

Dynamic Capital Allocation with Irreversible Investments

Daniel Bauer*, Shinichi Kamiya[†], Xiaohu Ping[‡], George Zanjani*

April 13, 2016

Abstract

This paper studies the capital allocation rule in a dynamic portfolio optimization problem with irreversible investments. We find that marginal cost pricing can still be connected to capital allocation in this setup, although the basis for allocation is different from that found in static problems. The investment decision for an opportunity presented today is made on the basis of an expected future marginal cost of risk associated with that opportunity. The capital allocated for today's opportunity is a probability-weighted average of the product of the marginal value of capital in future states of the world and the amount of capital consumed by today's opportunity in those future states. In addition, our numerical examples show that failure to explicitly model uncertainties regarding future opportunities can lead to misstatement of the marginal cost of risk, and, as a result, over-investment or under-investment in current opportunities.

Keywords: capital allocation, diversification, real-option, portfolio optimization

JEL classification: C61, D81, D92, G11.

*Daniel Bauer (dbauer@gsu.edu) and George Zanjani (gzanjani@gsu.edu) are in the Department of Risk Management and Insurance at J. Mack Robinson College of Business, Georgia State University

[†]Shinichi Kamiya (skamiya@ntu.edu.sg) is in the Division of Banking and Finance at Nanyang Business School, Nanyang Technology University

[‡]Xiaohu Ping (xiaohu.ping@ge.com, +1 (518) 387 4555) is a Quantitative Finance Researcher in the Global Research Center of General Electric, 1 Research Circle, Niskayuna, NY, 12309.

ACKNOWLEDGMENTS

The authors thank our colleagues from Georgia State University and Nanyang Technology University who provided insight and expertise that greatly assisted the research. We also thank the discussants in European Group of Risk and Insurance Economists (EGRIE) Annual Conference for comments that greatly improved the manuscript.

The authors also declare that they have no relevant or material financial interests that relate to the research described in the paper.

Daniel Bauer

Shinichi Kamiya

Xiaohu Ping

George Zanjani

Dynamic Capital Allocation with Irreversible Investments

April 13, 2016

Abstract

This paper studies the capital allocation rule in a dynamic portfolio optimization problem with irreversible investments. We find that marginal cost pricing can still be connected to capital allocation in this setup, although the basis for allocation is different from that found in static problems. The investment decision for an opportunity presented today is made on the basis of an expected future marginal cost of risk associated with that opportunity. The capital allocated for today's opportunity is a probability-weighted average of the product of the marginal value of capital in future states of the world and the amount of capital consumed by today's opportunity in those future states. In addition, our numerical examples show that failure to explicitly model uncertainties regarding future opportunities can lead to misstatement of the marginal cost of risk, and, as a result, over-investment or under-investment in current opportunities.

Keywords: capital allocation, diversification, real-option, portfolio optimization

JEL classification: C61, D81, D92, G11.

1 Introduction

Practitioners have long wrestled with the problem of how to allocate a financial institution’s capital to the various risks within its portfolio for purposes of pricing and performance measurement. Over the past 15 years, a number of papers have explored the problem of allocating capital based on marginal risk contributions in different settings via different conceptual approaches (see Bauer and Zanjani (2013) for a review). However, these techniques are largely restricted to static portfolios, where the exposures are either set in advance or chosen simultaneously in an optimization problem.¹

Unfortunately, the real world application of these concepts is more complicated. Specifically, allocation for pricing purposes is typically done in an environment where the portfolio is not fixed in advance and is instead in the process of being constructed. Opportunities for underwriting or investments do not all arrive at the same time. If they did, or if positions in existing opportunities could be costlessly modified as new opportunities arrived, existing allocation techniques could be applied without modification. More typically, however, opportunities arrive sequentially in time, and, moreover, 1) decisions must be made about those opportunities when they arrive and 2) those decisions are typically irreversible, or reversible only with significant cost. This characterization applies to insurance companies, private equity firms, banks, and other institutions dealing in illiquid investments. How should risk be priced in such an environment? And can risk pricing be connected to capital allocation?

This paper tackles these questions by studying the usual portfolio optimization problem—where expected returns are optimized subject to a risk measure constraint—with three complications motivated by the discussion above. First, investment opportunities arrive sequentially. Second, the risk and return properties of future

¹While some papers do study “dynamic” approaches to allocation (Tsanakas (2004); Laeven and Goovaerts (2004)), the “dynamics” are introduced by evolving information about the components of a previously chosen portfolio rather than changes to the composition of the portfolio itself.

opportunities are uncertain. Third, investment decisions are irreversible.

We find that marginal cost pricing can indeed be connected to capital allocation in this setup, although the basis for allocation is different from that found in static problems. In static problems, the marginal value of capital is known with certainty, and capital is allocated to each risk based on how it consumes that capital at the margin. In this dynamic problem, the ultimate marginal value of a dollar of capital will not be known until all opportunities have been presented and decisions have been made; similarly, it is also not known what demands a given investment will make on the firm's capital until the portfolio is set. What *is* known—assuming that the uncertainty in the risk and return structure of future opportunities is understood—is the future distribution of possibilities. To elaborate, we know the possible paths ahead and their probabilities, so we know what the value of a marginal unit of capital will be in various future states of the world, and we also know in each of those states how that marginal unit of capital is consumed by the risks in the portfolio.

Thus, the investment decision for an opportunity presented today is made on the basis of an expected *future* marginal cost of risk associated with that opportunity. Expected future marginal costs can alternatively be expressed in terms of an allocation of the firm's capital cost to current investments and future opportunities. The basis of the allocation for today's opportunity is a probability-weighted average of the product of the marginal value of capital in future states of the world and the amount of capital consumed by today's opportunity in those future states of the world.

Part of the logic and intuition of the model is related to the literature on real options (e.g., Dixit and Pindyck (1994)). The fundamental problem in this literature entails the impact of uncertainty on investments that are either partially or completely irreversible. The classic intuition is that simple internal rate of return calculations, using a firm's cost of capital as a hurdle rate, lead to overly aggressive investment because they fail to account for the option value of waiting for more information on

the productivity of the investment. As a result, hurdle rates need to be adjusted upward, and greater uncertainty about productivity can be shown to lead to even higher hurdle rates (e.g., McDonald and Siegel (1986)). To the extent that the amount of investment is discretionary, current investment is delayed or reduced as there is value created by waiting.

Counteracting diversification effects, however, are introduced by the portfolio context of the problem studied in this paper. Since the ultimate productivity of any investment is determined by the final full portfolio of investments, early opportunities may receive relatively heavy investment relative to the attractiveness of their risk and return properties when viewed in the context of the current portfolio. The firm may invest heavily in anticipation of later additions to the portfolio, which will make the early opportunity more attractive in the end. Opportunities that arrive late in an underwriting season, however, will tend to be judged in the context of the portfolio as it exists at that point in time, as the firm has much less flexibility in building a portfolio around the opportunity in question.

The rest of this paper is organized as follows. We start in Section 2 by sketching basic theory in a two period model to build intuition. A multiperiod generalization is provided in an Appendix, which is then specialized to a Markovian setting in Section 3 to produce numerical examples showing how future uncertainties can affect optimal decisionmaking today. Section 4 concludes.

2 General Dynamic Capital Allocation

We consider first the case of two investment opportunities in two periods, and then discuss the generalization of the analysis to multiple investment opportunities in the Appendix.

The financial institution receives two opportunities in sequence. Upon receipt of

the first opportunity, it must decide on the quantity q_0 to invest in the opportunity. The second opportunity is known to be coming, but there is uncertainty about its risk and return properties. This uncertainty is denoted by the random variable α , which takes values in the set $\{\alpha_1, \dots, \alpha_m\}$ with probabilities $\{p(\alpha_1), \dots, p(\alpha_m)\}$. The uncertainty is resolved upon receipt of the second opportunity, at which point the firm decides on the quantity q_1 to invest in the second opportunity.

Let the return on the first opportunity be a random variable \mathbf{r}_0 and the return on the second have a distribution that depends on the realization of α , as in $\mathbf{r}_1(\alpha_i)$. Thus, different realizations of α will affect the prospective distribution of outcomes for the second opportunity. Let the firm's cost of risk bearing be τ , with risk capital needed to support the exposures being captured by the measure $\rho(q_0, q_1^i; \alpha_i)$.² If we consider the firm's problem at decision time on the first opportunity, we can frame the optimization problem as choosing a quantity for the first investment and a set of plans—contingent on different realizations of α —for the second investment:

$$\max_{q_0, \{q_1^1, \dots, q_1^m\}} q_0 \mathbb{E}[\mathbf{r}_0] + \sum_{i=1}^m p(\alpha_i) (q_1^i \mathbb{E}[\mathbf{r}_1(\alpha_i)] - \tau \rho(q_0, q_1^i; \alpha_i)), \quad (1)$$

subject to a set of constraints restricting the total amount of risk capital to a maximum level K :

$$\rho(q_0, q_1^i; \alpha_i) \leq K \quad \forall i = 1, \dots, m. \quad (2)$$

The first order conditions can be expressed as:

$$\mathbb{E}[\mathbf{r}_0] - \sum_{i=1}^m (p(\alpha_i)\tau + \lambda_i) \frac{\partial \rho(q_0, q_1^i; \alpha_i)}{\partial q_0} = 0, \quad (3)$$

$$p(\alpha_i) \mathbb{E}[\mathbf{r}_1(\alpha_i)] - (p(\alpha_i)\tau + \lambda_i) \frac{\partial \rho(q_0, q_1^i; \alpha_i)}{\partial q_1^i} = 0, \quad i = 1, \dots, m, \quad (4)$$

²It is also possible to write the objective function as an expected return minus frictional costs associated with capital, as in τK , so that the firm is burdened with a cost of capital rather than a cost of risk. Results of a largely similar flavor are obtained in this case.

where λ_i is the Lagrange multiplier associated with constraint i . Equation (3) balances the marginal benefit of exposure to the first risk (the expected return) with its marginal cost, while equation (4) does the same for the second risk in each of the possible states of the world.

Marginal cost is driven by the per unit cost of risk capital τ , as well as the costs relating to the maximum risk constraint, which are captured by the Lagrange multipliers λ_i . Note in each state that λ_i reflects an additional shadow cost or marginal value of state-contingent risk capital (i.e., at time zero, what an additional dollar of capital in state i —and only in state i —would be worth to the firm, net of the risk bearing cost τ). When (4) holds for an interior solution, we may write:

$$\lambda_i = p(\alpha_i) \frac{\mathbb{E}[\mathbf{r}_1(\alpha_i)] - \tau \frac{\partial \rho(q_0, q_1^i; \alpha_i)}{\partial q_1^i}}{\frac{\partial \rho(q_0, q_1^i; \alpha_i)}{\partial q_1^i}} \quad \forall i = 1, \dots, m. \quad (5)$$

It is possible for this shadow cost to be zero in states where the total risk constraint is not binding; in such states, the firm has additional capacity for risk and does not have use for it any more.

Thus, the marginal cost of exposure in each case is equal to the marginal amount of risk capital used (calculated as the partial derivative of the risk measure with respect to the quantity of the risk) times the sum of the frictional cost of risk τ and any additional shadow cost of risk capital used to support that risk. In the case of the first exposure, risk capital is affected in all possible future states of the world since the firm must commit to the exposure before the resolution of uncertainty, as seen in equation (3). The exact amount consumed as a result of the first underwriting decision, however, is uncertain and determined by the later underwriting decision.

Marginal cost can be reinterpreted as an allocation of risk capital times a cost if

the risk measure is homogeneous. Working with the second terms in (3) and (4):

$$\sum_{i=1}^m (p(\alpha_i)\tau + \lambda_i) \underbrace{\left[\frac{\partial \rho(q_0, q_1^i; \alpha_i)}{\partial q_0} q_0 + \frac{\partial \rho(q_0, q_1^i; \alpha_i)}{\partial q_1^i} q_1^i \right]}_{=\rho(q_0, q_1^i; \alpha_i)} = \tau \bar{\rho} + \sum_{i=1}^m \lambda_i K, \quad (6)$$

where $\bar{\rho} = \sum_i p(\alpha_i) \rho(q_0, q_1^i; \alpha_i)$ is the *expected value* of risk capital used. It is useful to consider two polar cases. First, if the total risk constraint *never* binds, then the sum of marginal costs times exposures contemplated in (6) reduces to $\tau \bar{\rho}$. In this case, total capacity K is never fully utilized; the *expected* risk $\bar{\rho}$, however, is fully allocated to each exposure through marginal cost pricing, with each unit of allocated risk being charged the frictional cost τ . Second, if the total risk constraint *always* binds, then total risk costs tally up to $(\tau + \sum_{i=1}^m \lambda_i) K$, which is naturally interpreted as a cost of capital (in parentheses) times total capital. In intermediate cases, the risk cost that is allocated is the *expected* total risk assumed times the frictional cost, $\tau \bar{\rho}$, plus the total shadow costs of risk capital, $\sum_{i=1}^m \lambda_i K$. Thus, marginal cost has a capital allocation interpretation that “adds up.”

It is important to contemplate the nature of the allocation to the first risk. At the time of decision making, the marginal cost of the first risk, and the capital allocated to the first risk, may differ from the marginal cost of that risk after uncertainty associated with α has been resolved and the portfolio completed. This *ex ante* allocation is an average of possible final cost allocations in future states of the world. To illustrate, by using equations (5) and (6), we see that the *ex ante* cost allocation to the first risk can be expressed as:

$$\sum_{i=1}^m p(\alpha_i) \left[\tau + \frac{\mathbb{E}[\mathbf{r}_1(\alpha_i)] - \tau \frac{\partial \rho(q_0, q_1^i; \alpha_i)}{\partial q_1^i}}{\frac{\partial \rho(q_0, q_1^i; \alpha_i)}{\partial q_1^i}} \right] \left[\frac{\partial \rho(q_0, q_1^i; \alpha_i)}{\partial q_0} q_0 \right].$$

Thus, for each state, we multiply the probability of state i times the marginal value of risk capital in state i (the first bracketed term) times the risk capital allocated to

the first risk in state i (the second bracketed term); the total cost allocation is then the sum of these figures over all states.

A few remarks on the two-period case are in order:

1. Note that if the uncertainty about the second opportunity concerns only the expected return and not the risk properties, then the allocation $\frac{\partial \rho(q_0, q_1^i; \alpha_i)}{\partial q_1^i} q_1^i$ will be the same in all states of the world. In other words, capital allocation will not be affected by the uncertainty. This is an artifact of a two-opportunity setup and is not likely to generalize to more than two opportunities.
2. The total marginal cost of risk capital (including the shadow cost) will not generally equal τ in every future state. Since capacity is inherited and fixed, the firm will typically face constraints on its risk-bearing, leading the marginal value of risk capital to differ from its carrying cost τ . To illustrate, the RORAC expression for the first risk can be obtained from equation (3):

$$\frac{\mathbb{E}[\mathbf{r}_0]}{\sum_{i=1}^m p(\alpha_i) \frac{\partial \rho(q_0, q_1^i; \alpha_i)}{\partial q_0}} = \frac{\sum_{i=1}^m p(\alpha_i) \left[\tau + \frac{\mathbb{E}[\mathbf{r}_1(\alpha_i)] - \tau \frac{\partial \rho(q_0, q_1^i; \alpha_i)}{\partial q_1^i}}{\frac{\partial \rho(q_0, q_1^i; \alpha_i)}{\partial q_1^i}} \right] \frac{\partial \rho(q_0, q_1^i; \alpha_i)}{\partial q_0}}{\sum_{i=1}^m p(\alpha_i) \frac{\partial \rho(q_0, q_1^i; \alpha_i)}{\partial q_0}}.$$

The expression on the right hand side is the hurdle rate, which is equal to the *expected* marginal value of risk capital in future states of the world, gross of the frictional cost τ . If the total risk bearing constraints were never binding, then the right hand side would collapse to τ . In general, however, the hurdle rate corresponding to the target return on risk capital will generally be higher than τ , depending on the extent to which the capacity constraints bind. Extension to multiple opportunities is conceptually straightforward though notationally tedious. We provide this extension in the Appendix.

3 Portfolio Optimization in a Markovian Setting

3.1 Multiple Risks

Suppose the financial institution with a certain initial wealth W has access to multiple potential opportunities (projects) from a risk pool. In each period $t \in \{0, 1, \dots, n\}$, one of m possible opportunities becomes available, or no opportunity emerges. We assume the probabilities are independent and identical for each period. More precisely, we assume the probability of opportunity i emerging is p_i and, consequently, the probability of no underwriting opportunity is $p_0 = 1 - \sum_{i=1}^m p_i$ for each period.

When one particular opportunity emerges, the firm makes a decision on the amount of risk that it will take on in its portfolio, $q_i(t)$. Moreover, let $\pi_i(t)$ represent the total cumulative amount of risk i in the company's portfolio acquired until time t . Clearly, we have $\pi_i(n) = \sum_{t=0}^n q_i(t)$, and $\pi_i(0) = 0, \forall i = 1, 2, \dots, m$.

Each opportunity generates a certain expected return per unit. In general, this return $r_i(t, \pi_i(t), q_i(t), W(t))$ may depend on the time t , quantity of the risk, and the company's wealth (as measured by initial assets and the expected value of profits associated with underwritten policies) at that time. However, to keep things simple, here we assume the expected return is exogenous r_i . Hence the total wealth at time t is $W(t) = W + \sum_{i=1}^m r_i \pi_i(t)$.

On the other hand, taking on additional risk will affect the institution's portfolio risk, which will influence the decision in two ways: 1) The risk corresponds to costly risk capital that the company has to hold, where we assume this amount is given via a *risk measure* and the frictional cost of capital is constant at the rate τ ; 2) the total capital that the company can access is constrained in that it cannot exceed the level K . We also introduce the portfolio risk measure at time t , $\rho(\boldsymbol{\pi}(t))$, where $\boldsymbol{\pi}(t) = [\pi_1(t), \dots, \pi_m(t)]$.

Therefore, the firm faces the following problem:

$$\max_{\mathbf{q}[0,n]} \mathbb{E}[W(n) - \tau\rho(\boldsymbol{\pi}(n))],$$

subject to:

$$\rho(\boldsymbol{\pi}(n)) \leq K.$$

Note that this problem is equivalent to the optimization problem elaborated in the Appendix, which extended the two-period model of Section 2 to a multiperiod generalization. The objective function is identical to the objective function (14), and the constraint lines up with equation (15). In words, the firm seeks to maximize the expected profits associated with underwriting while respecting the constraint that total risk capital never exceeds capacity K .

To solve this problem, we write the objective function as a combination of a final risk measure at time n and cumulative instantaneous revenues, given the initial data (π_0, t_0) :

$$V[\mathbf{q}, \boldsymbol{\pi}, n](\boldsymbol{\pi}_0, t_0) = \sum_{i=1}^m \sum_{s=t_0}^n [r_i q_i(s) \mathbb{1}_{\{\text{Risk}(s)=i\}}] - \tau\rho(\boldsymbol{\pi}(n)). \quad (7)$$

Thus, the expected value function is:

$$v^*[\boldsymbol{\pi}, t] = \max_{\mathbf{q}[t,n]} \{ \mathbb{E} [V[\mathbf{q}, \boldsymbol{\pi}, n](\boldsymbol{\pi}, t) | \boldsymbol{\pi}(t) = \boldsymbol{\pi}] \}, \quad (8)$$

with the portfolio $\boldsymbol{\pi}(t) = \boldsymbol{\pi}$ at time t . By the dynamic programming principle, we can rewrite the value function as:

$$v^*[\boldsymbol{\pi}, t] = \max_{\mathbf{q}(t)} \left\{ \sum_{i=1}^m p_i r_i q_i(t) + p_i v^*[\boldsymbol{\pi} + q_i(t), t + 1] \right\}.$$

Thus, in each period t , the firm faces m possible problems, depending on the risk

$i \in \{1, 2, \dots, n\}$ that emerges:

$$\max_{q_i(t)} r_i q_i(t) + v^*[\boldsymbol{\pi} + q_i(t), t + 1],$$

subject to:

$$\rho(\boldsymbol{\pi} + q_i(t)) \leq K.$$

The capital condition needs to be satisfied only after the decision at the final time n . However, at any point in time, there is a non-zero chance that no future opportunities will emerge. Thus, the constraint indeed has to be satisfied at every time t .

Solving the first order conditions of the above problem, we obtain the decision for each risk i at each period t from the implicit function:

$$r_i = -\frac{\partial v^*[\boldsymbol{\pi} + q_i(t), t + 1]}{\partial q_i(t)}, \quad (9)$$

assuming this decision $q_i(t)$ does not violate the risk constraint $\rho(\boldsymbol{\pi} + q_i(t)) \leq K$. Embedded in (9) is a notion of risk capital cost allocation. This can be seen more easily when mapping the Markovian problem into the more general framework of the Appendix to rewrite (19) as:

$$r_i = \tau \mathbb{E} \left[\frac{\partial \rho(\boldsymbol{\pi}(n))}{\partial q_i(t)} \middle| \boldsymbol{\pi}(t) = \boldsymbol{\pi} \right],$$

which we write, for purely pedagogical purposes, under the assumption that $\boldsymbol{\pi}(n)$ reflects the corresponding period maximizers and that the constraint never binds. In this extreme case, the marginal cost of risk can be written as the frictional cost of risk capital τ times an allocation of capital per unit of exposure (the bracketed term on the right-hand side). For the more general case where the constraint binds in some

states of the world, the result (19) yields:

$$r_i = \mathbb{E} \left[(\tau + \lambda) \frac{\partial \rho(\boldsymbol{\pi}(n))}{\partial q_i(t)} \middle| \boldsymbol{\pi}(t) = \boldsymbol{\pi} \right],$$

where λ is a random variable reflecting the value of relaxing the risk measure constraint in every possible future state of the world. Thus, the current marginal return on an investment is again equated to an expected product of a marginal cost of risk capital $(\tau + \lambda)$ and a marginal allocation of risk capital.

3.2 Numerical Illustration: A Two Risk Problem

Suppose the investor has the opportunity to invest in two independent risks, each having an expected return r_i , $i = 1, 2$. In each period $t \in \{0, 1, \dots, n\}$, there are three possibilities: 1) risk 1 emerges with a probability p_1 , 2) risk 2 emerges with a probability p_2 , or 3) neither of the risks emerges with a probability $p_0 = 1 - p_1 - p_2$. The homogeneous risk measure we will use in this subsection is the Standard Deviation:

$$\rho(\boldsymbol{\pi}) = \sqrt{\pi_1^2 \sigma_1^2 + \pi_2^2 \sigma_2^2}.$$

This is especially convenient for Normal distributed marginals since advanced risk measures have an analogous form (see Embrechts, Frey, and McNeil, 2005). Other risk measures can be considered similarly. Several general features of this problem are worth observing.

First, it is possible for investment to be trivially zero if the cost of risk capital is too high. Specifically, it can be shown that zero investment will be optimal if $\tau^2 > \frac{r_1^2}{\sigma_1^2} + \frac{r_2^2}{\sigma_2^2}$, even if we removed uncertainty from the problem and allowed the investor to set exposures to both risks without constraint.

Second, assuming that positive investment can be supported (i.e., if $\tau^2 < \frac{r_1^2}{\sigma_1^2} + \frac{r_2^2}{\sigma_2^2}$),

neither risk is viable on its own if:

$$\tau^2 > \max \left\{ \frac{r_1^2}{\sigma_1^2}, \frac{r_2^2}{\sigma_2^2} \right\}. \quad (10)$$

Profitability in this case depends on combining the risks to take advantage of diversification. Thus, when the first opportunity presents itself, the investor's calculus in deciding on the amount to invest depends crucially on the chances of the second opportunity emerging. It should be noted that this diversification calculus may drive decisions in a different direction than the real option effect typical in sequential investment problems: Real option logic typically encourages the investor to delay investment and wait for better opportunities early in the process, while becoming more aggressive later in the process when fewer opportunities remain. Diversification logic, on the other hand, encourages the investor to stake out large initial positions early in the game, when diversification opportunities are likely to arrive later on, but be more cautious later on the game, when those diversification opportunities are less likely to ever come.

Finally, if the investor faced no uncertainty and were allowed to set the exposures freely, the optimal ratio of investment is given by:

$$\frac{\pi_1}{\pi_2} = \frac{r_1 \sigma_2^2}{r_2 \sigma_1^2}, \quad (11)$$

which, together with the total risk constraint:

$$\rho(\boldsymbol{\pi}) = \sqrt{\pi_1^2 \sigma_1^2 + \pi_2^2 \sigma_2^2} \leq K, \quad (12)$$

can be used to determine the best possible combination of investments that could possibly be achieved by the investor.

We seek numerical solutions of the dynamic optimization problem by implement-

ing the Bellman equation from Section 3.1. If opportunity i emerges, the optimization problem is (for $i = 1, 2$):

$$\max_{q_i(t)} \{r_i q_i(t) + v^*[\boldsymbol{\pi} + q_i(t), t + 1]\},$$

subject to:

$$\sqrt{\pi_1^2 \sigma_1^2 + \pi_2^2 \sigma_2^2} \leq K.$$

If the two risks are dependent, then the standard deviation will take the form:

$$\rho(\boldsymbol{\pi}) = \sqrt{\pi_1^2 \sigma_1^2 + \pi_2^2 \sigma_2^2 + \phi \sigma_1 \sigma_2 \pi_1 \pi_2},$$

where ϕ is the correlation between the risks, so the constraint can be modified to:

$$\sqrt{\pi_1^2 \sigma_1^2 + \pi_2^2 \sigma_2^2 + \phi \sigma_1 \sigma_2 \pi_1 \pi_2} \leq K.$$

3.2.1 Numerical Approach

In the numerical approach, we solve the dynamic optimization problem backwards. For the last period n , we can analytically solve the optimization problem for each of the three possible opportunities that may emerge, and obtain the expected value function. More details of the solution are shown later in this section. For each previous period, we look for the optimal amount of investment for each possible risk among a set of possible investment levels to maximize the sum of the return on the new investment and the next period's expected value function. Specifically, in each period, we need to find the optimal new investments q_1^* or q_2^* for all possible initial risk exposure levels π_1 and π_2 .³

We solve the problem using a grid-based approach with discrete support. First,

³In cases where negative investment is possible, the initial exposure to a risk will not affect the optimal investment in that risk. This will be discussed in greater detail in Section 3.2.2.

we find the possible range of risk exposures, from zero to the amount that binds the risk constraint without exposure to the other risk, i.e., K/σ_1 and K/σ_2 . Then, we parse each risk exposure range into an equally spaced set of points, with s_1 being the number of points for the first risk and s_2 the number for the second. These sets of points constitute the possible investments in the decision set. The finer the grid, the more possible investment decisions we provide, and the more accurate is the approximation. Optimal choices must respect the risk capital constraint, so that admissible π_1 and π_2 have to satisfy $\sigma_1^2\pi_1^2 + \sigma_2^2\pi_2^2 \leq K^2$. As a result, we have a 2-dimensional grid (matrix) of investment levels where the admissible part resembles an upper triangle.

We now illustrate the mechanics of the approach for a case where (10) holds, so that profitability depends on diversification. For the final period n , if no risk emerges, then the value matrix is calculated by $v^*(\boldsymbol{\pi}, n|0) = -\tau\sqrt{\pi_1^2\sigma_1^2 + \pi_2^2\sigma_2^2}$, where π_1 and π_2 exist for the admissible grid. If risk 1 emerges, first we calculate a candidate for the optimal level of exposure to the first risk, conditional on the fixed exposure to the second risk and ignoring the constraint, by solving the first order condition for q_1^* :

$$\pi_1 + q_1^* = \frac{(r_1/\sigma_1^2)\sigma_2\pi_2}{\sqrt{\tau^2 - r_1^2/\sigma_1^2}}.$$

Then we consider the risk measure constraint, so we update the new π_1 matrix by:

$$\pi_1 + q_1^* = \frac{1}{\sigma_1}\sqrt{K^2 - \sigma_2^2\pi_2^2},$$

for those locations where $\sigma_1^2\pi_1^2 + \sigma_2^2\pi_2^2 > K^2$. The constraint will bind in circumstances where:

$$\pi_2 \leq \frac{K}{\tau\sigma_2}\sqrt{\tau^2 - r_1^2/\sigma_1^2}.$$

The calculations are analogous if risk 2 emerges. Therefore, we can calculate the

conditional value matrices as:

$$v^*[\boldsymbol{\pi}, n|i] = \begin{cases} -\sigma_j \pi_j \sqrt{\tau^2 - r_i^2/\sigma_i^2} - r_i \pi_i & \text{if } \pi_j \leq \frac{K}{\tau \sigma_j} \sqrt{\tau^2 - r_i^2/\sigma_i^2}, \\ \frac{r_i}{\sigma_i} \sqrt{K^2 - \sigma_j^2 \pi_j^2} - \tau K - r_i \pi_i & \text{if } \pi_j > \frac{K}{\tau \sigma_j} \sqrt{\tau^2 - r_i^2/\sigma_i^2}. \end{cases}$$

Hence, the expected value matrix is:

$$v^*[\boldsymbol{\pi}, n] = p_1 v^*[\boldsymbol{\pi}, n|1] + p_2 v^*[\boldsymbol{\pi}, n|2] + p_0 v^*[\boldsymbol{\pi}, n|0].$$

Once we have the value matrix, we can proceed to analyze the previous period and work backwards through time: For each period before n , we know the value function of the next period, and based on the next period's value matrix, we find the optimal decisions in the current period, and subsequently the expected value matrix of this period. For example, at period $n - 1$, we know the expected value matrix at period n is $v^*(\boldsymbol{\pi}, n)$, and we will use it to find the optimal $q_1(n - 1)$ and $q_2(n - 1)$, as well as the expected value matrix $v^*(\boldsymbol{\pi}, n - 1)$.

If no risk emerges, the value matrix is:

$$v^*[\boldsymbol{\pi}, n - 1|0] = v^*[\boldsymbol{\pi}, n].$$

If risk i emerges, the investor's problem is:

$$\max_{q_i(n-1)} r_i q_i(n - 1) + v^*[\boldsymbol{\pi} + q_i(n - 1), n],$$

which is equivalent to:

$$\max_{q_i(n-1)} r_i \pi_i(n) + v^*[\boldsymbol{\pi}(n), n] - r_i \pi_i(n - 1). \quad (13)$$

We determine the corresponding maximizers by a grid search over the discrete choice

set. More precisely, in the case where positive and negative adjustments are permissible, for each given value of π_j , we solve for the optimal π_i , $j \neq i$, and subtract the current level according to Equation (13). In the case of nonnegative investments, we additionally check whether the optimal choice is admissible—and choose the boundary value otherwise.

3.2.2 Base Case

We start by considering the case of independence and allow investments to be partially reversible in the sense that the investor can choose negative exposure to an opportunity if and when it arises. Our “base case” thus considers two independent risks, and investments can be negative.

Table 1: Parameters in Base Case

Name	Meaning	Value
τ	cost of risk capital	4.5
K	maximum risk capital	60
r_1	expected return of risk 1	9
r_2	expected return of risk 2	8
σ_1	standard deviation of risk 1	3
σ_2	standard deviation of risk 2	2
p_1	probability of opportunity 1	0.5
p_2	probability of opportunity 2	0.3
p_0	probability of no opportunity	0.2
n	total periods	8
ϕ	correlation of two risks	0
s_1	number of possible investments in risk 1	2000
s_2	number of possible investments in risk 2	3000

We select parameters in Table 1 to insure that Condition (10) is satisfied so that diversification is essential. Risk 1 has higher return with a higher standard deviation and a higher probability of arrival. We consider an 8-period model.

First, we investigate the relationship between the decision on risk 1 given an initial risk 2 exposure π_2 , given that risk 1 is emerging. Since in the base case, negative investment is allowed, the initial risk 1 exposure π_1 will not affect the total optimized

amount of risk 1 exposure, $\pi_1 + q_1^*$, at the end of each period. We analyze how the decision curve changes with time, as shown in Figure 1.

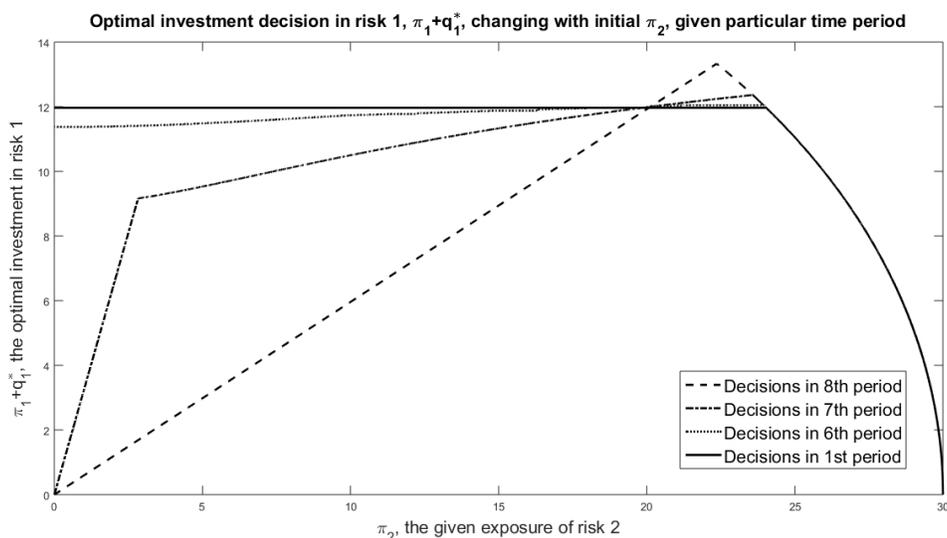


Figure 1: Optimal investment decisions in risk 1, $\pi_1 + q_1^*$, with different initial π_2

The vertical axis measures $\pi_1 + q_1^*$, the decision on risk 1 given that risk 1 emerges, as a function of the amount of risk 2 exposure (shown on the horizontal axis) at the beginning of different periods. The different lines show how the decision changes from the last period (dashed line) to the first period (solid line).

In this setting, given the risk capital constraint of 60, the best possible investment levels are $\pi_1^b = 12$ and $\pi_2^b = 24$. From this figure, we can see that the optimal choice $\pi_1 + q_1^*$ is set very close to this best possible level, regardless of the given level of risk 2 exposure, during early periods when the constraint is not binding. At high levels of risk 2 exposure, beyond the best-possible level of 24, the constraint forces the choice for $\pi_1 + q_1^*$ below the best-possible level of 12; however, such high levels of risk 2 exposure could only occur if they were inherited, as they would never be chosen by the investor. In later periods, the optimal choice for $\pi_1 + q_1^*$ drops at low levels of risk 2 exposure. At higher levels of risk 2 exposure, in the 20 to 24 range, the optimal choice for $\pi_1 + q_1^*$ rises above the best-possible level in the latest periods.

The intuition is fairly straightforward. Early on, the investor aims for the best-possible outcome, as there is plenty of time for the second risk to emerge, and, even if it does not, it is likely that the investor will have an opportunity to reverse the position later on through negative investment in risk 1. In the later periods, however, the possibility that the investor will not get another chance to adjust her exposure to either risk looms much larger in her calculus. As a result, she moves in the direction of assuming that the current level of risk 2 exposure will remain at that level, and the choice for $\pi_1 + q_1^*$ is therefore optimized relative to that level of risk 2 exposure. The latter late-period behavior can lead to investment in the first risk that is lower than the best-possible level if risk 2 exposure is either low or above 24 (in which case investment is made to the point where the risk constraint is binding). The late-period behavior can also lead to investment greater than the best-possible level of 12 when risk 2 exposure is close to, but short of 24; in this circumstance, the investor takes up remaining capacity that would ideally be dedicated to risk 2 if more time remained.

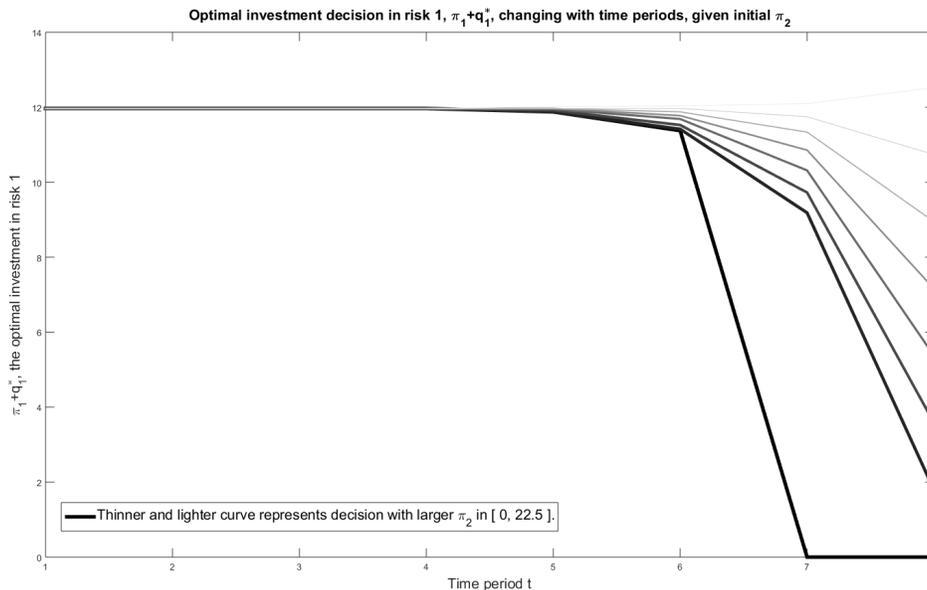


Figure 2: Optimal investment decisions in risk 1, $\pi_1 + q_1^*$, with different time periods

Figure 2 presents this information from another angle, showing how the risk 1

decision varies by time periods, for different levels of given risk 2 exposure. The horizontal axis shows time from period 1 to 8, and the vertical axis displays the investment decision on risk 1 given that it arrived. As the lines get thinner, the given level of π_2 increases from 0 (the thickest line) to the level of 22.5 (the thinnest line).⁴

As time progresses, it becomes more difficult to adjust positions, so the investment in risk 1 is pushed toward the optimal level for diversifying purposes (that is, the optimal level conditional on the current, though potentially nonoptimal, level of exposure to risk 2). For small levels of risk 2, the optimal investment on risk 1 is decreasing in time, while for high enough (though suboptimal) levels of risk 2, investment can increase in time as capacity that was reserved for risk 2 is soaked up by risk 1 in later periods. The key time period appears to be time period 4: In period 4 and prior, the investor chooses the best-possible level for risk 1, while the dream of perfection—if it has not been achieved—is abandoned in favor of compromise starting in period 5.

3.2.3 Dependent Risks

We now introduce correlation between the two risks. The six panels of Figure 3 show how the decisions change with time at correlations of 0.3, 0.4, 0.5, 0.6, 0.8, and 0.9.

The basic story is largely the same. As correlation ϕ increases from 0 to 0.5, the best-possible level of investment on risk 1 falls, but the firm still aims for perfection in the early periods on the assumption that adjustments can be made at a later date. Later periods push investment toward the optimal level for diversification purposes, given the level of exposure to risk 2. If correlation increases enough, investment is no longer profitable, and the investor will aim for zero investment in both lines. If

⁴The highest-possible level of π_2 is 24. When π_2 takes values in excess of 22.5, the optimal decision on risk 1 exposure is driven lower to satisfy the constraint. We show decisions with a maximum of $q_2 = 22.5$, corresponding to 75% of the maximum exposure to risk 2 in the grid, and similar rules are applied in the later figures.

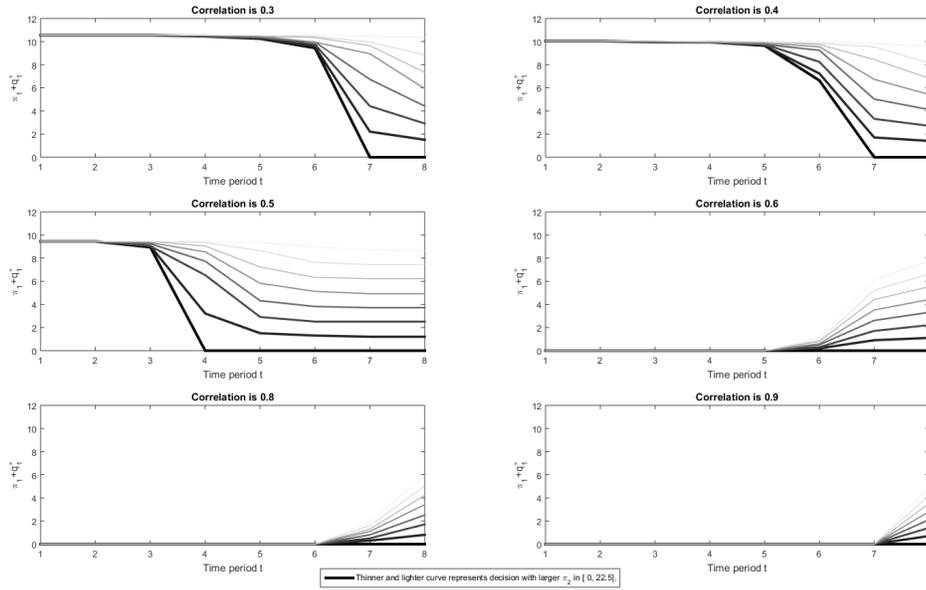


Figure 3: Optimal decision changes with time periods for different correlations

the investor has inherited an exposure to risk 2 that she has been unable to shed in the earlier periods, diversifying investments in risk 1 become optimal. This situation appears in this example with correlations above 0.5, as can be seen in Figure 3.

3.2.4 Nonnegative Investments

Restricting new investments to be nonnegative changes the early period calculus considerably. In the base case with unrestricted investment, the investor aimed for perfection in early periods with the knowledge that she would likely be able to reverse the investment if the diversification opportunity did not materialize. Making the investments irreversible eliminates the fall-back option.

The previously described numerical method is easily adapted to this case. For example, after we find a candidate for the optimal investment level in risk 1, we need to compare it with the current risk 1 exposure: if the candidate decision involves a decrease in the investment level, it will violate the nonnegativity rule, so the firm keeps the current investment on risk 1.

We now add the nonnegativity constraint to the base case. Introducing irreversibility complicates the visualization of results, since it makes the optimal decision for risk 1 depend on the current levels of exposures to *both* risk 1 and risk 2. Consequently, we cannot draw the decision structure over various periods (or various risk 2 exposures) in multiple curves on a 2-dimensional figure. To address this, we present results for a given risk 1 exposure of approximately 4 (corresponding to 20% of the maximum exposure to risk 1 in the grid).

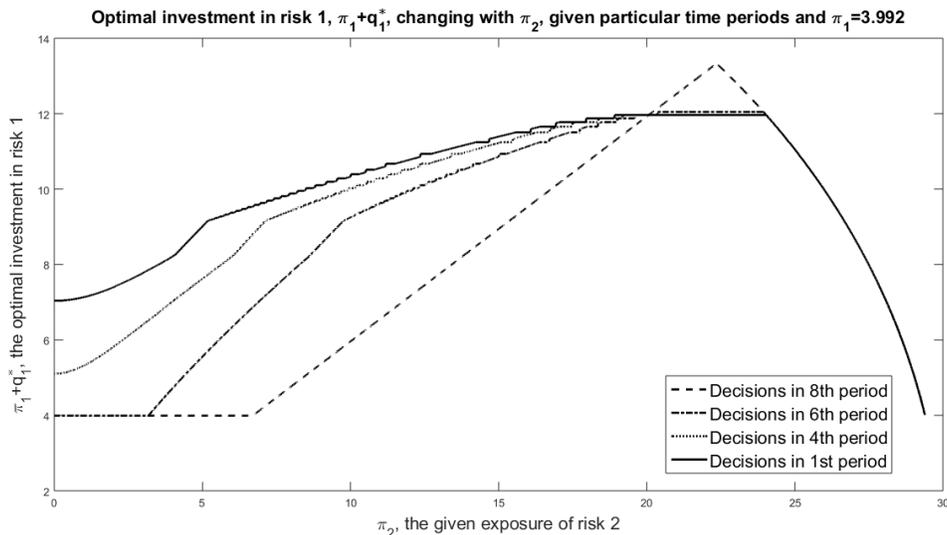


Figure 4: Optimal decisions on $\pi_1 + q_1^*$ with nonnegative investments

Figure 4 shows, given this initial risk 1 exposure, how optimal investment in risk 1 varies with an initial risk 2 exposure (displayed on the x -axis). The individual lines in the graph correspond to different time periods, with time shifting from late periods in dash line to early periods in more solid line. The impact of irreversibility is seen when comparing the early period decisions in Figure 4 with their counterparts in Figure 1. Investments in risk 1 in early periods are much smaller in the presence of irreversibility, as the possibility of being stuck with an undiversified investment looms larger in the profit calculus. This is also evident in Figure 5, which displays investment in risk 1 as a function of time period, with the various lines corresponding

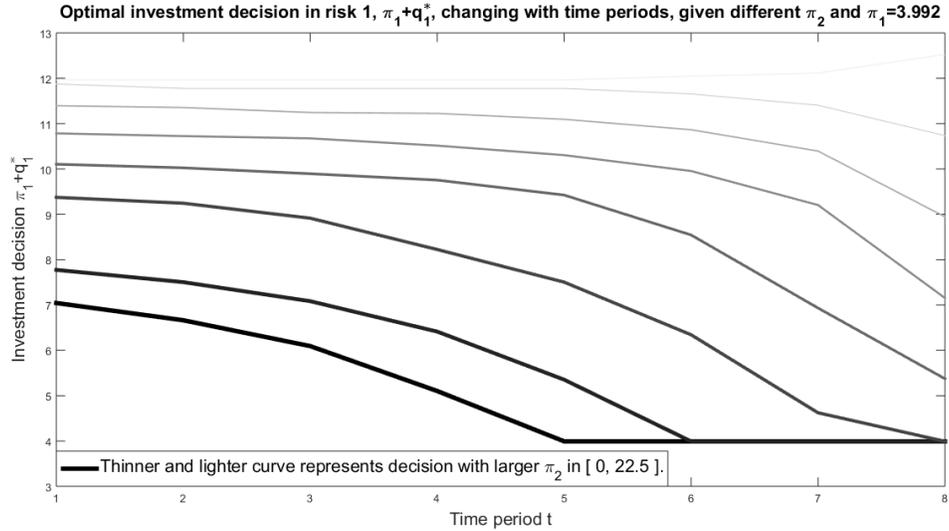


Figure 5: Optimal decisions on $\pi_1 + q_1^*$ with nonnegative investments

to different levels of initial risk 2 exposure, with the thinner lines indicating higher levels of initial exposure. In Figure 2, with reversibility, all levels of initial risk 2 exposure were associated with the same level of risk 1 investment in early periods, as the investor aimed for the best possible portfolio. In contrast, Figure 5 shows that progressively lower levels of initial risk 2 exposure are associated with progressively lower levels of risk 1 investment, even in early periods, as the investor must hedge the possibility of ending up with an undiversified portfolio.

3.2.5 Illustrating the Real Option Effect

The previous numerical examples all are driven primarily by the imperative of diversification, and this imperative yields decision rules that are somewhat unusual for sequential investment problems. The typical “real option” logic stresses patience early in the game, as the possible arrival of better opportunities leads to higher early period hurdle rates and more cautious early investment decisions than would be predicted by standard internal rate-of-return models. Diversification motives, on the other hand, dictate the opposite behavior: Early investment decisions are

more aggressive than later ones, as greater amounts of remaining time offer richer opportunities for diversification.

The “real option” effect—of restraining early investment to “leave room” for more attractive opportunities—can also be recovered in numerical analysis. To illustrate this effect, we consider several modifications to the base case. In addition to the modification of the parameters, which is shown in Table 2, we consider irreversible investment and demonstrate the decision curve with initial risk 1 exposure of approximately 4 as we did in Section 3.2.4.

Table 2: Parameters for Illustrating Real Option Effect

Parameter	Case 1	Case 2	Case 3	Case 4
τ	3	3	3	3
K	60	60	60	60
r_1	18	8	18	8
r_2	8	18	8	18
σ_1	3	2	3	2
σ_2	2	3	2	3
p_1	0.15	0.33	0.15	0.33
p_2	0.33	0.15	0.33	0.15
p_0	0.52	0.52	0.52	0.52
n	8	8	8	8
ϕ	0	0	0.8	0.8
s_1	2000	2000	2000	2000
s_2	3000	3000	3000	3000

We first consider, in Case 1, the situation where risk 1 is relatively infrequent (with the probability of emergence reduced from 0.5 to 0.15) but highly profitable (with a return doubled from 9 to 18). Risk 2 has similar parameters to those from the base case. We also modify the cost of capital τ from 4.5 to 3 so that each investment itself is profitable and diversification is not necessary for investment, although diversification is still beneficial. The best possible outcome is $\pi_1^* = \pi_2^* = 60/\sqrt{13} \approx 16.64$.

Figure 6 shows how the optimal investment in risk 1 evolves across time periods, with the individual lines representing different levels of given risk 2 exposure, with thinner lines representing higher levels of exposure.

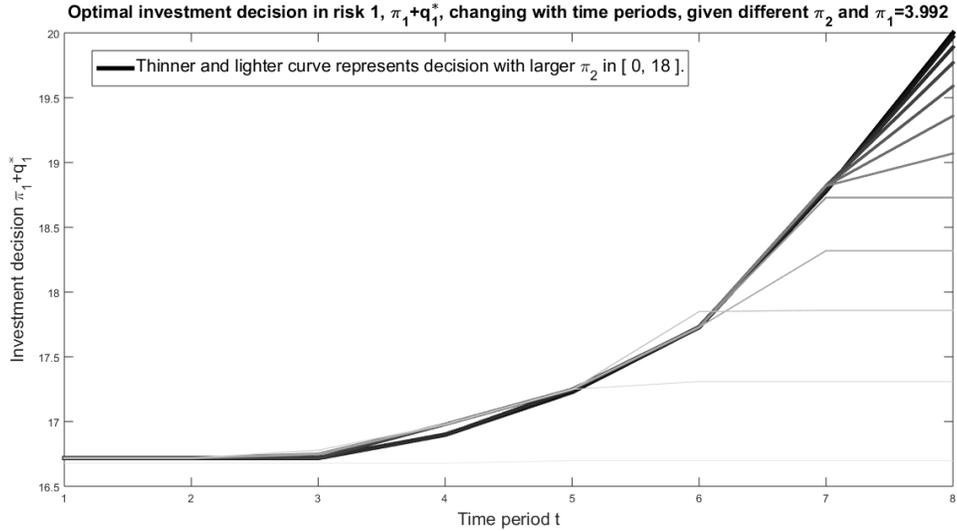


Figure 6: Optimal investment for high return and low frequency risk

From the figure, we can see that the risk 1 decision is increasing with time when the risk 2 exposure is small. In early periods, the firm targets an ideal level of investment, under the assumption that it will be able to diversify with risk 2. As time progresses, this hoped for situation is less likely to emerge, so it just soaks up more of its capacity with the first risk. This increasing decision rule is an illustration of the real option effect working in conjunction with the diversification effect. We show the results with initial risk 1 exposure at approximately 4, which is considerably lower than the best possible level 16.64.

Next, in Case 2, we reverse the relative positions of the two risks and look at the case where risk 1 is frequent and slightly profitable while risk 2 is infrequent and highly profitable. In other words, we exchange the probability and profitability parameters of risk 1 and 2. Figure 7 again shows how the optimal investment in risk 1 evolves across time periods, with the individual lines representing different levels of given risk 2 exposure, with thinner lines representing higher levels of exposure. In this case, the best possible investment is again $\pi_1^* = \pi_2^* \approx 16.64$, the same as in Case 1.

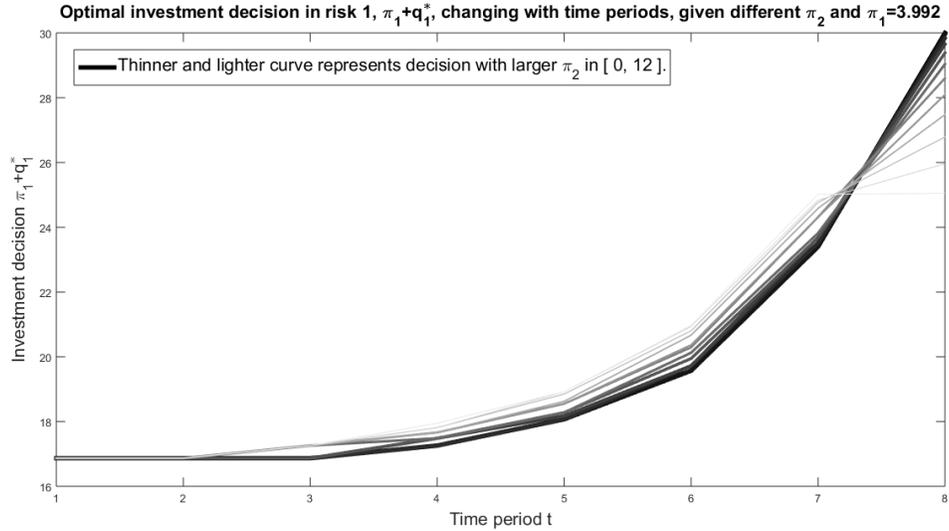


Figure 7: Optimal investment for low return and high frequency risk

The real option effect now appears more clearly. In the early periods, investment in the less profitable risk 1 is initially set at a level to accommodate a large investment in the more profitable risk 2. As time progresses and the arrival of risk 2 becomes less likely, the investor settles by soaking up remaining capacity with risk 1. While this capacity would ideally be devoted to the more profitable risk, using it on the less profitable risk is better than failing to use it at all.

Comparing Figure 6 and Figure 7, we find that in early periods, decisions are close to the first best investment. As time goes by, investments push beyond the first best level when risk 2 exposures are low, and the possibility that a risk 2 investment will never arrive starts to loom larger. The change is more dramatic in Case 2, as more room is left for the more highly profitable risk. In Case 1, the other opportunity is less important, and the firm starts investment in the profitable first risk at a relatively higher level in early periods.

We also investigate two more cases, which are slightly modified from Cases 1 and 2. In Case 3 and 4, we change the correlation from 0 to 0.8. In this circumstance, the diversification benefit becomes low. In Case 3, heavy investment is made in the

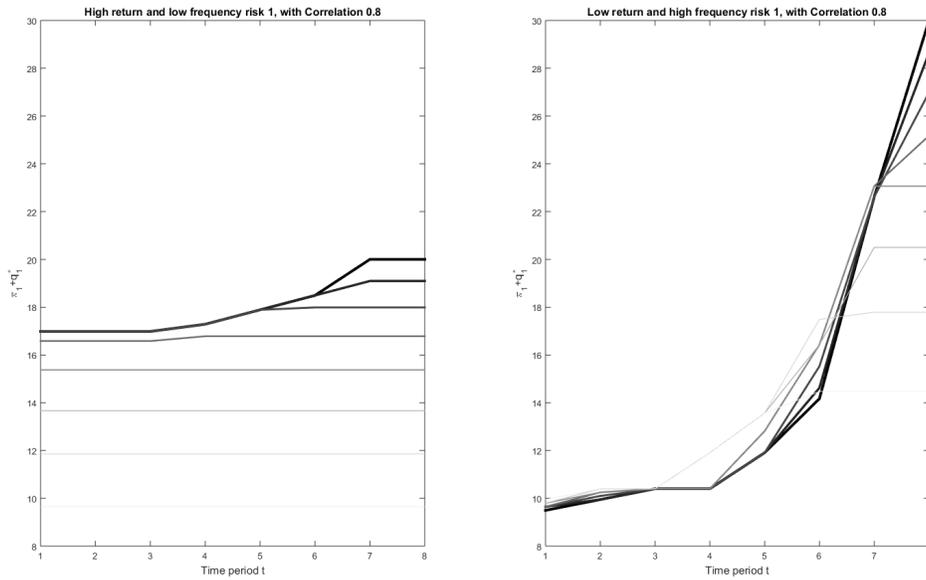


Figure 8: Decisions with various correlation, capital cost, profitability, and frequency highly profitable risk as soon as it appears; little room is left for the second risk, as it offers only modest diversification benefits. The slope of investment over time is thus slight. In contrast, Case 4 studies the decision on the less attractive risk. Significant room is left in hope for the arrival of the highly profitable risk, and the slope over time is even steeper than in Case 2.

4 Conclusion

Reinsurers, private equity firms, and other financial institutions investing in illiquid opportunities must take into account the expected risk-return profile of future opportunities when making current investments. Current capital allocation methods, which are largely developed in static full information settings, must be adjusted to this reality. RORAC evaluation of investments in this setting must consider an average of possible future capital allocations, rather than a single capital allocation based on a pro forma portfolio.

We have shown that such an adjustment is feasible if the uncertainties are understood. Capital allocation can proceed in a generalized form, with risk pricing reflecting averages of capital allocations in the future. Failure to explicitly model uncertainties about future opportunities can lead to misstatement of the marginal cost of risk, and, as a result, misallocation of capacity in current opportunities.

In numerical exercises, we show that the portfolio aspects of the problem introduce complexities into the economic assessment of this sequential investment problem. Consistent with the intuition of the real options literature, we find that the explicit consideration of uncertainty should lead firms to reduce investment in current opportunities and leave capacity for attractive future opportunities should they appear. However, diversification motivations can either reinforce the “waiting” effect or counteract it, as the firm anticipates the likely profitability of a current opportunity before a complete portfolio has been constructed. Specifically, a firm may invest more heavily in an opportunity presented early, as it will have ample chance to exploit its full potential by making use of subsequent diversification opportunities; that same opportunity attracts less interest late in the investment cycle, as the firm is no longer certain that it will be able to diversify.

Appendix - Dynamic Capital Allocation with Multiple Opportunities

Suppose the financial institution receives $n + 1$ opportunities in sequence from period 0 to period n . Upon receipt of the period-zero opportunity, it must decide on the quantity q_0 to invest in this known opportunity. Subsequent to this choice, n opportunities will arrive. The risk characteristics of the subsequent opportunities are uncertain, but the possible characteristics, as well as the probabilities associated with their appearance, are known to the decision maker. For the i th opportunity, uncertainty about the risk characteristics is resolved upon receipt of the i th opportunity, at which point the institution decides on the quantity q_i to invest in that opportunity. We denote the pool of possibilities for the i th opportunity as the set Ω_i , and denote the uncertainty in the i th period by a random variable α_i . The random variable representing all possible states of the world, covering all possible opportunities, is denoted $\alpha = \{\alpha_1, \dots, \alpha_n\}$.

As in the two-opportunity case, the cost of risk is determined by risk capital:

$$\rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha),$$

with the frictional cost of risk capital being $\tau\rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha)$. In every future state of the world, the measured risk of the firm's portfolio must not exceed a maximum level K . The complete set of possible realizations for α is represented by the set $\Omega^n = \Omega_1 \times \Omega_2 \times \dots \times \Omega_n$. The probability of any particular realization $\alpha \in \Omega^n$ is given by $p(\alpha)$.

If we consider decisionmaking at the time of the period-zero opportunity, the optimization problem is to find the optimal quantity for the period-zero investment and a set of plans—contingent on different states of the world—for all subsequent n

opportunities. Moreover, the quantity chosen for the i th opportunity must be adapted to the information available at the arrival of the i th opportunity, since the i th decision only depends on the realized opportunity in that period along with the realizations of previous periods and the decisions made up to that point. Denote the random variable reflecting possible state of the world at the arrival of the i th opportunity to be $\boldsymbol{\alpha}^i = \{\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2, \dots, \boldsymbol{\alpha}_i\}$, which takes values in the set $\Omega^i = \Omega_1 \times \Omega_2 \times \dots \times \Omega_i$, each realization having the corresponding probability $p_i(\alpha^i)$. Thus, the i th decision q_i must be a contingency plan for all possible realizations of the random variable $\boldsymbol{\alpha}^i$. Therefore, the choice variables at the decision time on the current opportunity are $\{q_0, q_i(\alpha^i), \forall i \in [1, \dots, n], \forall \alpha^i \in \Omega^i\}$. Thus, we can write the optimization problem as:

$$\max_{q_0, \{q_i(\alpha^i), \forall i \in [1, \dots, n], \forall \alpha^i \in \Omega^i\}} q_0 \left\{ \begin{array}{l} \mathbb{E}[\mathbf{r}_0] + \sum_{i=1}^n \left[\sum_{\alpha^i \in \Omega^i} p_i(\alpha^i) q_i(\alpha^i) \mathbb{E}[\mathbf{r}_i(\alpha^i)] \right] \\ -\tau \sum_{\alpha \in \Omega^n} p(\alpha) \rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha) \end{array} \right\}, \quad (14)$$

subject to a set of constraints that constrain risk for each possible realization of $\boldsymbol{\alpha}$:

$$\rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha) \leq K \quad \forall \alpha \in \Omega^n. \quad (15)$$

The first order conditions for the problem can be expressed as:

$$\mathbb{E}[\mathbf{r}_0] - \sum_{\alpha \in \Omega} (\tau p(\alpha) + \lambda(\alpha)) \frac{\partial \rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha)}{\partial q_0} = 0, \quad (16)$$

$$p_j(\alpha^j) \mathbb{E}[\mathbf{r}_j(\alpha^j)] - \sum_{\alpha \in \Omega^{-\alpha^j}} (\tau p(\alpha) + \lambda(\alpha)) \frac{\partial \rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha)}{\partial q_j(\alpha^j)} = 0, \quad (17)$$

where $\lambda(\alpha)$ is the Lagrange multiplier associated with constraint (15) for the particular state α , and $\Omega^{-\alpha^j}$ denotes the set of all possible states conditional on the

first j opportunities having been realized, in other words,

$$\Omega^{-\alpha^j} = \{\tilde{\alpha} | \tilde{\alpha}_1 = \alpha_1, \dots, \tilde{\alpha}_j = \alpha_j, \tilde{\alpha}_{j+1} \in \Omega_{j+1}, \dots, \tilde{\alpha}_n \in \Omega_n\}.$$

and equation (17) exists for all $\alpha^j \in \Omega^j, 1 \leq j \leq n$.

In particular, we obtain the first order condition for the last opportunity, when $j = n$, as follows:

$$p(\alpha)\mathbb{E}[\mathbf{r}_n(\alpha)] - (\tau p(\alpha) + \lambda(\alpha)) \frac{\partial \rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha)}{\partial q_n(\alpha)} = 0 \quad \forall \alpha \in \Omega^n.$$

Equation (16) balances the marginal benefit of exposure to the zero-period risk (the expected return) with its marginal cost, which is a risk cost relating to the risk measure constraint, while equation (17) does the same for the other risks in each of the possible states of the world given the previous path of realizations before the j th opportunity. Note in each state that $\lambda(\alpha)$ reflects the shadow price or marginal value of state-contingent risk capital (i.e., at time zero, what an additional dollar of risk capital in state α —and only in state α —would be worth to the firm, net of its associated frictional cost τ). If we assume that condition (17) always holds in the last period, we can write:

$$\lambda(\alpha) = p(\alpha) \frac{\mathbb{E}[\mathbf{r}_n(\alpha)] - \tau \frac{\partial \rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha)}{\partial q_n(\alpha)}}{\frac{\partial \rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha)}{\partial q_n(\alpha)}} \quad \forall \alpha \in \Omega^n. \quad (18)$$

Thus, the marginal cost of risk in each final state is equal to the marginal amount of risk capital consumed by the exposure (calculated as the partial derivative of the risk measure with respect to the quantity of the risk) times the value of that risk capital. In the case of exposures before the last opportunity, the marginal cost of risk at the time of their selection an *expected* marginal cost, since both the amount and the value of the risk capital that the exposure will ultimately consume depends

on the course of future opportunity realizations, as revealed in equation (16). To elaborate, we can use (17) and (18) to obtain:

$$\mathbb{E}[\mathbf{r}_j(\alpha^j)] = \mathbb{E} \left[\left(\tau + \frac{\mathbb{E}[\mathbf{r}_n(\alpha)] - \tau \frac{\partial \rho}{\partial q_n(\alpha)}}{\frac{\partial \rho}{\partial q_n(\alpha)}} \right) \frac{\partial \rho}{\partial q_j(\alpha^j)} \middle| \alpha \in \Omega^{-\alpha^j} \right]. \quad (19)$$

If the risk measure ρ is homogeneous, we get

$$\frac{\partial \rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha)}{\partial q_0} q_0 + \sum_{i=1}^n \frac{\partial \rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha)}{\partial q_i(\alpha^i)} q_i(\alpha^i) \quad (20)$$

$$= \rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha) \quad \forall \alpha \in \Omega^n, \quad (21)$$

and, thus, we have:

$$\begin{aligned} & \sum_{\alpha \in \Omega^n} (\tau p(\alpha) + \lambda(\alpha)) \left[\frac{\partial \rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha)}{\partial q_0} q_0 + \sum_{i=1}^n \frac{\partial \rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha)}{\partial q_i(\alpha^i)} q_i(\alpha^i) \right] \\ &= \sum_{\alpha \in \Omega^n} (\tau p(\alpha) + \lambda(\alpha)) \rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha) = \tau \bar{p} + \sum_{\alpha \in \Omega^n} \lambda(\alpha) K. \end{aligned}$$

Therefore, marginal cost has a risk capital allocation interpretation that “adds up” in the multiple opportunity case as well in a form similar to the two-period case. As before, if the total risk constraint *never* binds, then the sum of marginal costs times exposures contemplated reduces to $\tau \bar{p}$. In this case, total capacity K is never fully utilized, but the average risk \bar{p} is again fully allocated to each exposure through marginal cost pricing, with each unit of allocated risk being charged the frictional cost τ . On the other hand, if the total risk constraint *always* binds, then total risk costs tally up to $(\tau + \sum_{\alpha \in \Omega^n} \lambda(\alpha)) K$, which is again interpreted as a cost of capital (in parentheses) times total capital. In intermediate cases, the risk cost that is allocated is again the total risk capital times the frictional cost, $\tau \bar{p}$, plus the total shadow costs of risk capital, $\sum_{\alpha \in \Omega^n} \lambda(\alpha) K$.

If we consider the capital cost allocation to the j th risk, we have:

$$\sum_{\alpha \in \Omega^n} p_n(\alpha) \left[\tau + \frac{\mathbb{E}[\mathbf{r}_n(\alpha)] - \tau \frac{\partial \rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha)}{\partial q_n(\alpha)}}{\frac{\partial \rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha)}{\partial q_n(\alpha)}} \right] \left[\frac{\partial \rho(q_0, q_1(\alpha^1), \dots, q_n(\alpha^n); \alpha)}{\partial q_j(\alpha^j)} q_j(\alpha^j) \right].$$

For each final state of the world, we multiply the probability of the state α times the marginal value of capital in state α times the capital allocated to the j th risk in state α ; the total cost allocation is then the sum of these figures over all states in Ω^n .

Moreover, the capital allocation can be rewritten as follows:

$$\sum_{\alpha^j \in \Omega^j} \left\{ p_j(\alpha^j) \left[\sum_{\alpha \in \Omega - \alpha^j} \Pr(\alpha | \alpha^j) \left[\tau + \frac{\mathbb{E}[\mathbf{r}_n(\alpha)] - \tau \frac{\partial \rho}{\partial q_n(\alpha)}}{\frac{\partial \rho}{\partial q_n(\alpha)}} \right] \left[\frac{\partial \rho}{\partial q_j(\alpha^j)} q_j(\alpha^j) \right] \right] \right\}.$$

Given a particular path until the period j , we multiply the conditional probability of the remaining possible states times the marginal value of capital in the state times the capital allocation to the j th risk in the state. The total cost allocation for the particular path α^j is the sum of all the figures over all sub-states. The sum of all the resulting path-contingent capital allocations times the probabilities of the path is the total cost allocation.

In particular, the capital allocation to the period-zero risk can be expressed as:

$$\sum_{\alpha \in \Omega^n} p(\alpha) \left[\tau + \frac{\mathbb{E}[\mathbf{r}_n(\alpha)] - \tau \frac{\partial \rho}{\partial q_n(\alpha)}}{\frac{\partial \rho}{\partial q_n(\alpha)}} \right] \left[\frac{\partial \rho}{\partial q_0} q_0 \right].$$

As in the two opportunity case, we can recover a RORAC result:

$$\frac{\mathbb{E}[\mathbf{r}_0]}{\sum_{\alpha \in \Omega^n} p(\alpha) \frac{\partial \rho}{\partial q_0}} = \frac{\sum_{\alpha \in \Omega^n} p(\alpha) \left[\tau + \frac{\mathbb{E}[\mathbf{r}_n(\alpha)] - \tau \frac{\partial \rho}{\partial q_n(\alpha)}}{\frac{\partial \rho}{\partial q_n(\alpha)}} \right] \left[\frac{\partial \rho}{\partial q_0} \right]}{\sum_{\alpha \in \Omega^n} p(\alpha) \frac{\partial \rho}{\partial q_0}},$$

where the right-hand side reflects the expectation of the marginal value of risk capital.

As before, this will in general be larger than the carrying cost of risk capital τ , depending on the extent to which the risk constraints bind at the optimum.

References

- [1] Bauer, D. and G. Zanjani (2013), “Capital Allocation and Its Discontents” In: *Handbook of Insurance*, 2nd edition (Georges Dionne, ed.) Springer: New York, NY.
- [2] Dixit, A.K. and R.S. Pindyck (1994), *Investment Under Uncertainty*, Princeton University Press: Princeton, NJ.
- [3] Embrechts, P., R. Frey, and A. McNeil (2005), *Quantitative Risk Management*, Princeton Series in Finance, Princeton University Press: Princeton, NJ.
- [4] Kalkbrener, M. (2005), “An Axiomatic Approach to Capital Allocation” *Mathematical Finance* 15, 425-37.
- [5] Laeven, R.J. and M.J. Goovaerts (2004), “An Optimization Approach to the Dynamic Allocation of Economic Capital” *Insurance: Mathematics and Economics* 35, 299-319.
- [6] McDonald, R.L. and D. Siegel (1986), “The Value of Waiting to Invest” *Quarterly Journal of Economics* 101, 707-727.
- [7] Tsanakas, A. (2004), “Dynamic Capital Allocation with Distortion Risk Measures” *Insurance: Mathematics and Economics* 35, 223-243.