Collected Works of Stephen A. Ross: Some Highlights

Philip H. Dybvig
Washington University in Saint Louis

Minnesota

November 9,2007

Professor Stephen A. Ross, MIT Sloan School

- Franco Modigliani Professor Finance and Economics, 1998–present

- Yale University, 1977–1998

- University of Pennsylvania, 1970–1978

- PhD (Harvard Economics) 1969

- BS (CalTech Physics) 1965

- born February 3, 1944

Chester Spatt and I are editing Steve's collected works. Chester and I are both students and co-authors of Steve's. I was also Steve's colleague at Yale for seven years.

Papers

Steve is the greatest living finance scholar. Steve's vitae lists more than 100 papers, most of which are published in top journals. These publications cover many areas, including:

Agency Theory	APT	CAPM
Fixed Income	Growth Theory	Industrial Organization
Insurance	International Economics	Investments
Labor Economics	Option Pricing	Real Estate
Signalling	Survivorship Bias	Taxes

and more.

Steve's works also include some papers that cross areas and are hard to classify, including interesting "big picture" topics such as Financial Marketing, Forensic Finance, and Behavioral Finance.

Some (but not all!) of Steve's Pathbreaking Papers

Ross (AER 1973) - Agency Theory

Ross (JET 1976) - APT

Cox and Ross (JFE 1976) - Risk-neutral probabilities

Ross (QJE 1976) - spanning with options

Ross (BellJ 1977) - Signalling and Capital Structure

Ross (JB 1978) - Arbitrage and Linear Pricing Rule (Fundamental
Theorem of Asset Pricing)

Cox, Ross, Rubinstein (JFE 1979) - Binomial Option Pricing Model

Cox, Ingersoll, Ross (JF 1981) - Expectations Hypothesis

Cox, Ingersoll, Ross (Econometrica 1985a,b) - CIR model

Gibbons, Ross, Shanken (Econometrica 1989) - maximum-likelihood
CAPM test

Agency Theory

Ross, Stephen A. "The Economic Theory of Agency: The Principal's Problem," *American Economic Review* **63**, No. 2, May 1973, 134–139.

This paper introduced agency theory to economics and finance. It is a very short paper but very important. The method for analyzing agency problems is still widely used today.

In the law, an agency relationship is one in which one person (the principal) hires another person (the agent) to perform some task. Because the principal and agent have different preferences about unobserveable actions (for example over effort), there is a fundamental conflict that does not arise in the standard competitive model (in which all conflicts are internalized through the price system).

The Principal's Problem

Choose a fee schedule $\phi(Y)$ and the agent's optimal effort e^* to maximize

$$EU_P(Y(e^*,\varepsilon) - \phi(Y(e^*,\varepsilon)))$$

subject to the agent's participation constraint:

$$(i) \quad EU_A(\phi(Y(e^*,\varepsilon))) \geq U_R$$

and incentive-compatibility of effort:

$$(i) \quad \text{choosing } e = e^* \text{ maximizes } EU_A(\phi(Y(e,\varepsilon))).$$

Main Results

- Trade-off between incentives and risk-sharing

- If the agent is risk-neutral, the first-best is available

The trade-off between incentives and risk sharing is the normal case. For optimal incentives, the agent has a big exposure to the risk of the outcome. For optimal risk-sharing, the principal shares more of the risk.

As an exception, if the agent is risk-neutral, having the agent bear all the risk fills both needs. Specifically, in this case we choose e^* that maximizes $EY(e, \varepsilon)$ and then we choose the fee schedule $\phi(Y) = Y - k$ where $k = EY(e^*, \varepsilon) - U_R$. Because the same contract gives first-best incentives **and** first-best risk-sharing, there is no trade-off in this case.

Really Risk-neutral

Here, risk-neutral means really risk neutral (linear von Neumann-Morgenstern preferences over all wealth levels, not just positive wealth levels). Really risk neutral is equivalent to being indifferent between any gamble and its mean payoff. In recent years, it is common to call an agent with linear preferences risk neutral, even if consumption is only defined over positive wealth or some other restricted domain. I guess this started as a sloppy misapplication of the equivalence; this change in terminology might lead to a misunderstanding of Steve's result.

Arbitrage Pricing Theory (APT) Papers

Ross, S.A. "Mutual Fund Separation in Financial Theory – The Separating Distributions," Journal of Economic Theory 17, No. 2, April 1978, 254–286.

Ross, S.A. "Return, Risk and Arbitrage," I. Friend and J. Bicksler, eds., Risk and Return in Finance. (Cambridge: Ballinger), 1976, 189–217.

Ross, S.A. "The Arbitrage Theory of Capital Asset Pricing," Journal of Economic Theory 13, No. 3, December 1976, 341–360.

APT Assumptions and SML

The Arbitrage Pricing Theory (APT) extends the CAPM to a theoretical model pricing multiple factors. Suppose risky asset returns are given by a factor structure

$$x_i = \mu_i + \sum_{k=1}^{K} \beta_{ik} f_k + \varepsilon_i,$$

where μ_i is the mean return on asset i, f_k is a random mean-zero factor payoff, and β_{ik} is a constant giving the loading of asset i on factor k, and ε_i is a mean-zero error term uncorrelated across assets. Then the APT asserts that there exists a shadow riskfree rate r^f (equal to the actual riskfree rate if there is a riskfree asset) and factor risk premia λ_k such that, for all i,

$$\mu_i = r^f + \sum_{k=1}^{K} \beta_{ik} \pi_k.$$

This is the APT's Security Market Line (SML).

Properties of the CAPM and APT

- Market-level risk is priced (CAPM and APT)

- Idiosyncratic risk is not priced (CAPM and APT)

- Diversification pays (CAPM and APT)

- single source of priced risk, implying 2-fund separation (CAPM or 1-factor APT)

- multiple sources of priced risk, implying K-fund separation (K-factor APT)

The "market model" version of the CAPM is formally subsumed by the 1-factor APT, and in fact Steve has argued persuasively that most tests of the CAPM implicitly use the one-factor APT as the null hypothesis.

Origins of the APT: mutual fund separation

Steve (JET 1978) showed that having a factor structure for returns is a sufficient condition for K-fund separation (i.e., all risk averse agents would be happy to hold portfolios of the K "mutual funds" instead of a general portfolio). Specifically, if the vector of asset payoffs per dollar invested in a one-period model are given by $x = (x_1, ..., x_N)$, then a sufficient condition for K-fund separation is the existence of portfolios θ_1, ..., θ_K, such that, for each n, there exists weights w_{ik} such that

$$ x_n = \sum_{k=1}^{K} w_{ik}\theta_k + \varepsilon_i $$

where $E[\varepsilon_i|\theta_1, ..., \theta_K] = 0$. This is very close to the APT. Steve's paper shows that this characterization is necessary as well for 2-fund separation, and shows the route for proving necessity for K-fund separation.

Absence of Arbitrage and the Linear Pricing Rule

Ross, S.A. "A Simple Approach to the Valuation of Risky Streams," Journal of Business 51, No. 3, July 1978, 453-475.

Dybvig, P.H. and S.A. Ross "Arbitrage," J. Eatwell, M. Milgate and P. Newman, eds., The New Palgrave, A Dictionary of Economics, (London: The MacMillan Press, Ltd.). 1, 1987, 100-106.

Cox, J.C. and S.A. Ross "The Valuation of Options for Alternative Stochastic Processes" Journal of Financial Economics, 3, 1976, 145-166

Steve has said that absence of arbitrage is the unifying theme for all of finance. The results in these papers, are summarized in the *Fundamental Theorem of Asset Pricing* and the *Pricing Rule Representation Theorem*.

Fundamental Theorem of Asset Pricing

Theorem (Fundamental Theorem of Asset Pricing) The following are equivalent: (i) Absence of arbitrage
(ii) Existence of a positive linear pricing rule that prices all claims
(iii) Existence of a hypothetical agent who prefers more to less and has an optimal portfolio

The equivalence of (i) and (ii) was the point of Steve's landmark paper "A Simple Approach to the Valuation of Risky Streams." That was the heavy lifting. Inventing the name and adding the third property (iii) is something I did when for teaching doctoral finance at Princeton. It is obvious that (iii) implies (i), which reminds students why we are interested in absence of arbitrage in the first place. The other direction is more subtle and very useful. For example, if someone asks whether our no-arbitrage result is consistent with equilibrium, the answer is yes by (i) implies (iii).

Pricing Rule Representation Theorem

When arbitrage pricing goes to work, it is useful in different forms in different contexts.

Theorem (Pricing Rule Representation Theorem) The following are equivalent: (i) Existence of a positive linear pricing rule
(ii) Existence of consistent risk-neutral probabilities
(iii) Existence of a positive state-price density
(iv) Existence of positive state prices

The most famous of these is (ii), which is often called the martingale approach. I think many people do not know that this approach originated in Steve's work, especially Cox-Ross [1978]. The state-price density (iii) is also known as the stochastic discount factor or pricing kernel.

Using the various representations

Using general linearity is useful for proving a general distribution-free result such as put-call parity or Modigliani-Miller. The values add at maturity; linearity implies that values add now, too: $S+P = L(S) + L((X-S)^+) = L(S + (X-S)^+) = L((S-X)^+ + X) = L((S-X)^+) + L(X) = C + B$. Using risk-neutral probabilities, (the martingale approach) is most useful for pure valuation problems for which risk pricing is an unnecessary complication. The state-price density, also known as the stochastic discount factor or pricing kernel, is especially good for portfolio problems because then reward for risk-taking does matter, and in fact the state-price density is proportional to the marginal utility of consumption in the standard problem. State prices are most useful for economic analysis, especially when there are state claims in discrete state spaces and agents have different subjective probabilities.

Arbitrage pricing: more examples

Using the martingale approach is often convenient for pure asset pricing problems:

$$p_0 = E^*[\frac{1}{1+r}p_1]$$

Using the state-price density is good for consumption problems:
Choose c to
maximize $Eu(c)$
s.t. $E[\rho c] = w_0$

foc: $u'(c) = \lambda \rho$

Conclusion

Steve Ross has had a great influence on research and practice in finance. We have had a look at only a few of his important papers.

Epilogue: some current working papers by Ross's student

Theoretical Corporate: "Consensus in Diverse Corporate Boards" with Nina Baranchuk (UT-Dallas), forthcoming RFS

Theoretical Corporate: "Renegotiation-proof Contracting, Disclosure, and Incentives for Efficient Investment" with Nina Baranchuk (UT-Dallas) and Jun Yang (Indiana)

Empirical Corporate: "Money Grab in China," joint with Yingxue Cao (Tsinghua) and Joseph Qiu (Temple)

Theoretical Investments: "Lifetime Consumption and Investment: Retirement and Constrained Borrowing," with Hong Liu (Wash U)

Portfolio Delegation: "Portfolio Performance and Agency" with Heber Farnsworth (NISA) and Jennifer Carpenter (NYU)

http:dybfin.wustl.edu/research