A DYNAMIC MODEL FOR RISK PRICING IN GENERATION INVESTMENTS

John Parsons, MIT Sloan School of Management and
MIT Center for Energy and Environmental Policy Research

Fernando de Sisternes, MIT Engineering Systems Division

Seminar of the Chaire European Electricity Markets (CEEM),
Université Paris-Dauphine

March 21, 2014

Outline

- Motivation and Perspective
- The Big Picture – sketch of the overall plan and the main results.
- Detail
    - #1. The Market Price of Risk – Using a Stochastic Discount Factor.
    - #2. The Electricity Price Model.
    - #3. Executing the Valuation of an Asset.
- Discussion

Motivation and Perspective

Current Approaches to Risk Valuation

- The past number of years has seen significant interest in the role of risk in the valuation of electricity generating technologies.
- One approach leans on the now widespread availability of computing to generate large Monte Carlo distributions of payoffs to different assets or for the same asset financed with different contract.
    - Usually the different distributions are compared on the basis of means and variances. For example, fixing the mean, a distribution with a higher variance is considered worse than a distribution with a lower variance.

Shortcoming

- One shortcoming of this approach is its failure to connect with the standard tools of modern valuation and asset pricing.
  - #1 This approach ignores the key insight from portfolio theory that expected return is not a function of total variance, but rather of the component of variance that is correlated to macroeconomic variables.
  - #2 It also ignores the key insight from derivative pricing that variance in the final payoff is a poor tool for ranking risk.
    - The non-linearity of many payoffs makes the problem more difficult than is acknowledged in a simple mean-variance framework.

- This disconnect undermines the reliability of many conclusions drawn from these Monte Carlo simulations, and it undermines the confidence we might have in the specific values calculated using the simulations.

Our Contribution

- We show how to incorporate standard risk pricing principles into the popular Monte Carlo simulation analysis.

- Our methodology has many conservative advantages.
  - The foundation is identical with core principles of valuation and asset pricing.
  - The structure is a transparent generalization of traditional DCF.
  - The structure is consistent with widely applied Monte Carlo approaches.

- Our methodology has one key radical advantage.
  - It makes explicit demands on the modeler to be precise about the critical elements of risk and the price of risk.
    - "Whereof one cannot speak, thereof one must be silent." – Ludwig Wittgenstein, *Tractatus Logico-Philosophicus*

The Price and Value of Risk

"A cynic is a man who knows the price of everything, and the value of nothing."
— Oscar Wilde, *Lady Windermere's Fan*

- Two distinct questions about an asset's risk:
  - What is the price of risk? …the market price.
  - What is the value of risk? …to our company, in particular
- The CAPM and other asset pricing models are all about the first question.
  - Total risk is never the right variable. Only non-diversifiable risk matters.
  - The market price of the cash flows from an asset are independent of who owns the asset. There are no portfolio gains to be had.
  - Hedging is a zero NPV action. Risk is bought and sold at a fair price.
- Theories of hedging are all about the second question.
  - Total risk can matter. Diversifiable risk matters. For some companies.
  - Where risk lies can matter. Capital markets have friction, and more risk means more encounters with those frictions feeding back to cash flow.
- This paper tackles the first question only, the market price of risk.
  - Leave it to later to address the value of risk to a specific company.

The Big Picture

Step #1: The Market Price of Risk

- We employ the classic single factor model with the returns to a diversified stock portfolio as the underlying priced risk factor.
  - We assume the returns evolve as a random walk—arithmetic Brownian motion.

- This implies a set of stochastic discount factors that can be used to value cash flows received contingent on the different states of the market.
  - This is essentially what lies behind the Black-Scholes-Merton derivative pricing formulas.

- This is also the model that validates the CAPM and standard risk-adjusted discounting formulas for a subset of problems—for any asset with a risk structure linear in the market risk factor and where risk grows linearly with time.

- This is the simplest model. One could get fancy and use a different model of the underlying market risk factor or complicate it with multiple factors.

Step #2: Overlay a Model of Electricity Price Risk

- Assume that electricity demand growth is stochastic, but correlated with returns on the stock market.

- Yields a stochastic electricity price as follows…
  - surprisingly large demand growth leads to electricity price increases, but,
  - capacity additions cap the price;
  - drops in demand cause a drop in the electricity price.

- The correlation between demand growth and the underlying risk factor is translated into a correlation between the electricity price and the underlying market risk factor.
  - But the translation is not linear, neither in a single period, nor through time.
  - Cash flows tied to the price of electricity inherit some priced market risk.

Step #3: Value Electricity Assets

- Model the stochastic cash flows using an explicit state space approach.
  - The defining feature of a state is the underlying risk factor.
  - Each state has a unique discount factor – the stochastic discount factor.
  - Determine the expected cash flows for each state.

- Value the cash flows using
  - (1) the probability of the state,
  - (2) the unique discount factor for market risk in that state, and
  - (3) discount for the time value of money.

- Practitioners are familiar with summarizing risk via variances and co-variances. – e.g., Beta.
  - Training often fails to warn about the dangers and problems with these summary statistics.
  - Our method does not rely on faulty summary statistics. It's back to basics.

Demo #1: An Electricity Price Derivative

- Consider an electricity swap that pays the floating price of one unit of electricity every year…
- We can calculate the expected annual cash flow: $21.
- We can calculate the value using the stochastic discount factors associated with the market risk factor: $432.
- We can back out an implied average risk-adjusted discount rate: r_e = 4.9%
  - r_e is not fixed; it depends upon whether the initial electricity price is high or low.

The Risk Premium on an Electricity Price Derivative

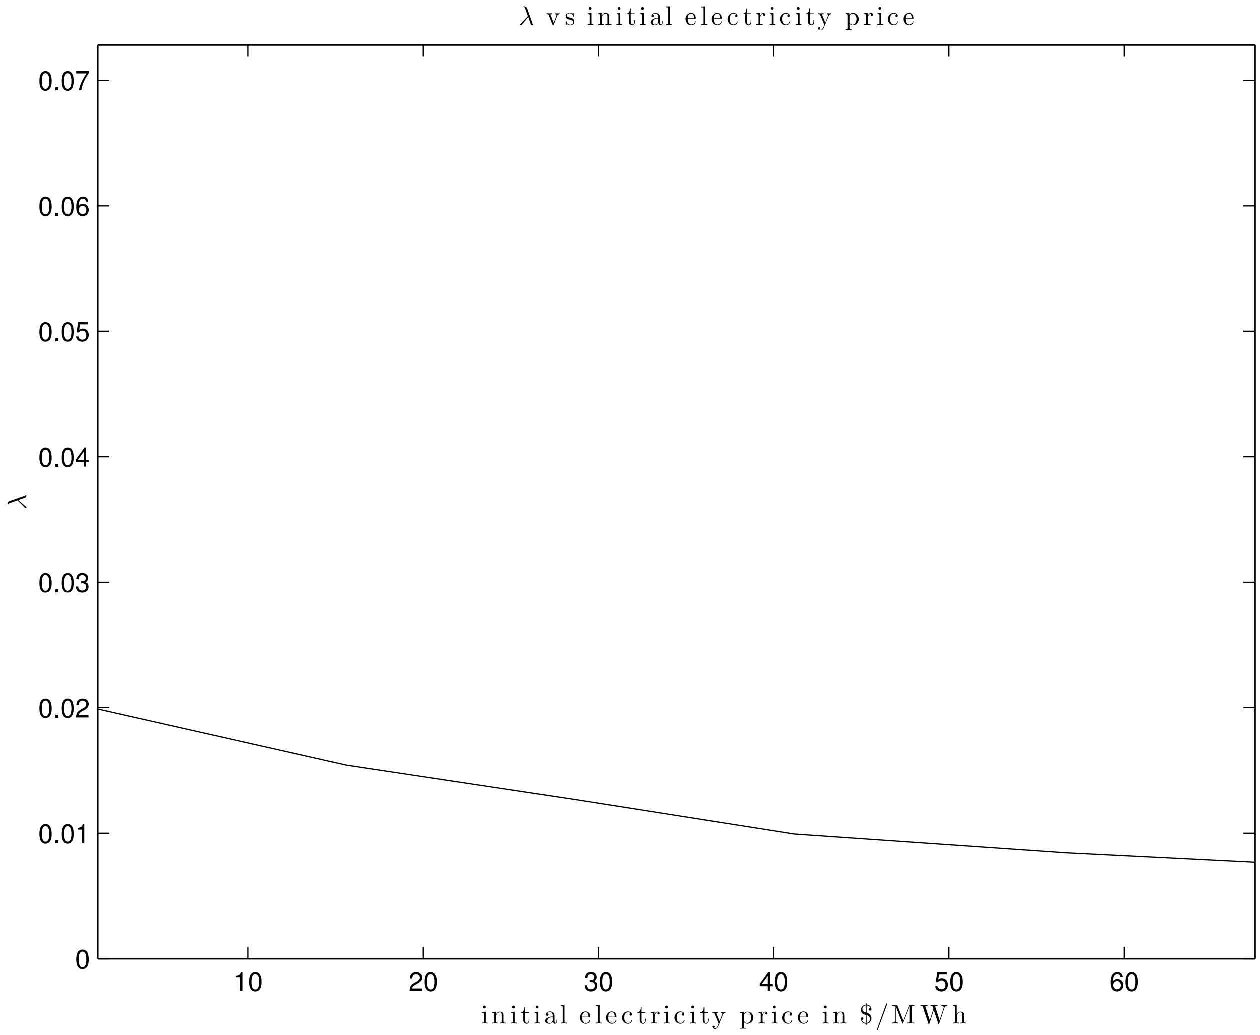

Demo #2: An Electricity Generation Plant

- Suppose we have an already installed generation plant.
- Model how its production varies with the electricity price.
  - Simplest version: produce whenever price is above marginal cost.
- We can calculate the expected annual cash flow: $20.
- We can calculate the value using the stochastic discount factors associated with the market risk factor: $317
- We can back out an implied average risk-adjusted discount rate: r_e = 6.3%
  - r_e is not fixed; it depends upon whether the initial electricity price is high or low; it depends upon where we are in the life of the plant.

The Risk Premium on an Electricity Generation Plant

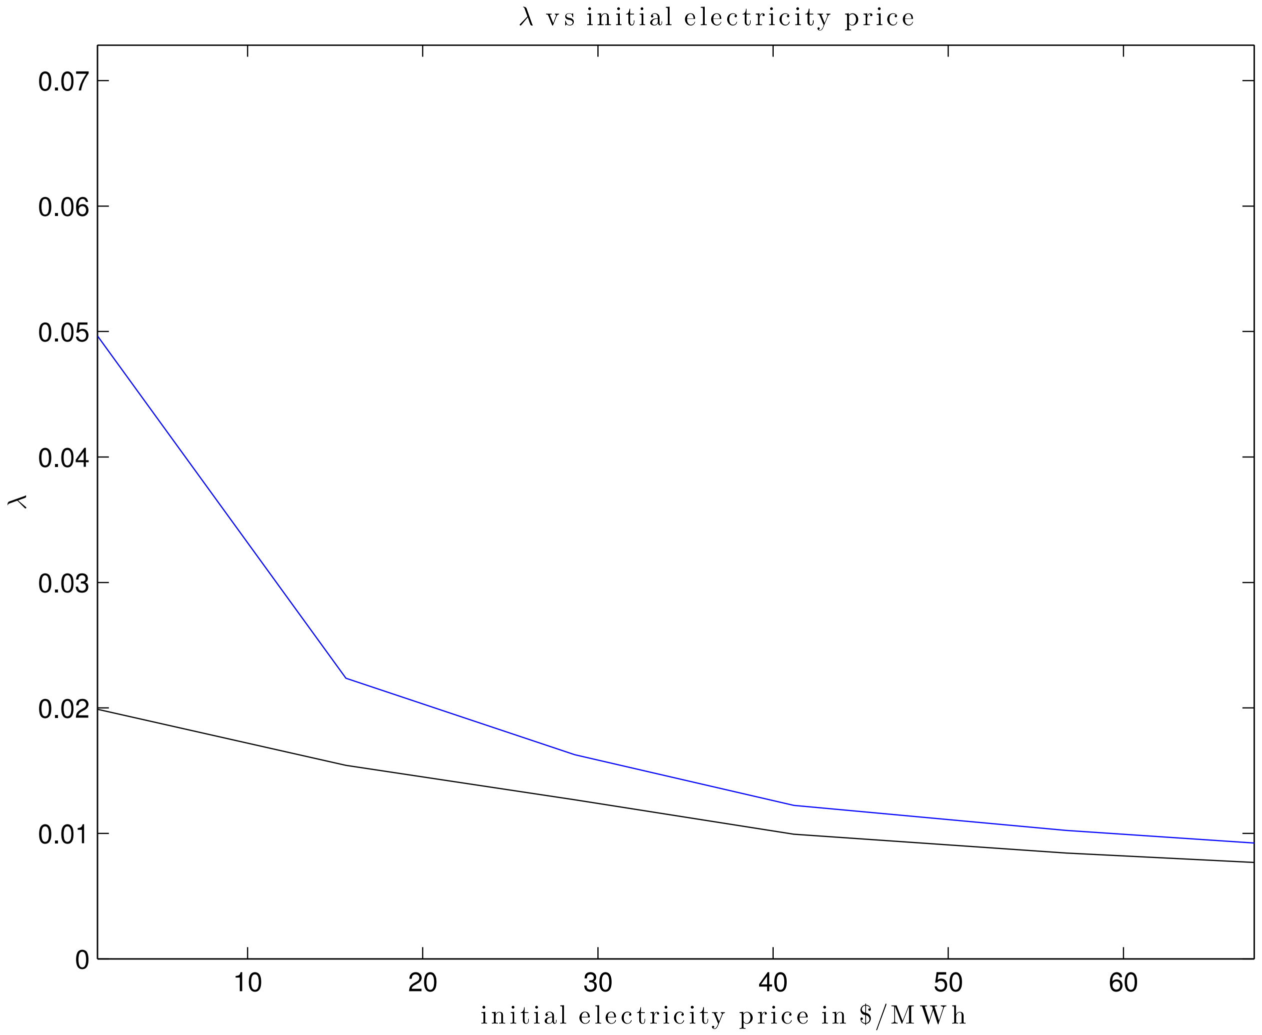

Demo #3: A Different Electricity Generation Plant

- Consider a plant with a higher marginal cost of operation.
  - Call it a peaker.
- Is the implied average risk-adjusted discount rate higher or lower than for the base-load plant?
- We can calculate the expected annual cash flow: \$31.
- We can calculate the value using the stochastic discount factors associated with the market risk factor: \$375.
- We can <u>back out</u> an implied <u>average</u> risk-adjusted discount rate: r_e = 8.3%
  - r_e is not fixed; it depends upon whether the initial electricity price is high or low; it depends upon where we are in the life of the plant.
- The peaker is riskier, in the sense of market price of risk.

The Risk Premium on a Peaker Plant

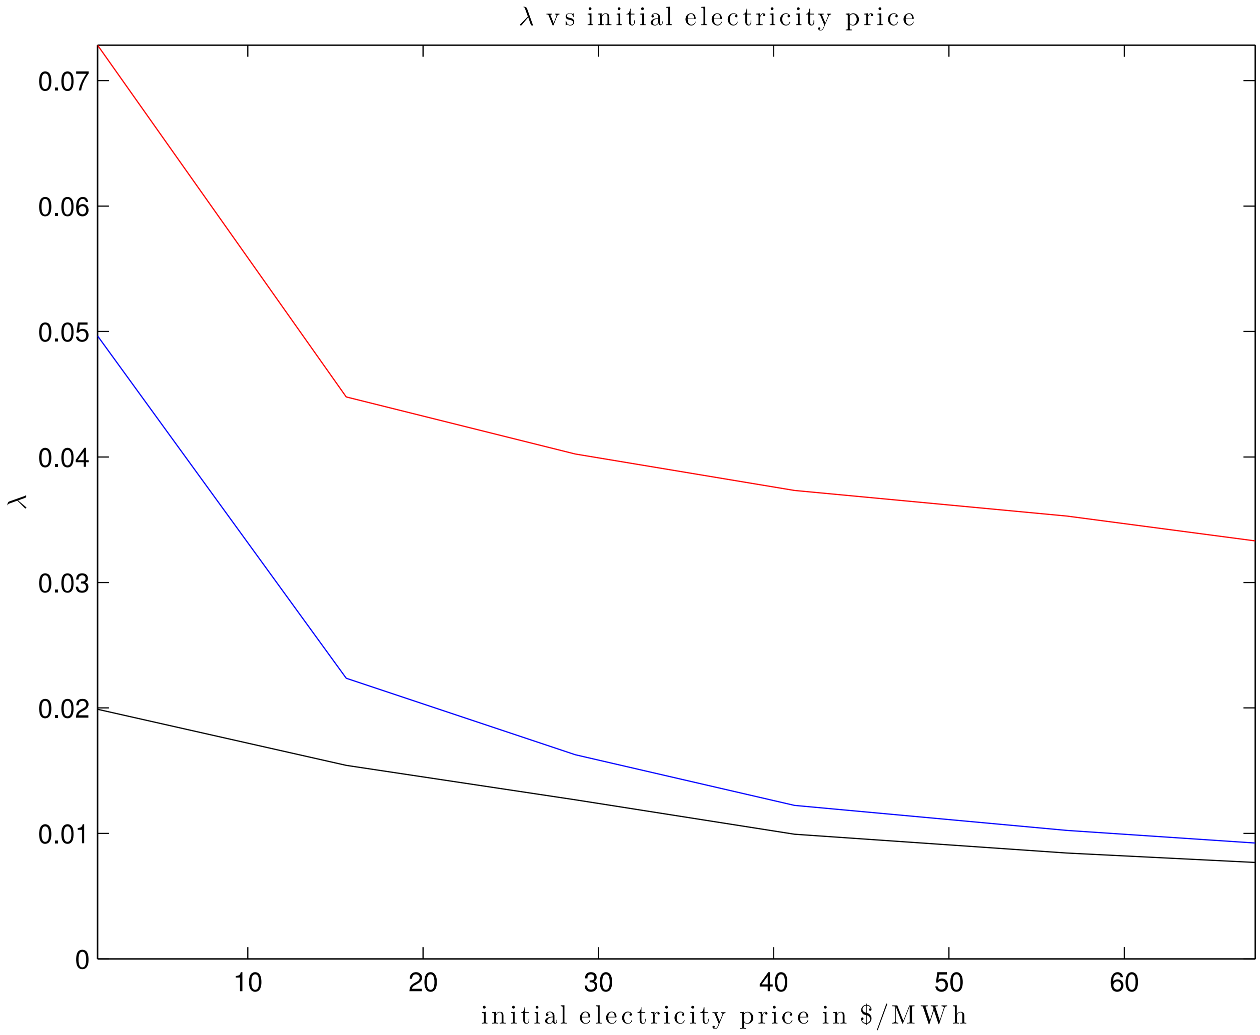

Detail #1:
The Market Price of Risk – Using a Stochastic Discount Factor

A Reminder of What's Behind a Risk-Adjusted Discount Rate: The Stochastic Discount Factor

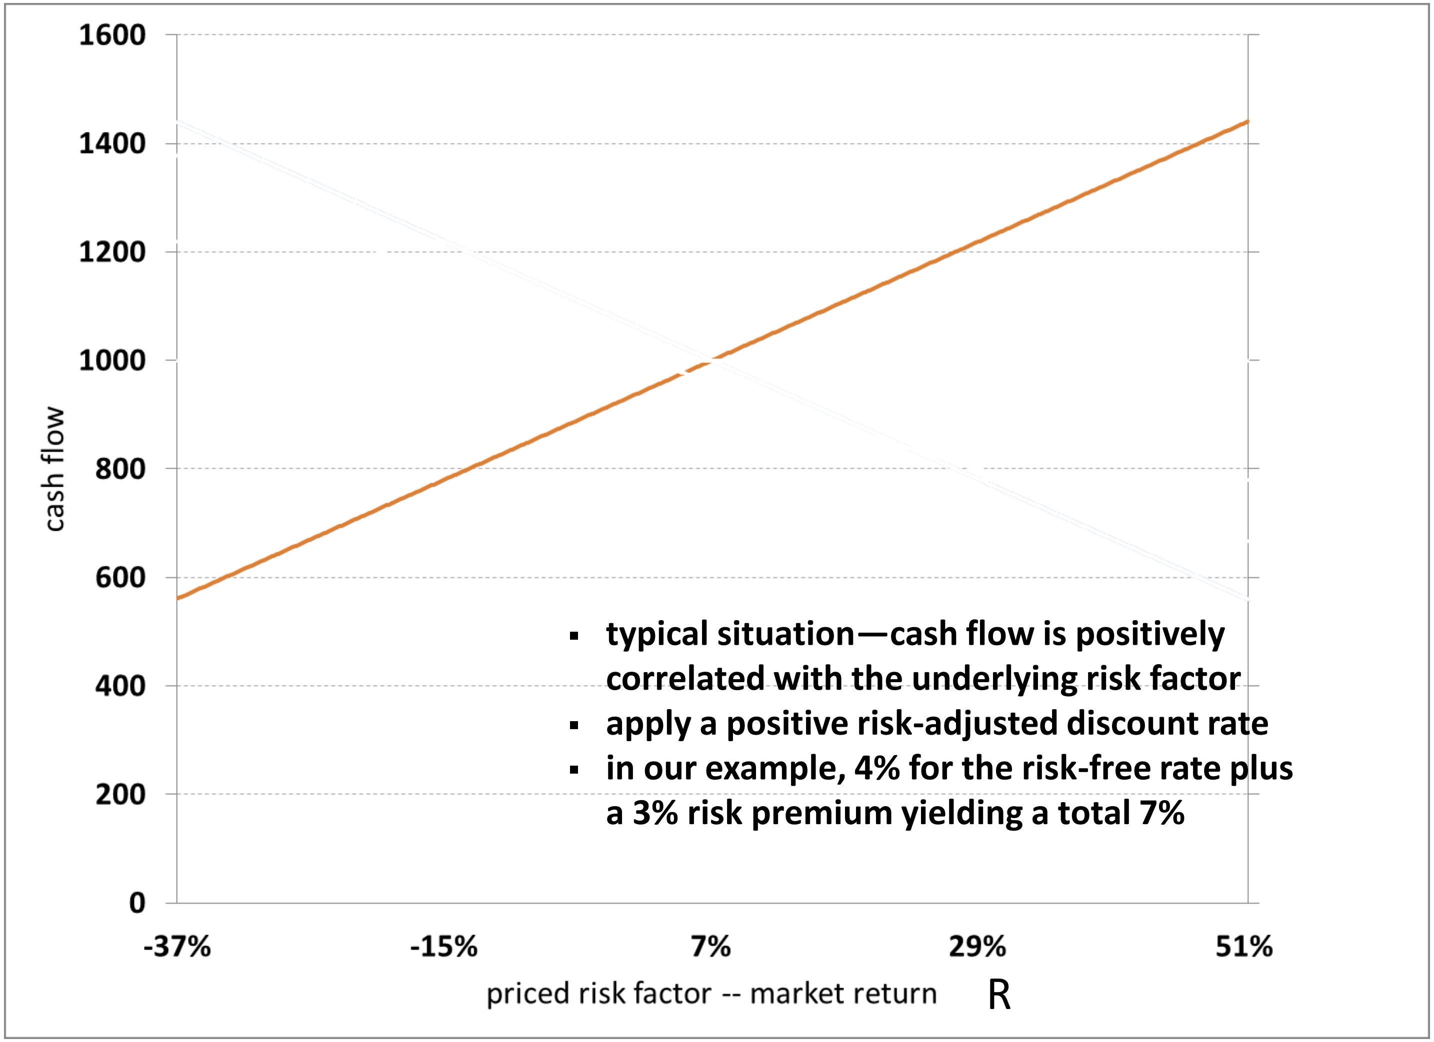

- typical situation—cash flow is positively correlated with the underlying risk factor
- apply a positive risk-adjusted discount rate
- in our example, 4% for the risk-free rate plus a 3% risk premium yielding a total 7%

A Reminder of What's Behind a Risk-Adjusted Discount Rate: The Stochastic Discount Factor (2)

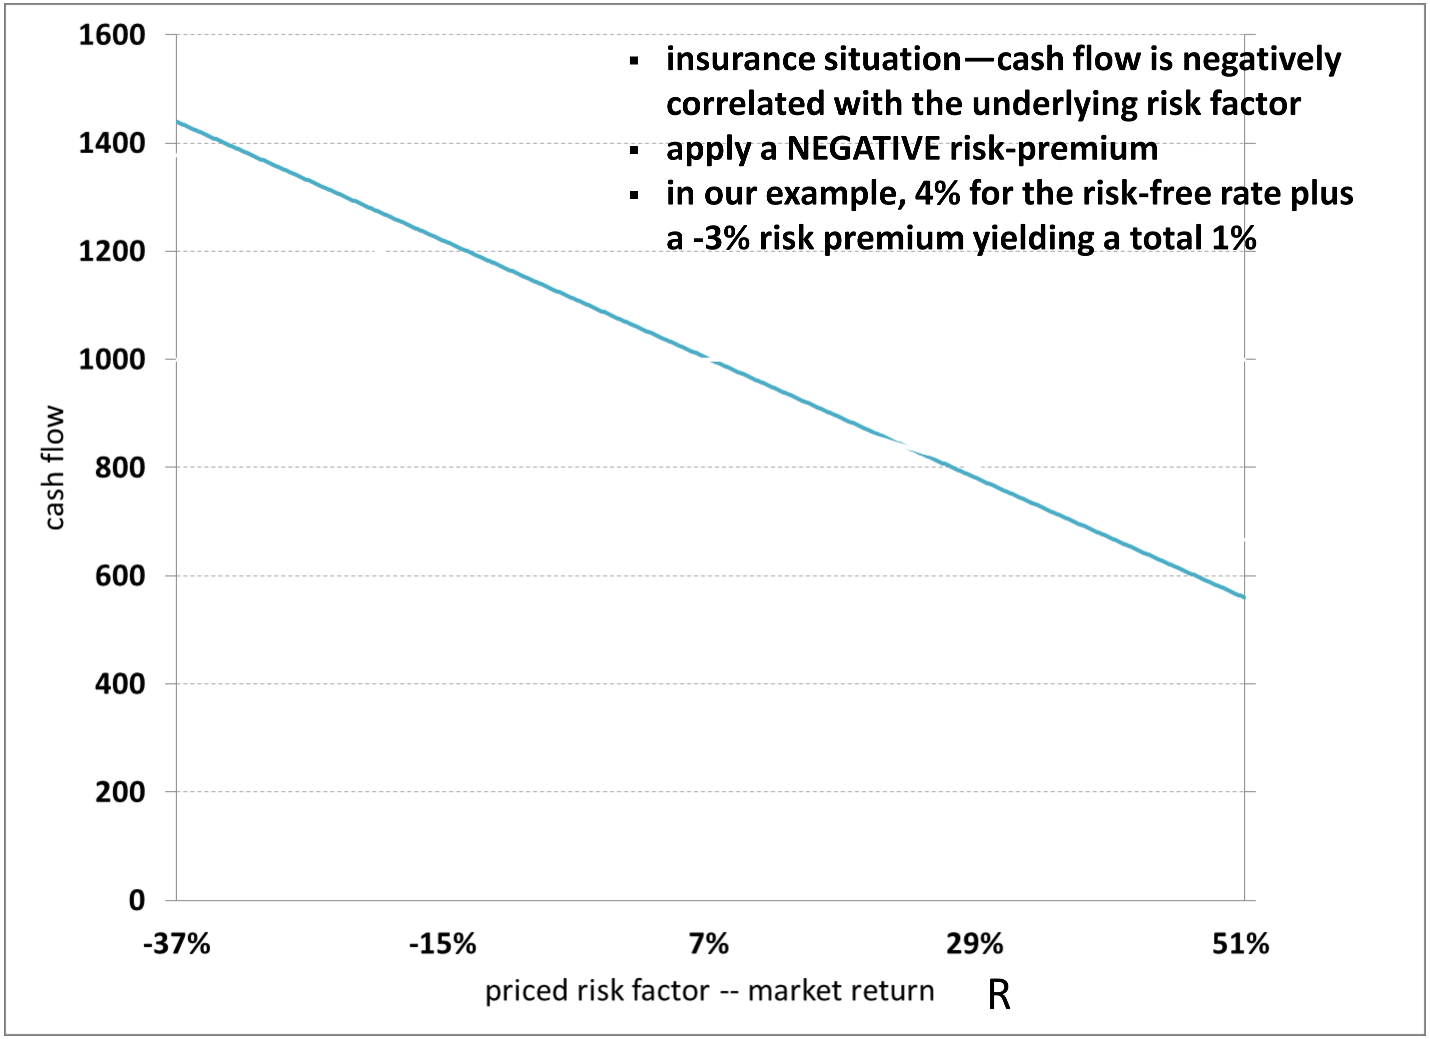

- insurance situation—cash flow is negatively correlated with the underlying risk factor
- apply a NEGATIVE risk-premium
- in our example, 4% for the risk-free rate plus a -3% risk premium yielding a total 1%

A Reminder of What's Behind a Risk-Adjusted Discount Rate: The Stochastic Discount Factor (3)

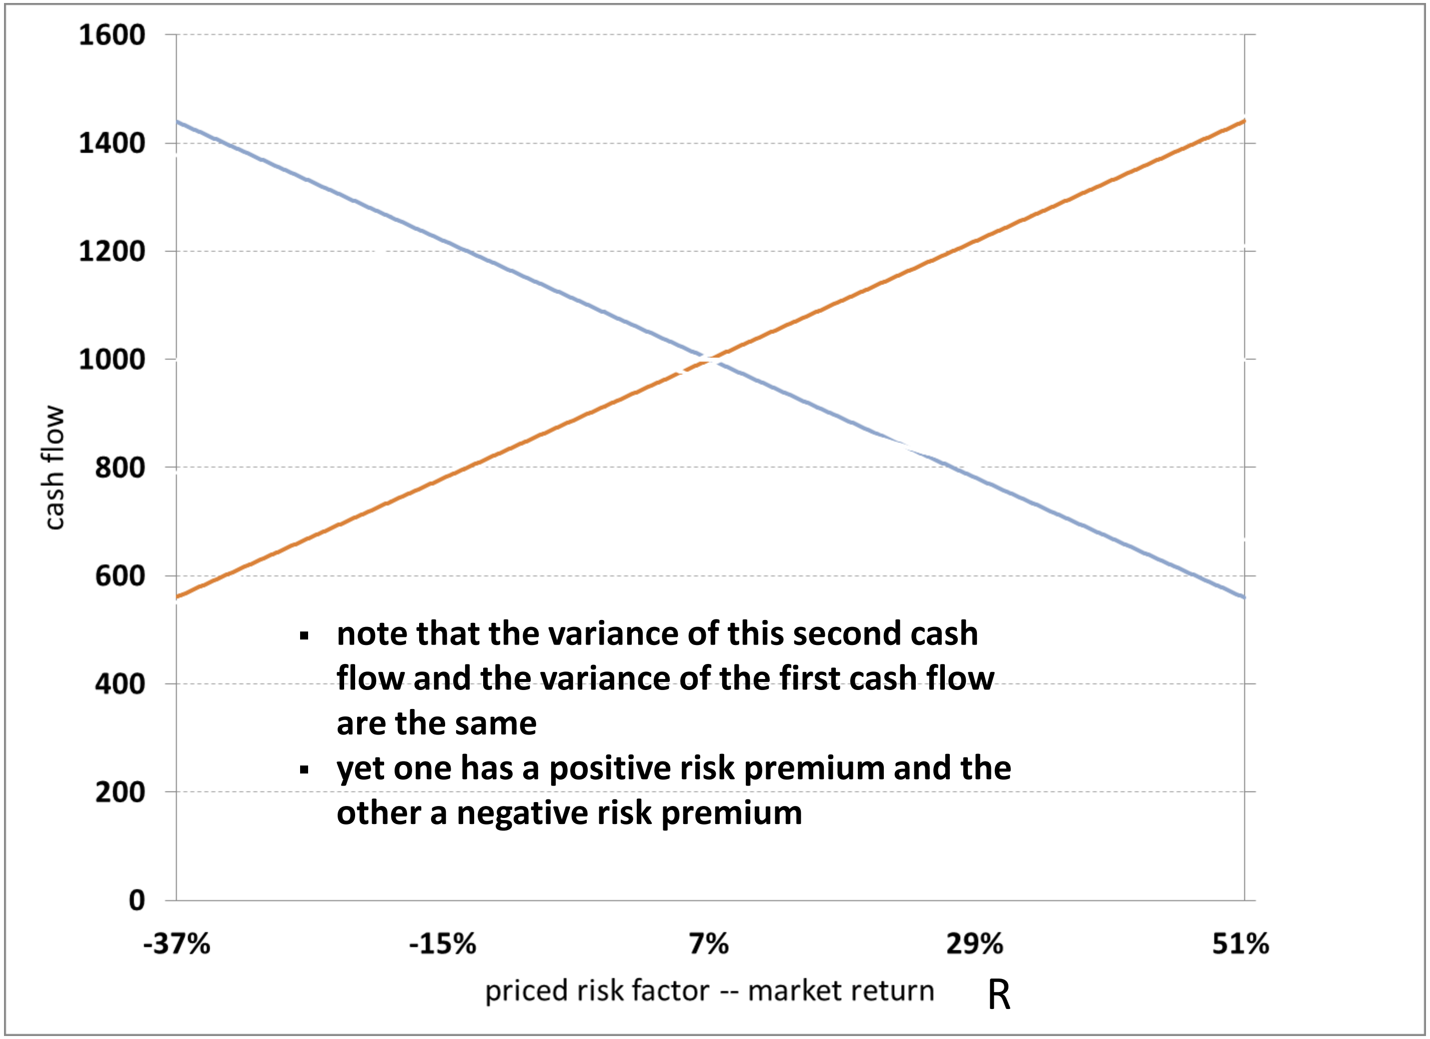

- note that the variance of this second cash flow and the variance of the first cash flow are the same
- yet one has a positive risk premium and the other a negative risk premium

A Reminder of What's Behind a Risk-Adjusted Discount Rate: The Stochastic Discount Factor (4)

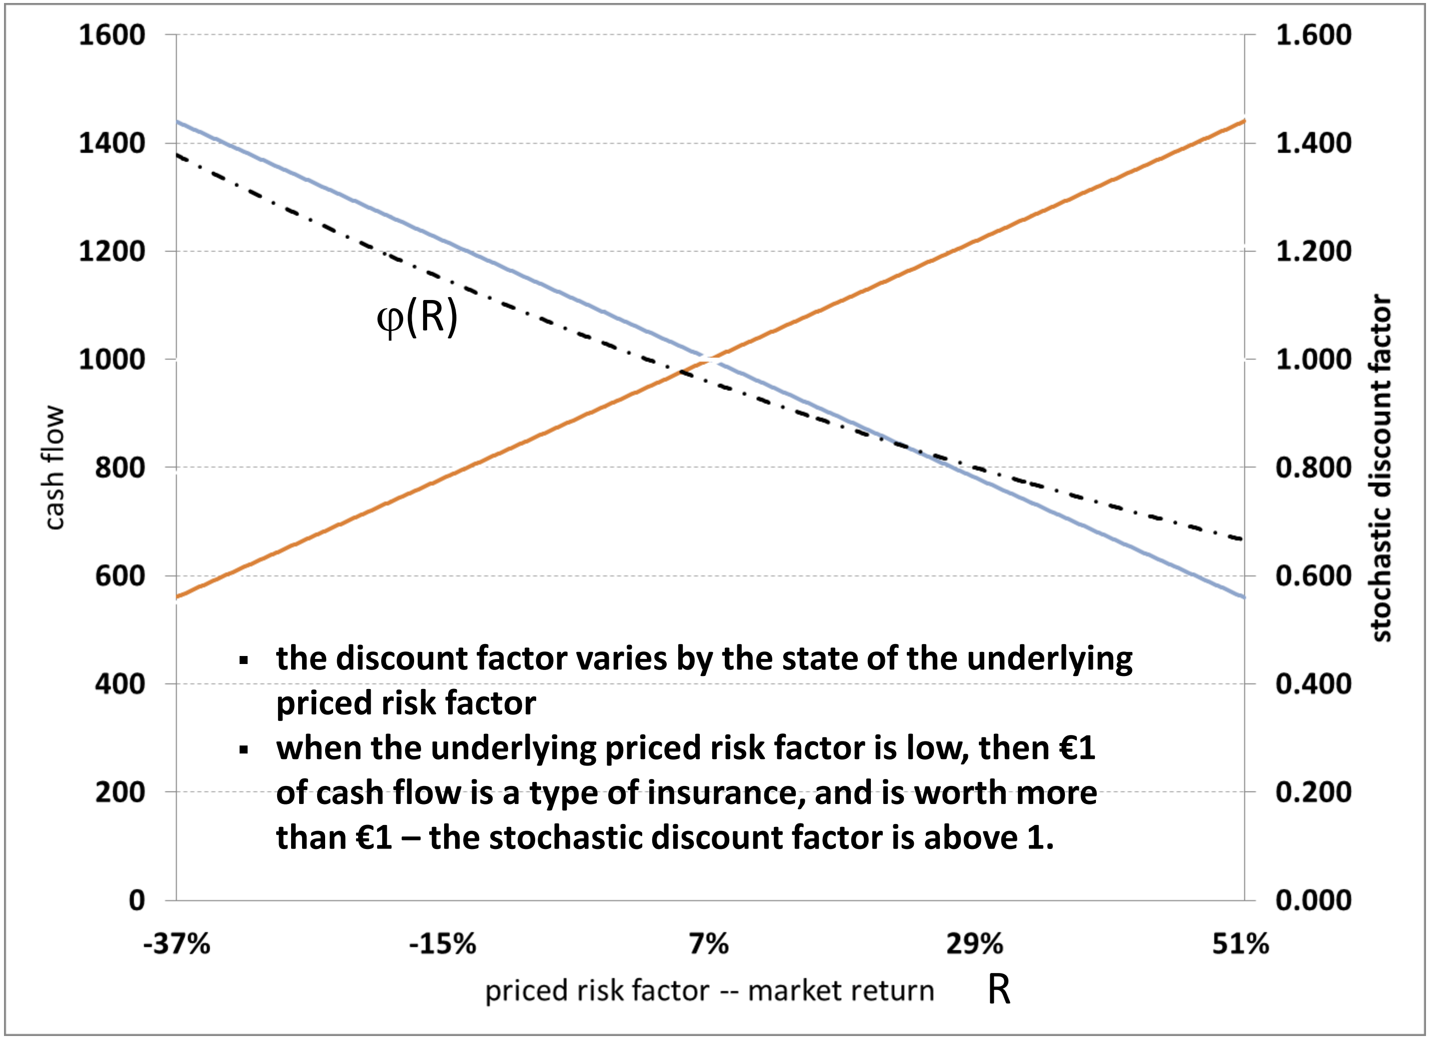

- **the discount factor varies by the state of the underlying priced risk factor**
- **when the underlying priced risk factor is low, then €1 of cash flow is a type of insurance, and is worth more than €1 – the stochastic discount factor is above 1.**

A Reminder of What's Behind a Risk-Adjusted Discount Rate: The Stochastic Discount Factor (5)

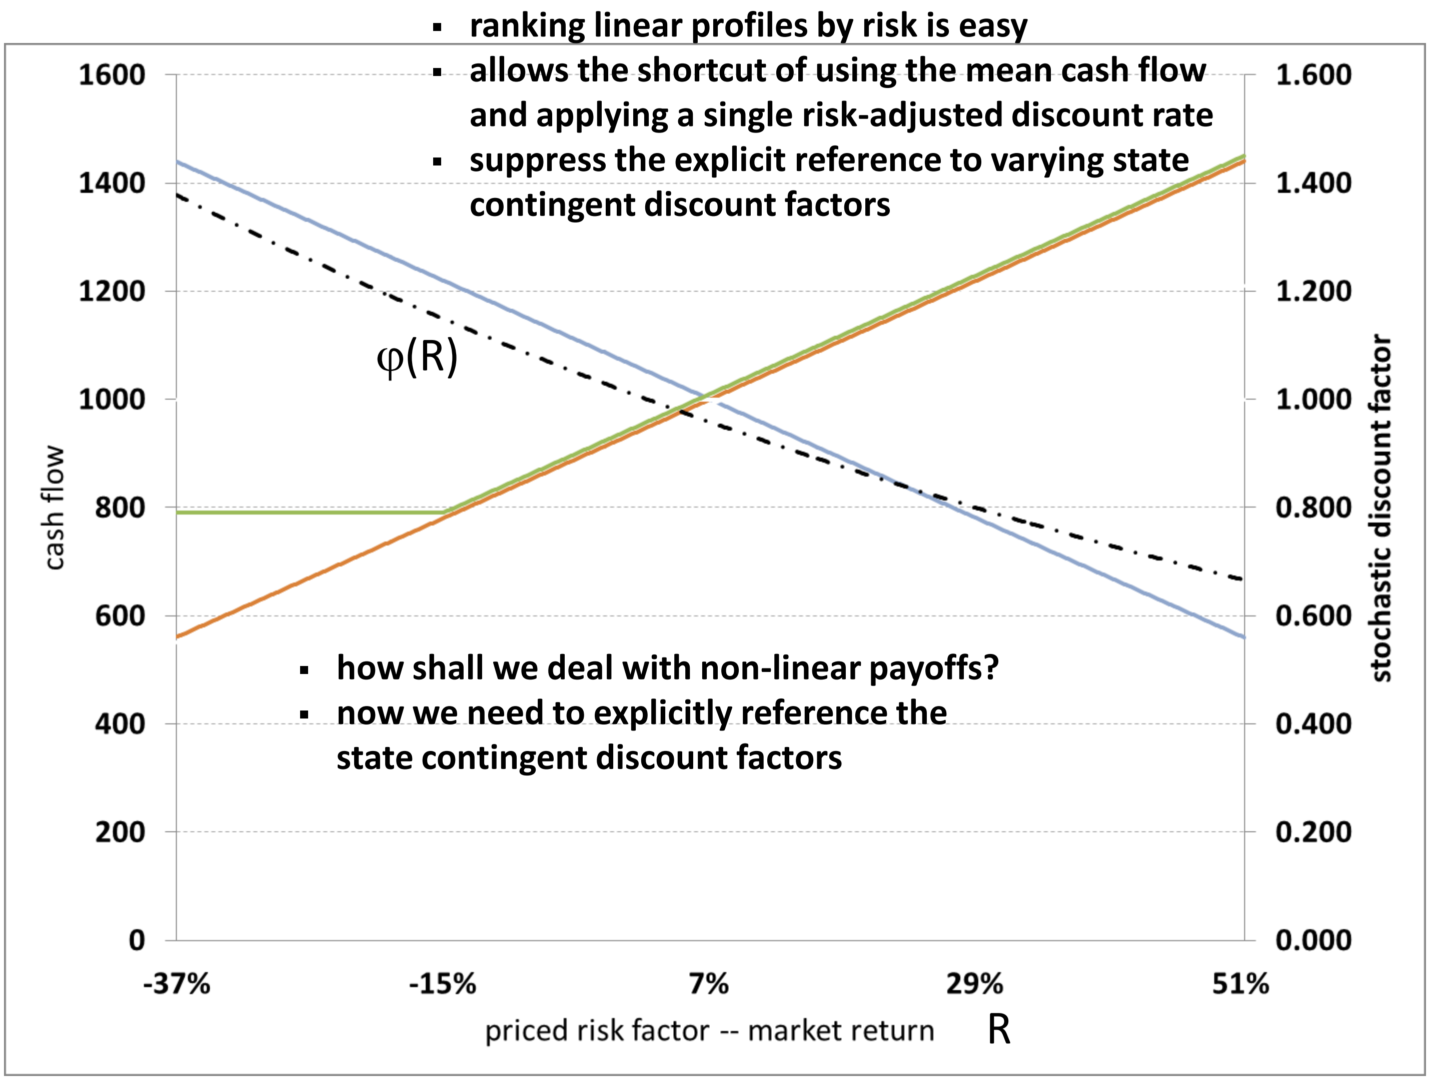

How Do We “Recover” the Stochastic Discount Factor

- Suppose that the stock market cumulative return, R, follows a random walk,
  - with mean return $\mu_m = r_f + \lambda_p$

- Then, the probability of any cumulative return at horizon T, $\pi\ (R_T)$, is given by the normal distribution function, with…
  - mean $\mu_m\ T$
  - variance $\sigma_m^2\ T$

- Then, the probability times the stochastic discount factor of any cumulative return at horizon T, $\pi(R_T)\varphi(R_T)$, is given by the adjusted normal distribution function, with…
  - mean $(\mu_m - \lambda_p)\ T$
  - variance $\sigma_m^2\ T$

  - aka, the risk-neutral probability distribution: $\pi^*(R_T) \equiv \pi\ (R_T)\ \varphi(R_T)$,
  - we can recover the stochastic discount factor, although not needed.

The "true" probability distribution, $\pi(R_T)$ -- assumed

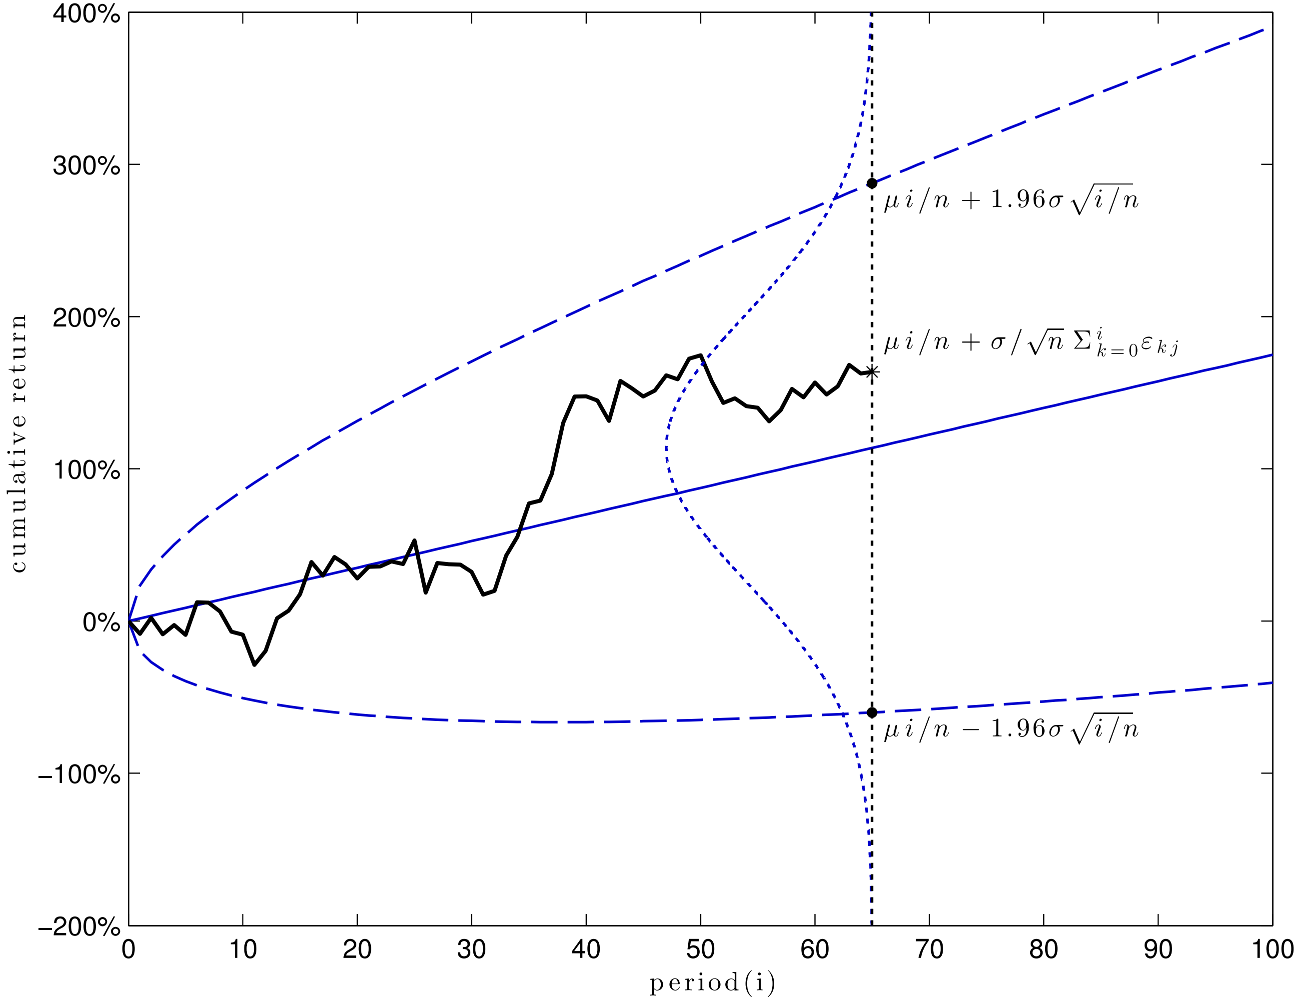

The “risk-neutral” probability distribution, $\pi^*(R_T)$ – implied

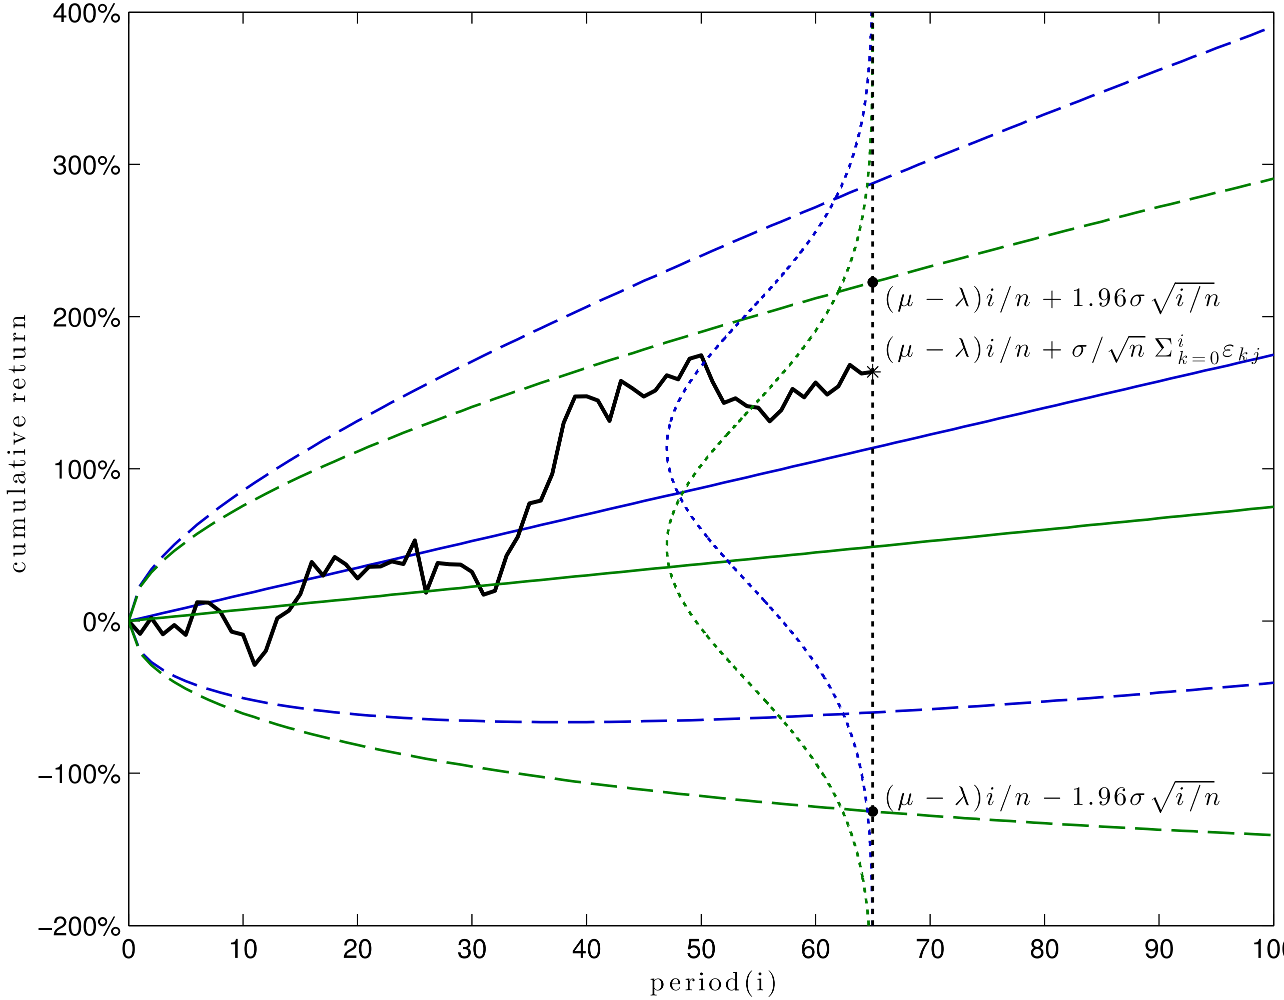

…and this implies the stochastic discount factor, $\varphi(R_T)$.

Detail #2:
Modeling Electricity Price Risk

Step #2: Overlay a Model of Electricity Price Risk

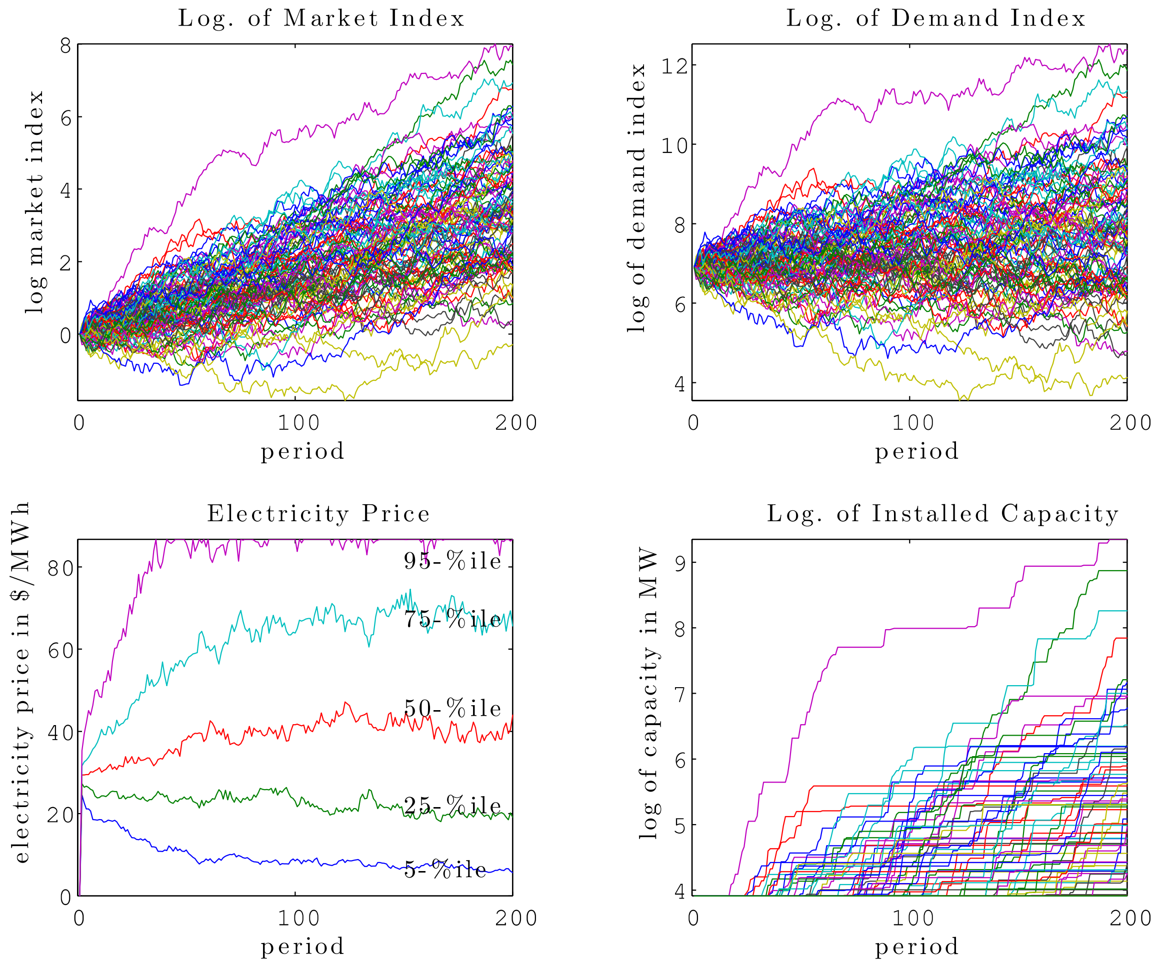

Detail #3:
Executing a Valuation

We're still doing DCF.

- The execution will operate much like the familiar DCF.
- Start with the familiar DCF equation,

$$PV = \sum_{t=0}^{T} \frac{E\left[CF_T\right]}{\left(1+R_a\right)^t}$$

- and focus on a single year's discounting.

$$PV = \frac{E\left[CF_1\right]}{1+R_a}$$

What's Old? What's New?

- Make 2 cosmetic changes:
  - use continuously compounded rates,
  - separate discounting for risk and discounting for time.

- Make 1 fundamental change – a generalization:
  - replace a single risk-adjusted discount rate with a state-contingent stochastic discount factor.

Continuously Compounded Rates

$$PV = \frac{E\left[CF_1\right]}{1+R_a} = E\left[CF_1\right]e^{-r_a}$$

Separate Discounting For Risk And Discounting For Time

$$PV = E\left[CF_1\right]e^{-r_a} = E\left[CF_1\right]e^{-\left(r_f + \lambda_a\right)} = E\left[CF_1\right]e^{-\lambda_a}\, e^{-r_f}$$

Preparing to Use the Stochastic Discount Factor

$$PV = E\left[CF_1\right]e^{-\lambda_a}\, e^{-r_f}$$

- Unpack the DCF formula by explicitly representing individual states of the underlying priced risk factor, indexed by j; here we shift to a discrete state space, …

$$PV = \left(\sum_{j=0}^{J} \pi_{1,j}\, CF_{1,j} \right) e^{-\lambda_a}\, e^{-r_f}$$

- Rearrange, moving the risk-adjusted discount factor inside the summation…

$$PV = \left(\sum_{j=0}^{J} \pi_{1,j}\, e^{-\lambda_a}\, CF_{1,j} \right) e^{-r_f}$$

Substituting in the Stochastic Discount Factor

$$PV = \left(\sum_{j=0}^{J} \pi_{1,j} \, e^{-\lambda_a} \, CF_{1,j} \right) e^{-r_f}$$

- Replace the single risk-adjusted discount factor with a state-contingent stochastic discount factor, …

$$PV = \left(\sum_{j=0}^{J} \pi_{1,j} \, \phi_{1,j} \, CF_{1,j} \right) e^{-r_f}$$

- Call the product of the probability and the stochastic discount factor the risk neutral probability, $\pi^*_{1,j} \equiv \pi_{1,j}\phi_{1,j}$ …

$$PV = \left(\sum_{j=0}^{J} \pi^*_{1,j} \, CF_{1,j} \right) e^{-r_f}$$

Discussion

Situating this Methodology in the Literature

- Real Options & Contingent Claims Valuation

- What is Different?
  - #1
    - real options is sold as a break with traditional DCF
    - we emphasize the continuity – DCF is a special case
  - #2
    - real options is sold as a magic trick
    - we open and demystify the central contribution
  - #3
    - real options solutions seem like one off special cases
    - we demonstrate a general methodology with an obviously flexible structure

Caveats

- This is more of a demonstration than a reliable set of values. The logic is correct, and the relationships shown are solid, but the specific values can only be consumed with the help of a heap of salt.
- The stochastic discount factor methodology puts great demands on risk modeling.
  - The underlying priced risk factor, and
  - The specific project representation … how the priced risk enters.
- The precision of the tools in principle far exceeds our ability to accurately calibrate them. So what is the point?

La méthode de simulation Monte-Carlo est mort, vive la méthode de simulation Monte-Carlo !