



Estimating the Cost of One-Sided Bets: How CBO Analyzes the Effects of Spending Triggers

Summary

All legislative proposals involve some uncertainty about their implementation and effects. In many cases, the uncertainty is limited, and the Congressional Budget Office can estimate the costs by focusing on a small range of possible outcomes. For some proposals, however, producing a meaningful estimate requires CBO to consider the probabilities of multiple outcomes with costs that might differ significantly. Accounting for such possibilities can have a significant impact on the estimate of the proposal's budgetary effect; in some cases, it can be the difference between an estimate of zero cost and one of a positive cost.

One key category of proposals for which CBO must consider the likelihood and costs of multiple outcomes comprises proposals with a particular structural feature: Their legislative language makes their effects sensitive to whether some uncertain variable—for example, the price of crude oil or an inflation or unemployment rate—passes a “trigger” threshold value. CBO refers to such trigger provisions as one-sided bets because they affect costs on only one side of the trigger value.

As illustrated in the examples presented below, CBO measures the cost of proposals involving new or existing one-sided bets—and of ongoing programs that involve such bets—as the estimated probability-weighted average of the costs of the full range of possible outcomes, also known as the expected-value cost. That is, CBO calculates its estimate of the cost of such a proposal or program by estimating the cost for each of the possible outcomes, assigning each outcome a weight on the basis of the probability that it occurs, and taking the weighted

average of those costs. The key step in the analysis is determining the appropriate probabilities of outcomes driven by one or more variables, each of which typically can take on a range of values. The examples illustrate the possibility that even if the trigger of a one-sided bet was not reached in CBO's baseline projections, the agency would nevertheless estimate that the bet would impose a cost because the trigger *could* be reached.

CBO's weighted-average estimates are not predictions of actual future costs. Even if the agency's analysis of all the possible outcomes and their probabilities was completely accurate, the actual cost would most likely still differ from the estimate because of the randomness of the uncertain variable. Consequently, CBO expects most weighted-average estimates to end up being either too high or too low. The agency's goal in making such estimates is to be correct on average.¹

The Need for Different Methods of Estimating the Costs of Legislative Proposals

CBO provides point estimates of the budgetary effects of proposed legislation because the Congress needs such estimates to enforce budgetary rules. The purpose of the estimates is to provide a concise story about a legislative proposal's likely effects on federal outlays and revenues.

Effects on revenues or mandatory spending—that is, spending generally determined by formulas and eligibility rules rather than by annual appropriations—are

1. This report updates Congressional Budget Office, *Estimating the Costs of One-Sided Bets: How CBO Analyzes Proposals With Asymmetric Uncertainties* (October 1999), www.cbo.gov/publication/11843.

measured in relation to CBO's baseline budget projections, which are a set of detailed year-by-year projections of what spending, revenues, deficits, and debt would be if current law generally remained unchanged. The costs of proposals that affect discretionary (that is, appropriated) spending are also estimated in relation to current law, which reflects only appropriations enacted to date.²

CBO's approach to estimating the budgetary effects of legislative proposals varies with the specific characteristics of the proposal under consideration. For example, some bills would authorize appropriations of specific amounts of money: CBO estimates how quickly those funds would be spent if appropriated, depending on their purpose. Other legislation specifies activities that agencies would be required to undertake: CBO estimates how much those activities would cost and projects when they would be carried out. Still other bills specify conditions under which a change in mandatory spending would occur: CBO projects—on the basis of its economic forecast and other factors underlying its baseline projections—whether and when those conditions would be met and how the bill would affect spending if they were.³

Most legislative proposals analyzed by CBO authorize specific amounts of money, involve relatively small costs or savings, or establish programs or activities that are similar to existing ones; such proposals generally involve little uncertainty. To estimate costs in those cases, CBO relies on feedback from agencies and other entities affected by the legislation and looks at historical spending for the affected program or others like it.

Other proposals are more complex, and their outcomes are more uncertain. Often, little or no data exist to allow CBO to assign weights to the budgetary effects associated with each possible outcome. In such cases, analysts consult the affected agencies and outside experts and research similar changes in other programs. After analyzing that information, CBO may perform a sensitivity analysis to ensure that, given the range of views about the proposals' effects and the uncertainties about key factors, the agency's point estimate is close to the median—in

other words, that it is about as likely to be too high as it is to be too low.⁴

For some legislative proposals, however, explicit analysis of the probabilities of various possible outcomes *is* feasible, and it plays an important role in formulating an appropriate cost estimate. Those more analytically complex proposals involve one or more one-sided bets. A related category of proposals includes those whose costs depend on a future administrative or judicial decision that could go one of two ways (see Box 1).

Estimating the Costs of Proposals and Programs With One-Sided Bets

To analyze proposals and ongoing programs with one-sided bets on one or more uncertain factors that can take on a range of values—such as commodity prices, interest rates, and inflation rates—CBO estimates the probabilities of potential outcomes for those factors and uses the estimates to calculate a weighted-average cost. A one-sided bet may take the form of a “floor” (a minimum threshold) or a “ceiling” (an upper limit), and a single proposed or existing program may include multiple bets with different trigger levels.

The significance of one-sided bets is illustrated by a simple example about a legislative proposal that would provide \$100 million for a hypothetical new job-training program for unemployed workers. That spending would be triggered only if the unemployment rate exceeded the rate underlying CBO's baseline projections in a specified year (see Figure 1, top panel).

Simply using the baseline unemployment rate as the basis for estimating the cost of the legislation would suggest that the proposal had no cost—because if the actual unemployment rate equaled that baseline rate, the spending would *not* be triggered. But given the uncertainty in projections of the unemployment rate, there is a 50 percent chance that the actual rate would be higher than that baseline rate, in which case the cost of the proposal would be \$100 million. Taking that probability

2. For more information on the agency's baseline projections, see Congressional Budget Office, *How CBO Prepares Baseline Budget Projections* (February 2018), www.cbo.gov/publication/53532.

3. For an overview of CBO's methods for analyzing legislative proposals, see Congressional Budget Office, *How CBO Prepares Cost Estimates* (February 2018), www.cbo.gov/publication/53519.

4. The information available in such cases is insufficient to allow CBO to consider whether the estimate is likely to be correct on average—the question underlying the expected-value approach.

Box 1.**CBO's 50/50 Approach to Estimating the Budgetary Effects of Proposals With Two Distinct Outcomes**

Often, when it analyzes proposals whose costs hinge on a future decision with two possible outcomes (such as a court judgment or an agency decision to enact or abandon a proposed regulation), the Congressional Budget Office does not have enough information to identify either outcome as more likely than the other. Under those circumstances, CBO typically assigns probabilities of 50 percent to both outcomes. When CBO uses that 50/50 approach, its estimate of the budgetary effects is simply the average of the effects estimated for each outcome.

The 50/50 approach can be used in any context involving two possible outcomes that, given the available information, seem about equally likely. In cases involving significant uncertainty about whether a particular administrative action will be taken or how a court case will be judged, CBO may use the approach as a simplification, even if one of the two possible outcomes seems somewhat more likely than the other.¹

CBO used the 50/50 approach to prepare its estimate for the Coastal and Marine Economies Protection Act (H.R. 1941), which would have prohibited oil and gas leases in the Atlantic and Pacific regions of the Outer Continental Shelf (OCS).² Under its current five-year plan, which is set to expire in 2022, the Department of the Interior (DOI) does not authorize any

1. CBO's use of the 50/50 approach is discussed in Congressional Budget Office, letter to the Honorable John M. Spratt Jr. explaining how CBO accounts for anticipated administrative actions in its baseline projections (May 2, 2007), www.cbo.gov/publication/18615.

2. See Congressional Budget Office, cost estimate for H.R. 1941, the Coastal and Marine Economies Protection Act (July 12, 2019), www.cbo.gov/publication/55450.

into account, CBO would estimate that the proposal would cost \$50 million.⁵

To estimate the costs of legislation involving more complex one-sided bets, CBO would use the same basic

5. The arithmetic would be the same as in the 50/50 approach discussed in Box 1, although the rationale here does not involve a future court ruling or regulatory decision. The distribution shown is a simplification for illustrative purposes; it does not reflect the fact that unemployment rates can take on relatively high values, which would imply a longer right "tail."

auctions for leases in the Atlantic or Pacific OCS, but a key uncertainty that CBO faced when estimating the cost of the bill was whether DOI would, after that plan expires, authorize such auctions to occur under current law.

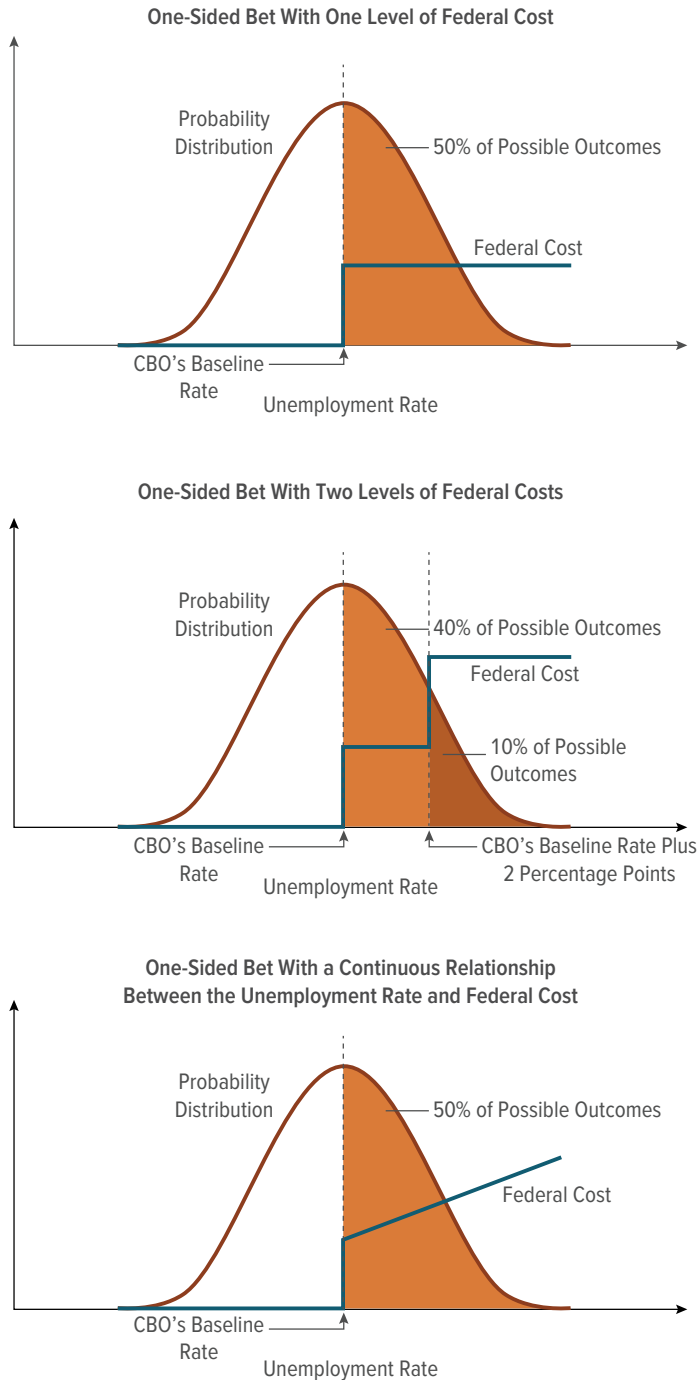
Lacking any basis for determining what DOI would do, CBO modeled two outcomes under current law, each of which was estimated to have a 50 percent chance of occurring. In one case, no auctions are held for Atlantic or Pacific OCS leases, and in the other, such auctions do occur. In the first case, the prohibition in the bill would have no budgetary effect; in the second, if the proposal was enacted, the winning auction bids that the government would have collected under current law would instead be forgone. To account for that second possibility, CBO estimated the receipts from those potential future auctions under current law. Projecting that roughly 300 leases would be auctioned and that the average winning bid would be \$1.3 million (an estimate based on the average value of leases acquired in the Gulf of Mexico in 2018), CBO estimated that the auctions would generate receipts of \$400 million. Thus, given the two possible outcomes, the agency estimated that the cost of prohibiting such auctions would be 50 percent of \$400 million, or \$200 million.

As the example illustrates, estimates derived from the 50/50 approach are expected to be too high under one of the two possible outcomes and too low under the other. CBO's goal with that approach is to produce estimates that are correct on average; that means that regardless of which outcome occurs, the agency's estimate would be expected to differ from the actual cost by the same amount.

approach of calculating probability-weighted averages. Take, for example, a proposal that included a provision that would increase funding for job training to \$200 million if the unemployment rate was 2 percentage points or more higher than CBO's baseline estimate of that rate (see Figure 1, middle panel). That proposal involves two one-sided bets—one on whether the unemployment rate will exceed CBO's baseline rate and another on whether it will exceed that baseline rate by more than 2 percentage points. The illustrative probability distribution indicates that the chance that the unemployment rate was at least 2 percentage points greater than the baseline

Figure 1.

The Effect of a One-Sided Bet on the Estimated Cost of Three Illustrative Proposals for a Job-Training Program for Unemployed Workers



The first proposal would provide \$100 million for a new job-training program if the unemployment rate in a specified year exceeded CBO's baseline projection of that rate. As shown here, there is a 50 percent chance that the proposal's cost would be zero and a 50 percent chance that it would be \$100 million. CBO's probability-weighted average cost estimate of the proposal would be \$50 million ($0.5 \times 0 + 0.5 \times \100 million).


The second proposal differs from the first in that it would provide \$200 million in funding if the unemployment rate in a specified year exceeded CBO's projection by 2 percentage points or more. In effect, this proposal has two one-sided bets—one for each cost threshold. There is a 50 percent chance that the cost of the proposal would be zero, a 40 percent chance that it would be \$100 million, and a 10 percent chance that it would be \$200 million. CBO's probability-weighted average estimate of the proposal's cost would thus be \$60 million ($0.5 \times 0 + 0.4 \times \100 million + $0.1 \times \$200$ million).

Under the third proposal, the funding provided for the job-training program would be at least \$100 million if the unemployment rate exceeded CBO's baseline projection, with proportionally higher funding for higher rates. The calculations involved in estimating the costs of the proposal would be more complex than those for the above versions, but they would reflect the same basic approach of calculating the probability-weighted average cost.

Source: Congressional Budget Office.

Figure 2.

Illustrative Probability Distributions



Probability distributions, which underlie CBO's analysis of one-sided bets, may take a variety of forms, including bell-shaped, asymmetric, and multipeaked. In its analyses, CBO typically approximates a continuous distribution by using a large number of discrete values (as shown here) or simulates the distribution by drawing a large number of random samples.

Source: Congressional Budget Office.

rate is 10 percent and that the chance that it was greater than the baseline rate by less than 2 percentage points is 40 percent. Thus, CBO would estimate the proposal's cost to be \$60 million—the sum of zero times its probability of 0.5, \$100 million times its probability of 0.4, and \$200 million times its probability of 0.1.

Most one-sided bets that CBO encounters in legislative proposals have costs that can take on a full range of values, not just a few discrete values. For the job-training program example, that possibility is illustrated by a proposal that would provide funding of \$100 million or more if the unemployment rate exceeded the rate underlying CBO's baseline projections; the total amount of funding would be determined by a formula that depends on the unemployment rate (see Figure 1 on page 4, bottom panel).

To estimate costs for proposals, such as that one, that involve a continuous probability distribution and a range of possible costs, CBO uses one of two computational approaches. In some cases, the agency creates an approximation of the continuous distribution—which can take a variety of forms—by assigning probabilities to a large number of discrete values (see Figure 2). In others, CBO randomly draws a large number of values for the uncertain variable while taking into account the shape of the

continuous distribution (so that more likely values occur more frequently in the sample).

One-sided bets appear in many legislative contexts. Two examples are discussed here—one involving projections of spending for existing agricultural support programs and the other involving an option for changing the student loan program. Both cases illustrate the use of random sampling to develop probability distributions that are consistent with historical data.

Whenever possible, CBO analyzes historical data to produce its probability-based estimates, but in some cases, sufficient data are unavailable. For example, to estimate the cost of the federal program that provides reinsurance for terrorist attacks, CBO had to consider the likelihood of terrorist attacks of different sizes in the United States—but the agency could not base its probability distribution on the few such attacks that have taken place.⁶ To model the probabilities of attacks of different sizes, CBO needed a downward-sloping distribution with a

6. The program involves multiple one-sided bets, including a liability cap and individual deductibles for each insurer with claims from a terrorist attack. See Perry Beider and David Torregrosa, *Federal Reinsurance for Terrorism Risk and Its Effects on the Budget*, Working Paper 2020-04 (Congressional Budget Office, June 2020), www.cbo.gov/publication/56420.

long right “tail” because larger attacks, though possible, are less likely than smaller attacks. To select a particular distribution from among the countless distributions with that general shape, CBO limited its choice to the “log-normal” type, and after consulting with experts, it identified the particular log-normal distribution that reflected a few key characteristics: the annual probability of an attack, the average annual insured loss from attacks, and the average annual probability of an attack as large as or larger than the attacks of September 11, 2001.⁷

Agricultural Support Programs

The Agricultural Act of 2014 (Public Law 113-79) created two new agricultural support programs: Price Loss Coverage (PLC) and Agricultural Risk Coverage (ARC). Both programs are available to producers of 23 specified commodities. CBO analyzes the programs together, in part because producers must choose one to participate in—they cannot be in both at the same time.⁸

The two programs protect producers in different ways: PLC protects them against low crop prices, and ARC protects them against low crop revenues per acre (calculated as price times yield per acre). Both programs represent one-sided bets. The PLC program, the simpler of the two, is a bet on crop prices: In years when a particular crop’s market price exceeds the program’s reference price (the trigger), the program incurs no costs for payments to producers of that crop. When the market price is below the reference price, participating producers receive payments based on the following formula: the shortfall in price times the producer’s registered acres for the crop (which may differ from the planted acres), times the specified payment yield for the crop, times 85 percent. Thus, all else being equal, the program’s costs are

greater the wider the gap between the market price and reference price (see Figure 3, top panel).⁹

CBO’s analysis of the two programs reflects 30 years of data on the correlations between each crop’s yields and its prices. (High yields tend to be associated with low prices and low yields with high prices.) Moreover, for six major crops, CBO’s analysis also reflects the historical correlations between the variables for different crops—for example, the positive correlation between the prices of soybeans and corn, two crops that are often grown in rotation.¹⁰

Though some details of PLC and ARC were changed by the Agriculture Improvement Act of 2018 (P.L. 115-334), CBO uses the same basic approach for its current baseline projections of the programs’ costs as it used to produce the cost estimates for the 2014 legislation. In brief, that approach starts with projected values of crop prices and average yields that reflect the latest available supply and demand data and the professional judgment of CBO’s staff and outside advisers. The agency then accounts for the range of possible outcomes by establishing probability distributions around the projected values. Next, it draws 1,000 sets of random samples from each distribution and adjusts them to reflect the historical correlations among the price and yield variables. CBO then calculates the PLC and ARC costs associated with each of the 1,000 sets of samples and averages the sampled costs to produce the final estimates.

To elaborate, the first step is to develop projections for the price and national average yield per acre for each of the 23 commodities during each year of the 10-year projection period.¹¹ To do so, CBO first develops preliminary projections, taking into account government data on several variables—including crop prices, acreage planted and harvested, yields, ending stocks, and

7. A random variable has a log-normal probability distribution if its logarithm is normally distributed. A normal distribution is a symmetric, bell-shaped distribution determined by a particular mathematical formula.

8. Initially, the program covered 22 crops and required producers to make a onetime choice between the two programs for the five years in which the bill would be in effect. The Bipartisan Budget Act of 2018 added seed cotton to the list of covered commodities; the 2018 farm act required producers to commit to one of the two programs for 2019 and 2020, but thereafter, it allows them to choose which program they will participate in annually.

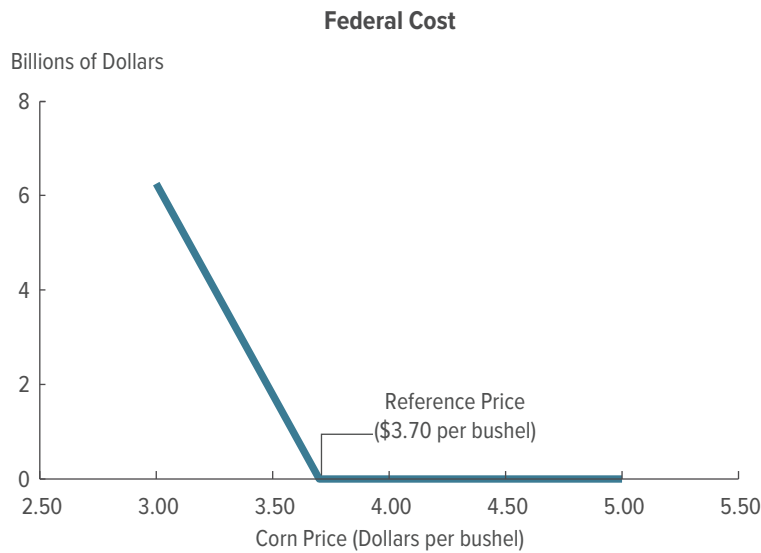
9. Under the Agricultural Act of 2014, reference prices were set by statute. Currently, under the Agriculture Improvement Act of 2018, the effective reference price may increase by up to 15 percent if the moving average of the annual market prices from the previous five years, excluding the lowest and highest of those five prices, exceeds the statutory reference price.

10. CBO models correlations between corn, soybeans, wheat, and cotton as well as the correlation between barley and oats.

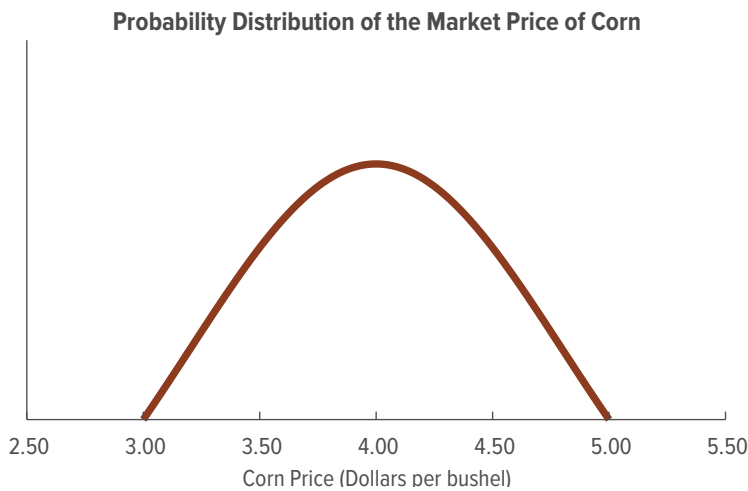
11. Specifically, the analysis projects the national average of county-level average yields per acre. The average of the county-level yields is the same as the national average itself, but the county-level figure exhibits more variability, and that variability affects estimates of the cost of the ARC program.

Figure 3.

Budgetary Effect of Price Loss Coverage for Agricultural Commodities



The Price Loss Coverage program is a one-sided bet: The government incurs costs when the crop price is below the legislatively determined reference price but does not reap savings when the crop price is above that reference price. In this example, the reference price for corn is \$3.70 per bushel. When market prices are less than that amount, the federal cost equals the reference price minus the market price, times the total number of acres registered for the program (here, 75 million), times the specified yield (140 bushels per acre), times 85 percent.



When estimating the costs of the program, CBO accounts for the probabilities of prices above and below the reference price. To do so, the agency models a symmetric distribution centered on its projected market price; that distribution is bounded to rule out unrealistically low or high prices. In this example, the market price of corn per bushel is projected to be \$4.00, but it could be as low as \$3.00 or as high as \$5.00.

Source: Congressional Budget Office.

demand from various sectors. Those preliminary projections reflect many interactions among the variables, such as the effects that an imbalance between supply and demand in one year has on the next year’s planting decisions. CBO then uses feedback from experts within the Department of Agriculture, the private sector, and academia to fine-tune the preliminary projections.

Next, CBO accounts for the uncertainty around those projections (associated with the variability of weather, among other factors) by establishing a symmetric probability distribution for each variable in each year with the

agency’s projected value as the peak. Specifically, CBO models normal distributions that have been truncated to rule out unrealistically low and high values.¹² (See Figure 3, bottom panel, which shows a simplified

12. For each distribution, both the points at which it is truncated and its standard deviation (which determines how much of the full normal distribution lies between the truncation points) are chosen on the basis of the variability observed in historical data. Currently, the range of possible values for each variable is proportionately narrower in the first year or two of the projection period than it is for years further out and remains constant in those later years.

distribution for corn prices centered at a projected level of \$4 per bushel and truncated at \$3 and \$5 per bushel.)

Then, for each of the 10 years of the projection period, CBO begins the sampling process by drawing 1,000 sets of random samples from the standard normal distribution; each set contains a separate draw for each of the 46 price and yield variables. The 46 draws in each set are independent of one another, but CBO uses the correlation data to adjust them so that they reflect the historical relationships among the variables (see Box 2 for details). Each of the adjusted draws then determines a value above or below the projected average value for the corresponding variable, and the values for the full set of variables determine the federal costs associated with that set of draws. For example, the cost to the federal government of the PLC program for any crop is zero for all of the sets that include a draw for the market price of that crop that is at least equal to the reference price (see Figure 3 on page 7, top panel). Finally, CBO averages the results for all 1,000 sets of draws to calculate the agency's expected-value estimate of federal costs for that year of the projection period.

Interest Rate Caps on Student Loans

The interest rates on student loans issued by the federal government (referred to as direct loans) are indexed annually to the rate on 10-year Treasury notes, subject to various caps.¹³ For example, the rate for undergraduate loans (subsidized and unsubsidized) is generally the 10-year Treasury rate plus 2.05 percentage points, but for Treasury rates above 6.2 percent, the loan rate is capped at 8.25 percent.¹⁴ Those rates are fixed for the life of the loans.

The caps produce one-sided bets. As required by the Federal Credit Reform Act of 1990 (FCRA), the federal cost of student loans is measured as the net present value of all cash flows calculated using Treasury interest rates from the fiscal year in which the loans

were disbursed as discount rates.¹⁵ A key determinant of the net present value of direct student loans is thus the difference between the interest rate on the loan and the Treasury rates used for discounting. As long as the 10-year Treasury rate is below the thresholds that trigger the caps, that difference is roughly constant.¹⁶ If the caps are triggered, however, the difference in the two rates becomes smaller and continues to decrease as the Treasury rates rise, driving federal costs up.

For its 2018 report *Options for Reducing the Deficit: 2019 to 2028*, CBO analyzed the budgetary effects of removing the cap on interest rates for student loans.¹⁷ The analysis was driven by two key factors:

- The Treasury interest rate is one of three variables—along with the unemployment rate and the rate of inflation in consumer prices—that strongly influence each other.
- Fluctuations in those variables tend to carry over from one period to the next.

An important consequence of that second factor is that if a variable is higher (or lower) than CBO's projection in one year, it is likely to be higher (or lower) than projected in the next year as well. Accordingly, CBO modeled the probability distributions for the three student loan variables as paths (rather than snapshots) that evolve together over time. (By contrast, fluctuations in the values of variables in the agricultural support programs are driven largely by weather; CBO models those fluctuations as uncorrelated from one year to the next.)

Using quarterly data from 1970 to 2018, CBO estimated the relationships among the interest rate on 10-year Treasury notes, the unemployment rate, and the inflation rate, as well as the extent of the random fluctuations that

13. Since the 2013–2014 academic year, interest rates on student loans have been based on the high yield of the 10-year Treasury note from the last auction before June 1 of the previous academic year.

14. For unsubsidized loans to graduate students, the rate is the Treasury rate plus 3.6 percentage points, with a cap of 9.5 percent. For PLUS loans (available to parents and graduate students), the rate is the Treasury rate plus 4.6 percentage points, with a cap of 10.5 percent.

15. A present value is a single number that expresses the flow of current and future payments or income in terms of an equivalent lump sum paid or received at a specified time. A present value depends on the rate of interest, or discount rate, used to translate a cash flow in a future year into current dollars.

16. If the difference in rates is large enough to offset the expected loan defaults—as it has generally been in the past—the budgetary cost of the loan program is negative.

17. Congressional Budget Office, “Remove the Cap on Interest Rates for Student Loans,” *Options for Reducing the Deficit: 2019 to 2028* (December 2018), www.cbo.gov/budget-options/2018/54723.

Box 2.**Details of the Sampling Process That CBO Uses to Estimate the Costs of Agricultural Support Programs**

The costs of agricultural support programs depend on whether crop prices and crop revenues per acre fall short of various trigger levels, and if so, by how much—that is, the programs involve what the Congressional Budget Office refers to as one-sided bets. The method that CBO uses to estimate the probability-weighted average costs of those programs is an example of one of the two broad approaches that the agency uses to analyze one-sided bets—the simulation of probability distributions through random sampling.¹

In the case of agricultural support programs, the sampling reflects not only the uncertainty surrounding CBO’s projections of crop prices and yields but also the correlations that exist between some of those variables—or more specifically, between the deviations of the variables from their projected levels. For example, a corn price above its projected value in a given year tends to be associated with a corn yield that is lower than projected and a soybean price that is higher than projected in that same year. Variables that are positively or negatively correlated with each other essentially have a combined, multivariable probability distribution. In its analysis, CBO uses random sampling to simulate multivariable distributions by drawing on the distributions for the individual variables in a way that reflects the correlations observed in the historical data.

For each year of the projection period, the process involves 1,000 repetitions of the following steps:

1. A set of random samples, one for each variable, is drawn from the standard normal distribution. Each draw can have any positive or negative value, but most (68 percent) fall between -1 and 1 , and almost all (95 percent) are between -2 and 2 .
2. The random draws are adjusted to reflect the historical data on the correlations among the prices and yields.

1. The other broad approach involves approximating a predetermined continuous probability distribution by assigning its probabilities to a large number of discrete values.

Specifically, a vector containing the random draw for each variable is multiplied by a matrix that represents the historical correlations among the variables.² After that adjustment, the draws of variables that are positively correlated are closer together in value, and those of negatively correlated variables are further apart. For example, original draws of -1.3 and 0.7 for two positively correlated variables might be adjusted to -1.0 and 0.5 , respectively.

3. The adjusted draws are converted to percentiles of the cumulative standard normal distribution. For example, an adjusted draw of zero would correspond to the 50th percentile (the midpoint) of that distribution, and adjusted draws of -1 and 0.5 would correspond roughly to the 16th and 69th percentiles.
4. The percentile values are mapped to the truncated normal distributions. The truncation points are symmetric, so the 50th percentile always corresponds to the peak of the distribution. The mapping of other percentile values depends on the truncation points and the spread of the distribution. For example, see Figure 3 on page 7, bottom panel; in that illustrative scenario, an adjusted draw at the 16th percentile would correspond to a price of about \$3.40 per bushel, and a draw at the 69th percentile would yield a sampled price of about \$4.30 per bushel.

The sampled values for all the price and yield variables together determine the estimated costs of the two programs in each of the 1,000 iterations of the process. To produce its estimates of the programs’ costs, CBO takes the averages of the results for all of those iterations. The entire process is repeated for each year of the projection period.

2. The matrix is a triangular decomposition of the matrix of correlations. Triangular decomposition is a mathematical operation analogous to the square root operation; in that analogy, the elements of the correlation matrix are like variances or covariances, and the elements of the triangular decomposition are like standard deviations.

affect any particular value of those variables. The agency then used those estimates to simulate 3,000 possible time paths for the three variables over the estimation period, each of which was affected by random shocks drawn from distributions based on the fluctuations observed in the historical data. The 3,000 paths for the 10-year Treasury rate provided the basis for a set of quarterly probability distributions for that rate, which in turn provided the

basis for the distributions for the student loan rates.¹⁸ To estimate the effect of the rate cap, CBO lowered the loan

18. One adjustment was required: The distributions of the rates on 10-year Treasury notes generated by the 3,000 paths were calibrated, using scaling factors, so that the average in each period matched CBO’s baseline forecast of the rate for that period. CBO uses the same basic method to estimate the costs of student loans in its baseline projections.

rates in the distributions that were above the cap to the level of the cap and recalculated the averages.

On the basis of that analysis, CBO estimated that eliminating the caps would result in a savings of \$16 billion over the 2019–2028 period because borrowers' interest payments to the government would be greater, on average, than they would be under current law.¹⁹ By contrast, a nonprobabilistic estimate—one that did not account for the uncertainty in CBO's 10-year economic projections—would have indicated no savings over the period, because the projected loan rates, based on the agency's projections of Treasury interest rates, did not exceed the caps.

19. The savings were estimated using the FCRA method. If the fair-value method, an alternative approach that accounts for market valuations of risky assets, had been used to calculate the cost, the estimated savings would have been \$12 billion.

This report was prepared to enhance the transparency of the work of the Congressional Budget Office and to encourage external review of that work. In keeping with CBO's mandate to provide objective, impartial analysis, the report makes no recommendations.

Perry Beider prepared the report with contributions from Tiffany Arthur, Kathleen FitzGerald, Erik O'Donoghue, and Jeffrey Perry and guidance from Joseph Kile. Robert Arnold, Michael Falkenheim, Ann Futrell, Sebastien Gay, Teri Gullo, Mark Hadley, Justin Humphrey, and Jeffrey Werling provided comments.

Jeffrey Kling and Robert Sunshine reviewed the report. The editor was Bo Peery, and the graphics editor was Robert Rebach. This report is available on CBO's website (www.cbo.gov/publication/56698).

CBO continually seeks feedback to make its work as useful as possible. Please send any comments to communications@cbo.gov.



Phillip L. Swagel
Director

