AGRICULTURAL COMPETITIVENESS: MARKET FORCES AND POLICY CHOICE

PROCEEDINGS
OF THE
TWENTY-SECOND
INTERNATIONAL CONFERENCE
OF AGRICULTURAL ECONOMISTS

Held at Harare, Zimbabwe
22–29 August 1994

Edited by
G.H. Peters, International Development Centre,
Queen Elizabeth House, University of Oxford, England
and
Douglas D. Hedley, Agriculture and Agri-Food, Canada

INTERNATIONAL ASSOCIATION OF
AGRICULTURAL ECONOMISTS
QUEEN ELIZABETH HOUSE
UNIVERSITY OF OXFORD

1995

Dartmouth

JOHN STRAUSS AND DUNCAN THOMAS*

Health and Labour Productivity: Sorting out the Relationships

INTRODUCTION

A very large literature has documented, time and time again, that there are substantial returns to investments in human capital in a wide variety of settings: in developed and developing countries and in both the agricultural and non-agricultural sectors. While the correlation between human capital and productivity is typically positive, the precise mechanisms underlying this correlation are still debated. In these studies, human capital is typically measured in terms of years of schooling or final qualification attained. Some studies have also examined the returns to both formal and informal training as well as on-the-job experience. Others have attempted to examine the impact of outputs from schooling, such as scores on cognitive tests.

Our goal is to argue that human capital is much broader than education, training and test scores but encompasses a wide array of skills and attributes, some of which are readily observable, while others are not. In particular, we highlight health status and its relationship with income. This is a relationship that has featured prominently in the economic history, labour and development economics literatures (see Behrman, 1993; Strauss, 1993; Strauss and Thomas, 1993; for reviews of the development literature) and the relationship lies at the heart of several efficiency wage theories (Leibenstein, 1957; see Dasgupta, 1993, for a review). This paper focuses on behavioural choices that affect health and labour productivity and, in particular, the relationship between the two. Specifically, we examine theoretical and empirical issues in the estimation of the impact of a broad array of health indicators on wages, farm output, profits and costs.

MEASUREMENT OF HEALTH

Just as there are many possible measures of educational attainment, it is not obvious how to measure 'health status'. First, it is multidimensional and the full extent of health problems is not likely to be captured in any single index. Second, health status varies over the life course and many indicators can change quickly over time. It is important, therefore, to distinguish stocks of

*Department of Economics, Michigan State University, and RAND and University of California at Los Angeles, respectively. The authors are grateful to Janet Currie for comments.

health from flows. Since different dimensions of health are likely to have different effects on labour market outcomes, it also seems critical that a range of health indicators be adopted in empirical analyses. These may include, for example, days ill over a period, specific morbidities, problems with physical functions, self-perceptions of health status and anthropometrics, all of which can be thought of as outputs from a health 'production function'. Inputs to the 'production function', such as nutrient intakes, may also be relevant.

Health information can be collected by clinical examination, by objective assessment in a survey or by respondent self-reports. Clinical evaluations of health status are very expensive and so have tended to be restricted to either small or selected samples and typically samples with little socioeconomic information. Moreover, some studies have drawn data from health facility records. But people who use these facilities are themselves a select group and not necessarily those in poorest health. In fact, they tend to be higher-income people and so it will be very hard to infer anything about the relationship between health and income in the whole population without at least having information on the mechanisms underlying the choice to visit a health care facility.

Some household surveys collect self-reported health information from respondents. These data can be subject to serious, systematic measurement error. For example, a respondent might be asked about illness, or specific problems (such as fevers, diarrhoea or respiratory problems) during a reference period. The answer, however, will be subjective and what is deemed an 'illness' or a 'problem' may not be the same thing for all respondents. It is not unusual for the poorest in a survey to appear to be the most healthy, according to these measures. Such a pattern may arise for several reasons, including how much use is made of the modern health system. A commonly used variant on self-reported illness is to ask whether any days of 'normal' activity were lost to ill-health. This indicator is also problematic because 'normal' is not a well-defined concept and those with a high opportunity cost of time will have less incentive to miss activities. For all of these measures, therefore, it will be very hard to separate the effect on labour market outcomes of health problems from the role of respondent perception.

As an alternative to questions on illness, a few recent household surveys have collected information on difficulties with physical functioning, such as walking a kilometre, climbing stairs or preparing food. These are probably better defined than 'being ill' and there is some evidence that they are prone to less measurement error than information on morbidity. Furthermore, in some cases, they can be cross-validated with direct observation by the enumerator; for example, by asking the respondent to perform specific tasks (such as lifting a weight).

In response to concerns with measurement, some studies have resorted to anthropometrics such as height and weight. They can be measured by a trained enumerator at the time of the survey and are thus more objective health indicators. Height may be directly related to productivity. But it may also reflect human capital investments during childhood. Weight varies in the short run and Body Mass Index (BMI), which is the ratio of weight (in kilograms) to height (in metres), has been shown to be related to maximum physical capacity independent of energy intake. Extremely low and high values of BMI have also been associated with adult mortality, although the causal mechanism underlying this association has not been established. The cost of objectivity in measurement is that, relative to information on specific morbidities and ailments, anthropometrics are rather blunt indicators of health problems and are likely to measure different dimensions of health than, say, specific morbidities.

Another 'objectively' measured variable that has received considerable attention in this literature is energy input. For example, there is evidence that calorie intake is associated with increases in maximum oxygen uptake, which is, in turn, related to maximum work capacity. This suggests a productivity–nutrient link, although few jobs require maximum physical effort, so it is not obvious that energy or other nutrient intakes should be correlated with either productivity or labour supply. Furthermore, if the body is able to adapt to changes over some range in energy intakes, so that functioning is unaffected, then it is only at extremely low levels of calorie intake that productivity or labour supply should suffer. This suggests that, if there is a relationship between calorie intake and labour outcomes, it is likely to be non-linear. There is evidence suggesting that labour productivity is also related to intakes of proteins, iron and other micronutrients.

The measurement of nutrient intakes is not straightforward, especially at the individual level. Even at the household level, intake recalls or weighing of food consumed (usually for a short period of time such as 24 hours) are likely to be 'noisy' because there is considerable variation in daily intakes of most individuals. On the other hand, indirect estimates of food availability, from purchase, production, sales and other disappearance data may be systematically biased, as discussed in more detail below.

While measurement of health status is a serious concern, there are several new and exciting possibilities on the health indicator horizon. Under certain conditions, it is conceivable that large-scale household surveys could collect saliva or urine samples and those samples could then be linked to a wide array of health problems including a series of nutrient deficiencies. There is also some recent evidence that the presence of HIV can be detected in saliva samples. It may thus become possible to evaluate the productivity costs of specific health problems (such as malaria or HIV) which can, in turn, provide key information for the assessment of the benefits associated with particular interventions. Experiments are also under way to conduct more in-field cross-validation of self-reported health problems, particularly by more judicious use of direct observation. These and further refinements in the measurement of health status are likely to play a key role in improving the understanding of the relationship between health and labour market outcomes.

CORRELATIONS BETWEEN HEALTH AND LABOUR OUTCOMES

Are health and productivity correlated? Using data on men in urban Brazil, from the 1974/5 *Estudo Nacional da Despesa Familiar* (ENDEF), Figure 1 presents non-parametric estimates of the bivariate relationship between hourly market wages and two dimensions of health: height and body mass index.[1]

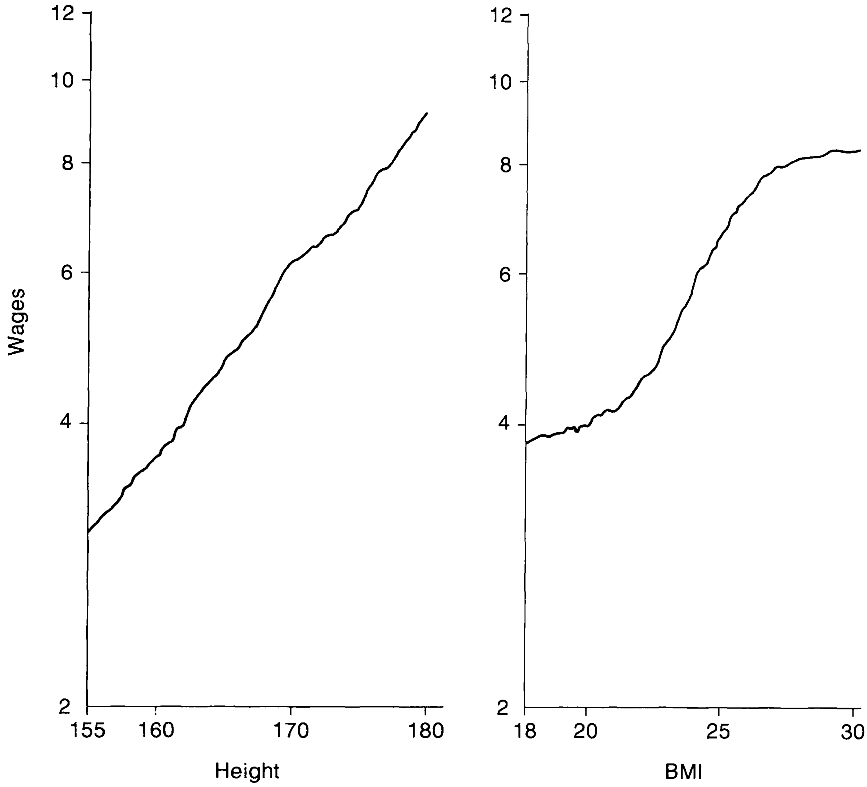

FIGURE 1 *Wages and health: Males in market wage sector in urban Brazil (non-parametric estimates)*

Source: *Estudo Nacional da Despesa Familiar* (ENDEF, 1974/5).

Since these non-parametric estimates impose no structure on the relationship between the variables, they truly allow the data to speak and provide a useful, simple description of the correlations.

The shapes in Figure 1 present clear evidence that health, or at least height and BMI, are correlated with labour outcomes. The scales are in logarithms, so a linear relationship implies a constant elasticity. This is the case for wages and height: taller men earn higher wages. (The same is true for women in these data.) The wage–BMI function is sigmoidal in shape. The curve is flat until BMI reaches roughly 20, then it rises steeply up to a BMI of about 26, after which it flattens. It is, perhaps, surprising that, below 20, higher BMIs are not associated with higher wages, since it is at these levels that BMI is thought to be associated with health risks, such as adult mortality. It may be that a threshold level of BMI is required in order for productivity improvements to show up or it may be that these bivariate correlations reflect (endogenous) behavioural choices. Before discussing the interpretation of these correlations in some detail, the next section considers the circumstances under which a relationship between health and labour outcomes might be observed in survey data.

IMPACT OF HEALTH ON LABOUR OUTCOMES

A correlation between health and wages may arise because better health improves productivity which is, in turn, reflected in higher wages. This presupposes that wages reward productivity directly. This will occur if productivity is directly and costlessly observed, as would be the case for the self-employed, for piece-rate workers and possibly for some workers who earn time rates. However, for many workers, especially workers on short-run contracts paid by time, productivity is not easily observed; these include workers paid daily, for whom monitoring is costly. Assume, for the moment, that healthier workers are more productive, but their productivity is not observable. Health will be reflected in the wages of these workers only if it (or something correlated with it) is observable to the employer. If health is not rewarded in the labour market, because of costly monitoring, say, then in the long run one would expect the development of alternative labour or land contractual forms (such as sharecropping or longer-term contracts) that minimize problems associated with observability of productivity. One would also expect healthier (and more productive) workers to move selectively into sectors that do reward health, such as self-employment or piece-rate work.

Thus a key factor determining whether better health is translated into higher market wages is the observability of both productivity and health. It was pointed out above that there are many potential indicators of health and they may have different effects on productivity. It should be apparent that they also differ in their degree of observability: for example, it may be difficult for employers to monitor nutrient intakes, even when they feed their workers (who may divert food to other family members by eating less food at home). Thus it is not obvious that nutrient intakes will have a positive impact on daily wage workers in an environment with costly monitoring. But outputs related to nutrient intakes, such as BMI, are readily observed and so, even among these workers, BMI may be rewarded either because it is correlated with nutrient intakes or because it has an independent effect on wages or both.

Assuming that health does have a positive impact on productivity which is reflected in wages (or output), then worker *full* income will be higher with better health status. Whether *labour* income is higher, however, depends crucially on the relationship between health and labour supply. Healthier people may choose to work less, in which case earnings may be negatively associated with health. Recently, considerable attention has been given to the impact of health on profits, particularly among farm households and those operating non-farm self-employment enterprises. See, for example, Antle and Pingali (1994) and other papers in the August, 1994, issue of the *American Journal of Agricultural Economics*.

However, even under the assumption that health does affect productivity, this may not be reflected in costs or profits: whether it does depends critically on the efficiency of markets (Pitt and Rosenzweig, 1986). Suppose labour markets are competitive, labour effort is observable and that health augments labour in units of effectiveness. The equilibrium market wage will reward an effective unit of labour (for example, a kilogram of maize harvested) and the demand for farm labour will be measured in effective units of labour. If family and hired labour are perfectly substitutable, in common efficiency units, then if a worker or family member falls ill, others can be hired to meet the effective labour demand. The health of family and hired workers will have no impact on profits or costs because any differences in units of effective labour will be exactly offset by the nominal wage differential. Profits and costs will remain invariant to health even if family and hired labour are not perfect substitutes, as long as at least one family member works off-farm and off-farm labour can be adjusted to meet the desired amount of farm family labour, measured in efficiency units.

This argument does *not* mean that health has no positive impact on labour productivity, farm output or full income. In this model, better health increases the units of effective labour and hence raises full income and output, but has no direct impact on farm profits or costs, because labour markets function well. However, if that assumption is false, say, for example, because health affects managerial ability (and markets do not exist for managers), then health may directly affect profits. It is important to recognize, therefore, that inferences based on studies of the impact of health on farm profits or costs involve *joint* tests of a health and productivity link as well as market failure.

SIMULTANEITY, UNOBSERVABLES AND MEASUREMENT ERROR

The fact that there is a correlation between health and labour outcomes in Figure 1 says nothing about causality. It may be that better health makes a worker more productive, but it may also be that higher income is spent on improving one's health. Without an understanding of the behaviours that underlie the production of health, it is impossible to distinguish these mechanisms. The simultaneous determination of health and labour outcomes will, therefore, result in ordinary least squares (OLS) estimates of the impact of current health status on labour outcomes being afflicted with simultaneity bias.

Furthermore, it is not just the effect of current health status that may be biased in these regressions. Consider height, for example. It is certainly predetermined by adulthood, but it may reflect previous health investments made primarily early in life. If these investments are correlated with other omitted characteristics that affect labour outcomes, such as entrepreneurial ability, then estimates of even predetermined health indicators will be biased because they are picking up the impact of unobserved, time-persistent heterogeneity. Health is not unique in this respect: the same issue arises with estimating the effect on labour outcomes of education, which is also predetermined by adulthood. Of course, estimates of the impact of current health status may also suffer from bias due to unobserved heterogeneity, in addition to 'true' simultaneity bias discussed above.

The issue of causality is clearly crucial and has become one of the dominant themes in the current literature on health and productivity. The two most commonly used empirical methods to address the issue are instrumental variables (IV) and fixed effects (FE) estimators. Neither is a panacea for all potential problems and, as with any empirical work, each case needs to be judged on its own merits. It is also important to take into consideration other potential empirical problems when choosing among estimators (and specifications). For example, some health indicators (such as nutrient intakes and ADLs) are likely to be measured with error. This can be addressed in an IV framework – although, as discussed below, it affects the choice of instruments. In a FE framework, however, the cure can be worse than the disease if a good deal of the signal in the health measure is swept out in the fixed effect (along with all unobserved heterogeneity) leaving behind mostly noise. It is, thus, possible for FE estimates to be even *more* biased than OLS.

Turning first to IV estimates, the conditions under which they are consistent are well known: the identifying instruments must be uncorrelated with any unobservables in the regression of interest, such as the wage function. But satisfying this statistical condition alone is not enough. First, there must be a *theoretical* justification for excluding instruments from the second-stage regression: clearly, the choice of those variables will depend on whether the regression of interest is a wage function, production function, profit or cost function. It will also be governed by the specific health measures under consideration.

There is an important second issue which has not received the attention it deserves. It is typically assumed that, if instruments are weakly correlated with the regressand in the first stage (health, say), then second-stage estimates are unbiased but inefficient; the standard errors are large and so the problem is readily detected by a lack of power. However, that assumption is false. Recent results demonstrate that, when first-stage predictions are poor, the second-stage coefficient estimates are not only biased but the associated t-statistics can be biased *upwards*. That is, it is possible for the t- statistics in the second stage to indicate a significant impact of health (in this case) on the labour outcome when, in fact, none exists: the association is completely spurious and driven by the weak correlation between the instruments and first-stage regressand. In this literature, and the human resource literature more generally, this is a concern of considerable import since outcomes, such as health, incorporate considerable heterogeneity and are difficult to predict well. How well is 'well enough' remains an open issue; Staiger and Stock (1993) offer some suggestions and guidance. At the very least, studies should routinely report test statistics for the joint significance of *identifying* instruments in first-stage regressions along with the associated partial R^2. Insignificance of the instruments should obviously cause the researcher to at least pause.

In principle, the prices of health inputs and outputs are potential identifying instruments for health. While explicit prices do not exist for many indicators (such as measures of physical functioning), the implicit prices of these outcomes will include the monetary prices of health care visits and time costs of travelling to (and waiting at) facilities. More generally, measures of the availability and quality of health services (such as clinics and hospitals) as well as health-related infrastructure (such as water and sanitation) in the community may serve as instruments. Interactions between infrastructure and household characteristics, such as education or assets, may also be good instruments if, for example, service quality varies with household resources. It is, however, important to control for other more general infrastructure (such as transport or industrialization) in the income function. Otherwise, all community infrastructure effects will be forced to operate through health for which there is little *a priori* rationale. Even then, this identification strategy is not without problems: it assumes that local health infrastructure is exogenous and rules out the possibility that programmes are systematically located (in, say, better endowed areas).[2]

Identification is possibly slightly easier in the case of nutrition-related inputs or current health outputs since food prices are natural instruments. Food prices affect consumption and thus nutrient intakes as well as anthropometric outcomes such as body mass; they are unlikely to be good instruments for longer-run indicators of nutritional status, such as height which is determined at an early age. It will be important to use food prices *relative* to an aggregate price index in order to control for cost of living differences, which would be reflected in nominal wages, profits and costs. (Thus the second-stage regressions should include a local price index for essentially the same reason general infrastructure should be included, as discussed above.) Again, interactions between household characteristics and relative food prices might serve as additional instruments.

The use of food (and health) prices for identification relies critically on the assumption that they do not belong in the second stage. Whether this is reasonable varies with the outcome: output, market wages, wages in self-employment, profits or costs.

The best case can be made for farm (or firm) production functions. Farm input and output prices, as well as prices of health inputs and outputs, are clear potential instruments and this is the strategy used by Strauss (1986) in his study of farm production in Sierra Leone. He takes care to include community characteristics that belong in the production function and are likely to be correlated with prices, such as agroclimatic and soil variables, to avoid inducing a spurious correlation between prices and the error term in the production function. Wage functions for market (or off-farm) labour can also be identified using prices and, possibly, household assets and household composition characteristics (Sahn and Alderman, 1988; Foster and Rosenzweig, 1992).[3] Relative food prices are also legitimate instruments for health in cost functions.

Wages from self-employment and profit functions are trickier.[4] Prices of non-health inputs and all outputs belong in these functions and, if the inputs or outputs include foods, then those prices will have a direct impact on wages and profits. They are, therefore, not valid instruments. Obviously, prices of goods that are not potential inputs or outputs in the self-employment enterprise, family firm or farm remain valid instruments. The same problem arises with labour supply and sectoral choice. In these cases, we are reduced to relying on local health prices and health infrastructure for identification; see, for example, Pitt and Rosenzweig (1986) who examine the effects of illness on the profits of Indonesian farmers. As noted, however, if community characteristics are not exogenous then even they are not valid instruments. Furthermore, since there is not necessarily a direct link between health status and these characteristics, they may only be weakly correlated, in which case the second-stage estimates may be seriously misleading, as discussed above.

Household characteristics, such as land owned, non-labour income and possibly household composition (such as the number of adult males and females) may also serve to identify health, although this involves stronger assumptions. For example, there can be no unobservables, such as unmeasured land quality or managerial ability, that affect these characteristics and also farm output (or profits or costs). If farmers who are more productive (controlling for all observables) tend to save more, have more land or live in bigger (smaller) households, the assumption will be violated. A similar caveat arises with a farmer's education: if it has an impact on technical efficiency then it belongs both in the health function and the production function. Likewise, if education raises allocative efficiency, it will belong in profit and cost functions. A similar argument can be made regarding household composition variables in the case of profit and cost functions: the farm household model must be recursive in production and consumption decisions, conditional on health, for household demographics to be properly excluded from the profit or cost functions (Benjamin, 1992).

The discussion thus far has considered only cross-section data. If prices vary over time, as well as across space, then, for health indicators that have a stock dimension, such as body mass or general health, lagged health and food prices (or their proxies) will also provide information about current health; they may, therefore, be additional instruments. This is one dimension in which longitudinal data can be very helpful in quantifying health–productivity relationships.

Panel data with multiple observations on the same individual offer several additional and potentially important advantages. Repeated observations make it possible to control for all individual-level unobservables that are fixed over time and thus concentrate on the impact of time-varying health indicators such as body mass, nutrient intakes and physical functioning. That is, fixed effects estimates place the spotlight on the relationship between health flows (rather than stocks) and changes in productivity and may be applied to wage functions, production functions, profit or cost functions. (See Deolalikar, 1988; Behrman and Deolalikar, 1989; Haddad and Bouis, 1991). While all time-invariant heterogeneity is swept out by the fixed effect in these models, time-varying heterogeneity is not. Without strong assumptions, therefore, the direction of causality is not obvious in these models.

For example, consider a farmer who experiences a surprisingly good year and spends some of the unanticipated income on improving his health (by eating more nutritious food, for example). It is not clear whether a positive association between changes in health and changes in productivity reflects better health causing productivity to be higher, whether causality is in the reverse direction, or whether the two are affected by a common unobservable (such as good rainfall). By combining fixed effects and IV estimates, it is possible to address (or at last test) this concern as long as there are plausible and good instruments. However, the data demands are not trivial: Haddad and Bouis (1991) attempt this strategy but do not report the estimates because of their fragility.

A third method that has been used to quantify the impact of health on productivity is to attempt first to measure the exogenous part of health, call it the health 'endowment', and to use it as a regressor in a wage or labour supply function. For example, Pitt *et al.* (1990) first estimate a production function for adult weight for height. Using the residual as an estimate of the individual's biological health endowment, they examine its impact on calorie consumption, household income and labour supply of men and women in Bangladesh.

Observing that women consume significantly fewer calories than men, they ask whether this reflects discrimination against women in household allocations. They demonstrate that, relative to other men, those with a better weight-for-height endowment have higher energy intakes and are also more likely to work in strenuous occupations. Furthermore, after controlling for household characteristics, the average household endowment of adult males is positively related to household income. This implies that better endowed, or healthier, men are more productive, or at least earn more. The results are certainly suggestive that intra-household allocation of resources may be efficient if not, at first blush, apparently equitable. Pitt *et al.* directly address the equity–efficiency trade-off and point out that men with better endowments are taxed, indicating that households do, in fact, care about equity.

Treating estimated residuals from a health production function as endowments is not without its problems. The authors are forced, by data constraints, to make several simplifying assumptions. These include estimating a single weight-for-height production function for all men and women and also including in the production function only contemporaneous inputs, such as individual calorie consumption and indicators for currently working in a strenuous occupation. However, weight for height is, in part, a stock measure of health. Thus the residuals will embody not only endowments but also the influence of all omitted factors that should enter the production function.

As discussed above, failure to control for simultaneity and unobservables is not the only pitfall in measuring health–productivity relationships. Measurement error poses another, related concern. Furthermore, measurement error may be either random or systematic and the implications of the two types of error for analyses are very different. Random measurement error in health status is likely to arise in several commonly used indicators. For example, nutrient intakes are often based on food recalls over the previous 24 hours. Recall data are notoriously 'noisy' and so some surveys have measured intakes by weighing food before preparation and taking account of all wastage afterwards. But diets are not the same every day and thus long-run food intakes will be measured with error even with this method. The impact of these errors may be minimized, at some expense, by multiple visits to the same household (over a week, for example) and calculating average intakes. Even weight and height measurements are not immune to problems if measuring instruments break down in the field (or are not recalibrated frequently enough); again, multiple measures (or good supervision, or both) should minimize the impact of these errors. Physical functioning, morbidities and functional problems are also prone to being measured with a good deal of (random) noise. In addition to using averages of multiple measures, instrumental variable methods along the lines discussed above are commonly adopted to purge estimates of bias due to measurement error.[5] Fixed effects estimates, however, will tend to exacerbate these biases and estimates of health residuals are also likely to be contaminated in the presence of measurement error.

Systematic measurement error is potentially a very serious problem for several classes of health indicators. The difficulties associated with self-reported illness or days missed of normal activities are discussed above. If higher wage earners are more likely to report themselves as being ill, then it will be very difficult to disentangle this systematic reporting error from an underlying association between poor health and wages. The same argument may be applied to all the 'subjective' measures of health status. Multiple reports by the same respondent do not help in this case (assuming the systematic error is time-invariant) and it is hard to think of instruments that are likely to be correlated with 'true' health status but not with wages. For example, if the reason people report themselves as being sick is that they are better informed, which is, in turn, a reflection of the availability of health facilities in the community, health infrastructure will be correlated with the measurement error and so is not a valid instrument.

It is not just subjective health indicators that are prone to systematic measurement error. For example, as discussed above, in many surveys, nutrient intakes are not directly measured but derived from estimates of food availability; that is, by combining data on market purchases of food with estimates of food consumed from own production. The latter are often estimated through disappearances: production less sales, less changes in stocks and so forth. These estimates tend to overpredict intakes for high-income households and under predict for low-income households, in part because higher-income households are more likely to give food to workers or guests and are likely to waste more (Bouis and Haddad, 1992). If wages and productivity are positively related to income, then this systematic measurement error will result in wages and productivity appearing to be related to nutrient intakes. Furthermore, since the measurement error is correlated with income and assets, they are not appropriate instruments in this case.

EMPIRICAL EVIDENCE

As an empirical example, the same Brazilian data presented in Figure 1 are used to examine the impact of a series of health indicators on (log) wages of urban men in a multivariate framework. In addition to height and BMI, the analysis includes two inputs into the health production function, calorie and protein intakes which are measured at the household level and thus are converted into a per capita basis.[6]

Height, like education, is treated as predetermined. While both may reflect investments in human capital prior to adulthood, along with unobserved family background characteristics, good instruments for these choices are hard to come by: the instruments should, for example, reflect the environment the respondent lived in while a child (such as the price or availability of education, health and nutrition) but have no direct effect on productivity. This information is not contained in these survey data and so the potential correlations between height (and education) and unobservables in the wage function are ignored.

However, BMI and nutrient intakes vary in the short run and may be simultaneously determined with wages; this is straightforward to test by comparing OLS with IV estimates, assuming the availability of good instruments. All three health indicators are nutrition-related and so relative food prices are candidates for instruments; prices of ten foods are included in the reduced-form estimates reported in Table 1. The regressions also control for an aggregate price index (so that food prices can be interpreted in relative terms),[7] the respondent's education, age and its square, non-labour income, non-labour income of other household members, dummies for the 13 months of the survey and the Northeast region which is the poorest part of the country. Table 1 reports only the price results.

TABLE 1 *Effect of relative food prices on health indicators*

	ℓn(BMI)		ℓn(calorie intakes)		ℓn(protein intakes)	
Covariates	β	*t*	β	*t*	β	*t*
Relative food prices						
beans	–0.04	(2.2)	0.07	(2.3)	0.11	(2.9)
rice	–0.01	(0.4)	0.36	(5.1)	0.09	(1.0)
tuber	–0.00	(0.1)	0.09	(2.4)	–0.12	(2.8)
wheat	–0.02	(1.0)	–0.23	(5.4)	–0.40	(7.9)
dairy products	0.06	(0.9)	–0.08	(0.7)	0.05	(0.4)
fruit/vegetables	0.09	(3.8)	0.25	(5.6)	0.25	(4.8)
meat	–0.08	(1.2)	–0.98	(8.0)	–0.60	(4.1)
fish	0.07	(3.2)	0.30	(7.5)	0.04	(0.8)
oils	–0.20	(4.6)	–0.38	(4.8)	–0.67	(7.1)
sugar	0.02	(0.5)	–0.00	(0.1)	–0.17	(1.9)
Aggregate price index	–0.00	(0.2)	–0.18	(5.9)	0.20	(5.2)
F tests (p value)						
All prices	13.52	(0.00)	14.09	(0.00)	23.74	(0.00)
All covariates	34.36	(0.00)	87.80	(0.00)	91.00	(0.00)
R^2	0.07		0.16		0.16	

Note: Sample size = 15 074.

In these regressions, prices reflect both substitution and income effects and so cannot be signed on *a priori* grounds: BMI, calorie and protein intakes are all significantly affected by at least four of the relative food prices and the effect of some prices is negative; for others, it is positive. The critical point, for our purposes, is that the ten relative food prices, which are the identifying instruments, explain a significant fraction of the variation in the three health measures as demonstrated by the F-statistic for 'All prices' towards the bottom of the table. Apparently, the instruments do have predictive power.

However, there is a good deal of heterogeneity in survey data and particularly in these kinds of indicators. Thus the R^2s are not large. While an F-statistic for all covariates is a statistically more meaningful measure of goodness of fit (and it is in all cases large and significant), the fact that only 7 per cent of the variation in BMI is explained by the regressors suggests that polynomials in BMI included in the wage function may not perform well, although non-linearities are clearly of interest in view of Figure 1. We may

TABLE 2 *Impact of health characteristics on ℓn (wages in market sector instrumental variable estimates*

Covariates	(1) Height only	(2) Add BMI	(3) Add calories	(4) Add protein	(5) All health
ℓn(height)	2.43	2.41	2.83	1.44	3.92
	(14.2)	(13.9)	(6.4)	(5.0)	(4.0)
ℓn(body mass index)		2.22			4.74
		(2.1)			(2.1)
ℓn(per capita calories)			88.76		163.76
			(2.5)		(2.2)
squared			–5.86		–10.96
			(2.5)		(2.2)
ℓn(per capita protein)				27.54	–28.85
				(2.0)	(1.0)
squared				–2.05	2.30
				(1.9)	(1.0)
Education					
(1) literate	0.39	0.34	0.26	0.20	0.22
	(16.5)	(10.3)	(4.0)	(3.1)	(2.9)
(1) elementary	0.80	0.71	0.64	0.48	0.52
	(33.1)	(14.2)	(7.2)	(5.7)	(5.1)
(1) secondary +	1.79	1.64	1.61	1.37	1.34
	(65.8)	(21.4)	(13.8)	(14.4)	(10.1)
Tests					
χ^2(education)	5643.11	675.28	692.87	483.49	146.75
p value	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
χ^2(calories)			6.10		7.78
p value			(0.05)		(0.02)
χ^2(protein)				21.59	9.68
p value				(0.00)	(0.01)
χ^2(nutrients)			6.10	21.59	25.33
p value			(0.05)	(0.00)	(0.00)
χ^2(BMI & intakes)					27.02
p value					(0.00)
χ^2(all health)	202.47	196.64	140.57	183.71	108.54
p value	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

Notes: [*t* statistic] below coefficient estimate; [*p* value] below x^2 test statistic. Sample size = 11 555.

hope for more from the nutrient intakes for which 16 per cent of the variance can be explained.

Table 2 presents IV estimates of the effect of health and education on (log) wages of men in the urban labour market (excluding the self-employed). In an attempt to account for the fact that one out of four men are working outside the market sector, the wage functions include the hazard of participating in that sector.[8]

The first column of Table 2 reports the effects of education and height, both of which capture human capital investments prior to adulthood. They are powerful predictors of productivity in the market sector (and this is also true among the self-employed). To put the estimates in perspective, an illiterate man would have to be 38 centimetres taller than a literate man (*ceteris paribus*), if they are paid the same wage. Many have debated the interpretation of education effects in wage functions: in part, they probably pick up the role of unobservables associated with family background. Some of these will be captured in height and, in fact, height and education are highly correlated; the inclusion of height in the wage function significantly reduces the observed impact of education, especially at the top of the education distribution.

The second column adds body mass, which is treated as endogenous. Conditional on stature, heavier men are paid more and the reward to both stature and mass is greater in the self-employed sector. This suggests that strength does enhance productivity. The fact that taller women are paid more but heavier women are not adds some credence to this interpretation.

It is one thing to worry that estimates may be biased because of unobserved heterogeneity; it is quite another to demonstrate its empirical importance. A Wu–Hausman test, which compares IV with OLS estimates, indicates that it is inappropriate to treat BMI intakes as exogenous (a t-test on the first-stage residual in the wage function is over ten). Comparing the OLS and IV coefficient estimates also demonstrates that failure to take account of the correlation between observed BMI and unobservables would result in seriously misleading conclusions: the OLS estimate of the impact of ln BMI is *half* the magnitude of the IV estimate.

Most empirical work in the health–productivity literature has focused on establishing relationships between the two: subtleties regarding the form of that relationship have seldom been addressed. However, there are good reasons to expect relationships to be non-linear and this has important implications for policy. In the medical literature that relates body mass to subsequent mortality, mortality risks are higher only among those at very low or very high levels of BMI and there is a suggestion in Figure 1 that similar patterns may be found in the BMI–productivity relation. While these shapes are observed in OLS regressions, our attempts to identify thresholds in IV estimates by including a quadratic and cubic in ln BMI have been fruitless: in view of the first-stage results in Table 1, this is not very surprising.

The biomedical literature also suggests that the relationship between nutrient intakes and productivity may be non-linear. There is considerable debate about the impact that moderate deficiencies in calories may have on a range of outcomes as, the argument goes, within some range, basal metabolic rates and efficiency of absorption may adjust to intake levels, in which case intakes and productivity would be unrelated. If this is true, calorie intakes are not likely to have much effect on productivity or activity levels unless body adaptation is incomplete either because calories have fallen below some critical threshold or because the body has not had time to adjust. While there is, in our view, little convincing evidence on the degree of successful body adaptation, it does make good sense as an empirical matter to investigate non-linearities to the extent the data allow. With respect to nutrient intakes, for example, it seems plausible that, among poorly nourished populations, additional energy intakes may be associated with greater energy output and higher productivity; but that gain will diminish as intakes rise and may even decline when intakes become very high.

The third and fourth columns of Table 2 show the impact of (quadratics) in (log) calorie and protein intakes. Both have a positive impact on wages at low intake levels but the effects dissipate with intakes so that when calorie intakes reach about 2000 per day (and protein intakes are greater than 85g per day), additional nutrients garner no gain in the labour market. Non-parametric estimates indicate that, above these levels, intakes and wages are uncorrelated (Thomas and Strauss, 1993). Using farm household data from Sierra Leone, Strauss (1986) reports a similar concave relation between labour efficiency and per capita calorie availability in the household.

The interpretation of these estimates is not unambiguous. For example, calories alone (in column 3) may be proxying for other health characteristics with which they are correlated, such as body mass. Similarly, protein and calories are likely to be highly correlated and it is of interest to determine whether they have independent effects on wages. However, few studies have simultaneously examined the effect of multiple health indicators on wages, in part because few surveys contain a wealth of information on health and fewer have sufficient heterogeneity to support the analysis of several health indicators simultaneously.

Using the ICRISAT Indian village-level survey data, Deolalikar (1988) finds that farm and wage productivity are affected by weight for height and not calories when the two are included in farm production and wage functions (also including fixed effects to control for unobserved heterogeneity). Haddad and Bouis (1991) also use fixed effects estimators with longitudinal data from Bukidnon, Philippines. They control for individual calorie intake, body mass and height and find a strong effect of height, but not of body mass or of calories. This is rather weak evidence in favour of a health effect, since the association with height may simply reflect past human capital investments when the worker was a child.

The wage function in the fifth column of Table 2 includes all four indicators of health. Taken together, all the health measures are significant ($\chi^2 = 109$). In addition, a Wu–Hausman test for the exogeneity of BMI and the nutrient indicators is rejected: $F_{5,11523} = 170$). Height remains a powerful determinant of wages: taller men (and women) earn more, even after controlling for education and other dimensions of health. Body mass is also rewarded in the labour market. Conditional on size, nutrient intakes have an independent effect on wages: apparently calorie and protein intakes are picking up more than just the effect of mass. There is a positive impact of additional energy intake only at very low intake levels (below 1800 calories). Controlling for body size and calorie intake, higher protein intake is indicative of a higher-quality diet and a better diet is rewarded in the labour market. Furthermore, additional protein has the greatest return at high levels of intakes. Evidence for self-employed males in urban Brazil indicates that, taken together, body mass and nutrient intakes do affect productivity and, furthermore, the effect of all four health measures are not significantly different in the market and self-employed sectors.

Health–productivity relationships are also likely to vary with the type of work performed, with, for example, strength being rewarded most in manual labour. To sidestep the fact that wages and occupational choice are jointly determined, the Brazilian sample is stratified by education level under the presumption that manual labourers typically are less well educated. Whereas, among urban men, body mass is associated with higher productivity on average, this effect turns out to be significant only among those with little education and not among the better educated. Furthermore, body mass has no impact on the productivity of women, although market wages are higher among women with no education if they have a larger body mass, suggesting it may be used as a signal by their employers. Behrman and Deolalikar (1989) find body mass affects market wages only for men in India, whereas female wages in the sample are unrelated to both body mass and calorie intake. Both these sets of results are suggestive that, at least for men, body mass is a proxy for strength which, in turn, is a productive asset in the labour market.

It was argued above that, in the context of labour markets with short-term contracts, even if productivity is enhanced by good health, it will be rewarded by higher wages only if the employer can observe both productivity and health (or something that is correlated with them which can be used as a marker). Foster and Rosenzweig (1992, 1993) directly address this issue by comparing the impact of health on piece rates, which are presumably a good indicator of productivity, and time wages of daily workers in Philippine agriculture. They exploit the fact that the two payment schemes coexist for harvesting at the same time of year and, furthermore, a sub-set of workers engages in both types of contract. By examining the effects of health on differences in the implicit wages of the same worker, they are able to control for all individual unobserved heterogeneity and place the spotlight on the role of observability of both productivity and health.[9] Foster and Rosenzweig (1992) show that calorie intakes have a significantly larger impact on piece rates than on time wages; they argue that, since employers cannot directly observe nutrient intakes, they are not fully rewarded in time wage contracts. Moreover, conditional on intake, body mass has a significantly bigger effect on time wages than on piece rates (Foster and Rosenzweig, 1993), which is consistent with the interpretation that BMI serves as a signal to employers for nutrient intakes as well as having an independent effect.

Foster and Rosenzweig (1993) provide additional evidence along these lines by examining how energy intakes and energy expenditures differ according to the nature of the labour contract. Unless monitoring is costless, daily labour contracts have incentive problems which can be avoided by piece-rate schemes and linkages with other markets, such as sharecropping, and which obviously do not arise when working one's own land. Effort, as measured by energy expenditures, will therefore be less among those working for day wages. The authors estimate a production function for body mass, which depends on energy intakes, lagged body mass and other inputs and, by examining the effect on body mass of days spent working on different types of labour contracts, they infer energy expenditures.[10] Relative to working for daily wages, working on one's own land, or for piece rate wages, reduces body mass significantly more. There is no significant difference in worker effort between these two groups and sharecroppers (which is consistent with perfect monitoring of effort by landlords). To explore whether calorie intake varies directly with work expenditure, Foster and Rosenzweig express changes in energy intake as a function of changes in food prices, illness and in days worked on different contracts. They again find that working more days for piece rates or on one's own farm is associated with more calorie intake than is daily wage work.

Related evidence is provided by Behrman and Deolalikar (1989) who allow the effect of health on wages to vary with seasons, using the ICRISAT Indian data. Differences across seasons could result from different work being performed, from different contractual types, or from different resource constraints at different times in the year. They find that calories have a larger impact during the peak labour demand season, while weight for height has a larger impact in the off-peak season. The peak season includes harvesting activities, which are largely paid by the piece, whereas off-peak season activities are not.

Taken together, this evidence may be viewed as supporting the hypothesis that worker effort and the dimensions of health that are hard to observe will have a bigger impact on labour productivity when contracts are incentive-compatible. Furthermore, Foster and Rosenzweig (1992) also find that more productive workers tend to self-select into piece-rate jobs. However, the interpretation relies crucially on assumptions about the costs of monitoring worker effort and observing health indicators and the generality of the results has yet to be established. For example, contracts are not immutable and, if day labourers are not rewarded for good health, then one would expect to see the development of sharecropping arrangements or longer-term wage contracts. This does not seem to characterize the rural labour market in the Bukidnon area of the Philippines. Moreover, in urban Brazil, the impact of nutrient intakes is actually stronger in the market sector relative to self-employment, suggesting that either observability of effort or of health is not a key issue there (Thomas and Strauss, 1993).

The focus, thus far, has been on wages, productivity and effort. We turn next to farm (or firm) profits. Using data on farmers in rural Indonesia, Pitt and Rosenzweig (1986) report that self-reported illness of men and women does not affect profits from their farming activities. But there are at least three reasons why this may be observed, two having to do with measurement and the third being a more conceptual issue. First, there may be biases inherent in self-reported morbidity with higher-profit (and presumably higher-income) farmers being more inclined to report themselves as ill. Second, the duration of many illnesses may be too short to have any effect on labour outcomes. Third, as discussed above, if labour markets function well, so that farmers are able to hire in healthy workers to replace ill family members, then farm profits will be unaffected by the incidence of illness in the family. Of course, this does not imply that health has no impact on productivity, labour supply or income. In fact, household full income will decline by the value of the time lost to ill-health, which is the value of hired-in labour.

The interactions between health, labour supply and sectoral choice have received very little attention in the literature although, as the populations of many developing countries age, knowledge about these relationships is likely to become increasingly important. Using data on a cross-section of rural and urban men age 50 to 80 in the Second Malaysian Family Life Survey (MFLS2), Figure 2 describes the age profile of participation in the market and self-employed sectors, distinguishing those who report themselves as being in good, fair or poor health. While the sample is small (670 men), the patterns are suggestive that health and participation are related. Roughly a third of the sample men are working in the market sector and another third are self-

FIGURE 2 *Employment probabilities and health status: Males age 50–80 years in Malaysia (non-parametric estimates)*

Source: Second Malaysian Family Life Survey.

employed. Participation rates decline with age and the rate of decline is considerably faster in the market sector. Putting aside the interpretation of self-reported 'general health status', those in poor or fair health in the market sector are less likely to participate than those in good health at all ages and, especially, over 65 (to the right of the dashed vertical bar). Among the self-employed, however, it is men in good and fair health who are equally likely to be working and they are much more likely to participate than those in poor health. Furthermore, not only do about 15 per cent of all the men in poor health tend to be self-employed, but this fraction is *constant* across the age distribution. In addition, among men older than 65, self-employment participation rates are significantly higher than in the market sector for *all* health classes. Since the pictures say nothing about causality, or even about changes over an individual's life course, we cannot say why. It may be that men in poor health move into the self-employed sector, perhaps in the absence of social nets. Or it may be that the self-employed remain attached to the labour market longer. These kinds of questions are sure to provide fertile ground for many future analyses.

CONCLUSIONS

There has been substantial progress recently in the analysis of relationships between health and labour outcomes in poor countries. Bliss and Stern (1978) and the first volume of the *Handbook of Development Economics* (see, for example, Rosenzweig, 1988) concluded there was little reliable evidence that health had an important impact on labour productivity or labour use. Today, that assessment would have to be amended. There is now a body of evidence based on careful empirical studies that have shown clear relationships between dimensions of health, labour productivity and the allocation of labour.

However, teasing these relationships out of data has proved to be difficult. Health and income clearly affect each other and both are related to many factors that are hard to measure. Interpretation of associations between health and wages is not unambiguous and blind examination of correlations or regressions is unlikely to take us very far. But analyses of health and labour outcomes that are judicious in the choice of assumptions, thoughtful in the choice of estimation and sensitive to issues of robustness promise to be very profitable.

The number of convincing studies remains small, partly because the data demands are considerable. Much remains to be learned about the measurement of health, which dimensions of health matter, under what conditions they matter and for which groups in the population they matter. Virtually nothing is known about the dynamics underlying the relationships. There can be little doubt that, while progress has been made in recent years, research on the interaction between labour markets and health remains a key area that has only just begun to be explored.

NOTES

[1]The estimates are locally weighted smoothed scatter-plots. Each observation is replaced by its predicted value from a weighted regression using the observations in a band around it. The weights are one at the point itself and decline to zero at the boundaries of the band; we adopt a weighting function which is a tricube in the distance to the neighbouring observation. There are 10 675 men in the sample and, for this exploratory purpose, results are presented with 10 per cent of the sample in each band.

[2]If programme placement is not random, but related to the labour outcome of interest, then community infrastructure may be correlated with unobservables in the second stage and estimates will be biased (see Rosenzweig and Wolpin, 1986). This identification strategy also rules out selective migration of more productive people to areas with better infrastructure.

[3]Using earnings, rather than wages, complicates the analysis, since earnings incorporate both productivity and labour supply. Thus factors that affect labour supply, such as household assets or composition, belong in the earnings function and therefore are not potential identifying instruments. If wages depend on hours then these characteristics also belong in the wage function.

[4]Endogenous health variables do not appear in reduced form cost and profit functions and so it is conditional (or restricted) functions that must be estimated to examine the impact of health. We will refer to them, for simplicity, as cost and profit functions.

[5]For example, in their study, Pitt *et al.* (1990) use estimated health production function residuals from other periods as instruments for the estimated 'endowment' from the current period in order to correct for random measurement error. This imposes the strong assumption that the error is independent over time.

[6]During the survey each household was visited on a daily basis for seven days and food that was to be eaten over the following 24 hours was weighed, along with wastage from the previous 24 hours. An enumeration of all people present at all meals was also recorded and this information is taken into account in the calculation of per capita intakes.

[7]Price indices are calculated with household budget data reported in the survey; using information on several hundred foods, sub-aggregates were constructed for foods as well as non-foods; these were, in turn aggregated into an overall cost of living index measured for each state, distinguishing urban from rural areas. Hence food prices in the first-stage health functions may be interpreted as relative prices. In the second stage, nominal wages are deflated by the cost of living index and so can be interpreted as real wages.

[8]65 per cent of urban men work exclusively in the market sector, 12 per cent are also self-employed and 23 per cent work only in the self-employed sector. In the first stage, a multinomial logit is estimated and the hazard associated with working in the market sector is computed. The identifying instruments are quadratics in own non-labour income and also non-labour income of other household members. They are significant predictors of working in the market sector (χ^2 statistics are 98 and 14, respectively). It may be argued that assets and income from them reflect previous labour market choices and may thus be correlated with unobservables in the wage function; in this case, the instruments are not valid. Some studies have used parental characteristics, such as occupation and education, as instruments, although this imposes the restriction that all the effects of these characteristics on productivity must operate through the choice to work.

[9]One important difference is that time rates tend to be used in harvesting maize (mostly on small farms), while piece rates tend to be used more for sugarcane (on larger farms).

[10]To account for both simultaneity and unobserved heterogeneity, they use community dummies (as proxies for prices), education, land owned and household composition as the instruments for health inputs.

REFERENCES

Antle, J. and Pingali, P. (1994), 'Pesticides, productivity and farmer health: a Philippine case study', *American Journal of Agricultural Economics*, 76 (3).

Behrman, J. (1993), 'The economic rationale for investing in nutrition in developing countries', *World Development*, 21 (11).

Behrman, J. and Deolalikar, A. (1989), 'Agricultural wages in India: the role of health, nutrition

and seasonality', in D. Sahn (ed.), *Seasonal Variability in Third World Agriculture: The Consequences for Food Security*, Baltimore: Johns Hopkins University Press.
Benjamin, D. (1992), 'Household composition, labour markets and labour demand: testing for separation in agricultural household models', *Econometrica*, 60 (2).
Bliss, C. and Stern, N. (1978), 'Productivity, wages and nutrition: the theory, some observations', *Journal of Development Economics*, 5 (4).
Bouis, H. and Haddad, L. (1992), 'Are estimates of calorie–income elasticities too high? A recalibration of the plausible range', *Journal of Development Economics*, 39 (2).
Dasgupta, P. (1993), *An Inquiry into Well-being and Destitution*, Oxford: Oxford University Press.
Deolalikar, A. (1988), 'Nutrition and labour productivity in agriculture: estimates for rural south India', *Review of Economics and Statistics*, 70 (3).
Foster, A. and Rosenzweig, M. (1992), 'Information flows and discrimination in labour markets in rural areas in developing countries', *Proceedings of the World Bank Annual Conference on Development Economics Review.*
Foster, A. and Rosenzweig, M. (1993), 'A test for moral hazard in the labour market: effort, health and calorie consumption', mimeo, Department of Economics, University of Pennsylvania.
Haddad, L. and Bouis, H. (1991), 'The impact of nutritional status on agricultural productivity: wage evidence from the Philippines', *Oxford Bulletin of Economics and Statistics*, 53 (1).
Leibenstein, H. (1957), *Economic Backwardness and Economic Growth*, New York: John Wiley.
Pitt, M. and Rosenzweig, M. (1986), 'Agricultural prices, food consumption and the health and productivity of farmers', in I. Singh, L. Squire and J. Strauss (eds), *Agricultural Household Models: Extensions, Applications and Policy*, Baltimore: Johns Hopkins University Press.
Pitt, M., Rosenzweig M. and Hassan, M.N. (1990), 'Productivity, health, and inequality in the intrahousehold distribution of food in low-income countries', *American Economic Review*, 80 (5).
Rosenzweig, M. (1988), 'Labour markets in low income countries', in H. Chenery and T.N. Srinivasan (eds), *Handbook of Development Economics, Volume 1*, Amsterdam: North-Holland.
Rosenzweig, M. and Wolpin, K. (1986), 'Evaluating the effects of optimally distributed programs', *American Economic Review*, 76 (3).
Sahn, D. and Alderman, H. (1988), 'The effects of human capital on wages and the determinants of labour supply in a developing country', *Journal of Development Economics*, 29 (2).
Staiger, D. and Stock, J. (1993), 'Instrumental variables regression with weak instruments', mimeo, Kennedy School of Government, Harvard University.
Strauss, J. (1986), 'Does better nutrition raise farm productivity?', *Journal of Political Economy*, 94 (2).
Strauss, J. (1993), 'The impact of improved nutrition on labour productivity and human resource development: an economic perspective', in P. Pinstrup-Andersen (ed.), *The Political Economy of Food and Nutrition Policies*, Baltimore: Johns Hopkins University Press.
Strauss, J. and Thomas, D. (1993), 'Human Resources: Empirical Modeling of Household and Family Decisions', in T.N. Srinivasan, and J. Behrman (eds), *Handbook of Development Economics, Volume 3*, Amsterdam: North-Holland.
Thomas, D. and Strauss, J. (1993), 'Health, wealth and wages of men and women in urban Brazil', mimeo, RAND, Santa Monica, California.