# Measuring and Understanding Behavior, Welfare, and Poverty[†]

*By* Angus Deaton*

The work cited by the Nobel committee spans many years, covers areas of economics that are not always grouped together, and involves many different collaborators. Yet, like the committee, I believe that the work has an underlying unity. It concerns well-being, what was once called welfare, and uses market and survey data to measure the behavior of individuals and groups and to make inferences about well-being. Often, little more than counting is involved, as in the estimation of the fraction of the population whose spending is below a cutoff, or the calculation of the fraction of newborn children who die before their first birthday. Measurement, even without understanding of mechanisms, can be of great importance in and of itself—policy change is frequently based on it—and is necessary if not sufficient for any reasoned assessment of policies, including the many that are advocated for the reduction of national or global poverty. We are wise to remember the importance of good data, and not to neglect the challenges that measurement continuously poses.

More ambitiously, estimation of behavior can elucidate mechanisms and causes, and help to make predictions about the effects of policy, providing a guide to policy improvement. Indeed, the analysis of consumer behavior with a view to measuring well-being has long been a basic task of economists. Although, as directed by the Committee, this article is about my own work, I shall try to set the work in the context to which it belongs, allowing myself liberal use of hindsight.

The link between measurement, behavior and policy is a running theme. So is the necessity of telling a coherent story of what we observe. Another key idea is the distinction between individuals and aggregates, what the committee refers to as "consumption, great and small." Aggregation needs to be seen, not as a nuisance, but as a hallmark of seriousness, as well as a source of hypotheses and understanding. The link between behavior and well-being, when it holds at all, holds for individuals, not for aggregates. While we often must focus on aggregates for macroeconomic policy, it is impossible to think coherently about national well-being while ignoring inequality and poverty, neither of which is visible in aggregate data. Indeed, and except in exceptional cases, macroeconomic aggregates themselves depend on distribution. These arguments are much more widely accepted today than they were 30 years ago.

Much of what follows is based on the traditional (in economics) premise that people know what is good for them and act in their own interest. People reveal (something about) their preferences in their behavior, which allows us to infer (something about) well-being from the choices that they make. The validity of revealed preference is currently being robustly challenged by behavioral economics, though no new general operating link between behavior and well-being yet exists; this is surely a key task for economics in the years ahead. I here stick to the traditional position, if only because of the many successes that approach has brought.

I start with household surveys, how they are used to document living standards, inequality, and poverty and, beyond that, to understand behavior. From there, it is a short step, in Section II, to demand analysis, which looks at how consumption patterns respond to prices and incomes. Understanding such effects is necessary for the design of tax and pricing policies and is useful for much else. Much of the early work on demand was concerned with single-period models, but the same set of methods were gradual extended to help understand the dynamics of behavior. Sections IIA and IIB consider those in turn.

## I. Using Household Surveys for Measurement and for Analysis

### A. *Documenting the Lives of the Poor*

The documentation of how people live, how much they spend, and on what, has long been used as a political tool, to make visible the living conditions of the poor to those in power, to shock, and to agitate for reform. According to Stigler (1954), the first surveys were those of Davies (1795) and Eden (1797) in England. The wave of social unrest in Europe in the 1840s brought a wave of budget studies, including Ducpetieaux (1855) which was a predecessor of Engel's (1857) famous study. Engel was also influenced by Quételet's arguments for the statistical analysis of social data, including Quételet's concept of *l'homme moyen*, an early avatar of the representative agent. At around the same time, the pioneers of social epidemiology were making parallel inquiries into the health and living standards of the working classes, for example Villermé (1830) on the geography of mortality and poverty in Paris, and perhaps most famously, Engels (1845), who documented mortality differences and living standards in Manchester and argued that the industrial revolution had immiserated the working classes, starting off a debate about the effects of the industrial revolution on well-being that continues to this day.

The descendants of these studies are today's randomly selected, stratified, and clustered household surveys that are run regularly by most statistical offices in the world. The statistical theory of random surveys was developed only after Neyman (1934), with important practical contributions in India, see Mahalanobis (1946), who established one of the first national sample surveys, see Seng (1951) and Smith (1976). Today's household surveys typically collect information on household incomes and/or (often detailed) expenditures, as well as demographic, geographical, and other characteristics of household members. Their official purpose is often to collect weights for consumer price indexes, but they are also used to calculate national and global estimates of poverty and inequality. The United States Agency for International Development funds the system of internationally comparable

Demographic and Health Surveys (www.dhsprogram.com), each of which documents the health, anthropometric status, and mortality rates of infants and children. Collectively, they provide much of the infrastructure for comparisons of health between countries. As was the case in the eighteenth and nineteenth centuries, these estimates of poverty and of mortality are used today by national and international organizations, aid agencies, and NGOs to enliven the consciences of the privileged of the world and to agitate for pro-poor policies.

Within countries both rich and poor, household surveys are the basis for estimates of poverty rates, of income inequality, and of changes in real wages across percentiles of the distribution. Sometimes, this information is sufficient to evaluate an approximation to the welfare effects of policy changes; a familiar example is the calculation of the compensating variation of an increase in a tobacco or gas tax by examining expenditures on tobacco or gas in a household survey. Better approximations require estimation of the response of purchases to prices, on which more below. Another example comes from countries where staple foods are both produced and consumed, with some households (farmers) producing, and others consuming. If, for example, the country is a net exporter, the imposition of an export tax will harm net producers and help net consumers and for each the welfare equivalent variation can be approximated by the value of *net* consumption, which can be read off from a household survey. In both cases, the distributional effects of the policy change can be read off from a national survey (see Deaton 1989a for the example of rice in Thailand).

Among the most difficult and pressing problems with household surveys is the quality of the data; in some cases, the problems are severe enough to threaten even the most basic understanding of growth, poverty, and inequality. India is perhaps the most important illustration. Per capita consumption estimated from household surveys has long been less than per capita consumption estimated in the national accounts statistics (NAS), even when adjustments are made for conceptual differences in coverage, for example surveys do not collect data on imputations for rents or financial intermediation, nor on the cost of publicly provided education and healthcare. This discrepancy has long concerned Indian economists, e.g., Minhas (1988), and has steadily worsened over time; in 1972–1973 the survey estimate of consumption fell short of the NAS estimate by five percent, while in 2009–2010, only one-half of national accounts consumption showed up in the surveys, rising to two-thirds after adjustment for differences in definition (Government of India 2014, Tables 3.2 and 3.3). That per capita consumption as measured in the NAS grows more rapidly than per capita consumption as measured in the surveys happens not only in India, but in many countries including, over some periods, the United States (Deaton 2005). While the sources of the discrepancies are largely obscure—itself a testament to the lack of attention devoted to the topic by national and international agencies—it is clear that the national accounts cannot be held blameless; again in India, there is likely exaggeration of the rate of growth in the national accounts—governments whose legitimacy depends on high growth tend to be unenthusiastic about any revision that decreases measured growth—and, on the other side, surveys cannot easily penetrate gated communities, nor capture the increasing share of expenditure outside of the household and unobserved by the single respondent that surveys rely upon.

These discrepancies block any coherent account of poverty, inequality, and growth. The historically high rates of growth of (national accounts) consumption in India since the 1980s would appear to warrant a much larger reduction in poverty than is shown by the surveys. When growth brings little poverty reduction, the usual inference would be that inequality is rising, which may well be true and likely accounts for some of the difference, but the main reason today is not a failure of trickle-down, but a failure of measurement. Those who choose to believe the national accounts, and disbelieve the surveys—except to note that they show limited increases in inequality—believe that the rate of poverty reduction is grossly understated in India and, beyond that, in the world as a whole given that India accounted for one-quarter of global poverty in 2012. By contrast, those who choose to believe the surveys believe that a large fraction of India's spectacular growth is illusory. The most basic economic data, the rate of growth, its distribution, and who is gaining and who is being left behind are inaccessible, so that people are free to choose their facts according to their political prejudices, unconstrained by reality, and a reasoned political debate about these central issues becomes impossible.

The measurement of global poverty has long been carried out by the World Bank, which, in its latest estimates that document 1984 to 2012 (World Bank 2016), uses information from more than 1,000 household surveys from 131 development countries and 21 high-income countries; 43 of the surveys for the 2012 estimates are from sub-Saharan Africa, which is the region where poverty is most prevalent. Yet household surveys in Africa are often weak, often outdated (the 2012 estimates use surveys as old as 2003), are sometimes inconsistent over time within countries, have nonmatching definitions—different reporting periods, or are surveyed at different times of year, either over time or over countries—so that it is extremely difficult to assess progress over time, or to make comparisons of poverty or inequality between countries.

In many countries and in the world as a whole, large numbers of people live in the vicinity of the national or global poverty line. In consequence, small changes in the position of the line, for example through the choice of price index for updating, or small changes in survey practice, can have dramatic effects on the number of people counted as poor. A spectacular example comes from India, where the sample survey organization ran a large randomized controlled trial on the effects of different lengths of the reporting (recall) period. People in one arm reported their consumption over the last 30 days, in the other, over the last 7 days (Visaria 2000). Because the reported flow rate is higher with the shorter reporting period, switching from 30 to 7 days reduced the national poverty rate in 1998 *by almost one-half* and "removed" 175 million people from poverty (Deaton 2001). Statistical solutions to poverty are easier than real solutions.

Education and health are two important categories that are inconsistently treated in different countries because private provision is included in the surveys while public provision is not, so that poverty measurement is hostage to local arrangements, which vary across countries; in sub-Saharan Africa, for example, the private share of healthcare spending ranges from 27 percent to 74 percent. An effective improvement in state provision of either education or health could do much to improve the lives of the poor, but if it crowds out private expenditures—which would be desirable in many settings where the quality of private care is poor—measured poverty would increase.

It is perhaps tempting to abandon measures of material well-being and move to other measures, such as anthropometrics, or mortality, and I believe that the material poverty measures have been assigned too much weight, given their inherent unreliability. Yet, while it is surely important to emphasize other aspects of poverty, particularly health and education, those other dimensions are not substitutes for measures of material deprivation. While different deprivations are often correlated, the extent of the correlations are different in different places and times; for example, Africans tend to do better than Indians on anthropometrics, but worse on income and on mortality, and in China, when economic growth brought massive poverty reduction after the mid-1970s, infant and child mortality, which had been falling rapidly, greatly slowed (Deaton 2013, p. 115). In the end, if we want to know about material poverty, we will have to measure it better.

Household surveys from around the world are the underlying source for measures of global poverty and global inequality, where the former is defined as counts of those living below some uniform cutoff. Inequality can refer to dispersion in income (or consumption) across all of the citizens of the world, or to dispersion of per capita incomes across nations, or to inequality of per capita incomes of nations weighted by population (see Milanovic 2007); the first can only be estimated with household survey data. These global poverty and inequality measures require that the data from each country be first converted to a common currency using purchasing power parity exchange rates. These PPPs are multilateral international price indexes which, although widely used by the research community through the Penn World Table and other databases, have properties that are not always well understood.

The International Comparison Project (ICP), started in the 1970s and continuing and developing today (Kravis, Heston, and Summers 1978); Summers and Heston 1991), is one of the outstanding intellectual achievements of economic measurement, and one that has never been recognized by the Nobel Committee (see the comments by Samuelson 2004). A primer on current practice is Deaton and Heston (2010); what is not widely realized is that there are many unresolved theoretical and practical questions, and that successive rounds of estimates have often been disconcertingly incompatible. The ICP produces benchmark estimates of PPPs in each round, most recently in 1985, 1993/1994, 2005, and 2011. Between rounds, PPPs are updated using domestic price indexes so that, for example, the PPP for consumption for India relative to the United States, is the benchmark PPP between India and the United States updated by the relative rates of increase in CPIs of the two countries since the benchmark. While this is an intuitive and sensible procedure, it is at best an approximation; even under ideal conditions, changes in multilateral indexes over time, which use weights from all countries, will not match the relative changes in national price indexes, which use weights from one country only (see, e.g., McCarthy 2013).

The problem here is not the existence of approximation errors, but that in the three most recent rounds, these changes have been large enough so as to seriously reconfigure the economic geography of the world. In both 1993 to 1994 and again in 2005, the consumption and GDPs of poor countries were revised downward relative to those of rich countries; recall that the United States is always the (arbitrarily chosen) numeraire, so the revisions are to relative, not to absolute standings. For example, ICP 2005 revised downward the GDPs of India and China by 36 and

39 percent relative to the United States; for some African countries the downward revision was much larger. In 2011, by contrast, the *price levels* of poorer countries were on average revised downward, so that per capita consumption and GDP levels were revised upward, offsetting at least some of the 2005 revisions, although not consistently so. These large and unpredictable revisions wreak havoc with attempts to measure global poverty and global inequality: see Deaton (2010) on the 2005 revisions, and on the difficulties that they caused for the World Bank poverty estimates, and Deaton and Aten (2014) and Inklaar and Rao (2014) for attempts to understand the 2011 revision. Ferreira et al. (2015) adapt the World Bank's poverty count in the light of the new numbers.

The reasons for the PPP revisions are not fully understood. Given that each round is done anew, there are often substantial methodological revisions; these are the most likely causes of change, but there is no well-defined procedure for measuring the effects of any particular revision. There is also variation that comes from the sampling of prices and from the choice of which goods to sample, from the choice of index formula, and from the structure of relative prices in any given year (Deaton and Dupriez 2011). One gauge of uncertainty is the ratio of Laspeyres to Paasche indexes for pairs of countries; in ICP 2005, these ratios were 9.6 and 5.1 for Tajikistan and Kyrgyzstan relative to the United States, and are 1.7 and 1.6 for important countries like India or China. (An unfortunate aspect of multilateral price indexes is that weak data in one place, such as Tajikistan, can in principle affect *all* PPP comparisons including, for example, those between the United States and China, or between France and Senegal.)

More troubling still are the conceptual questions. On the one hand, we need to compare like with like, using only goods and services that are close to identical in different countries. On the other hand, we also wish to capture what people actually spend, so that we want to use goods and services that are widely consumed and representative of actual purchases. These two requirements often stand in sharp opposition; in the extreme case where consumption bundles have nothing in common, there is no basis for comparisons of living standards. We need to be more humble about what PPPs can do, more cautious in using them in analysis—especially when very different countries are included—and more skeptical about the measures that depend on them, including international comparisons of GDP and consumption as well as calculations of global poverty and global inequality.

Even given a good set of PPP exchange rates, there are further hurdles in the way of calculating global poverty. One is how to set a global poverty line that can be used across the world, from Chad to Chile, from Colombo to Canberra. In the past, the World Bank has used the national poverty lines of the poorest countries, converted to US purchasing power equivalent, and averaged to give a global line. The idea is to aim for a destitution level of income that can serve as a cutoff for absolute poverty, and this is the origin of the famous $1-a-day line (World Bank 1990). Yet it is not always clear where those national lines come from, or what sort of intellectual or political legitimacy they should be accorded. Beyond that, those national lines should arguably be converted at PPPs that are tailored to the spending patterns of poor people, though this makes much less difference than might at first be thought (Deaton and Dupriez 2011). Paradoxes can arise: for example, India's high growth gave it an income that disqualified it from the poorest group when the

lines were reset after ICP 2005 (Chen and Ravallion 2010). However, it turns out that India's national line is lower than its national income would predict, so that when its national line was dropped from the average that defines the global line, the global line increased. At this new, higher, line, the world, including India, was estimated to be poorer; in effect, India became poorer because India had become richer (Deaton 2010). While such paradoxes are clearly undesirable, it is unclear how to avoid them, and new approaches need to be developed, perhaps using the new linear-programming methods being developed by the economic historian Allen (2016). There is also an argument for the multidimensional poverty indexes developed by Alkire and her coauthors (Alkire et al. 2015), in which material destitution is not given all of the weight.

In the meantime, both the World Bank and the US government are committed to eliminating global poverty by 2030, or at least to reducing it to below 3 percent of the world's population. Placing such a difficult to measure object at the center of international development policy seems ill advised, though, as always with such global and cosmopolitan measures, it is unclear who is actually responsible for meeting the target, or how (or indeed whether) anyone will be held accountable.

## B. *Surveys for Understanding Behavior and Welfare*

The behavioral analysis of household surveys with a view to welfare-improving policy often focuses on food, and goes back to Engel and his famous law, that the poorer the household, the larger share of its outlay must be expended for the procurement of food. Engel (1895) went far beyond the original law, and claimed that the food share is itself an indicator of family well-being independent of family composition. He used this identification assumption to measure the "costs" of children, a topic of continuing policy importance given that various public benefits, and indeed private settlements, such as those associated with divorce, are typically conditioned on family structure. If $w_f(x, z)$ is the budget share of food for a family with total expenditure $x$ and household composition $z$, Engel's assumption allows us to compare any household of interest, $h$ with a "reference" household, 0. If the household of interest has structure $z^h$, we can calculate how much $x$ it needs to be as well off as the reference household by finding that $x$ that gives it the same food share as the reference household, i.e., by solving $w_f(x^h, z^h) = w_f(x^o, z^o)$. The ratio $x^h/x^o$ is known as the "equivalence scale," so that if, for example, the reference household has two adults, the equivalence scale might be three for a household with two adults and two children; each child costs one-half of an adult.

That this beguiling procedure makes no sense was long ago noted by Nicholson (1976); if children consume mostly food, a fully compensated family with more children will still have a higher food share than one with fewer children and so would be overcompensated by Engel's procedure: see also Deaton and Muellbauer (1986). Even so, Engel's method is still used, perhaps because of its simplicity, perhaps because of a misunderstanding that Engel's Law implies Engel's later assertion, or perhaps because of the seemingly attractive idea that if people behave identically, they must be equally well off. For example, that equal food shares imply equal welfare is sometimes used to calculate the welfare effects of changes in *any* background circumstance $z$, for example changes in unmeasured quality of consumer goods, so

as to correct biases in consumer price indexes as suggested by Hamilton (2001) and Costa (2001). As with Engels' original procedure, these methods require additional assumptions to be valid; as noted by Pollak and Wales (1979), an (indirect) utility function of the form $\psi(x, p, z)$ for prices $p$, generates identical demands functions for each good in terms of $p$ and $z$, as does the utility function $F[\psi, (x, p, z), z]$ where $F$ is monotone increasing in its first argument. Yet the two utility functions give different levels of utility; circumstances $z$ can affect well-being without affecting observable market behavior—public goods whose costs are compensated could be an example—so that we cannot get from behavior to well-being without identification rules, essentially an exclusion restriction that $z$ affects well-being only through its observable effects on behavior. It is not clear how such exclusion restrictions can be justified, and we know from Nicholson's discussion of Engel's own example that the exclusion cannot always be right. Direct observation of utility through happiness surveys would resolve the conundrum, provided that such observations do indeed correspond to standard concepts of utility (McFadden 2014).

Household surveys continue to provide insights and to pose puzzles, particularly about food, which has been a consistent focus since the beginning. Here are some.

There is a large body of evidence, primarily from differential mortality, that suggests that, in at least some of the countries of the world, girls are discriminated against in favor of boys. Sen's (1989) documentation of "missing women" is the most famous. We might reasonably expect to see evidence of such discrimination in household surveys. The obvious place to look is expenditure on food, but surveys rarely attempt to find out who eats what. However, there is an indirect approach. Long ago, Rothbarth (1943) suggested that we might look at goods that were bought only by adults, tobacco and alcohol being the most obvious, and that expenditure on such items should indicate the expenditure needs of children. Given that children do not bring additional resources at birth, space needs to be made for their needs in the family budget, and so we might expect expenditure on adult goods to fall. In Deaton (1989b), I suggested that, in the presence of discrimination against girls, Rothbarth's method should reveal that parents make more space in the budget for boys than for girls, so that, to take a concrete example, controlling for total expenditure, Indian parents of boys might spend less on their *bidis* than Indian parents of girls. Yet I found no difference, indeed relatively precisely so, a result that has been replicated in a number of settings: see Deaton (1997). It might well be that discrimination lies, not in food provision, but in the provision of medical care, for example—in countries with dowries, girls are long-term liabilities and boys long-term assets—but the method itself could (as usual) be challenged on various econometric grounds.

Another food puzzle is related to economies of scale, the idea that two people can live more cheaply together than apart, so that members of larger households with the same per capita resources should be better off. Once again, it would seem that this idea, if true, should leave traces in the survey data. If we compare a larger with a smaller household at the same level of per capita income or total expenditure, then the presence of economies of scale implies that the larger household has been over-compensated, is better off, and thus should spend more per head on normal goods, such as food. Yet when Deaton and Paxson (1998) investigated using data from the United Kingdom, the United States, France, Taiwan, Thailand, Pakistan, and African households in South Africa, we found exactly the opposite, that per capita

food expenditures are *lower* in the supposedly overcompensated larger households, and that the difference is largest in the poorest countries, South Africa, Thailand, and Pakistan, where the need to spend more on food would seem to be the most urgent. There are a number of possible explanations, none very convincing, though there is a possible connection with another food puzzle about growth and calorie consumption.

In spite of the historically unprecedented rates of growth in India since 1980, in spite of upward-sloping calorie Engel curves, and in spite of its near record levels of child malnutrition, per capita consumption of calories and protein has been *falling* (Deaton and Drèze 2009). At the same time, anthropometric measures show that around one-half of India's children are severely malnourished; Indian women do not get enough to eat when they are pregnant (Coffey 2015), and Indian adults are among the shortest in the world. Hence, it seems obvious that with rapid economic growth, upward-sloping calorie Engel curves, and severe nutritional need, people should be eating more calories, not less. Once again, the puzzle is unresolved. Perhaps the best story is that heavy manual labor is declining along with rising living standards, so that the need for calories as fuel is diminishing, even while people remain malnourished; it is straightforward to build a model in which utility depends positively on consumption and negatively on effort, for which calories are a direct and necessary input. In an intermediate phase of development, during which people are getting better off and the need for heavy manual work is falling, the demand for calories can decline, at least temporarily. Eventually, once people are sedentary, calories stop being fuel, will yield positive net marginal utility, and consumption will rise. This account finds some support from the fact that the most rapid declines are in fuel-like commodities, such as grains and particularly the less-valued "coarse" cereals in India such as sorghum, millet, and maize, and also from the fact that India's more developed states are those with the lowest per capita calorie consumption. Manual labor may also help explain the economies of scale puzzle if there are overheads of labor—for example, fetching fuel or water, or farm tasks—that can be shared by additional household members.

In rich countries where transport is good, there is little spatial price variation in most commodities so that household surveys have limited or no ability to analyze how spending patterns change with prices. Yet price responses are required for a wide range of policy problems, and there is a literature that searches for assumptions that will allow bricks to be made without straw. One early attempt was by Pigou (1910), who noted that if preferences are additive (utility is a monotone increasing transformation of an explicitly additive function of single commodity subutilities), price elasticities are approximately proportional to income elasticities: see Deaton (1974a). Given this, only one additional piece of information—one price elasticity, or two cross-sectional surveys at different times with different prices—is required for budget surveys to identify price as well as income responses. Of course, additivity is a strong assumption, even for broad groups of goods, and it can be avoided if we have some other source of price variation. One such is the spatial variation in prices that exists in countries where transport is expensive. Many such countries, including India, collect data, not only on each household's expenditures, but also on the physical quantities purchased, at least for those goods, like many foods or fuels, where quantities are readily defined. The ratio of expenditure to quantity gives a unit

value for each good for each household. These unit values contain a quality component as well as a price component—indeed in their classic treatment of household surveys, Prais and Houthakker (1956) used the fact that unit values are higher for better off households to indicate welfare and to calculate equivalence scales—but, with a suitable theory of quality choice, price, and quality can be disentangled and price effects estimated (Deaton 1988, 1997). An important finding from such studies, at least for countries like India and Indonesia, is that the absolute values of the estimated own price elasticities tend to be (absolutely) large relative to those obtained from time series data. This is what might be expected if complementary capital goods—including possibly even tastes (Atkin 2013)—adapt slowly to prices, but it also raises issues about the relevant periods for policy analysis; certainly the dead-weight losses of distortions are likely to increase over time.

## II. Saving, Consumption, and Demand Analysis

### A. *Demand Analysis*

In the late 1960s and early 1970s, considerable attention was devoted to the construction of large-scale macroeconomic models, either for short-term forecasting and control, or for state planning that made a longer-term assessment of industrial structure, including manpower planning and investment. The Cambridge Growth Project (Stone and Brown 1962) used a model constructed around an input-output matrix, a consumption function, and a matched set of demand equations that used prices and incomes to map out the demand for commodities. The model used the linear expenditure system whose empirical analysis was pioneered in Stone (1954), perhaps the first case where the parameters of a utility function were estimated. Parenthetically, Stone seemed to be unaware of what is now (somewhat ironically) called the Stone-Geary utility function (Cobb Douglas with an affine shift of origin); the cost function was derived by Klein and Rubin (1948), with the direct utility function given in an accompanying note by Samuelson (1947–1948), and then rediscovered by Geary (1950).

When I was asked to work on the consumption part of the model in 1969, I was troubled by a number of issues, which eventually required or led to further developments. First, the linear expenditure system is linear in prices and income, but not in its parameters, and the Cambridge team had not managed to come up with satisfactory estimates. This was easily solved using the new methods then evolving in engineering and economics, particularly the algorithms by Marquardt (1963) and by Goldfeld, Quandt, and Trotter (1966). This was the computational frontier in economics in the 1960s. According to Cambridge legend, Prais and Houthakker (1956) was the first study in economics to use an electronic computer; in Stone (1954), the term is not computer, but computor, which refers to a person not a machine.

Second, the utility underlying the linear expenditure system is additive, and as such cannot handle the full generality of behavioral response that choice theory permits. Of course, strong restrictions were essential at a time when researchers were limited to a handful of time-series data points, but even so, assuming additivity from the start loses control of the choice of how much to measure and how much to assume, always a key to convincing applied work. Even then, there is more than one

kind of additivity, and different functional forms for each, and the question of how to choose between them, issues I tried to address in Deaton (1974b).

Additivity is more plausible for broad groups of goods, such as food or clothing, but how should such groups of goods be defined? By what principles are commodities and services to be combined into groups, and how might those principles be implemented in practice? Perhaps it would make sense to have a model of two-stage or hierarchic preferences, in which additive preferences are used to allocate income into broad categories, and some other rule used to break up the aggregates into finer categories. Just as aggregation over goods was a problem, so was aggregation over people. What justified the use of the representative agent in the linear expenditure system? Was this just an assumption, or an implication of such a utility function? And more broadly, why were demand functions not influenced by the distribution of income? In Cambridge in the 1960s and 1970s, we all listened to the Cambridge Keynesians, their denunciations of the validity of the nonsubstitution theorem, and their insistence on the interdependence of equilibrium prices and income distribution; could it really be true that a sweeping increase in income inequality would leave consumer demands unchanged?

The theory that answered most of those questions had been already developed by W. M. (Terence) Gorman; on aggregation in Gorman (1953), on two-stage budgeting in Gorman (1959), and on separability in Gorman (1968). These papers, which made extensive use of dual representation of preferences, were not easily comprehensible by those with a standard economics education at that time; they are classic examples of papers that were so far ahead of their time that their influence and usefulness for applied work was recognized only after the results had become familiar from other work. One strand of that other work came from Daniel McFadden, then at Berkeley, whose lectures had an enormous influence on those who heard them, and which were eventually published as McFadden (1978). McFadden's work influenced a generation of researchers on production and demand, including John Muellbauer who attended his lectures, and with whom I collaborated in England after 1970. Together we could combine Gorman and McFadden, and develop our combined knowledge into a coherent body of work that was relevant for the empirical work that we were both trying to do; this program eventually led to our book on consumer behavior, Deaton and Muellbauer (1980b). It also enabled us to address the aggregation issues.

There were two key sets of results. One was Muellbauer's (1975) work on aggregation, which extended Gorman's results to a class of preferences, price independent generalized linearity (PIGL), in which the distribution of income played an essential role in the demand functions. Of particular empirical interest was a logarithmic case (PIGLOG), by which the budget share of each good was a linear function of log income for each household, which previous studies had suggested gave a good fit to the data. PIGLOG preferences allow the aggregation of demands, and again the aggregate budget shares are linear in log income, but log income is adjusted for a measure of income inequality. The other key came from another student of McFadden; this was Diewert's (1971, 1974) concept of a flexible function form. The concept provided a solution to the problem of assumption versus measurement; a system of demand functions is a flexible functional form if the demand functions are unrestricted beyond the general restrictions implied by choice theory, adding

up, zero degree homogeneity in prices and income, and symmetry and negative semi-definiteness of the Slutsky matrix of compensated price derivatives. A flexible functional form is a utility-consistent set of demand functions that can provide a first-order approximation (at a point) to an arbitrary set of utility-consistent set of demand functions. Given that the demand functions are derivatives of the expenditure (or cost) function, the preferences underlying the flexible functional forms are second-order local approximations to arbitrary preferences. A number of notable flexible functional forms were developed, by Diewert himself, the Generalized Leontief system (Diewert 1971), and by Christensen, Jorgenson, and Lau (1975), the translog system, and later, the Almost Ideal Demand System (AIDS) of Deaton and Muellbauer (1980a). The AIDS, which is a member of the PIGLOG class, was the result of many months of tinkering to try to combine the best features of the earlier models; it allows linear estimation of its parameters, at least under an often reasonable approximation, it has one income parameter per commodity, which controls whether or not the good is a necessity or luxury, and a matrix of own and cross price responses of the budget shares on the logarithm of prices. It is a flexible functional form, with an explicit indirect utility function. The modesty of the "almost" refers to the fact that the quasiconcavity of the cost function cannot be globally imposed without destroying the flexibility of the functional form. Even so, its convenience and consistency with price theory, as well as the availability of a quadratic generalization, Banks, Blundell, and Lewbel (1997) has made it a widely used tool in work that requires inference from prices to welfare, for example in tax evaluation, regulatory, or antitrust work. Note that the AIDS, in spite of its incorporation of distribution, retains much of the representative agent; the model works by choosing a convenient utility function that, if possessed by everyone, would lead to a representative utility with adjustment only to income. It is a much harder undertaking to start from more realistic models of heterogeneous individuals, each of whom buys a different collection of goods, and then explicitly aggregate them, a procedure that, in general, will not lead to aggregate behavior that is in any way analogous to individual behavior, see for example, Houthakker (1955) for a startling example from production theory where Leontief technology for each firm turns into Cobb-Douglas for the economy as a whole.

## B. *Saving and Intertemporal Choice*

The two papers by Modigliani and Brumberg (1955a, 1955b) are the foundations for utility-based modeling of intertemporal choice. I had the good fortune to be sent to read both (one of which was then unpublished) when I was an undergraduate, and they seemed to me then (and now) a template for how to do economics. They proposed a simple theoretical structure for choice over time, which gave a clear way of thinking about an issue of the first order of importance, both for individuals and for society. The theory was used to bring order and a coherent interpretation to a mass of contradictory and previously disorganized empirical evidence from many studies using both cross-section and time series; the theory had to match everything, or it was nothing. It provided clear new predictions that could be tested. In a later paper, Modigliani (1970) extended the theory to cross-country evidence, and derived the famous rate of growth effects on national saving. Even if each individual has no

net saving over life, accumulating wealth in youth and running it down in old age, the economy will have a positive saving rate if there is either population growth or economic growth because the savers would be either more numerous or working on a larger scale than the dissavers. There is no representative agent, the aggregate is unlike any individual, and the explicit aggregation generates nonobvious hypotheses linking national growth and national saving. It is an abiding sadness of my career that it turns out that there is overwhelming evidence that these correlations, which are indeed in the data, and that Modigliani was the first to see, turned out *not* to be attributable to life-cycle rate of growth effects; there is simply not enough age-related life-cycle saving to drive the effects (Kotlikoff and Summers 1981; Carroll and Summers 1991), and fluctuations in national saving are driven more by within age-group effects rather than by changes in either the population or economic size of the cohorts themselves (Bosworth, Burtless, and Sabelhaus 1991; Paxson 1996; Deaton and Paxson 1997, 2000).

In the 1970s, a major innovation to life-cycle theory was the extension to consumption and labor supply simultaneously. Heckman's (1971) PhD thesis appears to have been the key here, and was followed by a book by Ghez and Becker (1975), as well as an empirical landmark paper by MaCurdy (1981). In these models, there is a subutility function for each period of life, in which current labor supply (or leisure) appears together with consumption, and lifetime utility is a discounted sum over all periods. Under certainty, labor supply and the commodity demands are functions of current wage rates and prices together with a quantity, interpretable as the marginal utility of lifetime wealth, which captures the budget constraint and links the periods together. Similar (though static) demand functions were first used by Frisch (1932) so that functions of prices, wages, and the marginal utility of money have come to be known as Frisch demand functions, in distinction to Hicksian demands (functions of utility and prices) or Marshallian demands (of income and prices). MaCurdy noted that with suitable functional forms, the unobservable marginal utility of money, which was constant over the lifetime, could be treated as a fixed effect in panel data and differenced away, while, under uncertainty, the marginal utility of money behaved as a martingale difference, so that its difference was an unpredictable shock.

Deaton (1985) and Browning, Deaton, and Irish (1985) extended this line of work. Panel data, such as those used by MaCurdy, were (and are) scarce while repeated (independently drawn) cross-sectional data are available in most countries. Deaton (1985) noted that it is possible to track birth cohorts through successive household surveys, so that, for example, the data on those aged 30 in year 1 can be pooled with those aged 31 in year 2, and so on. The validity of this depends on there being a fixed underlying birth cohort available for repeated independent sampling, which is problematic among older cohorts where mortality rates are high. Although the method cannot track individuals over time, it can track statistics of birth cohorts over time, not only means, but medians, variances, and higher moments, and although all such statistics are samples from the underlying birth cohort, the sampling scheme is known, and standard errors can be calculated following standard methods so that errors-in-variables (or instrumental variable) methods can be used to adjust estimators and correct attenuation bias (e.g., Fuller 1987). The method was first applied in Browning, Deaton, and Irish (1985)—henceforth, BDI—and has been widely used in many subsequent studies of life-cycle behavior.

BDI's findings, using British data, were not favorable for the model. Wage rates are hump-shaped over the life cycle, peaking in middle age. In the simplest form of the life-cycle consumption story, and as proposed in the original Modigliani and Brumberg papers, the starting hypothesis was that consumption should be flat over life, and the essence of the model is the saving and dissaving that permits consumption smoothing over life. Yet in our data, as in subsequent findings, consumption tracks wages and earnings, also peaking in middle age. We also tested, and faulted, the hypothesis that the way that labor supply responded to wages over the life cycle was consonant with the way that it responded over the business cycle, so that high frequency and low frequency smoothing did not appear to be governed by the same set of parameters.

There have been two responses to such findings. One is to maintain the strict version of the model, and conclude that the life-cycle story really does not explain very much (which of course does not mean that thinking about consumption, saving, and pensions should not be done within a life-cycle framework). The other is to generalize the model, noting, for example, that tastes or needs are different at different points in the life cycle, or to allow for precautionary motives: see, for example, Attanasio and Weber (2010) for a review of a large literature. A somewhat simpler story that is consistent with the findings in BDI is that consumption is much more closely tied to income than life-cycle theory predicts; for example, some fraction of people may live hand-to-mouth, consuming their earnings and whatever liquid assets they have on hand.

One way to think about hand-to-mouth consumption is to prohibit consumers from borrowing but leave other assumptions intact. For those who are patient enough to want to accumulate, the prohibition makes no difference because they do not want to borrow. Those who do want to borrow could simply consume their income, but can usually do better by accumulating and decumulating assets on their own, which helps smooth their consumption in the face of stochastic income or earnings; the classic example is a farmer in poor country, with little opportunity to borrow at affordable rates, and weather-driven stochastic income.

A formal model of such a consumer is mathematically identical to the classic model of speculative commodity storage as originally developed by Gustafson (1958), used by Newbery and Stiglitz (1981), and, in the consumption and savings literature, by Schechtman (1976), Schechtman and Escudero (1977), and by Zeldes (1989). In a similar vein, Deaton and Laroque (1992) had been further exploring (and challenging) the empirical usefulness of the commodity model, the consumption version of which appears in Deaton (1991). Together with the similar model by Carroll (1997), in which precautionary motives prevent consumers from ever wanting to borrow, these models are now referred to as "buffer stock" models of saving. Consumers typically maintain some assets so as to buffer consumption against random fluctuations in income just as a smallholder will keep a stock of grain so that, even with a bad harvest, there will be something to eat. Eventually, after a series of sufficiently bad draws of income, the consumer will spend everything she has—current income plus assets—because the marginal utility of money now is higher than the expected marginal utility of money going into the next period without a buffer. These asset "stock outs" can be rare, so that, the consumer nearly always has *some* cash on hand. In spite of never being able to borrow, and in spite of almost always

having assets, the behavior of these consumers is quite different from the behavior of a consumer who can borrow. In effect, the presence of the borrowing constraints changes the consumer's behavior—their policy function—even though the borrowing constraint almost never bites. These "hand-to-mouth" consumers are also different from those who simply spend their earnings each period; they do better at consumption smoothing and are better off with the same earnings process.

A final point is relevant for thinking about saving, development, and poverty. Many buffer stock savers would not be made better off if somehow they were made to save and "escape" from their low-asset buffer-stock stochastic equilibrium; if they are given assets, they would be better off, but will eventually return to the buffer stock stochastic equilibrium. In that sense, they do not want to be rich. If borrowing restrictions were eased, for example by the opening of local banks charging low interest rates, they would borrow to raise their current consumption and move to their optimal trajectory of falling consumption over time. One might easily mistake such an increase in consumption as evidence that better credit reduces poverty, but the mechanism has nothing to do with escaping from poverty by borrowing for productive assets, which these consumers do not want. Instead, we are looking at a consumption boom that is financed by borrowing that, in the long run, will produce lower consumption (higher poverty as usually measured), even though these people are indeed better off: see Fulford (2013).

The other decisive innovation of the 1970s was Hall's (1978) reworking of the permanent income hypothesis (PIH) in the light of rational expectations; Hall derived the stochastic Euler equation relationship between consumption in adjacent periods, an approach that has dominated most of the subsequent literature. Flavin (1981), under the assumption of certainty equivalent quadratic preferences, worked out a workhorse version of the PIH that gave an explicit form for the change in consumption in terms of the changes in expectations of current and future labor incomes. In this model, which is a special case of Hall's, consumers do not change their consumption from one period to the next unless current or expected future labor income changes, and when there are such changes, there is an explicit formula that depends on the interest rate and the time horizon that gives the change in consumption that is warranted by the changes in earnings prospects. The implication that consumption change should be zero in the absence of new information—the famous random walk of consumption hypothesis—seemed absurd at the time, given that standard consumption functions for the previous 25 years had regressed consumption on large numbers of lags of income and lagged consumption. That Hall could barely reject the hypothesis was almost as stunning as if he had accepted it, and radically changed subsequent research in the field. As investigations continued, it became evident that, in US quarterly data, the change in consumption was in fact correlated with the lagged change in labor income, which became known as the excess sensitivity finding.

The explicit link between innovations in labor income and consumption change in the PIH allowed for the first time a precise characterization of the dynamics of consumption. In particular, if we know the dynamics of earnings, we can solve for the dynamics of consumption. Today, the data on earnings are infinitely better, and the estimates more sophisticated (see Guvenen et al. 2015), but in the mid-1980s, the best practice approximation to mean quarterly labor income in the United States was a first-order autoregressive process in growth rates, with a positive autoregressive

parameter. When we plug this into the formula for the change in consumption, we get the counterintuitive result that consumption should respond more than one for one to innovations in earnings, so that consumption, far from being smoother than income, which is what the data say, and what the permanent income hypothesis had been designed to explain, should actually be *less* smooth than income. The PIH actually predicts the opposite of its most famous predictions and all the textbooks are wrong (Deaton 1987).

If innovations in earnings growth are positively autocorrelated, an unanticipated earnings gain is not only good in itself, but it signals that there is more to come next quarter. In consequence, people can spend, not only their Christmas bonus, but also the Easter bonus that it signals.

How can we escape this paradox? First, it is derived under the assumption of certainty equivalence, so that there is no precautionary motive for holding back given that the Easter bonus is far from certain. Second, and more promising, is the possibility that earnings innovations have to be paid back, so that the bonus in the first and second quarters has to be paid back eventually, or even more than paid back. Earnings may be tied to some predetermined path—for example, set by the personnel department on the day you join your company—so that the bonuses are always short-lived; people know this, and do not change their consumption by much. In the end, this possibility is hard to test, if only because the very long-run ripples from earnings innovations can only be observed on very long-run data.

A better resolution brings us back to one of the main themes of this article, which is to think about aggregation, about individuals versus aggregates, and about the amount of information that individuals might reasonably possess. The dynamic properties of average earnings come from averaging millions of individual earnings processes, and averaging does not preserve dynamics. In particular, the idiosyncratic component of individual earnings will be annihilated by the averaging and leave only any common or macro component. Yet this component may account for only a small fraction of the individual's earnings, small enough that she may not be aware of it. More generally, there can be a stochastic process, such as a random walk, or an autocorrelated growth process, that is common to all workers, but on top of which, each worker has their own idiosyncratic process, for example pure white noise. If each worker follows the PIH, consumption will be smoother than earnings for each individual. If the aggregate process accounts for only a small share of the variance for each individual, and if the individual cannot filter out the aggregate component her own earnings, then as formally shown by Pischke (1995), aggregate consumption will fail to satisfy the PIH formula; it will appear to be too smooth relative to the aggregate earnings process, and will be correlated with previous information, as in the excess sensitivity finding. The plumber, who worries most about the year-to-year fluctuations in his earnings, has little reason to be concerned about a common macro process that is part of his earnings, but that accounts for a tiny fraction of it, and would be difficult to detect, even if he knew that it existed.

Note that aggregation issues exist even in Hall's original Euler equation version. Even if each person's consumption follows a random walk, births and deaths will generate predictable growth if the young are systematically richer than the old. Similarly, unless people (or dynasties) live forever, aggregate consumption will not satisfy the individual Euler equations.

The permanent income hypothesis also has implications for inequality in consumption, income, and wealth; these were developed in Deaton and Paxson (1994). Suppose, as in Hall, that each consumer in a birth cohort is random walking. Unless their consumption changes are perfectly correlated, the cohort members' consumption levels will move further apart as the cohort ages. Consumption inequality will increase with age within a birth cohort. Aggregate inequality in the economy depends on demographics, as in Modigliani's growth effects for saving. Note that this does *not* require that earnings inequality increase with age, though it is entirely consistent with it doing so. We can imagine a stationary distribution of earnings over people, in which each person has a person-specific mean plus an individual-specific stationary process; if the innovations to these individual-specific processes are (at least partially) uncorrelated across people, and if people follow the certainty-equivalent PIH, then the changes in consumption will be less than perfectly correlated, and consumption inequality will widen over time within the birth cohort. This theory also implies that *income* inequality should also widen, whether or not *earnings* inequality is increasing (the difference between income and earnings being the return to assets) while wealth inequality will increase an order of magnitude more rapidly than consumption inequality, essentially because there is increasing inequality in the increments to wealth, so that individual asset levels are spiraling apart even more rapidly. Everything here is driven by the stochastic innovations to labor income, or luck, so that consumption, income, and wealth inequality can be thought of as the fossils of accumulated luck.

Paxson and I looked at the data on birth cohorts using the repeated cross-section method for three countries, the United States, Britain, and Taiwan, and found that consumption inequality did indeed increase with age. While it is always gratifying when a new prediction is confirmed in the data, Popper's curse is always lurking in the background; we learn most clearly by refutation, not confirmation, so that if one theory fits the data, so will others. In this case, the obvious alternative explanation is one in which income inequality is increasing within a birth cohort as some members are more successful than others, and consumption is tied to income according to a buffer stock model. An acid test of the model against this situation would be one in which the cross-sectional labor earnings were stationary, with no inequality increase, and where consumption inequality is still increasing. I know of no such test.

Even so, and as Paxson and I noted, the "fanning-out" of consumption with age is a measure of the extent to which society fails to provide consumption insurance to its members. With perfect insurance, which is typically undesirable because of moral hazard, individual idiosyncratic luck would be neutralized by redistribution across people, and there would be no fanning out of consumption. Of course, society provides its citizens with insurance in many ways, through families, through unemployment, disability, and pension payments, and even through national defense. The spread of consumption inequality is potentially informative about all of these, and so is of great interest beyond a test of the PIH. For example, as noted by Lucas (2003), the fanning out of consumption and of earnings, though consistent with both the PIH and with the buffer stock model, in either case indicates that there is less than perfect insurance; insurance could operate on consumption, conditional on income, or through income itself. The extent of insurance is also closely tied to

the other puzzles discussed above; Campbell and Deaton (1989) showed that Hall's excess sensitivity finding can be seen as another aspect of the excess smoothness finding, or vice-versa, and Attanasio and Pavoni (2011) show that both can be used to measure the extent of insurance. Blundell (2014) provides a general discussion and review of recent developments.

Explicitly modeling earnings at the individual level, separating out macro from idiosyncratic effects, assessing insurance, and thinking about macroeconomics as an aggregate of heterogeneous agents has become a (or perhaps even the) central topic in macroeconomics today, displacing the representative agent models that so troubled me at the beginning of my career. In one key paper Aiyagari (1994) developed a general equilibrium model with buffer stock consumers; although each consumer is saving and dissaving, aggregate variables are unchanging, and there is no representative agent. Krusell and Smith (1998) added macroeconomic shocks to this model, thus resurrecting a representative agent. Even so, they make a range of special assumptions, and there is good ongoing work today exploring more realistic cases where, once again, aggregated and individual behavior are sharply different: for example Kaplan, Violante, and Weidner (2014) and Kaplan, Moll, and Violante (2016). According to the opening words of a recent review paper by Heathcote, Storesletten, and Violante (2009, p. 319), "Macroeconomics is evolving from the study of aggregate dynamics to the study of the dynamics of the entire equilibrium distribution of allocations across individual economic actors." I am delighted if I have played a part in this evolution.

### III. Conclusion: On Discoveries

The Nobel committee asks that this article include a discussion of my discoveries and how they happened. There are certainly many things that I know now that I did not know when I started though many of them are things that other people knew all along—I am thinking in particular of McFadden and Gorman. My usual experience is that a "discovery" turns out to be wrong: it turns out to be a coding error, or a misinterpretation of theory or of data, or is not a discovery at all, but has long been known. Occasionally, discoveries are real, though most are personal, in the sense that they change what I think, but not what others think.

I cannot resist referring back to one of my first discoveries on consumption and saving. As a cash-strapped young father in Britain in the mid-1970s when the government had lost control of inflation (the British retail price index rose by 16, 24, and 17 percent in 1974, 1975, and 1976), I realized that, when I went shopping, mostly for one or two items at a time, I could not tell inflation from relative price increase, especially in the early months of the inflation, and especially for items like coffee which, for me then, was a luxury with a wildly fluctuating price. On the basis of this, I argued, in Deaton (1977), that in an economy with unanticipated inflation, there would be "involuntary" saving as each consumer, buying one good at a time, held off from that good on the mistaken supposition that it was only relatively expensive. I recall too that my Cambridge colleagues, who were kind to a young research assistant, thought this was interesting, but absurd; we all knew, after all, that inflation caused people to dissave, and that this was how hyperinflations worked. So I was as surprised as were my colleagues when the entirely unanticipated increase in the

household saving rate was announced, and later when there were similar increases in a number of other countries (Koskela and Virén 1982). That was also my first encounter with Popper's curse: if one story correctly predicts a new finding—however unexpected—many other stories can do so too, in this case wealth effects are candidates.

Beyond that, I look back with the greatest pleasure on three discoveries. None of these insights came from problems that I was working on, but from realizing that something apparently unconnected had implications elsewhere. The first is that it is possible to track birth cohorts through repeated cross-sectional surveys, and that this insight could be used to investigate life-cycle consumption and labor supply. The second is my work with Christina Paxson on the dynamic effects of luck. The way that people respond to luck is a mechanism that drives up consumption inequality within a birth cohort and that, in turn, allows us to assess the extent to which society insures its members. Third, in development, is my work with Jean Drèze and with Paxson on food puzzles, primarily in India. The economies of scale puzzle had its early origins in a conversation with Drèze on the effects on well-being of family size. In time, that led to our later work on food and nutrition in India, on how food Engel curves do not identify the effects of rising income on calorie consumption, and on how nutrition depends on many more factors than the intake of calories and protein. Even on the subject of food and well-being, one of the oldest topics in economics, much remains unresolved.

## REFERENCES

**Aiyagari, S. Rao.** 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics* 109 (3): 659–84.

**Alkire, Sabina, James Foster, Suman Seth, Maria Emma Santos, José Manuel Roche, and Paola Bellon.** 2015. *Multidimensional Poverty Measurement and Analysis.* Oxford: Oxford University Press.

**Allen, Robert C.** 2016. "Absolute Poverty: When Necessity Displaces Desire." http://www.economics.ox.ac.uk/Oxford-Economic-and-Social-History-Working-Papers/absolute-poverty-when-necessity-displaces-desire (last updated March 2016).

**Atkin, David.** 2013. "Trade, Tastes, and Nutrition in India." *American Economic Review* 103 (5): 1629–63.

**Attanasio, Orazio P., and Nicola Pavoni.** 2011. "Risk Sharing in Private Information Models with Asset Accumulation: Explaining the Excess Smoothness of Consumption." *Econometrica* 79 (4): 1027–68.

**Attanasio, Orazio P., and Guglielmo Weber.** 2010. "Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy." *Journal of Economic Literature* 48 (3): 693–751.

**Banks, James, Richard Blundell, and Arthur Lewbel.** 1997. "Quadratic Engel Curves and Consumer Demand." *Review of Economics and Statistics* 79 (4): 527–39.

**Blundell, Richard.** 2014. "Income Dynamics and Life-Cycle Inequality: Mechanisms and Controversies." *Economic Journal* 124 (576): 289–318.

**Bosworth, Barry, Gary Burtless, and John Sabelhaus.** 1991. "The Decline in Saving: Evidence from Household Surveys." *Brookings Papers on Economic Activity* (1): 183–241.

**Browning, Martin, Angus Deaton, and Margaret Irish.** 1985. "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle." *Econometrica* 53 (3): 503–43.

**Campbell, John Y., and Angus Deaton.** 1989. "Why Is Consumption So Smooth?" *Review of Economic Studies* 56 (3): 357–73.

**Carroll, Christopher D.** 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112 (1): 1–55.

**Carroll, Christopher D., and Lawrence H. Summers.** 1991. "Consumption Growth Parallels Income Growth: Some New Evidence." In *National Saving and Economic Performance,* edited by B. Douglas Bernheim and John B. Shoven. Chicago: University of Chicago Press.

**Chen, Shaohua, and Martin Ravallion.** 2010. "The Developing World Is Poorer Than We Thought, but No Less Successful in the Fight against Poverty." *Quarterly Journal of Economics* 125 (4): 1577–625.

**Christensen, Laurits R., Dale W. Jorgenson, and Lawrence J. Lau.** 1975. "Transcendental Logarithmic Utility Functions." *American Economic Review* 65 (3): 367–83.

**Coffey, Diane.** 2015. "Pregnancy Body Mass and Weight Gain During Pregnancy in India and Sub-Saharan Africa." *Proceedings of the National Academy of Sciences of the USA* 112 (11): 3302–3307.

**Costa, Dora L.** 2001. "Estimating Real Income in the United State from 1988 to 1994: Correcting CPI Bias Using Engel Curves." *Journal of Political Economy* 109 (6): 1288–1310.

**Davies, David.** 1795. *The Case of Labourers in Husbandry Stated and Considered in Three Parts*. Bath: Cruttwell for Robinson.

**Deaton, Angus.** 1974a. "A Reconsideration of the Empirical Implications of Additive Preferences." *Economic Journal* 84 (334): 338–48.

**Deaton, Angus.** 1974b. "The Analysis of Consumer Demand in the United Kingdom, 1900–1970." *Econometrica* 42 (2): 341–67.

**Deaton, Angus.** 1977. "Involuntary Saving through Unanticipated Inflation." *American Economic Review* 67 (5): 899–910.

**Deaton, Angus.** 1985. "Panel Data from Time Series of Cross-Sections." *Journal of Econometrics* 30 (1–2): 109–26.

**Deaton, Angus.** 1987. Life-Cycle Models of Consumption: Is the Evidence Consistent with the Theory?" In *Advances in Econometrics: Fifth World Congress*, Vol. 2, edited by Truman Bewley, 1221–12248. New York: Cambridge University Press.

**Deaton, Angus.** 1988. "Quality, Quantity, and Spatial Variation of Price." *American Economic Review* 78 (3): 418–30.

**Deaton, Angus.** 1989a. "Rice Prices and Income Distribution in Thailand: A Non-Parametric Analysis." *Economic Journal* 99 (395): 1–37.

**Deaton, Angus.** 1989b. "Looking for Boy-Girl Discrimination in Household Expenditure Data." *World Bank Economic Review* 3 (1): 1–15.

**Deaton, Angus.** 1991. "Saving and Liquidity Constraints." *Econometrica* 59 (5): 1221–48.

**Deaton, Angus.** 1997. *The Analysis of Household Surveys*. Baltimore: Johns Hopkins University Press.

**Deaton, Angus.** 2001. "Counting the World's Poor: Problems and Possible Solutions." *World Bank Research Observer* 16 (2): 125–47.

**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.** 2010. "Price Indexes, Inequality, and the Measurement of World Poverty." *American Economic Review* 100 (1): 5–34.

**Deaton, Angus.** 2013. *The Great Escape: Health, Wealth, and the Origins of Inequality*. Princeton, NJ: Princeton University Press.

**Deaton, Angus, and Bettina Aten.** 2015. "Trying to Understand the PPPS in ICP 2011: Why Are the Results So Different?" National Bureau of Economic Research Working Paper 20244.

**Deaton, Angus, and Jean Drèze.** 2009. "Food and Nutrition in India: Facts and Interpretations." *Economic and Political Weekly* 44 (7): 42–65.

**Deaton, Angus, and Olivier Dupriez.** 2011. "Purchasing Power Parity Exchange Rates for the Global Poor." *American Economic Journal: Applied Economics* 3 (2): 137–66.

**Deaton, Angus, and Alan Heston.** 2010. "Understanding PPPs and PPP-Based National Accounts." *American Economic Journal: Macroeconomics* 2 (4): 1–35.

**Deaton, Angus, and Guy Laroque.** 1992. "On the Behaviour of Commodity Prices." *Review of Economic Studies* 59 (1): 1–23.

**Deaton, Angus, and John Muellbauer.** 1980a. "An Almost Ideal Demand System." *American Economic Review* 70 (3): 312–26.

**Deaton, Angus, and John Muellbauer.** 1980b. *Economics and Consumer Behavior*. New York: Cambridge University Press.

**Deaton, Angus, and John Muellbauer.** 1986. "On Measuring Child Costs: With Applications to Poor Countries." *Journal of Political Economy* 94 (4): 720–44.

**Deaton, Angus, and Christina Paxson.** 1994. "Intertemporal Choice and Inequality." *Journal of Political Economy* 102 (3): 437–67.

**Deaton, Angus, and Christina Paxson.** 1997. "The Effects of Economic and Population Growth on National Saving and Inequality." *Demography* 34 (1): 97–114.

**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): 2012–25.

**Diewert, W. Erwin.** 1971. "An Application of the Shephard Duality Theorem: A Generalized Leontief Production Function." *Journal of Political Economy* 79 (3): 481–507.

**Diewert, W. Erwin.** 1974. "Applications of Duality Theory." In *Frontiers of Quantitative Economics,* Vol. 2, edited by M. Intriligator and D. Kendrick, 106–71. Amsterdam: North-Holland.

**Ducpetiaux, Édouard.** 1855. *Budgets* économiques *des classes ouvrières en Belgique*, *subsistances, salaires, population*. Brussels: Commission Centrale de Statistique.

**Eden, Frederick Morton.** 1797. *The State of the Poor or, an History of the Laboring Classes in England, from the Conquest to the Present Period.* Cambridge, UK: Cambridge University Press, 2011.

**Engel, Ernst.** 1857. "Die productions- und consumtionverhältnisse des Königreiches Sachsen." Reprinted with Engel, 1895.

**Engel, Ernst.** 1895. "Die Lebenskosten belgischer Arbeiter-Familien früher und jetzt." *International Statistical Institute Bulletin* 9 (1): 1–124.

**Engels, Friedrich.** 1845. *Die Lage der arbeitenden Klasse in England.* Leipzig: Wigand. (English translation: New York, 1887).

**Ferreira, Francisco H. G, Shaohua Chen, Andrew Dabalen et al.** 2015. "A Global Count of the Extreme Poor in 2012: Data Issues, Methodology and Initial Results." World Bank Policy Research Working Paper 7432.

**Flavin, Marjorie A.** 1981. "The Adjustment of Consumption to Changing Expectations About Future Income." *Journal of Political Economy* 89 (5): 974–1009.

**Frisch, Ragnar.** 1932. *New Methods of Measuring Marginal Utility.* Tubingen: J. C. B. Mohr.

**Fulford, Scott L.** 2013. "The Effects of Financial Development in the Short and Long Run: Theory and Evidence from India." *Journal of Development Economics* 104 (0): 56–72.

**Fuller, Wayne A.** 1987. *Measurement Error Models.* New York: Wiley.

**Geary, R. C.** 1950. "A Note on a 'Constant-Utility Index of the Cost of Living'." *Review of Economic Studies* 18 (1): 65–66.

**Ghez, Gilbert R., and Gary S. Becker.** 1975. *The Allocation of Time and Goods over the Life Cycle.* New York: National Bureau of Economic Research Inc.

**Goldfeld, Stephen M, Richard E. Quandt, and Hale F. Trotter.** 1966. "Maximization by Quadratic Hill-Climbing." *Econometrica* 34 (3): 541–51.

**Gorman, William M.** 1953. "Community Preference Fields." *Econometrica* 21: 63–80.

**Gorman, William M.** 1959. "Separable Utility and Aggregation." *Econometrica* 27: 469–81.

**Gorman, William M.** 1968. "The Structure of Utility Functions." *Review of Economic Studies* 35: 367–90.

**Government of India.** 2014. *Report of The Expert Group to Review the Methodology for Measurement of Poverty.* Planning Commission, June.

**Gustafson, Robert L.** 1958. *Carryover Levels for Grains.* US Department of Agriculture Technical Bulletin 1178.

**Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song.** 2015. "What Do Data on Millions of U.S. Workers Reveal About Life-Cycle Earnings Risk?." National Bureau of Economic Research Working Paper 20913.

**Hall, Robert E.** 1978. "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy* 86 (6): 971–87.

**Hamilton, Bruce W.** 2001. "Using Engel's Law to Estimate CPI Bias." *American Economic Review* 91 (3): 619–30.

**Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante.** 2009. "Quantitative Macroeconomics with Heterogeneous Households." *Annual Review of Economics* 1 (1): 319–54.

**Heckman, James J.** 1971. "Three Essays on the Supply of Labor and the Demand for Market Goods." PhD diss. Princeton University.

**Houthakker, Hedrik S.** 1955. "The Pareto Distribution and the Cobb-Douglas Production Function in Activity Analysis." *Review of Economic Studies* 23 (1): 27–31.

**Inklaar, Robert, and D. S. Prasada Rao.** 2014. "Cross-Country Income Levels Over Time: Did the Developing World Suddenly Become Much Richer?" Groningen Growth and Development Center Research Memorandum 151.

**Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante.** 2016. "Monetary policy according to HANK." National Bureau of Economic Research Working Paper 21897.

**Kaplan, Greg, Giovanni L. Violante, and Justin Weidner.** 2014. "The Wealthy Hand-to-Mouth." *Brookings Papers on Economic Activity* (1): 77–138.

**Klein, L. R., and H. Rubin.** 1948. "A Constant-Utility Index of the Cost of Living." *Review of Economic Studies* 15 (2): 84–87.

**Koskela, Erkki, and Matti Virén.** 1982. "Saving and Inflation: Some International Evidence." *Economics Letters* 9 (4): 337–44.

**Kotlikoff, Laurence J., and Lawrence H. Summers.** 1981. "The Role of Intergenerational Transfers in Aggregate Capital Accumulation." *Journal of Political Economy* 89 (4): 706–32.

**Kravis, Irving B., Alan W. Heston, and Robert Summers.** 1978. "Real GDP per Capita for More than One Hundred Countries." *Economic Journal* 88 (350): 215–42.

**Krusell, Per, and Anthony A. Smith, Jr.** 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy* 106 (5): 867–96.

**Lucas, Robert E., Jr.** 2003. "Macroeconomic Priorities." *American Economic Review* 93 (1): 1–14.

**MaCurdy, Thomas E.** 1981. "An Empirical Model of Labor Supply in a Life-Cycle Setting." *Journal of Political Economy* 89 (6): 1059–85.

**Mahalanobis, Prasanta Chandra.** 1946. "Recent Experiments in Statistical Sampling in the Indian Statistical Institute." *Journal of the Royal Statistical Society* 109 (Series A): 326–70.

**Marquardt, Donald W.** 1963. "An Algorithm for Least-Squares Estimation of Nonlinear Parameters." *Journal of the Society of Industrial and Applied Mathematics* 11 (2): 431–41.

**McCarthy, Paul.** 2013. Extrapolating PPPs and Comparing ICP Benchmark Results." In *Measuring the Real Size of the World Economy*, Chapter 18. Washington, DC: World Bank.

**McFadden, Daniel.** 1978. "Duality of Production, Cost, and Profit Functions." In *Production Economics: a Dual Approach to Theory and Applications,* edited by Mel A. Fuss and Daniel McFadden, 3–109. Amsterdam: North Holland.

**McFadden, Daniel.** 2014. "The New Science of Pleasure: Consumer Choice Behavior and the Measurement of Well-Being." In *Handbook of Choice Modeling*, edited by Stephane Hess and Andrew Daly, 7–48. Elgar.

**Milanovic, Branko.** 2007. *Worlds Apart: Measuring International and Global Inequality*. Princeton, NJ: Princeton University Press.

**Minhas, Bagicha S.** 1988. "Validation of Large-Scale Sample Survey Data: Case of NSS Estimates of Household Consumption Expenditure." *Sankhyā: the Indian Journal of Statistics* 50 Series B (3 Supplement): 1–63.

**Modigliani, Franco.** 1970. "The Life-Cycle Hypothesis and Intercountry Differences in the Saving Ratio." In *Induction, Growth and Trade: Essays in Honour of Sir Roy Harrod,* edited by W. A. Eltis, M. FG. Scott, and J. N. Wolfe. 197–225. Oxford. Oxford University Press.

**Modigliani, Franco, and Richard H. Brumberg.** 1955a. "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data." In *Post-Keynesian Economics,* edited by Kenneth K. Kurihara, 388–436. London: George Allen and Unwin.

**Modigliani, Franco, and Richard H. Brumberg.** 1955b. "Utility Analysis and Aggregate Consumption Functions: An Attempt at Integration." In *The Collected Papers of Franco Modigliani*: *Volume 2*, *The Life-cycle Hypothesis of Saving,* edited by Andrew Abel, 128–97. Cambridge, MA: MIT Press, 1990.

**Muellbauer, John.** 1975. "Aggregation, Income Distribution and Consumer Demand." *Review of Economic Studies* 42 (4): 525–43.

**Newbery, David M. G., and Joseph E. Stiglitz.** 1981. *Theory of Commodity Price Stabilization: Study in the Economics of Risk*. Oxford: Oxford University Press.

**Neyman, Jerzy.** 1934. "On the Two Different Aspects of the Representative Method: The Method of Stratified Sampling, and the Method of Purposive Selection." *Journal of the Royal Statistical Association* 97: 558–625.

**Nicholson, J. Leonard.** 1976. "Appraisal of Different Methods of Estimating Equivalence Scales and Their Results." *Review of Income and Wealth* 22: 1–11.

**Paxson, Christina.** 1996. "Saving and Growth: Evidence from Micro Data." *European Economic Review* 40(2): 255–88.

**Pigou, Arthur Cecil.** 1910. "A Method of Determining the Numerical Value of Elasticities of Demand." *Economic Journal* 20 (80): 636–40.

**Pischke, Jörn-Steffen.** 1995. "Individual Income, Incomplete Information, and Aggregate Consumption." *Econometrica* 63(4): 805–40.

**Pollak, Robert A., and Terence J. Wales.** 1979. "Welfare Comparisons and Equivalence Scales." *American Economic Review* 69(2): 216–21.

**Prais, Sigbert J., and Hendrik S. Houthakker.** 1956. *The Analysis of Family Budgets*. Cambridge, UK: Cambridge University Press.

**Rothbarth, Erwin.** 1943. "Note on a Method of Determining Equivalent Income for Families of Different Composition." In *War-time Pattern of Saving and Spending*, edited by Charles Madge, Appendix 4. Cambridge, UK: Cambridge University Press.

**Samuelson, Paul A.** 1947–1948. "Some Implications of 'Linearity'." *Review of Economic Studies* 15 (2): 88–90.

**Samuelson, Paul A.** 2004. "The First Fifteen Nobel Laureates in Economics, and Fifteen More Might-Have-Beens." In *Nobel Economists*, edited by Peter Badge. Foundation Lindau Nobel Laureates Meetings.

**Sen, Amartya K.** 1989. "Women's Survival as a Development Problem." *Bulletin of the American Academy of Arts and Sciences* 43 (2): 14–29.

**Seng, You Poh.** 1951. "Historical Survey of the Development of Sampling Theories and Practice." *Journal of the Royal Statistical Society* 114 Series A (2): 214–31.

**Schechtman, Jack.** 1976. "An Income Fluctuation Problem." *Journal of Economic Theory* 12 (2): 218–41.

**Schechtman, Jack, and Vera L. S. Escudero.** 1977. "Some Results on 'an Income Fluctuation Problem'." *Journal of Economic Theory* 16 (2): 151–66.

**Smith, T. M. Fred.** 1976. "The Foundations of Survey Sampling: A Review." *Journal of the Royal Statistical Society* Series A 139 (2): 183–204.

**Stigler, George J.** 1954. "The Early History of Empirical Studies of Consumer Behavior." *Journal of Political Economy* 62(2): 95–113.

**Stone, Richard.** 1954. "Linear Expenditure Systems and Demand Analysis: An Application to the Pattern of British Demand." *Economic Journal* 64 (255): 511–27.

**Stone, Richard, and Alan Brown.** 1962. *A Computable Model of Economic Growth: A Programme for Growth.* Vol. 1. London: Chapman and Hall.

**Summers, Robert, and Alan Heston.** 1991. "The Penn World Table (Mark 5): An Expanded Set of International Comparisons, 1950–1988." *Quarterly Journal of Economics* 106(2): 327–68.

**Villermé, Louis-René.** 1830. "De la mortalité dans les divers quartiers de la ville de Paris, et des causes qui la rendent très différente dans plusiers d'entre aux, ansi que les divers quartiers de beaucoup de grandes villes." *Annales d'Hygiène Publique et de Médecine Légale* 3: 294-341.

**Visaria, Pravin.** 2000. "Poverty in India during 1994–98: Alternative Estimates." Institute for Economic Growth, New Delhi.

**World Bank.** 1990. *World Development Report: Poverty*. New York: Oxford University Press.

**World Bank.** 2016. PovcalNet: an online analysis tool for global poverty monitoring. http://iresearch.worldbank.org/PovcalNet/ (accessed, January 20, 2016).

**Zeldes, Stephen P.** 1989. "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence." *Quarterly Journal of Economics* 104 (2): 275–98.