Understanding the Mechanisms of Economic Development

Angus Deaton

Angus Deaton is Dwight D. Eisenhower Professor of International Affairs and Professor of Economics and International Affairs at the Woodrow Wilson School of Public and International Affairs and in the Department of Economics, Princeton University, Princeton, New Jersey. His e-mail aggress is <deaton@princeton.edu>.

Understanding the mechanisms of economic development

Angus Deaton

April 2010

ABSTRACT

I argue that progress in understanding economic development (as in other branches of economics) must come from the investigation of *mechanisms*; the associated empirical analysis can usefully employ a wide range of experimental and non-experimental methods. I discuss three different areas of research: the life-cycle saving hypothesis and its implication that economic growth drives higher rates of national saving, the theory of speculative commodity storage and its implications for the time-series behavior of commodity prices, and the relationship between economic growth and nutritional improvement. None of these projects has yet been entirely successful in offering a coherent account of the evidence, but all illustrate a process of trial and error, in which although mechanisms are often rejected, unlikely theoretical propositions are sometimes surprisingly verified, while in all cases there is a process of learning about and subsequently modifying our understanding of the underlying mechanisms.

Economic development differs from most fields in economics because the study of low-income economies, and of people living in low-income economies, draws on all branches of economics. The particular study of economic growth is perhaps closest to being a field like labor or health economics, but what is currently referred to as development economics is broader, so that, although my examples involve important issues in economic development, my conclusions apply more broadly to other areas of applied economics. In an earlier paper, I have argued that learning about development requires us to investigate *mechanisms* (Deaton, 2009). Finding out about how people in low-income countries can and do escape from poverty is unlikely to come from the empirical evaluation of actual projects or programs, whether through randomized trials or econometric methods that are designed to extract defensible causal inferences, unless such analysis tries to discover *why* projects work rather than *whether* they work—however important the latter might be for purposes of auditing. By contrast, investigation, testing, and modification of mechanisms that can be widely applied, at least potentially, allows the integration of disparate empirical findings and comprises a progressive empirical research strategy.

In this paper, I discuss three lines of work that have elucidated mechanisms that are relevant for development: 1) connections between saving and growth; 2) the determinants of commodity prices, which are a key source of income for many developing countries; and 3) some unexpected puzzles that arise in considering the linkages between income and food consumption. In each case, my discussion illustrates what Cartwright (2007) calls the positivist approach to the hypothetico-deductive method. In this approach, mechanisms are proposed, key predictions derived and tested, and if falsified, the mechanisms are rejected or modified. If the predictions of a mechanism are confirmed, if they are sufficiently specific, and if they are hard to explain in other ways, we attach additional credence to the mechanism, albeit provisionally since

later evidence may undermine it. Sometimes the falsifications can be repaired by changing supplementary assumptions, and sometimes they involve long steps backwards when the model is abandoned, and there is often disagreement about which is the correct response. But the end result is an accumulation of useful knowledge and understanding.

Theory and Simple Observation: Savings and Growth

In a famous 1954 paper, Arthur Lewis (p. 155) wrote that "the central problem in the theory of economic development is to understand the process by which a community which was previously saving, and investing, 4 or 5 percent. of its national income or less, converts itself into an economy where voluntary saving is running at about 12 to 15 percent. of national income or more. This is the central problem because the central fact of economic development is rapid capital accumulation." Since that paper, economists have been proposing theories to explain saving and testing the implications of those theories in a research agenda that continues to pose questions for modern researchers. In this case, a number of the key tests involve straightforward observations that require no more than a cross tabulation or a two-dimensional graph—but what to put in the table or the graph was not at all obvious until the theory was developed to show the way.

At the same time Lewis (1954) was writing, Modigliani and Brumberg (1954a, 1954b [1990]) developed the life-cycle theory of saving, a mechanism that specifies how, why, and how much people will save. Their main focus was not economic development, but rather an attempt to provide a theoretically coherent account that could make sense of a mass of disorganized pre-existing empirical evidence from time-series and cross-sectional data. Their attempt was a brilliant success, and it still provides a useful framework for thinking about saving. It turned out

that the life-cycle theory also had important implications for the relationship between saving and growth, a link that was further developed by Modigliani (1966, 1970) who, in the second paper, took the predictions to international data. He showed that if people saved when they were young, and dissaved when they were old, both population growth and per capita income growth will increase the national saving rate; with population growth, there are more savers than dissevers, and with economic growth, the former are richer than the latter, so that their saving more than offsets the dissaving of their elders. If these predictions seem obvious now, they do not follow, for example, from a representative agent model which necessarily lacks a range of ages within a population. The theory further predicts that saving rates should be independent of the *level* of development, because everything is scaled in proportion to income, should be zero when growth is zero, because the saving of the young exactly matches the dissaving of the old, and that, in simple cases, the relationship between the saving rate and the growth rate should be a concave function whose parameters are determined by the ratio of the retirement to the work span, see Deaton (1997, pp 45–7) for an exposition.

In the 1970 paper, Modigliani drew on earlier work by Houthakker (1965), and pulled together the evidence available at that time, and showed that, indeed, saving rates are higher in more rapidly growing economies, a correlation that remains true today. He also confirmed that the *level* of development did not predict the saving rate, and that the size of the growth effect on saving matched what would be expected from earlier work fitting the model to time-series data from the United States. The life-cycle mechanism could thus offer a broad integration of widely dispersed evidence, from cross-sections, from time-series, from international comparisons, and from both rich and poor countries.

Of course, a correlation between growth and saving rates is hardly a surprise, and is predicted by other theories. Before Solow's (1956) paper, growth models predicted a simple mechanical link from saving to growth, which is what Lewis (1954) had in mind, and which Modigliani (1970) recognized and compared with the life-cycle account. After Solow, with the understanding that, in the long-run, growth rates do not depend on the saving rate, it is still the case that higher saving rates will generate higher growth along a potentially long-lived transitional path (Atkinson, 1969). Yet Modigliani's insight provided a new mechanism for an old correlation, and freed development economists to think of explanations for economic growth that did not depend on an increase in saving, with life-cycle saving providing the needed finance in response to the growth. His paper also illustrates the importance of something that is often neglected: the need to develop the theory, often with substantial additional work, until it delivers propositions that are non-obvious and that can readily and transparently be taken to the data. Modigliani's propositions were remarkable for their specificity. It is not simply that growth and saving should be correlated, but the relationship has a specific form, dependent on other observable quantities, and the relationship needs to be consistent with different relationships estimated on both cross-section and time-series data.

These predictions held up remarkably well for a long period of time, for two further decades, until they faced an empirical challenge based on a very simple observation. Carroll and Summers (1995) noted that, if the life-cycle theory were true, the cross-sectional age-profiles of consumption should rotate clockwise with the rate of economic growth. Younger cohorts of consumers have greater life-time resources than older cohorts of consumers, so that, whatever are people's preferences for the shape of their age profiles of consumption—provided only that those preferences are the same in all countries—the ratio of young people's to old people's

consumption in the momentary cross-section should be higher in more rapidly growing economies. In the simplest case, people might like their consumption to be the same at all ages, so that if we could track individuals, their consumption would not change as they age. If we compare a fast growing with a slow growing economy, young people in the former are lifetime richer relative to old people than are young people in the latter, so that the consumption of the young relative to the old will be higher in the more rapidly growing economy. If the preference is for consumption to grow with age, the age-profile of consumption will rise in a zero growth economy, but rotate clockwise as the rate of growth increases. That countries have the same preferences is a strong assumption, but is required in some form to get Modigliani's (1970) original correlation, so it entirely appropriate for testing the theory. Carroll and Summers' idea is easy enough to understand and describe, at least with hindsight, but was certainly not obvious in advance. Moreover, testing this prediction requires data only on consumption and age from a single household survey, and does not require microeconomic data on savings. (Just as the explorers who tried to find the source of the Nile were long frustrated by the *Sudd*, the giant impenetrable swamp that separates the sources of the Nile from its main channels in Sudan and Egypt, so the giant impenetrable data swamp that prevents us from connecting microeconomic and macroeconomic data on saving has long prevented the tracing of national saving back to its sources in household behavior.) Carroll and Summers' prediction required no more econometrics than a handful of graphs, which showed that the age-profiles of consumption are remarkably similar in slow growing and fast growing economies, which means that life-cycle saving cannot account for the positive association between growth and saving. Deaton (1997, pp 53–6) extends this argument to other poor and rich countries.

This is straightforward hypothetico-deductive methodology: a prediction is tested, the tests fail, and the hypothesis is wrong. There are other supporting contradictions. For example, the long decline in the U.S. saving rate (reversed only very recently) is not a consequence of slower growth; conversely, the rise in saving rates in several East Asian countries is not the result of their faster rates of growth. Neither can be attributed to the aggregation effects across birth-cohorts that drive the life-cycle correlation, and in both cases, households of all ages appear to have changed their saving behavior for reasons that we do not fully understand (Bosworth, Burtless and Sabelhaus, 1991; Paxson, 1996; Deaton and Paxson, 2000; Parker, 2000). Although the data on saving rates by age are weak—saving is the difference between two large numbers, both of which are hard to measure—it is clear that there is not enough life-cycle saving—and that it happens too late in life and too close to the age of dissaving—to permit changes in the rate of growth to have much of an effect on the average saving rate (Kotlikoff and Summers, 1991; Gourinchas and Parker, 2000).

I want to draw several methodological lessons from this story. Although I believe that the life-cycle mechanism has been falsified as a general proposition, the evidence should be interpreted as showing that its range of applicability is smaller than we thought, not that the mechanism is never relevant. People do save for retirement, and governments sometimes do so on their behalf, and those mechanisms have a role to play in understanding well-being over the life-cycle, as well as the supply of savings in general. But the life-cycle mechanism alone is not capable of explaining the international correlation between saving and growth. Indeed, the life-cycle model was always of doubtful relevance in poor countries, where the widespread existence of extended families reduces the need for formal life-cycle saving. The falsification of the hypothesis as a general proposition has been fruitful in stimulating new research on behavioral

6

approaches to saving in both rich and poor countries. So we are certainly not left with nothing, even if there are things that we thought we understood but no longer do. We have a more nuanced and qualified understanding of the life-cycle mechanisms, and a useful agenda for future research. Perhaps Lewis was broadly right, and we are certainly now free to consider causation from saving to growth as well as from growth to saving, and to look for other mechanisms; candidates might be habit formation by consumers, Carroll and Weil (1994), or low rates of investment in poor countries, Hsieh and Klenow (2007).

I also want to emphasize the importance of the work that is intermediate between theory and data and that was required to find specific and transparently testable predictions, like Modigliani's (1970) derivation of the link between saving, growth, and the length of work and retirement spans, and Carroll and Summers' (1995) work on age-profiles of consumption. This kind of intermediate work is not econometrics, nor is it theory as usually understood, but it is work that can only be done by those with a familiarity with the theory and an understanding of the evidence that is available or might be made available.

Structural Modeling: Understanding Commodity Prices

Many low-income countries, especially in sub-Saharan Africa, depend on primary commodities for a large share of their exports. For these countries, fluctuations in the world prices of commodities pose difficult problems of macroeconomic management. The dynamics of commodity prices are characterized by long periods of quiescence, interrupted by dramatic

7

upward flares and downward plunges. In theory, a country might save when prices are high against times when prices are low, but commodity prices are highly positively autocorrelated, so that strategies to smooth government or household expenditure would often requiring periods of saving and dissaving that are too long to be within the borrowing or lending capacity—or indeed the political patience—of the countries involved. Indeed, there have been many spectacular disasters in history. The boom in cotton prices during the American Civil War funded a huge and not easily reversible increase in state expenditure in Egypt that eventually led to the collapse of the government and to the country's takeover by Britain, (Issawi, 1966). Nor have international organizations been better at understanding and predicting commodity prices; during the 1970s and 1980s, as commodity prices collapsed, the World Bank consistently revised *upwards* its forecasts of copper and cotton price, until prices a few years ahead were predicted to be four to five times their current or subsequently realized prices, encouraging producer countries to indebt themselves to an eventually unsustainable degree (Powell, 1991; Deaton, 1999).

A coherent mechanism for commodity price fluctuations was first developed by Gustafson (1958). The model is one of speculative storage. There is an underlying process of agricultural supply and demand, primarily driven by supply shocks that are either independent and identically distributed or at most mildly autocorrelated. However, the effects of these shocks on market prices are filtered by the actions of profit-maximizing and risk-neutral speculators who use a costly storage technology to buy when commodities are relatively cheap in the expectation of selling when they are relatively dear. In normal times, when speculators are holding stocks, the price of the commodity must be expected to rise by enough to cover the speculators interest and storage costs; otherwise they would hold nothing. This is consistent with price fluctuations but, over periods when stocks are being held, price will rise on average.

However, when prices get too high—above the "stockout" price—speculators drop out of the market, and no inventory is carried forward. When that happens, there is no storage buffer to prevent a price spike in the event of a harvest failure in the next period. If a series of bad harvests occurs—as in the seven lean years of the Bible—inventories will eventually be exhausted, and prices will reach historically high levels until normal times return. In spite of the trend increase in price when inventories are held, there will always come a time when the price is too high to support speculation, and the price will collapse; the model predicts no long run trend in price.

In a series of papers, Deaton and Laroque (1992, 1996) have investigated the properties of this mechanism, and whether it fits the data for a range of commodity prices. The analysis starts from the simplest case where the shocks to the harvest are normal, independently and identically distributed; then numerically solves for the (policy) function that links the prices to the amount of inventory on hand; and finally simulates the time-series and distributions of prices. The calculated series show some of the properties of actual prices, such as the evident nonlinearity and the price flares, but the patterns of autocorrelation are wrong, and it is the autocorrelation that is the central feature of the policy problem for commodity exporting states.

Many of the standard criticisms of structural modeling apply to this work. It is difficult to disentangle the auxiliary assumptions from the central core that we want to test, the computations are time-consuming and error-prone, the substance of the problem tends to be lost in the sometimes byzantine complexity and programming of the estimation, and it is hard to get a sense of why the results are what they are. But the calculations yielded a crucial insight that could have obtained from the beginning, had we been able to see it. In normal times, when stocks are being held, price is expected to rise, and autocorrelation is high. However, when prices rise to the cutoff point at which speculators stop holding inventories, the price is expected

to stay where it is. If we plot the expected value of tomorrow's price against today's price, the graph—the autoregression function—is a straight line through the origin with slope greater than 1 up to some critical price—the stockout price—after which it is constant. We first discovered this pattern from graphs of simulations based on estimation, and then proved the result more generally. This prediction allows a transparent test of (one aspect of) the model, a prediction that is falsified on the data; there is no evidence that the actual autoregression function flattens out at high levels of prices. When commodity prices are high, a simple first-order linear autoregressive model makes better predictions than does the speculative storage model (Deaton and Laroque, 1996, Figure 1). Indeed, it is precisely because the autocorrelation is so high at all levels of prices that it is so hard for governments of commodity exporting countries or their advisors to smooth their incomes.

As in the saving and growth example, we finish up without a satisfactory model, but once again we have learned something about when the mechanism works and when it does not. Our failure points to good future directions such as exploring the role of demand for commodities, which is driven by income and thus tends to be strongly autocorrelated, something that is taken up by Dvir and Rogoff's (2009) study of the historical price of oil. One methodological lesson that I want to draw here is that the simulation and estimation of tightly specified models may reveal regularities and behaviors that are worth investigating as possible generalizations and that can lead, in favorable cases, to "acid tests" of the theory. In effect, we are using the computer to think for us. Structural estimation is useful, not only for the estimates (whose credibility is often undercut by the panoply of supporting assumptions that are required to obtain them), but for understanding the empirical predictions of the theory.

Measurement in Search Of Theory: Two Food Puzzles

My third example contrasts with the first two because it begins from observations, rather than from theory, and illustrates the interplay of mechanisms and data from a different starting point.

When many people think about global poverty, they think about hunger, about people who regularly do not have enough to eat because they are too poor to buy food, about children being abnormally susceptible to disease because they are undernourished, about those children growing up abnormally thin (wasted) and abnormally short (stunted), and about the long-term consequences for their physical and mental health, and ability to earn a living. It seems obvious enough that higher income is the solution to hunger and indeed, when we look across richer and poorer households at any moment of time, households with higher incomes get more calories, more protein, and more fat. This is true whether we compare the average household in India with the average household in the US—where we typically worry about too much consumption, not too little—or whether we compare rich and poor households within low-income countries, like India. By the same token, food consumption seems like a good proxy for living standards more broadly, and there is a long tradition of research, going back to Ernst Engel in the middle of the nineteenth century, that attempts to use food consumption to make inferences about living standards, particularly about how we might infer individual welfare from household data by adjusting incomes for the numbers and demographic composition of households. However, the data on food consumption presents several challenges to this seemingly straightforward framework.

A first puzzle arises out of the measurement of household economies of scale. As was long ago documented by Simon Kuznets (1976), virtually all household surveys, in rich and poor

countries alike, show that *total* household expenditure rises with household size while *per capita* household expenditure falls with household size. In consequence, when comparing whether larger households are poorer than smaller households on a per capita basis, or whether children are more likely to be in poverty than adults, the welfare comparison across households will depend on the degree of economies of scale. One way to tackle this comparison problem was originally suggested by Jean Drèze in personal conversation. If we compare larger with smaller households *who have the same per capita income*, the larger households will be better off in the presence of economies of scale. Better-off people spend more on food, so the increase in per capita food expenditure for the larger households should give us an idea of the income equivalent of economies of scale. But when Christina Paxson and I did the empirical work in 1998, we found that larger households actually spend *less* per head on food than do smaller households with the same per capita expenditure or income. On the face of it, better-off people spend less per person on food! This pattern holds true for a wide range of rich and poor countries. Even more strangely, the extent of the reduction was largest in the poorest countries where food is scarcest, and the slope of the overall Engel curve relationship between income and quantity of food consumed is the steepest.

Clearly, something is seriously wrong with the simple theory. Exactly what is unclear. One possible set of answers considers how household size or even household income may adjust to food scarcity, or to third factors that are correlated with food, but these approaches have not offered much of a handle on the puzzle. Note that economies of scale in food preparation, which almost certainly exist, make the puzzle worse because they lower the effective price of food and encourage its consumption. Another possible resolution is that the price elasticity of food is relatively high, so that the equivalent price reduction from the economies of scale may cause

enough substitution away from food to offset the income effect. Yet this explanation is surely implausible for people who are poor and malnourished. This example is perhaps better classed as a puzzle or a paradox than a falsification, but only because the mechanism on which it is based seems so obvious, that malnourished people will consume more food when they have more money.

A related and even starker puzzle, investigated by Deaton and Drèze (2009), is that in spite of recent startlingly high rates of growth of per capita income and per capita consumption in India, per capita *calorie* consumption has been *falling* for quarter of a century. In rural India, per capita intake fell from 2,240 calories per person per day in 1983 to 2,047 in 2004–05 while, in urban India, calorie consumption remained constant at around 2,000 per person. Although the evidence is less clear, per capita calorie consumption appears also to have fallen in China, in spite of even more rapid rates of growth (Du, Lu, Zhai and Popkin, 2002). Clark, Huberman and Lindert (1992) argue that the same probably happened in Britain during the industrial revolution from 1770 to 1850. Although the amount of poverty reduction in India is modest compared with what might be expected from its growth in national income, there have been real gains throughout the income distribution, in rural as well as urban areas. As a result, the calorie decline cannot be attributed to massive redistribution from poor to rich, nor to some generalized—but unmeasured—impoverishment among the rural poor. Most of the decline is accounted for by cereal consumption; the trend reduction in "coarse" cereals—sorghum, millet, maize—is more than 50 years old, but in recent years it has been supplemented by declines in consumption of rice and wheat. Reductions in protein match reductions in calories, and although there has been an increase in per capita consumption of fats, this has not been sufficient to increase total per

capita calorie consumption. (Indians, unlike Americans, consume less fat than is desirable for their health.)

The calorie reduction does not reflect an adequacy of nutritional status in India. Indian women are among the shortest in the world (Deaton, 2007), and 35.6 percent of adult women have a body-mass index (calculated as weight in kilograms divided by height in meters squared) below the World Health Organization minimum of 18.5, a fraction exceeded only by Eritrea, and higher than any other country in Asia or sub-Saharan Africa, where per capita income levels are much lower than India's. The rates of child wasting and stunting in India are also among the highest in the world, although there has been improvement in all of these nutritional indicators over time. One indicator of comparative rates of improvement is the rate at which adult height is increasing over time. This is about 0.56 cm per decade for men in India, but only 0.18 for women—a rate of improvement that is about half the rate in China, or the historical rate of improvement in the high-income European countries. (Interestingly, these Chinese and European rates of growth in adult height are actually lower than the contemporary rate of growth in Kerala, a state in southwestern India, where there is also no differential between men and women.) So the general rate of improvement in nutritional status in India, although real, has been relatively slow, in spite of the unprecedented growth in per capita national incomes over the period.

What mechanism could account for these outcomes? It is certainly not an increase in the relative price of food or of calories, which has declined during most of the period. I suspect, as others have earlier suspected (Rao, 2000, 2005), that the reduction in calorie intake in India, and perhaps elsewhere, comes from a reduction in physical activity levels associated with some movement out of agriculture along with the increasing mechanization of agriculture and of transportation—replacing human with mechanical power, facilitated by large improvements in

road infrastructure. It may also be associated with the improvements in health associated with a huge expansion over the last 25 years in the availability of piped water so that fewer calories are lost to diarrheal disease and to intestinal parasites.

It is not difficult to construct a simple model of this process, in which people use calories for two purposes: their own nutrition, and for heavy physical labor, fueled mostly by cereals. Laborers have different abilities to work, through health or strength, and those who work harder use more calories for work, and earn more money. At a moment in time, in a single labor market with a common wage, the relationship from strength to income through work and calorie consumption will induce a positive relationship between income and calorie consumption, the familiar positively sloped calorie Engel curve. But as real wages rise, either across space or over time, physical labor will decline and calorie consumption for work will fall, even as the calorie consumption for improving nutritional status is increasing but still inadequate. Over time, as people become rich enough so that manual labor makes minimal nutritional demands, there will be no further fall from this cause, and increases in income will be associated with increases in calories and nutritional status. Eventually, among the richest inhabitants of poor countries, as among the rich inhabitants of rich countries, increases in income will lead to increased intake and eventually to "over"-nutrition and its associated health risks.

That said, there is not much direct evidence for this mechanism, and it is controversial, among those who interpret the calorie decline as evidence of impoverishment (for example, Patnaik, 2007). In favor of the theory, the time series decline in caloric consumption is matched by a negative association between per capita calorie consumption and per capita total expenditure at the state level so that, for example, the lowest per capita calorie consumption is in relatively well-off and well-nourished states like Kerala. Children have better anthropometrics

when they live in parts of India where per capita calorie consumption is lower. There is also a pronounced decline over time in reports of hunger, and there is a *positive* correlation between hunger and per capita calorie consumption. All of this is indirect evidence, and we have no direct measure of the fall in physical effort in agriculture that would be required to uphold this argument, given that the movement out of agriculture has been small. Indeed, in the latest round of the Gallup World Poll, while 80 percent of rural Indians said that their work involved a lot of physical effort, fully 41 percent said that their work involved *more* effort than five years previously. Nor is there any obvious link to the first puzzle, in the form of evidence that larger households can achieve the same per capita income with less physical labor by each person, for example through economies of scale in production that outweigh the economies of scale in consumption.

This work is at a relatively early stage. We have anomalous observations that are replicated in a number of settings and that are hard to explain. Alternative mechanisms are only just being developed, and are some way from delivering sharp non-obvious predictions that are not part of what we were trying to explain in the first place. Indeed, in this case we are trying to work back from the data to a theory, not the other way round. But the interplay between theory and data is the same, and it makes no difference where we join the cycle.

Conclusions: What Do We Learn About Methods?

The general methodology that fits all of these cases is the hypothetico-deductive method; hypotheses are formed, and predictions deduced that can, at least in principle, be falsified on the data. In its strict form, as formulated by Karl Popper (1934), the method tests predictions, and

16

looks for falsification; there is no possibility of confirmation, and falsification is the only way that we learn. In the weaker or positivist form, we formulate hypotheses and derive predictions, what I have referred to here as "acid tests," that are hard to explain if the theory is not true, so that we seem to learn at least *something* when the predictions are confirmed. The best example here of an acid test is the precisely-specified and parameter-free relationship between growth and saving rates that comes out of the stripped-down life-cycle model, and whose match with the existing evidence was an uncanny piece of wizardry. Yet as this case illustrates, such corroborations are strictly provisional and are subject to revision in the light of new evidence, or of a deeper analysis of the theory. Any fact that is consistent with one theory is consistent with an infinite number of theories. We can also learn from falsification about which parts of the theory are wrong, which supplementary assumptions need to be modified, or under what circumstances the theory does not hold. We are making progress, learning more about mechanisms as we go.

Nothing in this process is unfamiliar to applied economists. We argue in this style all the time, arguing that a regression coefficient or a stylized fact is consistent with one model and inconsistent with another; there is nothing difficult or unfamiliar about this way of working though spectacular examples such as the life-cycle model, that shift whole areas of enquiry, are necessarily rare. While individual findings are often easily explained by other theories, or contradictions easily removed by changing auxiliary assumptions, I do not think there can be real doubt that the profession as a whole makes progress. The endless cycling of fashion in development thought among policymakers is more an indication of the bankruptcy of the aid enterprise than of the underlying scientific understanding. None of this means that anything goes, or that good methods have no payoff. I am sure we make most progress when theory and

17

empirical work are closely articulated, not necessarily in the same person, but at least when different people with different skills read and talk to one another. Good tests require deep understanding of models, and the ability to manipulate them into delivering predictions that are not obvious and that are specific enough to the model to be informative about it. At the same time, good theories, or good modifications of existing theories, require theorists who are familiar with and pay attention to historical and empirical evidence. My main concern with current practice in development economics is that these links are weak and that much empirical work makes no attempt to investigate mechanisms.

Instrumental variables and randomized trials can play a role in uncovering the mechanisms of development. Indeed, an instrumental variable strategy is often readily interpretable in the hypothetico-deductive framework; if the instrument works as designed, it should affect the outcome variable, often in a surprising and interesting way, an "acid test" in its own right. Randomized trials have a powerful ability to isolate one mechanism from another; in particular, an experiment will often allow us to short-circuit the often difficult process of developing theoretical mechanisms to the point where they can be convincingly tested on non-experimental data. At the same time, the routine use of instrumental variable methods and of randomized controlled trials for project evaluation is often uninformative about *why* the results are what they are, and in such cases, nothing is learned about mechanisms that can be applied elsewhere.

I want to end with a final note on measurement, especially descriptive measurement motivated by low-level theory. Few students of economic development will learn much about measurement in their graduate courses, which is a pity. The basic facts of economic development, such as the growth rates of GDP, come from measures that ought to be much more deeply debated than is the case. GDP is a poor measure of welfare to start with, and many

important developing countries, China and India being only the leading examples, have weak statistical systems that are most likely overstating their rates of growth (Deaton, 2005). The microeconomic data on income and consumption from the surveys are also weak, and in many countries show growth rates of mean consumption that are considerably lower than the corresponding estimates from the national accounts, which raises a huge barrier to understanding the relationship between economic growth and poverty reduction. Back to the *Sudd*.

The macroeconomic data in the Penn World Tables and the World Development Indicators rest on the calculation of price indexes based on comparing prices in different countries, a task that is accomplished on an irregular basis by the International Comparison Program (ICP), whose latest round gathered information for 2005. These data underpin a large fraction of what development economists think they know about economic growth, yet the assumptions that go into these price indexes are not widely debated or even understood. In particular, we do not know how to make comparisons between countries whose patterns of relative prices and consumption are very different from one another. For example, the ratio of US to Indian per capita GDP is 1.6 times larger in Indian prices than it is in American prices, and while the ICP resolves that difference, essentially by splitting the difference, it is unclear why such an average is economically interesting. Beyond that, the ICP protocols require the pricing of closely matched goods in disparate countries, e.g. Kellogg's cornflakes or 2003 Bordeaux *superieur* in Kenya, Senegal, Cameroon, Japan, and Britain. Again it is unclear whether these are the right comparisons, and whether matching precisely specified goods holds quality constant in a way that yields useful welfare comparisons. Yet again, when comparing two countries, the ICP constructs price indexes by weighting individual prices by what is effectively the average of the budget shares in both, so that a very expensive good that is consumed by almost no one in a poor

country can have a large downward effect on the measured per capita consumption or GDP in international currency. And these are only a few of the issues that pose unsolved practical or theoretical problems, see Deaton and Heston (2010), Deaton (2010). Yet our measures of global poverty and inequality, as well as of progress over time in growth and poverty reduction, depend on the as yet elusive solution to the problems of making such comparisons. The same is true for at least some of the literature that uses the Penn World Table for running cross country regressions, Johnson, Larson, Papageorgiou, and Subramanian (2009). These problems cannot be resolved without a close understanding of both theory and data.

Acknowledgements

References

Atkinson, Anthony B., 1969, "On the timescale of economics models: how long is the long run?" *Review of Economic Studies*, 36, 137–52.

Bosworth, Barry, Gary Burtless, and John Sabelhaus, 1991, "The decline in saving: evidence from household surveys," *Brookings Papers on Economic Activity*, 183–241.

Carroll, Christopher D. and Lawrence H. Summers, 1991, "Consumption growth parallels income growth: some new evidence," in B. Douglas Bernheim and John B. Shoven, eds, *National saving and economic performance*, Chicago, Il. Chicago University Press.

Carroll, Christopher D., and David N. Weil, 1994, "Saving and growth: a reinterpretation," *Carnegie-Rochester Conference Series on Public Policy*, 40: 133–92.

Cartwright, Nancy, 2007, "Are RCTs the gold standard?" *BioSocieties*, 2, 11–20.

Clark, Gregory, Michael Huberman, and Peter Lindert, 1992, "A British food puzzle 1770–1850," *The Economic History Review*, 48:2, 215–37.

Deaton, Angus, 1997, *Understanding consumption*, Oxford. Clarendon Press.

Deaton, Angus, 1999, "Commodity prices and growth in Africa," *Journal of Economic Perspectives*, 13:3, 23–40.

Deaton, Angus, 2005, "Measuring poverty in a growing world (or measuring growth in a poor world), *Review of Economics and Statistics*, 87:1, 1–19.

Deaton, Angus, 2007, "Height, health, and development," *Proceedings of the National Academies of Science*, 104:33, (August 14), 13232–7.

Deaton, Angus, 2009, "Instruments of development: randomization in the tropics, and the search for the elusive keys to economic development," Keynes Lecture, British Academy, *Proceedings of the British* Academy, forthcoming and available at http://www.princeton.edu/~deaton/downloads/Instruments%20of%20development%20v1 d_mar09_all.pdf

Deaton, Angus, 2010, "Price indexes, inequality, and the measurement of world poverty," *American Economic Review*, 100:1, forthcoming.

Deaton, Angus and Jean Drèze, 2009, "Food and nutrition in India: facts and interpretations," *Economic and Political Weekly*, 44:7 (February 14), 42–65.

Deaton, Angus, and Alan Heston, 2009, "Understanding PPPs and PPP-based national accounts," *American Economic Journal Macro*, forthcoming. Available at http://www.princeton.edu/~deaton/downloads/Understanding_PPPs_complete_with_abstract_14nov08.pdf

Deaton, Angus and Guy Laroque, 1992, "On the behavior of commodity prices," *Review of Economic Studies*, 59:1, 1–23.

Deaton, Angus and Guy Laroque, 1996, "Competitive storage and commodity price dynamics," *Journal of Political Economy*, 104:5, 896–923.

Deaton, Angus, and Christina Paxson, 1998, "Economies of scale, household size, and the demand for food," *Journal of Political Economy*, 106:5, 897–930.

Deaton, Angus, and Christina Paxson, 2000, "Growth and saving among individuals and households," *Review of Economics and Statistics*, 82:2, 212–25

Du, Shufa, Bing Lu, Fenying Zhai, and Barry M. Popkin, 2002, "A new stage in the nutrition transition in China," *Public Health Nutrition*, 5:1A, 169–74.

Dvir, Eyal, and Kenneth S. Rogoff, 2009, "Three epochs of oil," NBER Working Paper No. 14927.

Gourinchas, Pierre-Olivier, and Jonathan A. Parker, 2002, "Consumption over the life cycle," *Econometrica*, 70:1, 47–89.

Gustafson, R. L, 1958, *Carryover levels for grains*, Technical Bulletin 1178, US Department of Agriculture.

Houthakker, Hendrik S., 1965, "On some determinants of saving in developed and underdeveloped countries," in Robinson, E.A.G., editor, *Problems in Economic Development*. London, Macmillan.

Hsieh, Chang-tai, and Peter J. Klenow, 2007, "Relative prices and relative prosperity," *American Economic Review*, 97:3, 562–85.

Issawi, Charles, ed., 1966, *The economic history of the middle east 1800–1914: a book of readings*, Chicago. Chicago University Press.

Johnson, Simon, William Larson, Chris Papageorgiou, and Arvind Subramanian, 2009, "Is newer better? Penn World Table revisions and their impact on growth estimates," Cambridge, MA. NBER Working Paper No. 15455 (Oct.)

Kotlikoff, Laurence, and Lawrence H. Summers, 1981, "The role of intergenerational transfers in aggregate capital accumulation," *Journal of Political Economy*, 89:4, 706–32.

Kuznets, Simon, 1976, "Demographic aspects of the size distribution of income: an exploratory essay," *Economic Development and Cultural Change*, 25:1, 1–94.

Lewis, Arthur, 1954, "Economic development with unlimited supplies of labor," *Manchester School of Economics and Social Studies*, 22, 139–91.

Modigliani, Franco, and Richard H Brumberg, 1954a, "Utility analysis and the consumption function: an interpretation of the cross-section data," in Kenneth K. Kurihara, ed, *Post-Keynesian Economics*, New Brunswick. Rutgers University Press, 388–436.

Modigliani, Franco, and Richard H. Brumberg, 1954b, [1990], "Utility analysis and aggregate consumption functions: an attempt at integration," in Andrew Abel, ed., *The collected papers of Franco Modigliani: Volume 2: The life-cycle theory of saving*, Cambridge, MA., MIT Press, 128–97.

Modigliani, Franco, 1966, "The life-cycle hypothesis of saving, the demand for wealth, and the supply of capital, *Social Research*, 33:2, 160–217.

Parker, Jonathan, 2000, "Spendthrift in America: two decades of decline in the US saving rate," in Ben Bernanke and Julio Rotemberg, eds, *NBER Macroeconomics Annual 1999*, 14, 317–70.

Paxson, Christina, 1996, "Saving and growth: evidence from micro data," *European Economic Review*, 40:2, 255–88.

Patnaik, Utsa, 2007, "Neoliberalism and rural poverty in India," *Economic and Political Weekly*, 42:30, (July 28), 3132–50.

Powell, Andrew, 1991, "Options to alleviate the costs of uncertainty and stability: a case study of Zambia," in Louis Phlips, ed., *Commodities, futures, and financial markets*, Dordect, Kluwer.

Rao, C.H. Hanumantha, 2000, "Declining demand for foodgrains in rural India: causes and implications," *Economic and Political Weekly*, 35:4 (January 22), 201–6.

Rao, C. H. Hanumantha, 2005, *Agriculture, food security, poverty, and environment: essays on post-reform India*, New Delhi. Oxford University Press.

Schultze, Charles and Christopher Mackie, 2002, *At what price? Conceptualizing and measuring cost-of-living and price indexes*, Washington, DC. National Research Council.

Solow, Robert M., 1956, "A contribution to the theory of economic growth," *Quarterly Journal of Economics*, 70:1, 65–94.