

Ending Extreme Poverty and Promoting Shared Prosperity Could There Be Trade-offs Between These Two Goals?

**Kathleen Beegle, Pedro Olinto, Carlos Sobrado,
Hiroki Uematsu, and Yeon Soo Kim,
with Maximillian Ashwill**

In April 2013, the Joint Ministerial Committee of the Boards of Governors of the World Bank Group and the International Monetary Fund endorsed two ambitious goals to galvanize international development efforts by donors, governments and the international development community, generally. These goals are (1) to end extreme poverty (defined as reducing the number of people earning \$1.25 per day or less to no more than 3 percent



Photo by Dorte Verner

Could a focus on shared prosperity turn global attention away from the extremely poor?

shared prosperity, has some important implications for global development efforts. Notably, this new goal is a refinement of a longer-standing, implicit focus on growth, which is widely con-

The authors would like to acknowledge the helpful peer review comments of Francisco Ferreira, Peter Lanjouw and Martin Ravallion. The authors also thank Joao Pedro Wagner De Azevedo, Shaohua Chen, Jose Cuesta, Carolina Diaz-Bonilla, Leonardo Lucchetti, Rinku Murgai, David Newhouse, and Minh Cong Nguyen for their suggestions on the draft and the poverty profile calculations.

globally) by 2030, and (2) to promote “shared prosperity” by boosting the incomes of the poorest 40 percent of the population in every country. This added focus on shared prosperity marks a major shift in the World Bank Group’s¹ mission, which in the late 1940s was to rebuild Europe, but evolved in the 1970s toward poverty reduction. Ending extreme poverty has remained the primary focus of the World Bank Group for the past few decades, but the introduction of the goal to promote shared prosperity introduces a second primary focus.

The inclusion of this second goal, to promote


Raj Chetty: On Increasing Opportunities and Improving Income Mobility

What keeps people at the bottom of the income ladder from climbing up and out of poverty?

Raj Chetty, an award-winning economist and Harvard professor, uses an innovative blend of empirics and theory, with help from big data, to produce work with important implications for global inequality, including showing which societal conditions allow people to more easily climb the income ladder.

His work on the importance of U.S. teachers—and how the quality of teaching at all levels of education can substantially improve the outcomes and opportunities for students—has been cited by President Barack Obama, among others.

In an interview with *Inequality in Focus*, Chetty discusses these topics and others and the feasibility of adapting them to a broader, international context. | **Page 7.**

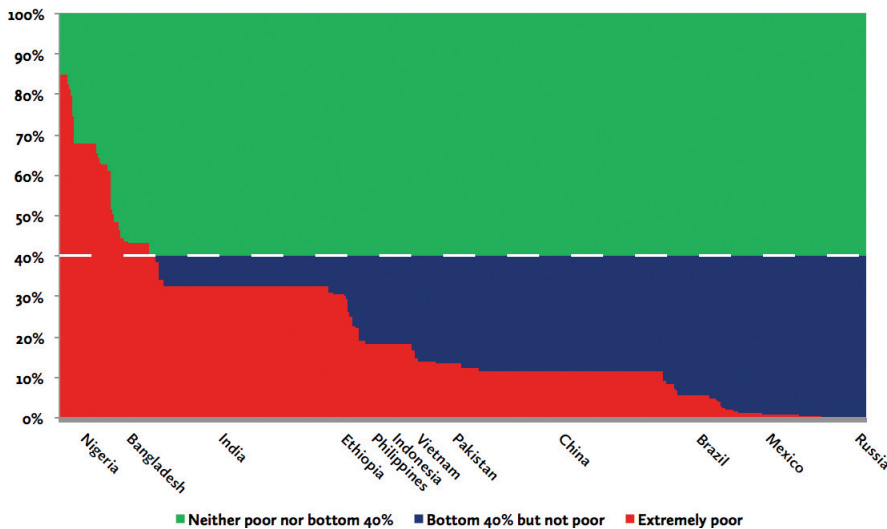


sidered a necessary condition for poverty reduction. The growth elasticity of poverty reduction—the percentage reduction in poverty rates associated with a percentage change in mean (per capita) income—is likely not high enough in many countries to achieve the goal of ending extreme poverty by 2030.² Focusing efforts to raise incomes for the lower part of the income distribution, which is what shared prosperity does by focusing on the bottom 40 percent, can increase the growth elasticity of poverty reduction. The implication, then, is that the shared prosperity goal is instrumental for reducing extreme poverty. However, to our knowledge, it has never been stated clearly whether promoting shared prosperity is a secondary tool to reduce extreme poverty, or if the two goals are prioritized equally. In fact, the discussion around extreme poverty and shared prosperity—and how these concepts should be interpreted and expected to interact with one another—seems to be far from settled.³

This raises some important questions. Could the shared prosperity goal compete with the extreme poverty goal over the allocation of development efforts and resources? Considering that the extremely

Figure 1

Distribution of the extreme poor, non-poor bottom 40 percent, and neither poor nor bottom 40 percent based on data from 2010. The area of the figure represents the total population of the developing world (5.9 billion people). Countries are listed along the horizontal axis in descending order from the highest extreme poverty rate on the left to the lowest on the right. A vertical bar represents each individual developing country used in this sample, with the width of the bar corresponding to the country's population, so, for example, India and China have the widest bars.



Source: PovcalNet and WDI.

poor are not distributed evenly across countries, how will the new mission alter the global development agenda? Within countries, what are the implications? Could this weaken policy makers' emphasis on ending extreme poverty? In other words, are there trade-offs between boosting the incomes of the bottom 40 percent in every developing country and ending extreme poverty globally? Our analysis suggests that there might indeed be trade-offs.

To better understand this, we must first answer the following two questions: (1) How much importance will be placed on each goal? (2) Are the extremely poor and the bottom 40 percent of income earners different? If so, how? If the two populations are identical, and the bottom 40 percent are also extremely poor, then the first question becomes irrelevant. In this case, boosting the incomes of the bottom 40 percent is equivalent to fighting poverty, and the two goals are well aligned. There is little or no trade-off.

However, if the two populations are different, it then becomes crucial to specify how much weight will be placed on each of these two goals. For instance, will gains in shared prosperity be considered a substitute for making progress toward ending extreme poverty? Or is promoting shared prosperity simply a secondary instrument to help reduce extreme poverty, which would remain the ultimate goal? These questions will need to be addressed if developing country governments and international donors are to efficiently allocate resources aimed at achieving these goals.

The first question—how much importance will be placed on each goal?—is impossible to answer with data. In this article, therefore, we will take a look at the second question—are the two populations indeed different? We do this by comparing the characteristics of the world's extremely poor to the bottom 40


percent of income earners in each developing country. We rely on three sources of data. The first is the PovcalNet database,⁴ an online poverty measurement tool maintained by the World Bank Group's Development Research Group. The second are the World Development Indicators,⁵ or the WDI, the World Bank Group's primary collection of development indicators. The third is the International Income Distribution Database or I2D2,⁶ a new dataset recently created by the World Bank Group. The I2D2 is a globally harmonized database that draws on more than 600 nationally representative household surveys. By examining these datasets we can create profiles of five relevant groups. These groups, visualized in figure 1, are:

1. The red group: The extremely poor (those who earn less than \$1.25 per day). As figure 1 shows, in some countries (to the left) this population is much larger than 40 percent and in other countries (to the right) this population is much smaller than 40 percent. Before the addition of the shared prosperity goal, this was the World Bank Group's main target population.

2. The bottom 40 percenters: The bottom 40 percent of income earners in each developing country, which can include both extremely poor individuals and those who are not extremely poor, depending on the country (these people are located below the dotted line in figure 1).

Figure 2

Composition of the extremely poor (red group), bottom 40 percenters, non-poor bottom 40 percent (blue group), and the new target population (red group plus blue group) in the world in 2010.



Legend: Low Income (purple), Lower Middle I. (w/o India) (orange), Upper Middle I. (w/o China) (green), China (dotted), India (hatched)

Source: PovcalNet and WDI.



Photo by Dorte Verner

In UMICs (Upper Middle Income Countries without China), the introduction of the shared prosperity goal could shift the World Bank Group's target population away from rural areas—and the agriculture sector—toward urban areas and likely the service and manufacturing sectors.

3. **The blue group:** These are individuals who are not extremely poor (earn more than \$1.25 per day) but fall within the bottom 40 percent. With the inclusion of shared prosperity this blue group represents people who were not previously a part of the main target population for the World Bank Group and its partners, but will be now.
4. **The new target population:** This group combines the extremely poor (the red group)—traditionally the World Bank Group's main target group, with the bottom 40 percent in every developing country who are not poor (the blue group)—the World Bank Group's additional target group. With the introduction of shared prosperity these two groups merge to become the World Bank Group's new target population.
5. **The green group:** This group encompasses the top 60 percent of income earners in each developing country who are not extremely poor. This does not include the individuals who belong to the top 60 percent of the income distribution but who are also extremely poor. These individuals are in red but above the 40 percent dotted line, toward the left side of figure 1.

In this article we use the term “extremely poor” to refer to those living with an income of less than \$1.25 per day as measured in 2005 Purchasing Power Parity (PPP) dollars. We also separate developing countries into one of the following five categories of country income classes: India, China, Low Income Countries (LICs), Lower Middle Income Countries without India (LMICs), and Upper Middle Income Countries without China (UMICs). The UMIC class also includes developing countries that have recently graduated to High Income status, such as Chile, Poland, Russia, and Uruguay. The LIC and LMIC income classes


correspond to the World Bank Group's latest income classification from July 1, 2013.⁷ In this classification, countries were defined as LIC if their per capita annual Gross National Income (per capita GNI) in 2012 was below \$1,035 per year. LMICs are countries in which the per capita GNI is between \$1,036 and \$4,085. UMICs, as defined here, are countries in which the per capita GNI is more than \$4,086.

Where in the world are the extremely poor (red group) and the bottom 40 percenters?

How will the introduction of the new shared prosperity goal shift the World Bank Group's target population across country income classes? Figure 2 shows the distribution by country income class of the world's (a) 1.2 billion extremely poor people (red group), (b) 2.4 billion bottom 40 percenters, (c) 1.3 billion bottom 40 percenters who are not extremely poor (blue group), and (d) 2.5 billion people of the “new target population” (red group plus blue group). As can be seen, the introduction of the shared prosperity goal shifts the distribution of the target population away from LICs and India toward UMICs and China. The weight of LMICs remains the same at 22 percent. We can also see that while 62 percent of the world's extremely poor (red group) live in LICs and India, only 11 percent of the blue group live in those countries. On the other hand, while only 17 percent of the red group live in UMICs and China, 67 percent of the blue group live in those countries. As a result, compared to when ending extreme poverty was the only goal, the new target population of 2.5 billion is distributed more or less evenly across the four top country income classes, each hosting approximately 20 percent, with the LICs hosting only 16 percent.

Figure 3


Share of agricultural workers among workers who are poor, non-poor bottom 40%, and neither poor nor bottom 40%, by country income class.



Source: I2D2.

Figure 4

Share of children among the poor, non-poor bottom 40%, and neither poor nor bottom 40%, by country income class.



Source: I2D2.

How important is agriculture for the extremely poor (red group) and the bottom 40 percenters who are not extremely poor (blue group)?


Figure 3 shows the importance of agriculture for the red group, the blue group, and the green group. As can be seen, agriculture is an important sector for the majority of the red group and the blue group in all country income classes, except for UMICs. In these countries, less than half of the red group, and slightly more than one-third of the blue group work in farming. More important is the difference between the red and blue groups within UMICs. Compared with the red group, approximately 30 percent fewer of the blue group work in agriculture. Therefore, in UMICs, the introduction of the shared prosperity goal shifts the target population away from rural areas—and the agriculture sector—toward urban areas and likely the service and manufacturing sectors.

Does promoting shared prosperity turn the focus away from children?

It is hard to say with certainty. What can be said is that children represent a smaller share of the additional target population (blue group) than they did of the old target population (the red group). Figure 4 shows that, for every country income class, there is a higher share of children under 13 in the red group than in the blue group. Therefore, the new target population that is created by introducing the shared prosperity goal is

Figure 5

School-aged children out of school among children who are poor, non-poor bottom 40%, and neither poor nor bottom 40%, by country income class.



Source: I2D2.

comparatively older than the old target population. Figure 5 shows the proportion of school-aged children (under 13) that are not in school. It can be seen that there are more children out of school from the red group than there are from the blue group, except in LICs. These results suggest that, at least in terms of enrollment, children's education is a more pressing issue among the extremely poor than it is for the blue group. Therefore, promoting shared prosperity may diminish the previous attention received by children from the international development community.

Do the bottom 40 percent who are not extremely poor (blue group) have better access to public utilities than the extremely poor (red group)?

Figure 6 shows the difference in access to electricity between the red group and the green/blue groups. In both LICs and LMICs, the red group has much less access to electricity than do the non-poor (blue and green groups). For UMICs and China, the two populations have similar access. Thus, for the new target population, access to electricity is less of an issue than it is for the extremely poor. This would likely hold true for other basic public utilities, such as access to drinking water and proper sanitation.

The political economy of shared prosperity: Could the political willingness to end poverty wane?

Looking again at figure 1, we see the distribution across countries of 1.2 billion extremely poor people (red group), 1.3 billion non-poor people in the bottom 40 percent (blue group), and those 3.4 billion people who are neither extremely poor nor of the bottom 40 percent (green group). We can see in the figure that the blue area is about the same size as the red area. This means that the introduction of the shared prosperity goal will (approximately) double the World Bank Group's target population. Also, with anticipated population growth and sustained poverty reduction in the developing world, it can be expected that the blue group will continue to grow in relation to the red group. Could this diminish the political will of governments to end extreme poverty? Will advocacy for the poor be weakened in countries where the poor are a small minority?



Photo by Dorte Verner

Children represent a large share of the world's extremely poor population.

The political willingness of governments to end extreme poverty is unlikely to be impacted by the new goal in low-income/high-poverty countries such as Nigeria and Bangladesh, where more than 40 percent of the population is extremely poor. In those countries and others like them, promoting shared prosperity may in fact refocus attention on the poorest segments of the population because the bottom 40 percenters are often poorer than the average poor person. In India, where approximately 33 percent of the population lives in extreme poverty, the additional focus on the bottom 40 percenters might slightly dilute the focus on the poor, but probably not by much.


The story is different for China and the UMICs such as Brazil, Mexico, and Russia. In those countries, the introduction of the shared prosperity goal substantially expands the World Bank Group's target population. Because the extremely poor are a very small, and shrinking, minority in these countries, a refocusing of the development dialogue away from ending poverty and toward boosting the income of the bottom 40 percent may weaken the political willingness to end extreme poverty. Only 4 percent of the world's extremely poor people reside in UMICs, but they still represent a large share of the population—approximately 45 million people. Would a greater focus on shared prosperity mean that those people can be more easily ignored?

Concluding remarks

For the first time in more than 30 years, the World Bank Group has enhanced its mission from ending extreme poverty to also raising the incomes of the bottom 40 percent in every developing country. Our analysis indicates that unless the World Bank Group

Figure 6

Access to electricity among population that is poor, non-poor bottom 40%, and neither poor nor bottom 40%, by country income class.



Source: I2D2.

and the international development community names ending extreme poverty as the higher priority, there might be a shift in focus away from LICs and India, where the majority of people in the red group currently live, toward UMICs and China, where the majority of people in the blue group currently reside.

In terms of strategy within country income classes, we find that in LICs and LMICs, the extremely poor (red group) and the bottom 40 percenters are very much the same people. Therefore, at least in the short run in those countries, the introduction of the shared prosperity goal will not be likely to shift attention away from the reduction of extreme poverty. In these countries, what is good for shared prosperity is also good for ending poverty and vice versa.

In UMICs, however, the World Bank Group and the interna-

tional community need to be cautious in their dialogues with client governments when they introduce the goal of shared prosperity. If there is a sense that gains in the income of the bottom 40 percent are as important as poverty reduction, then focusing on the extremely poor (and children) in these countries may become less important. The needs, location, and livelihoods of the red group can be very different from those of the blue group. As a consequence, by targeting the blue group in addition to solely the red group, the focus and energy spent on the red group might abate.

In order to accelerate the rate of poverty reduction to a level that ensures that the goal to end extreme poverty is achieved by 2030, the World Bank Group and international development community may need to be explicit in how they prioritize the two goals. If it is made clear that gains in shared prosperity will be celebrated only if we are on track to end extreme poverty by 2030, then the addition to the World Bank Group's target population of 1.3 billion people from the blue group, should not detract from the efforts to lift out of poverty the 1.2 billion people currently living on less than \$1.25 per day. However, if an increase in the incomes of the bottom 40 percent in every country is seen as a substitute for progress in reducing extreme poverty, then the introduction of the second goal might work against the world's poorest people.

Notes

1 The World Bank Group consists of five international organizations, including: the International Bank for Reconstruction and Development (IBRD), the International Finance Corporation (IFC), the International Development Association (IDA), the International Centre for Settlement of Investment Disputes (ICSID), and the Multilateral Investment Guarantee Agency (MIGA).

2 Yoshida, Nobuo, Hiroki Uematsu, and Carlos E. Sobrado (2014). *Is Extreme Poverty Going to End? An Analytical Framework to Evaluate Progress in Ending Extreme Poverty*, Number 6740, The World Bank Group, Washington, DC.

3 For example, in early discussions about the inclusion of the shared prosperity goal, there was some agreement that the two goals would be lexicographically ordered. This means that not only does extreme poverty have a much larger weight than share prosperity; but that no gains in shared prosperity should be allowed to compensate for a lack of progress in poverty reduction. This implies that the marginal dollar of development resources is applied only to shared prosperity once it is assured that the goal to reduce extreme poverty will be achieved.

4 The PovcalNet database can be accessed at <http://iresearch.worldbank.org/PovcalNet/index.htm>.

5 Population data used in figures 1 and 2 are obtained from WDI (<http://data.worldbank.org/data-catalog/world-development-indicators>) using `wbopendata` command in Stata (<http://data.worldbank.org/developers/apps/wbopendata>). Note that we use the latest population data available in WDI as of the time of data collection (March 1, 2014), which results in small discrepancies between our estimates and those in PovcalNet in terms of poor population and total population.

6 For this article and all I2D2-related figures we use data from a sample of countries from each income class, but not from all countries. For more details about the I2D2 dataset, which is not yet public, refer to Olinto, Pedro, Kathleen Beegle, Carlos Sobrado, and Hiroki Uematsu (2013). "The State of the Poor: Where Are The Poor, Where Is Extreme Poverty Harder to End, and What Is the Current Profile of the World's Poor?" World Bank Group—Economic Premise 125 (2013): 1–8. The table below summarizes the coverage in terms of number of countries, total population, and total poor population in the statistics presented in Figures 3 through 6 based on the I2D2. In each figure, we selected the latest available survey for all available countries in the I2D2, with a cut-off year of 2000. For example, we use a total of 64 surveys in the I2D2 to calculate the share of workers in agriculture in figure 3. These countries explain 83 percent of the total population and 89 percent of the total poor population as of 2010.

Coverage*	Figure 3	Figure 4	Figure 5	Figure 6	Total**
Poor population in millions	1,032	1,061	897	1,030	1,165
(% of total)	89%	91%	77%	88%	100%
Population in millions	4,653	4,838	3,390	4,507	5,574
(% of total)	83%	87%	61%	81%	100%
Number of countries	64	71	68	51	115
(% of total)	56%	62%	59%	44%	100%
Median survey year	2006.5	2006	2007	2007	2009

*Coverage is calculated assuming that each survey in I2D2 is nationally representative.

**Based on countries included in 2010 poverty estimates in PovcalNet for which the survey year is not earlier than 2000.

7 The complete list of countries according to the latest World Bank Group classifications is available at: <http://data.worldbank.org/about/country-classifications/a-short-history>.

The views expressed in this article are solely those of the authors and do not necessarily reflect the views or positions of the World Bank Group or peer reviewers. The authors are solely responsible for any errors or omissions.




Photo of Raj Chetty courtesy of the John D. & Catherine T. MacArthur Foundation.

An Interview with Professor Raj Chetty

Increasing Opportunities and Improving Income Mobility for the Poor

Raj Chetty, the Bloomberg Professor of Economics at Harvard University, is producing work with very strong practical and theoretical implications for global inequality, including the widely publicized Equality of Opportunity Project and research on the long-term impacts of teachers that was cited by U.S. President Barack Obama in his 2012 State of the Union address.

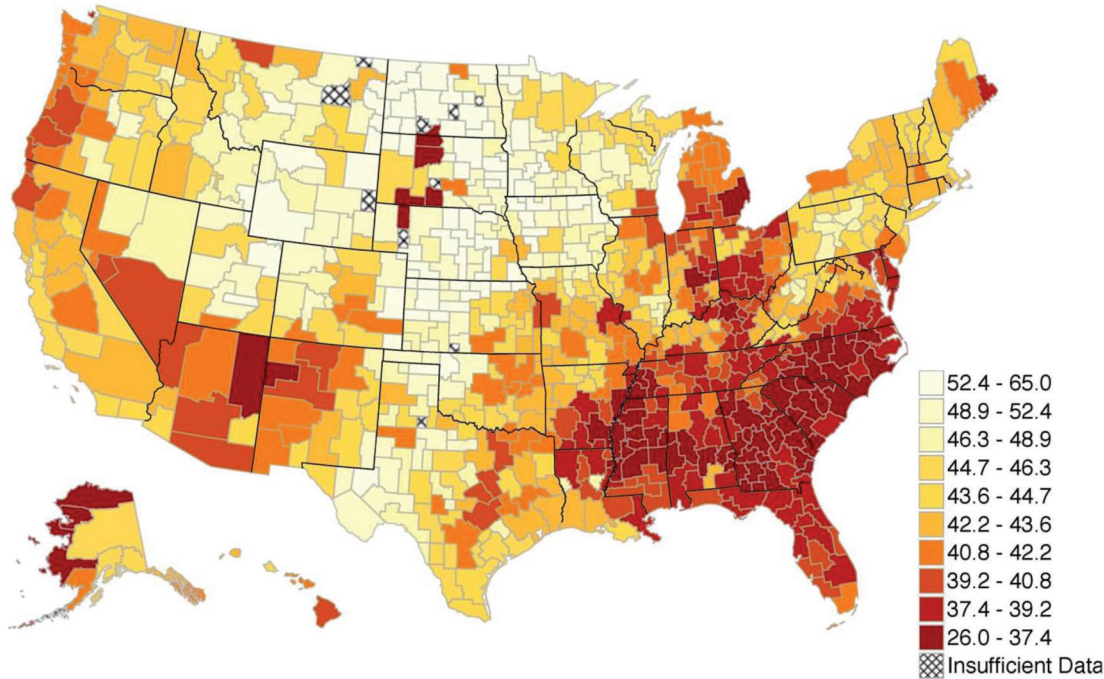
Members of the Inequality in Focus staff recently interviewed Chetty, who at 34 has already won two of the three most prestigious awards in economics, including the MacArthur Fellowship, better known as the “genius grant,” in 2012, and the John Bates Clark Medal in 2013. Earlier, the Economist magazine recognized him as one of the world’s eight brightest young economists. Chetty received his PhD from Harvard at 23, and at 29 became one of the youngest tenured professors in Harvard’s history. His research is an

innovative blend of empirics and theory, and he has often tapped into large administrative data sets, or big data, for his work.

Inequality in Focus: In your “Equality of Opportunity” research,¹ you find that upward income mobility (that is, the capability for those in lower income classes to move up the economic ladder) varies substantially within the United States [see figure 1], and the areas with the greatest upward income mobility tend to have five characteristics: less segregation, less income inequality, better schools, greater social capital, and more two-parent families. Interestingly, you show, first, that a larger middle class has a greater positive influence on upward mobility than does having a lower concentration of wealth in the top 1 percent of incomes. Second, that racially segregated areas have lower economic mobility for all races (whites as well as blacks or Hispanics) than non-

Figure 1

This map shows the average percentile rank of children who grow up in below-median income families across areas of the United States (absolute upward mobility). Lighter colors represent areas where children from low-income families are more likely to move up in the income distribution.



Source: www.equality-of-opportunity.org.

segregated areas. Third, that areas with more two-parent-households have higher upward mobility but having two parents is not the reason for this. Could you explain the causal mechanisms behind these three findings?

Raj Chetty: As you know, it is hard to figure out exactly what is causing what because all of these characteristics are correlated with many other omitted factors. But I can speculate to their causes based on some theories that have been discussed in prior literature.

One is the idea that the size of the middle class, and the amount of inequality, really affects people's prospects for upward mobility. The visual analogy that people sometimes give is to think about climbing the income ladder: If a lot of the rungs of the ladder are missing in the middle of the distribution, and the space between the lowest rungs of the bottom percentiles and the highest rungs of the top percentiles is very large, intuitively it is harder to climb up. This makes sense because there must be medium-skilled labor positions available for people with a high school education in order to give that family a prospect for upward mobility. However, if the only jobs available are those of high-skilled engineers and low-skilled service workers there is much less scope for moving up. That type of explanation would also explain, as you noted, why we find a much weaker correlation between extreme upper tail inequality and upward mobility. To take an example from California, San Jose and the Bay Area have a large middle class and a lot of tech billionaires. In these cases, the existence of wealthy tech billionaires does not seem to be holding everyone else back from moving up and there is still a high level of upward income mobility. So, the ladder

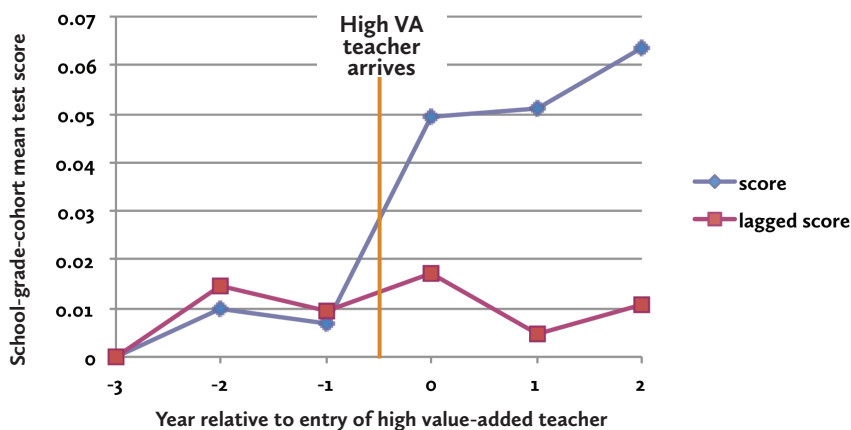
analogy maybe explains why the size of the middle class matters. I can't say we know for sure if that's the explanation, but it would be consistent with the data.

The second factor you discussed is segregation. There, I think, the sociologists have produced many intuitive theories that would explain why segregation would be correlated with mobility. One example is peer effects. If you're around other high skilled, or high-achieving peers, you will also have role models, social networks, and connections that may improve your aspirations and help you see what is feasible in your own life. When people are cut off from such peers, the community evolves very differently. Another possible mechanism for why segregation might matter is the funding of public goods; if we all live in an area where there are both high- and low-income people, the low-income people are going to benefit from the better schools and the property tax payments of the high-income families, but if your community is very separated you might not benefit from those public goods. Also, when we say "segregation" we mean not just segregation by race but also by income. So places that are racially segregated also tend to be more segregated by income, meaning the low- and high-income people are not living in proximity to each other. This would explain why segregation tends to negatively affect both white and black people, because regardless of race, poor people in segregated areas are less likely to see the public good or peer-effect benefits from a more integrated community.

Finally, a third comment on the two-parent-households. Theories in sociology suggest that family structure or the strength of the commu-

Figure 2

This figure shows that when a high value-added (top 5%) teacher enters a school, the end-of-school-year test scores in the grade he or she teaches rise immediately.



Source: http://obs.rc.fas.harvard.edu/chetty/value_added.html.

nity might matter. A key caveat here is that the fraction of two-parent households in an area matters even if your own parents are married, which is to say that having more single-parent families in your neighborhood is correlated with poor outcomes for kids even if their own parents are married. This suggests that it's not the direct mechanism of your own parents being married that matters but something about the community; the idea that more stable communities produce better outcomes for kids. This could also be consistent with other factors, for example, maybe places with more two-parent households tend to have better schools or lower teen birth rates. But again, this is hard for us to disentangle.

Inequality in Focus: Your research² shows that students assigned to higher-value-added teachers (a teacher's value-added rating is defined as the average test-score gain for his or her students) are more likely to attend college, earn higher salaries, live in better neighborhoods, save more for retirement, and are less likely to have children as teenagers [see figure 2]. You also show that replacing a teacher whose value-added rating is in the bottom 5 percent of the value-added distribution with one of average quality would generate cumulative earnings of \$80,000 per student, or more than \$1.4 million for the average classroom. But, your results also show that the impact from better teachers is greater for richer kids than for poorer kids. If so, wouldn't providing free, high-quality education to all increase inequality, since by subsidizing richer parents they can make complementary investments in their kids' education, via tutors, extracurricular activities and so on? Or, do you think that there are enough positive externalities to justify the public financing of education for rich kids?

Raj Chetty: First of all, I agree to some extent with the points you make about the complementarity of public and private education. In levels, that's exactly what we find in the data; the high-income kids, in dollar terms, actually benefit more from better teachers so it does look consistent with complementarity between inputs at home and inputs at school. So when you just improve the quality of schools for everyone, it's not necessarily true that you will reduce inequality. That is not to say

that it is not valuable to improve the quality of schools for all, because if we improve everyone's outcomes, and bring some people above the poverty line, it is, at least in my view, a good thing even if inequality is not reduced.

In terms of externalities, we have to remember that an important part of kids doing well—especially the more affluent—comes back to society as a benefit. This benefit can be in the form of tax dollars, important innovations, or job creation. We don't have direct evidence that precisely shows this, but it would go along with the idea that better teachers improve kids' outcomes through a variety of dimensions and will have externalities that benefit both high- and low-income families. To sum, I do think that there is a strong argument for improving the quality of education for all and providing public subsidies to increase the value across the distribution.

Inequality in Focus: But some of the externalities you mention, for example, innovation from high-income kids who receive a subsidy, may not be that common. Isn't it more likely that these more affluent kids will end up going to law school or medical school, in which case wouldn't these benefits then be internalized?

Raj Chetty: Benefits are certainly partly internalized but I don't think they are entirely internalized. As I was saying, in the United States, 30 to 40 percent of your income is going to taxes, so by definition that's a fiscal externality on the government that students would not be internalizing. But also there is a question, and actually something we are doing some work on now: To what extent are the returns to innovation captured by the individual as opposed to captured by other members of society? Certainly things like research are benefits that largely accrue to others. So you are right, for some things like law and medicine they would be internalized, but this isn't universally true.

Inequality in Focus: Currently, policy makers seem to be emphasizing the importance of early-childhood education while paying less attention to later education. This certainly seems to be the case at the World Bank. In your work, however, you show that

the quality of teachers has a significant impact on all grades and not just on the youngest children.

Raj Chetty: Let me first say that I think efforts to invest in early childhood education are valuable because I do think early childhood education matters. Still, I strongly believe, after having seen more and more evidence, that there are significant returns to improving the quality of education much later—well beyond ages 5, 10 or 15—to even when people are in college, and by no means should we be giving up on kids once they reach 5 or 8 years old. There is good evidence that you continue to improve kids' outcomes in the long run if you improve education at all age levels. This is extremely important when structuring policy because it means that, while investments in early childhood education are valuable, we should continue investing and trying to improve the quality of education at all points in time.

Inequality in Focus: Much of your work uses large administrative data sets or big data. For example, in The Equality of Opportunity Project, you compile statistics from millions of anonymous earnings records. So far, such big data has been instrumental in providing clarity to microeconomic questions, but because of limitations on what can be tested it provides less clarity on macroeconomic questions. Do you think there is a role for big data, first, in economics generally, and, second, in providing clarity to long-standing macroeconomic questions, like how to overcome a recession or manage interest rates?

Raj Chetty: I think big data is going to play a fundamental role in economics because it allows us to approximate experiments, which is the hallmark of success in science. By having enormous data sets we are basically able to find experiments in the data from which we can study, learn, and generalize. You already see this to a large extent in publications from the top economic journals today—70 or 80 percent of the empirical papers are using large administrative data sets. And the era of survey-based research, at least in developed countries, is ending; there are simply fewer and fewer people writing papers with those data sets. In developing countries, survey data continue to be very important because those are the best data we have, but even there with the advent of technology, like mobile phones and information technology, I think we'll start to have more large administrative data sources that will be useful.

The question of how well big data can answer macroeconomic questions is difficult to answer because the types of methods that we are getting good at developing—causal experimental methods—lend themselves well to cases where you can manipulate things at the micro level. Now, those micro level interventions can aggregate up to have big macro impacts. For instance, changes in social safety nets, or changes in unemployment benefit policies, do have important macroeconomic effects, but their nature is that they can be analyzed at the micro level. Things like a recession, interest rates, or the level of inflation, are harder to manipulate at the micro level, which then makes them harder to analyze. So, in my view, it's not a question of what important macroeconomic questions big data can help answer, it is more a question of what macro questions lend themselves to being studied with big data at a microeconomic level. My sense is that we are going to have within the next 10 to 15 years a much better understanding empirically and not just theoretically of how these things work. Thereby we will be able to build better macro models and have more agreement. But still, there are going to be other questions that are purely general-equilibrium

macro issues that are going to be harder to study. Maybe by building up from micro models we will have better macro models as well to tackle those questions, but I still think that it's going to be some time before we have a better empirical sense of those issues.

Inequality in Focus: A lot of your work has very strong implications for U.S. national policy but how applicable is it for other countries?

Raj Chetty: One of the approaches we try to take with these data sets is to extract lessons that are not only going to apply to the specific context we are studying, but will have broader applicability. For instance, take the issue we were discussing about teachers mattering not just at young ages, but at older ages. I intuitively think of that as a pattern that is likely to generalize. If it's true in the U.S., it is likely to be true in other developed countries, and perhaps in developing countries as well. I think there is more low-hanging fruit in developing countries, meaning that policy solutions are perhaps much simpler and more powerful than in the developed world and can have larger returns. For example, in the United States there is a complex debate on how to improve the quality of the teachers—if we should use value-added measurements of teaching or other methods, if we should train teachers, and so on. In many developing countries, the issue is much simpler. In India, for example, a problem is that teachers often do not show up for work, so the policy response is not a sophisticated intervention to improve the quality of teaching; it is simply getting teachers to show up. We demonstrate, using data from the United States, that teachers matter and it is likely that teachers matter everywhere.

While it is true that in poor countries the problems are bigger and inequality is greater because you are farther away from the frontier in terms of growth, it is also true that even if you can't close that huge gap, a 1 percent improvement in India, in a sense is much more important than a 1 percent improvement in the U.S., because there is so much more to be done.

I think, while recognizing the magnitude of the problem in developing countries, it is important to recognize that making even some progress in these countries is incredibly valuable given from where you are starting.

Inequality in Focus: In a recent presentation³ you gave at the World Bank, you questioned whether the United States can still be hailed as the "land of opportunity." You show evidence to the contrary by comparing mobility between the U.S. and Denmark. But Denmark has a much narrower income distribution than the United States, which makes it easier for someone in a low-income percentile to move to a higher-income percentile. Wouldn't this explain why you would likely see greater income mobility in Denmark?

Raj Chetty: Yes, absolutely. Let me restate what you just said because I agree with it: In Denmark it takes less to move across percentiles of the income distribution in absolute dollars because, as you said, the income distribution is more compressed. So going from, say, the 25th to 75th percentile is a smaller dollar movement in Denmark than it is in the United States. This is precisely why I think the comparison within the U.S., like what we focused on, is preferable to cross-country analyses that previous work has focused on. When we compare mobility across countries there are many different factors to take into account, which includes not only the income distribution, but also national levels of equality and the different institutions that exist in each country. So I think it is harder to learn across countries than it is from

within the United States where we are able to hold more things fixed. If you wanted to more accurately compare the two countries, or any others, you would need to look at global ranks. You could synthetically combine Denmark and the United States' income distribution and hold the ranks fixed. My sense is that you will find less of a difference in mobility between the two countries when you do that, but I think you will still find greater mobility within Denmark. I don't know exactly how it would break down.

Inequality in Focus: To conclude, we wanted to ask a more personal question. Your father was an economist and you accomplished and contributed a lot in the field of economics at a very young age, so we wanted to know, first of all, how your upbringing influenced your personal trajectory, and, secondly, did it influence some of your research ideas?

Raj Chetty: One of the great things about this research is seeing how it fits with my own introspection. Both of my parents are academics; my dad is an economist, as you mention, and my mom is a physician who does research. I learned a lot from being around that environment and it also made me interested in research. So I think family played a huge role for me, but so did my teachers. In high school, for instance, I remember having a teacher who said to me “you are going to be a great economist someday” and I remember at that point really wanting to do biomedical science. I wasn't thinking about doing economics at all, and I thought “oh well, there is no way I am going to do that.” Sometimes it seems like your teachers know you better than you know yourself. Besides teachers, I think environment also plays a huge role and luck to some extent. I am very lucky to have had a great background and a lot of people investing in my coming up.

Although I obviously had no sense of how test-score-based, value-added measures were going to capture things in my research or what patterns would emerge in the data, my background certainly contributed to my interest in studying these questions.

To give another anecdote, I have many cousins—who actually have the same name as me; we were all named Nadarajan after my grandfather—and it so happens that, while they are all living fulfilling lives, we had very different outcomes. One cousin is doing difficult labor in a temple in Singapore; others are working in manufacturing jobs. Very different outcomes.

I trace one of the reasons for this back one generation to our parents. It so happens that my mom was the one person in her family who received a higher education and ended up going to medical school. Same thing with my dad. At that time in India, my grandparents did not have money to invest in the educations of all of their kids. You see the lasting impacts of those differences among the grandchildren and even the next generation. To me, this illustrates the importance of opportunities. My cousins and I were quite similar to begin with. The fact that we ended up in really different places further motivates me to understand why this happens and what we can do to give all children an equal opportunity to succeed.

This interview took place on March 7, 2014, between Professor Raj Chetty and the Inequality in Focus editorial team: Pedro Olinto, Maximillian Ashwill, Julie Barbet-Gros, and Fernanda Luchine. The interview was recorded and transcribed by Fernanda Luchine. The above dialogue is an interpretation of the conversation written by Maximillian Ashwill and Pedro Olinto.

Notes

1 The Equality of Opportunity Project can be found online at www.equality-of-opportunity.org.

2 This research, “The Long-Term Impact of Teachers: Teacher Value-Added and Student Outcomes in Adulthood,” can be found online at http://obs.rc.fas.harvard.edu/chetty/value_added.html.

3 This presentation can be viewed online at <http://live.worldbank.org/improving-equality-opportunity>.

The *Inequality in Focus* series aims at informing the public debate on equity, inequality of opportunity, and socioeconomic mobility. It features articles written by World Bank Group staff, as well as researchers and policy makers from the broad development community. The views and interpretations in the articles are those of the authors and do not necessarily represent the views of the World Bank Group, its Executive Directors, or the countries they represent.

The *Inequality in Focus* series is not copyrighted and may be reproduced with appropriate source attribution.

Executive Editor: Pedro Olinto
Managing Editor and Staff Writer: Maximillian Ashwill
Copy Editing and Design: Mary Anne Mulligan
Research Staff: Fernanda Luchine Ishihara
and Julie Barbet-Gros
Production: Anna Reva
Web Designer: Maura K. Leary

Sector Manager: Christina Malmberg Calvo



Poverty Reduction and Equity Department
Poverty Reduction and Economic Management Network (PREM)