Comments and Discussion

COMMENT BY

DEBORAH LUCAS As much of my discussion is critical of the approach taken in this paper, I want to start on a positive note and sincerely applaud the authors for their contributions to the discussion of pension fund sustainability. Their careful projections of the cash flows associated with future state and local pension liabilities are particularly valuable. It's said that imitation is the sincerest form of flattery, and so I'll note that we have a project underway at the MIT Golub Center for Finance and Policy to project the probability distribution of future cash flows and funding levels for a larger universe of public plans. I can attest to how much work it takes to create each new data point.

I also want to highlight a few of the authors' conclusions where I am in agreement. The first is that there is unlikely to be a widespread crisis caused by state and local pension plans being unable to meet their liabilities, at least absent a major meltdown in the stock market. Even if some plans run out of funds much sooner than suggested by the analysis in this paper (and I believe some will), exhaustion dates are likely to be staggered over time. Hence, it is incorrect to view pervasive underfunding as a harbinger of an impending national crisis. Second, I concur with the lack of an imperative to quickly or fully close funding gaps, although as explained below, the theoretical and practical considerations that lead me to those conclusions are distinctly different from the ones suggested in this paper.

However, the authors go much further in challenging the conventional wisdom. They argue that not only is there no looming nationwide crisis but that there is no reason to be concerned about the vast majority of individual plans. They further provocatively suggest that it might well be desirable

to divert funds earmarked for pension funding to other uses and that such policies are likely to be sustainable.

Those sanguine conclusions are the result of several questionable modeling choices that include a trip down the dangerous rabbit hole of $r - g$ accounting; abstracting from uncertainty; and treating the information in financial market prices as irrelevant to public finance. I believe that any analysis that avoided those shortcomings would conclude that there is a strong link between funding levels and sustainability. The remainder of my comments elaborate on the following observations.

The authors' calculations rest on accounting identities that critically depend on the assumed return on assets, r, and the growth rate of GDP, g. The authors take these variables to be deterministic and only loosely justify the assumed relation between them. In fact, the restriction that r must be greater than g is a robust implication of any standard, deterministic, general equilibrium model. When growth rates and asset returns are stochastic, a welfare analysis based on comparing average r and average g has no theoretical foundation. Meaningful evaluation of fiscal policy alternatives requires incorporating the effects of uncertainty and its associated costs.

The assumptions of perpetually low interest rates and asset returns, which are necessary to conclude that waiting to fill funding gaps makes little difference, rest on shaky empirical grounds.

Pension plan sustainability is equated in this analysis to projecting a finite long-run debt-to-GDP ratio in a deterministic world. A more conventional view of sustainability is that it requires a plan to be able to withstand the risk of adverse shocks such as lower than expected asset returns or unanticipated population loss over some extended period of time.

Under a definition of sustainability that focuses on the likelihood of asset exhaustion over a specified horizon, increasing funding always improves sustainability because it increases a pension plan's distance to default. A larger number of plans would be classified as unsustainable under this definition.

The authors suggest that increases in pension fund contributions have reduced public investment. It is important to recognize that there is no economic trade-off between the two uses of funds. If there is crowding out, it is a political phenomenon, and the potentially high costs of relaxing pension funding requirements suggest it is probably better addressed by other means.

THEORETICAL FOUNDATIONS In recent years, the (short-term, risk-free) interest rate, r, has been lower than the GDP growth rate, g. This has generated optimism among some commentators that accumulating high levels

of government debt is sustainable. If relatively low interest rates were to persist indefinitely, GDP growth would exceed debt service cost growth, and tolerating high levels of debt-to-GDP would place a small or even shrinking debt service burden on future generations. Put differently, if g is permanently greater than r, then it is better to borrow and invest in real growth rather than to pay down government debt; there is essentially an arbitrage opportunity.

If policymakers are considering running a high-debt policy based on these observations, it is important to ask first whether it is reasonable to expect interest rates to remain below growth rates indefinitely. The prediction of a standard, deterministic, equilibrium model suggests the answer is no.

The logic that r will exceed g in a deterministic equilibrium can be illustrated very simply in a two-period setting that can be shown to generalize to an infinite horizon. Assume that aggregate income grows at a constant rate g. Agents maximize utility through the choice of how much first-period income to save, S, subject to the usual lifetime budget and wealth constraints:

$$\max_{S} U(C_1) + \beta U(C_2),$$

such that $C_1 = Y_1 - S$ and $C_2 = Y_2 + S(1 + r)$. Differentiating with respect to S implies the optimality condition:

$$\frac{1}{1+r} = \beta \frac{U'(C_2)}{U'(C_1)} = \beta \frac{U'(C_1(1+g))}{U'(C_1)}.$$

The interest rate, which must satisfy the optimality condition in order to clear the market for borrowing and lending, depends on g. For any concave utility function, the market-clearing value of r increases in g. When $U(C) = ln(C)$, the equation simplifies to $(1 + r) = (1 + g)/\beta$. If $\beta < 1$, then $r > g$ (e.g., $\beta = .98$, $g = 1$ percent implies $r = 3.06$ percent). With log utility, $r - g$ is roughly the subjective rate of time preference.[1]

What if interest rates and growth rates are uncertain? A few years before Samuelson (1958) presented the accounting identity used in this analysis, Arrow and Debreu (1954) introduced the elegant and powerful concept of state prices—the recognition that a dollar in different future states of the world will have a different utility value relative to a dollar today,

1. For constant relative risk aversion (CRRA) preferences, the higher the coefficient of relative risk aversion, the stronger is the preference for consumption smoothing over time, and hence the larger is the positive difference between r and g.

depending among other things on the future strength of the economy: a claim on resources in bad times is worth more than the same claim in good times. That idea is at the core of modern finance. It explains why the expected return on risky assets like stocks, whose payoffs are positively correlated with aggregate consumption, exceeds the return on safe government debt. These insights about how risk influences value are as critical for assessing the welfare implications of fiscal policy choices as they are for understanding private sector asset returns.

We can now turn back to the relation between r and g in a similar setup to the one described above but incorporating uncertain growth. In that setting, an induced precautionary demand for savings may imply an average risk-free rate that is less than the average growth rate of consumption. However, as just explained, the value today of uncertain future output can no longer be inferred by discounting expected payoffs at the risk-free rate. Hence it is impossible to conclude on the basis of average g exceeding average r that it is desirable for the government to postpone paying debt or to issue more of it in order to invest in growth assets. To do so treats risk premiums as arbitrage opportunities rather than as market-determined compensation for bearing costly aggregate risk.

This relates directly to the authors' calculations, some of which assume a fixed return on assets that is higher than the risk-free rate that affects the future value of liabilities. In this case the analysis has a built-in arbitrage opportunity. Had the authors also taken into account that a pension plan (or the government sponsor) could borrow risk free to invest in plan assets, the accounting identity would show this to be a money machine that makes it costless to eliminate underfunding.

Perhaps anticipating this criticism, the authors also look at a "certainty-equivalent" case where investment assets are assumed to earn the risk-free rate. While the certainty-equivalent methodology can be a convenient tool for inferring the market value of a risky asset, using it here to extrapolate future outcomes is problematic. It is noteworthy that one of the most provocative claims of the analysis—that there isn't much gain to closing funding gaps sooner rather than later—rests on this assumption. That's because when asset returns are low and unfunded liabilities grow at the same low rate, acting sooner has little advantage. Viscusi (2007) points this out in the context of explaining the paradoxes that arise when the social discount rate is taken to be zero. However, if the goal is to forecast average outcomes, simple algebra implies that assuming typical plan portfolio choices, higher funding levels would reduce the average debt service burden on future generations.

If the reader is wondering what I am recommending to replace the authors' assumptions about interest rates and asset returns, it is to do a stochastic simulation that takes into account the higher expected return on assets and also the higher risk. This is discussed in the context of sustainability below.

The insights of state pricing are also critical for evaluating costs to future generations and whether they will be equitably distributed. A pension plan is more likely to run out of money when it is highly underfunded to begin with and during a sustained economic downturn. In order to cover a funding shortfall, the sponsoring government may be forced to cut other spending. Those spending cuts, which are especially costly because they occur when the economy is depressed, will be largely borne by the unlucky current residents of that locale. Hence, under realistic assumptions about government behavior in response to cash shortfalls, the costs of underfunding will fall disproportionately on generations that are least able to afford it. This is in sharp contrast to the assumption of complete smoothing of costs across generations through the mechanism of debt financing, and the reason why I think that sort of analysis cannot be used to draw conclusions about the welfare consequences or distributional impacts of underfunding.[2]

On the issue of fairness between generations, spreading underfunding evenly over future generations may not be perceived as equitable. In fact, a standard argument for fully prefunding accruing benefits is on equity grounds: the primary beneficiaries of current public services then bear the full cost. There is a stronger case for spreading legacy underfunding across multiple generations, but there are nuances involving relative wealth and incentive effects that may weigh against it.

MORE ON PARAMETER CHOICES As the calculations and conclusions of the paper are difficult to interpret because of the foundational issues just discussed, I only have a few comments on specific parameter choices.

Unlike most of the variables, which are justified with reference to historical data, the risk-free rate r is loosely based on the low-rate conditions of 2020. Historical data suggest that r, and also $r - g$, has been lower recently than it has been on average. It is not clear why the authors choose to deviate from the standard practice of equating fixed parameter values to historical

2. Citing an unpublished CBO paper, the authors also claim that "to the extent the risk premium reflects business cycle risk, the government can lower that risk by spreading it across future generations." Even if the government redistributes risk within the population so as to reduce the welfare costs of business cycles, the equilibrium equity premium reflects the cost of the risk that remains. That premium is relevant to assessing the cost of risk associated with the government's fiscal policies.

Figure 1. Federal Debt Held by the Public, 1900 to 2050

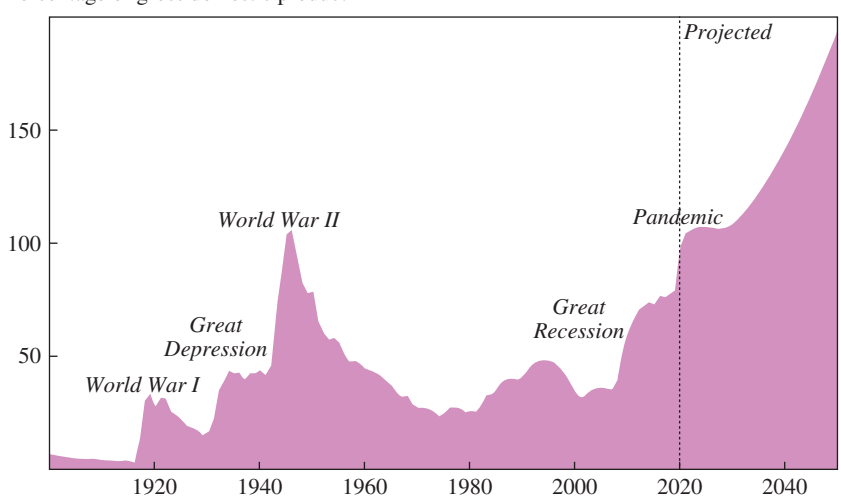

Source: CBO (2020).

averages when it comes to *r*, but that choice is important for the conclusion that there is little gain from closing the funding gap more quickly.

A minor quibble is that the authors justify some parameter choices on the basis that CBO makes similar assumptions. CBO's assumptions are not a natural reference point because it makes current law projections used in baseline exercises, not statistical forecasts. Over long horizons the difference between the two can be significant.

SUSTAINABILITY REVISITED Whether or not the ratio of long-run debt-to-GDP can be stabilized (i.e., it remains finite as time is taken to infinity) in a deterministic world seems largely orthogonal to whether public pension funds are sustainable for at least two reasons. First, a stable debt ratio does not imply a manageable debt service burden. Second and more fundamentally, sustainability requires resilience to adverse events, and that can only be assessed in a stochastic framework.

With regard to level effects, a stable but high debt-to-GDP ratio would probably be viewed as unsustainable, or at least very undesirable, when it entails high debt service costs paid for with distortionary taxes. In assessing the burden, the cost of servicing federal and other public debt also has to be factored in. As shown in figure 1, the federal debt-to-GDP ratio is projected to reach unprecedented levels, doubling over the next thirty years and showing no sign of stabilizing. It is hard to imagine that adding to

those projected debt ratios by deliberately increasing pension underfunding would be considered sustainable or wise.

Focusing on the level rather than the stability of the debt ratio also makes the risk of future interest rate hikes more salient to the question of sustainability. As the debt ratio rises, the debt becomes riskier and investors require a higher credit spread to buy it, pushing up the cost of debt service. In extreme cases, it will be impossible to roll over the debt and public services or pension benefits will have to be cut. Notice that the properly measured prospective social cost of such adverse outcomes exceeds the average increase in interest payments and services lost because the costs tend to be realized when the economy is weak. The policy of tolerating higher debt ratios transfers costly market risk to future generations. This is another example of where the policy-induced cost of market risk remains hidden when analysts think only in terms of average cash flows.

Long-run stability is not a sufficient condition for sustainability. An underfunded pension plan, despite being characterized as stable in the long run, may run out of available resources to pay promised benefits due. A pension plan with a high chance of becoming insolvent is not sustainable as the term is commonly understood. There's an analogy in the Social Security system. Whatever is projected about future taxes and benefits, if the trust fund hits zero, there is no budget authority to pay current benefits if they exceed current payroll tax revenues; the system is unsustainable without political action. This is true even though the trust fund is an accounting mechanism and not a repository of financial assets. In the same way, if a public pension plan's assets fall to zero and current contributions are insufficient to cover current benefits, it will trigger a crisis that will require legislative or executive action to address.

I would therefore propose an alternative definition of sustainability: a fund is sustainable if it can meet its contractual obligations with sufficiently high probability over a specified period of time. Operationalizing this definition would require choosing a threshold probability and time horizon, for example, a system could be classified as sustainable if there is a 95 percent probability that assets will not be exhausted at any point over the next two decades. This alternative definition recognizes that assessing sustainability requires modeling funding levels as stochastic. It also requires computing the time path of the distribution of funding outcomes, not just its average.

Under that definition, higher funding levels unambiguously improve sustainability, including in low-rate environments. While many sources of uncertainty affect funding levels, arguably the largest source of year-to-year volatility is a plan's risky asset holdings. Assuming that asset returns are

normally distributed (and not too highly correlated with other determinants of funding levels), the formula for the probability of first passage times for a normally distributed variable implies that the probability of assets going to zero over a specified horizon is monotonically decreasing in the level of assets. This is analogous to the concept of distance to default used in standard analyses of default risk for businesses. Closing a funding gap more quickly also always improves sustainability under this definition.

For a given funding level, incorporating more realistic assumptions about expected asset returns and volatility has partially offsetting effects on sustainability. A higher average return increases sustainability while volatility reduces it. The prior literature on this that the authors cite, and my own recent work (Lucas and Smith 2020) on whether a collective defined contribution system can deliver a fairly safe and adequate benefit without boosting contribution rates (we find it can't), suggests that taking stochastic returns into account would cause more plans to be classified as unsustainable, again under a definition based on the likelihood of asset exhaustion. The authors have the data to simulate the probability distribution of exhaustion dates for each of the plans in their database when asset returns are stochastic, and it would be very informative to see those results.

FURTHER OBSERVATIONS ON THE COSTS AND BENEFITS OF PREFUNDING As noted at the beginning of my remarks, I agree with the authors that there is no imperative to quickly or fully close funding gaps, but for distinctly different reasons than the ones they emphasize.

The authors suggest that a benefit of relaxing pension funding requirements is to avoid crowding out public investments such as in infrastructure and education. However, because most state and local governments have access to capital markets or bank financing, from a purely economic perspective there is no reason for the two goals to compete. Governments can borrow to turn unfunded pension liabilities into funded ones, or to pay for investments, or both. Total pension liabilities might rightly influence the perceived affordability of other spending plans, but fundamental affordability should be invariant to pension funding status. That reasoning applies the logic of Modigliani and Miller's (1958) famous capital structure irrelevance theorem to public pension plans, and its implications are explored in more detail in Lucas (2017). In practice, crowding out may occur, but if so, it is a political phenomenon that can be addressed in other ways. For example, a state legislature could weaken its self-imposed balanced budget requirement.

Relatedly, the main reason that I think a narrow focus on full funding is misplaced is the fact that funding gaps can be reduced using borrowed funds, and hence increased funding may have no impact on current or future

fiscal policies. The authors make the same point when they say, "More broadly, governments typically hold debt, and unfunded pension liabilities are simply a form of (implicit) debt." Although we agree on that fact, it doesn't alleviate my concerns about the consequences of high consolidated debt levels. Evaluating the sustainability of the public finances of a city or state requires also taking into account other current and projected liabilities, and the possible paths of future spending and revenues. This is essentially what municipal bond rating agencies do.

There are other legitimate reasons one might favor higher funding levels beyond concerns about the possibility of a funding crisis or eventual limits on the capacity to borrow. Those include imposing fiscal discipline, making opaque liabilities more transparent and salient to the public, discouraging governments from making unsustainable benefit promises, and ensuring that the cost of services falls on the current beneficiaries and not on future generations.

REFERENCES FOR THE LUCAS COMMENT

Arrow, K. J., and G. Debreu. 1954. "Existence of an Equilibrium for a Competitive Economy." *Econometrica* 22, no. 3: 265–90.

CBO (Congressional Budget Office). 2020. "The 2020 Long-Term Budget Outlook." Washington. https://www.cbo.gov/publication/56516.

Lucas, Deborah J. 2017. "Towards Fair Value Accounting for Public Pensions: The Case for Delinking Disclosure and Funding Requirements." Prepared for Harvard Kennedy School Conference, Gathering Storm: The Risks of State Pension Underfunding, October 19–20.

Lucas, Deborah J., and Daniel Smith. 2020. "How Much Can Collective Defined Contribution Plans Improve Risk-Sharing?" *Journal of Investment Management* 18, no. 3.

Modigliani, F., and M. H. Miller. 1958. "The Cost of Capital, Corporation Finance, and the Theory of Investment." *American Economic Review* 48, no. 3: 261–97.

Samuelson, Paul A. 1958. "An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money." *Journal of Political Economy* 66, no. 6: 467–82.

Viscusi, W. Kip. 2007. "Rational Discounting for Regulatory Analysis." *University of Chicago Law Review* 74, no. 1: 209–46.

COMMENT BY

JOSHUA RAUH The question of the sustainability of pension promises is a good one. It is well known that most state and local government pension promises are underfunded, but what kind of adjustments must

governments make to their contributions in order to prevent these unfunded liabilities from causing fiscal instability and debt crises?

If the accounting for pension promises were based on measuring liabilities at their actual financial value rather than at an artificially low value based on expected return, then one could very well entertain arguments that the optimal pension funding ratio is not 100 percent, just as the optimal amount of public debt is unlikely to be zero. The authors are correct to point out that the goal of moving from the currently underfunded status of pension promises to full prefunding is a goal that would be analogous to paying down government debt, as opposed to stabilizing it as a percentage of GDP.

There is an important caveat here, however, which is that when pension sponsors can measure unfunded liabilities using expected returns on plan assets, they may be incentivized to ignore risk in the setting of pension funding policy. While I agree in principle that achieving a stable ratio of properly measured government debt to GDP is a sufficient goal for financial stability, the narrower goal of targeting a stable value of unfunded pension liabilities as a share of GDP may not be sufficient for financial stability—especially when sponsors' measurement of those liabilities biases them toward taking risk in order to meet expected returns dictated by a political process. Even though the authors' analysis targets a debt ratio at a correctly measured discount rate, pension funds still must set an expected rate of return and an asset allocation that targets that rate of return.

That setup by itself may introduce instability. Indeed, partially funded or even fully unfunded PAYGO can be sustainable under the right conditions. However, the potential for instability arises in this context because of the possibility that a pension system would have to start paying benefits out of current resources, requiring suddenly much higher draws on current resources than under a 7.25 percent return assumption. The authors effectively eliminate this possible instability by assuming plans can borrow after exhausting their assets. It is the concern that municipal credit markets might view this risk as substantial that creates instability, since that would then lead to large increases in borrowing rates and challenges to the ability of municipalities to access credit markets.

One of the features of the analysis in this paper is that the discount rate and the expected return play separate roles. In the context of the simulation-based approach adopted by the authors later in the paper, this is potentially fruitful, as of course the discount rate for a fixed pension promise and the expected return on a risky portfolio of assets should be different. For deterministic analysis, however, it makes little sense to choose rates of return that deviate from the risk-free rate. Based on the analysis by Costrell and

McGee (2020) of the authors' model, the deterministic model seems to have the feature that for any given steady-state value of assets to payroll, the steady-state contribution rate is independent of the discount rate and depends only on the expected return. In fact, what seems to happen is that lower discount rates increase the normal cost but (holding the expected return fixed) allow the difference between expected return and risk-free rate to offset that additional cost.

Under the fundamental theorem of finance, the risk-adjusted expected return on assets is the risk-free rate. So if one has to pick only one point to represent the distribution, then this is the right one. Under this scenario, pension debt to nominal GDP would rise indefinitely without fiscal adjustments. For these reasons, the most appropriate single parameterization is one with $r = \delta = 0$ percent. The thirty-year TIPS yield is –0.15 percent as of June 14, 2021, so there is no reason to use something higher. A real yield of –0.15 percent is below the lowest real return assumption in table 2 of the paper, but it is close to the 0.5 percent parameterization, so for most of the goals examined in table 2, I think of the correct percentages as a few points higher than those shown in the 0.5 percent return case. Taking that 0.5 percent scenario at face value, the required contribution increase to get implicit debt back to today's levels is 14.7 percent of payroll if they start today. Relative to the current weighted contribution rate of 24 percent in the authors' estimation sample, the 14.7 percent of payroll hike amounts to an increase in contributions by over 60 percent. This is substantial.

Appropriately, the authors have introduced a stochastic analysis that allows them to consider the distribution of outcomes generated by the risk-loaded investment strategies of public pension funds. Here it does make sense to consider this distribution of possible outcomes. As is well known in finance, the likelihood of exceeding the risk-free rate of return is high under standard lognormal return distribution assumptions, and the distribution of outcomes shown by the authors reflects this.

There are a number of points that must be considered in interpreting the paper's simulation analysis. First, it must be recognized that the states of the world where bad return outcomes occur are high marginal utility states. This is precisely the reason why the stock market has an expected return that is higher than the risk-free rate. Under risk-neutral probabilities, which put more weight on bad states of the world, the expected return remains the risk-free rate.

Second, for their main stochastic simulation (figure 13 in the paper), the authors are assuming contribution rates based on a strategy in which contributions are set under a 2.5 percent real return assumption, while actual

returns on assets in the simulation average a 4.5 percent real return (e.g., 6.7 percent nominal and 2.2 percent inflation). The authors' caveat that if the plans set contributions based on a deterministic real rate of return of 4.5 percent, decreasing contributions by 3 percent instead of increasing them by 7 percent, "the outcomes are less sanguine." But of course, public pension plans are setting their contribution rates today based on their expected rates of return, not based on lower rates of return. The introduction of any strategy like targeting a debt level in thirty years equal to that of today would be unlikely to move public pension systems toward more conservative return assumptions in setting their contribution rates. Therefore, this less sanguine distribution of outcomes, shown in online appendix figure A7, seems a more appropriate illustration of the likely impact of these policies.

Third, the risk-taking necessary to target high returns exposes pension systems to significant volatility—not only a distribution of long-term return outcomes but also a range of possible paths that lead to those outcomes, including some that could lead to near-term insolvency. The issue is not only the possibility of ending the thirty-year time period with negative assets but of having to avoid crossing over to negative assets by conducting large debt issues in the interim.

This last problem arises as well in the main deterministic analysis. Waiting thirty years leads to a smaller required adjustment. Fifteen plans (37.5 percent of those in the sample) are insolvent before thirty years, under current contribution rates (online appendix table A2). The authors rely on the idea that plans can issue debt to "smooth through the period of peak pension cash flow demand," yet add any volatility and there is a chance that even more could become insolvent. It seems critical that benefit growth really would slow dramatically. It also seems that waiting would stabilize the funding ratio at a much lower level. The intuition for why this is ultimately less costly seems to be that GDP is higher (and normal costs are slightly lower) at the time when the adjustment would start.

This gets to the other main driver of the authors' results for the goal of stabilizing implicit debt to GDP under the (nearly) risk-free rate scenario, and that is that the discount rate (and rate of return) is less than the rate of economic growth ($\delta = r < g$). There are many reasons why such an assumption may not hold. Not least of these is that the growth rate g is the growth rate of public payrolls, which cities and states might well lower if they need to increase contributions.

A further consequence of this modeling is that the authors' recommended contribution rates are actually well below the normal cost rate (Costrell and McGee 2020). An online appendix table provided by the authors

(table A9) allows the reader to see this. Baseline normal costs measured at the 0.5 percent rate are 51.3 percent of payroll in 2017 and 47.5 percent in 2047. This highlights the fact that the goal of stabilizing the implicit debt-to-GDP ratio involves the recommendation of contributing less than the normal cost. The authors do not seem too worried about this, but to me it highlights the fundamental dynamics in the paper: due to the assumptions about payroll growth and returns, the authors are concluding that plans can contribute less than the present value of newly accrued benefits and still achieve stability. Since online appendix table A9 shows that normal costs are not really declining by that much over time, this must involve plans bearing substantial financial risk.

In sum, this paper provides a wealth of calculations and scenarios that will be useful to anyone studying this issue. My takeaways from the analysis are somewhat different from the conclusions of the authors: stabilization would require quite substantial contribution increases of 50 percent or more on average, and much more for specific plans; scenarios that involve waiting to address the issue or goals of stabilizing debt to GDP depend heavily on the $\delta < r < g$ assumption, which may not be appropriate; state pricing requires us to put more weight on bad states of the world when assessing scenario analyses; and the volatile paths that assets might take require future modeling to address the possibility of interim insolvency if pension systems are going to be content with goals less aggressive than paying down unfunded liabilities.

REFERENCE FOR THE RAUH COMMENT

Costrell, Robert M., and Josh McGee. 2020. "Sins of the Past, Present, and Future: Alternative Pension Funding Policies." Paper prepared for the Municipal Finance Conference, Hutchins Center on Fiscal and Monetary Policy, Brookings Institution, Washington, July 13–14.

GENERAL DISCUSSION Henry Aaron remarked that Deborah Lucas's comments suggested that volatility may make pensions unsustainable by increasing the risk that assets will be exhausted. Aaron questioned whether, taking risks into account, pension managers could take steps to insure against or reduce the likelihood of exhaustion, and he wondered how administratively feasible or effective such insurance would be. He asked if the relative optimism of the paper would be restored if such steps were feasible.

Deborah Lucas responded to Aaron's comments by noting that what is salient from her analysis of the defined benefit contribution plan is that

it found that it was not possible to guarantee people a reasonable benefit using investments in a mix of risky and safe assets and to also have the existing contribution rates.[1] She noted that this implies that if a pension invests in enough risky assets to create the desired benefit levels with the existing contribution levels, the pension is relying on an equity premium with a risk that the fund will run out. Lucas said that if a pension bought insurance, then the cost of the insurance is the difference between the risk-free liability and the risky assets. Governments could pre-commit to diverting revenues intended to pay for other spending to make pension benefits safer, but that would be politically difficult.

Janice Eberly noted that both Lucas and Joshua Rauh recommended adding uncertainty to the analysis in their comments. She wondered whether the work that Lucas had already done changed the authors' view on heterogeneity. Eberly remarked that she could see how plans could be different because of investment profiles and asked if there were other significant sources of variation that would add to that heterogeneity besides the ones the authors had already identified.

Rauh responded to Eberly's question, emphasizing that dealing with heterogeneity requires using a stochastic simulation such as a Monte Carlo simulation. He noted that standard error bounds are not as effective due to the asymmetry in the way that asset returns are typically modeled (such as in the Black-Scholes-Merton lognormal model used in his comments), and that the median outcome is much below the expected return in this type of model.

Eberly clarified that she was asking if adding risk would provide different results for different plans in excess of the heterogeneity from investment profiles. She thought that Rauh's comment about asymmetry might imply that the answer was yes.

In response to Eberly's question, Louise Sheiner commented that the paper took a public finance approach and thought about state and local pensions as government debt. The analysis smoothed through heterogeneity in order to take a broad look, similar to looking at the unified budget instead of the Social Security Trust Fund. She noted that stochastic analysis would be both helpful and relevant to the analysis, but that it was more complex than accounting for heterogeneity in returns. Sheiner observed that there are multiple factors such as wages, demographics, and GDP that affect returns but also have a direct impact on pensions. She also commented that the

1. Deborah Lucas and Daniel Smith, "How Much Can Collective Defined Contribution Plans Improve Risk Sharing?," *Journal of Investment Management* 18, no. 4 (2020).

authors thought about sustainability analysis similarly to work done for the federal government using projected cash flows. She mentioned that since a lot of revenues for the federal government come from capital income, expected values are used. Sheiner said that while there is a lot of uncertainty surrounding this type of analysis, most people seem to be comfortable with this methodology for the federal government. She also pointed out that a state like California has a lot of similarities to the federal government, particularly in the fact that it can respond if something goes poorly.

Caroline Hoxby noted that there is literature that suggests that public sector employees would be willing to accept less than the present value of their future pensions if they got more of it in current income. Hoxby also mentioned that there are political economy reasons why government employees tend to receive compensation in pensions rather than in salaries or present-day income. She asked if the authors had considered this when thinking about stabilizing pensions. David Wessel asked if unfunded liabilities are providing misleading metrics for policymakers, since the unfunded liability approach does not fully incorporate changes in pension benefits for new hires.

Sheiner responded to Hoxby's point by saying that the authors have not considered the literature mentioned, but she noted that plans may be able to lower compensation in the future, since the current competitive labor market equilibrium implies higher wages. She mentioned that interest rates have come down significantly, which has increased the value of defined benefit plans, but that this is not something that will continue going forward. Sheiner also mentioned that approximately 25 percent of state and local government employees are not covered under Social Security and that they would be willing to receive higher compensation instead of Social Security benefits.

In response to Wessel's question on misleading metrics, Sheiner said that the authors believed that the value of the paper was in the similarities of the public pension space to the lowering of rates of retirement twenty years in the future during the Greenspan Commission. From a long-run public perspective, Sheiner noted that smoothing out was necessary, and while assets might dwindle now, there would be relief later, unlike the Baby Boom, when benefits went up and stayed up.

Rauh, in agreement with Sheiner on her response to Wessel, said that he believed that the unfunded liability ratio was not any more or less misleading than the debt-to-GDP ratio. Sheiner noted that the measurement looks at current stock, not flows to project, and that the authors would like to extend this analysis for state and local governments as a whole.