FIN 539 MATHEMATICAL FINANCE
Lecture 4: FTAP, valuation, and the one-shot approach

Philip H. Dybvig
Washington University in Saint Louis

Fundamental Theorem of Asset Pricing (FTAP)

The following are equivalent:

- Absence of riskless arbitrage
- Existence of a consistent positive linear pricing rule
- Existence of an optimal choice for some hypothetical agent who prefers more to less

Originally in
Dybvig, Philip H., and Stephen A. Ross, 1987, "Arbitrage," a contribution to The New Palgrave: a Dictionary of Economics 1, New York: Stockton Press, 1987, 100-106.

This exposition follows
Dybvig, Philip H., and Stephen A. Ross, 2003, "Arbitrage, State Prices, and Portfolio Theory," Handbook of the Economics of Finance: Asset Pricing (Volume 1B), George M. Constantinedes, Milton Harris, and Rene M. Stulz, ed., North Holland, 605–637.

FTAP – notation

N: number of securities
Ω: number of states of nature
$W \in \Re$ initial wealth
$C \in \Re^{\Omega+1}$ consumption vector
$P \in \Re^{N}$: vector of security prices
$\Theta \in \Re^{N}$: vector of portfolio choices
$X \in \Re^{\Omega \times N}$: matrix of security payoffs

Budget constraint:

$$C = \begin{bmatrix} W \\ 0 \end{bmatrix} + \begin{bmatrix} -P' \\ X \end{bmatrix} \Theta$$

The first row has cash flows at time 0, and the the remaining rows have cash flows across states at time 1.

FTAP – arbitrage

An arbitrage is a money pump: something for nothing.

A net trade η, the change in portfolio choice from Θ to $\Theta + \eta$, gives us a change in consumption

$$\Delta C \; = \begin{bmatrix} W \\ 0 \end{bmatrix} + \begin{bmatrix} -P' \\ X \end{bmatrix} (\Theta + \eta) - \left(\begin{bmatrix} W \\ 0 \end{bmatrix} + \begin{bmatrix} -P' \\ X \end{bmatrix} \Theta \right) = \begin{bmatrix} -P' \\ X \end{bmatrix} \eta$$

An arbitrage opportunity is a net trade that increases consumption in some contingency and never reduces consumption:

$$\begin{bmatrix} -P' \\ X \end{bmatrix} \eta > 0$$

my notation for vector inequalities:
$X \geq Y$: $(\forall i) X_i \geq Y_i$
$X > Y$: $X \geq Y$ and $X \neq Y$
$X >> Y$: $(\forall i) X_i > Y_i$

FTAP – choice problem and pricing

Generic problem Choose Θ to
maximize $U(C)$ s.t.

$$C = \begin{bmatrix} W \\ 0 \end{bmatrix} + \begin{bmatrix} -P' \\ X \end{bmatrix} \Theta$$

We are interested in *strictly increasing* preferences. The utility function $U : \Re^{\Omega+1} \rightarrow \Re$ is called *strictly increasing* if $(\forall c, c')((c > c') \Rightarrow (U(c) > U(c'))$.

Pricing:

$$P' = p'X$$

$L(x) = p'X$ is a consistent linear pricing rule. We are interested in a consistent positive linear pricing rule, $p >> 0$.

Fundamental Theorem of Asset Pricing (FTAP): statement

The following are equivalent:

(i) Absence of riskless arbitrage: $(\nexists\eta)\left(\begin{bmatrix} -P' \\ X \end{bmatrix}\eta > 0\right)$

(ii) Existence of a consistent positive linear pricing rule: $(\exists p >> 0)(P' = p'X)$

(iii) Existence of a hypothetical agent who prefers more to less and has an optimal choice: there exists strictly increasing U and W such that the generic problem has a solution.

Proof: (i)⇒(ii) separation theorem
(ii)⇒(iii) by construction
(iii)⇒(i) by contradiction

Note: This is true as stated in finite dimensions, but requires more structure in general.

Pricing Rule Representation Theorem

The positive linear pricing rule can be represented equivalently using

(i) an abstract linear function $L(c)$ that is positive: $(c > 0) \Rightarrow (L(c) > 0)$

(ii) positive state prices $p >> 0$: $L(c) = \Sigma_{\omega=1}^{\Omega} p_\omega c_\omega$

(iii) positive risk-neutral probabilities π_i^* summing to 1 with associated shadow risk-free rate r^*: $L(c) = (1+r^*)^{-1} E^*[c_\omega] = (1+r^*)^{-1} \Sigma_{\omega=1}^{\Omega} \pi_\omega^* c_\omega$

(iv) positive state-price densities $\xi >> 0$: $L(c) = E[\xi c]$ (also called stochastic discount factor or pricing kernel)

Complete markets

When the pricing rule is unique, we say markets are complete. This is the case in which the one-shot approach is simplest, especially if we use the state-price density (stochastic discount factor) approach.

Choose $\{c_\omega\}$ to
maximize $\Sigma_{\omega=1}^{\Omega} \pi_\omega u(c_\omega)$
subject to $\Sigma_{\omega=1}^{\Omega} \pi_\omega \xi_\omega c_\omega = W_0$

FOC: $u'(c_\omega) = \lambda \xi_\omega$

In many periods (time separable vN-M utility):

Choose adapted $\{c_t\}$ to
maximize $E[\Sigma_{t=0}^{T} \delta^t u(c_t)]$ subject to $E[\Sigma_{t=0}^{T} \xi_t c_t] = W_0$

FOC: $\delta^t u'(c_t) = \lambda \xi_t$

The portfolio strategy solves an option replication problem.

Mini math review: multidimensional Itô's lemma

Let $H : \Re^d \times [0,T] \to \Re$ with continuous partial derivatives $H_x(x,t)$, $H_{xx}(x,t)$, and $H_t(x,t)$. Let $dX_t = g(t)dt + G(t)dZ_t$, where X_t is a d-dimensional process and Z_t is an m-dimentional standard Wiener process. Then $Y_t \equiv H(X_t, t)$ is an Itô process with stochastic differential

$$dY_t = H_t dt + H_x' dX + \frac{1}{2}\mathrm{tr}(GG'H_{xx})dt$$

where, for any symmetric matrix A, $\mathrm{tr}(A)$ denotes the trace, which is the sum $\Sigma_i A_{ii}$ of its diagonal elements.

Note: if H takes values in $\Re^k$, we can apply the result elementwise.

What is the trace?

The trace $\mathrm{tr}(A)$ of the square matrix A is the sum of its diagonal elements, $\Sigma_i A_{ii}$. The trace equals the sum of the eigenvalues.[1] For matrices A $i \times j$ and B $j \times i$, then $\mathrm{tr}(AB) = \mathrm{tr}(BA)$. For matrices C $i \times j$, D $j \times k$, and E $k \times i$, $\mathrm{tr}(CDE) = \mathrm{tr}(DEC) = \mathrm{tr}(ECD)$. If F and G both $n \times n$, $\mathrm{tr}(F + G) = \mathrm{tr}(F) + \mathrm{tr}(G)$ and $\mathrm{tr}(F') = \mathrm{tr}(F)$. Also, if X is $n \times n$, $d(\mathrm{tr}(X))/dX = I_{n \times n}$ and $d(\mathrm{tr}(AB))/dA = B'$.

[1] It is also useful to know that the determinant is the product of the eigenvalues.

Stochastic discount factor in continuous time

In continuous time, the stochastic discount factor (or state-price density or pricing kernel) is an adapted process ξ_t such that for "all" reinvested portfolios[2] having a price process P_t, we have that for $s < t$,

$$\mathrm{E}_s\left[\frac{\xi_t}{\xi_s}P_t\right] = P_s,$$

or equivalently, since ξ_s is known at time s,

$$\mathrm{E}_s[\xi_t P_t] = \xi_s P_s.$$

Therefore, $\xi_t P_t$ is a martingale for all re-invested marketed assets. Suppose that randomness is driven by an underlying k-dimensional Wiener process Z_t. The asset returns are given by $dS_{it}/S_{it} = \mu_{it}dt + \gamma_{it}dZ_t$, for $i = 0, ..., N$. Asset 0 is the riskless asset where $\mu_{0t} = r_t$ is the riskfree rate and $\gamma_{0t} = 0$. Since $\xi_t P_{it}$ is a martingale for all reinvested portfolioss, $\mathrm{E}[d(\xi_t P_{it})] = 0$.

[2] To make this rigorous, we would have to specify a set of feasible trading strategies to rule out bubbles. A simply but unappealing choice (because ξ is endogenous) is the set of assets for which $\mathrm{E}[\xi P]$ is a martingale.

Deriving the stochastic discount factor

Let's suppose the stochastic discount factor follows the process

$$d\xi_t = \xi_t(\mu_\xi dt + {\gamma_\xi}' dZ_t).$$

Now, we can apply the multivariate Itô's lemma, letting $X = (\xi, P_i)'$, and $H(X) = H(\xi, P_i) = \xi P_i$, then $f = (\xi\mu_\xi, P_i\mu_i)'$ and $G = (\xi\gamma_\xi, P_i\gamma_i)'$:

$$\begin{aligned}
0 &= E[d(\xi_t P_{it})] \\
&= E[P_{it} d\xi_t + \xi_t dP_{it} + \frac{1}{2}\mathrm{tr}\,(GG' H_{xx})\,dt] \\
&= P_{it}\xi_t\,(\mu_\xi + \mu_{it})\,dt + \mathrm{tr}\left(\left(\begin{array}{cc} \xi_t^2 {\gamma_\xi}'\gamma_\xi & \xi_t P_i {\gamma_\xi}'\gamma_i \\ \xi_t P_i {\gamma_i}'\gamma_\xi & P_i^2 {\gamma_i}'\gamma_i \end{array}\right)\left(\begin{array}{cc} 0 & 1 \\ 1 & 0 \end{array}\right)\right) dt \\
&= P_{it}\xi_t\,(\mu_\xi + \mu_{it})\,dt + \frac{1}{2}\mathrm{tr}\left(\begin{array}{cc} \xi_t P_i {\gamma_\xi}'\gamma_i & \xi_t^2 {\gamma_\xi}'\gamma_\xi \\ P_i^2 {\gamma_i}'\gamma_i & \xi_t P_i {\gamma_i}'\gamma_\xi \end{array}\right) dt \\
&= P_{it}\xi_t\,(\mu_{\xi t} + \mu_{it} + {\gamma_i}'\gamma_\xi)\,dt
\end{aligned}$$

Deriving the stochastic discount factor: continued

Since $(\forall i)(\mu_{\xi t}+\mu_{it}+\gamma_i{}'\gamma_\xi = 0)$, we can use the bond $n=0$ (with $\mu_0 = r$ and $\gamma_0 = 0$) to infer that $\mu_\xi = -r$. Then we have $(\forall n)(\mu_{it}-r+\gamma_i{}'\gamma_\xi = 0)$. As a vector equation (omitting $n=0$), we have

$$\mu - r\mathbf{1} + \Gamma\gamma_\xi = 0,$$

where

$$\Gamma = (\gamma_1|\gamma_2|\ldots|\gamma_N)'.$$

Now Γ is an $N \times k$ matrix. If $N = k$ and Γ is invertible, then markets are locally complete and

$$\gamma_\xi = -\Gamma^{-1}(\mu - r\mathbf{1}).$$

$$d\xi/\xi = -rdt - (\mu - r\mathbf{1})'(\Gamma')^{-1}dZ_t$$

Assuming these processes are not too wild, this will mean that ξ is unique given the initial condition $\xi_0 = 1$, and markets are complete.

Univariate stochastic discount factor: fixed coefficients

Our derivation of the stochastic discount factor is consistent with the riskfree rate r, vector of mean returns μ, and risk loadings Γ being adapted processes. However, the case of constants leads to the "lognormal model" which is interesting and useful. We will further specialize to the case of a single risky asset. Assuming the riskless asset has a constant return r and the risky asset has a constant mean return μ and constant risk exposure σ (so that $dS/S = \mu dt + \sigma dZ_t$), we have

$$d\xi_t/\xi_t = -rdt - \kappa dZ_t,$$

where $\kappa \equiv (\mu - r)/\sigma$ is the Sharpe ratio. This implies that

$$\xi_t = \xi_0 \exp((-r - \kappa^2/2)t - \kappa Z_t),$$

which is lognormal, since $\log(\xi_t/\xi_0) \sim N((-r - \kappa^2/2)t, \kappa^2 t)$. The stochastic discount factor ξ_t is lognormal in the multi-asset case as well, and can be used for calculations.

Stochastic discount factor and the stock price

The stock price $S_t = S_0 \exp((\mu - \sigma^2/2)t + \sigma Z_t)$ is also lognormal with the same underlying noise Z_t, so we can write ξ as a function of the stock price and time:

$$\begin{aligned} \log\left(\frac{\xi_t}{\xi_0}\right) &= \left(-r - \frac{\kappa^2}{2}\right) t - \frac{\kappa}{\sigma}\left(\log\left(\frac{S_t}{S_0}\right) - \left(\mu - \frac{\sigma^2}{2}\right) t\right) \\ &= -\frac{\kappa}{\sigma}\log\left(\frac{S_t}{S_0}\right) + \left(-r - \frac{\kappa^2}{2} + \frac{\kappa}{\sigma}\left(\mu - \frac{\sigma^2}{2}\right)\right) t \end{aligned}$$

or equivalently,

$$\frac{\xi_t}{\xi_0} = e^{ht}\left(\frac{S_t}{S_0}\right)^{-\kappa/\sigma}$$

where

$$h \equiv -r - \frac{\kappa^2}{2} + \frac{\kappa}{\sigma}\left(\mu - \frac{\sigma^2}{2}\right)$$

Pricing of options or a re-invested wealth process

The equation $\mathrm{E}[d(\xi_t P_t)] = 0$ has to hold for options and reinvested wealth processes as well as for the traded assets. In particular, suppose we have an option price or re-invested wealth process of the form $\mathcal{O}(S_t, t)$ where $\mathcal{O}()$ is smooth and S_t is one-dimensional. Since $dS_t = \mu S_t dt + \sigma S_t dZ_t$ we have $d\mathcal{O}(S_t, t) = (\mathcal{O}_t + \mu S_t \mathcal{O}_S + (\sigma^2/2) S_t^2 \mathcal{O}_{SS}) dt + \sigma S_t \mathcal{O}_S dZ_t$. Also, we have derived that $d\xi = -r\xi dt - \kappa \xi dZ_t$. Consequently the formula $0 = \mu_{\xi t} + \mu_{it} + {\gamma_i}' \gamma_\xi$ we derived for asset n becomes

$$0 = -r + \frac{\mathcal{O}_t + \mu S \mathcal{O}_S + (\sigma^2/2) S^2 \mathcal{O}_{SS}}{\mathcal{O}} - \frac{\kappa \sigma S \mathcal{O}_S}{\mathcal{O}}$$

Since $\kappa = (\mu - r)/\sigma$, this simplifies to

$$0 = -r\mathcal{O} + \mathcal{O}_t + rS\mathcal{O}_S + \frac{\sigma^2}{2} S^2 \mathcal{O}_{SS},$$

which is the Black-Scholes differential equation.

One-shot approach (Pliska (1986)[3])

If markets are complete, setting $\xi_0 = 1$, we can restate our standard portfolio problem as:

Given w at time 0,
choose adapted c_t and w_t to
maximize $\mathrm{E}[\int_{t=0}^{T} e^{-\rho t} u(c_t) dt + e^{-\rho T} b(w_T)]$
st $\mathrm{E}[\int_{t=0}^{T} \xi_t c_t dt + \xi_T w_T] = w$.

The first-order condition for the maximum is existence of λ such that $e^{-\rho t} u'(c_t) = \lambda \xi_t$ and $e^{-\rho T} b'(w_T) = \lambda \xi_T$. The solution is $c_t = I_u(\lambda \xi_t e^{\rho t})$ and $w_T = I_b(\lambda \xi_T e^{\rho T})$. For $0 \leq t \leq T$, we can compute the wealth w_t at time t from

$$\xi_t w_t = \mathrm{E}_t[\int_{s=t}^{T} \xi_s c_s ds + \xi_T w_T],$$

and compute the corresponding portfolio strategy by matching coefficients.

[3] Pliska, Stanley R, 1986, A Stochastic Calculus Model of Continuous Trading: Optimal Portfolios, Mathematics of Operations Research 11, 371-382. Popularized by Cox and Huang (1989).

One-shot approach: simple example

For $b(w) = w^{1-R}/(1-R)$, consider the terminal horizon problem ($u(c) \equiv 0$) in the case of a single risky asset and fixed coefficients. Then $\xi_t = e^{ht}(S_t/S_0)^{-\kappa/\sigma}$, and we have the following problem:

Given w at time 0,
choose adapted w_T to
maximize $\mathrm{E}[\frac{w_T^{1-R}}{1-R}]$
st $\mathrm{E}[e^{hT}(S_T/S_0)^{-\kappa/\sigma}w_T] = w$

The first-order condition is

$$w_T^{-R} = \lambda e^{hT}(S_T/S_0)^{-\kappa/\sigma}$$

which implies

$$w_T = \lambda^{-1/R}e^{-hT/R}(S_T/S_0)^{\kappa/(\sigma R)}.$$

One-shot approach: simple example, continued

Now, we have that

$$
\begin{aligned}
w_t &= \mathrm{E}_t\left[\frac{\xi_T}{\xi_t} w_T\right] \\
&= \mathrm{E}_t\left[e^{h(T-t)}\left(\frac{S_T}{S_t}\right)^{-\kappa/\sigma} \lambda^{-1/R} e^{-hT/R}\left(\frac{S_T}{S_0}\right)^{\kappa/(\sigma R)}\right] \\
&= \left(\frac{S_t}{S_0}\right)^{\kappa/(\sigma R)} \mathrm{E}\left[\lambda^{-1/R} e^{h(T-t)-hT/R}\left(\frac{S_T}{S_t}\right)^{\kappa(1-R)/(\sigma R)}\right] \\
&= Q(t)\left(S_t/S_0\right)^{\kappa/(\sigma R)}
\end{aligned}
$$

for some function $Q(t)$, since S_T/S_t is independent of S_t. If we want to, we can compute $Q(t)$ exactly (and also then λ from the expression for w_0), since $\log(S_T/S_t) \sim N((\mu - \sigma^2/2)(T-t), \sigma^2(T-t))$ and $\log((S_T/S_0)^{\kappa(1-R)/(\sigma R)}) = (\kappa(1-R)/(\sigma R)) \log(S_T/S_0)$.

One-shot approach: simple example, continued 2

Matching the change in wealth to what would be implied by a risky asset investment θ_t, we have

$$\begin{aligned} dw_t &= d(Q(t)(S_t/S_0)^{\kappa/(\sigma R)}) \\ &= w_t((\ldots)dt + \frac{\kappa}{\sigma R}\sigma dZ_t) \\ &= rwdt + \theta((\mu - r)dt + \sigma dZ_t) \end{aligned}$$

so that matching the coefficients of dZ_t implies that $\theta = \frac{\kappa}{\sigma R} w$.

One-shot approach: financial engineering tips

The common standard utility functions (CARA, CRRA, HARA) all have closed forms for the inverse marginal utility function, so they are good candidates for the one-shot approach. So do the GOBI utility[4] used in the first homework set and its close relative SAHARA utility[5]. I also like using piecewise HARA utility:

$$u(c) = \begin{cases} a_0 + b_0 \frac{c^{1-R_0}}{1-R_0} & \text{for } c \le c_0 \\ a_1 + b_1 \frac{c^{1-R_1}}{1-R_1} & \text{for } c_0 < c \le c_1 \\ \vdots & \\ a_n + b_n \frac{c^{1-R_n}}{1-R_n} & \text{for } c_{n-1} < c \end{cases}$$

For all i, choose $b_i > 0$ and $R_i > 0$, and match the derivatives to make $u(c)$ continuous and differentiable at the boundaries c_i.

[4]Dybvig, Philip H., and Fang Liu, 2018, On Investor Preferences and Mutual Fund Separation, Journal of Economic Theory 174, 224–260.

[5]Chen, An, Antoon Pelsser, and Michel Vellekoop, 2011, Modeling non-monotone risk aversion using SAHARA utility functions, Journal of Economic Theory 146, 2075–2092