

NBER WORKING PAPER SERIES

**RISK MANAGEMENT, CAPITAL
BUDGETING AND CAPITAL STRUCTURE
POLICY FOR FINANCIAL INSTITUTIONS:
AN INTEGRATED APPROACH**

**Kenneth A. Froot
Jeremy C. Stein**

Working Paper 5403

**NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 1996**

This research is supported by the Global Financial System Project at HBS (Froot), the International Financial Services Research Center at MIT (Stein), and the National Science Foundation (Stein). Thanks to Maureen O'Donnell for help in preparing the manuscript, and to Fischer Black, Bob Merton, Andre Perold, Rene Stulz and seminar participants at NYU and the NBER for helpful comments. This paper is part of NBER's research programs in Asset Pricing and Corporate Finance. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

© 1996 by Kenneth A. Froot and Jeremy C. Stein. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

**RISK MANAGEMENT, CAPITAL
BUDGETING AND CAPITAL STRUCTURE
POLICY FOR FINANCIAL INSTITUTIONS:
AN INTEGRATED APPROACH**

ABSTRACT

We develop a framework for analyzing the capital allocation and capital structure decisions facing financial institutions such as banks. Our model incorporates two key features: i) value-maximizing banks have a well-founded concern with risk management; and ii) not all the risks they face can be frictionlessly hedged in the capital market. This approach allows us to show how bank-level risk management considerations should factor into the pricing of those risks that cannot be easily hedged. We examine several applications, including the evaluation of proprietary trading operations, and the pricing of unhedgeable derivatives positions.

**Kenneth A. Froot
Harvard Business School
Morgan Hall-390
Soldiers Field
Boston, MA 02163
and NBER**

**Jeremy C. Stein
Massachusetts Institute of Technology
E52-448, Sloan School of Management
50 Memorial Drive
Cambridge, MA 02142-1347
and NBER**

I. Introduction

One of the fundamental roles of banks and other financial intermediaries is to invest in illiquid financial assets--assets which, because of their information-intensive nature, cannot be frictionlessly traded in the capital markets. The standard example of such an illiquid asset is a loan to a small or medium-sized company. A more modern example is the credit-risk component of a foreign exchange swap. Even if the currency risk inherent in the swap can be easily laid off by the dealer bank, the same is not likely true of the credit risk.

At the same time that they are investing in illiquid assets, most banks also appear to engage in active risk management programs. Holding fixed its capital structure, there are two broad ways in which a bank can control its exposure to risk. First, some risks can be offset simply via hedging transactions in the capital market. Second, for those risks where direct hedging transactions are not feasible, the other way for the bank to control its exposure is by altering its investment policies. Therefore, with illiquid risks, the bank's capital budgeting and risk management functions become linked.

To see this point more clearly, return to the example of the foreign exchange swap. If a dealer bank is considering entering into such a transaction, its own aversion to currency risk should not enter into the decision of whether or not to proceed. After all, if it doesn't like the currency risk embodied in the swap, it can always unload this risk in the market on fair terms. Thus with respect to the tradeable currency risk, the risk management and investment decisions are separable. The same is not true, however, with respect to the illiquid credit-risk component of the swap. If the bank is averse to this risk, the only way to avoid it is by not entering into the swap in the first place.

This reasoning suggests that if the bank is asked to bid on the swap, its pricing should have the following properties. First, the pricing of the swap should be independent of the bank's own attitudes toward currency risk--the bank should evaluate currency risk just like any other market participant, based only on the risk's correlation with systematic factors that are priced in the capital market. Second, however, the swap's pricing should depend on the bank's own attitude toward the credit risk. Thus if the bank already has a portfolio of very highly correlated credit risks, it might bid less aggressively for the swap than another institution with a different balance sheet, all else equal. This should hold true even if the credit risk is uncorrelated with factors that are priced in the capital market.

Although this sort of approach to the pricing of bank products may sound intuitively reasonable, it differs substantially from the dominant paradigm in the academic literature, which is based on the classical finance assumptions of frictionless trading and absence of arbitrage. In the specific case of pricing the credit risk on a swap, the classical method boils down to a contingent-claims model of the sort pioneered by Merton (1974).¹ This type of model--like any classical pricing technique--has the implication that the correct price for the swap is the same for any dealer bank, independent of the bank's pre-existing portfolio. Of course, this is because the classical approach by its very nature assumes away exactly the sorts of imperfections that make the bank's problem challenging and relevant. Indeed, it is only appropriate if either: i) the bank can frictionlessly hedge all risks--including credit risks--in the capital market; or ii) the Modigliani-Miller theorem applies, so that the bank has no reason to care about risk

¹For recent applications to the pricing of credit risk on derivatives, see, e.g., Cooper and Mello (1991) and Jarrow and Turnbull (1995).

management.

Perhaps because the classical finance approach does not speak to their concerns with risk management, practitioners have developed alternative techniques for capital budgeting. One common method uses the ratio of the return on an investment to some measure of the capital allocated to that investment. For example, the RAROC (risk-adjusted return on capital) method divides returns by "capital at risk", often defined as the investment's maximum loss within a prespecified confidence interval that is computed based on total volatility. Thus in contrast to the classical approach, RAROC discounts individual investments according to their total risk, rather than their priced systematic risk. So, for the foreign-exchange swap, RAROC would put a non-zero price on its credit risk, regardless of the credit risk's correlation with priced factors.

Although the RAROC method has some intuitive appeal, it is not clear that it is the optimal technique for dealing with the sorts of capital budgeting problems facing financial institutions. That is, RAROC--as currently applied--is not derived from first principles to address the objective of shareholder value maximization. Consequently, it has some features that might be considered troublesome, and it leaves other issues potentially unresolved. For example, why is "capital at risk" a function of an investment's volatility, as opposed to some sort of covariance measure? And once one has determined the amount of capital at risk, what is the right cost to assign to this capital? Should the cost of capital at risk depend on the strength of the bank's balance sheet, or other related variables?

Our primary goal in this paper is to develop a conceptual framework for capital budgeting that blends some of the most desirable features of both the classical approach and the RAROC-style bank-practitioner approach. To accomplish this goal, we build a model that--like in the

classical approach--is squarely rooted in the objective of maximizing shareholder value in an efficient market. However, the model also incorporates two other key features: i) there is a well-founded concern with risk management; and ii) not all risks can be hedged in the capital market. This allows us to capture the important intuition that bank-level risk-management considerations should enter into the pricing of those risks that cannot be hedged. In principle, this framework can be applied to any company attempting to manage the risks associated with illiquid assets. Nonetheless, we think it is particularly well-motivated by the sorts of financial-industry problems discussed above.²

In order to endogenize banks' concern with risk management, we assume that there are increasing costs to raising new external funds. Thus if a bank were to be hit with a negative shock that depleted its capital, it would incur some costs in rebuilding its balance sheet. In an effort to avoid these costs, the bank will behave in a risk averse fashion. This basic rationale for risk management is essentially the same as that presented by Froot, Scharfstein and Stein (1993) in the context of non-financial corporations. It is also closely related to the banking models of Kashyap and Stein (1995), Stein (1995), and Greenwald, Levinson and Stiglitz (1991), all of which emphasize the costs that banks face in raising non-deposit external finance.³

²This is because getting the cost of capital right for any given instrument is very likely to be a first-order consideration for a financial institution. In contrast, for many industrial companies doing capital budgeting, the uncertainties associated with projecting cashflows on their physical assets may be so large as to swamp any modifications in discount rates that our method might suggest.

³In what follows, we ignore any risk-taking incentives that might arise with government-provided deposit insurance. Thus our model is most literally applicable either to financial institutions that are not insured commercial banks (e.g. investment banks like Salomon Brothers) or to commercial banks that are sufficiently well-capitalized that one can for practical purposes safely ignore the incentive effects of deposit insurance. (e.g., JP Morgan, which is AAA-rated.)

As will become clear, one key feature of this modelling approach is that it highlights a tradeoff between: i) managing risk via ex ante capital structure policy; vs. ii) managing risk via capital budgeting and hedging policies. Aside from engaging in hedging transactions, a bank has two other methods for controlling the risk of being caught short of funds. First, it can adopt a very conservative capital structure. Or second, it can invest less aggressively in (i.e., charge a higher price for bearing) non-hedgeable risks. If there are no costs to holding a lot of capital, this will be the preferred way of dealing with the problem. In the limit when the bank holds a very large capital buffer, risk-management concerns will no longer enter investment decisions, and the model will converge back to the classical paradigm. In contrast, when holding capital is costly (e.g., due to tax or agency effects) risk-management concerns will have a meaningful impact on capital budgeting policies. The bottom line is that in our framework, optimal hedging, capital budgeting and capital structure policies are jointly and endogenously determined.⁴

The remainder of the paper is organized as follows. In Section II, we lay out the basic timing and assumptions of our model. In Section III, we analyze the model's implications for hedging, capital budgeting and capital structure decisions. Section IV gives several examples of how the basic framework can be applied to different sorts of capital budgeting problems facing financial institutions. Section V concludes.

⁴Other papers which investigate the linkages between capital budgeting and risk management include Merton and Perold (1993), and Merton (1995a, b).

II. The Model: Timing and Assumptions

The model has three time periods, 0, 1, and 2. In the first two periods, the bank chooses its capital structure, and then makes capital budgeting and hedging decisions. These two periods are at the heart of our analysis. The last period is needed to close the model--to give the bank a well-founded objective function that incorporates both shareholder-value maximization as well as a concern for risk management.

A. Time 0: bank chooses its capital structure

The bank enters time 0 with an initial portfolio of exposures. This portfolio will result in a time-2 random payoff of Z_p . The random variable Z_p is normal, and can be written as $Z_p = \mu_p + \epsilon_p$, where μ_p is the mean and ϵ_p is a mean-zero disturbance term. For simplicity, we assume that this portfolio of exposures requires no net financing--i.e., it is a zero-net-wealth portfolio.

The only decision facing the bank at time 0 is how much equity capital to hold. Specifically, the bank can raise an amount of capital K , and invest the proceeds in riskless Treasury bills. Holding "financial slack" in this manner involves direct deadweight costs. These costs might in principle arise from a number of sources; for concreteness, it is useful to think of them as being driven by taxes. Thus the deadweight costs of holding an amount of capital K are given by τK , where τ is the effective net tax on cash holdings.⁵

Although it involves deadweight costs, we will show below that banks typically opt to

⁵The tax cost of holding equity-financed slack is just the mirror image of the tax advantage of debt finance.

hold non-zero levels of capital. This is because holding capital allows banks to tolerate risks better, and thereby price their products more aggressively.

To summarize the set-up at time 0: On the balance sheet, the bank holds cash of K , which is entirely equity-financed. In addition, the bank also has an "off-balance-sheet" (in the sense of requiring no net financing) portfolio of exposures of Z_p .

One point here deserves further comment. Although we have assumed that there are costs associated with a given stock of capital K , there is no sense in which banks have trouble adjusting K at time 0. This runs counter to the spirit of many models of financing under asymmetric information (e.g., Myers and Majluf 1984), where, loosely speaking, costs are incurred not by having equity on the balance sheet per se, but rather by having to raise new external funds.⁶ Indeed, we will assume momentarily that there are exactly such flow costs of new external finance at time 2. Moreover, these flow costs will be a convex function of the amount raised, and they will be the driving force behind the bank's desire to manage risk.

Thus it is clearly a shortcut to ignore the potential for convex costs of external financing at time 0. We do so in order to better focus on the question we ultimately wish to address with the time-0 analysis: what is the appropriate long-run "target" capital structure for the bank? In other words, how should the bank be seeking to position itself over the long haul--as a AAA credit, a BBB credit, or something in between? We recognize that, if at any point in time, the bank is far away from its ideal target, it may face costs of adjustment in getting to the target quickly, but it is nonetheless interesting to ask the question of what the target should be.

⁶In other words, our time-0 analysis of bank capital structure is of the static-tradeoff variety, as opposed to the dynamic pecking-order behavior implied by models of asymmetric information.

B. Time 1: bank invests in new products and makes hedging decisions

At time 1, the bank faces two decisions. First, it has the opportunity to invest in a new product.⁷ The new product offers a random payoff of Z_N at time 2. This payoff is normal, and can be written as $Z_N = \mu_N + \epsilon_N$, where μ_N is the mean and ϵ_N is a mean-zero disturbance term. The magnitude of the bank's exposure to the new product is a choice variable, and is given by α . As before, it is assumed for simplicity that the exposure to the new product does not require any cash to be put up at time 1. For example, the exposure may represent the assumption of a forward position where no money changes hands at the time the position is put on.⁸

The second decision to be made at time 1 is how to hedge both the initial and new exposures. We will have much more to say about the hedging technology later on; for now we just establish a piece of notation by defining the aggregate time-2 hedging position as H , and the associated payoff on this hedging position as Z_H .

Given the assumptions that have been made so far, the bank's realized internal wealth

⁷We begin by focusing on the case where there is only one new product at time 1 for simplicity. However, it is easy to generalize the results to the case of multiple new products, as we do below.

⁸Given that we are working with zero-net-wealth forward positions, it is more natural to specify the dollar payoff on these positions, as opposed to a percentage rate of return. However, it is easy to reinterpret our notation in terms of the implied percentage returns on the associated underlying assets. Think of the forward position as being on an underlying asset whose time-1 price is normalized to one. Then if z is the rate of return on the underlying asset, the dollar payoff on the forward Z that we are defining satisfies $Z = z - r$, where r is the riskless rate of interest between time 1 and time 2. Thus all the statements we make below about required dollar payoffs on various forward positions can be trivially reinterpreted as statements about required rates of return on the underlying assets, simply by adding r .

at time 2, w , is given by:

$$w = Z_p + \alpha Z_N + Z_H + K(1 - \tau) \quad (1)$$

In words, the amount of cash the bank has on hand at time 2 will depend on the realizations on its old exposures, on its new product, and on its hedging positions, as well as on the amount of capital raised at time 0.

C. Time 2: bank reacts to its cashflow realization

As noted above, we need to build into the model a reason for the bank to care about risk management--that is, to care about the distribution of w . To do so, we follow closely Froot, Scharfstein and Stein (1993), henceforth FSS. In particular, we assume that after w is realized, the bank has a further, non-stochastic investment opportunity--e.g., it might be able to extend some new loans. This investment requires a cash input of I , and yields a gross return of $F(I)$, where $F(I)$ is an increasing, concave function. The investment can either be funded out of internal sources, or funds can be raised externally in an amount e . Thus $I = w + e$. The hitch is that there are convex costs to raising external finance; these costs are given by $C(e)$.⁹

Denote by $P(w)$ the solution to the bank's time-2 problem:

⁹In FSS, we give a number of microeconomic rationales--based on agency and/or information problems--to justify this sort of specification for the $C(e)$ function. We also show explicitly how such a convex functional form arises in a specific optimal contracting setting, a variant of the costly-state verification model due to Townsend (1979) and Gale-Hellwig (1985). Moreover, Stein (1995) generates a similar formulation in a banking model where non-deposit liabilities are subject to adverse selection problems.

$$P(w) = \max F(I) - I - C(e), \text{ subject to } I = w + e. \quad (2)$$

FSS demonstrate that $P(w)$ is in general, an increasing concave function, so that $P_w > 0$, and $P_{ww} < 0$. It is the concavity of the $P(w)$ function that generates a rationale for risk management. This concavity in turn arises from the convexity of $C(e)$, interacting with the concavity of $F(I)$. (See FSS for more explicit details.) Loosely speaking, the more difficult it is for the bank to raise external funds on short notice at time 2, the more risk averse it will be with respect to fluctuations in its time-2 internal wealth w .

D. Some observations about the structure of the model

If one were to leave time 1 out of the model, and keep only the parts corresponding to time 0 and time 2, we would be left with a very standard pecking-order-type story of corporate investment and financing. At time 2, the presence of increasing costs of external finance can lead to underinvestment--i.e., a level of I that is less than the first-best. This distortion can be partially alleviated to the extent that financial slack can be built up at time 0. Thus it may be desirable to hold such slack on the balance sheet, even if there are some costs to doing so. This is very much in the spirit of Myers (1984), Myers and Majluf (1984), and the large literature that has followed these papers.

What we have added to this basic pecking-order story are the ingredients that come into play at time 1. Specifically, the bank now has two other tools at its disposal--beyond simply holding financial slack--that it can use to offset the underinvestment distortions caused by costly

external finance: i) the bank can adjust its risk exposure by undertaking hedging transactions H ; and ii) the bank can also adjust its risk exposure by varying the amount α it invests in the new product. Thus overall, the bank optimizes by picking the right combination of three policy variables: K , H and α . It is in this sense that capital structure, hedging and capital budgeting decisions are linked to one another.

III. Analysis of the Model

To understand the properties of the model, it is easiest to work backwards. We have already seen that any given realization of w at time 2 can be mapped into a non-stochastic payoff $P(w)$, once we specify $F(I)$ and $C(e)$. The next step is to ask: from the perspective of time 1, when w is uncertain, what is the market value of the bank? And given this valuation function, how should a bank that seeks to maximize its value set its hedging and capital budgeting policies at time 1--i.e, how should it choose H and α ?

A. Valuing the bank at time 1

From the perspective of time 1, the ultimate payoff $P(w)$ to a bank shareholder is a random variable. In order to value this random variable, we need a pricing model. Without any real loss of generality, we can assume that asset prices are determined by a simple one-factor model. In this setting, the value of the bank's shares V , will be:

$$V = \{EP(w) - \gamma \text{cov}(P(w), M)\} / (1+r) \quad (3)$$

where M (for "market") is the one priced factor, γ is the equilibrium excess return per unit of variance for bearing M risk, and r is the riskless rate of interest between time 1 and time 2.

B. Optimal hedging policy at time 1

Now suppose the bank designs its risk management policy so as to maximize shareholder value V . What should it do? To build intuition, let us first consider a simple case where all the bank's risks are "perfectly tradeable". What we mean by this is that these risks can be frictionlessly unloaded in the market, on terms that are just fair, given the correlation of the risks with the priced factor.

In the appendix, we prove the following:

Proposition 1: In the case where all risks are perfectly tradeable, the bank maximizes value by hedging completely. That is, it picks its hedging position H so that the payoff on the hedging position $Z_H = -\epsilon_P - \alpha\epsilon_N + k$, where k is a constant.

This result is fairly intuitive. It is a generalization of a result presented in FSS, where we considered the more restrictive case in which investors in the capital market were risk-neutral and hence where there were no priced factors. The one subtlety is that while it is tempting to conclude from the proposition that the bank simply behaves "like a risk-averse individual", this is not quite correct. In general, a risk-averse individual will not wish to completely shun systematic risk, as this involves a reduction in expected return. Rather, an individual will typically opt for an interior solution in which he bears some systematic risk. However, this is

not true for a publicly traded bank in our set-up. A bank does not reduce shareholder value by sacrificing return in exchange for a reduction in risk, so long as the terms of trade are set in an efficient market--i.e., so long as the hedging transactions are zero-NPV. Since there is no cost to reducing risk, we are left with the pure effect that, because of the concavity of $P(w)$, risk reduction on fair terms is always desirable.

Of course, by assuming that all risks can be frictionlessly sold off in the capital market, we have trivialized the bank's problem. As was stressed in the Introduction, the very existence of intermediaries such as banks is testimony to the fact that certain risks are somewhat information-intensive, and hence cannot be traded perfectly liquidly. To take a first cut at capturing this notion, we make the following decomposition. We assume that a bank's exposures can be classified into two categories: i) perfectly tradeable exposures, which as above, can be unloaded frictionlessly on fair-market terms; and ii) completely non-tradeable exposures, which must be retained by the bank no matter what.

In terms of our previous notation, this amounts to decomposing both ϵ_P and ϵ_N as follows:

$$\epsilon_P = \epsilon_P^T + \epsilon_P^N \tag{4}$$

$$\epsilon_N = \epsilon_N^T + \epsilon_N^N \tag{5}$$

where ϵ_P^T is the tradeable component of ϵ_P , ϵ_P^N is the non-tradeable component, and so forth.

We assume that the priced factor M is fully tradeable, so $\text{cov}(\epsilon_P^N, M) = \text{cov}(\epsilon_N^N, M) = 0$.¹⁰

There are a number of different examples that help illustrate what we have in mind with this decomposition, and we will develop several of these examples in more detail shortly. For the time being, it may be helpful for concreteness to think of the new product as being an investment in a company whose stock is not publicly traded.¹¹ Clearly, some of the risk associated with such an investment may be tradeable, to the extent that it is correlated with, say, macroeconomic conditions and hence can be hedged with some sort of contract on an aggregate variable. However, some of the idiosyncratic exposure associated with the investment cannot be laid off, at least not frictionlessly. This is what we are trying to capture.

In this environment, a simple extension of Proposition 1 can be proven:

Proposition 2: A bank will always wish to fully hedge its exposure to any tradeable risks. That is, it picks its hedging position H so that the payoff on the hedging position $Z_H = -\epsilon_P^T - \alpha \epsilon_N^T + k$, where k is a constant.

C. Capital budgeting policy at time 1

1. The case of a single investment decision

Now that we have established how the bank sets its hedging policy, we can turn to the

¹⁰It is straightforward to relax the assumption that M is tradeable. We discuss this briefly below, in Section IVB.

¹¹To be absolutely literal in terms of our notation, the investment should be thought of as a forward position in the non-traded company. However, as mentioned in footnote 5 above, it is trivial to reinterpret our results in terms of required returns on the underlying assets.

capital budgeting question, namely: what should the bank's desired investment α in the new product be at time 1, and how does this investment depend on the new product's expected return and its risk characteristics?

To attack this question, we begin by noting that, given the results of Proposition 2 on optimal hedging policy, we can rewrite time-2 wealth w as:

$$w = \mu_p + \epsilon_p^N + \alpha(\mu_N + \epsilon_N^N) + k + K(1-\tau) \quad (6)$$

Equation (6) reflects the fact that all the tradeable risk components have been hedged out of w , leaving only the non-tradeable components. More structure can be put on this expression by pinning down k , which is the expected return on the hedging transactions. Given that these transactions are zero-NPV, the expected return just offsets the systematic risk, so:

$$k = -\gamma(\text{cov}(\epsilon_p^T, M) + \alpha\text{cov}(\epsilon_N^T, M)) \quad (7)$$

Substituting (7) into (6), we obtain:

$$w = \mu_p + \epsilon_p^N + \alpha(\mu_N + \epsilon_N^N) - \gamma(\text{cov}(\epsilon_p^T, M) + \alpha\text{cov}(\epsilon_N^T, M)) + K(1-\tau) \quad (8)$$

Thus the bank's objective function is to maximize its value V as given by (3), subject to the constraint that w satisfy (8). To keep things simple, we assume that the bank views itself as a price-taker--i.e., it takes μ_N as a fixed parameter, and chooses the quantity invested α

accordingly.¹² Thus the first order condition is simply that $dV/d\alpha = 0$. In the appendix, we show that this condition reduces to:¹³

$$\alpha^* = \{(\mu_N - \gamma \text{COV}(\epsilon_N^T, M)) - G \text{COV}(\epsilon_N^N, \epsilon_P^N)\} / G \text{var}(\epsilon_N^N) \quad (9)$$

where $G \equiv -EP_{ww}/EP_w$ is a measure of the bank's effective risk aversion.

Unlike in the case of an individual decisionmaker, the bank's risk aversion G is an endogenous variable. In particular, G will depend on the amount of capital K that the bank holds. It is easy to see that in the limiting case where K becomes infinitely large, G converges to zero, for any arbitrary specification of the underlying $C(e)$ function. In other words, with infinite capital, the bank becomes risk-neutral, as in the classical setting. This is because the probability of it ever having to seek costly external finance falls to zero. Moreover, it can be shown that in the sort of optimal contracting set-up considered in FSS, this convergence is always monotonic--i.e., $dG/dK \leq 0$ everywhere. We will assume that this "declining risk aversion" property holds in what follows.

Note that in the polar case where $G = 0$, the bank will invest an infinite amount in anything with a return that exceeds the market risk premium. However, when $G > 0$, and investment requires the assumption of non-tradeable risk, the bank will be more conservative. The greater is the contribution of the new non-tradeable risk to the variance of the bank's overall

¹²It is straightforward to generalize to the case where the bank faces less-than-perfectly elastic demand for its products.

¹³ In doing so, we assume that the second-order conditions with respect to α are satisfied.

portfolio of non-tradeable risk, the more pronounced is the conservatism.

When $G > 0$, the investment behavior embodied in (9) cannot in general be summarized in terms of a single hurdle rate, as is the common convention in most corporate capital budgeting applications. This is because the bank's risk aversion makes the required return an increasing function of the amount invested in the new product. However, in the limiting case where α goes to zero, there is a simple hurdle rate representation. We can interpret this case as one where the bank must make an accept/reject decision on an investment opportunity of small fixed size. For such an investment, the hurdle rate is given by:

$$\mu_N^* = \gamma \text{COV}(\epsilon_N^T, M) + G \text{COV}(\epsilon_N^N, \epsilon_P^N) \quad (10)$$

Equation (10) can be thought of as a simple two-factor pricing model. The first factor is the standard market-risk factor. The novel twist lies in the addition of the second factor. To the extent that a bank takes on a non-tradeable risk, this risk should be priced based on its correlation with the bank's pre-existing portfolio of non-tradeable risks ϵ_P^N . Moreover, the unit price of non-tradeable risk is given by G . Thus ultimately, the extent to which non-tradeable risk gets priced depends on the factors which shape the curvature the $P(w)$ function; namely the $F(I)$ and $C(e)$ functions, as well as the initial amount of capital K held by the bank.

2. Multiple investments, interdependencies, and decentralization

Thus far, we have assumed that the bank only has the opportunity to invest in one new product at time 0. However, it is straightforward to extend the results to the case where the

bank can invest in multiple new products simultaneously. Suppose there are n new products, indexed by i . To streamline the notation slightly, it will be helpful to work with the expected payoff on each product in excess of the market risk premium. Thus we define:

$\pi_i \equiv \mu_{Ni} - \gamma \text{COV}(\epsilon_{Ni}^T, M)$. The first order conditions now become:

$$\alpha_i^* = \{\pi_i - \text{GCOV}(\epsilon_{Ni}^N, \epsilon_P^N + \sum \alpha_j^* \epsilon_{Nj}^N)\} / \text{Gvar}(\epsilon_{Ni}^N), \quad (11)$$

for all i and $j \neq i$. Equation (11) is of exactly the same form as (9). The only difference is that the appeal of the i th new product now depends not only on its covariance with the bank's pre-existing non-tradeable exposures, but also on its covariance with any other new non-tradeable exposures that are taken on at time 1. There is one condition for each new product, and these conditions must be solved simultaneously to yield the optimal portfolio mix of new products.

The solution for this set of equations is just:

$$\alpha^* = \Omega^{-1} \{\pi - \text{GC}_{NP}\} / G \quad (12)$$

where α^* and π are now both $n \times 1$ vectors, Ω is the $n \times n$ variance-covariance matrix for the non-tradeable exposures on the new products (i.e., $\Omega_{ij} = \text{COV}(\epsilon_{Ni}^N, \epsilon_{Nj}^N)$), and C_{NP} is an $n \times 1$ vector whose i th element is $\text{COV}(\epsilon_{Ni}^N, \epsilon_P^N)$.

Equation (12) can be given a simple interpretation. The term $\pi - \text{GC}_{NP}$ can be thought of as a vector of net returns for the i products, where the netting takes into account both the market risk of these products and their covariance with the bank's pre-existing portfolio. Once

this netting has been done, the desired mix of the new products follows from a standard mean-variance optimization, using the net returns as the means.

Equations (11) and (12) make it clear that things are more complicated in this multi-investment setting than in the usual corporate capital budgeting framework. In the usual framework, investment decisions are independent of one another--holding fixed their cashflows, the appeal of project i does not depend on whether or not project j is undertaken. Here, this no longer holds true. In fact, there are two distinct sources of interdependence.

First, there is what might be termed a "covariance spillover" effect. Holding fixed G , investment in any product i will be less (more) attractive to the extent that there is also a significant investment made in another product j with positively (negatively) correlated non-tradeable risk. Thus with non-zero covariances across the new products, α_i typically depends not only on own-project return π_i , but also on the π 's of all the other products.

Second, and somewhat more subtly, there is what might be termed a "bank-wide cost of capital" effect. Even if all the covariances across the new products are zero--i.e., Ω is a diagonal matrix--investment decisions are in general interdependent because they can all influence the value of G . Thus for example, if the bank takes a large, very risky position in one product, even if this position is completely orthogonal to all others, this might raise G and thereby make the bank less willing to take on any other risks.

These interdependencies imply that in order for the bank to make optimal investment decisions, these decisions must be centralized. If one thinks of individual product managers as observing the π 's of their own products but not of others, one cannot simply delegate the investment decisions to these managers, even under the strong assumption that there are no

agency problems and that the managers would therefore act in the bank's best interests with the information that they have. Rather, the information of the individual managers must be pooled.

As a practical matter, however, such centralized decisionmaking may present its own set of difficulties. This is likely to be especially true when decisions are made at very high frequencies, in which case the costs and delays associated with transmitting new information to headquarters each time an investment is considered made may be prohibitive. To take an extreme example, think of a bank with several hundred different traders who reevaluate their positions on an almost continuous basis. Clearly, in this sort of polar situation, complete centralization of decisionmaking is impossible.

This raises the question of whether one can approximate the full-information centralized solution in a decentralized setting where individual product managers cannot condition on the contemporaneous π 's of other new products. Analogous to the one-new-product case, things become simpler if one is willing to entertain the limiting case where the α 's approach zero. In this case, equation (12) reduces to $\pi^* = GC_{NP}$, which is just a vector version of equation (10). In other words, in the limiting case where all the investments in the new products are small, one can use the same two-factor hurdle rate approach independently for multiple investments that one would use if there was only a single investment. Of course to the extent that the investments in the new products are not small, this decentralized approximation will be an imperfect one.

Overall, this line of reasoning suggests that the right question is not whether or not the bank should centralize its decisionmaking, but rather how often headquarters should gather information and use this pooled information to help guide investment decisions. Loosely speaking, what we have in mind is a dynamic version of the model wherein each time

headquarters gathers information, it can update its estimate of both G , and the stochastic characteristics of the pre-existing portfolio. These updated estimates can then be passed back to individual product managers, who will use them to form hurdle rates and thereby do their best to approximate optimal incremental investment decisions on a decentralized basis over the interval of time before the next round of information-pooling.

Although we have not analyzed such a dynamic model formally, we suspect that following basic tradeoff would emerge: on the one hand, shortening the interval between rounds of information-pooling should lead to smaller deviations from the full-information centralized solution. On the other hand, this will also clearly increase the costs of information transmission. The task is then to properly balance these two competing considerations.¹⁴

D. Optimal capital structure at time 0

We are now in a position to fold back to time 0 and solve for the optimal capital level K . There is a simple tradeoff at work: on the one hand, as noted above, a higher K reduces the bank's effective risk aversion G . From an ex ante perspective, this allows the bank to invest more aggressively in products that promise an above-market return at time 1--i.e., products with high values of π . On the other hand, a higher K also involves deadweight costs of τK .

To illustrate the first part of the tradeoff most transparently, consider a simple one-

¹⁴Again, we should emphasize that this informal story completely ignores any agency issues associated with delegating investment decisions to individual product managers. In reality, these considerations are likely to be very important. For example, for a given information set, a product manager may have a tendency to take what the bank would view as excessive risks, because his reward structure is inherently a convex function of outcomes. In this case, decentralization may involve not only setting appropriate prices--i.e., hurdle rates-- but also imposing position limits or capital constraints on individual managers.

product case where the investment in question is a small one. In this setting, a natural question to ask is how the bank's hurdle rate--as given by equation (10)--changes with K:

$$d\mu_N^*/dK = \text{cov}(\epsilon_N^N, \epsilon_P^N)dG/dK = (\pi_N^*/G) dG/dK \quad (13)$$

Since dG/dK is negative, equation (13) says that if the hurdle rate is initially above the market-required return--i.e., if $\pi_N^* > 0$ --the hurdle rate will fall smoothly as the amount of capital K is increased.

As of time 0, the bank's objective function is to pick K so as to maximize $V - K$, recognizing that $V = V(w(\alpha^*(K), K))$. In words, K affects w directly through the amount of financial slack that will be available at time 2, as well as indirectly through its influence on the optimal investment strategy α^* . One can use the envelope theorem to show that the solution to this problem can be written simply as:

$$EP_w = 1/(1-\tau) \quad (14)$$

Equation (14) has an intuitive interpretation. The bank has two choices: it can hold more slack K as a buffer at time 0, or it can be forced (in an expected sense) to seek more financing later on, with the attendant costs. It should optimally set K so that the expected shadow value of external funds at time 2 just balances the cost of holding more capital at time 0. In the limiting case where $\tau = 0$, so that there are no deadweight costs of holding capital, the bank holds an arbitrarily large amount. This drives the expected shadow value of external funds EP_w

to one, which in turn implies that G converges to 0. Thus the bank behaves in a classical manner, doing capital budgeting according to a purely market-based model of risk and return. In contrast, as τ increases above 0, the bank holds less capital, thereby raising its effective risk aversion G , and amplifying the deviations from textbook capital budgeting principles.

IV. Examples

A. A private investment/lending group

As noted above, the most literal interpretation of our tradeable/non-tradeable risk decomposition would be to think of a group in a bank that invests in non-public companies, or slightly more generally, companies whose stock is sufficiently illiquid as to preclude easy hedging of the bank's position. To take a concrete illustration, a merchant banking group may have the opportunity to invest in a portion of the equity in a private company. How much should the group be willing to pay for an stake of a given size? Or equivalently, for a given price, how large a stake should the group be willing to take?

As in any valuation exercise of this sort, the group will have to do a discounted cash flow analysis of the company in question. The cashflow projections will be done in the usual manner. However, our results suggest that if risk management is a serious concern for the group's parent bank, the discount rates should differ from those used in a classical setting. Specifically, to calculate an appropriate discount rate, the group should start with the classical rate (e.g., a CAPM-based weighted average cost of capital) and then add a premium π given by the appropriate variant of equations (9)-(12).

For example, suppose the purchase price for the company in question has already been

set at \$1000, and the group is deciding how many shares it wants to bid for. Its internal projections suggest that the company will have level cashflows of \$200/year forever--i.e., it perceives the internal rate of return on the investment as being 20%. Moreover, the group's best estimate of the classical discount rate for this company is 17%. Thus from the group's perspective, the deal offers a return premium π of 3%. If this is the only decision facing the bank at this time, one can now determine the optimal stake α in the company from equation (9).

B. A proprietary trading desk

A slightly less literal, but nonetheless useful application of our framework is to a proprietary trading operation located inside a larger financial institution. For the time being, suppose we are thinking of a desk that trades actively in simple "linear" instruments such as futures and forwards. At first glance, it might appear that our approach would be of little use in thinking about such a desk. To the extent that all the instruments it deals in a relatively liquid, it can in principle hedge any risk it faces, and therefore our tradeable/non-tradeable risk decomposition would seem to have little bite.

However, one needs to be a bit careful with the interpretation of the words "non-tradeable". Even if all the risks facing the desk are hedgeable in principle, this obviously cannot be what the desk does in practice--if it did hedge out all of its risks, it would have no business. In other words, being a trading desk by definition requires intentionally assuming certain exposures. Ostensibly, these exposures are justified by the desk's ability to earn a positive return on average, even after adjusting for market-wide risk factors. The presence of such positive (subjective) risk-adjusted returns makes such exposures "non-tradeable" in our sense.

Seen in this light, our framework can be helpful in thinking about two closely related questions facing the managers of a trading desk: First, there is the ex ante capital budgeting question: given a particular directional "view" about an asset, how aggressively should the desk invest in that asset? Second, there is the ex post performance measurement question: how can one evaluate whether the desk made enough of a profit to compensate for the risks it imposed on the bank as a whole?

In the case where the trading desk hedges out all the risk associated with the priced market factor M , our previous results apply directly, so that one can again use the appropriate variant of equations (9)-(12) to answer both of these questions. For example, if the desk is making a go/no-go decision on a trade ϵ_N of relatively small size that is uncorrelated with M , equation (10) says that the decision should be to go ahead only if the trade offers a subjective expected return that exceeds $G\text{cov}(\epsilon_N, \epsilon_P^N)$.

A slight added wrinkle arises in the case where the trading desk also has a subjective view on the priced factor M , and therefore chooses not to hedge out M -risk, even though this is in principle feasible. (Recall that in deriving our capital budgeting results, we relied on Proposition 2, which said that without such a subjective view, the bank would always hedge M -risk completely.) In the appendix, we show that the relevant analogue to the hurdle-rate result in equation (10) is now given by:

$$\mu_N^* = \gamma\text{cov}(\epsilon_N, M) + G^M\text{cov}(\epsilon_N, \epsilon_P^N) \quad (15)$$

where G^M is a slightly modified version of the risk-aversion parameter G that takes account of

the complication that the bank is now bearing priced M-risk. (We define G^M formally in terms of primitive variables in the appendix.) Thus the basic logic is exactly the same as before; the only change is that the risk aversion parameter is calculated a bit differently.

C. Pricing non-hedgeable derivatives positions

Finally, our approach may be useful in helping to price derivatives positions that a bank cannot hedge cost-effectively. For concreteness, suppose the bank is acting as a dealer and has been asked to write a put option on the equity of another firm. If the option can be effectively delta-hedged by trading in the underlying equity, we are in the case where all the risk is tradeable. Thus the option should be priced using standard methods--i.e., a Black-Scholes approach.

However, suppose instead that the firm in question is either privately-held, or only very thinly traded. More precisely, the current market value of the firm is observable, but because of either trading costs or short-selling constraints, it is infeasible to hedge the option. Thus if the bank writes the option, it must bear the associated exposure. What price should the bank now charge for the option?

To attack this problem, we need to make a few further assumptions. First, the market value of the firm's stock S , follows a lognormal diffusion process with drift θ (for simplicity, there are no dividends) and instantaneous variance σ :

$$dS = \theta S dt + \sigma S dz \tag{16}$$

Second, we also need to assume that we are in the limiting case where the investment under consideration is very small. This allows us to express the bank's required return on the stock μ_s^* as a fixed constant, in the form of equation (10).¹⁵ Note that if the stock is positively correlated with the rest of the bank's portfolio, $\mu_s^* > \theta$ --the bank's required return exceeds the market's. This is because of the bank's risk aversion. If the bank were forced to hold the underlying stock directly, it would value it at a discount relative to the market. As before, we can define the premium that the bank requires on the stock (above and beyond the market required return) as $\pi_s^* \equiv \mu_s^* - \theta = \text{Gcov}(dS/S, \epsilon_P^N)$.¹⁶

With these definitions in hand, the appendix shows that the value of the option from the bank's perspective F , satisfies the following partial differential equation:

$$(r - \pi_s^*)SF_s + F_t + \frac{1}{2}\sigma^2S^2F_{ss} - rF = 0 \quad (17)$$

This is exactly the same valuation equation that one would obtain in a classical pricing setting where the underlying stock paid a proportional dividend of π_s^* . Thus to adjust the option's value for non-hedgeability, all one needs to do is take a standard pricing model and augment the dividend by yield by π_s^* .¹⁷ (Note that as in the standard setting, the market's

¹⁵The only distinction is that because we are now talking about the required return on a stock, instead of a forward position, one must add the riskless rate r to the expression in equation (10) to obtain μ_s^* . See footnote 5.

¹⁶In this continuous-time setting, ϵ_P^N should now be interpreted as the instantaneous innovation in the rate of return on the non-tradeable component of the pre-existing portfolio.

¹⁷The one slight complication has to do with early exercise. In the example above, the holder of the put option will generally attach a different value to the option than the bank. And the holder's decision of when to exercise will be determined so as to optimize value from his

expected return on the underlying stock θ , does not enter into consideration.)

In the context of our example of a bank writing a put, equation (17) implies that the bank will charge a higher price than suggested by a standard model. The intuition is straightforward: the put is written on a stock which the bank discounts at a rate of π_s^* more than the market. Thus from the bank's perspective, the underlying stock is worth less than in the open market, and accordingly, the put is worth more. The upward effect on put price is the same as would occur if the option was fully hedgeable but was written on a stock that paid a dividend of π_s^* .

More generally, the same sort of method can be used to value a wide range of illiquid derivatives positions. For example, one might wish to value illiquid credit risks of the sort discussed in the Introduction. Following Merton (1974), one could take the general approach of modelling a bank's credit exposure to a firm as being equivalent to a short put position in the firm's market value. Of course, this general type of model can be tailored along a number of dimensions, according to how one wants to treat issues of priority in bankruptcy, etc. But whatever the specific variant of the perfect-markets pricing model is deemed appropriate, our results suggest that it can be adjusted for illiquidity very simply, once the correct value of π_s^* has been established.

V. Conclusions

This paper has stressed the linkages between banks' risk management, capital budgeting, and capital structure policies. In our model, all three of these policies are shaped by two

perspective. This exercise strategy must then be incorporated when valuing the option from the bank's perspective.

related primitive frictions: first, it is costly for banks to raise new external funds on short notice; and second, it is also costly for banks to hold a buffer stock of equity capital on the balance sheet, even if this equity is accumulated over time through retained earnings.

Given these frictions, bank-value maximization implies the following: i) Banks should hedge any risks that can be offloaded on fair-market terms. ii) Banks should also hold some capital as a device for absorbing those illiquid risks which cannot be hedged, but the optimal amount of capital is limited. iii) Finally, given limited capital, banks should value illiquid risks much as an individual investor would--i.e., according to their impact on overall portfolio risk and return--with the degree of risk aversion being a decreasing function of the amount of capital held.

Because capital budgeting in our model is driven by an exposure's impact on portfolio risk and return, decentralized implementation becomes a problem if the pre-existing portfolio cannot be treated as fixed for all time. Thus for example, if several new investment opportunities arrive at once, the optimal allocation to each must be jointly determined. This requires a very well-informed "central planner" operating out of bank headquarters--it is not in general optimal to let individual product managers or traders choose their positions independently of one another.

The whole issue of centralization vs. decentralization strikes us as very important in a banking context, where asset portfolios can in principle be adjusted quite rapidly, and hence where fully centralized decision making could become extremely unwieldy. The fact that our model has had little of substance to say on this front suggests that it is missing certain key ingredients. As discussed above, the two most glaring omissions in this regard would seem

to be: i) the lack of an explicitly dynamic element, whereby information about investment opportunities evolves continuously over time, but where headquarters might choose to collect data only at discrete intervals; and ii) the absence of any agency problems at the level of individual product managers. Incorporating one or both of these considerations into our general framework would be a natural next step, and might help shed some light on a variety of practices--such as the widespread use of capital or position limits for individuals or small units within a bank--that seem to be important in the real world but which cannot be understood in the context of the current version of our model.

VI. Appendix

A.1 Proof of Propositions 1 and 2:

Since the bank may face numerous tradeable exposures, we prove Propositions 1 and 2 allowing for multiple exposures. Thus, we denote the bank's hedging policy as a set of coefficients h_i , which weight a set of L tradeable factors, x_i , $i = 1, \dots, L$. This implies that $Z_H = \sum h_i x_i$. We assume that the last of these x_i 's is just the market factor, i.e., $x_L = M$. We show below that the value-maximizing set of h_i 's is that which emerges from a simple regression of pre-hedged wealth, $Z_P + \alpha Z_N + K(1 - \tau)$, on the x_i 's. This hedging policy is complete, in that it minimizes the variance of wealth by making post-hedged wealth at time 2 uncorrelated with each tradeable factor.

To derive this, note that the bank chooses its hedging policy to maximize value, V .

This implies that:

$$\frac{dV}{dh_i} = \frac{dE(P(w))}{dh_i} - \gamma \frac{dcov(P(w), M)}{dh_i} = 0, \quad (\text{A.1})$$

where dh_i represents a change in the hedge ratio for the i th exposure. Using the definition of covariance, we can write:

$$\frac{dV}{dh_i} = E(P_w) E \left[\frac{dw}{dh_i} \right] + \text{cov} \left[P_w, \frac{dw}{dh_i} \right] - \gamma \frac{dcov(P(w), M)}{dh_i}. \quad (\text{A.2})$$

Next, given that the components of w are normally distributed, we can use the fact that, for

normally distributed x and y , $\text{cov}(f(x),y) = E(f_x)\text{cov}(x,y)$. Applying this yields:

$$\frac{dV}{dh_i} = E(P_w) E\left[\frac{dw}{dh_i}\right] + E(P_{ww})\text{cov}\left[w, \frac{dw}{dh_i}\right] - \gamma \frac{d(E(P_w)\text{cov}(w, M))}{dh_i}. \quad (\text{A.3})$$

We then do the differentiation indicated in the last term, and observe that, if the bank performs its hedging in an efficient market, changes in risk and return must be related by:

$$E\left[\frac{dw}{dh_i}\right] = \gamma \frac{d\text{cov}(w, M)}{dh_i} \quad (\text{A.4})$$

Thus, noting that $dw/dh_i = x_i$, we are left with:

$$\frac{dV}{dh_i} = E(P_{ww})\text{cov}(w, x_i) - \gamma E(P_{ww}x_i)\text{cov}(w, x_i) = 0 \quad (\text{A.5})$$

Given the concavity of $P(w)$, the solution to these L equations involves setting each h_i so as to make its associated factor uncorrelated with internal wealth, i.e., setting $\text{cov}(w, x_i) = 0$ for all i . This can only be accomplished by stripping out all of the tradeable exposures from internal wealth. Note that this logic applies regardless of whether or not all the risks are tradeable. Thus we have proven both Propositions 1 and 2.

A.2 Derivation of Equation (9):

The bank's first-order condition is:

$$\frac{dV}{d\alpha} = \frac{dE(P(w))}{d\alpha} - \gamma \frac{dcov(P(w), M)}{d\alpha} = 0 \quad (\text{A.6})$$

This expression can be rewritten as:

$$\text{cov} \left(P_w, \frac{dw}{d\alpha} \right) + E(P_w) E \left(\frac{dw}{d\alpha} \right) - \gamma \text{cov}(w, M) E \left(P_{ww} \frac{dw}{d\alpha} \right) - \gamma E(P_w) \frac{dcov(w, M)}{d\alpha} = 0 \quad (\text{A.7})$$

Using again the fact that for normally distributed x and y , $\text{cov}(f(x), y) = E(f_x) \text{cov}(x, y)$, and rearranging terms, this expression can be rewritten as:

$$\begin{aligned} \frac{dV}{d\alpha} &= E(P_w) \left[E \left(\frac{dw}{d\alpha} \right) - \gamma \frac{dcov(w, M)}{d\alpha} \right] + \\ E(P_{ww}) &\left[\text{cov} \left(w, \frac{dw}{d\alpha} \right) - E \left(\frac{dw}{d\alpha} \right) \gamma \text{cov}(w, M) \right] - \\ E(P_{www}) &\left[\gamma \text{cov}(w, M) \text{cov} \left(w, \frac{dw}{d\alpha} \right) \right] = 0 \end{aligned} \quad (\text{A.8})$$

Next note that:

$$\begin{aligned} \frac{dw}{d\alpha} &= \mu_N + \epsilon_N^N - \gamma \text{cov}(\epsilon_N^T, M) \\ E \left(\frac{dw}{d\alpha} \right) &= \mu_N - \gamma \text{cov}(\epsilon_N^T, M) \\ \text{cov} \left(w, \frac{dw}{d\alpha} \right) &= \text{cov}(\epsilon_N^N, \epsilon_p^N) + \alpha \text{var}(\epsilon_N^N) \\ \frac{dcov(w, M)}{d\alpha} &= \text{cov}(\epsilon_N^N, M) \end{aligned} \quad (\text{A.9})$$

Substituting these expressions into equation (A.8), and rearranging yields a generalization of equation (9) in the text:

$$\mu_N = \gamma \text{cov}(\epsilon_N^T, M) + \gamma \text{cov}(\epsilon_N^N, M) + G^M (\text{cov}(\epsilon_N^N, \epsilon_p^N) + \alpha \text{var}(\epsilon_N^N)), \quad (\text{A.10})$$

where

$$G^M = \frac{-(E(P_{ww}) - \gamma E(P_{www}) \text{cov}(w, M))}{E(P_w) - \gamma E(P_{ww}) \text{cov}(w, M)} \quad (\text{A.11})$$

One last step is required to generate equation (9) in the text. Equation (A.10) above is slightly more general, in that it explicitly allows for the possibility that $\text{cov}(\epsilon_N^N, M) \neq 0$ and hence $\text{cov}(w, M) \neq 0$, even after the bank has hedged. As we discuss in section IV.B of the text, a trading desk may wish to leave itself exposed to the market factor M in this way, if it has a subjective view on the expected return to this factor. However, in equation (9) in the text, we consider the simpler case, in which hedging completely removes the market factor, so that $\text{cov}(\epsilon_N^N, M) = \text{cov}(w, M) = 0$, and G^M reduces to $G = -E(P_{ww})/E(P_w)$.

A.3 Derivation of Equation (17):

By Ito's lemma, the instantaneous expected change in the value of the option is given by:

$$E(dF) = \theta S F_s + F_t + \frac{1}{2} \sigma^2 S^2 F_{ss} \quad (\text{A.12})$$

This must be equal to $\mu_F^* F$, where μ_F^* is the bank's required return on the option. By equation (10) we have:

$$\mu_F^* = r + \gamma \text{cov}\left(\frac{dF}{F}, M\right) + \text{Gcov}\left(\frac{dF}{F}, \epsilon_P^N\right) \quad (\text{A.13})$$

By the linearity of the covariance operator, this can be rewritten as:

$$\begin{aligned} \mu_F^* &= r + \frac{F_s S}{F} (\gamma \text{cov}\left(\frac{dS}{S}, M\right) + \text{Gcov}\left(\frac{dS}{S}, \epsilon_P^N\right)) \\ &= r + \frac{F_s S}{F} (\mu_s^* - r) \end{aligned} \quad (\text{A.14})$$

Thus the valuation equation is given by:

$$\theta S F_s + F_t + \frac{1}{2} \sigma^2 S^2 F_{ss} = F \left(r + \frac{F_s S}{F} (\mu_s^* - r) \right) \quad (\text{A.15})$$

Substituting $\pi_s^* = \mu_s^* - \theta$ and rearranging gives equation (17) in the text.

VII. References

- Cooper, Ian A., and Antonio S. Mello, 1991, The Default Risk of Swaps, *Journal of Finance* 46, 597-620.
- Froot, Kenneth A., David S. Scharfstein, and Jeremy C. Stein, 1993, Risk Management: Coordinating Corporate Investment and Financing Policies, *Journal of Finance* 48, 1629-1658.
- Gale, Douglas, and Martin Hellwig, 1985, Incentive-Compatible Debt Contracts I: The One-Period Problem, *Review of Economic Studies* 52, 647-664.
- Greenwald, Bruce C., Alec Levinson, and Joseph E. Stiglitz, 1991, Capital Market Imperfections and Regional Economic Development, in Giovannini, ed.: *Finance and Development: Lessons and Experience* CEPR, 65-93.
- Jarrow, Robert A., and Stuart M. Turnbull, 1995, Pricing Derivatives on Financial Securities Subject to Credit Risk, *Journal of Finance* 50, 53-85.
- Kashyap, Anil, and Jeremy C. Stein, 1995, The Impact of Monetary Policy on Bank Balance Sheets, *Carnegie-Rochester Conference Series on Public Policy* 42, 151-195.
- Merton, Robert C., 1974, On the Pricing of Corporate Debt: The Risk Structure of Interest Rates, *Journal of Finance* 29, 449-470.
- Merton, Robert C., 1995a, A Model of Contract Guarantees for Credit-Sensitive Opaque Financial Intermediaries, Harvard Business School working paper.
- Merton, Robert C., 1995b, A Functional Perspective on Financial Intermediation, *Financial Management* 24, 23-41.
- Merton, Robert C., and Andre F. Perold, 1993, The Theory of Risk Capital in Financial Firms, *Journal of Applied Corporate Finance* 5, 16-32.
- Myers, Stewart C., 1984, The Capital Structure Puzzle, *Journal of Finance* 39, 575-592.
- Myers, Stewart C., and Nicholas Majluf, 1984, Corporate Financing and Investment Decisions When Firms Have Information that Investors Do Not Have, *Journal of Financial Economics* 13, 187-221.

Stein, Jeremy C., 1995, An Adverse Selection Model of Bank Asset and Liability Management with Implications for the Transmission of Monetary Policy, NBER working paper # 5217.

Townsend, Robert M., 1979, Optimal Contracts and Competitive Markets and Costly State Verification, *Journal of Economic Theory* 21, 265-293.