
Health, Nutrition, and Economic Development

Author(s): John Strauss and Duncan Thomas

Source: *Journal of Economic Literature*, Jun., 1998, Vol. 36, No. 2 (Jun., 1998), pp. 766-817

Published by: American Economic Association

Stable URL: <https://www.jstor.org/stable/2565122>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



is collaborating with JSTOR to digitize, preserve and extend access to *Journal of Economic Literature*

JSTOR

Health, Nutrition, and Economic Development

JOHN STRAUSS
and
DUNCAN THOMAS¹

1. Introduction

Over the past 20 years, investment in human resources has taken center stage in the study of developing economies. A voluminous set of wage function estimates provides the basis for calculating market returns to education for virtually every country in the world. Studies have also looked at the effects of schooling on nonmarket outcomes. Prominent among those outcomes is the health of children and adults. Since health, like schooling, is a form of human capital, one might expect it to also be related to labor market success. That link has received much less attention in the empirical literature, although in recent years there have been substantial advances in our understanding of the complex interrelationships between health, nutrition, and economic devel-

opment. This paper reviews some of the evidence.

There are many reasons why the relationship between health and labor market outcomes in developing economies should be of special interest. First, there is a long tradition of theoretical models of nutrition-based efficiency wages in the development literature. Harvey Leibenstein (1957) hypothesized that, relative to poorly nourished workers, those who consume more calories are more productive, and that at very low levels of intake, better nutrition is associated with increasingly higher productivity. Such nonconcavities, which lie at the heart of efficiency wage models, have powerful implications for the level and composition of employment. Employers have an incentive to raise wages above the minimum supply price of labor, and to exclude those in poorest health from the labor market because they are too costly to hire. While many variants of the model have appeared in the theoretical literature, the models have been subjected to little direct empirical scrutiny. We describe the empirical conditions that need to be satisfied if these models are true, and assess the case for the exist-

¹ Strauss: Michigan State University; Thomas: RAND and UCLA. Acknowledgments: Financial support from NICHD Grant P01 HD28372, NIA Grant P01 AG08291 and the College of Business, Michigan State University is gratefully acknowledged. The comments of John Pencavel, four anonymous referees, Harold Alderman, Partha Dasgupta, Robert Fogel, Elizabeth Frankenberg, Mark Gersovitz, Kathleen McGarry, Eileen Miech, David Neumark and James P. Smith have been very helpful.

tence of health-based efficiency wages. We conclude the case is weak, at least in terms of evidence from modern developing economies.

If, however, we were to focus only on the specific nonlinearities that are implied by efficiency wage models, we would miss a very rich set of insights that emerge from studying health and labor outcomes. There are good reasons to expect improved health to result in improved functionality and productivity, which has important implications for an array of behaviors and investments. And yet this most basic question has received serious attention only recently.

This is all the more surprising in a developing country context, since the marginal productivity of health is likely to be higher in those settings relative to higher-income industrialized economies. Not only are levels of health far lower in developing countries, but the incidence and nature of diseases tend to be different. Notably, there is a much higher prevalence of malnutrition and infectious diseases (many of which are preventable) which often interact, causing even more serious damage (World Bank 1993; Dean Jamison et al. 1993; Per Pinstrup-Andersen et al. 1993).² These differences are important because they result in an age distribution of ill health that is tilted toward infants and preschool children, as manifest by these groups comprising a much higher fraction of total deaths than in higher-income settings (World Bank 1993). An important implication for our purposes is that adults in poorer economies are more likely to be afflicted with health

problems, some of which stem from early childhood, and the functional consequences of ill health are likely to be felt throughout the life cycle (and not primarily at older ages as in advanced societies). In addition, the structure of employment in lower income economies is such that work often relies more heavily on strength and endurance and, therefore, on good health.

Knowledge of the nature and extent of links between health and labor market outcomes is also important for policy. The health sector accounts for a sizeable fraction of the public purse in most countries. If public investment in health infrastructure and interventions yields benefits in terms of higher productivity and economic growth, then those benefits belong in evaluations of health programs.

The possibility that the income-generating capacity of the poorest is enhanced more by some health sector investments relative to others raises issues revolving around the distributional effects of policies. The literature indicates that some health problems (such as malnutrition) can become debilitating at extremes, and the biomedical evidence highlights the potential importance of thresholds, below which poor health can have dire consequences for functionality. This suggests that in a developing country context, the labor market consequences of poor health are likely to be more serious for the poor, who are more likely to suffer from severe health problems and to be working in jobs for which strength (and therefore good health) has a payoff.

Empirical studies of the links between health and labor outcomes in the context of developing countries have provided more than documentation on the nature and extent of the relationships. They have also proven to be fertile ground for addressing a broad set of

² For example, Alan Lopez (1993) estimates that 45 percent of all deaths in developing economies in 1985 can be attributed to infectious and parasitic diseases such as diarrhea and malaria; these diseases account for about 4.5 percent of deaths in industrial market economies.

questions that economists have been grappling with for many years. There is, for example, evidence that credit market constraints are binding for low-income households in several settings, particularly in very poor, rural areas. These constraints may result in health investments being below efficient levels. Several recent studies have exploited this insight to shed new light on the operation of such constraints. Others have used observations on health to compare worker effort under different labor contracts and to draw inferences about incentives underlying alternative contractual forms. Interactions between health and labor have also provided opportunities to peek inside the household and examine the allocation of health-related inputs to different household members, in light of differential activity levels and returns to health among members. Insights from this literature are highlighted below.

Section 2 presents correlations between health and labor outcomes. We show that, in the United States, men who are taller and heavier (given height) earn higher wages. We also show that the magnitudes of the correlations in the United States are dwarfed by the magnitudes in low-income settings like Brazil. This is followed by a discussion of both economic theory and econometric issues involved in drawing causal interpretations from the evidence in Section 2. In contrast with schooling, health varies over the life course and is the outcome of behavioral choices both during childhood and in later life. If better health is associated with improved functionality, and therefore productivity gains, then individuals, families and even society will invest more income in health than would be implied by its value in purely utility terms. As productivity increases, so will income, and that additional income may

be invested in health. This would generate potentially important feedbacks between health and income. Disentangling the direction of causality will clearly be key if we are to interpret behaviors. We discuss these issues in detail in Section 3.

Although there are many commonalities between health and education, Section 4 highlights two aspects of health that distinguish it from most other human capital measures. First, health is fundamentally multidimensional, and it is important to differentiate among these dimensions. Second, measurement of health is difficult, and in many cases, measurement error is likely to be correlated with outcomes of interest like income. We draw out the implications for interpreting empirical evidence on the links between health and labor market outcomes.

Section 5 reviews the current state of the empirical evidence regarding the relationship between health and labor outcomes, taking into account both experimental and nonexperimental studies. In addition to setting out the facts regarding a causal impact of health on labor outcomes, we discuss studies that have addressed broader questions regarding the functioning of the economy, and we review the evidence on efficiency wages. We conclude by drawing some lessons and discussing gaps in the literature.

2. *Correlations between Health and Labor Outcomes*

Drawing on historical series on the stature of adults in the United States and Europe, Robert Fogel (1992, 1994) makes a compelling case for linking aggregate movements in adult height to long-run changes in standards of living, including income, mortality, and possibly morbidity. Having demonstrated, for example, that the average American




Figure 1. Adult Stature by Birth Cohort

(or European) today is a good deal taller than his counterpart of two centuries ago, Fogel argues that stature is a useful index of the well-being of a population, supplementing data on wages, incomes, and economic activity. (See the survey by Richard Steckel 1995.)

Following this argument, Figure 1 presents evidence on the relationship between stature of adults and date of birth for the United States and three contemporary low-income populations: Brazil (in Panel B), the Côte d'Ivoire (Panel C) and Viet Nam (Panel D). Each panel covers a span of between 45 and 55 years.³ Note that the scales for

men (on the left) and women (on the right) are directly comparable within each panel, but not across panels.

Among these people, American males are the tallest. The average American

³ All the figures are based on samples of prime-age adults (ages 23 through 60) and present LOWESS estimates (Cleveland 1979) of the relationship between height (measured by an anthropometrist) and exact birth date (or birth year in the Côte d'Ivoire). The United States data are drawn from the National Health and Nutrition Ex-

amination Survey (NHANES) II (1976–80); we restrict attention to only the native born (who are taller than immigrants). There are 9,323 men and women in the sample. The *Estudo Nacional da Despesa Familiar* (ENDEF) was conducted in Brazil in 1974–75, and we have 23,107 adult respondents in the sample. The Viet Nam Living Standards Survey (VLSS), which was collected between October 1992 and October 1993, covers 4,800 households randomly drawn from the entire country; the sample included in this figure contains 10,235 adult respondents. The Côte d'Ivoire Living Standards Study (CILSS) (1986) contains information on 3,744 native born adults. In the CILSS, most adult birth dates, or ages, are reported in complete years, and there is a good deal of stacking on even digits. These age errors, along with the relatively small sample size, give us less confidence in the estimated shapes for this survey, relative to the other three.

male born in 1950 is nearly 177 cm tall; his Ivorian counterpart is about 170 cm, a Brazilian male is slightly shorter and a Vietnamese male born that year is almost 15 cm shorter than the American. It would be naive to attribute all of these differences to differences in the standards of living of these people when they were children, since height also reflects genotype influences. However, the differences and the changes over time within each country do provide a wealth of information about health and development over the long term.

First, it is immediately clear from the figure that there have been substantial increases in attained height by maturity in all four populations during this century. Average growth rates lie between .75 cm and 1.5 cm per decade. For example, an Ivorian man born in 1950 would be, on average, about 169 cm tall, which is almost 3 cm taller than his father would have been had he been born in 1930.

Second, underlying these average growth rates, there is considerable heterogeneity over time. In the United States, growth was very rapid for men and women until around the 1935 birth cohort; growth has substantially tapered off among later cohorts. A roughly similar pattern emerges for Brazil, with growth in heights of men and women tracking each other closely and tapering off since World War II. The Côte d'Ivoire provides a stark contrast. There has been no tapering off of growth among post-War cohorts, and changes in female height have not tracked the male profile very closely. The stature of adult males increased substantially and significantly faster than that of females until around 1945, when the gender gap reached its peak; since then, growth rates in stature have been roughly equal for men and women.

The Vietnamese case is striking, and

highlights a key advantage of using height data as an indicator of social and economic development over long periods of time in low-income populations where data are scarce. Between 1925 and 1955, heights of adult men and women increased very rapidly (by between 1.6 and 1.8 cm each decade). But for birth cohorts since the early 1950s, adult stature has remained unchanged. The 25-year period 1950–75, and especially from 1965 to 1975, was obviously a period of tremendous upheaval and dislocation in Viet Nam. Charles Hirschman, Samuel Preston, and Vu Manh Loi (1995) estimate that approximately 1 million (North and South) Vietnamese were killed during the intense war period 1965–75. Panel D of Figure 1 indicates that during the entire 25 years of upheaval, the health of children suffered. The figure also suggests that those born toward the end of the period suffered the most (since their parents were taller than earlier birth cohorts, and child height is positively correlated with parental height).⁴

While the secular increases in heights are sizable in all four countries, the gaps in height between them are even larger. However, comparison of means hides the very large dispersion within each country. In order to isolate the cross-section variation, we focus on the 1950 birth cohort in each country. Among men, the interquartile range of

⁴ Our interpretation of the evidence will be wrong if taller men were more likely to die during the war. There are three reasons why we do not think this form of mortality selection fully explains the evidence. First, the same pattern is found for Vietnamese women, who were much less likely to be soldiers. Second, mortality selection would have to be increasing over time (to explain the fact that the effect on height becomes cumulatively larger during the period). Third, wars typically ravage those at the bottom of the socioeconomic distribution, and we will present evidence below that this is a fair characterization of the Viet Nam War, particularly in the North.

TABLE 1
ANNUAL RATES OF GROWTH OF ADULT STATURE: VIETNAMESE MALES

Birth Cohort:	1925-55		1956-70	
	North	South	North	South
10 th Percentile	0.251 [*] (0.03)	0.164 (0.03)	-0.086 (0.05)	-0.041 (0.05)
Mean	0.189 [*] (0.02)	0.150 (0.02)	-0.008 (0.03)	-0.060 (0.04)
90 th Percentile	0.134 (0.03)	0.129 (0.03)	0.039 [*] (0.05)	-0.075 (0.05)

Notes: Coefficients from piecewise-linear regressions of height (in cm) on exact birth date (measured in years) for least squares (Mean) and quantile regressions (at 10th and 90th percentiles). Standard errors in parentheses. Quantile regression standard errors calculated using bootstrap. * Denotes significant difference between North and South at 5 percent level.

height is 9.2 cm in the United States, slightly higher in Brazil and lower in the Côte d'Ivoire and Viet Nam. The gap in heights of women is slightly smaller both across countries and within countries; for example, the interquartile range of female height of women born in 1950 in the U.S. is 6.9 cm. This within-country heterogeneity does not solely reflect genetic variation, but is related to socioeconomic conditions (John Strauss and Duncan Thomas 1995b).

The *variation in changes* in adult stature within a country yields insights into how the benefits of growth have been distributed within the population. Viet Nam serves as an example. Table 1 reports annual changes in stature among men born in the North and South of the country; the periods before and after 1955 are also distinguished. Among the 1925-55 cohort, adult stature increased 1.9 cm each decade in the poorer North, which is about 25 percent faster than in the South. Moreover, the shortest (and the poorest) benefited most during this period, especially in the North, where average height increased by 2.5 cm each decade—over

50 percent faster than among the shortest in the South, and nearly twice as fast as among those in the top decile of height. This evidence suggests there was improvement in the standard of living of all Vietnamese and a decline in inequality in terms of height (and, perhaps, economic conditions). In the post-1955 era, however, growth in attained height stopped for the average Vietnamese male in both the North and South. But in the North, growth appears to have declined among the shortest while increasing slightly among those at the top of the height distribution, resulting in an increase in height inequality during the period. This evidence suggests that not only did growth falter, but in the North the worst off were hit hardest by the disruption from the war period.

Turning to data at the household level, there is a good deal of direct evidence indicating that income and health are correlated. Not only does the demand for health services typically rise with income, but there is abundant evidence that the poorest are typically in the worst health. (See Jere Behrman and Anil Deolalikar 1988, and Strauss

and Thomas 1995a for discussions.) For example, child height tends to rise with income in many low-income societies, and well-nourished children are about the same height throughout the world (except in South and East Asia where well-nourished children are shorter but height is rising rapidly across generations).

The central question for this study is whether health and productivity are correlated at the individual level. Panel A1 of Figure 2 presents nonparametric estimates of the bivariate relationship between height and (log) hourly wages of men in the United States and urban Brazil.⁵ There is a powerful association between height and wages in Brazil. Taller men earn more: a 1 percent increase in height is associated with an almost 8 percent increase in wages. While this dwarfs the magnitude of the correlation in the United States, even taller American men earn higher wages.⁶ This ranking probably reflects differences in the extent of poor nutrition in the two countries as well as differences in the nature of work that is commonplace in each society, since manual labor—and thus reliance on physical strength—is far more important in Brazil.⁷

⁵ Data on U.S. men are drawn from the 1992 wave of National Longitudinal Survey of Youth for Panel A of Figures 2 and 3. The figures are based on wages in 1992 of 1,955 white males aged 27 to 35 years. Data from NHANES II are used in Panel B of Figure 2. (The NHANES data do not record wages but they do report height and education of all cohorts.) The Brazilian data are from ENDEF.

⁶ In an OLS $\ln(\text{wage})$ regression, the coefficient on $\ln(\text{height})$ is 7.7 (standard error is 0.2) in Brazil; the coefficient is 1.0 (standard error is 0.4) in the United States.

⁷ There are other reasons that could explain the difference. The U.S. sample is restricted to relatively young men (aged 26 to 35) whereas the Brazilian sample includes men ages 25 through 60, and so cohort differences may be confounding the correlation between height and wages in Brazil. The direction is not obvious. On one hand, as demonstrated above, adult stature has risen over time in Brazil, on the other hand, wages tend to

Height, which is but one indicator of investments in human capital made during childhood, is also related to educational attainment. Panel B of Figure 2 presents the correlation between adult stature and years of schooling for two cohorts in the United States (B1) and Brazil (B2). Taller men tend to be better educated in both countries, although the correlation is substantially larger in Brazil. Specifically, among 25 to 34 year olds, a 10-cm gap in height is associated with an additional year of schooling in the United States and 1.5 years in Brazil; because levels of schooling are very different, this translates into an 8 percent increase in the U.S. and a 25 percent increase in Brazil.

Panel B of Figure 2 raises a legitimate question: Is height simply proxying for education in the wage function in Panel A1? The answer is provided in Panel A2, in which relationships between height and wages are presented separately for two groups of Brazilian men: those with no education and those who have some schooling. While a good part of the observed positive association between height and wages can indeed be attributed to the role of education, even for those with no education the correlation persists and is large in magnitude. For them, a 1 percent increase in height is associated with a 4 percent increase in wages. After controlling for education, the elasticity remains large and, it turns out, is largest (4.6) for men who completed secondary schooling.⁸

increase over the life course. Empirically, it is the latter that dominates. Controlling for age (by including a dummy for each birth year), the correlation between height and wages is slightly larger. In fact, restricting the sample to men aged 30 to 35, the elasticity rises from 7.7 to 8.3 in Brazil.

⁸ A similar pattern is evident in the NLSY. Controlling for educational achievement, the effect of log height on log wages of white males is reduced by about one-third (to 0.62 with a standard error

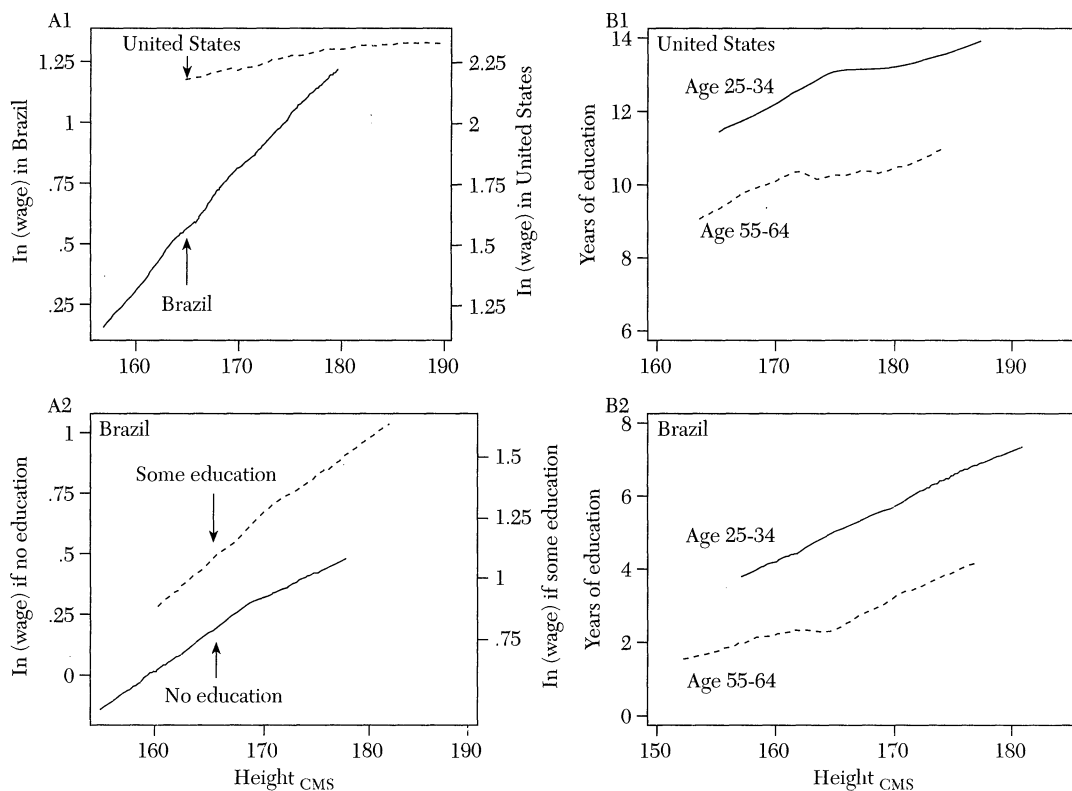


Figure 2. Wages, Education, and Height of Males in Brazil and the United States

Height may be a proxy not just for the quantity of schooling but also for its *quality*, another dimension of human capital investments. While this concern does not apply to those who have no schooling, there may be other human capital investments, such as parental time allocation, which are correlated with height and underlie the positive correlation between height and wages. It is important, therefore, when interpreting these correlations, to recognize that they may reflect both rewards to human capital investments early in life and strength or robustness as an adult.

In contrast with height, body mass in-

dex (BMI)⁹ varies over the life course and thus may capture both longer- and shorter-run dimensions of nutritional status and health. Clearly, BMI is related to energy intake, net of output; it has also been shown to be related to maximum oxygen uptake during physical work (VO_2 max), which is, in turn, related to aerobic capacity and endurance, independent of energy intake (G.B. Spurr 1983, 1988; Reynaldo Martorell and Guillermo Arroyave 1988). Whether this is an important pathway through which health may influence productivity is not obvious, since many jobs do not require sustained physical effort. Moreover, energy can be stored

of 0.34) and the correlation between height and wages is largest for those white men who have some college education.

⁹ Weight (in kilograms) divided by height (in meters) squared.

in the body and expended when needed, which implies that BMI may be affected by contemporaneous movements in incomes or prices, so that the dynamics linking them, BMI and productivity, may be very complicated. Furthermore, current BMI partly reflects previous health and human capital investments, and so a correlation between BMI and productivity may be capturing the influence of those prior investments.

With these caveats in mind, Figure 3A presents the relationship between wages and BMI for United States and Brazilian males. In the United States, the function is an inverted U with optimal BMI (in terms of maximizing wages) being around 24. (This translates into a weight of 177 lbs for a man who is 6 ft. tall.) The magnitude of the differences in wages across the BMI distribution in the United States is dwarfed by the magnitude in Brazil, where the shape is also quite different. Among men whose BMI is less than 27, wages rise dramatically as BMI increases, particularly for those above 22. Wages are essentially unrelated to BMI for the 13 percent of men whose BMI is above 27.

It is instructive to examine the wage–BMI curve within education groups (Fig. 3B). In contrast with height, only a small part of the correlation with BMI appears to be capturing the influence of education, further suggesting that nutritional status and wages are related over and above other human capital investments. The wage–BMI curve flattens at low BMI only for the poorest (those with no education), suggesting that poor nutrition (or poor health) takes its heaviest toll on the most vulnerable. The positive correlation between wages and BMI persists for those with no education even at high levels of BMI, but is zero when BMI exceeds 25 among the better educated. It is plausible that

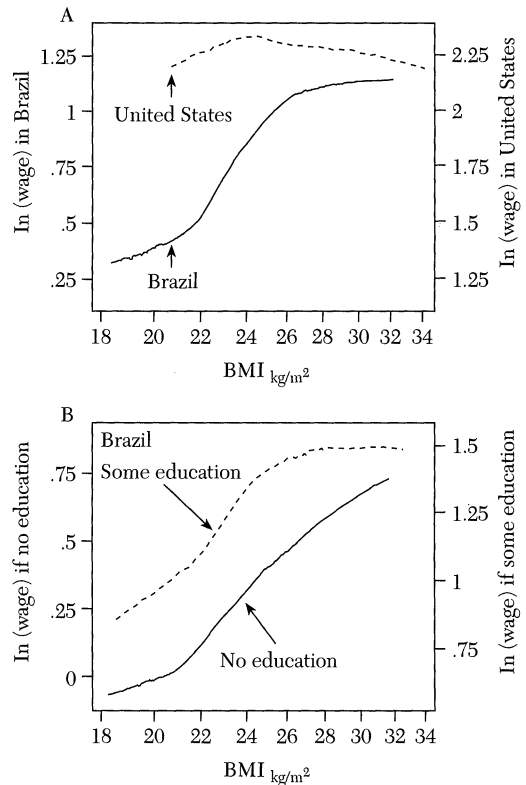


Figure 3. Wages and BMI in Brazil and the United States

among men with no education, elevated BMI is associated with greater physical strength, which is of value for manual labor, but that strength is of less value among the better educated (who are more likely to have sedentary occupations).

Figures 2 and 3 include only working males. But labor force participation choices may also be related to health. The relationship between the fraction of urban men in Brazil who are not working and their nutritional status is reported in Figure 4. It is clear that shorter men not only earn less, they are also less likely to be working. Over 10 percent of men who are 154 cm tall were not working at the date of the survey, but among those who are about 167 cm tall, the fraction is only 5 percent. For men taller than 167 cm, there is a




Figure 4. Height, BMI, and Percent Not Working (Males in Urban Brazil)

slight decrease in the probability of not working as height increases, but the slope is not significantly different from zero.¹⁰ A similar pattern emerges for BMI: the probability that a man is not working decreases with BMI until around 24, at which point it is essentially flat.

Other dimensions of nutritional status are also correlated with labor outcomes in the Brazilian data. Wages rise as calorie intakes increase (until around 2400 calories per day) and as protein intakes increase. In fact, wages rise as diet quality increases (as meas-

¹⁰ Looking only at men below 167 cm in height, the slope of the participation–height function is steepest among men with no schooling; a centimeter increase in height is associated with a .5 percent decline in the probability of not working.

ured by the fraction of calories from protein sources). The preponderance of evidence suggests that health and labor outcomes are, indeed, correlated among men in Brazil; similar conclusions emerge for women (Thomas and Strauss 1997).

There are, however, very good reasons to be cautious about giving these correlations a causal interpretation. Thus, before discussing what has been learned about the impact of health on productivity in the empirical literature, we outline a simple model in the next section to guide our assessment of that literature and discuss issues that arise in interpreting the evidence.

3. Economic Framework

It is useful to begin with a simple household production model in order to highlight the main issues in the literature (Gary Becker 1965); several extensions are discussed below. Assume a single-person household that maximizes utility over one period. Let worker productivity be perfectly observed by firms and assume the labor market is competitive.

Our central focus is on the link between labor outcomes and health. It is important to distinguish health outcomes, such as height, body mass, disease incidence or physical functioning, from health inputs that might include nutrient intakes, exercise, smoking, and utilization of preventive or curative health care.¹¹ Following the seminal work of Michael Grossman (1972), let there be a generic health production function for an individual:

¹¹ The distinction is not always clean-cut since some outputs may also be intermediate inputs into other outputs, such as disease incidence for body mass and vice-versa. We abstract from these complications here. Note also that some health outputs may be jointly produced, while others may be produced solely.

$$H = H(N, L; A, B', D, \mu, e_h) \quad (1)$$

where H represents an array of measured health outcomes. They depend on a vector of health inputs, N , and labor supply, L , both of which are under the control of the individual. We assume that actual health is increasing in inputs, except for labor supply, which consumes energy and may in other ways be taxing on health. The technology underlying the health production function is likely to vary over the life course and, possibly, with other socio-demographic characteristics, A , such as gender; in addition to family background, B' , such as parental health, and environmental factors, including local public health infrastructure and treatment practices, as well as the disease environment, D .

From an econometric point of view, there are two sources of unobservables in the health production function. First, the inherent healthiness of the individual, μ , is generally not observed in socioeconomic data, but may be known, at least in part, to the individual by adulthood. Second, there is a part which is unknown to both the individual and the econometrician, e_h , which includes measurement error. For now, we will be agnostic regarding the sources of unobservables, and return to them in detail below.

Assume that an individual's real wage is equal to his marginal product, which is costlessly observed by the employer. A person's (log) real wage, w , varies with health outputs, H , as well as other traits including schooling human capital, S , and family background, B (which now includes, for example, parents' schooling):

$$w = w(H; A, S, B, I, \alpha, e_w). \quad (2)$$

Apart from the inclusion of health, this is a standard earnings function that is the bread and butter of labor economics. As discussed below in Section 4, health out-

puts (such as body mass) may affect wages through better physical or mental health and through strength and endurance. Local community infrastructure, I , such as electrification or road density, may be related to labor demand or to various work characteristics that are valued separately from wages. Wages will also be influenced by unobserved factors, α , such as ability or school quality. Random fluctuations in wages—and measurement error—are captured in e_w .¹²

Discussions of how to interpret estimates from wage or earnings function like (2) have long been central in the labor economics literature (see, for example, Zvi Griliches 1977, and Orley Ashenfelter and Alan Krueger 1994). The most common and well-known argument is that estimated effects of education are biased because of omitted variables, such as unobserved individual ability, which are correlated both with years of schooling and with wages. Precisely the same problem arises with measures of health, even time-invariant ones such as height. There may be (we would argue that there are likely to be) omitted factors in the wage regression that are correlated, both with wages and with health factors. For example, school quality is typically unobserved in survey data. If higher quality education is associated with both higher wages and better allocative decisions by the individual during adulthood (or by the parents during childhood) with regard to health inputs and behaviors, then estimated effects of health, as well as schooling,

¹² An alternative model assumes wages are paid by the task, again costlessly observed, with piece-rate wages being competitively determined. Wages per unit time can be thought of as a piece-rate wage, w^p , multiplied by a function, $\gamma(H, S, B, I, \alpha)$ that relates an hour of clock time to efficiency units. This efficiency units (production) function depends on the same factors as the Mincerian wage function.

will be upward biased. This point was made in the discussion of Figure 2 above.

3.1 *Feedbacks between Health and Income*

For the case of health, there are two additional substantive issues that are likely to be of considerable import: direct contemporaneous feedback, or simultaneity, discussed in this subsection, and measurement of health, which is discussed in the next section. Feedbacks between current health and income will arise, for example, if better health results in higher productivity, more hours at work, or a higher probability of working, and if, in turn, higher incomes are invested in health by spending more on health-augmenting inputs. The worker may either buy more inputs (purchase more calories, for example) or buy higher-quality inputs (such as better health care); there is evidence suggesting investments are made on both the quality and quantity margin, especially among the very poor. Investments in training provide a direct analogy in the education literature.

Closing the household production model by allowing for behavioral choices helps clarify the role that unobserved heterogeneity plays and suggests some empirical methods for dealing with it. Assume that an individual's (or a household's) welfare depends on labor supply, L , and consumption of purchased goods, C , (some of which, such as foods, may be health inputs). Utility may be conditioned on health outputs, H , observable characteristics such as schooling, S , and family background measures, B , as well as unobserved characteristics including tastes, ξ .

$$U = U(C, L, H; S, A, B, \xi) \quad (3)$$

Allocations are confined by budget and time constraints in addition to the health

production function (1). Suppose that the individual works for a wage, w , and that asset or nonlabor income is V . The budget constraint is

$$p_c C^* + p_n N = wL + V \quad (4)$$

where, for expositional simplicity, consumption, C , has been decomposed into two elements, a vector of health inputs, N , with prices p_n , and nonhealth consumption, C^* with prices p_c . In practice, the division between these groups is not always clear, and not all health inputs may be valued directly in the utility function.

Letting λ be the marginal utility of income, and assuming interior solutions, we see several points emerge from the first-order condition with respect to usage of the j th health input:¹³

$$\frac{\partial U}{\partial H} \frac{\partial H}{\partial N_j} = \lambda \left(p_n - L \left[\frac{\partial w}{\partial H} \frac{\partial H}{\partial N_j} \right] \right) \quad (5)$$

First, if health inputs raise wages through improving health outcomes, then the shadow price of health-augmenting inputs declines, inducing greater use of those inputs. Second, the degree of decline of the health input shadow prices is likely to be greater for those in worse health. Biomedical evidence suggests that, in some instances, links between health inputs and outputs are nonlinear. For example, it has been argued that higher calorie intakes have no beneficial effect above some threshold but are associated with improved health below the threshold. If calories affect productivity in a similar way, then the shadow price of calories will decline among those people whose intakes are below the threshold, enhancing their incentive to consume additional calories. This has important

¹³ An input, N_j , such as calorie intake may also have a direct effect on utility; this effect is ignored in (5).

distributional consequences if, as is generally thought, the poor are most likely to have low intakes.

The second point is apparent from the reduced form demand function for health input j :

$$N_j = N_j(p_n, p_c, A, S, B, V, D, I, \xi, \mu, \alpha) \quad (6)$$

which depends on health input prices, p_n ; consumption prices, p_c ; demographic characteristics, A ; schooling human capital, S ; family background, B ; nonlabor income, V ; the disease and health environment, D ; and underlying nonhealth determinants of wages, I . All of these are observable. Unobservables also affect demand for inputs. These include tastes, ξ , innate healthiness, μ , and unobservables such as ability, α , which affect wages.¹⁴

The presence of α in both the wage function (2) and the health input demand function (6) captures the simultaneity problem that is central to the difficulty in disentangling the causal effects of health on productivity. A comparison of the functions does suggest instrumental variables. Conditional on health outcomes, H , and levels of local infrastructure, I , productivity should not be affected by health input prices, p_n , or the disease environment, D . Both sets of characteristics are, therefore, potential instruments for health in wage functions (2).

Studies have used food prices as instruments for nutritional inputs, such as calorie intakes, and health outcomes, such as body mass (Strauss 1986; David Sahn and Harold Alderman 1988).¹⁵ In many poor countries, some health in-

puts, such as health care in the public sector, are heavily subsidized and can be purchased at either zero cost or for very low (monetary) prices. Thus, financial outlays may be a small part of the full cost of health care; the dominant cost is often travel time to service providers and waiting time once there (Jan Acton 1975). Potential instruments therefore include time costs and their proxies such as distance to providers or the availability of health services in the community. Along the same lines, the quality of health services may also be valid instruments for health. The instruments, D , may include the local disease environment and health infrastructure such as water quality or sanitation services. One might expect measures of underlying community disease incidence to be potential instruments. For example, in tropical areas, rainfall is often related to the incidence and severity of malarial diseases. However, rainfall (or the lack of it) may also affect labor demand, and thus wages, in which case it is *not* a valid identifying instrument. More generally, if instrumental variable (IV) estimates of the effect of health on wages are going to be consistent, it is key that local labor demand conditions and infrastructure, I , be included in the wage function. If they are not, and if they are correlated with health prices or health infrastructure, which seems plausible, then the instruments will be correlated with the unobservables, α , in the wage function and will be invalid.¹⁶

There are, however, costs associated with using instruments measured at the community level. First, they may be

¹⁴ There is generally no reason to expect that measurement error in wages, e_w , will be correlated with health input demand, thus it does not appear in (6).

¹⁵ Because the cost of living varies across space and time, it is important to use *real* wages (or control for cost of living in the wage function) so that it is *relative* health prices that provide identification in the model.

¹⁶ It is important to recognize that the appropriate instruments using prices or infrastructure are those *available* in the community and not those actually *paid or used*, since the latter may be chosen by the individual and hence are correlated with α .

only weakly correlated with individual health input utilization or health outcomes, which can result in biased inferences in small samples, the bias being in the same direction as OLS (Charles Nelson and Richard Startz 1990; Douglas Staiger and James Stock 1997).¹⁷ Second, several studies have argued that local infrastructure may be endogenous. This could occur for two reasons: individuals may choose their location partly based on public health infrastructure (Mark Rosenzweig and Kenneth Wolpin 1988) or the infrastructure itself may be placed selectively by public policy, perhaps in response to local health conditions (Rosenzweig and Wolpin 1986). Selective migration may not be a major issue in this context, because migration would have to be correlated with unobserved factors that are correlated with health in a location, such as the availability of a clinic, over and above other measures of infrastructure included in the wage function. This does not seem to be the most likely cause of migration in poor economies, compared to the pull of wage differentials, for example. Selective program placement, however, may be a more serious issue for health programs (Mark Pitt, Rosenzweig, and Donna Gibbons 1993; Elizabeth Frankenberg 1995). One strategy for addressing this concern may be to model the process underlying public investments in infrastructure, although to date there has been little empirical experience with this strategy (but see Timothy Besley and Anne Case 1995).

¹⁷ This problem may be alleviated somewhat if "price" elasticities vary with individual characteristics, in which case interactions between those characteristics and "price" proxies may yield improved predictions in the first stage regression. (See, for example, Paul Gertler and Jacques van der Gaag 1990, who argue that demand for health care price elasticities are greatest for the poorest.)

Comparing (2) with (6) indicates that the value of assets, such as land owned or nonlabor income, V , are also potential instruments for health. However, assets may not perform well empirically, since in many surveys a large fraction of households reports having none, and these households tend to be the poorest. Furthermore, as we shall see below, there are theoretical reasons to be wary of relying on this restriction for identification. First, assets and wealth may be systematically correlated with measurement error in health. Second, in a dynamic setting, the evolution of assets is a choice and becomes endogenous. To the extent that assets reflect saving from prior earnings, it is likely that they will be correlated with unobservables in the wage function.

3.2 Effects of Health on Labor Supply

Our primary focus thus far has been on the relationship between health and market wages; we turn now to examine the impact of health on labor supply, which includes the decision to work, sectoral choice, and hours of work. In low-income economies, analyzing time allocation choices between self-employment and market wage work and self-employment productivity may be of special interest, since self-employment is overwhelmingly dominant, particularly in rural areas.

For households without productive (nonlabor) assets such as land, labor force participation decisions are made in the usual way, by comparing the marginal rate of substitution (MRS) between goods consumption and leisure when no work is performed to the potential market wage and choosing to work if the wage is greater. Both the market wage and the marginal rate of substitution depend on health outcomes, provided that health affects pro-

ductivity and that health is also valued directly. Better health may raise the MRS between goods and leisure if health and leisure are complements. Thus, even if health raises market productivity, it is an empirical question whether better health raises labor force participation rates on the margin.

To see how health affects labor supply, it is useful to define a “structural” labor supply function that is conditional on health outcomes and wages:

$$L = L(H, p_c, w; S, A, B, I, \alpha, e_w, S, A, B, V, \xi) \quad (7)$$

Health, like schooling, will affect labor supply by influencing offered wages, with resulting substitution and income effects, and also by affecting the marginal rate of substitution between goods and leisure. This latter effect, which is independent of any health effects on wages, results from our assumption that health is also directly valued (appears in the utility function).

The reduced form market labor participation and labor supply equations that would be derived from substituting for w and H in (7) will depend on the same factors that appear in the health input demands (6). These include health input prices, p_n , and the local disease and public health environment, D . Since these factors will appear in the reduced form if there is a direct effect of health operating through preferences on the MRS between goods and leisure, or if health affects wages, which in turn affect labor supply, the reduced form cannot provide an unambiguous determination of whether health does influence productivity. It can only give us an assessment of the (total) effect of health prices and the health environment on labor supply.

In contrast, estimation of the “structural” equation (7) (and its analogue for participation) provides an opportunity to identify separately the role of prefer-

ences from the productivity-enhancing effects of health. Estimation of the equation, however, poses the same problems of endogeneity and measurement discussed with regard to market wage equations. The prices of health inputs that do not appear in the utility function, health infrastructure, and the disease environment do not belong in the structural equation. They are potential identifying instruments and are subject to all the caveats raised above (Pitt and Rosenzweig 1986).

Thus far, the discussion of identification of health effects has focused exclusively on analyses conducted at the level of a single individual. In many contexts, studies examine data on households with multiple members, and identification is substantially more complicated. The essential nature of the problem is that, in general, prices of health inputs will not vary across family members, but levels of health outcomes or inputs will (Pitt and Rosenzweig 1990; Pitt 1995). We can see the issue clearly by considering labor supply of an individual conditional on health, H . It is not just the individual’s own health that will affect time allocation; the health status of all other household members will also, and instruments are then needed for each element of this vector H . A natural solution might be to use individual-specific health input prices or health infrastructure. But market prices of purchased inputs and infrastructure—such as the distance to facilities—do not typically vary among individuals within a household, and so it is easy to imagine quickly running out of instruments, particularly in large, complex households as is the case in many low-income settings. Including the health of only a subset of household members in the labor supply function will not solve the problem, since health prices will then belong in the function

(except under very strong separability assumptions).¹⁸

3.3 *Effects of Health on Sectoral Time Allocation and on Self-Employment Input Use, Outputs, Costs and Returns*

In order to model time allocation between self-employment and market labor, and the effect of health on self-employment productivity, it is useful to extend the discussion and include home production of a good that is both consumed and sold by the individual. With the addition of self-employment, a distinction needs to be drawn between market and total (including both market and self-employment) labor supply. The individual will work exclusively in the market sector if the market wage is at least as great as the marginal rate of substitution (MRS) between goods and leisure and the value marginal product (VMP) of labor in self-employment, both evaluated at zero hours of work. Exclusive self-employment will be chosen if the value marginal product is at least as great as the MRS between goods and leisure, at zero hours of work, and greater than the market wage at the point at which labor's VMP equals the MRS. Finally, the individual may choose to work in both the market and self-employment sectors, in which case there will be equality of the market wage, MRS and VMP in equilibrium. The impact of improved health on time allocation will clearly depend on how health differentially affects productivity in the market and self-employment work, com-

pared to how it affects the rate of substitution between leisure and income.

Both total and *market* labor supply functions can be defined that are conditional on health and wages, similar to (7); they will also depend on self-employment input and output prices and technology factors. Identification of health in these functions is no different from the case where there is no self-employment.

Time allocated to self-employment should be viewed as being derived from demand for labor in self-employment. Thus, modeling health effects on labor allocated to self-employment is analogous to modeling health effects on self-employment output, costs, or profits. There are two alternative strategies: using primal or dual functions. One could estimate either production functions or input demand, cost or profit functions.¹⁹ Production functions have the obvious advantage of being able to estimate the production effect directly, although not allocative efficiency effects, if any. As will become clear, production effects may exist without there being any consequent effects on input demands, costs, or profits.

We assume, first, that hired and family labor are imperfect substitutes, perhaps because of greater moral hazard for hired labor. If the individual spends time working in both the market and self-employment sectors, and if health does not affect labor efficiency, the household's optimization problem is recursive (or separable) in production and consumption (Inderjit Singh, Lyn Squire, and Strauss 1986).²⁰ That is, the household acts as though its deci-

¹⁸ Incorporating multiple household members in a model of health and wages does not complicate estimation if the wage of one member, conditional on own health, is not affected by the health of another household member. This rules out externalities (such as the ill health of one member being transmitted to another) and assumes wages do not vary with hours (if the ill health of one member alters time allocation of others).

¹⁹ See, for example, Strauss (1986), Pitt and Rosenzweig (1986), Deolalikar (1988), John Antle and Prabhu Pingali (1994), and Behrman, Andrew Foster, and Rosenzweig (1997).

²⁰ It is also necessary for competitive markets to exist for other commodities that are both produced and consumed (*e.g.*, foods) and that no corner solutions are chosen for them.

sions can be divided into two steps: first, maximize self-employment profits subject to the available technology, and second, maximize utility subject to the budget constraint augmented by the value of profits.

In this case, reduced form solutions for self-employment family labor (as well as other input demands, output supplies, cost and profits) depend only on prices of self-employment outputs and inputs, on quasi-fixed production factors (including human capital), and on any relevant community infrastructure (Dwayne Benjamin 1992). Taking the profit function as an example:

$$\pi = \pi(p_f, p_m, w_h, w_f, I, F, S, \theta, e_\pi) \quad (8)$$

where p_f represents price of the output (say food); p_m , the price of nonlabor inputs; w_h , the wage of hired labor; w_f , the wage of family labor; I , local infrastructure; F , fixed factors, if any; θ , unobserved factors affecting the production of the self-employment output; and e_π , measurement error in profits. Note that these functions do *not* depend on health input prices (except insofar as they are also prices of self-employment inputs or outputs, such as food prices), nor can functions conditional on health be derived.

Even if health does affect labor productivity, production and consumption decisions may still be recursive, and dual functions equivalent to (8) will be valid, provided wages are paid per efficiency unit. This would be the case if productivity were observable and wages were paid on a piece-rate basis. (Efficiency units of labor may, of course, depend on health as well as other human capital factors.) Intuitively, with a market for effective labor, if a family member falls ill, time can be reduced from market work. As long as market work is not reduced to zero, the opportunity cost of family labor remains the market piece-rate wage, which is unchanged by

the illness. Hence health has no impact on inputs, costs, or profits (Pitt and Rosenzweig 1986). What *does* change, if a family member falls ill, is the household's full income, which declines by the value of family labor time lost.

However, if it is optimal to allocate no time to market work, the model is not recursive and health input prices *will* affect reduced form self-employment profit, cost, and input demand functions. In this case, there is a shadow wage for family labor that depends on all prices and infrastructure, as well as on general taste parameters, ξ , in the utility function (3). Assuming that the shadow family labor wage is not observed, or cannot be solved analytically, then reduced form self-employment profit, cost, and input demands will be a function of all the determinants of the shadow family wage, including health input prices and the disease environment. Profit functions that are conditional on health and wages are similar in structure to those for labor supply (7), and do not depend on health input prices, or on the disease environment. Thus, these latter variables are potential identifying instruments. As is the case for labor supply, discussed above, if health enters the utility function, the effects of health prices in reduced forms, or health outcomes in conditional functions, do not allow anything to be necessarily inferred about the effects of health on productivity.

Even in this case, however, the possibility that profits and other dual functions are influenced by health is dependent on model assumptions. As pointed out by Pitt and Rosenzweig (1986), if the model is changed so that effective family and hired labor are perfect substitutes in the production technology, then even if no family labor is sold on the wage labor market, health will not affect self-employment inputs, outputs,

and profits. This follows since the opportunity cost of effective family time is the market piece-rate wage, which is invariant to any individual's health.

3.4 *Context of a Household*

As noted above, if we move from thinking about individual behavior to household choices, the analytics are more complicated, particularly in the case of labor supply. Because of its importance, we return to the general issue of treating households as collections of individuals. Although this raises a rich set of questions regarding the allocation of consumption across different household members or member types, it comes at a price since it is also necessary to grapple with the question of how allocations are made within households.

The nutrition-wage literature suggests a good, simple example. Suppose there is a minimum health threshold below which a person cannot find work (but can stay alive). Consider an extremely poor household whose income is so low that it would always be insufficient to allow everyone to consume at this threshold level, even when all consumption is devoted to health inputs. Even if everyone in the household has the same endowment, it will be optimal to have unequal allocations among members so that at least one person can earn an income.²¹

A large literature has now emerged regarding differential (and possibly unequal, in some sense) allocations of health outcomes across different types of individuals: men and women, adults and children (see Behrman 1992 and 1997 for good summaries). Young children and pregnant and lactating women

are often thought to be at particular risk of deprivation in health-input allocations. Economists have tended to attribute different investments to differential economic returns (Rosenzweig and T. Paul Schultz 1982). For example, if households maximize income and pay no attention to consumption levels of individuals, then better-endowed members would be allocated more inputs if the returns to those inputs in terms of their wages were greater than returns to increased levels of health inputs on the wages of the less well-endowed in the household (see Pitt, Rosenzweig, and Md. Nazmul Hassan 1990). Under more general models of household welfare maximization, other considerations come into play, including the weights associated with the well-being of each individual and aversion to unequal outcomes. In that case, health investments will depend on health endowments (Behrman 1988).

Differences in health input allocations among household members may also reflect different preferences and degrees of influence over decision making between household members (Nancy Folbre 1984).²² This suggests a model that falls outside the domain of the usual assumption in the economics of the family that the household may be treated as a unit with a set of well-behaved preferences. A discussion of alternative models of the household that allows heterogeneity in preferences of members would take us far afield and well outside the scope of this paper. What is important here, however, is to note that theoretical work that explores applications to the household setting of cooperative and noncooperative game

²¹ The threshold is not critical; the argument carries through if returns to health are nonconcave so that they increase at very low levels of health and decrease at higher levels. See James Mirrlees (1975) and Joseph Stiglitz (1976).

²² Additional evidence in support of this interpretation is provided by, for example, the finding that income in the hands of a mother has a bigger impact on the health of her children than income in the hands of a father (Thomas 1990).

theory (Marjorie McElroy and Jean Horney 1981; Shelly Lundberg and Robert Pollak 1993) or that examines the implications of Pareto-efficiency in household allocations (Pierre-Andre Chiappori 1988, 1992) has not yet been adopted in a sophisticated way in this literature. There are several obvious implications of these models. For example, investments in health will depend not only on household resources but also on the distribution of those resources. Moreover, a better understanding of the determinants of power within the household has the potential of yielding useful insights into possible instruments for individual health status in labor outcome functions of individuals in multiple-member households.

3.5 *Dynamic Issues*

The discussion so far has been in terms of static models and cross-sectional data. Yet as Grossman (1972) pointed out, health has critical dynamic components, which provide additional important insights. After reviewing some of the theoretical issues, we discuss empirical difficulties and opportunities that arise in this context.

Health evolves over the life course. While some health indicators are fixed during adulthood (such as height), most change over time and thus contain both a stock and flow component. These flows, or changes in health, may reflect investments in health or they may be the result of unexpected shocks; thus, some changes will be anticipated whereas others will not.

In this context, health can be thought of as a durable good: investments now can reap benefits in the future. These investments in health-augmenting inputs can, therefore, be thought of as an alternative form of savings or as a consumption-smoothing device. Whether health does serve this role in low-in-

come settings is an interesting but relatively unexplored issue. One might expect health considerations to affect consumption-smoothing strategies in low-income economies because the levels of income and health risk are so high. Further, because consumption-smoothing mechanisms are often less adequate for the very poor, the potential consequences of income shortfalls on the poor can be grave (see Partha Dasgupta 1993, and Robert Townsend 1995, for reviews).

As an example of the potential mechanisms, consider a simple two-period, perfect foresight model, in which future productivity depends on consumption today. It can be shown (Mark Gersovitz 1983) that very-low-income persons may have little incentive to invest in financial savings, because the loss of future wages due to lower labor efficiency may more than offset any returns to savings. This generates financial saving rates that depend positively on income. One can build on this, allowing current consumption to affect survival probabilities, which will also provide incentives for the poor to consume more today, rather than save (Gersovitz 1983), and to distinguish consumption from health investments so that portfolio choice between investments in health and in financial assets can be considered (Gerhard Glomm and Michael Palumbo 1993).²³

²³ Glomm and Palumbo have a more complicated model in which they build a dynamic, but deterministic, health production function and allow for fluctuating but exogenous income. They argue that if income is fixed over the life course, it is optimal for poor persons to diversify wealth into health and to increase consumption early in life (relative to a simple life-cycle model) in order to do so. With income uncertainty added, it is still optimal to invest in health early in life, but financial savings become far more important as a method of smoothing consumption later in life, and thus reducing the probability of a disaster that results in death.

Incorporating health dynamics in the model has important empirical as well as theoretical implications. It is natural to start with a dynamic health production function in which current health status, H_t , depends on current and *past* health inputs, N , labor supplies, L , and the disease or public health environment, D (Grossman 1972):

$$H_t = H(N_t, N_{t-1}, \dots, N_0, L_t, L_{t-1}, \dots, L_0, D_t, D_{t-1}, \dots, D_0, A_t, B', \mu, \mu_t, \mu_{t-1}, \dots, \mu_0, e_{ht}). \quad (9)$$

We have assumed that individual background variables, B' , are time invariant; it is, in principle, straightforward to allow them to change over time, as we have for age, A_t . We have also assumed that the unobserved, individual-level health endowment comprises two components: one that is time invariant, μ , and is known to the individual, and another that varies over time, μ_t . Whereas past realizations of μ_t will also be known to the individual, future realizations can only be forecast with error. We assume that only concurrent measurement error, e_{ht} , affects the observed realization of current health.

If health contains both a stock and a flow component, then it is conceivable that labor productivity will depend on lagged as well as current health status in addition to schooling, S , and infrastructure, I :

$$w_t = w(H_t, H_{t-1}, \dots, H_0, S; A_t, B, I, \alpha, \alpha_t, e_{wt}). \quad (10)$$

As an example, productivity may be affected by the duration of an illness, in addition to its incidence and severity. Alternatively, levels of health and changes in health may interact in their influence on productivity; for example, the effect of a loss in body mass due to a crippling bout of diarrhea is likely to differ depending on the initial level of mass (health) of the individual. Once again, it is useful to decompose unobserved heterogeneity, α , into a component that is

time invariant and known, and a part that varies over time and is only known when wages are revealed.²⁴

Inspection of (10) immediately yields several important new questions about the impact of health and health changes on labor outcomes. For example, to what extent do current episodes of ill-health affect labor market outcomes in the future? Does it matter whether these episodes are expected or unexpected? How important is previous health, conditional on current health, in determining labor market outcomes over the life course? By allowing for separate effects of current and past health on wages in (10), it is possible to examine the impact of *changes* in health status, or of the joint effects of levels and changes in health, on labor market outcomes. To date, however, very little evidence has been amassed to answer these sorts of questions.

There are several possible avenues that might be followed to estimate the wage function (10). The decomposition of unobservables into time-varying and time-invariant components suggests the use of a fixed effects procedure (Deolalikar 1988; Lawrence Haddad and Howarth Bouis 1991). Use of an individual fixed effect will sweep out the individual unobservable time-invariant component, α , so long as it enters the model linearly. The fixed effects transformation will, however, also sweep out observable time-invariant health indicators (like height), whose impact may be of considerable interest. Although it is very convenient, the linearity assumption is both *ad hoc* and not necessarily innocuous. For

²⁴ The potential effects of worker health investments on worker or firm investments in on-the-job training or related human capital are ignored here because, to date, they have not been considered in the empirical literature. Incorporating them would add both richness and complexity to the analysis.

example, in a conditional self-employment profit function, the time-invariant component of α may represent managerial ability (Yair Mundlak 1963). If it does, and if the profit function is represented by a flexible form approximation such as translog, there is no theoretical reason why managerial ability, just because it is not observed in the data, should not interact with other quasi-fixed inputs or with input or output prices. Other observable quasi-fixed factors would be interacted in this way, which would violate the linearity assumption.

In order to consistently estimate the impact of (changes in) health on wages, differences in the time-varying errors ($\alpha_t - \alpha_{t-1}, e_{wt} - e_{wt-1}$) must be uncorrelated with differences in time-varying health characteristics. If the time-varying errors only contain random measurement error (e_{wt}), the assumption seems plausible. It is less appealing if the errors also comprise omitted characteristics, α_t , which affect wages since part of these wage shocks may be invested in health. This will arise, for example, if a worker experiences a surprisingly good year and spends some of the unanticipated income on health-augmenting inputs (perhaps by eating more nutritious foods).

A natural solution to this problem is to combine fixed effects with instrumental variables. In a dynamic optimization framework, in which households maximize expected utility, reduced form health input demands will be a function of all *lagged and expected future* health input prices (and possibly their interactions with various exogenous household-level variables). These prices are, therefore, additional instruments, provided they can be measured. As an example, Behrman, Foster, and Rosenzweig (1997) use lagged food prices to estimate health effects in a dy-

namic, health-conditioned profit function.

3.6 *Observability of Effort and Health*

Another important issue relevant to our ability to detect effects of health on market wages is the role of imperfect information by the employer. We have assumed that real wages equal marginal productivity. This may be a reasonable characterization for self-employed and piece-rate workers; it may also adequately characterize workers on time rates whose productivity is observed at little or no cost. However for many workers paid on a time basis, productivity may not be easily observed if monitoring is costly.

Assume, for the moment, that healthier workers are more productive, but their productivity is not observable. Health will then be reflected in the wages of these workers only if it (or some factor known to be correlated with it) is observable to the employer. If health is not rewarded in the labor market, because of costly observability and monitoring, say, then in the long run one would expect the development of alternative labor or land contractual forms (such as sharecropping or longer-term labor contracts) that minimize problems associated with observability of productivity. One would also expect healthier (and more productive) workers to selectively move into sectors that do reward health, such as self-employment or piece-rate work (Foster and Rosenzweig 1996).

Thus, a key factor determining whether better health is translated into higher market wages is the observability of both productivity and health. As we will explore below, there are many potential indicators of health and they may have different effects on productivity. These measures also differ in their degree of observability: for exam-

ple, it may be difficult for employers to monitor total nutrient intakes, even when they feed their workers (because workers may substitute food to other family members by eating less at home). So it is not obvious that higher nutrient intakes will be rewarded in higher wages in an environment with costly monitoring (Foster and Rosenzweig 1994). But outputs related to nutrient intakes, such as body mass, are readily observed and so may be rewarded, either because they are correlated with nutrient intakes, or because they have independent effects on wages, or both.

3.7 General Equilibrium Implications: Nutrition Efficiency Wages

Workers in poor health may be rationed out of the job market because they are too expensive to hire. This is the key insight underlying nutrition-based efficiency wage models.²⁵ In all of these models, a healthier worker is a more efficient worker, and this relationship is not concave, being characterized either by a region of increasing returns before decreasing returns set in, or by a minimum health threshold that is necessary in order to work at all. This hypothesized nonconcavity has critical implications for the theory. In particular, if the reservation (time) wage faced by firms is sufficiently low, it pays all firms to raise wages above that level, because the gain in profits from increased worker productivity outweighs the costs.

Assume that all workers are homogeneous, that hours of labor supply are fixed for those who get jobs, and that all consumption is spent on health inputs, so that health varies directly with the wage paid. In order to maximize profits,

firms will set the value marginal product of labor *in efficiency units* to the daily wage paid per efficiency unit of labor, the *efficiency wage*, where efficiency units per hour worked can be defined as $\gamma(w)$ and the efficiency wage as w/γ . Equivalently, firms will set daily wages so as to minimize the efficiency wage, subject to the daily wage being at least as great as the reservation wage of workers. As long as the reservation wage constraint is not binding, the equilibrium condition for minimizing efficiency wages will involve equating the average efficiency wage to the marginal wage cost of increasing labor by an efficiency unit. The latter is simply the reciprocal of the marginal efficiency unit gains of an increase in wages, γ_w . Hence in equilibrium we have the familiar condition: $1/\gamma_w = w/\gamma$. Since this interior solution assumes that the offered time wage is above the reservation wage of workers, some will be unemployed. Indeed, without transfers or nonlabor income of some form, those who are unemployed will starve. This is a striking, if implausible, implication of the theory. As mentioned above in the context of intrahousehold allocations, if there are multiple members in a household, consumption will differ among members, with some members having first claims on consumption and others possibly starving.²⁶

There are several additional problems with the theory. For example, if investments in health now reap productivity benefits in the future, then this externality can only be captured by writing long-term labor contracts. We would therefore expect these contracts to exist in very-low-income settings, but in fact in many rural areas of developing

²⁵ The major papers that exposit this theory include Leibenstein (1957), Dipak Mazumdar (1959), Mirrlees (1975), Stiglitz (1976), Christopher Bliss and Nicholas Stern (1978a), and Dasgupta and Debraj Ray (1986, 1987).

²⁶ Leibenstein did allow for a nonstarvation solution, specifically, social norms that would prevent starvation by inducing firms to give up profits by lowering their wages and increasing employment.

economies, short-term, even daily, contracts are the norm (Rosenzweig 1988).

As noted by T.N. Srinivasan (1994), Dasgupta and Ray's (1986, 1987) model offers a solution to the appropriability problem, while also deriving the general equilibrium properties more completely than other models. In their model, piece-rate (or efficiency) wages are paid, where the piece-rate wage is determined in a competitive equilibrium in which the amount of efficient units of labor demanded (derived from an aggregate production function as the marginal product of efficient labor) is equated to the number of efficient units supplied. People in the model are allowed to be heterogeneous in their asset holdings, but are otherwise identical. There are two sources of income: returns on assets (for landowners) and wages (for the employed). As in the other nutrition efficiency wage models, all income is consumed in the form of health inputs. Because of the heterogeneity in asset income, one can determine for each person (or persons having a unique level of assets) a distinct efficiency wage: that is, a wage rate per efficiency unit that minimizes the cost *per* efficiency unit, $w/\gamma(w + \rho Z)$, of hiring them; subject to their reservation wage constraint. Here ρ is the rate of return on assets, Z , equal in competitive equilibrium to the marginal product of aggregate capital.

If the market equilibrium piece-rate wage, w^p , is lower than any individual's cost-minimizing efficiency wage, $w^*/\gamma(w^* + \rho Z)$, that worker will not be hired, even if the time-wage equivalent is above the worker's reservation wage, since it is not profitable for any employer to do so. The model thus predicts there *may* be involuntary unemployment, and if there is, those who have lower levels of assets and are thus in poorer health are less likely to be

employed.²⁷ If healthier people also have higher reservation wages (because they have higher total incomes and value leisure more), then some may choose not to work in the labor market, and so there will be voluntary employment among them. In a competitive equilibrium, efficiency wages for the employed will be bid up to the equilibrium level, even though workers with more assets than the marginal worker could have been hired for less. Thus time wages will not be equal across workers, but will be higher for persons with greater assets.²⁸

One of the more interesting empirical possibilities derived from the model is that some persons who are alike in all respects (including being landless) may differ in their employment experiences. One of the possible equilibria entails that the minimum efficiency wage of the landless just equals the market equilibrium efficiency wage. If so, it is possible that the demand for efficient units of labor is less than the supply, requiring rationing of jobs. Thus not only does one's asset position matter for employability, but so does luck.

However, contrary to some views, one implication that is not necessarily implied by the Dasgupta-Ray model is fixity of the market equilibrium efficiency, or piece-rate, wage. In the case when the market efficiency wage is so high that no involuntary labor exists, this implication is not surprising. However, it also holds when the market piece-rate is below the efficiency wage of the landless. In this case, too, if demand for labor shifts out, the market piece-rate wage will rise, and the more healthy of the unhealthy unemployed

²⁷ Of course, if equilibrium piece-rate wages are sufficiently high, no involuntary unemployment need ensue.

²⁸ If the labor market equilibrium were monopsonistic, this implication would be reversed.

will now get jobs. So long as the new market piece-rate continues to fall below the efficiency wage of landless workers, the least healthy will remain unemployed. It is only in the rationing equilibrium described above, in which the market piece-rate equals the efficiency wage of the landless, that market piece-rates will not respond to small shifts in labor demand.

The Dasgupta–Ray representation, then, provides predictions regarding who will be employed and who will be destitute. Given the historical importance of the nutrition efficiency wage models in the economic development literature, it is of considerable interest to empirically test their implications.²⁹ In order to consider the empirical evidence, we must consider both how to measure health and the effect mismeasurement may have on interpretation of empirical work.

4. Measurement of Health

Whereas economists have largely focused on understanding the behaviors that underlie associations between health and labor market outcomes, physicians and epidemiologists have devoted considerable effort to measuring health. This section highlights issues revolving around measurement, both because it is important for interpreting empirical relationships with labor outcomes and because of its relative neglect by economists.³⁰ Specifically, we highlight two aspects of health that

²⁹ Dasgupta (1997) argues that the theory should not be judged solely by its sharp predictions, such as the potential starvation of the unemployed landless, but rather should be interpreted more loosely as implying a cycle of destitution that may be initiated by low assets and bad luck.

³⁰ One may well argue there is a need to discuss measurement of labor market success; since much has been written about it in the economics literature, this does not seem like an appropriate place to repeat the arguments.

set it apart from other indicators of human capital like education: first, it is multidimensional and second, measurement error in health is likely to be related to income and labor market outcomes.

First, there is a consensus in the literature that the number of years of schooling is a reasonably good indicator of educational attainment. No similar agreement exists for health, in part because health is fundamentally multidimensional (John Ware, Allyson Davies-Avery, and Robert Brook 1980).³¹ Moreover, different dimensions of health are likely to have different effects on one's productivity or labor supply, and these effects may well vary over the life course or wage distribution. We argue therefore that it makes good sense to examine the relationship between labor outcomes and *multiple* health indicators simultaneously, wherever possible.

Second, it is typically assumed that measurement error in schooling is random (Griliches 1977). In contrast, many health indicators are measured with error that is systematically related to demand for health and other behaviors which are in turn related to wages, labor supply, and other socioeconomic characteristics. This complicates interpretation of empirical relationships between health and labor outcomes and seriously compromises the value of standard fix-ups, such as instrumental variables. In addition, the extent and nature of errors are likely to vary from measure to measure. (See Anita Stewart and Ware 1992, for a comprehensive and thoughtful discussion.)

³¹ In fact, education is more than the highest grade completed, but may also be characterized by years in school (or grade repetition), test scores, quality of schooling, post-school training and experience. Several studies have included these measures when calculating returns to education (for example, Joseph Altonji 1995).

In this section, a series of different health indicators that have been used in the empirical literature are discussed. We begin with self-reported general health status, followed by self-reported morbidities, limitations to normal activities, and measures of physical functioning. We then turn to nutrition-based health measures and include both health outputs (anthropometrics) and inputs (nutrient intakes). For each group of indicators, we discuss the nature of the information they are likely to contain along with the kinds of measurement error that are likely to be important. The section ends with evidence on relationships among different health measures to highlight the importance of multidimensionality and a discussion of the implications of random and systematic measurement error for empirical work.

4.1 *General Health Status*

Comprehensive clinical evaluations of health status, for example as undertaken in the United States National Health and Nutrition Examination Surveys, are far too expensive to be included in a typical socioeconomic survey, except on a small scale or with very selected samples; a fortiori in lower income settings where health service infrastructure is relatively weak. Most household surveys have, therefore, relied on interviews with respondents who provide an assessment of their own health.

Within the class of self-evaluations, general health status (GHS) is probably the most widely used indicator in the empirical literature in the United States. Some advocates have argued that it is the best single health index available, citing, for example, the fact that GHS is correlated with subsequent morbidity and mortality (see, for example, Ware, Davies-Avery, and Cathy Donald 1978).

GHS, however, suffers from two key drawbacks. First, a respondent is typically asked to rate his health in one of four or five discrete categories ranging from excellent to poor health. Relying on such a small number of discrete categories cannot possibly do justice to the complexity and diversity of health status of individuals; this is one reason we argue for using multiple indicators simultaneously.

The second problem has to do with measurement. "Good" health may not mean the same thing to all people, and respondents are not provided with an established metric against which to compare their own health. Indeed, few surveys provide a clear definition of a reference health status, and there is seldom an explicit reference group. Because questions about GHS are typically vague, we have no idea whether the respondent is rating his health relative to the national average, to his neighbor, or to whom.

Moreover, self-evaluations reflect perceptions of health. While important, perceptions are likely related to values, background, beliefs, and information, all of which are systematically related to socioeconomic characteristics, including wages and income. For example, information about own-health status is almost surely correlated with the extent of concurrent and prior use of health care, because people who have used the health care system are likely to be better informed. Since most people assume they are in good health unless they have information to the contrary, it is plausible to suppose that, conditional on a level of health status, those who have little exposure to the health system are likely to report themselves as being in better health. If this is true, given that lower-income people are less likely to use health care, particularly in poor societies, measurement error in GHS will

be systematically related to income (and wages). There is evidence in the RAND Health Insurance Experiment (HIE) that suggests this is a legitimate concern, at least for those people who, in the baseline, were in poor health and in the bottom quintile of the income distribution. Among these people, those randomly assigned to receive free care used more health care and were, at the end of the experiment, in better health as measured by clinically evaluated outcomes (such as blood pressure) and risk of subsequent death. But, according to their own evaluation of their health, measured by GHS, it actually *worsened* (Joseph Newhouse et al. 1993). Similar evidence is reported from a health price experiment in Indonesia, where people used more care in those places where prices were lower; their health improved, but their self-reported GHS was worse (William Dow et al. 1997).

We conclude that GHS does provide information in that it predicts later health problems. But, in the context of models of labor outcomes, there are good reasons to be concerned that it is contaminated by measurement error which is correlated with socioeconomic characteristics, including income.

4.2 *Self-Reported Morbidity, Illness and "Normal" Activity*

Disease-oriented definitions of health status, favored by many clinicians and some epidemiologists,³² have the advantage of a foundation in medical practice. However, even in a clinical setting, it is difficult to diagnose all problems unambiguously, especially in the presence of multiple, interrelated problems. Furthermore, from a social science and public health perspective, it is often the

³² See the discussion, for instance, in Jamison et al. (1993) and World Bank (1993).

functional consequences of ill health that are of primary interest, and those consequences typically cut across diseases and are exacerbated by interactions among different diseases.

As a low-cost alternative to clinical evaluations of health status in a household survey, some studies have drawn data from health facility records. But in low-income countries, a significant proportion of the population does not use these facilities, and those that do tend to be a select group that is not necessarily those in the worst health. In fact, they tend to be higher-income people, and so it is very hard to infer anything about the relationship between health and well-being in the whole population without at least information on the mechanisms underlying the choice to visit a health care facility.³³

Several surveys have asked questions about illness or specific symptoms (such as fevers, diarrhea, respiratory problems) during a reference period. As with GHS, these self-evaluations are difficult to interpret if what is deemed an illness or a symptom is not the same thing for all respondents. In fact, in surveys from low-income countries, it is not unusual for the poorest to appear to be the most healthy by this metric! For example, in Ghana and the Côte d'Ivoire, the propensity for adults to report being ill in the last four weeks is positively associated with own education (Schultz and Aysit Tansel 1997) and with per capita household expenditures (Mead Over et al. 1992). As with GHS, these indicators are likely to be measured with error that is correlated with use of the health system (and thus

³³ For example, Gertler and van der Gaag (1990) report that in the Côte d'Ivoire only 40 percent of persons who said they were ill in the last month sought medical care at a modern facility, and they tended to be from higher-income households. A similar pattern is observed in many countries.

with income and the price of health services). In addition, other price incentives may influence self-reported morbidity. For instance, an individual may claim to suffer from an illness in order to become eligible for health-related benefits (Donald Parsons 1982; John Bound 1991).

Another commonly used variant of self-reported illness is to ask whether any days of "normal" activity were lost to ill health. Some argue that this measure is less likely to be contaminated by systematic respondent error. Apart from the fact that "normal" is not well-defined, people whose opportunity cost of time is high (the better educated, say) will have less incentive to miss activities. By this metric, they will appear to be in *better* health than people with a lower value of time (conditional on a particular "true" health status). This is in contrast with the argument above whereby the better educated report themselves as being in *worse* health because of better knowledge or greater exposure to health services. While the net impact is unclear, there is no reason to expect a closer correspondence between normal days lost with "true" health than with days ill or some other indicator. As it turns out, reported correlations between days lost and education have both been negative (Allan Hill and Masuma Mamdani 1989), and positive (Schultz and Tansel 1996). Moreover, people with acute health problems are likely to make lifestyle and occupation choices in response to these problems, making it very difficult to interpret "normal."

In general, for all of these measures and for GHS, it will be very difficult to separate true health status from measurement error. The key point for our purposes is that this measurement error is likely to be correlated with socioeconomic behaviors and outcomes, such as

labor force participation, productivity, and wages.

4.3 *Self-Reported Physical Functioning*

Several household surveys have collected information on difficulties with physical functioning that are considered normal activities for people in good health; these might include walking a specified distance, lifting a particular weight, bending, or climbing stairs. While the notion of "difficulty" is subjective, questions about specific activities of daily living (ADLs) are more precisely defined than "being ill" or "normal activities," and there is some evidence that, relative to GHS and morbidity, ADLs are less prone to the type of measurement error discussed above. (See the discussion in Strauss et al. 1993; Michael Schoenbaum 1995, and Dow et al. 1997, present corroborating evidence.) This is an important advantage of ADLs in a survey setting. However, they have one key drawback: the limitations in physical activities that ADLs typically capture are frequently due to physical health problems such as shortness of breath, joint problems, or back problems (Stewart et al. 1978). Few prime-age adults have difficulty with these activities, and so ADLs may not be as useful in studies of health and labor outcomes as they have proved to be in the gerontological literature.

4.4 *Nutrition-Based Indicators:*

Nutrient Intakes

Nutrient intakes play a central role in the efficiency wage literature, and so it is not surprising that nutritional status has been prominent in the associated empirical literature. It is presumably net energy intake that is related to productivity, but since energy expenditures are difficult to measure, studies have tended to focus on nutrient intakes.

Biomedical evidence indicates that calorie intake is associated with increases in maximum oxygen uptake (Spurr 1983, 1988), which suggests there may be a link between nutrition and productivity. Interpretation, however, is made somewhat ambiguous by different body maintenance needs, which vary by weight and height. In the absence of measures of energy output, this suggests examining the link between nutrient intakes and labor outcomes conditional on, say, height and BMI. This provides one example of the usefulness of examining health measures in combination.

Measurement of energy consumption is not nearly as straightforward as might be thought. Several different methods have been adopted and each involves its own strengths and weaknesses. Calorie *availability* is computed by converting food quantities (purchases and consumption from own production) into nutrient intakes, using standard food composition tables. This has the advantage of being relatively easily calculated using data commonly collected in many household expenditure and farm production surveys. But the method suffers from several potentially important sources of systematic bias.

First, it assumes no food is wasted: everything that is available is converted into nutrients. It is plausible that very low-income (or low-wage) households waste less than those that are better off, in which case nutrient intakes will tend to be upward biased and the bias will increase with income. Second, it is very difficult in consumption and production surveys to take into account all meals that are given to guests or employees and all meals that are received in-kind. For example, the members of a household that entertained many guests during the survey period will appear to have consumed more calories than they

actually did consume.³⁴ If the (net) receipt of in-kind food (including transfers and gifts) declines with income and the probability of having guests rises with income, then nutrient intakes will be biased and the bias will, also, be positively correlated with income and wages. As a third example, it is very difficult to measure nutrient intakes for meals eaten away from home; typically, it is assumed that those meals have the same calorie content as meals at home. Clearly, this need not be true. For example, low-income workers are often given food at work; if it is more nutritious than the food eaten at home, measurement error will be negatively correlated with income.

There is *prima facie* evidence that these sorts of leakages (from wastage, food given or received in kind, and meals away from home) are important and result in systematic errors in estimated nutrient intakes. In most studies that use estimates based on availability, a substantial fraction of the population appears to be consuming at unrealistically low or high levels. Srinivasan (1992), for example, reports that in the 1976 National Sample Survey of India, over 5 percent of rural households consume less than 1,500 calories per capita *per day* and almost 20 percent consume over 4,000 calories per capita daily!

An alternative method of collecting nutrient consumption is to obtain information on *intakes* rather than on *availability*. One method weighs ingredients prior to each meal, and wastage after it, and then converts measured consumption into nutrients. Although meals eaten away from home are not captured, this method is probably the most

³⁴ While uncommon, it is possible to draw up an entire list of all people at all meals for a week (in a consumption survey), but that is not a realistic option over a longer survey period (such as in a farm household survey).

accurate one used to date. But the method is prohibitively expensive to field in large-scale household surveys and has therefore been employed very infrequently.³⁵

Thus, the most common strategy used to collect nutrient intakes has been to ask respondents to recall ingredients that went into meals consumed, usually over the previous 24 hours. While this approach has the advantage of potentially excluding leakages (such as meals for guests, transfers of food, and wastage of food), it comes at a substantial cost in terms of survey time. Moreover, since there is considerable variation in eating habits, 24-hour recalls are likely to be very noisy; extending the recall period raises concerns about recall bias (which is thought to rise rapidly in this context), and multiple visits further raise the cost of collecting these data.

Systematic evaluation of the accuracy of these different methods of collecting nutrient data is scarce, and there are virtually no surveys that contain more than one measure. A survey in Bukidnon, Philippines, conducted by Bouis and Haddad (1992), is an important exception. It has two different measures: intakes from 24-hour recalls, and calorie availability from a recall of purchases for the month. The survey was conducted over four rounds, and the

household averages of reported daily intakes over the four rounds are very similar: 2,439 per adult equivalent for availability and 2,358 for 24-hour recall, although the availability measures have somewhat fatter tails. Key for our purposes is that the difference between availability and intakes for the same household is positively correlated with household characteristics, including income, expenditure, and education of the head. Among low-income households, availability tends to be less than intake, but higher-income households report higher availability than intakes. This is all consistent with the idea that leakages are missed in availability and tend to be systematically associated with income.

It is possible to get some idea of the relative importance of random noise by regressing intake on availability and vice-versa. If 24-hour recalls are measured without error, then the coefficient when it is the regressand should be 1: it is very close (0.97 with a standard error of 0.01). But, as suggested by the fatter-tailed distribution, availability seems to be measured with more noise: when it is the regressand, the slope is significantly smaller (0.86 with a standard error of 0.01). In these data, availability appears to be a worse indicator of calorie intakes from the point of view of both random and systematic errors.

It is apparent, therefore, that all nutrient intakes are likely to be subject to some random measurement error that is probably greatest for recall-based methods. In addition, availability-based measures suffer from systematic errors that are probably correlated with income, wages, and productivity. Since the nature of the errors and the correlations with labor outcomes are likely to differ depending on how intakes are measured, it seems safe to conjecture that interpretation of the impact of nutrient

³⁵The method was used in the Brazilian ENDEF survey. To account for the fact that there is a good deal of variation in daily intakes, each of 55,000 households was visited on a daily basis for a week. All food that would be prepared in the following 24 hours and wastage from the previous 24 hours was measured at each visit. The enumerator also listed at each visit all household members and guests present at every meal during the previous 24 hours. It is therefore possible to control for a good part of the leakages discussed above. Of interest is the fact that among the poorest households, there is evidence of a decline in intakes after the first day or two; the survey staff attributed this to an attempt by the respondents to impress the enumerators (or downplay their own poverty) (Marcello Vasconcellos, personal communication).

intakes on wages will not be independent of the way the data have been collected.

4.5 Nutrition-Based Indicators: *Anthropometrics*

Having discussed nutrient intakes, which are inputs in a health-production function, we turn next to outputs, specifically, anthropometrics such as height, weight, and BMI. It is reasonable to wonder how much information on health can be embodied in anthropometry. As discussed in Section 2, adult stature has been profitably used as an indicator of well-being in both the economic history and the development literatures. Moreover, child height has proven to be an informative indicator of the nutritional status of children and is viewed as a longer-run indicator of nutritional status (see, for example, John Waterlow et al. 1977; Frank Falkner and James Tanner 1986; and Waterlow 1988). While height is clearly determined by the time an individual reaches adulthood (apart from shrinking later in life), there is some debate in the literature about the extent to which adult stature is completely determined by the time the child has progressed beyond early childhood. Nevertheless, to the extent that height does affect labor outcomes, it will clearly reflect returns to human capital investments made during early childhood and, perhaps, returns to strength.

Whereas height is predetermined by adulthood, weight varies in the short run and so provides a more current indicator of nutritional status. Since a light person may also be small, and thus not underweight given height (and, conversely, a heavy, tall person may not be overweight), nutritionists have found it convenient to analyze weight given height. There are many potential ways of expressing this ratio; one that has

been commonly used for adults is body mass index (BMI), the ratio of weight (in kilograms) to height (in meters) squared. On average, a prime-age male in the United States has a BMI of about 25; BMIs are considerably lower in poor countries (ranging, on average, between 21 and 23 in the three countries in Figure 1).³⁶ Extremely low (below 18) and high (above 30) values have been associated with higher adult mortality (Hans Waaler 1984; Fogel 1994), but BMI and mortality are essentially unrelated between that range. The relationship between adult BMI and subsequent mortality, the so-called "Waaler curve," is U-shaped with a long, flat bottom between approximately 20 and 30. The causal mechanism underlying this association has not been established, and the precise definition of extreme values remains debated.

Because they are relatively inexpensive to collect, in many surveys, height, weight, and sometimes arm circumference have been measured in the field by an anthropometrist. While the measures may be subject to random error (which can be reduced by good field procedures), a key virtue of these health indicators is the absence of measurement error that is systematically correlated with respondent characteristics (such as income). (But see Strauss and Thomas 1996 for a discussion of the National Longitudinal Survey in the United States, which relies

³⁶ Moreover, the fraction of the population at either extreme varies dramatically across countries. In the United States, for example, only 1 percent of men have BMIs below 18, 6 percent below 20, and 13 percent above 30. In Viet Nam, where the BMI of the average male is below 20, 22 percent are below 18, 66 percent are below 20, and less than 1 percent have BMIs above 30. In all four countries, the distributions for women are more fat-tailed: 3 percent are below 18, and 18 percent are above 30 in the United States with the proportions being 24 percent and 1 percent, respectively, in Viet Nam.

on self-reported height and weight, measured with error that is related to income.)

4.6 Multidimensionality of Health Status

We have argued above that because of the multidimensionality of health, it is often useful to examine several indicators simultaneously. Data from the Indonesia Family Life Survey (IFLS) are used to illustrate this point: Figure 5 presents the relationship between height and the ability to carry a heavy load easily, a commonly used ADL or functional limitation. In addition, we have included only those people who report their GHS as being “fair” (accounting for almost three-quarters of the adult respondents). The figure demonstrates three points. First, there is heterogeneity in both height and functional limitations even within the group of people in fair health; relying only on GHS is throwing away information. Second, the positive slopes in the figure indicate that taller people have less difficulty carrying a heavy load (a fact that carries through if we do not control for GHS). Third, comparing the two panels, younger adults (aged 25 to 40) have less difficulty than older adults (40–55), and comparisons within panels indicate that men have less difficulty than women. The evidence is suggestive that height is capturing more than background (or investments in early childhood) but also reflects differences in strength and robustness. In any case, the fact that height predicts functional limitations (even after controlling for general health status, age, and gender) points to the potential usefulness of employing multiple indicators of health in addition to (ordinal) measures of general health status.

More generally, in some instances, it is not at all clear how to interpret health effects without additional con-

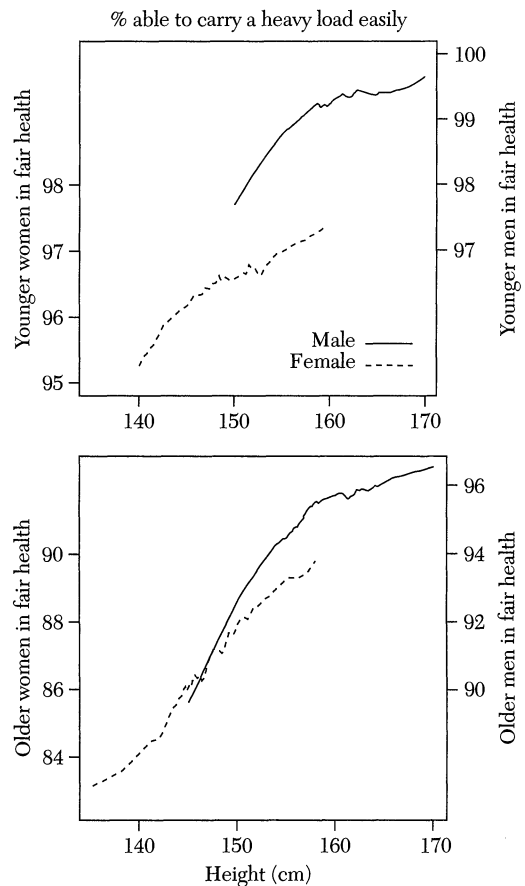


Figure 5. Anthropometrics and Physical Functioning, Conditional on General Health Status

trols, such as the effects of increased nutrient intakes without controlling for body mass (height and weight, say). In other cases, the inclusion of multiple indicators changes the interpretation of the effects. For example, holding calorie intakes constant, increasing protein intake implies an increase in the protein density of the diet and hence an increase in diet quality.

4.7 Empirical Implications of Measurement Error in Health Status

Throughout the discussion, we have taken care to distinguish random from systematic errors in measurement of

health. This is important because estimated effects of health on labor outcomes will vary with the type of error, as will strategies that seek to ameliorate the influence of measurement error. We therefore discuss each in turn, although it is important to recognize that a particular health indicator may be affected by both types of error, albeit to different degrees.

Random measurement error in health will bias estimated effects towards zero; this is a case of classical errors-in-variables. It is difficult to think of any health indicator that is not subject to at least some random measurement error. Even weight and height measured by a trained anthropometrist are not immune to error. For example, subjects may not be aligned properly on measuring boards, scales may not be recalibrated, and so forth. Moreover, many indicators vary over the course of a day (like weight) or across days (like nutrient intakes or functional problems). In all of these cases, repeated measures can ameliorate the impact of random error in regression models.

Systematic measurement error poses thornier problems. Consider the example discussed above in which poor health is more likely to be reported if the respondent has more contact with a modern health practitioner. If health care utilization rises with income, then higher-wage individuals are more likely to report themselves as ill, given a particular level of underlying health status: the impact of health on wages will be negatively biased.

In contrast with random errors, averaging multiple reports by the same respondent is not likely to help reduce the impact of systematic error. Moreover, it is not easy to think of instruments that are likely to be correlated with "true" health status, but uncorrelated with wages. For example, if the

reason people report themselves as being sick is because 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. Assuming that people's propensity to report themselves as ill changes slowly over time, repeated observations on the same individual may help, since examining differences in health status (transitions) will reduce the influence of systematic measurement error—but at the expense of increasing the relative magnitude of random error.

In general, measurement error in health status or inputs is likely to be partly random and, in many cases, partly systematic. Obviously, the best way to reduce measurement error is to pay more attention to measurement. In some cases this involves taking averages across repeated measures (to reduce random noise), examining differences (to reduce systematic error), taking special care in fieldwork to avoid systematic errors (by carefully measuring "leakages" when collecting nutrient intake data, for example), or by avoiding certain indicators that are especially prone to error. Recent innovations in survey methodology offer some potentially exciting opportunities in this regard. For example, a small number of household surveys have experimented with greater reliance on direct observation, such as assessing the incidence of anemia (based on hemoglobin counts), tuberculosis (using sputum), and hypertension (based on blood pressure). In addition, studies have sought to measure glucose levels (with saliva), net energy intake (with labeled water), and lung capacity (using peak flow meters), and also to directly observe functional limitations through timed moves (such as walking a particular distance, stand-

ing from a sitting position, and so on). Whether these measures will prove to be fruitful in studies of the links between health and productivity remains to be seen.

5. *The Relationship Between Health and Productivity*

Nutritionists, physicians, and a handful of economists deserve the credit for much of the early empirical work on health and productivity. Most of these studies were interested in documenting correlations between physical productivity and various health indicators, particularly those that are nutrition-oriented or that measure the incidence of specific diseases. By and large, these studies were oblivious to problems associated with time-invariant and time-varying unobserved heterogeneity. (See Ralph Andriano and Thomas Helminiak 1988; Behrman 1993; and Strauss 1993, for reviews.) Most of the nonexperimental studies that have attempted to sort out causality have also focused on documenting whether or not health affects productivity and remarkably little attention has been paid in these studies to potential biases that might arise from mismeasurement of health. A small number of experimental and quasi-experimental studies are important because they have sought to address concerns revolving around both causality and measurement.

Over the last few years, however, there have been substantial advances in the field, and the literature has progressed beyond looking for relationships. Attention is now often focused on more economically interesting, and subtle, questions. These include, for example, drawing distinctions between the effects of different health and nutrition indicators; comparing the effects of health depending on an individual's

type of activity and type of contractual arrangement in order to test hypotheses about the efficiency of markets, including labor markets; testing conjectures in the literature about the importance of thresholds in health and their links with poverty and deprivation; and drawing inferences about how resources are allocated within families.

To begin our review of the evidence, we return to the association between height and wages presented in Section 2. We next discuss the experimental and nonexperimental evidence on the nature of the effects of different health indicators on productivity and labor supply. The discussion distinguishes morbidities, limitations, and ADLs from nutrition-related health indicators. We do not stand at one end of the spectrum on the virtue of experimental relative to nonexperimental studies but see them as complements with each contributing value added to our understanding of the relationships.³⁷ On balance, our reading of the literature suggests that some dimensions of health status and some health inputs do affect labor supply and worker productivity. In several cases, the effects tend to be largest for the poorest. With this in mind, the section continues with a discussion of the shapes of the relationships and how these shapes vary by type of work, season, gender and also the observability of health and work effort. Many of the results provide insights into the functioning of markets and asymmetries of information in low-income settings.

³⁷ Our goal is to summarize the evidence. In virtually every nonexperimental study, one could reasonably quibble about specifications and identification. Since these issues were covered in detail in the previous section, we prefer to focus here on results and to discuss details only when new insights emerge. However, because it has not yet been discussed, experimental design is treated somewhat asymmetrically in this section and is described in more detail.

Finally, we review the rather scant evidence that relates to the potential general equilibrium implications of health-based efficiency wage models.

5.1 Height and Labor Outcomes

As Figure 2 indicates, taller men earn higher wages in Brazil and the United States; this is not a new result. It has been observed in a wide range of countries across the world (Haddad and Bouis 1991; Foster and Rosenzweig 1993b; Schultz 1996; Thomas and Strauss 1997; Peter Glick and Sahn 1997). Although it has been demonstrated less frequently, taller women also earn higher wages, and furthermore, both taller men and women are more likely to participate in the labor force. The link between height and wages has also appeared in the historical literature, where, for example, Robert Margo and Steckel (1982) demonstrate that sales values of slaves in Mississippi were influenced by height (and also weight). While among slaves, part of the value of height reflects the benefits of longevity (Gerald Friedman 1982; Meredith John 1988), today employers are more likely to pay for better health and robustness, and, possibly, strength.

More generally, height reflects investments made in the worker during childhood, and so we might interpret it as an indicator of human capital much along the lines of education. Indeed, as shown in Section 2, height and schooling are positively correlated. In Brazil, adding height to a wage function for urban, self-employed males results in the estimated returns to schooling falling by nearly 20 percent for all levels of schooling and, in market wage equations, cuts the magnitude of a significantly positive hazard term (controlling for selection into market work) by 60 percent, causing it to become insignifi-

cant (Thomas and Strauss 1997). Schooling also provides a useful benchmark against which to evaluate the magnitude of the correlation between height and wages. For example, in urban Brazil, for a literate man (who has not completed elementary school) and an illiterate man to earn the same wage, with their only difference being height, then the illiterate man would have to be 30 cm taller than the literate man; this is 16 standard deviations in the data!

5.2 Morbidities, ADLs, Health Limitations, GHS and Labor Outcomes

Nutrition-related health indicators are discussed in the next subsection; we focus on all other indicators in this subsection. In an experimental study that sought to measure the impact of a specific disease on output, sugarcane workers on an irrigated estate in Tanzania were divided into two groups at the time of a baseline: those diagnosed with schistosomiasis and those who were not. Schistosomiasis is caused by parasitic worms that live in slow-moving water, and repeated exposure can result in fevers, aches, and often fatigue. Earnings of workers on the estate were paid on the basis of the amount of sugarcane cut in a day. This indicated that, at the baseline, the productivity of those suffering from schistosomiasis was lower than those who were not infected. In the experiment, those infected were randomly divided into two groups. The treatment group members were given chemotherapy, and their earnings increased but did not fully make up the gap with the uninfected group. There was no change in the earnings of the control group, the infected workers who were not given chemotherapy. The study clearly suggests that schistosomiasis does affect productivity (A. Fenwick and B. Figenschou 1972).

In contrast, however, a similar experimental study of sugarcane workers in the Cameroons found no effect of the same illness on output. In that study, however, there was no difference in the baseline output of workers infected with schistosomiasis and those who were not infected. Moreover, treating the sick with chemotherapy had no impact on later productivity, and in the follow-up survey there was no difference in output of the treatment group, the control group, or the workers who were not sick at all (C. Gateff et al. 1971).³⁸

These studies highlight some of the strengths and weaknesses of experiments in this literature. By design, both studies address concerns regarding the endogeneity of health, since the evaluation is based on the impact on productivity of changes in health status that result from a randomized experiment. Moreover, because they are very focused, the studies are able to measure health status with great confidence: in these cases, they exploit the fact that it is straightforward to test for schistosomiasis with blood or urine samples. The downside is that experiments tend to be small, have a short time-frame, and involve relatively homogenous populations, and therefore may not have great power. It is difficult to know whether low power explains the absence of an observed productivity effect in the Cameroons.

In contrast with the experimental research, nonexperimental studies have paid little attention to issues of measurement, and those that have tend to focus on *random* measurement error in health. The fact that *systematic* measurement error has been largely ignored is troubling, since these studies use

survey data in which general health, morbidities, and ADLS are all self-reported; thus the measures carry with them the problems of interpretation discussed in the previous section.

The nonexperimental evidence on whether these health measures affect wages or labor productivity is, at best, also mixed. On one hand, a study of Ivorian males reported that those who missed a day of work in the four weeks prior to the survey earned 18 percent lower wages than those who were not absent. But in a companion analysis using comparable data from Ghana, there is no link between wages of men and days ill. Moreover, women's wages are unaffected by days ill in both countries (Schultz and Tansel 1997). A study of farmers in Indonesia demonstrates that farm profits are invariant to whether the head of the household or spouse report themselves as being ill during the previous week although, as we shall see, that does not necessarily imply the illness did not affect the individual's productivity (Pitt and Rosenzweig 1986).

In sharp contrast, these and other household surveys do indicate that poor health reduces hours of labor supply. In Pitt and Rosenzweig's Indonesia study, the magnitude of the estimated effect is large: those men who reported being ill the previous week worked 70 hours less than other men. The authors explain the magnitude by arguing that only severely sick respondents reported themselves as being ill. This seems reasonable, given that less than 3 percent of farmers report themselves as sick, although it does raise the question of what being ill means in the survey.

It is worth taking a short detour to discuss an important implication of these results, since they speak to a question broader than the link between health and labor outcomes. While poor health is associated with reduced labor

³⁸Other studies are reviewed in Nicholas Prescott (1979).

supply, the fact that farm profits are not affected does not imply the farmer's own productivity was unaffected but, instead, that labor was available that could substitute for him during this time of illness. Pitt and Rosenzweig conclude that since this labor was provided by hired wage earners, labor markets work efficiently in rural Indonesia, an argument that is corroborated by other evidence on efficiency of markets in Indonesia (Benjamin 1992). This is a good example of how studies in this literature have moved beyond documenting relationships, providing empirical insights into questions that economists have been grappling with for a long time.

Returning to health and labor supply, it would be remarkable if workers who reported missing a day of work during the reference period did not also report lower hours. Schultz and Tansel present evidence in support of this intuition for men and women in both Ghana and Côte d'Ivoire. It is, however, not entirely clear how to interpret the result since, as discussed in Section 4, one might interpret work days lost or days of limited activity as being time allocation choices, and thus intimately related to labor supply, rather than as being only reflections of health status.

Other studies have used measures of general health to explain labor force participation for prime-age adults in Jamaica (Victor Lavy, Palumbo and Steven Stern 1995) and for men and women over age 60 in Taiwan (Schoenbaum 1995). Both studies report that labor force participation is significantly lower among those in poorer health. Arguing that health is endogenous and measured with random error, both studies predict an ordinal index of general health status with education, income or assets, and with ADLs. If, as suggested may be the case in Section 4, measure-

ment error in self-reported health status is correlated with education or income, the estimated health effects will be biased. Moreover, it seems likely that an individual who is inclined to report himself as being in poor general health will also be inclined to report difficulties with some ADLs (conditional on "true" health). Thus, we would expect there to be common reporting error across these multiple health indicators for the same individual. In that case, the ADLs will be poor instruments for general health.³⁹

A social experiment in Indonesia, the Indonesian Resource Mobilization Study (IRMS), provides an opportunity to examine the importance of reporting error among these different indicators. In the experiment, user fees at public health centers were raised in randomly selected "treatment" districts, while prices were held constant (in real terms) in neighboring "control" districts. A baseline household survey was conducted prior to the intervention, and the same households were resurveyed two years later. The goal of IRMS was to assess the effect of the price increase on utilization: it declined in the treatment areas, relative to controls (Gertler and Jack Molyneaux 1996). Labor force participation also declined in the treatment areas, relative to controls, with effects being particularly large (and significant) for men and women at the bottom of the education distribution, those whom we would expect to be the most vulnerable. Dow et al. (1997) estimate a conditional labor supply function that relates *changes* in health outcomes to *changes* in labor force participation, thus differencing out any time-invariant unobserved het-

³⁹Schoenbaum also finds substantial negative effects on labor force participation of difficulties with ADLs when they are included directly, both in the form of an index and as separate covariates.

erogeneity. Changes in health status are instrumented with experimental treatment status, which is, by virtue of the design, random. In addition to these advantages, IRMS contains information on all the self-reported health measures discussed in this subsection. This turns out to be key, since the estimated impact of health differs substantially depending on the particular measure. Worsening of ADLs, such as greater difficulty in walking 5 kilometers, and those measures related to labor supply, such as more days in bed, significantly depress male labor supply. However, reduced labor supply was also recorded for those men who reported their general health status had *improved* during the hiatus between the surveys. Looking a little more closely at the data, ADLs and days of limited activity changed for the worse among individuals in treatment areas, relative to controls. But, counterintuitively, GHS and morbidities actually *improved* for treatments relative to controls. Since control households faced lower prices and used more health care during this period, a reasonable interpretation of the evidence is that self-reported health status is likely to be worse among those people who have greater exposure to the health system and, presumably, more information about themselves. If this is correct, then reporting error in GHS and morbidities will be systematically related to use of health care and, therefore, in most nonexperimental studies, to income and education as well as health-seeking behaviors. The evidence suggests that this problem is less pernicious in the case of ADLs, perhaps because they are more "objective."

In sum, poor health, as measured by the indicators discussed here, does appear to reduce labor supply. However, the evidence that it affects productivity and wages is more ambiguous. The ab-

sence of a consistent impact on wages in the nonexperimental literature may be because the health indicators used in those studies tend to reflect shorter-term health problems (morbidities or days of limited activities in the last few weeks) but wages tend to adjust relatively slowly. Hours of work, in contrast, can vary in the short run. We suspect that the differences in the estimated impacts on wages and labor supply also reflect differential impacts of systematic reporting or measurement error that is correlated with information and exposure to the health system. Thus, reporting error will be more highly correlated with wages and income than with hours of labor supply.

5.3 *Nutrient Intakes and Labor Market Outcomes*

The relationship between nutrition-based health measures and labor outcomes has received substantially more attention, particularly in the development literature. There are several experimental and quasi-experimental studies demonstrating that nutrition-based interventions improve health outcomes, particularly among children. Although it has been more difficult to make the connection with improvements in productivity, a small number of studies with experimental designs do provide evidence that low nutrient intakes have a deleterious impact on productivity.

Until fairly recently, it was thought that poor nutrition largely reflects inadequate energy, or protein, intakes. Recent work suggests that this focus has resulted in a failure to fully recognize the key role that may be played by several micronutrients. The central problem lies in disentangling the separate effects of different macro- and micronutrients, all of which are typically consumed in combination. Experimental

designs have a clear edge in this regard, as they have the luxury of implementing specific nutritional interventions.

A particularly compelling example is a longitudinal study of 302 male rubber tree tappers and weeders in Indonesia (Samir Basta et al. 1979). Baseline health measures indicated that about half of these men were anemic. Their baseline productivity (measured by kilograms of latex collected by tappers per day and the area of trenches dug by weeders) was about 20 percent lower than the productivity of nonanemic workers. In the experiment, workers were randomly assigned to one of two groups (irrespective of their anemia status). The treatments were given a special iron supplement for 60 days; the controls were given a placebo. Workers received an incentive payment to take the pills as scheduled.⁴⁰ At the end of the period, blood hemoglobin, aerobic capacity (measured by a step test), and output of those who were initially anemic and received the treatment increased to nearly the levels of the nonanemic workers (whose biological indices did not change). Among those in the control group who were anemic, productivity and blood hemoglobin levels also rose, although the increase was substantially smaller than among those in the treatment group. (This is attributed by Basta et al. to the effect of the incentive payment, a claim which is corroborated by a comparison of dietary intakes before and after the experiment: they point out that, during the experiment, anemics in the control group spent

⁴⁰ Nevertheless, attrition was substantial. This is a concern that plagues all experiments and longitudinal data collection more generally. In their comparisons of productivity, Basta et al. use only 77 of the 302 males enrolled in the experiment. The key question, to which we do not know the answer, is whether the attrition is random with respect to unobservables that also affect productivity.

more on leafy greens and other foods that provide higher amounts of iron.)

The importance of adequate iron in the diet is a robust result that has been reported in several experimental studies of both productivity of adults and cognitive achievement of children (see Ernesto Pollitt 1997 for a review).⁴¹ Results from less specific food supplementation interventions have been more varied for several reasons. Maarten Imminck and Fernando Viteri (1981a,b) examine the effects of calorie supplementation among sugarcane cutters in Guatemala, where assignment to groups was by village and not randomized within villages. Productivity rose in both the treatment villages and the control villages. Because of the design, it is difficult to know whether this is because calorie supplementation had no effect or because of changes in the villages that differed between treatments and controls; that is, we cannot separately identify an experimental effect from village-time effects.⁴² This pitfall was avoided in a study of 47 Kenyan road construction workers who were randomly assigned to calorie supplementation or placebo (June Wolgemuth et al. 1982); calorie supplementation

⁴¹ Behrman (1996) provides an excellent review of the more general issue of links between health and education of children. He concludes that "good health and nutrition have more nuanced and qualified effects on schooling success than is often recognized."

⁴² A much-cited intervention of pregnant and lactating women and their children in Guatemala suffers from the same design problem. Participants in treatment villages were given a high calorie drink (*Atole*) and controls were given a placebo. Perhaps the best evidence that the food supplement resulted in improved health outcomes and enhanced cognitive achievement is derived from a "difference-in-difference" type of estimates. Children who were in the poorest households (and thus were most malnourished) benefited most from *Atole*, substantially reducing the gap with their better-off peers in the treatment villages. No similar effect was observed in the control villages (Pollitt et al. 1993).

has a small positive impact on dirt dug per hour. In contrast, in the Indonesian iron-supplement study, productivity does not vary by calorie intake. There are several interpretations of these results. First, it may be that what appears to be a beneficial effect of additional calories in the Kenyan study is, in fact, the effect of iron or other micronutrients. Second, it is possible that energy supplements are more important for road diggers whose tasks are physically more strenuous than rubber tappers. Third, levels of protein-energy malnutrition may be higher among the Kenyan road diggers, in which case they will benefit more from additional calories, assuming the calorie-productivity relationship is not linear. Fourth, several of the food supplement studies that have focused on children suggest that the benefits of protein and energy interventions tend to emerge only after an extended intervention period. Finally, it is difficult to know how much of a calorie supplement actually translates into higher intakes among the treatments. It is reasonable to suppose that the treatments will consume less calories at other meals, resulting perhaps in an increase in calorie intakes by other family members. This suggests that individual-specific food supplementation interventions may be difficult to implement in an experimental setting and is one dimension in which nonexperimental household survey data has a clear advantage.

Nonexperimental studies of the impact that higher nutrient intake has on labor outcomes can also do better than studies of self-reported illness or health, because there are multiple ways of measuring intakes and something is known about measurement error in each case. This is important since the results differ substantially, depending on whether intakes are measured by availability (from food consumption

data), by individual 24-hour recalls (or averages of recalls) or by weighing food consumed. We discuss the evidence, starting with intakes of calories and protein, and then draw in the role of body size.

Several studies indicate that per capita household calorie availability has a powerful impact on farm output and on wages. For example, Strauss (1986) estimates that among households in Sierra Leone, the elasticity of farm output with respect to calorie consumption of all family workers is 0.33.⁴³ Sahn and Alderman (1988) report that while wages of women in rural Sri Lanka are not responsive to calories, the wages of men are: the elasticity is 0.2. Although the populations of both countries are extremely poor, and levels of health are very low in Sierra Leone, the estimated effects in both settings are very large.

Based on the discussions in the previous sections, there are three reasons why the estimates may be biased. First, if higher-income households tend to waste more food (give more food away or have more guests at their meals), calorie availability will be an error-ridden measure of intakes, and the error will be positively correlated with income (and therefore output or wages). The estimated calorie effects will be upward biased. Second, calories may be proxying for other health indicators with which they are correlated, such as body size. Since wages rise with stature (and perhaps mass), the estimated calorie effects in these studies will be upward biased. Third, calorie availability is measured at the household level, but wages are recorded for individuals; if calories are not shared

⁴³ This implies, for example, that a worker who consumes 1,500 calories daily (at the third decile in the distribution) would be only 60 percent as productive as a worker consuming 2,400 calories, the sample average.

equally among all household members, the calorie effects will be biased, although the direction is a priori ambiguous and depends on how resources are allocated within the household. We will discuss issues revolving around measurement error first.

Instead of calorie availability, some studies have used 24-hour food recalls measured at the individual level to estimate nutrient intakes. As noted in Section 4, 24-hour recalls may suffer less from bias caused by income-related waste and transfers, but may suffer more from random measurement error because of day-to-day variation in diets. For example, Deolalikar (1988) uses 24-hour food recall data to examine the effect of calorie intakes on market wage and farm production functions in rural south India with data collected by the International Crops Research Institute for the Semi-Arid Tropics (ICRISAT). The effect of calorie intakes on wages in rural Philippines is estimated by Haddad and Bouis (1991). Both studies exploit the longitudinal nature of their data and include an individual fixed effect to control for all time-invariant unobserved heterogeneity that affects both intakes and wages or farm production. In both studies, the effect of calories in that specification is small and not significant. However, calorie intakes measured in this way are known to be very noisy and the fixed-effect estimates may well be biased downward by measurement error. In fact, in the Indian study, the estimated effect of calories is positive and sizable, if imprecise, when the model replaces the fixed effect with a random effect.⁴⁴ Similarly, in an OLS

specification, Haddad and Bouis report that individual-specific calorie intakes have a large, positive impact on market wages. The evidence in both studies is, therefore, consistent with the interpretation that, after controlling for unobserved heterogeneity, there is no effect of calorie intakes on wages. But, as Haddad and Bouis conjecture is likely, the evidence is also consistent with measurement error in 24-hour recalls swamping the impact of intakes on wages.

Alok Bhargava (1997) also uses calorie intakes based on 24-hour food recalls, but he takes the average intake over seven consecutive days for the entire household. This is important because averaging is likely to reduce random measurement error. His result is strikingly different: calorie intakes are a significant determinant of the probability that a man in rural Rwanda engages in "strenuous" (and usually higher-wage) activities.

Calorie intakes based on weighing and measuring food prepared and food wasted over seven days are likely to also minimize random noise. In addition, these estimates of intakes should not be affected by the biases that may arise in availability measures. This class of estimates is therefore probably subject to the least measurement error of all. Using these sorts of measures, Thomas and Strauss (1997) report that wages of workers in urban Brazil are positively and significantly affected by calories at low intake levels.

Moreover, additional protein is also associated with higher wages until relatively high levels of intakes. If calorie intakes are held constant, then increases in protein intakes imply that the price (or quality) of calories increases. This is because a greater fraction of calories consumed will be rich in protein which would typically mean a diet

⁴⁴ The random effects estimates may also be biased because no correction is made for correlations between covariates (calories) and time-invariant unobserved factors (such as height). It is difficult to draw firm conclusions regarding the nature of the bias.

with more meat and other animal products. When calories and protein are both included in the wage function, the relationship between protein and wages switches from being concave to convex. This suggests that better-quality diets are not only associated with higher wages but that there are *increasing returns* to diet quality in the labor market in urban Brazil.

We think, therefore, that measurement error in nutrient intakes has important implications for understanding the links between intakes and labor outcomes. Whereas the magnitudes of the effects of calorie availability on wages may be upward biased (because of systematic measurement error), estimates based on intake recalls for the previous 24 hours seem to be downward biased (because of random error). Taking this into account, we conclude that the balance of evidence points to a positive effect of elevated nutrient intakes on wages, at least among those who are malnourished.

5.4 *Body Size and Labor Outcomes*

It is not obvious how to interpret the result that additional calories are associated with higher productivity if higher productivity workers are stronger and consume more calories because they have higher needs. In the regressions based on calorie intakes discussed above, height and either weight or BMI are included in the regressions as proxies for body needs, so that the interpretation of calorie consumption is closer to net energy intakes. In addition, the effects of body size on productivity are of interest in and of themselves.

This section began by noting that height has a positive effect on wages. Conditional on height, weight is also associated with higher wages or output in many, but not all, studies (such as Deolalikar 1988 and Thomas and Strauss

1997). Moreover, higher BMI appears to raise the propensity of a worker to engage in work that is more physically demanding (Bhargava 1997). Pitt, Rosenzweig, and Hassan (1990) go a step further and ask whether “weight-for-height endowment” is associated with labor outcomes. They report that men with higher endowments are more likely to engage in strenuous work in rural Bangladesh, and that the average endowment of all men in the household is positively related to income.⁴⁵

To summarize our conclusions thus far, height does seem to have an impact on market wages, although whether this represents health, strength or family background is not clear. Body mass seems to positively impact the productivity of men, at least those active in physically demanding jobs. The effect of nutrient intake is more uncertain. Micronutrient deficiency (particularly iron) has a deleterious impact on productivity, and there is some evidence pointing to a positive impact when calorie intake is well measured. All of these conclusions refer to linear effects of health on labor outcomes. There are important additional insights to be gleaned from examining nonlinearities in the effects of health, including interactions with the type of job, gender,

⁴⁵ “Endowments” are measured by residuals from a weight-for-height production function. The idea is that, after controlling for all inputs, the residual is a noisy measure of health endowment, $\mu + e_h$, in (1). For each individual, production functions are estimated for four different periods within a year, and so in any period, endowment measures in other periods are valid instruments. Under the assumption that measurement error, e_h , is not correlated across time, these IV estimates will be purged of measurement error. Of course, it is critical that all inputs be controlled in the production function, since otherwise the residual will reflect those as well as endowment. This is very demanding of data in the case of weight-for-height, since it depends on past as well as current health inputs; few surveys contain detailed information on past health inputs, and the survey used by Pitt, Rosenzweig, and Hassan is no exception.

time of year, and the nature of the labor contract. We turn to these issues next.

5.5 *Nonlinearities in Relationships Between Nutrition and Labor Outcomes*

Most empirical work in this literature has focused on establishing relationships between health and productivity. Subtleties regarding the form of these relationships, and in particular the role of nonlinearities, have seldom been addressed. This is remarkable in view of their prominence in theories of nutrition efficiency wages, and also in light of evidence in the biomedical literature which suggests that thresholds in health indicators may be important. For example, as discussed above, mortality risk is elevated only among those with very low or very high BMI, and a similar pattern is observed in the relationship between BMI and other health indicators. Consider nutrient intakes as a second example. It seems plausible that among very poorly nourished populations, additional energy intake 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. An additional but controversial argument posits that the body is able to adapt to moderately low (or high) intakes, and so there is no link between intakes and output over some range (except over the short term because the body has not had time to adjust); however, once a lower threshold has been passed, adaptation mechanisms are likely to be incomplete. The extent to which the body is able to successfully adapt to nutritional stress is clearly an empirical matter.

We would argue, therefore, that there are good reasons to investigate the shape of relationships between health and labor outcomes. The evi-

dence, however, is thin. In his analysis of farm production in Sierra Leone, Strauss (1986) finds evidence for mild concavity in the relationship between labor efficiency and per capita calorie availability in the household although the turning point is around 5,000 calories, which seems implausibly high. In contrast, Thomas and Strauss (1997) find a positive relationship between wages and calories for low per capita intakes, which flattens out around 2,000 calories per day and remains flat at higher intakes. The differences between these results may reflect differences in the nature of work in the two samples, since the heavy hoe-oriented agriculture dominant in rural Sierra Leone is more energy intensive than the average urban job in Brazil. We suspect, however, that the different results reflect differences in measurement of calories. There is too much spread in calorie availability in the Sierra Leone sample to be plausible, with too many people apparently consuming 4,000 calories or higher, and too many consuming 1,500 or lower; Srinivasan (1992) pointed out the same problem in comparable Indian NSS data. In the Sierra Leone and Indian samples, it seems very likely that poor households receive more food transfers and waste less than higher-income households. In contrast, the Brazil survey took special care to measure these leakages and wastage. Underestimation of calorie intakes of lower-income households is likely to lead to a downward bias in the estimated effect of calories on wages at the bottom of the calorie distribution; while overestimates of intakes among those with higher incomes would lead to upward-biased calorie estimates at the top. The net effect would be to twist the productivity-calorie relationship counterclockwise.

A number of other “nonlinearities”

have been explored in the literature and clarify our understanding of the relationship between dimensions of health and productivity. First, the returns to health are likely to vary with the type of work performed, with, for example, strength being rewarded most in manual labor. To sidestep the fact that wages and occupational choice are jointly determined, Thomas and Strauss stratify the Brazilian sample by education level under the presumption that manual laborers have less education. For the average urban Brazilian male, body mass is associated with higher productivity but that effect is significant *only* among those with little education. For them, strength or aerobic capacity is presumably more important at work than for the better educated, who apparently substitute brain for brawn as they climb the education ladder.

A second stratification that has proved informative distinguishes health impacts by gender. Behrman and Deolalikar (1989) find that body mass affects market wages of men in India, but that female wages are unrelated to body mass and calorie intake. Pitt, Rosenzweig, and Hassan (1990) and Bhargava (1997) find that while BMI positively influences the probability of engaging in strenuous work for men, it has a much smaller effect on women. Pitt, Rosenzweig, and Hassan also find that a higher average "health endowment" of males in the household is associated with increased household income, but not so for higher female "endowments." Finally, Thomas and Strauss (1997) show that body mass has a positive effect on productivity only among those women in Brazil who have no education at all (and are likely to be employed in jobs such as domestic service); for Brazilian men, the impact of BMI declines with education. All of these results are strongly suggestive that body mass is a

proxy for strength or aerobic capacity, which are productive assets in the labor market of men, and possibly women, when working in more menial jobs.⁴⁶

5.6 *Health and Constraints to Consumption Smoothing*

It has been possible to draw inferences regarding barriers to consumption smoothing by examining differences in the impacts of health, and in consumption of health inputs, across time (see Townsend 1995 for a general discussion). The existence of such barriers would imply a socially inefficient level of investments in health inputs. To date, the stratification that has proved most informative on this point, particularly in rural samples, distinguishes seasons. For instance, using the ICRISAT Indian data, Behrman and Deolalikar (1989) find that calories have a significant impact on wages during the peak labor demand season (an elasticity of 0.29), but not during the slack season. While weight-for-height has a larger impact on wages in the off-peak season. This may reflect differences in tasks performed across seasons, but it may also reflect the absence of perfect interseasonal consumption smoothing. If body fat is used to store energy for later use (rather than financial forms of saving), then BMI should have a larger average impact during the off-peak season.⁴⁷ An interesting but unanswered question is whether using the body as a store of energy savings across seasons is

⁴⁶ This gender difference in the link between BMI and wages is in stark contrast with the association between education and wages, which has as high, or even higher, marginal returns for women than men in many developing economies (Schultz 1993).

⁴⁷ It is difficult to know whether this inference is correct given the empirical result that peak labor demand season does not correspond to a low-income or "hungry" season (see Barbara Harriss 1995).

the most efficient method in the absence of barriers to smooth consumption.⁴⁸

Other evidence based on drawing comparisons across seasons provides even stronger evidence that consumption smoothing is not complete. Using data from Pakistan, Behrman, Foster, and Rosenzweig (1997) estimate an explicitly dynamic model that distinguishes harvest from planting seasons and takes into account the influence of assets and income on health, as well as the impact of health on farm profits. Their harvest-season profit function is conditioned on assets, harvest period prices, and planting period inputs. One of those inputs is planting period labor, which is augmented by contemporaneous household per capita calorie intake. They find that while calories have a positive, concave impact on farm profits, the magnitude of the effect is small. However, an increase in per capita calorie consumption throughout the planting season nearly pays for itself in terms of enhanced labor productivity. This suggests that more calories ought to be consumed during planting season. Instead, calorie consumption is shown to rise during harvest season, which, together with the first finding, strongly suggests an absence of consumption smoothing mechanisms that could be used to raise calorie consumption during planting season.

5.7 Observability of Productivity and of Health

We turn next to the issue of observability of health. It was argued above that in the context of labor markets with short-term contracts, even if productivity is enhanced by good health, it will be rewarded by higher

⁴⁸ Given high food storage losses in some rural environments it is not so clear, see Philip Payne (1989) for a detailed discussion and simulations.

wages only if the employer can observe either productivity or health (or something correlated with them that can be used as a marker). The most convincing studies of this issue are of daily workers in Philippine agriculture (Foster and Rosenzweig 1993a, 1994). The impact of health on time wages is compared with its impact on piece rates, which are presumably a good indicator of productivity. Exploiting the fact that both payment schemes coexist for harvesting at the same time of year, and that some workers engage in both types of contract simultaneously, Foster and Rosenzweig examine the effects of health on *differences* in the implicit time wages between time and piece work for the same worker. They are thereby able to control for all individual unobserved heterogeneity and place the spotlight on the role of observability of both productivity and health.⁴⁹ Calorie intakes have a significantly larger impact on piece rates than on time wages. Foster and Rosenzweig (1993a) argue that employers cannot directly observe nutrient intakes on daily workers and so intakes 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 1994), which is consistent with the interpretation that BMI serves in part as a signal to employers for nutrient intakes.⁵⁰

⁴⁹ 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).

⁵⁰ Consistent with this evidence, recall that Behrman and Deolalikar (1989) report that in India calories have a significant impact on wages during the peak season (harvest), but not at other times. It is the case that piece rates dominate during the harvest season in the ICRISAT villages, while daily wages are more common at other times. It should also be true that calorie intake raises productivity among self-employed farmers in the same villages; yet it does not (Deolalikar 1988).

These differences may, however, reflect differences in the behavior of each worker as he or she chooses how much effort to expend on time-rate and piece-rate work. Unless monitoring is costless, daily labor contracts have incentive problems that may not arise in piece-rate schemes, when working one's own land, or when contracts involve linkages with other markets, such as sharecropping arrangements. Specifically, workers are more likely to shirk and expend less effort when paid on a time basis. Foster and Rosenzweig (1994) point out that by comparing energy intakes and energy expenditures across different labor contracts, it is possible to assess the extent to which there is shirking (or moral hazard in the labor market). Effort is measured by energy expenditures which are deduced from changes in body mass controlling for calorie intake. They estimate a body-mass production function that depends on lagged body mass, on current individual energy intakes, and on whether the worker is ill. Controlling for these inputs, the impact on body mass of time worked in different labor contracts is examined.⁵¹ They report no significant differences in worker effort between those who toil more for piece rates and those working more time on their own land. But daily contract workers expend significantly less effort than the others (that is, they lose less weight), and Foster and Rosenzweig conclude that there is substantial moral hazard associated with time-wage employment. Corroborating evidence is provided by the fact that calorie intake is positively associated with piece-rate work or with working on one's own farm, holding constant food prices, body mass, and illness.

⁵¹ BMI inputs, including the time allocation variables, are treated as endogenous, using height, land owned, season, and age as identifying variables.

The ideas underlying this paper are clever, and provide excellent examples of how thinking carefully about the relationship between health and labor outcomes yields novel empirical tests of ideas that have played a central role in many models of development and labor markets. However, as discussed at length in Section 4, the credibility of results in this literature depends critically on the exclusion restrictions. One can certainly raise questions about some of these. Height, season of the survey, and land owned are all included in the set of identifying instruments: one may well argue that season (independent of net energy intakes) may affect the biology of BMI. The latter exclusion is perhaps key, since, as Foster and Rosenzweig tell us, piece rates are predominantly paid in the harvest season. These concerns, however, may be tempered by the fact that Foster and Rosenzweig find consistent results with different tests that make different assumptions.

In sum, the evidence may be viewed as supporting the hypothesis that worker effort and the dimensions of health that are hard to observe will have more impact on labor productivity when contracts are incentive compatible. Costs that either give rise to imperfect information by employers or that impede the emergence of incentive-compatible contracts thus induce social inefficiencies regarding investments in health, leading presumably to underinvestment in health inputs. This interpretation, however, relies crucially on assumptions about the costs of monitoring worker effort and observing health indicators. For example, contracts are not immutable, and if day laborers 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 labor market in Bukidnon, the area of the Philippines wherein the data were collected. What does exist are contracts in which meals served are part of the remuneration.⁵² By serving meals at the job site, employers can be sure that a minimum nutrient intake is reached and thus reduce the observability problem with respect to nutrient intakes.⁵³ Foster (1995) presents results indicating that more productive Filipinos work on jobs that provide meals, suggesting that meals lead to higher productivity. The generality of these results remains to be established. For example, 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.⁵⁴

5.8 *Efficiency Wage Hypothesis: Empirical Evidence*

Finally, we turn to the empirical implications of the nutrition efficiency wage hypothesis which, as mentioned, posits that there is a nonlinear link between health and productivity and that the function is convex at very low levels of health. This implies that there may be some people who are so poor and so unhealthy that they are too costly to be employed.

The evidence in support of the hypothesis is thin (see Bliss and Stern 1978b, and Rosenzweig 1988 for discussions). One implication for which there

⁵² Serving meals is fairly common in many rural areas in low-income economies.

⁵³ However, serving meals may result from other considerations: Saving on transaction costs of workers traveling long distances to and from home is one; using meals to build team spirit when team production is involved is another.

⁵⁴ Glick and Sahn (1997) find the opposite result for the wages of urban male workers in Guinea. However, they also find that body mass of males has a larger impact on market than on self-employment wages.

is no evidence is that, in equilibrium, there is widespread starvation, and presumably death, in low-income settings. While starvation and death does occur in times of famine due to natural, political, or economic shocks, this is not an equilibrium condition.

Nevertheless, a number of results in the literature are consistent with the efficiency wage hypothesis. For example, as noted above, people in better health (measured by BMI, for example), are more likely to undertake strenuous tasks (Pitt, Rosenzweig, and Hassan 1990; Bhargava 1997). But that evidence is also consistent with a simple model of investment in human capital without involuntary equilibrium unemployment, as discussed above.⁵⁵ Similarly, the efficiency wage hypothesis predicts that poor households will allocate health inputs unequally among members to ensure that at least one member is fit for work. There is evidence that household resources are allocated unequally toward workers, which is consistent with the hypothesis; but the evidence is not consistent with the additional implication that workers should receive preferential treatment in allocations. For example, in Bangladesh men are more likely to be workers, and those with better health “endowments” incur the heaviest health tax in terms of within-family allocations (Pitt, Rosenzweig, and Hassan 1990).⁵⁶

Several studies indicate that poor health reduces labor force participation. According to the efficiency wage hypothesis, the effect should be greatest

⁵⁵ Recall that such an equilibrium is possible in the Dasgupta–Ray model.

⁵⁶ They show that men who have better “endowments” tend to be fed more, relative to other household members. However, the effect on health of those better-endowed men working in more energy-intensive jobs outweighs the effect of having larger calorie intakes. This represents a “tax” on the health of better-endowed men.

among the poorest. In most studies, it is not; see, for example, Schoenbaum's (1995) examination of older men and women in Taiwan. Dow et al. (1997) find that changes in health have a bigger impact on changes in participation among the poor in Indonesia, although they argue that this is likely to be because health care price increases took their greatest toll on the budgets of the poorest. Again, the evidence is more consistent with a human capital interpretation of the effect of health on labor market outcomes than with an efficiency wage interpretation.

One might argue that standards of living in Taiwan and Indonesia are too high to detect a labor market equilibrium with rationing based on low nutrition or health. India is a more promising context for finding such an equilibrium. However, Anand Swamy (1997) has documented that rural wage levels in much of India are simply too large to be consistent with there being a nutrition-based constraint to labor force participation.⁵⁷ He shows, using several data sets covering both casual and permanent workers, that daily wages would purchase roughly 5,000 calories, assuming the income is spent in the same way an average poor household allocates its expenditures. While these calories might well be fed to multiple family members, the key point is that an unemployed worker could offer to work for half the wage, given in the form of meals (which are often part of the wage payment). As long as the resultant labor efficiency is greater than half that of current workers, it would be cheaper for employers to hire that worker. Using calorie-wage estimates from Behrman and Deolalikar (1989), Swamy estimates that worker efficiency would be

only 36 percent greater at 5,000 calories than it would be at 2,500. Hence it would appear unlikely that these wage levels, plus low assets, are responsible for large-scale involuntary unemployment among the poor. Moreover, there is little evidence that involuntary unemployment is widespread in rural India (see the discussion in Rosenzweig 1988, for example). Within the Dasgupta-Ray framework discussed earlier, it would appear that the state of the world is normally one in which equilibrium piece-rate wages are above the minimum wage per unit of work (subject to being above the reservation wage) of landless (or low-income) workers.

The efficiency wage hypothesis speaks to a market adjustment in response to specific nonlinearities in the relationship between health and productivity. It is important to recognize that failure to find empirical support for that hypothesis says nothing about whether or not health affects productivity at the individual level. As we have discussed, there is abundant evidence that a range of dimensions of health do reap a reward in the labor market. Moreover, there is some evidence indicating that households allocate their resources in response to the fact that poor health may impede some people from engaging in certain types of work, and that there are complex interrelationships between the nature and observability of health and of the work performed.

It has been difficult to establish that poor health results in catastrophic labor market outcomes. However, moderate to severe shortfalls in health do appear to have negative impacts on productivity and to reduce labor supply. Not only have these patterns emerged in the empirical literature summarized above, they also appear in the historical literature. Fogel (1992), for example, inter-

⁵⁷ Shankar Subramanian and Angus Deaton (1996) note the same fact, also using Indian data.

prets the decline of European mortality in the 19th and 20th centuries as resulting primarily from improvements in nutrition among the moderately undernourished.

6. Conclusion

There is a general consensus that income has a strong effect on demand for health. Until recently, there was less agreement on the reverse relationship: the effect of health on income or, more generally, labor outcomes. Even fairly recent reviews of this literature conclude that there is little reliable evidence that health has an important impact on labor productivity or labor use (Rosenzweig 1988). Today that assessment would have to be amended: In recent years, substantial progress has been made in documenting the existence of a causal impact of health on wages and productivity in low-income settings using both experimental and nonexperimental methods. A small number of studies suggest that health has a larger return at very low levels of health and (perhaps) in jobs requiring more strength. With economic development, these types of jobs will shrink, and one might expect the labor market impact of improved health to decline, especially relative to the impact of education and skill acquisition.

However, the critical reader will agree that teasing these relationships out of data has proved to be very difficult. Health and income clearly affect each other and are related to many factors that are hard to measure. In the nonexperimental literature, interpretations of associations between health, wages, and income are not unambiguous. But analyses of these sorts of data that are judicious in their choice of assumptions, thoughtful in the choice of estimation, and sensitive to issues of ro-

bustness promise to be very profitable. Many studies have been tripped up by inadequate data. Investments in improved data, particularly longitudinal surveys with good measures of health status, are likely to yield large returns.

How to improve our health measurements, while not constituting economic analysis in itself, has very important ramifications for economic research. Work remains to be done on which dimensions of health matter, which population groups are most affected, and under what labor market (or other) conditions health emerges as important. Questions of the impact of health dynamics, particularly in response to negative economic shocks and aging, have barely been touched in low-income environments. The same is true of the important issue of whether inefficiencies exist in household and individual health input allocations, given the estimated productive impacts. Thus, we conclude that, while considerable progress has been made in recent years, research on the interactions between labor markets and health remains a very promising area in the study of developing economies.

REFERENCES

- ACTON, JAN. "Nonmonetary Factors in the Demand for Medical Service: Some Empirical Evidence," *J. Polit. Econ.*, June 1975, 83(3), pp. 595-614.
- ALTONJI, JOSEPH. "The Effects of High School Curriculum on Education and Labor Market Outcomes," *J. Human Res.*, Summer 1995, 30(3) pp. 409-38.
- ANDRIANO, RALPH AND HELMINIAK, THOMAS. "Economics, Health and Tropical Diseases: A Review," in *Economics, health and tropical diseases*. Eds.: ALEJANDRO HERRIN AND PATRICIA ROSENFELD. Manila, Philippines: U. of the Philippines, School of Economics, 1988, pp. 19-72.
- ANTLE, JOHN AND PINGALI, PRABHU. "Pesticides, Productivity and Farmer Health: A Philippine Case Study," *Amer. J. Agr. Econ.*, Aug. 1994, 76(3), pp. 418-30.
- ASHENFELTER, ORLEY AND KRUEGER, ALAN. "Estimates of the Economic Return to School-

- ing From a New Sample of Twins," *Amer. Econ. Rev.*, Dec. 1994, 84(5), pp. 1157-73.
- BASTA, SAMIR ET AL. "Iron Deficiency Anemia and Productivity of Adult Males in Indonesia," *Amer. J. Clinical Nutrition*, Apr. 1979, 32(4), pp. 916-25.
- BECKER, GARY. "A Theory of the Allocation of Time," *Econ. J.*, Sept. 1965, 75(299), pp. 493-517.
- BEHRMAN, JERE R. "Intrahousehold Allocation of Nutrients in Rural India: Are Boys Favored?" *Oxford Econ. Pap.*, Feb. 1988, 40(1), pp. 32-54.
- . "Intrahousehold Allocation of Nutrients and Gender Effects: A Survey of Structural and Reduced Form Estimates," in SADDIQ R. OSMANI 1992, pp. 287-321.
- . "The Economic Rationale for Investing in Nutrition in Developing Countries," *World Devel.*, Nov. 1993, 21(11), pp. 1749-72.
- . "The Impact of Health and Nutrition on Education," *World Bank Res. Observer*, Feb. 1996, 11(1), pp. 23-37.
- . "Intrahousehold Distribution and the Family," in *Handbook of population and family economics*, Vol. 1A, Eds.: MARK R. ROSENZWEIG AND ODED STARK. Amsterdam: North Holland Press, 1997, pp. 125-87.
- BEHRMAN, JERE R. AND DEOLALIKAR, ANIL B. "Health and Nutrition," in *Handbook of development economics*, Vol. 1. Eds.: HOLLIS CHENERY AND T. N. SRINIVASAN. Amsterdam: North Holland Press, 1988, pp. 631-711.
- . "Agricultural Wages in India: The Role of Health, Nutrition and Seasonality," in *Seasonal variability in third world agriculture*. Ed.: DAVID SAHN. Baltimore: Johns Hopkins U. Press, 1989, pp. 107-17.
- BEHRMAN, JERE R., FOSTER, ANDREW AND ROSENZWEIG, MARK R. "The Dynamics of Agricultural Production and the Calorie-Income Relationship," *J. Econometrics*, Mar. 1997, 77(1), pp. 187-208.
- BENJAMIN, DWAYNE. "Household Composition, Labor Markets, and Labor Demand: Testing for Separation in Agricultural Household Models," *Econometrica*, Mar. 1992, 60(2), pp. 287-322.
- BESLEY, TIMOTHY AND CASE, ANNE. "Incumbent Behavior: Vote Seeking, Tax Setting and Yardstick Competition," *Amer. Econ. Rev.*, Mar. 1995, 85(1), pp. 25-45.
- BHARGAVA, ALOK. "Nutritional Status and the Modeling of Time in Rwandese Households," *J. Econometrics*, Mar. 1997, 77(1), pp. 277-95.
- BLISS, CHRISTOPHER AND STERN, NICHOLAS. "Productivity, Wages and Nutrition: Part I: The Theory," *J. Devel. Econ.*, Dec. 1978a, 5(4), pp. 331-62.
- . "Productivity, Wages and Nutrition: Part II: Some Observations," *J. Devel. Econ.*, Dec. 1978b, 5(4), pp. 363-98.
- BOUIS, HOWARTH E. AND HADDAD, LAWRENCE J. "Are Estimates of Calorie-Income Elasticities Too High? A Recalibration of the Plausible Range," *J. Devel. Econ.*, Oct. 1992, 39(2), pp. 333-64.
- BOUND, JOHN. "Self-Reported Objective Measure of Health in Retirement Models," *J. Human Res.*, Winter 1991, 26(1), pp. 106-38.
- CHIAPPORI, PIERRE-ANDRE. "Rational Household Labor Supply," *Econometrica*, Jan. 1988, 56(1), pp. 63-89.
- . "Collective Labor Supply and Welfare," *J. Polit. Econ.*, June 1992, 100(3), pp. 437-67.
- CLEVELAND, WILLIAM S. "Robust Locally-Weighted Regression and Smoothing Scatterplots," *J. Amer. Statist. Assoc.*, Dec. 1979, 74(368), pp. 829-36.
- DASGUPTA, PARTHA. *An inquiry into well-being and destitution*. Oxford: Oxford U. Press, 1993.
- . "Nutritional Status, the Capacity for Work and Poverty Traps," *J. Econometrics*, Mar. 1997, 77(1), pp. 5-38.
- DASGUPTA, PARTHA AND RAY, DEBRAJ. "Inequality as a Determinant of Malnutrition and Unemployment: Theory," *Econ. J.*, Dec. 1986, 96(384), pp. 1011-34.
- . "Inequality as a Determinant of Malnutrition and Unemployment: Policy," *Econ. J.*, Mar. 1997, 97(385), pp. 177-88.
- DEOLALIKAR, ANIL B. "Nutrition and Labor Productivity in Agriculture: Estimates for Rural South India," *Rev. Econ. Statist.*, Aug. 1988, 70(3), pp. 406-13.
- DOW, WILLIAM ET AL. "Health Care Prices, Health and Labor Outcomes: Experimental Evidence." Labor and Population Program Working Paper Series 97-01, RAND, Santa Monica, CA, Jan 1997.
- FALKNER, FRANK AND TANNER, JAMES M. *Human growth: A comprehensive treatise*, Vol. 3. New York: Plenum Press, 1986.
- FENWICK, A. AND FIGENSCHOU, B. M. "The Effect of *Schistosoma Mansoni* on the Productivity of Cane Cutters on a Sugar Estate in Tanzania," *Bull. World Health Organ.*, Sept. 1972, 47(5), pp. 567-72.
- FOGEL, ROBERT "Second Thoughts on the European Escape From Hunger: Famines, Chronic Malnutrition and Mortality Rates," in SADDIQ R. OSMANI 1992, pp. 243-86.
- FOGEL, ROBERT. "Economic Growth, Population Theory and Physiology: The Bearing of Long-Term Processes on the Making of Economic Policy," *Amer. Econ. Rev.*, June 1994, 84(3), pp. 369-95.
- FOLBRE, NANCY. "Market Opportunities, Genetic Endowments and Intrafamily Resource Distribution: Comment," *Amer. Econ. Rev.*, June 1984, 74(3), pp. 518-20.
- FOSTER, ANDREW. "Nutrition and Health Investment," *Amer. Econ. Rev.*, May 1995, 85(2), pp. 148-52.
- FOSTER, ANDREW AND ROSENZWEIG, MARK R. "Information Flows and Discrimination in Labor Markets in Rural Areas in Developing Countries," *Proceedings of the World Bank An-*

- nual Conference on Development Economics 1992, Mar. 1993a, pp. 173–204.
- . “Information, Learning and Wage Rates in Low-Income Rural Areas,” *J. Human Res.*, Fall 1993b, 28(4), pp. 759–90.
- . “A Test for Moral Hazard in the Labor Market: Effort, Health and Calorie Consumption,” *Rev. Econ. Statist.*, May 1994, 76(2), pp. 213–27.
- . “Comparative Advantage, Information and the Allocation of Workers to Tasks: Evidence From the Agricultural Labour Market,” *Rev. Econ. Stud.*, July 1996, 63(3), pp. 347–74.
- FRANKENBERG, ELIZABETH. “The effects of access to health care on infant mortality in Indonesia,” *Health Transitions Review*, 1995, 5(2), pp. 143–62.
- FRIEDMAN, GERALD. “The Heights of Slaves in Trinidad,” *Soc. Sci. Hist.*, Fall 1982, 6(4), pp. 482–515.
- GATEFF, C. ET AL. “Influence de la Bilharziose Vesicale sur la Rentabilite Economique d’une Population Adult Jeune du Cameroon,” *Annales de la Societe Belge de Medicin Tropical*, 1971, 51(3), pp. 309–24.
- GERSOVITZ, MARK. “Savings and Nutrition at Low Incomes,” *J. Polit. Econ.*, Oct. 1983, 91(5), pp. 841–55.
- GERTLER, PAUL AND MOLYNEAUX, JACK. “Pricing Public Health Services: Results From a Social Experiment in Indonesia.” Mimeo, RAND, Santa Monica, CA, 1996.
- GERTLER, PAUL AND VAN DER GAAG, JACQUES. *The willingness to pay for medical care: Evidence from two developing countries.* Baltimore: Johns Hopkins U. Press, 1990.
- GLICK, PETER AND SAHN, DAVID. “Health and Productivity in a Heterogeneous Urban Labor Market.” Mimeo, Cornell U., Mar. 1997.
- GLOMM, GERHARD AND PALUMBO, MICHAEL. “Optimal Intertemporal Consumption Decisions Under the Threat of Starvation,” *J. Devel. Econ.*, Dec. 1993, 42(2), pp. 271–92.
- GRILICHES, ZVI. “Estimating the Returns to Schooling: Some Econometric Problems,” *Econometrica*, Jan. 1977, 45(1), pp. 1–22.
- GROSSMAN, MICHAEL. “On the Concept of Health Capital and the Demand for Health,” *J. Polit. Econ.*, Apr. 1972, 80(2), pp. 223–55.
- HADDAD, LAWRENCE J. AND BOUIS, HOWARTH E. “The Impact of Nutritional Status on Agricultural Productivity: Wage Evidence From the Philippines,” *Oxford Bull. Econ. Statist.*, Feb. 1991, 53(1), pp. 45–68.
- HARRISS, BARBARA. “Gender Bias in Intrahousehold Nutrition in South India: Unpacking Households and the Policy Process,” in *Intrahousehold resource allocation in developing countries: Methods, models and policy*. Eds.: LAWRENCE HADDAD, JOHN HODDINOTT, AND HAROLD ALDERMAN. Baltimore: Johns Hopkins U. Press, 1997.
- HILL, ALLAN AND MAMDANI, MASUMA. “Operational Guidelines for Measuring Health Through Household Surveys.” Mimeo, Centre for Population Studies, School of Hygiene and Tropical Medicine, U. of London, 1989.
- HIRSCHMAN, CHARLES; PRESTON, SAMUEL AND LOI, VU MANH. “Vietnamese Casualties During the American War,” *Population Devel. Rev.*, Dec. 1995, 21(4), pp. 783–812.
- IMMINCK, MAARTEN AND VITERI, FERNANDO. “Energy Intake and Productivity of Guatemalan Sugarcane Cutters: An Empirical Test of the Efficiency Wage Hypothesis—Part 1,” *J. Devel. Econ.*, Oct. 1981a, 9(2), pp. 251–71.
- . “Energy Intake and Productivity of Guatemalan Sugarcane Cutters: An Empirical Test of the Efficiency Wage Hypothesis—Part 2,” *J. Devel. Econ.*, Oct. 1981b, 9(2), pp. 273–87.
- JAMISON, DEAN ET AL., eds. *Disease control priorities in developing countries*. Oxford: Oxford U. Press, 1993.
- JOHN, MEREDITH. “Plantation Slave Mortality in Trinidad,” *Population Stud.*, July 1988, 42(2), pp. 161–82.
- LAVY, VICTOR; PALUMBO, MICHAEL AND STERN, STEVEN. “Health Care in Jamaica: Quality, Outcomes, and Labor Supply.” Living Standards Measurement Study Working Paper No. 116, World Bank, 1995.
- LEIBENSTEIN, HARVEY. *Economic backwardness and economic growth: Studies in the theory of economic development*. New York: Wiley & Sons, 1957.
- LOPEZ, ALAN. “Causes of Death in Industrial and Developing Countries: Estimates for 1985–1990,” in DEAN JAMISON ET AL. 1993, pp. 35–50.
- LUNDBERG, SHELLY AND POLLAK, ROBERT. “Separate Spheres Bargaining and the Marriage Market,” *J. Polit. Econ.*, Dec. 1993, 101(6), pp. 988–1010.
- MCELROY, MARJORIE AND HORNEY, JEAN. “Nash-bargained Household Decisions: Toward a Generalization of the Theory of Demand,” *Int. Econ. Rev.*, June 1981, 22(2), pp. 333–49.
- MARGO, ROBERT AND STECKEL, RICHARD. “The Heights of American Slaves: New Evidence on Slave Nutrition and Health,” *Soc. Sci. Hist.*, Fall 1982, 6(4), pp. 516–38.
- MARTORELL, REYNALDO AND ARROYAVE, GUILLERMO. “Malnutrition, Work Output and Energy Needs,” in *Capacity for work in the tropics*. Society for the Study of Human Biology Symposium 26. Eds.: K. COLLINS AND F. ROBERTS. Cambridge: Cambridge U. Press, 1988, pp. 57–76.
- MAZUMDAR, DIPAK. “The Marginal Productivity Theory of Wages and Unemployment,” *Rev. Econ. Stud.*, June 1959, 26(3), pp. 190–97.
- MIRRELES, JAMES. “A Pure Theory of Underdeveloped Economies,” in *Agriculture in development theory*. Ed.: LLOYD G. REYNOLDS. New Haven: Yale U. Press, 1975, pp. 84–108.
- MUNDLAK, YAIR. “Estimation of Production and Behavioral Functions From a Combination of Time Series and Cross Section Data,” in *Mea-*

- surement in economics: *Studies in mathematical economics and econometrics*. Ed.: CARL CHRIST ET AL. Stanford: Stanford U. Press, 1963, pp. 138–66.
- NELSON, CHARLES R. AND STARTZ, RICHARD. "The Distribution of the Instrumental Variables Estimator and Its t-ratio When the Instrument is a Poor One," *J. Bus.*, Jan. 1990, 63(1) Part 2, pp. 125–40.
- NEWHOUSE, JOSEPH AND THE INSURANCE EXPERIMENT GROUP. *Free for all? Lessons from the RAND Health Insurance Experiment*. Cambridge: Harvard U. Press, 1993.
- OSMANI, SIDDIQ R., ed. *Nutrition and poverty*. Oxford: Clarendon Press, 1992.
- OVER, MEAD, ET AL. "The Consequences of Adult Ill-Health," in *The health of adults in the developing world*. Eds.: RICHARD FEACHEM ET AL. Oxford: Oxford U. Press, 1992, pp. 161–208.
- PARSONS, DONALD. "The Male Labor Force Participation Decision: Health, Reported Health and Economic Incentives," *Economica*, Feb. 1982, 49(193), pp. 81–91.
- PAYNE, PHILIP. "Public Health and Functional Consequences of Seasonal Hunger and Malnutrition," in *Seasonal variability in Third World agriculture: The consequences for food security*. Ed.: DAVID SAHN. Baltimore: Johns Hopkins U. Press, 1989, pp. 19–46.
- PINSTRUP-ANDERSEN, PER ET AL. "Protein-energy Malnutrition," in DEAN JAMISON ET AL. 1993, pp. 391–420.
- PITT, MARK, M. "Specification and Estimation of the Demand for Goods Within the Household," in *Intrahousehold resource allocation in developing countries: methods, models and policy*. Eds.: LAWRENCE HADDAD, JOHN HODDINOTT, AND HAROLD ALDERMAN. Baltimore: Johns Hopkins U. Press, 1997.
- PITT, MARK M. AND ROSENZWEIG, MARK R. "Agricultural Prices, Food Consumption and the Health and Productivity of Indonesian Farmers," in INDERJIT SINGH, LYN SQUIRE, AND JOHN STRAUSS 1986, pp. 153–82.
- . "Estimating the Intrahousehold Incidence of Illness: Child Health and Gender-Inequality in the Allocation of Time," *Int. Econ. Rev.*, Nov. 1990, 31(4), pp. 969–80.
- PITT, MARK M., ROSENZWEIG, MARK R. AND HASSAN, MD. NAZMUL. "Productivity, Health and Inequality in the Intrahousehold Distribution of Food in Low-Income Countries," *Amer. Econ. Rev.*, Dec. 1990, 80(5), pp. 1139–56.
- PITT, MARK M., ROSENZWEIG, MARK R. AND GIBBONS, DONNA. "The Determinants and Consequences of the Placement of Government Programs in Indonesia," *World Bank Econ. Rev.*, Sept. 1993, 7(3), pp. 319–48.
- POLLITT, ERNESTO ET AL. *Early supplementary feeding and cognition*. Monographs of the Society for Research on Children, Vol. 58, No. 7, Chicago: U. of Chicago Press, 1993.
- POLLITT, ERNESTO. "Iron Deficiency and Educational Deficiency," *Nutritional Reviews*, 1997, 55(4), pp. 133–40.
- PRESCOTT, NICHOLAS. "Schistosomiasis and Development," *World Devel.*, Jan. 1979, 7(1), pp. 1–14.
- ROSENZWEIG, MARK R. "Labor Markets in Low Income Countries," in *Handbook of development economics*, Vol. 1. Eds.: HOLLIS CHENERY AND T. N. SRINIVASAN. Amsterdam: North Holland Press, 1988, pp. 713–62.
- ROSENZWEIG, MARK R. AND SCHULTZ, T. PAUL. "Market Opportunities, Genetic Endowments and Intrafamily Resource Distribution: Child Survival in Rural India," *Amer. Econ. Rev.*, Sept. 1982, 72(4), pp. 803–15.
- ROSENZWEIG, MARK R. AND WOLPIN, KENNETH I. "Evaluating the Effects of Optimally Distributed Public Programs: Child Health and Family Planning Interventions," *Amer. Econ. Rev.*, June 1986, 76(3), pp. 470–82.
- . "Migration Selectivity and the Effects of Public Programs," *J. Public Econ.*, Dec. 1988, 37(3), pp. 265–89.
- SAHN, DAVID AND ALDERMAN, HAROLD. "The Effects of Human Capital on Wages, and the Determinants of Labor Supply in a Developing Country," *J. Devel. Econ.*, Sept. 1988, 29(2), pp. 157–83.
- SCHOENBAUM, MICHAEL. "The Health Status and Labor Force Behavior of the Elderly in Taiwan." Comparative Study of the Elderly in Asia, Research Report No. 95–32. Population Studies Center, U. of Michigan, 1995.
- SCHULTZ, T. PAUL. "Returns to Women's Education," in *Women's education in developing countries: Barriers, benefits and policies*. Eds.: ELIZABETH M. KING AND M. ANNE HILL. Baltimore: Johns Hopkins Press, 1993, pp. 51–99.
- . "Human Capital and Economic Development," in *Agricultural competitiveness: Market forces and policy choices*. Proceedings of the 22nd International Conference of Agricultural Economists. Eds.: GEORGE H. PETERS AND DOUGLAS D. HEDLEY. Aldershot: Dartmouth, 1995, pp. 523–39.
- . "Wage Rentals for Reproducible Human Capital: Evidence from Two West African Countries." Mimeo, Department of Economics, Yale U., 1996.
- SCHULTZ, T. PAUL AND TANSEL, AYSIT. "Wage and Labor Supply Effects of Illness in Cote d'Ivoire and Ghana," *J. Devel. Econ.*, Aug. 1997, 53(2), pp. 251–86.
- SINGH, INDERJIT, SQUIRE, LYN AND STRAUSS, JOHN, eds. *Agricultural household models: Extensions, applications and policy*. Baltimore: Johns Hopkins U. Press, 1986.
- SPURR, G. B. "Nutritional Status and Physical Work Capacity," *Yearbook of Physical Anthropology*, 1983, 26, pp. 1–35.
- . "Body Size, Physical Work Capacity and Productivity in Hard Work: Is Bigger Better?" in JOHN C. WATERLOW 1988, pp. 215–44.
- SRINIVASAN, T. N. "Undernutrition: Concepts,

- Measurements and Policy Implications," in SIDDIQ R. OSMANI 1992, pp. 97-120.
- . "Destitution: A Discourse," *J. Econ. Lit.*, Dec. 1994, 32(4), pp. 1842-55.
- STAIGER, DOUGLAS AND STOCK, JAMES. "Instrumental Variables Regression With Weak Instruments," *Econometrica*, May 1997, 65(3), pp. 557-86.
- STECKEL, RICHARD. "Stature and the Standard of Living," *J. Econ. Lit.*, Dec. 1995, 33(4), pp. 1903-40.
- STEWART, ANITA AND WARE, JOHN. *Measuring functioning and well-being: The Medical Outcomes Study*. Durham, NC: Duke U. Press, 1992.
- STEWART, ANITA ET AL. *Conceptualization and measurement of health status for adults in the Health Insurance Study: Vol. II. Physical health in terms of functioning*. R-1987/2-HEW, RAND, Santa Monica, CA, 1978.
- STIGLITZ, JOSEPH. "The Efficiency Wage Hypothesis, Surplus Labor and the Distribution of Income in LDCs," *Oxford Econ. Pap.*, July 1976, 28(2), pp. 185-207.
- STRAUSS, JOHN. "Does Better Nutrition Raise Farm Productivity?" *J. Polit. Econ.*, Apr. 1986, 94(2), pp. 297-320.
- . "The Impact of Improved Nutrition on Labor Productivity and Human Resource Development," in *The political economy of food and nutrition policies*. Ed.: PER PINSTRUP-ANDERSEN. Baltimore: Johns Hopkins Press, 1993, pp. 149-72.
- STRAUSS, JOHN AND THOMAS, DUNCAN. "Human Resources: Empirical Modeling of Household and Family Decisions," in *Handbook of development economics*, Vol. 3A. Eds.: T. N. SRINIVASAN AND JERE BEHRMAN. Amsterdam: North Holland, Press, 1995a, pp. 1883-2023.
- . "Health, Nutrition and Economic Development." Labor and Population Program Working Paper Series 95-23, RAND, Santa Monica, CA, 1995b.
- . "Measurement and Mismeasurement of Social Indicators," *Amer. Econ. Rev.*, May 1996, 86(2), pp. 30-34.
- STRAUSS, JOHN ET AL. "Gender and Life-Cycle Differentials in the Patterns and Determinants of Adult Health," *J. Human Res.*, Fall 1993, 28(4), pp. 791-837.
- SUBRAMANIAN, SHANKAR AND DEATON, ANGUS. "The Demand for Food and Calories," *J. Polit. Econ.*, Feb. 1996, 104(1), pp. 133-62.
- SWAMY, ANAND. "A Simple Test of the Nutrition-Based Efficiency Wage Model," *J. Devel. Econ.*, June 1997, 53(1), pp. 85-98.
- THOMAS, DUNCAN. "Intra-Household Resource Allocation: An Inferential Approach," *J. Human Res.*, Fall, 1990, 25(4), pp. 635-64.
- THOMAS, DUNCAN AND STRAUSS, JOHN. "Health and Wages: Evidence on Men and Women in Urban Brazil," *J. Econometrics*, Mar. 1997, 77(1), pp. 159-86.
- TOWNSEND, ROBERT. "Consumption Insurance: An Evaluation of Risk-Bearing Systems in Low-Income Economies," *J. Econ. Perspectives*, Summer 1995, 9(3), pp. 83-102.
- WAALER, HANS. "Height, Weight and Mortality: The Norwegian Experience," *Acta Medica Scandinavia*, 1984, supplement no. 679.
- WARE, JOHN, DAVIES-AVERY, ALLYSON AND BROOK, ROBERT. *Conceptualization and measurement of health status for adults in the Health Insurance Study: Vol. VI. Analysis of relationships among health status measures*. R-1987/6-HEW, RAND, Santa Monica, CA, 1980.
- WARE, JOHN; DAVIES-AVERY, ALLYSON AND DONALD, CATHY. *General health perceptions*. R-1987/5-HEW, RAND, Santa Monica, CA, 1978.
- WATERLOW, JOHN C., ed. *Linear growth retardation in less developed countries*. Nestle Nutrition Workshop Series Vol. 14, New York: Raven Press, 1988.
- WATERLOW, JOHN C. ET AL. "The Presentation and Use of Height and Weight Data for Comparing the Nutritional Status of Groups of Children Under the Age of 10 Years," *Bulletin of the World Health Organization*, July 1977, 55(4), pp. 489-98.
- WOLGEMUTH, JUNE ET AL. "Worker Productivity and the Nutritional Status of Kenyan Road Construction Laborers," *Amer. J. Clinical Nutrition*, July 1982, 36(7), pp. 68-78.
- WORLD BANK. *World Development Report 1993*, Oxford: Oxford U. Press, 1993.