Financial Modeling and Optimization under Uncertainty

Sixth Summer School on Optimization and Financial Modeling

Garich, Munich, October 9–11, 2002.

Location:

Munich University of Technology, Center for Mathematical Sciences, Mathematical Finance, Boltzmannstr. 3, 85748 Garching, Germany.

Special Events:

Wednesday, 6pm: A Guided Tour of Munich.

Thursday, 3pm: Guest lecture by Dr. Rudi Zagst: "Managing Interest Rate Risk".

Thursday, 7pm: Summer School Dinner, location to be announced.

Acknowledgements:

The GAMS Summer School would like to thank the Center for Mathematical Sciences, University of Munich, and Professor, Dr. Rudi Zagst for hosting the School.

Main Course Materials:

PFO: Zenios: *Practical Financial Optimization* Book
LIB: Nielsen/Zenios: *A Library of Financial Optimization Models*
FO: Zenios (ed): *Financial Optimization*
FinNotes: Nielsen (notes on CD: `fin_notes.pdf`, a brief overview)
Models: found on the CD in directory `Models`, a few in `gamsfinance`.

Course Outline:

We meet at 9:00am. Lunch 12–13. Coffee 10:30–11 and 14:30–15. Finish around 16:30.

Wednesday: Fundamentals. Classical Models. The GAMS System.

- Financial Risks and Optimization
- Dedication Model: Starting simple: variables, constraints, objective
- PC-Lab. GAMS: Installing, the GUI, the language, solving models.
- Borrowing and Lending, Lot-size constraints, Transaction costs
- LUNCH
- Immunization: Theory and models
- Mean-Variance (or Markowitz) Model.
- PC-Lab: Exercises on the models covered today.

Text: LIB: Chapter 1, Sections 2.1, 2.2, Chapter 3. PFO: Chapters 2, 3 and 4; FinNotes: Sections 2.1–2.3 and 5.
Models: dedication, BondModel, FinCalc_c, FinCalc_d, Trade, Immunization, Factor, FactDir, MeanVar, MeanVarMIP, MeanVarShort.

Thursday: Utility Theory. Modeling using Scenarios. Value at Risk

- Scenario Dedication
- Horizon Return,
- Mean-Absolute Deviation (MAD)
- Utility Theory. Models maximizing Expected Utility.
- LUNCH
- Value at Risk (VaR) and Conditional VaR (CVaR).
- **Dr. Rudi Zagst's Guest Lecture.**

Text: LIB: Chapter 4. PFO: Chapter 5; FinNotes 2.4
Models: Horizon MAD, MADTrack, Util, CVaR, VaR.

Friday: Stochastic Programming; Case Study

- Stochastic Programming, the Scenario Tree, Formulations,
- PC-Lab: Winvest: Developing a Stochastic Program Based on Windata.gms, follow the directions in `WinQuestions.pdf`.
- Some theory: Formal Mathematical Formulation; Generating Interest-Rate Scenarios.
- LUNCH
- PC Lab: The Winvest Model, continued: Working with Risk-Attitudes.
- A Larger Case Study: Danish Mortgage Model.
- A Larger Case Study (if time): International Asset Allocation.
- Participants who would like to present their own models for discussion.

Text: LIB: Chapter 5, Sections 7.5 and 7.2. Edinburgh.pdf, Mortgage_paper.pdf; PFO: Chapter 6; FinNotes: Sections 3 and 4.
Models: Windata (to cheat, see: Winvest), DanCase, IntlAssets.

Financial Risks: A Classification

- Different investors worry about different risks.
- Different models address different risks.
- Investors *trade* risks, e.g. through derivatives.

But what are the risks involved? Zenios/Meeraus/Dahl gives a classification:

Market Risk: Movement of an entire market, e.g., the US S&P-500, or the Europaean Bond Market, measured by some broad bond index

Shape Risk: (Fixed-Income) Changes in the shape of a country's term structure

Volatility Risk: (Options) A market's volatility directly influences options prices, and is perceived as a risk in other investments.

Sector Risk: Risk of co-movements of all securities within a sector: The Oil industry; the mortgage bond market; technology stocks, ...

Currency Risk: Risk coming from international exposure: international trade, multinational corporations, international portfolios

Credit Risk: Risk that a borrower may default, or that credit instruments may be downgraded (lose value).

Liquidity Risk: The risk of being unable to buy or sell when desired due to lack of market liquidity. A concern for large, actively managed portfolios or portfolios of thinly traded instruments.

Residual Risk: "All the rest": company-specific risks etc.

As *investors* we need to identify which risks we accept, and which risks we wish to protect (hedge) against.

As *modelers* we need to model the different risk types, and appropriately weigh risk-return characteristics of investments.

The Role of Modeling and Optimization

Almost all financial instruments are *packages* of risks:

Bonds: Sector Risk (interest rates), shape risk (except zero-coupon), credit risk (especially corporate bonds), ...

Stocks: Market risk, sector risk, liquidity risk, company-specific risk, ...

Many *derivatives* are attempts to "unpackage" risks to allow trading or hedging them separately. Examples:

1. options on the S&P 500 index to offset market risk.
2. currency futures to offset currency risk.
3. options on T-bond futures contracts to offset interest-rate risk

Modeling helps uncover the *underlying factors* that influence financial returns and risks, the qualitative nature of these factors, and thus correlations among financial instruments.

Optimization then helps compose diversified portfolios where different instruments offset (or cancel) each others' risks, or that in other ways seek *optimal* characteristics with respect to risk, expected return, deviation from target, ...

<u>but be careful...</u>

"Optimization makes a good portfolio manager better,
and a bad one worse"

– Maximizing returns necessarily maximizes risks (efficient markets).

– Eliminating controlled risks and then maximizing returns *maximizes* uncontrolled risks!

A Deterministic Model: Dedication.

A company must pay out a future *liability stream*: L_{2000} in year 2000, L_{2001} the year after, and so on (assumed known, i.e., deterministic). Examples: Pensions, lottery payouts.

They would like to forget about these future liabilities by (1) buying a risk-free portfolio (T-bonds) whose cash-flows will cover the liabilities, or (2) selling the liability stream to some financial corporation at a *fair* price.

How do we find the cheapest *dedicated* portfolio? How do we determine the fair price?

Setting up the notation:

$T = \{0, ...m\}$: The set of time periods, for instance measured in years, from $t = 0$ ("now") to $t = m$, the **horizon**.

$U = \{1, ..., n\}$: The *Universe* of assets (bonds) under consideration for inclusion in the portfolio.

$F_{i,t}$: The cash flow arising from asset i at time t.

L_t: The liability due in period $t \geq 1$.

P_i: The price of one unit of asset i

T and U are *sets* that specify the dimension and size of the model: We model m time periods, and we consider a universe of n bonds.

Simplification: (1) Assume that the assets are risk-free bonds, (2) with one coupon payment per year, (3) whose timing coincides with the payment of the liability.

Mathematical Formulation:

Decision Variables:

x_i : Number of bonds $b \in U$ to purchase.

v_0 : Up-front payment, or lump sum, to pay out today.

Model 0.1 Portfolio Dedication (without Borrowing).

$$\text{Minimize} \quad v_0 \tag{1}$$

$$\text{subject to} \quad v_0 - \sum_{i \in U} P_i x_i = v_0^+, \tag{2}$$

$$\sum_{i \in U} F_{it} x_i + (1 + \rho_{t-1}) v_{t-1}^+ = L_t + v_t^+, \qquad t = 1, \ldots, T, \tag{3}$$

$$x, v^+ \geq 0. \tag{4}$$

♠

Comments:

1. *Constraints:* For each year after the first, the (incoming) cash flows must be at least as large as the (outgoing) liability.

   The first year is special: Here, the cash flows are actually negative (bond purchase prices), and we must put in a lump sum, λ, to cover these purchases.

2. *Objective:* Naturally, to minimize the initial lump sum we must pay.

3. *Non-negativity:* All x_b are non-negative. Without this constraint we would allow *short-selling* of bonds, which is actually possible in certain contexts.

GAMS Notation

See `dedication.gms` in the `Models` directory.

Comments:

1. The GAMS names are different from the "mathematical" ones - different naming styles in mathematics and programming/modeling. For instance, years are 0, 1, 2, ... or 2001, 2002, 2003, ...
2. The set `time` has an index `t`, associated with it, and `bond` has the index `i` assocated with it. This is a way to distinguish (unlike GAMS) between a *set* and an *index* into the set.
3. Note the use of `tau(t)` to map time indices to calendar years.
4. Note the use of the asterisk in `/2002 * 2006/`, to abbreviate `/2002, 2003, 2004, 2005, 2006/`.
5. Note `cf(b,t)`, the Cash Flow from one unit of bond b in year t – The first year's cash flow is negative, namely the bond's purchase price.
6. The bonds' cash flows are calculated by GAMS based on fundamental data. Good principle!

Including Short-Term Reinvesting and Borrowing

An unfortunate effect of dedicate.gms is that year t's liabilities can only be covered by a bond that matures in year t, or by coupon payments from bonds maturing later. This basically requires that we purchase as many different bonds as there are years in our horizon.

Model files: BondData.inc, BondModel.gms.

We may get a better (cheaper) solution if we allow surplus incoming cash to be saved (reinvested) until next year, or allow borrowing against future cash flows.

More Decision Variables:

$\mathbf{surplus}_t$: Amount saved in year t, to be available (with interest) in year $t+1$,

$\mathbf{borrow}_t$: Amount borrowed in year t, to be paid back (with interest) in year $t+1$.

We let ρ_t be the savings interest rate for year t, and $\rho_t + \gamma$ be the borrowing interest rate for year t (γ is called `spread` in the model).

Comments:

1. Left-hand sides of constraints are *incoming* cash (except purchases), right-hand sides are *outgoing* cash.
2. Borrowing is allowed from one time period to the next, except ing in the last time period (why?). How is this implemented in the model?
3. ... but we can invest money (put into savings) in the last year. Does this make sense? Is it wrong to allow it?
4. Are any numbers legal for the interest rate and spread, $\rho, +\gamma$ (`rho` and `spread`)? For instance, what happens if the borrowing rate, $\rho + \gamma$ (`rho + spread`), are very low?
5. Will later incorporate tradability considerations: x_b integers, or multiples of 100, and transactions costs.
6. "surplus", "borrow" strange names in mathematics but fine in modeling!

Letting GAMS do the Cash Flow Calculations

Good principle: *Only specify the most basic data to the model.* For bonds, these are: Price, Maturity, and Coupon rate (assuming one coupon payment per year):

```
* calculate the ex-coupon cashflow of Bond i in year t:
  F(i,t) =        1      $ (tau(t)  = Maturity(i))
            +  coupon(i) $ (tau(t) <= Maturity(i) and tau(t) > 0);
```

Also a good example of using the $-operator.

For more examples of GAMS used as a "Financial Calculator", see FinCalc_c.gms, FinCalc_d.gms.

Tradeability and Transaction Costs

Now we explicitly model some considerations that occur in real-life trading, such as lot-sizes and transaction costs.

Model files: Trade.gms and TransCosts.inc.

1. Trade.gms, model `TransCost1` models "Even Lot-size Constraints": Bonds can only be purchased in multiples of `LotSize`, $1000. This is modeled using integer variables, `Y(i)`.
2. Trade.gms, model `TransCost2` models Fixed and Variable transactions costs: Every purchase (no matter the size) has a `FixedCost`, there is also a proportional (or variable) cost, `VarblCost`.

   The fixed cost is modeled using binary variables, `Z(i)`.

Comments:

1. Have now introduced *Mixed-Integer Programs (MIP).*
2. For integer models: Suggest using

   `OPTION optcr = 0, reslim = 1800, iterlim = 999999;`

   and explicit bounds on all integer variables.

Immunization

Real-life example from the 80's:

1985: An insurance company sells a GIC (Guaranteed Investment Contract) worth \$1,000,000, maturing in 1992. They promise to pay an annual rate of 5% during these 7 years.

They invest the money, \$1,000,000, invested in a secure bond with the highest yield available: The 30-year zero-coupon Treasury STRIP, with a yield of 7%. The face value amount of this investment is $\$1{,}000{,}000 \cdot 1.07^{30} = \$7{,}612{,}300$.

1985-92: Interest rates rise by a modest 2%. In particular, the 30-year zero yielding 7% becomes a 23-year zero yielding 9%.

1992: The company pays back the GIC, now worth $\$1{,}000{,}000 \cdot 1.05^{7} = \$1{,}407{,}100$.

To do this, the company liquidates its zero-coupon bonds, which are now worth $\$7{,}612{,}300 \cdot 1.09^{-23} = \$1{,}048{,}800$.

The company faces a loss of \$358,300, even though it *paid* 5% and invested at 7%!

What happened? Over a 7-year period, interest rates increased from about 7% to about 9% - not very dramatic. Yet, the value of the company's assets fell drastically, but the value of its liabilities didn't change.

The asset-liability balance was *duration-mismatched*. The change in value of the asset-liability balance as a function of changes in interest rates were grossly uneven.

Present Value and Duration

Assumptions of the Immunization Model: We have a known (deterministic) liability stream. We wish to buy an asset portfolio with the same *Present Value*, and with the same sensitivity to general changes in interest rates (parallel shifts in the term structure)[1].

$t \in T$**:** Time indices where bond cash flows occur

y_t**:** The yield of a t-year investment

$F_{i,t}$**:** Bond i's cash flow occurring at time t

P_i**:** Present value of bond i's cash flows

k_i**:** Dollar duration of bond i's cash flows

Present Value of bond i's Cash Flows:

$$P_i = \sum_{t \in T} F_{i,t}(1 + y_t)^{-t} \tag{5}$$

Present Value of Liability:

$$P_L = \sum_{t \in T} L_t(1 + y_t)^{-t} \tag{6}$$

How do the present values change with a parallel shift of c in interest rates? Write

$$P_i(c) = \sum_{t \in T} F_{i,t}(1 + y_t + c)^{-t} \tag{7}$$

and consider the **Dollar Duration**, or the *derivative of P_i with respect to c*:

$$k_i = \frac{dP_i}{dc} = \sum_{t \in T} -t \cdot F_{i,t}(1 + r_i)^{-(t+1)}. \tag{8}$$

The liability's dollar duration is:

$$k_L = \frac{dP_L}{dc} = \sum_{t \in T} -t \cdot L_t(1 + r_i)^{-(t+1)}. \tag{9}$$

[1]Notes, 5.1, where D_i^{dol} is used for dollar duration

The Immunization Model

We want a model that builds a portfolio with[2]:

1. Equal present value on the asset and liability sides (NPV = 0),

$$\sum_{i \in \mathcal{U}} P_i x_i = P_L,$$

2. Equal Dollar Duration on the asset and liability sides ("NDD" = 0)

$$\sum_{i \in \mathcal{U}} k_i x_i = k_L$$

Then the total position is "immunized" against small, parallel shifts in the term structure.

Note: The quantities P_i, P_L, k_i and k_L are calculated "outside" of the optimization model, in GAMS. In the immun2 model (GAMS Financial model library), P_L and k_L are calculated using continuous compounding.

[2] notes, 5.2

The Immunization Model's Objective

What should we optimize in this model?

Minimize the Portfolio's Purchase Price: In general meaningless because the purchase price, in an efficient market, equals the Present Value of the portfolio - which is fixed to equal P_L.

Maximize the Portfolio's Yield to Maturity: A first-order approximation to the portfolio's YTM is:

$$YTM = \frac{\sum_{i \in U} k_i r_i x_i}{\sum_{i \in U} k_i x_i}$$

where r_i is the YTM of bond i. The denominator equals P_L by constraint, so we just maximize the (negative of the) numerator:

Immun-1: *(Simple Immunization Model)*

$$\begin{array}{llll}
\underset{x \in U}{\text{Maximize}} & -\sum_{i \in U} k_i \cdot r_i \cdot x_i & & \\
\text{Subject to} & \sum_{i \in U} P_i \cdot x_i & = P_L & \\
& \sum_{i \in U} k_i \cdot x_i & = k_L & \\
& x_i & \geq 0 & i \in U
\end{array}$$

Comments:

1. The model has only *two constraints*. In general, this means that the solution will contain only *two bonds*!

   The solution is often a "barbell" portfolio: One very long-maturity bond (high yield), and a very short-maturity bond (low duration). This portfolio is, unfortunately, *maximally exposed to non-parallel shifts* in the term structure!

2. Another duration measure is *Macaulay-duration*, which can be viewed as the "average timing of the cash flows". For instance, a t-year zero coupon bond has Macaulay-duration t. For optimization purposes, Dollar- and Macaulay-duration are virtually equivalent[3]. A third duration measure is *Modified Duration*.

This is a good example of unintentionally maximizing uncontrolled risks, in this case, shape risk.

The classical "fix" is *Convexity Matching*, but a more rigorous and direct approach is to use *Factor Immunization*.

[3] Notes, 5.3

Convexity Matching

Convexity is the *second* derivative of P_i or P_L with respect to c:

$$Q_i = \frac{d^2 P_i}{dc^2}, = \sum_{t \in T} t(t+1) F_{i,t} (1+y_t)^{-(t+2)},$$

$$Q_L = \frac{d^2 P_L}{dc^2}, = \sum_{t \in T} t(t+1) L_t (1+y_t)^{-(t+2)}.$$

Convexity measures (roughly) sensitivity to non-parallel, and to large, shifts in the term structure. The "barbell" portfolio has large convexity, and we thus wish to keep convexity *small.*

But, convexity is also the curvature of the Present Value graph as a function of parallel term structure shifts. Hence, if Net Convexity is non-negative, then the Asset Present Value will increase faster than the Liability Present Value when interest rates fall, and decrease slower when interest rates increase! Hence, we also want to keep Net Convexity non-negative:

Immun-2: *(Immunization Model with Convexity)*

$$\begin{array}{llll}
\underset{x \in U}{\text{Minimize}} & \sum_{i \in U} Q_i \cdot x_i & & \\
\text{Subject to} & \sum_{i \in U} P_i \cdot x_i & = P_L & \\
 & \sum_{i \in U} k_i \cdot x_i & = k_L & \\
 & \sum_{i \in U} Q_i \cdot x_i & \geq Q_L & \\
 & x_i & \geq 0 & i \in U
\end{array}$$

But still, convexity matching is somewhat "ad hoc".

Factor Immunization

We have immunized against parallel shifts in the term structure. It might also be desirable to be immunized against:

- changes at the short end (or long end),
- non-parallel changes (short end down, long end up or vice-versa),
- changes to the curvature of the yield curve

Each of these possible types of change is modeled as a **factor**.

Let the term structure be represented by a vector

$$y_t = (y_1, y_2, y_3, ..., y_30)^T$$

Then a parallel shift (factor $j = 1$) is represented by the vector

$$a_{1,t} = (1, 1, 1, ..., 1)^T$$

because any parallel change in y_t can be modeled by adding some (positive or negative) multiple, F_j[4], of a_1 to y_t.

Similarly, a non-parallel shift could be represented by:

$$a_{2,t} = (-15, -14, -13, ..., -1, 0, 1, ..., 13, 14)^T$$

and a curvature change by:

$$a_{3,t} = (0, 0.5, 1, ..., 4, 4.5, 4, 3.5, ..., 0.5, 0)^T$$

[4] The notation is somewhat confusing: $F_{i,t}$ denotes cash flows; F_j denotes a factor level

The Factor Immunization Model

The model explicitly hedges against the types of changes represented by the factors, $j \in J$. As before, the Present Value of bond i is:

$$P_i = \sum_{t \in T} F_{i,t}(1 + y_t)^{-t} \tag{10}$$

Consider the derivative of P_i with respect to the y_t:

$$dP_i = -\sum_{t \in T} F_{i,t} \cdot t \cdot (1 + y_t)^{-(t+1)} \cdot dy_t$$

If the term structure y_t is modified by F_j units of factor j, then it changes by:

$$dy_t = a_{j,t} \cdot dF_j$$

which, substituted into the above yields:

$$dP_i = -\sum_{t \in T} F_{i,t} \cdot t \cdot (1 + y_t)^{-(t+1)} \cdot a_{j,t} \cdot dF_j$$

which defines the bond's *loading factor* for factor j:

$$f_{i,j} \doteq \frac{dP_i}{dF_j} = -\sum_{t \in T} a_{j,t} \cdot F_{i,t} \cdot t \cdot (1 + y_t)^{-(t+1)}$$

This is the sensitivity of the i'th bond's present value to changes of the term structure, as specified by the j'th factor.

Immun-3: *(Factor Immunization Model)*

$$\begin{array}{llll}
\underset{x \in U}{\text{Maximize}} & -\sum_{i \in U} k_i \cdot r_i \cdot x_i & & \\
\text{Subject to} & \sum_{i \in U} P_i \cdot x_i & = P_L & \\
 & \sum_{i \in U} f_{i,j} \cdot x_i & = f_{L,j}, & j \in J \\
 & x_i & \geq 0 & i \in U
\end{array}$$

Comments:

1. The objective is, again, to maximize the portfolio yield.
2. What happens if we include 30 factors, where factor j has a 1 in position j, 0 elsewhere (i.e., one for each y_t)?

The Mean/Variance (Markowitz) Model

Modeling fixed-income instruments (such as bonds) comes down to modeling interest rates. The Mean/Variance, or Markowitz, model addresses **stocks**. It does so using entirely historical stock returns (in the simplest case)[5].

The idea is to build a portfolio:

1. that *maximizes the expected return* of the investment,
2. that *minimizes the risk*, measured by the variance of returns, of the investment

There's an inherent conflict: How can one simultaneously minimize risk and maximize expected return?

[5] see FO Chapter 1, section 3.5

Notation and Derivations

Assume that we have observed the value of an investment in a unit holding of stock $i \in U$ at time points $t \in T$, observed every Δt years (e.g., monthly ($\Delta t = 1/12$), semi-yearly ($\Delta t = 1/2$), etc.).

Then

$$r_{i,t} = \frac{S_{t+1} - S_t}{\Delta t \cdot S_t} \qquad \text{or} \qquad r_{i,t} = \frac{1}{\Delta t} \ln \frac{S_{t+1}}{S_t}$$

is the (annualized) return of the investment in time period t, and

$$\mu_i = \sum_{t \in T} r_{i,t}$$

is the average historically observed return.

In addition,

$$\sigma_{i,j} = \sum_{t \in T} (r_{i,t} - \mu_i) \cdot (r_{j,t} - \mu_j)$$

is the covariance of returns of stocks i and j. In particular,

$$\sigma_i^2 = \sigma_{i,i}$$

is the historical variance of returns of stock i.

For a portfolio P consisting of an x_i investment in stock $i \in U$, the total portfolio variance is then

$$\mathrm{Var}_P = \frac{1}{2} \sum_{i \in U} \sum_{j \in U} x_i \cdot \sigma_{i,j} \cdot x_j$$

and the total portfolio average (or expected) return is

$$\mathrm{ER}_P = \sum_{i \in U} \mu_i \cdot x_i$$

The Model:

There are several equivalent version of the model. A simple one is:

Mean-Variance, or Markowitz Model:

$$\begin{array}{ll} \underset{x \in U}{\text{Maximize}} & (1-\lambda) \cdot \left(\sum_{i \in U} \mu_i \cdot x_i \right) - \lambda \cdot \left(\frac{1}{2} \sum_{i \in U} \sum_{j \in U} x_i \cdot \sigma_{i,j} \cdot x_j \right) \\ \text{Subject to} & \sum_{i \in U} x_i = 1 \; ; \qquad x_i \geq 0, i \in U \end{array}$$

where $\lambda \in [0,1]$ is a parameter.

Comments:

1. The model scales the total dollar amount of the invest to equal 1. Hence, the x_i can be interpreted as *fractions* of the total investment that go into each stock.
2. The model uses variance as the only measure of risk. A more general approach is to use utility.
3. The conflict between *maximizing expected return* and *minimizing portfolio variance* is resolved by weighing these two objectives by the factor λ: $\lambda = 0$ means to focus entirely on risk, $\lambda = 1$ means focus only on return.
4. The model relies entirely upon historical data. In practice one would incorporate analysts' estimates and judgements about individual firms, or use company "betas" in a "Mean-Variance Factor Model".
5. The model is non-linear - more precisely, it is linearly constrained with a quadratic objective. The objective is convex; hence, it's still an "easy" model to solve, almost as easy as an LP.
6. Individual investors have different risk-preferences, indicated by wishing different optimal portfolios, for different values of λ. By plotting the (Var, ER) points for optimal portfolios corresponding to all values of λ, one gets the *efficient frontier*:

Note: We will work with the "cent1.gms" model. Copy it, but erase the model at the bottom - we want to build our own! The resulting file contains the expected returns and covariances, $\mu_i, \sigma_{i,j}$.

A standard criticism of the Mean-Variance models is that it penalizes up-side and down-side risk equally, whereas most investors don't mind up-side "risk". The use of utility theory is one way to address this problem[6].

[6]See FO Chapter 1, section 3.6 for a slightly different formulation

Introduction to Utility Theory

How do we behave when faced with a risky situation – like an investment?

Utility Theory attemps to answer this question by using simple behavioral models. Simple *games* (investments == games) are used to explore people's behavior.

Why not simply maximize *expected return*?

Consider two simple games:

You HAVE to play one of the games! Which one do you prefer?

G1 has an expected return of Eur 50. G2 has an expected return of Eur 5. Most people prefer G2 – we do *not* simply maximize expected return. (If not convinced, try "multiplying" G1 by 100!)

Losses hurt more than Gains please!

One way to explain this behavior is to assume that a loss *hurts more* than an equivalent gain *pleases*:

An equal chance of winning X Eur or losing X Eur is altogether negative, not "neutral".

The figure to the right shows a simple *utility function*: "Happiness" as a function of final wealth (financial position).

Let us evaluate G1 and G2 but this time *maximize expected utility, EU* (read off from graph): EU(G2) is greater than Eu(G1) – this might explain why most people prefer G2.

Shape of the Utility Function

How should the utility function look? How can we get closer to an answer? And what does the mathmetics look like?

We will investigate these questions in two very different ways: (1) through simple gambles, and (2) through a simple investment.

Simple gamble:

p chance of winning EUR 100 and of losing 0. What would you pay to play? p is varied from 0% to 100%; we *exhibit* our own, empirical utility function:

A Simple Investment Example

Assume we invest EUR 1000 (or any other amount) for one year. The return is uncertain: After one year we have EUR $1000 \cdot \tilde{r}$, where $\tilde{r}$ is the uncertain return.

After one year we *reinvest* our $1000 \cdot \tilde{r_1}$ in a similar investment, and so on for n year. We cannot invest *more* money, nor take money out. In every year, the return has the same the same uncertain return, $\tilde{r_t}$, but is independent for each year: (Mossin, 1968)

- Returns are identically distributed each year
- Returns are independent from year to year
- No transactions costs

Since the time horizon is long, we will *maximize the expected return*[7] after n years, which is:

$$1000 \cdot (1 + \tilde{r_1}) \cdot ... \cdot (1 + \tilde{r_n})$$

or: Chose an investment (i.e., select a "good" $\tilde{r_t}$ to:

$$\text{Maximize E}\left\{1000 \cdot (1 + \tilde{r})^n\right\}.$$

The following maximization problems are then equivalent (have same optimal solution/investment):

$$\text{Maximize E}\left\{(1 + \tilde{r_t})^n\right\},$$

$$\text{Maximize E}\left\{\exp(n \cdot \log(1 + \tilde{r_t}))\right\},$$

$$\text{Maximize E}\left\{n \cdot \log(1 + \tilde{r_t}\right\},$$

$$\text{Maximize E}\left\{\log(1 + \tilde{r_t}\right\}.$$

The last problem says to *select the portfolio which has the highest expected log of returns.*

Or, in other words: Maximize expected utility where the utility function is the logarithm!

Note: We will never chose an investment which risks losing *all* our money (that is, we can assume that $\tilde{r_t} > -1$ a.s.)

[7] Why do we suddenly accept to maximize expected return? Because it's a repeated gamble/investment...

Reconsider G1, G2: Technical problem, cannot take the log of negative numbers (violates $\tilde{r}_t > -1$) – so look at G1', G2' by adding EUR 2000 to G1 and G2: Now,

EU(G1) = 1/2 log (3100) + 1/2 log (1000) = 7.47

EU(G2) = 1/2 log (2020) + 1/2 log (1990) = 7.60.

So G2 is indeed the most attractive game!

By two very different (ad hoc) arguments, we have discovered that a reasonable utility function has these properties:

- It is *increasing*: More is better.
- It is *concave*: Additional wealth becomes less and less important, additional losses hurt more and more.
- The logarithm seems to have special significance in Finance.

Generalization

We have derived the fact that the logarithm is *growth-optimal*, in that it maximizes the *long-term, expected return* under mild conditions.

Does this mean we should always seek investments (even for a single period, a "one-shot" investment) that maximize the expected logarithm?

NO! Different people have different risk-attitudes. Some are very risk-averse, some are almost risk-neutral.

The *Iso-elastic utility functions* is a family of functions that are useful in Finance (for deeper reasons than we can cover here; one is the concept of *Constant Relative Risk-Aversion, CRRA* (see later)).

These functions have a *risk-attitude* parameter, α, and include risk-neutrality, the logarithm, and even more risk-averse behavior as special cases.

Utility Models

Define $U(x)$: The *utility* of a financial position, or return, x. We assume these properties:

1. $U(\cdot)$ is continuously differentiable over $\Re_+$
2. $U'(x) = dU/dx > 0$ ("more is better"),
3. $U''(x) = d^2U/dx^2 < 0$ ("more matters less the more you already have", or *risk aversion*).

The basic idea is to

Construct a portfolio that *maximizes the Expected Utility of returns.*

Financially relevant utility functions

The most-often used utility functions in finance are the so-called
Iso-elastic utility functions:

$$U_\alpha(x) = \begin{cases} \frac{1}{1-\alpha}(x^{1-\alpha} - 1) & \text{for } \alpha \neq 1, \\ \log(x) & \text{for } \alpha = 1, \end{cases}$$

$\alpha \geq 0$ is the risk-preference parameter (like λ in Mean-Variance): Higher values of α indicates higher risk-aversion. If $\alpha = 0$, this models risk-neutrality (maximizing expected return).

The Iso-elastic utility functions model "Constant Relative Risk Aversion": No matter how much money you invest, the relative composition of the optimal portfolio (for a given α) is the same. Small and big investors invest in the same portfolio if they have the same risk preferences.

The case $\alpha = 1$, or $U_1(x) = \log x$, has special significance: In a multi-period setting, where the financial position after each period is the initial investment in the next period, the log utility function, when applied in each separate period, maximizes the long-run return — it is *growth-optimal*[8].

[8] Under certain assumptions - Mossin 1968.

The Certainty-Equivalent

Having optimized a utility model, we have an optimal *Expected Utility*, e.g.:

$$\text{Maximize } EU = \sum_{s \in s} p^s U(W^s)$$

How good is it? What does the optimal EU mean?

An answer can be obtained by calculating the *Certainty-Equivalent*, defined as:

> The Certainty-Equivalent amount CE is the monetary amount which has the *same utility* as the optimal, expected utility:
>
> $$CE = U^{-1}(EU)$$

It is the amount such that a decision maker would be *indifferent* between the (uncertain) investment having the specified expected utility EU, and receiving the (certain) amount CE.

Example

Assume a decision maker's risk attitudes are accurately represented by the log utility function. He faces an investment that returns either \$1000, or \$2000, with equal (50%) chance.

The expected utility is $\frac{1}{2}(\log 1000 + \log 2000) = 7.254$. Then $CE = \exp 14.509 = \$1414.21$.

Another interpretation: The investor would be willing to pay up to \$1414.21 for this investment, but not more. He would prefer to receive the certain amount \$1415 for sure rather than receiving the investment.

This is true only for an invester that has the log as his/her utility function!!! For others, the specific numbers will differ.

Note: The expected return, $ER = 1500.00$. Risk-aversion is characterized by always having $CE < ER$. The *Risk-Premium*, $ER - CE$, measures the degree of risk-aversion. It can be shown to be roughly proportional to the variance of the gamble, which can be taken as a defense for the Mean-Variance model.

Mean-Variance vs. Utility

Review the Mean-Variance Model in view of Utility Theory:

What is the MV-utility function (if any)? Is it concave?

Value at Risk,

Contitional Value at Risk,

and Coherent Risk Measures

Maximizing *expected return* is fine if you are *risk-neutral.* Nobody is. Maximizing *expected utility* is fine if you understand it.

But Will I Lose Money?

Real-life Portfolio Managers ask the question

> What is the risk that I lose money? How much can I lose? What is the chance that I lose even more?

The **Value at Risk (VaR)** concept attempts to answer these questions and provides a way to control the risk of losses.

Let:

Ω be a set of scenarios, indexed by l, having probabilities p^l.

$x = (x_i), i = 1, ..., n$ be the investment,

V_0 be the amount invested

$\tilde{P} = (\tilde{P}_i), i = 1, ..., m$ be the final value of one unit of investment i

$L(x, \tilde{P}) = V_0 - \sum_i x_i \cdot \tilde{P}_i$ be the investment *loss*

$\Psi(x, \zeta)$ be the *probability of losing no more than the amount* ζ

We have (the CDF of losses).

$$\Psi(x, \zeta) = \sum_{\{l \in \Omega | L(x, \tilde{P}) \leq \zeta\}} p^l$$

Definition: Value at Risk:

$$\mathrm{VaR}(x, \alpha) = \min\{\zeta \in \Re \mid \Psi(x, \zeta) \geq \alpha\}.$$

What does this mean?

The General Picture:

An Actual Example

Consider this investment: You invest V_0. Then there's:

96% chance of gaining 20%, 4% chance of losing it all!

Expected Return = 15.2%.

VaR answers the question: What is the largest loss I can suffer, with, say, 5% (α) chance?

- 100%-VaR =V_0
- 96%-VaR = V_0
- 95%-VaR = 0
- 80%-VaR = 0

Coherence

Consider now an enterprise that has 2 *identical, independent* investments like this one: 96% chance of gaining 20%, 4% chance of losing it all! They invest the same amount, V_0, in each.

- 92.16% chance of losing nothing,
- 7.68% chance of losing V_0
- 0.16% chance of losing $2V_0$

Now:

- 100%-VaR $=2V_0$
- 95%-VaR $=V_0$
- 80%-VaR $= 0$

So: Two independent investments that each has a 95%-VaR equal to 0 may have a joint 95%-VaR that is NOT 0!

Problems with VaR

1. VaR is not "coherent".
2. Consider a "Russian roulette" investment: With a very small probability, say 0.0001%,, your company gets wiped out!

   α-VaR $= 0$ for any $\alpha \leq 99.9999\%$ — Doesn't reveal the real risk.

Coherent Risk Measures

Consider two investments, X and Y, with uncertain returns $\tilde{X}$ and $\tilde{Y}$. Let $\rho(\cdot)$ be a risk measure so the risk of X is $\rho(\tilde{X})$ and the risk of Y is $\rho(\tilde{Y})$.

Definition: Coherent Risk Measure $\rho(\cdot)$ is called "coherent" if:

1. Sub-additivity : $\rho(\tilde{X}+\tilde{Y}) \leq \rho(\tilde{X}) + \rho(\tilde{Y})$,

2. Homogeneity : $\rho(\lambda\tilde{X}) = \lambda\rho(\tilde{X})$,

3. Monotonicity : $\rho(\tilde{X}) \leq \rho(\tilde{Y})$ if $\tilde{X} \leq \tilde{Y}$ a.s

4. Risk-free Condition : $\rho(\tilde{X} - nr_f) = \rho(\tilde{X}) - n$

Comments:

1. The risk of joint investments in multiple projects is bounded by the sum of the individual project risks (VaR violates this)
2. A λ times as large investment has a λ times as large risk
3. if X always has a smaller loss than Y, then X has a smaller risk-measure than Y
4. The risk can be reduced by investing in a risk-free asset

Conditional Value at Risk (CVaR)

CVaR addresses the shortcomings of VaR. It answers the question:

If I lose money, then what is my *expected loss*?

So, when chosing an investment, we can control the risk of losing money at all (through VaR), and in addition control the expected loss *when* we lose money (through CVaR).

In addition, CVaR is a coherent risk measure, which allows for *enterprise-wide, distributed* risk management (due to 1. above).

Definition: Conditional Value at Risk:

$$\begin{aligned}
\mathrm{CVaR}(x,\alpha) &= E[L(x,\tilde{P}) \mid L(x,\tilde{P}) > \zeta] \\
&= \frac{\sum_{\{l\in\Omega \mid L(x,P^l)>\zeta\}} p^l \cdot L(x,P^l)}{\sum_{\{l\in\Omega \mid L(x,P^l)>\zeta\}} p^l} \\
&= \frac{\sum_{\{l\in\Omega \mid L(x,P^l)>\zeta\}} p^l \cdot L(x,P^l)}{1-\alpha}
\end{aligned}$$

The last equation follows when the risk is controlled (through VaR) such that $\Psi(x,\zeta) = \alpha$.

Finlib Models:

CVaR.gms, CVaRMIP.gms.

Stochastic Programming

All the models seen so far have been *static models*: A decision is made, then not further modified. They have been essentially *single-period models*, since there's only one decision to be made, for the first period.

Real-life decision processes are more complicated:

1. Although we must make an initial decision now, there will be many opportunities to adjust down the road
2. We do not today have a complete decision basis for future decisions – the *future is unknown.*

Stochastic Programming Models hence are *dynamic*, covering multiple time periods with associated, separate decisions, and they account for the *stochastic decisions process.*

How does it work? The main features of SP are:

1. **Scenarios:** The uncertainty about future events are captured by a set of *scenarios*; a representative and comprehensive set of possible realizations of the future.
2. **Stages:** SP recognizes that future decisions happen in *stages:* A first-stage decision now. Then, after a certain time period, a second-stage decision, which depends upon (1) the first-stage decision, and (2) the events that occurred during the time period. Possible third-, fourth- etc. stage decisions.

The Scenario Tree

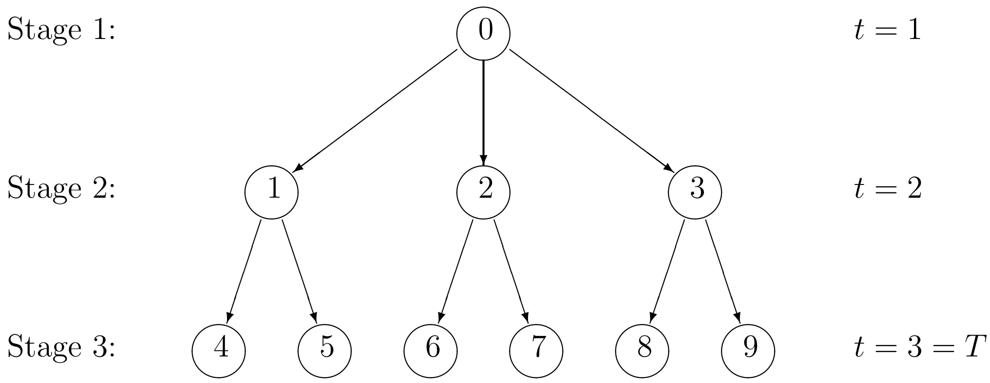

Scenario tree for a 3-stage program ($T = 3$) having 6 scenarios. In this example, ξ_2 has three possible realizations. ξ_3 has two possible realizations for each realization of ξ_2.

The T-stage stochastic program:

$$
\begin{array}{lll}
[\mathbf{MS}] & \min\limits_{x_1} \Big\{ c_1 x_1 + \mathrm{E}_{\xi_2} \Big[\min\limits_{x_2} \Big(\mathbf{c}_2 x_2 + \mathrm{E}_{\xi_3|\xi_2} \Big(\min\limits_{x_3} \mathbf{c}_3 x_3 + \cdots + \mathrm{E}_{\xi_T|\xi_2,\ldots,\xi_{T-1}} \min\limits_{x_T} \mathbf{c}_T x_T \Big) \Big) \Big] \Big\} & \\
 & \text{s.t.} \quad A_1 x_1 & = b_1, \\
 & \quad \mathbf{B}_2 x_1 + \mathbf{A}_2 x_2 & = \mathbf{b}_2, \\
 & \quad \quad \mathbf{B}_3 x_2 + \mathbf{A}_3 x_3 & = \mathbf{b}_3, \\
 & \quad \quad \quad \ddots & \vdots \\
 & \quad \quad \quad \quad \mathbf{B}_T x_{T-1} + \mathbf{A}_T x_T & = \mathbf{b}_T, \\
 & \quad 0 \le x_t \le u_t, \quad \text{for } t = 1, \ldots, T, &
\end{array}
$$

where

$$\xi_t = (\mathbf{A}_t, \mathbf{B}_t, \mathbf{b}_t, \mathbf{c}_t) \text{ for } t = 2, \ldots, T$$

are random variables, i.e., $\mathcal{F}_t$-measurable functions $\xi_{\mathbf{t}} : \Omega_t \mapsto \Re^{M_t}$ on some probability spaces $(\Omega_t, \mathcal{F}_t, P_t)$.

What does it mean? Where do the scenarios come from?

The Winvest Case

The Winvest Case introduces several important concepts:

1. Scenarios: The case consists of 4 interest rate scenarios.
2. Network Modeling: Each scenario is a network problem - very common in financial settings.
3. Stages: The model is a two-stage, stochastic program — or, actually, three-stages.
4. Risk Attitudes: Different objective functions illustrate risk-neutrality, growth-optimality, extreme risk-aversion.

Winvest Data

There's an interesting story about MBSs, Mortgage-Backed Securities[9]

```
 set Cscen /uu, ud, dd, du/;                alias(s, Cscen);
 set Cmbs  /io2, po7, po70, io90/;          alias(i, Cmbs);
 set Ctime /t0, t1, t2/;                    alias(t, Ctime);

table yield(i,t,s)
                       UU          UD          DD          DU
      IO2 .T0    1.104439    1.104439    0.959238    0.959238
      IO2 .T1    1.110009    0.975907    0.935106    1.167817
      PO7 .T0    0.938159    0.938159    1.166825    1.166825
      PO7 .T1    0.933668    1.154590    1.156536    0.903233
      PO70.T0    0.924840    0.924840    1.167546    1.167546
      PO70.T1    0.891527    1.200802    1.141917    0.907837
      IO90.T0    1.107461    1.107461    0.908728    0.908728
      IO90.T1    1.105168    0.925925    0.877669    1.187143 ;

table cash_yield(t, s)

                       UU          UD          DD          DU
           T0    1.030414    1.030414    1.012735    1.012735
           T1    1.032623    1.014298    1.009788    1.030481 ;

 table liab(t,s)
                       UU          UD          DD          DU
           T1   26.474340   26.474340   10.953843   10.953843
           T2   31.264791   26.044541   10.757200   13.608207 ;

 parameter val(s)
   / uu = 47.284751, ud = 49.094838, dd = 86.111238, du = 83.290085/;

 parameter transcost; transcost = 0.01
```

[9]Nielsen and Zenios: A Stochastic Programming Model for funding Single Premium Deferred Annuities.

Network Scenarios

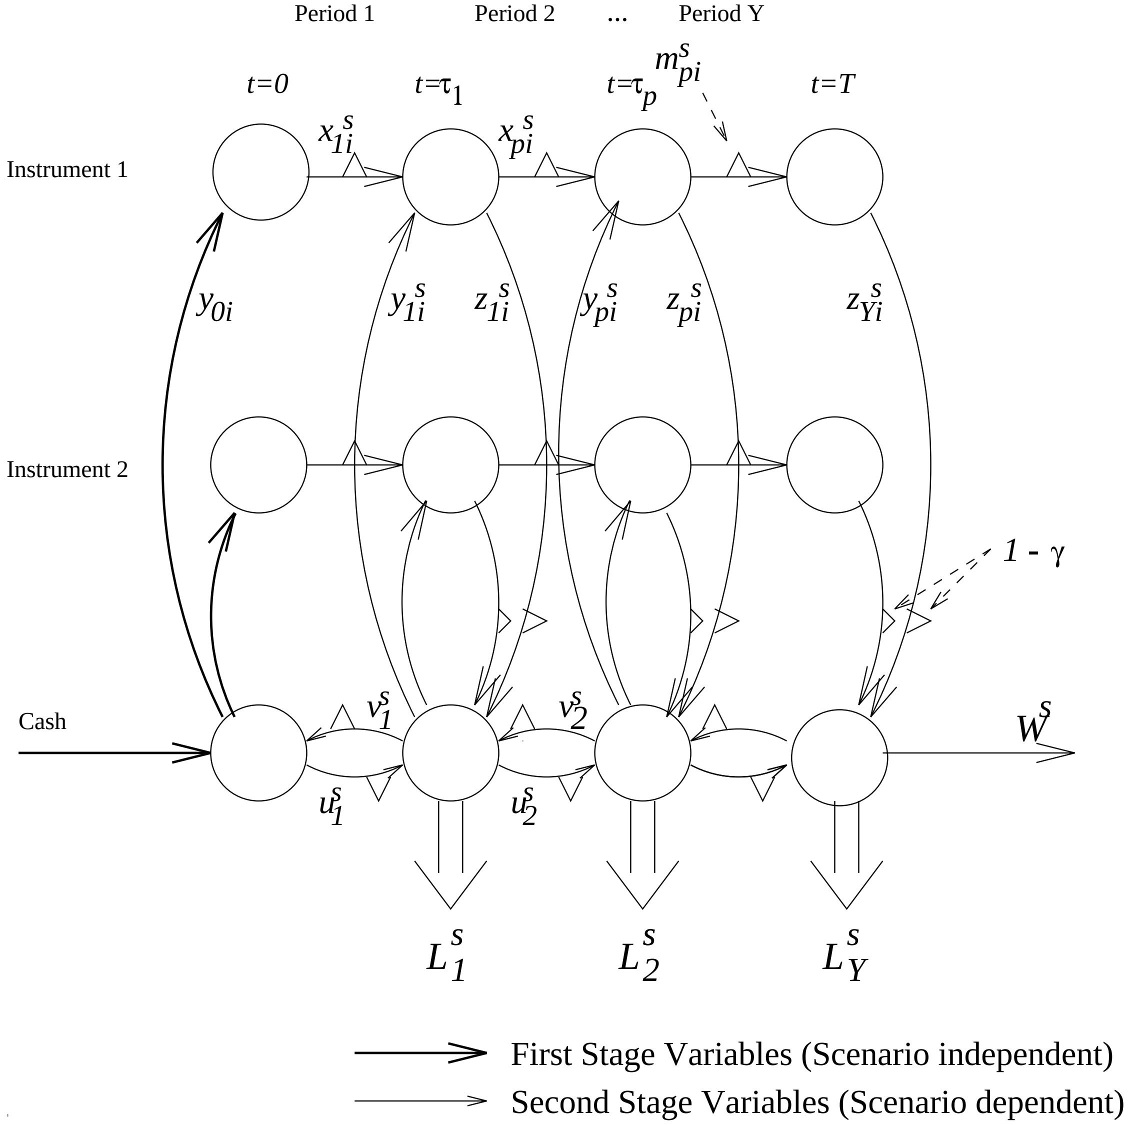

Network model underlying the two-stage, stochastic network model. This figure includes two instruments, and depicts a 4-period model. Stochastic quantities are denoted by a superscript, s.

The Algebra Behind the Network

1. *Arcs* are decisions: How big is the flow?
2. *Nodes* are constraints: Flow in = Flow out.
3. *Multipliers* indicate gains/losses: One dollar invested (entering an arc) may be more or less when realized (exiting the arc).
4. *Scenarios:* There's only one set of first-stage variables, but second-stage (and 3rd, ...) variables for each scenario

Recommendation: Model a single scenario first. Then add (1) a scenario index to all data, variables, and constraints, and (2) *non-anticipativity constraints* to force first-stage variables to assume the same values across scenarios, etc.

Winvest: First Steps

1. Copy the winvest.gms file. Erase everything after the first 42 lines!
2. Define decision variables corresponding to each type of arcs.
3. Define constraints corresponding to each node (there are two types: Instruments and cash). Model a single scenario. When referencing data, use "UU" as the scenario index.
4. Maximize the final wealth, on the W arc. For now, *ignore* liabilities and transactions costs.
5. Verify that the model returns the optimal solution
6. Now introduce transaction costs: Every time a security is sold (not when purchased), we only receive 99% of its value. The rest vanishes as transaction costs. Does this change the optimal solution?

Note: The GAMS \$-operator comes in handy here, to get an elegant, compact GAMS formulation. We also need to understand the *ord* and *card* functions.

Winvest: Making it an SP

A few small modifications will make the model a true Stochastic Program:

1. Include a scenario index, s on all variables and constraints. Use it when referencing data instead of "UU".
2. We now have a final wealth variable for *each scenario*, $W(s)$. What do we optimize?

   One possibility: Define the expected final wealth,

   $$EW = \sum_{s \in S} p^s W(s)$$

   and maximize that. In Winvest, all $p^s = 1/4$.
3. Run the model; verify that each scenario is optimal.
4. BUT - the first-stage decisions are all different! *We still don't know how to invest!*

   Add constraints that force the first-stage decisions to be the same across scenarios (or, equivalently, remove the scenario index from the first-stage variables. This may be more complicated to manage).
5. To be strict about it, the second-stage decisions must also agree between the UU and the UD scenario (because all we know at the second stage is that interest rates went up - we don't yet know what they do next). Similarly for the DU and DD scenarios. Add the appropriate non-anticipativity constraints. Now we have a true three-stage SP!
6. The objective suggested above corresponds to a risk-neutral attitude. Would a real-life investor feel happy about the optimal solution? Or is the worst outcome just a little too bad for comfort?

Working with Risk-Attitudes

1. The worst outcome in the risk-neutral case is pretty bad. Implement an objective that *maximizes the final wealth under the worst outcome.*

   It's convenient to introduce a variable, Worst, whose value is equal to the worst of the 4 outcomes, then maximize it.

   Hint: This variable has a simple property:

   $$Worst \leq W(s); \text{ for all } s \in S$$

   What happens to the expected final wealth? Is the worst outcome better? What happens to the initial portfolio — for instance, diversification?

2. Finally, try a little utility theory. Find the portfolio that maximizes the expected utility, using $U(x) = \log x$.

Two-Stage SP: Formal Model

Modellerer en to-trins beslutningsprocess:

- $t = 0$: **Stage One:** A decision is made today, accounting for future uncertainty.
- $t = \tau$: **Stage Two:** A new, corrective (recourse) decison. The decision at time τ is conditional upon (1) the initial decision, and (2) the observed realization of uncertainties between $t = 0$ and τ.

Matematical Formulation:

[**SNLP**]

$$\begin{array}{ll} \text{Minimize} & f(x) + \mathcal{Q}(x) \\ \text{Subject to} & Ax \quad = b \\ & 0 \leq x \leq u^x \end{array}$$

where

$$\mathcal{Q}(x) = \ \mathrm{E}\{Q(\mathbf{g}, \mathbf{r}, \mathbf{B}, \mathbf{v}, \mathbf{C} \mid x)\},$$

$$Q(g, r, B, v, C \mid x) = \quad \begin{array}{ll} \text{Minimize} & g(y) \\ \text{Subject to} & By = r - Cx \\ & 0 \leq \text{y} \ \leq v \end{array}$$

- x, y: First, resp. second stage decision
- The objective minimizes the first-stage (deterministic) costs, plus the expected second-stage (stochastic) costs, assuming that y is chosen optimally, given x and the realization of uncertainties.

Scenario-based formulation

In practice, the uncertainty in SP is represented by *scenarios*. A scenario $s \in \S$ is a complete, joint realization of all stochastic parameters, an *event*.

Then SP can be formulated (Wets, 1974) by its *deterministic equivalent:*

$$
\begin{array}{llll}
\underset{x \in \mathbb{R}^{n_0}, y^s \in \mathbb{R}^{n_1}}{\text{Minimize}} & f(x) + \sum_{s=1}^{S} p^s g^s(y^s) & & \\
\text{Subject to} & Ax & = & b \\
 & C^s x + B^s y^s & = & r^s \text{ for all } s \in <S> \\
 & 0 \leq x & \leq & u^x \\
 & 0 \leq y^s & \leq & v^s \text{ for all } s \in <S>
\end{array}
$$

Observation: For a given, fixed first-stage decision x, the problem decomposes into $\mid S \mid$ separate, independent scenario-problems. This forms the basis for (1) the "split-variable" formulation and (2) most specialized algorithms for SP.

The Scenario problem:

Given $\hat{x}$ which satisfies $A\hat{x} = b, 0 \leq \hat{x} \leq u^x$ the problem for scenario s is:

$$
\begin{array}{llll}
\text{Minimize} & g^s(y^s) & & \\
\text{Subject to} & B^s y^s & = & r^s - C^s \hat{x} \\
 & 0 \leq y^s & \leq & v^s
\end{array}
$$

C^s is the *tecnology matrix*, which "transmits" information about the first-stage decison to the scenario problem.

The Split-Variable Formulation

Original 3-stage scenario tree:

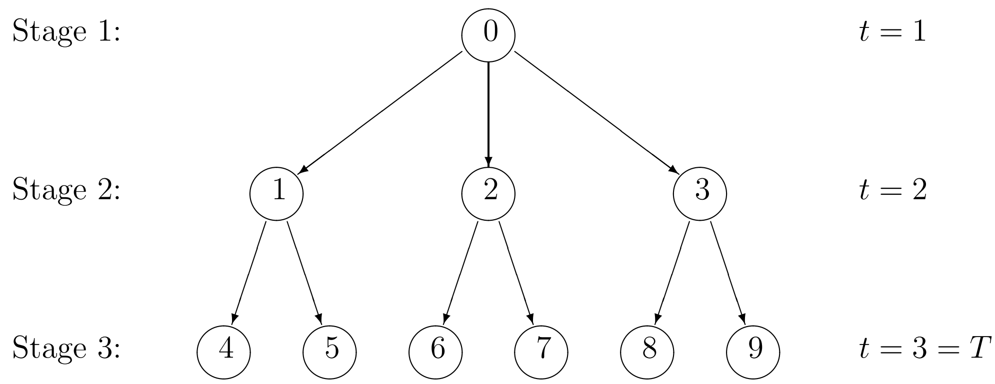

"Split" tree:

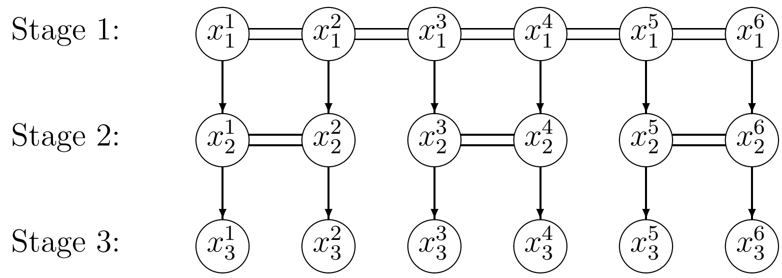

- Make a copy of *all* variables for each scenario
- Add "non-anticipativity constraints" to force logically identical variables (all but last stage) to agree across scenarios.
- Winvest had non-ancitipativity constraints to enforce agreement among all $t = 0$ variables across all scenarios, and among $t = 6$ months variables between (UU, UD) and between (DU, DD) scenarios.

The Split-Variable Formulation

The deterministic equivalent 2-stage problem has the following split-variable formulation:

$$
\begin{array}{lllll}
\underset{x^s \in \mathbb{R}^{n_0}, y^s \in \mathbb{R}^{n_1}}{\text{Minimize}} & \sum_{s=1}^{S} p^s (f(x^s) + g^s(y^s)) & & & \\
\text{Subject to} & A x^s & = & b & \text{for all } s \in <S> \\
 & C^s x^s + B^s y^s & = & r^s & \text{for all } s \in <S> \\
 & 0 \le x^s & \le & u^s & \text{for all } s \in <S> \\
 & x^1 & = & x^s & \text{for all } s \in <S> \\
 & 0 \le y^s & \le & v^s & \text{for all } s \in <S>
\end{array}
$$

– compare to the original problem: –

$$
\begin{array}{lllll}
\underset{x \in \mathbb{R}^{n_0}, y^s \in \mathbb{R}^{n_1}}{\text{Minimize}} & f(x) + \sum_{s=1}^{S} p^s g^s(y^s) & & & \\
\text{Subject to} & A x & = & b & \\
 & C^s x + B^s y^s & = & r^s & \text{for all } s \in <S> \\
 & 0 \le x & \le & u^x & \\
 & 0 \le y^s & \le & v^s & \text{for all } s \in <S>
\end{array}
$$

Binomial Lattices for One-factor Models

The Cox-Ross-Rubinstein (CRR) model of stock prices discretizes the continuous-time process[10]

$$dS = \mu S dt + \sigma S dz$$

where:

1. μ is the *drift*, or short-term expected return,
2. σ is the *volatility*,
3. dz is a Wiener process.

The discrete process exists only at times $t = 0, \Delta t, 2\Delta t, ..., n\Delta t$. Starting from the stock prices S_0 at $t = 0$, the process at time t moves to

$$S_{t+1} = S_t \cdot u$$

or to

$$S_{t+1} = S_t \cdot d$$

with probability p and $1 - p$, respectively, where

1. $u = \exp(\sigma\sqrt{\Delta t})$
2. $d = \exp(-\sigma\sqrt{\Delta t})$
3. $p = \frac{\exp(r\Delta t) - d}{u - d}$
4. r is the risk-free interest rate

These numbers are chosen so that the continuous-time and the discrete-time processes have the same first- and second-order statistics (means and variances).

A popular model for interest rates leading to a similar binomial lattice is the Black-Derman-Toy model.

[10]See Hull: Options, Futures and other Derivatives; Chapters 10 and 15

Generating Scenarios:

1. Sample the process directly, in discrete time:

$$S_{t+1} = \mu_t \cdot S_t \cdot \Delta t + \sigma_t \cdot S_t \cdot \sqrt{\Delta t} \cdot \epsilon_t,$$

   ϵ_t standard normal;

2. Sample (randomly) the binomial tree (next page)

3. Importance sampling; Nielsen [1995]

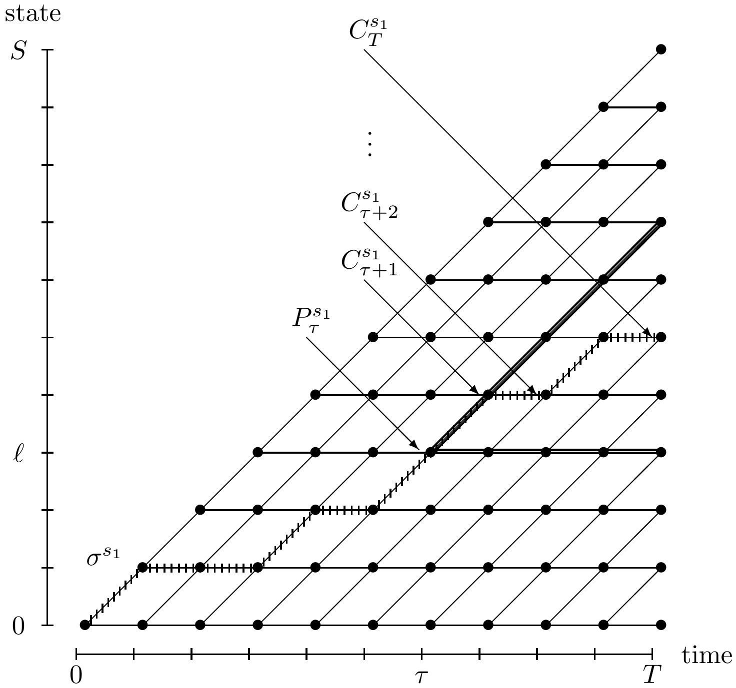
state
S
ℓ
0
$C_T^{s_1}$
$\vdots$
$C_{\tau+2}^{s_1}$
$C_{\tau+1}^{s_1}$
$P_\tau^{s_1}$
σ^{s_1}
time
0
τ
T

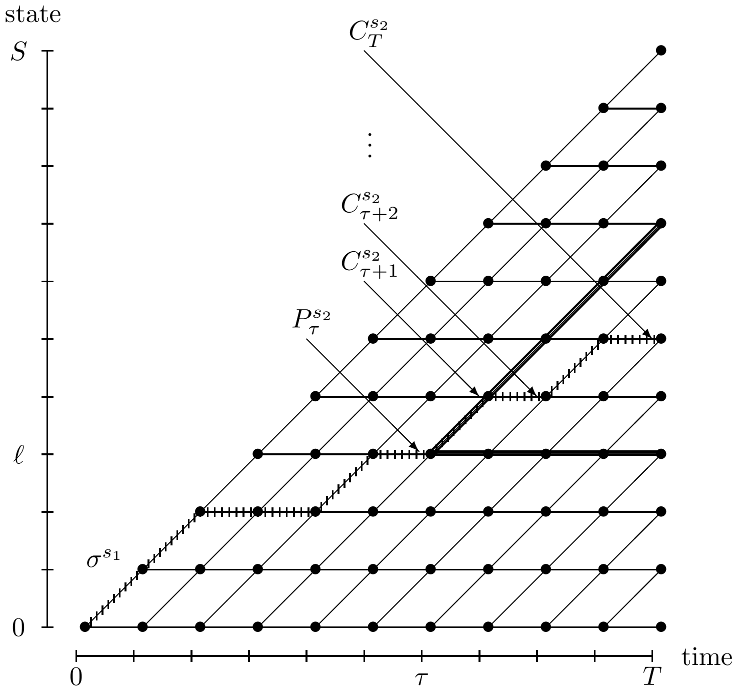
state
S
ℓ
0
$C_T^{s_2}$
$\vdots$
$C_{\tau+2}^{s_2}$
$C_{\tau+1}^{s_2}$
$P_\tau^{s_2}$
σ^{s_1}
time
0
τ
T

Interest-rate Scenarios (Nielsen-Zenios SPDA paper)

Spot Rates, Percent

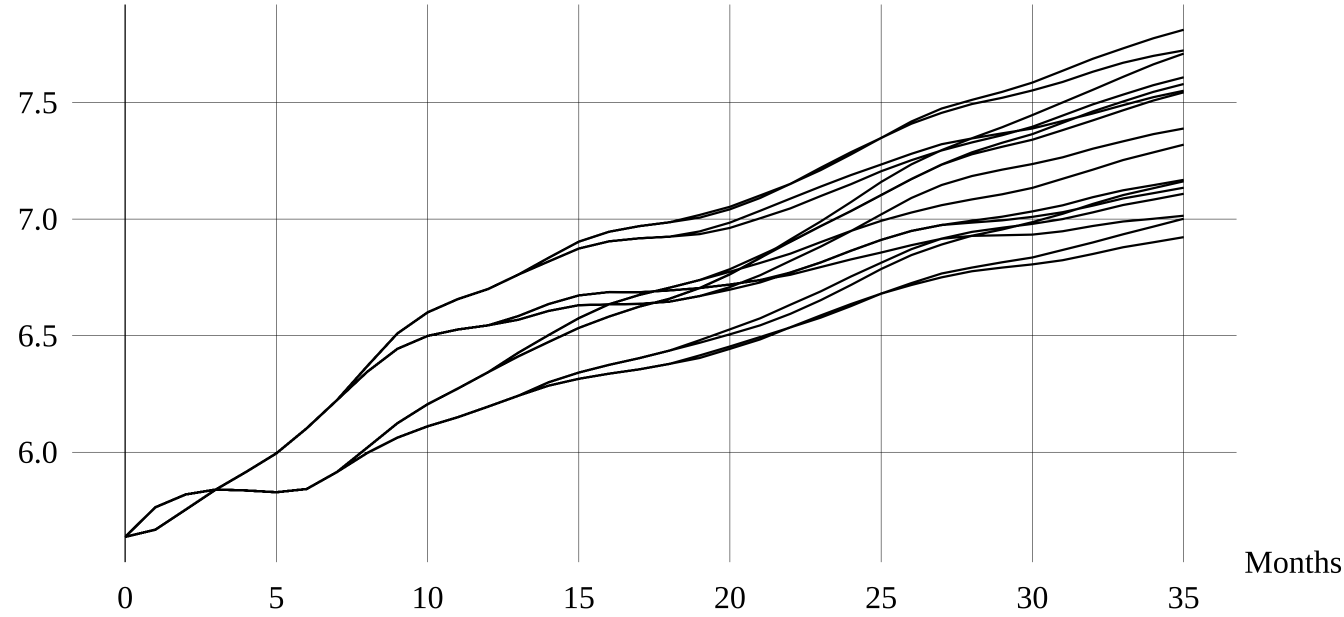

Spot Rates, Percent