Volume Title: Tax Policy and the Economy, Volume 26

Volume Author/Editor: Jeffrey Brown, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-92345-2 (paper); 978-0-226-92344-4 (cloth); 978-0-226-92345-1 (paper)

Volume URL: http://www.nber.org/books/brow11-2

Conference Date: October 6, 2011

Publication Date: July 2012

Chapter Title: The Excess Burden of Government Indecision

Chapter Author(s): Francisco J. Gomes, Laurence J. Kotlikoff, Luis M. Viceira

Chapter URL: http://www.nber.org/chapters/c12562

5

The Excess Burden of Government Indecision

Francisco J. Gomes, *London Business School and CEPR*
Laurence J. Kotlikoff, *Boston University and NBER*
Luis M. Viceira, *Harvard University and NBER*

Executive Summary

Governments are known for procrastinating when it comes to resolving painful policy problems. Whatever the political motives for waiting to decide, procrastination distorts economic decisions relative to what would arise with early policy resolution. In so doing, it engenders excess burden. This paper posits, calibrates, and simulates a life cycle model with earnings, life span, investment return, and future policy uncertainty. It then measures the excess burden from delayed resolution of policy uncertainty. The first uncertain policy we consider concerns the level of future Social Security benefits. Specifically, we examine how an agent would respond to learning in advance whether she will experience a major Social Security benefit cut starting at age 65. We show that having to wait to learn materially affects consumption, saving, labor supply, and portfolio decisions. It also reduces welfare. Indeed, we show that the excess burden of government indecision can, in this instance, range as high as 0.6% of the agent's economic resources. This is a significant distortion in and of itself. It is also significant when compared to other distortions measured in the literature. The second uncertain policy we consider concerns marginal tax rates. We obtain similar results once we adjust for the impact of tax rates on income.

I. Introduction

Virtually all US policy makers, budget analysts, and academic experts agree that the United States faces a very serious, if not a grave, long-term fiscal problem. Yet few policy makers want to be the bearer of bad news and say how or when they would fix it. Delaying the resolution of fiscal imbalances comes at two costs. First, it leaves a larger bill for a smaller number of people to pay. Second, and of primary interest here, it perpetuates uncertainty, leading economic agents to make saving, investment, labor supply, and other decisions that are suboptimal from an ex post perspective. Take, as an example, the prospect facing the baby boomers of having their Social Security retirement benefits cut.

The likelihood of this outcome may be leading them to save more, work more, and invest in safer assets than would otherwise be true.

Whatever are the political gains to government indecision and whatever are the decisions being deferred, it is clear that delays in policy making distort economic choices and, as such, engender excess burden. This paper provides some sense of the magnitude of this excess burden. Specifically, we posit, calibrate, and simulate a realistic life cycle model featuring optimal consumption, portfolio choice, and labor supply decisions in the face of uncertainty in earnings, life span, investment returns, and government policy. We then measure the welfare gain of early resolution of policy uncertainty. The size of this gain is also the size of the excess burden associated with delayed policy resolution.

Our life cycle model builds on the literature. As in Cocco, Gomes, and Maenhout (2005), it includes investment in risky stocks and safe bonds. As in Gomes and Michaelides (2005), it incorporates housing as well as age-related expenditures. And, as in Gomes, Kotlikoff, and Viceira (2008), it incorporates (as an extension) variable labor supply.

Our model is by necessity stylized. However, our goal is not to arrive at a precise understanding of the magnitude of US economic inefficiency arising from policy indecision. Doing so would require considering all policies that might be changed in the future as well as the impact of policy uncertainty on all types of American households. Our goal is simply to understand whether policy indecision could generate a reasonably large excess burden for typical middle-class households.

The main indecision we consider involves the level of future Social Security benefits. Specifically, we examine how agents respond to learning prior to age 65 whether or not they will experience a major Social Security benefit cut starting at age 65. We show that having to wait to learn materially affects behavior. Most important, it reduces welfare. Indeed, the excess burden of government indecision, in this instance, can exceed more than 0.5% of agents' resources. Note that we are comparing two scenarios with the same expected Social Security income and the same ex ante uncertainty. Hence, this welfare loss stems from neither changes in agents' expected incomes nor income risks. Rather it comes exclusively from delay in resolving policy uncertainty, that is, from government's indecision. Excess burdens in the range of 0.5% of resources, measured as the prereform certainty-equivalent level of consumption, are significant in and of themselves.[1] They are also significant in comparison to other distortions measured in the public finance literature such as those arising from maintaining an inefficient tax structure (see Auerbach and Kotlikoff 1987).

We also consider delays in determining/announcing future tax rate changes. Scaled appropriately, the quantitative findings here are similar to those with respect to Social Security benefit indecision. So too are the welfare costs of joint indecision over benefits and taxes. Each of our calculated indecision costs is highly sensitive to the degree of risk aversion, the number of years of indecision, and the size and probability of policy changes.

Our study appears to be the first to identify and measure the cost of this form of government risk spreading. Previous work has examined the impact on consumption and saving of early resolution of uncertainty (e.g., Blundell and Stoker 1999; Eeckhoudt, Gollier, and Treich 2001). Blundell and Stoker point out that consumption growth will be faster when young if earnings uncertainty occurs primarily when young rather than when old. And Eeckhoudt et al. examine how earlier resolution of uncertainty affects the current level of consumption. But neither study examines how delay in government decision making affects levels of expected utility in a multiperiod, realistically parameterized model. Pastor and Veronesi (2011) study the impact of uncertainty about government decisions, regarding both their implementation and their impact on the economy, on asset prices. Croce, Nguyen, and Schmid (2011) show that this form of uncertainty can generate nontrivial risk premia in the economy. Luttmer and Samwick (2011) document survey evidence that households expect significant cuts in future Social Security benefits but with considerable uncertainty and estimate a willingness to forgo 4%–6% of promised benefits in order to remove the political uncertainty associated with future benefits. Judd (1989), Chari, Christiano, and Kehoe (1994), Diamond (1997), Bohn (1998), Auerbach and Hassett (2002), and others study the manner in which governments should optimally spread risk across generations. But this is different from our question, which is not how to spread risk, but when to resolve it.

Before proceeding, we reference several issues omitted from our analysis. To our minds, these issues are not germane to our specific focus, namely, the distortion arising from the delay in making and announcing policy decisions. The first such issue involves the political process leading to indecision. There are, no doubt, many different explanations for why politicians do not make timely decisions, just as there are many explanations for why politicians fail to enact the least distortionary set of taxes.[2] But our focus is not on the reasons for indecision, but rather on its costs. Any specific model of why governments delay making decisions would pose two questions, namely, why we chose that particular formation and why the political process we considered would ever

change its behavior. If government policy cannot be changed, there is no scope for efficiency gains and, thus, no excess burden to be calculated. This is no less true of standard analyses of excess burden, such as the burden of distortionary taxation. As scores of excess burden studies confirm, modeling government decision making is not a prerequisite to measuring the excess burden of distortionary taxation. Indeed, precise modeling of the politics leading to a particular configuration of distortionary taxes presumes away the ability to eliminate distortions using alternative taxes.

The second issue is that policy changes, no matter when announced, may affect more than one generation as the government proceeds to satisfy its intertemporal budget. For example, cutting baby boomers' future Social Security benefits might be associated with lowering the payroll taxes of generations born in the future, say after 2030. While this is true, the timing of when the boomers' benefit cuts are determined and announced is a distinct question that can be analyzed separately from the question of who else will be affected by such cuts. Moreover, the government's decision to announce sooner rather than later how it will treat a given generation need not affect the timing of policy announcements made to other generations. For example, the government's decision whether to tell boomers in 2007 or in 2015 about benefit cuts that will occur in 2020 will not affect the uncertainty about this policy experienced by generations born after 2020 since the policy will be resolved by the time they are born.

A third issue is whether we need to consider how policy indecision affects aggregate capital formation and the evolution of wage and interest rates. The answer is no. Any proper excess burden calculation requires that one fully compensate all agents for all first-order income/incidence effects, including those arising from general equilibrium factor price changes. Were we to specify a general equilibrium model, we would have to introduce a compensation policy to undo all the general equilibrium incidence in order to measure excess burden.[3]

Our paper proceeds as follows. Section II uses a very simple model to show how government indecision reduces welfare. Sections III and IV introduce and calibrate our life cycle model assuming, initially, fixed labor supply. Section V shows how saving and investment are affected by delay in retirement policy resolution and reports the excess burden arising from waiting to resolve this uncertainty. Section VI repeats these analyses except it assumes that uncertainty involves future tax rates rather than retirement benefit levels. Section VII jointly considers uncertainty about future retirement income and tax rates and calibrates the benefits of joint early resolution of these uncertainties. Section VIII adds

variable labor supply, showing that this addition makes no significant difference to the results. Finally, Section IX presents a summary and conclusions.

II.   Understanding the Welfare Costs of Policy Delay

Consider an agent who lives between time 0 and time T and has initial assets A_0 and preferences described by a power utility function of instantaneous consumption. There is one riskless investment instrument. The agent's time preference rate and riskless rate of return are both zero. There are no borrowing constraints. The agent learns at time $L \le T$ whether she gets high or low benefits, B, per period in retirement. The receipt of these benefits begins at time R.

We solve this model via backward recursion starting from the point where the agent learns the size of her future benefits. Let A_L be assets accumulated by the agent at the time of the announcement of future benefits. Since there is no uncertainty about future benefits after L and the time preference and interest rates are equal, the agent's consumption is constant between times L and T. Wealth is simply assets plus the present value of retirement benefits, $A_L + B(T - R)$.

Therefore, optimal consumption is given by

$$A_L + \overline{B}(T - R) = \overline{C}(T - L) \tag{1}$$

when retirement benefits $B = \overline{B}$ and

$$A_L + \underline{B}(T - R) = \underline{C}(T - L) \tag{2}$$

when retirement benefits $B = \underline{B}$.

We now solve for C, optimal consumption before time L. Given our assumptions about the interest rate and the time preference rate, the agent will choose a constant level of consumption C prior to learning the outcome of B. The agent sets C to maximize expected utility at time 0, which is given by

$$EU = \frac{C^{1-\gamma}}{1-\gamma}L + (T - L)\left[p\frac{\overline{C}^{1-\gamma}}{1-\gamma} + (1-p)\frac{\underline{C}^{1-\gamma}}{1-\gamma} \right], \tag{3}$$

where p is the probability of a high benefit. This maximization is subject to the constraint that assets at time L satisfy

$$A_L = A_0 - CL. \tag{4}$$

Substitution of the budget constraint (4), (1), and (2) into (3) gives

$$
EU = \frac{C^{1-\gamma}}{1-\gamma}L + (T-L)\left(p\frac{\{[A_0 - CL + \overline{B}(T-R)]/(T-L)\}^{1-\gamma}}{1-\gamma}\right.
$$

$$
\left. + (1-p)\frac{\{[A_0 - CL + \underline{B}(T-R)]/(T-L)\}^{1-\gamma}}{1-\gamma}\right).
$$

The first-order condition is

$$
C^{-\gamma} = p\overline{C}^{-\gamma} + (1-p)\underline{C}^{-\gamma}. \tag{5}
$$

Equations (1), (2), (4), and (5) determine optimal consumption C between 0 and L.

The derivative of expected utility with respect to L is given by

$$
\frac{\partial EU}{\partial L} = \frac{\gamma}{1-\gamma}[C^{1-\gamma} - p\overline{C}^{1-\gamma} + (1-p)\underline{C}^{1-\gamma}] < 0 \quad \text{for all } \gamma. \tag{6}
$$

That is, early resolution of uncertainty about the future value of B is unambiguously welfare improving.

To see this, insert (5) in (6). This yields

$$
\frac{\partial EU}{\partial L} = \frac{\gamma}{1-\gamma}\{[p\overline{C}^{-\gamma} + (1-p)\underline{C}^{-\gamma}]^{1-1/\gamma} - p\overline{C}^{1-\gamma} + (1-p)\underline{C}^{1-\gamma}\}. \tag{7}
$$

The expression in brackets is a function of the form $f(x) = x^{1-1/\gamma}$, whose second derivative is negative for $\gamma > 1$ and positive for $\gamma < 1$.[4] Therefore, by Jensen's inequality we have that $f(E(z)) > Ef(z)$ for $\gamma > 1$ and $f(E(z)) < Ef(z)$ for $\gamma < 1$. A direct application of this result to the term in braces in equation (7), with $E(z) \equiv p\overline{C}^{-\gamma} + (1-p)\underline{C}^{-\gamma}$, implies that this term is negative when $\gamma < 1$ and positive when $\gamma > 1$, thus implying $\partial EU/\partial L < 0$.

It remains to show that $\partial EU/\partial L < 0$ in the special case $\gamma = 1$. In that case, expected utility is given by

$$
EU = L \log C + (T-L)[p \log \overline{C} + (1-p) \log \underline{C}],
$$

and the first-order condition for consumption is

$$
C = p\overline{C}^{-1} + (1-p)\underline{C}^{-1},
$$

which implies that

$$\frac{\partial EU}{\partial L} = -\log\left[p\overline{C}^{-1} + (1-p)\underline{C}^{-1}\right] - \left[p\log\overline{C} + (1-p)\log\underline{C}\right].$$

Note that $-\log(x)$ is a convex function for which $f(E(z)) < Ef(z)$. Therefore, we have

$$-\log\left[p\overline{C}^{-1} + (1-p)\underline{C}^{-1}\right] < p\log\overline{C}^{-1} + (1-p)\log\underline{C}^{-1}$$

$$< -\left[p\log\overline{C} + (1-p)\log\underline{C}\right],$$

which implies that $\partial EU/\partial L < 0$ when $\gamma = 1$.

Clearly, the sooner an agent learns about her future benefits, the sooner she can make the consumption and saving decisions appropriate to that information. The longer she is forced to wait, the longer she must consume and save defensively, thereby making more ex post mistakes. Understanding the economic costs of these mistakes requires a more realistic framework, to which we now turn.

III. A Life Cycle Model of Policy Delay

We now construct a more realistic model to study the costs of policy delay and behavioral impacts of policy delay. We begin with two economic behaviors—consumption choice and asset allocation—taking labor supply as fixed. We then allow labor supply to respond to the economic environment, showing that this response has only minor effects on our results.

In making their decisions, households face borrowing constraints and have to meet "off-the-top" housing expenses. They also experience changes over time in household size due to the initial presence and subsequent absence of children. But they are mainly concerned with the following four types of uncertainty: earnings, longevity, returns, and government policy. Policy uncertainty involves the future levels of retirement benefits and labor income taxes. This uncertainty is resolved either prior to or at retirement. By varying the timing of when agents learn about their future government treatment, we can see the degree to which early policy resolution matters.

A. Model Specification

1. Time Parameters and Preferences

Let t denote age and assume that agents work their first K periods and live for a maximum of T periods. Let p_t denote the probability that the

investor is alive at date $t + 1$, conditional on being alive at date t. Preferences over consumption and leisure are given by

$$\mathrm{E}_1 \sum_{t=1}^{T} \delta^{t-1} D_t \left(\prod_{j=0}^{t-2} p_j \right) \frac{C_t^{1-\gamma}}{1 - \gamma}, \tag{8}$$

where $\delta < 1$ is the discount factor, C_t is time t consumption, L_t is time t leisure, $\gamma > 0$ is the coefficient of relative risk aversion, and

$$D_t = \begin{cases} 1 & \text{for } t < 18 \\ \overline{D} & \text{for } t \geq 18. \end{cases} \tag{9}$$

The term D_t captures the change in household size when adult children leave the household. We calibrate $\overline{D}$ to produce a 30% drop in household consumption at age 45 ($t = 18$). As can be seen in Section VIII, this preference structure can be described as incorporating leisure whose value at each point in time is exogenously set to one.

2.   The Labor Income Process and Retirement Income

With labor supply fixed, age t labor income Y_t ($t \leq K$) is exogenous. Its logarithm is taken to equal the sum of a deterministic component—which we calibrate to capture the hump shape of earnings over the life cycle—and two random components, one transitory and one permanent.[5] More precisely,

$$\ln(Y_t) = f(t) + v_t + \varepsilon_t \quad \text{for } t \leq K. \tag{10}$$

The function $f(t)$ controls for the age-earnings profile. The term v_t is a permanent component given by

$$v_t = v_{t-1} + u_t, \tag{11}$$

where u_t is distributed as $N(0, \sigma_u^2)$, and ε_t is a transitory shock uncorrelated with u_t, which is distributed as

$$\begin{cases} N(0, \sigma_\varepsilon^2) & \text{with probability } 1 - \pi \\ \ln(0.1) & \text{with probability } \pi \end{cases} \tag{12}$$

and captures the probability of a large negative income shock as in Deaton (1991), Carroll (1992), and Heaton and Lucas (1997).

Government retirement benefits equal a fraction λ of permanent labor income in the last working year. Specifically,

$$\log(Y_t) = \log(\lambda) + f(K) + v_K \quad \text{for } t > K. \tag{13}$$

Making retirement benefits proportional to final earnings eases computation by reducing the number of state variables by one. Section V specifies the process governing λ.

Ignoring the transitory income shock, ε_t, setting benefits at λ times preretirement earnings is equivalent to setting benefits at λ times the sum of all earnings realizations from age t_0 until age K. To see this, note that v_t follows a random walk. Hence, labor income in the period before retirement equals $\sum_{t=t_0}^{K} u_t$. One can also describe retirement benefits as proportional to average lifetime income by multiplying λ by $K - t_0$ and dividing lifetime income by $K - t_0$.

3. Financial Assets

There are two assets—one risky and one riskless. *Bonds*, the riskless asset, are held in quantity B_t and yield a fixed gross real return of $\overline{R}_f$. *Stocks*, the risky asset, are held in quantity S_t and yield a gross real return R_t given by

$$\ln(R_t) \sim N(\mu + \overline{r}_f, \sigma_R^2), \tag{14}$$

where $\overline{r}_f = \ln(\overline{R}_f)$, μ denotes the expected log return on stocks in excess of the log return on bonds, and σ_R^2 denotes the volatility of log stock returns. Stock returns are correlated with innovations to the shock to permanent labor income (u_t) with correlation coefficient ρ.

Agents face the following borrowing and short-sales constraints:

$$B_t \geq 0, \tag{15}$$

$$S_t \geq 0. \tag{16}$$

Let α_t denote the proportion of assets invested in stocks at time t. Then these constraints imply that $\alpha_t \in [0, 1]$ and that wealth is nonnegative.

4. Taxes

We assume proportional taxes for all sources of income to preserve the scalability/homogeneity of the model. We assume that labor income is

taxed at a rate τ_L, that retirement income is taxed at a rate τ_R, and that asset income is taxed at a rate τ_C. As discussed in Section IV.B, we calibrate these tax rates to match the effective income tax rates currently faced by typical US households.

B.   The Investor's Problem

The investor starts the period with wealth W_t. Then labor income Y_t is realized. The purchase of durable goods and homes, referenced as *housing expenditures*, constitutes off-the-top spending. Although endogenizing durable purchases is beyond the scope of the paper, including this spending on an exogenous basis is important because it affects the likelihood that borrowing constraints will bind. We model the percentage of household income that is dedicated to housing expenditures (h_t) as an exogenous process and subtract it from disposable income.

Following Deaton (1991), period t *cash on hand*, X_t, obeys

$$X_t = W_t + (1 - h_t)(1 - \tau)Y_t,$$

where $\tau = \tau_L$ during working life and $\tau = \tau_C$ during retirement.

The investor must decide how much to consume, C_t, and how to allocate the remaining cash on hand (savings) between stocks and bonds. Next-period wealth, before earning period $t + 1$'s labor income, is given by

$$W_{t+1} = R^p_{t+1}[W_t + (1 - h_t)(1 - \tau)Y_t - C_t], \tag{17}$$

where R^p_{t+1} is the net return on the portfolio held from period t to period $t + 1$:

$$R^p_{t+1} \equiv 1 + (1 - \tau_C)[\alpha_t R_{t+1} + (1 - \alpha_t)\overline{R}_f - 1]. \tag{18}$$

The control variables are $\{C_t, \alpha_t\}^T_{t=1}$. The state variables are $\{t, X_t, v_t\}^T_{t=1}$. Given the setup, the value function is homogeneous with respect to current permanent labor income. This scalability lets us normalize v_t to one and reduce by one the dimensionality of the state space.

The Bellman equation for this problem, which we solve numerically via backward induction, is

$$V_t(X_t) = \max_{C_t \geq 0, 0 \leq \alpha_t \leq 1} [U(C_t) + \delta p_t E_t V_{t+1}(X_{t+1})] \quad \text{for } t < T, \tag{19}$$

where

$$X_{t+1} = (1 - h_t)(1 - \tau)Y_{t+1} + (X_t - C_t)\{1 + (1 - \tau_C)[\alpha_t R_{t+1} \\ + (1 - \alpha_t)\overline{R}_f - 1]\}.$$

IV.   Model Calibration

A.   Labor Income Process

The labor income profile is taken from Cocco et al. (2005). The authors estimate age profiles for three different education groups (households without a high school education, households with a high school education but without a college degree, and college graduates); we take the weighted average of the three. Cocco et al. also estimate the fraction of permanent income replaced by retirement income λ to be 83%. In Sections V and VII, where λ is uncertain, we set the probability distribution of λ so that its median is equal to 0.8. In Section VI, where there is no uncertainty in future retirement income, we fix λ at a constant 0.8.

We set the probability of a large negative income shock at 2.0%. Following Heaton and Lucas (1997), we set the magnitude of the shock at 10% of the household's expected income. The values of σ_u and σ_ε are 10.95% and 13.89%, respectively.[6] Finally, we set the correlation between stock returns and innovations in the permanent component of income (ρ) equal to 0.15 (Campbell et al. 2001), and we assume the same housing expenditure profile ($\{h_t\}_{t=1}^{T}$) as in Gomes and Michaelides (2005).

B.   Other Parameters

Agents are initially age 28, retire at 65, and die with probability one at age 100. Prior to 100 we use the mortality tables of the National Center for Health Statistics to set conditional survival probabilities, p_j, for $j = 1, \ldots, T$. We chose age 28 as the initial age of adulthood to roughly match the age at which working Americans marry and start having children.

We set the discount factor δ to 0.95 and the coefficient of relative risk aversion γ to 5. This baseline choice of γ is often used in the life cycle investment literature and lies well within upper-bound estimates (Mehra and Prescott 1985). Nevertheless, we routinely include sensitivity analysis with respect to this parameter.[7] Section VIII relaxes the assumption of fixed labor supply and discusses the calibration of θ.

The mean equity premium (in levels) is set at 4.00% per year, the risk-free rate is set at 1.00% per year, and the annualized standard deviation of innovations to the risky asset is set at 20.5%. This equity premium is lower than the historical equity premium on the basis of a comparison of average stock and Treasury-bill returns, but it is in line with the forward-looking estimates reported in Fama and French (2002). Also, a higher premium generates unrealistically high equity portfolio shares.

Finally, in the baseline case we set the tax rate on labor income (τ_L) to 30% and the tax rate on retirement income (τ_R) to 15%. Asset income is taxed at a 20% rate (τ_C). These rates roughly match the effective income tax rates currently faced by a typical household (Kotlikoff and Rapson 2005). In Section VI we incorporate uncertainty in these rates.

V.   The Benefits of Early Resolution of Uncertainty about Retirement Income

We start our analysis of the excess burden of policy delay by examining the costs of government indecision over retirement benefits.

A.   Retirement Income Uncertainty

Equation (13) shows that retirement income equals the product of a fixed replacement ratio (λ) and permanent income at age 65. We assume throughout this section that at age 28 the household does not yet know the value of λ. She knows only that the realization of λ is governed by the following distribution:

$$\lambda = \begin{cases} \overline{\lambda} & \text{with probability } 1 - p \\ \overline{\lambda}(1 - \xi) & \text{with probability } p, \end{cases} \tag{20}$$

where ξ represents a potential percentage cut in the (expected) replacement ratio.

In the baseline calibration we set the probability p of a cut in the replacement ratio at one-third and the magnitude ξ of this cut at 0.3. With this choice of parameter values, $\overline{\lambda}$ represents the median replacement ratio; the mean is about $0.9 \cdot \overline{\lambda}$. As mentioned, $\overline{\lambda}$ equals 0.8, which is close to the 83% value in Cocco et al. (2005). We assume that this uncertainty continues until age A (with $A \leq K$, naturally), when the exact value of the replacement ratio is revealed. The relevant variables determining the uncertainty in retirement benefits will then be ξ (the level of uncertainty) and A (inversely related to the duration of uncertainty).

Our consideration of a 30% future cut in US Social Security retirement benefits seems plausible given the magnitude of the current-value budgetary shortfall facing Social Security as reported in the 2011 Social Security Trustees Report (http://www.ssa.gov/oact/tr/2011/). Indeed, this report's table BIV6 indicates that benefits would need to be cut immediately by 22% under what appear to be relative optimistic "immediate" assumptions to achieve present value budget balance. A delay in dealing with Social Security's fiscal problem could well result in a 30% benefit cut, particularly if the Social Security actuaries' "high-cost" assumptions materialize.

In fact, we have seen benefit cuts of this magnitude enacted in the past. Munnell (2003) points out that the combination of non-inflation-indexed adjusted gross income thresholds for determining the income taxation of Social Security benefits, the ongoing rise in the age of normal retirement, and the increase in Medicare Part B premium payments (which are subtracted from Social Security benefits) entails benefit cuts of 26% for median-income workers retiring in 2030. For upper-income workers, these cuts, which were, in part, promulgated as far back as 1981, appear to range as high as one-third. Shoven and Slavov (2006) also document major changes in Social Security benefits, which they ascribe to "political risk."

Although we view a one-third chance of a 30% future cut to Social Security benefits as eminently reasonable, we recognize that cuts of different magnitudes occurring with different probabilities are also feasible. To accommodate these alternative possibilities, we show results for a range of benefit cuts. We do not report results for different probabilities of a benefit cut because they are very similar; that is, a higher chance of a cut of a given size or the same chance of a larger-sized cut yields very similar results. Again, our objective is not to offer a precise estimate of the excess burden from government indecision about Social Security policy. We simply seek to understand the potential welfare costs of government indecision.

A final point is worth making with respect to the potential for Social Security benefit cuts. As Persson and Svensson (1989) point out, Social Security is, of course, just one part of a comprehensive fiscal policy. Indeed, as discussed in Kotlikoff (2002) and Green and Kotlikoff (2006), what benefits and taxes one labels as "Social Security" are economically arbitrary. From this perspective, the probability and magnitude of Social Security benefit cuts (or what are labeled Social Security benefit cuts) ultimately depend on the need for fiscal adjustments to satisfy the government's intertemporal budget constraint. As Gokhale and Smetters (2005)

demonstrate, the US fiscal gap is stunningly large, raising the possibility that Social Security benefit cuts will be needed to fix more than just what is labeled as Social Security.

B.   Optimal Consumption and Portfolio Choice with Late Resolution of Uncertainty about Retirement Income

Our baseline model assumes that uncertainty about retirement income—or, more precisely, the fraction λ of permanent income to be replaced in retirement—is not resolved until the household retires at age 65. That is, we set $A = 65$. Our choice of age 65 is illustrative, but it can be disputed. But partial price indexing of benefits is an example of a policy that would affect actual beneficiaries were it enacted.

Figure 1 shows the life cycle pattern of financial wealth, income, and optimal consumption generated by the baseline model. The units are thousands of 1992 dollars—the year to which Cocco et al. (2005) calibrate their income profiles. Figure 2 shows the percentage portfolio allocation to stocks. More precisely, these figures plot the average life cycle profile of wealth, income, consumption, and portfolio allocations based on 10,000 simulations of the model.

Figures 1 and 2 show that household consumption, wealth accumulation, and portfolio allocations exhibit a hump-shaped pattern with two humps. The consumption profile early in the life cycle is typical for a liquidity-constrained investor. Optimal consumption grows until age 45, when it falls sharply as adult children leave the household. Through roughly age 40, consumption remains below labor earnings as the household saves a small fraction of its income for precautionary reasons. Figure 2 shows that this asset accumulation is also associated with an increasing allocation to stocks until about age 35. By age 40, accumulated assets exceed annual consumption.

The pronounced, short-lived decline in consumption at age 45 reflects a one-time reduction in the size of the household. Consumption starts growing again after this event and keeps doing so until retirement at age 65. However, the household chooses not to consume all available income between age 45 and age 58. Instead, it chooses to save for retirement in the context of uncertainty about retirement benefits, which will not be resolved until the household retires at age 65. This saving, coupled with an aggressive allocation to stocks, allows the household, on average, to accumulate assets rapidly.

Interestingly, figure 2 shows that the allocation to stocks reaches a peak of about 98% at age 47 and declines until age 65, when the household
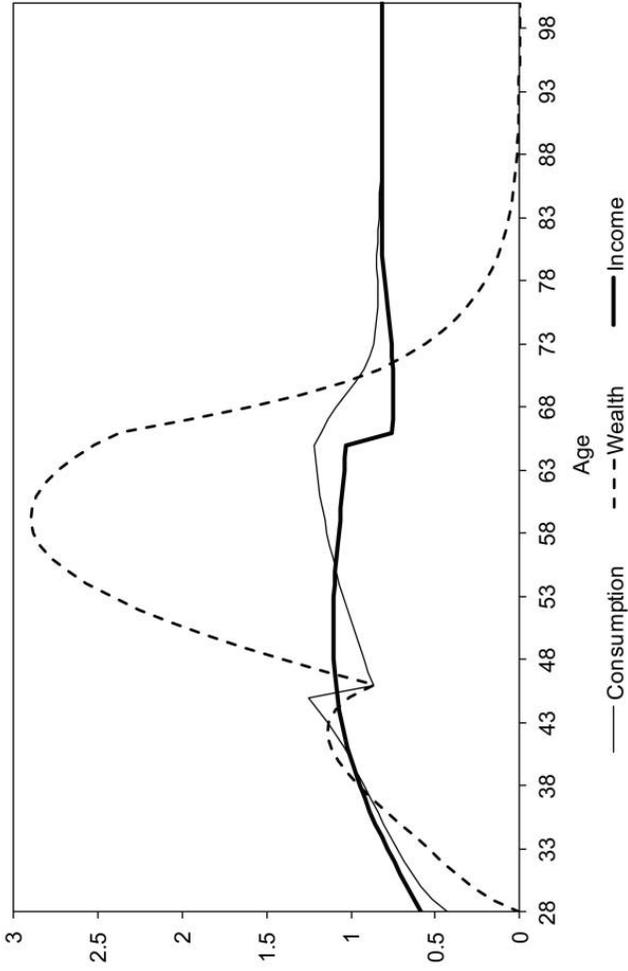
Fig. 1. Consumption, wealth, and income. The figure shows the life cycle pattern of financial wealth, income, and optimal consumption generated by the baseline model, whose parameter values are given in Section IV. The units in this figure are thousands of 1992 dollars—the year to which Cocco et al. (2005) calibrate their income profiles. The patterns are averages based on 10,000 simulations of the baseline model.
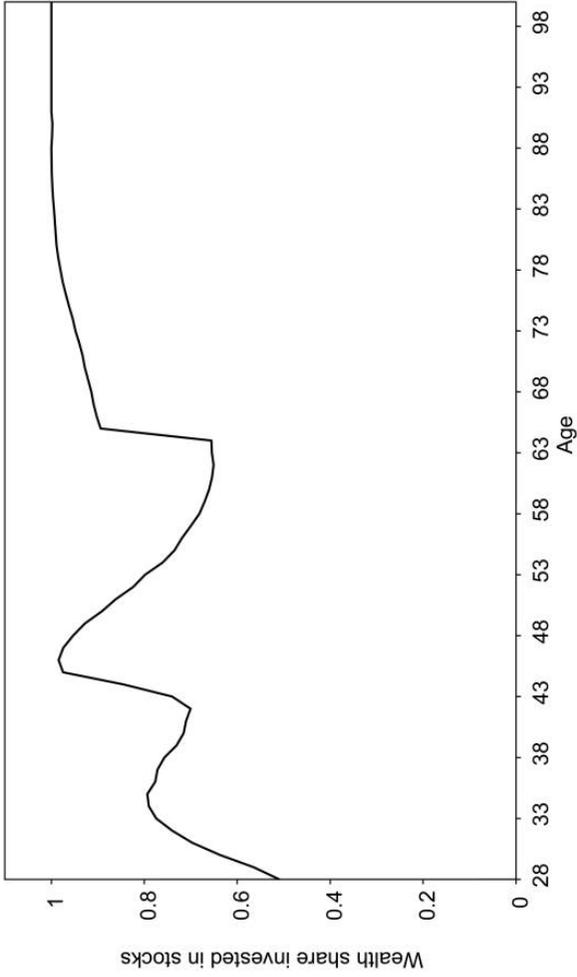
Fig. 2. Portfolio allocation. The figure shows the life cycle pattern of the percentage portfolio allocation to stocks generated by the baseline model, whose parameter values are given in Section IV. The pattern is an average based on 10,000 simulations of the baseline model.

allocates about 65% of its financial wealth to stocks. It then increases the stock allocation again, which eventually reaches 100% of assets and stays there until death. Changes in the resource share of the household's human capital, which from a financial perspective is quite similar to an implicit investment in bonds, explain this pattern. The rapid accumulation of financial wealth that starts at age 45 raises the household's resource share of financial assets and correspondingly reduces the resource share represented by human capital. Thus, as the relative weight of "bond-like" human capital in remaining household resources declines, the household optimally starts allocating a smaller fraction of financial wealth to stocks. This trend continues until retirement, when the household starts depleting assets to finance consumption in retirement. Consumption declines and approaches retirement income as assets dwindle. The declining trend in assets during retirement increases the relative weight of riskless retirement benefits in total wealth, which leads the household to optimally increase its allocation to stocks.

C.   Impact on Welfare, Optimal Consumption, and Portfolio Choice of Early Resolution of Uncertainty about Retirement Income

1.   Welfare Analysis

We next explore the impact on welfare, optimal consumption, and portfolio decisions of letting households learn in advance the size of their retirement benefits. The results are shown in table 1. The table reports welfare gains for different values of A, the age at which the household learns about the retirement income it will receive at age 65, relative to the case in which the household learns at age 65.

The table has five panels, each of which entails a change from the baseline model along a different dimension. In particular, we consider welfare gains as we vary the coefficient of relative risk aversion γ (panel A), the potential percentage cut ξ in the replacement ratio (panel B), the volatility to shocks to permanent income σ_u (panel C), the ability to invest in equities (panel D), and a combination of some of the previous cases (panel E).

The welfare calculations are standard consumption-equivalent variations. For each case (i.e., for each value of A) we compute the constant consumption stream that makes the household as well off in expected utility terms as under the consumption stream that it will actually obtain. Relative utility gains are measured as the change in this equivalent consumption stream relative to the case $A = 65$. Thus we can interpret the

Table 1
Welfare Gains from Early Resolution of Uncertainty about Future Retirement Benefits (%)

	Age of Resolution of Uncertainty (A)					
	35	40	45	50	55	60
A. Relative risk aversion (γ):						
$\gamma = 7$	.345	.327	.298	.262	.211	.131
$\gamma = 5$ (baseline)	.117	.115	.111	.100	.084	.056
$\gamma = 3$	.012	.012	.012	.012	.011	.009
B. Cut in replacement ratio (ξ):						
$\xi = 45\%$	.294	.290	.274	.243	.204	.153
$\xi = 40\%$	.251	.248	.237	.213	.176	.114
$\xi = 30\%$ (baseline)	.117	.115	.111	.100	.084	.056
$\xi = 20\%$	.042	.042	.040	.037	.031	.021
$\xi = 15\%$	.021	.021	.020	.019	.017	.014
C. Volatility of shocks to permanent income (σ_u):						
$\sigma_u = 15\%$	.175	.169	.158	.143	.121	.080
$\sigma_u = 10.95\%$ (baseline)	.117	.115	.111	.100	.084	.056
$\sigma_u = 7\%$	.028	.028	.028	.027	.024	.018
D. Assets available to household:						
Only bonds (nonequity investors)	.141	.140	.135	.122	.101	.066
Bonds and equities (baseline)	.117	.115	.111	.100	.084	.056
E. Combined scenarios:						
$\gamma = 7$, $\xi = 40\%$, nonequity investors	.733	.691	.624	.542	.430	.262
$\gamma = 7$ and $\xi = 40\%$	.666	.630	.571	.499	.398	.244
$\gamma = 5$, $\xi = 30\%$, bonds and equities (baseline)	.117	.115	.111	.100	.084	.056

Note: The table reports welfare gains for different values of A, the age at which the household learns about the retirement income it will receive at age 65, relative to the case in which the household learns at age 65. The welfare calculations are standard consumption-equivalent variations: For each case (i.e., for each value of A), we compute the constant consumption stream that makes the household as well off in expected utility terms as under the consumption stream that it will actually obtain. Relative utility gains are measured as the change in this equivalent consumption stream relative to the case $A = 65$. Thus we can interpret the numbers in the table as the percentage annual consumption loss that a household is willing to accept in order to learn at age A about the replacement income ratio it will receive at age 65 when it retires. The baseline row in each table reports results for the basic experiment, whose parameter values are given in Sec. IV. The table has five panels, each of which entails a change from the baseline model along a different dimension indicated in the heading of the panel. Panel A examines the effect of considering different coefficients of relative risk aversion. Panel B considers the effect of changing the size of the cut in the replacement ratio. Panel C considers the effect of changing the volatility of shocks to permanent income. Panel D considers the effect of allowing the investor to invest only in bonds. Finally, panel E considers a combination of cases.

numbers in the table as the percentage annual consumption loss that a household is willing to accept in order to learn at age A about the replacement income ratio it will receive at age 65. The appendix describes the procedure used to compute our welfare metric.

The rows labeled "baseline" in table 1 report annual welfare gains from learning at age A, instead of age 65, the exact realization of the replacement ratio λ in our baseline model. As expected, the gains are larger the earlier the household learns about its retirement income. Most important, these gains are economically significant. For example, our baseline household is willing to pay, each year, 0.117% of consumption in order to learn at age 35 what income replacement rate it will experience at retirement. With consumption averaging about $30,000 per year in 1992 prices in our baseline model, this is equivalent to a one-time fee of about $906 at age 35 in 1992 prices—or about $1,371 in today's prices. It is also roughly equivalent to the annual expense ratio on a typical index mutual fund. Even at ages as old as 50 and 55, this household is still willing to pay 0.084% and 0.056% of its annual consumption in order to eliminate its uncertainty about retirement income.

Panel A in table 1 shows that the benefit of learning early about future retirement income changes dramatically with risk aversion. Our baseline case assumes a coefficient of relative risk aversion equal to 5. A household with a coefficient of relative risk aversion of 7 is willing to reduce consumption by almost three times as much as our baseline household in order to learn its replacement rate at age 35. By contrast, a household with a coefficient of relative risk aversion of 3 is willing to pay only a tenth of what our baseline household is willing to pay to learn its retirement income at age 35. Although not reported in table 1 to save space, the willingness to pay is essentially zero for a household with a unit relative risk aversion coefficient (i.e., log utility). There is considerable controversy in the literature as to what risk aversion coefficient best captures modal behavior in single-agent models. We report results for a wide range of values of the coefficient of relative risk aversion to reflect this disagreement but also to capture the considerable heterogeneity across the population with respect to attitudes toward risk.

We have repeated the calculations in panel A replacing the baseline value of $\overline{\lambda}$ with .6 instead of .8. The welfare gains are somewhat smaller for the case with a relative risk aversion coefficient of 7 and slightly larger for the cases with relative risk aversion coefficients of 5 and 3.[8] We have also considered the case of a unit relative risk aversion coefficient (i.e., log utility). In this case the resulting excess burden is essentially zero.

Panel B shows that welfare gains from learning early increase dramatically with the magnitude of the potential cut in retirement benefits. For example, when the size of the potential cut in benefits is 40% of the replacement income ratio instead of 30% (when the replacement ratio is 48% of permanent income instead of 56%), the welfare gains from

learning about the cut in retirement benefits in advance are at least twice as large as in the baseline case for all ages A. When the size of the potential cut in benefits is 45% instead of 30%, a household is willing to pay 0.394% of annual consumption in order to learn at age 35 the income replacement ratio it will obtain at retirement. This fee is more than three times the 0.117% fee the household is willing to pay when the size of the potential cut is 30%. Conversely, welfare gains decrease dramatically when the size of the potential cut decreases.

Panel C explores the effect of changes in the volatility of shocks to permanent income (σ_u) on the household's willingness to pay to learn about its retirement benefits in advance. Since the level of retirement benefits depends on the level of permanent income at retirement, an increase in the volatility of shocks to permanent income makes retirement income more uncertain. This uncertainty compounds with the uncertainty about the replacement ratio that will be applied to the level of permanent income to determine actual retirement benefits. Panel C shows that the welfare gains from learning early about the income replacement ratio are increasing in σ_u. These gains are between 43% and 50% larger when σ_u is 15% than when it is 10.95%.

The welfare gains reported thus far are based on the assumption that the household can adjust both its consumption and its asset allocation in response to learning early about retirement benefits. In practice, however, many households do not participate in the stock market. While our model is not designed to explain optimal nonparticipation in the stock market, it is still interesting to explore within the model the welfare gains from early resolution of uncertainty when the household is fully invested in bonds at all times and can adjust consumption—or, equivalently, saving—only in response to learning early about retirement benefits. Panel D explores this scenario. It shows that welfare gains are about 20% larger when the household is unable to invest in stocks. In other words, there is a 20% marginal benefit of being able to invest in both bonds and stocks. The marginal benefit of being able to modify the investment policy, while large, is not as large as the effect of being able to modify the level of consumption and saving.

Panel E shows that the effects of higher risk aversion and a larger cut in benefits interact. A household with a coefficient of relative risk aversion of 7 facing a potential 40% cut in the income replacement ratio at retirement is willing to pay 0.666% of annual consumption to learn about the replacement ratio at age 35. These welfare gains are very large, both in absolute terms and relative to the cases that consider changes in each factor in isolation. They are 165% larger than the effect when risk aversion is

5 and 93% larger than the effect when the potential cut is 30%. Of course, limiting the access of the household to the stock market increases the welfare gains even more. For example, with risk aversion at 7 and a 40% potential benefit cut, the welfare gain to learning at age 35, rather than age 65, whether or not the cut will happen is the equivalent of 0.733% of annual consumption.

2.   Effect on Consumption and Portfolio Choice

Table 1 shows that households are willing to pay a nontrivial fraction of their resources in order to eliminate uncertainty about their future retirement income. Early learning is advantageous for them because they can modify their consumption and asset allocation plans in response to news about their future retirement income. We now examine these changes for the cases considered in table 1.

Table 2 summarizes the effect on optimal consumption—or, equivalently, saving—of early resolution of uncertainty about retirement income. To facilitate interpretation, the table reports the percentage change in consumption for a household that does not learn about the income replacement ratio in retirement until age A relative to a household that learns about it at the earliest possible age (i.e., age 28). For example, the number in the baseline row corresponding to $A = 55$ indicates that, relative to a household that knows the exact value of its retirement income replacement ratio by age 28, a household facing uncertainty about retirement benefits until age 55 will, on average, consume 0.1% less per year between age 28 and age 55; of course, after that age the uncertainty has been resolved for both households.[9] Since there are two possible realizations of the income replacement ratio, we compute optimal consumption under each and then average across the two using their probabilities. The table reports the results for our baseline case as well as the cases we examine in table 1.

Our results indicate that households respond optimally to a delay in the resolution of uncertainty by reducing consumption. The magnitude of the effect increases with the delay in the resolution of uncertainty. In our baseline model, a household that does not learn its income replacement ratio until age 55 does not change consumption significantly relative to a household that learns at age 28. However, a household that learns only at age 65, when it retires, will consume between age 28 and age 65 about 0.4% less per year than a household that learns at age 28. A 0.4% reduction in annual consumption over 37 years is not extremely large, but it is still economically significant.

Table 2
Effect on Consumption of Early Resolution of Uncertainty about Future Retirement Benefits (%)

	Age of Resolution of Uncertainty (A)						
	35	40	45	50	55	60	65
A. Relative risk aversion (γ):							
$\gamma = 7$	.0	.0	−.1	−.3	−.4	−.6	−.9
$\gamma = 5$ (baseline)	.0	.0	.0	.0	−.1	−.2	−.4
$\gamma = 3$	.0	.0	.0	.0	.0	.0	.0
B. Cut in replacement ratio (ξ):							
$\xi = 45\%$	.0	.0	.0	−.2	−.4	−.6	−1.3
$\xi = 40\%$	.0	.0	.0	−.1	−.2	−.4	−.9
$\xi = 30\%$ (baseline)	.0	.0	.0	.0	−.1	−.2	−.4
$\xi = 20\%$	.0	.0	.0	.0	.0	.0	−.1
$\xi = 15\%$	.0	.0	.0	.0	.0	.0	−.1
C. Volatility of shocks to permanent income (σ_u):							
$\sigma_u = 15\%$	.0	.0	.0	−.1	−.1	−.2	−.5
$\sigma_u = 10.95\%$ (baseline)	.0	.0	.0	.0	−.1	−.2	−.4
$\sigma_u = 7\%$	.0	.0	.0	.0	.0	.0	−.2
D. Assets available to household:							
Only bonds (nonequity investors)	.0	.0	.0	.0	−.1	−.2	−.5
Bonds and equities (baseline)	.0	.0	.0	.0	−.1	−.2	−.4
E. Combined scenarios:							
$\gamma = 7$, $\xi = 40\%$, nonequity investors	−.1	−.2	−.3	−.5	−.8	−1.2	−2.0
$\gamma = 7$ and $\xi = 40\%$	−.1	−.1	−.3	−.4	−.7	−1.1	−1.8
$\gamma = 5$, $\xi = 30\%$, bonds and equities (baseline)	.0	.0	.0	.0	−.1	−.2	−.4

Note: The table reports the percentage change in consumption for a household that does not learn about its income replacement ratio in retirement until age A, relative to a household that learns about it at the earliest possible age (i.e., age 28). The table reports percentage changes in consumption only to the first decimal digit because digits beyond that are indistinguishable from approximation error. The table reports results for our baseline case as well as the cases we examine in table 1.

The reduction in consumption caused by a delay in resolution of uncertainty increases with risk aversion, the size of the potential cut in benefits, and the volatility of shocks to permanent income. Increasing risk aversion leads to a reduction in consumption even for small delays in the resolution of uncertainty and to a large reduction for long delays. A household with a coefficient of relative risk aversion of 7 that does not learn until age 65 reduces consumption by about 0.9% per year relative to a household that learns at age 28. Increasing the size of the potential cut in benefits does not have a large impact for small delays in the resolution of uncertainty, but it leads to a large reduction in consumption for long delays. A household that does not learn until age

65 reduces consumption by about 1.3% per year relative to a household that learns at age 28 when the potential cut in the replacement ratio is 45%. Finally, panel E shows that the combination of increased risk aversion coefficients and increased size of the potential cut in the retirement benefits leads to the largest reductions in consumption, which are significant even for small delays in the resolution of uncertainty.

We have also examined the impact on portfolio allocations of a delay in the resolution of uncertainty. Consistent with our finding of the relatively low marginal value of being able to modify the investment policy, we find that these effects are all very small.[10]

VI.   The Benefits of Early Resolution of Uncertainty about Taxes

Section V has considered the impact on welfare, optimal consumption, and portfolio choice of early resolution of uncertainty about retirement income. But households also face uncertainty about future tax rates. This section examines the impact on welfare, optimal consumption, and portfolio choice of early resolution of uncertainty about future labor income tax rates. We assume throughout that the income replacement ratio λ is known in advance and equals 0.8.

A.   Labor Income Tax Uncertainty

We assume that the new tax rate takes effect at age 50, but as in Sections V.A and V.B, the exact value is unknown until some age A (with $A \le 50$). The tax rate uncertainty applies to the tax rates both during working life and during the retirement period. Both face the same level of uncertainty and both are revealed at the same time.

We first consider a symmetric uncertainty case:

$$\tau_L = \begin{cases} \tau_L(1 + \xi) & \text{with probability } p \\ \tau_L & \text{with probability } 1 - 2p \\ \tau_L(1 - \xi) & \text{with probability } p \end{cases} \tag{21}$$

and the same for τ_R.

In our baseline experiment the labor income tax rate might increase or decrease with an equal 25% probability ($p = 0.25$) or remain constant with a 50% probability. As in Section V.B, we are interested in measuring the welfare costs and distortions associated with delaying the announcement of the (new) tax rate. In our calibration, the labor income tax rate is 30% during working life (τ_L) and 15% during retirement (τ_R).

As a result, in the baseline uncertainty case, the labor income tax rate may increase (decrease) from 30% to 39% (21%) between ages 50 and 65 and from 15% to 19.5% (10.5%) between ages 65 and 100 with a 25% probability.

B. Welfare Analysis and Consumption Distortions

Table 3 reports welfare gains from learning early about the future change in the labor income tax rate, and table 4 reports the percentage change in optimal consumption for a household that does not learn

Table 3

Welfare Gains from Early Resolution of Uncertainty about Future Labor Income Tax Rates (%)

	Age of Resolution of Uncertainty (A)		
	35	40	45
A. Relative risk aversion (γ):			
$\gamma = 7$	.084	.067	.037
$\gamma = 5$ (baseline)	.030	.028	.020
$\gamma = 3$	.003	.003	.000
B. Change in labor income tax rate:			
$\xi = 45\%$	.070	.066	.048
$\xi = 30\%$ (baseline)	.030	.028	.020
$\xi = 15\%$	.007	.007	.005
C. Assets available to household:			
Only bonds (nonequity investors)	.030	.029	.022
Bonds and equities (baseline)	.030	.028	.020
D. Asymmetric change in tax rate:			
Asymmetric 30% change	.037	.030	.017
Symmetric 30% change (baseline)	.030	.028	.020

Note: The table reports welfare gains for a household from learning early about the future change in the baseline labor income tax rate of 30% (during working life) and 15% (in retirement) for a household that does not learn about the change until age A, relative to a household that learns about it at the earliest possible age (i.e., age 28). The baseline row in each table reports results for the basic experiment. In this table, households know in advance the replacement ratio for the permanent labor income they get in retirement. We consider four variants of the baseline experiment, which are reported in panels A–D. Panel A examines the effect of considering different coefficients of relative risk aversion. Panel B considers the effect of changing the size of the shock to the labor income tax rate. Panel C considers the effect of allowing the investor to invest only in bonds. Finally, panel D considers the case in which the uncertainty about the income tax rate is asymmetric. In the asymmetric case, the tax rate can increase (decrease) by 30% from 30% to 39% (19.5%) during working life (retirement) with a probability of only one-third. This case is similar in spirit to our case of a change in retirement benefits, where we have considered only cuts in retirement income.

Table 4
Effect on Consumption of Early Resolution of Uncertainty about Future Labor Income Tax Rates (%)

	Age of Resolution of Uncertainty (A)			
	35	40	45	50
A. Relative risk aversion (γ):				
$\gamma = 7$	.0	−.1	−.1	−.3
$\gamma = 5$ (baseline)	.0	.0	.0	−.2
$\gamma = 3$	.0	.0	.0	−.1
B. Change in labor income tax rate:				
$\xi = 45\%$	.0	−.1	−.1	−.4
$\xi = 30\%$ (baseline)	.0	.0	.0	−.2
$\xi = 15\%$	.0	.0	.0	−.1
C. Assets available to household:				
Only bonds (nonequity investors)	.0	.0	.0	−.1
Bonds and equities (baseline)	.0	.0	.0	−.2
D. Asymmetric change in tax rate:				
Asymmetric 30% change	.0	.0	.0	−.1
Symmetric 30% change (baseline)	.0	.0	.0	−.2

Note: The table reports the percentage change in optimal consumption for a household from learning early about the future change in the baseline labor income tax rate of 30% (during working life) and 15% (in retirement) for a household that does not learn about the change until age A, relative to a household that learns about it at the earliest possible age (i.e., age 28). The table reports percentage changes in consumption only to the first decimal digit because digits beyond that are indistinguishable from approximation error. The table reports results for the same cases we examine in table 3.

about the labor income tax rate change until age A relative to a household that learns about it at the earliest possible age (i.e., age 28).[11]

The baseline row in each table reports results for the basic experiment. We consider four variants of the baseline experiment, which are reported in panels A–D. Panel A examines the effect of considering different coefficients of relative risk aversion. Panel B considers the effect of changing the size of the shock to the labor income tax rate. Panel C considers the effect of allowing the investor to invest only in bonds. Finally, panel D considers the case in which the uncertainty about the income tax rate is asymmetric. In the asymmetric case, the tax rate can increase (decrease) by 30% from 30% to 39% (19.5%) during working life (retirement) with a probability of only one-third. This case is similar in spirit to our case of a change in retirement benefits, where we have considered only cuts in retirement income.

Table 3 shows that the baseline welfare gains from early resolution of labor income tax uncertainty increase as we consider earlier announcement dates. In the baseline case the welfare loss is as high as 0.03% of annual consumption, and this increases to 0.084% for the more risk-averse

households. These magnitudes are commensurate with the welfare gains from knowing early about future retirement benefits shown in table 1. To see this, note that a household is willing to pay an annual fee equivalent to 0.03% of consumption between age 35 and age 50 for eliminating at age 35 a 30% uncertainty over 30% of its income between ages 50 and 65, and 15% of its income between ages 65 and 100. This is essentially equivalent to eliminating a 9% uncertainty over 100% of the household's income between ages 50 and 65 and a 4.5% uncertainty over 100% of the household's income between ages 65 and 100.

By contrast, the uncertainty that is eliminated in the retirement benefits case is proportionally much larger since it implies a 30% uncertainty over 100% of the household's income between ages 65 and 100. Accordingly, the welfare gain from knowing early is also proportionally larger—about 0.12%. Thus, for the same level of uncertainty, the benefits from early resolution of uncertainty in labor income taxes are similar to those derived from early resolution of uncertainty about future retirement benefits.

Table 3 also shows that, similarly to the retirement benefits case, welfare gains are most sensitive to changes in the coefficient of relative risk aversion and to the magnitude of the uncertainty about the future labor income tax change, while preventing households from accessing the stock market does not have a large impact. Interestingly, panel D shows that considering an asymmetric change in the labor income tax rate in lieu of a symmetric change does not have a significant effect on welfare gains. This is an important result, and it highlights, once again, that the welfare costs are due to the late resolution of uncertainty and not to the nature of uncertainty.

The changes in optimal consumption reported in table 4 are also commensurate with the changes in consumption reported in table 2. As expected, a delay in learning about future changes in labor income tax rates causes households to reduce their consumption. The reduction is largest for those households that bear the longest delay in learning.

C. Capital Income Tax Uncertainty

Finally, we also explore the implications of early resolution of uncertainty about future capital income tax rates. We find that the welfare gains from early resolution of this type of uncertainty are very small. Figure 1 is helpful to understand why moderate uncertainty about future capital income tax rates is less costly than uncertainty about future labor income tax rates or retirement income. This figure shows that labor earnings finance most of the consumption of the typical household

and suggests that the household uses wealth, particularly during its working life, to smooth the impact of income shocks on consumption.

Recall that our baseline interest rate is 1%, and our baseline real return on equities is 5%. Our baseline 20% capital income tax rate implies that the after-tax return on the household portfolio is between 0.8% and 4%. A 30% degree of uncertainty in this tax rate implies that the after-tax return on the household portfolio could be further reduced by an amount ranging from 6 basis points to 30 basis points. This is small compared to a potential reduction of 9% of labor income caused by the chance of a 30% change in the 30% baseline labor income tax rate. Thus, a change in the labor income tax rate or a change in retirement income is likely to have a more significant impact on future consumption than a change of similar magnitude in the effective capital income tax rate.

VII. Simultaneous Uncertainty about Retirement Income and Tax Rates

Our exercises so far have considered the welfare costs, and their impact on consumption and portfolio choice, of delays in several policy decisions each considered in isolation. It is of course possible for these policy indecisions to manifest simultaneously, particularly if they result from situations of fiscal crises. Accordingly, we next explore the implications of early joint resolution of uncertainty about both future tax rates and the benefit replacement rate.

We consider a scenario in which there is a one-third probability of a 30% simultaneous decline in the baseline 30% labor income tax rate at age 50 and in the 80% benefit replacement rate at age 65. There is a two-thirds probability that they do not change. We assume that, in the default scenario, the household does not learn whether these changes actually occur until age 50.

We also consider a special case in which the uncertainty about future tax rates and the uncertainty about future retirement income are uncorrelated. In particular, we consider the welfare gains from early resolution of uncertainty about future labor income tax rates when the income replacement ratio λ is uncertain and not known until age 65 (Sec. VI explores the case in which λ is known to the household in advance).

A. Welfare Analysis

Table 5 reports the welfare gains from resolving this uncertainty at an earlier age A relative to learning at age 50. Table 6 reports the change in

Table 5
Welfare Gains from Joint Resolution of Uncertainty about Future Retirement Benefits and
Labor Income Tax Rates (%)

	Age of Resolution of Uncertainty (A)		
	35	40	45
A. Relative risk aversion (γ):			
$\gamma = 7$	.247	.194	.107
$\gamma = 5$ (baseline)	.076	.069	.047
$\gamma = 3$	.004	.004	.004
B. Size of cut in retirement benefits and			
increase in labor income tax rate:			
$\xi = 45\%$	.246	.220	.140
$\xi = 30\%$ (baseline)	.076	.069	.047
$\xi = 15\%$	.013	.012	.008
C. Assets available to household:			
Only bonds (nonequity investors)	.086	.080	.055
Bonds and equities (baseline)	.076	.069	.047
D. Uncorrelated uncertainty:			
Uncorrelated ξ	.038	.033	.022
Baseline	.076	.069	.047

Note: The table reports the welfare gains from joint resolution of uncertainty about future
tax rates and the replacement ratio of permanent income in retirement at age A relative to
learning at age 50. We consider a scenario in which there is a one-third probability that both
the baseline 30% labor income tax rate increases by 30% at age 50 and the 80% replacement
ratio of permanent income in retirement declines by 30% at age 65. There is a two-thirds
probability that they do not change. We assume that, in the default scenario, the household
does not learn whether these changes actually occur until age 50. The baseline row in
each table reports results for the basic experiment, whose parameter values are given in
Sec. IV. The table has four panels, each of which entails a change from the baseline model
along a different dimension indicated in the heading of the panel. Panel A examines the
effect of considering different coefficients of relative risk aversion. Panel B considers the
effect of changing the size of the cut in the replacement ratio and the size of the increase
in the labor income tax rate. Panel C considers the effect of allowing the investor to invest
only in bonds. Finally, panel D considers a special case in which the uncertainty about
future tax rates and the uncertainty about future retirement income are uncorrelated. In
particular, it considers the welfare gains from early resolution of uncertainty about fu-
ture labor income tax rates when the income replacement ratio is uncertain and not
known until age 65.

optimal consumption relative to a household that resolves this uncertainty
at age 28. We omit the results for the changes in optimal portfolios because
they are, once again, quantitatively small. Both tables also explore some
variants of the basic exercise: panel A examines scenarios in which the
household coefficient of relative risk aversion changes relative to the
baseline model, panel B considers changes in the size of the potential
percentage increase in the income tax rate and decrease in the income
replacement ratio, panel C considers households that can invest only

Table 6
Effect on Consumption of Joint Resolution of Uncertainty about Future Retirement Benefits and Labor Income Tax Rates (%)

	Age of Resolution of Uncertainty (A)			
	35	40	45	50
A. Relative risk aversion (γ):				
$\gamma = 7$	–.8	–1.4	–1.9	–2.5
$\gamma = 5$ (baseline)	–.1	–.3	–.5	–1.3
$\gamma = 3$	.0	.0	.0	–.5
B. Size of cut in retirement benefits and increase in labor income tax rate:				
$\xi = 45\%$	–.2	–.6	–1.1	–2.5
$\xi = 30\%$ (baseline)	–.1	–.3	–.5	–1.3
$\xi = 15\%$	.0	–.1	–.2	–.5
C. Assets available to household:				
Only bonds (nonequity investors)	–.1	–.3	–.5	–1.4
Bonds and equities (baseline)	–.1	–.3	–.5	–1.3
D. Uncorrelated uncertainty:				
Uncorrelated ξ	.0	–.2	–.4	–1.1
Baseline	–.1	–.3	–.5	–1.3

Note: The table reports the change in optimal consumption from joint resolution of uncertainty about future tax rates and the replacement ratio of permanent income in retirement at age A relative to a household that resolves this uncertainty at age 28. The table reports percentage changes in consumption only to the first decimal digit because digits beyond that are indistinguishable from approximation error. The table reports results for the same cases we examine in table 5.

in bonds, and, finally, panel D considers the case with uncorrelated uncertainty.

Table 5 shows that the welfare changes from adding uncertainty about retirement benefits to uncertainty about the income tax rate—shown in table 3—are very large, both when uncertainty about retirement benefits is correlated with uncertainty about the future income tax rate and when it is uncorrelated. For example, a household that learns at age 40 about the income tax rate prevailing at age 50 instead of learning at age 50 experiences a utility gain of 0.028% per year. If the household is also uncertain about its future retirement benefits and first learns about them at age 40, the welfare gain is more than twice as large—0.069% per year.

The welfare gains are proportionally larger when the household is more risk averse or faces more uncertainty. In both cases, the gains are about three times as large relative to the case in which there is uncertainty only about the labor income tax rate. Interestingly, the gains for nonequity investors are now larger than the gains for households that can invest in both bonds and equities. This suggests that

the importance of being able to change portfolio allocations increases as we consider scenarios with more uncertainty about noncapital income.

B.   Consumption Distortions

Consistent with our findings about welfare gains, table 6 shows that early resolution of uncertainty about future labor income tax rates and retirement income has a large impact on optimal consumption. Households that do not resolve this uncertainty until age A consume significantly less per year than households that resolve it early in their life cycle—at age 28. For example, in our baseline model a household that does not resolve this uncertainty until age 50 consumes about 1.3% per year less than an identical household that learns at age 28. For a household with a coefficient of relative risk aversion of 7, the reduction in optimal consumption is 2.5%. Even at age 35, where the delay is only one of 7 years, the reduction in consumption is 0.1% per year in the baseline case and 0.8% per year in the case with a coefficient of relative risk aversion of 7. In general, the fall in consumption is largest for households that are more risk averse and for households that face large possible reductions in income.

VIII.   The Costs of Policy Delay with Flexible Labor Supply

The costs of policy delay we have presented so far are based on the assumption that households can react to policy indecision only by modifying their savings and asset allocation decisions. Another important dimension along which households might be able to buffer the costs of policy delay is through their labor supply decisions. By varying the hours they work, households have more flexibility to react to adverse shocks to their income and wealth, thus making less costly the effect of a late resolution of uncertainty. On the other hand, early knowledge of future policy changes is more valuable when there is an extra dimension along which households can adjust, thus making the benefits of early resolution of uncertainty more valuable.

   We now explore which of these two effects dominates in the context of a realistically calibrated life cycle model with flexible labor supply and policy uncertainty. Our work builds on Gomes et al. (2008), which extends life cycle models of consumption and portfolio choice to allow for labor supply decisions.[12]

   We continue to assume that agents work their first K periods and pay wage taxes in proportion to their labor earnings and capital income

taxes in proportion to their asset income. But we now allow leisure L_t to be a choice variable that enters the instantaneous utility function of the household. Preferences over consumption and leisure are now given by

$$\mathrm{E}_1 \sum_{t=1}^{T} \delta^{t-1} D_t \left(\prod_{j=0}^{t-2} p_j \right) \frac{(C_t L_t^{\theta})^{1-\gamma}}{1-\gamma}, \qquad (22)$$

where L_t is time t leisure, θ is a leisure preference parameter, and the rest of the variables and coefficients are defined as in (8). Households can choose how much leisure L_t they want to consume each period or, equivalently, how much they want to work. Leisure is measured as a fraction of total available time and satisfies $L_t \in [\underline{L}, 1]$, where $\underline{L}$ is minimum leisure time, which we set to one-third.

To calibrate θ and the wage process, we follow Gomes et al. (2008). We choose θ so that the average labor supply over the life cycle matches the average male hours of work per year reported in the Consumer Expenditure Survey—2,080 hours per year. The process for the wage rate takes the same form as the process (10)–(11) for total labor income. Of course, total labor earnings now equal optimally chosen labor supply times the wage rate. The rest of the parameters of the model follow the same calibration as the model with fixed labor supply of Sections III–VII. The appendix describes in full detail the wage process and the calibration of the model.

To keep our exposition of results as simple as possible, we report only welfare gains in this section of the paper and instead briefly describe here the impact of early resolution of uncertainty about retirement income replacement ratios on labor supply decisions. We find that early resolution of uncertainty leads to a greater adjustment of labor supply when the household is closer to retirement. Relative to our benchmark case in which uncertainty is resolved at age 65, labor supply changes by 0.03% per year at age 30 and changes by 0.46% per year at age 60. When uncertainty is resolved early, the household needs to increase labor supply much less than when uncertainty is resolved late and the household has less time to adjust.

Table 7 shows the welfare gains from early resolution at different ages of uncertainty about retirement income replacement ratios. Because utility now depends on leisure as well as consumption, we can no longer measure welfare changes in terms of changes in certainty-equivalent consumption as it is standard in this literature. Instead, we report the percentage of first-year labor earnings the household

Table 7
Welfare Effects of Early Resolution of Uncertainty: The Case of Variable Labor Supply (%)

	Age of Resolution of Uncertainty (A)					
	35	40	45	50	55	60
Fixed labor supply	.673	.658	.633	.574	.482	.313
Variable labor supply	.678	.665	.641	.593	.540	.420

Note: The table reports welfare gains from early resolution at different ages of uncertainty about retirement income replacement ratios. For each of the two cases reported in the rows of the table, it reports the percentage of first-year labor earnings the household would be willing to give up in order to resolve the uncertainty about the retirement income replacement ratio at age A, relative to the case in which the household learns about this ratio at a retirement age of 65. The gains are based on our baseline case, described in table 3 and in Secs. IV and VIII.

would be willing to give up in order to resolve the uncertainty about the retirement income replacement ratio at age A, relative to the case in which the household learns about this ratio at a retirement age of 65. For ease of comparison, we also show the welfare gains for the fixed labor supply model calculated on the same basis.

There are two interesting results emerging from table 7. First, to the extent one can really compare the variable and fixed labor supply models, the welfare gains are larger in the case of variable labor supply. Thus being able to adjust labor supply makes early knowledge of future policy changes more valuable to households. Second, the difference in welfare gains from resolving uncertainty early is small at younger ages and becomes larger at later ages. The ability to adjust labor supply in response to policy news is not very valuable when the household can already adjust its savings and portfolio allocation and has a long horizon until retirement. However, this extra degree of flexibility is much more valuable when there are fewer years until retirement.

There is, however, a counterargument for why variable labor supply would lead to a smaller gain from early resolution, namely, that being able to adjust labor supply is similar to having an insurance policy against receiving smaller benefits than expected; if the insurer is risk neutral, its premium will not differ whether it learns the policy outcome early or late since its payouts will be the same. Thus, if the part of the utility function involving leisure were separable, linear, and not time discounted, households would, in effect, represent risk-neutral insurance companies. They would adjust at age 65 to the news about their benefits by altering their labor supply and earnings, not their consumption. And since they would make the same adjustments no matter whether they learned early or late about the benefit change, learning

early would not change expected utility. With our preferences, which are nonlinear in leisure, being able to respond to shocks makes learning early more valuable even though the shock is of smaller consequence, for any given date of learning, because of the ability to adjust labor supply.

IX.   Conclusions

This paper has explored the effects on life cycle saving, asset allocation, labor supply, and welfare of delays in the resolution of uncertainty. Specifically, the paper conducts an analysis of the excess burden of government indecision, an issue that has not, to our knowledge, been directly addressed in prior research. The paper evaluates the effects of delaying the resolution of uncertainty about future government-provided retirement benefits and government-imposed taxes using a realistically calibrated model of life cycle consumption, saving, and portfolio choice.

We find that households respond optimally to a delay in the resolution of uncertainty by reducing their consumption. This reduction in consumption is increasing in the size of the potential cuts in disposable income (arising from retirement benefit cuts or tax hikes), risk aversion, and the volatility of shocks to labor earnings. A delay in the resolution of uncertainty has no significant effect on asset allocation decisions.

We also find that households experience sizable welfare gains from learning early about future changes in benefits and tax rates regardless of their attitudes toward risk or the uncertainty they face about their own labor earnings. For example, our baseline household is willing to pay an annual fee equivalent to 0.12% of annual consumption in order to learn at age 35 the Social Security benefit income replacement ratio that it will experience at retirement. Welfare gains are most pronounced for more risk-averse households, households facing more uncertainty about future labor earnings, and households facing either larger potential cuts in benefits or increases in marginal tax rates.

Finally, being able to adjust one's labor supply in addition to consumption and portfolio allocations exacerbates the excess burden of policy indecision. Having early knowledge of future changes in policy is more valuable the larger the number of dimensions along which the household can adjust in advance of the policy change. The marginal gain from elastic labor supply is relatively small when the household has a long horizon before the policy change along which it can adjust its savings and asset allocation. But it becomes very valuable when the household has a shorter horizon to adjust.

This paper has highlighted the importance of the economic benefits of early resolution of uncertainty in the context of two specific policy changes. There are other types of policy indecisions whose costs might be worth exploring, including switching tax structure (e.g., from income to consumption taxation), cutting back on health care benefits, or printing money to pay public debt, leading potentially to very high rates of inflation and overall financial and economic instability.

There is always a natural resistance to deliver bad news and to postpone difficult decisions, particularly if there are short-term benefits from delaying the announcement of bad news. Governments are not different in this respect. Whatever gains governments might obtain from delaying the announcement of policy changes, delay also fosters and exacerbates economic uncertainty. In this paper, we have begun to model and quantify the excess burden of government indecision. As we have stressed, this excess burden arises not from implementing specific policies but from delaying their determination and announcement. We have shown that this efficiency loss can be large depending on the precise policy in question, the degree of risk aversion, and access to capital markets.

Appendix

Welfare Metric in the Model with Fixed Labor Supply

The welfare calculations are done in the form of standard consumption-equivalent variations: for each case (i.e., value of A) we compute the constant consumption stream that makes the investor as well off in expected utility terms as the expected consumption stream that she or he will actually obtain. Relative utility losses are then obtained by measuring the change in this equivalent consumption stream when deviating from the optimal rule toward the rule considered.

More precisely, we first solve the optimal consumption/savings problem for an agent for a given value of A. Denoting the optimal consumption stream for this problem by $\{C_t^A\}_{t=1}^T$, we then compute the corresponding expected lifetime utility:

$$V^A \equiv E_1 \sum_{t=1}^T \delta^{t-1} \left(\prod_{j=0}^{t-1} p_j \right) \frac{C_t^{A^{1-\gamma}}}{1-\gamma}. \qquad (A1)$$

Note that this is just the value function from the maximization problem.

We can convert this expected discounted lifetime utility into consumption units by computing the equivalent constant consumption stream $\{\overline{C}^A\}_{t=1}^T$ that leaves the investor indifferent between this and the consumption stream $\{C_t^A\}_{t=1}^T$. This is equivalent to solving

$$V^A = \sum_{t=1}^T \delta^{t-1} \left(\prod_{j=0}^{t-1} p_j \right) \frac{\overline{C}^{A^{1-\gamma}}}{1-\gamma}. \tag{A2}$$

Therefore,

$$\overline{C}^A = \left[\frac{(1-\gamma)V^A}{\sum_{t=1}^T \delta^{t-1}(\prod_{j=0}^{t-1} p_j)} \right]^{1/(1-\gamma)}. \tag{A3}$$

With $A = 65$ taken as our baseline case, the utility gain from changing A can then be simply computed as the percentage loss in equivalent consumption:

$$\frac{\overline{C}^A - \overline{C}^{65}}{\overline{C}^{65}} = \frac{V^{A[1/(1-\gamma)]} - V^{65[1/(1-\gamma)]}}{V^{65[1/(1-\gamma)]}}. \tag{A4}$$

Adding Flexible Labor Supply

We start by noting that under the preference structure (22), labor supply is invariant to secular changes in the real wage in accord with US experience. Under the assumptions of Section VIII, the investor's financial wealth at the end of working period t is given by

$$W_{t+1} = R_{t+1}^p[W_t + (1-h_t)(1-\tau_L)w_t N_t - C_t], \tag{A5}$$

where w_t is the time t wage and N_t is labor supply as a fraction of total available time, that is, $N_t = 1 - L_t$.

The log of wages follows the process

$$\ln w_t = f(t) + v_t + \varepsilon_t, \tag{A6}$$

where $f(t)$ is a deterministic function of age, v_t is a permanent component given by

$$v_t = v_{t-1} + u_t, \tag{A7}$$

u_t is distributed as $N(0, \sigma_u^2)$, and ε_t is a transitory shock uncorrelated with u_t, which is distributed as $N(0, \sigma_\varepsilon^2)$. The innovation to the permanent component of the wage rate (u_t) can be correlated with the return to equity R_t, with coefficient ρ.

We calibrate the wage income process (A6)–(A7) as follows: first, we use the wage profile reported in Fehr, Jokisch, and Kotlikoff (2007) to calibrate the deterministic age-dependent component of wages.[13] Second, we use the estimates of σ_u and σ_ε of 10.95% and 13.89% reported in Cocco et al. (2005) to calibrate the stochastic component of wages. Following Carroll (1997), we divide the estimated standard deviation of transitory income shocks by 2 to account for measurement error. The implied wage growth rates over the life cycle generated by this function exhibit an inverted-U shape and are comparable to average total income growth rates in the Panel Study of Income Dynamics data.

We set λ exactly as in the case with fixed labor supply and uncertainty about the replacement ratio for retirement income; that is, we assume that the log of Social Security income is a fraction λ of the average lifetime labor earnings that the agent would have obtained had he worked full-time during his working life. That is,

$$\ln(Y) = \lambda \frac{\sum_{t=1}^{K}[f(t) + v_t]}{K} \overline{N}, \tag{A8}$$

where $\overline{N}$ denotes full-time labor supply. As noted in Gomes et al. (2008), this simplified assumption is a first-order approximation to the incentives built into the Social Security system, which make Social Security income depend on the individual's average earnings in her 35 highest-earnings years. Letting Social Security income depend on past labor supply decisions—specifically, average past labor supply—introduces a computationally costly extra state variable but makes little difference to the results.

The rest of the calibration follows the calibration in the model with fixed labor supply.

Endnotes

1. This is the level of consumption the agent would need to receive for sure each year to have a sure level of utility equal to the level of utility the person can expect, on average, to enjoy.

2. One explanation for delaying decisions is that older generations, which control polit-ical outcomes and thus policy decisions, seek to wait as long as possible in determining how much to expropriate younger generations because the longer they wait the more they can learn about the level of income that the young are able to generate in the labor market. See Altonji, Hayashi, and Kotlikoff (1997) for an analysis of how uncertainty in the labor earn-ings of adult children will lead their parents to delay deciding how well to treat them.

3. See Kotlikoff (2002) for a demonstration of this point.

4. Note that

$$\frac{\partial f(x)}{\partial x} = \left(1 - \frac{1}{\gamma}\right) x^{-1/\gamma}$$

and

$$\frac{\partial^2 f(x)}{\partial x^2} = -\frac{1}{\gamma}\left(1 - \frac{1}{\gamma}\right) x^{-(1+1/\gamma)}.$$

5. This is the same process as in Carroll (1997) and Gourinchas and Parker (2002). Hubbard, Skinner, and Zeldes (1995) replace the permanent shocks with a very persistent first-order autoregressive process.

6. Following Carroll (1997), we divide the estimated standard deviation of transitory income shocks by 2 to take into account measurement error.

7. We do not include a sensitivity analysis with respect to the discount rate because our welfare results are not significantly different for a sensible range of variation in this pa-rameter. Results are available on request from the authors.

8. We do not have good intuition for why the welfare losses would rise slightly with lower $\overline{\lambda}$ for these values of the relative risk aversion coefficient.

9. We report percentage changes in consumption only to the first decimal digit because digits beyond that are indistinguishable from approximation error.

10. We do not report them for this reason, but they are readily available on request.

11. Once again, the effect on portfolio choice is very small, and accordingly we do not report it.

12. French (2005) and Low (2005) also explore the labor supply decision in a life cycle model of consumption, but they ignore portfolio choice.

13. Specifically, we use their earnings function $E(a, 2)$, given in eq. (9) of their paper, with parameter λ equal to zero. In this function, the argument a denotes age, and 2 de-notes the middle income class.

References

Altonji, Joseph, Fumio Hayashi, and Laurence J. Kotlikoff. 1997. "Parental Al-truism and Inter Vivos Transfers: Theory and Evidence." *Journal of Political Economy* 105 (December): 1121–66.

Auerbach, Alan J., and Kevin A. Hassett. 2002. "Fiscal Policy and Uncertainty." *International Finance* 5 (Summer): 229–49.

Auerbach, Alan J., and Laurence J. Kotlikoff. 1987. *Dynamic Fiscal Policy*. Cam-bridge: Cambridge University Press.

Blundell, Richard, and Thomas M. Stoker. 1999. "Consumption and the Timing of Income Risk." *European Economic Review* 43:475–507.

Bohn, Henning. 1998. "Risk Sharing in a Stochastic Overlapping Generations Economy." Manuscript (January), University of California, Santa Barbara.

Campbell, John Y., João F. Cocco, Francisco J. Gomes, and Pascal J. Maenhout. 2001. "Investing Retirement Wealth: A Life Cycle Model." In *Risk Aspects of Investment-Based Social Security Reform*, ed. John Y. Campbell and Martin Feldstein. Chicago: University of Chicago Press.

Carroll, Christopher. 1992. "The Buffer-Stock Theory of Saving: Some Macro-economic Evidence." *Brookings Papers on Economic Activity*, no. 2:61–156.

———. 1997. "Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112, no. 1:3–55.

Chari, V. V., Lawrence J. Christiano, and Patrick J. Kehoe. 1994. "Optimal Fiscal Policy in a Business Cycle Model." *Journal of Political Economy* 102, no. 4:617–52.

Cocco, João, Francisco Gomes, and Pascal Maenhout. 2005. "Consumption and Portfolio Choice over the Life-Cycle." *Review of Financial Studies* 18, no. 2:491–533.

Croce, Max, Thien Nguyen, and Lukas Schmid. 2011. "The Market Price of Fiscal Uncertainty." Manuscript, University of North Carolina at Chapel Hill.

Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica* 59:1221–48.

Diamond, Peter A. 1997. "Macroeconomic Aspects of Social Security Reform." *Brookings Papers on Economic Activity* (Fall): 1–66.

Eeckhoudt, Louis, Christian Gollier, and Nicolas Treich. 2001. "Optimal Consumption and the Timing of the Resolution of Uncertainty." Manuscript, University of Toulouse, LEERNA-INRA.

Fama, Eugene F., and Kenneth R. French. 2002. "The Equity Premium." *Journal of Finance* 57:637–59.

Fehr, Hans, Sabine Jokisch, and Laurence J. Kotlikoff. 2007. "Will China Eat Our Lunch or Take Us Out to Dinner? Simulating the Transition Paths of the United States, the European Union, Japan, and China." In *Fiscal Policy and Management in East Asia*, ed. Takatoshi Ito and Andrew K. Rose. Chicago: University of Chicago Press (for NBER).

French, Eric. 2005. "The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behaviour." *Review of Economic Studies* 72:395–427.

Gokhale, Jagadeesh, and Kent Smetters. 2005. "Measuring Social Security's Financial Problems." Working Paper no. 11060, NBER, Cambridge, MA.

Gomes, Francisco, Laurence J. Kotlikoff, and Luis M. Viceira. 2008. "Optimal Life-Cycle Investing with Flexible Labor Supply: A Welfare Analysis of Life-Cycle Funds." *American Economic Review Papers and Proceedings* 98:297–303.

Gomes, Francisco, and Alexander Michaelides. 2005. "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence." *Journal of Finance* 60:869–904.

Gourinchas, Pierre-Olivier, and Jonathan Parker. 2002. "Consumption over the Life Cycle." *Econometrica* 70:47–89.

Green, Jerry, and Laurence J. Kotlikoff. 2006. "On the General Relativity of Fiscal Language." Working Paper no. 12344, NBER, Cambridge, MA.

Heaton, John, and Deborah Lucas. 1997. "Market Frictions, Saving Behavior, and Portfolio Choice." *Macroeconomic Dynamics* 1:76–101.

Hubbard, Glenn, Jonathan Skinner, and Stephen Zeldes. 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy* 103:360–99.

Judd, Kenneth. 1989. "Optimal Taxation in Dynamic Stochastic Economies." Manuscript, Hoover Institution, Stanford, CA.

Kotlikoff, Laurence J. 2002. *Generational Policy*. Cambridge, MA: MIT Press.

Kotlikoff, Laurence J., and David Rapson. 2005. "Comparing Average and Marginal Tax Rates under the Fair Tax and the Current System of Federal Income Taxation." Working Paper no. 11831, NBER, Cambridge, MA.

Low, Hamish W. 2005. "Self-Insurance in a Life-Cycle Model of Labour Supply and Savings." *Review of Economic Dynamics* 8:945–75.

Luttmer, Erzo F. P., and Andrew A. Samwick. 2011. "The Costs and Consequences of Perceived Political Uncertainty in Social Security." Manuscript, Dartmouth College, Department of Economics.

Mehra, Rajnish, and Edward C. Prescott. 1985. "The Equity Premium: A Puzzle." *Journal of Monetary Economics* 15:145–61.

Munnell, Alicia H. 2003. "The Declining Role of Social Security." Just the Facts no. 6, Boston College, Center for Retirement Research.

Pastor, Lubos, and Pietro Veronesi. 2011. "Uncertainty about Government Policy and Stock Prices." Manuscript, University of Chicago.

Persson, Torsten, and Lars E. O. Svensson. 1989. "Why a Stubborn Conservative Would Run a Deficit: Policy with Time-Inconsistent Preferences." *Quarterly Journal of Economics* 104:325–45.

Shoven, John B., and Sita N. Slavov. 2006. "Political Risk versus Market Risk in Social Security." Working Paper no. 12135 (March), NBER, Cambridge, MA.