



Article from

In the Public Interest

September 2016

Issue 13

Understanding the Argument for Market Valuation of Public Pension Liabilities

By Kent Smetters and Andrew G. Biggs

This article first appeared on www.aei.org. It is reprinted here with permission of the American Enterprise Institute. The views expressed in this article are solely those of the author. The Society of Actuaries takes no position on the views of the author.

Most public-sector employees participate in traditional defined-benefit pension plans, which promise them a fixed monthly retirement benefit for life. These benefits are generally calculated as some percentage of the employee's final salary multiplied by the number of years of employment. Defined-benefit pension plans differ from defined-contribution ones such as 401(k) and 403(b) plans that are common in the private sector, in which the employer contributes to the employee's investment account each year but makes no promises regarding the actual benefit the employee will receive at retirement.

Accounting for the finances of defined-benefit pension plans requires comparing the assets the plan holds today with a stream of benefits that can extend decades into the future. Making such comparisons requires "discounting" future benefit liabilities to the present, a process that subtracts annual interest from the future dollar amount until a "present value" is determined. The policy debate regards the appropriate discount rate to utilize in making such calculations. A higher discount rate will reduce the present value of plan liabilities and, all other things equal, portray a plan as being better funded. Likewise, lower discount rates generate higher measured liabilities and lower levels of plan funding.

Determining the appropriate discount rate to use is a function of the goals of pension policy as a whole. Pension accounting and funding rules should be designed to help plan stakeholders better achieve their policy goals. These stakeholders can include pension managers, elected officials, public employees and retirees, holders of state and municipal bonds, and taxpayers, all both present and future. In the public pensions accounting debate, however, these policy goals are often left unstated. Making these goals explicit illustrates the deficiencies in the current pension accounting rules and points the way toward better methods.

The goals of pension plan stakeholders constitute what economists refer to as an "objective function." The pensions literature

indicates that two goals are widely shared across most public pension plans. First, in a defined-benefit pension plan, the benefits retirees and other beneficiaries receive should be free of risk. Unlike defined-contribution plans, for which value changes from day to day with returns in the market, defined-benefit plans are intended to be guaranteed. This goal is embraced at many levels, from the design of benefit formulas, to communications with public employees, to laws and constitutional provisions that protect accrued pension benefits.

Second, each generation of taxpayers should fully fund the benefits accruing to public employees at that time, rather than shifting those costs to future generations. Although pension benefits may not be payable for years or decades in the future, those benefits compensate services rendered by public employees today. Because today's taxpayers are the beneficiaries of those services, current taxpayers should fully fund the compensation of public employees who provide them. The Governmental Accounting Standards Boards (GASB) refers to this standard as "interperiod" or "intergenerational equity," saying that it means that "taxpayers of today pay for the services that they receive and the burden of payment for services today is not shifted to taxpayers of the future."¹ Inter-period equity requires that future taxpayers be insulated or, in economists' terms, "immunized" against the risk of being forced to pay for pension benefits accrued in the past.

We can thus say that a pension plan might be considered "fully funded" if it satisfies both of these goals—if it can provide guaranteed retirement benefits to plan participants without imposing additional costs on future generations. Lacking these two criteria, the very idea of being fully funded loses meaning. If a pension plan counts toward its funding the right to at any time reduce benefits to retirees or return to the taxpayer for additional funding, then any public plan should be considered "fully funded." This reasoning makes little sense.

However, under current pension accounting rules promulgated by the GASB, a pension could consider itself fully funded even if its current assets had less than a 50 percent probability of being sufficient to pay for the benefits the plan has currently promised. Under these circumstances, even a supposedly fully funded plan either would be unable to guarantee payment of full benefits or would have to impose a contingent liability on future taxpayers to supplement plan funding should the need arise. By focusing on the lower standard of the mere expectation of being able to pay benefits while ignoring the fact that those benefits must be paid, current pension accounting standards fail to assist pension stakeholders in achieving the widely held policy goals of producing guaranteed benefits to retirees without shifting costs to future taxpayers.

In recent years, the economics profession has challenged current pension accounting rules, arguing that they fail to capture the full value of pension liabilities and therefore encourage plans to become underfunded. Most of the economists involved in these challenges are employed as academics in think tanks or in government agencies where they are not directly involved with managing plans or setting pension accounting standards. Individuals more closely tied to public pension plans, including the GASB, public pension actuaries, and employee interest groups, have pushed back against this challenge.

However, economists' underlying arguments are poorly understood, with the result that the public often misconstrues the issues at hand in the pension accounting debate. For instance, many believe that when economists argue that public pensions should not discount their liabilities at the 8 percent interest rate that pensions assume for their investments, they are claiming that actual investment returns will be lower than 8 percent. Similarly, some believe that because economists advocate using lower discount rates to value pension liabilities, they are arguing that public pensions should hold only low-risk investments in their portfolios. Neither belief is correct.

In fact, most economists believe that pension liabilities should be valued using low discount rates even if pensions continue to invest in stocks and other risky but high-return assets and even if their assumptions regarding future average investment returns are accurate. However, the theories and assumptions underlying a fair-market valuation approach are so ingrained among economists that, in many cases, the proponents fail to make these background arguments explicit. This allows misunderstandings regarding the pension accounting debate to continue. Properly understood, the economic argument motivating the movement for fair-market valuation differs from the way it is often portrayed in the news media.

This paper first reviews how public pensions value their liabilities under current GASB rules. Next, we outline the standard approach to valuing liabilities from an economic point of view and what this market-based approach implies for public-sector pensions and their funding levels. Following that, we provide examples designed to better convey the qualitative principles regarding the economic approach to pension liability valuation.

The emphasis here is not on detailed calculations of how fair-market valuation would affect pension funding in states and cities around the country, nor the increased budgetary burden the pensions might impose.² Likewise, the emphasis is not on how defined-benefit pensions might be reformed in light of information conveyed via more accurate accounting rules.

Rather, the intent is to provide readers with a better handle on the simple intuition that lies behind the economists' call for

GASB accounting rules help pension plans calculate a “best guess” ... a roughly 50 percent probability of being able to meet its benefit obligations ...

fair-market valuation of public pension liabilities. Those who follow the debate are aware that economists argue for using lower discount rates to value public pension liabilities but often are unaware of why economists believe what they do. This paper aims to better articulate those beliefs.

BACKGROUND ON PUBLIC EMPLOYEE PENSION PLANS

Most state and local governments provide a defined-benefit pension plan for public employees as part of their overall compensation.³ These plans generally provide for retirement, disability, and survivors' benefits and may either supplement or substitute for Social Security benefits. Around three-quarters of all state and local government employees take part in an employer-provided pension plan, with coverage rates among full-time employees being significantly higher. Around 80 percent of public employees have only a defined-benefit pension, with 14 percent having only a defined-contribution pension and 6 percent having both.⁴

State and local pensions work similarly to defined-benefit pensions in the private sector. Once vested, an employee becomes entitled to a benefit based on a percentage of final salary—ordinarily an average of earnings over the last three to five years—multiplied by the number of years of service. Public pensions typically pay benefits equal to around 2 percent of final earnings per year of employment, although these replacement factors can differ from place to place, in particular based on whether the employee also participates in Social Security. Some public-sector employees take part in Social Security while others are not covered by Social Security and receive their principal retirement income from their employer's program.⁵

Unlike private-sector pensions, which are managed under the federal Employee Retirement Income Security Act, state and local government pensions do not in general fall under federal law. For that reason, protections afforded to accrued pension benefits vary state to state. In practice, however, once earned, pension benefits are in general considered to be very safe. Most states grant accrued pension benefits legal protections under contract law or state constitutions.⁶

In a number of states, not only are accrued benefits protected but so is the right to accrue future benefits; as a result, the current terms on which benefits are accrued also may not be altered. In a 2012 case from Arizona, even the government's ability to increase employee contributions was restricted based on the idea that the full terms of the pension plan in place at the time an employee was hired may not be altered in future years. More broadly, the often-substantial political power of public-sector employees—along with basic precepts of fairness—generally precludes reducing or eliminating their benefits after they have earned them. In this paper, when we refer to pension liabilities, we are for the most part referring to benefits already accrued rather than the right to accrue future benefits.

Public-sector pensions are financed through a combination of employee and employer contributions and investment earnings. Contribution rates vary from program to program, as do the criteria by which rates may be altered. In some cases, contribution rates are set by law, while in other cases contributions are automatically adjusted based on regular actuarial valuations of plan financing. The average employee contribution rate as of 2009 was 6.4 percent of wages, according to the Public Plans Database, although contributions vary significantly from plan to plan.⁷ In addition, some public employees have their formal contributions “picked up” by their employers, a fact not captured in the Public Plans Database. One aspect of the controversial reforms passed in Wisconsin in 2011 was a prohibition on government employers picking up employee pension contributions.

According to the most recent asset data available via the Public Fund Survey, 52 percent of pension funds are allocated to equities, 26 percent to fixed investments such as bonds, 6 percent to real estate, 11 percent to “alternative” investments such as hedge funds and private equity, and 3 percent to absolute/real return investments.⁸ The share of assets held in equities has risen significantly from the early 1980s, when only around one-third of pension portfolios were in stocks. A more recent trend is the shift to alternative investments, which promise higher returns than equities but with greater risk. The increase in the share of risky investments held by pension funds highlights the need for an improvement in accounting rules governing public-sector pensions.

PLAN VALUATION UNDER GASB ACCOUNTING RULES

Pension plans measure their financial health by comparing the value of their assets to that of their liabilities. The difference between assets and liabilities is referred to as the unfunded liability, while the ratio of the two is referred to as the funding ratio. Public sector pensions perform these calculations using guidelines issued by the GASB. Although GASB rules are not legally binding, a government must disclose if its calculations do not follow them.

Because a benefit payment made in the future is less valuable and less costly to finance than a payment made today, it is necessary to discount future benefit payments to the present to make them comparable to the value of plan assets. Discounting is a process similar to compound interest. While compound interest begins with a current dollar amount and adds interest to determine the future value, discounting begins with the future value and subtracts interest each year until a present value is arrived at.

The present value of a plan's liabilities depends on the interest rate at which the liabilities are discounted. Under standard actuarial accounting as outlined by GASB, a public pension plan discounts its liabilities using the return it expects the portfolio of assets it holds to generate. The average expected return on assets used in such valuations is currently slightly below 8 percent. The discounted value of plan liabilities is then compared to the value of assets to calculate the plan's funding ratio (assets divided by liabilities) and its unfunded liability (assets minus liabilities).⁹

Present values of plan liabilities also are used to calculate the plan's annual required contribution. This amount reflects the contribution the plan sponsor would need to make in a given year to both fund benefits accruing to employees in that year and to gradually pay off any unfunded liabilities from prior years.

The effects of changes in the discount rate can be dramatic. Under GASB accounting, the same plan, with the same assets and future benefit payments, could reduce its measured liabilities by nearly one-fifth by shifting from a portfolio with an expected return of 7 percent to one with an expected return of 8.5 percent.

When public pension liabilities are discounted at an 8 percent average interest rate, plans were around 77 percent funded, on average, as of mid-2010, and unfunded liabilities were equal to about three-quarters of \$1 trillion. And the situation could appear significantly worse if a different interest rate had been used in the calculation.

THE FAIR-MARKET VALUATION CRITIQUE

GASB pension accounting rules are essentially designed as a guide to funding. If a pension plan accurately predicts the average return on its portfolio over very long periods, and if the annual ups and downs of investment returns even out over the shorter periods of several decades in which pension funding is most relevant, then a plan that makes its annual required contributions calculations using the average return on the plan's portfolio will be able to precisely meet its benefit obligations over time.¹⁰

However, GASB accounting standards have been challenged by a movement among economists that is sometimes referred to as fair-market valuation. This approach argues that liabilities for public-sector pension plans should be valued in ways consistent

with economic theory and with how pension liabilities and other liabilities are valued in private financial markets.

The problem, economists point out, is that while there are several “ifs” in the description of how GASB accounting rules guide pension funding, there are few if any “ifs” with regard to the payment of benefits. Public employee pension benefits are intended to be guaranteed, they are described to employees as being guaranteed, and in most states they carry strong legal protections. Pension liabilities calculated using rules that include so many “ifs” will not be accurate or helpful to policymakers when the pension plan in the sponsoring government has no discretion with regard to paying benefits.

Put another way, GASB accounting rules help pension plans calculate a “best guess” annual contribution that gives the plan a roughly 50 percent probability of being able to meet its benefit obligations. Assuming the plan has correctly pegged the ultimate return on its investments, over any given period of time—even decades—the annual fluctuations in returns mean that the return actually received is almost sure to be either above or below the projection. Thus, even a “fully funded” plan has only a roughly 50–50 chance of generating returns sufficient to pay full benefits. But legally guaranteed pension benefits must be paid with 100 percent certainty. Thus, under GASB accounting rules, there is a mismatch between the plan’s legal requirement to pay benefits and its probability of being able to do so.

Economists argue that the discount rate used to value future pension liabilities should reflect the fact that pension benefits are guaranteed, even if the returns on a pension’s investments are not. More formally, the discount rate applied to the liability should be based on the risk of the liability, not the risk of any assets used to fund that liability. While there is some disagreement regarding how risky accrued public pension benefits actually are, economists are united in believing that the appropriate discount rate is a function of the liability rather than of the assets.

Moreover, this is how financial markets value liabilities. If a pension sought to transfer its liabilities to a private insurance company—something common in the United Kingdom, though for tax reasons not in the United States—the insurer would base its price on the size and risk of the liabilities, without reference to the risk or expected return of the asset backing those liabilities. The reason is that the investment portfolio can be changed at any time to any combination of risk and return the pension chooses, but the liabilities, if guaranteed, must be paid regardless.

Some academic research has concluded that accrued public pension benefits are actually less risky than explicit debt issued by state and local governments. In a 2009 paper, Jeffrey Brown and David Wilcox argue that, in the cases of New York City in the 1970s and Orange County, California, in the 1990s, pension

benefits continued to be paid even as those governments essentially defaulted on their explicit debts.¹¹ More recently, the city of San Bernardino, California, delayed repayment on so called “pension obligation bonds,” which are issued to the public as a means of financing pension benefits, but continues to pay benefits to retired public employees.¹² Likewise, Stockton, California, is currently in bankruptcy proceedings and plans to impose a repayment penalty on municipal bondholders while maintaining full benefit payments to public retirees.

In some cases, however, governments have sought to reduce accrued benefits via reductions in annual cost-of-living adjustments or other postretirement benefit increases. In nearly all cases, however, public employees have sued over such reductions, and most of these cases await resolution. If such reductions prove to be legally permissible, then a somewhat higher discount rate might be warranted. For instance, the bond rating agency Moody’s has recently announced that it will value public pension liabilities using the yield on high-quality corporate bonds, a similar standard to that applied to private-sector pension benefits.

Regardless of how the legal suits over cost-of-living reductions are resolved, public-sector pension benefits indisputably carry much less risk than the assets used to finance these benefits. Thus, while economists have not reached an agreement on what the right discount rate is to apply to public pension liabilities, they are united in believing that the 8 percent assumed return on pension assets is the wrong discount rate.

THE MYTH OF TIME DIVERSIFICATION

The most common argument made by investment managers for using a high discount rate is that the long time horizon of pension funds makes it possible to take on more stock market exposure with little additional risk. Indeed, a widely held belief is that time itself can help diversify risk. For example, as American Funds, one of the largest mutual fund providers, puts it: “In general, the more time you have to invest, the more risk you can afford to take.”¹³ Or, as even the US Securities and Exchange Commission states: “An investor with a longer time horizon may feel more comfortable taking on a riskier, or more volatile, investment because he or she can wait out slow economic cycles and the inevitable ups and downs of our markets.”¹⁴ Economists widely reject this view, however.

... while there are several “ifs” in ... how GASB accounting rules guide pension funding, there are few ... “ifs” with regard to the payment of benefits.

Economists argue that the discount rate used to value future pension liabilities should reflect the fact that pension benefits are guaranteed. ...

For example, as Nobel Laureate Paul Samuelson elegantly argues:

Invest for the long term, the theory goes, and the risk lessens. Is the dogma true as told? Alas, no. . . . Most of the time the buy-and-hold common stock investors do beat their more cautious neighbors; and, as the time horizon N becomes larger, the odds do grow that the bold holders of stock will win the duel. But it is also true that a longer time horizon brings bigger losses when an inevitable loss does occur. . . . Ask yourself: Will stepping down toward a poverty level, when that rarely but inevitably does happen, outweigh for me the pleasures that occur in those likely outcomes when my equity nest egg does increase?¹⁵

Or, as Zvi Bodie and Rachele Taqqu recently put it: “If you hear the catchphrase that an investment portfolio has a 90 percent chance of getting to your goal, remember to flip the statement on its head. Ask how much you stand to lose. Then weigh the consequences.”¹⁶

One impetus for the belief in time diversification comes from the belief that younger people should hold relatively more stocks and move toward bonds as they age. Although it is true that economic theory suggests that younger people should generally invest more in risky assets relative to older people, the actual reason has nothing to do with time horizons: it is because younger people actually hold a large asset that is not necessarily highly correlated with stock returns, namely their human capital, which is equal to the present value of their future wages. In essence, holding more stocks in their accumulated assets at an early age allows them to maintain a fairly constant amount of risk across their entire net worth, which includes the present value of their human capital. As they age, more of their human capital is converted into accumulated assets, so they should shift away from stocks in their accumulated assets to maintain a constant amount of risk across their entire net worth.¹⁷ The actual amount of risk depends on the importance of investors’ goals and their tolerance if they do not achieve them.

But this argument has nothing to do with time. Moreover, it does not apply to a pension system, which does not face a depreciating human capital asset like an individual saver. Instead, for time diversification to work, stock prices must manifest some prop-

erty like mean reversion, where low prices are predictably followed by high prices and vice versa. But that could be true only if market prices were inefficient. Although some potential evidence exists of mean reversion over short trading periods, most analyses fail to reject standard random models consistent with efficient pricing.¹⁸ If mean reversion were persistent enough to justify the “time diversification” hypothesis, there would be considerable market pressures to eliminate it, as it effectively leaves a lot of money on table.

A second potential impetus behind the time diversification hypothesis comes from focusing on random probabilities alone while ignoring the true value of the risk. Specifically, stakeholders—be they investors or the taxpayers who must make up for pension losses—place more weight on dollars received when they are poor than when they are rich. For example, suppose someone gave you \$1,000. Ask yourself: would you be happier to receive that additional \$1,000 if you had no money previously or if you had \$1 million? Of course, the vast majority of people would choose the former, because an additional dollar is least valuable when you already have a lot of money.

Economists more formally express this real-world preference with the idea of “diminishing marginal utility.” Specifically, marginal utility captures the additional happiness from receiving one extra dollar. That should be quite high when you have no money to start with. But as you make more money, the additional utility of receiving an extra dollar, while still positive, should decrease.

Now ask yourself, suppose that you have \$1 million today in assets, but there is a 5 percent chance that you could lose it all, maybe in the market or because of a personal liability. Suppose that you could transfer some of that \$1 million to a time in the future where you otherwise would have no money. How many dollars would you be willing to forfeit from your \$1 million today to ensure that you have at least \$1 after a loss that otherwise would leave you penniless? Empirically, most people would be willing to give up over \$50 today to ensure that they have \$1 available after the loss. Of course, due to diminishing marginal utility, that “intrapersonal exchange rate” declines with subsequent dollars transferred between these two situations: a person might be willing to pay only \$49.95 today to ensure that a second dollar is available after the loss.

Hence, when we think about risk, we should not think about probabilities alone: we ought to incorporate the value of risk aversion. Indeed, if marginal utility did not increase as people lose money, then the entire discussion about whether pension funds should invest in stocks would actually be irrelevant. A market full of risk-neutral people would buy equities whenever the expected return exceeded the safe bond yield, thereby eliminating any potential gain from investing in stocks. In other words,



LIVING to 100

Jan. 4–6, 2017
Orlando, Florida

SOCIETY OF ACTUARIES
INTERNATIONAL SYMPOSIUM

Save the Date

Registration for the 2017 Living to 100 Symposium will open soon. This prestigious event on longevity brings together a diverse range of professionals, scientists and academics to discuss:

- How and why we age;
- Methodologies for estimating future rates of survival;
- Implications for society, institutions and individuals;
- Changes needed to support an aging population increasing in size;
- Applications of existing longevity theories and methods for actuarial practice.



the so-called “equity premium” is actually just a “risk premium,” which compensates people for risk.

To see the importance of risk aversion, let’s consider a slightly more detailed example. Suppose that you have \$100,000. You have two options for what to do with that money. First, you can simply hold it, not investing it and not taking any risk.¹⁹ Or you could try to increase your wealth by rolling a standard six-sided die where the values of 1, 2, 3, 4, 5 and 6 correspond to the rates of return.

To incorporate the potential for both gains and losses, you earn a percentage rate of return on your wealth in a given year equal to whatever the die produces minus three, all times 10 percent. So,

for example, if you roll a 4 then you get $(4 - 3) \times 10\% = 10\%$, but if you roll a 2, you earn $(2 - 3) \times 10\% = -10\%$. The potential outcomes are shown in Table 1.

Since the expected average die roll is 3.5, the expected return from the outcomes shown in Table 1 is 5 percent, larger than the 0 percent return from taking no risk. In other words, the “risk premium” is 5 percent per year. The standard deviation, a measure of investment risk, is about 17.1 percentage points. Altogether, therefore, the risk shown in Table 1 is similar to stock investing.²⁰

Naturally, you are quite happy if you can double your money to \$200,000. But you will also be quite miserable if you lost half of your wealth and ended up with only \$50,000. If you are like most risk-averse people, the utility you gain from doubling your money will be much smaller than the utility that you lose from losing half of your wealth. Doubling your money might allow you to travel the world; losing half of it requires you to cut back on your food budget, keep your house colder during the winter, or even declare bankruptcy.

We can think of \$200,000 as your “bliss point” and \$50,000 as your “ruin point.” In the case of a pension fund, obtaining its bliss point will allow it to pay a larger level of benefits than it otherwise could afford; hitting ruin might require sudden benefit cuts or tax increases, often during a time when the economy is also doing poorly.

Given these facts, should you roll the die? At first glance, you might say that it all depends on your time horizon. Since the roll of the die produces an expected additional 5 percent return each year, then the “law of large numbers” should allow you to buffer the low returns with the higher expected returns. But that reasoning is incorrect once we properly incorporate the importance of risk aversion.

Table 2 shows the results from 5,000 simulations over a 100-year period from rolling the die. The third column shows that the chance of obtaining the \$200,000 bliss point without first falling to the \$50,000 ruin (bankruptcy) point does indeed increase across longer holding periods. For example, over the next decade, there is a 26.3 percent chance of hitting the bliss point. For an investing period over the next 20 years, this probability increases to 51.6 percent. Over enough time, the probability increases and converges to a value around 84 percent.

However, notice that the probability of hitting the ruin point also increases over time, since there is also more chance for failure. There is a 4.4 percent chance of ruin over the next decade, but that increases to around 13 percent at longer time horizons.

Although the probability of success at longer time horizons is much higher than the probability of ruin, a person who is suffi-

Table 1

Illustrative Outcomes From Roll of a Single Die	
Die Outcome	Return in Year
1	-20%
2	-10%
3	0%
4	+10%
5	+20%
6	+30%

Table 2

Potential Outcomes from Repeated Rolls of a Single Die		
Time Period (Number of Investing Years)	Chance of Obtaining Ruin over Period (%)	Chance of Obtaining Bliss without First Ruining (%)
10	4.4	26.3
20	8.7	51.6
30	10.5	64.2
40	11.5	72.0
50	12.3	76.7
60	12.6	79.7
70	12.8	82.2
80	12.9	83.3
90	13.0	84.0
100	13.1	84.4

ciently risk averse might nonetheless prefer to not take this gamble, instead keeping her original \$100,000 safe. She places more utility weight on a dollar at the ruin point than on a dollar at the bliss point. In other words, stocks do not become safer over longer holding periods. The potential for bliss and ruin both grow.

Moreover, because stock prices are correlated with the rest of the economy, success and failure in stock investments will not happen in isolation from what is happening in terms of economic growth, tax revenues, unemployment, and other factors. As Washington State's actuary has written with regard to its own pension plans' experiences: "Weak economic environments were correlated with weak investment returns. Lower investment returns created the need for increased contributions at a time when employers and members could least afford them."²¹ Likewise, good times for the pension fund's investments will be correlated with good times in the economy, when everyone else is flush and the value of an additional dollar is low.

ILLUSTRATING CONTINGENT PENSION LIABILITIES

For these reasons, risk—or the lack of risk—should be factored into the pension valuation process. Exactly how much value investors place on risk is already reflected in market prices. For example, the average return to equities above bonds, known as the "equity premium," reflects the expected increased compensation the investors demand if they are to accept the additional risk of stocks. If investors did not care about risk, then the equity premium would disappear. Properly accounting for risk, therefore, requires using fair-market valuation that incorporates market prices.

How exactly does fair-market valuation help policymakers measure whether they are properly accounting for risk in pensions? An example helps illustrate.

Consider a pension that owes a guaranteed lump sum payment of \$1 million in 15 years' time. Under GASB accounting rules, if the plan invests \$301,194 today—the present value of \$1 million discounted at an 8 percent interest rate²²—it can call itself fully funded. This investment path is illustrated using the blue line in Figure 1.

But according to fair-market valuation, if the pension's payment is indeed guaranteed, it should be discounted at a riskless interest rate. If the riskless return is 4 percent, the true value of the liability is \$548,812, almost twice as much upfront as is required under the actuarial approach using the 8 percent rate. This is represented as the gray line in Figure 1. This illustration should demonstrate why most pension interests—governments, public employees, plan managers, and so on—prefer the actuarial approach.

If the pension's assets have an expected return of 8 percent, then investing \$301,194 today will deliver an expected payoff of \$1 million in 15 years. The problem is that assets with an expected return of 8 percent cannot produce such a return with certainty,

meaning that the portfolio's value after 15 years will almost surely end up being higher or lower than the desired \$1 million. In other words, rather than a single blue line in Figure 1 representing investment in risky assets, a better representation is a blue area showing a range of possible outcomes, roughly half of which exceed the \$1 million goal, with the remaining half falling short.

No matter how well a pension plan manages its investments, it cannot generate 8 percent returns with certainty. The actual return the plan receives is based on the luck of the draw. Given that the benefits must be paid 100 percent of the time, a plan that has only a 50 percent chance of being able to meet its obligations is not "fully funded" in the way that most laymen would interpret the term.

In reality, a plan seeks to neither overshoot nor undershoot its goal. If the plan's investments exceed their projected return, that means the initial contribution could have been smaller. In other words, the plan "overcharged" current taxpayers. Alternately, if the investments come up short of their goal, the plan will not be able to pay what it owes and must turn to the taxpayers for a bailout. In that case, the plan will overcharge future taxpayers to compensate past public employees who did not work for them. In either case, the principle of intergenerational equity will have been violated.

However, financial products called "options" provide a solution. A "call option" allows the pension plan to sell off any surplus if the plan's investment turns out to be worth more than \$1 million. A plan that sells a call option can use the proceeds to offset the cost of the initial investment, thereby eliminating the costs of overshooting the pension's goal and protecting today's taxpayer against being overcharged.

Likewise, a "put option" can be purchased to top up the difference between the assets' actual value and \$1 million if the investment comes up short. The put option protects against outcomes in which the plan's investments fall short and so protects tomorrow's taxpayer. Barring some truly catastrophic collapse of financial markets, the plan will always be able to pay exactly the promised \$1 million, with no wasted money, if it first invests \$301,194 in stocks or other high-returning assets and then sells a call option to dispose of any surplus and purchases a put option to cover any shortfall.

This means that the cost of truly fully funding the \$1 million future liability—meaning funding it so that it is guaranteed to be paid without recourse to a taxpayer bailout and without any wasted surplus—is the \$301,194 initial investment minus the proceeds from selling the call option, which total \$11,436, plus the cost of purchasing the put option, which is \$259,053. The net cost is \$548,812, precisely the same as if the liability had been discounted and funded using the 4 percent riskless rate of return.²³

The difference between the liability's value when discounted at the 8 percent rate used under current GASB standards and when discounted at a riskless interest rate represents the value of the contingent liabilities that have been placed on future taxpayers based on funding decisions made today. This cost is not a worst-case scenario, as some seem to believe. Rather, it represents the price that future taxpayers would willingly pay to rid themselves of the risk of being called on to make good on promises that were made by, and should have been paid for by, today's taxpayers.

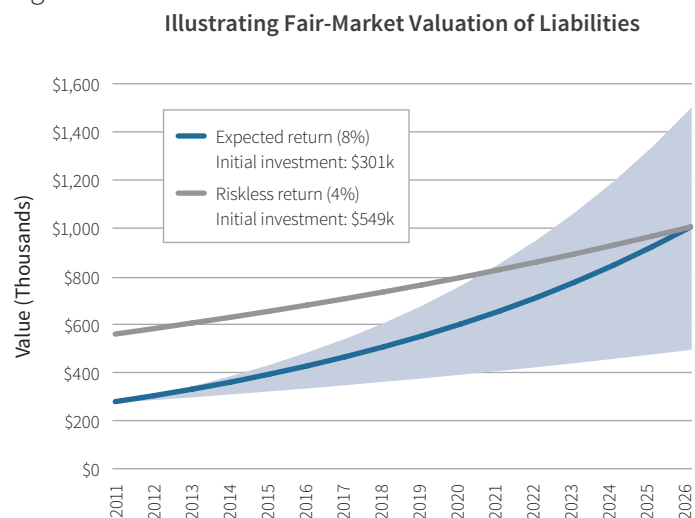
A more conservative pension might invest larger amounts in safer assets, increasing costs for current taxpayers but leaving smaller contingent liabilities on future generations. Alternately, a more aggressive plan might make smaller upfront contributions but invest them in riskier assets. This reduces costs today but generates a matching increase in the value of the contingent liability on future taxpayers. In either case, though, the total cost of the liability will be the same regardless of how the plan chooses to invest.²⁴

This example of a pension plan required to fund a \$1 million future liability also illustrates that fair-market valuation is not an academic exercise with no relevance to the actual investments public pensions make. When economists say that a pension should apply a discount rate of, say, 4 percent to its liabilities, they are not assuming that the plan invests in safe assets that yield a 4 percent return. The cost of the put and call options is determined in the market and based on the riskless return available in the market and the risk of the investments the plan holds. In other words, it makes sense to discount riskless pension liabilities using a riskless interest rate even when the plan itself invests in riskier and higher-returning assets. The fair-market valuation approach captures both the full range of possible investment outcomes and the welfare values that real-world individuals place on those outcomes. It is current actuarial standards, in which a risky investment portfolio is assumed to earn a constant rate of return, that differ from reality.

Finally, this example again illustrates that the value of the liability does not shrink because public pensions invest over long time horizons, as pension and public pension actuaries sometimes claim. If "time diversification" made stocks less risky over long periods, then put options protecting an investor against

poor market returns would become cheaper as the exercise date on the option was pushed out. In fact, the opposite is the case: long-dated options are more expensive than short-dated options. This reflects that fact that while the variation in average stock returns becomes smaller over long periods, the variation in the actual dollar amounts that investors receive grows larger. As the blue shaded area in Figure 1 shows, the longer a stock is held, the wider the possible range of outcomes the investor may face.

Figure 1.



Source: Authors' calculations.

CONCLUSION

The fair-market value of a pension's liabilities represents the answer to the question posed at the outset of this paper: what is the cost of fully funding future benefits such that there is risk neither to beneficiaries nor to future taxpayers? Lower costs are possible if elected officials wish to acknowledge risk to either of these parties—that is, if they allow that full funding includes the possibilities that benefits may be cut or that the plan may return to future taxpayers for a bailout when planned funding falls short. But in such a scenario, any pension plan may be considered fully funded, rendering the term meaningless for policymakers and the general public.

A better approach takes full funding to mean that accrued benefits can be paid without extracting additional resources from the taxpayer. The fair-market valuation method tells us the costs of achieving full funding. And it incorporates a number of truths from economics and finance: that stocks pay higher returns because they are more risky; that the risk of stock investments does not disappear over long holding periods; and that individuals, either as investors or taxpayers, do not value gains equally with losses. Finally, valuing guaranteed pension liabilities using a riskless discount rate does not imply that pensions must invest only in bonds. Rather, it merely shows that the value of a guaranteed

A better approach takes full funding to mean that accrued benefits can be paid without extracting additional resources from the taxpayer. ...

pension benefit is independent of the returns on risky investments used to fund that benefit.

Public pension accounting is undergoing changes as the GASB looks to revise its rules through the recently introduced Statements 67 and 68, even if these do not alter the basic logic—or illogic—of how public pensions value their liabilities. Moreover, public plans themselves are being reformed in a number of states, principally through higher employee contribution rates and lower benefits for newly hired employees. But pension financing will not truly be stabilized until plans first adopt better standards for determining how much they truly owe. ■

Andrew G. Biggs is a resident scholar at the American Enterprise Institute (AEI).

Kent A. Smetters is the Boettner Professor of Business Economics and Public Policy at the Wharton School, University of Pennsylvania, and an adjunct scholar at AEI.

ENDNOTES

- ¹ Governmental Accounting Standards Board, “The User’s Perspective: Interperiod Equity and What It Means to You,” June 2009, www.gasb.org/cs/ContentServer?c=GASBContent_C&pagename=GASB%2FGASBContent_C%2FUsersArticlePage&cid=1176156731381.
- ² On the latter point, see Robert Novy-Marx and Joshua D. Rauh, “The Revenue Demands of Public Employee Pension Promises,” *American Economic Journal: Economic Policy* (forthcoming).
- ³ Defined-benefit plans base retirement benefits on a formula deriving from the employee’s earnings and years of service; any investment risk is borne by the plan sponsor.
- ⁴ Alicia H. Munnell and Mauricio Soto, “State and Local Pensions Are Different from Private Pensions,” Center for Retirement Research, November 2007.
- ⁵ For individuals spending part of their careers in public employment not covered by Social Security and part under Social Security–covered employment, the Government Pension Offset and Windfall Elimination Provision may affect the Social Security benefits they or their spouses are eligible to receive.
- ⁶ Amy Monahan, “Public Pension Plan Reform: The Legal Framework,” *Education, Finance & Policy* 5 (2010).
- ⁷ See the Public Plans Database at <http://pubplans.bc.edu/pls/apex/f?p=1988:3:0>.
- ⁸ Survey data are available at www.publicfundsurvey.org.
- ⁹ Pensions currently use a value known as actuarial assets in calculating unfunded liabilities and funding ratios. Actuarial assets differ from the market value of assets in that rates of return applied to the assets are smoothed, generally over a five year period. In revisions to accounting rules published in June 2012, the GASB announced that plans must use the market value of assets for such calculations. Doing so would, as of the present at least, reduce measured funding ratios by around 10 percentage points.
- ¹⁰ This is true only if the plan discounts its liabilities using the geometric mean, or compound return, on its investments. However, many plans base their discount rates on their investments’ arithmetic mean return, which for risky assets is always higher than the geometric mean. Although it is a distinct point from the financial economics critique of pension discounting, this error means that even a plan that achieves its projected investment returns may be unable to meet its accrued obligations in full.
- ¹¹ Jeffrey R. Brown and David W. Wilcox, “Discounting State and Local Pension Liabilities,” *American Economic Review* (May 2009).
- ¹² Arleen Jacobius, “San Bernardino Votes to Delay City’s Pension Bond Payments,” *Pensions & Investments*, July 25, 2012.
- ¹³ American Funds, “Develop an Investment Strategy,” <http://americanfundsretirement.retire.americanfunds.com/planning/strategy/what-is-asset-allocation/time-horizon.htm>.
- ¹⁴ US Securities and Exchange Commission, “Beginners’ Guide to Asset Allocation, Diversification, and Rebalancing,” www.sec.gov/investor/pubs/assetallocation.htm.
- ¹⁵ Paul Samuelson, “Dogma of the Day,” *Bloomberg Personal Finance* (January/February 1997).
- ¹⁶ Zvi Bodie and Rachele Taqqu, *Risk Less and Prosper: Your Guide to Safer Investing* (Hoboken, NJ: Wiley, 2012), XVII. Emphasis in original.
- ¹⁷ See Zvi Bodie, Robert C. Merton, and William F. Samuelson, “Labor Supply Flexibility and Portfolio Choice in a Life-Cycle Model,” *Journal of Economic Dynamics and Control* 16 (1992): 427–49.
- ¹⁸ The strongest case for mean reversion was provided by James M. Poterba and Lawrence H. Summers, “Mean Reversion in Stock Prices: Evidence and Implications,” *Journal of Financial Economics* 22, no. 1 (1988): 27–59. Subsequent papers found little evidence, especially when an outlier period was removed. See, for example, Myung Jig Kim, Charles R. Nelson, and Richard Startz, “Mean Reversion in Stock Prices? A Reappraisal of the Empirical Evidence,” *Review of Economic Studies* 58, no. 195 (1991): 515–28.
- ¹⁹ In reality, of course, you could earn a risk-free short-term bond return that covers inflation and some small real yield merely to compensate you for being patient.
- ²⁰ If anything, it is a bit more favorable than stock investing since the standard deviation on US stocks is closer to 20 percent and also includes some fatter negative tails (negative skew).
- ²¹ Office of the State Actuary, “Washington State 2009 Actuarial Valuation Report,” October 2010; and Office of the State Actuary, “2010 Risk Assessment: Moving Beyond Expectations,” August 31, 2010.
- ²² Throughout the example, present values are calculated using continuous discounting for consistency with options pricing; differences from periodic discounting are small. The present value equals the size of the future payment divided by the exponential of $(r \cdot n)$, where r is the annual discount rate and n is the number of years until the future payment will be made.
- ²³ The listed numbers contain a one-dollar discrepancy, reflecting rounding error.
- ²⁴ This result is based on a principal known as “put-call parity.” See H. R. Stoll, “The Relationship between Put and Call Option Prices,” *Journal of Finance* 24 (December 1969): 801–24.