HEALTH ECONOMICS
Health Econ. **16**: 911–928 (2007)
Published online 2 August 2007 in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/hec.1260

SPENDING TO SAVE? STATE HEALTH EXPENDITURE AND INFANT MORTALITY IN INDIA

SONIA BHALOTRA*

Department of Economics & CMPO, University of Bristol, UK

SUMMARY

There are severe inequalities in health in the world, poor health being concentrated amongst poor people in poor countries. Poor countries spend a much smaller share of national income on health expenditure than do richer countries. What potential lies in political or growth processes that raise this share? This depends upon how effective government health spending in developing countries is. Existing research presents little evidence of an impact on childhood mortality. Using specifications similar to those in the existing literature, this paper finds a similar result for India, which is that state health spending saves no lives. However, upon allowing lagged effects, controlling in a flexible way for trended unobservables and restricting the sample to rural households, a significant effect of health expenditure on infant mortality emerges, the long run elasticity being about −0.24. There are striking differences in the impact by social group. Slicing the data by gender, birth order, religion, maternal and paternal education and maternal age at birth, I find the weakest effects in the most vulnerable groups (with the exception of a large effect for scheduled tribes).

Received 11 October 2006; Revised 17 May 2007; Accepted 18 May 2007

KEY WORDS: public spending; health; poverty; infant mortality; India

INTRODUCTION

Motivation and context

Inequalities in life expectancy across and within countries are created mainly by variation in childhood mortality. In poor countries, 30% of deaths are amongst children, compared with less than 1% in rich countries (Cutler *et al*., 2005, p. 15). As many as 10 million children under the age of five die each year, mainly from preventable (or curable) conditions that seldom kill children in rich countries (Jones *et al*., 2003; Black *et al*., 2003). Yet, most of the relevant interventions, such as immunisation or oral rehydration therapy, are very low cost (Deaton, 2006). This suggests that it is not just a question of raising incomes, but of the effective delivery of publicly provided health services. In this paper, the effectiveness of public intervention (state health expenditure) is measured in terms of its impact on infant mortality or death in the first year of life.

Analyses of the historical decline in childhood mortality rates in today's industrialised countries suggest that important drivers of this decline were improved nutrition, public health and medical technological progress (see Fogel, 2004; Cutler and Miller, 2005; Cutler *et al*., 2006). Improved nutrition tends to be associated with growth in income. Medical progress may, in principle, diffuse across geographic boundaries with no tight connection to incomes or public expenditure. Improvements in education, water and sanitation, immunisation and targeted programmes against diseases like malaria

*Correspondence to: Department of Economics, University of Bristol, 8 Woodland Road, Bristol BS8 1TN, UK.
E-mail: s.bhalotra@bristol.ac.uk

and diarrhoea tend to be associated with growth in public spending. In assessing the role that public spending might play in bringing down childhood death in poor countries, it is important to disentangle its effects from those of other trended variables, in particular, income growth and scientific progress.

This is done here by investigating the impact of fluctuations in health expenditure around a state-specific trend. Although the conventional wisdom is that fiscal policy should be counter cyclical (Lane, 2003), smoothing the effects of income shocks, in practice it is often procyclical in developing countries (Woo, 2005). At the same time, aggregate income volatility is much greater in poor than in rich countries (Pritchett, 2000; Koren and Tenreyro, 2007). In these circumstances, we may expect that mortality is counter-cyclical, with adverse shocks to household income being reinforced by cuts in social expenditure. For the Indian sample analysed here, this is the case (Bhalotra, 2007). This paper isolates the impact on mortality of changes in state health expenditure, holding constant state income. The effect I identify is therefore the effect of changes in the share of state income that is dedicated to health. This may vary, for example, in response to health shocks (natural disasters, rainfall variation, epidemics), inequality (Woo, 2005), the political climate in the state and the salience of public health.

The analysis is conducted for India, which accounts for one in four of under-five deaths, one in three of the poor and one in six of the population in the world. On account of its size, it has the highest child death toll in the world: 2.4 million under-five deaths (Black *et al.*, 2003), and infant deaths account for more than two-thirds of under-five deaths (my calculations). Infant mortality is regarded as a sensitive indicator of the availability, utilisation and effectiveness of health care, and it is commonly used for monitoring and designing population and health programmes (The Tribune, 2002). Like the United States, India has a federal political structure, and health is a 'state subject', which means that the level and allocation of health expenditure are decided at the state level.

Individual data on mortality are derived from retrospective fertility histories recorded in a national sample survey and merged by birth cohort with a 29-year panel of data on state health expenditure, income and other variables. The individual data are, in this way, 'nested' in a state-year panel. The main contributions of this paper over the existing literature lie in its exploiting sub-national panel data on health expenditure to identify its impact, and its use of individual data on mortality to investigate heterogeneity in this impact by social group. Let me elaborate each. Most previous studies use a single section of cross-country data. They are therefore unable to control for unobservable trends in medical technology which have been important in driving mortality reduction, and omission of which will tend to bias the estimated effects of health expenditure.[1] Cross-country regressions are also prone to other forms of correlated heterogeneity which, in a panel, are absorbed by state fixed effects. A further advantage of using a panel and, especially, a long panel, is that dynamics can be explored. No previous research in this domain appears to have explored dynamics and, here, I show that this is critical. This is also the first study in this area that controls for correlated weather shocks, omission of which will generate spurious co-variation of mortality and health expenditure. This study investigates heterogeneity in health expenditure by observed individual and family characteristics. Examining heterogeneity is interesting in itself but it also provides insight into the mechanisms by which health expenditure has an impact, if any. Some previous studies have examined the impact of health expenditure on health outcomes by (simulated) income groups (Gupta *et al.*, 2003; Bidani and Ravallion, 1997).

Using specifications similar to those in the existing literature, this paper finds a similar result for India, which is that state health spending saves no lives. However, restricting the sample to rural households (more than two-thirds of all) and conditioning upon state-specific trends, a significant effect

[1] Deaton and Paxson (2004), for example, emphasise the importance of controlling for time-varying unobservables in identifying the impact of income on mortality.

DOI: 10.1002/hec

emerges that is driven by the third lag of health expenditure. The long run marginal effect is -0.023, and the elasticity is -0.24. A one standard deviation (0.48) increase in log health expenditure per capita at a given level of state income is therefore estimated to reduce the risk of mortality by 0.01. Taking a UN estimate of live births in India in 1990 of 26.3 million, this amounts to saving 0.26 million lives.

I find striking differences in the impact of health expenditure by social group. The impact is greater for rural and scheduled tribe households than for urban or higher-caste households. However, slicing the data by gender, birth order, religion, mother's and father's education and maternal age at birth, I find the weakest effects in the most vulnerable groups. I argue that this may be related to the way in which health expenditure is used. Previous studies that have looked at the distribution of health expenditure effects have focused on income. The effects I find suggest that attitudes and information, which may not be strongly correlated with income, mediate the effects of state spending.

Related literature

Paxson and Schady (2005) show that, when Peru's GDP fell in 1987–1990 by 30%, government health expenditure fell by 58%, its budget share falling from 4.3 to 3%. At the same time, infant mortality spiked, rising by 2.5 percentage points. While this analysis of trends broken by a big exogenous shock is persuasive, it is difficult to generalise from. In particular, changes in health expenditure might impact mortality only when they are very large. There is limited evidence in the literature on the health effects of regular fluctuations in state health spending, most previous studies having used a single section of cross-country data.

In an influential study, Filmer and Pritchett (1999) investigate the effect of government health expenditure on infant and under-five mortality using cross-sectional data on 98 developing countries in 1992/1993. They find a very small and statistically insignificant effect. They show that 95% of the variation in mortality between countries is explained by income per capita, income inequality, female education, ethnic fractionalisation and whether the country is more than 90% Muslim. This is an important study with striking results. But the results are not incontrovertible. Using data for 50 developing and transition countries observed in 1994, Gupta *et al.* (2002) find some evidence that health expenditure reduces childhood mortality, but argue that this effect is not very robust. Using cross-sectional data for 22 developing countries in 1985, Anand and Ravallion (1993) find that health expenditure raises life expectancy and that, conditional upon this, income has no effect. All of these studies suffer two important limitations. First, as the authors recognise, data on both mortality and government health expenditure are unlikely to be comparable across countries. Second, the estimates in these studies are subject to bias on account of unobserved heterogeneity that might be correlated with the variable of interest (see Durlauf *et al.*, 2005). The present study addresses the first problem by using sub-national data, and the second problem by using panel data on state health expenditure and income.

There is some relevant recent work for India (Deolalikar, 2005). Using a state panel for 1980–1999, this study finds no effect of current health expenditure on mortality rates once state fixed effects and a linear time trend are included in the model. I find a similar result (see below). Anil Deolalikar further investigates the relationship for a reduced sample of four years and 14 states ($N = 56$) for which information on female literacy is available. For this sample, an interaction term between health expenditure and state income is included and the results suggest a negative effect of health expenditure but only in the poor states. In a complementary analysis of micro-data for the period 1994–1998, he finds the opposite – that the effects of health expenditure are weaker in the poor states. My state-specific estimates show no very clear relation between the effects of health expenditure and the per capita income of the state. However, there are differences between the two studies in sample period, data and

DOI: 10.1002/hec

estimator. In particular, this study uses a longer time period and a more flexible specification of trends, it conditions upon rainfall shocks, and it investigates lagged effects. It also investigates heterogeneity in the impact of health expenditure, not only by state, but also by social group.

The rest of this paper is organised as follows. The second section describes the data and the third section presents relevant descriptive statistics. An empirical model is set out in the fourth section and results are discussed together with a range of robustness checks in the fifth section. The six section concludes.

THE DATA

The micro-data are derived from the second round of the National Family Health Survey of India (NFHS-2).[2] This contains complete fertility histories for ever-married women aged 15–49 in 1998–1999, including the time and incidence of child deaths. I use these to construct individual-level indicators of infant mortality. The children in the sample are born in 1961–1999. It is unusual to have annual data on mortality rates over a period as long as this (see Pritchett and Summers, 1996). An advantage of using the individual data on mortality is that I can explore heterogeneity in the effect of health expenditure by population sub-groups. This is important since a genuine relation for one group (e.g. rural) can be obscured by pooling when there is heterogeneity in the slopes. This is an unusual facility, the existing literature relating social expenditures and outcomes being dominated by cross-country data analysis. The NFHS provides a rich set of individual-level controls, which are included in the model to control for heterogeneity in death risk. There are no data on household incomes over the 28-year period analysed. Permanent income at the household level is proxied by parental education, and also caste and religion.[3] The estimates control for aggregate income at the state level, which finances state health expenditure.

As the micro-data are constructed from retrospective fertility histories, they are wedge-shaped, there being fewer observations for children born earlier in time. Moreover, the thinning of the data does not occur randomly, but is a function of maternal age at birth. It is therefore important to condition upon maternal age at birth. I also drop children born in the 1960s, since this is where the data are both most thin and most skewed. Another issue that arises with retrospective data is that the mother may have migrated between states between the birth of the index child and the date of interview. However, the survey asks the mother how long she has lived in her current location. Using this information, the analysis is restricted to the 85.1% of births that occurred in the mother's current location, so that we can be confident that infant mortality risk is related to health expenditure in the state in which the child was born. As a (rough) check on whether this sample selection is endogenous, I compared estimates on the restricted and unrestricted samples, and found that they were not significantly different.

The conventional definition of infant death is death before the first birthday of the child. Since mother's reports of age at death of their children exhibit age heaping at six-monthly intervals, infancy is defined here to include the 12th month. The results are not sensitive to this difference, but the inclusive definition is retained since this increases the ratio of ones to zeroes in the dependent variable. To ensure that every child is allowed full exposure to the risk of infant mortality, births that occur in the 12 months preceding the survey are excluded. The estimation sample contains more than 150 000 children of more than 59 000 mothers born in 1970–1998 across the 15 major Indian states. These micro-data are merged by state and year of birth with a panel of data on health expenditure and other relevant

[2] For details on sampling strategy and context, see IIPS and ORC Macro (2000). The data are available at www.measuredhs.com.
[3] These variables are recorded at the time of the survey. The educational attainment of parents varies by cohort of child and, in this way, it varies over time.

DOI: 10.1002/hec

statistics for the 15 Indian states.[4] To obtain descriptives, I aggregated the individual data to the state level using sampling weights. The aggregation is done by birth cohort, yielding a straightforward panel in which state mortality rates can be related to state health expenditure. A potential advantage of using the panel for the analysis is that unobserved heterogeneity at the child and mother level is averaged out. As a robustness check, I report results obtained using the panel (fifth section).

State health expenditure includes expenditure from state revenue (85%) and central government health allocations to the state (15%), the latter often being tied to public health and family welfare programmes. I use actual as opposed to budgeted revenue expenditure, even if this makes it more likely that health expenditure is endogenous. State health spending covers rural and urban public health services; medical education, training and research; general administration; water supply and sanitation; and family welfare. The expenditure series is cast in per capita terms and deflated by the consumer price index for agricultural workers. State income is measured as net domestic product and is subject to the same deflators.

DESCRIPTIVE STATISTICS

Figures 1–4 reveal the dispersion in levels and trends across the states in mortality, state health expenditure, the share of health expenditure in state income and income.[5] These figures show that the rate of increase in both the level and the share of health expenditure has slowed since about the mid-1980s, even as the increase in state income has accelerated. Regressing state health expenditure on state income, a lagged dependent variable (instrumented with two further lags), year and state dummies and state-specific trends, I find a long run income elasticity of health spending of -0.41, which is identified off within-state variation. An elasticity smaller than one indicates that the share of health expenditure in income is decreasing in income. So, growth in India in the period analysed was not used to finance proportional increases in health expenditure. (I find a similar elasticity for state education expenditure; results available on request.)

Government health expenditure in India was, on average, 1.3% of GDP in 1990, and this had declined to 0.9% in 1999 (NRHM, 2005). India devotes a smaller share of its income to health spending than, for example, Bangladesh (1.4%) or Sri Lanka (1.8%) (Deolalikar, 2005, Chapter 2; these are figures for the year 2000), and it spends a disproportionate part of its health budget on (curative) hospital services which are less propoor than (preventive) public health expenditures (Peters *et al.*, 2002).

Figure 5 shows that the ‘raw’ relationship between mortality and health expenditure is generally negative, as one might expect. Figure 6 plots these data again, after removing state-specific trends. What is striking is that, in the de-trended data, there is little evidence that increases in health expenditure are systematically associated with decreases in mortality. The rest of this paper explores whether these simple associations in the data persist after conditioning upon other covariates, and after allowing for lagged effects.[6]

[4] I am grateful to Tim Besley and Robin Burgess for letting me use their state-level panel (see Besley and Burgess, 2004). Detailed definitions of the state-level variables used in the analysis can be found at http://sticerd.lse.ac.uk/eopp/research/indian.asp. The health expenditure series were kindly given to me by Juan Pedro Schmid, who gathered them from Reserve Bank of India publications. Juan made the series consistent before and after 1985, the year in which the published categorisation of health expenditure was changed. Before 1985, state health expenditure included expenditure on medical and public health, family planning and water supply and sanitation. From 1985 onwards, family planning and water–sanitation expenditures appear separately in the accounts and need to be added in.

[5] The figures are drawn using state-level aggregates of the individual data on mortality.

[6] Growth rates of the main variables by state (Table A1) and summary statistics for all variables in the model (Table A2) are in an online appendix available at www.efm.bris.ac.uk/www/ecsrb/bhalotra.htm. Figures A1 and A2 in this appendix describe inequality in the levels of mortality and health expenditure across the states.

DOI: 10.1002/hec

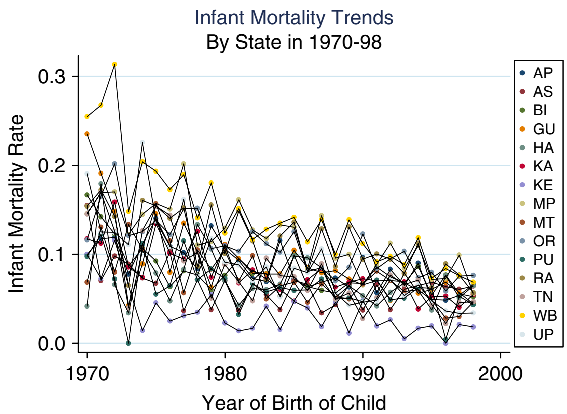

Figure 1. Infant mortality trends

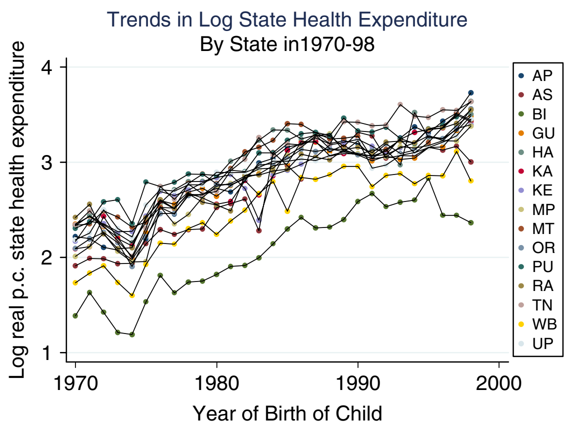

Figure 2. Trends in log state health expenditure

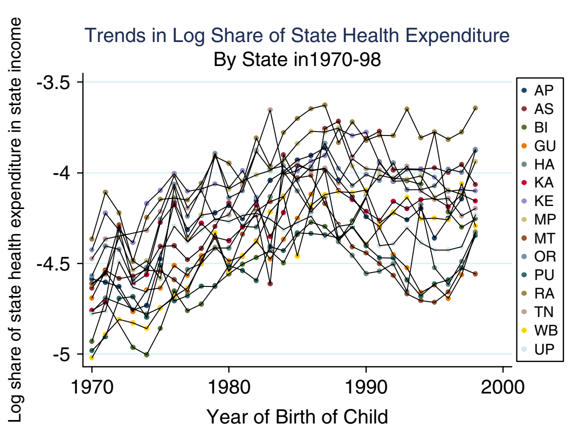

Figure 3. Trends in log share of state health expenditure

DOI: 10.1002/hec

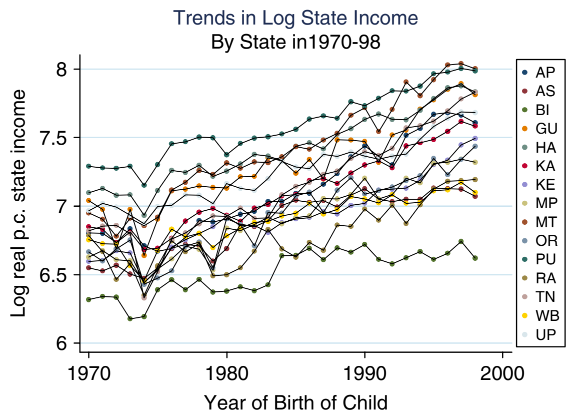

Figure 4. Trends in log state income

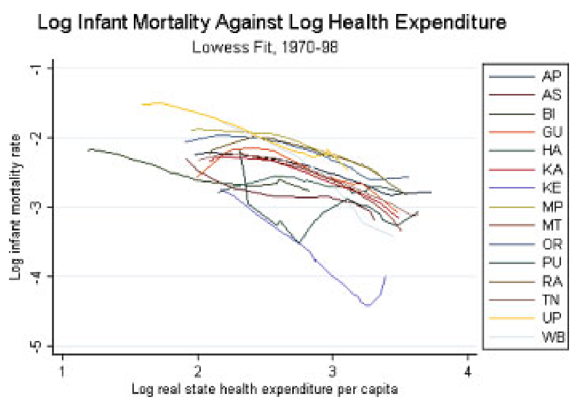

Figure 5. Log infant mortality against log health expenditure

THE EMPIRICAL MODEL

The baseline model is

$$M^*_{\text{imst}} = \alpha_0 + \alpha_{\text{s}} + \alpha_{\text{t}} + \mu_{\text{st}} + \beta \ln H_{\text{st}} + \gamma \ln Y_{\text{st}} + +\eta'_{\text{s}} R^{\text{f}}_{\text{st}} + \gamma'_{\text{s}} R^{\text{d}}_{\text{st}} + \lambda' X_{\text{imst}} + \phi' Z_{\text{st}} + \epsilon_{\text{imst}}$$

Subscripts s and t indicate state and year and i and m indicate the individual child and mother, respectively, ln denotes logarithm. The individual data are nested in a state panel. M^* is a latent variable measuring the probability of infant death. H is per capita real health expenditure, Y is per capita real net domestic product and β is the parameter of interest. X is a vector of variables observed at the child or mother level, Z is a vector of state-level controls and R^{f} and R^{d} are vectors of positive and negative state-year-specific rainfall shocks, the effects of which are allowed to be state-specific (superscripts f and d denote 'flood' and 'drought', respectively). To avoid clutter, I do not show dynamics or interaction (and quadratic) terms, though these are investigated, and discussed in the Results section. α_{s} and α_{t} are state and year fixed effects and μ_{st} are state-specific trends.

The model is estimated as a probit. All reported standard errors are robust and clustered by state. These adjustments allow for conditional heteroskedasticity and for conditional autocorrelation within states (see Bertrand *et al.*, 2004; Cameron and Trivedi, 2005, p. 788). Note that adjusting

DOI: 10.1002/hec

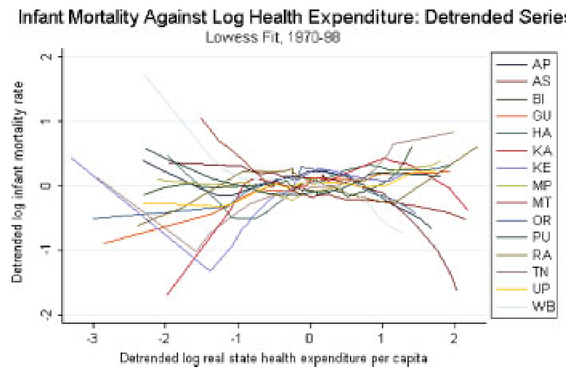

Figure 6. Infant mortality against log health expenditure: detrended series

for clustering at the state level takes care of any lower-level clustering such as at the community or mother level. Identification of β relies upon there being independent fluctuations in health expenditure within states. The relatively long time dimension of the data makes it more likely that this is the case.

X includes child gender, dummies for birth month of the child, age of mother at birth of the child, levels of education of each of mother and father and ethnicity and religion of the household. These characteristics have been shown to be significant predictors of mortality risk in a number of previous studies, and also in India (e.g. Bhalotra and van Soest, 2006). Z includes income inequality measured as the log of the Gini coefficient for each of the rural and urban sectors, poverty measured as the log of the sector-specific headcount ratio, the ratio of agricultural to non-agricultural income in the state, inflation in consumer prices and a quadratic in newspaper circulation per capita.

Rainfall shocks are calculated as the absolute deviation of rainfall in each state year from the 30-year state mean. A positive shock is then defined as equal to this deviation when the deviation is positive, and equal to zero otherwise. A negative shock is symmetrically defined. These are the terms that appear in the regressions. However, to allow the effects of rainshocks on mortality to be different in different states, each of these two indicators is interacted with the 15 state dummies, so that R is a vector of 30 variables. The richness of this specification is justified by the results. When rainshocks are restricted to have the same effect in every state, they are insignificant. Once state-specific coefficients are allowed, they are jointly significant at 1%. The results also indicate that it is restrictive to force positive and negative deviations to have the same effect.[7]

The panel aspect of these sub-national data offers some clear advantages relative to previous cross-sectional or time-series analyses, as well as relative to analyses of cross-country panels. An important advantage is that it permits controls for time-varying unobservables. There is little doubt that medical technological progress has contributed to secular reductions in mortality. Failing to control for this will result in an over-estimation of the effects of any included trended variables (health expenditure, income). The assumption that time dummies capture technology trends is not innocuous when the regions are countries,[8] or at least is more plausible for states within a country. This assumption is further relaxed by including state-specific trends in the model. The time dummies will also capture

[7] A natural alternative to using absolute deviations is to use the z-score of rainfall which normalises deviations with respect to the standard deviation in the state. The specification used here allows a big deviation in rainfall to impact infant mortality as much in a state that often experiences rainfall fluctuations as it would in a state with a more stable weather pattern. This seems to me the more relevant specification, but I have confirmed that using z-scores does not alter the main results of this analysis.

[8] Temple (1999), for example, shows that countries have different rates of technical progress in growth regressions, casting doubt that technology is a public good. This said, diffusion of health technology across countries may occur more effectively than diffusion of production technology.

DOI: 10.1002/hec

common (all-India) shocks such as famines, floods or epidemics, and the state-specific trends will capture not only state-specific components of health technology but also other omitted trends, for example, in fertility or public services.

The state effects, α_s, control for all forms of time-invariant unobserved heterogeneity specific to a state. In this context, this is likely to include sluggish political institutions, ethnic composition, geography, and initial conditions, including the initial level of mortality in the state. They will also pick up any persistent differences across the states in accounting conventions.

Since health expenditure varies across state year, we cannot, of course, include state year dummies that would control comprehensively for state-specific health shocks. As a result, health expenditure remains potentially endogenous in this model. Consider, for example, that a particular state suffers an epidemic, a flood or a famine. Suppose that, as a result, more infants die, and the state raises health expenditure. This will show as a (spurious) positive relationship between infant mortality and health expenditure. Since most infant deaths occur in rural areas and, amongst rural households, rainfall shocks are probably the most important sorts of shocks, the analysis controls in a flexible way for rainshocks. To investigate this directly, I estimated an auxiliary panel data model in which state health expenditure is the dependent variable. Controlling for state income, fixed effects and state-specific trends, I find that rainshocks are jointly significant at the 1% level.

Controlling for other state-level variables (Z) contributes to reducing the chances that health expenditure is endogenous. I also replace current health expenditure with its first four lags so as to protect the expenditure coefficient from being biased by any contemporaneous correlation of mortality and health expenditure induced by state-specific shocks that are not controlled for. Lags are also useful to explore because health expenditure might take more than a year to take effect on the field. Lagged effects may alternatively arise on account of state dependence in mortality within families (e.g. Arulampalam and Bhalotra, 2007).

Since the key regressor (state health expenditure) varies at the state and not the individual level, the data are a bit scarce for estimation of state-specific models. However, to gain at least an indicative sense of the state-specific relationships, I also estimate the following simple linear model for each state ($T = 28$):

$$M^{*}_{\mathrm{imt}} = \phi_0 + \eta_t + \chi \ln H_{\mathrm{t}} + v \ln Y_{\mathrm{t}} + \eta'_{\mathrm{s}} R^{\mathrm{f}}_{\mathrm{t}} + \gamma'_{\mathrm{s}} R^{\mathrm{d}}_{\mathrm{t}} + \lambda' X_{\mathrm{imt}} + \phi' Z_{\mathrm{t}} + u_{\mathrm{imt}}$$

RESULTS

Henceforth, *health expenditure* refers to the logarithm of real per capita state health expenditure and *income* refers to the logarithm of real per capita net domestic product of the state.

Static models

Table I presents marginal effects estimated from the probit model of infant mortality displayed in the fourth section. The rainshocks and the micro-demographic variables are jointly significant, but dropping them from the model does not alter the elasticity of interest (results available upon request). Some of the state-level controls are significant but, again, conditioning upon them does not make a significant difference to the health expenditure effect. This effect is sensitive to conditioning upon time dummies and state-specific trends. For this reason, results are reported with and without these controls. Table I explores alternative functional forms – a (log)linear term in current health expenditure, a quadratic and a first lag.

Consider the results for rural households presented in Table I(A). Current health expenditure has a significant marginal effect of −0.015 in column 1. However, once time dummies are included, this effect is close to zero, and insignificant (columns 2 and 3). Using a quadratic in expenditure, I find some evidence of nonlinearity in column 4. Mortality risk is hump-shaped in health expenditure, the relation being positive at low levels of expenditure and then turning negative. The marginal effect at the

DOI: 10.1002/hec

Table I. Probit estimates of infant mortality using alternative specifications of health expenditure: (A) rural sample and (B) urban sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(A)									
Health expenditure	**−0.015**	0.007	0.000	0.014	0.031	0.105			
	[1.92]	[0.57]	[0.03]	[0.66]	[1.22]	[1.28]			
Square of health expenditure				***−0.006***	−0.005	−0.018			
				[1.87]	[1.26]	[1.27]			
First lag of health expenditure							−0.012	0.015	0.008
							[1.63]	[1.58]	[0.91]
Income	**−0.051**	−0.028	**−0.037**	**−0.046**	−0.023	**−0.045**	**−0.055**	−0.026	**−0.038**
	[6.00]	[1.26]	[2.73]	[6.08]	[1.14]	[3.82]	[6.70]	[1.28]	[2.91]
State dummies	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year dummies	×	✓	✓	×	✓	✓	×	✓	✓
State-specific trends	×	×	✓	×	×	✓	×	×	✓
(B)									
Health expenditure	−0.002	0.013	0.003	−0.022	−0.008	−0.055			
	[0.22]	[0.88]	[0.18]	[0.48]	[0.12]	[0.86]			
Square of health expenditure				0.004	0.004	0.010			
				[0.47]	[0.36]	[0.90]			
First lag of health expenditure							−0.008	0.019	0.011
							[1.07]	[1.69]	[0.85]
Income	**−0.042**	−0.023	−0.032	**−0.044**	−0.026	−0.030	**−0.034**	−0.023	−0.033
	[3.70]	[1.24]	[1.29]	[3.64]	[1.36]	[1.16]	[2.94]	[1.19]	[1.30]
State dummies	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year dummies	×	✓	✓	×	✓	✓	×	✓	✓
State-specific trends	×	×	✓	×	×	✓	×	×	✓

Notes: The number of observations (number of live births) is 117 088 in the rural sample and 35 783 in the urban sample. These are marginal effects from a probit; significant coefficients are in bold. Standard errors are robust and clustered at the state level. Absolute *t*-statistics are in parentheses. State health expenditure and state income are real per capita measures cast in logarithms. Every equation includes state-specific positive and negative rainfall shocks and micro-demographic controls (dummies for child gender and birth month, age of mother at birth of the child, level of education of each of mother and father and ethnicity and religion of the household).

mean is −0.020. A possible explanation of this tendency for health expenditure to lower mortality only when it is relatively high is that, at low levels of expenditure, most of it goes to politically prioritised areas such as curative care in urban areas, while bigger budgets extend to lower-priority areas such as preventive care, water supply or sanitation that are more likely to impact mortality at the margin (see Lanjouw and Ravallion, 1999). However, conditional on time effects, this relationship is also rendered insignificant. The third panel in Table II (columns 7–9) reports results of replacing current expenditure with its first lag. The results are essentially similar to those obtained with current expenditure. I also investigated a specification in which I interacted health expenditure and income (as in Deolalikar, 2005). The interaction term was negative but insignificant; these results are not displayed.

Marginal effects of income are also reported in Table I(A). In the absence of controls for omitted trends, income has a significant marginal effect of −0.05. Although this effect vanishes upon including time dummies (as did the health expenditure effect), it re-establishes itself (ME of −0.04) upon inclusion of state-specific trends (which health expenditure did not).[9] Dropping income raises the marginal effect of health expenditure but does not alter its significance level.

[9] This result is consistent with the state-specific trends capturing omitted variables that are positively correlated with mortality and health expenditure, but negatively correlated with income, for example, fertility. Alternatively, they might reflect trends in technology or in the delivery of public services. These might be negatively correlated with mortality and health expenditure and positively correlated with income, producing a similar configuration of results.

Table II. Probit estimates of infant mortality: distributed lag model

	(1)	(2)	(3)
	Baseline	Add rainshocks	Add state controls
L (ln health expenditure)	0.012	0.008	0.006
	[1.38]	[0.96]	[0.75]
*L*2 (ln health expenditure)	0.003	0.002	−0.005
	[0.23]	[0.16]	[0.44]
*L*3 (ln health expenditure)	**−0.022**	**−0.020**	**−0.020**
	[2.39]	[1.94]	[1.84]
*L*4 (ln health expenditure)	−0.007	−0.005	−0.008
	[0.71]	[0.45]	[0.64]
Long run marginal effect	−0.014	−0.015	**−0.027**
	[0.96]	[0.98]	[1.73]
Health expenditure elasticity	−0.148	−0.158	−0.285
L (ln income)	−0.018	−0.019	−0.027
	[1.10]	[1.24]	[1.62]
*L*2 (ln income)	−0.006	−0.005	−0.007
	[0.47]	[0.44]	[0.56]
*L*3 (ln income)	**−0.056**	**−0.056**	**−0.052**
	[4.14]	[4.73]	[5.52]
*L*4 (ln income)	**0.031**	**0.034**	**0.031**
	[3.00]	[3.92]	[3.73]
Long run marginal effect	**−0.048**	**−0.047**	**−0.054**
	[2.07]	[2.06]	[2.55]
Income elasticity	−0.506	−0.496	−0.570

Notes: Rural sample, $N = 117\,088$. Standard errors are robust and clustered at the state level. Absolute t-statistics are in parentheses. L denotes lag. Every equation includes state and year dummies, state-specific trends and the micro-demographic controls listed in Notes to Table I. Column 2 further includes positive and negative state-specific rainfall shocks. The additional controls in Column 3 are the ratio of agricultural to non-agricultural output in the state, inflation of consumer prices for agricultural and industrial workers, the log poverty headcount ratio and the log of the Gini coefficient for each of the rural and urban sectors and a quadratic in per capita newspaper circulation.

Now consider the more stark results for urban households that are displayed in Table I(B). In no column is health expenditure significant. Income is significant only if there are no time effects in the model. Infant mortality rates are lower in urban households and they are, on average, richer and will tend to have higher private expenditure on health and nutrition. Although the further specifications discussed below were estimated for both samples, there is no case in which health expenditure is significant for urban households. From here on, all reported results are therefore for the rural sample.

The results are not sensitive to the choice of estimator. The linear probability model yields a similar pattern of results. Adjusting the standard errors for clustering by state increases them by about 43% (see Table A4 in the online appendix – refer footnote 6). This said, in specifications that include time dummies, health expenditure is insignificant even before clustering is allowed. Each of the sets of state dummies, year dummies and state-specific trends is jointly significant at the 1% level in every specification in which they appear. I find that the restriction that the time dummies can be replaced by a linear trend is not satisfied by the data.

In sum, health expenditure appears to have no effect on mortality once common time-varying unobservables are removed, and this result is not sensitive to controlling for rainfall shocks and individual heterogeneity.

State-specific estimates of the static model

It is possible that these negative results conceal some significant state-specific slopes. To investigate this, I estimated state-specific models using the time series. Results are in Table A3 in the online appendix. Health expenditure has a significant negative effect in three of the 15 states (Assam, Maharashtra and

DOI: 10.1002/hec

West Bengal). However, if I drop state-specific trends, health expenditure has a mortality-reducing effect in five states: Karnataka, Maharashtra, Tamil Nadu, Uttar Pradesh and West Bengal. These five states do not form a natural group in terms of being, for example, poorer, or more politically liberal. In contrast, the states that show a negative effect of income are the poor high mortality states.

Distributed lag model

Although controlling for rainfall shocks removes a possible source of correlated unobserved heterogeneity, it remains possible that there are other state-specific health shocks that raise (or lower) both infant mortality and health expenditure, as a result of which the estimated coefficient on health expenditure will tend to carry a positive bias. It is possible that this is dominating an underlying negative causal effect in the results we have seen so far. I therefore investigated a distributed lag model with four lags of health expenditure. Using lags breaks any contemporaneous correlation between mortality and expenditure that is driven by an omitted variable, and it also allows for the possibility that causal effects take time to play out. It is natural to allow the same lags for income.

Refer Table II (covariate effects are in Table A5 in the online appendix). Every column includes micro-demographics, rainfall shocks, state and time dummies and state-specific trends. Results are displayed with and without controlling for state-level variables. The third lag of health expenditure and the third and fourth lags of income are significant. This specification is asking rather a lot of data. Recall that with aggregate and state trends in the model, we are identifying the effect of expenditure from variation around these trends. And the lagged terms are likely to be collinear. In Table III, I therefore report estimates of a more parsimonious model that retains only the significant terms from the fourth-order lag specification. In this, the preferred model, the marginal effect of health expenditure is −0.023. The long run elasticity is −0.24. The long run income elasticity, at −0.28, is bigger. The effects of both health expenditure and income are sensitive to exclusion of the state-specific trends (see Table III).

Replacing current with lagged values does not make a dramatic difference to the long run income effect, but it makes an important difference to the health expenditure effect (compare Tables I(A) and III). Health expenditure appears insignificant in most standard specifications, consistent with much of the existing literature. However, a sufficiently flexible model reveals a highly significant effect driven by the third lag of expenditure. What might explain this? Most infant death occurs in the first month and

Table III. Probit estimates of infant mortality: parsimonious model with significant lags

	(1)	(2)	(3)	(4)
	No state trends		Include state trends	
		Add state X		Add state X
$L3$ (ln health expenditure)	−0.009	−0.016	−0.020	−0.023
	[0.67]	[1.46]	[2.17]	[2.86]
Health expenditure elasticity	−0.095	−0.169	−0.211	−0.243
$L3$ (ln income)	−0.059	−0.058	−0.059	−0.057
	[2.89]	[3.66]	[4.31]	[5.30]
$L4$ (ln income)	0.039	0.034	0.032	0.031
	[3.45]	[3.16]	[3.28]	[2.93]
Long run marginal effect	−0.020	−0.023	−0.027	−0.026
	[0.83]	[1.38]	[2.47]	[2.64]
Income elasticity	−0.212	−0.247	−0.283	−0.277

Notes: Rural sample, $N = 117\,088$. Standard errors are robust and clustered at the state level. Absolute t-statistics are in parentheses. L denotes lag. Full results, showing the marginal effects of all covariates, are in Table A5 of the online appendix referred in footnote 6. Every equation includes state and year fixed effects, rainshocks and micro-demographic controls. Columns 2 and 4 also include state-level controls; see Notes to Tables I and II.

DOI: 10.1002/hec

even the first week of life, and it is well known that the proximate cause of this is low birth weight, the predominant cause of which is poor maternal health. So one lag may simply denote the importance of health expenditure in the year before birth (e.g. antenatal care). Lagged effects can also be explained in terms of state dependence in mortality within families. A drop in state expenditure three years ago may have killed a sibling of the index child who, at the time, was an infant. This, in turn, may have a causal effect on the death risk of the index child (see Arulampalam and Bhalotra, 2007). Alternatively, it may take longer than a year for increases in health spending to reach the ground. The next section reports specification checks on the robustness of the results.

Robustness

To summarise the results so far, it is only when we restrict the sample to rural households, allow lags, and condition upon state-specific trends that a significant impact of health expenditure emerges. A possible explanation is that health expenditure is endogenous, and that this endogeneity is being limited by factoring out state-specific trends, and by lagging health expenditure. As for the rural–urban difference, it is well known that failing to allow for heterogeneity can obscure important relationships in sub-populations.

The rest of this section reports the results of further specification checks. The individual mortality information (0/1) is aggregated to the state level by cohort using sampling weights. The equation is then estimated by within-groups on the state panel, weighting by the number of births (see Table IV). The health expenditure effect is a bit larger, but insignificantly different from that obtained in the analogous model run with micro-data on mortality. This suggests that the micro-demographic covariates in the model capture individual heterogeneity sufficiently well. Panel regressions in which the dependent variable is the log of infant mortality produce broadly similar results. I use the level rather than the log

Table IV. Panel data estimates for the infant mortality rate

	(1) 4 lags	(2) Significant lags
L (ln health expenditure)	0.005	
	[0.51]	
*L*2 (ln health expenditure)	−0.014	
	[0.87]	
*L*3 (ln health expenditure)	−0.022	−0.029
	[1.90]	[2.97]
*L*4 (ln health expenditure)	−0.009	
	[0.92]	
Long run marginal effect	−0.040	−0.029
	[1.86]	[2.97]
Health expenditure elasticity	−0.425	−0.306
L (ln income)	−0.017	
	[0.83]	
*L*2 (ln income)	0.007	
	[0.44]	
*L*3 (ln income)	−0.054	−0.055
	[4.71]	[4.14]
*L*4 (ln income)	0.028	0.028
	[2.07]	[2.06]
Long run marginal effect	−0.037	−0.027
	[1.55]	[1.50]
Income elasticity	−0.392	−0.282

Notes: Columns 1 and 2 correspond to Tables II (column 3) and Table III (column 4). These equations include state and year fixed effects, state-specific trends, rainshocks, micro-demographic and state-level controls.

DOI: 10.1002/hec

because it is more directly comparable with the individual data model.[10] I allowed two lagged dependent variables to capture persistence in mortality but these were insignificant. To investigate the hypothesis that the third lag is in itself not meaningful but is proxying current health expenditure, I also estimated a panel model using the IV-Systems estimator, in which infant mortality is modelled as a function of current health expenditure with this being instrumented by its second and third lag. The marginal effect is −0.015 but it is insignificant, consistent with the results in Table I.

Since the only significant results are for the rural sample, I replaced total state income with alternative measures of average income that are specific to the rural sector. The marginal effect of health expenditure is larger, and remains significant. Relative to the benchmark model where the long run marginal effect is −0.020, it is −0.025 (agricultural income), −0.029 (mean consumption), −0.030 (rural wage) using the alternative measures.

We expect that infant mortality is lowered by improvements in education. We have already controlled for the educational levels of parents. I also included state education expenditure in the model. This is relevant to the extent that state education expenditure is correlated with both infant mortality and state health expenditure. The coefficient on education expenditure is negative but insignificant, and the marginal effect of health expenditure is not altered.

The estimates in this paper are likely to be conservative for the following reasons. The survey only records births of mothers who survive until the survey date. If these are relatively healthy mothers, then we have a selectively low-risk sample of children. Also, the survey records only live births and these may, again, be a selectively healthy sample. Also, there may be effects of health expenditure (for example, on antenatal care) that affect survival until birth but, with these data, we cannot measure these.[11]

Having found significant heterogeneity by sector (rural/urban) in the health expenditure effect, heterogeneity by social class (micro-demographics) was further investigated for the sample of rural households. We now turn to a discussion of these results.

Heterogeneity by social group

The model in column 3 of Table III is re-estimated using different slices of the data. Not only is it interesting to consider differences in impact according to the attributes of the child or household, but looking at these differences provides insight into the ways in which health expenditure works – where it does – to reduce mortality. Results are in Table V. Every slice produces a significant difference in the health expenditure effect by sub-group. A general – and surprising – pattern that emerges is that health expenditure is *less* effective in reducing infant mortality in more vulnerable sections of society, that is, sections with relatively high mortality rates.[12]

For example, the marginal effect is larger, and significant, for boys but not girls, for high- but not for low-caste children, for Muslim but not Hindu children, for children other than the first born, for children of educated but not uneducated mothers, and for children born when the mother is in the relatively safe age range of 19–30 years. These differences are, of course, even larger when we look at the elasticity at the mean rather than the marginal effect (see Table V).

[10] If the individual-level mortality equation displayed in the fourth section is cast as a linear probability model, aggregation to the state level will produce a specification in which the level (not log) of mortality is the dependent variable. Deaton (2006) argues that the interesting question is whether or not *income growth* causes the *level of mortality* to decline. He shows that the evidence of such a relationship in cross-country data is much weaker than evidence of a relationship of income growth with *proportional* changes in mortality. If the same arguments apply when income is replaced with health expenditure (or share of), the specification estimated in this paper is the more conservative one.

[11] UN statistics on mortality rates are also calculated with reference to live births.

[12] Mortality rates and the sample contribution of each group are in Table V. The reported percentages of children in each group will differ from, for example, census proportions of these social groups to the extent that there is differential fertility across groups. Also note that these are figures for rural India.

DOI: 10.1002/hec

Table V. Heterogeneity in the health expenditure effect by population sub-group

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Sector		Gender		Caste			Religion		
	Rural	Urban	Boys	Girls	High caste	Low caste	ST	Hindu	Not Hindu	Muslim
*L*3 (health expenditure)	**−0.020**	−0.010	**−0.026**	−0.016	***−0.020***	−0.016	**−0.089**	**−0.019**	−0.024	**−0.044**
	[2.17]	[1.06]	[2.20]	[1.14]	[1.87]	[1.16]	[3.12]	[1.75]	[1.45]	[2.33]
Elasticity	−0.211	−0.108	−0.271	−0.168	−0.196	−0.158	−0.817	−0.189	−0.434	−0.694
Income	−0.027	0.051	−0.025	−0.032	−0.019	−0.027	0.035	−0.031	−0.013	−0.043
	[2.47]	[2.59]	[1.89]	[2.36]	[0.97]	[2.05]	[0.90]	[2.14]	[0.27]	[0.80]
Elasticity	−0.283	0.537	−0.257	−0.342	−0.193	−0.315	0.375	−0.311	−0.173	−0.548
Mean of dep. var.	0.093	0.059	0.0939	0.0958	0.0835	0.0997	0.100	0.0987	0.0735	0.079
N	117 088	35 783	61 002	56 086	38 360	77 225	13 820	98 884	18 204	13 136
% of group	69	31	52.1	47.9	33.2	66.8	12.0	84.4	15.6	11.2

	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)	(21)
	Birth order		Mother's education			Father's education			Maternal age at birth		
	First born	Other	None	Some	Higher	None	Some	Higher	9–18	19–30	31–49
*L*3 (health expenditure)	−0.009	**−0.023**	−0.014	**−0.035**	**−0.001**	−0.009	−0.009	−0.019	−0.006	**−0.028**	0.011
	[0.31]	[2.75]	[1.54]	[2.41]	[2.65]	[0.31]	[0.31]	[1.38]	[0.26]	[3.02]	[0.68]
Elasticity	−0.091	−0.245	−0.138	−0.738	−4.858	−0.091	−0.091	−0.402	−0.047	−0.403	0.160
Income	−0.034	−0.028	−0.031	−0.024	0.065	−0.018	−0.034	−0.020	0.000	−0.045	−0.078
	[1.57]	[1.94]	[2.37]	[1.24]	[0.46]	[0.72]	[2.10]	[0.48]	[0.01]	[3.92]	[1.09]
Elasticity	−0.353	−0.302	−0.299	−0.358	1.889	−0.173	−0.392	−0.284	−0.003	−0.511	−0.859
Mean of dep. var.	0.0968	0.0942	0.1042	0.0685	0.0346	0.0968	0.0968	0.0709	0.1223	0.0878	0.0904
N	26 737	90 351	86 305	30 783	4017	26 737	26 737	20 505	22 993	83 818	10 277
% of group	22.8	77.2	73.7	26.3	3.5	22.8	22.8	17.5	19.6	71.6	8.8

Notes: The specification estimated is that in column 3 of Table III. Standard errors are robust and clustered at the state level. Absolute *t*-statistics are in parentheses. *L*3 denotes the third lag of log health expenditure p.c. The reported marginal effect for income is the long run effect derived from its third and fourth lag. Elasticities are calculated at the mean for the sub-group; these means are shown in the table. Controls include state and year effects and state-specific trends, micro-demographics and rainfall shocks. Except in the case of column 2, the sample is restricted to rural households. The last row shows the sample percentage of each sub-group. The category 'Not Hindu' includes Muslim, but I further show results for Muslims alone. Higher education is defined as completion of secondary or higher. The samples are created separately for mother's and father's education. In the sample of children whose fathers have no education, 93% of mothers have no education. However, in the sample whose mothers have no education, only 51% of fathers have no education.

What is especially interesting about these divisions is that at least some cannot be explained by variation in the reach or the quality of public services. Children of different gender and birth order co-reside, so they are exposed to the same family and 'policy' environment. Technological differences in health production would predict larger effects of health spending on girls and first borns, the groups that have higher baseline mortality rates. But we find the opposite. A possible explanation of this is that it reflects the way in which the household *uses* its policy environment. In the case of gender, the results are consistent with the widely documented fact of son preference in India. In particular, Basu (1989) shows that, conditional upon being sick, boys are more likely to be taken to a treatment centre than are girls. In the case of birth order, the results can be rationalised in terms of learning. If the first born dies of diarrhoea, the mother is more likely to learn about Oral Rehydration Therapy and use it to avert death for subsequent children.

The effects of maternal education, age at birth and religion can be explained in a similar way. Educated mothers are likely to be better informed and so to extract a greater marginal advantage from a given level of health expenditure (see Jalan and Ravallion, 2003), prime-aged mothers might be more aware than teenage mothers, and Muslim mothers might, for example, exercise higher standards of sanitation within the home if regular prayer is associated with the requirement of regular washing. The complete absence of any health expenditure effect for women with no education is striking because maternal education creates especially large differences in mortality risk: average infant mortality of children of uneducated mothers is 10.4%, falling to 6.9% for mothers with some (non-zero) education, and to 3.5% for mothers with secondary or higher education.

There are two deviations from the pattern described so far, that is, cases in which the more vulnerable group is more responsive to health expenditure. This, of course, is what we would expect on account of diminishing returns, and because better-off groups can afford to protect themselves against infant mortality even when state health services are weak. One case, that we have already encountered, is that health expenditure is more effective in rural than in urban areas. Mortality risk is 3.4 percentage points or 60% higher for rural as compared to urban children. Even if health services are more sparse and variable in rural areas, there is greater scope for bringing down mortality. The other deviation is evident only when the low-caste group is sub-divided into its three components, which are scheduled castes (SC), scheduled tribes (ST) and 'other backward classes' (OBC). State health expenditure has only small and insignificant effects on the SC and OBC groups, but it has a large negative effect on children of ST. Indeed, this is the largest marginal effect of any sub-group, about four times as large as the average effect in rural areas. The ST group are about 12% of the entire sample and 18% of the low-caste group. The infant mortality rate in the ST group is 10%, in contrast with 8.35% amongst high-caste Hindus. ST are thought to be the least integrated social group, historically having been isolated from community life, and tending to live in relative geographic isolation. This result is therefore quite striking.[13]

Some previous studies have found bigger impacts of state health expenditure on the poor (e.g. Bidani and Ravallion, 1997; Gupta *et al*., 2003). As discussed in the second section, we do not have household income data.[14] Since rural and ST households are clearly relatively poor, there is some support in these data for the view that state health expenditure is, at the margin, more beneficial to the poor. However, uneducated rural women are poor, and we find that health expenditure has no effect on the infant death risks they face. Father's education may be a better indicator of the permanent income of the household. But we find no significant variation in the health expenditure effects by father's education. Overall, with the exception of ST, it seems that the most poor (rural and uneducated) and the better off (urban) do not benefit as much as the group in the middle (rural but educated). Given the difficulties with measuring income for poor, often self-employed, households, it is useful to look at heterogeneity by these other, more stable, indicators of social class.

It is interesting that the pattern of income effects is not the same as the pattern of health expenditure effects (Table V). Indeed, in most cases, the differences are reversed! (Negative) income effects are larger for the more vulnerable groups: rural children, girls, low caste, Hindu, first borns and children of mothers with no education. Recall that the effects of each of health expenditure and income are obtained conditional upon the other. This contrast between their distributional impact is consistent with complementarities between state health expenditure and personal attributes (education, information) that bias its effectiveness away from those individuals who need it most.

CONCLUSIONS

This paper has identified a significant impact of variations in state health expenditure, given state income, on the risk of infant mortality in rural India, the long run elasticity being -0.24. We are unable to identify a corresponding effect amongst urban households. The analysis shows that failing to allow

[13] ST distinguish themselves from other social groups (including the SC) in having higher infant mortality rates for boys as compared to girls. This may be pertinent, although why exactly is unclear.

[14] Nor do the two studies cited here. They estimate the distribution of effects under sometimes strong assumptions – discussed in Gupta *et al*. (2003).

for heterogeneity, lagged effects and state-specific trended unobservables results in under-estimation of the beneficial effects of health expenditure, and I have argued that this might explain some of the negative findings in the literature. The identified effect is robust to controls for state-specific rainfall shocks and other state-level variables including inequality and media prevalence (an indicator of information and awareness). Although it is encouraging to find an impact, we also note that the share of health expenditure in national income is only about 1% and it appears to be decreasing in income.

There is considerable heterogeneity in these effects. Strikingly, health expenditure appears to bring no benefit to some of the most vulnerable sections of society. I have argued that this may reflect complementarities between public and private (parental) inputs (e.g. Jalan and Ravallion, 2003). It may also be related to the fact that the composition of state health expenditure in India is non-progressive. It can be made more progressive by shifting allocations in favour of public health, water and family welfare programmes in rural areas and, within rural areas, by improving information and access for politically and socially disadvantaged groups.

The effectiveness of health expenditure varies across the states, displaying a pattern that bears no evident relation to initial levels of mortality or income. As state governments control delivery of health expenditure, this raises questions concerning public service delivery (e.g. World Bank, 2003), and of the political economy issues that this involves (see Besley, 2006). There is some evidence that these problems are quite severe in India, possibly more so in the poorer states (Public Affairs Centre, 2002). A recent initiative of the central government of India, the National Rural Health Mission, with horizon 2005–2012, holds promise. Public spending on health over this period is expected to rise to '2–3%' of GDP (see NRHM, 2005). More importantly, it aims to undertake 'architectural correction' of the health system, promoting service delivery, for example, by increasing decentralisation to the village level and instituting a female health activist in each village. The analysis in this paper needs to be repeated six years from now!

ACKNOWLEDGEMENTS

I am grateful to Frank Windmeijer and Ellen van der Poel for helpful discussions. Comments from the Editor, Owen O'Donnell, and two anonymous referees have greatly improved the paper.

REFERENCES

Anand S, Ravallion M. 1993. Human development in poor countries: on the role of private incomes and public services. *Journal of Economic Perspectives* **1**(3): 133–150.

Arulampalam W, Bhalotra S. 2007. Sibling death clustering in India: state dependence vs unobserved heterogeneity. *Journal of the Royal Statistical Society*, *Series A* **169**(4): 829–848.

Basu AM. 1989. Is discrimination in food really necessary for explaining sex differentials in childhood mortality? *Population Studies*: *A Journal of Demography* **43**(2): 193–210.

Bertrand M, Duflo E, Mullainathan S. 2004. How much can we trust difference-in-difference estimators?. *The Quarterly Journal of Economics* **February**: 249–275.

Besley T. 2006. *Principled Agents*?: *The Political Economy of Good Government*. Oxford University Press: Oxford.

Besley T, Burgess R. 2004. Can labour regulation hinder economic performance?: evidence from India. *Quarterly Journal of Economics* **119**(1): 91–134.

Bhalotra S. 2007. Income volatility and infant death. Mimeograph, University of Bristol.

Bhalotra S, van Soest A. 2007. Birth spacing, fertility and neonatal mortality in India: dynamics, frailty and fecundity, *Working Paper 07/168*, Centre for Market and Public Organisation, University of Bristol.

Bidani B, Ravallion M. 1997. Decomposing social indicators using distributional data. *Journal of Econometrics* **77**(1): 125–139.

Black R, Morris S, Bryce J. 2003. Where and why are 10 million children dying every year? *The Lancet* **361**: 2226–2234.

Cameron C, Trivedi P. 2005. *Microeconometrics: Methods and Applications*. Cambridge University Press: Cambridge.

DOI: 10.1002/hec

Cutler D, Deaton A, Lleras-Muney A. 2006. The determinants of mortality. *Journal of Economic Perspectives* **Summer 20**(3): 97–120.

Cutler D, Miller G. 2005. The role of public health improvements in health advances: the twentieth century United States. *Demography* **42**(1): 1–22.

Deaton A. 2006. Global patterns of income and health: facts, interpretations and policies. *WIDER Annual Lecture*, September. WIDER: Helsinki.

Deaton A, Paxson C. 2004. Mortality, income and income inequality over time in Britain and the US. In *Perspectives in the Economics of Aging*, Wise D (ed.). University of Chicago Press: Chicago.

Deolalikar A. 2005. *Attaining the Millennium Development Goals in India*: *How Likely and What Will it Take to Reduce Infant Mortality*, *Child Malnutrition*, *Gender Disparities and Hunger-Poverty and to Increase School Enrollment and Completion*? Oxford University Press: New Delhi.

Durlauf S, Johnson P, Temple J. 2005. Growth econometrics. In *Handbook of Economic Growth*, vol. 1A, Chapter 8, Aghion P, Durlauf SN (eds). North-Holland: Amsterdam, 2005; 555–677.

Filmer D, Pritchett L. 1999. The impact of public spending on health: does money matter? *Social Science and Medicine* **49**(10): 1309–1323.

Fogel R. 2004. *The Escape from Hunger and Premature Death, 1700–2100 – Europe, America and the Third World*. Cambridge University Press: Cambridge.

Gupta S, Verhoeven M, Tiongson E. 2002. Does higher government spending buy better results in education and health care? *European Journal of Political Economy* **18**(4): 717–737.

Gupta S, Verhoeven M, Tiongson E. 2003. Public spending on health care and the poor. *Health Economics* **12**: 685–696.

IIPS and ORC Macro. 2000. *National Family Health Survey* (*NFHS-2*) 1998–1999: *India*. International Institute for Population Sciences (IIPS): Mumbai.

Jalan J, Ravallion M. 2003. Does piped water reduce diarrhoea for children in rural India. *Journal of Econometrics* **112**(1): 153–173.

Jones G, Steketee R, Black R, Bhutta Z, Morris S, the Bellagio Child Survival Study Group. 2003. How many child deaths can we prevent this year? *The Lancet* **362**: 65–71.

Koren M, Tenreyro S. 2007. Volatility and development. *Quarterly Journal of Economics* **February**: 243–287.

Lane P. 2003. The cyclical behavior of fiscal policy: evidence from the OECD. *Journal of Public Economics* **87**: 2661–2675.

Lanjouw P, Ravallion M. 1999. Benefit incidence, public spending reforms and the timing of program capture. *World Bank Economic Review* **May 13**(2): 257–273.

NRHM. 2005. Mission documents of the National Rural Health Mission, Government of India, available at http://mohfw.nic.in/nrhm.html.

Paxson C, Schady N. 2005. Child heath and economic crisis in Peru. *World Bank Economic Review* **19**(2): 203–223.

Peters D, Yazbeck A, Sharma R, Ramana G, Pritchett L, Wagstaff A. 2002. *Better Health Systems for India's Poor*, *Human Development Network*. The World Bank: Washington, DC.

Pritchett L. 2000. Understanding patterns of economic growth: searching for hills among plateaus, mountains and plains. *The World Bank Economic Review* **14**(2): 221–250.

Pritchett L, Summers LH. 1996. Wealthier is healthier. *Journal of Human Resources* **31**(4): 841–868.

Public Affairs Centre. 2002. *Benchmarks for the New Millennium*: *State of India's Public Services*. Public Affairs Centre: Bangalore, India.

Temple JRW. 1999. The new growth evidence. *Journal of Economic Literature* 1999; **37**(1): 112–156.

The Tribune. 2002. Reducing the infant mortality rate: a big challenge. *The Tribune*. Perspective feature on the Editorial Page, Sunday, 16 June, Chandigarh, India.

Woo J. 2005. The behaviour of fiscal policy: cyclicality and discretionary fiscal decisions. Mimeograph, Kellstadt Graduate School of Business, DePaul University.

World Bank. 2003. *World Development Report 2004*: *Making Services Work for Poor People*. The World Bank: Washington, DC.