Should Risky Firms Offer Risk-Free DB Pensions?

David A. Love*
Paul A. Smith†
David Wilcox‡

October 1, 2008

[PRELIMINARY DRAFT: Please do not cite without permission]

Abstract

We develop a simple model of pension financing to study the effect of pension risk on shareholder value. In the model, firms minimize costs, total compensation must clear the labor market, and a government pension insurer guarantees a portion of promised benefits. Shareholders value the pension risks according to their market values, while undiversified, risk-averse workers value the risks according to a certainty-equivalent utility approach. We find that in the absence of mispriced pension insurance, plan sponsors (i.e., firms) can maximize shareholder value by contributing the present value of promised benefits to the pension trust and by allocating those funds entirely to the risk free asset. Thus, in our framework, the optimal pension strategy under most specifications is to immunize all sources of market risk. Mispriced pension insurance, however, gives firms the incentive to introduce risk into their pension promises, offering an explanation for the observed prevalence of risky pensions in the real world.

JEL classification: G11

Keywords: Pensions; bankruptcy

*Dept. of Economics, Williams College, Williamstown, MA 01267, david.love@williams.edu.

†Federal Reserve Board, 20th and C St., NW, Washington, DC 20551, paul.a.smith@frb.gov.

‡Federal Reserve Board, 20th and C St., NW, Washington, DC 20551, david.wilcox@frb.gov. The views expressed herein are those of the authors and do not necessarily reflect those of the Board of Governors or the staff of the Federal Reserve System.

1 Introduction

This paper asks whether corporations can increase shareholder value by fully insulating their pension promises from the risks of market fluctuations and bankruptcy. Although traditional pensions have declined relative to 401(k)-style accounts, they remain a substantial force in the U.S. economy, with about 42 million participants[1] and assets totaling \$2.4 trillion.[2] Insulating these plans from market risks would represent a sharp departure from current practice, as many plans in the U.S. are underfunded and heavily invested in risky assets.[3] The extent of risk involved in the pension industry as a whole became apparent during the market decline of 2001. In a single year, a combination of falling stock prices and low interest rates helped to convert an \$18 billion surplus among the 100 largest pension plans into a \$163 billion deficit (Milliman Consultants and Actuaries, 2007). Over the same year, the Pension Benefit Guarantee Corporation (PBGC), the federal insurer of these plans, experienced a balance-sheet reversal from a \$7.7 billion surplus to a \$3.6 billion deficit (PBGC, 2003). Given that financial distress in the pension industry can have a substantial impact on employees, the PBGC, and the national economy more broadly, a natural question is whether the firms themselves actually benefit from offering their workers a risky pension promise.

Firms could remove most of the risk from their pension promises by contributing a sufficient amount of resources to the pension trust and by investing the trust in a portfolio that delivers cash flows that closely match those of future pension liabilities (Bodie, 1990). This funding and investment strategy would effectively "defease" the liabilities associated with promised benefits.[4] But while firms *could* eliminate much of this risk if they chose to, we are unaware of any firms in the U.S. that currently make their pension promises close to risk-free.[5] Using a simple model of the U.S. pension system, we demonstrate that the current practice of offering risky pensions is financially inefficient; if employees are informed

[1] Private Pension Plan Bulletin: Abstract of 2005 Form 5500 Annual Reports

[2] Flow of Funds Accounts of the United States (2008). State and local government employee pension funds hold an additional \$3.3 trillion of assets. We focus on the pension promises of private firms.

[3] At the end of 2006, 55 of the largest 100 corporate pension plans held assets less than their projected obligations, and about 70 percent of their aggregate assets were held in equities and other risky investments (Milliman Consultants and Actuaries, 2007).

[4] Defeasance involves purchasing assets that produce cash flows that match the timing and amount of the outflow of liabilities. Because of the tax-advantaged nature of pension plans, the ideal vehicle for defeasing would be a fully taxed, sufficiently long-dated security with normal liquidity properties and no default risk. While no security possesses all of these characteristics, there are some that provide a close match. One possibility, for instance, would be a bundle consisting of LIBOR-based eurodollar deposits combined with swap contracts, structured to convert the variable-rate EDs into synthetic fixed-rate securities with the right time to maturity. Note that defeasance is a more stringent requirement than simply immunizing the pension liabilities, which entails using duration-matched fixed-income securities to protect against changes in the term structure of interest rates.

[5] In England, Boots Pharmaceutical immunized most of its pension risk from 2001 to 2004 (Ralfe and Palin, 2004).

about the risk properties of the pension, they will demand extra compensation for taking on additional risk, and this extra compensation will lead to a net increase in the firm's total costs (as long as pension insurance is fairly priced).

In our model, firms seek to minimize total compensations costs, workers operate in a competitive labor market, and a government insurance agency guarantees some fraction of benefits (in some cases, zero). The starting point for our analysis is Bulow's (1982) observation that if compensation is subject to constant renegotiation, benefits and salary should equal the value of the employee's marginal product in each period.[6] As Bodie has argued (1988; 1990), an important implication of this market clearing condition is that workers must be compensated for any additional risk in the firm's pension plan, and that this transfer can be sufficiently large as to make risky pensions an inefficient form of compensation.

The primary contribution of our paper is to formally develop these insights in a simple model of optimal pension financing. Bader (2004) provides an especially clear description of the central argument. Asking one to think of the pension promise as an "employer bond" subject to default risk, he notes that "[u]nlike the investors who determine market prices, the employee cannot diversify the company-specific risk to which he is already overexposed, so he would not pay the full market price. Nor would it be rational for him to give up enough salary to cover the full market value of the risky pension" Even if the pension itself represents the only source of nondiversifiable risk, employees would still discount the "bond" at a higher rate than would investors purchasing the bond on the open market. Using an equilibrium model of pension financing, we examine the proposition that firms cannot increase shareholder value by issuing employees what is essentially a risky, nontransferable bond.

In earlier work that abstracted from the investment decisions of the firm and the employee, we demonstrated that it is optimal for a zero-beta firm to fully fund its pension promise (Love, Smith, and Wilcox, 2007). That paper, however, stopped short of incorporating certain aspects of risk, insurance, and portfolio choice that may be important in the real world. Our current paper introduces some of these missing elements by allowing for non-zero correlation between the firm and the market, optimal portfolio allocation decisions on the part of the worker and the firm, firm-specific human capital, and the presence of (possibly imperfect) of government pension insurance.

For most specifications of the model, regardless of our assumptions about risk aversion, the firm's beta, and initial values of the firm's financial health, we find that the firm's optimal pension decision is a corner solution: the firm contributes the present value of

[6] Ippolito (1985a,b) contrasts the competitive labor market view, which he calls the "legal interpretation," with an alternative model in which pensions form part of an implicit contract with workers (as discussed in Treynor (1977)). We discuss how the implicit contract view might affect our analysis in Section 5 below.

pension benefits, and it allocates the trust fund entirely in the risk-free asset. Further, when we compare the firm's pension cost savings (due to reduced pension contributions) from a risky promise to the cost of compensating employees for the extra risk, we generally find that the difference between them—the "net value" of a risky promise—grows more negative as the pension becomes riskier. Both underfunding (i.e., contributing less than the present value of promised benefits) and investing the pension trust fund in the risky asset lead to more negative net values, with underfunding generating the larger effect.[7] The one specification where we find some scope for risky pensions is, perhaps tellingly, one in which the government insurance agency charges less than market value for its pension guarantee. In this case, the firm and the employee can effectively split the surplus gained at the expense of the government pension insurance, and by extension, at the expense of U.S. taxpayers.

Our paper fits into a long-established body of pension research. Munnell (1982) provides a history of corporate pension plans in the U.S.[8] Early authors noted that bankruptcy gives firms a put option on the difference between promised benefits and the pension assets (Black, 1980; Sharpe, 1976; Treynor, 1977). Since the value of the put rises with the volatility of the underlying asset, shareholders benefit if the pension fund is invested in risky assets (Treynor, 1977). Sharpe (1976), however, pointed out that the additional required compensation to employees would eliminate the gains from a risky pension.[9]

But missing from these earlier studies is a recognition that employees are typically unable to hedge away the risks associated with default. Employees are generally either restricted from short-selling company stock, or they lack the financial sophistication to do so. This puts them in a similar position as risk-averse executives who receive part of their compensation in the form of company stock. Hall and Murphy (2002) and Cai and Vijh (2005) argue that such executives should place a lower value on stock options than traditional valuation methods (e.g., Black-Scholes) would predict. Although the pension risks facing employees differ from the stock-option risks facing executives, we obtain a similar result.

A few recent studies have examined other factors affecting pension funding and investment decisions. Rauh (2007) finds that, since solvent firms must make up any pension

[7]While greater investment risk increases the possibility that employees will receive less than promised benefits in the event of bankruptcy, low funding levels practically guarantee it.

[8]In this paper we take the history as given, and we focus on *how* firms fund their pension promises, rather than why they offer pensions at all or why employees value them. In particular, we are interested in how employees and firms value risk in the pension promise, and under what conditions firms would choose to offer risky, rather than risk-free, pension benefits.

[9]These results ignore distortions introduced by taxation and pension insurance. In the absence of default risk the tax-exempt nature of pension contributions favors bonds and full funding (Black, 1980; Tepper, 1981). With default risk, however, equities can dominate bonds as long as the option value of putting the shortfall to the employee during bankruptcy outweighs the tax advantages of bonds (Harrison and Sharpe, 1983).

shortfall, firms with poorly funded pension plans tend to invest their pensions in safer assets. In this case, the "risk management incentive" outweighs the put option of shifting risks to the employee. In the other direction, Lucas and Zeldes (2006) find that as long as wage growth and equity returns are sufficiently correlated, firms can benefit from investing the fund in equities as a hedge against future pension obligations (which are linked to wages through the benefits formula).

Finally, a number of authors, including Bader (2003), Bodie (1988, 1990), Gold (2003), Ralfe, Speed, and Palin (2004), and Wilcox (2006), have recently argued that pension funds should hold only riskless assets for various reasons, including immunization against interest rate movements, reduced balance-sheet risk, improved pension security for employees, improved tax efficiency, and reduced administrative costs. However, none of these papers offers an explicit model of pension funding and investment in the context of an equilibrium model of employee compensation.

2 Institutional Background and Empirical Findings

Firms that compensate employees with pension promises are generally required to maintain a separate pension trust fund for the benefit of the eventual recipients. Minimum contributions are required under the Employee Retirement Income Security Act (ERISA) and the tax code, which sets out the conditions that must be met for the trust fund (and contributions thereto) to receive favorable tax treatment. Investment of the trust fund is not heavily regulated under ERISA, though investments must be made for the benefit of the employees and must be "prudent."[10]

Plans that become underfunded must generally make up funding shortfalls over time.[11] However, bankruptcy generally relieves the firm of its obligation to make up funding shortfalls. In bankruptcy, a firm with an underfunded pension generally terminates the plan and turns over both its assets and liabilities to the PBGC, which invests the assets and assumes the responsibility of paying out promised benefits (up to a limit—for 2008, the maximum benefit is $51,750 for a worker retiring at age 65). In exchange for this insurance coverage, the PBGC charges solvent firms a premium of $33 per participant (in 2008), plus

[10]One specific regulation is that no more than 10 percent of assets may be invested in employer stock.

[11]These rules were tightened by the Pension Protection Act of 2006, which requires higher funding targets and quicker replenishment of shortfalls. Prior to the the Act, sponsors were required to fund 90 percent of pension liabilities. Brown (2008) lists several ways that firms could circumvent this requirement in practice. Plan sponsors were permitted to (1) compute liabilities using smoothed interest rates, (2) value plan assets using smoothed actuarial values rather than market values, (3) substitute accounting credits for required contributions, even in the event of underfunding, (4) amortize changes in plan benefits over a 30-year period, and (5) increase promised benefits even when the plan is underfunded. The Pension Protection Act fixed some of these shortcomings, but the funding rules still permit considerable latitude on the part of plan sponsors.

0.9 percent of any unfunded benefits.[12]

As noted above, many pension plans are underfunded and most invest heavily in risky assets. Table 1 shows the funding status of the largest 100 corporate pension plans from 2001 to 2006. Line 1 shows that the aggregate funding status, defined as total plan assets divided by total projected liabilities,[13] deteriorated substantially from full funding in 2001 to about 83 percent funding in 2002. Funding levels improved a bit the following year, remaining at about 90 percent until 2006, when they once again reached full funding. But, as shown by lines 2 through 4, this aggregate pattern masks significant heterogeneity. Even in the "fully funded" year of 2001 about half of the plans were underfunded, and 13 percent were significantly (more than 20 percent) underfunded. In the "crisis" year that followed, about three-quarters of plans were underfunded and 44 percent were significantly underfunded. By 2006, when the aggregate funding ratio reached 100 percent, 55 of the largest 100 plans remained underfunded, and 10 remained significantly underfunded.

These funding gyrations were caused by exposure to volatility in both stock returns and interest rates, but also by changes in contributions. As shown in line 5, the market downturn forced many firms to make large contributions to help close the gap—aggregate contributions jumped from $11 billion in 2001 to $34 billion the following year and $56 billion in 2003.[14]

Finally, Table 2 reports the share of pension assets held in equity and other investments in 2006.[15] As shown in line 1, about 71 percent of the aggregate portfolio was invested in risky assets in 2006. Lines 2 through 4 show that about 10 percent of plans hold more than 80 percent of assets in equities and other risky investments, while about 11 percent invest less than 60 percent in risky assets. Most firms hold about 70-80 percent of assets in risky investments.

3 The model

Our model describes the behavior of a representative firm, an employee, and, in some variations of the model, a government insurer of defined-benefit pension plans. The firm seeks to maximize shareholder value by minimizing total compensation costs, which include both wages and the cost of providing a defined-benefit-type pension. The employee must

[12]The Deficit Reduction Act of 2005 also established a "termination premium" of $1,250 per participant in the year a bankrupt firm terminates its plan and each of the following two years.

[13]The liability concept used here is the projected benefit obligation (PBO).

[14]Because funding rules allow firms to spread "deficit-reduction contributions" over several years, there is often a lag between market changes and contribution changes.

[15]We interpret this to be a rough measure of exposure to "risky assets." It is quite rough because "other investments" may include cash and other low-risk investments in addition to real estate and other risky investments.

decide how much of first-period compensation to save and how to allocate that saving between a risk-free asset and the equity market. The firm must provide a compensation package that is capable of generating a reservation level of utility assuming that the employee solves the relevant saving and investment problems optimally. Finally, in the extended model, a government agency insures part or all of the promised pension benefit.

The model is essentially a two-period affair. In return for the compensation package that provides the reservation level of utility, the employee supplies one unit of labor in the first period and a fraction of a unit of labor in the second period. (The wage payment in the second period exposes the worker to the potential loss of firm-specific human capital in the event of bankruptcy.) As compensation for that labor, the firm pays its employee wages w_1 in the first period, w_2 in the second period, and it promises to pay a pension benefit, b, in the second period. As of $t = 1$, the pension promise may be risky: The firm's ability to pay the full benefit depends on the amount it contributes to the pension trust fund in the first period, the returns it earns on the trust fund assets (which depend on the allocation of the assets between the risk-free asset and the market portfolio of equities), and whether the firm declares bankruptcy in the second period. We begin by considering the firm's and the employee's problems, and introduce the government insurer later.

3.1 The firm's problem

In the first period, the firm pays a wage, w_1, contributes F_1 to its pension trust fund, and chooses an allocation of the pension fund between a risk-free asset and a risky asset that can be thought of as the market portfolio of equities. The risk-free asset pays a gross return of R, while the risky asset generates a stochastic return of R_m. Letting λ denote the share of the trust fund invested in equities, the value of the trust fund in the second period is given by $\tilde{F}_2 = [R + (R_m - R)\lambda]F_1$.

If $\tilde{F}_2$ exceeds the promised level of benefits, the surplus reverts to the firm regardless of whether the firm declares bankruptcy in the second period. The structure of this payoff is equivalent to a conventional European call option on the excess in the trust fund over the promised benefit. We denote the value of this call $q^c(\tilde{F}_2, b)$. If, on the other hand, $\tilde{F}_2$ falls short of the promised level of benefits, the firm must contribute enough to eliminate the shortfall, provided the firm escapes bankruptcy in the second period. Thus, in addition to being long a conventional European call option on the value of the trust-fund assets, the firm is short a European put; crucially, however, the put is exercised on the firm only when the firm survives in the second period. We denote the value of this non-standard put option as $q^p(\tilde{F}_2, b|s = 1)$, where s is an indicator variable for whether the firm survives.

The firm declares bankruptcy when the value of the firm's non-pension assets, $\tilde{V}_2$, falls below a threshold equal to its liability to creditors, D, plus the shortfall, if any, in the pension trust fund.[16] The value of the firm's non-pension assets evolves according to a stochastic rate of return R_v, which may, in general, be correlated with the return on the market portfolio, R_m. To allow for such correlation, we follow Rubinstein (1994) and assume that the log returns on the market portfolio and the company stock follow a joint binomial distribution. Details of this specification are provided in Appendix 1. Letting s be an indicator variable for the survival of the firm, the bankruptcy threshold is given as follows:

$$s = \begin{cases} 1 & \text{if } \tilde{V}_2 \geq D + \max(b - \tilde{F}_2, 0) \\ 0 & \text{if } \tilde{V}_2 < D + \max(b - \tilde{F}_2, 0). \end{cases} \tag{1}$$

The bankruptcy condition allows for the possibility that the firm becomes insolvent with a non-pension market value in excess of its non-pension present value debt (so that $\tilde{V}_2 > D$). We assume that any remaining assets in bankruptcy stay with the firm and do not go to either the workers or the government insurer of the plan.[17]

In total, the firm's compensation costs are given by:

$$C = w + F_1 + q^p(\tilde{F}_2, b|s = 1) - q^c(\tilde{F}_2, b), \tag{2}$$

where w denotes the sum of first-period wages, w_1, and the market value of the second-period wage payment, w_2, which is paid as long as the firm survives.

Note that the value of the options will depend on the degree of correlation between the market portfolio and the firm's own assets. Positive correlations imply relatively low expected returns on the market portfolio during states of bankruptcy. Since in those states of the world the firm only covers part of the shortfall (the lesser of the shortfall and remaining assets after debts have been paid), a positive correlation reduces the value of the put option.

[16] We do not include the pension surplus in the bankruptcy condition for two reasons. First, the sponsor does not possess a legal right to the surplus and therefore cannot use it to meet obligations to creditors in the short term. Second, the employee is indifferent as to whether the firm declares bankruptcy with a fully funded pension, or remains solvent; in either case, the employee receives the promised pension benefit in full.

[17] There is some question about who holds a claim on the remaining assets of the firm after the non-pension debts have been paid. In the event of a distress termination or involuntary termination, the plan sponsor is liable to the PBGC for the shortfall (liabilities less assets with interest) at the point of termination (ERISA § 4062(b)(1); 29 U.S.C. § 1362(b)(1)). Whether the PBGC actually recovers this amount, however, depends on its ability to establish priority in bankruptcy court, either as an administrative expense or as an unsecured creditor. The ERISA rules grant the PBGC a claim on the unfunded liabilities of the pension up to 30 percent of the sponsor's net worth. While ERISA grants the PBGC a lien on the unfunded liabilities, with the same priority as a tax lien in bankruptcy proceedings, the Bankruptcy Code interferes with this claim. Brannick (2004) provides a detailed discussion of this conflict, noting that "[t]he PBGC faces a termination/liability paradox—its lien cannot arise until after termination, but termination in bankruptcy prevents perfection of the PBGC's lien." Also, because shareholders are not entitled to remaining firm assets when a shortfall triggers either a distress or involuntary termination, there is a moral hazard problem. The shareholders can extract any remaining firm value before the termination by increasing dividends or other means (Brannick, 2004).

Intuitively, a positive correlation means that the firm expects to avoid the largest shortfall obligations through bankruptcy. Conversely, the value of the put is higher if the firm and market are negatively correlated because a large shortfall is unlikely to be escaped via bankruptcy.

After we have computed the values of $q^p(\tilde{F}_2, b|s = 1)$ and $q^c(\tilde{F}_2, b)$, we can determine whether a risky funding strategy can be employed to make both the firm and the employee at least as well off as they would be in the case of a risk-free pension.

3.2 The employee's problem

The employee receives a wage, w_1, in the first period. If the firm avoids bankruptcy, the employee receives a wage in the second period, w_2, and promised pension benefits, b. One can think of the second period as spanning the tail end of the working life through retirement. If the firm enters bankruptcy in the second period, the employee has a claim on the lesser of promised benefits and the value of assets held in the trust fund. In addition, the employee is assumed to have some firm-specific human capital at stake. If the firm goes bankrupt, the worker receives a wage in period 2 of χw_2, where $\chi \in [0, 1]$.[18] The existence of a second-period wage is not required to derive our main results; we include it, however, to allow for the possibility that the worker is exposed to some non-pension risk in the event of bankruptcy. From the employee's perspective, second-period compensation can be treated as a random variable, $\tilde{x}$, with the following payoffs:

$$\tilde{x} = \begin{cases} b + w_2 & \text{if the firm remains solvent} \\ \min(\tilde{F}_2, b) + \chi w_2 & \text{if the firm enters bankruptcy.} \end{cases} \tag{3}$$

The employee chooses consumption and the portfolio share in the risky asset in order to maximize the following two-period lifetime utility function:

$$\max_{c,\alpha \geq 0} \left\{ u(c) + \delta E u(\tilde{R}[w_1 - c] + \tilde{x}) \right\}, \tag{4}$$

where $u(.)$ is a constant relative risk aversion utility function with coefficient γ, and $\tilde{R} = R + (R_m - R)\alpha$ is the uncertain return on the worker's investment portfolio. Note that in evaluating the pension from the employee's perspective, we are no longer working with risk-neutral rates of return.

There are two notable aspects of the employee's problem. First, the allocation decision

[18] The reduction in the wage reflects the decreased marginal product of labor the worker would generate in a new job. We abstract from the question of how the costs of firm-specific training are shared between the worker and the employer, a topic considered in the seminal works of Becker (1962) and Hashimoto (1981), and in more recent work by Acemoglu and Pischke (1999).

is an essential feature of the model. An employee restricted from holding the optimal share of the market portfolio outside of the pension plan would greatly value a marginal increase in exposure to any asset that offered a higher expected return than the risk-free rate.[19] In contrast, an employee at an interior portfolio decision problem would value marginal increments of the safe and risky assets at the same rate. Thus, if we forced the employee in our model to save only in the risk-free asset, she would tend to place more value on the risky pension relative to an employee able to hold both assets outside the pension.

Second, because the employee cannot short-sell company stock, she is a nondiversified holder of a risky pension promise. Therefore, while risk-neutral pricing may be appropriate for valuing the pension options of the firm, it would be inappropriate for solving the problem of the nondiversified, risk-averse employee. To quantify the value of a risky pension promise, we adopt a certainty equivalence approach (see Hall and Murphy, 2002) and solve for the first-period cash award that would make the employee just as well off (in utility terms) as she would be with a risk-free pension plan. In general, the net value of a firm's pension plan will be a complicated nonlinear function of the parameters governing the correlated stochastic return process on the risky assets and the risk aversion of the household, as well as the choice variables characterizing the firm's and employee's funding and portfolio decisions. Since there is no closed-form analytic solution, we compute the differences numerically for a wide range of parameters and present a representative selection of the results.[20]

4 The optimal pension strategy

4.1 Baseline: no pension insurance

To solve for the optimal pension strategy, we first compute the certainty-equivalent first-period payment, z, that would make the worker indifferent between a risky pension promise and a safe pension promise.[21] (We call a pension strategy "risky" if $F_1 < b/R$ and/or

[19]Cai and Vijh (2005) make the same point in the context of pricing executive options.

[20]Solving the model numerically requires that we specify a parameterization. Because we are dealing with a simple two-period framework, it is not possible to generate a realistic calibration of preference parameters, bankruptcy risk, etc. Instead, our approach is to experiment with a wide range of values to ensure that our qualitative results are not an artifact of a narrow specification. While the results in this paper are based on a particular subset of the parameter space, they are quite robust to different specifications. The baseline parameters are: $V_1 = 250{,}000$; $D = 0.72 \times V_1$; $\sigma_m = 0.18$; $\sigma_v = 0.20$; $b = 250$; $w_1 = 1{,}000$; $w_2 = 250$; $E(R_m - R) = 0.04$; and $\chi = 0.90$. Unless otherwise stated, we set the coefficient of relative risk aversion to 3 and the correlation between the market and the firm to 0. We vary these parameters in the next section. The ratio of the value of the firm to non-pension debt generates an average bankruptcy probability of 5% when the firm defeases its pension. Our main findings hold for much smaller values of the firm (equal to the size of wages), larger and smaller pension benefits, almost the entire range of bankruptcy probabilities (with special cases occurring at 0 and 100 percent), and different wage structures.

[21]In the current version of this paper, we do not allow the firm to short the market by setting $\lambda < 0$. Future versions will relax this constraint, which may bind for some (non-optimal) amounts of first-period

$\lambda \neq 0$, and we call a pension strategy "risk-free" if $F_1 = b/R$ and $\lambda = 0$.) Since z measures the utility cost to the employee of a risky pension promise relative to a safe one, we can compare z to the benefit the firm receives by offering a risky pension, holding wages constant. The benefit to the firm, B, consists of two pieces: the reduction in first-period pension contributions and the increase in the value of the pension options. That is:

$$B = b/R - F_1 + q^p(\tilde{F}_2, b|s = 1) - q^c(\tilde{F}_2, b). \tag{5}$$

By construction, a risk-free pension strategy will generate zero benefits, so that $B = 0$. If the benefit from offering a risky pension is less than the additional required compensation, then it will not be optimal for the firm to make a risky pension offer.

Figure 1 shows the net values, $B - z$, generated by different combinations of pension contributions and asset allocations for the case of a zero-beta firm. The surface bends up toward a maximum net value of zero at the upper right corner of the figure, which corresponds to a pension contribution equal to the present discounted value of promised benefits and an allocation entirely in the risk-free asset. We call this the "risk-free" pension strategy since the worker receives promised benefits in all states of the world. The shape of the surface in Figure 1 indicates that the level of contributions has a much larger impact on net values than does asset allocation. (A risky allocation exposes the worker to a potential loss in the event of bankruptcy, but a low funding level actually guarantees it.) In order to highlight the effect of allocation on funding, the right-hand panel of Figure 1 zooms in on the net values and focuses on a section of the surface that is closer to 100 percent funding. The surface clearly shows that the firm can obtain the highest net value by investing the trust fund entirely in the risk-free asset.

The surfaces in Figure 1 correspond to our baseline specification of zero correlation and a risk aversion coefficient of 3. Table 3 shows how the net values change as we vary the coefficient of relative risk aversion and the degree of correlation between the firm and the market. The left half of Table 3 reports results assuming a risk aversion coefficient $\gamma = 3$, which is well within the commonly accepted range of values. The right half displays results for $\gamma = 5$, a value that is frequently used in models of portfolio choice (where higher degrees of risk aversion are required to accommodate the equity premium puzzle). Since the results depend on the degree of correlation between the market portfolio and firm-specific returns, we divide the table into a zero-correlation firm, a positive-correlation firm, and a negative-correlation firm. Within each panel, we present the results for different choices of period-1 contributions to the trust fund and for different allocations of the trust fund to the risky asset. In each panel, the upper-right cell represents a riskless pension, which is fully funded

contributions.

in period 1 and invested entirely in the riskless asset. The riskless pension generates a net value of zero, since both z and B are identically equal to zero for the riskless strategy (i.e., employees demand no additional compensation and the firm's options have no value).

Looking at the zero-correlation case in the top panel, the negative values in the non-upper-right entries imply that at these levels of household risk aversion, moving away from the riskless strategy cannot be optimal from the firm's perspective since the amount it would have to compensate its employees would exceed the value of the risky pension to the firm. The intuition behind this result is straightforward: the curvature of the household's utility function implies that employees will demand compensation for a riskier pension strategy that exceeds the expected benefit to the firm. Indeed, comparing the net values for the two values of risk aversion, $\gamma = 3$ and $\gamma = 5$, we see that the cost to the firm is quite sensitive to the employee's tolerance for risk. Moving to the higher risk aversion coefficient more than doubles the size of most of the net values in Table 3.[22]

How does the degree of correlation between the firm and the market affect the net value associated with a firm's pension strategy? Intuitively, higher degrees of correlation make it more likely that a worker will experience a reduction in pension benefits since poor equity returns will tend to coincide with poor performance at the firm level. Employees should therefore demand additional compensation when the correlation is positive. The results in Table 3 suggest that this is indeed the case. As the allocation in the risky asset increases above zero, the (negative) net values for the positive correlation case grow larger relative to the zero-correlation case.

A negative correlation between the market portfolio and the firm's value would seem to provide some scope for the firm to gain from a risky pension strategy. Intuitively, a negative correlation means that employees are likely to receive their promised pension benefits even if the firm goes bankrupt because market returns are likely to be high. Since the worker is already exposed to own-firm risk in the form of firm-specific human capital, a negative correlation reduces the cost to the worker of a pension allocated in the risky asset. The results in Table 3, however, indicate that the firm cannot benefit from increasing the allocation when the pension is fully funded at the present value of promised benefits. At lower levels of funding, in contrast, the firm can increase the net values, holding contributions constant, by investing some of the pension in equities. An underfunded pension grants the worker a portion of the upside equity risk in the event of bankruptcy. Since equities are likely to have a high return in states of bankruptcy and since some human capital is at stake, the worker would like to have a bigger equity position in bankruptcy states. The

[22] Despite the sensitivity of the net values to risk aversion, we obtain negative net values for CRRA coefficients as low as 0.5 (the lowest we tried) as long as the employee remains at an interior solution in her portfolio problem. At low levels of risk aversion, an interior solution entails a highly leveraged position in the market portfolio.

equity investment in the pension allows the worker to achieve such a position. Importantly, though, while the firm can increase value by investing the pension fund in equities in some cases of underfunding, it does *not* pay for the firm to make the pension risky—all the net values in the table are still negative.

4.2 Effect of pension insurance

Results from the baseline model without pension insurance reported above indicate that it is optimal for the firm to offer a risk-free pension promise. We now expand the model to include federal pension insurance, a key feature of the pension landscape in the U.S. To see how pension insurance affects the sponsor's decisions in our model, we introduce a government agency that guarantees pension benefits up to an amount, $g \leq b$, with a strict equality indicating full coverage. In return for providing this insurance, the agency charges the firm a premium, denoted p, that may or may not be economically fair.[23]

The introduction of pension insurance alters the underlying problem in two ways. First, the cost of the risky pension is now given by:

$$C = w + F_1 + q^p(\tilde{F}_2, b|s = 1) - q^c(\tilde{F}_2, b) + p, \tag{6}$$

where p is the premium charged by the government insurance agency.

Second, the employee's problem is altered by a new guarantee on the pension outcome. The employee's payoff in period 2 is given by:

$$\tilde{x} = \begin{cases} b + w_2 & \text{if the firm survives} \\ \min(b, \max[\tilde{F}_2, g]) + \chi w_2 & \text{if the firm enters bankruptcy,} \end{cases} \tag{7}$$

where the bankruptcy payoff, $\min(b, \max[\tilde{F}_2, g])$, reflects the coverage provided by pension insurance. Note that in the bankruptcy state the employee receives b if second-period assets are sufficient to cover benefits or pension insurance provides full coverage. The employee receives $\max[\tilde{F}_2, g] < b$ if second-period assets are insufficient *and* pension insurance provides less than full coverage.

[23] In the real world, PBGC premiums are assessed annually, and may be paid either out of plan assets or directly by the employer. In our model, the firm pays the insurance premium only once, in the first period. The firm does not need to pay another premium at $t = 2$ because by that point, uncertainty about the firm's ability to pay the promised pension benefit has been resolved. We could have required the firm to make premium payments between $t = 1$ and $t = 2$ but chose not to in order to align the timing of the single premium payment with the timing of the single contribution into the trust fund. Our key results would still hold if the firm paid an insurance premium at every re-pricing of the assets held by the trust fund; the main influence of this alternative assumption would be to reduce the risk of the firm to the insurance entity. Also, we assume that the firm pays the premium directly rather than out of plan assets, but this decision is entirely arbitrary in the sense that all the key results we demonstrate below would hold under the opposite assumption.

How pension insurance affects the firm's equilibrium compensation decision depends on the extent of insurance coverage (complete or partial) and the structure of pricing (economically fair or mispriced). We take each of these cases in turn.

4.2.1 Case 1: Complete insurance and an economically fair premium

We begin with the case of complete coverage, $g = b$, and an economically fair premium. The first thing to notice is that complete coverage effectively removes the employee from the firm's decision. No matter how the firm funds its pension, the employee will always receive promised benefits. As a result, the employee will no longer impose discipline on the firm by demanding appropriate compensation for risk. However, an economically fair premium will impose its own discipline. A reasonable definition of an economically fair insurance premium is the market value of the obligation to cover any shortfall in the state of bankruptcy, $\max[b - \tilde{F}_2, 0]$. Again using the option value as an an approximation of the shortfall obligation, the fair market insurance premium is given by:

$$p^* = q^p(\tilde{F}_2, b|s = 0). \tag{8}$$

Setting the premium equal to the conditional put has the desirable property that the marginal change in premiums will always equal the marginal increase in the value of the put and call options. Thus, any change in pension strategy will cause insurance premiums to adjust in tandem with the change in the market value of the pension risk. In particular, any increase in the riskiness of the pension promise will result in a higher premium.[24]

We can formalize this result by first recognizing that the sum of an economically fair premium and the value of the firm's obligation to cover shortfalls if it survives is simply an unconditional put on the market portfolio with a strike price of b:

$$\begin{aligned} p^* + q^p(\tilde{F}_2, b|s = 1) &= q^p(\tilde{F}_2, b|s = 0) + q^p(\tilde{F}_2, b|s = 1) \\ &= q^p(\tilde{F}_2, b), \end{aligned} \tag{9}$$

where the last line follows from the fact that an unconditional put equals the sum of the conditional puts in bankruptcy and solvency.

Substituting equation (9) into the firm's cost function, we have:

$$C = w + F_1 - q^c(\tilde{F}_2, b) + q^p(\tilde{F}_2, b). \tag{10}$$

We find that the put associated with the shortfall is no longer conditional and now depends

[24] As pointed out by Brown (2008), CBO (2005), and others, this desirable feature is missing from the actual structure of PBGC premiums.

only on a single underlying asset—the market portfolio. Notice, however, that by put-call parity,

$$q^p(\tilde{F}_2, b) - q^c(\tilde{F}_2, b) = b/R - F_1.$$

Substituting this expression back into (10), it is clear that regardless of the firm's contribution and allocation decisions, total costs will always be the same: $w + b/R$.

Any solution is possible in this case, since the firm and the government insurance agency place identical values (though with different signs) on any change in the pension's risk characteristics. Thus, complete pension insurance with economically fair premiums could result in firms offering a riskless pension, an extremely risky pension, or anything in between. The indeterminacy disappears, however, if there are any additional administrative fees associated with the government-provided insurance. After adding such fees, $p > q^p(\tilde{F}_2, b|s = 0)$, and it will again be optimal for the firm to offer a risk-free pension. That is, if pension insurance is complete and the premium is even slightly greater than the economically fair value, the firm's optimal decision will be to fully fund its pension and invest the trust fund solely in the risk-free asset.

4.2.2 Case 2: Incomplete insurance and an economically fair premium

We now consider the case of imperfect insurance, where the guarantee is less than promised benefits: $g < b$. If the insurance is economically fair, $p^* = q^p(\tilde{F}_2, g|s = 0)$, and the firm receives some benefit from offering risky insurance as long as employees do not demand additional compensation.[25] In bankruptcy, the firm still owns a call on any pension surplus, but it effectively pays for only part of the put covering any shortfall. Nevertheless, because the employee understands the implications of the firm's pension strategy, she will end up demanding additional compensation that more than offsets the benefit to the firm.

Figure 2 shows the result for the case of a zero-beta firm with a pension guarantee equal to 75 percent of promised benefits. The front edge of the figure, which corresponds to different levels of contributions holding the equity investment at zero, shows that the net values are essentially flat for funding levels below the maximum guaranteed level of benefits.[26] As the funding level rises above the benefit guarantee, the front edge of the net value surface rises sharply as the worker concedes a higher fraction of wages in exchange for a less risky pension.

The back edge of Figure 2 shows a similar increase in net values near the point of

[25] Note that the strike price of the put option for the fair insurance premium is the maximum level of insured benefits, g, and not promised benefits, b.

[26] The net values are not entirely flat because low levels of funding make the firm more likely to declare bankruptcy and therefore increase the likelihood that the worker will lose some firm-specific human capital. The worker demands additional compensation in exchange for this risk, reducing the net values slightly at lower funding levels.

the guarantee, but the transition is less abrupt due the smoothing effect of uncertain asset returns. When the pension is invested entirely in the safe asset, the worker knows for certain that she will receive guaranteed benefits if the firm goes bankrupt. By contrast, when some of the pension is invested in the market portfolio, a favorable equity return can push the fund above the guaranteed amount, a possibility that increases the value of the risky pension to the worker. Again, however, it is worth emphasizing that the net value surface reaches a clear maximum at the point of full funding and zero equity risk. (Although we only present the figure for the zero-beta firm, we obtain similar figures for our other parameterizations of correlation, risk aversion, and bankruptcy risk.) Thus, if pension insurance is incomplete and the premium is at least the economically fair value, the firm's optimal decision will be to fully fund its pension and invest the trust fund solely in the risk-free asset.

4.2.3 Case 3: Complete insurance and a mispriced premium

Mispricing of insurance can occur along any or all dimensions of risk: bankruptcy, pension funding, or pension allocation. One way to characterize mispricing along different dimensions of risk is to consider a linearized version of the premium function in equation (8):

$$p(V_1, F_1, \lambda, \Omega) \approx \underbrace{p(V_1, b/R, 0, \Omega)}_{\text{fixed}} + \underbrace{\frac{\partial p(V_1, b/R, 0, \Omega)}{\partial F_1}(F_1 - b/R) + \frac{\partial p(V_1, b/R, 0, \Omega)}{\partial \lambda}\lambda}_{\text{variable}},$$

where Ω is a vector of model parameters, and the first-order Taylor approximation of the pension premium is taken around the point of full defeasance (i.e., $F_1 = b/R$ and $\lambda = 0$). Note that from the insurer's perspective, company risk only matters if the pension is either underfunded or invested in risky assets so that around the point of full defeasance, V_1 has no first-order effect on the value of the premium (i.e., $\partial p(V_1, b/R, 0, \Omega)/\partial V_1 = 0$).

With complete insurance and mispriced insurance, the firm can gain by exploiting the difference between the economic costs of underfunding and the insurance premium. Setting the fixed portion of the premium to zero,[27] the difference between the actual premium and the economically fair premium can be approximated as:

$$p - p^* \approx (p_F - p_F^*)(F_1 - b/R) + (p_\lambda - p_\lambda^*)\lambda,$$

which we can rearrange to solve for the approximate premium:

$$p \approx p^* + (p_F - p_F^*)(F_1 - b/R) + (p_\lambda - p_\lambda^*)\lambda,$$

[27] Without loss of generality, we can absorb the fixed portion in the firm's non-pension debt, D.

where p_F and p_λ denote the partial derivatives of the premium function with respect to F_1 and λ. The firm's costs are therefore approximately given by

$$C \approx w + F_1 - q^c(\tilde{F}_2, b) + q^p(\tilde{F}_2, b|s = 1) + p^* + (p_F - p_F^*)(F_1 - b/R) + (p_\lambda - p_\lambda^*)\lambda.$$

Using the definition of the economically fair premium, $p^* = q^p(\tilde{F}_2, b|s = 0)$, and the fact that $q^p(\tilde{F}_2, b) = q^p(\tilde{F}_2, b|s = 1) + q^p(\tilde{F}_2, b|s = 0)$, we can rewrite the costs as

$$C \approx w + F_1 - q^c(\tilde{F}_2, b) + q^p(\tilde{F}_2, b) + (p_F - p_F^*)(F_1 - b/R) + (p_\lambda - p_\lambda^*)\lambda.$$

Applying put-call parity, we arrive at the following expression for the firm's costs:

$$C \approx w + b/R + (p_F - p_F^*)(F_1 - b/R) + (p_\lambda - p_\lambda^*)\lambda.$$

Since w and b/R are unaffected by its decision, the firm will minimize costs by choosing F_1 and λ to make $(p_F - p_F^*)(F_1 - b/R)$ and $(p_\lambda - p_\lambda^*)\lambda$ as small as possible. If insurance is underpriced along the funding dimension, so that $p_F^* < p_F$,[28] the firm will optimally set $F_1 = 0$. If insurance is underpriced along the investment allocation dimension, so that $p_\lambda^* > p_\lambda$, the firm will optimally invest the pension trust entirely in the risky asset and set $\lambda = 1$.[29] If insurance is mispriced along both dimensions, the firm will only contribute a positive amount to the trust fund if it is legally required to do so, and it will invest in the risky asset in this case. Of course, the Taylor approximation only applies to a local neighborhood around the point of full defeasance. To see how different types of mispricing affect the firm's optimal pension decisions for the entire range of contributions and allocations, we return to our numerical model.

We consider two specific types of mispricing. The first assumes that the pension insurer correctly prices risks due to insufficient contributions but fails to account for the fraction of the fund invested in equities. That is, the insurer effectively assumes that the firm always invests 100 percent of the fund in equities and charges a "fair" premium based on that generally faulty assumption. The second type of mispricing is meant to approximate the variable pricing scheme of the PBGC, which charges a linear premium for underfunding, but essentially ignores credit and allocation risk.

Figure 3 shows how the first type of mispricing affects the net value of the firm's pension when the insurer guarantees 100 percent of promised benefits. As we might expect, the net values spike at the back corner of the figure, where contributions equal the present value of promised benefits and the pension fund is invested entirely in the risky asset. Since

[28] Since the premium should rise when funding falls or when the risky allocation rises, $p_F < 0$ and $p_\lambda > 0$.

[29] If the firm is able to take a leveraged position in the portfolio, the optimal value of λ may be greater than 1.

benefits are fully guaranteed by the insurer, the worker is indifferent to the firm's pension decisions, except to the extent that they affect the probability of bankruptcy (and therefore the potential loss of firm-specific human capital). Not having to worry about compensating the employee for pension risk, the firm is free to exploit the mispricing implicit in the structure of the insurance premium. It does this by maximizing equity exposure in the pension. The resulting increase in the volatility of the underlying asset (the pension fund) increases the value of the firm's pension options (the call less the conditional put) and therefore the net value of offering a risky pension. That is, heads the firm wins; tails the pension insurer loses.

What happens when we move to a type of mispricing that is closer to the variable premium structure of PBGC insurance? Under this new form of mispricing, the insurance agency charges a premium based solely on underfunding, so that $p = a \times \max(b/R - F_1, 0)$, where a is the variable premium rate (e.g., $a = 0.009$ in the PBGC's current pricing scheme). Figure 4 shows that the effect of this type of mispricing on net values depends on the amount of the variable premium. In the left panel of the figure, the pension insurer charges a premium of $a = 0.05$ (which is over six times the size of the actual PBGC premium). At this price, the firm maximizes the net value of its pension by contributing the present value of promised benefits and investing those funds entirely in the risky asset. As was the case with the first type of mispricing, neither the firm nor the worker bear any costs associated with allocation risk.

When the linear variable premium falls from 0.05 to 0.02, as shown in the second panel of Figure 4, a different optimal strategy emerges. At this lower premium, the firm can exploit the mispricing by completely underfunding its pension. Of course, in the real world, ERISA and the tax code impose funding restrictions that limit the extent of plan underfunding. If the firm in our model were required to fund its pension at a certain fraction of promised benefits, the results in Figure 4 suggest that such restrictions would be superfluous for some choices of the premium but not for others. But even if a sufficiently high linear premium acts as a *de facto* funding restriction, the firm in our model will still exploit the mispricing of allocation risk by investing the (fully funded) pension entirely in the risky asset.

4.2.4 Case 4: Incomplete insurance and an underpriced premium

In reality, pension insurance is both incomplete and mispriced.[30] If insurance is incomplete and the firm deviates from the risk-free strategy, the employee will be exposed to some degree of pension risk and will therefore demand additional compensation from the firm. In

[30] The maximum insurance benefit provided by the PBGC is set by law. For plans ending in 2008, the maximum annual benefit is \$51,750 for participants retiring at age 65 (the amount differs by age of retirement and elective survivor benefits).

this case, the firm will trade off any benefits from mispricing against the discipline imposed by the employee in the form of increased wage demands.

We again consider two specific forms of mispricing: one where the insurer fails to price allocation risk and another that represents a simplified version of the PBGC's linear variable rate premium on underfunding. Figure 5 shows the net value surface for the first type of mispricing with a guarantee equal to 75 percent of promised benefits ($g = 0.75 \times b$). The first thing to note about the surface is that the front edge, which corresponds to an investment allocation entirely in the safe asset, looks exactly like the front edge in Figure 2, which shows the net values in the presence of incomplete, but fairly priced, insurance. Thus, if the firm were constrained to invest only in the safe asset, the optimal strategy would again be to fully fund the pension by contributing the present value of promised benefits. In this case, if the firm reduced contributions below full funding, the pension insurer would charge an economically fair amount for the risks associated with underfunding, and the worker would demand additional compensation in exchange for hazarding a portion of promised benefits. The combined cost of insurance and the wage demand in this case would exceed any benefit associated with underfunding.

If the firm is not constrained to invest only in the risk-free asset, however, an alternative strategy presents itself. Examining the back spine of Figure 5, we can see that the net values reach a maximum at the point where contributions are equal to the discounted present value of insured benefits (75 percent of promised benefits) and the pension is allocated entirely in the risky asset. Here, the firm can maximize value by exploiting the insurance agency's mispricing of allocation risk. While it seems intuitive that the firm can exploit this type of mispricing by investing the fund in equities, why does the net value surface reach a maximum along the ridge where contributions equal the present value of guaranteed benefits? This point is special because both the worker and the firm are entirely protected against the consequences of downside market risk. If the market drops even a fraction, benefits remain unchanged. If, on the other hand, the market rises, the worker will receive more than the guaranteed level of benefits. The worker still demands additional compensation for being forced to take on the risk of losing pension benefits in the event of bankruptcy, but this compensation actually falls with increased allocation risk as contributions approach the guaranteed amount. On the other side of the ridge, where contributions fall below the guarantee, the net values decrease rapidly. In this region, the firm has to pay both the worker, in the form of compensating wage demands, and the insurance agency for the risks of underfunding.

Mispricing that resembles the PBGC's variable premium can generate extreme solutions to the firm's optimal pension policy. The net value surface in the left panel of Figure 6 corresponds to the case of a relatively high linear premium equal to 5 percent of unfunded guaranteed benefits. The right panel of the figure shows the surface for a lower premium equal to 2 percent of underfunding. The net value surfaces in both panels display a marked "V" shape, with values rising on either side of the guaranteed level of funding. The surface rises for contributions above the guarantee because the worker offers wage concessions in exchange for a less risky pension. Below the guarantee, the employee becomes increasingly indifferent to the funding level and allocation of the pension. Here, the net values respond primarily to the change in the value of the firm's pension options (the call less the conditional put) relative to the linear premium on underfunding. When the premium is high, as in the left panel of the figure, the firm prefers to fully fund its pension and invest the fund entirely in the risky asset. (Again, with a linear premium, there is no marginal cost of allocation risk.) However, when the premium is sufficiently low, as in the right panel of the figure, the firm can best take advantage of the mispricing by completely underfunding its pension, or at least funding it at the lowest amount allowable by law.

Since the real world features both incomplete insurance and underpriced premiums, our findings in this case help explain why so many firms offer risky pension plans rather than risk-free promises.

5 Discussion and Interpretation of Results

Unless pension insurance is underpriced, we find that firms will increase shareholder value by offering its employees a risk-free pension promise, because the cost to the firm of compensating the employee for risk exceeds the benefit to the firm from lower expected contributions. Of course, our stylized model omits taxes, labor market frictions, regulations, and imperfect information. We briefly discuss how these issues might affect our benchmark finding that a risk-free pension is optimal from the firm's perspective.

5.1 Taxes

In general, the tax treatment of pensions provides additional incentive for firms to invest the trust fund in the risk-free asset. Trust-fund income is tax-exempt, while other corporate income is not, and stock returns are generally taxed at the lower capital gains rate while fixed-income returns are taxed at the higher regular rate. As demonstrated in Tepper (1981) and Black (1980), it is therefore optimal from a tax-minimizing perspective to invest the trust fund in the higher-taxed asset (i.e., fixed-income) and concentrate the lower-taxed asset (i.e., stocks) outside the trust fund. Thus, including taxes in the model would only strengthen our benchmark finding that firms should optimally invest the trust fund in the risk-free asset.

5.2 Labor market frictions

The assumption of a competitive labor market is central to our model. Firms take required total compensation as given and allocate wages and pension benefits accordingly. But, in a world with labor market frictions related to search costs and firm-specific human capital, workers might accept less-than-competitive compensation and still remain at the firm.[31] In such a case, would a risk-free pension still be an optimal strategy? It would, because the central factors of the model would not change—the risk-averse employee's inability to hedge bankruptcy risk still makes the pension promise an inefficient way to transfer risk. In this case, a firm with market power in the labor market is better off minimizing its costs by reducing wages than by introducing risk into the pension.

A related extension is to consider the "implicit contract" model of Ippolito (1985a,b), in which firms optimally pay younger workers less than their marginal product, and older workers more than their marginal product, as a way to encourage workers to stay with the firm. In this case, the firm takes on the risk that rapid wage growth will lead to rapid growth in the pension promise, which is often tied to wages.[32] Note that in this extended model, the pension promise becomes stochastic to the firm as well as the employee. Lucas and Zeldes (2006) have argued that if firms face this kind of risk it can be optimal to invest a portion of the pension trust fund in equities, to the extent that stochastic equity returns are correlated with stochastic wage growth.

We make two observations on this point. First, while a competitive firm takes the utility of compensation as given, it is free to arrange a pension contract that reduces risk to itself by weakening the dependence of the pension promise on a stochastic variable such as wages. That is, the firm is not forced to accept risk for itself in its pension promise; it could eliminate this risk simply by promising a fixed, rather than wage-derived, pension benefit. Second, and more importantly, in our framework the firm would still not find it optimal to invest the trust fund in the risky asset, because again the central factors—a risk-averse employee facing undiversifiable bankruptcy risk—have not changed. That is, even if the firm and employee faced an uncertain pension promise whose realization were correlated with the risky asset, it would not be optimal from the employee's perspective to compound the pension-promise uncertainty by adding on pension-*funding* uncertainty. The employee will always prefer any given pension promise to be funded risklessly. As a result, the firm will be better off funding the expected pension benefit risklessly and hedging its wage-derived pension-promise risk by investing in the correlated risky assets *outside* the

[31]For example, Acemoglu and Pischke (1999) demonstrate that firms can exploit the wedge between internal and external marginal products by investing in both general and specific firm-sponsored training.

[32]Modeling this extension formally would require moving beyond our current two-period model, in which there is no wage growth because wages are paid only in the first period.

trust fund, topping up the trust fund, in the high-wage-growth state of the world, in the later period. This "secondary" (and non-tax-exempt) trust fund helps protect the firm from wage risk, while the primary (tax-exempt) trust fund protects the employee from funding risk.

5.3 Financing considerations and regulations

Our model implicitly assumes that firms are indifferent between internal and external sources of finance, and that pension contributions have no effect on a firm's ability to raise resources for capital investment. If, however, firms face financial constraints, pension contributions may crowd out a portion of internal funds available for new capital projects. Rauh (2006) exploits variation in minimum funding requirements to show that mandatory pension contributions have a strong and negative effect on capital expenditures and that the effect is strongest among younger firms with the greatest need for cash flows. Since underfunding and investing the trust fund in the risky asset can reduce expected contributions, a wedge between the costs of internal and external finance could explain why many firms pursue riskier pension strategies than our model predicts.

But again, in our framework, this result would not hold. While pension costs can be reduced by underfunding and investing the trust fund in risky assets, *total* costs would increase, due to the additional required compensation for risk. Thus, rather than freeing up resources for investment, adding risk to the pension promise would reduce investment resources.

5.4 Employees, employers, shareholders, and taxpayers

If taxes, labor-market frictions, and financing constraints do not provide a clear rationale for offering risky pensions, what else might explain the current practice of providing risky rather than risk-free pension promises? One possibility is informational and institutional barriers affecting employees, employers, shareholders, or taxpayers. For example, employees in our model are assumed to fully understand and value the effect of exposure to pension risk on their utility. If employees do not understand or value the pension risk, they will not demand additional compensation for it, and firms could reduce their cost in expected value by adopting a risky funding strategy.[33] Until recently, this scenario was not necessarily far-fetched. With relatively low incidence of pension failures (and the failures that did occur limited to particular industries such as steel and airlines), it would not be too surprising for employees to consider their pension promises to be essentially risk-free, despite the fact that the trust funds might be underfunded and heavily invested in risky assets. In recent

[33] Mitchell and Utkus (2003) examine the role of own-company stock in 401(k) plans and find evidence that plan participants systematically underestimate their exposure to own-firm risk.

years, however, following a wave of high-profile pension failures, freezes, and terminations, it seems likely that more employees are aware of the risk inherent in their pension plans. In any case, the possibility of informational barriers to employees raise the important public policy issue of whether employees should receive less compensation simply because they fail to fully understand the risk properties of their pension benefits.

Employers, for their part, may face market imperfections which limit their ability to perfectly defease pension risk (such as the lack of sufficient long-dated fixed-income instruments). Nonetheless, market instruments exist that allow firms to come fairly close to full defeasance, such as synthetic fixed-rate securities based on swaps contracts or eurodollar deposits.

Another potential informational barrier may arise among shareholders. Franzoni and Marín (2006) find that the market systematically misprices the cash flow and earnings implications of underfunded pensions. One reason shareholders may have a difficult time valuing pension plans is because of the historically untransparent set of accounting rules that applied to a plan's assets and liabilities. Under FASB guidelines, firms are essentially able to book the equity premium on trust fund assets as a source of income, with no adjustment for the risk in reported earnings. Even though the market values of assets and liabilities were reported in footnotes, Coronado and Sharpe (2003) and Coronado, Mitchell, Sharpe, and Nesbitt (2008) show that markets generally responded only to the unadjusted pension accounting in the main text of the statements. With improved accounting transparency, one might expect the market to more accurately price the riskiness of a firm's pension plan.[34]

The final source of barrier we consider is among taxpayers. Our model shows that in the presence of mispriced pension insurance, firms *and* employees can benefit from risky pension decisions, but only at the expense of the insurance agency (i.e., the PBGC). But the only way the PBGC can continue to offer underpriced pension insurance without itself becoming insolvent is if it receives an implicit guarantee from taxpayers. Fully informed taxpayers would impose discipline on firms via public policy, but there may be political economy considerations that hamper this process of accountability.

6 Conclusion

We find that the combination of competitive labor markets and risk averse employees who cannot hedge bankruptcy risk makes a risky pension promise an inefficient form of compensation from the standpoint of shareholders. In our model, the optimal pension strategy, in the absence of mispriced pension insurance, is to contribute the present value of benefits

[34]FASB has announced new guidelines requiring firms to bring the fair market value of pension trust funds onto their corporate balance sheets (FASB, 2006), and is currently working on revisions that could require firms to recognize the fair value of pension changes in their earnings statements.

to the pension trust fund and to invest those resources entirely in the risk-free asset. An important caveat, however, is that in the presence of underpriced pension insurance, this result breaks down, because firms and employees can effectively benefit from dividing a surplus gained at the taxpayers' expense. Since PBGC insurance is often characterized as underpriced (Brown, 2008; CBO, 2005), this provides an explanation for the large difference between our benchmark results and real-world behavior. Additional explanations are also available, including informational and institutional barriers that slow or prevent the benchmark economic forces from resulting in the theoretical equilibrium.

Nevertheless, the simple framework of our model highlights the economic tradeoffs facing regulators, employers, pension trust fiduciaries, and employees. For instance, to the extent that current funding practices are motivated by underpriced pension insurance, opaque accounting standards, or imperfect information on the part of the employees, regulators may choose to address these imperfections by fairly pricing pension insurance, closing funding loopholes, and improving accounting transparency. The Deficit Reduction Act of 2005, the Pension Protection Act of 2006, and FASB's ongoing project to revisit pension accounting standards all represent movement in this direction, but it remains to be seen whether these changes will be sufficient to prevent promised benefits from being exposed to considerable risk.

References

ACEMOGLU, D., AND J.-S. PISCHKE (1999): "The Structure of Wages and Investment in General Training," *The Journal of Political Economy*, 107(3), 539–572.

BADER, L. (2003): "The Case agaist Stock in Corporate Pension Funds," *Pension Section News (Society of Actuaries)*, pp. 17–19.

——— (2004): "Pension Deficits: An Unnecessary Evil," *Financial Analysts Journal*, 60(3), 22–32.

BECKER, G. S. (1962): "Investment in Human Capital: A Theoretical Analysis," *The Journal of Political Economy*, 70(5), 9–49.

BLACK, F. (1980): "The Tax Consequences of Long-Run Pension Policy," *Financial Analysts Journal*, 36, 21–29.

BODIE, Z. (1988): "Pension Fund Investment Policy," *NBER Working Paper No. 2752.*

——— (1990): "The ABO, the PBO, and Pension Investment Policy," *Financial Analysts Journal*, 41, 10–16.

BRANNICK, N. J. (2004): "At the Crossroads of Three Codes: How Employers Are Using ERISA, the Tax Code, and Bankruptcy to Evade Their Pension Obligations," *Ohio State Law Journal*, 65, 1577–1626.

BROWN, J. R. (2008): "Guaranteed Trouble: The Economic Effects of the Pension Benefit Guaranty Corporation," *Journal of Economic Perspectives*, 22(1), 177–198.

BULOW, J. (1982): "What Are Corporate Pension Liabilities?," *Quarterly Journal of Economics*, 97, 435–452.

CAI, J., AND A. VIJH (2005): "Executive Stock and Option Valuation in a Two State-Variable Framework," *Journal of Derivatives*, 12, 9–27.

CBO (2005): *The Risk Exposure of the Pension Benefit Guarantee Corporation.* Washington.

CORONADO, J. L., O. S. MITCHELL, S. A. SHARPE, AND S. B. NESBITT (2008): "Footnotes Aren't Enough: The Impact of Pension Accounting on Stock Values," NBER Working Paper 13726.

CORONADO, J. L., AND S. A. SHARPE (2003): "Did Pension Plan Accounting Contribute to a Stock Market Bubble?," *Brookings Papers on Economic Activity*, 1, 323–359.

FASB (2006): "Statement of Statement of Financial Accounting Standard No. 158 Standards No. 158: Employers' Accounting for Defined Benefit Pension and other Postretirement Plans," .

FRANZONI, F., AND J. M. MARÍN (2006): "Pension Plan Funding and Stock Market Efficiency," *Journal of Finance*, 61(2), 921–956.

Gold, J. (2003): "Creating Value in Pension Plans (Or, Gentleman Prefer Bonds)," *Journal of Applied Corporate Finance*, 15(4), 51–57.

Hall, B. J., and K. J. Murphy (2002): "Stock Options for Undiversified Executives," *Journal of Accounting and Economics*, 33(1), 3–42.

Harrison, J., and W. Sharpe (1983): "Optimal Funding and Asset Allocation Rules for Defined Benefit Pension Plans," in *Financial Aspects of the U.S. Pension System*, pp. 92–106. University of Chicago Press.

Hashimoto, M. (1981): "Firm-Specific Human Capital as a Shared Investment," *The American Economic Review*, 71(3), 475–482.

Ippolito, R. (1985a): "The Economic Function of Underfunded Pension Liabilities," *Journal of Law and Economics*, 28(3), 611–51.

——— (1985b): "The Labor Contract and True Economic Pension Liabilities," *American Economic Review*, 75(5), 1031–43.

Jarrow, R., and A. Rudd (1983): *Option Pricing*. Richard D. Irwin, Inc.

Love, D. A., P. A. Smith, and D. Wilcox (2007): "Why Do Firms Offer Risky Defined-Benefit Pension Plans?," *National Tax Journal*, 60(3), 507–520.

Lucas, D., and S. Zeldes (2006): "Valuing and Hedging Defined Benefit Pension Obligations—The Role of Stocks Revisited," Working Paper, Northwestern University.

Milliman Consultants and Actuaries (2006): "Milliman 2005 Pension Study," .

——— (2007): "Milliman 2006 Pension Study," .

Mitchell, O. S., and S. P. Utkus (2003): "The Role of Company Stock in Defined Contribution Plans," in *The Pension Challenge: Risk Transfers and Retirement Income Security*, ed. by O. S. Mitchell, and K. Smetters. Oxford University Press.

Munnell, A. H. (1982): *The Economics of Private Pensions*. The Brookings Institution.

Pension Benefit Guaranty Corporation (2003): *2002 Annual Report*.

Ralfe, John, C. S., and J. Palin (2004): "Pensions and Capital Structure: Why Hold Equities in the Pension Fund?," *North American Actuarial Journal*, 8, 103–113.

Ralfe, J., C. Speed, and J. Palin (2004): "Pensions and Capital Structure: Why Hold Equities in the Pension Fund?," *North American Actuarial Journal*, 8, 103–113.

Rauh, J. (2006): "Investment and Financing Constraints: Evidence from the Funding of Corporate Pension Plans," *Journal of Finance*, 61(1), 33–71.

——— (2007): "Risk Shifting versus Risk Management: Investment Policy in Corporate Pension Plans," *NBER Working Paper No. 13240*.

Rubinstein, M. (1994): "Return to OZ," *Risk*, 4, 63–66.

Sharpe, W. (1976): "Corporate Funding Policy," *Journal of Financial Economics*, 3, 183–193.

Tepper, I. (1981): "Taxation and Corporate Pension Policy," *Journal of Finance*, 36(1), 1–13.

Treynor, J. (1977): "The Principles of Corporate Pension Finance," *Journal of Finance*, 32(2), 627–638.

Wilcox, D. (2006): "Reforming the Defined-Benefit Pension System," *Brookings Papers on Economic Activity*, 1, 235–304.

Tables and Figures

Table 1: Funding of the Largest 100 Corporate Pension Plans

	2001	2002	2003	2004	2005	2006
1. Agg. Funding Ratio (%)	102.2	82.6	88.8	89.9	91.4	99.7
Distribution of Funding Ratios (# of Plans):						
2. <100%	51	74	79	75	77	55
3. <90%	30	61	61	48	47	29
4. <80%	13	44	26	22	20	10
5. Contribs (billion $)	10.6	34.3	56.9	42.7	44.9	35.7

Ratio of the market value of assets to the projected benefit obligation. Source: Authors' calculations from Milliman Consultants and Actuaries (2006) and Milliman Consultants and Actuaries (2007).

Table 2: Asset Allocation of the Largest 100 Corporate Pension Plans in 2006

	Share in Risky Assets*
1. Aggregate (%)	71
Distribution of Risky Asset Allocations (# of Plans):	
2. >80%	10
3. 70-80%	50
4. 60-70%	29
5. <60%	11

Source: Authors' calculations from Milliman Consultants and Actuaries (2006) and Milliman Consultants and Actuaries (2007).
*Risky Assets defined as equities plus "other investments.".

Table 3: Net Value ($\times 100$) to the Firm of a Risky Pension Strategy*

Allocation λ	CRRA $\gamma = 3$ $F_1\%$= 90	99	100	CRRA $\gamma = 5$ 90	99	100
	Corr. between market and firm $\rho = 0.0$					
0.00	-19.22	-1.31	0.00	-33.41	-2.23	0.00
0.01	-19.25	-1.33	-0.09	-33.45	-2.25	-0.15
0.50	-25.60	-8.17	-6.95	-44.22	-13.95	-11.88
1.00	-40.80	-21.56	-19.98	-71.73	-37.62	-34.84
	Corr. between market and firm $\rho = 0.5$					
0.00	-21.22	-1.39	0.00	-36.59	-2.39	0.00
0.01	-21.62	-1.59	-0.19	-37.28	-2.74	-0.34
0.50	-49.51	-20.91	-18.52	-85.65	-36.20	-32.11
1.00	-96.39	-60.79	-57.50	-168.39	-106.00	-100.27
	Corr. between market and firm $\rho = -0.5$					
0.00	-13.27	-0.90	0.00	-23.47	-1.57	0.00
0.01	-13.02	-0.78	-0.01	-23.00	-1.33	-0.02
0.50	-6.63	-1.03	-0.78	-11.21	-1.70	-1.29
1.00	-6.73	-2.25	-1.97	-11.52	-3.81	-3.34

*This table reports net values for different combinations of relative risk aversion, γ, and correlations between the market portfolio and the firm's assets, ρ. The net value is defined as the value of the firm's benefit from offering a pension less the compensation premium to the worker. $F_1\%$ is the percent of the first-period contribution relative to "full funding": $100 \times F_1 R/b$. The other parameters are as follows: $V_1 =$ 250,000; $D = 0.72 \times V_1$; $\sigma_m = 0.18$; $\sigma_v = 0.20$; $b = 250$; $w_1 = 1{,}000$; $w_2 = 250$; and $\chi = 0.90$. The parameters imply a 5% probability of bankruptcy at the point of full funding.

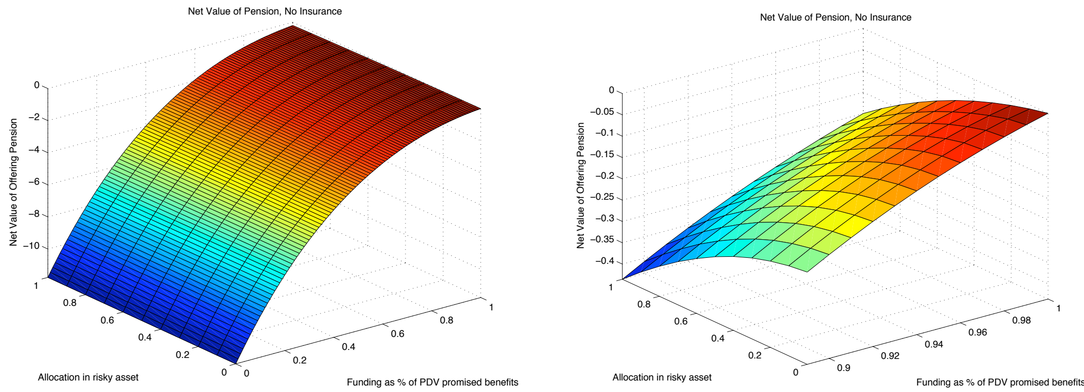

(a) Baseline, no insurance

(b) Baseline: a closer look

Figure 1: Net value surface for contributions ranging from 0 to 100% of PDV promised benefits, b/R, and for pension allocation from 0 to 100% in the market portfolio. The net value equals the estimated value of the firm's benefit from offering a pension, holding wages constant, minus the compensation premium to the worker. The model specification is as follows: $V_1 = 250{,}000$; $D = 0.72 \times V_1$; $\sigma_m = 0.18$; $\sigma_v = 0.20$; $b = 250$; $w_1 = 1{,}000$; $w_2 = 250$; and $\chi = 0.90$. The parameters imply a 5% probability of bankruptcy at the point of full funding.

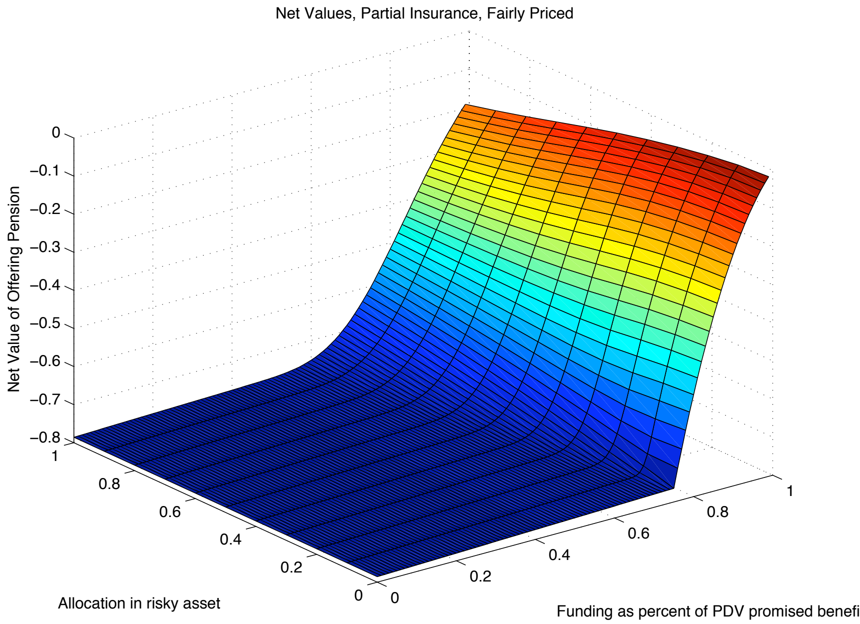

Figure 2: Effect of fair insurance guaranteeing 75% of promised benefits. Net value surface for contributions ranging from 0 to 100% of PDV promised benefits b/R and for pension allocation from 0 to 100% in the market portfolio. The net value equals the estimated value of the firm's benefit from offering a pension, holding wages constant, minus the compensation premium to the worker. The model specification is as follows: $V_1 = 250{,}000$; $D = 0.72 \times V_1$; $\sigma_m = 0.18$; $\sigma_v = 0.20$; $b = 250$; $w_1 = 1{,}000$; $w_2 = 250$; and $\chi = 0.90$. The parameters imply a 5% probability of bankruptcy at the point of full funding.

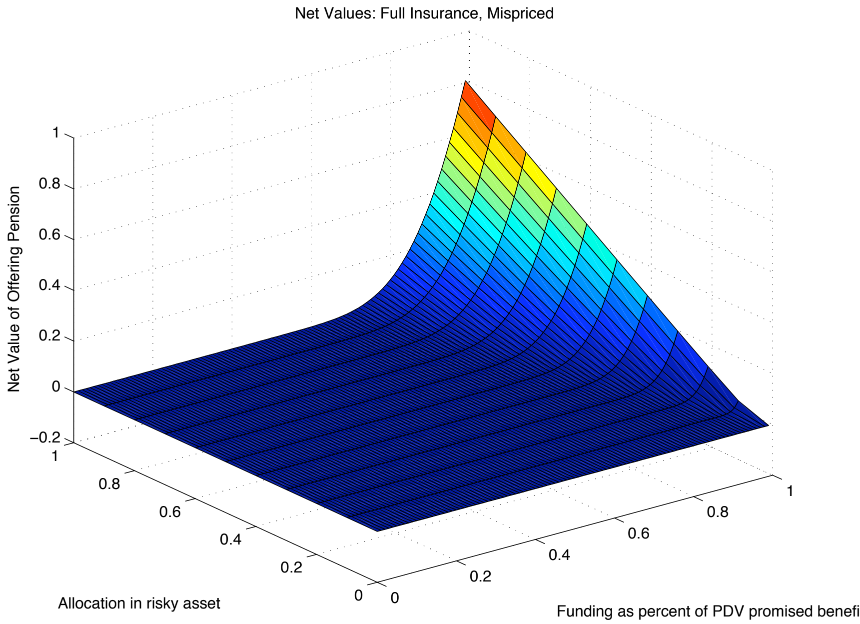

Figure 3: Effect of mispriced insurance guaranteeing 100% of promised benefits. The mispricing takes the form of charging the fair insurance premium that would obtain if the allocation were invested entirely in the safe asset (regardless of the actual allocation). Net value surface for contributions ranging from 0 to 100% of PDV promised benefits b/R and for pension allocation from 0 to 100% in the market portfolio. The net value equals the estimated value of the firm's benefit from offering a pension, holding wages constant, minus the compensation premium to the worker. The model specification is as follows: $V_1 = 250{,}000$; $D = 0.72 \times V_1$; $\sigma_m = 0.18$; $\sigma_v = 0.20$; $b = 250$; $w_1 = 1{,}000$; $w_2 = 250$; and $\chi = 0.90$. The parameters imply a 5% probability of bankruptcy at the point of full funding.

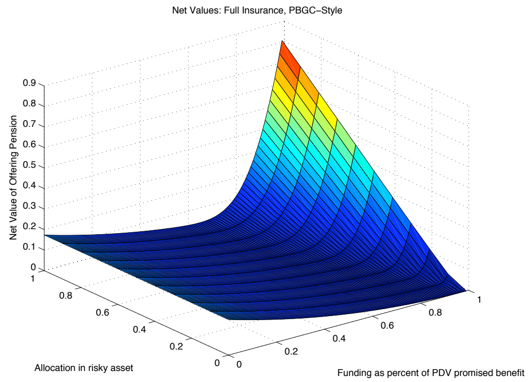

(a) Premium $= 0.05 \times \max(b/R - F_1, 0)$

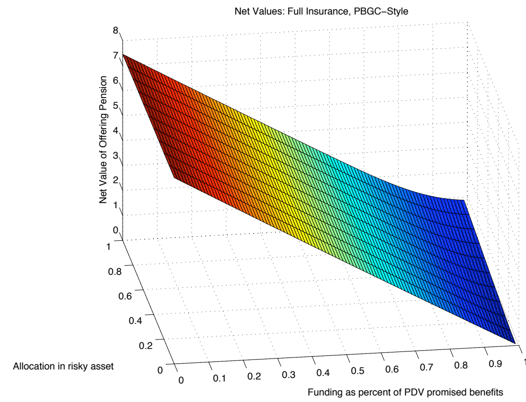

(b) Premium $= 0.02 \times \max(b/R - F_1, 0)$

Figure 4: Effect of PBGC-style mispricing: full guarantee. The mispricing takes the form of charging a linear premium based solely on underfunding (note that the premium in the left panel is over 6 times the actual PBGC variable premium of 0.009 per dollar of underfunding, while the premium in the right panel is over twice as large). Net value surface for contributions ranging from 0 to 100% of PDV promised benefits b/R and for pension allocation from 0 to 100% in the market portfolio. The net value equals the estimated value of the firm's benefit from offering a pension, holding wages constant, minus the compensation premium to the worker. The model specification is as follows: $V_1 = 250{,}000$; $D = 0.72 \times V_1$; $\sigma_m = 0.18$; $\sigma_v = 0.20$; $b = 250$; $w_1 = 1{,}000$; $w_2 = 250$; and $\chi = 0.90$. The parameters imply a 5% probability of bankruptcy at the point of full funding.

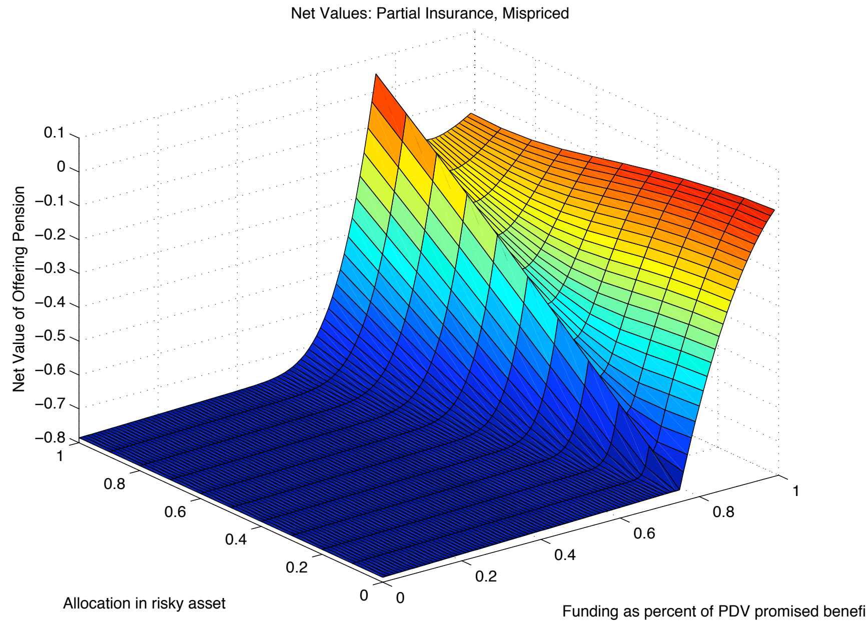

Figure 5: Effect of mispriced insurance guaranteeing 75% of promised benefits. The mispricing takes the form of charging the fair insurance premium that would obtain if the allocation were invested entirely in the safe asset (regardless of the actual allocation). Net value surface for contributions ranging from 0 to 100% of PDV promised benefits b/R and for pension allocation from 0 to 100% in the market portfolio. The net value equals the estimated value of the firm's benefit from offering a pension, holding wages constant, minus the compensation premium to the worker. The model specification is as follows: $V_1 = 250{,}000$; $D = 0.72 \times V_1$; $\sigma_m = 0.18$; $\sigma_v = 0.20$; $b = 250$; $w_1 = 1{,}000$; $w_2 = 250$; and $\chi = 0.90$. The parameters imply a 5% probability of bankruptcy at the point of full funding.

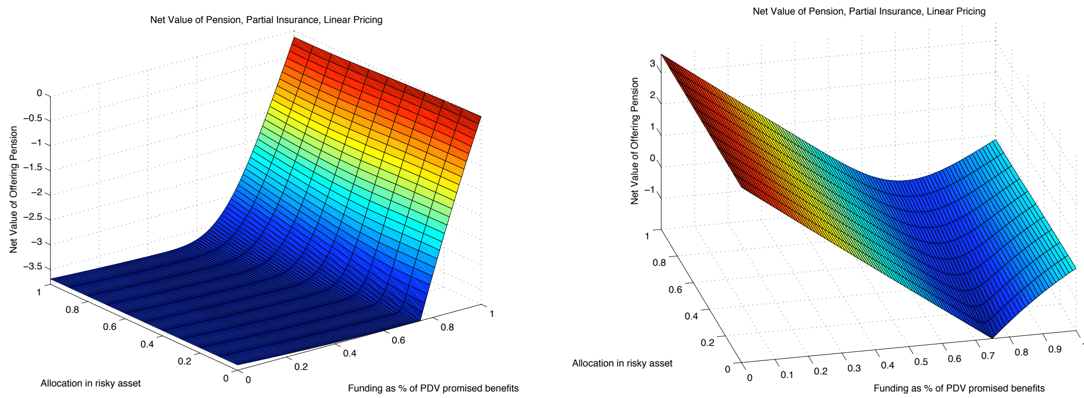

(a) Premium $= 0.05 \times \max(g/R - F_1, 0)$ (b) Premium $= 0.02 \times \max(g/R - F_1, 0)$

Figure 6: Effect of PBGC-style mispricing: incomplete guarantee. The mispricing takes the form of charging a linear premium based solely on underfunding relative to the guarantee (note that the premium in the left panel is over 6 times the actual PBGC variable premium of 0.009 per dollar of underfunding, while the premium in the right panel is over twice as large). Net value surface for contributions ranging from 0 to 100% of PDV promised benefits b/R and for pension allocation from 0 to 100% in the market portfolio. The net value equals the estimated value of the firm's benefit from offering a pension, holding wages constant, minus the compensation premium to the worker. The model specification is as follows: $V_1 = 250{,}000$; $D = 0.72 \times V_1$; $\sigma_m = 0.18$; $\sigma_v = 0.20$; $b = 250$; $w_1 = 1{,}000$; $w_2 = 250$; and $\chi = 0.90$. The parameters imply a 5% probability of bankruptcy at the point of full funding.

Appendix: Valuation of Options with Correlated Returns

This appendix provides a brief description of Rubinstein's (1994) method for approximating a correlated return process as a bivariate binomial distribution and explains how we adapt it to account for the possibility of bankruptcy. Letting h denote the length of each time interval we consider and T denote the total number of time periods, there are $n+1$ unique values of the firm return and the market return, where $n = T/h$. Since we are pricing European, and not American, options, we only need to keep track of the terminal values of the options at the end of the T periods.

We construct the $n+1$ values of our return variables by first considering the terminal realizations of a pair of standard binomial random variables, u and v. To allow for correlation, we can define $w = u$ and $z = \rho u + v\sqrt{1-\rho^2}$, which are themselves binomial random variables. For $j = 0, ..., n$ and $k = 0, ..., n$ moves, each pair of standardized realizations, $(w_{j,k}, z_{j,k})$, is given by:

$$w_{j,k} = [j - (n-j)]/\sqrt{n} \tag{A-1}$$

$$z_{j,k} = \rho[j-(n-j)]/\sqrt{n} + \sqrt{1-\rho^2}[k-(n-k)]/\sqrt{n}, \tag{A-2}$$

with associated probabilities given by the binomial pdf.[35]

We adopt the Jarrow-Rudd Jarrow and Rudd (1983) binomial approximation of the returns, where an "up" movement equals $e^{(r-\sigma/2)h+\sigma\sqrt{h}}$ and a "down" movement equals $e^{[(r-\sigma/2)h-\sigma\sqrt{h}]}$. Note that after n moves, our standard binomial distribution will return $w_{j,k} = n/\sqrt{n} = \sqrt{n}$ for $j = n$. After n "up" moves, the Jarrow-Rudd model returns a value of $e^{[(r-\sigma/2)hn+\sigma\sqrt{h}n]}$. Thus, we can convert the standard binomial $w_{j,k}$ values to the Jarrow-Rudd values as:

$$R_m(j,k) = e^{[(r-\sigma_m/2)T+\sigma_m\sqrt{T}w_{j,k}]}, \tag{A-3}$$

for the risk-neutral return on the market portfolio, and

$$R_v(j,k) = e^{[(r-\sigma_v/2)T+\sigma_v\sqrt{T}z_{j,k}]}, \tag{A-4}$$

for the risk-neutral return on the firm's assets. (Note that since the market return values depend only on the realization of $w_{j,k}$, itself a function only of j, we can write $R_m(j,k) = R_m(j)$. To obtain the returns corresponding to the "natural" probability space, we replace the risk-free rate of return r with the mean log rate of returns on the market and firm assets: μ_m and μ_v.

Bankruptcy occurs when the value of the firm drops below a threshold that depends on non-pension debt and the pension shortfall. In general, the probability of bankruptcy will therefore differ depending on the realization of the market return. Associated with each of our $n+1$ discrete values of the market return is a discrete list of second-period firm values ($V_2(i, 0, ..., n) = V_1 R_v(i, 0, ..., n)$) and their associated cumulative probability distribution. By interpolating linearly across this list, we can closely approximate the probability that a

[35] The formulas for $w_{j,k}$ and $z_{j,k}$ follow from a direct application of the formula for a standard distribution. Here, we subtract off the mean of a binomial random variable, $pn = 0.5n$, and divide by the standard deviation, $\sigma = \sqrt{np(1-p)} = 0.5\sqrt{n}$. Thus, for $w_{j,k}$, we have $(j - 0.5n)/(0.5\sqrt{n}) = (2j-n)/\sqrt{n} = [j-(n-j)]/\sqrt{n}$.

combination of debt, D, and pension shortfall, $\max(b - F_2, 0)$, will lead to bankruptcy. Thus, for each combination of pension funding, F_1, and pension asset allocation, λ, we can create a list of $n+1$ bankruptcy probabilities (risk neutral, for the market's valuation; "natural," for the worker's problem) corresponding to each possible j realization of the market return. Let $bp(j)$ denote the jth value of this list.

The risk-neutral payoff of the call option on the pension surplus is given by

$$q^c(F, b) = \frac{1}{R} \sum_{j=0}^{n} \max([R + (R_m(j) - R)\lambda]F_1 - b, 0). \tag{A-5}$$

Similarly, the risk-neutral payoff of the conditional put option on the pension surplus is given by

$$q^p(F, b \mid s = 1) = \frac{1}{R} \sum_{j=0}^{n} bp(j) \max(b - [R + (R_m(j) - R)\lambda]F_1, 0), \tag{A-6}$$

where the payoffs are each multiplied by the probabilities, $bp(j)$, obtained under the risk-neutral distribution.