

Notes: Exchange Economies 1

September 17, 2004

Start of class:

- Did everyone get my email? If not, let me know so I can get your address right.
- Problem set? Hopefully easy.
- Auditors: do what you want, but if you're not actively participating (doing problems, presentations), my guess is you're wasting your time.
- Preferences: our examples had two major simplifications, additive separable preferences and identical homothetic (define — increasing function of $hd1$). The former comes naturally with expected utility. The latter is sometimes useful, sometimes not, but has the following implication: prices don't depend on θ . Since agents have the same preferences across goods, prices don't depend on distribution.
- Presentation and paper: aimed at initiating a research project. Think about what issues interest you, come see me if you want to discuss or bounce around some ideas. I thought we'd start today with some examples.

Issues

We will typically start with some facts: What are we doing the theory for? What are we trying to explain? Today we'll use dynamic exchange economies to interpret a number of issues in international macroeconomics:

- Primary: consumption volatility, correlations across countries
- Secondary: capital flows, portfolios

In each case, we want to be concrete: What data? Where do we get it? How do we compute the statistics we're talking about?

Benchmark environment

Models are laboratories that we can use to generate comparison data. As far as possible, we want to do the same thing to data generated by our model as we do to real data.

Our benchmark for thinking about aggregate consumption data is a dynamic stochastic model with one good and two or more countries. Let's go through our list of what we need to describe the physical environment:

- List of commodities. Uncertainty. Draw event tree. Define events, states/histories. One commodity for each node in the tree.
- List of agents. Several countries, each represented by a single agent. Endowments are $y_i(s^t)$ (one element for each node in the tree — each s^t). Preferences come from discounted expected utility based on “true” probabilities:

$$U_i(c_i) = \sum_t \sum_{s^t \in S^t} \beta^t \pi(s^t) u_i[c_i(s^t)].$$

Ie, the only difference is u_i , each of which is increasing and (strictly) concave.

Economic system: complete markets — one market for every commodity.

Static risk-sharing (practice)

One date, S “states.” [Draw tree.]

Pareto problem:

$$\max_{\{c_i(s)\}} \sum_i \theta_i \sum_s \pi(s) u_i[c_i(s)]$$

subject to

$$\sum_i c_i(s) \leq \sum_i y_i(s) \equiv y(s) \quad \text{for each } s.$$

Set up Lagrangean using multipliers $P(s) = \pi(s)p(s)$ (why?). Problem decomposes into separate ones for each s :

$$\mathcal{L} = \sum_s \pi(s) \sum_i (\theta_i u_i[c_i(s)] + p(s)[y(s) - c_i(s)])$$

The foc’s imply

$$\theta_i u'_i[c_i(s)] = p(s) \quad \text{for all } i, s.$$

The key to what we’re doing: θ does not depend on s . For a particular s , note that allocations $c_i(s)$ are increasing in $y(s)$ and prices are decreasing. [Both follow from u' decreasing....] Furthermore, the allocation rules are the same for all states (we’re solving the same problem). That means that consumption allocations are monotonic functions of each other. [Graph c_1 v. c_2 .]

Example (identical power utility). Let $u_i = c^{1-\alpha}/(1-\alpha)$ with $\alpha > 0$ (identical homothetic preferences) (what is α ? why divide?). The foc is

$$\theta_i [c_i(s)]^{-\alpha} = p(s).$$

Solve for c :

$$c_i(s) = \theta_i^{1/\alpha} / p(s)^{1/\alpha}.$$

Sum over i :

$$\sum_i c_i(s) = y(s) = \sum_i \theta_i^{1/\alpha} / p(s)^{1/\alpha}.$$

Substitute into resource constraint to get $p(s) = (\sum_i \theta_i^{1/\alpha})^\alpha / y(s)^\alpha$, a decreasing function of y . Now plug into the allocation to get:

$$c_i(s) = \theta_i^* y(s),$$

where $\theta_i^* \equiv \theta_i^{1/\alpha} / \sum_j \theta_j^{1/\alpha}$. (Note that these transformed Pareto weights are positive and sum to one.)

Asset pricing. We can also use the prices to value arbitrary state-dependent claims. A claim to an arbitrary “dividend” $d(s)$ is worth

$$\sum_s P(s)d(s) = \sum_s \pi(s)p(s)d(s) = \left(\sum_i \theta_i^{1/\alpha} \right)^\alpha \sum_s \pi(s)y(s)^{-\alpha}d(s).$$

You can see the usual finance thing: value depends on covariance of the payoff (d) and “the market” (p). You can convert this to consumption- β form by using the foc to substitute for p .

Digression: solving equations

Much of modern economics is based on numerical methods: we solve models on the computer and describe their properties quantitatively. We’ll have occasional digressions to review some basic numerical methods. This one is on solving equations, which mathematicians call root-finding (think polynomials).

Solving a scalar equation: find x such that $f(x) = 0$. You can use Excel or Mathematica, but it’s useful to know some standard methods:

- Bisection. Draw a picture of a function on an interval $[a, b]$ such that $f(a)f(b) < 0$. (Why?) The idea is to divide the interval in two (bisection) and use the midpoint to replace either a or b . The only tricky part is deciding which. Here’s an algorithm:
 1. Find a and b . [This often takes some work.]
 2. Compute $\bar{x} = (a + b)/2$.
 3. If $|f(\bar{x})| < \varepsilon$ (a small number), we’re done. Otherwise, continue to the next step.
 4. If $f(a)f(\bar{x}) > 0$, replace a with \bar{x} . Otherwise, replace b with \bar{x} .

This method is much slower than the next one, but it works.

- Newton’s method. Take first-order approximation to f around a specific point and solve it exactly. If the point is x_k , the first-order approximation is

$$f(x) \approx f(x_k) + f'(x_k)(x - x_k) = 0.$$

We solve for x and use it as our next guess:

$$x_{k+1} = x_k - f(x_k)/f'(x_k).$$

The algorithm: starting with some x_0 , repeat this last step until $f(x_k)$ is sufficiently close to zero. Typically fast if you have a good initial value, and it generalizes naturally to more than one dimension. What could go wrong? Sometimes it wanders off; unlike bisection, it's not bounded. And if $f' = 0$, it bombs. [Think of imaginary roots to polynomials.]

- Secant method. An approximation to Newton's method, using

$$f'(x_k) \approx \left(\frac{f(x_k) - f(x_{k-1})}{x_k - x_{k-1}} \right)$$

Works almost as well as Newton's method and saves you the trouble of computing the derivative.

I use the Newton and secant methods regularly.

Example (heterogeneous power utility). Let $u_i(c) = c^{1-\alpha_i}/(1-\alpha_i)$ with $\alpha_1 = 1$ and $\alpha_2 = 2$. Aggregate endowment takes on the values of $2*(0.5, 0.75, 1.0, 1.25, 1.5)$, each with probability 0.2. The average per capita endowment is therefore one. In each state, we need to solve the equation:

$$f(c_1) = \theta_1 c_1^{-\alpha_1} - \theta_2 (y - c_1)^{-\alpha_2} = 0.$$

Let us say that $\theta_1 = \theta_2$ (arbitrary but simple), which leaves us with the single unknown c_1 . I computed the optimal allocation using Newton's method, with the starting value $c_1 = 1$. See the Matlab program `riskshare1.m`. [Note that this is vectorized.] The solutions are denoted by *s in Figures 1 and 2; the dashed lines represent the 50-50 split we'd get if the risk aversion parameters were the same. Note that c_2 varies less than c_1 . Why? The social planner allocates most of the risk to the less risk averse agent. (What happens in the picture if we increase θ_1 ? What if $\alpha_1 = 0$?)

Digression: log-linear approximation

Another approach to finding decision rules is to do some kind of approximation. In our generic problem, let us say we have a vector of decision variables x , a second vector of given state variables y , and a vector of functions (equal in dimension to x) $f(x, y) = 0$ that characterize the solution (foc's, etc).

The simplest, and probably the most common, approach in macroeconomics is based on first-order Taylor series and consists of these steps:

1. Choose a point. A Taylor series is computed at a point — which one? You could choose many points, but it's helpful to use one that's easily computed and likely to be toward the middle of the range of behavior. We commonly choose the solution to a deterministic version of our problem with exogenous random variables set equal to their means. In this case, set y equal to its mean (\bar{y} say) and solve $f(\bar{x}, \bar{y}) = 0$ for \bar{x} .

2. Approximate and solve. The same idea as the iterations in Newton’s method, but we do it only once. Two variations:

- Linear approximation. A first-order Taylor series is

$$f_x(\bar{x}, \bar{y})(x - \bar{x}) + f_y(\bar{x}, \bar{y})(y - \bar{y}) = 0.$$

A linear approximation to the decision rule is therefore

$$x = \bar{x} + [f_x]^{-1} f_y (y - \bar{y}),$$

a linear approximation to the decision rule.

- Log-linear approximation. For inexplicable aesthetic reasons, I prefer an approximation in terms of logs of variables. If we let $X = \log x$, $\bar{X} = \log \bar{x}$, and $\hat{x} = X - \bar{X} = \log(x/\bar{x})$, the function f to be solved can be expressed:

$$f(e^X, e^Y) = 0.$$

A linear approximation in terms of (X, Y) is therefore

$$\bar{x} f_x(\bar{x}, \bar{y}) \hat{x} + \bar{y} f_y(\bar{x}, \bar{y}) \hat{y} = 0.$$

You can solve this for \hat{x} in terms of \hat{y} . [The notation is a little sloppy if x and y are vectors: some of these products are element by element. If better notation crosses your mind, let me know.]

Example (continued). For the risk-sharing problem, the relevant equations are the foc’s and the resource constraint. With power utility, the foc’s are naturally log-linear, even if the curvature parameters are different. They lead to: $-\alpha_i \hat{c}_i = \hat{p}$. (Where does θ go? It’s built into \bar{c} .) Evidently, the less risk averse agent has more volatile log-consumption and absorbs more of the risk.

The difficulty is the resource constraint, which is linear, not log-linear. We use what I call “Taylor’s rule” to generate a log-linear approximation:

$$\sum_i e^{\log c_i} = e^{\log y}.$$

Our f is

$$\log \left(\sum_i e^{C_i} \right) - Y = 0.$$

A linear approximation in these log variables around the deterministic solution is $\sum_i \gamma_i \hat{c}_i = \hat{y}$, where $\gamma_i = \bar{c}_i / \bar{y}$. In our 2-agent example, $\gamma_1 = \gamma_2 = 1/2$.

Solving the foc’s and the resource constraint together gives us the (approximate) allocation rules:

$$\hat{c}_i = \left(\frac{\alpha_i^{-1}}{\sum_j \gamma_j \alpha_j^{-1}} \right) \hat{y}.$$

In our 2-agent example, this reduces to $\hat{c}_1 = [\alpha_1^{-1} / (\alpha_1^{-1}/2 + \alpha_2^{-1}/2)] \hat{y}$. You can get a sense of the approximation, which compares it to the exact solution for the example. Why an approximation? Because it only approximates the resource constraint. You can see the (slight) difference in Figure 2, where the dark line differs a little from the exact solution for values of y far from the mean.

Dynamic risk-sharing

We're going to make two points, one substantive the other conceptual. The substantive point is that the solution to the static risk-sharing problem carries over pretty much unchanged to dynamic environments with additive preferences. The conceptual point is that we can approach dynamic models in two ways: as "date-0" problems in which we solve everything at $t = 0$ and "sequence" problems in which decisions are made period by period. We'll take on faith that the two lead to the same solution; you can get a sense of what's involved in a proof from LS ch 8.

Date-0 approach. The Pareto problem is:

$$\max \{c_i\} \sum_i \theta_i U_i(c_i) = \sum_i \theta_i \sum_{t=0}^{\infty} \sum_{s^t \in S^t} \beta^t \pi(s^t) u_i[c_i(s^t)]$$

subject to the resource constraints

$$\sum_i c_i(s^t) \leq \sum_i y_i(s^t) \equiv y(s^t) \quad \text{for all } s^t.$$

Set up a Lagrangean using multipliers $P(s^t) = \beta^t \pi(s^t) p(s^t)$. As in the static case, this divides into a bunch of similar problems, one for each history:

$$\mathcal{L} = \sum_t \sum_{s^t} \beta^t \pi(s^t) \sum_i \left(\theta_i u_i[c_i(s^t)] + p(s^t) [c_i(s^t) - y(s^t)] \right).$$

The solution follows the same route as the static problem and leads to allocation rules that are increasing in the aggregate endowment. In a competitive equilibrium, we find θ 's that satisfy date-0 budget constraints,

$$\sum_t \sum_{s^t} P(s^t) c_i(s^t) \leq \sum_t \sum_{s^t} P(s^t) y_i(s^t)$$

for all agents i . In short, this approach does nothing to exploit the time element in the commodity space.

Sequence approach. We can take a recursive view of both the equilibrium and optimum problems. An equilibrium is again defined by: (a) agents maximize subject to budget constraints and (b) markets clear. In the sequence definition, we use sequence budget constraints and expand the set of markets to include assets. Specifically, consider Arrow securities, purchased at date t in history s^t that pay one unit of the good at a specific state (s^t, s_{t+1}) next period. If the price of such a claim is $q(s^t, s_{t+1})$ and the quantity purchased is $a(s^t, s_{t+1})$, the agent must obey the sequence of budget constraints,

$$c(s^t) + a(s^t) \leq y_i(s^t) + \sum_{s_{t+1}} q(s^t, s_{t+1}) a_i(s^t, s_{t+1}),$$

one for each history s^t . These constraints have initial conditions: we endow agents with initial security positions that sum to zero, $\sum_i a_i(s_0) = 0$. There are also terminal conditions. If the horizon is $T < \infty$, the terminal conditions are $a_i(s^{T+1}) \geq 0$ (time can't end with

debts). Analogous conditions apply with an infinite horizon. The additional market-clearing conditions are that markets for Arrow securities must clear at all nodes: $\sum_i a_i(s^t) = 0$. This is a little sloppy (we should explain exactly what the nodes are), but gives you the idea.

The key result here is due to Arrow: this set of securities is sufficient to replicate the date-0 equilibrium. The prices q are relative prices of state-contingent claims from one date to the next. To compare them to the date-0 prices, we multiply them together. Going the other way around, the Arrow prices are connected to the date-0 prices by

$$q(s_t, s_{t+1}) = P(s^t, s_{t+1})/P(s^t) = \beta\pi(s_{t+1}|s^t)p(s^t, s_{t+1})/p(s^t).$$

In short, we can easily go from one to the other.

The optimum problem is typically solved for stationary Markov environments: let s be the state each period and the conditional probability of a state next period depends only on the current state: $\pi(s'|s)$. This switch from unconditional to conditional probabilities reflects the sequence nature of the problem. The optimum problem can then be summarized by the Bellman equation:

$$J(s) = \max_{c_i} \sum_i \theta_i u_i[c_i(s)] + \beta \sum_{s'} \pi(s'|s) J(s')$$

subject to the resource constraint

$$\sum_i c_i(s) \leq y(s).$$

The idea: we're now in state s , and must choose an allocation. If we use the multiplier $p(s)$ on the resource constraint, we get the same answer we've gotten all along. Moreover, the answer has the same form for all states s . The value function doesn't matter — there's nothing dynamic about this problem. Moreover, the risk-sharing feature of this model carries over to many models with nontrivial dynamics.

Comparing theory to fact

This model has a number of quantitative properties that we could compare with similar properties of data:

Consumption correlations. We might say the model predicts perfect correlation of consumption across countries. That's not quite right: if preferences are not identical homothetic, the relation can be nonlinear, and correlation is a linear measure of "comovement." But the correlation is certainly positive (nonlinearity can't change the sign) and probably high. Another fine point: a correlation is defined for time series data only with (covariance) stationarity, so we need to decide how to filter the data to ensure this. In international macroeconomic data, of course, the correlations are not only less than one, they're typically less than correlations of output.

Consumption volatility. Roughly speaking, log-consumption has about as much volatility as aggregate output, with some adjustment for differences in curvature parameters (more

risk averse agents have less volatility). But how does consumption vol compare to income/endowment vol? Methods similar to what we used for the resource constraint imply $\hat{y} = \sum_i \omega_i \hat{y}_i$ with $\omega_i = \bar{y}_i / \bar{y}$. “Typically” $Var(\hat{y}_i) > Var(\hat{y})$ and we can say consumption is less volatile than income. [Counterexample: Two countries, perfectly correlated income process, one with variance 4, the other with variance 1. If they are the same size on average, so that $\omega_1 = \omega_2 = 1/2$, then $Var(\hat{y}) = 9/8 > Var(\hat{y}_2) = 1$. You could magnify this by giving agent 2 lower risk aversion.] An open question: Are there some natural conditions that deliver lower consumption vol? Independence across countries works, but is not realistic. Ditto symmetry. In any case, consumption is generally less volatile than output for rich countries. For poor countries, they’re similar.

Trade balance dynamics. As a matter of arithmetic, any model that gives you less volatile consumption than income will give you procyclical net exports: since $nx_i = y_i - c_i$, $Cov(nx_i, y_i) = Var(y_i) - Cov(c_i, y_i)$. This can’t be negative unless consumption is more volatile than output, and even then it takes some work to make the covariance negative. [Is it possible?] This model, therefore, tends to generate procyclical net exports, which is not what we see in the data. What happens typically is that in bad states (ie, when you’re endowment is low), other countries send you goods to eat, and in good states you send goods to them. [Should we be sending more goods to Argentina?] Risk-sharing, in other words. I’d argue that the issue is the absence of investment — more on this later. An open question: sufficient conditions that preclude countercyclical trade balance.

Portfolios. How would you even address home bias in equity portfolios in a model with Arrow securities, whose payoffs have no connection with any particular agent? We could, of course, change the set of assets, but is it then sufficient to support a Pareto optimum? Consider Lucas’s asset pricing mode: there a representative agent “trades” a single asset, a claim to the endowment stream (a “tree”). In this case the tree suffices — we don’t need the complete set of Arrow securities. But what if we had two agents? Typically the same trick won’t work. It does, however, if we have identical homothetic preferences. In this case, a Pareto optimum allocates the same fraction of the aggregate endowment to an agent in all states. We can reproduce this if all agents own the same fraction of every tree: portfolios that are completely diversified across agents/countries. This prediction is at odds with the evidence, which shows that even people in rich countries fail to hold more than a small amount of foreign assets.

Discussion

Is this a useful model? What does it explain? What does it not explain? What alternative environments do you think would provide better descriptions of the aggregate dynamics of consumption, net exports, and portfolio choice?

Figure 1. Relation between consumption allocations




Figure 2. Consumption allocation rule for agent 1

