



Scientific Background on the Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2015

ANGUS DEATON: CONSUMPTION, POVERTY AND WELFARE

The Committee for the Prize in Economic Sciences in Memory of Alfred Nobel

The consumption of goods and services is a fundamental determinant of human welfare. The distribution of consumption among individuals has a bearing on many important issues – including inequality and poverty – in society’s economic, political and social domains. In most countries, aggregate consumption is the largest component of aggregate demand and, as such, accounts for much of the time variation in economic activity. For a given level of income, consumption determines savings and thus investment through the supply of capital. It is thus quite natural that consumption has been at the center of economic research throughout the last century.

Over the last three to four decades, the study of consumption has progressed enormously. While many scholars have contributed to this progress, Angus Deaton stands out. He has made several fundamental and interconnected contributions that speak directly to the measurement, theory, and empirical analysis of consumption. His main achievements are three.

First, Deaton’s research brought the estimation of demand systems – i.e., the quantitative study of consumption choices across different commodities – to a new level of sophistication and generality. The Almost Ideal Demand System that Deaton and John Muellbauer introduced 35 years ago, and its subsequent extensions, remain in wide use today – in academia as well as in practical policy evaluation.

Second, Deaton’s research on aggregate consumption helped break ground for the micro-econometric revolution in the study of consumption and saving over time. He pioneered the analysis of individual dynamic consumption behavior under idiosyncratic uncertainty and liquidity constraints. He devised methods for designing panels from repeated cross-section data, which made it possible to study individual behavior over time, in the absence of true panel data. He clarified why researchers must take aggregation issues seriously to understand total consumption and saving, and later research has indeed largely come to address macroeconomic issues through microeconomic data, as such data has increasingly become available.

Third, Deaton spearheaded the use of household survey data in developing countries, especially data on consumption, to measure living standards and poverty. In so doing, Deaton helped transform development economics from a largely theoretical field based on crude macro data, to a field dominated by empirical research based on high-quality micro data. He showed the value

of using consumption and expenditure data to analyze welfare of the poor, and identified shortcomings when comparing living standards across time and place.

Deaton's research has addressed issues of great practical significance, and his contributions have influenced policymaking in developing and developed countries. His work covers a wide spectrum, from the deepest implications of theory to the grittiest detail of measurement. The common themes are connecting theory and measurement, and linking micro and macro data by using relevant statistical methods.

This document presents the three pillars of Deaton's research. Section 1 surveys his analysis of demand across commodities at a given point in time, with a focus on Deaton and Muellbauer's Almost Ideal Demand System. Section 2 discusses his most important contributions on aggregate consumption across time, with a focus on the use of individual (or household-level) data and careful handling of the aggregation problem. Section 3 deals with Deaton's contributions to welfare measurement in developing countries, with a focus on measurement and analysis of the living standards of the poor. Section 4 briefly mentions a couple of related contributions.

1. Demand analysis

1.1 Background

When studying consumption, individual decisions are often decomposed in various ways. A common approach is to distinguish consumption decisions across commodities in a given time period, be it a month or a year, from consumption decisions across time, i.e., how much to spend on consumption today and how much to save. Under some conditions – e.g., intertemporal additive separability – such a two-stage approach is equivalent to a more general approach, where an array of different commodities at a given time and total spending at that time are chosen simultaneously, by maximizing the expected value of an intertemporal utility function, with consumed quantity q_{jt} of each commodity j at each date t as arguments (Deaton and Muellbauer, 1980b, chapter 5).

This section deals with consumption choices across commodities at a given point in time. These choices can be represented by a *demand system*, a system of equations to describe the relationships between quantities demanded on the one hand, and prices and total expenditure on

the other. Such a system can be estimated on individual data, assuming we have data on prices and quantities consumed from H individuals and J commodities (or commodity groups), or aggregate data, assuming we have data on prices and aggregate quantities consumed over T time periods.

In any given time period t , a demand system for household h can be written (for the sake of exposition, drop the time subscript)

$$q_{jh} = g_{jh}(p, c_h) \quad j = 1, \dots, J, \quad h = 1, \dots, H, \quad (1)$$

where q_{jh} is the demand of household h for commodity j as a function $g_{jh}(\cdot)$ of a vector of prices $p = (p_1, \dots, p_J)$ and household expenditure, c_h , which in this section is taken as exogenous. Thus, demand for commodity j depends on its own price, on the prices of all other commodities, and on total expenditure. This system shows how the household's demand for each commodity changes in response to changes in prices and total expenditures. An aggregate demand system can be written in a similar fashion simply by dropping the subscripts h in (1). The relation between individual demand q_{jh} and market demand q_j is the essence of the so-called aggregation problem, which is discussed further in Section 1.3

A number of researchers used parameterized versions of such a model back in the 19th century. Perhaps the most famous early example is the estimation of Engel curves, named after the German statistician Ernst Engel. An Engel curve for household h connects the nominal expenditure $p_j q_{jh}$, or budget share $w_{jh} \equiv \frac{p_j q_{jh}}{c_h}$, spent on a particular good j to total expenditures (household income in a static case).

Parameterizing a model like (1) means taking a stand on the exact relationship between prices, expenditures and demand in the data. The approach chosen by most economists before Deaton was to assume that (1), written for aggregate data, represented the demand of a single rational individual (or household). In this case, rationality means that demand is derived from the maximization of that single individual's (or household's) utility function. This assumption provided structure on the relations that connect prices (p) and total expenditures (c) to total market demand, q_j .

Assuming that demand is derived from a rational consumer's welfare maximization is not only convenient from a purely statistical point of view. It also allows the researcher to use the powerful tools of consumer theory to derive useful welfare measures, such as consumer surplus, equivalent and compensated variations in income, and theoretically founded price indices (often referred to as cost-of-living indices). Moreover, it provides structural estimates of behavioral responses that can be used to forecast changes in demand due to prospective policy changes.

An early attempt to link statistical analysis to economic theory was made by 1984 Economics Laureate Sir Richard Stone. Stone (1954a, b) developed the so-called Linear Expenditure System (LES) where the relationship between market demand, prices and total expenditures (in a given time period t) is given by

$$q_j = \gamma_j + \frac{\beta_j}{p_j} (c - \sum_k \gamma_k p_k), \quad j = 1, \dots, J. \quad (2)$$

Demand system (2) solves the problem of maximizing a specific individual (Stone-Geary) utility function subject to a linear budget constraint. The parameters β_j and γ_j are part of that utility function: γ_j indicates a necessary subsistence quantity, while β_j indicates the share of total expenditure the (representative) individual will spend on good j (in excess of γ_j), given her necessary expenses on all goods.

Three implications of consumption theory Standard consumption theory states that a rational consumer has demand functions with three specific properties:

- i. Homogeneity of degree zero in nominal variables: if we double all prices p_j and total expenditure c_h , then the demand for any good q_{jh} does not change – a property sometimes referred to as “no money illusion”.
- ii. Symmetry: the derivative of the compensated, or Hicksian, demand for commodity i with respect to the price of commodity j should equal the derivative of the Hicksian demand for commodity j with respect to the price of commodity i .
- iii. Negative semi-definiteness: the $J \times J$ substitution, or Slutsky, matrix of all derivatives of Hicksian demands should be “negative semi-definite” – this implies, among other things, that a higher price of commodity j should not increase its Hicksian demand.

On top of these, an adding-up (or feasibility) constraint says that the sum of all outlays $p_j q_{jh}$ on different goods must amount to total expenditure c_h . This constraint is already embodied in the data, in the sense that total expenditure is defined as the sum of all outlays.

In the LES (2), properties i-iii are satisfied by definition. When estimating (2), regardless of whether we use individual or aggregate real-world data, we are thus imposing the three properties on the data. Therefore, we cannot use the model to test whether properties i-iii are indeed satisfied. That is, the model cannot be used to verify that the assumption underlying the demand-system specification is consistent with actual behavior.

To investigate whether observed demand data are consistent with utility maximization, one needs a more general model. The first attempt to test whether or not properties i-iii can be rejected was carried out by Barten (1967, 1969), who estimated a generalized version of the LES, called the Rotterdam model (after the Dutch economists who first formulated it), using aggregate Dutch data. A key finding in Barten (1969) was that all three properties could be rejected. In other words, the data seemed to clash with the theory of rational consumers.

1.2 Deaton's critique

Barten's somewhat discouraging result was challenged by Deaton (1974a). In this paper, Deaton estimated a few alternative formulations of the Rotterdam model (and similar models) using aggregate UK data from 1900 to 1970. While his estimation results were similar to Barten's, a careful analysis indicated that rationality – in particular homogeneity – was not strongly rejected. Deaton further argued that the rejection of rationality could be interpreted in two ways. Either people are not rational, or people are rational but the empirical model is misspecified.

The latter point was elaborated in Deaton (1974b). Specifically, Deaton highlighted two related problems. First, the theory of consumer choice is formulated for an individual and does not necessarily hold at the aggregate level, even if all individuals are rational.¹ Second, the assumed functional forms of the LES/Rotterdam models may be too restrictive to properly capture the patterns in real-world data. That is, the Rotterdam model provided a vehicle for testing

¹ In particular, Sonnenschein (1973) had shown that even if all individuals display demand functions that satisfy conditions i-iii, the sum, or average, of all these demands need not satisfy the conditions. See Rizvi (2006) for a discussion of this issue.

rationality, but at the same time imposed too many implausible restrictions on consumer behavior to be suitable for such a test.

The problem Deaton highlighted was a serious one. For example, using demand-system analysis to evaluate the effects of tax and subsidy policies relies critically on the system accurately capturing the response of expenditure on a particular commodity j to changes in income and prices. Deaton (1974a, b) convincingly showed that the existing demand systems did not fulfill those objectives.

Deaton ends these papers in an agnostic way: we need more research based on other countries and other time periods, and maybe also on models of consumer demand that are less restrictive than the LES and Rotterdam models.

In 1980, the research on demand systems was brought to a higher level by Deaton's and John Muellbauer's Almost Ideal Demand (AID) System, to be further described below. This system was flexible enough to permit testing for rationality, it was simple enough to be easily estimated, and it allowed for expenditure on commodity j to vary non-linearly with total expenditures.

1.3 *Aggregation*

As mentioned above, it was common to postulate that the demand system (1) was derived from the choices of a single individual. If the system is derived from that agent's rational welfare maximization, theory tells us that it satisfies conditions i-iii. But an important aspect of demand-system analysis is the aggregation issue, i.e., the question of whether the demand system of different individuals can be summed up to produce a market-demand system that also satisfies those conditions.

Ever since Gorman's (1953, 1961) seminal work, the research community had looked for minimum restrictions on preferences under which one can interpret aggregate consumer behavior as the result of the behavior of a single representative consumer. If different individuals have identical preferences, but different total expenditures c_h , an aggregate model could sometimes be interpreted as describing the behavior of a representative consumer with demand q_j (the economy-wide average) and expenditure level c_r (equal, or not, to average expenditure). With Stone-Geary preferences, aggregate demand and individual demand will both have the functional form in the LES (2), but with individual demands replaced by economy-wide average demand

and individual spending c_h replaced by average spending c . This is a particularly simple form of aggregation.

In the 1970s the research on aggregation took a big leap forward, when Muellbauer (1975, 1976) derived a set of minimum restrictions on preferences to allow aggregation. Muellbauer studied PIGL (Price-Independent Generalized Linearized) preferences. A subclass of PIGL that turned out to be especially useful in demand analysis is the PIGLOG (Price-Independent Generalized Logarithmic) preferences, which have the property that each individual's budget shares $w_{jh} \equiv \frac{p_j q_{jh}}{c_h}$ can be written on the form

$$w_{jh} = a_j(p) + b_j(p) \log\left(\frac{c_h}{k_h}\right). \quad (4)$$

Here, the coefficients a_j and b_j are functions of prices, while parameter k_h could be specific to the individual household. A particularly attractive interpretation of k_h is based on the notion of equivalence scales for households of different sizes; then c_h/k_h is the household's spending per "adult equivalent".

Preferences that conform to (4) allow expenditure $p_j q_{jh}$ to be non-linear in c_h , which gives considerable flexibility compared to earlier demand systems (multiply both sides of (2) by p_j to see that the LES implies expenditure $p_j q_{jh}$ to be linear in c). In fact, Working (1943) and Leser (1963) had already shown expenditure on individual commodities to be non-linear in total expenditure.

PIGL or PIGLOG preferences guarantee that there exists a "representative" expenditure level c_r , such that an individual with expenditure $c_h = c_r$ exhibits the same budget shares as the aggregate economy. The representative expenditure c_r is equal to the economy-wide average \bar{c} in the case of the LES but, in general, depends on higher moments of the distribution of c_h across individuals. The representative agent's demand functions also satisfy conditions i-iii, i.e., the behavior of the hypothetical representative agent satisfies the rationality postulates.

An attractive feature of (4) is that individuals do not have to have identical preferences; the individual demand functions could differ somewhat, as captured by parameter k_h , but there would still be a representative consumer.

1.4 The Almost Ideal Demand System

The AID System of Deaton and Muellbauer (1980a) builds on PIGLOG preferences. With a particular choice of functional forms for parameters $a_j(p)$ and $b_j(p)$, such preferences give rise to aggregate budget shares of the form

$$w_j = \alpha_j + \sum_k \gamma_{jk} \log p_k + \beta_j \log \left(\frac{c_r}{P} \right). \quad (5)$$

Here, P is a price index, while parameter β_j indicates whether good j is a luxury ($\beta_j > 0$) or a necessity ($\beta_j < 0$), and γ_{jk} shows how demand for good j responds to the price of good k .

The AID System (5) has a number of advantages relative to earlier systems. It allows expenditure $p_j q_j$ to be non-linear in c . It can be used to test for rationality; homogeneity, symmetry and negative semi-definiteness imply restrictions on parameter vectors α , β and γ , and by estimating (5) one can check if these restrictions are satisfied in real-world data. The AID system aggregates over consumers under weaker assumptions about preferences than did earlier models. The average budget share w_j for good j can be regarded as the demand by a representative consumer with expenditure

$$c_r = \exp \left(\frac{\sum_h c_h \log \left(\frac{c_h}{k_h} \right)}{\sum_h c_h} \right),$$

where subscript h indicates an individual consumer or household and where the k_h denote the parameters already discussed in connection with (4). Although all individual consumers are assumed to have preferences of the same functional form, their preferences can differ by parameter k_h . As the latter can be interpreted as demographic indicators, the system allows for heterogeneity with respect to age composition or household size.

Related research At about the same time, another group of researchers – Jorgensen, Lau and Stoker (1982) – introduced the Transcendental Logarithmic Demand (Translog) Model.²

Translog has many properties in common with the AID System, including expenditure shares that are linear in log expenditure, and has similar aggregation properties. While the Translog Model and the AID System have both been extensively used in the literature, the latter has

² The Translog model builds on earlier work by Jorgenson, Lau, and Stoker (1980). The indirect Translog utility function was introduced by Christensen, Jorgenson, and Lau (1975).

attracted more attention. This popularity reflects, at least partly, the fact that the AID System is easy to estimate: it has a very convenient approximate form which may be easily estimated by OLS, a property that was particularly important when computation was less developed than it is now. Subsequent research has built more general demand systems that nest both the AID System and Translog Model as special cases.³

Estimating the AID System Deaton and Muellbauer obtained statistical results for aggregate UK data that were consistent with earlier studies: the three restrictions imposed by theory (homogeneity, symmetry, negative semi-definiteness) were rejected. In line with Deaton (1974a,b), they reached a different conclusion to the earlier literature: the reason for the rejection is more likely to be a misspecified model rather than non-rational consumers. At the end of the paper, and in a textbook published at the same time (Deaton and Muellbauer, 1980b), they highlight a number of possible misspecifications of the model – and of all existing demand models at the time. Specifically, they note that the basis of a satisfactory theoretical and empirical analysis of consumer demand should be the careful aggregation of heterogeneous consumers. They also note that other variables than current prices and current total expenditure must be added to the demand system (e.g., possible credit constraints) to account for observed demand patterns in a theoretically coherent and empirically robust way. Deaton and Muellbauer’s discussion of the AID System thus emphasizes its shortcomings and the need to develop it further. They conclude that their system “...with its simplicity of structure, generality, and conformity with theory, offers a platform on which such developments can proceed” (Deaton and Muellbauer, 1980a, p. 323). This turned out to be a very good forecast. The AID System spurred a large body of work, both by the authors themselves and by numerous followers. The versatility and usefulness of the system is attested to by an enormous literature since its publication. Even after 35 years, it remains a cornerstone of demand estimation around the world, whether that estimation is based on aggregate, individual or household-level data.⁴

³ See Lewbel (1989) for a review.

⁴ The Deaton and Muellbauer paper was picked out by Arrow et al. (2011) as one of the 20 most influential articles published by the *American Economic Review* in its first 100 years.

1.5 *Subsequent research*

Evidence from numerous studies indicates that expenditure shares vary nonlinearly with total expenditure. For some commodity groups, like food, expressing their shares as linear functions of log expenditures – as in the original AID System (5) – is an accurate approximation. For other commodity groups, however, Engel curves display other types of curvature and a more flexible model is needed. Extending the model to Engel curves with a quadratic term in income – the “Quadratic Almost Ideal Demand System” (the QAID System) of Banks *et al.* (1997) – has been shown to provide such flexibility.

There is also clear evidence that expenditure shares vary systematically with observable demographic covariates of the household, including the number of children and the age of its members. Nowadays, a proper treatment of such heterogeneity is at the core of demand-system analysis (see Blundell and Stoker, 2005, for a review).

The analysis of consumer demand has also been extended to deal with the aggregation problem – i.e., how can one draw inferences about the structural relationship between average demand, average income, and possibly average demographic composition from a flexible demand model, which is estimated on individual data with heterogeneous consumers (again, see Blundell and Stoker, 2005). A related problem is how to assess possible bias in the estimates of structural price and income coefficients from aggregate data (Blundell, et al. 1993).

The AID System, and its further extensions, like the QAID System and the quasi-AID System (Lewbell and Pendakur, 2009), have been used in a large number of applications. For instance, it has been very influential in the area of agricultural economics. It has also been highly influential in applied work on CPI measurement (e.g., Hamilton, 2001, and Costa, 2001), international welfare comparisons (e.g., Almås, 2012, and Neary, 2004), and estimates of within-country inequality and poverty (Pendakur, 2002). Its flexible functional form combined with the desirable properties of the PIGL or PIGLOG class has also proved to be very useful in a macroeconomic (dynamic general equilibrium) context – see, e.g., Boppart (2014).

One can use the AID (or QAID) System to analyze the responses and distributional effects of tax reforms – like changes of VAT, Pigouvian taxes, and trade taxes.⁵ Demand systems incorporating demographic factors are also helpful in understanding how changes in the composition of the population affect consumer demands over time. Practical applications abound, in policy guidance from statistical agencies, in reports by think-tanks, as well as in policy papers by individual researchers. This way, the system has had a significant influence on policy.

To sum up, Deaton has been a transformative scholar in the research on demand systems. He has identified important limitations in earlier systems. Together with Muellbauer, he has developed a model that overcomes some of these limitations, a model that forms the basis of modern demand analysis. He has also highlighted many limitations in the AID System – and, in so doing, he has set a rich agenda for research on consumer demand.

2. The study of consumption over time

2.1 Background

The demand-system approach discussed in Section 1 implicitly regards the individual's consumption decision as a two-stage procedure. First, she decides how much to spend on consumption at each point in time. Having set the budget – total expenditure c – she then decides how much to spend on different commodities, or commodity groups. While the latter decision can be captured by the demand system (5), the determination of c is taken as given.

In a series of contributions during the 1980s, Deaton set out to study the first-stage decision: how much of one's income or wealth to spend on consumption (or, equivalently, how much to save) at each point in time.

In the 1930s, when Keynes considered the evolution of consumer spending (regarded as an aggregate good) across time, he assumed that people consume a constant share of a marginal change in income. In the 1950s and 1960s, aggregate consumption analysis became more

⁵ On tax policy, see e.g., Crawford, Keen and Smith (2010) in the Mirrlees Review of the UK tax system. On trade reforms, see e.g., Atkin (2013).

realistic with the introduction of the Permanent Income Hypothesis (PIH) of 1976 Economics Laureate Milton Friedman, and the closely related Life-Cycle (LC) Model of 1985 Economics Laureate Franco Modigliani. These were able to explain the empirical finding that savings rates vary with short-run income fluctuations, but not with long-term income trends. The explanation was that consumers save part of a temporary increase in income, so as to smooth consumption over time. Similarly, the new models could explain why savings increase in income, within a cross-section of individuals. By the 1970s, the general view in the economics profession was that these models matched the data well, and they were used extensively in the macroeconomic literature. While the models were formulated in terms of individual choice, they were empirically evaluated by studying the properties of aggregate data for consumption and income.

In the late 1970s, research on consumption and saving was rejuvenated by a new generation of scholars. Building on the rational-expectations revolution, and explicitly modeling the household's intertemporal optimization problem, Hall (1978) and MaCurdy (1981) showed that at the optimum, the marginal utility of consumption is equated to the marginal utility of wealth, and that the latter follows a martingale (a generalized random walk).

The Euler equation is a necessary condition for optimality, when individuals maximize a discounted sum of per period utility. For any asset a that the agent uses to transfer resources between periods, the Euler equation states:

$$u'(c_t) = \beta E_t[u'(c_{t+1})(1 + r_{t+1}^a)].$$

As before, c_t denotes total consumption expenditure in period t , and $u'(c_t)$ is marginal utility in that period t . Further, β is the subjective discount factor, E_t denotes expectations conditional on information in period t , while r_{t+1}^a is the return on asset a between periods t and $t+1$. Given (possibly stochastic) asset returns, this condition can be used to derive implications for the stochastic properties of consumption. Conversely, given a stochastic path of consumption, the condition has implications for asset returns. These latter implications include the Consumption CAPM model developed by, among others, Douglas Breeden (1979) and 1995 Economics Laureate Robert Lucas (1978).

Robert Hall (1978) focused on the implications for consumption, under the assumption that individuals have access to a single asset with constant return r . Denoting the forecast error, $E_t[u'(c_{t+1})] - u'(c_{t+1})$, by v_{t+1} the Euler equation can be written as

$$u'(c_t) = \beta(1+r)(u'(c_{t+1}) + v_{t+1}). \quad (6)$$

By the assumption of rational expectations, the forecast error v_{t+1} is zero in expectation and independent of all information known at t . In other words, no variable known in period t should predict v_{t+1} . Writing (6) as

$$\frac{1}{\beta(1+r)}u'(c_t) - u'(c_{t+1}) = v_{t+1},$$

one can easily compute the unobservable v_{t+1} from the observable variables c_t and c_{t+1} and check whether it really is unpredictable.

This insight stimulated a large amount of empirical research on consumption and labor supply in the 1980s. Deaton was a key contributor to this field, and his work was crucial in shifting the focus from aggregate to individual consumption data.

2.2 Theory and evidence on consumption over time

Blinder and Deaton (1985) was one of several papers – including the pioneering studies by Flavin (1981) and Hall and Mishkin (1982) – to investigate whether aggregate consumption and income data could be reconciled with the representative-agent version of the Hall (1978) PIH/LC model. They confirmed the earlier results that the random-walk implication of that model could be rejected with a high degree of confidence. More specifically, contrary to the implications of the PIH/LC model, real-world aggregate spending growth responded strongly to expected movements in aggregate income, wealth, and other variables. Thus, the error term v_{t+1} (6) could be predicted, a violation of the Euler equation under the assumption of rational expectations and of a constant rate of return r .

Permanent vs. current income and consumption In addition to its specific conclusions, the Blinder-Deaton study made Deaton realize theoretical implications of the PIH/LC model that had not been noted before. One such insight appeared in Deaton (1987): under rational expectations, the PIH/LC model applied to aggregate data implies that consumption should be *less* smooth than income. This was startling to economists, who had thought that consumption should be

smoother than income as it was, and is, in real-world data. The reasons were two. First, concave utility functions imply consumption smoothing if income varies randomly. Second, permanent income is close to average income over time, and thus transitory income fluctuations should even out when computing permanent income. The latter, however, turned out to be an incorrect conclusion when one takes the PIH/LC model seriously in combination with *aggregate* real-world income data. Specifically, Deaton (1987) argued that it is unreasonable, theoretically as well as empirically, to assume that income always tends back to a never-changing deterministic trend. In other words, income shocks have a permanent component. The simplest example is when income follows a random walk – i.e., income shocks are uncorrelated over time. In this case, the best guess about future income is current income, meaning that permanent income is equal to current income. If consumption is equal to permanent income, the two must have the same variance.

Deaton (1987) also demonstrated that aggregate consumption data did not respond to changes in the interest rate in a way that is consistent with the theory of a representative consumer.

Allowing time-varying safe interest rates, r_t , the Euler equation becomes:

$$u'(c_t) = \beta(1 + r_t)E_t[u'(c_{t+1})].$$

This equation clearly implies that expected consumption growth should be positively related to the interest rate. But Deaton showed that such a relation does not hold in the aggregate data. He then provided a simple, yet striking, example of why this cannot be taken as evidence against *individual* data being consistent with a positive relation between consumption growth and interest rates. Consider two hypothetical economies, each consisting of individuals with finite lives. In every period, old people die and young people are born. Both economies have constant aggregate income over time, but one has a higher interest rate. By construction, aggregate income and consumption are constant, but individual income can grow faster in the high-interest economy: newborn agents start out with a lower consumption level than in the low-interest economy, but let it grow faster. Clearly, aggregate data cannot recover the relation between individual consumption growth and the interest rate.

Variability in aggregate consumption and income The logic behind the relation between income and consumption variability was further developed and empirically evaluated by Campbell and Deaton (1989), who still relied on aggregate data. Their analysis is straightforward. Consider a

setting with infinite horizon and a constant discount rate $r > 0$. Define permanent income y^p as the annuity of the discounted present value of all expected future incomes:

$$y_t^p = \frac{r}{1+r} E_t \left[\sum_{s=1}^{\infty} \left(\frac{1}{1+r} \right)^s y_{t+s} \right]. \quad (7)$$

Here, E_t is an expectations operator in period t and y_{t+s} is current income in period $t + s$. Next, suppose income follows the stochastic process $A(L)y_t = \varepsilon_t$, where $A(L)$ is a lag polynomial and ε_t is an i.i.d. process. Then (7) can be written as

$$y_t^p = \frac{r}{1+r} \cdot \frac{y_t}{A\left(\frac{1}{1+r}\right)}. \quad (8)$$

We can now illustrate the importance of specifying the income process in a correct way. Suppose first that the income process is AR(1):

$$y_t = \lambda y_{t-1} + \varepsilon_t \quad (9)$$

for $\lambda > 0$. In this case, $A(L) = 1 - \lambda L$, and from (8) permanent income becomes:

$$\bar{y}_t^p = \frac{r}{1+r} \cdot \frac{y_t}{1 - \frac{\lambda}{1+r}} = \frac{r}{1+r-\lambda} y_t. \quad (10)$$

Clearly, for $\lambda < 1$ permanent income is less variable than current income, since then $\frac{r}{1+r-\lambda} < 1$.

If consumption equals permanent income, consumption becomes less variable than current income, as in aggregate consumption and income data. There is a crux, however. Campbell and Deaton argued that the AR(1) process (9) does not properly describe the evolution of aggregate income observed in the data, and that this has very important consequences for the evolution of (unobservable) permanent income. A better approximation than (9) of aggregate actual income is to model its growth rate as

$$\Delta y_t = \lambda \Delta y_{t-1} + \varepsilon_t,$$

with $0 < \lambda < 1$ – i.e., periods of high (low) income *growth* tends to be followed by periods of high (low) but less high (low) income growth. In this case, $A(L) = 1 - (1 + \lambda)L + \lambda L^2$ and applying formula (8) for permanent income yields

$$y_t^p = \frac{r}{1+r} \cdot \frac{y_t}{1 - (1 + \lambda)\left(\frac{1}{1+r}\right) + \lambda\left(\frac{1}{1+r}\right)^2} = \frac{1+r}{1+r-\lambda} y_t. \quad (11)$$

The implication when permanent income is given by (11) rather than (10) is very different. Now, $\lambda > 0$ means that permanent income varies *more* than current income (as $\frac{1+r}{1+r-\lambda} > 1$). The intuition is clear: if income growth is unexpectedly high in period t , temporarily high growth is expected to continue for some time. Because income will keep growing faster than it was expected to before t , the shock to permanent income is larger than the shock to current income.

In the quarterly data used by Campbell and Deaton (1989), $\lambda = 0.442$ implying that $\frac{1+r}{1+r-\lambda}$ is around 1.8 for plausible values of r . Permanent income is thus almost twice as variable as current income – and if consumption is equal to permanent income, it is almost twice as variable as current income. However, the data show that consumption is less variable than (current) income.

The Campbell-Deaton result came as a great surprise to the economics profession. But the result was theoretically impeccable and empirically founded on the actual time-series properties of income data. The apparent excess smoothness of aggregate consumption, relative to PIH/LC theory, became known as the Deaton paradox. The real impact of this discovery was not so much to convince everybody that real-world aggregate consumption was too smooth. Rather, it was to convince everybody that the time-series properties of the income processes individuals face are crucial to understanding intertemporal consumption patterns.

Changing the focus to individual data The empirical rejection of the PIH/LC model on aggregate data can be explained in three ways. First, the theoretical model could be wrong, consumers may not be rational. Second, working with aggregate data and the notion of a representative consumer is problematic; individuals may be rational, but the conditions for aggregation are not satisfied in reality. Third, rational individuals may be confronted with many constraints not accounted for in the models; one example is borrowing constraints, which have obvious implications for behavior and aggregation. In his summary of the literature on consumption, Deaton (1992) argued that to separate these possibilities, “progress is most likely to come when aggregation is taken seriously and when macroeconomic questions are addressed in a way that uses the increasingly plentiful and informative microeconomic data”. Thus, he argued that one should study individual consumption, where observed behavior may reflect a degree of consistent rational choice, and to explicitly address the aggregation issues. We see a clear parallel here to the insights in Deaton

and Muellbauer (1980b) about the avenue to future progress in the research on demand systems discussed in Section 1.

The enduring influence of this insight is seen in the sheer volume of research that macroeconomists now devote to the analysis of income dynamics. Nowadays, macroeconomics is not only about studying the dynamics of aggregate variables, it is about studying the dynamics of “the entire equilibrium distribution of allocations across individual economic actors.”⁶

Although we are still far from definite answers, the Deaton (1987) and Campbell and Deaton (1989) papers were instrumental in starting this transition.

2.3 Individual consumption over time

The study of individual data posed new challenges. To take proper account of intertemporal aspects, one would ideally need panel data where individual households can be followed over time. Also, to understand a rational individual’s intertemporal choice in a more realistic environment, one would need to introduce other constraints than the standard budget constraint, and add microeconomic labor-income uncertainty. On both these accounts, Deaton made important contributions in the 1980s and 1990s.

Pseudo panels True panel data for consumption are scarce today – and were even scarcer in the 1980s. However, many countries had repeated cross-sectional household surveys. Because new samples were drawn for each survey, one could not use these data to track individual households over time. In an important breakthrough, Deaton (1985) showed how to construct a “pseudo panel” made up of cross-sectional individual data for cohorts of individuals or households of the same age. His idea was that with enough data, successive surveys would generate successive random samples of individuals from each cohort. Summary statistics from these random samples generate a time series from which one can infer behavioral relationships for the cohort as a whole, just as if panel data were available. Deaton also showed that pseudo panels do not necessarily produce worse results than true panel data, as they do not suffer from the attrition problems that plague most panel-data surveys. Households move, and may quit the panel in systematic ways, which introduces bias in the data. On top of that, there are measurement errors.

⁶ Quote from Heathcote et al. (2009). See also Attanasio and Weber (2009) for a review of the literature.

In fact, studying cross-section data for different cohorts might mitigate the measurement-error problem when studying the dynamics of consumption over time.

Deaton's work on panel data had practical implications. In the mid-1980s, the World Bank had decided to invest in a large-scale household survey program and a key question was whether to collect panel data or not. In line with discussion above, Deaton argued that repeated cross-sectional data are not only cheaper and easier to collect than panel data, but may also be preferable in practice. Thus, depending on the purpose, panel data may sometimes not be worth the trouble and the cost of collection. This work also illustrates one of Deaton's major virtues – he studied a practical problem rigorously in order to inform an issue of great interest to policy-makers (here, the World Bank).

Individual income and liquidity constraints Another challenge to studying individual consumption was theoretical, namely how to solve for individual consumption in the presence of idiosyncratic uncertainty and liquidity constraints. Deaton (1991) addressed this challenge in an influential paper. This paper nicely combines the issues of optimal consumption under liquidity constraints, and the aggregation of individual behavior under idiosyncratic as well as aggregate shocks. It also offers a key to the solution of the Deaton paradox.

The Euler condition discussed earlier in this section provides a prediction for consumption levels in adjacent periods. Little information is required to test this prediction. An econometrician does not need to know the income and assets of the individual, but only needs data for her consumption. This is a key advantage of the approach – but a strong disadvantage if the aim is to understand how consumption relates to income and assets. In that case, theory has to be used to derive a consumption function that maps exogenous processes (like income) and state variables (like assets) into consumption choices.⁷ It was known that such analytical consumption functions could not be derived for realistic assumptions about risk and preferences.

Zeldes (1989) had used numerical methods to derive consumption functions and showed that, for a household with little financial wealth relative to income, labor-income risk strongly affects consumption through a precautionary savings motive (see Leland, 1968, and Sandmo, 1970, for

⁷ Of course, treating income as exogenous is a simplification. In later research, the challenge has been to allow for endogenously determined income processes where individuals choose labor supply in a general-equilibrium framework.

early treatments of precautionary savings). That income risk created deviations from the PIH was thus known to the profession around 1990, but the precise implications were unknown. Zeldes (1989) had assumed that the subjective discount rate was equal to the interest rate on savings and borrowing, i.e., $\beta = 1/(1 + r)$, so that $\beta(1 + r) = 1$. Under this assumption, households have an incentive to accumulate assets through (precautionary) savings, so as to move away from any borrowing constraints.

Deaton (1991) notes that existing empirical evidence showed that consumption responds to predictable changes in income, violating the Euler equation. He also notes that many U.S. households have little financial wealth and are therefore unlikely to be able to finance consumption by borrowing. A key assumption behind the Euler equation (6) – namely that all households can save or borrow at a given rate r – may thus be violated for many households. Deaton aimed at deriving a consumption function in the presence of borrowing constraints. Such constraints interact with the precautionary motive, creating stronger motives for saving. If a large share of consumers do not have many financial assets despite these strong savings incentives, Deaton (1991) argued that a substantial share of households may be less patient than when $\beta(1 + r) = 1$. Therefore, he focused on the case when $\beta(1 + r) < 1$, which would imply that many households choose a spending level close to the borrowing constraint despite the possibility of moving away from that constraint by accumulating assets.

Using theoretical results that would appear in Deaton and Laroque (1992), Deaton (1991) relied on numerical methods to derive consumption functions for realistic aggregate-income processes, showing that theoretically implied consumption behavior was grossly at odds with aggregate consumption data. He concluded that the representative-agent framework had to be replaced by explicit modeling of individual consumption before aggregation.

In particular, Deaton stressed that individual income shocks are less persistent than aggregate ones. While some consumers see their incomes rise in a particular time period, others see them fall. Such idiosyncratic shocks wash out in the aggregate, so that aggregate income is more persistent than individual income. This implies that aggregate consumption, derived as the sum of individual consumption responses to individual income – according to the standard theory – behaves very differently than if it were derived from the behavior of a hypothetical representative household facing the aggregate income process. Deaton (1991) showed that the

individual consumption behavior implied by an income process with both aggregate and idiosyncratic components, when consumers also face borrowing constraints, had much better potential to match the observed patterns in micro, as well as in macro data. These results offered one key to the resolution of the Deaton paradox. Although it has proven quite difficult to exactly pin down the stochastic properties of individual income, the literature on this topic is still lively (see e.g., Guvenen 2007 and 2009).

Deaton (1991) was a partial-equilibrium analysis, which took the interest rate to be exogenous.⁸ This was in line with Deaton's research interest: understanding individual consumption. However, analyzing the economy's general equilibrium in a setting with idiosyncratic uninsurable income shocks and borrowing constraints has since become a central part of modern macroeconomics (see e.g., Imrohoroğlu, 1989, Huggett, 1993, Krusell and Smith, 1998, Heathcote et al. 2009).

Individual consumption and insurance Townsend (1994) and Deaton and Paxson (1994) initiated an important literature, which used individual consumption data to infer how much income insurance individuals have access to. This issue is important in understanding consumption, but it is also important for welfare and policy analysis. If informal insurance arrangements exist and are efficient, the need for publicly provided insurance is smaller and introducing such programs may crowd out existing arrangements.

Deaton and Paxson's idea is that whenever individuals can fully insure against idiosyncratic income shocks, the marginal utilities of all individuals should evolve in parallel over time. Assuming sufficiently stable preferences, individual consumption paths should also evolve in parallel. On the other hand, if insurance is missing, the permanent component of idiosyncratic income shocks will lead to increasing consumption inequality as the members of a cohort age. This implication was strikingly verified in Deaton and Paxson's empirical study of data for the U.S., U.K. and Taiwan. They also noted that a study of the joint characteristics of individual income and consumption can help evaluate the availability of insurance households against income shocks. Neither the complete insurance assumption nor its opposite – complete absence of insurance – is realistic. As insurance comes in many forms, formal as well as informal,

⁸ Bewley (1977) provided an early general-equilibrium analysis of an economy with borrowing-constrained consumers.

through friends and family, a structural modelling of these insurance mechanisms is difficult. Deaton and Paxson (1994) thus envisioned “the construction and testing of market models with partial insurance” on panel data of consumption and income to be a fruitful approach – see e.g., Heathcote et al. (2007) and Blundell et al. (2008) for subsequent examples of such analyses.

To sum up, the hallmarks of Deaton’s research on consumer behavior over time is the interplay between theory and data and the interplay between individual behavior and aggregate outcomes. His work on consumption and saving showed that aggregate data could only be reconciled with theory if one took individual consumption behavior based on individual income processes seriously. In this way, he played a fundamental role in shifting the macroeconomic literature on consumption and savings towards micro-based empirical models.

3. Welfare in developing countries

3.1 Background

In the 1980s, many scholars thought that development economics had hit a major roadblock. The theories at hand to explain development, or lack thereof, lacked the rigor that characterized most other fields of economics. Only scarce data were available to test development theories and assess policies for how to best bring people out of poverty. Moreover, the empirical analyses of consumption allowed by existing data – essentially national accounts – did not help to gauge the economic wellbeing of households, especially poor households. Basic questions on the extent of poverty, and how poor households respond to their economic and physical environments, could not be answered credibly due to the lack of data and appropriate analytical methods.

This situation has now changed. Today the subject is dominated by microeconomic research based on high-quality household data (and other micro data). Research over the last two decades has improved our understanding of the essential mechanisms at play in developing countries, and has provided a large number of insights that are highly relevant for policy-making.

Angus Deaton has played a fundamental role in this transformation. Just like his contributions to demand analysis and to aggregate consumption and savings, his work on development spans a wide spectrum of methodological and substantial contributions: from details on how to measure

poverty – in specific countries and globally – to pioneering empirical analyses of micro data from household surveys.

3.2 Empirical analysis of household survey data

Household survey data In the 1980s, Deaton championed the use of household surveys as a means to measure poverty and living standards, and to identify their determinants. He shaped the new agenda through his work with the World Bank. In the early 1980s, the Bank was contemplating a major effort in collecting household survey data – the Living Standards Measurement Study (LSMS). Deaton played an important role in this endeavor, writing several background papers for the project, including how to measure household welfare, the most important issue in the design of the household-survey program (Deaton, 1980, 1981). He was also instrumental in ensuring that expenditure data were collected as the basis for welfare measurement. Consumption – as opposed to income – remains the core variable in poverty measurement in low-income countries, because it is often easier to measure and provides a more accurate measure of material welfare when income varies seasonally throughout the year (see e.g., Deaton and Grosh, 2000, for a review).

As household-survey data became increasingly available in the late 1980s, Deaton became a pioneer in empirical research, often grounded in solid demand analysis that was based on such data. Two examples, Subramanian and Deaton (1996) and Deaton (1989), are discussed further below. In Deaton (1997) – a standard reference for household survey-based research on poverty and development – he showcases how household survey data can offer a useful guide to individual behavior in areas crucial for human wellbeing. The book discusses a number of Deaton’s methodological and substantive contributions and lays a foundation for much of the subsequent household-survey research

Income and calories Subramanian and Deaton (1996) use traditional Engel-curve analysis to investigate the relationship between income and nutritional status, as measured by calories consumed, in low-income countries. Such an investigation is important for at least three reasons. First, poverty is closely related to whether people get enough to eat; documenting the living standard of the poor becomes a question of nutritional adequacy, which is well assessed by household surveys. Second, knowledge about the relationship between nutrition and income is important for designing policies to reduce poverty. If the elasticity of calories with respect to

income is high, then economic policies that foster growth (assuming the distribution of income does not dramatically change) will raise the nutritional status of the poor and eliminate hunger. On the other hand, if the elasticity of calories with respect to income is low or close to zero, as a number of papers in the early 1990s had suggested (individuals may consume tastier but not necessarily more nutritious food as income rises), policymakers may want to shift focus from growth to providing basic needs. Third, at least since the work of Leibenstein (1957), economists have been interested in the reverse link from nutritional status to productivity, and thereby income. Specifically, the theory of nutritional-based efficiency wages postulates that productivity depends positively but non-linearly on nutrition. If this is so, those who do not get enough food (calories) may not be productive enough to be employed even below the current market wage. As a result, they are stuck in an unemployment/poverty trap, with low productivity and low income.

Using detailed household survey data from India, Subramanian and Deaton (1996) showed the total food elasticity to be close to 0.8, while the calories elasticity was about half that number and declined slowly with income. In other words, as income rises the consumption of food increases, but as households substitute products like cereals with dairy and meat products, the cost of calories increases. Despite this substitution effect, the policy conclusion is clear: policies to foster income growth among the poor will reduce malnutrition.

While their main line of inquiry takes nutrition to be determined by income, Subramanian and Deaton (1996) also provide indirect evidence on the strength of the reverse link from malnutrition to poverty. Specifically, they document that the calories necessary for a day's activity cost less than 5 percent of the daily wage. This makes it quite implausible that malnutrition explains poverty, as suggested by the nutritional-based efficiency-wage theory. Rather, malnutrition is the consequence of poverty.

Within-family discrimination Another example where household-survey data can shed light on a question of high policy relevance is Deaton's (1989) work on discrimination within the household. While there is strong evidence that sons are favored over daughters in many developing countries – the phenomenon of “missing women” (see e.g., Anderson and Ray, 2010) is probably the starkest example – the mechanisms through which this discrimination comes about are unclear. One possible mechanism is that girls are systematically provided with fewer resources than boys. Testing this mechanism is difficult, however, because household survey

data rarely contain information on each individual household member's consumption. To overcome this measurement problem, Deaton proposed an ingenious way to use household consumption data to indirectly estimate whether girls are given less resources than boys. The idea was simple. When a child is born, the household – in effect – becomes poorer, simply because there is then one more mouth to feed. Measuring how the consumption of “adult goods”, such as adult clothing, alcohol, and tobacco, drops when a child is born thus provides an indirect estimate of the “cost” of a child. If the household cuts its expenditure on adult goods by less when the child is a girl rather than a boy, this is evidence that the household discriminates against girls. Using household surveys from several developing countries, Deaton was not able to find any systematic differences under normal circumstances. Later work has confirmed this finding, but also found evidence of clear discrimination in spending patterns when households face adverse circumstances.⁹

3.3 *Measuring poverty*

Data on consumption is normally collected at the household level, while poverty is naturally measured at the individual level. This raises the important question of how to compare individuals in households of different sizes and compositions, when consumption data are aggregated at the household level. In principle, this could be handled by applying a system of weights to create the number of adult equivalents, by counting children of a certain age as some fraction of an adult. Constructing appropriate weighting scales to account for different household structures has a long history in economics, but there is little consensus on the best approach. At the same time, the choice of equivalence scale may largely affect measured poverty. If individual welfare is measured as total household expenditures per capita, as is often the case, and if children require less of most things than do adults, the extent of poverty among children (or among households with children) is overstated by per-capita estimates. In a series of papers, including Deaton and Muellbauer (1986), Deaton made important contributions to this literature.

Deaton and Muellbauer (1986) deals with the comparison of welfare levels between households of different sizes and analyzes the measurement of child costs. They show that two of the simplest and most popular methods – the Engel and Rothbarth methods – of measuring child living costs yield rather different empirical estimates, and provide both theoretical and empirical

⁹ See Duflo (2012) for a review.

arguments that the cost plausibly lies between those obtained by these two traditional approaches.¹⁰ Using Sri Lankan and Indonesian data, and a modification of the Rothbarth method, their results suggest that child costs are about 30-40 percent of per capita adult expenditures.

Another crucial issue for welfare measurement is how to treat goods with different prices and qualities. Prices of the poor's consumption basket are often not available in household surveys, nor is information on the quality of the goods they consume. However, household surveys typically collect data on clusters of households in the same village, and households within such a cluster can be assumed to face the same market prices. Moreover, many household surveys collect data on expenditures and physical quantities – dividing the two provides unit values. These unit values suggest substantial spatial price variation in many developing countries, a result that may be explained by high transport costs. Can one use these unit values as direct substitutes for true market prices to analyze demand patterns and measure poverty? The general answer is no, as individuals consume not only different quantities but also different qualities of goods, and unit values are also measured with error. To investigate the question, Deaton (1988) begins from a theoretical model of quality and quantity choice. He then applies the model predictions in a concrete setting to household-survey data from the Republic of Côte d'Ivoire, and shows how such unit values, when corrected for quality effects and for measurement error, provide information on the underlying spatial variation in prices. Variation in local prices, in turn, can be matched to variation in demand patterns so as to estimate price elasticities. A key methodological finding of the paper is thus to show that unit values can be used to back out information about local market prices when the former are available and the latter are not.

Deaton (1988) has had a large influence on practical poverty measurement. India – which hosts one third of the world's extreme poor (Chen and Ravallion, 2010) – revised existing price estimates in the mid-2000s, partly in response to Deaton's critique of the previous methodology (Deaton and Tarozzi, 2000; Deaton, 2003b, c, 2008). The new methodology for obtaining price estimates follows directly from Deaton (1988). As a result of the revised method, rural poverty was found to be significantly higher than in previous estimates, leading to a politically charged

¹⁰ The theory of equivalence scales is reviewed in the textbook by Deaton and Muellbauer (1980b).

debate about the method for poverty measurement and the relationship between poverty and growth, a debate in which Deaton has actively participated.

3.4 *Welfare comparisons across time and across countries*

One of Deaton's central interests has been to compare welfare, across time and across countries. This agenda is at the heart of many practical policy debates.

Growth, poverty and welfare Does economic growth reduce poverty? If growth raises income for everyone in a proportional way, the fraction of the population below some fixed poverty line must fall, with the rate of absolute poverty reduction depending on the share of the population near the poverty line. If, on the other hand, economic growth is unequally distributed, then growth may do little to reduce poverty.

The empirical research on this topic has reached somewhat conflicting results. The findings in Deaton (2005) help us understand why. He discusses an important stylized fact: aggregate consumption measured in household surveys, which is used to measure poverty, grows less rapidly than aggregate consumption measured in national accounts. The household-survey data thus give the impression that poverty has not fallen very much. Deaton points out a number of reasons why consumption (and income) based on household surveys and national accounts differ and why neither of the data sources may accurately capture the true growth in consumption. For example, national accounts consumption excludes services that are not exchanged in a market (e.g., food preparation). But many of these services tend to be replaced by exchanged services as people become better off and markets for production in the home develop, causing the measured growth rate of consumption to be overestimated. Survey-based estimates of consumption, on the other hand, will understate mean consumption (and overstate the fraction of people in poverty) because rich households are less likely to participate in household surveys than poor people. Deaton concludes that current statistical procedures in low-income countries probably understate the global rate of poverty reduction and overstate the global growth rate.

World poverty Related to these measurement issues is the question of how to best provide accounts of global poverty. Deaton and Dupriez (2011) construct purchasing power parity (PPPs) exchange rates for the poor (PPPPs) from household surveys in 62 developing countries. They then use these poor-specific exchange rates to calculate global poverty lines and poverty counts

using specific weights corresponding to the consumption baskets of the poor. The results show that the difference between regular PPPs and PPPPs is not of great importance for estimating global poverty.

Deaton devoted his Presidential Address to the American Economic Association (Deaton, 2010) to the measurement of world poverty, summarizing the issues he had studied in the past decade. The most recent estimates of world poverty (Chen and Ravallion, 2010) show that there has been a dramatic reduction in poverty over the last 25 years. But updated poverty numbers, based on a larger set of household surveys and PPP exchange rates from the 2005 round of the International Comparison Project (ICP), indicate that the number of global poor has increased by nearly half a billion compared to previous estimates (for the same years). Deaton shows that this increase in the number of globally poor had little to do with the revision of the PPP exchange rates, as many had believed. Rather, it was largely driven by a higher global poverty line, because India – with a relatively low poverty line, even compared with countries with much lower per-capita incomes – was dropped from the set of poor countries forming the basis for the global poverty line. As the paper says: “In effect, India and the world have become poorer because India has become richer” (Deaton, 2010, p.16). Deaton concludes his address with a plea for greater use of self-reported wellbeing measures.

To sum up, Deaton championed the use of household survey data in developing countries, especially data on consumption, to measure living standards and poverty. He showed how such data can shed light on important development issues. Careful microeconomic analysis has become a cornerstone of modern development economics and Deaton has been an important driver of this change. His work has had immense practical significance and greatly increased our understanding of what can be learnt from poverty comparisons across time and place.

4. Related contributions

While consumption certainly produces wellbeing, it is not its sole determinant. In some of his most recent work, Deaton has considered other determinants as well.

Health People’s health is not only instrumental in determining their income and consumption possibilities, but is also an important element of wellbeing in its own right. Thus, to measure and

understand wellbeing requires grasping the relationship between health and income, as well as the distributions of these in the population. Deaton has made important contributions here, including his analysis of the connections between inequality and health (Deaton, 2003a).

Subjective wellbeing As mentioned in Section 3, Deaton (2010) concludes with a plea for greater use of self-reported wellbeing measures. As an expert for the Gallup World Poll, he has been involved in the design of surveys for this purpose. Moreover, he has used the Gallup data to analyze how subjective measures of wellbeing vary across people in different societal groups and in different countries.¹¹

5. Final remarks

Angus Deaton has made major improvements to the theory and measurement within three fields: consumption demand systems, the fluctuations of consumption over time, and the measurement of consumption and poverty in the developing world. In each field, he has written a major textbook or monograph – Deaton and Muellbauer (1980b), Deaton (1992), and Deaton (1997) – to share his deep insights with the research community. These three books – and the research they draw on – have served as “door-openers”, each setting a clear agenda that subsequent research has largely followed. While the three fields are quite different, Deaton’s research on them shares the same general themes. He has consistently tried to bring theory and data closer together through his mastery of measurement and statistical methods. He has consistently tried to bring the analysis of individual and aggregate outcomes closer together by attending to issues of aggregation. Few scholars have employed such a diverse set of methods in their research and, at the same time, helped us better understand the determinants of consumption and thereby human welfare.

¹¹ For example, Stone, Schwartz, Broderick and Deaton (2010) and Steptoe, Deaton and Stone (2015) analyze how different dimensions of wellbeing vary across countries and age groups. Deaton and Stone (2014) discuss whether having children increases wellbeing. Since there is a great deal of observed and unobserved heterogeneity among households, and having children is a choice that depends on this heterogeneity, it is difficult to say how the “treatment” of having children causally affects wellbeing. In the same spirit, Deaton and 2002 Economics Laureate Daniel Kahneman (2010) discuss the classical question of whether higher income increases life satisfaction – this too, is an issue where patterns in the data are obscured by observed and unobserved heterogeneity among individuals.

References

- Almås, I. (2012), “International Income Inequality: Measuring PPP Bias by Estimating Engel Curves for Food”, *American Economic Review* 102(1), 1093-1117.
- Anderson, S. and D. Ray (2010), “Missing Women: Age and Disease”, *Review of Economic Studies* 77(4), 1262-1300.
- Atkin, D. (2013), “Trade, Tastes and Nutrition in India”, *American Economic Review* 103(5), 1629-1663.
- Attanasio, O. P. and G. Weber (2010), “Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy”, *Journal of Economic Literature* 48(3), 693-751.
- Arrow, K.J., B.D. Bernheim, M.S. Feldstein, J.M. Poterba, and R.M. Solow (2011), “100 Years of the American Economic Review: The Top 20 Articles”, *American Economic Review* 101(1), 1-8.
- Banks, J., R. Blundell, and A. Lewbel (1997), “Quadratic Engel Curves and Consumer Demand”, *Review of Economics and Statistics* 79(4), 527-539.
- Barten, A.P. (1967), “Evidence on the Slutsky Conditions for Demand Equations”, *Review of Economics and Statistics* 49(1), 77-84.
- Barten, A.P. (1969), “Maximum Likelihood Estimation of a Complete System of Demand Equations”, *European Economic Review* 1(1), 7-73.
- Bewley, T.F. (1977), “The Permanent Income Hypothesis: A Theoretical Formulation”, *Journal of Economic Theory* 16(2), 252-92.
- Blinder, A. and A. Deaton (1985), “The Time Series Consumption Function Revisited”, *Brookings Papers on Economic Activity* 1985(2), 465-521.
- Blundell, R., P. Pashardes, and G. Weber (1993), “What Do We Learn About Consumer Demand Patterns from Micro Data?”, *American Economic Review* 83(3), 570-597.
- Blundell, R. and T. M. Stoker (2005), “Heterogeneity and Aggregation”, *Journal of Economic Literature* 43(2), 347-391.
- Blundell, R., L. Pistaferri, and I. Preston, (2008), “Consumption Inequality and Partial Insurance”, *American Economic Review* 98(5), 1887-1921.
- Boppart, T. (2014), “Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences”, *Econometrica* 82(6), 2167-2196.
- Breeden D.T. (1979), “An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities”, *Journal of Financial Economics* 7(3), 265-296.

- Campbell, J.Y. and A. Deaton (1989), "Why Is Consumption so Smooth?", *Review of Economic Studies* 56(3), 357-373.
- Chen, S. and M. 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), 1577-1625.
- Christensen, L.R., D.W. Jorgensen, and L.J. Lau (1975), "Transcendental Logarithmic Utility Functions", *American Economic Review* 65(3), 367-383.
- Costa, D. (2001), "Estimating Real Income in the United States from 1888 to 1994: Correcting CPI Bias Using Engel Curves", *Journal of Political Economy* 109(6), 1288-1310.
- Crawford, I., M. Keen, and S. Smith (2010), "Value Added Tax and Excises" in Mirrlees, J., S. Adam, T. Besley, R. Blundell, S. Bond, R. Chote, M. Gammie, P. Johnson, G. Myles, and J. Poterba (eds.), *Dimensions of Tax Design: The Mirrlees Review*, Oxford University Press. Oxford, for Institute for Fiscal Studies.
- Deaton, A. (1974a), "The Analysis of Consumer Demand in the United Kingdom, 1900-1970", *Econometrica* 42(2), 341-367.
- Deaton, A. (1974b), "A Reconsideration of the Empirical Implications of Additive Preferences", *Economic Journal* 84 (3), 338-348.
- Deaton, A. (1980), "Measurement of Welfare: Theory and Practical Guidelines", LSMS Working Paper No. 7, The World Bank, Washington, D.C.
- Deaton, A. (1981), "Three Essays on a Sri Lanka Household Survey", LSMS Working Paper No. 11, The World Bank, Washington, D.C.
- Deaton, A. (1985), "Panel Data from Times Series of Cross-Sections", *Journal of Econometrics* 30(1-2), 109-126.
- Deaton, A. (1987), "Life-Cycle Models of Consumption: Is the Evidence Consistent with the Theory?" in Bewley, T. (ed.): *Advances in Econometrics, Vol II*, North-Holland, Amsterdam.
- Deaton, A. (1989), "Looking for Boy-Girl Discrimination in Household Expenditure Data", *World Bank Economic Review* 3(1), 1-15.
- Deaton, A. (1991), "Savings and Liquidity Constraints", *Econometrica* 59(5), 1221-1248.
- Deaton, A. (1992), *Understanding Consumption*, Oxford University Press, Oxford.
- Deaton, A. (1997), *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*, Johns Hopkins University Press, Baltimore, MD.
- Deaton, A. (1988), "Quality, Quantity and Spatial Variation in Price", *American Economic Review* 78(3), 418-430.
- Deaton, A. (2003a), "Health, Inequality, and Economic Development", *Journal of Economic Literature* 41(1), 113-158.

- Deaton, A. (2003b), “Adjusted Indian Poverty Estimates for 1999–2000”, *Economic and Political Weekly* 38(4), 322-326.
- Deaton, A. (2003c), “Prices and Poverty in India, 1987-2000”, *Economic and Political Weekly* 38(4), 362-368.
- Deaton, A. (2005), “Measuring Poverty in a Growing World (or Measuring Growth in a Poor World)”, *Review of Economics and Statistics* 87(1), 1-19.
- Deaton, A. (2008), “Price Trends in India and Their Implications for Measuring Poverty”, *Economic and Political Weekly*, 43-49.
- Deaton, A. and J. Campbell (1989), “Why Is Consumption So Smooth?”, *Review of Economic Studies*, 56(3), 357-373.
- Deaton, A. (2010), “Price Indexes, Inequality, and the Measurement of World Poverty”, *American Economic Review* 100(1), 5-34.
- Deaton, A. and J. Drèze (2009), “Food and Nutrition in India: Facts and Interpretations”, *Economic and Political Weekly* 44(7), 42-65.
- Deaton, A. and O. Dupriez (2011), “Purchasing Power Parity Exchange Rates for the Global Poor”, *American Economic Journal: Applied* 3(2), 137-166.
- Deaton, A. and M. Grosh (2000), “Consumption”, in Grosh, M. and P. Glewwe (eds.) *Designing Household Survey Questionnaires for Developing Countries: Lessons from 15 Years of the Living Standards Measurement Study*, Oxford University Press, New York, NY, for the World Bank.
- Deaton, A. and D. Kahneman (2010), “High Income Improves Evaluation of Life but not Emotional Well-being”, *Proceedings of the National Academy of Sciences of the United States of America* 107(38), 16489-16493.
- Deaton, A. and G. Laroque (1992), “On the Behavior of Commodity Prices”, *Review of Economic Studies* 59(1), 1-23.
- Deaton, A. and J. Muellbauer (1980a), “An Almost Ideal Demand System”, *American Economic Review* 70(3), 312-326.
- Deaton, A. and J. Muellbauer (1980b), *Economics and Consumer Behavior*, Cambridge University Press, Cambridge, UK.
- Deaton, A. and J. Muellbauer (1986), “On Measuring Child Costs: With Applications to Poor Countries”, *Journal of Political Economy* 94(4), 720-744.
- Deaton, A. and C. Paxson (1994), “Intertemporal Choice and Inequality”, *Journal of Political Economy* 102(3), 437-467.
- Deaton, A. and A. Stone (2014), “Evaluative and Hedonic Wellbeing among those with and without Children at Home”, *Proceedings of the National Academy of Sciences of the United States of America* 111(4), 1328-1333.

- Deaton, A. and A. Tarozzi (2000), “Prices and Poverty in India”, Princeton Working Paper.
- Diewert, E. (1974), “Applications of Duality Theory”, in Intriligator, M. and D. Kendrick (eds) *Frontiers of Quantitative Economics Vol. II*, North-Holland, Amsterdam.
- Duflo, E. (2012), “Women’s Empowerment and Economic Development”, *Journal of Economic Literature* 50(4), 1051-1079.
- Flavin, M. (1981), “The Adjustment of Consumption to Changing Expectations about Future Income”, *Journal of Political Economy* 89(5), 974-1009.
- Gorman, W.M. (1953), “Community Preference Fields”, *Econometrica* 21(1), 63-80.
- Gorman, W.M. (1961), “On a Class of Preference Fields”, *Metroeconomica* 13(2), 53-56.
- Guvenen, F. (2007), “Learning Your Earning: Are Labor Income Shocks Really Very Persistent?” *American Economic Review* 97(3), 687-712.
- Guvenen, F. (2009), “An Empirical Investigation of Labor Income Processes” *Review of Economic Dynamics* 12(1), 58-79.
- Hall, R.E. (1978), “Stochastic Implications of the Life-Cycle-Permanent Income Hypothesis: Theory and Evidence”, *Journal of Political Economy* 86(2), 971-987.
- Hall, R.E. and F.S. Mishkin (1982), “The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households”, *Econometrica* 50(2), 461-481.
- Hamilton, B. (2001), “Using Engel’s Law to Estimate CPI Bias”, *American Economic Review* 91(3), 619-630.
- Heathcote, J., K. Storesletten, and G.L. Violante (2009), “Quantitative Macroeconomics with Heterogeneous Households”, *Annual Review of Economics* 1, 319-354.
- Heathcote, J., K. Storesletten, and G.L. Violante (2014), “Consumption and Labor Supply with Partial Insurance: An Analytical Framework”, *American Economic Review* 104(7), 2075-2126.
- Huggett, M. (1993), “The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies”, *Journal of Economic Dynamics and Control* 17(5-6), 953-969.
- Imrohoroglu, A. (1989), “Cost of Business Cycles with Indivisibilities and Liquidity Constraints”, *Journal of Political Economy* 97(6), 1364-1383.
- Jorgenson, D.W., L.J. Lau, and T.M. Stoker (1980), “Welfare Comparison under Exact Aggregation”, *American Economic Review* 70(2), 268-272.
- Jorgenson, D.W., L.J. Lau, and T.M. Stoker (1982), “The Transcendental Logarithmic Model of Aggregate Consumer Behavior”, in Basman R.L. and G. Rhodes (eds.) *Advances in Econometrics*, Vol. 1, JAI Press, Greenwich.
- Krusell, P. and A.A. Smith (1998), “Income and Wealth Heterogeneity in the Macroeconomy”, *Journal of Political Economy* 106(5), 867-896.

- Leibenstein, H.A. (1957), *Economic Backwardness and Economic Growth: Studies in the Theory of Economic Development*, Wiley, New York, NY.
- Leland, H. (1968), "Saving and Uncertainty: The Precautionary Demand for Saving", *Quarterly Journal of Economics* 82(3), 465-473.
- Leser, C. (1963), "Forms of Engel Functions", *Econometrica* 31(4), 694-703.
- Lewbel, A. (1989), "Nesting the AIDS and the Translog Demand Systems", *International Economic Review* 30(2), 349-356.
- Lewbel, A. and K. Pendakur (2009), "Tricks with Hicks: The EASI Demand System", *American Economic Review* 99(3), 827-863.
- Lucas, R.E. (1978), "Asset Prices in an Exchange Economy", *Econometrica* 46(6), 1429-1445.
- MaCurdy, T.E. (1981), "An Empirical Model of Labor Supply in a Life-Cycle Setting", *Journal of Political Economy* 89(6), 1059-1085.
- Muellbauer, J. (1975), "Aggregation, Income Distribution and Consumer Demand", *Review of Economic Studies* 42(4), 525-543.
- Muellbauer, J. (1976), "Community Preferences and the Representative Consumer", *Econometrica* 44(5), 979-999.
- Neary, J.P. (2004), "Rationalizing the Penn World Table: True Multilateral Indices for International Comparisons of Real Income", *American Economic Review* 94(5), 1411-1428.
- Pendakur, K. (2002), "Taking Prices Seriously in the Measurement of Inequality", *Journal of Public Economics* 86(1), 47-69.
- Rizvi, S.A.T. (2006), "The Sonnenschein-Mantel-Debreu Results after Thirty Years", *History of Political Economy* 38(suppl.), 228-245.
- Sandmo, A. (1970), "The Effect of Uncertainty on Saving Decisions", *Review of Economic Studies* 37(3), 353-360.
- Sonnenschein, H. (1973), "Do Walras' Identity and Continuity Characterize the Class of Community Excess Demand Functions?", *Journal of Economic Theory* 6(4), 345-354.
- Steptoe, A., A. Deaton, and A. Stone (2015), "Subjective Wellbeing, Health, and Ageing", *Lancet* 385(9968), 640-648.
- Stone, R. (1954a), *The Measurement of Consumers' Expenditure and Behaviour in the United Kingdom, 1920-1938*, Vol. I, Cambridge University Press, Cambridge, UK.
- Stone, R. (1954b), "Linear Expenditure Systems and Demand Analysis: An Application to the Pattern of British Demand", *Economic Journal* 64(255), 511-527.
- Stone, A., J. Schwartz, J. Broderick, and A. Deaton (2010), "A Snapshot of the Age Distribution of Psychological Wellbeing in the United States", *Proceedings of the National Academy of Sciences of the United States of America* 107(22), 9985-9990.

Subramanian, S. and A. Deaton (1996), "The Demand for Food and Calories", *Journal of Political Economy* 104(1), 133-162.

Townsend, R.M. (1994), "Risk and Insurance in Village India", *Econometrica* 62(3), 539-591.

Working, H. (1943), "Statistical Laws of Family Expenditure", *Journal of the American Statistical Association* 38(221), 43-56.

Zeldes, S.P. (1989), "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence", *Quarterly Journal of Economics* 104(2), 275-298.