Micro Data and General Equilibrium Models*

Martin Browning
Institute of Economics
Copenhagen University
Copenhagen, Denmark
email: Martin.Browning@econ.ku.dk

Lars Peter Hansen
University of Chicago
Chicago, IL
email: l-hansen@uchicago.edu

James J. Heckman
University of Chicago
Chicago, IL
and Affiliated with The American Bar Foundation
email: jjh@uchicago.edu

September 2, 1998

[*]This chapter is prepared for the Handbook of Macroeconomics edited by John Taylor and Michael Woodford, to be published by North Holland, 1999. We thank Marco Cagetti, John Heaton, Jose Scheinkman, John Taylor, Edward Vytlacil and Noah Williams for comments. Hansen and Heckman gratefully acknowledge funding support by the National Science Foundation.

Introduction

An extensive literature in macroeconomics and public finance uses dynamic stochastic general equilibrium models to study consumption, savings, capital accumulation, and asset pricing and to analyze alternative policies. Except for a few special cases, the economies studied cannot be analyzed using "paper and pencil" style analysis. It is often difficult to produce general theorems that are true for all parameter values of dynamic general equilibrium models. This is a general feature of nonlinear dynamic models in economics as well as in the physical sciences. For such models, knowing which parameters govern behavior is essential for understanding their empirical content and for providing quantitative answers to policy questions. For the numerical output of a dynamic equilibrium model to be interesting, the inputs need to be justified as empirically relevant. There are two sources of information that are commonly used in rationalizing parameter values. One is the behavior of time series averages of levels or ratios of key variables. These time series averages are often matched to the steady state implications of versions of the models that abstract from uncertainty. The other input is from microeconomic evidence. In this essay we discuss the use of evidence from both sources, concentrating mostly on microeconomic evidence. See King and Rebelo (1998) and Taylor (1998) for extensive discussions of calibrating real business cycle and staggered contract models, respectively.

It was once believed to be a simple task to extract the parameters needed in general equilibrium theory from a large warehouse of stable micro empirical regularities. Indeed, Prescott (1986) argued that:

> The key parameters of growth models are the intertemporal and intratemporal elasticities of substitution. As Lucas (1980) emphasizes, 'On those parameters, we have a wealth of inexpensive available data from census and cohort information, from panel data describing market conditions and so forth.'

While this Lucas-Prescott vision of *calibration* offers an appealing defense for building models with microeconomic foundations, implementing in practice it exposes major discrepancies between the micro evidence and the assumptions on which the stylized dynamic models are based. The microeconomic evidence is often incompatible with the macroeconomic model being calibrated. For example, a major finding of modern microeconomic data analysis is that preferences are heterogeneous. For reasons of computational tractability, dynamic general equilibrium model-builders often abstract from this feature or confront it in only a limited way.

This chapter explores the discordance between micro evidence and macro use of it and suggests ways in which it can be diminished. Our chapter raises warning flags about the current use of micro evidence in dynamic equilibrium models and indicates the dangers in, and limitations of, many current practices. It also exposes the weak micro empirical foundations of many widely used general equilibrium modeling schemes. The decision to incorporate micro evidence in an internally consistent manner may alter the structure and hence the time series implications of the model. While steady state approximations may be useful for some purposes, compositional changes in labor supply or in market demand, alter the microeconomic elasticities that are relevant for macroeconomics.

Like several of the other contributions to this Handbook, ours is more of a guide for future research than a summary of a mature literature. Because the micro empirical literature and the macro general equilibrium literature have often moved in different directions, it is not surprising that they are currently so detached. The main goal of this essay is to foster the process of reattachment. Macro general equilibrium models provide a framework within which micro empirical research can be fruitfully conducted. At the same time, dynamic general equilibrium theory will be greatly enriched if it incorporates the insights of the micro empirical literature. The micro foundations of macroeconomics are more fruitfully built on models restructured to incorporate microeconomic evidence. This essay explores three challenges for closing the gap between empirical microeconomics and dynamic macroeconomic theory:

- *Heterogeneity*: Any careful reading of the empirical microeconomics literature on consumption, saving and labor supply reveals quantitatively important heterogeneity in agent preferences, constraints, in dimensions of labor supply and skill, and in human capital accumulation processes. Accounting for heterogeneity is required to calibrate dynamic models to microeconomic evidence.

- *Uncertainty*: Modern macroeconomics is based on models of uncertainty. Aggregating earnings across members in a household and across income types may create a disconnect between uncertainty as measured by innovations in time series processes of earnings and income equations and actual innovations in information. Government or interfamily transfers provide insurance that should be accounted for. Alternative risk components such as risks from changes in health, risks from unemployment and job termination, and risks from changes in family structure, have different degrees of predictability and are difficult to quantify. Measuring the true components of both micro-uncertainty and macro-uncertainty and distinguishing them from measurement error and model misspecification remains an empirical challenge that is just beginning to be confronted.

- *Synthesis*: Synthesizing evidence across micro studies is not a straightforward task. Different microeconomic studies make different assumptions, often implicit, about the economic environments in which agents make their decisions. They condition on different variables and produce parameters with different economic interpretations. A parameter that is valid for a model in one economic environment cannot be uncritically applied to a model embedded in a different economic environment. Different general equilibrium models make different assumptions and require different parameters, many of which have never been estimated in the micro literature.

In order to be both specific and constructive, in this essay we limit ourselves to two prototypical general equilibrium models: (a) a stochastic growth model and (b) a perfect foresight overlapping generations model. The first model is sufficiently rich to enable us to explore implications of uncertainty, market structure and some forms of heterogeneity in the preferences and opportunities of microeconomic agents. The second model introduces explicit life cycle heterogeneity and demographic structures in appealing and tractable ways. We consider a recent version of the second model that introduces human capital formation, heterogeneity in skills and comparative advantage in the labor market. These attributes are introduced to provide a framework for analyzing labor market policies, to account for a major source of wealth formation in modern economies, and to account for the phenomenon of rising wage inequality observed in many countries.

The plan of this paper is as follows. We first present two basic theoretical models analyzed in this chapter and the parameters required to implement them. We summarize the main lessons from the micro literature that pertain to each model and their consequences for the models we consider. The models are presented in Parts I and II, respectively with some accompanying discussion of the relevant micro literature. Part III presents further discussion of the micro evidence on intertemporal substitution elasticities.

I. Stochastic Growth Model

This part of the chapter presents alternative variants of a Brock-Mirman (1972) stochastic growth model and discusses the parameters needed to calibrate them. We explicitly consider the consequences of heterogeneity for the predictions of this model and for the practice of research synthesis. It is often not the median or "representative" preferences that govern behavior asymptotically; rather it is the extreme. The agents with the smallest rates of time preference or smallest risk aversion may dominate the wealth accumulation process, but not the supply of labor. Understanding the source and magnitude of the heterogeneity is required before microeconomic estimates can be "plugged" into macroeconomic models. We also explore the measurement of microeconomic uncertainty needed to quantify the importance of precaution in decision-making and to calibrate equilibrium models with heterogeneous agents.

We use the Brock-Mirman (1972) stochastic growth model as a starting point for our discussion because of its analytical tractability. Moreover, it is the theoretical framework for the real-business cycle models of Kydland and Prescott (1982) and Hansen (1985) and for subsequent multiple consumer extensions of it by Aiyagari (1994), Krusell and Smith (1998) and others. Our use of the stochastic growth model is not meant as an endorsement of its empirical plausibility. Much is known about its deficiencies as a model of fluctuations (*e.g.*, see Christiano (1988), Watson (1993), and Cogley and Nason (1995)) or as a model of security market prices implied by a Lucas-Prescott (1971) type of decentralization (*e.g.*, see Hansen and Singleton (1982,1983), Mehra and Prescott (1985), Weil (1989), Hansen and Jagannathan (1991), and Heaton and Lucas (1996)). Nevertheless the Brock-Mirman model and its extensions provide a convenient and widely used starting point for investigating the difficulties in obtaining plausible parameter configurations from microeconomic data and from aggregate time series data.

1. Single Consumer Model

Suppose that there is a single infinitely lived consumer. This consumer supplies labor and consumes in each period, evaluating consumption flows using a von-Neumann-Morgenstern discounted utility function:

$$E\sum_{t=0}^{\infty}\beta^t U(c_t)$$

where c_t is consumption at date t, U is an increasing concave function and $0 < \beta < 1$ is a subjective discount factor. Labor h_t is either supplied inelastically, or else preferences are modified to incorporate the disutility of work (utility of leisure):

$$E\sum_{t=0}^{\infty}\beta^t U(c_t, h_t).$$

Production takes place according to a two input production function:

$$c_t + (k_t - \lambda k_{t-1}) = d_t f(k_{t-1}, h_t) \tag{1.1}$$

where k_t is capital and d_t is a technology shock, which is a component of a Markov process $\{x_t\}$. The depreciation rate is $1-\lambda$. Associated with this Markov process is a sequence of information sets $\{I_t\}$. In what follows we sometimes adopt the common and convenient Cobb-Douglas specification of production:[1]

$$f(k, h) = k^{\theta} h^{1-\theta}. \tag{1.2}$$

1.1. Parameterizations

We first present the basic utility functions that have been used in empirical work and in many versions of the Brock-Mirman model. We briefly review the micro-econometric evidence on preferences, paying particular attention to the interactions between consumption and labor supply. This evidence is discussed more extensively in Part III. For convenience, some models abstract from the labor supply decision and use an iso-elastic one-period utility function defined

[1] In Part II we will study deficiencies of this Cobb-Douglas specification. In particular, labor is not homogeneous and an efficiency units assumption to adjust labor to homogeneous units is inconsistent with the evidence from factor markets. Comparative advantage and sectoral choices by agents are key features of modern labor markets.

over a single nondurable consumption good:[2]

$$U(c) = \frac{c^{1-\rho} - 1}{1-\rho} \tag{1.3}$$

for $\rho \geq 0$. This specification is used in part because, given intertemporal additivity of preferences, it is homothetic and hence leads to simple steady state characterizations.

To obtain a more interesting model of economic fluctuations, including fluctuations in total or average hours worked, Kydland and Prescott (1982) introduced leisure into the preferences of a Brock-Mirman model (see also King, Plosser and Rebelo (1988a,b) and Cooley and Prescott (1995)). Most subsequent investigators assume that the one-period utility function can be written in the form:

$$U(c,h) = \frac{\left(c^{\sigma}\left[\psi(h)\right]^{1-\sigma}\right)^{1-\rho} - 1}{1-\rho} \tag{1.4}$$

where h is hours of work and ψ is decreasing and positive and σ is in the interval $(0,1)$.[3] When

[2] Some consumption models allow for many goods. For example, many dynamic international trade models follow the tradition of static models and allow that 'traded' and 'non-traded' goods enter the utility function differently; see, for example, Backus, Kehoe and Kydland (1995) and Stockman and Tesar (1995).

[3] For some of dynamic equilibrium models, consumption and labor supply are composites. For example, Kydland and Prescott (1982) have preferences defined over a weighted sum of current and lagged labour supply and Eichenbaum and Hansen (1990) and Hornstein and Praschnik (1994) define consumption as a CES aggregator of the flow of services from durables and non-durables. Auerbach and Kotlikoff (1987) use a CES version of (1.4) in their overlapping generations model.

$\rho = 1$, we obtain the additively separable model:

$$U(c, h) = \sigma \log(c) + (1 - \sigma) \log[\psi\,(h)].$$

1.2. Steady States

With specification (1.4), the marginal rate of substitution between consumption and work is:

$$mrs = \frac{(1 - \sigma)\,\psi'(h)c}{\sigma\psi(h)} \tag{1.5}$$

and hence is linear in consumption. Suppose that there is geometric growth in the technology process $\{d_t\}$. Given the Cobb-Douglas production function (1.2), a steady state exists in which hours worked, the consumption-capital ratio and the implied firm expenditure share on labor costs are constant. Steady state calibration proceeds as follows. Steady states and steady state ratios are measured by time series averages. The production function parameter θ is pinned down by labor's share of output, and the depreciation factor for capital from the steady state investment capital ratio. For a given ψ, say $\psi = 1 - h$, the parameter σ may be solved out by equating minus the marginal disutility of work (1.5) with the marginal product of labor. This yields:

$$(1 - \theta)\,\frac{c + i}{c} = \frac{-\,(1 - \sigma)\,\psi'(h)h}{\sigma\psi(h)}$$

where i is steady state investment.[4]

An important question this theory has to confront is whether the functional forms for preferences of consumption and labor supply used to interpret aggregate time series data as steady-states are compatible with microeconomic evidence on the functional form of preferences. The time series evidence measures the fraction of available time an 'average' worker spends in market work. The claim in the real business cycle literature is that per capita leisure has remained relatively constant in the post-war period while real wages have been rising at the same rate as output. However, this stability in average hours worked per capita masks divergent trends for males and females. A central finding from the empirical micro literature is that the time series of the labor supply behavior of men and women is different and neither is close to a stationary time series. (See Pencavel (1986) and Killingsworth and Heckman (1987)).

If preference parameters are to be based on microeconomic evidence, two questions have to be answered. First, do the functional forms found in the micro literature produce growth steady states? Second, given the changes in the composition of the labor force, whose labor elasticities should be used in calibrating a macroeconomic model? The answers to these questions rely in

[4]Consideration of household production and the possibility of substituting work time in the home for expenditures on goods lead authors such as Benhabib, Rogerson and Wright (1991), Greenwood and Hercowitz (1991) and Greenwood, Rogerson and Wright (1995) to allow for derived utility functions over consumption and market hours that are somewhat more general than the class of models considered here. Their home production specification introduces technological progress into the "reduced from" depiction of the preferences for consumption and labor supply and loosens the restrictions needed for the existence of a steady state of the desired form. See Eichenbaum and Hansen (1990) for a similar development for modeling the preferences for durable and nondurable consumption goods.

part on the relative quality of the aggregate time series and the microeconomic evidence. Durlauf and Quah (1998) raise serious doubts about the validity of the steady state approximation as an accurate factual description of modern economies. Note further that the functional form restrictions required for the conjectured steady states apply to a fictitious composite household model of consumption and leisure. In practice the microeconomic evidence is extracted separately for men and women using preference specifications that are outside the form given in (1.4). For example, MaCurdy (1983) reports that a specification of male labor supply with

$$u(c,h) = \frac{\left(\sigma \frac{c^{1-\alpha_c}}{1-\alpha_c} - \frac{h^{1+\alpha_h}}{1+\alpha_h}\right)^{1-\rho} - 1}{1-\rho}$$

is consistent with monthly male labor supply data from the U.S. where $\rho = .14$, $\alpha_c = .66$ and $\alpha_h = .16$, and these parameters are precisely estimated. Note in particular that $\alpha_c \neq 1$. The marginal rate of substitution between consumption and work is:

$$mrs = -\frac{h^{\alpha_h}}{\sigma c^{-\alpha_c}},$$

and this empirical specification is not consistent with steady state growth because $\alpha_c \neq 1$. It is, however, consistent with the well known observation that male hours of work per capita have declined over time.

1.3. Micro evidence

Our more detailed discussion of the microeconomic evidence presented in Part III establishes the following additional empirical conclusions:

- Most of the responsiveness of labor supply with respect to wage change is due to entry and exit from employment; yet most of the micro evidence for intertemporal labor supply elasticities is presented for continuously working, continuously married prime age males – the demographic group least responsive to wage changes, especially at the extensive margin.[5]

- There is evidence that consumption is complementary with male labor supply while the evidence is mixed on the interaction between consumption and female labor supply. At present there are no reliable estimates of this interaction. Yet the difference between male and female labor supply highlights the problem of pooling the labor supply of diverse groups into one aggregate.

- The elasticity of intertemporal substitution ($eis = -1/\rho$) as determined from consumption is usually poorly determined. If constancy across the population is imposed on this

[5] Rogerson (1988) designed an aggregate model of labor supply that focuses exclusively on the extensive margin. Individuals are allocated randomly to jobs that require a fixed number of hours. The number of jobs fluctuates over time but not the number of hours per job. Hansen (1985) adapted this framework to the Brock-Mirman stochastic growth model. While these models successfully emphasize the extensive margin, they are not well suited to capture differential labor supply responses between men and women. We discuss this model further in Part III.

elasticity, then there is no strong evidence against the view that this elasticity is slightly *above* minus one. There is, however, evidence that the *eis* varies both with observable demographics and with the level of wealth so that the homothetic iso-elastic form is rejected. The same evidence suggests that low wealth households are relatively more averse to consumption fluctuations than are high wealth households.

- For leisure, the elasticity of intertemporal substitution is between 0.1 and 0.4 for annual hours for men and 1.61 for women. There is evidence that these elasticities are larger for shorter units within a year. Because these labor supply elasticities ignore the entry and exit decision, they provide only part of the information needed to construct the aggregate labor supply curve.

2. Multiple Agents

Heterogeneity in preferences, discount rates, and risk aversion parameters is found in numerous micro studies. As a step towards achieving a better match between micro economic evidence and dynamic stochastic economics, it is fruitful to explore macro general equilibrium models with explicit heterogeneity. Such models are of considerable interest in their own right and often produce rather different outcomes than their single consumer counterparts. Adding heterogeneity enriches the economic content of macro models, and calls into question current practices for obtaining parameter estimates used in general equilibrium models.

We start with a very simple specification. Consider a large population with J types of agents indexed by j. We abstract from labor supply as in the Brock-Mirman (1972) stochastic growth model and we also ignore human capital accumulation. Instead we suppose initially that labor is supplied inelastically. Following Aiyagari (1994), we adopt the simple version of the Brock-Mirman model in which individual agents confront stochastic productivity shocks $y_{j,t}$ to their labor supply. This scheme produces idiosyncratic shocks in labor income in spite of the presence of a common wage (per unit productivity) and leads to a simple analysis. Later on, we explore complications caused by the addition of the labor supply decision.

2.1. Preferences

We follow the common practice of using preferences with a constant elasticity of intertemporal substitution but we let the *eis* and the subjective rate of time discount differ among individuals:

$$E \sum_{t=0}^{\infty} \left(\beta_j\right)^t \frac{(c_{j,t})^{1-\rho_j} - 1}{1 - \rho_j} \tag{2.1}$$

where (β_j, ρ_j) differ by consumer type. The evidence discussed both here and in Part III documents that such heterogeneity is empirically important.

2.2. Labor Income

Assume that $\{x_t\}$ is a Markov process governing aggregate shocks and that the productivity for type j at time $t+1$, $y_{j,t+1}$, is a component of a person-specific state vector $s_{j,t+1}$. The probability distribution of $s_{j,t+1}$ given current period state vectors $s_{j,t}$ and x_t is denoted by $F_j(\cdot \mid s_{j,t}, x_t)$. The income of person j at time t is $w_t y_{j,t}$ where w_t is the endogenously determined wage rate at time t. Aggregate or average labor supply corrected for efficiency is given by:

$$h_t^a = \frac{1}{J}\sum_{j=1}^{J} y_{j,t}$$

where J is the number of agents in the economy and the individual labor supply is normalized to be unity. In equilibrium, wages satisfy the aggregate marginal product condition

$$w_t = (1-\theta)d_t(k_{t-1}/h_t^a)^{\theta}.$$

2.3. Market Structure

Depending on the institutional environment confronting agents, different interactions among them may emerge. Market interactions can be limited by informational constraints and the ability to commit to contractual arrangements. Here, we consider a variety of market contexts and their implications for behavior and calibration. We initially explore a complete market model as a benchmark. Suppose consumers trade in a rich array of security markets. We introduce a common sequence of information sets and use I_t to denote information available at date t. Consumers can make state contingent contracts conditioned on information available in subsequent time periods. Given the ability to make such contracts, we obtain a large array of equilibrium security market prices. Moreover, in studying consumption allocations, we may simplify the analysis by exploiting the implications of Pareto efficiency. Although our interest is in economies with heterogeneous consumers, we begin our exposition of Pareto efficient models by first considering agents with homogenous preferences but heterogeneous endowments.

2.4. Preference Homogeneity

Suppose initially that consumers have common preferences (β, ρ). Endowments may differ; in this case, preferences aggregate in the sense of Gorman (1953). At a mechanical level, this can be checked as follows. The intertemporal marginal rates of substitution are equated so:

$$m_{t+1,t} = \beta \left(\frac{c_{j,t+1}}{c_{j,t}} \right)^{-\rho}.$$

Thus

$$c_{j,t} \left(\frac{m_{t+1,t}}{\beta} \right)^{1/(-\rho)} = c_{j,t+1}.$$

Averaging over the consumption of each type, we find that

$$c_{a,t}\left(\frac{m_{t+1,t}}{\beta}\right)^{1/(-\rho)} = c_{a,t+1} \tag{2.2}$$

where $c_{a,t}$ denotes consumption averaged across types. We may solve (2.2) to obtain an alternative expression for the common marginal rate of substitution:

$$m_{t+1,t} = \beta\left(\frac{c_{a,t+1}}{c_{a,t}}\right)^{-\rho}.$$

This result is due to Rubinstein (1974). Under the stated conditions, there exists an aggregate based on preferences that *look like* the common individual counterparts. An alternative way to view this economy is as an example of Wilson's (1968) theory of syndicates. With the marginal rates of substitution equated across consumers, we are led to a solution whereby individual consumption is a constant fraction of the aggregate over time:

$$c_{j,t} = \kappa_j c_{a,t},$$

or equivalently that the allocation *risk-sharing rules* are linear.[6]

[6]The reference to this as a risk-sharing rule is misleading. Consider economies in which endowments of individuals oscillate in a deterministic way, but the aggregate endowment is constant. Then consumption allocations will be constant as implied by the linear allocation rule, but there is no risk.

Armed with this result, the general equilibrium of this model can be computed as follows. Given the aggregate endowment process and the capital accumulation process, we may solve for the optimal aggregate consumption process. This may be thought of as special case of a Brock-Mirman style stochastic growth model in which the fictitious consumer has preferences that coincide with those of the individual identical agents. The solution to this problem yields the equilibrium processes for aggregate consumption, aggregate investment and aggregate capital stock.[7] Notice that we can compute the aggregate quantities without simultaneously solving for the equilibrium prices. It is not necessary to ascertain how wealth is allocated across consumers because we can construct well defined aggregates.[8]

Given our assumption of homothetic preferences, the equilibrium allocation of consumption assigns a constant (over time and across states) fraction of aggregate consumption to each person. Each consumer is endowed with an initial asset stock along with his or her process for consumption endowments. To determine the equilibrium allocation of consumption across people we must solve for the equilibrium valuation of the consumption endowments. With this valuation in hand, the individual consumption assignments are readily deduced from the intertemporal budget constraint. Then we can consider the equilibrium pricing of state-contingent claims to consumption.

[7] See also Lucas and Prescott (1971).

[8] The simplicity here is overstated in one important respect. In characterizing the aggregate endowment behavior, we either must appeal to a cross sectional version of the Law of Large Numbers, or we must keep track of the idiosyncratic state variables needed to forecast individual endowments.

Following Rubinstein (1974) and Hansen and Singleton (1982), pricing implications for this economy may be obtained by using the equilibrium consumption vector and forming the equilibrium intertemporal marginal rates of substitution: the equilibrium versions of $\{m_{t+1,t}\}$. From this process we can construct the pricing operator: $\mathcal{P}_t$. Let z_{t+1} represent a claim to consumption at time $t+1$. For instance, z_{t+1} may be the payoff (in terms of time $t+1$ consumption) to holding a security between dates t and $t+1$. For securities with longer maturities than one time period, we can interpret z_{t+1} as the liquidation value of the security (the dividend at time $t+1$ plus the price of selling the security at $t+1$). The consumption claim z_{t+1} may depend on information that is only observed at date $t+1$ and hence is a random variable in the information set I_t. The equilibrium restrictions of our model imply that the price at time t can be expressed as:

$$\mathcal{P}_t(z_{t+1}) = E(m_{t+1,t} z_{t+1} | I_t)$$

where $\mathcal{P}_t(z_{t+1})$ is the date t price quoted in terms of date t consumption. Thus the pricing operator $\mathcal{P}_t$ assigns time t equilibrium prices to these contingent consumption claims in a linear manner.[9] The intertemporal marginal rate of substitution, $m_{t+1,t}$, acts like a state-contingent discount factor. Since it is stochastic, in addition to discounting the future, it captures risk

[9]We are being deliberately vague about the domain of this operator. Since $m_{t+1,t}$ is positive, any nonnegative payoff can be assigned an unambiguous but possibly infinite value. To rule out value ambiguities that arise when the positive part of the payoff has value $+\infty$ and the negative part $-\infty$, additional restrictions must be imposed on payoffs that depend on properties of the candidate discount factors.

adjustments for securities with uncertain payouts.

For this economy we may extract preference parameters using Euler equation estimation. Let Z_{t+1} denote a vector of one-period (gross) payoffs and Q_t the corresponding price vector. The preference parameter vector (β, ρ) can be identified from the unconditional moment restriction:

$$E\left[\beta\left(\frac{c_{t+1}}{c_t}\right)^{-\rho} Z_{t+1} - Q_t\right] = 0. \tag{2.3}$$

The asset payoffs may consist of multiple security market returns, or they may be synthetic payoffs constructed by an econometrician.[10] We obtain the same preference parameters when the model is estimated using aggregate or individual data on consumption. In principle, this provides a way to test the assumed preference specification given the market environment if there is access to micro data. As noted by Rubinstein (1974), Hansen (1987), and, in effect, Gorman (1953), this result is special, even under preference homogeneity. It is not applicable to any concave increasing one-period utility function unless the consumers have the same initial wealth.

Given knowledge of the technology parameters, the stochastic process for aggregate capital stocks may be deduced as in Brock and Mirman (1972) by endowing the representative agent

[10] Synthetic payoffs and prices are constructed as follows. Multiply a single gross return, say $1 + r_{t+1}$, by instrumental variables in the conditioning information set I_t. Since this conditioning information is available to the consumer at purchase date t, the price of the scaled return is given by the instrumental variable used in the scaling. By creating enough synthetic securities, unconditional moment condition (2.3) can be made to imititate the conditional pricing relation (see Hansen and Singleton, 1982, and Hansen and Richard, 1987).

just with preference parameters and labor supply h_t^a. From this solution we may solve for equilibrium aggregate consumption, equilibrium stochastic discount factors (from intertemporal marginal rates of substitution), equilibrium wages (from marginal products), the initial wealth distribution and hence the sharing parameters (the κ_j's).

2.5. Risk Aversion or Intertemporal Substitution?

By estimating Euler equations, econometricians may identify preference parameters without having to solve the decision problem confronting individual agents. In particular, parameters can be identified without any need to measure microeconomic uncertainty or wealth. Of course if preferences are misspecified by an econometrician, estimated Euler equations will not recover the true preference parameters of individual agents. We now consider a misspecification of particular substantive interest.

The parameter ρ is associated with two conceptually distinct aspects of preferences: risk aversion and intertemporal substitution along certain paths. The link between attitudes towards risk and intertemporal fluctuations in consumption over time is indissolubly embedded in models that simultaneously assume separability over time and over states (see Gorman, 1968).[11] It is the latter type of separability that is the key assumption in expected utility models. Hall (1988) and Epstein and Zin (1989, 1991) argue that it is fruitful to disentangle attitudes toward risk from

[11] The parameter ρ also governs precautionary savings or prudence (see Kimball, 1990). This latter link is readily broken by adopting a more flexible parameterization of expected utility.

intertemporal substitution as they correspond to two different aspects of consumer behavior. Concern about intertemporal substitution comes into play even in economies with deterministic movements in technologies. These considerations led Epstein and Zin (1989) to use a recursive utility formulation due to Kreps and Porteus (1978) in which preferences are represented using "continuation utility" indices, which measure the current period value of a consumption plan from the current period forward. The continuation utility index, $V_{j,t}$, for person j is obtained by iterating on the recursion:

$$V_{j,t} = \left[(1-\beta)\left(c_{j,t}\right)^{1-\rho} + \beta \mathcal{R}_t(V_{j,t+1})^{1-\rho}\right]^{1/(1-\rho)} \tag{2.4}$$

where $\mathcal{R}_t$ makes a risk adjustment on tomorrow's utility index:

$$\mathcal{R}_t(V_{j,t+1}) = \left(\left[E\left(V_{j,t+1}\right)^{1-\alpha} |I_t\right]\right)^{1/(1-\alpha)}. \tag{2.5}$$

Observe that the utility index today is homogeneous of degree one in current and future (state-contingent) consumption. This specification of preference nests our specification (2.1) with a common ρ provided that $\alpha = \rho$. By allowing α to be distinct from ρ we break the connection between risk aversion and intertemporal substitution. The parameter α is irrelevant in an environment without uncertainty, but the intertemporal elasticity of substitution $(-1/\rho)$ parameter is still important. The parameter α makes an incremental risk adjustment, which is absent from the standard von Neumann-Morgenstern formulation.

Some of our previous analysis carries over directly to this recursive utility formulation. The efficient allocation for individual consumptions and individual utility indices satisfies:

$$\begin{aligned} c_{j,t} &= \kappa_j c_{a,t} \\ V_{j,t} &= \kappa_j V_{a,t} \end{aligned}$$

for some numbers κ_j where $V_{a,t}$ is constructed using the process for the representative consumer $\{c_{a,t+k} : k = 0, 1, ...\}$ in place of $\{c_{j,t+k} : k = 0, 1, ...\}$. With this modification, the procedure we have previously described for solving out Brock-Mirman economies applies for preferences of the type (2.4).

The intertemporal marginal rates of substitution are, however, altered, and this complicates the construction of the one-period stochastic discount factors. With the Kreps-Porteus utility recursion,

$$m_{t,t+1} = \beta \left(\frac{c_{a,t+1}}{c_{a,t}} \right)^{-\rho} \left[\frac{V_{a,t+1}}{\mathcal{R}_t(V_{a,t+1})} \right]^{\rho - \alpha}$$

which now depends on the continuation utility $V_{a,t+1}$. The same formula works if individual consumptions and continuation utilities are used in place of the aggregates. Micro or macroeconomic estimation procedures based on Euler equations that erroneously assume that $\alpha = \rho$ generally fail to produce a usable estimate of either α or ρ unless they are equal. Even if a risk-free asset is used in estimating Euler equation (2.3), the intertemporal substitution parameter will not be identified. The presence of risk aversion parameter $\alpha(\neq \rho)$ will alter the Euler equation if there is any uncertainty affecting the consumption decisions of the individual (see Epstein and Zin, 1991). For this case, a correctly specified Euler equation contains the continuation utility ($V_{j,t+1}$) and its conditional moment.[12] If the risk adjustment is logarithmic ($\alpha = 1$), then a logarithmic version of the Euler equation will recover ρ provided the return on the wealth portfolio is used as the asset return instead of the risk free return (see equation 18 in Epstein and Zin, 1991).[13]

Since continuation utilities now enter correctly specified Euler equations, one way to modify the unconditional moment restrictions used in estimation is to solve recursion (2.4) for values of the preference parameters. This solution requires knowledge of the equilibrium consumption process for either individuals or the aggregate. Thus it is no longer possible to separate the estimation of preference parameters from the estimation of the other features of the model as is conventional in the standard Euler equation approach. In particular, it is necessary to specify the underlying uncertainty individuals confront. Given that this explicit computation of the full

[12] Epstein and Zin (1991) present a clever solution to this problem whereby they derive an alternative Euler equation that depends instead on the one-period return on a hypothetical wealth portfolio. In practice it is difficult to construct a reliable proxy that is compatible with the observed consumption data. In their Euler-equation estimation using aggregate data, Epstein and Zin (1991) used the value-weighted return on the New York Stock Exchange, but this proxy only covers a component of wealth in the actual economy.

[13] When the risk adjustment is made using the negative exponential counterpart to (2.5), then Euler equation (2.3) continues to apply but with an endogenously determined distorted conditional expectation operator. This risk adjustment is equivalent to inducing a specific form of pessimism. See Hansen, Sargent and Tallarini (1998) for a discussion of this point.

model is required, there are other, more direct, approaches to estimation than plugging solved continuation utilities into Euler equations in a two-stage procedure.[14] Barsky, Juster, Kimball and Shapiro (1997) pursue an alternative way of measuring risk preferences (independent of intertemporal substitution) that is based on confronting consumers with hypothetical employment gambles. We discuss their work below.

2.6. Preference Heterogeneity

Even when we restrict ourselves to state-separable power-utility functions, once we allow for heterogeneity in preferences we must modify our aggregation theorem and our method for solving the general equilibrium model. We continue to impose the complete market structure, but drop back to the simple additively-separable preference specification. It is again convenient to pose the equilibrium problem as an optimal resource allocation problem for the purpose of computing equilibrium quantities. To accomplish this in a world of heterogeneous preferences we use a method devised by Negishi (1960) and refined by Constantinides (1982), Lucas and Stokey (1984) and others.

Using standard Pareto efficiency arguments, and assuming interior solutions, consumers

[14] An interesting question is what security data are needed to identify the risk adjustment in the utility index. Epstein and Melino (1995) address this question without imposing parametric restrictions on the risk adjustment. Not surprisingly, the degree of identification depends on the richness of the security market returns used in the investigation. When investors have access to more security markets and the resulting returns are observed by an econometrician, the range of admissible risk adjustments shrinks. Hansen, Sargent and Tallarini (1998) illustrate this point in the parametric context of a permanent income model with an exponential risk adjustment.

equate their intertemporal marginal rates of substitution

$$(\beta_j)^t \left(\frac{c_{j,t}}{c_{j,0}}\right)^{-\rho_j} = (\beta_1)^t \left(\frac{c_{1,t}}{c_{1,0}}\right)^{-\rho_1}.$$

For each individual j, we assign a time t Pareto weight $\omega_{j,t}$ with a deterministic equation of evolution:

$$\omega_{j,t} = (\beta_j/\beta_1)\omega_{j,t-1}. \tag{2.6}$$

Equating intertemporal marginal rates of substitution we obtain:

$$\omega_{j,t}(c_{j,t})^{-\rho_j} = \omega_{1,t}(c_{1,t})^{-\rho_1}.$$

We may thus characterize Pareto efficient allocations by combining evolution equations (2.6) with the solution to the static deterministic optimization problem:

$$\max_{c_1,c_2...,c_J} \sum_{j=1}^{J} \omega_j \frac{(c_j)^{1-\rho_j} - 1}{1-\rho_j} \text{ subject to } \frac{1}{J}\sum_{j=1}^{J} c_j = c. \tag{2.7}$$

The solution to problem (2.7) is obtained from the following argument. Let μ denote the common marginal utility across individuals:

$$\omega_j(c_j)^{-\rho_j} = \mu.$$

Then,

$$c_j = \left(\frac{\mu}{\omega_j}\right)^{1/(-\rho_j)} \tag{2.8}$$

By first averaging this expression over individuals, we may compute μ by solving the nonlinear equation:

$$c = \frac{1}{J}\sum_{j=1}^{J}\left(\frac{\mu}{\omega_j}\right)^{1/(-\rho_j)}.$$

Plugging the solution for the common marginal utility μ back into (2.8), we obtain the allocation equations:

$$c_j = \phi_j(c;\omega), \qquad j = 1, .., J,$$

where ω denotes the vector of Pareto weights. The allocation rules ϕ_j are increasing and must average out to the identity map. Substituting these rules back into the original objective function we construct a utility function for aggregate consumption:

$$U(c;\omega) = \sum_{j=1}^{J}\omega_j\left[\frac{\phi_j(c;\omega)^{1-\rho_j} - 1}{1-\rho_j}\right]$$

It is straightforward to verify that $U(c,\omega)$ is concave and strictly increasing. By the Envelope Theorem,

$$\frac{\partial U(c;\omega)}{\partial c} = \omega_j\left[\phi_j(c;\omega)\right]^{-\rho_j} = \mu, \tag{2.9}$$

which is the common marginal utility. The "mongrel" function U will generally depend on the Pareto weights in a nontrivial manner.

We may use this constructed utility function to deduce optimal allocations in the following way. Given any admissible initial ω_0, solve the optimal resource allocation problem using the preference ordering induced by the von Neumann-Morgenstern "mongrel" utility function:

$$E \sum_{t=0}^{\infty} (\beta_1)^t U(c_t; \omega_t)$$

subject to the equation of motion (1.1) and the evolution equation for vector of Pareto weights (2.6). If resources are allocated efficiently, we obtain an alternative (to Gorman) justification of the representative consumer model. This justification carries over to the derived pricing relations as well. Prices may be deduced from the marginal rates of substitution implied by the mongrel preference ordering. This follows directly from the link between individual and aggregate marginal rates of substitution given in (2.9).

This construction justifies a two-step method for computing efficient allocations of resources when preferences are heterogeneous. In the first step we compute a mongrel utility function for a fictitious representative consumer along with a set of allocation rules from a static and deterministic resource allocation problem. The mongrel utility function may be used to deduce equilibrium aggregate consumption and investment rules and equilibrium prices. The static allocation rules may be used for each state and date to allocate aggregate consumption among the individual consumers. These computations are repeated for each admissible choice of Pareto weights. To compute a particular general equilibrium, Pareto weights must be found that satisfy the intertemporal budget constraints of the consumer with equality.

This economy has the following observable implications. First, under discount factor homogeneity, the procedure just described can be taken as a justification for using a representative agent model to study aggregate consumption, investment and prices. Microeconomic data are not required to calibrate the model provided that the mongrel preferences used to compute the general equilibrium are not used for welfare analyses. Second, one can use microeconomic and macroeconomic data together to test whether individual consumption data behave in a manner consistent with this model.

If discount factors are heterogenous, and if this economy runs for a long period of time, in the long run the consumer with the largest discount factor essentially does all of the consuming.[15] This follows directly from equation (2.6). Thus it is the discount factor of the (eventually) wealthiest consumer that should be of interest to the calibrator of a representative agent model, provided that the aim is to confront aggregate time series data and security market data. Since in the U.S. economy, 52% of the wealth is held by 5% of the households, and there is evidence,

[15] Lucas and Stokey (1984) use this as a criticism of models with discount factors that are independent of consumption.

discussed below, that wealthy families have lower discount rates, this observation is empirically relevant. A research synthesis that uses a *median* or *trimmed mean* discount factor estimated from micro data to explain long run aggregate time series data would be seriously flawed. This raises the potential problem that the estimated extreme value may be a consequence of sampling error or measurement error instead of genuine preference heterogeneity.

If the aim is to evaluate the impact of macroeconomic policies on the welfare of the person with the *median* discount rate, then the welfare evaluation should be performed outside the representative consumer framework used for matching aggregate time series data. To be accurate it would have to recognize the diversity of subjective discount rates in the population.

With discount factor homogeneity, individual consumption will be a time invariant function of aggregate consumption. This occurs because evolution equation (2.6) implies that the Pareto weights are invariant over time. In spite of this invariance, if there is heterogeneity in intertemporal substitution elasticities, in an economy with growth it may still be the case that one type of the consumer eventually does most of the consumption because of nonlinearity in the allocation rule. To demonstrate this we follow Dumas (1989) and, suppose that we have two types of consumers, both facing a common discount factor, but

$$\rho_1 < \rho_2.$$

This difference in the intertemporal substitution elasticity is sufficient for consumers of type 1 to eventually do most of the consumption in the economy. Simulations are depicted in Figures I-1 and I-2 for a Dumas-style economy. With all consumers facing a common β, to explain the data on aggregate quantities and security market prices, it is again the preferences of the eventually wealthiest consumer that matter and not the preferences of the average or median person.

When we add labor supply into this setup, the lessons for picking parameters is different. Suppose for simplicity, we assume individual preferences are additively separable between consumption and labor supply so that for an individual of type i:

$$U_i(c,h) = \sigma_i \frac{c^{1-\rho_i}-1}{1-\rho_i} + (1-\sigma_i)\frac{\psi_i(h)^{1-\rho_i}-1}{1-\rho_i}.$$

Let the preference parameter for the first type satisfy $\rho_1 = 1$ and assume $\rho_1 < \rho_2$. Consider a social planner allocating consumption and hours across individuals in a Pareto efficient manner. Provided that hours can be converted into an efficiency units standard and the hours allocations for each individual are away from corners (assumptions we question in Parts II and III), we can derive an aggregation result for the disutility of aggregate hours using the same techniques just described for aggregate consumption. While person 1 eventually does most of the consuming in the economy, the leisure preferences of person 2 will figure prominently in the aggregate (mongrel) preference ordering. Thus to construct the mongrel representative agent for this economy requires that the analyst use the intertemporal consumption elasticity of the rich person, but an intertemporal elasticity for labor supply that recognizes both agent types. In the presence of heterogeneity in preferences, it will sometimes be necessary to apply different weighting schemes across the population for consumption elasticities than for labor supply elasticities in order to construct an aggregate that fits the data.[16]

Tests of the complete-market model usually focus on linear allocation rules, whereas accounting for preference heterogeneity leads one to recognize important nonlinearities in allocation rules.[17] In models with endowment uncertainty, but heterogeneity in discount rates and intertemporal elasticities of substitution, the efficient allocation is that individual consumption is a deterministic function of aggregate consumption alone. Even this implication can be altered, however, if von Neumann-Morgenstern preferences are replaced with a more general recursive formulation in which (as in Epstein and Zin 1991) future utilities are adjusted for risk. In this case the evolution of the Pareto weights is stochastic (see Dumas, Uppal and Wang, 1997 and Anderson, 1998). As a consequence, allocation rules depend, not only on current aggregate consumption, but also on the past history of the aggregate state variables (the capital stock and x_t). The deterministic relationship between individual consumption and aggregate allocations will be altered if time nonseparabilities are introduced in the form of habit persistence or dura-

[16] In a somewhat different setting, Kihlstrom and Laffont (1979) use preference heterogeneity to build a model in which more risk averse individuals become workers and less risk averse individuals become entrepreneurs.

[17] See Attanasio (1998) for a survey of the literature testing for linear allocation rules.

bility in consumption goods. In these cases, past histories of consumption also enter into the allocation rules. Thus the fix up of mongrel preferences has to be substantially altered when we consider more general preference specifications.

We now present our first bit of evidence on the empirical importance of preference heterogeneity. This evidence demonstrates that our concerns about this issue are not purely aesthetic.

2.7. Questionnaire Evidence on the Scale and Distribution of Risk Aversion

In an innovative paper, Barsky, Juster, Kimball and Shapiro (1997) elicit risk preferences from hypothetical questions administered to a sample of respondents in the University of Michigan Health and Retirement Survey (HRS). The aim of this study was to extract the degree of relative risk aversion without linking it to the elasticity of intertemporal substitution. Respondents were asked about their willingness to participate in large gambles of various magnitudes. For example, adult respondents are asked to imagine themselves as the sole earners in their families and are asked to choose between a job with their current annual family income guaranteed for life versus a prospect with a 50-50 chance of doubling family income and a 50-50 chance of reducing family income by a third. Respondents who take the gamble offered in the first question are then asked if they would take an option with the same gain as offered in the first question, and the same probability of gain, but a greater loss that cuts income in half. Respondents who decline the gamble offered in the first question are offered a second option with the same gain (and probability of gain) as the first option but the loss is reduced to 20 percent.

Answers to these questions enable one to bound the coefficient of relative risk aversion. The results of this hypothetical exercise are reported in Table I-1 for a variety of demographic groups. The two notable features of this table are: (a) the substantial proportion of risk averse people and (b) the heterogeneity in risk aversion both across demographic groups and within them.

There are serious questions about the relationship between actual risk-taking behavior and the responses elicited from questionnaires. In addition, there are serious questions about the magnitude of the gambles in these hypothetical choice experiments. Preferences that exhibit constant relative risk aversion, link the aversion to small gambles to the aversion to large gambles and hence justify calibrating risk aversion to large gambles. Behavior responses to small bets may be different from the responses to large ones, and the bets studied in this survey are indeed substantial. In fact, Epstein and Melino (1995) provide empirical evidence that the risk aversion may be much larger for small gambles than large ones. Nonetheless, the results summarized in Table I-1 are very suggestive of considerable heterogeneity in the population. We present further evidence on preference heterogeneity in Part III.

We next consider more general market environments without full insurance, and the empirical challenges that arise in constructing and calibrating general equilibrium models in such environments.

3. Incomplete Markets

While the multiple consumer, Pareto optimal economy is pedagogically convenient, it assumes the existence of a rather large collection of markets. Moreover, it eliminates many interesting policy questions by assumption such as those having to do with borrowing constraints or adverse selection. We now consider what happens when most of the markets assumed to be open in the Pareto optimal economy are closed down. Following Miller (1974), Bewley (1977), Scheinkman and Weiss (1986), Aiyagari (1994) and Krusell and Smith (1998), among others, we suppose that consumers can only purchase and sell shares of the capital stock and are not permitted to trade claims to their future individual productivities. Moreover, only nonnegative amounts of capital can be held. This is an economy with a "borrowing" constraint and other forms of market incompleteness that hold simultaneously. We discuss how these constraints can arise later on.

In an environment with incomplete markets, we can no longer exploit the convenient Pareto formulation. The economy we consider is one in which prices and quantities must be computed simultaneously. Under the efficiency wage assumption, the current period wage rate satisfies a standard marginal productivity condition. The gross return to holding capital must satisfy:

$$1 + r_{t+1} = \theta d_{t+1} (k_t / h^a_{t+1})^{\theta - 1} + \lambda.$$

For the special case of Aiyagari's economy, there is no aggregate uncertainty. As a consequence, the equilibrium rate of return to capital will be riskless. In Krusell and Smith (1998), there is aggregate uncertainty but this uncertainty is sufficiently small so that there is little difference between the risky return on capital and a riskfree security.[18] Given that only nonnegative amounts of capital can be held, the familiar consumption Euler equation is replaced by:

$$E\left[\beta_j\left(\frac{c_{j,t+1}}{c_{j,t}}\right)^{-\rho_j}(1+r_{t+1})\,|I_t\right]\geq 1 \tag{3.1}$$

where equality holds when the consumer chooses positive holdings of the capital stock. The familiar interior solution Euler equation no longer characterizes some of the agents in the economy (see, for example, Zeldes, 1989).

In an attempt to explain the wealth distribution, Krusell and Smith (1998) introduce a form of discount factor heterogeneity modeled as a persistent stochastic process. They do not, however, make specific use of the preference heterogeneity measured from microeconomic consumption studies in their calibration. Furthermore, as we will see, calibrating to microeconomic data in environments with less than full insurance requires more than just extracting preference parameters; it requires measuring the uninsured uncertainty facing agents.

[18] As a consequence, like the original Brock-Mirman model, theirs is a poor model of the aggregate return on equity measured using stock market data.

3.1. Microeconomic uncertainty

Aiyagari (1994) and Krusell and Smith (1998) attempt to quantify the impact of the precautionary motive for savings on both the aggregate capital stock and the equilibrium interest rate, assuming that the source of uncertainty is in individual labor market productivity. In order to accomplish this task, these analysts require a measure of the magnitude of microeconomic uncertainty, and how that uncertainty evolves over the business cycle. Euler equation estimates of preference parameters must be supplemented by measures of individual uncertainty. This introduces the possibility of additional sources of heterogeneity because different economic agents may confront fundamentally different risks. To calibrate the macroeconomic model it becomes crucial to measure the distribution of individual shocks. The income of person j at time $t+1$ is $w_{t+1}y_{j,t+1}$ and its distribution depends in part on the aggregate state variable x_t. In practice, household income arises from many sources with possibly different risks and people in different occupations face quantitatively important differences in the uncertainty they confront (see the evidence in Carroll and Samwick, 1997). Aggregating income from all sources or pooling agents in different risk classes is a potentially dangerous practice that may mask the true predictability of the individual components. Aggregates of income sources may not accurately represent the true economic uncertainty facing agents. The persistence in the idiosyncratic shocks and the manner in which aggregate state variables shift the distributions of idiosyncratic shocks are known to have an important impact on consumption allocations in incomplete market models (see Mankiw, 1986 and Constantinides and Duffie, 1996). Aggregating across risk components can alter the measured predictability.

We now present evidence on the time series processes of labor earnings and wage innovations drawing on research by MaCurdy (1982) and Abowd and Card (1989). Hubbard, Skinner and Zeldes (1994) consider measures of uncertainty for other sources of household income. We summarize the econometric evidence on the form of the forecasting equation represented as an ARMA model, and the magnitude of the innovation variance, which is often used as a measure of uncertainty.

3.1.1. Estimated Processes for Wages and Earnings

There is agreement among micro studies of nationally based representative samples that differences in the residuals of male log earnings or wage rates from a Mincer earnings function are adequately represented by either an MA(2) process or an ARMA (1,1) process. There is surprisingly little information available about the time series processes of female earnings and wage rates. The representations for the change in log earnings and wage rates for married males that receive the most support in the studies of MaCurdy (1982) and Abowd and Card (1989) are:

$$\Delta u_{i,t} = \varepsilon_{i,t} + m_1 \varepsilon_{i,t-1} + m_2 \varepsilon_{i,t-2} \tag{3.2}$$

or

$$\Delta u_{i,t} = a\Delta u_{i,t-1} + \varepsilon_{i,t} + m_1\varepsilon_{i,t-1} \tag{3.3}$$

where

$$\Delta u_{i,t} = u_{i,t} - u_{i,t-1}$$

and $u_{i,t}$ is the residual of a Mincer regression for log earnings or wage rates for person i. (See Section II-2 for a discussion of Mincer earnings models.) Estimates of the parameters of these models are presented in Table I-2 (taken from MaCurdy, 1982). He reports that he is unable to distinguish between these two representations of the time series process of residuals. In Table I-3 we report MaCurdy's (1982) estimates when the autoregressive unit root specification is not imposed in (3.3). The freely estimated autoregressive coefficients are close to but slightly less than one, and they are estimated with enough accuracy to reject the unit root model using statistical tests.

The analysis of Abowd and Card (1989) is generally supportive of the results reported by MaCurdy (1982), except that MaCurdy finds that the coefficients m_1, m_2, are constant over time whereas Abowd and Card report a rejection of the overall hypothesis of stationarity for the model. There is no necessary contradiction between the two studies because MaCurdy does not require that the variances of the $\varepsilon_{i,t}$ be constant, nor does he report evidence on the question.

However, he uses the assumption of constancy in the innovation variances of earnings processes to report the innovation variances given in the final column of Tables I-2 and I-3.

From the vantage point of macroeconomics, time series variation in the innovation variances is of interest, especially the extent that the variances fluctuate over the business cycle. Figures I-3 and I-4 demonstrate how the conditional variance of Δu_{it} changes over the business cycle. In periods of rising unemployment, the innovation variance in log wage equations increases. This evidence is consistent with the notion that microeconomic uncertainty is greater in recessions than booms.[19]

As we have previously noted, the models described in this section take households as the decision unit. This led researchers such as Heaton and Lucas (1996) and Hubbard, Skinner and Zeldes (1994) to present estimates of pooled family income processes.[20] They do not report separate estimates of the earnings processes for husbands and wives. In samples for earlier periods, there is evidence of negative covariance between spousal earnings (Holbrook and Stafford, 1971). In later samples, there is evidence of positive covariance (Juhn and Murphy, 1994).[21] We question whether pooling household earnings processes is a sound practice for extracting

[19] The evidence for professional and educated workers reported by Lillard and Weiss (1979) and Hause (1980) suggests the presence of a person-specific growth trend. This growth trend is not found to be important in national samples across all skill groups and for the sake of brevity we do not discuss specifications of the earnings functions with this growth trend.

[20] Also, to obtain a better match with their model, Heaton and Lucas (1996) look at individual income relative to aggregate income. In effect they impose a cointegration relation between individual and aggregate log earnings.

[21] However, one should not make too much of this difference. The Holbrook and Stafford study reports a relationship for panel data; the Juhn and Murphy study is for cross sectional data.

components of risk facing households in models designed to track the wealth distribution of the economy. Each income source is likely to have its own component of uncertainty. Further research on this topic would be highly desirable.

3.1.2. Missing Risks

Most of the microeconomic evidence is based on samples of annual earnings or average hourly wages for continuously working, continuously married males. Risks associated with long term job loss due to job displacement, illness or marital disruption are typically ignored. On these grounds the estimated innovation variances to narrowly defined earning processes are likely to *understate* the risks confronted by agents. Many labor force and life-cycle sources of risk are abstracted from in these studies. Carroll (1992), Hubbard, Skinner and Zeldes (1994), and Lillard and Weiss (1997) make this point, and estimate additional components of uncertainty.

3.1.3. Statistical Decompositions

Statistical decompositions of wage and earnings processes are intrinsically uninformative about the information available to economic agents. As in Friedman (1957), all of the components of (3.2) or (3.3) could be known and acted on by agents. Estimated innovation variances include (a) measurement error components; (b) factors known to the agents and unknown to the econometrician; and (c) true components of uncertainty. On these grounds, the estimates of the variances in empirical earnings equations likely *overstate* the true uncertainty facing agents.

To demonstrate the value of the cross-equation restrictions connecting consumption and earnings in identifying the innovation in earnings, consider the following example based on the permanent income model of consumption. Suppose that the first difference of the level of labor income (earnings) e_t evolves according to:[22]

$$\Delta e_t = \varphi + \sum_{j=0}^{\infty} \alpha_j \cdot \eta_{t-j}$$

where $\{\eta_t\}$ is a nondegenerate stationary, multivariate martingale difference sequence with a finite second moment. We impose as a normalization that the covariance of η_t is the identity matrix. The α_j's are vectors of moving-average coefficients. The martingale difference sequence $\{\eta_t\}$ is adapted to the sequence of information sets available to the consumer. The multiple components are introduced into η_t to capture the multiple sources of uncertainty in, say, household income.

For simplicity, we assume the preferences for consumption are quadratic, as in Flavin (1981) or Hansen (1987), and that the real interest r is constant.[23] Then from those analyses we know that the change in consumption from date $t-1$ to t, $c_t - c_{t-1}$, is just the change in the flow of

[22] This specification includes ones used by Friedman in his Ph.D thesis, see Friedman and Kuznets (1945). In contrast to the processes fit by MaCurdy (1982), this specification is depicted in terms of first-differences of income levels instead of first-differences of logarithms.

[23] See Hansen (1987) for a general equilibrium interpretation of this model.

discounted current and future income from t to $t-1$:

$$c_t - c_{t-1} = \left[\sum_{j=0}^{\infty} \left(\frac{1}{1+r}\right)^j \alpha_j\right] \cdot \eta_t. \tag{3.4}$$

Not only is consumption a martingale as noted by Hall (1978), but the composite income-consumption process must be present-value neutral (see Hansen, Roberds and Sargent, 1991). That is, consumption and income must respect a present-value budget constraint for all realizations of the shock vector η_t. Relation (3.4) respects this constraint by taking account of the fact that the shock η_t alters income in future time periods.

Even in the absence of measurement errors in consumption and income, the entire innovation vector η_t cannot necessarily be identified by an econometrician using data on income alone. For instance, if η_t has more than one entry, fitting a univariate time series process to income will not reveal this vector of new information pertinent to the consumer. By looking instead at consumption, η_t can at least be partially identified because consumption is a martingale adapted to true information set of the consumer. It follows from (3.4) that the first-difference in consumption reveals one linear combination of the shock vector η_t.[24]

[24] Hansen, Roberds and Sargent (1991) use this result and generalizations of it to deduce testable implications of present value budget balance.

Suppose now that measured earnings are

$$e_t^* = e_t + \nu_t$$

where the measurement error $\{\nu_t\}$ is mean zero, finite variance and is independent of the process $\{ \eta_t\}$, but the serial correlation properties of the $\{\nu_t\}$ are unrestricted. Then from the autocovariances of $\{\Delta e_t^*\}$ one cannot identify the moving-average coefficients (the α_j's) even if $\{\eta_t\}$ is a scalar process. Thus the strategy of estimating the variance of the innovations in information from the covariances of the error processes of measured earnings equations, as used by Carroll (1992), Hubbard, Skinner and Zeldes (1994) and others, fails, because it is impossible to separate the measurement error from earnings innovations. Again consumption data is informative because agents respond to innovations in earnings but not to the measurement error.

Suppose that consumption is also contaminated by measurement error. Thus we write:

$$\Delta c_t^* = \left[\sum_{j=0}^{\infty} \left(\frac{1}{1+r} \right)^j \alpha_j \right] \cdot \eta_t + \upsilon_t$$

where $\{c_t^*\}$ is measured consumption and $\{\upsilon_t\}$ is a measurement error process that is independent of $\{\eta_t\}$ and $\{\nu_t\}$. We may now use the cross covariances between the $\{\Delta c_t^*\}$ and $\{\Delta e_t^*\}$ to obtain at least partial identification of the income information structure confronting the consumer. This follows from the formula:

$$Cov(\Delta e^*_{t+k}, \Delta c^*_t) = \alpha_k \cdot \left[\sum_{j=0}^{\infty} \left(\frac{1}{1+r} \right)^j \alpha_j \right],$$

which uses the fact that η_t has an identity as its covariance matrix and is valid for $k \geq 0$. Given prior knowledge of the constant interest rate r, we may use this formula and formula (3.4) to deduce the variance of the Δc_t:

$$\begin{aligned} Var(\Delta c_t) &= \left[\sum_{k=0}^{\infty} \left(\frac{1}{1+r} \right)^k \alpha_k \right] \cdot \left[\sum_{j=0}^{\infty} \left(\frac{1}{1+r} \right)^j \alpha_j \right] \\ &= \sum_{k=0}^{\infty} \left(\frac{1}{1+r} \right)^k Cov(\Delta e^*_{t+k}, \Delta c^*_t). \end{aligned}$$

From these calculations then we may again infer how the true income process responds to a linear combination of the η_t shock vector.[25] If the $\{\eta_t\}$ process is scalar, then we have full identification of the information structure confronting the consumer that is pertinent for the evolution of income. Thus the use of consumption data, even if measured with error, can help to identify the true income uncertainty confronting economic agents.

We conclude this section with a brief discussion of the sources of market incompleteness.

[25] Hansen, Roberds and Sargent (1991) use a closely related argument to show that present-value budget balance has testable implication when consumption is a martingale. They do not, however, explore the ramifications of measurement error in consumption and income

3.2. Limited Commitment and Private Information

In our incomplete markets model, we made no attempt to justify the form of market incompleteness. Two common justifications include problems of enforceability and observability. Measurements of microeconomic uncertainty are critical ingredients to economies that explicitly account for limited commitment and private information.

3.2.1. Limited Commitment

Kehoe and Levine (1993) and Kocherlakota (1996) propose the following alternative to the incomplete markets models we have considered thus far. Suppose consumers are permitted to walk away from obligations, but when they do so they are excluded from future participation in markets. Instead they are restricted to use their own backyard storage technologies or simply consume only their labor income. As a consequence, in equilibrium, consumers are guaranteed a lower bound on their discounted utilities at each date and state. Consumption allocations are obtained by solving Pareto problems subject to utility lower bounds implied by the utility threat points: the points at which consumers are indifferent between honoring their obligations and defaulting. While there is no default in equilibrium, allocations are altered by the presence of the utility threat points.

Alvarez and Jermann (1998) consider an economy with limited commitment in which consumers have no access to backyard storage technologies and hence are punished by constraining future consumption to equal future income, period by period. They show that the resulting allocations may be decentralized by introducing person-specific solvency or borrowing constraints.[26] Euler equations are replaced by Euler inequalities as in Luttmer's (1996) work on asset pricing when investors face solvency constraints. Alvarez and Jermann (1998) are able to imitate asset pricing features of an economy with solvency constraints, but with different predictions about when agents will be up against financial market constraints. Individual Euler equations linking consumption to asset prices continue to characterize the behavior of some individuals in each time period. When and which consumers are constrained in their ability to borrow can be ascertained by computing the utility threat points, which in turn depend on the microeconomic uncertainty they would be forced to confront in the absence of risk sharing. Thus the same problem of measuring uncertainty discussed for economies with incomplete markets applies to economies in which limited commitment is the only source of financial market frictions.

3.2.2. Private Information

In the limited commitment economies considered by Kehoe and Levine (1993) and by Alvarez and Jermann (1998), idiosyncratic endowment shocks are publicly observed. Suppose instead that they are only known to the individual agents and not to the public. Again, the options

[26] The person-specific nature of the solvency constraint stretches a bit the notion of a decentralized economy. As an alternative, we may view these limited commitment economies as (constrained) efficient benchmarks to which we might compare alternative institutional arrangements.

and information available to agents matter. For instance, suppose that a capital accumulation technology is only available to the society as a whole and not to individuals. In this case, individual consumption contracts can be enforced because the agents are unable to privately transfer consumption from one period to the next. There is a substantial body of work on optimal resource allocation subject to incentive constraints that relies on the enforceability of consumption contracts (see Green, 1987; Phelan and Townsend, 1991; Atkeson and Lucas, 1992 and others). In general, the efficient allocations do not have decentralizations that look like the incomplete market structure previously described. Economies with a simple security market structure are not Pareto efficient even after accounting for the incentive constraints (*e.g.* see Lucas, 1992). Of course, security markets could be supplemented by other institutions designed to reduce or eliminate the efficiency wedge. In contrast, when a capital accumulation technology is privately available, individual agents can hide their consumption from the public. Thus individual consumption contracts are no longer enforceable. For some special versions of these environments, Allen (1985) and Cole and Kocherlakota (1997) show that the incomplete security market economy we described previously fully decentralizes the Pareto efficient allocations subject to incentive constraints. Even when there is an efficiency wedge, the specification of microeconomic uncertainty is a critical ingredient in both the decentralized economy and in the Pareto efficient economy subject to incentive constraints. The problems of measuring microeconomic uncertainty arise in private information economies as well.