Valuing Government Obligations When Markets Are Incomplete *

Jasmina Hasanhodzic[†] and Laurence J. Kotlikoff[‡]

October 11, 2014

Abstract

Valuing future government spending commitments and tax receipts, whether they are sure or risky, is critical to assessing the sustainability of fiscal policy. If markets were complete it would be an easy matter to determine if the present value of projected spending exceeded the present value of projected receipts. But such markets don't exist for numerous reasons including the inability of the living to trade with the unborn. Nor can one use prevailing security prices since they don't span long-term states of nature let alone cover all relevant contingencies in the short and medium runs. How then should one value state-contingent government payments and taxes, or, equivalently, discount their expected values?

The approach taken here posits and simulates a ten-period overlapping generations model and uses it to determine the immediate payments—the compensating variations—agents would require to forego promised government future net payments. "Agents" include future generations who are assumed to discount their utility while alive for the number of years it will take for them to be born. Our metric for pricing the compensating variations is agents' expected utility functions. I.e., we determine what immediate payment will deliver the same expected utility as the government's promised net payments.

The sum total of compensating differentials represents a risk-adjusted fiscal gap—the present value of what the government has promised to pay current and future generations in transfer payments net of what it has promised to take from them in taxes. In cases where the compensating differentials are positive, they can be expressed as the present values of net promises. Consequently, they can be used to tell us how to discount future government net payments.

We find that the appropriate fiscal discount rates to be applied to promises of sure future payments depend on the agent's age (which, as indicated, can be negative) and the state of the economy. They also depend on the size of the payments and whether the promises incorporate the general equilibrium effects of the promised payments. For infinitesimal payments, which have no general equilibrium feedback effects, the discount rates are, surprisingly, remarkably close to the economy's prevailing safe short-term rate of return. This finding provides some support for standard government practice of discounting firmly promised future benefits and taxes, such as those associated with the U.S. Social Security system, at a fixed rate.

*We thank Kenneth Judd, Stefania Garetto, Jianjun Miao, Alejandro Rivera, and the participants of the University of Chicago Theory Workshop for helpful comments. Hasanhodzic gratefully acknowledges the financial support of the NBER Pre-Doctoral Fellowship on the Economics of an Aging Workforce.

[†]Boston University Department of Economics, jah@bu.edu.

[‡]Boston University Department of Economics, kotlikoff@gmail.com.

1 Introduction

Assessing fiscal sustainability is tricky business. Governments commit to providing specific goods and services over time. They also promise to make transfer payments to current and future citizens. Some of these obligations are described as official. Others are labeled unofficial. But their economic costs don't depend on their titles. Instead they depend on their size, timing, and reliability. Thus, unofficial commitments can have larger present values than equally sized official commitments if the unofficial commitments are more likely to be paid.

Take U.S. federal debt service and Social Security payments. U.S. debt repayment to its worldwide holders is backed by lofty words, namely "the full faith and credit of the U.S. government." Social Security benefits, in contrast, are backed by the political power of tens of millions of current and near-term American retirees, most of whom routinely vote. Furthermore, Uncle Sam's debt service is a nominal dollar commitment, and Uncle Sam can print money at will. His Social Security payments, in contrast, are real (inflation-indexed) commitments.

The decision to label government obligations as unofficial, though economically arbitrary, is hardly innocent. It's used to keep liabilities off the book and, thereby, understate what governments truly owe. Indeed, by making specific obligations, like future pension payments or medical benefits, non-fungible, governments ensure that these I.O.U.s can't readily be priced.

How then can we determine the degree to which a government has promised more than it can deliver? Stated differently, how should we discount future government obligations (including negative obligations, i.e., taxes), no matter their title, when markets are incomplete? Relying on a social welfare function, i.e., on a social discount rate, is a frequent practice as recently evidenced by the Stern Report's (2006) valuation of the costs and benefits of climate-change abatement. However, the sustainability of government policy is a matter of economic feasibility, not ethics. And whether any self-proclaimed social planner, applying his/her own "social" discount rate, feels the expected path of government spending exceeds its cost will not illuminate whether the policy is affordable.

This paper's answer is to determine the current wealth equivalent of future government promises using a familiar metric—the change in remaining expected lifetime utility converted into current consumption units. Doing so is relatively straightforward for current generations. But what about future generations? Our approach is to treat future generations as if they are already alive, but simply have no utility in periods not just after, but also before their physical existence. Specifically, we treat a future generation that will materialize in x years as currently alive with an age of minus x. We also assume this age-minus-x generation discounts (applies its time preference rate to) utility when alive for another x years—i.e., by the number of years it will take for the generation to be born.

Making these calculations requires fully characterizing the economy's general equilibrium in the presence and absence of the government's promises. Our vehicle is a 10-period variant of Hasanhodzic and Kotlikoff's (2013) 80-period life-cycle model with aggregate risk. Like that paper, this study applies a numerical method originally proposed by Marcet (1988) and operationalized by Judd, Maliar, and Maliar (2009, 2011) to overcome the curse of dimensionality. Their means to ban the curse is to solve for decision functions only in states that fall within the economy's ergodic space; i.e., for states that the economy will frequent, not those it will never visit.

Our life-cycle model features standard Cobb-Douglas production in capital and exogenously supplied labor. There are random shocks to total factor productivity and either random shocks to capital's rate of depreciation or adjustment costs as in Hasanhodzic (2014). Preferences over the single commodity are time separable and isoelastic. There is also a one-period bond market. The model is intentionally highly stylized. Our goal is not to make realistic calculations, but to suggest, in the clearest and simplest setting, how one might make such calculations.

We proceed by positing and then solving our model in the absence of any government policy. Next we consider a sudden government promise at time zero to make ongoing payments to retirees. Finally we solve the model assuming, counterfactually, that the government has an external funder for these payments, named, well, God. The question then is how much the government would need to extract at time zero from each living and future generation to compensate it for not carrying through with its God-assisted policy. The sum total of all these compensation amounts represents the country's fiscal gap. It tells us precisely the total current resources the government is short when it comes to meeting, as in discharging, all its unfunded commitments.

Our God-assisted general equilibrium assumes that God makes her payments as they come due. Clearly, prepayment by God would produce a different dynamic stochastic equilibrium economy, different changes in lifetime expected lifetime utility, and a different assessment of what the government owes. This might suggest a shortcoming in our approach since it's inherently sensitive to the economy policy that the government promises, with God's help, to deliver. The sensitivity holds, but it's an intrinsic and inevitable feature of fiscal sustainability analysis. One simply cannot value government promises without knowing all they entail, including their general equilibrium effects.

In this regard, an interesting focal equilibrium is one in which the government is implicitly promising expected utility increases arising only from the provision of old age payments and not inclusive of any general equilibrium feedback effects. We calculate the value of these partial equilibrium promises as well.

When our calculated values of wealth-equivalent changes in expected utility are positive, they can be expressed as equaling the discounted present value of the payment promises, with the discount rate chosen to equate the present value of the promises to the compensating variation. The derived discount rates capture general equilibrium changes when such changes are assumed to be incorporated in the government's promises. But in partial equilibrium our implied discount rates lets us understand whether the short-term (one-period) bond rate comes close to the right discount rate for valuing longer term government promises.

We find that the discount rates for policies involving safe payments each period to the elderly aren't uniform over agents of different cohorts when a new general equilibrium is part of what the government is promising. In particular, discount rates decline with age for the young and increase with age for the elderly. The general equilibrium feedback effects are the key to these results. When agents anticipate future benefits, they save less and the capital stock declines. The resulting decline in wages hurts the young the most, lowering the value of the policy to them. If the general equilibrium effects are sufficiently adverse, young and future generations can be worse off. In this case, the compensating differential of the generation in question is negative and there is no discount rate that can be used to value future promised benefits, since these benefits, inclusive of their general equilibrium effects, have a negative present value.

Discount rates, when they are well defined, aren't uniform across states of the world either. Indeed, for each age group, the discount rates are higher in bad states of the world when the capital stock and productivity are low. In such states, the rate of return to capital is high, i.e., the price of future consumption is low. This induces agents to allocate more consumption into the future. Hence, the extra consumption, in the form of the benefit, does not matter as much to them.

When the benefit is small, the general equilibrium effects are less pronounced and there is less variation in discount factors across agents or states of the world. In fact, when the benefit is infinitesimal, the discounting is very close to the risk-free rate regardless of the age of the agent or the timing of the payment.

The paper is organized as follows. Sections 2 and 3 describe the model and its calibration. Section 4 presents results and Section 5 evaluates the accuracy of solutions. Section 6 concludes. The solution algorithm is relegated to the Appendix.

2 The Model

The model features G overlapping generations with shocks to total factor productivity and either capital depreciation shocks or capital adjustment costs. Each agent works full time through retirement age R, dies at age G, and maximizes expected lifetime utility. If there are adjustment costs, firms maximize their financial value, i.e., the present value of their revenue flow, otherwise they maximize static profits.

2.1 Endowments and Preferences

The economy is populated by G overlapping generations that live from age 1 to age G. All agents within a generation are identical and are referenced by their age g and time t. Each cohort of workers supplies 1 unit of labor each period. Hence, total labor supply equals the retirement age R. Utility is time-separable and isoelastic, with risk aversion coefficient γ:

$$u(c) = \frac{c^{1-\gamma} - 1}{1 - \gamma}. \tag{1}$$

2.2 Technology

Production is Cobb-Douglas with output Y_t given by

$$Y_t = z_t K_t^{\alpha} L_t^{1-\alpha}, \tag{2}$$

where z is total factor productivity, α is capital's share of output, and L_t is labor demand, which equals R, labor supply. Equilibrium factor prices are given by

$$w_t = z_t(1-\alpha)\left(\frac{\sum_{g=1}^{G} \theta_{g,t-1}}{R}\right)^{\alpha}, \tag{3}$$

$$r_t = z_t\alpha\left(\frac{\sum_{g=1}^{G} \theta_{g,t-1}}{R}\right)^{\alpha-1} - \delta_t, \tag{4}$$

where depreciation $\delta_t \sim \mathcal{N}(\mu_\delta, \sigma_\delta^2)$, as in Ambler and Paquet (1994).

With capital adjustment costs, r is given by

$$r_t = \frac{z_t \alpha (\frac{K_t}{R})^{\alpha-1} + \frac{m}{2}(\frac{I_t}{K_t})^2 + q_t - q_{t-1}}{q_{t-1}}, \tag{5}$$

where $q_t = 1 + m\frac{I_t}{K_t}$ is the price of capital, $K_t = \frac{\sum_{g=1}^{G} \theta_{g,t-1}}{q_{t-1}}$ is the capital stock, $I_t = K_{t+1} - K_t$ is the investment at time t, and m is the adjustment cost parameter (see, e.g., Auerbach and Kotlikoff, 1987 and the references therein).

Total factor productivity, z, obeys

$$\ln(z_{t+1}) = \rho \ln(z_t) + \epsilon_{t+1}, \tag{6}$$

where $\epsilon_{t+1} \sim \mathcal{N}(0, \sigma^2)$.

2.3 Financial Markets

Households save and invest in either risky capital or one-period safe bonds. Investing 1 unit of consumption in bonds at time t yields $1 + \bar{r}_t$ units in period $t + 1$. The safe rate of return, $\bar{r}_t$, is indexed by t since it is known at time t although it is received at time $t + 1$. The total demand for assets of household age g at time t is denoted by $\theta_{g,t}$, and its share of assets invested in bonds is denoted by $\alpha_{g,t}$. Households enter period t with $\theta_{g-1,t-1}$ in assets, which corresponds to the total assets they demanded the prior period. Since investment decisions are made at the end of the period, the aggregate supply of capital in period t, K_t, is the sum of assets brought by the households into period t, i.e.

$$K_t = \sum_{g=1}^{G} \theta_{g,t-1}, \tag{7}$$

normalized by q_{t-1} in the case of adjustment costs. Bonds are in zero net supply, hence by being short (long) bonds, households are borrowing (lending) to each other.

2.4 Government

In the model with government policy, each retiree receives a benefit equal to a fixed share ξ of either the average wage each period. The benefits are assumed to be externally financed.

2.5 Household Problem

Households of age g in state (s, z, δ) maximize expected remaining lifetime utility given by

$$V_g(s, z, \delta) = \max_{c,\theta,\alpha} \{u(c) + \beta \mathrm{E}\left[V_{g+1}(s', z', \delta')\right]\} \tag{8}$$

subject to

$$c_{1,t} = \ell_1 w_t - \theta_{1,t} + (1 - \ell_1)H_t, \tag{9}$$

$$\begin{aligned} c_{g,t} = {} & \ell_g w_t + \left[\alpha_{g-1,t-1}(1 + \bar{r}_{t-1}) + (1 - \alpha_{g-1,t-1})(1 + r_t)\right]\theta_{g-1,t-1} - \theta_{g,t} \\ & + (1 - \ell_g)H_t, \\ & \text{for } \; 1 < g < G, \; \text{ and} \end{aligned} \tag{10}$$

$$\begin{aligned} c_{G,t} = {} & \ell_G w_t + \left[\alpha_{G-1,t-1}(1 + \bar{r}_{t-1}) + (1 - \alpha_{G-1,t-1})(1 + r_t)\right]\theta_{G-1,t-1} \\ & + (1 - \ell_G)H_t, \end{aligned} \tag{11}$$

where $c_{g,t}$ is the consumption of a g-year old at time t and H_t is the benefit given to the elderly, and (9)–(11) are budget constraints for age group 1, those for age groups between 1 and G, and that for age group G.

2.6 Equilibrium

At time t, the economy's state is (s_t, z_t, δ_t), with $s_t = (\theta_{1,t-1}, \ldots, \theta_{G-1,t-1})$ denoting the set of age-specific asset holdings. Given the initial state of the economy s_0, z_0, δ_0, where $s_0 = (\theta_{1,-1}, \ldots, \theta_{G-1,-1})$, the recursive competitive equilibrium is defined as follows:

Definition. The recursive competitive equilibrium is governed by the collection of the value functions and the household policy functions for total savings $\theta_g(s, z, \delta)$, the share of savings invested in bonds $\alpha_g(s, z, \delta)$, and consumption $c_g(s, z, \delta)$ for each age group g, the choices for the representative firm $K(s, z, \delta)$ and $L(s, z, \delta)$, as well as the pricing functions $r(s, z, \delta)$, $w(s, z, \delta)$, and $\bar{r}(s, z, \delta)$ such that:

1. Given the pricing functions, the value functions (8) solve the recursive problem of the households subject to the budget constraints (9)–(11), and θ_g, α_g, and c_g are the associated policy functions for all g and for all dates and states.

2. Wages and rates of return on capital satisfy (3) and either (4) or (5), i.e. at each point, for given w and r the firm maximizes profits if there are no adjustment costs and maximizes the firm value otherwise.

3. All markets clear: Labor and capital market clearing conditions are implied by $L_t = R$ and (7). Since bonds are in zero net supply, bond market clearing requires

$$\sum_{g=1}^{G} \alpha_g(s, z, \delta)\theta_g(s, z, \delta) = 0. \tag{12}$$

   Market clearing conditions in labor, capital, and bond markets and satisfaction of household budgets imply market clearing in consumption.

Finally, for all age groups $g = 1, \ldots, G-1$, optimal intertemporal consumption and investment choice satisfies

$$1 = \beta \mathrm{E}_z \big[\big(1 + r(s', z', \delta')\big) \frac{u'(c_{g+1}(s', z', \delta'))}{u'(c_g(s, z, \delta))}\big] \tag{13}$$

and

$$0 = \mathrm{E}_z \left[u'(c_{g+1}(s', z', \delta'))(\bar{r}(s, z, \delta) - r(s', z', \delta')) \right], \tag{14}$$

where E_z is the conditional expectation of z' given z. Note that the endogenous part of the state next period, s', is determined by the asset demands chosen the period before.

3 Calibration

The parameters are calibrated as follows.

3.1 Endowments and Preferences

We set γ to 2. Agents work for 7 periods and live for 10. Hence, each period represents 6 years. We set the quarterly subjective discount factor, β, at 0.99, as is standard in the macroeconomics literature.

3.2 Technology

Quarterly values for ρ and σ are 0.95 and 0.01, respectively, as estimated in the empirical literature (see, e.g., Hansen (1985) or Prescott (1986)). Capital share of output, α, equals 0.33. In the model with stochastic depreciation, quarterly values for μ_δ and σ_δ are 0.0123 and 0.0026, respectively. In the model with adjustment costs, as well as in the base model without either stochastic depreciation or adjustment costs, depreciation is zero. The adjustment cost parameter m is set to 10.

3.3 Government

The wage share, ξ, equals 40 percent when the benefit is large, 10 or 20 percent when the benefit is medium-sized, and 2 percent when it is small.

4 Results

4.1 General Equilibrium Discount Factors

Age	Benefit Size: Medium		Small	
	State: Good	Bad	Good	Bad
	Annual Discount Rates (%)			
1	3.59	4.62	3.52	4.10
2	3.15	3.80	3.29	3.64
3	2.82	3.40	2.68	3.38
4	2.56	3.06	2.46	3.16
5	2.30	3.05	2.24	2.98
6	2.34	3.07	2.02	2.86
7	2.39	3.87	2.15	3.05
8	4.13	6.37	2.58	3.68
9	5.30	10.37	3.02	4.27
	Risk-free rate			
	3.06	4.39	3.25	4.41

Table 1: Annual discount factors for each cohort in good (high capital, high TFP shock) and in bad (low capital, low TFP shock) states of the world, for a medium-sized sure benefit (20 percent of the average wage) and a small one (2 percent of the average wage).

Table 1 shows annual discount factors for each cohort in good and bad states of the world, for different retirement benefit sizes. The compensating variations that form the basis of determining these discount rates are those needed to achieve the same remaining lifetime expected utility in the no-policy general equilibrium as one would enjoy in the general equilibrium arising under the externally funded policy.

In the good state the capital stock and the productivity shock are both high, while in the bad state they are both low. The medium and the small benefits amount to 20 percent and 2 percent of the average wage, respectively. Since the average wage is estimated in advance (from a previous run of the simulation), both benefits are sure.

Recall that to compute the discount factors we first solve our model in the absence of any government policy. Next we consider a sudden government promise at time zero to pay ongoing, externally financed benefits to retirees. Finally, for each generation, we find the amount by which we need to increase his time 0 assets to equalize his expected lifetime utility in the world without policy to that in the world with policy. From this we compute that generation's implied discount rate.

More precisely, let F_a denote the compensation an agent age $a > 0$ requires at time 0 when the policy is being evaluated in order to give up sure benefits in retirement (ages 8, 9, and 10). Then $F_a = \frac{\Delta \text{Expected Lifetime Utility}}{u'(c_a, 0)}$, where the change in expected lifetime utility is between the models without and with government policy. Dividing by $u'(c_a, 0)$ converts changes in utils into changes in the economy's single good.

For an agent who is age $a < 0$ at time 0 when the policy is being evaluated, the compensation $F_a = \frac{\Delta \text{Expected Lifetime Utility}}{\beta^{-a} E[u'(c_1, -a) \prod_s (1+r_s)]}$, where $s = 0, ..., -a$ and r_s is the random return on capital. For those age 8 or under, the implied discount rate δ is obtained by solving $F_a(0) = B/(1+\delta)^{8-a} + B/(1+\delta)^{9-a} + B/(1+\delta)^{10-a}$, where B is the amount of the sure benefit received in retirement (age 8, 9, and 10). For those age 9 and 10, there are 2 and 1 terms, respectively.

The table shows that the discount rates for policies involving safe payments each period to the elderly aren't uniform over agents of different cohorts. In particular, discount rates decline with age for the young and increase with age for the elderly. The general equilibrium effects are the key to these results. When agents anticipate future benefits, they save less and the capital stock declines. The resulting decline in wages hurts the young the most, lowering the value of the benefits to them. The further away the young are from the benefits, the more they are hurt. The older agents have fewer periods of higher rates of return on capital to enjoy, and the value of the benefit to them is smaller the older they are. Hence, the discount rate of the initial elderly age 9 exceeds the discount rate of the initial elderly age 8 and that rate is higher than the discount rate of those initially age 7. The medium size benefit is large enough to materially reduce the marginal utility of consumption of the elderly. This explains why we see very high discount rates for those initially age 9—rates that are much higher than the risk free rate.

Discount rates aren't uniform across states of the world either. Indeed, for each age group, the discount rates are higher in bad states of the world, where capital stock and productivity are low. In such states, the rate of return on capital is high, i.e., the price of future consumption is low. This induces agents to allocate more consumption into the future, in which case the extra consumption, in the form of the benefit, does not matter as much to them.

Table 1 also shows that when the benefit is smaller, there is less variation in discount factors across agents or states of the world. The reason is that with smaller benefits, the general equilibrium effects are less pronounced.

4.2 Partial Equilibrium Discount Factors

Tables 2 and 3 present discount factors in partial equilibrium, for different cohorts and different timing of the benefits. In Table 2 the policy is being evaluated in a typical state of the world, while in Table 3 it is evaluated in an atypical state, characterized by low capital stock and adverse productivity shock. The (row i, column j) cell of each table is the discount factor associated with the amount an i-year-old agent is willing to give up to get a sure benefit when he is i+j years old. These discount factors are computed as follows: Let F_a denote the amount an agent age $a > 0$ is willing to give up at time t to get a benefit at time $t + k$. Then $F_a = \frac{\beta^k E u'(c_{a+k}(t+k))}{u'(c_a(t))}$.

This calculation is capturing how much agents require today in order to give up a sure

		Periods Till Benefit Received								
		1	2	3	4	5	6	7	8	9
Current Age	1	3.20	3.24	3.27	3.30	3.33	3.35	3.36	3.38	3.39
	2	3.23	3.25	3.28	3.31	3.33	3.34	3.37	3.38	
	3	3.23	3.25	3.28	3.30	3.32	3.34	3.36		
	4	3.22	3.26	3.27	3.29	3.30	3.33			
	5	3.22	3.25	3.26	3.27	3.29				
	6	3.22	3.24	3.25	3.27					
	7	3.21	3.24	3.26						
	8	3.22	3.25							
	9	3.21								
		Risk-Free Rate								
		3.21								

Table 2: Annual discount factors for each cohort starting in a **typical state of the world**, when the benefit is sure but infinitesimal, so that there are no general equilibrium feedback effects. The (row i, column j) cell is the discount factor associated with the amount an i-year-old agent is willing to give up to get a sure benefit when he is i+j years old. The prevailing risk-free rate is 3.207 percent.

		Periods Till Benefit Received								
		1	2	3	4	5	6	7	8	9
Current Age	1	4.39	4.21	4.13	4.04	3.97	3.90	3.84	3.78	3.75
	2	4.41	4.23	4.14	4.04	3.97	3.90	3.82	3.77	
	3	4.52	4.31	4.19	4.08	4.01	3.91	3.84		
	4	4.57	4.35	4.21	4.11	4.01	3.92			
	5	4.58	4.35	4.21	4.07	3.99				
	6	4.54	4.35	4.10	4.00					
	7	4.76	4.35	4.14						
	8	4.42	4.22							
	9	4.27								
		Risk-Free Rate								
		4.57								

Table 3: Annual discount factors for each cohort starting in an **atypical (bad) state of the world**, when the benefit is sure but infinitesimal, so that there are no general equilibrium feedback effects. The (row i, column j) cell is the discount factor associated with the amount an i-year-old agent is willing to give up to get a sure benefit when he is i+j years old. The prevailing risk-free rate is 3.207 percent.

benefit in a specified period (not necessarily periods 8, 9, and 10) in the future, analogous to the general equilibrium calculation above. However, unlike in the general equilibrium case where the benefit is sizeable and received in each period in retirement, the partial equilibrium benefit is infinitesimal and received only at a given point in time. Hence, the derivative of expected utility in the period in question provides our measure of the change in expected utility. The implied discount rate δ is obtained by solving $F_a = 1/(1+\delta)^k$.

Table 2 shows that in partial equilibrium the discounting is very close to the risk-free rate of 3.207 percent in the typical state of the world. The results hold regardless of the age of the agent or the timing of the payment. Table 3 shows that discount rates can deviate from the risk-free rate even in partial equilibrium when the policy is being evaluated in an atypical state. In this case, at each cohort age, the larger is the number of periods till the benefit is received, the smaller is the discount factor, i.e., the more the benefits are valued by agents. This is intuitively what one might expect. Giving agents something safe further out in time when their consumption is less certain would be more highly valued than giving agents something safe sooner when there is less uncertainty. But, in addition, the interest rate can be expected to fall through time in Table 3 given we are starting the calculations out in a bad steady state. So the results here may be reflecting a decline in the expected term structure of returns. This Table 3 factor—that interest rates will likely fall—does not arise in Table 2 where interest rates are likely to rise.

4.3 Fiscal Gap

Table 4 shows the compensation amounts (F_a's) that the live and unborn cohorts require in general equilibrium. The sum total of these amounts is reported in the last column of the table as share of GDP. It represents the country's fiscal gap, i.e., the total current resources the government is short when it comes to meeting all its unfunded commitments. The results are shown for a good and a bad state of the world, as well as for sure benefits of different size. The small and large benefits amount to 10 and 40 percent of the average wage, respectively.

Benefit	State	Compensation Amounts (as Share of GDP)												Fiscal Gap (as Share of GDP)
		Current Age												
		-10	-9	...	-1	1	2	...	7	8	9	10		
Small	Bad	-0.01	-0.01	...	0.13	0.40	0.65	...	1.99	2.23	1.65	1.04		13
	Good	-0.01	-0.01	...	0.18	0.48	0.72	...	1.73	1.80	1.29	0.70		12
Large	Bad	-0.07	-0.08	...	0.24	1.15	2.09	...	5.87	5.40	4.54	2.71		38
	Good	-0.05	-0.07	...	0.49	1.63	2.65	...	6.02	5.74	4.43	2.53		42

Table 4: Present value of compensation amounts in total and as share of GDP in general equilibrium. Benefits are provided in periods 8, 9, and 10. Small benefit is 10 percent of the average wage. Large benefit is 40 percent of the average wage.

The fiscal gap is larger when the benefit is larger, as is intuitive. It also depends on the state of the world the economy is in: When the benefit is large, it is larger when the policy is being evaluated in the good state than in the bad state, and vice versa when the benefit is small.

The table also shows that for some unborn agents, the compensation amounts are negative. In other words, these agents are willing to pay to get from down under the fiscal burden. This makes sense, since the general equilibrium effects crowd out investment and reduce future wages. The transition path takes time and the agents born in the future suffer the most from this drop in wages. In this case, the discount factors are not well defined. This suggests that discounting at the risk-free rate for these agents is inappropriate, and that the approach taken here for valuing the fiscal gap may be more reasonable.

5 Accuracy of Solutions

A satisfactory solution requires the generation-specific Euler equations (13) hold out of sample, i.e., on a set draws for the shocks not used to compute the equilibrium decision rules. Hence, for each model considered, the accuracy of solutions is tested on a fresh sequence of z's and δ's that is 60 times longer than the 640-period sequence used in the original simulation. This test entails simulating the model forward on the new path of shocks, using the original asset demand functions, θ_g, and clearing the bond market in each period.[1] The out-of-sample deviations from full satisfaction of the Euler equations,

$$\epsilon(s,z,\delta) = \beta \mathrm{E}_z \big[\big(1 + r(s',z',\delta')\big) \frac{u'(c_{g+1}(s',z',\delta'))}{u'(c_g(s,z,\delta))} \big] - 1, \tag{15}$$

are computed for each period in the newly simulated time path and for each generation $g \in 1, \ldots, G-1$.[2] Finally, the average, across time, of the absolute value of the deviations from Euler equations is computed for each generation. The largest deviation is less than 1 percent.

6 Conclusion

The proper manner to discount government commitments when markets are incomplete and the general equilibrium is held fixed is a longstanding question in economics. Of particular interest is whether the prevailing short-term real interest rates are appropriate for discounting longer-term, government sure net payment promises—promises of future payments that are independent of the economy's state of nature. We find that the discount rates depend on the age of the agent, the state of the economy, and the size and riskiness of the government promises. A related question is what discount rate to use to value uncertain government promises, such as payments to the elderly that are dependent on the state of the economy. We leave this for future work.

Our method of assessing fiscal sustainability starts from a position of no existing policy and considers the costs arising from new government payment promises that have no funding source. But how could this method be used to assess fiscal sustainability in models, as well as actual economies, with existing policies? The answer is to value state-specific payments to

[1] For details of the solution method, including the bond market clearing algorithm, see the Appendix.

[2] The out-of-sample test does not apply to (14) since the inner loop is rerun, i.e. (14) will hold by construction.

agents that have no clear funding source. Take, for example, a large cohort that is scheduled to retire and receive payments from a much smaller cohort of young, but those payments, because of their size, aren't collectable. The size of the fiscal imbalance can be measured either in terms of what it would cost to a) immediately compensate the elderly for her promised benefits to them or b) endow young people with enough resources to compensate them for having to fund these payments through time. There is nothing requiring these two present values to be identical since with incomplete markets agents can differentially value a given payment being received and made.

In principle, our method of determining discount rates applicable to specific government payments, be they positive or negative (i.e., taxes) or safe or risky, could be used to improve fiscal gap accounting. Governments can use a realistic version of this model to value their obligations and quantify how much they would hurt different generations if they were to renege on their promises to them. Alternatively, they can understand what burdens they will inflict on generations they make pay these obligations.

A Computational Appendix

The model is solved using Hasanhodzic (2014), which is the first to incorporate stochastic depreciation and adjustment costs into the setting like this paper's. That algorithm, at the high level, closely follows that of Hasanhodzic and Kotlikoff (2013). However, it also resolves some low-level execution issues which arise because getting the model to converge is more challenging in the presence adjustment costs and stochastic depreciation. For completeness, the high-level structure is outlined below.

The algorithm consists of an inner loop and an outer loop. The outer loop solves for the asset demand functions of each age group by porting Judd, Maliar, and Maliar's (2009, 2011) generalized stochastic simulation algorithm (GSSA) to the OLG setting. It starts by making an initial guess of generation-specific asset demand functions θ_g as polynomials in the state variables. Next it draws a path of the shocks for T periods and runs the model forward over those periods using the guessed asset demand functions to compute the state variables in each period. Then, for each age group, g, it evaluates the Euler equation (13) to determine what age group g's asset demand should be in each period t. Finally, it regresses these time series of generation-specific asset demands on the state variables, and uses the regression estimates to update the corresponding polynomial coefficients. It repeats these steps using the same path of shocks until asset demand functions converge.

The inner loop is the extension of GSSA by Hasanhodzic and Kotlikoff (2013) that allows for the bond market. It consists of a binary search algorithm which determines the risk-free rate $\bar{r}$ that satisfies (12). In this binary search, the evaluation of the net bond demand is achieved by using another binary search to determine the unique bond shares that satisfy the first order conditions (14).

The following is the step-by-step description.

Initialization:

- Set $\bar{z} = 1$, $\bar{\delta} = \mu_\delta$, and solve for the nonstochastic steady state asset demands of each age group without bond, $\bar{s} = (\bar{s}_1, \ldots, \bar{s}_{G-1})$. Let $(s_0, z_0, \delta_0) = (\bar{s}, \bar{z}, \bar{\delta})$ be the starting point of the simulation.

- Approximate $G-1$ asset demand functions by polynomials in the state variables: $\theta_1(s, z, \delta) = \phi_1(s, z, \delta; b_1), \ldots, \theta_{G-1}(s, z, \delta) = \phi_{G-1}(s, z, \delta; b_{G-1})$, where $b_1, \ldots b_{G-1}$ are polynomial coefficients. We use degree 1 polynomials. To start iterations, we use the following initial guess for the coefficients: $b_1 = (0, 0.9, 0, \ldots, 0, 0.1\bar{s}_1, 0), \ldots, b_{G-1} = (0, 0, \ldots, 0, 0.9, 0.1\bar{s}_{G-1}, 0)$. Note that for all $g \in \{1, \ldots, G-1\}$, the initial b_g is such that $\bar{s}_g = \phi_g(\bar{s}, \bar{z}, \bar{\delta}; b_g)$.

Outer loop:

- Take draws of the path of z's and δ's for T years. We set T to 640.

- Simulate the model forward for $t = 0, \ldots, T$. More precisely, at time t, for each age group g, calculate its asset demand $\theta_g^{(p)}$ given the current guess for the coefficients $b_g^{(p)}$, where the subscript (p) denotes the current iteration of the outer loop. I.e., $\theta_{g,t}^{(p)}$ equals the inner product of the vector $(1, s_t, z_t, \delta_t)$ with the vector of coefficients $b_g^{(p)}$, where $s_t = (\theta_1^{(p)}(s_{t-1}, z_{t-1}, \delta_{t-1}), \ldots, \theta_{G-1}^{(p)}(s_{t-1}, z_{t-1}, \delta_{t-1}))$. Then the state at time $t+1$ and iteration p is given by $(s_{t+1}, z_{t+1}, \delta_{t+1}) = (\theta_1^{(p)}(s_t, z_t, \delta_t), \ldots, \theta_{G-1}^{(p)}(s_t, z_t, \delta_t), z_{t+1}, \delta_{t+1})$, where z_{t+1} given z_t is determined by (6).

- Inner loop:
  - Use binary search to solve (12) for $\bar{r}_t$, for all $t = 0, \ldots, T$. To start, make an (arbitrary) initial guess for the value of $\bar{r}_t$.
  - For all $t = 0, \ldots, T$, given $\bar{r}_t$, for all $g = 1, \ldots, G-1$, solve (14) for $\alpha_{g-1,t}$ using another binary search (evaluate the expectation in (14) using Gaussian quadrature).

- Use $\alpha_{g-1,t}$ found above for all g and for all t to calculate (12) and update $\bar{r}_t$ for all t.

- Note that for each age group g and each state (s_t, z_t, δ_t), $t = 1, \ldots, T$, (13) implies

$$\theta_g(s_t, z_t, \delta_t) = \beta \mathrm{E}_z \big[\theta_g(s_t, z_t, \delta_t)\big(1 + r(s_{t+1}, z_{t+1}, \delta_{t+1})\big) \frac{u'(c_{g+1}(s_{t+1}, z_{t+1}, \delta_{t+1}))}{u'(c_g(s_t, z_t, \delta_t))}\big] \quad \text{(A.1)}$$

for equilibrium asset demands θ_g. Denote the right-hand-side of (A.1) by $y_g(s_t, z_t, \delta_t)$ and evaluate the expectation using Gaussian quadrature.

- For each age group g, regress $y_g(s_t, z_t, \delta_t)$ on (s_t, z_t, δ_t) and a constant term using regularized least squares with Trikhonov regularization (see Judd, Maliar, and Maliar, 2011 for details). Denote the estimated regression coefficients by $\hat{b}_g^{(p)}$.

- Check for convergence: If

$$\frac{1}{G-1} \sum_{g=1}^{G-1} \frac{1}{T} \sum_{t=1}^{T} \left| \frac{\theta_g^{(p-1)}(s_t, z_t, \delta_t) - \theta_g^{(p)}(s_t, z_t, \delta_t)}{\theta_g^{(p-1)}(s_t, z_t, \delta_t)} \right| < \epsilon,$$

end. Otherwise, for each age group g update the coefficients as $b_g^{(p+1)} = (1-\xi) b_g^{(p)} + \xi \hat{b}_g^{(p)}$, for $\xi = 0.01$, and return to the beginning of the outer loop.

References

Ambler, S. and A. Paquet, 1994, "Stochastic Depreciation and the Business Cycle", *International Economic Review* 35, 101–116.

Auerbach, A. J. and L. J. Kotlikoff, 1987, *Dynamic Fiscal Policy*, Cambridge University Press.

Hansen, G. D., 1985, "Indivisible Labor and the Bussiness Cycle", *Journal of Monetary Economics* 16, 309–327.

Hasanhodzic, J., 2014, "Increasing Borrowing Costs and the Equity Premium", SSRN Working Paper 2425330, Social Science Research Network.

Hasanhodzic, J. and L. J. Kotlikoff, 2013, "Generational Risk—Is It a Big Deal: Simulating an 80-Period Model With Aggregate Shocks", NBER Working Papers 19179, National Bureau of Economic Research.

Judd, K., L. Maliar, and S. Maliar, 2009, "Numerically Stable Stochastic Simulation Approaches for Solving Dynamic Economic Models", NBER Working Papers 15296, National Bureau of Economic Research.

Judd, K., L. Maliar, and S. Maliar, 2011, "Numerically Stable Stochastic Simulation Approaches for Solving Dynamic Economic Models", *Quantitative Economics* 2, 173–210.

Marcet, A., 1988, "Solution of Nonlinear Models by Parameterizing Expectations", Technical report, Carnegie Mellon University.

Prescott, E. C., 1986, "Theory Ahead of Business-Cycle Measurement", in "Carnegie-Rochester Conference Series on Public Policy", volume 25, 11–44.

Treasury, U., 2006, *Stern Review on the Economics of Climate Change*, London: U.K. Treasury.