

Justice and Gini coefficients

Theodore J. Everett and Bruce M. Everett

Politics Philosophy Economics published online 21 April 2014

DOI: 10.1177/1470594X14528653

The online version of this article can be found at:

<http://ppe.sagepub.com/content/early/2014/04/20/1470594X14528653>

Published by:



<http://www.sagepublications.com>

On behalf of:

The Murphy Institute of Political Economy

Additional services and information for *Politics, Philosophy & Economics* can be found at:

Email Alerts: <http://ppe.sagepub.com/cgi/alerts>

Subscriptions: <http://ppe.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

>> [OnlineFirst Version of Record](#) - Apr 21, 2014

[What is This?](#)

Justice and Gini coefficients

Theodore J. Everett

SUNY-Geneseo, USA

Bruce M. Everett

Tufts University, USA

Abstract

Gini coefficients, which measure gross inequalities rather than their unfair components, are often used as proxy measures of absolute or relative distributive injustice in Western societies. This presupposes that the fair inequalities in these societies are small and stable enough to be ignored. This article presents a model for a series of ideal, perfectly just societies, where comfortable lives are equally available to everyone, and calculates the Gini coefficients for each. According to this model, inequalities produced by age and other demographic factors, together with reasonable choices under equal opportunity, can raise the Gini coefficients for perfectly just societies to levels at least as high as those of any current Western country, and can as easily account for differences in Gini coefficients between such societies or within one such society over time. If Gini coefficients at these levels are possible for ideal societies without distributive injustice, then they should not be used as proxy measures of distributive injustice in real societies.

Keywords

distributive justice, injustice, inequality, Gini coefficient, income, wealth, age distribution, choice

The Gini coefficient as a measure of distributive injustice

Let us begin with a basic distinction among three forms of economic inequality:

Gross inequality ignores distinctions among individuals that are morally relevant to distributive justice.

Corresponding author:

Theodore J. Everett, Department of Philosophy, SUNY-Geneseo, Geneseo, NY 14454, USA.

Email: everett@geneseo.edu

Fair inequality results entirely from morally acceptable distinctions among individuals.

Unfair inequality results in part from morally unacceptable distinctions among individuals. By any measure, unfair inequality equals gross inequality minus fair inequality.

It should be obvious that gross inequality is different from unfair inequality, hence that distributive injustice cannot be inferred from gross inequality alone. In order to infer anything useful about distributive injustice in any society from a total measure of its inequality, we would also need to know, or at least to have good evidence about, how much of that total inequality is fair and how much is unfair. This fact is quite well known, at least in principle, to everyone who seriously thinks about distributive justice.

Nevertheless, a simple statistical measure of gross inequality called the Gini coefficient has been and continues to be widely used as a *proxy* measure for distributive injustice in both academic and popular discussion on the topic. The Gini coefficient, invented in 1912 by Italian statistician Corrado Gini, is a simple statistical measure of the relationship between an actual distribution and an equal distribution. In income or wealth analysis, the Gini coefficient is calculated as the difference between the actual income or wealth distribution in a society and a perfectly equal distribution in which every member of society has identical income or wealth.¹ In a totally equal income distribution, 10 percent of the people would have 10 percent of the income, 50 percent of the people would have 50 percent of the income, and so on. In any actual society's income distribution, the bottom 10 percent of the people have less than 10 percent of the income and the bottom 50 percent of the people have less than 50 percent of the income, while the top 10 percent of the people have more (typically much more) than 10 percent of the income. If we represent these distributions on a graph, the two curves will look as shown in Figure 1.

In Figure 1, the shaded area represents the difference between the actual income distribution and the perfectly equal distribution. The Gini coefficient (we will usually just call it the Gini) is the ratio of that shaded area to the whole triangular area beneath the straight line for perfect equality. Thus, if there were no inequality at all in a society (the curve matches the line, so that there is no area between them), the Gini would be 0. If, on the other hand, all of the wealth or income in a society were in the hands of a single person (which pushes the curve as far away from the straight line as possible), the Gini would be 1. In any real society, the Gini will of course be somewhere in between.

Gini coefficients have lately become popular in academic, journalistic, and political discussions of distributive justice – particularly with respect to the US, where the Ginis most commonly cited for income and wealth are higher than many people believe they ought to be, higher than they are for other Western countries, and higher than they were in the US itself several decades ago. Yet there are a number of ambiguities and limitations, well known to experts, that make Gini coefficients hard to interpret clearly and consistently.

First, there are different ways of defining income and wealth. The US Census Bureau (2011a), for example, which is the authoritative source on income statistics for the US, develops income data through a household survey called the Annual Social and

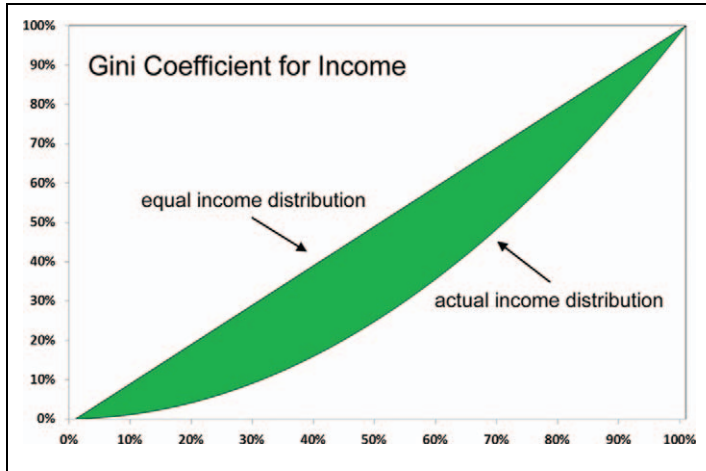


Figure 1. Gini coefficient for income.

Economic Supplement. The Census Bureau income series measures pretax income, thus excluding income taxes as well as Social Security and Medicare taxes.² Alternatively, the Organization for Economic Cooperation and Development (OECD, 2008), an inter-governmental organization of the world's industrial democracies, uses a posttax definition of income, which reduces apparent inequality. Comparisons across countries thus require extreme care to ensure identical definitions. In 2010, for example, the Census Bureau pretax definition yields a Gini of 0.47 for the US, while the OECD definition indicates 0.38 for the same year. Similarly varied Gini coefficients are calculated for wealth defined in terms of accumulated assets instead of annual income. The OECD calculation is a better intuitive measure of the inequality that is relevant to injustice, since it is closer to a measure of overall well-being or utility and seems to be used more commonly in discussions of distributive justice.

A second main problem is that Gini coefficients are indifferent to absolute levels of income or wealth. Intuitively, there is more distributive injustice in a society where annual income ranges between \$1000 and \$100,000 than in a society where income ranges between \$1,000,000 and \$100,000,000, given the lower marginal utility of money for people who are already wealthy. But the Gini coefficient would be identical for these two societies, since it measures only relative differences in income or wealth, and it would seem wrong to count them as equally distributively unjust, all things considered (Garrett, 2010).

A third main problem, and the one that we will focus on here, is that the Gini coefficient in its ordinary use is only a *static* measure of inequality, whereas a *dynamic* measure is required for a careful application of income or wealth statistics to questions of justice. That is, it takes a 'snapshot' of income or wealth distribution at a particular moment in time and does not consider changes in the income of individuals over the course of their lives.³ This makes far more difference to considerations of distributive justice than is generally realized, as we will show.

According to the OECD definition, the current (2008) Gini for income in the US is about .38 and has grown over the past four decades, from about .32 in the mid-1970s. According to Credit-Suisse (2013), the Gini for wealth in the US is .85. Other Western countries have generally lower Ginis for both income and wealth: for Finland, .27 and .66, respectively; for Sweden, .23 and .80; for Germany, .30 and .77; for the UK, .34 and .68; and for Canada, .32 and .73. Viewed in isolation, such numbers can look like strong, even conclusive evidence that the US is a distributively unjust society, more unjust than other Western countries, and increasingly unjust over time. In fact, Ginis and similar statistics are now often treated as if they were direct measures of distributive injustice (indeed, injustice between economic *classes*) in themselves, when they are at best very indirect (and potentially very misleading) evidence for such injustice.⁴

There are three possible assumptions on which these Ginis might still be reasonably used as proxies for distributive injustice. The first is that the fair inequalities in relevant societies are so small as to be negligible compared to the unfair inequalities, so that any seemingly *high* Gini coefficient is an effective proof of absolute distributive injustice.⁵ The second is that some nonnegligible baseline Gini coefficient can be established as the measure of fair inequality in a society, so that any Gini above that level is necessarily *too high* to be compatible with absolute distributive justice.⁶ The third is that in comparing societies across space and time, the level of fair inequality – whatever that level might be – is constant, or close enough to constant, so that a *higher* Gini is an effective evidence of *more* distributive injustice.⁷ Based on these often unspoken assumptions, high, too high, or relatively high Gini coefficients are then typically explained in terms of patently unjust causes such as social class divisions and unequal opportunity, while the factors of fair inequality like age distribution and free career choice are ignored or set aside as trivial.⁸

We believe that the three assumptions that support this way of interpreting Ginis are all false, or at least easily *could* be false for all that is generally known. Nobody has established that the fair component of the gross inequality in Western societies is negligibly small, or that it can be set at some known level over which all inequalities are necessarily unfair, or that it's constant between different societies or over time. And we will not attempt to establish for ourselves how much of the gross inequality in Western societies is fair. But we will argue that minimal common assumptions about economic justice imply Gini coefficients for ideal, perfectly just societies that are surprisingly high – indeed, high enough in theory to account potentially for all of the gross inequalities in actual Western societies as well as for all of the gross differences in inequality between Western societies or within one society over time. While no one (including us) believes that all of the *substantive* inequalities in Western societies are actually fair, the fact that our *statistical* inequalities *could* all be fair invalidates the frequent use of Ginis as a proxy measure of distributive injustice.

A simple model for a perfectly just society

In order to make the gap between gross inequality and unfair inequality clear in absolute terms, let us define a series of ideal model societies that are entirely just by almost any reasonable standard and then calculate their Gini coefficients. Of course, the notion of a

perfectly just society depends on many factors over which reasonable people disagree. Is it fair for people with different natural abilities to have different incomes? Is it fair to let some children inherit wealth from their parents while other, equally deserving children receive nothing? Should women's greater average longevity be reflected in lower annual retirement incomes? We will not consider controversial questions like these in our constructions. We will include only inequalities that almost everyone in the discussion of distributive justice plainly accepts as fair in principle (and also in practice among themselves) and then compute the Gini coefficients for societies with only those inequalities. This will provide no more than a *minimum* baseline of fair inequality in these model societies, but that minimum will be sufficient for our argument.

The two ingredients of the first and simplest perfectly just society we have in mind are that every individual lives exactly the same economic life as every other (so that class structure and other discriminatory inequalities are completely ruled out) and that all individuals live in the comfortable, free conditions common to contemporary, upper-middle-class Western professionals such as tenured philosophers and social scientists. To this end, we take for granted the following two moral principles:

The principle of identity: If two people in a society live identical lives in every morally relevant respect and receive identical benefits, then there is no distributive injustice between them.

The principle of adequacy: The lifestyle of upper-middle-class professional Westerners is not intrinsically unjust.

This model society will be one in which everyone is guaranteed what most of us currently seem to desire for our own children and friends, namely, a successful professional life with a high income and a comfortable retirement. In this ideal society, there is no reason to discriminate between one person and another, because the citizens are all utterly alike in every morally relevant way. The only factor relevant to wealth or income that we allow to vary from one person to another is each person's age at any moment. Thus, we are deliberately leaving out potentially important difference factors like hard work, perseverance, honesty, and talent in order to isolate age as a single factor in our ideal society. If you like, you can imagine that our ideal citizens are all exact clones, indistinguishable in every way other than their age.

It might well be argued against our principle of identity that injustice can still exist between age cohorts, so that even if people have identical whole careers, it is wrong at each moment that those who are at one point in their careers should have more or less than those at any other point. A strict egalitarian in particular might want to assert that every adult should have the same net worth and income at all times regardless of age, though many others who consider themselves egalitarians would be satisfied with equality of outcome over people's whole lives. It is certainly the case that both income and wealth vary widely over the financial lives of ordinary middle-class Americans, with both the highest income and the greatest accumulation of wealth ordinarily gained by those approaching retirement.⁹ It is also undoubtedly true that many young workers feel oppressed by having to make largely regressive transfer payments through Social Security and Medicare to retirees who never had to carry such a huge burden themselves,

while some seniors still feel that they are being treated unfairly by younger people who are generally better off.¹⁰ But such intergenerational resentments would presumably not occur in an ideal system where whole lives are treated equally and where people enjoy upper-middle-class comforts over their entire lives. In real life, they seem to arise mainly between competing cohorts in societies that change over time and distribute different total packages to people of different generations. Thus, much of the present resentment of young workers toward Social Security recipients is based on their (sadly, reasonable) expectation that they will both pay disproportionately for the current retirees' benefits and fail to receive such benefits themselves when they retire. By contrast, in our idealized steady-state model, everyone both expects and receives exactly the same treatment as everybody else of the same age. So, for our simplified model society, it seems reasonable to treat this sort of intergenerational injustice as negligible according to the principle of adequacy.¹¹ In any case, this sort of social distinction is not the usual focus of discussion on distributive justice and not what Gini coefficients and other statistical measures of inequality are typically thought to characterize. Instead, the bulk of our discussion of distributive injustice is concerned with persistent class distinctions between different people, not between different stages of similar people's lives.

Here are some initial assumptions to define our ideal upper-middle-class society, including arbitrary round number values for several important factors in income and wealth distribution:

- Every person becomes independent and starts working at the same age, say 21, with no assets or debts.¹²
- Every person retires at the same age, say 65.
- Every person lives to the same age, say 80.
- Every person has the same initial salary, say \$50,000 after taxes.¹³
- Every person's income increases at the same rate, say 2 percent (real) per year.¹⁴
- Every person saves the same percentage of his/her annual income, say 10 percent, in personal retirement accounts, home equity, or other reasonable investments.
- All savings earn a constant return, say 2 percent (real) per year.
- Every person retires on an annuity based on his/her savings, which are exhausted just at the time he/she dies (so there is no inheritance).¹⁵
- There is a demographic steady state, that is, no immigration and no natural population change.

According to our two principles, this model society is perfectly distributively just, if also rather dull, because everybody lives exactly the same comfortable modern economic life as everybody else. But consider what happens statistically when we create this model society with the values we've suggested. First, looking at income, we see that salary increases from \$50,000 at 21 years of age exponentially to \$129,813 at the retirement age of 65, at which point an income of \$41,852 from annuities kicks in until death at the age of 80.¹⁶ This produces a Gini of .208 for income (using the usual, static OECD calculation). When we look at wealth instead of income, the inequality is considerably greater (as it is in all actual Western societies), with a Gini of .400, resulting from older workers' accumulation of compounded salary increases and interest on regular savings over time. Thus, with age as the sole variable among otherwise identical upper-middle-

class lives, and with moderate values plugged in for interest rates and the other variables, we derive Ginis for income and wealth in our ideal society that are approximately half those of the current US.¹⁷

It is easy to understand how there can be substantial Gini coefficients in this perfectly just society. Since its individuals earn and own different amounts of money at different points in their financial lives, a crosssection of our model society at any moment will include young people with good starting incomes but little or no wealth, middle-aged people with much higher incomes and considerably more wealth, and retirees with lower but still comfortable incomes and gradually declining wealth. The Ginis for this society might be taken superficially as evidence of inequalities between rich and poor *classes* of people, the way that Ginis tend to be interpreted in much of the current discussion. But that cannot be right, since in our model the apparent economic 'classes' comprise *exactly the same people*, just at different points in their careers. What we have defined, then, is an absolutely classless Western-style society with considerable gross inequality, all of which is perfectly fair according to our assumptions.

Our model is robust with respect to its general features. The Gini coefficients that it yields are naturally sensitive to the arbitrary values we plug in for all the variables, but no values that seem plausible make much difference to our basic argument, other than suggesting what the range of Ginis for ideal Western societies might be like. For example, raising or lowering our initial salary of \$50,000 makes no difference at all to the inequalities involved, though raising or lowering the annual *increase* in salary raises or lowers the Ginis for both income and wealth accordingly. Thus, if we halved the annual increase in our initial model from 2 percent to 1 percent, the resulting Ginis would be lowered from .208 and .400 to .161 and .383, respectively. If we doubled it to 4 percent, the Ginis would increase to .301 and .436.

The more of their income people in our model save, the less gross income inequality results. Our initial model's Gini of .208 for income assumes that everybody puts aside 10 percent of his/her annual earnings for investments of one sort or another. If we raised that arbitrary number to 20 percent, the result would be a lower Gini of .145, because each person's income in retirement would be closer to his/her average income while working.¹⁸ If we lowered it to 5 percent, the coefficient would go up to .275.¹⁹ Changing the savings rate has no effect on the model's Gini for wealth of .400, because, with everyone still saving the same percentage over his/her whole careers, relative wealth at any age remains the same.

Raising or lowering the retirement age in our model society changes the Ginis for income and wealth in opposite directions. For example, if everyone in this society retired at 60 instead of 65, the income Gini would increase from .208 to .268, because fewer people would be earning very high preretirement incomes at their jobs, while the wealth Gini would decrease from .400 to .383, because each person would have had 5 fewer years to accumulate his/her savings. Alternatively, if everyone retired at 70 instead of 65, we would see the income Gini drop from .208 to .165, largely because there would be fewer 'poor' retirees with relatively small incomes, while the wealth Gini would increase, from .400 to .417, largely because each older person would accumulate more years of savings and interest. So, here we have a trade-off between inequalities of income and wealth: the later people retire, the less income inequality and the more wealth inequality are liable to result.²⁰

The same is true of other variables like longevity and appreciation on savings. Thus, if longevity increased from 80 to 85 years in our original ideal society, the Ginis for income and wealth would change from .208 and .400 to .251 and .385, respectively. If longevity were reduced from 80 to 75, the results would be .160 and .407. Similarly, halving the real rate of appreciation on savings from 2 percent to 1 percent per year results in increased income inequality and decreased wealth inequality (Ginis of .235 and .384, respectively), while doubling the rate to 4 percent results in changes the other way around (Ginis of .163 and .434, respectively).

There are other desirable features of upper-middle-class economic life that have not been represented in our very simple model but that can lead to greater gross inequality. One is that young people do not always head straight from college to their first career-level jobs but typically spend several years beforehand in some combination of graduate study, interning, volunteering, traveling, living in big cities for adventure, and otherwise exploring life before settling down. To represent this in our model, we might suppose that each person lives for 5 years, say, at something over minimum wage, say \$20,000 income per year, and then starts his/her \$50,000 career job at the age of 26, with all other parameters being the same. Adding this feature to our original model would raise the income Gini from .208 to .254 and the wealth Gini from .400 to .440. Stipulating less time between college and career would produce intermediate results.²¹

In sum, we can see that Gini coefficients are sensitive to several real-life technical factors that have nothing apparently to do with injustice.²² We have not tried to model more than a skeleton of modern economic life here, obviously, and we have opted for simplicity over strict realism in the values chosen for the basic variables, so we are in no position to say that any particular Gini for either income or wealth is the right baseline coefficient for actual societies like ours. The best that we might do is to suggest a range of coefficients for plausible-looking sets of values in demographically steady state, Western-style societies, say setting retirement age between 60 and 70 years, longevity between 75 and 85 years, salary increases between 1 percent and 4 percent (real) per year, savings rates of between 5 percent and 20 percent per year, savings interest and home appreciation of between 1 percent and 4 percent (real) per year, and postgraduate education of between 0 and 5 years. This range of values yields a range of 'fair-inequality' coefficients for income of anywhere from .126 to .487 (which would include the current income Ginis for all actual Western societies) and for wealth of .358 to .534 (which still falls somewhat short of actual Western Ginis for wealth). These are all possible Ginis for a modern society of perfect justice – in which, again, the *only* difference among its equally comfortable citizens would be their present age.

Fair inequality by choice

There are several plausibly fair or neutral distinctions among individuals that we have left out of our simple model but that are worth considering for somewhat more realistic distribution models that permit some diversity in people's lives. Again, it is controversial whether distinctions in income and wealth depending only on differences in natural ability are fair, but we all agree that equally talented upper-middle-class people can desire

different things in life. Some value prestige more than high income; some value leisure time more than either; some particularly value excitement, others comfort, others security, and so on. Such choices typically have predictable consequences for an individual's lifelong levels of income and wealth, so they can reasonably be seen as elements of 'package deals'. And almost everyone believes that some income and wealth differentials among reasonable people of the same age can be perfectly fair if they depend entirely on free and equal choice. To reflect this in a more realistic model of ideal societies, let us retain the principle of adequacy, so that we can speak about distinctions that are fair within an otherwise acceptable upper-middle-class society. And let us still presuppose that all individuals in our model have identical natural abilities and economic opportunities. But let us replace the general principle of identity with this:

The principle of choice: If two people in a society differ economically only as a result of the foreseeable consequences of reasonable choices equally available to each, there is no distributive injustice between them.

We should be clear that we are not saying that distributive justice *demands* that all inequalities be attributable to reasonable choice. The principle of choice is not intended to commit us to this or any other positive theory of distributive justice. It is only intended to state a very minimal thesis that almost any theory would entail, namely that distributive justice *permits* predictable (that is, not luck-dependent) inequalities that are entirely attributable to free, equal, and reasonable choice.²³

Including difference factors due to free choice with foreseeable consequences naturally raises the potential gross inequality in our model upper-middle-class societies, driving their maximum Gini coefficients for income and wealth considerably higher than in our initial model. In combination, they increase the range of Ginis even further. Here are a few examples.

Fair choice in household partnership

So far we have been speaking only about individual income and wealth. But the US Census statistics typically used for Gini calculations are given for household, not personal, income and wealth. This obviously makes little difference for people who are single, but modern upper-middle-class society includes increasingly many two-earner households that ordinarily have much more income and wealth. Among upper-middle-class Western professionals, it is ordinarily a matter of free choice whether people enter into marriage or other household partnerships, with or without children or other dependents. But the gross statistical consequences of different choices are significant. If, for example, we put a random half of the members of our original uniform society into two-income households (such couples then constituting one third of all households, each with twice the income of the remaining single households), the Gini for household income in the new model society rises from .208 to .279 and the Gini for household wealth increases from .400 to .443.²⁴

Fair choice in working conditions

Some people choose jobs that require longer hours than others; some agree to work in less pleasant environments; some will travel further to their places of work; some accept particularly arduous, stressful, or dangerous duties, and the like, while others choose more pleasant and comfortable jobs closer to home. If we assume that any such hardship factor justifies some extra increment in pay, and we suppose again that such choices are equally available to all, then we must accept a corresponding increment of gross inequality as fair. Thus, if we return to our original model and now add a differential pay ratio of, say, 2:1 between high-hardship work and low-hardship work, however these may be distinguished, and imagine equal numbers of people choosing each, then the result is a more unequal but still perfectly just model society, with a Gini for income that has risen from .208 to .279 and a Gini for wealth that has increased from .400 to .443.

Fair choice in job location

In real societies like the US, the cost of living can vary substantially from urban to rural settings or from one region of the country to another. For example, public school teachers in cities like New York can make about twice the salary of teachers in the rural Midwest, but it can cost them about twice as much to maintain an equivalent standard of living.²⁵ Other things being equal, we should think of such differences as morally neutral – but they have misleading consequences for statistical gross inequality. For a rough idea of how misleading this potentially could be, suppose we let half of our original model society live in cities and half in the country, with the urban workers earning twice the nominal salary but spending twice as much to live as all their rural peers. This would have the same statistical effect as any other 2:1 difference in starting pay, with Ginis for income and wealth rising again from .208 to .279 and from .400 to .443, respectively.

Fair choice in education

In the modern West, people from comfortable backgrounds sometimes choose to go to work directly after college, while others choose to spend several years going to graduate school or doing other pre-career activities that pay less than their career starting salaries. Even if we suppose that such activities make no difference to initial salaries, gross inequality in both income and wealth necessarily increases as a result of some people delaying their careers. If we suppose that half our initial model population still enters the workforce with an annual salary of \$50,000 straight from college at age 21, but the other half chooses to enter at 26 with the same starting salary, but after earning only, say, \$20,000 per year in the meantime, the Ginis for income and wealth increase from .208 and .400, respectively, to .235 and .424. If we suppose instead that the delayed but more educated workers start at a compensatory higher salary of, say, \$75,000, then the Gini for income rises even further to .253, while the Gini for wealth stays the same .424. It seems intuitively to be a morally neutral choice whether someone starts his/her career straight from college or postpones it for the sake of further education, then recouping just enough from higher salaries to end up with the same wealth and income at retirement.

But this neutral choice among peers increases the gross income inequality of our model society in two ways: first, by counting people as relatively 'poor' while they are in graduate school; and second, by counting them as relatively 'rich' while they are working at compensatory higher salaries.

In real life, most such educational choices are not wealth neutral, of course, and can result in substantially above- or below-average overall financial rewards, which further increase the gross inequality of the society in which they occur. For example, every year roughly the same pool of graduating college students chooses between attending academic graduate programs and attending law school, which decision ordinarily results in greatly different subsequent financial lives. Academic PhD programs take about twice as long as law school (roughly 6 as opposed to 3 years), and professors make much lower average salaries than lawyers, but many people choose this freely so that they can follow their own interests in their work and live a comfortable, generally quiet life with reasonably high prestige.²⁶ If we were to split the postgraduate half of the population in our ideal society into two further halves, one of which starts only at the college-level salary of \$50,000 after 6 years of graduate school ('professors') and the other at \$100,000 after 3 years ('lawyers'), the Ginis for income and wealth would increase to .290 and .454, respectively.²⁷

Fair choice in productivity

Even people with the very same jobs can choose to work at them in different ways. Some people focus almost all of their energies on their careers, work harder, take fewer vacations, attend conferences, study at night and on weekends to improve their performance, cultivate useful relationships, and constantly strive for raises and promotions at their jobs while seeking better positions elsewhere. Others prefer to enjoy their lives with friends and family, pursue sports, travel, gardening, or other hobbies, or to improve their minds with studies not related to their careers. Both can be perfectly reasonable choices, depending on the different things we want out of life as individuals.²⁸ But they result in very different career trajectories with different rates of salary increase. If we imagine another 2:1 differential, with half of our initial ideal workforce increasing their salary by 2 percent per year, and the other, more ambitious half by 4 percent, the resulting Ginis for income and wealth will be .296 and .436, respectively.²⁹

Fair choice in savings rate

In our original model, we assumed a uniform savings rate of 10 percent for the entire working population. But we know that people of the same age in similar comfortable circumstances save money at different rates, some choosing to enjoy themselves more while they are young, others putting more away for their future lives. Within a wide range of reasonable choices, the resulting differences in wealth among initial peers are intuitively neutral with respect to social justice: one person may retire somewhat more comfortably now, but another had a nicer car and more vacations back then, and these were perfectly voluntary choices, so neither person has a reasonable claim against the other. But when we add this differential to our calculations, greater gross inequalities

result, particularly for wealth. This can be especially misleading, because the new statistics will show the first person's greater maximum wealth and retirement income, while they ignore the other person's big house, fancy cars, and pricey annual vacations. To get a notion of the superficial change that this can make, imagine a difference in savings rate between two random halves of our initial population, with one half saving, say, 5 percent and the other 10 percent of their annual income toward retirement in a new model society. This will increase the Gini for income from .208 to .245 and the Gini for wealth from .400 to .443. As above, a greater or lesser difference in choice of savings rates will yield greater or lesser changes in both coefficients.

Fair choice in pension plans

Western middle-class workers are sometimes given a choice these days between the so-called defined benefit and defined contribution retirement plans (or, more commonly and less directly, employers make the choice of which to offer and the workers choose among employers). Defined contribution is what we have been assuming for everyone in our model so far: each person keeps his/her savings in a personal account, accrues further net worth through compound interest, and then trades that personal wealth for an annuity. Defined benefit programs promise workers a certain annuity up front, typically a given percentage of their final salary, while keeping the money to fund it in a common pool. Ideally, the retirement income from both types of plan should be the same, so income Ginis should be unaffected by this choice. But the savings that fund defined benefit plans are invisible to Ginis for personal or household wealth, since they are held by the employers. Instead of personal wealth that counts in calculating Ginis, workers have essentially the same guaranteed income that the same savings would fund in either system, but in a form that is opaque to these statistics. If we imagine that half of the people in our first model society switched from a defined contribution to an equivalent defined benefit plan, and that these savings constituted half of their total investment portfolio (that is, 5 percent of their income per year), this would have no effect on the original income Gini of .208, but it would make half of their savings disappear from personal into common accounts, with the result that the Gini for wealth would increase from .400 to .443, just the same as if they'd never saved that money at all. If we supposed that people had all of their savings in their retirement plans, so that for half of the society zero savings would be visible to Gini calculations, the Gini for wealth would rise all the way to .695.³⁰

Fair choice in savings risk

Upper-middle-class people with 401k plans, Individual Retirement Accounts (IRAs), and other modern retirement accounts are commonly required to choose how much investment risk they will accept in their portfolio. High-risk investments like stocks tend to pay better (about double, on average) over time than low-risk investments like bonds or certificates of deposit but carry a greater risk of loss in bad economies.³¹ On top of this, many adventurous upper-middle-class people choose to

invest their savings in particular stocks that they believe will beat the market or real estate investments, restaurants, and other small-business endeavors, while their security-minded peers keep all their non-retirement savings in the bank. If we imagine another 2:1 differential in return on savings between successfully aggressive and conservative investors of all sorts, so that half of our model society achieves 4 percent growth per year on savings while the other half gains our standard 2 percent, the Gini for wealth rises from .400 to .436, just as it did with differential productivity above, while the Gini for income actually goes down from .208 to .192. This is a rather perverse result, because this lower overall income inequality is just a function of the fact that the successful investors' much higher income in retirement, at \$76,525 per year as opposed to the usual \$41,852, is now close to the average income of people who are working, so that the relative 'poverty' of retirees as a group is diminished.

Fair choice in retirement age

Among healthy middle-class people in similar financial condition, some choose to retire earlier and others later than average, some never retiring at all. As everyone facing retirement knows, this has a great effect on how much income we receive during retirement. But again, we do not consider the resulting inequalities to be unjust, provided they are made as totally unforced decisions in a reasonable way. In our models so far, we have assumed a uniform retirement age of 65. But suppose that a random half of our society retires at 60 and the other half works until 70. The additional fair inequality that results raises the Gini for income from .208 to .226 and the Gini for wealth from .400 to .428.

Combining several such choice-based distinctions can produce almost arbitrarily high Ginis for both income and wealth. For an extreme example, let us imagine an upper-middle-class society in which one half of the population freely chooses to marry, to work in high-cost areas at high-salary advanced-degree jobs with long hours, to strive for advancement, to save at a high rate in 401k or IRA accounts, to invest aggressively (all with our arbitrary 2:1 differentials), and to retire at 70, while the other half chooses to remain single, to work in low-cost areas at relatively low-salary jobs, to advance, save, and invest at ordinary rates with defined benefit retirement plans, and to retire at 60. The resulting Ginis in this model society would be .743 for income and .818 for wealth.³² Both of these coefficients are higher than those for any actual Western country (the income Gini far higher) – all in a model society that is still perfectly just according to our principles.³³ It is unrealistic, of course, to imagine a society composed entirely of polar opposites like this, with one half conservative and relatively placid workers and the other half highly ambitious, risk-taking careerists; in real life, almost everyone falls somewhere in between.³⁴ Our point is only that this is a *possible* society composed of people like ourselves in which there is no distributive injustice at all, but which nevertheless has higher gross inequality as measured by Gini coefficients than any actual Western society.

Comparing Ginis across places and times

It might be argued that even though Gini coefficients tell us little or nothing by themselves about absolute distributive injustice, they can still be used as a guide to relative levels of injustice between different societies or as a means of telling whether a given society has gotten more or less distributively unjust over time. It is this relative interpretation of Gini statistics that we find most prominent in arguments for the distributive inferiority of the current US, both to other Western countries now and to the US itself of several decades ago.³⁵ But relative gross inequality as measured by Gini coefficients can be just as misleading as absolute gross inequality. For one thing, changes or differences in age structure obviously matter. Other things being equal, one society will have a lower Gini for income than another if it has fewer very-low-earning or very-high-earning workers, which might just mean fewer younger or older citizens, perhaps through lower fertility or less emphasis on the medical extension of old age. One society will tend to have a lower Gini than another for both income and wealth if it has fewer people in their 50s and 60s, perhaps because more people drink alcohol or smoke tobacco and die young as a result. Such basic demographic facts need to be taken into account before a reasonable judgment of relative injustice can be made from Gini coefficients. All of the other factors we have mentioned must be accounted for as well. One country will display a higher Gini for income than another if the first has greater longevity in general; if it has a two-earner household partnership rate closer to a peak of about 60 percent than does the second; if it has a lower average retirement age than the second or has more variation in retirement age; if it has more varied working conditions; if it comprises people who start their careers later in life than the second, or has more varied starting ages than the second; if it has a lower savings rate or more difference among savings rates; if it has greater differences in individual productivity increase or level of investment risk; or if it has greater differences in cost of living between urban and rural settings.³⁶ Other things being equal, one society will have a higher Gini coefficient for wealth in the same circumstances, except when the first has higher longevity, a lower retirement age, or more variation in retirement age, in which case the Ginis for wealth will be lower. Also, as we have seen, large differences in Ginis for wealth can result from preferences between equally reasonable defined-benefit and defined-contribution pension systems. To the extent that any of these differences are based on people's voluntary individual or social choices – possibly influenced by culture, geography, or other morally neutral conditions – they must be accounted for in inferences concerning relative as well as absolute distributive injustice.

The same thing goes for those societies in which the Gini coefficients change over time. We cannot infer that increasing Ginis over time indicates greater unfair inequality, hence possibly more distributive injustice, unless we rule out changes in fair, baseline inequality that might explain the changing Ginis. Since the US and other Western countries have benefitted from 7 or 8 years of increased average longevity over the past four decades, for example, we should expect to see increases in their Ginis for income, other things being equal, because more people will be living in retirement. Since there has also been a trend with the advance of working women toward both more single households

and more two-earner married households, we should expect to see further increases in Ginis for both income and wealth as a result. We cannot estimate how much of the noted increase in Ginis for the US and elsewhere over recent decades ought to be attributed to such morally acceptable causes. But it is important to note that such explanations are possible and must be accounted for in any analysis intended to derive increased distributive injustice from increased Gini coefficients.

Another possible difference between countries, or within countries as different times, is the rate of immigration, particularly Third World immigration. If one society chooses to admit large numbers of immigrants who are less skilled or educated than the native population, younger on average, or more fertile at younger ages than average – all typical of large Third World populations entering into Western countries – then this will raise income and wealth inequality in three connected ways. First, it will add a number of people to the society without the education necessary to make upper-middle-class incomes. Second, even if all immigrants are able to earn starting salaries equal to those of natives with a college education, the new immigrants will alter the standing age structure to increase the ratio of younger to older workers. Third, by having more children at younger ages, the new subpopulation will increase in size until it is assimilated into the native fertility structure.³⁷ Thus, any society more open to immigration of younger, less skilled, and more fertile workers than another (for example, the US compared to any other Western country) will tend to have proportionately higher Gini coefficients, even if the immigrants are treated individually just like natives of the same age.³⁸ Yet most of us think it is a good, not a bad, thing for distributive justice to let large numbers of Third World immigrants join our prosperous and free societies, since even below-average incomes in the West can be substantially higher than average incomes in developing countries.³⁹

Another sort of choice with consequences for Gini coefficients is the extent to which a society pursues policies that tend to even out the incomes or wealth of its people over the course of their own lives. For example, a country might institute a mandatory retirement age, or raise or lower an existing one, in order to reduce unemployment among young workers or perhaps to resolve a shortage of labor. Countries might also opt to choose favorites among other normally free choices, for example, to encourage more postgraduate education or higher savings rates, in order to promote economic growth or other socially valued ends. A just, democratic society might even choose to reduce gross income inequality between older and younger workers for its own sake, if the majority of its citizens expected such a policy to produce a happier nation overall. This might be a very good utilitarian choice for homogeneous societies like the Nordic countries, with their established high prosperity, well-tended elderly, and unworried youth. At the same time, more diverse, immigrant-rich societies like the United States might reasonably favor long-term economic growth or other goals at some cost to average utility over the short run. Obviously, different policies along these lines will produce quite different Ginis, and deliberately so. But, at least within some reasonable range of policy options, this will be a social choice in income or wealth redistribution, not a requirement of justice.

Conclusion

We have argued that economic justice in modern, Western-style societies like ours is consistent with considerable gross inequality of wealth and income as measured by Gini coefficients, both as a function of natural differences in age and as a result of voluntary differences in savings rates, retirement plans, and other ordinary choices at the individual or social level. We have not tried to show that distributive injustice does not exist in the US or other Western countries. There is little doubt that such injustice exists in the West just as it does everywhere else in the world, whether it is due to human cruelty and indifference, to economic exploitation, to ethnic, gender, and other forms of unfair discrimination, or to well-meant but ineffective or destructive social and economic policies; and there is little doubt that some societies are more unjust than others. But we have argued that gross inequality statistics like Gini coefficients for income and wealth do not constitute good *evidence* that of such absolute or relative injustice, since even the highest actual Ginis could in theory be explained in terms of morally acceptable differences within and among perfectly just societies. Therefore, these Gini coefficients ought not to be used as proxy measures of distributive injustice.

Acknowledgements

We would like to thank Christopher Annala, James Bennett, Robert R. Everett, Carlo Filice, Harry Howe, Kevin McCarthy, Stephen Sullivan, and three anonymous referees for helpful comments on this paper. We are also grateful for the patient and detailed advice that we received from the editors of PPE.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Notes

1. By way of illustration, assume a society of 100 people with a total income of \$1 million. To calculate the Gini coefficient, we would create two data series. The first series is the cumulative income shares of the society if each member had an identical annual income of \$10,000. The first data point would be set at zero, and the second would be the income of the poorest person as a share of the total. Since in a perfectly equal society each member has an income of \$10,000 (\$1 million divided by 100), the second data point would be 1 percent. The third data point would be the total income of the two poorest people or 2 percent and so forth. The final data point in this series would be 100 percent. Plotted on a graph, this series would be a straight line rising from 0 to 100 percent. The second data series would be the cumulative income shares of the society as they actually are. Again, the first data point in the series is set at zero. If the poorest person in the society had an annual income of \$1000, the second data point in this series would be 0.1 percent (\$1000 divided by \$1 million). If the second poorest person had an annual income of \$1200, the third data point would be 0.22 percent and so forth. Once again, the final data point would be 100 percent.
2. The Annual Social and Economic Supplement defines 'money income' for each household member over the age of 15 to include earnings (wages and salaries), unemployment

compensation, workers' compensation, Social Security, supplemental security income, public assistance, veterans' payments, survivor benefits, disability benefits, pension or retirement income, interest, dividends, rents, royalties, and estates and trusts, educational assistance, alimony, child support, financial assistance from outside of the household, and other income. The series does not include income in kind, such as food stamps, employer-paid health insurance, employer-held retirement plans, and subsidized housing or farm products consumed by farmers themselves. This definition is not in itself unreasonable, but relatively small adjustments can have significant impacts on the Gini. For example, in the year 2003 (the last year for which the census Bureau has made these detailed calculations), the money income Gini for the US was 0.450. If we eliminate government transfer payments, the Gini increases 0.507. However, if we then subtract payroll taxes and state and federal income taxes, the Gini falls back to 0.475. Adding in government health care benefits (Medicare and Medicaid) plus means-tested government transfer payments and the imputed return on home equity, the Census Bureau Gini can be as low as 0.390, very close to the Organization for Economic Cooperation and Development measure of 0.38.

3. This is presumably because the most complete and readily available data sets comprise anonymous, census-style survey data and would not support dynamic analysis of the lifetime changes of individuals. The US Treasury Department (2008) addresses the problem differently by surveying the tax returns of individuals. Their report offers very different insights into inequality but has been largely ignored in the public debate. Blomquist (1981) comes closest to what we are after here, comparing actual annual incomes with modeled lifetime incomes.
4. For one very prominent example, Branko Milanovic, a lead economist in the World Bank research group, takes Ginis for granted as measures of inequality between rich and poor classes of citizens, with no reference to age distribution and other background factors of fair inequality (Milanovic, 2010).
5. This is what we find in many journalistic treatments of the subject, for example Kristof (2012).
6. Norton and Ariely (2011) offer a Rawlsian approach to baseline inequality, which avoids consideration of detailed moral and economic issues. In effect, they just ask people in surveys how much inequality they would consider fair behind the 'veil of ignorance', and take their answers as a good proxy for their intuitions about distributive justice. They conclude from their surveys that a distribution in which the top quintile of the society had about three times as much wealth as the bottom quintile is what most Americans consider fair, which suggests a baseline Gini for wealth of between about .20 and .23. A detailed critique of this approach is beyond the scope of this article, but we think that a proper Rawlsian experiment ought to distinguish carefully between total gross inequality and gross inequality within age cohorts or other groups of peers and to provide some minimal test for the coherence of subjects' responses.

Kristof's (2012) column, among many other popular sources, spreads a remarkable 'meme' from Norton and Ariely's article (2011: 9–12), allegedly based on their surveys, to the effect that Americans generally prefer Sweden's wealth distribution over that of the US. The claim is remarkable because it is actually based on a comparison between the American *wealth* distribution and the Swedish *income* distribution, a fact that the authors acknowledge only in an endnote. A genuine, apples-to-apples comparison would have made no sense because the US and Sweden have very similar wealth inequalities.

7. See, for example, Fischer (2011). Even the comedian Jon Stewart on his popular *Daily Show* (18 August, 2011) gets into the issue, proclaiming that 'The United States is not a Third World

country by any measure – except, perhaps, income inequality’, while showing a list of Gini income coefficients with the US toward the bottom.

8. See, for example, Therborn (2011). Similarly, Page and Jacobs (2009) take the recent increase in the US Gini for wealth as evidence of a ‘frantic sprint by the super rich’ to gain at the expense of others. And it is not just social scientists and political commentators who are using the Gini coefficient in this way. Even business executives are taking recent changes in the US Gini for income as a measure of the ‘shrinking middle class’, according to Byron (2011).

9. For some indication, here are the most recent statistics available to us. Median household income by age of head in 2010 was as follows:

<25: 29.1
 25–34: 51.4
 35–44: 63.4
 45–54: 64.3
 55–64: 58.9
 65–74: 41.0
 ≥75: 26.2

in thousands of 2011 dollars (US Census Bureau, 2011b). Median household net worth distribution by age of head in 2010 was as follows:

<35: 65.3
 35–44: 217.4
 45–54: 573.1
 55–64: 880.5
 65–74: 848.3
 ≥75: 677.8

in thousands of 2010 dollars (Bricker et al., 2012). These statistics are for the whole US population, but we assume the upper-middle-class component is at least roughly similar, given the professional career patterns typical of people in this group.

10. In fact, professional people seem to prefer having their highest incomes and most of their wealth when they are getting older, seeing this as some compensation for losing the freedom and energy they had when they were young, and in any case appropriate to the greater financial responsibilities that come with raising families and other obligations that we take on as we age. This sense of fair exchange may diminish as intergenerational inequalities increase, with young middle-class people in the US facing the massive student loan debts now common, together with their parents’ increasing longevity, which delays and diminishes their prospects for inheritance.
11. The issue of intergenerational equity has been discussed extensively in the literature (see, for example, Kerlie, 2001). There may hypothetically be cases where people in the society live identical lives, but all suffer poverty during some period in those lives, for example, during old age. Such cases may indeed give rise to a kind of distributive injustice, though not the class-based injustice that some people infer from Gini statistics. The principle of adequacy assumes that this is not an issue in our model society, even though incomes do vary over members’ lifetimes.
12. This is imagined to include free higher education and to exclude the student loans that burden many of our graduates today.

13. All salary and other income in our model is meant to be understood as net, disposable income, after all taxes are subtracted and any transfer payments added in.
14. This is imagined to cover any normal raises, promotions, or moves to better positions during each person's career.
15. For home equity investments, this could take the form of a reverse mortgage.
16. In any real Western society, this would be income on top of what is guaranteed by tax and transfer, such as through the US Social Security system.
17. The fact that these baseline Ginis are about half of the actual US Ginis should not be taken as evidence that the US is only half as distributively unjust as critics claim it is. As with a person's body temperature and his/her health, it could turn out that all the interesting 'action' takes place only within a fairly narrow range of Ginis. It is worth noting, though, that this baseline Gini for wealth is already twice what Norton and Ariely (2011) claim that most Americans believe is fair.
18. Ginis for income will go down as the saving rate goes up, until the point where retirement income is the same as average working income. Beyond that point, the retirees will have increasingly higher than average incomes, so the Gini for income will go up again. Note that we are considering only gross income here, not the net income that would result from subtracting savings. Looking at net incomes would show a more equal distribution at moderate savings rates, but this is not how Ginis are ordinarily computed.
19. The income Gini will continue to increase as the savings rate decreases to a maximum of .356 when savings reaches zero, after which the retirees get nothing.
20. We do not mean to imply that Ginis are sensitive to changes only at the extremes of wealth or poverty. Income or net worth in the middle quintiles can also increase or decrease relative to the extremes, resulting in higher or lower Ginis as well.
21. For what it's worth, the authors of this article both took considerably longer than 5 years each to finish their educations and made considerably less than \$20,000 real incomes in the meantime. We are assuming that most of today's young people are more prudent. As with college in the basic model, we are ignoring the large debts that many young people accumulate these days in graduate school, and assuming that they all start their careers with zero net worth.
22. Others (for example, Kwok, 2010) have previously noted the Gini coefficient's sensitivity to demographic factors.
23. Whether choices that do involve luck, such as gambling with money (or, say, one's career), are always distributively just is another, and a rather heated, question. For a recent discussion of 'luck egalitarianism', see Williams (2013).
24. Two-earner households typically have higher costs as well, especially respecting child care, but these are ignored in the usual statistics about gross household inequality.
25. For example, last year the cost of living in Brooklyn, NY, was 181 percent of the national average and the average public teacher salary was \$70,000, while the cost of living in Jefferson city, MO, was 97 percent of the national average and the average teacher made \$37,000. (Council for Community and Economic Research, 2011; US Bureau of Labor Statistics, 2011). On average, the difference in both salary and cost of living between urban and rural settings in the US is about 50 percent.
26. Note that many ABDs and PhDs end up going to law school anyway after trying an academic career, and these people lose several years of income and savings that cannot be replaced. We still think of these highly educated people not as victims of distributive injustice but as free

people following their own desires. The same is true for would-be professional artists or musicians who end up as teachers or in other steady jobs. These are reasonable choices for people like ourselves to make in full view of their likely consequences.

27. This ignores the much higher debts that burden American law school graduates today, which most but not all lawyers can pay back comfortably. It is still a far better 'package deal' financially for most equivalent students to choose law rather than an academic career.
28. Even in unionized industries, the most ambitious and career-focused workers tend to rise to higher paying management positions, sometimes within the union itself.
29. The additional exponential increase in annual salary yields a top salary for 'strivers' of \$296,958 as opposed to our model's standard \$129,813, with a retirement income of \$66,216 as opposed to \$41,852. Wealth tops out at \$850,830 instead of the usual \$537,762.
30. Defined-benefit plans have largely been replaced by defined-contribution plans within the private sector in the US over the past few decades, and the public sector will probably follow suit, with many government employees currently being given a choice. This may account for some of the widely noted increase in US Gini coefficients for wealth over the past 40 years.
31. For some indication of these differences, as of October 2009 the average annual return on stocks (Standard & Poor's 500) since 1926 was 9.8 percent excluding dividends, while the average annual return on government bonds was about 5.5 percent, and on Treasury bills, about 3.7 percent (Sommer, 2009).
32. Assuming further that the less ambitious group keeps not just half but all their savings in defined benefit plans, the wealth Gini rises to .829.
33. It is easy to imagine that a society like this would quickly become unjust for future generations, since the wealthier-by-choice have much greater resources to spend on their children's education and other unearned advantages over the children of the relatively-poor-by-choice, plus much more money to pass on to them in their estates. Note again, though, that our ideal state model presupposes, in effect, 100 percent inheritance taxes and free higher education available to all, in order to preserve the exact same choices for everyone in every generation.
34. Still, we can easily imagine a career-driven heart surgeon with an aggressive 401k married to a colleague and living in Manhattan, whose unmarried brother is a teacher with a defined benefit pension plan, living in rural Maine on one-sixteenth of her household income, with an even smaller fraction of her nominal savings, but both having made their financial choices freely under similar initial conditions, and both quite happy with their decisions. The reader probably knows people from upper-middle-class origins who have freely chosen equally divergent economic lives as well as many others in between.
35. See note 8 for examples.
36. In the US, for example, the most expensive city is Manhattan, with a cost of living index of 216.7 percent of the national average (US Census Bureau, 2012) with several other major cities at around 130 percent to 140 percent. By comparison, Stockholm's cost of living is only 115.8 percent of the Swedish national average, and Helsinki's is only 106.6 percent of the Finnish national average. Together with corresponding differentials in salary, this superficial, morally neutral difference alone will tend to make these Nordic countries appear more equal than the US, according to Ginis for income, hence to casual observers more distributively just.

37. The age factors involved in immigration are hard to represent within our simple model, but we might arbitrarily adjust our initial model to include, say, 10 percent immigrants who begin work at 21 with a pay rate of, say, \$20,000 per year, with the same other factors as everybody else in our original model. Even ignoring age-structure changes, this by itself raises the income Gini for that society from .208 to .238 and the wealth Gini from .400 to .418. If that portion of the population then doubles in size relative to the native population over time, those Ginis will increase to .260 and .432, respectively.
38. Again, Kwok (2010) and others have already noted that Gini coefficients mask such demographic changes. Our point here is that they are important to understanding of the Gini's misleading moral implications.
39. It might be argued that a Western society with higher Gini coefficients based on immigration is still in fact more distributively unjust with respect to its own post-immigration population, while at the same time the global 'society', for which the Gini coefficients would have decreased, has become less distributively unjust. The proper argument would then move from the question of whether distributive inequality is always unjust to the question of whether national distributive injustice is always a morally bad thing. The immigration example seems to show that global distributive *justice* might well justify local distributive *injustice* understood this way.

References

- Blomquist N (1981) A comparison of distributions of annual and lifetime income: Sweden around 1970. *Review of Income and Wealth* 27: 243–264.
- Bricker J, Kennickell A, Moore K, et al. (2012) Changes in U.S. family finances from 2007 to 2010: evidence from the survey of consumer finances. *Federal Reserve Bulletin* 98: 1–80.
- Byron E (2011) As middle class shrinks, P&G aims high and low. *Wall Street Journal*, 12 September (updated 14 September). Available at: <http://online.wsj.com/news/articles/SB10001424053111904836104576558861943984924> (accessed 1 February 2014).
- Council for Community and Economic Research (2011) *ACCRA Cost of Living Index* 44(3). Available at: <http://www.cedbr.org/content/2011/ACCRA.pdf> (accessed 1 February 2014).
- Credit-Suisse Research Institute (2013) *Global Wealth Databook 2013*, p. 101. Available at: <https://publications.credit-suisse.com/tasks/render/file/?fileID=1949208D-E59A-F2D9-6D0361266E44A2F8> (accessed 1 February 2014).
- Fischer M (2011) Map: U.S. ranks near bottom on income inequality. *The Atlantic*, 19 September. Available at: <http://www.theatlantic.com/international/archive/2011/09/map-us-ranks-near-bottom-on-income-inequality/245315> (accessed 1 February 2014).
- Garrett T (2010) US income inequality: it's not so bad. *Inside the Vault* (U.S. Federal Reserve, St. Louis) 14: 1–3.
- Kerlie D (2001) Justice between the young and the old. *Philosophy & Public Affairs* 30: 152–177.
- Kristof N (2012) Why let the rich hoard all the toys? *New York Times*, 4 October: A35
- Kwok K (2010) Income distribution of Hong Kong and the Gini coefficient. *The Government of Hong Kong, China*. Available at: <http://www.eabfu.gov.hk/en/pdf/income.pdf> (accessed 1 February 2014).
- Margonelli L (2013) The formula. *Pacific Standard*, 22 January. Available at: <http://www.psmag.com/magazines/magazine-feature-story-magazines/gini-coefficient-index-poverty-wealth-income-equality-51413> (accessed 1 February 2014).

- Milanovic B (2010) *The Haves and the Have-Nots: A Brief and Idiosyncratic History of Global Inequality*. New York, NY: Basic Books.
- Norton M and Ariely D (2011) Building a better America – one wealth quintile at a time. *Perspectives on Psychological Science* 6: 9–12.
- Organization for Economic Co-operation and Development (OECD) (2008) *Growing Unequal? Income Distribution and Poverty in OECD Countries*. Available at: http://www.keepeek.com/Digital-Asset-Management/oced/social-issues-migration-health/growing-unequal_9789264044197-en#page1 (accessed 1 February 2014).
- Page B and Jacobs LR (2009) *Class War? What Americans Really Think about Economic Inequality*. Chicago, IL: University of Chicago Press.
- Sommer J (2009) In this 10-year bonds win by a mile. *New York Times*, 25 October, BU4.
- Therborn G (2011) Global inequality: the return of class. *Global Dialogue* 2, 30 September. Available at: <http://isa-global-dialogue.net/global-inequality-the-return-of-class> (accessed 1 February 2014).
- US Bureau of Labor Statistics (2011) *May 2011 National Occupational Employment and Wage Estimates*. Available at: http://www.bls.gov/oes/2011/may/oes_nat.htm (accessed 1 February 2014).
- US Census Bureau (2011a) *Historical Income Tables H10AR_2011*. Available at: <http://www.census.gov/hhes/www/income/data/historical/household> (accessed 1 February 2014).
- US Census Bureau (2011b) *Income, Poverty, and Health Insurance Coverage in the United States: 2010*. Available at: <http://www.census.gov/prod/2011pubs/p60-239.pdf> (accessed 1 February 2014).
- US Census Bureau (2012) *Statistical Abstract of the United States: 2012*. Available at: <http://www.census.gov/compendia/statab/2012/tables/12s0728.pdf> (accessed 1 February 2014).
- US Treasury Department (2008) *Income mobility in the U.S. from 1996 to 2005*. Available at: <http://www.treasury.gov/resource-center/tax-policy/Documents/Income-Mobility-1996to2005-12-07-revised-3-08.pdf> (accessed 1 February 2014).
- Williams A (2013) How gifts and gambles preserve justice. *Economics and Philosophy* 29: 65–85.

Author biographies

Theodore J. Everett is a professor of philosophy at SUNY-Geneseo. He has published articles on a variety of topics in *The Journal of Philosophy*, *Philosophy and Phenomenological Research*, *The Australasian Journal of Philosophy*, *Synthese*, and elsewhere.

Bruce M. Everett is an adjunct associate professor of international business at the Fletcher School of Law and Diplomacy, Tufts University. He has published a variety of articles in *Issues in Science and Technology*, *Foreign Policy*, and elsewhere and writes on energy and economics at bmeverett.wordpress.com.