Comment on "Game Over: Simulating Unsustainable Fiscal Policy" by Richard W. Evans,
    Laurence J. Kotlikoff, and Kerk L. Phillips
Douglas W. Elmendorf
Congressional Budget Office
May 2012

This paper by Rick Evans, Larry Kotlikoff, and Kerk Phillips is clever and thought-provoking, and I am pleased to have the opportunity to discuss it. I certainly agree with the opening lines of the abstract: "Fiscal sustainability is one of the most pressing policy issues of our time. Yet it remains difficult to quantify." The question of how best to quantify fiscal sustainability is one that my colleagues and I at the Congressional Budget Office (CBO) spend a fair amount of time thinking about. Therefore, I try to do two things in my remarks: First, I make some specific comments about the approach to quantification and simulation used in this paper. Second, I discuss CBO's approach to quantifying the fiscal challenges facing the U.S. federal government—describing what we do now, some limitations of what we do, and ways we plan to strengthen our analysis.[1]

The Approach to Quantification and Simulation Used in the Paper

The paper by Rick, Larry, and Kerk examines the probability and timing of insolvency of a government program that transfers resources from young to old, as in the U.S. Social Security program. The key elements of the approach are as follows:

- The program has benefit payments that are fixed and tax collections that vary with the size of the economy. There is uncertainty about the size of the economy and thus about the amount of taxes collected. With fixed benefits and uncertain taxes, there is a possibility of the program going bankrupt.

- When the program goes bankrupt, it can shift to one of two regimes depending on the simulation: One is a complete shutdown of the economy, and the other is a permanent shift to a program in which the benefit payments equal a fixed share of the wages of the young. Naturally, the second regime seems to me a more plausible and interesting one.

- The model incorporates two periods with overlapping generations and rational expectations. The financial system of the economy has both government bonds and claims to risky capital, so the simulation results include an equity premium. The imbalance in the government's intertemporal budget is labeled the "fiscal gap," and the simulation results also include that gap.

- The model is calibrated along some dimensions to the U.S. economy.

This is an interesting setup. It lets the authors explore the interplay of fiscal policy and asset prices in an uncertain world, and the simplicity of the model makes its workings fairly

---

[1] The discussion of CBO's analysis refers to the agency's analytic tools and estimates as of December 2011 when a preliminary version of these remarks was presented at an NBER conference.

transparent. A number of the qualitative conclusions seem sensible: Higher precautionary savings extend the time to bankruptcy, and higher transfers reduce the time to bankruptcy. In cases when the bankruptcy of the program leads to a complete shutdown of the economy, the equity premium is roughly in line with the historical premium. Both the fiscal gap and the equity premium increase as the program approaches bankruptcy; therefore, those indicators provide useful signals that bankruptcy is approaching.

However, there are several problems in trying to link the quantitative results of this modeling to the actual U.S. economy. One is that the modeling shows that a drop in wages reduces taxes paid into the Social Security-like program but leaves benefits unaffected. The true situation is not as stark as the paper indicates, however, because aggregate wages affect Social Security benefit payments as well as its tax receipts: When one claims Social Security benefits for the first time, the amount one receives is indexed to average wages. After that initial calculation, however, one's benefits rise with prices rather than wages. So a drop in wages affects the program's tax revenues more than its benefit payments, and indeed hurts the finances of the program—but not as much as suggested by the modeling.

Another problem in linking this model to the U.S. economy is that there are other sources of uncertainty besides aggregate productivity growth. In CBO's stochastic modeling for Social Security, we allow for variation in productivity growth, but also in other economic variables and in demographic outcomes. In a paper we published half a dozen years ago, we estimated that productivity growth was indeed one of the largest sources of uncertainty about Social Security's finances. However, there was also significant uncertainty stemming from fertility, immigration, mortality, and various economic factors.

A third problem is that the U.S. government operates other intergenerational transfer programs besides Social Security. If one cares about the unsustainability of current U.S. fiscal policies, one should care particularly about health care programs for older Americans; outlays for those programs are roughly as large as outlays for Social Security today and are growing faster. One could view those programs as being like Social Security in some respects but with an additional critical source of uncertainty—the growth in health care costs per beneficiary of those programs. All told, uncertainty about productivity is not all or even most of the uncertainty associated with U.S. fiscal outcomes.

A fourth problem is that allowing for only two periods of life limits people's ability to smooth consumption by trading with people from other generations. I am concerned that this limitation might distort the estimated equity premium. For example, since the consumption of the young can be driven to zero, should we not expect that the state prices for consumption would be even higher than they are? Are they held down because the young cannot trade with anybody?

Addressing all of these issues and others would be even more challenging than what the authors have already accomplished, so I am not suggesting that they have overlooked straightforward alternatives. Still, given these considerations, I think the model's value is primarily in suggesting issues and relationships to have in mind rather than in providing a realistic appraisal of the risks facing U.S. fiscal policy.

Indeed, the quantitative estimates may seem surprisingly benign to readers of Larry's other writings or CBO's projections. The authors estimate that the expected time to bankruptcy of the Social Security-like program is about 100 years, with a 35 percent chance of such bankruptcy occurring in 30 years. The fiscal gap—the difference between the net present value of expected revenues and expected benefits—is less than 4 percent of GDP in most of the scenarios and is actually negative in some scenarios (because of the regime shift).

In contrast, CBO's long-term budget projections imply even more significant risks and even larger fiscal gaps under current U.S. policies.[2] We publish long-term projections of federal debt under two scenarios. The "extended baseline scenario" is an extension of our regular budget projections, which are based on current law. Under current law, the expiration of the tax cuts enacted since 2001, the growing reach of the alternative minimum tax (AMT), the tax provisions of the recent health care legislation, and the way in which the tax system interacts with economic growth would result in steadily higher revenues relative to GDP. Revenues would reach 23 percent of GDP by 2035—much higher than has typically been seen in recent decades—and would grow to larger percentages thereafter. At the same time, under this scenario, spending on everything other than the major health care programs, Social Security, and interest on the debt would decline to the lowest percentage of GDP since before World War II. That significant increase in revenues and decrease in the relative magnitude of other spending would offset much—though not all—of the rise in spending on health care programs and Social Security. As a result, debt held by the public as a share of GDP would increase only slowly from its current high level.

However, CBO estimates that the budget outlook is much bleaker under an "alternative fiscal scenario," which reflects what one might think of as current policies. In particular, this scenario incorporates several changes to law that are widely expected to occur or that would modify some provisions of law that might be difficult to sustain for a long period. In this scenario, the tax cuts are extended; the reach of the AMT is restrained to stay close to its historical extent; over the longer run, tax law evolves further so that revenues remain near their historical average of 18 percent of GDP; and certain spending programs deviate from current law. Under those policies, federal debt would grow very rapidly. Debt held by the public as a share of GDP would exceed its historical peak of 109 percent by 2023 and would approach 190 percent in 2035. The fiscal gap in this scenario is estimated to be between 4 percent and 5 percent of GDP over the next 25 years (equivalent to about $700 billion this year) and more than 8 percent of GDP over the next 75 years as a whole.

Clearly, current policies are unsustainable, and they appear more unsustainable in CBO's projections than in the simulation results of this paper.

The Approach to Quantifying Fiscal Policy Used by CBO

Let me now turn from the paper by Rick, Larry, and Kerk to discuss what we do at CBO to quantify the fiscal challenges facing the U.S. federal government. We currently use four different analytic approaches:

---

[2] See Congressional Budget Office, *CBO's 2011 Long-Term Budget Outlook*, June 2011.

- First and most important, CBO regularly constructs projections of spending, revenue, deficits, and debt. Three times a year, we publish projections looking ahead 10 years; once a year, we publish projections that extend 75 years, although we focus on the first 25. As I just mentioned, those projections show that current U.S. fiscal policies would keep federal debt on an unsustainable trajectory.

- Second, CBO regularly uses its long-term projections to estimate the fiscal gap. We define the gap as the present value of revenues over a given period minus the present value of noninterest outlays over that period, adjusted to keep federal debt at its current percentage of GDP.[3] As I just mentioned, the estimated fiscal gap based on current policies exceeds 4 percent of GDP over the next 25 years and 8 percent of GDP over the next 75 years. Because revenue has averaged 18 percent of GDP and spending 21 percent, a gap of that magnitude requires a large change in policies.

- Third, CBO regularly quantifies the effects of delay in closing the fiscal gap. For example, we estimated that if policymakers wait about a decade to change policies, the gap rises from 4 percent or 5 percent of GDP to around 8 percent—not even counting the feedback effects on the economy.

- Fourth, CBO sometimes estimates the distributional impact across generations of waiting to resolve the long-term budget imbalance. In a report in December 2010, we estimated that stabilizing the ratio of debt to output in 2025 instead of 2015 would benefit the average person over age 55 today, hurt people not yet born, and have small effects in both directions on people in intermediate cohorts.[4] Of course, in weighing distributional burdens, policymakers also need to take into account the progression of underlying living standards.

Those are the analytic approaches that CBO is currently using to quantify fiscal challenges. Let me offer three observations about what I see as the limitations of those approaches and some directions for improvement.

The first observation is that all of the approaches listed above involve point estimates and do not explicitly address uncertainty. Indeed, almost all of our analysis involves point estimates rather than ranges or probabilities.

There are some good reasons for that predilection: One is that the Congressional budget process operates with point estimates. Committees are given allocations of funds, and those are expressed as point values. Another reason for our focus on point estimates is that our methodologies do not readily yield measures of uncertainty. CBO's projections for the economy and the budget do not generally come from formal probability models, so ranges and probabilities do not fall out naturally in the projection process; instead, they would need to be constructed separately. A further reason we focus on point estimates is that communicating

---

[3] CBO's calculations use a discount rate equal to the average interest rate on federal debt held by the public, which is projected to be 2.7 percent on an inflation-adjusted basis in the long term.

[4] The estimates depended in part on how the debt was stabilized—by raising marginal tax rates or by reducing federal transfer payments (which go mainly to older people).

uncertainty in an effective way without obscuring the basic results is difficult. When we report ranges for our estimates, it is common for people who would prefer that our estimate be smaller to quote the bottom of our range and for people who would prefer that our estimate be larger to quote the top of our range, which muddies the public discussion of our estimates at least as much as it illuminates it. Thus, the practical gains from our analysis of uncertainty are often smaller than one would hope.

That said, we think it is important that policymakers understand the uncertainty of our methodologies and our estimates, so they can take this uncertainty into account in their decision-making. Therefore, we think and write about uncertainty when we can, and I will say more about that in a moment.

My second observation is that, in constructing point estimates, CBO aims to be in the middle of the distribution of possible outcomes. I have used that phrase repeatedly when talking with Members of Congress.

I am not usually explicit, though, about whether the word "middle" refers to the mean or the median of the distribution. In many contexts, the mean and median are probably fairly close to each other, so the distinction is not important. However, for distributions in which the median is noticeably different from the mean—say, distributions with long tails on one side—the best way for CBO to proceed is not clear. One example is our approach to projecting the unemployment rate. Our current approach captures, we think, the normal ebbs and flows of business cycles, but it may not adequately capture the risk of a severe slump like the Great Depression or the current downturn; therefore, we discuss the chance of such a slump as a risk to the long-run budget and economic outlook. Incorporating the possibility of such an event in our numerical projections could make those projections more accurate, on average, over the long run. Yet, the estimates would be too pessimistic almost all of the time and still far too optimistic on the rare occasions when a severe slump occurs. Would such a change in our estimates make the estimates more useful to the Congress, or is it more useful to continue with our current approach?

Another concern related to our reporting the middle of the distribution of possible outcomes is whether we are aggregating different sources of uncertainty in the most effective way. For example, when we choose expected values for two variables and then construct a third variable from them, our projection of that third value will not be its expected value if the two underlying variables interact in a nonlinear way. Many of our estimates involve nonlinearities, so we think about how to cope with this challenge—but we do not have a good general way of dealing with it.

My third observation is that CBO is working to be more explicit about the uncertainty in its budget and economic projections. Let me mention four examples:

- One example is to show ranges of effects based on different parameter assumptions. We have been doing this for our estimates of the effects of fiscal policies on the economy, regarding both the near-term impact through changes in aggregate demand and the medium-term and long-term impacts through changes in potential output. Specifically, we have published ranges of estimates for various policies corresponding to different

short-term multipliers, different crowding out of investment by government debt, and different elasticities of labor supply with respect to marginal tax rates.

- Another example is our ongoing efforts to extend our stochastic analysis of Social Security to the rest of the budget. For Social Security, we have allowed most of the key demographic and economic factors that underlie the analysis—including fertility and mortality rates, interest rates, and the growth rate of productivity—to vary on the basis of historical patterns of variation, and we sometimes publish 80-percent confidence regions for our projections. In a recent document, for example, we projected that the Social Security trust funds would be exhausted in 2038 but that there was a 10 percent chance of exhaustion in 2030 or earlier and a 10 percent chance of exhaustion in 2059 or later. To extend this approach to the rest of the budget, we are strengthening the health care aspects of the microsimulation model we use in our long-term projections and then will try to quantify the uncertainty about health care spending per beneficiary under current policies.

- A further example is our descriptions of our projections. Our long-term outlook for the budget, which we update each year, now includes a section on the budgetary risks posed by recessions and financial crises, changes in interest rates on federal debt, changes in demographics, changes in health status and health care, long-term changes in productivity, and catastrophic events or major military actions. We also wrote a separate issue brief about the risk of a fiscal crisis, which we defined as investors losing confidence in a government's ability to manage its budget and the government thereby losing its ability to borrow at affordable rates.

- The last example I will mention is analysis of the ways that alternative policies expose the government budget to more or less risk. We are engaged in a project now with Debbie Lucas and Steve Zeldes about the different amounts of risk in Social Security when benefits are calculated according to different formulas. We are also examining the effects of proposals to transform certain federal health care programs from defined-benefit programs to defined-contribution programs. Of course, policy choices that insulate the federal budget from risk may achieve that insulation by shifting the risk to benefit recipients or taxpayers, and our analyses will make that clear as well.

Conclusion

In sum, this paper by Rick Evans, Larry Kotlikoff, and Kerk Phillips illuminates some key relationships between fiscal policy and the economy that are important for judging the sustainability of that policy. Their paper, the projections of CBO, and research by other budget and economic analysts show that the current policies of the U.S. federal government have put federal debt on an unsustainable path and that the adjustments needed to achieve sustainability are very large. Illustrating the consequences of such an unsustainable fiscal policy is critically important work.