NBER WORKING PAPER SERIES

THE EXCESS BURDEN OF GOVERNMENT INDECISION

Francisco J. Gomes
Laurence J. Kotlikoff
Luis M. Viceira

Working Paper 12859
http://www.nber.org/papers/w12859

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2007

The Excess Burden of Government Indecision
Francisco J. Gomes, Laurence J. Kotlikoff, and Luis M. Viceira
NBER Working Paper No. 12859
January 2007
JEL No. H2,H21,H55,H6

ABSTRACT

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, lifespan, 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, 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 percent of the agent's economic resources. This is a significant distortion in of itself. It's 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.

Francisco J. Gomes
London Business School
Regent's Park
London NW1 4SA
UK
fgomes@london.edu

Laurence J. Kotlikoff
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
and NBER
kotlikoff@bu.edu

Luis M. Viceira
Baker Library 367
Graduate School of Business Administration
Harvard University
Boston, MA 02163
and NBER
lviceira@hbs.edu

1 Introduction

Virtually all U.S. policymakers, budget analysts, and academic experts agree that the U.S. faces a very serious, if not a grave, long-term fiscal problem. Yet few policy makers will publicly say how or when they would fix it, perhaps because they fear being the bearer of bad news and getting voted out of office. Delaying the resolution of fiscal imbalances incurs 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, 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, buy more life insurance, 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's clear that delays in policymaking 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 life cycle model featuring optimal consumption and portfolio choice in the face of uncertainty in earnings, lifespan, 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, which builds on the recent literature on life-cycle consumption and investing modeling, is designed to capture those features of life cycle models shown to have a first-order impact on optimal saving. In particular, we include an asset allocation decision between a riskless asset—the standard assumption in life-cycle models—and a risky asset, as in Cocco, Gomes and Maenhout (2005). We also account for housing expenditures as in Gomes and Michaelides (2005), and for household size. Our model is by necessity stylized. However, our goal is not to arrive at a precise understanding of the magnitude of U.S. economic inefficiency arising from policy indecision. Doing so would require considering all policies that might be changed in the future as well the impact of policy uncertainty on

1

all types of American households. Our goal is much more modest – simply to understand whether policy indecision could generate a reasonably large excess burden for typical middle class households.

The main policy we consider concerns 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 their consumption, saving, and portfolio decisions. Most important, it reduces welfare. Indeed, the excess burden of government indecision, in this instance, can exceed more than .5 percent of agents' resources. In considering this finding, note that we are comparing two scenarios with the *same* expected social security income and the *same* ex-ante uncertainty. Therefore this welfare loss stems not from a change in the agent's expected income, or even from a change in the agent's future income risk. Rather it comes exclusively from delay in resolving policy uncertainty, i.e. from government's indecision. Excess burdens in the range of .5 percent of resources are significant in of themselves. They are also significant in comparison to other distortions measured in the public finance literature such as those reported by Auerbach and Kotlikoff (1987) arising form maintaining an inefficient tax structure.

In addition to considering delays in determining/announcing future Social Security benefit policy, we also consider delays in determining/announcing future tax rate changes. The quantitative findings here are similar to those with respect to Social Security once one scales for the magnitude of the net income at stake. So too are the welfare costs of simultaneous indecision with respect to both future Social Security benefit and tax rate levels. Another central finding is that the excess burden from government indecision is highly sensitive to the degree of risk aversion, the number of years one must wait to have the policy uncertainty resolved (captured in our framework by the agent's age), and the size and probability of policy changes.

Our study appears to be the first to identify and measure the excess burden of government indecision. Previous work has examined the impact on consumption and saving of

early resolution of uncertainty (e.g., Blundell and Stoker, 1999 and Eeckhoudt, Gollier, and Treich 2001) and the manner in which governments should optimally spread risk across generations (e.g., Judd, 1989; Chari, et. al., 1994; Diamond, 1997; Bohn, 1998; and Auerbach and Hassett, 2002).

Before presenting our model we want to reference several issues that we have chosen not to incorporate in 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 that leads to indecision. There are, no doubt, many different explanations for why politicians don't 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 in not on the reasons for indecision, but rather on its costs. Any specific model of why governments delay making decisions would beg two questions, namely why we chose that particular formation and why the political process we considered would ever change its behavior. If government policy can't be changed, there is no scope for efficiency gains and thus, no excess burden to be calculated with respect to current policy. This is no less true of standard analyses of excess burden, such as the burden of distortionary taxation: One does not need a model of government determination of marginal tax rates to measure the excess burden of distortionary taxation. Precise modeling of the politics leading to a particular configuration of distortionary taxes implicitly assumes away the ability to eliminate distortions using alternative taxes.

The second issue is the point 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

---

[2] One explanation for delaying decisions is that older generations, who control political 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 (2007) for an analysis of how uncertainty in the labor earnings of adult children will lead their parents to delay deciding how well to treat them.

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]

A final issue is that of variable labor supply, which we have not included in our model, although hope to include in future work. Including variable labor supply is a non-trivial task, presenting its own computation and calibration challenges. It is possible that this would reduce the magnitude of the excess burden from government indecision because it would give agents additional ways to adjust to late resolution of uncertainty. However, since our main goal is to get a sense of the maximum size of the potential distortion from government indecision, the inclusion of variable labor supply would not alter our findings concerning the potential magnitude of excess burden in question since we would be free, in such an expanded model, to choose consumption-leisure preference parameters that returned us to our current setup, i.e., that entailed very little response of labor supply to expected or unexpected changes in income or incentives.

---

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

We proceed in section 1 by laying out a simple model of policy indecision and clarifying the source of its excess burden. Section 2 introduces our life cycle model. Section 3 discusses its calibration. Section 4 simulates the model assuming uncertainty about future retirement income, 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 5 repeats these analyses except it assumes the uncertainty involves future tax rates rather than retirement benefit levels. Section 6 jointly considers uncertainty about future retirement income and tax rates and calibrates the benefits of joint early resolution of these uncertainties. Finally, section 7 summarizes and concludes.

2   A Stylized Model of Policy Delay

Consider an agent who lives between time 0 and time T and has initial assets A_0 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 $L \leq T$ whether she gets high or low benefits, B, per period in retirement. The receipt of these benefits begin at time R. Initial assets are A_0, and consumption preferences are CRRA.

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, L, of the announcement of future benefits. Since there is no uncertainty about future benefits after L and the time preference and interest 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

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

5

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

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

The first order condition is

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

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

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

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

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

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

The expression in parenthesis 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 equation (7), with $E(z) \equiv p\overline{C}^{-\gamma} + (1-p)\underline{C}^{-\gamma}$, implies that $\partial EU/\partial L < 0$ when $\gamma < 1$, and positive when $\gamma > 1$.

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

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

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

$$
\begin{aligned}
-\log\left(p\overline{C}^{-1} + (1-p)\underline{C}^{-1}\right) &< p \log \overline{C}^{-1} + (1-p)\underline{C}^{-1} \\
&< -\left(p \log \overline{C} + (1-p)\underline{C}\right),
\end{aligned}
$$

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.

---

[4]Note that
$$\frac{\partial f(x)}{\partial x} = \left(1 - \frac{1}{\gamma}\right)x^{-1/\gamma},$$
and
$$\frac{\partial^2 f(x)}{\partial x^2} = -\frac{1}{\gamma}\left(1 - \frac{1}{\gamma}\right)x^{-(1+1/\gamma)}.$$

3   A Life-Cycle Model of Policy Delay

As indicated, our calibrated model features four types of uncertainty – earnings, longevity, returns, and government policy.   The policy uncertainty, which involves the level of government-provided retirement income as well as the level of labor income taxation, is resolved either prior to or at retirement. We use this model to study the effects on consumption, portfolio choice, and welfare of changes in the age at which uncertainty about future retirement income, uncertainty about future tax rates, or uncertainty about both future retirement income and future tax rates is resolved. In our model agents not only choose how much to save each period.   They also decide how much of their savings to allocate each period to a riskless bond and a risky asset. These decisions are made subject to liquidity constrainst and take into account "off-the-top" housing expenses and changes over time in household size due to the initial presence in the household of children and their subsequent departure from the household.

3.1   Model Specification

3.1.1   Time parameters and preferences

Let t denote age and assume agents work their first K periods and live for a maximum of T periods. We allow for lifespan uncertainty in the manner of Hubbard, Skinner and Zeldes (1995). Let p_t denote the probability that the investor is alive at date $t+1$, conditional on being alive at date t. The investor's preferences over consumption are given by

$$E_1 \sum_{t=1}^{T} \delta^{t-1} D_t \left(\prod_{j=0}^{t-2} p_j \right) \frac{C_t^{1-\gamma}}{1-\gamma}, \tag{8}$$

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

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

The term D_t captures the change in household size when adult children leave the household. We calibrate $\overline{D}$ to produce a 30 percent drop in household consumption at age 45 ($t = 18$).

3.1.2 The labor income process and retirement income

Before retirement ($t \leq K$), age-t labor income, Y_t, is exogenously given by the sum of a deterministic component that is calibrated to capture the hump shape of earnings over the life cycle and two random components, one transitory and one permanent.[5] More precisely,

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

where $f(t)$ is a deterministic function of age, and 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

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

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

We model government-provided retirement income as a fraction λ of permanent labor income in the last working-year. Specifically,

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

This specification facilitates the model's solution by reducing by one the number of state variables. Section 5 introduces uncertainty about λ and studies the implications for welfare, optimal consumption, and optimal portfolio choice of learning about λ at different ages.

It is important to note that our formulation (13) of the process for government-provided retirement income is equivalent to setting retirement income as a fraction of lifetime earnings. Since v_t follows a random walk, this implies that labor income at retirement is set

---

[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.

equal to a fraction λ of $\sum_{t=t_0}^{K} u_t$. Thus, ignoring the transitory income shocks, retirement income is equal to a fraction of all individual income realizations from age t_0 until age K. Alternatively, one can view retirement income in our model as proportional to average lifetime income by simply multiplying λ by $(K - t_0)$ and dividing lifetime income by $(K - t_0)$.

3.1.3 Financial assets

There are two assets – one risky and one riskless. The riskless asset, which we call *Treasury bills* or *bonds* indistinctly, has a constant gross real return of $\overline{R}_f$. We denote the dollar amount of T-bills the investor has at time t by B_t. The risky asset, which we will call *stocks*, has 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. We allow stock returns to be correlated with innovations to the shock to permanent labor income (u_t), and write the correlation coefficient as ρ. We denote the dollar amount the investor holds in stocks at time t by S_t.

We assume the investor faces the following borrowing and short-sales constraints

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

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

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

3.1.4 Taxes

We assume flat 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. We calibrate these tax rates to match the effective income tax rates currently faced by a typical household. Section 4.2 discusses the calibration of these tax rates. Section 6 introduces uncertainty about future tax rates and studies the implications for welfare, consumption, and portfolio choice of learning about future tax rates at different ages.

3.2   The investor's optimization problem

The investor starts the period with wealth W_t. Then labor income Y_t is realized. The purchase of durable goods and homes early in life constitutes "off-the-top" spending from the perspective of our model. Although modelling durable purchases directly is beyond the scope of the paper, it's important to include this spending because it affects the likelihood that the household will be liquidity constrained. We model the percentage of household income that is dedicated to housing expenditures (h_t) as an exogenous process and subtract it from the measure of disposable income.

Following Deaton (1991) we denote *cash-on-hand* in period t by X_t:

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

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 T-bills. 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 of the problem are $\{C_t, \ \alpha_t\}_{t=1}^T$. The state variables are $\{t, \ X_t, v_t\}_{t=1}^T$. Given the set up, the value function is homogeneous with respect to current permanent la-

11

bor income. Exploiting this scalability allows us to normalize v_t to one and to reduce the dimensionality of the state space.

The Bellman equation for this problem is given by:

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

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

We solve this problem numerically via backward induction.

4 Calibration

4.1 Labor Income Process

The labor income profile is taken from Cocco, Gomes and Maenhout (2005). They estimate age profiles for three different education groups (households without high school education, households with high school education, but without a college degree, and finally college graduates), and we take the weighted average of the three. Cocco, et. al. (2005) also estimate the fraction of permanent income replaced by retirement income λ to be 83 percent. In sections 5 and 7, where λ is uncertain, we set the probability distribution of λ so that its median is equal to 0.8, which is close to their value. In section 6, where there is no uncertainty about future retirement income, we fix λ at a constant 0.8.

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

---

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

(2005).

4.2 Other Parameters

We assume that agents are initially age 28, retire at 65, and die with probability one at age 100. Prior to this age we use the mortality tables of the National Center for Health Statistics to parameterize the conditional survival probabilities, p_j for $j = 1, ..., T$. We choose age 28 as the initial age of adulthood to roughly match the age at which working Americans marry and start having children. For example, the average U.S. ages of first marriage are 27 for males and 25 for women. Age 25 is also the average age of first birth. Our results are not sensitive to the starting age of the optimization, as long as agents accumulate the same levels of wealth by the time that uncertainty is resolved.

We set the discount factor δ to 0.95 and the coefficient of relative risk aversion γ to 5. This choice of γ is well within what is considered a reasonable upper bound for this parameter (Mehra and Prescott 1985), and consistent with the choice of this parameter in the literature on life-cycle investing. Nevertheless, we include in every table of our calibration a sensitivity analysis with respect to this parameter, which considers values above and below our baseline value.[7]

The mean equity premium (in levels) is set at 4.00 percent per annum, the risk-free rate is set at 1.00 percent per annum, and the annualized standard deviation of innovations to the risky asset is set at 20.5 percent. This equity premium is lower than the historical equity premium based on a comparison of average stock and T-bill returns, but it's 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 percent and the tax rate on retirement income (τ_R) to 15 percent during retirement (τ_R). Asset income

---

[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 parameter. Results are available upon request from the authors.

is taxed at a 20 percent rate (τ_C). These rates roughly match the effective income tax rates currently faced by a typical household (Kotlikoff and Rapson, 2005). In section 6 we incorporate uncertainty in these rates.

5   Model with uncertainty about retirement income

5.1   Retirement income uncertainty

As indicated in (13), 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 only knows that the realization of λ is governed by the following distribution

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

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 $1/3$ 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 previously mentioned, we set $\overline{\lambda}$ equal to 0.8, which is close to the 83 percent value in Cocco et al. (2005). We assume that this uncertainty remains 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 choice of a 30 percent future cut in future U.S. Social Security retirement benefits in our baseline calibration seems quite plausible given the magnitude of the current value budgetary shortfall facing Social Security as reported in the 2006 Social Security Trustees Report.[8]   Indeed, this report indicates that benefits would need to be cut immediately

---

[8]See http://www.ssa.gov/OACT/TR/TR06/index.html.

by roughly one fifth under what appear to be highly 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 percent benefit cut, particularly if the Social Security Actuaries' "high cost" assumptions materialize.

Although Social Security is viewed as the "third rail" of politics—any administration that touches it may face prospects of a sure defeat in the next election—, the reality is that major Social Security cuts have been enacted in the past. Indeed, 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 cuts of 26 percent for median-income workers retiring in 2030. For upper income workers, these cuts, that were promulgated in the past, 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 percent future cut to Social Security benefits as eminently reasonable, we want to make one thing clear: our objective is not to precisely estimate the excess burden from government indecision about Social Security policy. Our objective is much more limited. It is simply to understand whether government delay in reaching major policy decisions could potentially generate a significant excess burden. In saying this 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 sizes of benefit cuts. We do not report results for different probabilities of a benefit cut because they are very similar; i.e., a higher chance of a cut of a given size, or the same chance of a higher sized cut are very close substitutes.

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" is economically

15

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 U.S. fiscal gap is stunningly large, raising the possibility that Social Security benefit cuts will be needed to fix more than just what is labelled as "Social Security."

5.2 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 $A = 65$. Figure 1 shows the life-cycle pattern of financial wealth, income and optimal consumption generated by the baseline model. The units in this figure 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.

Figure 1 shows that both household consumption and portfolio allocations exhibit an inverted 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 captures the effect on household consumption of 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 and in response to the uncertainty about retirement benefits, which won't be resolved until the household retires at age 65. This saving, coupled with an aggressive allocation to stocks, allows the household to accumulate assets rapidly.

Interestingly, figure 2 shows that the allocation to stocks reaches a peak of about 98 percent at age 47 and declines until age 65, when the household allocates about 65 percent of its financial wealth to stocks. It then increases the stock allocation again, which eventually reaches 100 percent 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.

5.3 The benefits of early resolution of uncertainty about retirement income

5.3.1 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

17

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 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 gives a detailed description of the procedure we use to compute our welfare metric.

The rows labelled "Baseline" in table 1 report the 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 an annual fee equivalent to 0.117 percent of annual consumption in order to learn at age 35 the income replacement ratio it will experience at retirement. With consumption averaging about $30,000 per annum 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 –, or an annual fee which is similar to the annual expense ratio on a typical index mutual fund. Even at ages as late as 50 and 55, this household is still willing to pay 0.084 percent and 0.056 percent of its annual consumption in order to eliminate the 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 at age 35 its income replacement ratio. 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.

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 percent of the replacement income ratio instead of 30 percent — or in other words, when the replacement ratio is 48 percent of permanent income instead of 56 percent in the event of a cut —, 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 percent instead of 30 percent, a household is willing to pay 0.394 percent 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 percent fee the household is willing to pay when the size of the potential cut is 30 percent. 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, which range between 43 percent and 50 percent are larger when σ_u is 15 percent than when it is 10.95 percent.

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 non-participation 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 only adjust consumption — or equivalently saving — in response to learning early about retirement benefits. Panel D explores this scenario. It shows that welfare gains are about 20 percent larger when the household is unable to invest in stocks. In other words, there is a 20 percent 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 percent cut in the income replacement ratio at retirement is willing to pay 0.666 percent of annual consumption in order 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 percent larger than the effect when risk aversion is 5, and 93 percent larger than the effect when the potential cut is 30 percent. 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 percent potential benefit cut, the welfare gain to learning at age 35, rather than 65, whether or not the cut will happen is the equivalent of 0.733 percent of annual consumption.

5.3.2   Effect on consumption and portfolio choice

Table 1 shows that households are willing to pay a non-trivial 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 who 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 who 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 percent 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 who does not learn its income-replacement ratio until age 55 does not change consumption significantly relative to a household who learns at age 28. However, a household who learns only at age 65, when it retires, will consume between age 28 and age 65 about 0.4 percent *per year* less than a household that learns at age 28. A 0.4 percent reduction in annual consumption over 37 years is not extremely large, but it is still economically significant.

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 who does not learn until age 65

---

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

reduces consumption by about 0.9 percent per year relative to a household who 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 percent per year relative to a household who learns at age 28 when the potential cut in the replacement ratio is 45 percent. 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 reduction 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]

6 The benefits of early resolution of uncertainty about tax rates

Section 5 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. Throughout this section we assume that the income replacement ratio λ is known in advance and equal to 0.8.

6.1 Labor income tax uncertainty

We assume that the new tax rate takes effect at age 50 but, as in sections 5.1 and 5.2, the exact value is unknown until some age A (with $A \leqslant 50$). The tax rate uncertainty applies

---

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

to both the tax rate during working life and during the retirement period. Both face the same level of uncertainty and both are revealed at exactly the same time.

We first consider a symmetric uncertainty case:

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

and the same for τ_R.

In our baseline experiment the labor income tax rate might increase or decrease with an equal 25 percent probability ($p = 0.25$) or remain constant with a 50 percent probability. As in section 5.2, 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 percent during working life (τ_L) and 15 percent during retirement (τ_R). As a result, in the baseline uncertainty case, the labor income tax rate may increase (decrease) from 30 percent to 39 percent (21 percent) between ages 50 and 65 and from 15 percent to 19.5 percent (10.5 percent) between ages 65 and 100 with a 25 percent probability.

6.2   Welfare analysis and consumption distortions

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

---

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

D considers the case in which the uncertainty about the income tax rate is asymmetric. In the asymmetric case, the tax rate can only increase (decrease) by 30 percent from 30 percent to 39 percent (19.5 percent) during working life (retirement) with a probability of 1/3. 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 percent of annual consumption, and this increases to 0.084 percent 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 percent of consumption between age 35 and age 50 for eliminating at age 35 a 30 percent uncertainty over 30 percent of its income between ages 50 and 65, and 15 percent of its income between ages 65 and 100. This is essentially equivalent to eliminating a 9 percent uncertainty over 100 percent of the household's income between ages 50 and 65, and a 4.5 percent uncertainty over 100 percent of the household's income between ages 65 and 100.

By contrast, the uncertainty that it is eliminated in the retirements benefits case is proportionally much larger, since it implies a 30 percent uncertainty over 100 percent of the household's income between ages 65 and 100. Accordingly, the welfare gain from knowing early is also proportionally larger—about 0.12 percent. 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, similar 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 E 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

24

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 who bear the longest delay in learning.

6.3   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 the its working life, to smooth the impact of income shocks on consumption.

Our baseline interest rate is 1 percent, and our baseline real return on equities is 5 percent. Our baseline 20 percent capital income tax rate implies that the after-tax return on the household portfolio is between 0.8 percent and 4 percent. A 30 percent 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 percent of labor income caused by a chance of a 30 percent change in the 30 percent baseline labor-income tax rate. Thus, a change in the labor income tax rate or a change in retirement income are likely to have a more significant impact on future consumption than a change of similar magnitude in the effective capital-income tax rate.

7 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 both be manifest, particularly if they result from situations of fiscal crises. Hence, it's interesting to explore the implications for welfare, optimal consumption, and portfolio choice 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 that both the baseline 30 percent labor-income tax rate increases by 30 percent at age 50 and the 80 percent replacement ratio of permanent income in retirement declines by 30 percent 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. (Section 6 explores the case where λ is known to the household in advance).

7.1 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 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

26

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 who can only invest in bonds; 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 percent per annum. 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 at large – 0.069 percent per annum.

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 non-equity investors are now larger than the gains for households who 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 non-capital income.

7.2 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 who do not resolve this uncertainty until age A consume significantly less per year than households who resolve it early in their life cycle—at age 28. For example, in our baseline model a household that does not solve this uncertainty until age 50 consumes about 1.3 percent per annum less than an identical household who learns at age 28. For a household with coefficient of relative risk aversion of 7, the reduction in optimal consumption is 2.5 percent. Even at age 35, where the delay is only one of seven

years, the reduction in consumption is 0.1 percent per annum in the baseline case and 0.8 percent per annum in the case with a coefficient of relative risk aversion of 7. In general, the fall in consumption is largest for households who are more risk averse and for households who face large possible falls in income.

8 Conclusions

This paper has explored the effects on life-cycle saving, asset allocation, 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 equivlaent to 0.12 percent 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, or households facing either larger potential cuts in benefits or increases in marginal tax rates.

There are at least two extensions to our initial analysis of this excess burden that seem worth pursuing. The first is including variable labor supply, including the choice of retirement dates. Doing so, as mentioned in the introduction, would likely reduce the size

of the distortion we've measured since households would have more ways of responding to the added uncertainty associated with policy determination delay. The second extension is to consider other types of policy uncertainties, including the possibilities of the government switching tax structures (e.g., from income to consumption taxation), cutting back on healthcare benefits, and printing money to pay for its bills, leading potentially to very high rates of inflation and overall financial and economic instability.

To conclude, no one likes to deliver bad news, least of all governments. The likelihood of being the messenger that gets shot leads many politiicans to postpone difficult decisions. Whatever politicians gain from this behavior, it fosters and exacerbates economic uncertainty. In this paper, we've begun to model and quantify the excess burden of government indecision. As we've stressed, this excess burden arises not from implementing specific policies, but from delaying their determination and announcement. We've 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.

9 References

Altonji, Joseph, Fumio Hayashi, and Laurence J. Kotlikoff. 2007. "Parental Altruism and Intergenerational Transfers: Theory and Evidence," *The Journal of Political Economy*, December.

Auerbach, Alan J. and Kevin A. Hassett. 2002. "Fiscal Policy and Uncertainty," *International Finance*, v5(2,Summer),229-49.

Auerbach, Alan J. and Laurence J. Kotlikoff, 1987. *Dynamic Fiscal Policy*. Cambridge, England: 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," University of California, Santa Barbara, January.

Campbell, John, Joao Cocco, Francisco Gomes, and Pascal Maenhout, 2001. "Investing Retirement Wealth: A Life Cycle Model." In *Risk Aspects of Social Security Reform* (The University of Chicago Press).

Carroll, Christopher. 1992. "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence." *Brookings Papers on Economic Activity* 2: 61-156.

Carroll, Christopher. 1997. "Buffer stock saving and the life cycle/permanent income hypothesis." *Quarterly Journal of Economics* 112, 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*, August, 617-52.

Cocco, Joao, Francisco Gomes and Pascal Maenhout. 2005. "Consumption and Portfolio Choice over the Life-Cycle." *The Review of Financial Studies*, vol. 18 (2), 491-533.

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, LEERNA-INRA, University of Toulouse.

Fama, Eugene F. and Kenneth R. French, 2002, "The Equity Premium," *Journal of Finance*, Vol. 57, pp. 637-659.

Gokhale, Jagadeesh and Kent Smetters. 2005. "Measuring Social Security's Financial Problems," NBER working paper no. 11060.

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," NBER working paper no. 12344, June.

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, 360-399.

Judd, Kenneth. 1998. "Optimal Taxation in Dynamic Stochastic Economies." Hoover Institution.

Kotlikoff, Laurence J., *Generational Policy*, MIT Press, 2002.

Kotlikoff, Laurence J. and David Rapson, "Comparing Average and Marginal Tax Rates Under the FairTax and the Current System of Federal Income Taxation," NBER working

paper 11831, December 2005.

Munnell, Alicia H. 2003. "The Declining Role of Social Security," JFT #6, The Center for Retirement Research at Boston College.

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*, pp. 325-45.

Shoven, John B. and Sita N. Slavov, "Political Risk Versus Market Risk in Social Security," NBER working paper no. 12135, March 2006.

Appendix: Welfare Metric

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/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 towards 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 life-time utility:

$$V^A = E_1 \sum_{t=1}^T \delta^{t-1} \left(\prod_{j=0}^{t-1} p_j \right) \frac{C_t^{A^{1-\gamma}}}{1-\gamma}, \tag{22}$$

Note: This is just the value function from the maximization problem.

Then 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,

$$V^A = E_1 \sum_{t=1}^T \delta^{t-1} \left(\prod_{j=0}^{t-1} p_j \right) \frac{\overline{C}^{A^{1-\gamma}}}{1-\gamma}. \tag{23}$$

Therefore:

$$\overline{C}^A = \left[\frac{(1-\gamma)V^A}{\sum_{t=1}^T \delta^{t-1} \left(\prod_{j=0}^{t-1} p_j \right)} \right]^{\frac{1}{1-\gamma}}. \tag{24}$$

Taking $A = 65$ as our baseline case, the utility gain from changing A can then be

obtained simply computed as the percentage loss in equivalent consumption,

$$\frac{\overline{C}^A - \overline{C}^{65}}{\overline{C}^{65}} = \frac{V^{A\frac{1}{1-\gamma}} - V^{65\frac{1}{1-\gamma}}}{V^{65\frac{1}{1-\gamma}}}.$$ (25)

Figure 1. Consumption, Wealth and Income

Figure 1 shows the life-cycle pattern of financial wealth, income and optimal consumption generated by the baseline model, whose parameter values are given in section 4. 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.

**Figure 2. Portfolio Allocation**

Figure 2 shows the life-cycle pattern of the percentage portfolio allocation to stocks generated by the baseline model, whose parameter values are given in section 4. The pattern is an average based on 10,000 simulations of the baseline model.

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)						
$\gamma = 7$	0.345%	0.327%	0.298%	0.262%	0.211%	0.131%
$\gamma = 5$ (Baseline)	0.117%	0.115%	0.111%	0.100%	0.084%	0.056%
$\gamma = 3$	0.012%	0.012%	0.012%	0.012%	0.011%	0.009%
B. Cut in replacement ratio (ξ)						
$\xi = 45\%$	0.294%	0.290%	0.274%	0.243%	0.204%	0.153%
$\xi = 40\%$	0.251%	0.248%	0.237%	0.213%	0.176%	0.114%
$\xi = 30\%$ (Baseline)	0.117%	0.115%	0.111%	0.100%	0.084%	0.056%
$\xi = 20\%$	0.042%	0.042%	0.040%	0.037%	0.031%	0.021%
$\xi = 15\%$	0.021%	0.021%	0.020%	0.019%	0.017%	0.014%
C. Volatility of shocks to permanent income (σ_u)						
$\sigma_u = 15\%$	0.175%	0.169%	0.158%	0.143%	0.121%	0.080%
$\sigma_u = 10.95\%$ (Baseline)	0.117%	0.115%	0.111%	0.100%	0.084%	0.056%
$\sigma_u = 7\%$	0.028%	0.028%	0.028%	0.027%	0.024%	0.018%
D. Assets available to household						
Only bonds (Non-equity investors)	0.141%	0.140%	0.135%	0.122%	0.101%	0.066%
Bonds and equities (Baseline)	0.117%	0.115%	0.111%	0.100%	0.084%	0.056%
E. Combined scenarios						
$\gamma = 7$, $\xi = 40\%$, non-equity investors	0.733%	0.691%	0.624%	0.542%	0.430%	0.262%
$\gamma = 7$ and $\xi = 40\%$	0.666%	0.630%	0.571%	0.499%	0.398%	0.244%
$\gamma = 5$, $\xi = 30\%$, bonds and equities (Baseline)	0.117%	0.115%	0.111%	0.100%	0.084%	0.056%

Table 1 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 section 4. 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.

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%	0.0%	-0.1%	-0.3%	-0.4%	-0.6%	-0.9%
$\gamma = 5$ (Baseline)	0.0%	0.0%	0.0%	0.0%	-0.1%	-0.2%	-0.4%
$\gamma = 3$	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
B. Cut in replacement ratio (ξ)							
$\xi = 45\%$	0.0%	0.0%	0.0%	-0.2%	-0.4%	-0.6%	-1.3%
$\xi = 40\%$	0.0%	0.0%	0.0%	-0.1%	-0.2%	-0.4%	-0.9%
$\xi = 30\%$ (Baseline)	0.0%	0.0%	0.0%	0.0%	-0.1%	-0.2%	-0.4%
$\xi = 20\%$	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	-0.1%
$\xi = 15\%$	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	-0.1%
C. Volatility of shocks to permanent income (σ_u)							
$\sigma_u = 15\%$	0.0%	0.0%	0.0%	-0.1%	-0.1%	-0.2%	-0.5%
$\sigma_u = 10.95\%$ (Baseline)	0.0%	0.0%	0.0%	0.0%	-0.1%	-0.2%	-0.4%
$\sigma_u = 7\%$	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	-0.2%
D. Assets available to household							
Only bonds (Non-equity investors)	0.0%	0.0%	0.0%	0.0%	-0.1%	-0.2%	-0.5%
Bonds and equities (Baseline)	0.0%	0.0%	0.0%	0.0%	-0.1%	-0.2%	-0.4%
E. Combined scenarios							
$\gamma = 7$, $\xi = 40\%$, non-equity investors	-0.1%	-0.2%	-0.3%	-0.5%	-0.8%	-1.2%	-2.0%
$\gamma = 7$ and $\xi = 40\%$	-0.1%	-0.1%	-0.3%	-0.4%	-0.7%	-1.1%	-1.8%
$\gamma = 5$, $\xi = 30\%$, bonds and equities (Baseline)	0.0%	0.0%	0.0%	0.0%	-0.1%	-0.2%	-0.4%

Table 2 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 who 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 undistinguishable from approximation error. The table reports results for our baseline case as well as the cases we examine in Table 1.

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)			
$\gamma = 7$	0.084%	0.067%	0.037%
$\gamma = 5$ (Baseline)	0.030%	0.028%	0.020%
$\gamma = 3$	0.003%	0.003%	0.000%
B. Change in labor income tax rate			
$\xi = 45\%$	0.070%	0.066%	0.048%
$\xi = 30\%$ (Baseline)	0.030%	0.028%	0.020%
$\xi = 15\%$	0.007%	0.007%	0.005%
C. Assets available to household			
Only bonds (Non-equity investors)	0.030%	0.029%	0.022%
Bonds and equities (Baseline)	0.030%	0.028%	0.020%
D. Asymmetric change in tax rate			
Asymmetric 30% change	0.037%	0.030%	0.017%
Symmetric 30% change (Baseline)	0.030%	0.028%	0.020%

Table 3 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 who does not learn about the change until age A, relative to a household who 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 their permanent labor income they get in retirement. We consider four variants of the baseline experiment, which are reported in Panels A through 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 only increase (decrease) by 30% from 30 percent to 39% (19.5%) during working life (retirement) with a probability of 1/3. 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)				
$\gamma = 7$	0.0%	-0.1%	-0.1%	-0.3%
$\gamma = 5$ (Baseline)	0.0%	0.0%	0.0%	-0.2%
$\gamma = 3$	0.0%	0.0%	0.0%	-0.1%
B. Change in labor income tax rate				
$\xi = 45\%$	0.0%	-0.1%	-0.1%	-0.4%
$\xi = 30\%$ (Baseline)	0.0%	0.0%	0.0%	-0.2%
$\xi = 15\%$	0.0%	0.0%	0.0%	-0.1%
C. Assets available to household				
Only bonds (Non-equity investors)	0.0%	0.0%	0.0%	-0.1%
Bonds and equities (Baseline)	0.0%	0.0%	0.0%	-0.2%
D. Asymmetric change in tax rate				
Asymmetric 30% change	0.0%	0.0%	0.0%	-0.1%
Symmetric 30% change (Baseline)	0.0%	0.0%	0.0%	-0.2%

Table 4 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 who does not learn about the change until age A, relative to a household who 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 undistinguishable from approximation error. The table reports results for the same cases we examine in Table 3.

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)			
$\gamma = 7$	0.247%	0.194%	0.107%
$\gamma = 5$ (Baseline)	0.076%	0.069%	0.047%
$\gamma = 3$	0.004%	0.004%	0.004%
B. Size of cut in retirement benefits and increase in labor income tax rate			
$\xi = 45\%$	0.246%	0.220%	0.140%
$\xi = 30\%$ (Baseline)	0.076%	0.069%	0.047%
$\xi = 15\%$	0.013%	0.012%	0.008%
C. Assets available to household			
Only bonds (Non-equity investors)	0.086%	0.080%	0.055%
Bonds and equities (Baseline)	0.076%	0.069%	0.047%
D. Uncorrelated uncertainty			
Uncorrelated ξ	0.038%	0.033%	0.022%
Baseline	0.076%	0.069%	0.047%

Table 5 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 1/3 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 2/3 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 section 4. 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 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 E considers a special case where 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 future labor income tax rates when the income replacement ratio is uncertain and not known until age 65.

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)				
$\gamma = 7$	-0.8%	-1.4%	-1.9%	-2.5%
$\gamma = 5$ (Baseline)	-0.1%	-0.3%	-0.5%	-1.3%
$\gamma = 3$	0.0%	0.0%	0.0%	-0.5%
B. Size of cut in retirement benefits and increase in labor income tax rate				
$\xi = 45\%$	-0.2%	-0.6%	-1.1%	-2.5%
$\xi = 30\%$ (Baseline)	-0.1%	-0.3%	-0.5%	-1.3%
$\xi = 15\%$	0.0%	-0.1%	-0.2%	-0.5%
C. Assets available to household				
Only bonds (Non-equity investors)	-0.1%	-0.3%	-0.5%	-1.4%
Bonds and equities (Baseline)	-0.1%	-0.3%	-0.5%	-1.3%
D. Uncorrelated uncertainty				
Uncorrelated ξ	0.0%	-0.2%	-0.4%	-1.1%
Baseline	-0.1%	-0.3%	-0.5%	-1.3%

Table 6 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 who resolves this uncertainty at age 28. The table reports percentage changes in consumption only to the first decimal digit, because digits beyond that are undistinguishable from approximation error. The table reports results for the same cases we examine in Table 5.