Presidential Address

Spreading the Wealth Around: Reflections Inspired by Joe the Plumber

N. Gregory Mankiw

Department of Economics, 223 Littauer Center, Harvard University, Cambridge, MA 02138, USA.

E-mail: ngmankiw@fas.harvard.edu

This essay discusses the policy debate concerning optimal taxation and the distribution of income. It begins with a brief overview of trends in income inequality, the leading hypothesis to explain these trends, and the distribution of the tax burden. It then considers the normative question of how the tax system should be designed. The conventional utilitarian framework is found to be wanting, as it leads to prescriptions that conflict with many individuals' moral intuitions. The essay then explores an alternative normative framework, dubbed the Just Deserts Theory, according to which an individual's compensation should reflect his or her social contribution.

Eastern Economic Journal (2010) 36, 285–298. doi:10.1057/eej.2010.22

Keywords: income distribution; utilitarianism; optimal taxation

JEL: H2; D3

One of my favorite recent moments in political theater was when “Joe the Plumber” posed a question to candidate Barack Obama during the presidential campaign of 2008.¹ As you may recall, Joe was an aspiring small business owner, and he asked then-Senator Obama about his proposal to raise taxes on high-income households. The candidate responded, in part, “It’s not that I want to punish your success. I just want to make sure that everybody who is behind you, that they’ve got a chance at success, too ... I think when you spread the wealth around, it’s good for everybody.”

The reason I like this particular moment is that it focused public attention on one of the defining differences among competing economic philosophies. Indeed, I don’t think it is an exaggeration to say that the single most important difference between the political left and the political right is over the questions of whether, and to what extent, “spreading the wealth around” is a proper function of government.

Looking ahead, I fully expect the issue to remain at the center of political debate. One reason is that the tax cuts signed into law by President Bush in 2001 and 2003 will expire next year unless Congress takes action to extend them.

Another, perhaps more important, reason is that the US federal government is running a large budget deficit and faces an ominous fiscal gap looming on the horizon. As the baby boom generation retires and starts claiming Social Security and Medicare, government spending will slowly and steadily continue to rise as a share of the economy. It is possible that Congress will suddenly read Milton Friedman’s book *Capitalism and Freedom*, become committed classical liberals (in the 19th century use the term), and decide to scale back the size and scope of government. But, more likely, Congress will find past entitlement promises hard to break, and so it will have little choice but to raise taxes to levels unprecedented in US history, which naturally raises the question: *Whose* taxes should go up?

I should say at the outset that the issues I will discuss with you here involve not only economics but also some political philosophy. Because I am not a political

philosopher by training, I hope you will forgive me if my occasional philosophical ruminations seem like those of an amateur. If I am right that the issue of redistributive justice will be at the heart of the coming policy debate, it will be hard to leave the topic to the philosophical experts. And in light of the inextricable linkages between philosophy and economics that characterize this topic, I hope it is possible that those experts might learn something from humble economists like me.

FACTS ABOUT THE INCOME DISTRIBUTION

A good place to start in thinking about this issue (and many others) is with the facts. One reason that President Obama and Joe the Plumber were focused on “spreading the wealth” is that the distribution of income has changed dramatically over the past few decades. So let’s start our discussion not with political philosophy but on the firmer ground of economic statistics.

Figure 1 is taken from my Harvard colleagues Claudia Goldin and Larry Katz [2007]. The figure uses census data for two time periods to show the growth in income for the five income quintiles, as well as for the top 5 percent of the population. The result is striking.

For the early period, from 1947 to 1973, the bars all have about the same height. That is, all income groups experienced growth of about 2–3 percent per year. There is, however, some variation. Because the bottom quintile had the highest average growth rate and the top quintile had the lowest, the overall income distribution became more compressed.

After 1973, however, this trend reversed itself, as the average growth rates for low-income groups were markedly below those for higher income groups. For the top 5 percent, the average growth rate continued at about 2 percent per year. For the lowest quintile, the average growth rate has been approximately zero. This fact explains why US poverty rates, after a long period of decline, have remained largely unchanged over the past couple of decades.

Another source that documents the recent widening in the income distribution is data from tax returns, as has been extensively studied by Thomas Piketty and

Figure 1. Growth in real mean family income by income quintile.

Source: Goldin and Katz [2007].

Figure 2. Share of income going to the top 1 percent.
Source: Piketty and Saez [2003] and web updates.

Figure 3. Share of income going to the top 0.01 percent.
Source: Piketty and Saez [2003] and web updates.

Emmanuel Saez [2003]. I should point out that tax return data are highly imperfect for this purpose, as changes in the tax code over time can alter the incentive to receive and report income in any particular form. For example, the tax law influences the choice many small businesses make about whether to organize as a C-corporation or as an S-corporation, and this decision in turn influences where their income appears in tax data. Nonetheless, despite this problem, the tax data may be our best lens into the lives of the truly rich.

Figure 2 shows the percentage of total income received by the top 1 percent of the income distribution. To get into this group today, your household has to make more than \$400,000 a year. As the figure illustrates, the share of income accruing to this group has more than doubled since the mid-1970s.

Figure 3 shows the income share of an even more elite group — the top 1 percent of the top 1 percent. To get into this group, your family needs annual income of more than \$11 million. Since the mid-1970s, this group's share of total income has increased about sixfold.

So those are the facts about the changing income distribution. While one can argue to some degree about their reliability, they are probably the least controvertible things I am going to address in this essay.

WHY IS INEQUALITY GROWING?

The natural next question to address, at least for an economist, is why such changes in the income distribution have been occurring.

The best diagnosis so far comes from Goldin and Katz in their recent book *The Race Between Education and Technology*. Their bottom line is that “the sharp rise in inequality was largely due to an educational slowdown.”

According to Goldin and Katz, for the past century technological progress has been a steady force not only increasing average living standards, but also increasing the demand for skilled workers relative to unskilled workers. Skilled workers are needed to apply and manage new technologies, while less skilled workers are more likely to become obsolete.

For much of the 20th century, however, skill-biased technological change was outpaced by advances in educational attainment. In other words, while technological progress increased the demand for skilled workers, our educational system increased the supply of them even faster. As a result, skilled workers did not benefit disproportionately from economic growth.

But recently things have changed. Over the last several decades, technological advance has kept up its pace, while educational advancement has slowed down. The cohort of workers born in 1950 averaged 4.67 more years of schooling than the cohort born in 1900, representing an increase of 0.93 years of schooling in each decade. By contrast, the cohort born in 1975 had only 0.74 more years of schooling than that born in 1950, an increase of only 0.30 years per decade. That is, the pace of educational advance has fallen by 68 percent.

Because growth in the supply of skilled workers has slowed, their wages have grown relative to those of the unskilled. This is evident in Goldin and Katz’s estimates of the financial return to education. In 1980, each year of college raised a person’s wage by 7.6 percent. In 2005, each year of college yielded an additional 12.9 percent. Over this time period, the rate of return from each year of graduate school has risen even more — from 7.3 to 14.2 percent.

While education is the key to understanding broad inequality trends, it is less obvious whether it can help explain the incomes of the superrich. Simply going to college and graduate school is hardly enough to join the top echelons making millions a year.

But neither is education irrelevant. We don’t know much about the demographic characteristics of the superrich, but it is a good bet that they are on average highly educated. A good education is not a guarantee of great riches, but for many highly paid career paths it may be a prerequisite.

Perhaps, advanced degrees are like Willie Wonka’s famous chocolate bars. A few of them come with golden tickets that give you opportunities almost beyond imagination. Over the past several decades, as the return to education has increased, the value of those golden tickets has increased as well. But even if you aren’t lucky enough to get a golden ticket, you can still enjoy the chocolate, which by itself is well worth the price.

Table 1 Federal taxes as a percentage of income

	<i>Average pretax income (\$)</i>	<i>Effective federal tax rate (%)</i>
Lowest quintile	17,200	4.3
Second quintile	39,400	10.3
Middle quintile	60,700	14.2
Fourth quintile	89,500	17.6
Highest quintile	248,400	25.8
Top 1 percent	1,743,700	31.2

Source: Congressional Budget Office. Data are for 2006.

FACTS ABOUT THE TAX BURDEN

All the numbers I have shown you so far describe before-tax incomes. But these facts are clearly related to the public debate over tax policy. If policymakers want to change the distribution of economic well-being, the federal system of taxes and transfers is the most direct and arguably most powerful tool they have.

There is no doubt that these facts are politically charged. The question, “Do the rich pay their fair share in taxes?” is one of defining issues of our time. One way we can start thinking about this question is by looking at how much people in different parts of the income distribution pay.

During the presidential campaign of 2008, at a fund-raiser for Hillary Clinton, the billionaire investor Warren E. Buffett said that rich guys like him weren’t paying enough. Mr. Buffett asserted that his taxes in the previous year equaled only 17.7 percent of his taxable income. By contrast, he said that his receptionist paid about 30 percent of her income in taxes.

Mr. Buffett was echoing a refrain that is popular in some circles. A few years ago, Robert B. Reich, labor secretary during the Clinton administration, wrote on his blog that “middle-income workers are now paying a larger share of their incomes than people at or near the top We have turned the principle of a graduated, progressive tax on its head.”

These claims are enough to get populist juices flowing — which perhaps was the motivation for them. The problem is that these assertions don’t hold up under close examination.

The best source for objective data on the distribution of the tax burden is the Congressional Budget Office (C.B.O.). The C.B.O. goes beyond anecdotes and bald assertions to provide hard data on who pays taxes. One can argue about some aspects of its methodology, but there is no doubt that it is non-partisan and that its tax analysts are some of the best in the business.

The C.B.O.’s calculations of federal tax rates, displayed in Table 1, show a highly progressive system. (The numbers are based on 2006 data, the most recent year available, but the tax code has not changed much since then.) The poorest fifth of the population, with average annual income of \$17,200, pays only 4.3 percent of its income in federal taxes. The middle fifth, with income of \$60,700, pays 14.2 percent. And the top fifth, with income of \$248,400, pays 25.8 percent.

For the very top of the income distribution, the C.B.O. reports even higher tax rates. The richest 1 percent, with an average income of \$1,743,700, forks over 31.2 percent of its income to the federal government.

One might wonder how Mr. Buffett gets away with a tax rate of only 17.7 percent, while a typical millionaire is paying so much more. Most likely, part of the answer is that Mr. Buffett's income is made up largely of dividends and capital gains, which are taxed at only 15 percent. By contrast, many other top earners pay the maximum ordinary income tax rate of 35 percent on their salaries, bonuses, and business income.

The distinction is crucial for understanding how much the rich pay. Indeed, the share of top incomes coming from capital is much lower now than it has been historically. According to the Piketty and Saez data, for the very richest Americans — those in the top 0.01 percent of the distribution — the percentage of income derived from capital fell from 71 percent in 1929 to 33 percent in 2007. If your image of the typical rich person is someone who collects interest and dividend checks and spends long afternoons relaxing on his yacht, you are decades out of date. The leisure class has been replaced by the working rich.

Another piece of the puzzle is that Mr. Buffett's tax burden is larger than it first appears, because he is a major shareholder in Berkshire Hathaway. When the C.B.O. studies the tax burden, it includes all federal taxes, including individual income taxes, payroll taxes, and corporate income taxes. In its analysis, payroll taxes are borne by workers, and corporate taxes by the owners of capital. For the richest 1 percent of the population, 10.4 percentage points of their 31.2 percent tax rate comes from the taxes that corporations have paid on their behalf. The corporate tax would undoubtedly loom large if the C.B.O. were to calculate Mr. Buffett's effective tax rate.

So it is simply wrong to say we don't have a progressive tax system. The best analysis shows that average federal tax rates rise steeply with income.

None of these calculations, however, say whether the rich are paying their fair share. At the 2000 Republican National Convention, candidate George W. Bush stated clearly his view: "On principle, no one in America should have to pay more than a third of their income to the federal government." As judged by the C.B.O. data, he accomplished his goal. Eight years later, during the campaign of 2008, candidate Barack Obama made clear that he thought the rich weren't paying enough. And if his proposed policies become law in the coming years, the richest Americans will end up paying over a third of their income in federal taxes.

What I want to do now is turn to the question of how we might begin to evaluate these alternative points of view. So far, most of what I have said here has safely fallen within the realm of positive economics. I now want to take a bold leap into normative analysis.

AGNOSTICISM AND UTILITARIANISM

Here is the question I would like to consider: What can economists contribute to the policy debate over inequality and the distribution of the tax burden?

One plausible answer is, not much. It is tempting to take the position that the optimal redistribution of income is really a question only for political philosophers. We economists can try to estimate the cost of redistribution — that is, the negative impact on efficiency that comes with attempts to achieve more equality. But in the end, picking the best point on the tradeoff between efficiency and equality comes from policy preferences about which we, as economists, must be agnostic.

The academic literature on optimal taxation, however, takes a different tack. Rather than remaining agnostic, work in this area adopts a particular philosophical

perspective — utilitarianism [see, e.g., Mirrlees 1971]. That is, it is assumed not only that individuals maximize their own utility (the standard assumption in micro-economics), but that society should choose a tax-and-transfer policy to maximize the sum of utility over all individuals (or perhaps a more complicated, nonlinear social welfare function of individual utilities).

For economists, this approach is quite natural. One reason is that utilitarians and economists share an intellectual tradition. The early utilitarians, such as John Stuart Mill, were also among the early economists. Another reason is that utilitarianism seems to be a simple and natural extension of the economist's model of individual decision making to the societal level. Moreover, utilitarian political philosophy allows economists to use the tools of analysis we find useful in many other contexts. Running a society becomes a problem of constrained optimization.

Nonetheless, the utilitarian approach is fraught with several problems. One classic problem is the interpersonal comparability of utility. We can infer an individual's utility function from his choices. But from this revealed-preference perspective, utility is not inherently measurable, and it is impossible to compare one person's utility with another person's.

Philosophers entertain other objections to utilitarianism. Undergraduate philosophy courses are filled with hypothetical examples to test the limits of this perspective. In a classic example, you are standing on a bridge over some trolley tracks, and you see an out-of-control trolley car hurtling toward three innocent children. There is, however, a fat man standing next to you, and you can save the three children by pushing the fat man off the bridge and onto the tracks below. (You are assumed to be too thin to save the children by throwing yourself off the bridge.) The question is, would you kill the fat man to save the children? A utilitarian would have no trouble doing so, but many people are left uncomfortable with that conclusion.

These kinds of examples can be greatly entertaining, but I am skeptical of their practical usefulness. Testing the implications of theories against our moral intuitions makes sense, but I am wary of hypothetical examples so far outside of any actual experience. Maybe we resist pushing the fat man not because we reject utilitarianism but because the facts of the case are too outlandish to take seriously. I for sure have never found myself standing on a bridge with the certain knowledge that I could sacrifice a fat man and save three innocent children. To evaluate whether we really believe utilitarianism, it might be better to consider actions or policies that we might actually be in a position to implement.

UTILITARIANISM ON A GLOBAL SCALE

Let me propose a somewhat less graphic hypothetical. Imagine a candidate for president campaigned on a platform of imposing a one-third tax on the average American's income and transferring the entire proceeds of the tax to poor nations around the world. Would you be inclined to support this candidate?

I am confident that most voters would not. I say this because I know of no political candidate who has proposed something even remotely like this. Moreover, the foreign aid that the United States does provide to the world's poor is far smaller than this and, even so, tends to be wildly unpopular. Even proponents of increased foreign aid, such as Jeff Sachs, would not have the temerity to suggest sending a third of our GDP abroad.

But if we were truly utilitarian in our policy design, why not? According to utilitarian logic, we impose tax rates of one-third or more on the residents of Palm Beach, Florida, and Greenwich, Connecticut, because they are richer and therefore have lower marginal utility than the average American. By the same logic, however, the average American is vastly richer than the average citizen of the world and should, therefore, have lower marginal utility.

The logic of utilitarianism as a theory of justice provides no reason to give a special role to national boundaries. If this political philosophy gives the moral authority to the federal government to tax rich Americans for the benefit of poorer Americans, it should authorize the United Nations to impose similar tax rates on rich nations such as the United States for the benefit of poorer nations.

My guess, however, is that most supporters of progressive taxation within the United States would not readily embrace a similar global system. Which raises the question, why not?

THE OPTIMAL TAXATION OF HEIGHT

Let me give you another example of how utilitarianism leads to policy conclusions about which many people are uncomfortable. This one is based on recent work of mine (Mankiw and Weinzierl [2010]).

More than a century ago, Francis Y. Edgeworth [1897] pointed out that a utilitarian social planner with full information will be completely egalitarian. More specifically, the planner will equalize the marginal utility of all members of society; if everyone has the same separable preferences, equalizing marginal utility requires equalizing after-tax incomes as well. Those endowed with greater than average productivity are fully taxed on the excess, and those endowed with lower than average productivity are subsidized to bring them up to the average.

William S. Vickrey [1945] and James A. Mirrlees [1971] emphasized a key practical difficulty with Edgeworth's solution. The government cannot observe innate productivity. Instead, it observes income, which is a function of productivity and effort. The social planner with such imperfect information has to limit his utilitarian desire for the egalitarian outcome, recognizing that too much redistribution will blunt incentives to supply effort. The Vickrey–Mirrlees approach to optimal nonlinear taxation is now standard in the academic literature on public finance.

Vickrey and Mirrlees assumed that income was the only piece of data the government could observe about an individual. That assumption, however, is far from true. In practice, a person's income tax liability is a function of many variables beyond income, such as mortgage interest payments, charitable contributions, health expenditures, number of children, and so on. George Akerlof [1978] called these variables "tags" and suggested that they might be used to identify individuals whom society deems worthy of special support.

In a recent paper, Weinzierl and I used the Vickrey–Mirrlees utilitarian framework to explore the potential role of another variable — the taxpayer's height. This inquiry is supported by two legs, one theoretical and one empirical.

The theoretical leg is that, according to the theory of optimal taxation, any exogenous variable correlated with productivity should be a useful indicator for determining an individual's optimal tax liability. Intuitively, such a variable allows society to tax its more productive members, on average, without incurring the efficiency costs that come with taxing income.

The empirical leg is that a person's height is strongly correlated with his or her income. For example, Anne Case and Christina Paxson [2008] report that "for both men and women ... an additional inch of height [is] associated with a one to two percent increase in earnings." This fact, together with the canonical approach to optimal taxation, suggests that a person's tax liability should be a function of his height. That is, a tall person of a given income should pay more in taxes than a short person of the same income.

My recent research with Weinzierl shows that, according to a conventional utilitarian calculus, the optimal height levy is sizeable. We calculate optimal taxes for adult white males in the United States, whom we divide into three height groups — tall (above 72 inches), medium (between 70 and 72 inches), and short (below 70 inches). The optimal average tax on the tall is about 7.1 percent of average tall income, while the average tax on the medium is about 3.8 percent of average medium income. These taxes pay for an average transfer to the short of more than 13 percent of average short income. Expressed in a more tangible way, Table 2 shows the optimal tax schedule by height group. A tall person making \$50,000 should pay about \$4,500 more in taxes than a short person making the same income.

Many people, however, will not quickly embrace the idea of levying higher taxes on tall taxpayers. Indeed, when first hearing the proposal, most people either recoil or express amusement. That reaction is precisely what makes the policy so intriguing. A tax on height follows inexorably from the standard utilitarian approach to the optimal design of tax policy coupled with a well-established empirical regularity. If we reject the conclusion, we must reconsider the assumptions.

JUST DESERTS

Unlike pushing the fat man off the bridge, the two policies I just considered are ones that we could actually implement if we wanted to. The United States could turn over a third of our GDP to the world's poorer nations, and the US Congress could pass a sizeable tax based on height. My guess, however, is that relatively few people would endorse these two policies with enthusiasm. There are many possible reasons why this might be the case, but what I would like to propose is that our moral intuitions are not, fundamentally, utilitarian.

Among philosophers, there is a long tradition of rejecting utilitarianism and similar approaches to distributive justice. Robert Nozick wrote the following in his famous 1974 book *Anarchy, State, and Utopia*:

We are not in the position of children who have been given portions of pie by someone who now makes last minute adjustments to rectify careless cutting. There is no *central* distribution, no person or group entitled to control all the resources, jointly deciding how they are to be doled out. What each person gets, he gets from others who give to him in exchange for something, or as a gift. In a free society, diverse persons control different resources, and new holdings arise out of the voluntary exchanges and actions of persons.

Despite the prominence of Nozick in general and this book in particular, it is hard to see much influence of this philosophical perspective in recent analytic work among economists.

Table 2 Optimal tax schedules by height

<i>If your taxable income is closest to ...</i>				<i>If your taxable income is closest to ...</i>			
<i>And you are —</i>				<i>And you are —</i>			
<i>Short</i>	<i>Medium</i>	<i>Tall</i>		<i>Short</i>	<i>Medium</i>	<i>Tall</i>	
69 inches or less	70–72 inches	73 inches or more		69 inches or less	70–72 inches	73 inches or more	
<i>Your tax is —</i>				<i>Your tax is —</i>			
5,000	–22,697	–20,546	–20,137	105,000	33,947	36,919	38,280
10,000	–19,136	–16,741	–16,391	110,000	36,859	39,704	41,406
15,000	–16,107	–13,488	–13,062	115,000	39,771	42,488	44,532
20,000	–13,248	–10,413	–9,962	120,000	42,682	45,273	47,658
25,000	–10,581	–7,563	–7,061	125,000	45,594	48,058	50,784
30,000	–7,992	–4,882	–4,319	130,000	48,506	50,843	53,559
35,000	–5,549	–2,274	–1,671	135,000	51,289	53,628	55,930
40,000	–3,201	327	860	140,000	53,290	56,244	58,300
45,000	–882	2,920	3,420	145,000	55,291	58,344	60,671
50,000	1,411	5,444	5,976	150,000	57,292	60,444	63,041
55,000	3,599	7,746	8,368	155,000	59,204	62,481	65,412
60,000	5,810	10,044	10,788	160,000	60,694	64,500	67,615
65,000	8,867	12,350	13,766	165,000	62,184	66,519	69,658
70,000	11,931	14,828	16,744	170,000	63,674	68,538	71,701
75,000	15,264	18,151	19,722	175,000	65,163	70,556	73,743
80,000	18,622	21,506	22,715	180,000	66,653	72,575	75,778
85,000	21,979	24,861	25,819	185,000	68,143	74,594	77,722
90,000	25,211	28,216	28,922	190,000	n/a	76,613	79,665
95,000	28,123	31,349	32,028	195,000	n/a	78,632	81,609
100,000	31,035	34,134	35,154	200,000	n/a	80,651	83,552

Source: Optimal taxes for short, high-income taxpayers are not available because of data limitations. Mankiw and Weinzierl [2010].

Perhaps the reason is that a viable alternative to utilitarianism is far from obvious. That is, if we reject utilitarianism as the basis for optimal tax policy, what can economists and other policy analysts put in its place as a normative framework?

Let me propose the following principle: *People should get what they deserve.* A person who contributes more to society deserves a higher income that reflects those greater contributions. Society permits him that higher income not just to incentivize him, as it does according to utilitarian theory, but because that income is rightfully his. This perspective is, I believe, what Robert Nozick, Milton Friedman, and other classically liberal writers have in mind. We might call it the Just Deserts Theory.

I am drawn to this approach in part by reflecting on some of the public anger that we see over some very high incomes. My sense is that people are rarely outraged when high incomes go to those who obviously earned them. When we see Steven Spielberg make blockbuster movies, Steve Jobs introduce the iPod, David Letterman crack funny jokes, and J.K. Rowling excite countless young readers with her Harry Potter books, we don't object to the many millions of dollars they earn in the process. The high incomes that generate anger are those that come from manipulating the system. The CEO who pads the corporate board with his cronies and the banker whose firm survives only by virtue of a government bailout do not seem to deserve their multimillion dollar bonuses. The public perceives them (correctly or incorrectly) as getting more than they contributed to society. That is, if we take public attitudes as a gauge of our innate moral intuitions, then in evaluating distributive justice, we should focus not on the marginal utility of different individuals but on the congruence between their contributions and their compensation.

One implication of the Just Deserts Theory is that it gives a new normative interpretation of the equilibrium of a competitive market economy. Under a standard set of assumptions, a competitive economy leads to an efficient allocation of resources. But we economists often say that there is nothing particularly equitable about that equilibrium. Perhaps we are too hasty in reaching that judgment. After all, it is also a standard result that in a competitive equilibrium, the factors of production are paid the value of their marginal product. That is, each person's income reflects the value of what he contributed to society's production of goods and services. One might easily conclude that, under these idealized conditions, each person receives his just deserts.

Another way to view this claim of the inherent equity of the competitive market equilibrium is to consider an exchange economy. Suppose that nature endows individuals with various combinations of goods (which we can view as time and talent), and these individuals engage in voluntary and mutually advantageous trade with one another. From the perspective of classical liberalism, it is natural to presume that any individual, or group of individuals, should be allowed to leave the large society to live on their own and form smaller communities. They exercise this right if they feel their contributions are insufficiently rewarded — that is, if they can do better on their own. This freedom ensures that the resulting allocation of resources will be in what game theorists call the core. This constraint is significant: Gerard Debreu and Herbert Scarf [1963] proved that, as the number of players gets large, the core of such games converges to the competitive equilibria. Thus, if the freedom to exit a society is taken as axiomatic, then the only permissible allocations of resources are the competitive market equilibria. For any other allocation, some group will exercise their right to leave because they are not getting their just deserts.

OPTIMAL TAXATION REVISITED

Let me now turn to what this approach suggests about the distribution of the tax burden. One might be tempted to conclude that if people are earning their just deserts, there is no room for a progressive system of taxes and transfers. One might think it would point us toward a system without government, or perhaps lump-sum taxes to finance minimal government, which would certainly be a radical departure from current practice. But I don't think that is necessarily correct. There are various ways in which the real world differs from a classical competitive economy free of market imperfections, and these real-world elements move the Just Deserts policy toward what we observe.

First of all, Pigovian taxation fits naturally into the Just Deserts Theory. If a person earns income in part by imposing negative externalities on bystanders, there is little doubt that he is getting more than he contributed. A corrective tax not only moves us toward economic efficiency, but it also helps align individual compensation with social contribution. I have long advocated increasing the use of Pigovian taxation [Mankiw 2009].

Similarly, Pigovian subsidies are a logical implication of the Just Deserts Theory. If some activities, such as the basic research of scientists, exhibit positive externalities, the government should help support them. A corrective subsidy, financed by taxes on the beneficiaries of the externality, not only enhances efficiency but also brings individual compensation closer to social contribution.

Next, there is the issue of public goods, such as national defense, police, and the court system. These expenditures have to be funded somehow, and if people are to get what they deserve, these public goods should be paid for by those people who benefit most from them. In their classic work on the financing of public goods, Knut Wicksell and Erik Lindahl made much the same suggestion.

Public goods and Pigovian subsidies lead naturally to a tax system in which higher income individuals pay more in taxes. Surely, those with higher income and greater property benefit more from a governmental system that protects property rights. Moreover, the monetary value attached to other public goods (such as parks and playgrounds) and to positive-externality activities (such as basic research) very likely rises with income as well. Indeed, if the income elasticity of demand for these services exceeds one, as is plausible, a progressive tax system is perfectly consistent with the Just Deserts Theory.

What about transfer payments to the poor? These can be justified along similar lines. As long as people care about others to some degree, antipoverty programs are a type of public good. [Thurow 1971] That is, under this view, the government provides for the poor not simply because their marginal utility is high but because we have interdependent utility functions. Put differently, we would all like to alleviate poverty. But because we would prefer to have someone else pick up the tab, private charity can't do the job. Government-run antipoverty programs solve the free-rider problem among the altruistic well-to-do.

This perspective can, I believe, potentially resolve the two puzzles I posed earlier. The reason Americans are more ready to vote for transfer payments to the American poor than for foreign aid is simply that they care more about their own neighbors than they do about the poor abroad. As a result of these preferences, caring for the American poor is more of a public good than caring for the poor in other nations. This perspective might also explain why most people are averse to the idea of a height tax. If the benefits of government services rise with income, as opposed to

innate ability, there is no reason to use proxies for ability, such as height, when designing a tax system.

A NEW SET OF QUESTIONS

In the end, I don't think the Just Deserts Theory necessarily calls for radical changes in policy toward taxes and income distribution. It does, however, suggest that we focus on a different set of questions when thinking about policy design. A utilitarian asks how quickly marginal utility falls as income rises and how much people respond to the disincentive effects of redistributive tax policy. A Just Deserts Theorist admits that questions regarding utility functions and incentive effects may enter into the analysis, but they are the wrong place to start. Rather, he begins by asking whether people's compensation reflects the contributions they make to society and how much they benefit from government actions.

There may be no way to decide which of these approaches to tax policy is right, at least as judged by the criteria economists usually apply in evaluating theories. The issue is not one of positive economics, so data alone cannot settle the matter. My guess is that people will have different moral intuitions about which approach makes more sense. It is very possible that if we had the opportunity to ask them, Barack Obama and Joe the Plumber would offer different answers about this fundamental question.

Acknowledgements

I am grateful to Joshua Bookin, Stacy Carlson, and participants at the EEA meetings for helpful comments.

Note

1. Presidential Address presented at the February 2010 meetings of Eastern Economic Association.

References

- Akerlof, George 1978. The Economics of 'Tagging' as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning. *American Economic Review*, 68(1): 8–19.
- Case, Anne, and Christina Paxson. 2008. Stature and Status: Height, Ability, and Labor Market Outcomes. *Journal of Political Economy*, 116(3): 499–532.
- Debreu, Gerard, and Herbert Scarf. 1963. A Limit Theorem on the Core of an Economy. *International Economic Review*, 4(3): 235–246.
- Edgeworth, Francis Y. 1897. The Pure Theory of Taxation. *Economic Journal*, 7: 46–70, 226–238, and 550–571 (in three parts).
- Friedman, Milton 1962. *Capitalism and Freedom*. Chicago, IL: University of Chicago Press.
- Goldin, Claudia, and Lawrence F. Katz. 2007. Long-Run Changes in the Wage Structure: Narrowing, Widening, and Polarizing. *Brookings Papers on Economic Activity*, 2007(2): 135–165.
- . 2008. *The Race between Education and Technology*. Cambridge, MA: Harvard University Press.
- Lindahl, Erik 1958. Just Taxation — A Positive Solution, in *Classics in the Theory of Public Finance*, edited by Richard Musgrave and Alan Peacock. London: Macmillan, 98–123.
- Mankiw, N. Gregory 2009. Smart Taxes: An Open Invitation to Join the Pigou Club. *Eastern Economic Journal*, 35(1): 12–23.
- Mankiw, N. Gregory, and Matthew Weinzierl. 2010. The Optimal Taxation of Height: A Case Study of Utilitarian Income Redistribution. *American Economic Journal: Economic Policy*, 2(1): 155–176.

- Mirrlees, James A. 1971. An Exploration in the Theory of Optimal Income Taxation. *Review of Economic Studies*, 38(2): 175–208.
- Nozick, Robert 1974. *Anarchy, State, and Utopia*. Basic Books.
- Piketty, Thomas, and Emmanuel Saez. 2003. Income Inequality in the United States, 1913–1998. *Quarterly Journal of Economics*, 118(1): 1–39, updates available on Saez’s website.
- Thurow, Lester 1971. The Income Distribution as a Pure Public Good. *Quarterly Journal of Economics*, 85(2): 327–336.
- Vickrey, William S. 1945. Measuring Marginal Utility by Reactions to Risk. *Econometrica*, 13(4): 319–333.
- Wicksell, Knut 1958. A New Principle of Just Taxation, in *Classics in the Theory of Public Finance*, edited by Richard Musgrave and Alan Peacock. London: Macmillan, 72–118.