

ABSTRACT. Agents increase their expected utility by using state-contingent transfers to share risk; many institutions seem to play an important role in permitting such transfers. If agents are suitably risk-neutral, then in the absence of any frictions the benchmark Arrow-Debreu model predicts that *all* risk will be shared, so that idiosyncratic shocks will have no effect on individuals; we call this full risk sharing. Real-world tests of full risk sharing tend to reject it; accordingly, researchers have devised models incorporating various frictions to try and explain the partial risk sharing evident in the data.

Risk sharing

Any two agents may be said to share risk if they employ state-contingent transfers to increase the expected utility of both by reducing the risk of at least one. A very wide variety of human institutions seem to play an important role in risk sharing, including insurance, credit, financial markets, and sharecropping in developing countries.

To be precise, consider a set of agents indexed by $i = 1, \dots, n$ each with von Neumann-Morgenstern utility function U_i and a finite set of possible states of the world $s = 1, \dots, S$, each of which occurs with probability $p(s)$. For simplicity, suppose that each agent i receives a quantity of a single consumption good $x_i(s)$ in state s , thus receiving expected utility

$$EU_i(x_i) = \sum_{s=1}^S p(s)U_i(x_i(s)),$$

where x_i denotes the random variable, $\{x_i(s)\}$ denotes its realizations, and E is the expectation operator. We assume that U_i is strictly increasing, weakly concave, and continuously differentiable for all $i = 1, \dots, n$, so that all agents are at least weakly risk averse. Define the risk faced by agent i to be a quantity

$$R_i(x_i) = U_i(Ex_i) - EU(x_i).$$

This cardinal measure orders probability distributions in the same manner as Rothschild and Stiglitz (1970). We say that i faces *idiosyncratic risk* if $R_i(x_i) > 0$ and $\text{corr}(U'_i(x_i), U'_j(x_j)) < 1$ for some j , where U'_j denotes j 's marginal utility. If any agent i bears idiosyncratic risk, then there exists a set of state-contingent transfers of the consumption

good between i and j , $\{\tau_i^j(s)\}$ which will strictly increase the expected utility of each, while strictly decreasing the risk of at least one of i and j . Implementing such transfers is risk sharing.

Full risk sharing

What might be termed *full risk-sharing* (Allen and Gale, 1988; Rosenzweig, 1988) is a situation in which all idiosyncratic risk is eliminated. While agents may still face risk, this risk is shared, so that marginal utilities of consumption are perfectly correlated across all agents. Full risk-sharing is a hallmark of any Pareto efficient allocation in an Arrow-Debreu economy, provided that agents have von Neumann-Morgenstern preferences, no one is risk-seeking, and at least one agent is strictly risk averse.

Let us establish the necessity of full risk-sharing for any interior Pareto efficient allocation in a simple multi-period endowment economy. The environment is similar to that described above, but agents consume in several periods indexed by $t = 1, \dots, T$, with agent i discounting future expected utility using a discount factor β_i . Different states of the world are realized in each period, with the probability of state $s_t \in \{1, \dots, S\}$ being realized in period t allowed to depend on the period, and so given by $p_t(s_t)$. Then consider the problem facing a social planner, who assigns state-contingent consumption allocations to solve

$$\max_{\{c_{it}(s)\}} \sum_{i=1}^n \lambda_i \sum_{t=1}^T \beta_i^{t-1} \sum_{s_t=1}^S p_t(s_t) U_i(c_{it}(s_t))$$

subject to the resource constraints

$$\sum_{i=1}^n c_{it}(s_t) \leq \sum_{i=1}^n x_{it}(s_t),$$

which must be satisfied at every period t and state s_t ; the planner takes as given the initial state s_0 and a set of positive weights $\{\lambda_i\}$. By varying these weights one can compute the entire set of interior Pareto efficient allocations (Townsend, 1987).

If we let $\mu_t(s_t)$ denote the Lagrange multiplier associated with the resource constraint for period t in state s_t , then the first order conditions for the social planner's problem are

$$(1) \quad \lambda_i \beta_i^{t-1} p_t(s_t) U_i'(c_{it}(s_t)) = \mu_t(s_t).$$

Since this condition must be satisfied in all periods and states for every agent, it follows that

$$U'_i(c_{it}(s_t)) = \frac{\lambda_j}{\lambda_i} \left(\frac{\beta_j}{\beta_i} \right)^{t-1} U'_j(c_{jt}(s_t))$$

for any period t , any pair of agents (i, j) and any state s_t , so that $\text{corr}(U'_i(c_{it}), U'_j(c_{jt})) = 1$, and we have full risk-sharing.

Thus far, we've only considered risk-sharing in the context of an endowment economy. However, the thrust of the claims advanced above hold much more generally. If we were, for example, to add production and some kind of intertemporal technology (e.g., storage), the first order conditions of the planner's problem with respect to state-contingent consumptions (1) would remain unchanged—the effect of these changes would be that the Lagrange multipliers $\{\mu_t(s_t)\}$ would change. This is an illustration of what is sometimes called “separability” between production and consumption, which typically prevails only when there is full risk sharing (cf. Benjamin, 1992).

Risk sharing can also be thought of as a means to smooth consumption across possible states of the world. This suggests a connection to the permanent income hypothesis, which at its core involves agents smoothing consumption across periods. And indeed, it's easy to show that full risk sharing in every period implies the kind of smoothing across periods implied by the consumption Euler equation. However, the consumption Euler equation doesn't imply full risk sharing.

Tests of full risk sharing

The insight that Pareto efficient allocation among risk-averse agents implies full risk sharing has led to tests of versions of (1). The usual strategy involves adopting a convenient parameterization of U_i , and then calculating the logarithm of both sides of (1). For example, if $U_i(c) = \frac{c^{1-\gamma}}{1-\gamma}$, with $\gamma > 0$, then this yields the relationship

$$(2) \quad \gamma \log c_{it}(s_t) = \log \frac{\mu_t(s_t)}{p_t(s_t)} - \log \frac{\lambda_i}{\beta_i} - t \log \beta_i.$$

This is a simple consumption function, which we would expect to be consistent with any efficient allocation. The quantity $\frac{\mu_t(s_t)}{p_t(s_t)}$ is related to the aggregate supply of the consumption good. Note that this is the only determinant of consumption which depends on the random state. This reflects the fact that the only risk borne by agents in an efficient allocation will be aggregate risk. The second term varies with neither the state nor the date, and can be thought of as depending on the levels of consumption that agent i can expect (in a decentralization

of this endowment economy, λ_i could be interpreted as a measure of i 's time zero wealth). The final term has to do with differences in agents' patience.

Now, suppose one has panel data on realized consumption for a sample of agents over some period of time. Letting \tilde{c}_{it} denote observed consumption for agent i in period t , (2) implies the estimating equation

$$(3) \quad \log \tilde{c}_{it} = \eta_t + \alpha_i + \delta_i t + \epsilon_{it},$$

where $\eta_t = \log \frac{\mu_t(s_t)}{\gamma p_t(s_t)}$, $\alpha_i = -\log \frac{\lambda_i}{\gamma \beta_i}$, $\delta_i = -\log \beta_i$, and ϵ_{it} is some disturbance term. Because this final disturbance term isn't implied by the model, it's typically motivated by assuming that it's either related to measurement error in consumption, or to some preference shock.

The reduced form consumption equation (3) can be straightforwardly estimated using ordinary least squares, but this doesn't constitute a test of full risk-sharing. To construct such a test, one typically uses data on other time-varying, idiosyncratic variables which would plausibly influence consumption under some alternative model which predicts less than full risk sharing. Perhaps the most obvious candidate for such a variable is some measure of income, for example the observed endowment realizations \tilde{x}_{it} referred to in the model above. Then one can add (the logarithm of) this variable to (3) as an additional regressor, yielding an estimating equation of the form

$$(4) \quad \log \tilde{c}_{it} = \eta_t + \alpha_i + \delta_i t + \phi \log \tilde{x}_{it} + \epsilon_{it}$$

(Mace, 1991; Cochrane, 1991; Deaton, 1992; Townsend, 1994). Then full risk sharing and an auxiliary assumption that ϵ_{it} is mean independent of the regressors implies the exclusion restriction $\phi = 0$, which can be easily tested.

Partial risk sharing

Restrictions along the lines of (4) have been used to test for full risk sharing in a wide variety of settings, including within-dynasty risk sharing (Hayashi et al., 1996) in the U.S., risk sharing across countries (Obstfeld, 1994), risk sharing within networks in the Philippines (Fafchamps and Lund, 2003), and risk sharing across households in India (Townsend, 1994), Africa, or the U.S. (Mace, 1991). A typical finding is that the estimated response of consumption to income shocks is small but significant, leading one to reject the null hypothesis of full risk sharing. In this case it is tempting to interpret the estimated relationship as determining the response of consumption to income. However, this is generally a mistake. By rejecting the hypothesis of full risk sharing one also rejects the model which generated the

hypothesis, so that theory no longer supports the interpretation of (4) as a consumption function.

Given this kind of evidence against full risk sharing, scholars have been led to devise and test alternative models in which some kind of friction leads to agents bearing some idiosyncratic income risk. Two promising frictions are private information and limited commitment. In the case of private information, realized or announced incomes may provide a useful signal regarding hidden actions or information, and thus an agent's consumption will optimally depend on this signal, leading to a balance between risk sharing and incentives (Holmström, 1979); Ligon (1998) tests this weaker risk-sharing hypothesis in three Indian villages, and is unable to reject it. In the case of limited commitment, an agent who receives an unusually large endowment realization may be tempted to renege on a pre-existing risk-sharing arrangement unless she receives a larger share of resources Kocherlakota (1996); a test of this model in the same three Indian villages by Ligon et al. (2002) finds that this model predicts a response of consumption to income of just the right magnitude. Still, the construction, estimation, and testing of well-specified models which predict only partial risk sharing remains in its infancy.

Ethan Ligon

REFERENCES

- Allen, F. and D. Gale (1988). Optimal security design. *The Review of Financial Studies* 1(3), 229–263.
- Benjamin, D. (1992). Household composition, labor markets, and labor demand: Testing for separation in agricultural household models. *Econometrica* 60(2), 287–322.
- Cochrane, J. H. (1991). A simple test of consumption insurance. *Journal of Political Economy* 99, 957–976.
- Deaton, A. (1992). *Understanding Consumption*. Oxford: Clarendon Press.
- Fafchamps, M. and S. Lund (2003). Risk-sharing networks in rural Philippines. *Journal of Development Economics* 71, 261–87.
- Hayashi, F., J. Altonji, and L. J. Kotlikoff (1996). Risk-sharing between and within families. *Econometrica* 64(2), 261–294.
- Holmström, B. (1979, Spring). Moral hazard and observability. *Bell Journal of Economics* 10, 74–91.
- Kocherlakota, N. R. (1996). Implications of efficient risk sharing without commitment. *The Review of Economic Studies* 63(4), 595–610.

- Ligon, E. (1998). Risk-sharing and information in village economies. *Review of Economic Studies* 65, 847–864.
- Ligon, E., J. P. Thomas, and T. Worrall (2002). Informal insurance arrangements with limited commitment: Theory and evidence from village economies. *The Review of Economic Studies* 69(1), 209–244.
- Mace, B. J. (1991). Full insurance in the presence of aggregate uncertainty. *Journal of Political Economy* 99, 928–956.
- Obstfeld, M. (1994). Risk-taking, global diversification, and growth. *American Economic Review* 84(5), 1310–1329.
- Rosenzweig, M. R. (1988). Risk, implicit contracts and the family in rural areas of low-income countries. *Economic Journal* 98, 1148–1170.
- Rothschild, M. and J. E. Stiglitz (1970). Increasing risk: I. A definition. *Journal of Economic Theory* 2, 225–243.
- Townsend, R. M. (1987). Microfoundations of macroeconomics. In T. F. Bewley (Ed.), *Advances in Economic Theory, Fifth World Congress*. Cambridge: Cambridge University Press.
- Townsend, R. M. (1994). Risk and insurance in village India. *Econometrica* 62(3), 539–591.