

Ethics and Epidemiology: The Income Debate

Gopal Sreenivasan*, Duke University

*Corresponding author: Gopal Sreenivasan, 201 West Duke Building, Box 90743, Durham NC USA 27708. Email: gopal.sreenivasan@duke.edu

This paper reviews the epidemiological debate between the relative income hypothesis and the absolute income hypothesis. The dispute between these rival hypotheses has to do with whether an adequate account of the relationship between income and life expectancy requires the definition of ‘income’ to include any comparative element. I discuss the evidence offered for the relative hypothesis (which answers, ‘yes’), as well as two important criticisms that have been levelled against this evidence. I also offer some critical reflections on the debate from a philosophical standpoint concerned with the ethics of population health. Both hypotheses agree that a redistribution of income towards the worst off will improve their life expectancy.

There are significant inequalities in health between countries and also within countries. Intuitively, these inequalities in health are seriously unjust; and, not surprisingly, claims of injustice have figured prominently in public calls for action to ‘close the [health] gap in a generation’ (Marmot *et al.*, 2008). Since an inequality’s being avoidable is very plausibly a necessary condition on its being an injustice, philosophical efforts to analyse the injustice of unjust health inequalities must begin from an understanding of the likely causes of inequalities in health—in particular, of the subset of *socially controllable* causes of inequalities in health. This means that ethics has to work hand in hand with epidemiology.

Income is one of the most widely researched candidates for the status of a social determinant of health. Within a given society, an individual’s life expectancy correlates highly and positively with his or her income (e.g., McDonough *et al.*, 1997). Of course, it is a separate question—a controversial, as well as an important question—as to what extent this correlation is causal. However, in the case of the correlation between income and health, this standard social scientific puzzle is further complicated by a prior controversy concerning how best to characterise the correlation itself. The debate has to do with which definition of ‘income’ is most adequate to capture the correlation(s) between income and health.¹

In this paper, I review the epidemiological debate about the definition of income, and then offer some critical reflections on the debate from a philosophical standpoint concerned with the ethics of population health.

1. The roots of the income debate lie in a widely remarked apparent paradox presented by the comparative international data on the relationship between income and health. As Richard Wilkinson—one of the foremost

protagonists in the debate—describes it, the apparent paradox is that ‘income is closely associated with health within countries but not between them’ (1996: 73). The claim that, between countries, income is not closely associated with health requires some interpretation. But Wilkinson’s point is easy to see on a so-called Preston curve (after Preston, 1975), such as the one presented in Figure 1.

Above a per capita gross domestic product (GDP) of (say) \$8000, there is tremendous variation in average national *income* with not much variation in national life expectancy, whereas below a per capita GDP of (say) \$6000 there is tremendous variation in national *life expectancy* with not much variation in average national income. It clearly makes a vital difference to national life expectancy whether per capita GDP is above \$6000–8000 or not. But when comparison is restricted to either side of the bend in the Preston curve, income is not closely associated with health—neither between developed nations nor between developing nations.

The apparent paradox focuses on *developed* societies, i.e., beyond the bend in the curve. There, as we have just seen, income and health are not closely associated *between* societies. On the other hand, *within* a given developed society, an individual’s health remains highly positively correlated with her income. How can this combination of very different characteristic patterns in developed societies be explained?

One prominent explanation is known as the *relative income hypothesis*, and it is the explanation of the apparent paradox advocated by Wilkinson himself. It holds, roughly, that an individual’s life expectancy is a function not simply of the ‘absolute’ (i.e., noncomparative) level of her income, but *also* of its ‘relative’ level (i.e., compared




Figure 1. International relationship between income and life expectancy.²

to others' income in her society). More specifically, the relative income hypothesis holds that once an individual crosses some threshold of noncomparative material deprivation, further gains in her noncomparative income make *no* significant contribution to her life expectancy. Above that threshold, the income that contributes significantly to an individual's life expectancy is her relative income. In particular, (greater) discrepancies in relative income make a (greater) negative contribution to the life expectancy of the less well off.³

The principal rival of the relative income hypothesis is known as the *absolute income hypothesis*. Roughly speaking, this alternative hypothesis holds that the contribution income makes to individual life expectancy is entirely a function of the individual's *noncomparative* income. Both hypotheses accept, then, that an individual's noncomparative income makes a significant (positive) contribution to her life expectancy. Causally, this contribution is widely understood to be mediated by material risk factors such as inadequate nutrition, lack of clean water and sanitation, and poor housing.

The debate between the rival hypotheses concerns where, if anywhere, on some noncomparative scale income *ceases* to contribute to individual life expectancy; and, more importantly, whether any of the contribution that income does make is mediated by an individual's *comparative* share of income in society. While the relative hypothesis affirms this last point, the absolute hypothesis simply denies it. To preview, my own position will be that there is as yet no evidence for the relative hypothesis, which is not to say it is false—merely completely unproven.

But let us begin by considering how the relative income hypothesis offers to explain the apparent paradox. This hypothesis is clearly consistent with the correlation between income and life expectancy *within* developed na-

tions. The challenge is therefore to explain why income and life expectancy are not closely associated *between* developed nations. Here we should bear in mind that a nation's per capita GDP is a measure of its inhabitants' absolute incomes. Correlations with per capita GDP thus have nothing to do with relative income.

Now suppose that the bend in the Preston curve (\$6000–8000 per capita GDP in 2000) corresponds to the point at which a great majority of a nation's inhabitants have crossed the individual-level threshold of noncomparative material deprivation. In that case, in comparisons between developed nations, the great majority of individuals on both sides will be beyond the point at which differences in their absolute incomes make any significant difference to their individual life expectancy. According to the relative income hypothesis, then, it only stands to reason that such comparisons can reveal significant differences in per capita GDP without significant differences in national life expectancy.

By contrast, when a developed nation is compared to a developing one, the comparison in effect straddles the threshold of noncomparative material deprivation. On the developing nation side of such comparisons, the great majority of individuals will fall below this threshold, and so will have the lowest individual life expectancies. Under these conditions, the differences in per capita GDP between the two nations will inevitably be closely associated with significant differences in national life expectancy. Of course, this explanation is not peculiar to the relative income hypothesis, since it appeals to the health effects of differences in *absolute* income. But it remains legitimate to invoke because of the significance the relative hypothesis assigns to absolute income below the threshold of material deprivation.

Evidently, the keys to this dissolution of the apparent paradox lie in the existence of a *threshold* of noncomparative material deprivation, as well as in the association of that individual-level threshold with the bend in the Preston curve. Wilkinson (1994) bolsters these points by appealing to certain features of the 'epidemiological transition.' According to this idea, a nation's rounding the bend on the Preston curve marks the stage at which the predominant burden of national mortality shifts from communicable diseases, traditionally associated with poverty, to noncommunicable diseases (such as cancer and cardiovascular disease): 'For the bulk of the population, the stranglehold of the absolute standard of living on health has been overcome' (Wilkinson, 1996: 45).

2. Whatever the merits of the relative income hypothesis as an explanation of the apparent paradox, we should distinguish between explanatory power and truth. Various kinds of evidence have been adduced as more

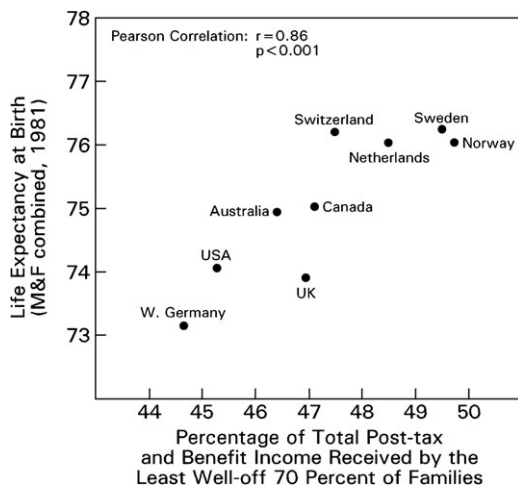


Figure 2. Income distribution and life expectancy in developed countries, 1981.⁴

direct confirmation of the hypothesis itself. The evidence Wilkinson (1992) initially adduced came from comparative international studies. Additional evidence was subsequently collected by other researchers from various domestic comparisons, notably among the 50 U.S. states. (A selection of these later papers is assembled in Kawachi *et al.*, 1999: chap. 5–8.) Let me very briefly review Wilkinson's evidence for the relative income hypothesis.

Drawing on data from nine developed nations, Wilkinson obtained a strong correlation between a nation's degree of income inequality and its average life expectancy—the greater the inequality, the lower the life expectancy. This result is exhibited in Figure 2, where 'income inequality' is measured as the share of total personal income received by the least well-off 70 per cent of families and 'life expectancy' combines male and female rates at birth.

Even more impressively, Wilkinson also obtained strong correlations between the *rate of change* in a nation's income inequality and the rate of change in its average life expectancy: Nations with greater improvements in income equality had greater improvements in average life expectancy. Figure 3 illustrates this result with data from 12 European nations between 1975 and 1985. Changes in income inequality are measured as changes in the proportion of the population living on less than half the national average disposable income.

As these data suggest,⁵ the central policy implication of the relative income hypothesis is that an equalising redistribution of income within a given society will improve national life expectancy. For example, Wilkinson (1992)




Figure 3. Annual rate of change in life expectancy and of proportion of population in relative poverty, 1975–1985.⁶

concludes that '[i]f Britain was to adopt an income distribution more like the most egalitarian European countries the slope of the regression equation suggests that about two years might be added to the population's life expectancy' (p. 167).

At bottom, the controversy surrounding the relative income hypothesis has to do with whether there is actually any evidence for it. But it is useful to distinguish two rather different criticisms lodged against the evidence we have just reviewed. On the one hand, there is what might be called a *straightforward empirical* critique. According to this critique, the strong correlations between income inequality and national life expectancy that have been adduced are not robust: As a result of problems with the quality of the data and the methodology employed, the correlations have proved very difficult to replicate. Kawachi *et al.* (1999: part II) contains a selection of papers making this critique, and of others responding to it. I shall not enter into the details here.

Still, it is fair to report an emerging consensus that there is *no reliable* evidence of a strong correlation between income inequality and average life expectancy, particularly not from international comparisons. In 2002, the *British Medical Journal*—where Wilkinson's original paper (1992) appeared—published an editorial on the subject reaching a negative summary verdict: 'Now that good data on income inequality have become available for 16 Western industrialised countries, the association between income inequality and life expectancy has disappeared' (Mackenbach, 2002: 1). Comprehensive reviews of the literature by Deaton (2003) and by Lynch *et al.* (2004) arrive at the same conclusion.

On the other hand, there is what might be called a *subtle logical* critique. This critique grants, for the sake of argument, that there is a strong correlation between income inequality and national life expectancy. Its claim




Figure 4. How nonlinearity explains the population health effects of income inequality.⁷

is rather that this correlation does *not constitute evidence* for the relative income hypothesis. According to the subtle logical critique, the existence of such a correlation is in itself perfectly consistent with the absolute income hypothesis, and hence incapable of functioning as evidence against it. The essentials of this critique, which recalls the utilitarian argument for resource egalitarianism, were spelled out almost 35 years ago by Preston (1975) himself, and again by Rodgers (1979). But I shall follow Gravelle (1998), since his illustration of the point is especially nice.

Let us assume that an individual's life expectancy is solely a function of her noncomparative income. Assume, furthermore, that the function is nonlinear: specifically, that there are diminishing marginal returns (in individual life expectancy) to increasing individual income.⁸ These assumptions suffice to explain a correlation between income inequality and average life expectancy in a society. To see this, consider two countries, A and B, with the same average income, \bar{y} . For simplicity, assume that, in each country, there are only two incomes—low and high, each earned by half the population. The basic difference between the countries is that the low income in B, y_{1B} , is lower than the low income in A, y_{1A} ; and the high income in B, y_{2B} , is correspondingly higher than the high income in A, y_{2A} . In other words, the distribution of income is more unequal in B than in A, even though average income is the same. These assumptions are illustrated in Figure 4, where the dark concave curve represents the individual-level relationship between income and life expectancy. Note that the vertical axis plots mortality risk, which is the *inverse* of life expectancy.

Average mortality in B, m_B , is higher than average mortality in A, m_A . That is, the country with greater income inequality (B) also has *lower* average life expectancy. In this respect, our two-country comparison duplicates the result that Wilkinson claimed to find. Yet, in Figure 4, it is clear to see that the higher average mortality in B is entirely due to the concavity in the curve (i.e., to the nonlinearity in the individual-level relationship between income and mortality).

Since average income is the same in both countries, the income difference between the high-income halves of the populations ($y_{2B} - y_{2A}$) exactly cancels out the difference between the low-income halves ($y_{1A} - y_{1B}$). However, the same cannot be said for the mortality differences entailed by these income differences. The 'mortality premium' paid by the poorer of the *high-income* halves ($m_{2A} - m_{2B}$) is *smaller* than the premium paid by the poorer of the *low-income* halves ($m_{1B} - m_{1A}$). That is because, while the income difference is constant, the implied mortality premium is determined differently in the two cases: In the high-income case, the premium is determined by the flatter right-hand segment of the curve; and in the low-income case, by the steeper left-hand segment. Of the two countries, the one with less income inequality (A) has greater average life expectancy *because* its population gains more life expectancy on its richer low-income half than it loses on its poorer high-income half ($(m_{1B} - m_{1A}) > (m_{2A} - m_{2B})$). But this fact is explained entirely by the shape of the curve, which conforms to the dictates of the absolute income hypothesis.

Now it is important to be clear about just what the subtle logical critique shows. It shows that a given aggregate-level association between income inequality and average life expectancy *can* be entirely explained by nonlinearity in the individual-level relationship between absolute income and mortality. That is to say, for any given aggregate-level association, there is a degree of nonlinearity in the individual-level relationship that would explain it entirely. Hence, no aggregate-level association, in itself, contradicts the absolute income hypothesis. But it remains an independent empirical question, which must be investigated separately, whether in a particular case the actual degree of nonlinearity in the individual-level relationship (i.e., the actual shape of the curve) suffices to explain the entire aggregate-level association at issue.⁹ The subtle logical critique therefore leaves open the possibility of finding evidence for the relative income hypothesis.¹⁰ What it excludes is that aggregate correlations between income inequality and average life expectancy, simply by themselves, count as such evidence.

Recall the apparent paradox that income and life expectancy are not closely associated between developed

nations, but are closely associated within them. We are now in a position to recognise that Wilkinson's explanation of the paradox is not, in fact, peculiar to the relative income hypothesis. What his explanation appeals to, fundamentally, is the proposition that the relationship between national income and national life expectancy depends on how that income is *distributed*. But this proposition is perfectly consistent with the absolute income hypothesis: If the relationship between absolute individual income and individual life expectancy is nonlinear (as in Figure 4), then, as we have seen, the contribution a given unit of national income makes to average life expectancy depends on *where* on the individual-level curve that unit of income is located.¹¹ In that case, it would not be surprising to find that per capita GDP is not closely associated with national life expectancy, since per capita GDP is not sensitive to the distribution of national income, while national life expectancy is. Thus, the weak association *between developed* nations simply conforms to what one should generally expect, even on the absolute income hypothesis.

From this perspective, it is actually the close association between income and life expectancy found when *developed* nations are compared to *developing* nations that calls for some explanation. But this was the association that Wilkinson himself explained by reference to the health effects of absolute income. Here it suffices, for example, to suppose that very low levels of national income compel the great majority of a nation's inhabitants to occupy the steepest left-hand segment of the individual-level income–life expectancy curve (cf. Figure 4).¹²

3. In several respects, the debate between the relative and absolute income hypotheses is puzzling. I shall describe two puzzles, both of which strike me as fairly acute, and then record two further, cautionary observations. The first puzzle is quite simple: Why do proponents of the relative income hypothesis even offer simple aggregate correlations between income inequality and average life expectancy as evidence for their hypothesis? Since such correlations are in themselves logically incapable of discriminating between the two hypotheses, they plainly cannot function as evidence for the relative hypothesis against the absolute one. While the subtle logical critique is decisive, it is not *that* subtle. More to the point perhaps, Wilkinson, at least, was always aware of its essentials (e.g., 1996: 102), as any reader of Preston (1975) or Rodgers (1979) was positioned to be.

The second puzzle concerns the point of the debate. Since the relative and absolute income hypotheses do contradict each other, there is evidently a coherent theoretical dispute that can be engaged. What is less clear is whether this dispute has any practical significance. As we

have seen, the central domestic policy implication of the relative income hypothesis is that an equalising redistribution of income will improve average life expectancy. But the same policy implication follows with equal force on the *absolute* income hypothesis, as sceptics about the relative hypothesis are often themselves at pains to emphasise (e.g., Gravelle, 1998: 385; Deaton, 2003: 118). So the point of the debate cannot have anything to do with the basic epidemiological merits of redistributing income. Curiously, Wilkinson himself appears to agree: 'The argument is about how, rather than whether, narrower income differences are related to better population health. The pathway does not alter the reality of the health benefits or the central policy implications' (1999: 956).

My own suspicion is that advocacy of the relative income hypothesis may be the surrogate expression of a conviction in the importance of *psychosocial* pathways in the social determination of health.¹³ The importance of material pathways is agreed on both sides of the income debate, at least for very low levels of absolute income (i.e., where there is absolute material deprivation). However, it may be that what proponents of the relative income hypothesis really want to insist upon is that psychosocial pathways are *also* important (Wilkinson, 1999; Marmot and Wilkinson, 2001; Marmot, 2004); and perhaps even, after a certain point, more important. Certainly, the efficacy of some psychosocial pathway between income and health is plausibly regarded as a necessary condition of the relative income hypothesis;¹⁴ and its advocates devote considerable attention to arguing for psychosocial pathways (see Kawachi *et al.*, 1999: part III).

Yet, even if this diagnosis is correct, it serves more to transform the puzzle than to dissipate it. On the one hand, the absolute income hypothesis says nothing to exclude psychosocial pathways between health and some social factor. It only denies that such pathways run through (*relative*) income in particular. On the other hand, the existence of a psychosocial pathway between health and (say) job control is plainly not a sufficient condition of the relative income hypothesis: it would still remain to trace the further causal connection between 'job control' and 'relative income.' If job control were established as a social determinant of health (see Marmot *et al.*, 1997), why would one insist on adding a further link to 'relative income,' given that the contribution made by job control would itself be mediated by a psychosocial pathway? On this telling, there is no answer.

My first observation is something philosophy might contribute to the epidemiological debate. Alternatively, and perhaps more realistically, it can be understood as a cautionary note for philosophers who look to the epidemiological literature for empirical lessons. The

debate between the relative and absolute income hypotheses assumes that the policy goal is to improve *average* life expectancy. Among other things, this seems to reflect the continuing influence of utilitarianism.¹⁵ From the standpoint of justice, however, it is at least unclear that improvements in average life expectancy have first call on our attention. More likely, it is simply false. Both egalitarians and prioritarists (Parfit, 2001), for example, will agree not only in rejecting this goal, but also in affirming the capital importance of improving the life expectancy of the worst off. It therefore pays to observe that some of the (apparently empirical) positions maintained in the epidemiological debate crucially depend upon the background assumption of a normatively controversial goal.

Let me give two examples. Suppose the domestic correlation between individual income and life expectancy is causal; and that its shape resembles Figure 4 (note 8). In that case, transferring income from the best off to the worst off *will improve* the life expectancy of the worst off. No one doubts this in principle. In practice, however, redistribution may also lower the society's average life expectancy: for instance, if the redistributive mechanism is inefficient (Deaton, 2002: 23). Still, this only licenses practical scepticism about the health merits of redistributing domestic income *if* our policy goal is to improve average life expectancy. No such scepticism follows if our goal is rather to improve the life expectancy of the worst off. Likewise, the straightforward empirical critique also focuses on average life expectancy (as the correlate of income inequality). Hence, its negative findings provide no reason to doubt the health efficacy of a domestic redistribution of income either—at least, not to egalitarians or prioritarists.

My second observation concerns the significance of the fact that national income's contribution to life expectancy is sensitive to the distribution of income. As we have seen, this is a fact on which the relative and absolute income hypotheses agree. But it requires some caution to be exercised in extending the central domestic policy implication of the individual income–life expectancy correlation to the international arena. In particular, it is not enough simply to transfer income from rich countries to poor *countries*. Even if the correlation is causal, such transfers may fail to realise anything like the maximum increase in life expectancy from a given increment of income. The maximum increase in life expectancy requires the income to be delivered to the bottom of the income distribution *within* a given country. An international redistribution of income that aims to improve health should therefore target the poorest individuals within poor countries.

The significance of this sensitivity to distribution is not confined to *transfers* of income from rich countries to poor ones. It applies to any kind of growth in the income of poor countries. Recent discussions of globalisation provide a case in point. One prominent argument holds that 'globalisation is good for your health' (Dollar, 2001; Feachem, 2001).¹⁶ It proceeds in two main steps: one step claims that growing participation in international trade by poor countries contributes to significant growth in their income, while the other step claims (along the lines of the absolute income hypothesis) that income growth contributes to significant improvement in individual health. However, even if both claims are correct (a matter of controversy), it does not follow that the health of the world's poorest inhabitants will improve significantly with significant growth in their nation's income: That conclusion requires a further step, demonstrating significant growth specifically in the income of the worst off.¹⁷

Notes

1. An extensive collection of relevant articles, including all of the most prominent early contributions to this debate, is reprinted in Kawachi *et al.* (1999).
2. Deaton (2003: 116). Figure reproduced by permission. Circles have a diameter proportional to population size.
3. Note that this formulation only defines a *family* of relative income hypotheses. Among other things, it omits a specific definition of 'relative income' (e.g., of the comparison group). This is a point on which the proponents of the relative income hypothesis have not been entirely clear. While its details do not matter for our purposes, the proponents are committed to there being at least one specific definition on which the claims made for the family defined in the text can be vindicated. For more precise discussion of alternative definitions, see Deaton (2003) and Wagstaff and van Doorslaer (2000).
4. Wilkinson (1992: 166). Figure reproduced by permission.
5. Suggest, that is, on the assumption that the correlation between income and life expectancy is causal.
6. Wilkinson (1992: 166). Figure reproduced by permission.
7. Gravelle (1998). Figure reproduced by permission.
8. This assumption has been confirmed empirically (Wolfson *et al.*, 1999, Figure 1; Deaton, 2003, Figure 2).
9. Using data from 52 nations, Preston himself later investigated whether national income makes any

contribution to national life expectancy above and beyond that made by the degree of nonlinearity in the actual individual level relationship (1980: 291–293). He concluded that it does not, i.e., that the ‘relations between mortality and income at the national level are indeed dominated by relations between mortality and income at the individual level.’ However, he added that his ‘result should be treated with great caution because of inaccuracy and incomparability in the measure of income distribution and because the log-linear functional form probably simplifies a more complex relationship.’

10. Using domestic data from the 50 U.S. states, Wolfson *et al.* (1999) found that an aggregate correlation between income inequality and average life expectancy remained even after the contribution made by nonlinearity in the actual individual level relationship was taken into account. Note, however, that these are almost the same authors, albeit in a different order, as Lynch *et al.* (2004), whose previously cited review concluded that there was no reliable evidence of an aggregate correlation between income inequality and average life expectancy among wealthy nations. In their later publication, these authors seem to take the view that the earlier domestic evidence reveals something idiosyncratic about the United States (pp. 81–82). Furthermore, Deaton (2003: 122) cites an unpublished study that found *no* such residual aggregate correlation between income inequality and mortality among the 50 U.S. states (using a different technique from Wolfson *et al.*, 1999).
11. Indeed, both Preston (1975: 241–242) and Rodgers (1979) framed what I have called the subtle logical critique precisely in terms of the significance of the *distribution* of national income.
12. This resembles, but is weaker than, Wilkinson’s assumption that the great majority of a developing nation’s inhabitants fall below some *threshold* of noncomparative material deprivation. In particular, nothing in the text posits, let alone locates, a ceiling on the absolute income effect.
13. For an overview of some of the most prominent candidates, see Sreenivasan (2008: sec. 4).
14. How would a material pathway connect to *comparative* income?
15. The emphasis on average life expectancy may also reflect the illusion that when redistribution improves the life expectancy of the worse off, while also *conserving average* life expectancy, then ‘everyone’ is better off (cf. Wilkinson, 1999: 957). Whether or not this makes for effective political rhetoric, we should acknowledge that it is just not true. Redistribution

takes income from the better off, who thereby lose both income and life expectancy.

16. Daniels *et al.*’s (2000) argument that ‘justice is good for our health’ affords an interesting comparison. Both arguments exploit the empirical connection between higher individual income and better health. On some counts, the argument from justice has the advantage. Notably, since it appeals (as far as income is concerned) to Rawls’ difference principle, its connection to improved income among the worst off is *definitional*. By contrast, the connection between that outcome and globalisation is doubly empirical, as we shall see. On the other hand, the argument from globalisation has the advantage that globalisation is already well underway.
17. I do not mean to imply that Dollar and Feachem are unaware of this requirement, though the evidence they offer for the extra step is limited in various respects.

Acknowledgements

For helpful comments, I would like to thank Dan Brock and, especially, Angus Deaton, who was generous beyond measure.

References

- Daniels, N., Kennedy, B. and Kawachi, I. (2000). Why Justice Is Good for Our Health. In Cohen, J. and Rogers, J. (eds), *Is Inequality Bad for Our Health?* Boston, MA: Beacon Press.
- Deaton, A. (2002). Policy Implications of the Health and Wealth Gradient. *Health Affairs*, **21**, 13–30.
- Deaton, A. (2003). Health, Inequality, and Economic Development. *Journal of Economic Literature*, **41**, 113–158.
- Dollar, D. (2001). Is Globalization Good for Your Health? *Bulletin of the World Health Organization*, **79**, 827–833.
- Feachem, R. (2001). Globalisation Is Good for Your Health, Mostly. *British Medical Journal*, **323**, 504–506.
- Gravelle, H. (1998). How Much of the Relation between Population Mortality and Unequal Distribution of Income Is a Statistical Artefact? *British Medical Journal*, **316**, 382–385.
- Kawachi, I., Kennedy, B. and Wilkinson, R. (eds). (1999). *The Society and Population Health Reader: Income Inequality and Health*. New York: New Press.
- Lynch, J., Smith, G. D., Harper, S., Hillemeier, M., Ross, N., Kaplan, G. and Wolfson, M. (2004). Is Income

- Inequality a Determinant of Population Health? Part I. A Systematic Review. *Milbank Quarterly*, **82**, 5–99.
- Mackenbach, J. P. (2002). Income Inequality and Population Health. *British Medical Journal*, **324**, 1–2.
- Marmot, M. (2004). *The Status Syndrome*. New York: Henry Holt.
- Marmot, M., Bosma, H., Hemingway, H., Brunner, E. and Stansfeld, S. (1997). Contribution of Job Control and Other Risk Factors to Social Variations in Coronary Heart Disease Incidence. *Lancet*, **350**, 235–239.
- Marmot, M. and Wilkinson, R. (2001). Psychosocial and Material Pathways in the Relation between Income and Health. *British Medical Journal*, **322**, 1233–1236.
- Marmot, M., Friel, S., Bell, R., Houweling, T. and Taylor, S. (2008). Closing the Gap in a Generation: Health Equity through Action on the Social Determinants of Health. *Lancet*, **372**, 1661–1669.
- McDonough, P., Duncan, G., Williams, D. and House, J. (1997). Income Dynamics and Adult Mortality in the United States, 1972 through 1989. *American Journal of Public Health*, **87**, 1476–1483.
- Parfit, D. (2001). Equality or Priority? In Harris, J. (ed.), *Bioethics*. Oxford: Oxford University Press, pp. 347–386.
- Preston, S. M. (1975). The Changing Relation between Mortality and Level of Economic Development. *Population Studies*, **29**, 231–248.
- Preston, S. M. (1980). Causes and Consequences of Mortality Declines in Less Developed Countries during the Twentieth Century. In Easterlin, R. (ed.), *Population and Economic Change in Developing Countries*. Chicago: University of Chicago Press, pp. 289–360.
- Rodgers, G. B. (1979). Income and Inequality as Determinants of Mortality: An International Cross-Sectional Analysis. *Population Studies*, **33**, 343–351.
- Sreenivasan, G. (2008). Justice, Inequality, and Health. *Stanford Encyclopedia of Philosophy*.
- Wagstaff, A. and van Doorslaer, E. (2000). Income Inequality and Health: What Does the Literature Tell Us? *Annual Review of Public Health*, **21**, 543–567.
- Wilkinson, R. G. (1992). Income Distribution and Life Expectancy. *British Medical Journal*, **304**, 165–168.
- Wilkinson, R. G. (1994). The Epidemiological Transition: From Material Scarcity to Social Disadvantage? *Daedalus*, **123**, 61–77.
- Wilkinson, R. G. (1996). *Unhealthy Societies: The Afflictions of Inequality*. London: Routledge and Kegan Paul.
- Wilkinson, R. G. (1999). Two Pathways, but How Much Do They Diverge? *British Medical Journal*, **319**, 956–957.
- Wolfson, M., Kaplan, G., Lynch, J., Ross, N. and Backlund, E. (1999). Relationship between Income Inequality and Mortality: Empirical Demonstration. *British Medical Journal*, **319**, 953–955.