On Efficiency and Distribution

Robert E. Lucas, Jr.

The Economic Journal, Vol. 102, No. 411 (Mar., 1992), 233-247.

Stable URL:
http://links.jstor.org/sici?sici=0013-0133%28199203%29102%3A411%3C233%3AOEAD%3E2.0.CO%3B2-B

The Economic Journal is currently published by Royal Economic Society.

THE ECONOMIC JOURNAL

MARCH 1992

The Economic Journal, **102** (*March* 1992), 233–247
Printed in Great Britain

ON EFFICIENCY AND DISTRIBUTION*

Robert E. Lucas, Jr.

It helps us to see the actual world
to visualize a fantastic world.

Wallace Stevens

This article is concerned with the possibilities for allocating resources in an economy with the following characteristics. There is a constant, perpetual endowment flow of a single, non-storable good. There are many households, all infinitely lived, and all with the same *ex ante* preferences over time paths of the consumption of this single good. Each of these households is subject to shocks to its preferences, unpredictable even to itself, that give it a high urgency to consume in some periods and a low urgency to consume in other periods. These shocks are independent from household to household, so with a large number of households they will average out in any given period, with urgent consumers just balanced out by the less urgent.

Allocating resources in this setting means just distributing the given endowment, at each date, across these different households. There is no production and no contemporaneous goods-for-goods exchange. But however this allocation is done it must be consistent with the fact – the crucial assumption of the analysis – that these individual shocks are purely private information. The only way anyone can obtain any information about any particular consumer's urgency to consume is to ask him, and there is no way to audit or verify the answer this consumer chooses to give.

Of course, these data about tastes, technology, and information are not in themselves sufficient to determine how resources are allocated. This is why I referred to the 'possibilities' for allocation a moment ago. One could take a normative point of view and consider how a hypothetical, beneficent planner would distribute the consumption good across households, under various assumptions about the information available to him. Or, one could allocate property rights to the endowment stream, set up a system of markets, let households trade and see what allocation they come up with. This approach,

* Given as the 1991 Harry Johnson Lecture. I thank James Mirrlees, A. R. Nobay, and the Royal Economic Society both for their hospitality and for this opportunity to honour the memory of a friend and an admired teacher and colleague. I am grateful to Andrew Atkeson, Nancy Stokey, and Bart Taub for useful discussions, and for research support under National Science Foundation grant number SES 8808835.

too, will give different answers depending on the assumptions one makes about trading possibilities. I will pursue all of these directions.

In each case, given an initial distribution of wealth (of entitlements to current and future consumption) each specific method for allocating resources will imply a complete description of the way this society's wealth distribution evolves over time. What is striking to me in comparing the distributional dynamics implied by different allocative mechanisms is how radically they differ. Starting from a position of *ex ante* equality, we will see everything from perfect equality in perpetuity to convergence to a stationary distribution to inequality that grows without bound.

All of these possibilities will be seen to arise in a society of essentially identical households, free of the issues of class and race that so complicate questions of distribution in actual societies. Perhaps it is a mistake to try to think about distributional questions at all in a context that abstracts from these distinctions. Certainly this choice will dictate more than a little caution in drawing conclusions from our theoretical analysis about distribution in actual societies. But the idea that a society's income distribution arises, in large part, from the way it deals with individual risks is a very old and fundamental one, one that is at least implicit in all modern studies of distribution. If we cannot think clearly about this issue in the abstract context I have described, what hope is there for dealing with more realistic situations?

I will illustrate these distributional possibilities with concrete examples drawn from research in which Andrew Atkeson and I are currently engaged (Atkeson and Lucas (1991)). Our work has many antecedents, but I find it substantively most instructive to view it as contributing to a line of inquiry initiated by Truman Bewley (1983) and Edward Green (1987). Bewley's paper was the first attempt to imagine in detail what an entire society would have to look like if the behaviour of individual households in it were to be consistent with Friedman's (1957) permanent income hypothesis. Of course, it is this step that converts Friedman's theory of individual behaviour into a general equilibrium theory of distribution. Bewley took a particular market structure, necessarily incomplete in the Arrow–Debreu sense, as a given. Green took matters a step further to take the information structure of the economy as given, and then derived the efficient allocation implied by this information structure under specific parametric assumptions about consumer preferences.

Since then many others, notably Taub (1990), Phelan and Townsend (1991), Marimon and Marcet (1990), and Thomas and Worrall (1990) have used theoretical or numerical methods or both to work out the implications of other specific assumptions on preferences and information structures. As empirical work by Townsend (1989), Mace (1991), Cochrane (1991), and others amply demonstrates, the Bewley–Green viewpoint leads to new and extremely interesting ways of interpreting data on household income and consumption expenditures. I think it has equally radical implications for the way we think about distributional dynamics.

I. THE MODEL ECONOMY

I will next provide a more detailed description of the exchange economy we will examine, and then outline the rest of the lecture. The economy has a constant, perpetual endowment of y (per capita) units of a single, non-storable consumption good at each date $t = 0, 1, \ldots$. There is a continuum of households, each with the preferences:

$$E\left[\sum_{t=0}^{\infty} (1-\beta)\,\beta^t U(c_t)\,\theta_t\right],$$

where c_t is consumption of the one good and $\{\theta_t\}$ is a sequence of iid taste shocks, independent across agents. Let the distribution of the shock be μ, normalised so that the mean shock value is unity. Throughout, I will assume that the current period utility function $U(c)$ takes the constant relative risk aversion (CRRA) form:

$$U(c) = (1/\gamma)\,(c^{\gamma}-1), \quad \gamma \leqslant 1.$$

When $\gamma = 0$, $U(c) = \ln(c)$.

The only decision this society has to make each period is how to distribute the fixed stock of goods y over these many consumers. Even if households are identical *ex ante*, they will receive different idiosyncratic taste shocks as time passes. Depending on the resource allocating mechanism used, these shocks will affect their current consumption and also their entitlements to future consumption. Our objective will be to see how the distribution of entitlements evolves under different assumptions about the way resources are allocated.

I consider five distinct resource allocation mechanisms. The first two, examined in Section II, will serve as benchmarks: autarky, and the full-information efficient allocations. There are many of both, for the same reason that there are many possible endowment points and many points on the contract curve in an Edgeworth Box. In Section III I consider a market system in which the only marketable asset is money, and in which all consumption must be financed with cash-in-advance. This is the pure currency economy (in Wicksell's terminology) studied in my (1978) paper. In this economy, each household owns the endowment stream y but this claim is not a marketable security. Then in Section IV I go to the opposite extreme, to study what Wicksell called a pure credit economy, in which claims to future endowment are perfectly marketable – tradeable without restriction – and in which money disappears from use.

Finally, in Section V, I will characterise the efficient allocations that are incentive compatible – that respect the fact that the idiosyncratic θ shocks are private information. This analysis, taken in its essentials from Atkeson and Lucas (1991), will give us a standard against which we can evaluate the two market allocations I have just described that is more germaine than the unattainable full-information ideal. Even so, we will see that both market allocations fall short in terms of welfare. Are there other market arrangements that do better than these two? Maybe so, but I will explain why Atkeson and I concluded that full efficiency, even respecting the privacy of information, is unattainable through any set of purely private arrangements.

II. THE AUTARKY AND FULL-INFORMATION ALLOCATIONS

By an autarky allocation, I mean one in which each household is allocated a share a (say) of the endowment and consumes at the constant rate ay forever, Since individual consumers' names are immaterial, I will just refer to 'household a', identifying each with its relative consumption level. With the normalisations I have adopted, the expected discounted utility enjoyed by household a, as of date 0 but prior to the realisation of the date 0 shocks, is just $U(ay)$. Feasibility requires that the shares average to unity, so the possible *utility distributions* under autarky are just the distributions ψ of $U(ay)$ that can be obtained with share distributions ϕ that satisfy $\int a d\phi = 1$. Obviously, the utility distributions attainable through autarky are not ranked in Pareto's sense: any share reallocation that improves someone's welfare must reduce someone else's.

At the opposite extreme, consider next the allocation problem faced by a hypothetical, beneficent social planner who has *complete* information about everyone's individual shocks. As with the autarky allocations, there are many possibilities, depending on how different households are to be treated. In order to keep track of these, I will continue to think of endowment shares distributed across households and identify each household by its share holdings a. This term 'share' suggests something traded in a securities market, but that is ahead of my story. For now, think of these shares as a record-keeping device used by a social planner to keep track of who is entitled to what.

Households in this economy will want to pool the idiosyncratic risk they all face, or have this done on their behalf by the planner. Think of the planner as doing this in the following way. First, all households are divided into groups according to their share values, all members of a given group having equal share holdings a. These groups are to remain forever separate from one another. Within any group a, consumers will receive a consumption allocation $c(a, \theta)$, depending on their taste shock, and they will receive an allocation $g(a, \theta)$ of end-of-period shares. These allocations must satisfy:

$$\int c(a, \theta) \, d\mu \leqslant ay \tag{1}$$

and

$$\int g(a, \theta) \, d\mu \leqslant a, \quad \text{for all } a, \tag{2}$$

since both goods and shares allocated must be met with the resources of each group a in isolation.

Next period, the subset of group a that received $g(a, \theta)$ apiece for any fixed θ is again treated in isolation, with consumption averaging $g(a, \theta)\, y$ and new shares averaging $g(a, \theta)$, and so on, *ad infinitum*. Since there is a continuum of households to begin with, we can imagine that all of these ever-proliferating subgroups have a continuum of agents, so there are always plenty of other households to pool risks with. Even so, with general utility, the restriction that the planner must treat each subgroup in isolation from the rest would be inconsistent with efficiency: People at different wealth levels may have different attitudes toward risk, and hence gain from exchange. The great convenience of the assumption of CRRA preferences is exactly that under this assumption there *are* no gains from any kind of risk pooling across households with different share holdings. However θ-risk is pooled, it can be done separately for each fixed level of a. The idea that with CRRA preferences distribution does not matter is surely not a novel one, but this claim does require proof, and one is provided in Atkeson and Lucas (1991). At present, however, I just want to work out some of its implications.

Using this simplifying idea, we can state a Bellman equation for the function $v(a)$: the expected utility from now on enjoyed by a household that holds a shares. It is:

$$v(a) = \max_{c,g} \int \{(1-\beta)\, U[c(a,\theta)]\,\theta + \beta v[g(a,\theta)]\}\, d\mu, \tag{3}$$

subject to the constraints (1) and (2). (The choice in (3) is over functions of θ only; I have used a notation emphasising the dependence of this choice on the value a.)

The first-order conditions for this variational problem are:

$$(1-\beta)\, U'[c(a,\theta)]\,\theta = \lambda(a), \quad \text{all } \theta, \tag{4}$$

and

$$\beta v'[g(a,\theta)] = \xi(a), \quad \text{all } \theta, \tag{5}$$

where λ and ξ are the multipliers associated with (1) and (2). Since θ does not enter into (5), it is evident from (2) that $g(a,\theta)$ equals a: Everyone ends with the shares he began with. This invariance of wealth to idiosyncratic shocks is, of course, what we mean by 'full insurance'. The condition for consumption – which just equates marginal utility across all consumers at a given share level – together with the resource constraint (1) implies:

$$c(a,\theta) = [\mathrm{E}(\theta^{\frac{1}{1-\gamma}})]^{-1}\,\theta^{\frac{1}{1-\gamma}} ay. \tag{6}$$

Notice that as the variance of θ approaches zero or as γ approaches the risk-neutrality value of unity, consumption approaches its autarchy level of ay. Otherwise, consumption is an increasing function of the urgency to consume.

Neither of these two allocations is of much interest in its own right. Autarky is easy to improve upon, while full insurance is impossible to attain under realistic assumptions about what one person can know about another. They will serve as benchmarks for more complicated, intermediate cases.[1] It is interesting to note, however, that these two very different allocations have a common implication for distributional dynamics: any initial distribution of relative utilities will be maintained over time. Thus if we begin with a utility distribution concentrated at a single point – either because we have a preference for equality or because we imagine risk averse agents selecting initial entitlements in some optimally-designed, pre-existing lottery – this distribution will remain equal for all time.

[1] Thomas and Worrall (1990) use these same two benchmark allocations in their construction of a solution to a dynamic incentive problem similar to that studied in Atkeson and Lucas (1991) and in Section V of this paper.

III. A PURE CURRENCY EQUILIBRIUM

Neither of the two allocations considered so far involves anything like realistic exchange. In the first, autarky, there is no trade at all. The second, full-information, allocations could be attained by trade in a full set of Arrow–Debreu markets, but only if everyone's private shock were a matter of public information. In this section, I want to turn to a specific market set-up that is consistent with the privacy of information. It was designed for thinking about the demand for money, but will serve to advance the present discussion as well.

Let every household in the economy have the same endowment stream y, but treat claims to future endowments as untradeable. To ensure that money will be held, assume that no household can consume its own endowment. Instead, the endowment must be sold for fiat currency, which can later be spent to purchase goods from other households. No exchange except contemporaneous money-for-goods trades is permitted. There are many ways – plain and fancy – to motivate these conventions, but my present interest is in their consequences, not in their rationale.[2]

In this situation, the state of the economy is determined by the distribution ϕ (say) of households by their beginning-of-period money holdings, and each household's situation is described by its individual cash holdings m and its current taste shock θ. So I will speak of household (m, θ) in economy ϕ. The decision problem facing each household can be described by a Bellman equation in the value $v(\phi, m, \theta)$ of the maximised objective function of household (m, θ) in economy ϕ. This equation is:

$$v(\phi, m, \theta) = \max_{c,\, m'} \left[(1-\beta)\, U(c)\, \theta + \beta \int v(\phi', m', \theta')\, d\mu \right] \tag{7}$$

subject to:

$$p(\phi)\, c \leqslant m, \tag{8}$$

$$m' = m + p(\phi)\,(y-c) \quad \text{and} \quad c, m' \geqslant 0, \tag{9}$$

where ϕ' denotes next period's cash distribution. A rational expectations equilibrium in this economy consists of a value function v, policy functions for consumption and end-of-period cash holdings, a market-clearing price function $p(\phi)$, and a law of motion for the cash distribution that is consistent with individual money demand behaviour. Given such an equilibrium, one could calculate the evolution of the cash distribution, and hence of the distribution of goods, from any given initial distribution.

This is a harder problem than I am able to solve. The assumption of CRRA preferences is designed to keep distributional effects at a minimum but the cash-in-advance constraint (8) precludes a simple solution even under this assumption. Let us retreat, and seek instead a *stationary* equilibrium: an equilibrium in which the cash distribution remains unchanged from period to period and hence in which the price level remains constant. In such an

[2] See, for example, the informal motivation in Lucas (1978) or the more explicit construction of Townsend (1987).

equilibrium, an individual can be identified with his *real* balances $z = m/p$, and his Bellman equation becomes:

$$v(z,\theta) = \max_{c,z'}\left[(1-\beta)\, U(c)\,\theta+\beta\int v(z',\theta')\, d\mu\right] \tag{10}$$

subject to:

$$c \leqslant z, \quad z' = z+y-c \quad \text{and} \quad c, z' \geqslant 0. \tag{11}$$

Under standard regularity assumptions, (10) has a unique solution v, corresponding to which there is a unique end-of-period real balances demand function $z' = g(z,\theta)$.[3] This function g is equal to y for sufficiently small z and high θ values: those values for which the cash constraint $c \leqslant z$ binds. For larger z or smaller θ or both, the cash constraint is slack and g is increasing in z and decreasing in θ. The behaviour of the typical household's cash holdings is governed by the difference equation:

$$z_{t+1} = g(z_t, \theta_t). \tag{12}$$

Hence real balances follow a first-order Markov process. One can show, under reasonable conditions, that this process has a unique invariant distribution ϕ, say. Given ϕ, one can calculate the average real money demand for the economy as a whole. It is:

$$\int z\, d\phi = \iint g(z,\theta)\, d\mu\, d\phi. \tag{13}$$

If there is a *per capita*, nominal money supply of M, then equating the magnitude (13) to M/p gives the equilibrium price level p. Of course, Walras Law implies that this same price equates the demand and supply of goods as well.

The difference equation (12) describes the evolution of any individual's cash holdings, given that society as a whole is described by the invariant distribution ϕ. It does not describe the process by which society converges to the invariant distribution from some given initial position, since if such convergence occurs the price level will be changing along the transition path whereas the function g describes optimal behaviour only if the price level is constant. Nonetheless, it seems to me a reasonable conjecture that the invariant distribution ϕ describes the limiting behaviour of the system starting from a wide variety of initial distributions. If so, then one could say that the ultimate degree of inequality in the system is independent of the initial inequality, determined jointly by the nature of consumer preferences and the distribution of the idiosyncratic shocks.

Is the resource allocation produced by this cash-in-advance monetary system an efficient one? Surely it is an improvement on autarky, for the ability to accumulate cash gives households some ability to insure against high taste shocks. Yet in any period there are some individuals who are both short on cash and eager enough to consume to be willing to borrow at interest against their future endowment income to do so. There are others with large cash holdings and less urgent consumption needs who would surely be willing to lend. The

[3] See Lucas (1978) or Stokey *et al.* (1989), Section 13.5.

presence of such gains from exchange suggests inefficiency, a conjecture that will be confirmed in Section V, when the efficient allocation under private information is constructed. Let us first see, however, if simply opening up a credit market will resolve this difficulty.

IV. A PURE CREDIT EQUILIBRIUM

At the other extreme from the monetary equilibrium we have just examined, consider a situation in which claims to future endowment are perfectly marketable securities and in which these securities are directly exchanged for goods, with no need for cash. I call this a pure credit economy. Let Q be the value, in terms of current consumption, of a claim to y units of goods in perpetuity, starting next period. Let ϕ denote the current distribution of households by claims held, with $\int a\, d\phi = 1$. In general, the current price Q would depend on this share distribution ϕ, but we will see that under the assumption of CRRA preferences, only constant equity prices need to be considered. In this case, we can consider the Bellman equation for the maximised utility $v(a)$ of a household beginning with shares a and the current shock value θ:

$$v(a,\theta) = \max_{c,\,a'}\left[(1-\beta)\, U(c)\,\theta + \beta\int v(a',\theta')\, d\mu\right] \tag{14}$$

subject to:

$$c + Qa' = a(Q+y), \quad c, a' \geqslant 0. \tag{15}$$

Equation (14) can be solved explicitly for any CRRA utility function, but we can economise on formulas by illustrating the situation for the case of log utility ($\gamma = 0$). In this case, the value function takes the form: $v(a,\theta) = A(\theta) + [(1-\beta)\,\theta+\beta]\ln(a)$, where the nature of the function $A(\theta)$ will not affect margins and need not concern us. The demands for goods and shares are:

$$c = \frac{(1-\beta)\,\theta}{(1-\beta)\,\theta+\beta}(Q+y)\,a, \tag{16}$$

$$a' = \frac{\beta}{(1-\beta)\,\theta+\beta}\left(\frac{Q+y}{Q}\right)a. \tag{17}$$

These equations describe optimal individual behaviour, given the price Q.

To determine the equilibrium value of Q, we use the fact that average share holdings both before and after trading must be unity. Then (17) implies:

$$1 = \beta \mathrm{E}[h(\theta)]\,\frac{Q+y}{Q}, \tag{18}$$

where I use $h(\theta) = [(1-\beta)\,\theta+\beta]^{-1}$. One can solve (17) for the equilibrium share price Q. Note that the function $h(\theta)$ is convex, so (17) implies:

$$Q \geqslant \frac{\beta}{1-\beta}y, \tag{19}$$

with strict inequality unless θ has zero variance. That is to say, the equilibrium interest rate is lower than the rate of time preference. This has nothing to do with the riskiness of equities – which have a riskless return stream – but represents a premium on assets because of their use in self-insuring against idiosyncratic risks that (by my assumptions) cannot be insured against directly.

Look at the distributional dynamics implied by this equilibrium! Combining (17) and (18), one obtains:

$$a' = \frac{h(\theta)}{\mathrm{E}[h(\theta)]}a,$$

so that taking logs of both sides, a household's share holdings follow the difference equation:

$$\ln(a_{t+1}) = \ln[h(\theta_t)] - \ln\{\mathrm{E}[h(\theta)]\} + \ln(a_t). \tag{20}$$

This is a random walk with drift, and since ln (x) is a concave function, the drift term has a negative expected value. From this, we conclude that the variance of the log of share holdings (and from (16) of the log of consumption, too) is growing linearly, without bound. We conclude as well (paralleling Thomas and Worrall (1990) exactly) that every household's consumption is going to zero with probability one. Yet mean consumption is constant at y. The situation is one of ever growing inequality, with wealth concentrated in an ever shrinking number of ever wealthier households.

These dynamics, which follow even from an initial situation of equality, surely do not conform to our customary images of a well functioning society. But are they economically inefficient? The answer, as we will see in the next section, is yes and no.

V. EFFICIENCY WITH PRIVATE INFORMATION

The two very different market allocations worked out in the last two sections both have the feature that an individual with a high current shock can obtain higher current consumption by surrendering claims to future consumptions. In the monetary equilibrium this could be done, up to a point, by running down cash holdings. In the credit economy the same end is achieved by selling off future endowment holdings. These two examples obviously do not exhaust the range of market opportunities one might postulate and explore, but in this section I will return instead to the hypothetical planning problem introduced in Section II. In this case, however, the planner is assumed not to have access to any source of information about individual shocks. This is the case studied in Atkeson and Lucas (1991).

A social planner who is unable to monitor individual shocks must either revert to an autarchy allocation or else base his consumption assignments $c(a, \theta)$ on consumers' *claims* about the shocks they have received. It is obvious that the full information allocation cannot be achieved under these circumstances, for in that allocation everyone is allocated current consumption that increases with θ and future consumption that does not depend on θ at all. If the planner has only one's own testimony to go on, why not claim the highest urgency all the time? If consumers are to be honest about their reported shocks, they must be made to pay for a claim of high current urgency with lower consumption later on.

To work out the details of this tradeoff, we return to the basic Bellman equation (3). To capture the fact that allocations must depend on unmonitored reports of shocks rather than the shocks themselves, we add to the constraints (1) and (2) an *incentive-compatibility* constraint:

$$(1-\beta)\, U[c(a,\theta)]\,\theta+\beta v[g(a,\theta)] \geqslant (1-\beta)\, U[c(a,z)]\theta+\beta v[g(a,z)] \quad (21)$$

for all a, θ, and z. The constraint (21) describes the behaviour of a consumer receiving the actual shock θ: This is the value that multiplies his utility on both sides of the inequality, and what he reports cannot change this. If he tells the truth, he receives the pair $c(a,\theta), g(a,\theta)$ from the planner. If he reports any other value z, he receives $c(a,z), g(a,z)$. The inequality says that the planner must select the functions c and g in such a way that people are always better off reporting truthfully.

Could the planner not as well base allocations on false reports? For example, what if everyone reported twice his true θ value and the planner knew this and simply divided all reports by two? There are any number of such possibilities, but it is an implication of the revelation principle of Myerson (1979) and Harris and Townsend (1978; 1981) that taking these additional reporting possibilities into account adds *no* new possibilities for the ultimate resource allocations that are based on these reports. If our concern is with the way goods are allocated, as opposed to the nature of the process that brings these allocations about, we lose nothing by restricting attention to allocation schemes based on truth-telling.

The truth-telling restriction (21) precludes one-period lies, but leaves open the possibility that an individual might gain by a more complicated pattern of false reports over time. My own economic instincts are so thoroughly Bellmanised that I can barely imagine this non-recursive possibility, but it is present and takes some work and some assumptions to rule it out.[4] For present purposes, I ask you to accept on faith that in this particular model, (21) expresses a preference for truth over any possible pattern of lies.

The Bellman equation (3) is more difficult with the incentive constraint (21) imposed than in the full-information case, but we have built up some intuition about it from the analyses of the full-information case and of the pure credit economy. Based on these cases, one would conjecture that for general CRRA preferences the solution $v(a)$ takes the form of a constant times a^{γ}, and that in the log case it takes the form $v(a) = A+B\ln(a)$. In either case it is reasonable to conjecture that both current consumption and end-of-period shares will be proportional to initial shares, or that efficient allocations will take the form $c(a,\theta) = r(\theta)\,a$ and $g(a,\theta) = f(\theta)\,a$ for some functions r and f of θ.[5] I will focus here on the log case only.

[4] See Green (1987), Lemma 2, or Atkeson and Lucas (1991), Lemma 3.2.

[5] In Atkeson and Lucas (1991) a Bellman equation is derived that holds for a much wider class of utility functions, and these features of the CRRA case are proved.

Inserting these guesses into (3), with log utility, one sees that the value function is $v(a) = A + \ln(a)$. The constant A must satisfy:

$$(1-\beta) A = \max_{r,f} \int \{(1-\beta) \ln [r(\theta)]\, \theta + \beta \ln [f(\theta)]\}\, d\mu, \tag{22}$$

where the functions r and f are chosen subject to the three constraints:

$$\int r(\theta)\, d\mu \leqslant y, \tag{23}$$

$$\int f(\theta)\, d\mu \leqslant 1, \tag{24}$$

and $$(1-\beta) \ln [r(\theta)]\, \theta + \beta \ln [f(\theta)] \geqslant (1-\beta) \ln [r(z)]\, \theta + \beta \ln [f(z)] \tag{25}$$

for all θ and z. That is to say, with log utility, if we solve the allocation problem for the share level $a = 1$ we can scale it up or down to obtain the solution for any other share level.

The problem (22) is just a standard (which is not to say easy) incentive problem, entirely static in form although we are interpreting one of the two goods as entitlements to future goods. To study this problem, I assume only two states for θ, θ_1 and θ_2. Then we can follow the analysis of Jacklin (1987) and put the whole problem into a two good (r and f) two person-type (θ_1 and θ_2) Edgeworth Box – displayed as Fig. 1.

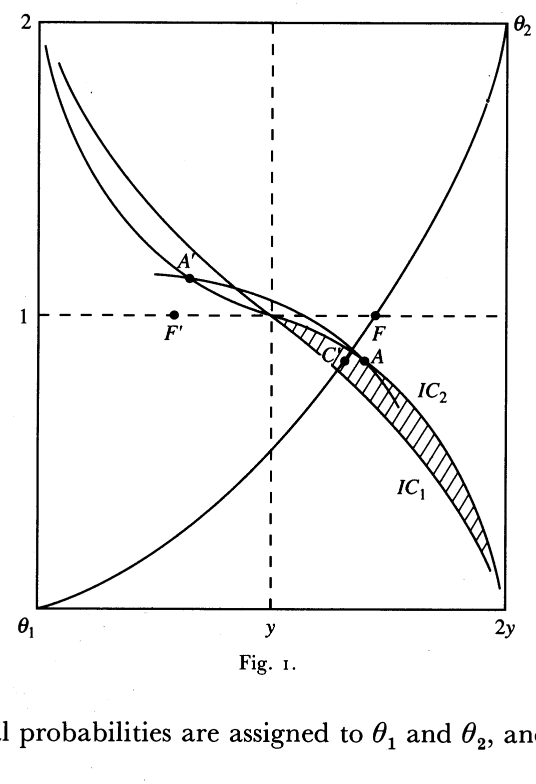

Fig. 1.

In Fig. 1, equal probabilities are assigned to θ_1 and θ_2, and $\theta_1 > \theta_2$, so that the most eager consumer has his origin in the southwest and the least eager consumer – θ_2 – is in the northeast. The boundaries of the box are given by the constraints (23) and (24): use the horizontal axis for goods, with total units $2y$, and the vertical axis for end-of-period shares, totalling two. The two origins are connected by the contract curve, which then must lie below the diagonal.

The centre of the box represents the autarky allocation: both agent-types consume y and end with one share each. The point F denotes the full information allocation: It lies on the contract curve, and gives both consumers one end-of-period share. The point C denotes the pure credit equilibrium: It is the point on the contract curve to which the parties would trade competitively from the centre of the box as the endowment point. All of these points are familiar from Sections II and IV.

The solution to the planning problem considered in the present section is displayed on Fig. 1 as well, but in order to read this figure, turn first to Fig. 2

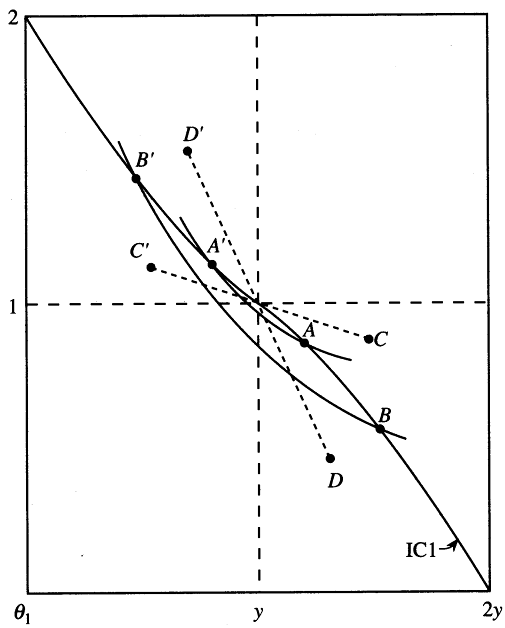

Fig. 2.

which refers to agent 1's problem only. The planner picks a point (r, f) in the box, say point A in Fig. 2, and offers it to all who declare themselves to be of type 1. This offer also makes available the point $(2y - r, 2 - f)$ to all who claim to be of type 2. This point is the reflection of A through the centre of the box, indicated as point A' on Fig. 2. Agent 1 can pick either A or A', by reporting his type either truthfully or falsely. As I have drawn agent 1's indifference curves and the points A and A' on the figure, agent 1 is indifferent between A and A'. Similarly, the point B is on the same indifference curve as its reflection, B'. These four points A, A', B and B' are all points on which agent 1's incentive constraint (25) is just binding. The curve on Fig. 2 labelled $IC1$ contains *all* the points for which (25) holds with equality for agent 1.

The point C on the figure is strictly preferred to its reflection, C'. Similarly, D' is strictly preferred to D. Thus (25) holds with strict inequality for agent 1 at C and D'. At C' and D, (25) is violated. Hence the set of points that satisfy (24) for agent 1 consists of the curve $IC1$ and all the points above this curve.

I have transferred this curve $IC1$ to Fig. 1. A similar curve can be constructed for agent 2, which is labelled $IC2$ on Fig. 1. The crescent-shaped shaded area in the southeast quadrant of Fig. 1 contains all the points that are incentive-compatible from the point of view of both agents. The planner's task is to choose a point in this crescent to maximise the sum of the two agents utilities. Obviously the full information point F does not have this incentive-compatibility property: agent 1 prefers it to F', but agent 2 prefers F'.

We cannot quite see the planner's choice on the diagram, since this is a purely ordinal picture and the objective is the cardinal one of maximising expected utility. But recall that the planner would *like* to pick point F, and would do so if he had full information. So he picks a point like A to get as close to F as he can. At this point, agent 2 is indifferent between truth-telling and lying – between the points A and A' – and agent 1 is strictly better off by truth telling. At this point A, the most eager consumer is awarded higher consumption, but at the cost of reduced claims to future consumption.

One cannot see the pure currency equilibrium point on Fig. 1: even with log utility, that case does not have the proportionality feature that permits the reduction of the dynamic problem to an essentially static one. But it follows from the analysis in Atkeson and Lucas (1991) that the monetary equilibrium is inefficient. The monetary mechanism described in Section III is feasible for a planner, and we show that he chooses to do something else.

It is a straightforward calculation to show that at the efficient allocation A, the marginal rate of substitution between current and future consumption is higher for agent 2 than for agent 1, which is to say that the point A lies below the contract curve.[6] This is enough to show that the pure credit equilibrium point C is inefficient. It is also enough to show that the efficient allocation cannot be maintained if agents are free to engage in unmonitored borrowing and lending: If A is treated as a post-insurance-payment endowment point and if people are free to trade from this point, agent 1 will exchange some of his current goods in return for claims to future goods.

There is some latitude in defining a market equilibrium in this context – what are the commodities assumed to be traded? – but I take the fact that A is off the contract curve to mean that the efficient allocation cannot be implemented through competitive exchange. It is true that a financial intermediary acting exactly as the hypothetical social planner of this section could deliver A to his clients, and with a continuum of agents, there is room for many such intermediaries. But such an intermediary would need to be able to monitor and prevent all exchange on the part of its clients. This capability seems to me well beyond that possessed by any actual private institution.[7]

The distributional dynamics of the efficient allocation are evidently identical in form to those of the pure credit equilibrium. Each consumer's share

[6] Atkeson and Lucas (1991). Section 9.

[7] This is the conclusion reached, by similar reasoning, in Hammond (1989).

entitlements follow (in logs) a random walk with negative drift: equation (20) with the function $h(\theta)$ replaced by $f(\theta)$. The problem with the pure credit equilibrium is not that inequality grows but only that it may not grow at the efficient rate. We are accustomed to thinking of insurance as a consumption equaliser, but that intuition is based on full-insurance models like that of Section II. With the partial insurance necessitated by private information, we can channel consumption to currently needy consumers in excess of market levels only by penalising their future claims in excess of market levels.

In summary, it is convergence to a finite-variance limiting distribution, as in the monetary equilibrium of Section III, that is a symptom of inefficiency, not the ever increasing inequality of the pure credit equilibrium of Section IV. The credit equilibrium *is* inefficient, but only because it provides inadequate insurance against a high consumption urgency. When this deficiency is remedied, necessarily by non-market means, the growing inequality of the credit equilibrium is not merely preserved but may even be accentuated.

VI. CONCLUSIONS

I said at the outset that my formulation would be a little too abstract for immediate application. Now that the exposition is complete, I doubt that any of you seriously disagrees. But not all abstractions are equally useful, and I want to conclude by asking whether the issues I have focused on this afternoon constitute a large part of what we mean by the problem of distribution, or a minor part, or no part at all.

In focusing on uninsurable risk and ways of dealing with it, I believe I have been discussing pretty much the whole problem. (In saying this, I refer of course to risks of all kinds – shocks to endowments as well as those I have called taste shocks.) If the children of Noah had been able and willing to pool risks, Arrow-Debreu style, among themselves and their descendants, then the vast inequality we see today, within and across societies, would not exist, and those whose ancestors had the talent and luck to participate most fully in the industrial revolution would be remitting a good part of their return to those whose did not. The study of distribution is, over a long enough time period, the study of social mobility, and one cannot discuss social mobility without reference to uninsured individual risks.

In order to view the processes analysed in this article as occurring slowly over long periods it is necessary, of course, to think of the typical household in these models as representing a family of successive generations. Viewed in this way, the positive analysis of the credit equilibrium rests on the assumption that each agent has unlimited ability to sell off the endowments of his heirs to meet his own current needs. He does so, I have assumed, altruistically, but subject to no externally imposed limits. In the same way, the normative analysis rests on an efficiency criterion that treats each currently alive person as the sole spokesman for his yet-to-be-born descendants. In ordering differing allocation schemes, the hypothetical social planner does not recognise members of future generations as distinct individuals whose preferences must be taken into account.

Perhaps it makes more sense to view the theory I have reviewed as applying to individuals during their own lifetimes, with the process starting anew for each generation. This interpretation would bring the positive theory I have offered into closer agreement with fact, since we do observe ever-increasing inequality over the lifetime of members of a given cohort. Could the normative analysis I have reviewed be reworked under an efficiency criterion that assigns independent weight to the members of future generations, without denying altogether the fundamental role of intrafamilial altruism? It is time, I think, for welfare economics to deal seriously with the economics of the family, and I will not be disappointed if the conclusions of the analysis I have described in this lecture are cited in support of this belief.

University of Chicago

Date of receipt of final typescript: August 1991

References

Atkeson, Andrew G. and Lucas, Robert E. Jr. (1991). 'On efficient distribution with private information.' University of Chicago working paper.

Bewley, Truman (1983). 'A difficulty with the optimum quantity of money.' *Econometrica* vol. 51, pp. 1485–504.

Cochrane, John (1991). 'A simple test of consumption insurance.' *Journal of Political Economy* vol. 99, pp. 957–76.

Friedman, Milton (1957). *A Theory of the Consumption Function*. Princeton: Princeton University Press.

Green, Edward (1987). 'Lending and the smoothing of uninsurable income.' In Prescott and Wallace (1987). *Contractual Arrangements for Intertemporal Trade*. Minneapolis: University of Minnesota Press.

Hammond, Peter J. (1989). 'On the impossibility of perfect capital markets.' Stanford University working paper.

Harris, Milton and Townsend, Robert M. (1978). 'Allocation mechanisms for asymmetrically informed agents.' Carnegie Mellon University working paper. Reprinted in Townsend (1990).

—— and —— (1981). 'Resource allocation under asymmetric information.' *Econometrica* vol. 49, pp. 33–64. Reprinted in Townsend (1990).

Jacklin, Charles J. (1987). 'Demand deposits, trading restrictions, and risk sharing.' In Prescott and Wallace (1987).

Laitner, John (1990). 'Random earnings differences, lifetime liquidity constraints, and altruistic intergenerational transfers.' University of Michigan working paper.

Lucas, Robert E. Jr. (1978). 'Equilibrium in a pure currency economy.' *Economic Inquiry* vol. 18, pp. 203–20.

Mace, Barbara J. (1991). 'Full insurance in the presence of aggregate uncertainty.' *Journal of Political Economy* vol. 99, pp. 928–56.

Marimon, Ramon and Marcet, Albert (1990). 'Communication, commitment and growth.' Carnegie Mellon University working paper.

Myerson, Roger B. (1979). 'Incentive compatibility and the bargaining problem.' *Econometrica* vol. 47, pp. 61–74.

Phelan, Christopher and Townsend, Robert M. (1991). 'Computing the optimal insurance-incentive tradeoff for multiperiod, information constrained economies.' *Review of Economic Studies* (forthcoming).

Prescott, Edward C. and Wallace, Neil (1987) eds. *Contractual Arrangements for Intertemporal Trade*, Minneapolis: University of Minnesota Press.

Stokey, Nancy L., Lucas, Jr., Robert E. and Prescott, Edward C. (1989). *Recursive Methods in Economic Dynamics*. Cambridge: Harvard University Press.

Taub, Bart (1990). 'The equivalence of lending equilibria and signalling-based insurance under asymmetric information.' *Rand Journal of Economics* vol. 21, pp. 388–408.

Thomas, Jonathan and Worrall, Tim (1990). 'Income fluctuation and asymmetric information: an example of a repeated principal-agent problem.' *Journal of Economic Theory* vol. 51, pp. 367–90.

Townsend, Robert M. (1987). 'Asset return anomalies in a monetary economy.' *Journal of Economic Theory* vol. 41, pp. 219–247.

—— (1989). 'Risk and insurance in village India.' University of Chicago working paper.

—— (1990 *Financial Structure and Economic Organisation*, Oxford: Basil Blackwell.