A000234

aggregation (econometrics)

The econometrics of aggregation is about modelling the relationship between individual (micro) behaviour and aggregate (macro) statistics, so that data from both levels can be used for estimation and inference about economic parameters. Practical models must address three types of individual heterogeneity – in income and preferences, in wealth and income risk, and in market participation. This entry discusses recent solutions to these problems in the context of demand analysis, consumption modelling and labour supply. Also discussed is work that uses aggregation structure to solve microeconometric estimation problems, and work that addresses whether macroeconomic interactions provide approximate solutions to aggregation problems.

Aggregation refers to the connection between economic interactions at the micro and the macro levels. The micro level refers to the behaviour of individual economic agents. The macro level refers to the relationships that exist between economy-wide totals, averages or other economic aggregates. For instance, in a study of savings behaviour refers to the process that an individual or household uses to decide how much to save out of current income, whereas the aggregates are total or per-capita savings and income for a national economy or other large group. The econometrics of aggregation refers to modelling with the individual–aggregate connection in mind, creating a framework where information on individual behaviour together with co-movements of aggregates can be used to estimate a consistent econometric model.

In economic applications one encounters many types and levels of aggregation: across goods, across individuals within households, and so on. We focus on micro to macro as outlined above, and our 'individual' will be a single individual or a household, depending on the context. We hope that this ambiguity does not cause confusion.

At a fundamental level, aggregation is about handling detail. No matter what the topic, the microeconomic level involves purposeful individuals who are dramatically different from one another in terms of their needs and opportunities. Aggregation is about how all this detail distils in relationships among economic aggregates. Understanding economic aggregates is essential for understanding economic policy. There is just too much individual detail to conceive of tuning policies to the idiosyncrasies of many individuals.

This detail is referred to as individual heterogeneity, and it is pervasive. This is a fact of empirical evidence and has strong econometric implications. If you ignore or neglect individual heterogeneity, then you can't get an interpretable relationship between economic aggregates. Aggregates reflect a smear of individual responses and shifts in the composition of individuals in the population; without careful attention, the smear is unpredictable and uninterpretable.

Suppose that you observe an increase in aggregate savings, together with an increase in aggregate income and in interest rates. Is the savings increase primarily arising from wealthy people or from those with moderate income? Is the impact of interest rates different between the wealthy and others? Is the response different for the elderly than for the young? Has future income for most people become more risky?

How could we answer these questions? The change in aggregate savings is a mixture of the responses of all the individuals in the population. Can we disentangle it to understand the change at a lower level of detail, like rich versus poor, or young versus old? Can we count on the mixture of responses underlying aggregate savings to be stable? These are questions addressed by aggregation.

Recent progress on aggregation and econometrics has centred on explicit models of individual heterogeneity. It is useful to think of heterogeneity as arising from three broad categories of differences. First, individuals differ in tastes and incomes. Second, individuals differ in the extent to which they participate in markets. Third, individuals differ in the situations of wealth and income risk that they encounter depending on the market environment that exists. Our discussion of recent solutions is organized around these three categories of heterogeneity. For deeper study and detailed citations, see the surveys by Blundell and Stoker (2005), Stoker (1993) and Browning, Hansen and Heckman (1999).

The classical aggregation problem provides a useful backdrop for understanding current solutions. We now review its basic features, as originally established by Gorman (1953) and Theil (1954). Suppose we are studying the consumption of some product by households in a large population over a given time period t. Suppose that the quantity purchased q_{it} is determined by household resources m_{it}, or 'income' for short, as in the formula:

$$q_{it} = \alpha_i + \beta_i m_{it}$$

Here α_i represents a base level consumption, and β_i represents household i's marginal propensity to spend on the product.

For aggregation, we are interested in what, if any, relationship there is between average quantity and average income:

$$\bar{q}_t = \frac{1}{n_t}\sum_{i=1}^{n_t} q_{it} \quad \text{and} \quad \bar{m}_t = \frac{1}{n_t}\sum_{i=1}^{n_t} m_{it}$$

where all households have been listed as $i=1,\ldots,n_t$. Let's focus on one version of this issue, namely, what happens if some new income becomes available to households, either through economic growth or a policy. How will the change in average quantity purchased $\Delta\bar{q}$ be related to the change in average income $\Delta\bar{m}$?

Suppose that household i gets Δm_i in new income. Their change in quantity purchased is the difference between purchases at income $m_{it}+\Delta m_i$ and at income m_{it}, or

$$\Delta q_i = \beta_i \cdot \Delta m_i$$

Now, the average quantity change is $\Delta\bar{q} = \sum_i \Delta q_i/n_t$, so that

$$\Delta\bar{q} = \frac{1}{n_t}\sum_{i=1}^{n_t} \beta_i \cdot \Delta m_i \tag{1}$$

In general, it seems we need to know a lot about who gets the added income – which i's get large values of Δm_i and which i's get small values of Δm_i. With a transfer policy, any group of households could be targeted for the new income, and their specific set of values of β_i would determine $\Delta\bar{q}$. A full schedule of how much new income goes to each household i as well as how they spend it (that is, Δm_i and β_i), seems like a lot of detail to keep track of, especially if the population is large. Can we ever get by knowing just the change in average income $\Delta\bar{m} = \sum_i \Delta m_i/n_t$?

There are two situations where we can, where a full schedule is not needed:

1. Each household spends in exactly the same way, namely, $\beta_I = \beta$ for all i, so that who gets the new income doesn't affect $\Delta\bar{q}$.
2. The distribution of income transfers is restricted in a convenient way.

Situation 1 is (common) micro linearity, which is termed *exact aggregation*. Another way to understand the structure is to write (1) in the covariance formulation:

$$\Delta\bar{q} = \bar{\beta}\cdot\Delta\bar{m} + \frac{1}{n_t}\sum_{i=1}^{n_t}(\beta_i - \bar{\beta})\cdot(\Delta m_i - \Delta\bar{m}) \tag{2}$$

where we denote the average spending propensity as $\bar{\beta} = \sum_i \beta_i / n_t$. With exact aggregation there is no variation in β_i, so that $\beta_i = \beta = \bar{\beta}$ and the latter term always vanishes. That is, it doesn't matter who gets the added income because everyone spends the same way. When there is variation in β_i, matters are more complicated unless it can be assured that the new income were always given to households in a way that is uncorrelated with the propensities β_i. 'Uncorrelated transfers' provide an example of a Situation 2, but that is a distribution restriction that is hard to verify with empirical data.

Under uncorrelated transfers, we can also interpret the relationship between $\Delta\bar{q}$ and $\Delta\bar{m}$, that is, the macro propensity is the average propensity $\bar{\beta}$. There are other distributional restrictions that give a constant macro propensity, but a different one from the parameter produced by uncorrelatedness. For instance, suppose that transfers of new income always involved fixed shares of the total amount. That is, household i gets

$$\Delta m_i = s_i \Delta\bar{m} \tag{3}$$

In this case, average purchases are

$$\Delta\bar{q} = \frac{1}{n_t}\sum_{i=1}^{n_t}\beta_i\cdot(s_i\Delta\bar{m}) = \tilde{\beta}_{wtd}\cdot\Delta\bar{m} \tag{4}$$

where $\tilde{\beta}_{wtd}$ is the weighted average $\tilde{\beta}_{wtd} \equiv \sum_i \beta_i s_i / n_t$. This is a simple aggregate relationship, but the coefficient $\tilde{\beta}_{wtd}$ applies only for the distributional scheme (3); it matters who gets what share of the added income. Aside from being a weighted average of $\{\beta_i\}$, there is no reason for $\tilde{\beta}_{wtd}$ to be easily interpretable – for instance, if households with low β_i's have high s_i's, then $\tilde{\beta}_{wtd}$ will be low. If your aim was to estimate the average propensity $\bar{\beta}$, there is no reason to believe that the bias $\tilde{\beta}_{wtd} - \bar{\beta}$ will be small.

Empirical models that take aggregation into account apply structure to individual responses and to allowable distributional shifts. Large populations are modelled, so that compositional changes are represented via probability distributions, and expectations are used instead of averages (for example, mean quantity $E_t(q)$ is modelled instead of the sample average $\bar{q}_t$). Individual heterogeneity is the catch-all term for individual differences, and they must be characterized. Distribution restrictions must be applied where heterogeneity is important. For instance, in our example structure on the distribution of new income is required for dealing with the heterogeneity in β_i, but not for the heterogeneity in α_i.

Progress in empirical modelling has come about because of the enhanced availability of micro data over time. The forms of behavioural models in different research areas have been tightly characterized, which is necessary for understanding how to account for aggregation. That is, when individual heterogeneity is characterized empirically, the way is clear to understanding what distributional influences are relevant and must be taken into account. We discuss recent examples of this below.

Some solutions to aggregation problems

Demand models and exact aggregation

It is well known that demand patterns of individual households vary substantially with whether households are rich or poor, and vary with many observable demographic characteristics, such as household (family) size, age of head and ages of children, and so on. As surveyed in Blundell (1988), traditional household demand models relate household commodity expenditures to price levels, total household budget (income) and observable household characteristics. Aggregate demand models relate (economy-wide) aggregate commodity expenditures to price levels and the distribution of income and characteristics in the population. Demand models illustrate exact aggregation, a practical approach for accommodating heterogeneity at the micro and macro levels. These models assume that demand parameter values are the same for all individuals, but explicitly account for observed differences in tastes and income.

For instance, suppose we are studying the demand for food and we are concerned with the difference in demands for households of small size versus large size. We model food purchases for household i as part of static allocation of the budget m_{it} to $j=1,\ldots,J$ expenditure categories, where food is given by $j=1$, and price levels at time t are given by $P_t=(p_{1t},\ldots,p_{Jt})$. Small families are indicated by $z_{it}=0$ and large families by $z_{it}=1$.

Expenditure patterns are typically best fit in budget share form. For instance, a translog model of the food share takes the form

$$w_{1it}=\frac{p_{1t}q_{1it}}{m_{it}}=\frac{1}{D(p_t)}\left[\alpha_1+\sum_{i=1}^{J}\beta_{1j}\ln p_{jt}+\beta_m\ln m_{it}+\beta_z z_{it}\right] \tag{5}$$

where $D(p_t)=1+\sum_{i=1}^{J}\beta_j\ln p_{jt}$. The parameters ($\alpha_1$ and all β's) are the same across households, and the price levels (p_{jt}'s) are the same for all households but vary with t. Individual heterogeneity is represented by the budget m_{it} and the family size indicator z_{it}. We have omitted an additive disturbance for simplicity, which would represent another source of heterogeneity. The important thing for aggregation is that model (5) is intrinsically linear in the individual heterogeneity. That is, we can write

$$w_{1it}=b_1(p_t)+b_m(p_t)\cdot\ln m_{it}+b_z(p_t)\cdot z_{it} \tag{6}$$

The aggregate share of food in the population is the mean of food expenditures divided by mean budget, or

$$W_{1t}=\frac{E_t(m_{it}w_{1it})}{E_t(m_{it})}=b_1(p_t)+b_m(p_t)\cdot\frac{E_t(m_{it}\ \ln\ m_{it})}{E_t(m_{it})} + b_z(p_t)\cdot\frac{E_t(m_{it}z_{it})}{E_t(m_{it})} \tag{7}$$

The aggregate share depends on prices, the parameters (α_1 and all β's) and two statistics of the joint distribution of m_{it} and z_{it}. The first,

$$S_{mt}=\frac{E_t(m_{it}\ln m_{it})}{E_t(m_{it})} \tag{8}$$

is an entropy term that captures the size distribution of budgets, and the second

$$S_{zt} = \frac{E_t\,(m_{it}z_{it})}{E_t\,(m_{it})} \tag{9}$$

is the percentage of total expenditure accounted for by households with $z_{it} = 1$, that is, large families.

The expressions (6) and (7) illustrate *exact aggregation* models. Heterogeneity in tastes and budgets (incomes) are represented in an intrinsically linear way. For aggregate demand, all one needs to know about the joint distribution of budgets m_{ti} and household types z_{it} is a few statistics; here S_{mt} and S_{zt}.

The obvious similarity between the individual model (6) and the aggregate model (7) raises a further question. How much bias is introduced by just fitting the individual model with aggregate data, that is, putting $E_t(m_{it})$ and $E_t(z_{it})$ in place of m_{it} and z_{it}, respectively? This can be judged by the use of *aggregation factors*. Define the factors π_{mt} and π_{zt} as

$$\pi_{mt} = \frac{S_{mt}}{\ln E_t\,(m_{it})} \quad \text{and} \quad \pi_{zt} = \frac{S_{zt}}{E_t\,(z_{it})}$$

so that the aggregate share is

$$W_{1t} = \frac{E_t(m_{it}\,w_{1it})}{E_t(m_{it})} = b_1(p_t) + b_m(p_t) \cdot \pi_{mt} \cdot lnE_t(m_{it}) + b_z(p_t) \cdot \pi_{zt} \cdot E_t(z_{it})$$

One can learn about the nature of aggregation bias by studying the factors π_{mt} and π_{zt}. If they are both roughly equal to 1 over time, then no bias would be introduced by fitting the individual model with aggregate data. If they are roughly constant but not equal to 1, then constant biases are introduced. If the factors are time varying, more complicated bias would result. In this way, with exact aggregation models, aggregation factors can depict the extent of aggregation bias.

The current state of the art in demand analysis uses models in exact aggregation form. The income (budget) structure of shares is adequately represented as quadratic in ln m_{it}, as long as many demographic differences are included in the analysis. This means that aggregate demand depends explicitly on many statistics of the income-demographic distribution, and it is possible to gauge the nature and sources of aggregation bias using factors as we have outlined. See Banks, Blundell and Lewbel (1997) for an example of demand modelling of British expenditure data, including the computation of various aggregation factors.

Exact aggregation modelling arises naturally in situations where linear models have been found to provide adequate explanations of empirical data patterns. This is not always the case, as many applications require models that are intrinsically nonlinear. We now discuss an example of this kind where economic decisions are discrete.

Market participation and wages

Market participation is often a discrete decision. Labourers decide whether to work or not, firms decide whether to enter a market or exit a market. There is no 'partial' participation in many circumstances, and changes are along the extensive margin. This raises a number of interesting issues for aggregation.

We discuss these issues using a simple model of labour participation and wages. We consider two basic questions. First, how is the fraction of working (participating) individuals affected by the distribution of factors that determine whether each individual chooses to work? Second, what is the structure of average wages, given that wages are observed only for individuals who choose to work? The latter question is of interest for interpreting wage movements: if average wages go up, is that because (*a*) most individual wages went up or (*b*) low-wage individuals become unemployed, or leave work? These two reasons give rise to quite different views of the change in economic welfare associated with an increase in average wages.

The standard empirical model for individual wages expresses log wage as a linear function of time effects, schooling and demographic (cohort) effects. Here we begin with

$$\ln w_{it} = r(t) + \beta \cdot S_{it} + \varepsilon_{it} \tag{10}$$

where $r(t)$ represents a linear trend or other time effects, S_{it} is the level of training or schooling attained by individual i at time t, and ε_{it} are all other idiosyncratic factors. This setting is consistent with a simple skill price model, where $w_{ti} = R_t H_{it}$ with skill price $R_t = e^{r(t)}$ and skill (human capital) level $H_{it} = e^{\beta S_{it} + \varepsilon_{it}}$. We take eq. (10) to apply to all individuals, with the wage representing the available or offered wage, and β the return to schooling. However, we observe that wage only for individuals who choose to work.

We assume that individuals decide whether to work by first forming a reservation wage

$$\ln w_{it}^{*} = s^{*}(t) + \alpha ln B_{it} + \beta^{*} \cdot S_{it} + \zeta_{it}$$

where $s(t)$ represents time effects, B_{it} is the income or benefits available when individual i is out of work at time t, S_{it} is schooling as before, and ζ_{it} are all other individual factors. Individual i will work at time t if their offered wage is as big as their reservation wage, or $w_{it} \geq w_{it}^{*}$. We denote this by the participation indicator I_{it}, where $I_{it} = 1$ if i works and $I_{it} = 0$ if i doesn't work. This model of participation can be summarized as

$$\begin{aligned} I_{it} &= 1\,[w_{it} \geq w_{it}^{*}] = 1[\ln w_{it} - \ln w_{it}^{*} \geq 0] \\ &= 1\,[s(t) - \alpha \ln B_{it} + \gamma \cdot S_{it} + v_{it} \geq 0] \end{aligned} \tag{11}$$

where $s(t) \equiv r(t) - s^{*}(t)$, $\gamma \equiv \beta - \beta^{*}$ and $v_{It} \equiv \varepsilon_{it} - \zeta_{it}$.

If the idiosyncratic terms ε_{it}, v_{it} are stochastic errors with zero means (conditional on B_{it},S_{it}) and constant variances, then (10) and (11) is a standard selection model. That is, if we observe a sample of wages from working individuals, they will follow (10) subject to the proviso that $I_{it} = 1$. This can be accommodated in estimation by assuming that ε_{it}, v_{it} have a joint normal distribution. That implies that the log wage regression of the form (10) can be corrected by adding a standard selection term as

$$\ln w_{it} = r(t) + \beta \cdot S_{it} + \frac{\sigma_{\varepsilon v}}{\sigma_v} \lambda \left[\frac{s(t) - \alpha \ln B_{it} + \gamma\, S_{it}}{\sigma_v} \right] + \eta_t. \tag{12}$$

Here, σ_v is the standard deviation of v and $\sigma_{\varepsilon v}$ is the covariance between ε and v. $\lambda(\cdot) = \phi(\cdot)/\Phi(\cdot)$ is the 'Mills ratio', where ϕ and Φ are the standard normal p.d.f. and c.d.f respectively. This equation is properly specified for a sample of working individuals – that is, we have $E(\eta_t | S_{it}, B_{it}, I_{it} = 1) = 0$. For a given levels of benefits and schooling, eq. (11) gives the probability of participating in work as

$$E_t[I_{it} | B_{it}, S_{it}] = \Phi \left[\frac{s(t) - \alpha \ln B_{it} + \gamma \cdot S_{it}}{\sigma_v} \right] \tag{13}$$

where $\Phi[\cdot]$ is the normal c.d.f.

For studying average wages, the working population is all individuals with $I_{it} = 1$. The fraction of workers participating is therefore the (unconditional) probability that $\alpha \ln B_{it} - \gamma \cdot S_{it} - v_{it} \leq s(t)$. This probability is the expectation of I_{it} in (11), an intrinsically nonlinear function in observed heterogeneity B_{it} and S_{it} and unobserved heterogeneity v_{it}, so we need some explicit distribution assumptions. In particular, assume that the participation index $\alpha \ln B_{it} - \gamma \cdot S_{it} - v_{it}$ is normally distributed with mean $\mu_t = \alpha E_t(\ln B_{it}) - \gamma E_t(S_{it})$ and variance

$$\sigma_t^2 = \alpha^2 Var_t(\ln B_{it}) + \beta^2 Var_t(S_{it}) - 2a\beta \cdot Cov_t(\ln B_{it}, S_{it}) + \sigma_v^2. \tag{14}$$

Now we can derive the labour participation rate (or one minus the unemployment rate) as

$$E_t[I_{it}] = \Phi\left[\frac{s(t) - \alpha E_t(\ln B_{it}) + \gamma E_t(S_{it})}{\sigma_t}\right] \tag{15}$$

where again $\Phi[\cdot]$ is the normal c.d.f. This formula relates the participation rate to average out-of-work benefits $E_t(\ln B_{it})$ and average training $E_t(S_{it})$, as well as their variances and covariances through σ_t. The specific relation depends on the distributional assumption adopted; (15) relies on normality of the participation index in the population.

For wages, a similar analysis applies. Log wages are a linear function (10) applicable to the full population. However, for participating individuals, the intrinsically nonlinear selection term is introduced, so that we need explicit distributional assumptions. Now suppose that log wage $\ln w_{it}$ and the participation index $\alpha \ln B_{it} - \gamma \cdot S_{it} - v_{it}$ are joint normally distribution. It is not hard to derive the expression for average log wages of working individuals

$$E_t[\ln w_{it} | I_{it} = 1] = r(t) + \beta \cdot E_t(S_{it} | I = 1) + \frac{\sigma_{\varepsilon v}}{\sigma_t} \lambda \left[\frac{s(t) - \alpha E_t(\ln B_{it}) + \gamma E_t(S_{it})}{\sigma_t}\right]. \tag{16}$$

This is an interesting expression, which relates average log wage to average training of the workers as well as to the factors that determine participation.

However, we are not interested in average log wages, but rather average wages $E_t(w_{it})$. The normality structure we have assumed is enough to derive a formulation of average wages, although it is a little complex to reproduce in full here. In brief, Blundell, Reed and Stoker (2003) show that the average wages of working individuals $E[w_{it} | I_{it} = 1]$ can be written as

$$\ln E[w_{it} | I_{it} = 1] = r(t) + \beta \cdot E_t(S_{it}) + \Omega_t + \Psi_t \tag{17}$$

where Ω_t, Ψ_t are correction terms that arise as follows. Ω_t corrects for the difference between the log of an average and the average of a log, as

$$\Omega_t \equiv \ln E_t(w_{it}) - E_t(\ln w_{it}) + \Omega_t.$$

Ψ_t corrects for participation, as

$$\Psi_t \equiv \ln E[w_{it} | I_{it} = 1] - \ln E_t(w_{it}).$$

Recall our original question, about whether an increase in average wages is due to an increase in individual wages or to increased unemployment of low-wage workers. That is captured in (17). That is, Ψ_t gives the participation effect, and the other terms capture changes in average wage $E_t(w_{it})$ when all are participating. As such, this analysis provides a vehicle for separating overall wage growth from compositional effects due to participation.

Blundell, Reed and Stoker (2003) analyse British employment using a framework similar to this, but also allowing for heterogeneity in hours worked. Using out-of-work benefits as an instrument for participation, they find that over 40 per cent of observed aggregate wage growth from 1978 to 1996 arises from selection and other compositional effects.

We have now discussed aggregation and heterogeneity with regard to tastes and incomes, and market participation. We now turn to heterogeneity with regard to risks and market environments.

Consumption and risk environments

Consumption and savings decisions are clearly affected by preference heterogeneity, as we discussed earlier. The present spending needs of a large family clearly differ from those of a small family or a single individual, the needs of teenage children differ from those of preschoolers, the needs of young adults differ from those of retirees, and on and on. These aspects are very important, and need to be addressed as they were in demand models above. Browning and Lusardi (1996) survey the extensive evidence on heterogeneity in consumption, and Attanasio (1999) is an excellent comprehensive survey of work on consumption.

We use consumption and savings to illustrate another type of heterogeneity, namely, that of wealth and income risks. That is, with forward planning under uncertainty, the risk environment of individuals or households becomes relevant. There can be individual shocks to income, such as a work layoff or a health problem, or aggregate shocks, such as an extended recession or stock market boom. Each of these shocks can differ in its duration – a temporary layoff can be usefully viewed as transitory, whereas a debilitating injury may affect income for many years. In planning consumption, it is important to understand the role of income risks and wealth risks. When there is no precautionary planning, such as when consumers have quadratic preferences, income risks do not become intertwined with other heterogeneous elements. However, when there is risk aversion, then the precise situation of individual income risks and insurance markets is relevant.

A commonly used model for income is to assume multiplicative permanent and transitory components, with aggregate and individual shocks, as in

$$\Delta \ln y_{it} = (\eta_t + \Delta u_t) + (\varepsilon_{it} + \Delta v_{it}).$$

Here $\eta_t + \Delta u_t$ is the common aggregate shock, with η_t a permanent component and Δu_t transitory. The idiosyncratic shock is $\varepsilon_{it} + \Delta v_{it}$, where ε_{it} is permanent and Δv_{it} transitory.

For studying individual level consumption with precautionary planning, it is standard practice to assume constant relative risk aversion (CRRA) preferences and assume that the interest rate r_t is small. This, together with the income process above, gives a log-linear approximation to individual consumption growth

$$\Delta \ln c_{it} = \rho r_t + (\beta + \varphi r_t)' z_{it} + k_1 \sigma_{At} + k_2 \sigma_{it} + \kappa_1 \eta_t + \kappa_2 \varepsilon_{it}. \qquad (18)$$

Here, z_{it} reflects heterogeneity in preferences, such as differences in demographic characteristics. σ_{At} is the variance of aggregate risk and σ_{it} s the variance of idiosyncratic risk (with each conditional on what is known at time $t-1$), so that these terms reflect precautionary planning. Finally, η_t and ε_{it} arise because of adjustments that are made as permanent shocks are revealed. At time $t-1$ these shocks are not possible to forecast, but then they are incorporated in the consumption plan once they are revealed. In terms of the level of consumption c_{it}, eq. (18) is written as

$$c_{it} = exp(\ln c_{it-1} + \rho r_t + (\beta + \varphi r_t)' z_{it} + k_1 \sigma_{At} + k_2 \sigma_{it} + \kappa_1 \eta_t + \kappa_2 \varepsilon_{it}).$$

This is an intrinsically nonlinear model in the following heterogeneous elements: $\ln c_{it-1}$, z_{it},σ_{it} and ε_{it}. For aggregation, it seems we would need a great deal of distributional structure.

Here is where we can see the role of the risk environment, or markets for insurance for income risks. That is, if there were complete markets with insurance for all risks, then all risk terms vanish from consumption growth. When complete insurance exists for idiosyncratic risks only, then the idiosyncratic terms σ_{it} and ε_{it} vanish from consumption growth, since less precautionary saving is needed. Otherwise, the idiosyncratic risk terms σ_{it} and ε_{it} represent heterogeneity that must be accommodated just like preference differences (and in other settings, participation differences).

In the realistic situation where risks are not perfectly insurable, we require distributional assumptions in order to formulate aggregate consumption. For instance, suppose that we assume that $(\ln c_{it-1}, (\beta + \varphi r_t)' z_{it}, \varepsilon_{it})$ is joint normally distributed with $E_t(\varepsilon_{it}) = 0$, and that idiosyncratic risks are drawn from the same distribution for each consumer (so $\sigma_{it} = \sigma_{It}$ for each i), and that a stability assumption applies to the distribution of lagged consumption. Blundell and Stoker (2005) show that aggregate consumption growth is

$$\Delta \ln E_t(c_{it}) = \rho r_t + (\beta + \varphi r_t)' E_t(z_{it}) + k_1 \sigma_{At} + k_2 \sigma_{It} + \kappa_1 \eta_t + \Lambda_t.$$

This model explains aggregate consumption growth in terms of the mean of preference heterogeneity, risk terms, and an aggregation factor Λ_t. The factor Λ_t is comprised of variances and covariances of the heterogeneous elements $\ln c_{it-1}$,z_{it} and ε_{it}. Thus, this model reflects how aggregate consumption will vary as the individual incomes become more or less risky, and captures how the income risk interplays with previous consumption values.

In overview, as micro consumption models are nonlinear, distributional restrictions are essential. On this point, an empirical fact is that the distribution of household consumption is often observed to be well approximated by a lognormal distribution, and so such lognormal restrictions may have empirical validity. Also relevant here is the empirical study of income and wealth risks, which has focused on earnings processes; see Meghir and Pistaferri (2004) for a recent contribution.

Micro to macro and vice versa

We now turn to two related uses of aggregation structure that have emerged in the literature.

Aggregation as a solution to microeconometric estimation

Consider a situation where the estimation of a model at the micro level is the primary goal of empirical work. Some recent work uses aggregation structure to enhance or permit micro-level parameter identification and estimation. Since aggregation structure provides a bridge between models at the micro level and the aggregate level, it permits all data sources – individual-level data and aggregate-level data – to be used for identification and estimation of economic parameters. Sometimes it is necessary to combine all data sources to identify economic effects (for example, Jorgenson, Lau and Stoker, 1982), and sometimes one can study (micro) economic effects with aggregate data alone (for example, Stoker, 1986). Recent work has developed more systematic methods of using aggregate data to improve micro-level estimates. In particular, one can match aggregate data with simulated moments from the individual data as part of the estimation process.

To see how this can work, suppose we have data on labour participation over several time periods (or groups). We assume that the participation decision is given by the model (11) with normal unobserved heterogeneity, as discussed above. We normalize $\sigma_v = 1$ and take $s(t) = \psi$, a constant, so that the unknown parameters of the participation model are α, γ and ψ. The data situation is as follows; for each group $t = 1, \ldots, T$, we observe the proportion of labour participants P_t and a random sample of benefits and schooling values, $\{B_{it}, S_{it}, i = 1, \ldots, n_t\}$. Given the (probit) expression (13), estimation can be based on matching the observed proportion P_t to the simulated moment

$$\bar{P}_t(\alpha, \gamma, \psi) = \frac{1}{n_t} \sum_{i=1}^{n_t} \Phi[\psi - \alpha \ln B_{it} + \gamma \cdot S_{it}].$$

For instance, we could estimate by least squares over groups, by choosing $\hat{\alpha}, \hat{\gamma}, \hat{\psi}$ to minimize

$$\sum_{t=1}^{T} (P_t - \bar{P}_t(\alpha, \gamma, \psi))^2.$$

Note that this approach does not require a specific assumption on the joint distribution of B_{it} and S_{it} for each t, as the random sample provides the distributional information needed to link the parameters to the observed proportion P_t.

It turns out that this approach for estimation is extremely rich, and was essentially mapped out by Imbens and Lancaster (1994). It has become a principal method of estimating demands for differentiated products, for use in structural models of industrial organization. See Berry, Levinsohn and Pakes (2004) for good coverage of this development.

Can macroeconomic interaction solve aggregation problems?

The basic heuristic that underlies much macroeconomic modelling is that, because of markets, individuals are very coordinated in their actions, so that individual heterogeneity likely has a secondary impact. In simplest terms, the notion is that common reactions across individuals will swamp any behavioural differences. This idea is either just wrong or, at best, very misleading for economic analysis. But that is not to deny that in real world economies there are many elements of commonality in reactions across individuals. Households face similar prices, interest rates and opportunities for employment. Extensive insurance markets effectively remove some individual differences in risk profiles. Optimal portfolio investment can have individuals choosing the same (efficient) basket of securities.

The question whether market interactions can minimize the impact of individual heterogeneity is a classic one, and by and large the answers are negative. However, there has been some recent work with calibrated stochastic growth models that raises some possibilities. A principal example of this is Krusell and Smith (1998), which we now discuss briefly. The Krusell–Smith set-up has infinitely lived consumers, with the same preferences within each period, but with different discount rates and wealth holdings. Each consumer has a chance of being unemployed each period, so there are transitory individual income shocks. Production arises from labour and capital, and there are transitory aggregate productivity shocks. Consumers can insure for the future by investing in capital only. Thus, insurance markets are incomplete, and consumers cannot hold negative capital amounts.

To make savings and portfolio decisions, consumers must predict future prices. To do this, each consumer must keep track of the evolution of the entire distribution of wealth holdings, in principle. This is a lot of information to know, just like what is needed for standard aggregation solutions as discussed earlier. Krusell–Smith's simulations show, however, that this forecasting problem is much easier than one would suspect. That is, for consumer planning and for computing equilibrium, consumers get very close to optimal solutions by keeping track of only two things: mean wealth in the economy and the aggregate productivity shock. This is approximate aggregation, a substantial simplification of the information requirements that one would expect.

The source of this simplification, as well as its robustness, is a topic of active current study. One aspect is that most consumers, especially those with lowest discount rates, save enough to insure their risk so that their propensity to save out of wealth is essentially constant. Those consumers also hold a large fraction of the wealth, so that saving is essentially linear in wealth. This means that there is (approximate) exact aggregation structure, with the mean of wealth determining how much aggregate saving is undertaken. That is, the nature of savings and wealth accumulation approximately solves the aggregation problem for individual forecasting. Aggregate consumption, however, does not exhibit the same simplification. Many low-wealth consumers become unemployed and encounter liquidity constraints. Their consumption is much more sensitive to current output than that of wealthier consumers.

These results depend on the specific formulation of the growth model. Krusell and Smith (2006) survey work that suggests that their type of approximate aggregation can be obtained under a variety of variations of the basic model assumptions. As such, this work raises a number of fascinating issues on the interplay between economic interaction, aggregation and individual heterogeneity. However, it remains to be seen whether the structure of such calibrated models is empirically relevant to actual economies, or whether forecasting can be simplified even with observed variation in saving propensities of wealthy households.

Future progress

Aggregation problems are among the most difficult in empirical economics. The progress that has been made recently is arguably due to two complementary developments. First is the enormous expansion in the availability of data on the behaviour of individual agents, including consumers, households, firms, and so on, in both repeated cross-section and panel data form. Second is the enormous expansion in computing power that facilitates the study of large data sources. These two trends can be reasonably expected to continue, which makes the prospects for further progress quite bright.

There is sufficient variety and complexity in the issues posed by aggregation that progress may arise from many approaches. For instance, we have noted how the possibility of approximate aggregation has arisen in computable stochastic growth models. For another instance, it is sometimes possible to derive properties of aggregate relationships with very weak assumptions on individual behaviour, as in Hildenbrand's (1994) work of the law of demand.

But is seems clear to me that the best prospects for progress lie with careful microeconomic modelling and empirical work. Such work is designed to ferret out economic effects in the presence of individual heterogeneity, and can also establish what are 'typical' patterns of heterogeneity in different applied contexts. Knowledge of typical patterns of heterogeneity is necessary for characterizing the distributional structure that will facilitates aggregation, and such distributional restrictions can then be refuted or validated with actual data. That is, enhanced understanding of the standard structure in the main application areas of empirical economics, such as with commodity demand, consumption and saving and labour supply, will lead naturally to an enhanced understanding of aggregation problems and accurate interpretation of aggregate relationships. There has been great progress of this kind in the past few decades, and there is no reason to think that such progress won't continue or accelerate.

Thomas M. Stoker

Bibliography

Attanasio, O. 1999. Consumption. In *Handbook of Macroeconomics*, vol. 1B, ed. J. Taylor and M. Woodford. Amsterdam: North-Holland.

Banks, J., Blundell, R. and Lewbel, A. 1997. Quadratic Engel curves, indirect tax reform and welfare measurement. *Review of Economics and Statistics* 79, 527–39.

Berry, S., Levinsohn, J. and Pakes, A. 2004. Differentiated products demand systems from a combination of micro and macro data: the new car market. *Journal of Political Economy* 112, 68–105.

Blundell, R. 1988. Consumer behaviour: theory and empirical evidence. *Economic Journal* 98, 16–65.

Blundell, R., Reed, H. and Stoker, T. 2003. Interpreting aggregate wage growth. *American Economic Review* 93, 1114–31.

Blundell, R. and Stoker, T. 2005. Aggregation and heterogeneity. *Journal of Economic Literature* 43, 347–91.

Browning, M., Hansen, L. and Heckman, J. 1999. Micro data and general equilibrium. In *Handbook of Macroeconomics*, vol. 1A, ed. J. Taylor and M. Woodford. Amsterdam: North-Holland.

Browning, M. and Lusardi, A. 1996. Household saving: micro theories and micro facts. *Journal of Economic Literature* 34, 1797–855.

Gorman, W.M. (Terence). 1953. Community preference fields. *Econometrica* 21, 63–80.

Hildenbrand, W. 1994. *Market Demand: Theory and Empirical Evidence*. Princeton: Princeton University Press.

Imbens, G. and Lancaster, T. 1994. Combining micro and macro data in microeconometric models. *Review of Economic Studies* 61, 655–80.

Jorgenson, D., Lau, L. and Stoker, T. 1982. The transcendental logarithmic model of aggregate consumer behavior. In *Advances in Econometrics*, vol. 1, ed. R. Basmann, and G. Rhodes. Greenwich, CT: JAI Press.

Krusell, P. and Smith, A. 1998. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy* 106, 867–96.

Krusell, P. and Smith, A. 2006. Quantitative macroeconomic models with heterogeneous agents. In *Advances in Economics and Econometrics, Proceedings of the Ninth World Congress of the Econometric Society*, ed. R. Blundell, W. Newey and T. Persson. Cambridge: Cambridge University Press.

Meghir, C. and Pistaferri, L. 2004. Income variance dynamics and heterogeneity. *Econometrica* 72, 1–32.

Stoker, T. 1986. Simple tests of distributional effects on macroeconomic equations. *Journal of Political Economy* 94, 763–95.

Stoker, T. 1993. Empirical approaches to the problem of aggregation over individuals. *Journal of Economic Literature* 31, 1827–74.

Theil, H. 1954. *Linear Aggregation of Economic* Relations. Amsterdam: North-Holland.

Index terms

Index terms not found: