Available online at www.sciencedirect.com
ScienceDirect
Journal of Health Economics 26 (2007) 1–24
JOURNAL OF HEALTH ECONOMICS
www.elsevier.com/locate/econbase

Inequality and mortality: Long-run evidence from a panel of countries

Andrew Leigh [a], Christopher Jencks [b,*]

[a] *Social Policy Evaluation, Analysis and Research Centre, Research School of Social Sciences, Australian National University, Australia*
[b] *Malcolm Wiener Center for Social Policy, Kennedy School of Government, Harvard University, United States*

Received 7 May 2005; received in revised form 23 July 2006; accepted 27 July 2006
Available online 11 September 2006

Abstract

We investigate whether changes in economic inequality affect mortality in rich countries. To answer this question we use a new source of data on income inequality: tax data on the share of pretax income going to the richest 10% of the population in Australia, Canada, France, Germany, Ireland, the Netherlands, New Zealand, Spain, Sweden, Switzerland, the UK, and the US between 1903 and 2003. Although this measure is not a good proxy for inequality within the bottom half of the income distribution, it is a good proxy for changes in the top half of the distribution and for the Gini coefficient. In the absence of country and year fixed effects, the income share of the top decile is negatively related to life expectancy and positively related to infant mortality. However, in our preferred fixed-effects specification these relationships are weak, statistically insignificant, and likely to change their sign. Nor do our data suggest that changes in the income share of the richest 10% affect homicide or suicide rates.
© 2006 Elsevier B.V. All rights reserved.

JEL classification: I12; N30

Keywords: Health; Inequality; Mortality; Homicide; Suicide

1. Introduction

Do changes in economic inequality lead to changes in mortality rates? More than 100 articles on this question have been published over the past two decades, but no consensus has emerged

* Corresponding author. Tel.: +1 617 495 0546.
E-mail addresses: andrew.leigh@anu.edu.au (A. Leigh), christopher_jencks@harvard.edu (C. Jencks).
URL: http://www.econrsss.anu.edu.au/∼aleigh/, http://www.ksgfaculty.harvard.edu/Christopher_Jencks.

0167-6296/$ – see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.jhealeco.2006.07.003

(Lynch et al., 2004a). One major reason has been the paucity of reliable historical data on income inequality. As a result, most studies have examined the relationship between inequality and mortality at a single point in time. Because income inequality and mortality are likely to have common causes that cannot all be measured, the cross-sectional relationship between inequality and mortality is unlikely to provide an unbiased estimate of how changes in income inequality affect mortality.

We investigate this issue using a new source of data on economic inequality: the share of personal income received by the richest 10% of adults in Australia, Canada, France, Germany, Ireland, the Netherlands, New Zealand, Spain, Sweden, Switzerland, the UK, and the US. We have annual data covering an average of 62 years per country. As a result, we can control both year and country fixed effects, thereby holding constant both stable country-to-country differences and annual changes in mortality that affect all countries in the same year.

Deaton (2003) and Lynch et al. (2004a) have recently surveyed the literature on economic inequality and mortality. Both reviews conclude that although there are plausible reasons for anticipating a relationship between inequality and mortality, the empirical evidence for such a relationship is weak. We have found only five studies that use time series from developed countries to analyze the inequality–mortality relationship. After examining changes in life expectancy in the UK during the 20th century, both Wilkinson (1989) and Sen (1999) conclude that longevity rose faster when the income gap between the rich and poor narrowed. However, their measures of income inequality are relatively inexact, and they do not try to take account of temporal variation in the effect of technological innovation. Focusing on the last decades of the 20th century in both the US and the UK, Wilkinson (1996) also argues that rising inequality during the 1980s was the main reason why the decline in infant mortality slowed between 1975 and 1985. Deaton and Paxson (2001), in contrast, find no systematic relationship between inequality and health in either the UK or the US from the mid-1970s to the mid-1990s. Likewise, when Lynch et al. (2004b) look at 100-year national trends and 30-year regional trends in the US, they find little evidence of a causal relationship between income inequality and mortality.

This paper extends previous studies by examining long time series for 12 of the world's richest countries rather than one or two. Our findings are consistent with those of Deaton and Paxson (2001) and Lynch et al. (2004b), not with those of Wilkinson (1989, 1996) or Sen (1999). In our preferred specifications we find only small and statistically insignificant relationships between income inequality and mortality. This holds true regardless of whether we measure mortality using life expectancy at birth, infant mortality, homicide, or suicide. It also holds true when we introduce lagged measures of inequality, when we focus exclusively on the period since 1960, when we control both the educational level of the population and health expenditures. Our findings suggest that the relationship between income inequality and mortality is either non-existent or too fragile to show up in a robustly estimated panel specification.

The paper is organized as follows. Section 2 presents a simple model of the relationship between inequality and health. Section 3 describes our data. Section 4 presents our results. Section 5 concludes.

2. A simple model of the relationship between inequality and health

Epidemiologists and social scientists have proposed numerous mechanisms by which income inequality might affect an individual's health. We can group these mechanisms under three broad headings: absolute income, relative income, and society-wide effects of income inequality. We discuss these mechanisms in turn.

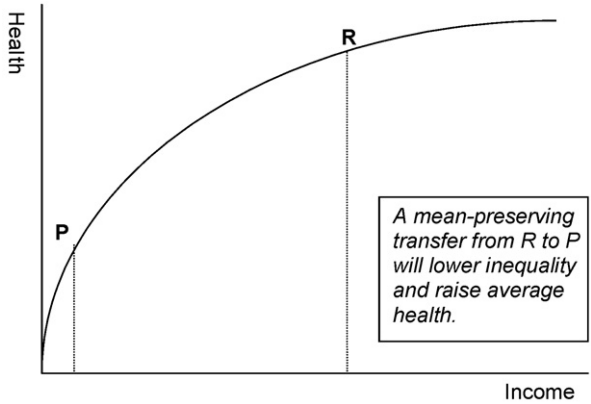

Fig. 1. A non-linear relationship between income and health.

2.1. The absolute income hypothesis

Even if health depends only on individual income, standard economic models predict that the health gains from an extra unit of income should diminish as income rises. Fig. 1 shows a stylized version of such a relationship. A mean-preserving transfer from the richer individual (R) to the poorer individual (P) raises the health of P by more than it lowers the health of R. Holding total income constant, therefore, a more equal distribution of income should improve population health.

When one compares countries, the relationship between average income and average life expectancy follows the pattern shown in Fig. 1 (Preston, 1975; Deaton, 2003). The relationship is almost flat among countries with incomes more than half the US average. When one compares individuals within the US, the relationship between family income and age-specific mortality again exhibits this pattern (Backlund et al., 1996). However, neither the comparisons of countries nor the comparisons of individuals take account of all the factors that could affect both income and mortality. Nor do they take account of reverse causation (poor health lowers income). Thus while there are strong theoretical reasons for expecting additional income to have less effect on health as income rises, the empirical evidence supporting this assumption is not conclusive.

2.2. The relative income hypothesis

Holding individual income constant, the income of others can affect people's health if they evaluate either their income or their lives as a whole by comparing themselves to others. The relative income hypothesis assumes that, at least in the economic domain, upward comparisons are more salient than downward comparisons and that upward comparisons are more likely to be stressful than soothing.[1] Wilkinson (1997) argues, for example, that if individuals assess their well-being by comparing themselves to others with more income than themselves, increases in economic inequality will engender "[l]ow control, insecurity, and loss of self esteem". When

[1] Upward comparisons can be soothing if they lead people whose current economic circumstances are stressful to think that their future circumstances could be better. We are not aware of any persuasive evidence on whether upward economic comparisons are more salient than downward comparisons or on whether they are more often upsetting than soothing.

upward economic comparisons are distressing, they are said to produce "relative deprivation". The most frequently suggested physical mechanism linking relative deprivation to mortality is chronic stress, which appears to lower resistance to many forms of disease in a variety of species.[2]

One potential objection to this hypothesis is that many studies of relative deprivation suggest that social comparisons are most stressful when they involve people who have a lot in common, such as co-workers, relatives, and neighbors.[3] Income differences within such reference groups are likely to be much smaller than differences between random members of national populations. Nonetheless, when income inequality changes in society as a whole, it is also likely to change in the same direction within reference groups composed of co-workers, relatives, or neighbors. If changes in inequality lead to parallel changes in chronic stress, increased inequality at the national level could increase mortality and decreased inequality could reduce mortality.[4]

2.3. Society-wide effects of inequality

(a) *Violent crime*: Fajnzylber et al. (2002) report a robust relationship between economic inequality and violent crime in a sample of rich and poor countries. Violent crime accounts for a tiny fraction of all deaths in the countries we study, but it could have larger second-order effects on mortality if it increases chronic stress among those who worry that they or their kin might become victims in the future.

(b) *Public spending*: if the Meltzer–Richard theorem is correct, greater economic inequality among voters should make the median voter more inclined to support government spending on health (Meltzer and Richard, 1981). Szreter (1988) shows, for example, that clean water was made available in much of the UK only after the franchise was extended to include the less affluent, for whom public spending on sanitation provided large health benefits at little cost to themselves.[5] However, Alesina et al. (1999) show that the average value of public goods to members of a community will decrease when heterogeneity increases. If income inequality makes voters' preferences more heterogeneous, it could lower government spending on health for that reason as well. In addition, increases in economic inequality may allow the rich buy more political influence, which could lead to reductions in government spending on health.

(c) *Social capital and trust*: comparing American states, Kawachi et al. (1997) find negative cross-sectional relationships both between inequality and social capital and between social capital and mortality. Other studies have also found that people in more unequal places tend to be less trusting (Knack and Keefer, 1997; Alesina and La Ferrara, 2002; Leigh, in press). Low trust may make voters more skeptical about the claim that public spending will improve health. Low trust may also be linked to thinner friendship networks, which are associated with higher age-specific mortality (Berkman and Syme, 1979).

[2] Marmot (2005) provides numerous references.

[3] Martin (1981) provides some relevant references for earnings.

[4] The relative income hypothesis also comes in a more extreme variant, where all that matters is ordinal rank, not the distance between ranks. In that variant any income hierarchy in which every individual has a unique income has the same effect as any other.

[5] The social capital and public expenditure channels do not posit any specific link between an individual's position in the income distribution and his or her health. Instead, they suggest that greater variance of incomes will adversely affect population health. This impact could conceivably affect individuals whose income is below, above, or at the mean.

Because this paper relies entirely on aggregate data about income inequality and mortality, it cannot distinguish empirically between the absolute income hypothesis, the relative income hypothesis, and hypotheses that involve society-wide effects of inequality. One way to see this is to combine the absolute and relative income hypotheses algebraically. Following Gravelle et al. (2002), we begin by assuming that absolute individual income, y, is the only factor affecting an individual's mortality risk, $m(y)$. The expression $m(y)$ can be expressed in terms of individual income y and mean income $\bar{y}$ through the following second order approximation:

$$m(y) \approx m(\bar{y}) + m'(\bar{y})(y - \bar{y}) + \tfrac{1}{2}m''(\bar{y})(y - \bar{y})^2 \tag{1}$$

We now introduce the relative income hypothesis. Deaton (2001) shows that under plausible assumptions an individual's experience of relative economic deprivation depends on the fraction of total income received by individuals' with higher incomes. He also shows that when that is the case the mean level of relative deprivation will be a linear function of the Gini coefficient for household income. Here, however, we follow Gravelle et al. and assume that the mean level of relative deprivation is a linear function of the variance of income, $V(y)$, rather than the Gini coefficient:

$$m(y) \approx \left\{m(\bar{y}) + m'(\bar{y})(y - \bar{y}) + \tfrac{1}{2}m''(\bar{y})(y - \bar{y})^2\right\} + \alpha V(y) \tag{2}$$

Taking expectations of each side:

$$Em(y) \approx m(\bar{y}) + m'(\bar{y})E(y - \bar{y}) + \tfrac{1}{2}m''(\bar{y})E(y - \bar{y})^2 + \alpha EV(y) \tag{3}$$

which simplifies to:

$$Em(y) \approx m(\bar{y}) + \tfrac{1}{2}m''(\bar{y})V(y) + \alpha EV(y) \tag{4}$$

Eq. (4) shows that:

- If both the second derivative of mortality with respect to individual income and α (the effect of income inequality via relative deprivation) are positive, there will be a positive relationship between mortality and inequality, $V(y)$.
- Unless we can estimate m'' using micro data, separating the contributions of absolute and relative income to the overall relationship between income inequality and mortality will be impossible.[6]

Our analysis therefore focuses on reduced-form models that estimate the net effect of a change in income inequality on various measures of mortality.

As Gravelle et al. (2002) point out, one recurrent problem when estimating the relationship between economic inequality and health is that we cannot be certain whether any particular measure of inequality captures the full effect of inequality on health. There is no consensus on which dimensions of economic inequality are most likely to affect population health, but most investigators prefer measures such as the Gini coefficient that are sensitive to changes in income

[6] Miller (2001) has shown that this argument holds only if the second-order approximation in Eq. (1) is exact. But while the second-order approximation is unlikely to be exact, we would need much better data than we have to distinguish the two effects.

shares at both the top and bottom of the distribution.[7] We use the income share of the top decile because it gives us far more observations per country than any previous study. However, this choice is not costless. As we shall see, our measure of inequality is a reasonable proxy for other changes in the top half of the distribution, but it is not a good proxy for changes in the bottom half of the distribution. We return to this issue in Section 3.

One final concern is worth noting. Since adults are less likely to work when they are sick or have to care for sick children, countries with worse overall health may have more unequal family incomes for that reason alone. The causal links between inequality and health can therefore run either from health to inequality or from inequality to health. With a long time series such as ours, we could use Granger causality tests to see whether lagged inequality affected current health or lagged health affected current inequality. But since we find no statistically significant relationship between inequality and health, we do not pursue this approach.

3. Data on inequality and health

Data quality has been a major problem in studies of the relationship between income inequality and health. As Judge, Mulligan and Benzeval (1998:569) note in their review of the literature:

> "Many of the studies use multiple sources of income distribution data and/or data from a wide range of years, which makes comparability between countries questionable. Only five of the studies use data based on a measure of equivalent disposable income. In fact, we believe it is the generally poor quality of the income data that poses the most serious weakness in most of the studies we have reviewed."

Most cross-national studies have used measures of inequality from the Deininger and Squire dataset or the World Income Inequality Database (WIID). Atkinson and Brandolini (2001) have shown that using higher-quality inequality data can substantially alter results based on these two sources. Judge et al. use data from the Luxembourg Income Study (LIS), which provides a more consistent measure of income for all countries and years, namely disposable household income adjusted for household size. They find no significant relationship between inequality and either life expectancy or infant mortality. However, their sample includes only 16 countries, and only 10 of these countries have data for more than 1 year. Even if there were a causal relationship between inequality and health, it would be hard to detect in samples this small, particularly if one tries to control other relevant factors.

In order to get more precise estimates we need more observations over longer periods of time in more countries. Annual estimates of the share of pretax income received by the richest 10% of the population (abbreviated as Share 10) are now available for 12 rich countries, with an average of 62 years per country. Table 1 shows the periods covered in each country. The estimates are derived by comparing the amount of income reported to the tax authorities by the richest 10% of individuals/households to an estimate of total personal income in the same year taken from each country's national accounts (see Leigh, 2006).[8] Both the numerator and the denominator of

[7] Exceptions are Waldmann (1992), who uses the share going to the richest 5% and Wilkinson (1992) who uses the share going to the poorest 60%.

[8] Leigh (2006) provides data for 13 countries. Here, we exclude Japan because information on the income share of the top 10% is unavailable. See also the comparable database in Atkinson and Piketty (in press).

Table 1
Country–years covered by the sample

Country	Number of years	Period(s)
Australia	61	1942–2002
Canada	60	1941–2000
France	81	1905 and 1919–1998
Germany	38	1961–1998
Ireland	30	1939–1943 and 1976–2000
The Netherlands	86	1914–1999
New Zealand	78	1925–2002
Spain	22	1981–2002
Sweden	87	1903–1920, 1930–1935 and 1941–2003
Switzerland	64	1933–1996
United Kingdom	52	1919 and 1950–2000
United States	85	1917–2001
Total	744	1903–2003

this ratio are, of course, subject to error, especially in the years before 1945, when estimates of national income are not very precise.

Our measure of inequality is particularly sensitive to changes at the top of the distribution. Regressing the share going to the top 10% on the share going to the top 1% in a specification that includes country and year fixed effects, the coefficient of the top 1% share is 1.35 ($t = 35.0$). However, changes in the share going to the top 10% also predict changes in the 90–50 ratio. Leigh (2006) regresses the income share of the top 10% on LIS measures of inequality, which are largely derived from household surveys and are then processed in a consistent way. When country and year fixed effects are included, the log of the top decile's share of pretax income is strongly related to the log of the 90/50 ratio for size-adjusted post-tax income in LIS ($\beta = 0.843$; $t = 6.7$). The logged share of the top 10% is also fairly strongly related to the log of the LIS Gini coefficient for size-adjusted post-tax income ($\beta = 0.610$; $t = 5.6$). However, the share going to the top 10% is not related to the log of the 50/10 ratio in LIS ($\beta = -0.077$; $t = -0.4$).

We use two standard measures of population health: life expectancy at birth and infant mortality. Life expectancy in a given year measures the weighted probability that individuals of different ages died in that year, with deaths at younger ages getting more weight than deaths at older ages. Infant mortality is the proportion of children under the age of one who died in a given year (excluding still-births). Because infant mortality asymptotes towards zero as population health improves, we use its logarithm as our dependent variable. Infant mortality and life expectancy are highly correlated, partly because infant mortality has a direct effect on life expectancy and partly because they have common causes. Regressing life expectancy on log infant mortality in a specification with country and year fixed effects, the coefficient on log infant mortality is −6.68 ($t = -11.5$). Without country and year fixed effects, R^2 is 0.89.

All our specifications control real per capita GDP, measured in 1990 International Geary-Khamis dollars. A few specifications also control the average educational attainment of the adult population, the log of per capita government spending on health spending, and the log of per capita private spending on health (in constant US dollars). The last three measures are not available until 1960.

Infant mortality, life expectancy, homicide, suicide, and our measure of inequality (Share 10) are interpolated when they are missing for short periods, but they are not extrapolated beyond the

Table 2
Summary statistics

	Mean	S.D.	*N*
Income share of richest 10%	33.693	6.100	744
Average life expectancy at birth (years)	70.740	6.569	744
Log infant mortality rate (per 1000 live births)	2.879	0.855	739
Log homicide rate (per 100,000 people)	0.226	0.719	536
Log suicide rate (per 100,000 people)	2.512	0.375	538
Real GDP per capita ($ 1000s)	11.406	5.345	744
Average years of education of adults aged 15+	9.081	1.682	455
Log real public health spending per capita (converted to 2003 $US at PPP)	6.878	0.582	438
Log real private health spending per capita (converted to 2003 $US at PPP)	5.970	0.791	447

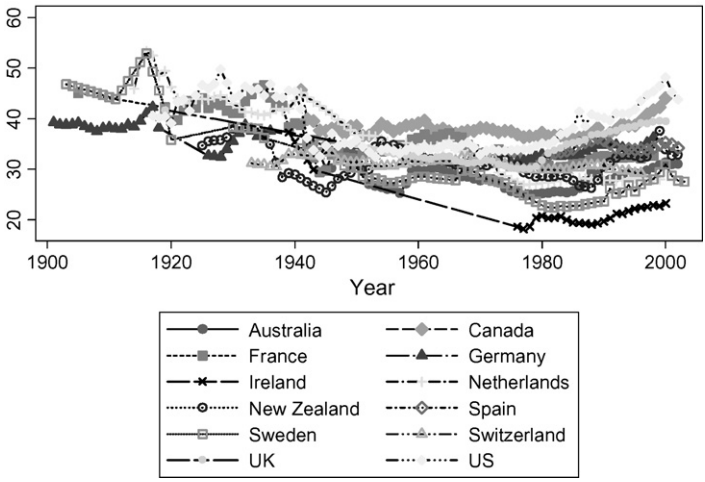

Fig. 2. Income share of richest 10%.

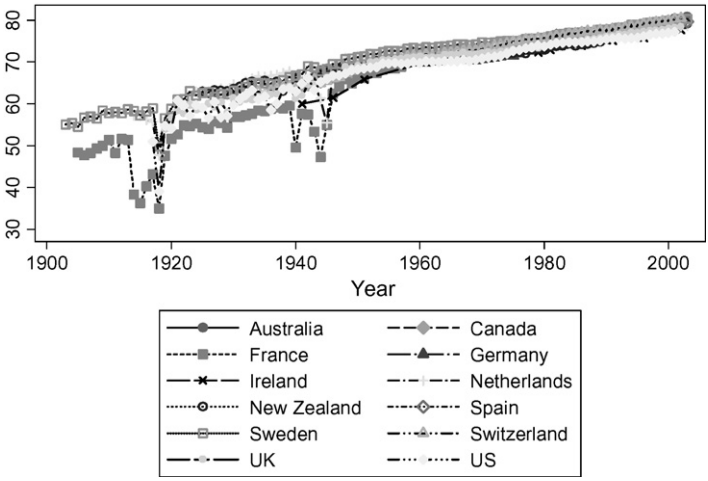

Fig. 3. Life expectancy at birth.

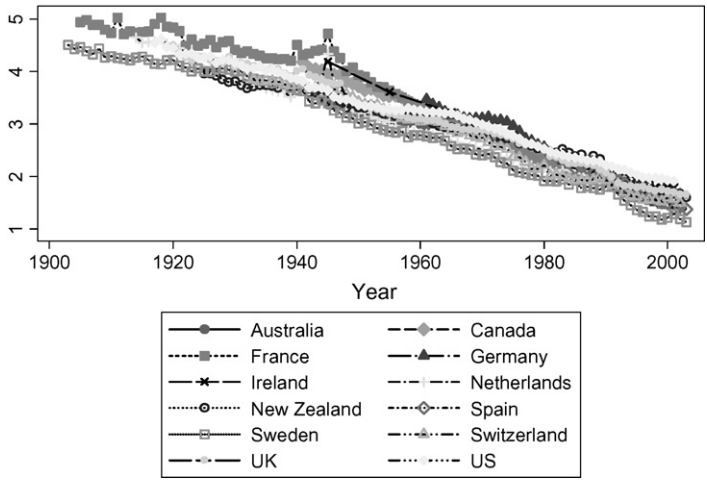

Fig. 4. Log infant mortality rate (per 1000 births).

earliest or latest available year. Our other control variables are not only interpolated when they are missing for short periods but extrapolated forward when they are missing in the last few years of our time series. Appendix A gives further details on variable construction. Table 2 presents summary statistics.

Figs. 2–4 show trends in inequality and health for each country. Share 10 varies more at the end of the 20th century than earlier, while life expectancy varies less at the end of the 20th century than earlier. The log of infant mortality shows no clear trend. Figs. 5 and 6 show the cross-sectional relationship between inequality and mortality in 1996, which is the last year for which all 12 countries had reported these statistics. For life expectancy, the cross-sectional gradient is negative but statistically insignificant. For infant mortality, the gradient is positive but statistically insignificant.

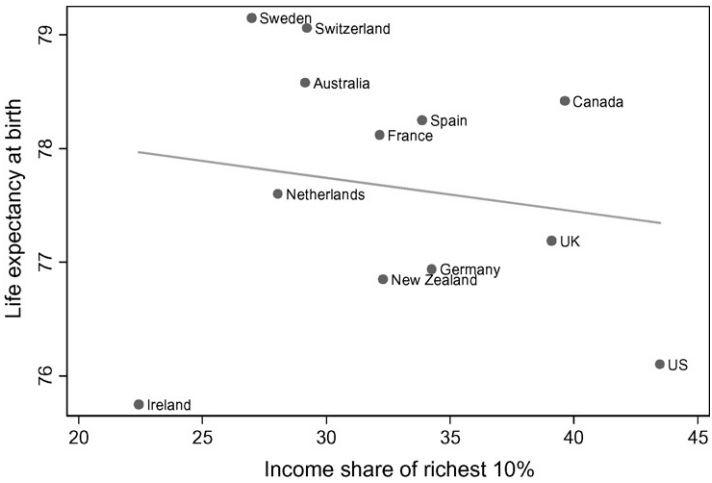

Fig. 5. Cross-sectional relationship between top incomes and life expectancy (1996).

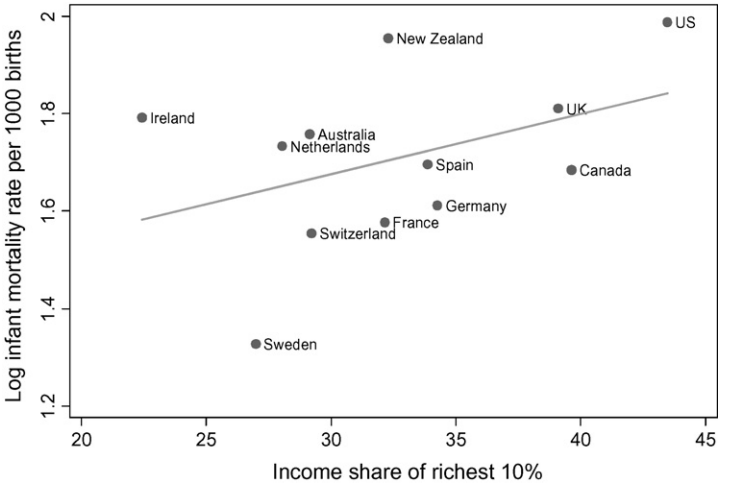

Fig. 6. Cross-sectional relationship between top incomes and infant mortality (1996).

4. Empirical strategy and results

Most of the existing literature relies on comparisons across countries at a single point in time or on changes over time within one or two countries. We begin by estimating an equation similar to the one often reported in this literature:

$$m_{jt} = \alpha + \beta(\text{Share 10})_{jt} + \gamma Z_{jt} + \varepsilon_{jt} \tag{5}$$

where m is the measure of mortality (life expectancy or infant mortality) for country j in year t, Share 10 the income share of the richest 10% of the population, Z the real GDP per capita, and ε is an error term. Standard errors are clustered at the country level to take account of possible within-country serial correlation.

Columns 1 and 2 of Table 3 presents results from this specification, which does not include country or year fixed effects. For both life expectancy and infant mortality, higher inequality is associated with significantly higher mortality. This accords with Waldmann (1992), who finds a strong positive relationship between infant mortality and the income share of the richest 5% of the population in a cross-sectional regression.

However, there are good reasons to think that we need to take account of country-specific and time-specific factors. We therefore move progressively towards estimating an equation of the following form:

$$m_{jt} = \alpha + \beta(\text{Share 10})_{jt} + \gamma Z_{jt} + \delta_j + \rho_t + \varepsilon_{jt} \tag{6}$$

in which δ is a country fixed effect, and ρ is a year fixed effect.

Country fixed effects capture stable mortality differences between countries. Column 3 of Table 3 shows that adding country fixed effects barely alters the coefficient of Share 10 in the life expectancy equation. Evidently the unobserved country characteristics that produce stable differences in life expectancy are not strongly correlated with economic inequality. However, Column 4 shows that introducing country fixed effects appreciably reduces the estimated effect of Share 10 on infant mortality. Because adding country fixed effects also reduces the standard errors in the infant mortality equation, the coefficient on Share 10 remains highly significant.

Table 3
Top 10% share and health: levels specification with and without fixed effects

	Dependent variable							
	[1] LE	[2] IM	[3] LE	[4] IM	[5] LE	[6] IM	[7] LE	[8] IM
Income share of richest 10%	-0.352^{***} [0.077]	0.041^{***} [0.009]	-0.382^{***} [0.042]	0.028^{***} [0.006]	-0.054 [0.065]	0.018^{**} [0.008]	0.016 [0.067]	0.006 [0.007]
Real GDP per capita ($ 1000s)	0.874^{***} [0.089]	-0.125^{***} [0.013]	0.884^{***} [0.068]	-0.138^{***} [0.010]	0.167 [0.169]	-0.02 [0.018]	0.356^{*} [0.168]	-0.036 [0.023]
Country FE?	No	No	Yes	Yes	No	No	Yes	Yes
F-Test			9.5^{***}	21.4^{***}			590000^{***}	51671^{***}
Year FE?	No	No	No	No	Yes	Yes	Yes	Yes
F-Test					264^{***}	573^{***}	1304^{***}	64^{***}
Observations	744	739	744	739	744	739	744	739
R^2	0.76	0.84	0.86	0.94	0.91	0.95	0.96	0.98

Note: Robust standard errors, clustered at the country level, in brackets. *, ** and *** denote statistical significance at the 10%, 5% and 1% levels, respectively. Dependent variables: LE is average life expectancy at birth, IM is the log of the infant mortality rate (per 1000 live births). *F*-Test is a test for the joint significance of the country fixed effects or year fixed effects.

Year fixed effects capture the influence of shocks that affect mortality in multiple countries at the same time. Examples might include major influenza epidemics, the spread of HIV/AIDS, the introduction of new vaccines, and the diffusion of antibiotics.[9] Adding year fixed effects makes the relationship between Share 10 and life expectancy both substantively and statistically insignificant (see Column 5). Adding year fixed effects also cuts the apparent effect of Share 10 on infant mortality, although the coefficient is still significant at the 5% level (see Column 6).

Columns 7 and 8 include both country and year fixed effects simultaneously. Both country and year fixed effects remain highly significant, but the coefficient on Share 10 is now small and insignificant for both life expectancy and infant mortality. Indeed, the coefficient on Share 10 in the life expectancy equation changes its sign. The coefficients on Share 10 do not become insignificant because their standard errors increase. (The standard errors in columns 7 and 8 are smaller than those in columns 1 and 2.) Rather, the coefficients on Share 10 become insignificant because they are very close to zero once we add country and year fixed effects. The coefficients without fixed effects are also well outside the 99% confidence intervals for the coefficients with fixed effects (compare columns 1 and 2 to columns 7 and 8).[10]

Table 3 shows that even with 744 country-year observations the relationship between income inequality and mortality is not robust to including both country and year fixed effects. This finding suggests that time-series correlations between inequality and mortality in studies of only one or two rich countries may be inflated by unobserved factors that affect both inequality and mortality in many rich countries simultaneously.[11]

Table 4 presents several sensitivity tests. Columns 1 and 2 of Table 4 are identical to columns 7 and 8 of Table 3. Columns 3 and 4 show that when we add GDP^2 to the basic specifications in columns 1 and 2, GDP^2 is significant for both life expectancy and infant mortality. Its sign indicates that the protective effect of additional GDP diminishes as GDP rises, which is consistent with the cross-sectional results in Preston (1975) and Deaton (2003). However, including GDP^2 does not change our conclusions about the effect of income inequality. The coefficient on Share 10 remains insignificant, and the sign is now reversed for both life expectancy and infant mortality.

One potential objection to the results presented thus far is that the time series may actually be "too long". We know that the proximate causes of mortality changed dramatically over the course of the 20th century, and the effects of economic inequality could have changed as a result. In England and Wales, for example, infectious diseases accounted for 55% of all deaths in 1901, but only 18% in 1961 (McKeown et al., 1975). This epidemiologic transition, which occurred

[9] Technological innovations do not, of course, reach all developed countries in exactly the same year. Deaton and Paxson (2001) argue, for example, that technological innovations tend to affect the UK about 4 years later than the US. However, we cannot allow year fixed effects to vary randomly by country, since they would then explain all of the variation we use to identify the effect of changes in inequality. However, our results are robust to excluding the years 1914–1919 and 1939–1945, which are the periods in which year effects vary most across countries.

[10] In principle, our use of year and country fixed effects is similar to that of Judge et al. (1998). Their Table 5 analyzes changes in the 10 countries for which they have two estimates of inequality: Australia 1981–1985, Belgium 1985–1988, Canada 1981–1987, Finland 1987–1990, France 1979–1984, Netherlands 1983–1987, Norway 1979–1986, Sweden 1981–1987, UK 1979–1986, and US 1979–1986. We have top incomes data for seven of these countries (Australia, Canada, France, Netherlands, Sweden, UK and US). We estimated the correlation between changes in life expectancy and changes in the top 10% share in these seven countries using the same time intervals as Judge et al. The correlation between ΔLE and ΔS10 was 0.227 (95% CI: −0.635 to 0.837), while the correlation between ΔIM and ΔS10 was −0.663 (95% CI: −0.945 to 0.180).

[11] For a discussion of the same issue in a different context, see Acemoglu et al. (2005).

Table 4
Top 10% share and health: levels

	Dependent variable							
	[1] LE	[2] IM	[3] LE	[4] IM	[5] LE	[6] IM	[7] LE	[8] IM
Income share of richest 10%	0.016 [0.063]	0.006 [0.006]	0.096 [0.089]	−0.003 [0.006]	0.054 [0.063]	−0.012 [0.010]	0.033 [0.057]	−0.010 [0.010]
Real GDP per capita ($ 1000s)	0.356** [0.156]	−0.036 [0.021]	1.199** [0.522]	−0.138*** [0.042]	0.199 [0.124]	−0.083** [0.032]	0.17 [0.119]	−0.075* [0.036]
Real GDP per capita squared ($ 1000s)			−0.027* [0.014]	0.003** [0.001]	−0.006 [0.004]	0.002** [0.001]	−0.005 [0.004]	0.002** [0.001]
Average years of education							−0.332 [0.288]	−0.050 [0.028]
Log real public health spending per capita							0.105 [0.234]	−0.049 [0.044]
Log real private health spending per capita							0.295 [0.285]	−0.050 [0.083]
Country and year FE?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	744	739	744	739	430	430	430	430
R^2	0.96	0.98	0.96	0.98	0.97	0.97	0.97	0.97

Note: Robust standard errors, clustered at the country level, in brackets. *, ** and *** denote statistical significance at the 10%, 5% and 1% levels, respectively. Dependent variables: LE is average life expectancy at birth, IM is the log of the infant mortality rate (per 1000 live births). Columns 5 and 6 are restricted to those country–years for which education and health spending variables are available.

throughout the developed world, accounted for most of the increase in life expectancy between 1901 and 1961. Once it was complete, increases in life expectancy slowed and became largely dependent on progress in reducing mortality from degenerative conditions like heart disease and cancer. If income inequality affected mortality from infectious and degenerative diseases differently, combining data for these two periods could be misleading.

Columns 5 and 6 drop the years prior to 1960, leaving only the 430 country–year observations between 1960 and 2003. Because annual changes in life expectancy are smaller after 1960 than before, the coefficients in Eq. (5) are both smaller and more precisely estimated than those in Eq. (3). Nonetheless, the coefficient of Share 10 is still insignificant and still has the "wrong" sign. Nor does variation in the rate of growth in GDP have a significant effect on life expectancy after 1960 in these rich countries.[12]

Because we log the infant mortality rate, its variance does not fall as the absolute rate approaches zero. The coefficient of Share 10 when predicting infant mortality is larger in the post-1960 sample than in the full sample, but the point estimate still suggests that increases in income inequality are more likely to be associated with reductions in infant mortality than with increases, and the confidence interval still includes zero. GDP and GDP^2 are also still significant. In short, we find no clear evidence that the effect of either GDP or income inequality on either life expectancy or infant mortality looks different when we focus on the years since 1960 than when we look at a longer time period.

Columns 7 and 8 include both public and private health expenditures, as well as the mean number of years of schooling completed by the adult population. Adding these potentially endogenous controls does reduce the absolute size of the coefficients on Share 10, especially when predicting life expectancy. But since the confidence intervals of all the relevant coefficients include zero, the most parsimonious explanation is that all these changes are due to random noise.

The models in Tables 3 and 4 treat the relationship between income inequality and mortality as if it were almost instantaneous. Although the literature on inequality and mortality often makes this assumption, it is not entirely plausible. Some of the hypotheses we have discussed do suggest that the lag between a change in inequality and a change in mortality could be quite short. If inequality were positively related to violent deaths, for example, the mortality effect might show up in the same year as the income change, although any second-order effects on levels could take longer to kill people. Likewise, if rising inequality leads to a decline in the absolute income of the poor, mortality might rise in the same year, especially among newborns.[13] But some of the hypotheses we have discussed suggest that the lag between a change in inequality and a change in mortality could be fairly long. If inequality affects health by weakening the social fabric or by intensifying feelings of relative deprivation, for example, these effects could take some years to influence mortality.

To assess the importance of lags we estimate a variant of Eq. (6) that includes inequality lagged by n years:

$$m_{jt} = \alpha + \beta(\text{Share 10})_{jt-n} + \gamma Z_{jt-n} + \delta_j + \rho_t + \varepsilon_{jt} \tag{7}$$

We estimate Eq. (7) using various combinations of lags, ranging from 1 to 5 years. Table 5 presents regressions using 1-year lagged inequality (columns 1 and 2), 1-year and 2-year lagged inequality

[12] The point estimate is actually negative once per capita GDP exceeds $17,000 in 1990 dollars—a level that most countries in this sample had achieved by 1990.

[13] Because our measure of income inequality is not a good proxy for the absolute income of the poor, we would not be likely to detect this particular effect even if it existed.

Table 5
Top 10% share and health: lags (1 year)

	Dependent variable					
	[1] LE	[2] IM	[3] LE	[4] IM	[5] LE	[6] IM
Income share of richest 10% ($t-1$)	0.011 [0.066]	0.006 [0.007]	0.177 [0.144]	0.002 [0.008]	0.162 [0.141]	0.001 [0.009]
Income share of richest 10% ($t-2$)			−0.188 [0.116]	0.003 [0.007]	−0.163** [0.073]	0.004* [0.002]
Income share of richest 10% ($t-3$)					−0.057 [0.041]	0.002 [0.004]
Income share of richest 10% ($t-4$)					0.013 [0.069]	0.003 [0.006]
Income share of richest 10% ($t-5$)					0.033 [0.112]	−0.007 [0.009]
Real GDP per capita ($ 1000s) ($t-1$)	0.339** [0.149]	−0.037 [0.023]	0.641* [0.348]	−0.006 [0.025]	0.601* [0.296]	−0.021 [0.023]
Real GDP per capita ($ 1000s) ($t-2$)			−0.321 [0.295]	−0.033 [0.028]	−0.173 [0.183]	0 [0.020]
Real GDP per capita ($ 1000s) ($t-3$)					0.005 [0.157]	0.003 [0.009]
Real GDP per capita ($ 1000s) ($t-4$)					0.098 [0.205]	−0.014 [0.011]
Real GDP per capita ($ 1000s) ($t-5$)					−0.28 [0.223]	−0.007 [0.024]
Country and year FE?	Yes	Yes	Yes	Yes	Yes	Yes
Observations	742	736	725	720	680	678
R^2	0.96	0.98	0.95	0.98	0.95	0.98
Sum of lagged inequality coefficients			−0.010 [0.068]	0.005 [0.007]	−0.11 [0.083]	0.003 [0.008]

Note: Robust standard errors, clustered at the country level, in brackets. * and ** denote statistical significance at the 10% and 5% levels, respectively. None of the sums of lagged coefficients are statistically significant. Dependent variables: LE is average life expectancy at birth, IM is the log of the infant mortality rate (per 1000 live births).

Table 6
Top 10% share and health: lags (GDP AND GDP^2)

	Dependent variable					
	[1] LE	[2] IM	[3] LE	[4] IM	[5] LE	[6] IM
Income share of richest 10% ($t-1$)	0.083 [0.090]	−0.004 [0.007]	0.263* [0.138]	−0.008 [0.005]	0.260* [0.128]	−0.012** [0.005]
Income share of richest 10% ($t-2$)			−0.201* [0.105]	0.003 [0.009]	−0.164* [0.082]	0.004 [0.003]
Income share of richest 10% ($t-3$)					−0.049 [0.042]	0.001 [0.004]
Income share of richest 10% ($t-4$)					0.019 [0.065]	0.003 [0.004]
Income share of richest 10% ($t-5$)					0.002 [0.094]	−0.005 [0.008]
Real GDP per capita ($ 1000s) ($t-1$)	1.112** [0.502]	−0.141** [0.047]	3.127** [1.023]	0.234*** [0.054]	3.120*** [0.872]	0.282*** [0.048]
Real GDP per capita ($ 1000s) ($t-2$)			−2.018** [0.654]	0.082 [0.058]	−1.359* [0.685]	0.113*** [0.033]
Real GDP per capita ($ 1000s) ($t-3$)					0.348 [0.349]	−0.036* [0.019]
Real GDP per capita ($ 1000s) ($t-4$)					−0.076 [0.609]	0.005 [0.022]
Real GDP per capita ($ 1000s) ($t-5$)					−1.017*** [0.200]	0.028 [0.027]
Real GDP per capita squared ($ 1000s) ($t-1$)	−0.025* [0.014]	0.003** [0.001]	−0.110** [0.036]	0.010*** [0.002]	0.106*** [0.030]	0.011*** [0.002]
Real GDP per capita squared ($ 1000s) ($t-2$)			0.089*** [0.027]	−0.007** [0.003]	0.055** [0.025]	0.005*** [0.001]
Real GDP per capita squared ($ 1000s) ($t-3$)					−0.014 [0.012]	0.002 [0.001]
Real GDP per capita squared ($ 1000s) ($t-4$)					0.007 [0.021]	−0.001 [0.001]
Real GDP per capita squared ($ 1000s) ($t-5$)					0.043*** [0.009]	−0.003*** [0.001]
Country and year FE?	Yes	Yes	Yes	Yes	Yes	Yes
Observations	742	736	725	720	680	678
R^2	0.96	0.98	0.96	0.98	0.96	0.98
Sum of lagged inequality coefficients			0.062 [0.085]	−0.005 [0.007]	0.068 [0.086]	−0.009 [0.007]

Note: Robust standard errors, clustered at the country level, in brackets. *, ** and *** denote statistical significance at the 10%, 5% and 1% levels, respectively. None of the sums of lagged coefficients are statistically significant. Dependent variables: LE is average life expectancy at birth, IM is the log of the infant mortality rate (per 1000 live births).

Table 7
Top 10% share and homicide/suicide

	Dependent variable					
	[1] HOM	[2] SUI	[3] HOM (1-year lag)	[4] SUI (1-year lag)	[5] HOM (5-year lag)	[6] SUI (5-year lag)
Income share of richest 10%	−0.043* [0.020]	−0.021 [0.012]	−0.045** [0.020]	−0.021 [0.012]	−0.042* [0.019]	−0.018 [0.011]
Real GDP per capita ($ 1000s)	−0.065 [0.060]	0.073 [0.063]	−0.074 [0.061]	0.073 [0.068]	−0.084 [0.048]	0.034 [0.051]
Real GDP per capita squared ($ 1000s)	0.002 [0.001]	−0.001 [0.001]	0.002 [0.001]	−0.001 [0.001]	0.002 [0.001]	0 [0.001]
Country and year FE?	Yes	Yes	Yes	Yes	Yes	Yes
Observations	536	538	538	541	531	538
R^2	0.91	0.85	0.91	0.85	0.91	0.84

Note: Robust standard errors, clustered at the country level, in brackets. * and ** denote statistical significance at the 10% and 5% levels, respectively. Dependent variables: HOM is the log of the homicide rate per 100,000 people; SUI is the log of the suicide rate per 100,000 people. For lagged specifications (columns 3–6), all independent variables are lagged over the same interval.

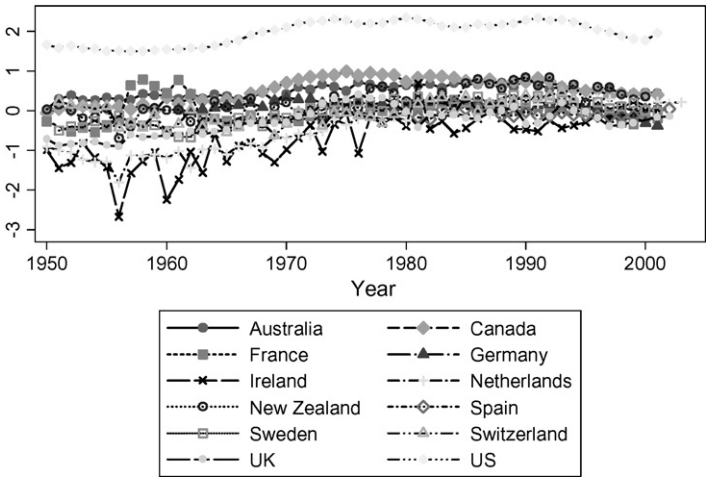

Fig. 7. Log homicide rate (per 100,000 people).

(columns 3 and 4) and 1-year through 5-year lagged inequality (columns 5 and 6). Table 6 repeats this exercise including lagged GDP and GDP^2. In each case we estimate the linear sum of the Share 10 coefficients, which is a summary measure of the overall impact of lagged inequality on current mortality. Using only 1-year lags, we find no statistically significant relationship between inequality and mortality. When we include multiple lagged terms, specific lags are sometimes statistically significant, but their implied effects are offset by the fact that other lags have the opposite sign. Summing 1-year and 2-year lags, or 1-year through 5-year lags, we found no significant relationship between inequality and mortality.

Overall, we find no robust relationship between changes in current mortality and changes in economic inequality over the past 5 years. The coefficients of the controls, in contrast, mostly accord with expectations. Increases in GDP are associated with increases in life expectancy

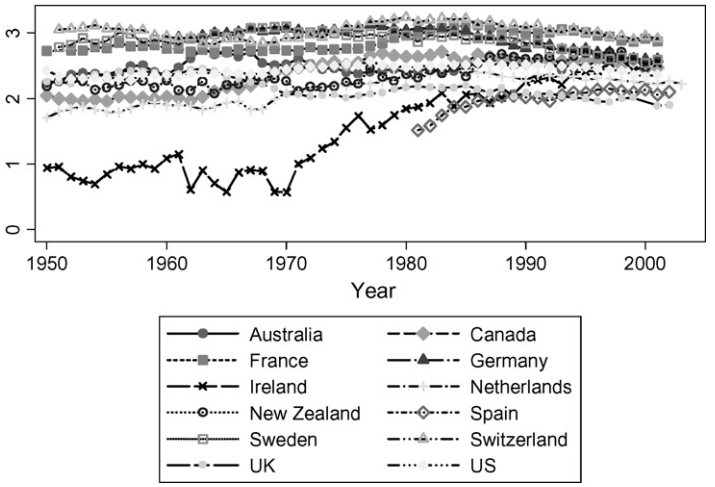

Fig. 8. Log suicide rate (per 100,000 people).

and reductions in infant mortality. Both effects diminish as GDP rises. Increases in public and private health spending are associated with lower mortality, but the coefficients are not statistically significant.[14]

Some researchers have suggested that even if there is no relationship between inequality and overall mortality, there may be a relationship between inequality and homicide (Deaton, 2003; Lynch et al., 2004a) or suicide (Lynch et al., 2004b). To test this hypothesis we calculated annual homicide and suicide rates for each country starting in 1950, which is the first year in which such data are available for our full sample. Appendix A provides details on variable construction. Figs. 7 and 8 depict our homicide and suicide series.

Table 7 shows the relationship of Share 10 to the homicide and suicide rates per 100,000 people. As with infant mortality, we log the homicide and suicide rates, since they must asymptote towards zero. We tested their relationship to Share 10 using fixed effects specifications, with levels, with 1-year lags, and with 5-year lags. The coefficient of inequality was always negative, suggesting that more inequality was associated with a lower suicide or homicide rate. For homicide, the negative relationship was statistically significant at the 5% or 10% level, depending upon the specification. Overall, the results in Table 7 are inconsistent with the hypothesis that higher top income shares lead to more homicide or suicide.

5. Conclusion

While there is a strong consensus in the literature that the correlation between income and health is positive, there is much less agreement about the relationship between income inequality and health. This paper has used a new measure of inequality – the income share of the richest 10% of the population – to test the relationship between inequality and mortality. Because we have longer time series for more countries than past studies, we have tried to circumvent some of the problems that plagued these studies by holding constant country and year fixed effects.

We have shown that higher GDP is associated with lower mortality, and that this effect declines as GDP rises. Without country and year fixed effects, we also find that more inequality is associated with higher mortality. But once we include country and year fixed effects the relationship between inequality and health becomes small and statistically insignificant. This finding is consistent with a number of other careful cross-country papers, such as Judge et al. (1998) and Deaton and Paxson (2001).

The confidence intervals around our estimates are sufficiently tight to make substantively important detrimental effects of inequality on population health unlikely. Consider the coefficients from the levels specifications controlling GDP and GDP^2, shown in Columns 3 and 4 of Table 4. For life expectancy, the 95% confidence interval for the effect of a 1% point increase in the income share of the top 10% includes no negative values larger than −0.1 years. This result implies that even a 10-point increase in the income share of the top decile would be unlikely to lower life expectancy by more than 1 year. For infant mortality, the upper bound of the 95% confidence interval is 0.01 (an additional 0.2 deaths per 1000, when evaluated at the mean). Our confidence intervals do not allow us to rule out the possibility that inequality *raises* life expectancy by a substantively significant amount. But because there is not much evidence for this hypothesis

[14] We find no systematic relationship between changes in adult educational attainment and changes in population health once we control GDP. However, year-to-year changes in the educational attainment of the adult population are inevitably small and may not be well measured.

elsewhere in the published literature, and because our confidence intervals include zero, a claim that growing inequality reduces mortality would be hard to defend.

The most important limitation of our findings is that we may not have measured the type of inequality that affects mortality. While changes in the top decile's share of pretax income are a reasonably good proxy for changes in both the 90–50 ratio and the Gini coefficient, changes in the top decile's share are not a good proxy for the 50/10 ratio. If the effects of inequality on mortality derive from the fact that increased inequality slows or reverses the long-term increase in living standards among the poor, our data would not detect this effect.

With the possible exception of the absolute income hypothesis, however, the theories we reviewed about how income inequality might affect mortality were not limited to households near the bottom of the income distribution. Relative deprivation, for example, can affect people anywhere in the income distribution except the very top, as long as long as they compare themselves to others who have even more money. The same is potentially true for the societal effects of violence or the effects of economic inequality on social networks and trust. If changes in inequality within the top half of the distribution do not affect mortality, these theories may need to be reconsidered. Three possibilities deserve consideration:

(1) The absolute income hypothesis, the relative income hypothesis, and the hypotheses linking inequality to society-wide changes that affect mortality are all false.
(2) The absolute income hypothesis is true only for those near the bottom of the income distribution in rich countries, and the other hypotheses are either false or have countervailing effects, so that their net impact on population health is negligible.
(3) The absolute and relative income hypotheses are both true, but the adverse effects of inequality on mortality that they predict are offset by other societal effects of inequality that reduce mortality.

We cannot distinguish between these possibilities with the data analyzed in this paper.

Acknowledgements

Thanks for valuable suggestions that helped improve earlier drafts to Philip Clarke, Andrew Clarkwest, David Cutler, Christian Dustmann, Richard Frank, Campbell Murray, Joseph Newhouse, Betsey Stevenson, S.V. Subramanian, and an anonymous reviewer, as well as seminar participants at the Australian National University, the University of Adelaide, the 2005 SOLE/EALE annual meetings, and the 2005 NBER Summer Institute.

Appendix A. Data appendix

A.1. Treatment of Germany

All series for Germany are for West Germany for 1956–1991, and for reunified Germany from 1992 onwards.

A.2. Sources of top incomes data

Top incomes data are from Leigh (2006), who adjusts top incomes series from 13 different papers to produce a comparable data set. Note that in Australia, Canada and Spain, the tax unit is the individual throughout the period of analysis, while in France, Ireland, the Netherlands, Switzerland and the United States, the tax unit is a married couple or single individuals throughout

the analysis. Germany has a hybrid system, with most taxpayers filing as tax units, and the very rich filing as individuals. New Zealand, Sweden and the United Kingdom all switched from household to individual filing. Where this appears to have caused substantial breaks in continuity, Leigh (2006) adjusts the top incomes series to take account of the changes.

A.3. Sources of life expectancy data

Most of our data on life expectancy at birth is taken from the Human Mortality Database (HMD), found at http://www.mortality.org. There are three exceptions:

- US data are from the National Vital Statistics Reports, Vol. 52, No. 14, February 18, 2004, Table 12 (found at http://www.cdc.gov/nchs/about/major/dvs/mortdata.htm). For 1900–1928, the figures are from death-registration states only. From 1929 onwards, they cover the entire US.
- Australian data are from Australian Bureau of Statistics, Australian Historical Population Statistics, ABS Catalog Number 3105.0.65.001, Table 48.
- Figures for Ireland are from Central Statistics Office (2004), Irish Life Tables No. 14, 2001–2003 (available at http://www.cso.ie).

The following should also be noted:

- In the case of New Zealand, life expectancy from the HMD is available for 1937 onwards for Maori, non-Maori, and the total population, and from 1876 onwards for non-Maori only. We use the ratio of Maori to non-Maori life expectancy in 1937 and 1938 to form a consistent life expectancy series for the entire population from 1876 to 1936. This method assumes that the ratio of Maori to non-Maori life expectancy was the same in the pre-1937 period as in 1937–1938.
- Although our inequality data cover the entire United Kingdom (including Ireland prior to 1921), the HMD only provides mortality figures for England and Wales (omitting Scotland and Northern Ireland). For the period 1999–2002, we update the HMD figures using National Statistics, "Life expectancy at birth by health and local authorities in the United Kingdom 1991–1993 to 2001–2003, including revised results for England and Wales 1991–1993 to 2000–2002" (available at http://www.statistics.gov.uk). For consistency, we continue to use only figures from England and Wales in 1999–2002.
- Life expectancy is linearly interpolated for missing years (but not extrapolated).

A.4. Sources of infant mortality data

The infant mortality rate is measured as probability that a baby born live does not survive until its first birthday. This figure is typically expressed as a rate per 1000 births, and we follow this convention.

Most infant mortality data is taken from the Human Mortality Database (HMD), found at http://www.mortality.org We use the tables Life Tables by Year of Death (1×1), and calculate infant mortality as $q(x) \times 1000$ for $x = 0$, where $q(x)$ is the probability of death between exact ages x and $x + 1$. There are three exceptions:

- US infant mortality is from the Statistical Abstract of the United States, Table No. HS-13. Live births, deaths, infant deaths, and maternal deaths: 1900–2001. Prior to 1960, this excludes Alaska and Hawaii. Beginning 1970, it excludes births to, and deaths of, non-residents of the United States.
- Australian data are from Australian Bureau of Statistics, Australian Historical Population Statistics, ABS Catalog Number 3105.0.65.001, Table 46.
- Irish data is from Vital Statistics, 2001 Annual, p. 137.

Additionally:

- New Zealand data prior to 1937 are adjusted in the same manner as for life expectancy.
- UK infant mortality data only covers England and Wales.
- Infant mortality for missing years is interpolated log-linearly (but not extrapolated).

A.5. Sources of homicide and suicide data

Homicide and suicide figures are from the World Health Organization Mortality Database (4 March 2005 update), available from http://www3.who.int/whosis/mort/. This database tabulates deaths by country back to 1950, classified according to the prevailing International Classification of Diseases system (ICD7–ICD10). Homicide and suicide rates are both expressed as rates per 100,000 people. Since the coding changes over time, it is useful to set out the precise ICD codes that were used here.

Homicide: A149 and B050 from ICD7; A148 and B050 from ICD8; B55 from ICD9; 1102–1103, X85–X99, and Y00–Y09 from ICD10.

Suicide: A148 and B049 from ICD7; A147 and B049 from ICD8; B54 from ICD9; 1101 and X60–X84 from ICD10

In missing years, homicide and suicide rates are linearly interpolated for all countries (but not extrapolated). For Australia and the US, we checked the WHO figures against those from the national statistical agencies, and found very little difference between the two series. WHO homicide figures for Switzerland are not credible for the period 1995–2001, so we drop these years from our analysis.

A.6. GDP

GDP is real GDP per capita (measured in 1990 International Geary-Khamis dollars), from Angus Maddison, The World Economy: Historical Statistics (2003). Downloaded from http://www.ggdc.net/Maddison/. Maddison's series end in 2001, and we extrapolate linearly for 2002–2003 where necessary.

A.7. Educational attainment

Educational attainment is the average number of years of schooling for the population aged 15 and over, from: Barro, R.J., Lee, J.W., 1993. International comparisons of educational attainment. Journal of Monetary Economics 32, 363–394; Barro, R.J., Lee, J.W., 1996. International measures of schooling years and schooling quality. American Economic Review 86, 218–223; Barro, R.J., Lee, J.W., 2000. International data on educational attainment: updates and implications. Center for International Development Working Paper 42. CID, Cambridge, MA.

Barro and Lee provide figures every 5 years from 1960 to 2000, and we linearly interpolate for intervening years (and linearly extrapolate after 2000). Data can be downloaded from http://www.cid.harvard.edu/ciddata/ciddata.html.

A.8. Health expenditure

Health expenditure is from OECD Health Data 2004 (updated September 24, 2004), downloaded from http://www.oecd.org/health/healthdata. We use two variables, public health expenditure per capita, and private health expenditure per capita (created as real total health expenditure per capita minus real public health expenditure per capita). Both are supplied by the OECD database in US$ (converted at purchasing power parity). We then adjust for inflation by converting these amounts into 2003 dollars using the CPI-U-RS.

References

Acemoglu, D., Johnson, S., Robinson, J., Yared, P., 2005. From education to democracy? NBER Working Paper 11204. NBER, Cambridge, MA.

Alesina, A., La Ferrara, E., 2002. Who trusts others? Journal of Public Economics 85, 207–234.

Alesina, A., Baqir, R., Easterly, W., 1999. Public goods and ethnic divisions. Quarterly Journal of Economics 114, 1243–1284.

Atkinson, A., Brandolini, A., 2001. Promise and pitfalls in the use of "secondary" data-sets: income inequality in OECD countries as a case study. Journal of Economic Literature 39, 771–799.

Atkinson, A., Piketty, T. (Eds.), in press. Top incomes over the twentieth century: a contrast between European and English-speaking countries. Oxford University Press, Oxford.

Backlund, E., Sorlie, P., Johnson, N., 1996. The shape of the relationship between income and mortality in the United States: evidence from the national longitudinal mortality survey. Annals of Epidemiology 6 (1), 12–20.

Berkman, L., Syme, L., 1979. Social networks, host resistance, and mortality: a nine-year follow-up study of Alameda county residents. American Journal of Epidemiology 109, 186–204.

Deaton, A., 2001. Relative deprivation, inequality, and mortality. NBER Working Paper 8099. NBER, Cambridge, MA.

Deaton, A., 2003. Health, inequality, and economic development. Journal of Economic Literature 41, 113–158.

Deaton, A., Paxson, C., 2001. Mortality, income, and income inequality over time in Britain and the United States. NBER Working Paper 8534. NBER, Cambridge, MA.

Fajnzylber, P., Lederman, D., Loayza, N., 2002. Inequality and violent crime. Journal of Law and Economics 45 (1), 1–40.

Gravelle, H., Wildman, J., Sutton, M., 2002. Income, income inequality and health: what can we learn from aggregate data? Social Science and Medicine 54, 577–589.

Judge, K., Mulligan, J., Benzeval, M., 1998. Income inequality and population health. Social Science and Medicine 46 (4–5), 567–579.

Kawachi, I., Kennedy, B.P., Lochner, K., Prothrow-Stith, D., 1997. Social capital, income inequality, and mortality. American Journal of Public Health 87 (9), 1491–1498.

Knack, S., Keefer, P., 1997. Does social capital have an economic payoff? A crosscountry investigation. Quarterly Journal of Economics 112 (4), 1251–1288.

Leigh, A., in press. Does equality lead to fraternity? Economics Letters.

Leigh, A., 2006. Using Panel Data on Top Income Shares to Analyze the Causes and Effects of Inequality. Australian National University, Mimeo.

Lynch, J., et al., 2004a. Is income inequality a determinant of population health? Part 1. A systematic review. Milbank Quarterly 82 (1), 5–99.

Lynch, J., et al., 2004b. US national and regional trends in income inequality and age- and cause-specific mortality. Milbank Quarterly 82 (2), 355–400.

Marmot, M., 2005. The Status Syndrome: How Social Standing Affects Our Health and Longevity. Times Books, Henry Holt, New York.

Martin, J., 1981. Relative deprivation: a theory of distributive injustice for an era of shrinking resources. Research in Organizational Behavior 3, 53–107.

McKeown, T., Record, R.G., Turner, R.D., 1975. An interpretation of the decline of mortality in England and Wales during the twentieth century. Population Studies 29 (3), 391–422.

Meltzer, A., Richard, S., 1981. A rational theory of the size of government. Journal of Political Economy 89 (5), 914–927.

Miller, D., 2001. Income Inequality and Mortality in the US: Aggregate Data and Micro Relationships. University of California, Berkeley, Mimeo.

Preston, S.H., 1975. The changing relation between mortality and level of economic development. Population Studies 29, 231–248.

Sen, A., 1999. Development as Freedom. Knopf, New York.

Szreter, S., 1988. The importance of social intervention in Britain's mortality decline c. 1850–1914: a reinterpretation of the role of public health. Social History and Medicine 1, 1–37.

Waldmann, R., 1992. Income distribution and infant mortality. Quarterly Journal of Economics 107 (4), 1283–1302.

Wilkinson, R.G., 1989. Class mortality differentials, income distribution and trends in poverty 1921–1981. Journal of Social Policy 18, 307–335.

Wilkinson, R.G., 1992. Income distribution and life expectancy. British Medical Journal 304, 165–168.

Wilkinson, R.G., 1996. Unhealthy Societies: The Affliction of Inequality. Routledge, London.

Wilkinson, R.G., 1997. Socioeconomic determinants of health: health inequalities: relative or absolute material standards? British Medical Journal 324 (7343), 978.