or GED graduates instead, that might not be the case anymore. Such parents may be able to spend a lot of time with their children, but each unit of time is likely less effective than if they were more educated.

QUESTION: *I'd like to complicate your problem by describing California. For many students, the government has a monopoly over access to education. The government spending produces the public school and you have to go to it. You don't have a choice. They're trying to expand choice, but the reality is, for many students, they have no choice. The teachers in the school are unionized and their dues are deducted automatically from their paycheck. The teachers' union has a gigantic cash flow, which they spend in elections. And they elect the people who they then bargain with or who control the schools. That puts you in an odd position. And there's been an interesting lawsuit in LA, which is catching attention, where California was sued for violating the Constitution and depriving kids of a decent education because of refusal to fire incompetent teachers. According to the rules, it's practically impossible to get rid of somebody. Now the governor and the attorney general are both going to contest this ruling by arguing that incompetent teachers have a constitutional right to ruin kids' lives. It's going to be wonderful to see how this plays itself out.*

SPENKUCH: Let me play devil's advocate—just for a little bit. As devil's advocate, I would argue that the government is really effective at reducing inequality. If everybody gets the same poor education, everybody ends up poor. Of course, this is not prudent policy. Linking your question to the model, our analysis would say, since the government provides only a deficient education, it is important to figure out what parents do in addition to that. For instance, you might imagine that many high-socioeconomic-status parents purchase, say, SAT tutoring for their children, or piano and violin lessons. Those are enrichment expenditures that governments typically do not provide.

Comment: *There's this implicit assumption that the government provides it. Maybe they spend equal dollars on everybody, and that's an equal education. The truth is the well-educated parents are much better at navigating the system both within a district*

and across districts. Talk to anybody who does special education and they know that the well-educated parents are in there arguing for every little thing their child is entitled to that's within a district, let alone the fact that they move to districts where the public schooling actually is pretty good.

So the idea that the government is giving everybody an equal deal and that education matters less when the government's providing it doesn't seem to be consistent with the facts. Wal-Mart is more similar in poor and rich neighborhoods than schools are. I always think that's a great comparison to make. I walk into my Wal-Mart, and it looks a lot like the Wal-Mart in another neighborhood that's a lot poorer, but the schools don't look anything like each other.

It's interesting because one of the big equalizers for people is the market. People get to take advantage of other people's skills. And when you do education in this way, you defeat that market mechanism. I can't use the skills of other people to help educate my kids. I've got to do it myself. I've got to know where to live, I've got to know how to navigate the system, how to get the best teachers, all of those things. It's actually raising inequality. And this fits into what Jörg was talking about earlier. How is this government spending really a complement or a substitute?

One of the things that happens in the market is you get all these changes and there's increasing demand for more skilled workers, but there's a supply response that benefits everybody, including those who aren't so educated, because they get to take advantage of what doesn't happen in a household. You're much more on your own in a household and that's why this complementarity is so important, and why government spending, by defeating the market in many ways, actually makes it worse rather than better.

QUESTION: *One question and two observations. On figures 4.3 to 4.5, you have the sons in particular deciles. That has to be at a*

point in time. So as you're laying that out, I'm wondering what point in time. People move between deciles a lot and so I'd just like one answer, if you'd hold it for the other two.

And what I find striking in the twenty-first century, even late twentieth century, is just looking at sons. I mean, there are daughters, right? Anyway, that's quite striking to me.

The other thing that strikes me is that your whole model is based on the idea that human capital is like other capital and there's not a big signaling element. If you have a big signaling element—I'm a blogger at EconLog with Bryan Caplan, he's writing a book on this. And he finds that signaling—he's already got a lot of data—signaling is huge. It's probably well over half. And what you're going to do is just have more rats running around the wheel if you put more money into it.

SPENKUCH: I completely agree with your last point, there may very well be signaling. If you wanted to incorporate signaling into the model, you could allow for parents who went to, say, Harvard to have an easier time getting their children into other elite universities. The model is actually quite flexible. It depends on what you want to call "human capital." Broadly construed, human capital might include all personal characteristics that raise wages. As long as there is still a complementarity between those characteristics and investments in children's human capital, our conclusions continue to go through.

To your earlier point about sons and daughters, there is an emerging literature that looks at the earnings of daughters and how they correlate with those of their parents. Econometrically, such an exercise is a lot harder to do than the corresponding one for sons because of selection into market work. Even in the twenty-first century, not all daughters work, and the ones who do are probably systematically different from the ones who do not. This may be the reason for why, even today, most of the literature focuses on sons.

Coming back to your very first point about how these graphs are constructed: the papers on which these graphs are based average earnings over a period of time, sometimes ten years or even longer. Ideally, one would like to look at lifetime earnings, but the available data now do not allow for that yet.

QUESTION: *I really like the paper and started thinking through some of the implications. And you touched on the credit constraints but I would use a different terminology. I would say tests for whether parental income matters for kids' investments holding kids' ability constant—put it that way because that's the only way to make sense in your model.*

My question for you is: if those tests fail, isn't that a rejection of your technology? And then the next thing I wanted to point out about credit constraints is that consumption data is very helpful. Even in your technology, I believe the consumption data is going to reveal—or fail to reveal—credit constraints. And it also solves your sons-and-daughters problem because maybe all daughters don't work but they all consume. And there's data on consumption. People have used it to look at these mobility issues. And so I think you can promote that literature a little bit.

SPENKUCH: I completely agree with you. Ultimately, people derive utility from consumption, and we should look at that. Regarding the earlier point: our model predicts a positive correlation between parental resources and the education of children, even after controlling for ability. The reason is that there are complementarities in production. The previous literature has interpreted these correlations as evidence of credit constraints. We, however, argue that such a conclusion is premature, at least if we think that complementarities are quantitatively important.

QUESTION: *Seems to me that today there are two classes of people, say, graduating from a Stanford or a prestigious university—those who have no debt because their parents have been able to pay*

for everything and those who come out, however bright and well-educated, they're way down in a hole. So if you're looking at their future income possibilities and so forth, in one case, people can't save at an earlier age and invest and complement their income, whereas others can, if they're wisely guided. They'll be able to start investing at a much earlier age. There's been an accumulation of huge amounts of student debt. How is that going to play into this? And to what extent is that debt, since a lot of it isn't going to get repaid, essentially a government subsidy?

SPENKUCH: Your point about student debt strikes me as very important. Most research is concerned with correlations in incomes or positions in the income distribution, but we fail to look at lifetime utility. It is quite plausible that someone who graduated from Stanford with a hundred and fifty thousand dollars of debt has much lower lifetime utility than someone else who is otherwise similar but does not have as much debt. If we could look at lifetime utility, we would probably find that parental resources have a much greater effect on utility than on educational attainment or salaries.

References

Becker, Gary, Kevin Murphy, Scott Kominers, and Jörg Spenkuch. 2014. "A Theory of Intergenerational Mobility," mimeographed, University of Chicago.

Becker, Gary, and Nigel Tomes. 1979. "An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility." *Journal of Political Economy* 87 (6): 1153–1189.

Becker, Gary, and Nigel Tomes. 1986. "Human Capital and the Rise and Fall of Families." *Journal of Labor Economics* 4 (3): S1–S39.

Belley, Philippe, and Lance Lochner. 2007. "The Changing Role of Family Income and Ability in Determining Educational Achievement." *Journal of Human Capital* 1 (1): 37–89.

Black, Sandra, and Paul Devereux. 2011. "Recent Developments in Intergenerational Mobility." In *Handbook of Labor Economics, vol. 4B,* edited by David Card and Orley C. Ashenfelter, 1487–1541. San Diego, CA: North Holland.

Carneiro, Pedro, and James Heckman. 2002. "The Evidence on Credit Constraints in Post-Secondary Schooling." *Economic Journal* 112 (482): 705–734.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, Emmanuel Saez, and Nicholas Turner. 2014. "Is the United States Still a Land of Opportunity? Recent Trends in Intergenerational Mobility." NBER Working Paper No. 19844.

Corak, Miles. 2013a. "Income Inequality, Equality of Opportunity, and Intergenerational Mobility." *Journal of Economic Perspectives* 27 (3): 79–102.

Corak, Miles. 2013b. "Inequality from Generation to Generation: The United States in Comparison." In *The Economics of Inequality, Poverty, and Discrimination in the 21st Century, vol. 1,* edited by Robert Rycroft. Santa Barbara, CA: Praeger.

Corak, Miles, and Andrew Heisz. 1999. "The Intergenerational Earnings and Income Mobility of Canadian Men: Evidence from Longitudinal Income Tax Data." *Journal of Human Resources* 34 (3): 504–533.

Duncan, Greg, and Richard Murnane. 2011. "Introduction: The American Dream, Then and Now," in *Whither Opportunity? Rising Inequality, Schools, and Children's Life Chances*, edited by Greg Duncan and Richard Murnane. New York: Russell Sage Foundation.

Guryan, Jonathan, Erik Hurst, and Melissa Kearney. 2008. "Parental Education and Parental Time with Children." *Journal of Economic Perspectives* 22 (3): 23–46.

Juhn, Chinhui, Kevin Murphy, and Brooks Pierce. 1993. "Wage Inequality and the Rise in the Returns to Skill." *Journal of Political Economy* 101 (3): 410–442.

Krueger, Alan. 2012. "The Rise and Consequences of Inequality in the United States." Speech at the Center for American Progress, Washington, DC.

Lochner, Lance, and Alexander Monge-Naranjo. 2012. "Credit Constraints in Education." *Annual Review of Economics* 4 (1): 225–256.

Lovenheim, Michael. 2011. "The Effect of Liquid Housing Wealth on College Enrollment." *Journal of Labor Economics* 29 (4): 741–777.

Mazumder, Bhashkar. 2005. "The Apple Falls Even Closer to the Tree than We Thought: New and Revised Estimates of the Intergenerational Inheritance of Earnings," in *Unequal Chances: Family Background and Economic Success,* edited by Samuel Bowles, Herbert Gintis, and Melissa Osborne Groves. New York: Russell Sage Foundation.

Murphy, Kevin, and Lawrence Katz. 1992. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *Quarterly Journal of Economics* 107 (1): 35–78.

OECD. 2011. *Education at a Glance 2011: OECD Indicators*. Paris: OECD Publishing.

Piketty, Thomas, and Emmanuel Saez. 2003. "Income Inequality in the United Sates, 1913–1998." *Quarterly Journal of Economics* 118 (1): 1–39.

Ramey, Valerie, and Garey Ramey. 2010. "The Rug Rat Race." *Brookings Papers on Economic Activity,* Spring 2010: 129–176.

Solon, Gary. 1999. "Intergenerational Mobility in the Labor Market." In *Handbook of Labor Economics, vol. 3A,* edited by David Card and Orley C. Ashenfelter. Amsterdam: Elsevier.

Corrigendum

Unfortunately, the formulae in the published chapter contain two typos. I had pointed them out in the page proofs, but, inexplicably, they were not corrected before the book went into print.

Specifically, on p. 75, y^* should be $y^* = (\frac{R_k}{r} - \kappa - \theta H_p)/(2\varphi)$, and γ^* ought to be defined as $\gamma^* \equiv -\theta^2/(4\varphi)$.