

The Cost of Uncertainty about the Timing of Social Security Reform*

Frank N. Caliendo[†] Aspen Gorry[‡] Sita Slavov[§]

This Version: April 9, 2015. First Version: November 5, 2013.

Abstract

We develop a methodology to study how uncertainty about the timing of the resolution of uncertainty affects economic decision making and welfare. As an example of the power of our framework, we analytically solve and simulate household decisions and welfare in a setting in which households don't know *how* or *when* Social Security is going to be reformed. Uncertainty about the timing and structure of Social Security reform is highly regressive: the very low income group from the 2013 Trustees Report (who earn 25 percent of the economy-wide average wage) experience a welfare loss that is about 3 times larger than those who earn the maximum taxable amount in each year of their career.

*We thank Shantanu Bagchi, Peter Bossaerts, Eric Fisher, Jim Feigenbaum, Carlos Garriga, Nick Guo, Tim Kehoe, Jorge Alonso Ortiz, Kerk Phillips, Jennifer Ward-Batts, Jackie Zhou, and seminar audiences at ITAM, UConn, Midwest Macro at Minnesota, Utah State University, BYU-USU Macro Workshop, University of Nevada Reno, and the 2014 WEAI conference for helpful comments and suggestions.

[†]Utah State University. Email: frank.caliendo@usu.edu.

[‡]Utah State University. Email: aspen.gorry@usu.edu.

[§]George Mason University. Email: sslavov@gmu.edu.

1. Introduction

How does uncertainty about the timing of the resolution of uncertainty affect economic decision making and welfare? Some researchers have considered the implications of uncertainty that is resolved at different time horizons, but the timing of the resolution of uncertainty is typically known in advance.¹ We are interested in situations in which decision makers not only face uncertainty about the outcome of a future event, but the timing of that event is also uncertain.² We use the term structural uncertainty to describe uncertainty over the outcome of a future event, and we use the term timing uncertainty to describe uncertainty about the timing of that future event. For example, uncertainty about Social Security reform has both a structural component and a timing component. Young Americans who must make consumption and saving plans for retirement don't know *how* or *when* Social Security will be reformed.

We develop a dynamic setting to study decision making in the face of both structural and timing uncertainty. Our setting has three main features:

First, with two layers of uncertainty instead of just one, our analysis differs from a related literature that seeks to understand how the timing of the resolution of uncertainty affects decision making. Rather than modeling timing uncertainty, this literature focuses on understanding how early versus late resolution of uncertainty affects decision making. In each case, the timing of the resolution of uncertainty is known in advance.³

¹Long-run risk models are an example. Epstein et al. (2014) provide an overview of this literature and they compute the “timing premium” under recursive preferences, which is the amount households would be willing to pay to resolve uncertainty about future consumption right away. But there is no timing uncertainty in their setting. Households always know in advance whether they are living in the early-resolution world or the late-resolution world, and then the experiment is to calculate how much households in the second world would pay to live in the first.

²Such situations arise naturally throughout economics, in the areas of policy uncertainty, sudden stops, regime switches, and more.

³Blundell and Stoker (1999) and Eeckhoudt et al. (2005) study the connection between optimal consumption and the timing of income risk. They compare the case in which income risk gets resolved early to the case in which income risk gets resolved late, but in each case the household knows when the income risk gets resolved. Likewise, Wright et al. (2014) extend these ideas to firm investment behavior. They seek to understand the differential effects of long and short term uncertainty on firm investment policy, but firms in their model don't face uncertainty about when this uncertainty gets resolved.

Second, we allow for any distribution over the resolution of timing uncertainty. This allows the framework to handle non-stationary distributions of the timing of events, which contrasts with the standard stochastic dynamic programming setting in which uncertainty is represented as a Markov process and hence is stationary and depends only on the state. Our methodology is flexible enough to not only handle the Markovian case in which the distribution over a shock date is time invariant but also the case in which such a distribution is a function of the date itself. For instance, it may be that one can characterize uncertainty about the timing of Social Security reform as a state dependent process, but it is convenient to specify the distribution of reform dates as a function of time itself. This allows us to quickly consider a wide variety of distributions without having to reconsider the appropriate state space.⁴ We do this by combining and generalizing existing tools from the two-stage optimal control literature that deal independently with either structural uncertainty or with timing uncertainty but not both at once.⁵

Third, our analysis is set in a rich enough environment to obtain quantitative answers to important policy questions. We assume time is continuous because we want decision makers in our setting to face structural uncertainty that can strike at any moment, as in the case of Social Security reform.

To show the power of this framework, we tackle the important policy example that we have alluded to already. It is well known that the Social Security Old Age, Survivors, and Disability Insurance (OASDI) program faces severe long run solvency concerns. The 2013 Social Security

⁴We are interested in a variety of possible distributions, such as the stationary case with a constant hazard rate of reform, the non-stationary case with a spike in the likelihood of reform near the date at which the trust fund is projected to run out of money, and the non-stationary case in which the likelihood of reform increases with each passing day.

⁵Examples of two-stage control problems with a stochastic regime switch date (timing uncertainty) can be found in the early studies on resource extraction (Dasgupta and Heal (1974)), operations research (Kamien and Schwartz (1971)), and environmental catastrophe (Clarke and Reed (1994)). And examples of problems with uncertainty about the characteristics of the new regime (structural uncertainty) appeared in the early resource extraction literature (Hoel (1978)) and later in the technology adoption literature (Hugonnier et al. (2006), Pommeret and Schubert (2009), Abel and Eberly (2012)). Our theoretical contribution is to build a bridge between these literatures by generalizing these specific examples and summarizing how to nest both layers of uncertainty together in a single control problem.

Trustees Report projects that the program’s trust fund will run out of money in the year 2033.⁶ This means that over the coming decades either promised retirement benefits must be cut or the payroll taxes used to fund them must be increased to keep the program solvent. Gokhale (2013) estimates that an immediate tax increase of 3.1% of taxable wages, or an immediate benefit cut of 21%, will keep the OASI part of the program (which we focus on in this paper) solvent for the infinite horizon.⁷

While there is very little uncertainty about the need for reform, there is a great deal of uncertainty surrounding its timing and structure (Sargent (2005)). And even though the feasibility/optimality of various Social Security reforms have been widely studied in macroeconomic models (e.g., McGrattan and Prescott (2014) and Kitao (2014) and the references therein), relatively little attention has been paid to the microeconomic welfare consequences of uncertainty surrounding the timing and structure of reform. The macroeconomics literature treats the timing and structure of reform as part of the household information set, which allows researchers to study the effects of particular reforms in general equilibrium, while we are interested in the separate issue of measuring welfare losses for households making decisions in an environment where general equilibria cannot be constructed because the timing and structure of reform are unknown.⁸

We use our method to characterize and simulate household consumption/saving decisions and welfare under uncertainty about the structure and timing of Social Security reform. We parameterize the model to match individual wage and survival profiles. The parameterized model allows us to assess the welfare cost of uncertainty surrounding the timing and structure of Social Security reform. We measure this cost as the fraction of lifetime consumption that a young individual facing reform uncertainty over the life cycle would be willing to give up to live in a separate world with no reform uncertainty and with endowed wealth that is equal to expected wealth over all

⁶Throughout the paper we refer to the 2013 Trustees Report because the 2014 Trustees Report doesn’t have replacement rates for hypothetical earners or the estimated benefit cut required to balance the long-term budget.

⁷If the disability component of the program is also included, then the required tax increase (assuming no behavioral response) is 4% and the required benefit cut is 23.9% as estimated in the 2013 Trustees Report.

⁸While Bütler (1999) studies a particular example of the implications of uncertainty about both the timing and structure of social security reform in Switzerland, our methodology allows us to handle any assumptions about the distributions of both structural and timing uncertainty. This flexibility allows us to study timing and structural uncertainty in a general way and to explore the distributional effects of reform in the US.

possible dates and structures of reform. Endowing the individual in the no-uncertainty world in this manner acts to net out changes in wealth that come from reform itself, leaving just the pure cost of uncertainty about reform.⁹

Because Social Security benefits and taxes vary with income, the impact of Social Security reform uncertainty may vary across income groups. To be more specific, Social Security is financed by a proportional payroll tax up to a taxable maximum wage (\$113,700 in 2013).¹⁰ But benefits are paid according to a progressive formula that gives individuals with low lifetime earnings higher replacement rates (benefits as a share of average lifetime earnings). Since all individuals face the same tax rate (at least up through the taxable maximum income), but receive different replacement rates, we can vary the replacement rate in our model to simulate the effect of Social Security reform uncertainty for individuals of different incomes.

Uncertainty about the timing of reform is considered first because timing uncertainty is harder to deal with theoretically and because less is known about its welfare consequences. In this scenario, the individual has full information about the structure of reform and we consider the possibility of either a known benefit cut or a known tax increase occurring at an unknown date.

Uncertainty about the timing of reform has moderate costs to individuals of average income, a result which holds regardless of whether reform takes the form of across-the-board tax increases or across-the-board benefit cuts. However, tax increases and benefit cuts have different effects for low-income and high-income individuals. If individuals know that the Social Security budget will

⁹Bütler (1999) also attempts to compute the welfare losses from uncertainty about the timing of social security reform. However, she compares the welfare of individuals who must live with timing uncertainty to individuals who live in a separate world with no timing uncertainty and with reform that is guaranteed to happen at the date that is mathematically expected in the first world. While such a comparison would certainly capture important differences in behavior, it does not actually pin down the cost of uncertainty. Instead, it confounds changes in wealth with the effects of uncertainty. In fact, with such a comparison it is theoretically possible to come to the mistaken conclusion that uncertainty actually is a good thing (because it increases an individual's expected lifetime wealth). For instance, if we assume that individuals in the first world are uncertain about when a benefit cut will strike and the mathematical expectation is that it will strike just before retirement, then individuals would want maximum variance around this expectation in order to create the possibility of collecting full benefits for at least some portion of the retirement period.

¹⁰The taxable maximum is \$118,500 in 2015, but again we focus on 2013 values throughout the paper because the 2013 Trustees Report is more comprehensive.

be balanced using an across-the-board benefit cut, uncertainty about the timing of reform harms low-income individuals more because Social Security accounts for a larger portion of their lifetime wealth. Alternatively, when the budget will be balanced using an across-the-board tax increase, high-income individuals earning the taxable maximum face larger welfare losses because they face the same percentage tax increase but rely less on the progressive benefit.

In addition to timing uncertainty, individuals also do not know the structure of future reform. The government may balance the budget through tax increases, benefit cuts, or some combination of the two. After adding this second layer of uncertainty to our model, the overall welfare cost of the combined reform uncertainty can be greater or less than the cost for the case where there is only timing uncertainty. However, we find that this double uncertainty consistently imposes much larger welfare costs on low-income individuals across all parameterizations of the model. The very low income group from the 2013 Trustees Report (who earn 25 percent of the economy-wide average wage) experience a welfare loss that is about 3 times larger than those who earn the maximum taxable amount in each year of their career.

This paper is related to Gomes et al. (2007), Benítez-Silva et al. (2007) and Luttmer and Samwick (2012) who also attempt to quantify the costs of Social Security reform uncertainty. Gomes et al. (2007) study a model environment in which there is uncertainty about whether a benefit cut will occur at a given future date. Their baseline household would be willing to give up 0.12% of annual consumption in exchange for learning about the cut to Social Security benefits at age 35 instead of age 65. While this result has a similar flavor to ours, their model likely overstates the result by placing the uncertainty at the worst possible date (i.e., at retirement) rather than modeling the full uncertainty about the timing of reform. Moreover, it is important to note that their welfare calculation—how much households would pay to resolve uncertainty earlier rather than right at retirement—does not capture how households behave in the face of timing uncertainty. Households in their model never actually face timing uncertainty because the date at which information is released to the public is known in advance.^{11,12}

¹¹Gomes et al. (2007) also consider the gains from early resolution of uncertainty about the magnitude of a future tax increase. But once again households know in advance the future date at which the government will release the information, so there is no timing uncertainty in their model.

¹²Evans et al. (2012) consider a model in which future transfers to the old are uncertain because transfers are

Similar to Gomes et al. (2007), Benítez-Silva et al. (2007) also compute the welfare loss from having to live with uncertainty about the future level of Social Security benefits. There is no timing uncertainty in their model either. The magnitude of the welfare losses from uncertainty about the structure of reform are similar across papers.

Luttmer and Samwick (2012) use survey data to elicit the degree of Social Security reform uncertainty that individuals perceive, and how costly such uncertainty is to them. They find that, on average, individuals between the age of 25 and 59 expect to receive only 60% of the Social Security benefits that they have been promised. Moreover, individuals would be willing to tolerate an additional 4-6% cut in benefits in exchange for certainty about their level.

Finally, we depart from a literature that assumes that individuals don't fully understand the rules and incentives of Social Security taxation and benefit calculations. In that literature, individuals have incomplete information because information is costly to acquire or because they lack the ability to process information, while the political process is stationary and the rules are knowable (see Liebman and Luttmer (2014) for additional discussion). In contrast, in our setting individuals are fully rational and have full information about everything that is knowable, which includes the distributions of random variables relating to political risks, while the realization of these random variables is unknowable. This implies that our estimates of the cost of uncertainty about Social Security reform are likely a lower bound.¹³

based on the stochastic wages of the young. If promised transfers are infeasible because of a negative productivity shock, the government switches to a more affordable transfer scheme. Households anticipate this risk, which in turn affects the equity premium. Likewise, van der Wiel (2008) considers the effect of uncertainty about future Social Security benefits on private savings, but similar to Gomes et al. (2007), the individual knows in advance that the government will announce the new level of benefits at the date of retirement, so there is no timing uncertainty.

¹³Our paper is related to the growing literature on policy uncertainty more generally. Uncertainty about future policy can affect investment, saving, and hiring decisions (Rodrik (1989), Bernanke (1983)). The impact of policy uncertainty on firms' decisions and profitability can be substantial, potentially resulting in large efficiency costs. Several of these studies estimate or calibrate models in which taxes, government spending, or other policies are uncertain, and then explore the economic impacts of this uncertainty (Fernández-Villaverde et al. (2011), Croce et al. (2012), Hassett and Metcalf (1999), Pastor and Veronesi (2011), Sialm (2006), Ulrich (2012)). Others use the timing of national elections as a measure of policy uncertainty, examining the impact of uncertain election outcomes on economic behavior (Belo et al. (2011), Boutchkova et al. (2012), Durnev (2010), Julio and Yook (2012), Pantzalis et al. (2000)).

2. Theory: Dynamic Problems with Timing Uncertainty and Structural Uncertainty

In this section we provide a self contained summary of how to solve a generic dynamic problem that features both timing uncertainty and structural uncertainty. Problem 1 features timing uncertainty only. And Problem 2 features both timing uncertainty and structural uncertainty.¹⁴

Time is continuous and indexed by t . Time starts at $t = 0$ and never ends. The planning interval of the decision maker, which begins at $t = 0$ and ends at $t = T$, is comprised of two regimes or stages. The first stage stretches from $t = 0$ to $t = t_1$, and the second stage stretches from $t = t_1$ to $t = T$. Each stage has a unique performance index and/or a unique state equation. The regime switch date t_1 is a continuous random variable, with probability density $\phi(t_1)$ and sample space $[0, \infty]$. The control variable $u(t)$ is unconstrained and the state variable $x(t)$ is constrained only at the initial and terminal points in time.

Problem 1 (Timing Uncertainty Only).

$$\max_{u(t)_{t \in [0, T]}} : J = \mathbb{E} \left[\int_0^{t_1} f_1(t, u(t), x(t)) dt + \int_{t_1}^T f_2(t, u(t), x(t)) dt \right], \quad (1)$$

subject to

$$\frac{dx(t)}{dt} = g_1(t, u(t), x(t)), \text{ for } t \in [0, t_1], \quad (2)$$

$$\frac{dx(t)}{dt} = g_2(t, u(t), x(t)), \text{ for } t \in [t_1, T], \quad (3)$$

$$x(0) = x_0, \quad x(T) = x_T, \quad (4)$$

$$t_1 \text{ random with density } \phi(t_1) \text{ and sample space } [0, \infty]. \quad (5)$$

¹⁴Our analysis builds on the two-stage control literature in which the switch date is deterministic and may be a choice variable or exogenous (Kemp and Long (1977), Tomiyama (1985), Amit (1986), Tahvonen and Withagen (1996), Makris (2001), Boucekkine et al. (2004), Dogan et al. (2011), Saglam (2011), Boucekkine et al. (2012), Boucekkine et al. (2013a), and Boucekkine et al. (2013b)).

It is helpful to lay out some assumptions and definitions before stating the necessary conditions.

Assumption 1 (Differentiability). The functions $f_1, f_2, g_1, g_2,$ and $\phi(t_1)$ are continuously differentiable in their arguments.

Definition 1 (Solution Notation). Let $(u_2^*(t|t_1, x(t_1)), x_2^*(t|t_1, x(t_1)))_{t \in [t_1, T]}$ be the optimal control and state path for any t *after* the realization of the random regime switch, conditional on the switch date t_1 and conditional on the stock of the state variable at that switch date $x(t_1)$. Similarly, let $(u_1^*(t), x_1^*(t))_{t \in [0, t_1]}$ be the optimal control and state path for any t *before* the realization of the random switch. Hence, the path that is actually followed, conditional on switch date t_1 , is $(u_1^*(t), x_1^*(t))_{t \in [0, t_1]}$ and $(u_2^*(t|t_1, x_1^*(t_1)), x_2^*(t|t_1, x_1^*(t_1)))_{t \in [t_1, T]}$.

Theorem 1 (Necessary Conditions to Problem 1). The necessary conditions can be derived recursively in two steps.¹⁵

Step 1. Solve the post-switch ($t = t_1$) subproblem:

The program $(u_2^*(t|t_1, x(t_1)), x_2^*(t|t_1, x(t_1)))_{t \in [t_1, T]}$ solves a fixed endpoint Pontryagin subproblem

$$\max_{u(t)_{t \in [t_1, T]}} : J_2 = \int_{t_1}^T f_2(t, u(t), x(t)) dt, \quad (6)$$

subject to

$$\frac{dx(t)}{dt} = g_2(t, u(t), x(t)), \text{ for } t \in [t_1, T], \quad (7)$$

$$t_1 \text{ given, } x(t_1) \text{ given, } x(T) = x_T. \quad (8)$$

Given the Hamiltonian function

$$\mathcal{H}_2 = f_2(t, u(t), x(t)) + \lambda_2(t)g_2(t, u(t), x(t)), \quad (9)$$

¹⁵We refer readers to Appendix A for a full derivation of the theorem.

the necessary conditions that must hold on the path $(u_2^*(t|t_1, x(t_1)), x_2^*(t|t_1, x(t_1)))_{t \in [t_1, T]}$ include $\partial \mathcal{H}_2 / \partial u(t) = 0$ and $d\lambda_2(t)/dt = -\partial \mathcal{H}_2 / \partial x(t)$. For convenience, change the time dummy t to z , and change the switch point t_1 to t and write the solution $(u_2^*(z|t, x(t)), x_2^*(z|t, x(t)))_{z \in [t, T]}$. Thus we have the optimal control and state paths for all points in time z greater than switch point t .

Step 2. Solve the pre-switch ($t = 0$) subproblem:

The program $(u_1^*(t), x_1^*(t))_{t \in [0, T]}$ solves a fixed endpoint Pontryagin subproblem with continuation function $\mathcal{S}(t, x(t))$:

$$\max_{u(t)_{t \in [0, T]}} : J_1 = \int_0^T \left\{ \left[\int_t^\infty \phi(t_1) dt_1 \right] f_1(t, u(t), x(t)) + \phi(t) \mathcal{S}(t, x(t)) \right\} dt, \quad (10)$$

subject to

$$\mathcal{S}(t, x(t)) = \int_t^T f_2(z, u_2^*(z|t, x(t)), x_2^*(z|t, x(t))) dz, \quad (11)$$

$$\frac{dx(t)}{dt} = g_1(t, u(t), x(t)), \text{ for } t \in [0, T], \quad (12)$$

$$x(0) = x_0, \quad x(T) = x_T. \quad (13)$$

Given the Hamiltonian function

$$\mathcal{H}_1 = \left[\int_t^\infty \phi(t_1) dt_1 \right] f_1(t, u(t), x(t)) + \phi(t) \mathcal{S}(t, x(t)) + \lambda_1(t) g_1(t, u(t), x(t)), \quad (14)$$

the necessary conditions that must hold on the path $(u_1^*(t), x_1^*(t))_{t \in [0, T]}$ include $\partial \mathcal{H}_1 / \partial u(t) = 0$ and $d\lambda_1(t)/dt = -\partial \mathcal{H}_1 / \partial x(t)$.

Corollary 1 (Sufficiency). If g_1 and g_2 are linear in $u(t)$ and $x(t)$ and the integrands of J_1 and J_2 are concave in $u(t)$ and $x(t)$, then the necessary conditions are sufficient (Mangasarian (1966)).

Checking the concavity of the integrand of J_2 is standard. But checking the concavity of the integrand of J_1 is more involved. This is because the integrand of J_1 depends on the optimal post-switch path. Thus, one must first derive the post-switch solution $(u_2^*(z|t, x(t)), x_2^*(z|t, x(t)))_{z \in [t, T]}$,

which depends on $x(t)$, and then insert this solution into \mathcal{S} before checking the concavity of the integrand of J_1 .

Finally, to conclude this section we note that in addition to stochastic *timing* of the regime switch, we can easily allow for the possibility that the *structure* of the new regime itself (the functional form of the post-switch state equation) is uncertain. Adding in this second layer of uncertainty is relatively easy and requires just a few adjustments to Problem 1 and Theorem 1.¹⁶

Problem 2 (Timing Uncertainty and Structural Uncertainty). Add the following to Problem 1: the uncertainty about the functional form of g_2 is summarized by the random variable α , with density $\theta(\alpha)$ and sample space normalized to $[0, 1]$, where $\theta(\alpha)$ is continuously differentiable and realizations of α and t_1 are uncorrelated.

Theorem 2 (Necessary Conditions to Problem 2). Follow Theorem 1 to recursively derive necessary conditions with the following slight modifications: in Step 1 use the notation $g_2(t, u(t), x(t)|\alpha)$ and $(u_2^*(t|t_1, x(t_1), \alpha), x_2^*(t|t_1, x(t_1), \alpha))_{t \in [t_1, T]}$ to emphasize dependence of the solution on the realization of α , and in Step 2 write the continuation function $\mathcal{S}(t, x(t), \alpha)$ and replace the last term in the integrand of J_1 with $\int_0^1 \theta(\alpha) \phi(t) \mathcal{S}(t, x(t), \alpha) d\alpha$.

3. Application: Uncertainty about the Timing of Reform

3.1. Notation

Age is continuous and is indexed by t . Households are born at $t = 0$ and pass away no later than $t = T$. The probability of surviving to age t is $\Psi(t)$. Retirement and benefit collection occur

¹⁶Models with uncertainty about the structure of the new regime appeared in the early resource extraction literature (Hoel (1978)) and then again in the more modern literature on technology adoption when future returns to technology are stochastic (Hugonnier et al. (2006), Pommeret and Schubert (2009), Abel and Eberly (2012)). Problem 2 and Theorem 2 show how to augment our previous control problem to include both layers of uncertainty.

exogenously at $t = t_R$ and labor is supplied inelastically.¹⁷ A given household collects wages at rate $w(t)$ during the working period.

The government's *current* policy is summarized by a tax rate on wage earnings and benefit annuity (τ_1, b_1) . The current policy is unsustainable and this is publicly known. Therefore households know that reform is coming, but they don't know when. The reform date t_1 is a random variable with probability density $\phi(t_1)$ and sample space $[0, \infty]$.

The post-reform policy is (τ_2, b_2) . For now, we assume households have full information about the nature of the reform (τ_2, b_2) , they just don't know when it will kick in. Thus we are dealing with an application of Problem 1 from the previous section and Theorem 1 applies. Later we will consider an application of Problem 2, where both the timing and structure of reform are uncertain.

To compress notation let $y_1(t)$ be disposable income before the reform and let $y_2(t)$ be disposable income after the reform,

$$y_1(t) = \begin{cases} (1 - \tau_1)w(t), & \text{for } t \in [0, t_R], \\ b_1, & \text{for } t \in [t_R, T], \end{cases} \quad (15)$$

$$y_2(t) = \begin{cases} (1 - \tau_2)w(t), & \text{for } t \in [0, t_R], \\ b_2, & \text{for } t \in [t_R, T]. \end{cases} \quad (16)$$

Consumption is $c(t)$ and savings is $k(t)$, which earns interest at rate r .

3.2. Household Problem

Period utility is CRRA with relative risk aversion σ , and utils are discounted at the rate of time preference ρ . The household solves a dynamic stochastic control problem, taking as given the disposable income functions $y_1(t)$ and $y_2(t)$ while treating the reform date t_1 as a random variable:

$$\max_{c(t)_{t \in [0, T]}} : J = \mathbb{E} \left[\int_0^T e^{-\rho t} \Psi(t) \frac{c(t)^{1-\sigma}}{1-\sigma} dt \right], \quad (17)$$

¹⁷Fixing the retirement and collection dates allows us to parameterize the model to empirically reasonable values for these choices while abstracting from natural and institutional complications that shape the incentives behind these choices, such as declining worker productivity with age, disability risk, the delayed retirement credit, the taxation of social security benefits, and the ability of individuals to claim benefits before or after retirement.

subject to

$$\frac{dk(t)}{dt} = rk(t) + y_1(t) - c(t), \text{ for } t \in [0, t_1], \quad (18)$$

$$\frac{dk(t)}{dt} = rk(t) + y_2(t) - c(t), \text{ for } t \in [t_1, T], \quad (19)$$

$$k(0) = 0, k(T) = 0, \quad (20)$$

$$t_1 \text{ random with density } \phi(t_1) \text{ and sample space } [0, \infty]. \quad (21)$$

We refer readers to Appendix B for a step-by-step derivation of the solution to this problem. Here we report just the end results. Using Theorem 1 as our guide, we solve this problem recursively: following Step 1 we find that the post-reform (after the shock has hit) consumption path is

$$c_2^*(t|t_1, k(t_1)) = \frac{k(t_1) + \int_{t_1}^T e^{-r(v-t_1)} y_2(v) dv}{\int_{t_1}^T e^{-r(v-t_1)+(r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} e^{(r-\rho)t/\sigma} \Psi(t)^{1/\sigma}, \text{ for } t \in [t_1, T]. \quad (22)$$

Note the dependence on the timing of the reform and on the stock of assets at the time of reform. Then, following Step 2 and using the post-reform solution and working backwards, we find that the pre-reform solution $(c_1^*(t), k_1^*(t))_{t \in [0, T]}$ solves the following system of differential equations and boundary conditions¹⁸

$$\begin{aligned} \frac{dc(t)}{dt} = & \left(\frac{c(t)^{\sigma+1}}{\Psi(t)} \left[\frac{k(t) + \int_t^T e^{-r(v-t)} y_2(v) dv}{\int_t^T e^{-r(v-t)+(r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} \right]^{-\sigma} e^{(\rho-r)t} - c(t) \right) \times \left[\frac{\sigma}{\phi(t)} \int_t^\infty \phi(t_1) dt_1 \right]^{-1} \\ & + \left[\frac{d\Psi(t)}{dt} \frac{1}{\Psi(t)} + r - \rho \right] \frac{c(t)}{\sigma}, \end{aligned} \quad (23)$$

$$\frac{dk(t)}{dt} = rk(t) + y_1(t) - c(t), \quad (24)$$

$$k(0) = 0, k(T) = 0. \quad (25)$$

Finally, once we have the stage-one solution $(c_1^*(t), k_1^*(t))_{t \in [0, T]}$, we can use it to obtain the explicit stage-two solution conditional on realization of reform date t_1 ,

$$c_2^*(t|t_1, k_1^*(t_1)) = \frac{k_1^*(t_1) + \int_{t_1}^T e^{-r(v-t_1)} y_2(v) dv}{\int_{t_1}^T e^{-r(v-t_1)+(r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} e^{(r-\rho)t/\sigma} \Psi(t)^{1/\sigma}, \text{ for } t \in [t_1, T]. \quad (26)$$

¹⁸This system must be solved numerically. We guess and iterate on $c(0)$ until all the constraints are satisfied.

3.3. Welfare

As a point of reference, consider the case where the individual faces no risk (NR) about future taxes and benefits. The individual is endowed at $t = 0$ with expected future income and solves

$$\max_{c(t)_{t \in [0, T]}} : \int_0^T e^{-\rho t} \Psi(t) \frac{c(t)^{1-\sigma}}{1-\sigma} dt, \quad (27)$$

subject to

$$\frac{dk(t)}{dt} = rk(t) - c(t), \quad (28)$$

$$k(0) = \int_0^T \phi(t_1) Y(t_1) dt_1 + \left[\int_T^\infty \phi(t_1) dt_1 \right] \int_0^T e^{-rv} y_1(v) dv, \quad k(T) = 0, \quad (29)$$

where

$$Y(t_1) \equiv \int_0^{t_1} e^{-rv} y_1(v) dv + \int_{t_1}^T e^{-rv} y_2(v) dv. \quad (30)$$

The solution is

$$c^{NR}(t) = \frac{\int_0^T \phi(t_1) Y(t_1) dt_1 + \left[\int_T^\infty \phi(t_1) dt_1 \right] \int_0^T e^{-rv} y_1(v) dv}{\int_0^T e^{-rv + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} e^{(r-\rho)t/\sigma} \Psi(t)^{1/\sigma}, \quad \text{for } t \in [0, T]. \quad (31)$$

The welfare cost of living with reform uncertainty Δ is the solution to the following equation

$$\begin{aligned} & \int_0^T e^{-\rho t} \Psi(t) \frac{[c^{NR}(t)(1-\Delta)]^{1-\sigma}}{1-\sigma} dt \\ &= \int_0^T \phi(t_1) \left(\int_0^{t_1} e^{-\rho t} \Psi(t) \frac{c_1^*(t)^{1-\sigma}}{1-\sigma} dt + \int_{t_1}^T e^{-\rho t} \Psi(t) \frac{c_2^*(t|t_1, k_1^*(t_1))^{1-\sigma}}{1-\sigma} dt \right) dt_1 \\ &+ \left[\int_T^\infty \phi(t_1) dt_1 \right] \int_0^T e^{-\rho t} \Psi(t) \frac{c_1^*(t)^{1-\sigma}}{1-\sigma} dt. \end{aligned} \quad (32)$$

Notice that Δ captures just the cost of uncertainty about reform, and not the cost of reform itself. It is the fraction of consumption that the individual would be willing to give up to live in a world with no timing uncertainty. By endowing the individual with expected wealth over all possible reform dates we are netting out the lost wealth that comes from reform itself, leaving only the cost of uncertainty about reform. In other words, if at date $t = 0$ the government were to announce a future reform date t_1 , then $\Delta = 0$.¹⁹

¹⁹Our welfare metric Δ is similar to what Epstein et al. (2014) call the “timing premium,” though in the

4. Parameterization

The model is parameterized to capture individual income levels and survival probabilities over the life cycle. The parameters to be chosen are the maximum lifespan T , the survival function $\Psi(t)$, the exogenous retirement date t_R , the real return on assets r , the individual discount rate ρ , the utility preference parameter governing risk aversion σ , the age-earnings distribution $w(t)$, the probability density over reform dates $\phi(t_1)$, and policy parameters capturing tax rates and benefit levels before and after reform $\{\tau_1, \tau_2, b_1, b_2\}$.

Our survival data come from the Social Security Administration’s cohort mortality tables. These tables contain the mortality assumptions underlying the intermediate projections in the 2013 Trustees Report. The mortality table for each cohort provides the number of survivors at each age $\{1, 2, \dots, 119\}$, starting with a cohort of 10,000 newborns. However, we truncate the mortality data at age 100, assuming that everyone who survives to age 99 dies within the next year. We assume individuals enter the labor market at age 25, giving them a 75-year potential lifespan within the model. In our baseline parameterization, we use the mortality profile for males born in 1990, who are assumed to enter the labor market in 2015. For this cohort, we construct the survival probabilities at all subsequent ages conditional on surviving to age 25.

We normalize time so that the maximum age in the model is $T = 1$. Thus $t = 0$ in the model corresponds to age 25, and $t = 1$ corresponds to age 100. Because the survival data are discrete (providing the probability of surviving to each integer age), we fit a continuous survival function

models that they consider people are willing to pay for early resolution of uncertainty because of the way utility is specified (Epstein-Zin), even though the early information cannot be used to reoptimize. This contrasts with our setting in which early resolution leads to welfare gains precisely because it allows for better optimization and smaller distortions to consumption/saving decisions. This distinction arises because they consider uncertainty over *consumption* streams whereas we consider uncertainty over *income* streams. If we were to use Epstein-Zin utility then presumably our timing premium would be larger because it would capture not just the distortions to consumption and saving caused by late resolution of uncertainty but it would also capture the innate desire to know one’s consumption outcomes in advance. Epstein et al. (2014) remain very skeptical, however, that people would be willing to pay very much to know their consumption outcomes in advance if there is nothing that can be done to change those outcomes.

that has the following form:

$$\Psi(t) = 1 - t^x. \tag{33}$$

After transforming the survival data to correspond to model time, with dates on $[0, 1]$, $x = 3.28$ provides the best fit to the data.

The fixed retirement age is assumed to occur at age 65, which corresponds to $t_R = \frac{40}{75}$ in the model.²⁰ We assume a risk-free real interest rate of 2.9% per year, which is consistent with the long-run real interest rate assumed by the Social Security Trustees. In our model, this implies a value of $r = 75 * 0.029 = 2.175$. Estimates of the individual discount rate ρ vary substantially in the literature, and values of $\rho < r$ are necessary to generate a hump shaped consumption profile in the model. In the baseline model we set $\rho = 0$, although we consider other values in robustness exercises. In the baseline calibration we also set $\sigma = 3$, with other values considered for robustness.

We import the individual income profile from Gourinchas and Parker (2002) with age normalized onto model time $[0, 1]$ and the maximum income normalized to one. The continuous-time wage function is approximated by fitting a fifth-order polynomial to the discrete-time wage data:

$$w(t) = 0.697 + 1.49t - 3.41t^2 + 19.08t^3 - 59.78t^4 + 52.70t^5. \tag{34}$$

Figure 1 shows the graphs of the wage profile and the survival probabilities.

There is not much evidence about the actual distribution of possible reform dates, as this depends on the political process. Although the Social Security trust fund runs out in 2033, uncertainty about reform may extend beyond that date. For example, policy makers may adopt a temporary fix as 2033 approaches, postponing major reform even further into the future. The Health and Retirement Study (HRS), an ongoing panel survey of older Americans, regularly asks

²⁰While the Social Security normal retirement age is 66 for cohorts born between 1943 and 1954, and will gradually rise to 67 for cohorts born in 1960 and later, we use 65 as the exogenous retirement age for a few reasons. First, income data from Gourinchas and Parker (2002) is only available until age 65. Second, many individuals stop working and claim an actuarially reduced Social Security benefit before the normal retirement age. Finally, this assumption can make our results easier to compare with previous research, as many prior studies specify a retirement age of 65. This assumption will also be important when setting replacement rates for individuals of different incomes. We will then use replacement rates corresponding to retirement at age 65, rather than normal retirement age.

respondents to rate the chances of a cut in Social Security benefits in general within the next 10 years. In the 2010 wave of the survey, the mean subjective probability of a benefit cut within the next 10 years is around 65%; however, there is much variance around this value.²¹

We assume that reform is a Weibull random variable,

$$\phi(t_1) = \frac{\mu}{\gamma} \left(\frac{t_1}{\gamma} \right)^{\mu-1} e^{-(t_1/\gamma)^\mu}, \text{ for } t_1 \in [0, \infty]. \quad (35)$$

We consider two special cases, the exponential density and the Rayleigh density. First, we assume $\mu = 1$ to generate a constant hazard rate of reform

$$\phi(t_1) = \frac{e^{-t_1/\gamma}}{\gamma}. \quad (36)$$

Because it seems unlikely that the individual will totally escape reform, we calibrate this function by assuming $\int_1^\infty (e^{-t_1/\gamma})/\gamma dt_1 = 1\%$, which implies $\gamma = 0.2171473$.

Alternatively, it is plausible that political pressure for reform will mount as the trust fund runs out of money by 2033. For this reason we will also consider a second calibration where the likelihood of reform rises as the trust fund exhaustion date approaches. To capture this, we compute $\phi'(t_1) = 0$ and set $t_1 = (2033 - 2015)/75$, which implies

$$\gamma = \frac{18}{75} \left(\frac{\mu}{\mu - 1} \right)^{\frac{1}{\mu}}. \quad (37)$$

The larger the value of μ , the greater the mass around 2033. Our computational procedure struggles with values of μ larger than 2, so we set $\mu = 2$ which then implies $\gamma = 0.3394$. Figure 2 shows the graphs of these two calibrations of $\phi(t_1)$. We will report welfare calculations for both calibrations.

The *current* Social Security policy (τ_1, b_1) in the model is parameterized to match the current

²¹The Survey of Economic Expectations also elicits information on household expectations about future social security benefits (Dominitz et al. (2003), Manski (2004)). While this survey does document substantial uncertainty, it does not specifically measure uncertainty about the timing of reform. We are likely understating the costs of uncertainty because we are assuming that social security will continue to exist no matter what, whereas in reality a large portion of young households are not even confident of that (Dominitz et al. (2003)).

policy in the US. Consistent with our modeling in earlier sections, we focus only on retirement insurance (the Old Age and Survivors, or OASI, program) and ignore disability insurance. The OASI payroll tax rate (combined employer and employee shares) is $\tau_1 = 10.6\%$. Benefits b_1 are chosen to match observed replacement rates for various income groups.

Social Security benefits are based on an individual's Average Indexed Monthly Earnings (AIME), calculated as average monthly earnings, indexed for economy-wide wage growth, over the highest 35 years of the individual's career. A progressive benefit formula is applied to AIME to arrive at an individual's Primary Insurance Amount (PIA), the monthly benefit payable if benefits are claimed at normal retirement age. Claiming before normal retirement age - for example, at age 65, as we assume in our model - results in an actuarial reduction to benefits. The progressive benefit formula implies that the replacement rate - the ratio of monthly benefits to AIME - falls with AIME.

The Social Security Trustees Report publishes replacement rates for several stylized workers, each earning a fixed multiple of the economy-wide average wage throughout their career. According to the 2013 Trustees Report, the very low income group, which earns 25% of the economy-wide average wage, receives a replacement rate of 67.5% of AIME if benefits are claimed at age 65. The low-income group, which earns 45% of the economy-wide average wage, has a replacement rate of 49.0% of AIME. The medium income group, which earns the economy-wide average wage, receives a replacement rate of 36.4% of AIME. The high income group, which earns 1.6 times the average wage, receives a replacement rate of 30.1% of AIME. Finally, workers who earn the maximum taxable amount in each year of their career receive a replacement rate of 24.0% of AIME. These replacement rates apply to the year 2055, when the 1990 cohort turns 65. To compute pre-reform benefits, b_1 , we apply these replacement rates to the AIME for the normalized life-cycle income profile in our model, which is 0.92927.²²

For the policy experiments considered in the paper we assume that the reform undertaken will balance the Social Security budget over the infinite horizon.²³ According to Gokhale (2013) the

²²In computing AIME for our life cycle model, we disregard wage indexation because the Gourinchas and Parker (2002) real income profiles are already adjusted not only for price inflation, but also for cohort and time effects.

²³A partial reform may extend the horizon of the timing uncertainty, leaving the welfare costs that we consider in the paper.

unfunded liabilities of the OASI program over the infinite horizon would require a permanent tax increase of 3.1% of payroll (assuming no behavioral response to the tax increase). Alternatively, all current and future benefits could be cut by 21%.²⁴ As an initial pass, we consider both of these possible reforms:

- **Full benefit cut.** An across-the-board reduction in benefits by 21% for all current and future retirees (no exemption for current retirees) and no change in taxes. Hence, the new policy is $(\tau_2, b_2) = (\tau_1, b_1 \times (1 - 21\%))$.
- **Full tax increase.** An across-the-board increase in the OASI tax rate by 3.1 percentage points for all taxpayers regardless of age and no change in benefits. Hence, $(\tau_2, b_2) = (\tau_1 + 3.1\%, b_1)$.

We consider each of these reforms separately and assume that individuals know the reform that will occur, but they are uncertain about the timing. The results from these two scenarios will clarify the intuition of how timing uncertainty influences individual consumption and savings decisions. We will later consider the case where individuals are also uncertain about the structure of reform, not knowing if it will be a benefit cut, a tax increase, or some combination of the two.

5. Results

In order to gain intuition about the effects of uncertainty about the timing of reform, we plot the results for the benefit cut and tax increase scenarios. Figure 3 shows four consumption profiles over the life cycle for the full benefit cut experiment, for an individual with average earnings who faces a constant hazard rate of reform: the consumption path during the first stage c_1^* ; a pair of hypothetical consumption paths during stage two, c_2^* , conditional on reform at $t_1 = 0.5$ and $t_1 = 0.7$ as an example; and the consumption profile from a world with no Social Security risk where the individual gets her expected lifetime income, c^{NR} , as a reference point. The path the

²⁴Of course, the required tax and benefit adjustments could be much different in a macroeconomic model that includes household behavioral responses and factor price responses to demographic and fiscal shocks. For instance, in Kitao (2014), the required tax increase is 6% and the required benefit cut is 33%. Our welfare costs will go up if we feed bigger policy shocks into our model.

individual actually follows is c_1^* up to the stochastic date of reform and then consumption drops down to c_2^* . There is a continuum of many c_2^* paths, we have simply plotted two possibilities for illustrative purposes. An important feature of the experiment to emphasize is that the individual is still subject to a benefit cut even after the retirement date, so the first stage consumption path takes this possibility into account. This figure clearly shows how uncertainty about the timing of benefit reform causes non-trivial distortions to consumption-saving decisions. If the reform shock is realized fairly early, then consumption falls below the level without Social Security risk, whereas a later reform shock leaves consumption above the no-risk path. The trouble of course is that the individual doesn't know when the shock will hit.

Figure 4 plots the same information but for the case of a tax increase, again for an individual with average earnings who faces a constant hazard rate of reform. We consider two alternative hypothetical reform dates, $t_1 = 0.1$ and $t_1 = 0.4$ just as an example. We don't plot reform dates after retirement because, even though reform may indeed strike after retirement, there is no distortion to consumption because the individual no longer pays Social Security taxes. Similar to the case of a benefit cut, consumption always drops the moment reform strikes.

The drop in consumption in either scenario is not because the individual is taken by surprise per se, but is instead the result of rational, forward-looking behavior in the face of uncertainty. Reconsider Figure 3 and suppose the individual is standing just to the left of $t = 0.7$. From this perspective, the individual knows the shock may happen at any time over the interval $[0.7, \infty]$ and therefore bases consumption on that expectation, which is rational ex ante. But if the shock hits at the next moment (i.e., at $t = 0.7$) then ex post the individual turns out to be a little poorer than anticipated the moment before and hence consumption must be revised down.

Our first set of results corresponds to the welfare loss for an individual with average earnings who faces a constant hazard rate of reform. We find that the magnitude of the welfare loss is moderate for this individual: 0.01% of total lifetime consumption for the case of a benefit cut. We also find that uncertainty about the timing of tax reform is twice as costly as uncertainty about the timing of benefit reform, again for an individual with average earnings who faces a constant hazard of reform.²⁵ This is because the exponential density puts much of the mass during the

²⁵While these numbers seem small, the fact that Social Security only taxes 10.6 percent of labor income implies that it is only a limited part of individuals overall lifetime consumption.

working period, so that if the individual is living in the benefit-reform world, then he actually faces somewhat less uncertainty about his lifetime wealth because he knows there is a good chance that benefits will be cut before he retires anyway.

We also find important distributional effects as we look beyond the average individual. Table 1 reports welfare losses associated with uncertainty about the timing of reform for different income levels (each relative to the average individual in the benefit reform scenario). Due to the progressivity of benefits, uncertainty about the timing of a benefit cut is more harmful to low income individuals than to high income individuals. For example, Table 1 shows that, for the case of constant hazard rate of reform (Panel A), the very low income group will experience a welfare loss that is 4 times larger than what is experienced by the highest income group that maximizes their social security contribution in each year of work. This asymmetry occurs because social security benefits are a larger share of total retirement income for the poor than for the rich, and uncertainty over something important is naturally going to be costly. Qualitatively, this result continues to hold when we replace the exponential distribution with the Rayleigh distribution (see Panel B).

But this distributional effect reverses its sign for the case of uncertainty about the timing of a tax increase. Essentially, now the progressivity argument works in the opposite way. Even though all income groups pay the same tax rate, uncertainty about the tax rate is more costly for the rich because social security benefits are smaller for them (relative to their wage) and hence they face uncertainty about a larger portion of their wealth than do the poor. This result holds across both assumptions about the distribution of reform shocks.

Table 1. Welfare Loss from Uncertainty about the Timing of Reform:

Panel A. Baseline Density Function with Constant Hazard Rate of Reform

Income Level	Repl. Rate	Full Benefit Reform (21% benefit cut)	Full Tax Reform (3.1 ppt increase)
very low	67.5%	2.39	1.70
low	49.0%	1.51	1.91
average	36.4%	1.00*	2.07
high	30.1%	0.78	2.16
max	24.0%	0.60	2.25

Panel B. Alternative Density Function with Mode at 2033

Income Level	Repl. Rate	Full Benefit Reform (21% benefit cut)	Full Tax Reform (3.1 ppt increase)
very low	67.5%	1.02	1.31
low	49.0%	0.73	1.49
average	36.4%	0.57	1.62
high	30.1%	0.50	1.69
max	24.0%	0.45	1.77

*All other welfare costs are relative to this case.

6. Double Uncertainty: Timing and Structural Uncertainty

We now assume that both the timing and structure of social security reform are uncertain. Let $\tilde{\tau}_2$ be the new tax rate that would be sufficient to balance the budget without any reduction in benefits, and likewise let \tilde{b}_2 be the new benefit level that would balance the budget without a tax

increase. Of course, $\tilde{\tau}_2 > \tau_1$ and $\tilde{b}_2 < b_1$. But the new policy that the government actually chooses (τ_2, b_2) is an uncertain, linear combination of these extremes. We will express the new tax policy as a function of a continuous random variable α (with density $\theta(\alpha)$ and sample space $[0, 1]$):

$$\tau_2(\alpha) = \tau_1 + \alpha(\tilde{\tau}_2 - \tau_1), \quad (38)$$

$$b_2(\alpha) = b_1 - (1 - \alpha)(b_1 - \tilde{b}_2), \quad (39)$$

and

$$y_2(t|\alpha) = \begin{cases} (1 - \tau_2(\alpha))w(t), & \text{for } t \in [0, t_R], \\ b_2(\alpha), & \text{for } t \in [t_R, T]. \end{cases} \quad (40)$$

Using Theorem 2, the post-reform consumption path is

$$c_2^*(t|t_1, k(t_1), \alpha) = \frac{k(t_1) + \int_{t_1}^T e^{-r(v-t_1)} y_2(v|\alpha) dv}{\int_{t_1}^T e^{-r(v-t_1) + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} e^{(r-\rho)t/\sigma} \Psi(t)^{1/\sigma}, \text{ for } t \in [t_1, T], \quad (41)$$

and the pre-reform Euler equation is

$$\begin{aligned} \frac{dc(t)}{dt} &= \left(\frac{c(t)^{\sigma+1}}{\Psi(t)} \int_0^1 \theta(\alpha) \left[\frac{k(t) + \int_t^T e^{-r(v-t)} y_2(v|\alpha) dv}{\int_t^T e^{-r(v-t) + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} \right]^{-\sigma} e^{(\rho-r)t} d\alpha - c(t) \right) \\ &\quad \times \left[\frac{\sigma}{\phi(t)} \int_t^\infty \phi(t_1) dt_1 \right]^{-1} \\ &\quad + \left[\frac{d\Psi(t)}{dt} \frac{1}{\Psi(t)} + r - \rho \right] \frac{c(t)}{\sigma}. \end{aligned} \quad (42)$$

Using this Euler equation, together with the law of motion and boundary conditions for the savings account, we can compute the stage-one solution $(c_1^*(t), k_1^*(t))_{t \in [0, T]}$.

Finally, for welfare comparisons, the no risk benchmark is

$$c^{NR}(t) = \frac{\int_0^1 \int_0^T \theta(\alpha) \phi(t_1) Y(t_1|\alpha) dt_1 d\alpha + [\int_T^\infty \phi(t_1) dt_1] \int_0^T e^{-rv} y_1(v) dv}{\int_0^T e^{-rv + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} e^{(r-\rho)t/\sigma} \Psi(t)^{1/\sigma}, \text{ } t \in [0, T], \quad (43)$$

where $Y(t_1|\alpha) \equiv \int_0^{t_1} e^{-rv} y_1(v) dv + \int_{t_1}^T e^{-rv} y_2(v|\alpha) dv$, and the welfare cost of reform uncertainty Δ

solves the following equation

$$\begin{aligned}
& \int_0^T e^{-\rho t} \Psi(t) \frac{[c^{NR}(t)(1-\Delta)]^{1-\sigma}}{1-\sigma} dt \\
= & \int_0^1 \int_0^T \theta(\alpha) \phi(t_1) \left(\int_0^{t_1} e^{-\rho t} \Psi(t) \frac{c_1^*(t)^{1-\sigma}}{1-\sigma} dt + \int_{t_1}^T e^{-\rho t} \Psi(t) \frac{c_2^*(t|t_1, k_1^*(t_1), \alpha)^{1-\sigma}}{1-\sigma} dt \right) dt_1 d\alpha \\
& + \left[\int_T^\infty \phi(t_1) dt_1 \right] \int_0^T e^{-\rho t} \Psi(t) \frac{c_1^*(t)^{1-\sigma}}{1-\sigma} dt.
\end{aligned} \tag{44}$$

In the absence of reliable data on expectations about the structure of future reform, which is ultimately a political decision that will reflect the preferences of policymakers, we assume $\theta(\alpha)$ is the uniform density. Table 2 documents the welfare costs of double uncertainty for various income levels, each relative to the average earner in the full benefit reform scenario with a constant hazard rate of reform. The main finding of interest is that double uncertainty hurts the poor more than the rich by roughly a factor of 3, regardless of which assumption we make about the distribution of reform dates.

Figures 5 through 7 show the effects of double uncertainty on consumption allocations over the life cycle. In Figure 5 we report five consumption profiles: the optimal consumption path given that reform has not yet happened, c_1^* ; the consumption path corresponding to a hypothetical world without reform risk, c^{NR} ; and three post-reform consumption profiles c_2^* conditional on reform striking at date $t_1 = 0.4$. Note that reform may happen at any time of course. We are just showing the optimal responses for this particular date. Furthermore, there is a continuum of post-reform consumption profiles even for a single reform date because the structure of reform, α , is a continuously distributed random variable as well. For illustrative purposes however, we plot optimal responses to just three particular realizations $\alpha \in \{0, 0.5, 1\}$, which correspond to full benefit reform, 50-50 tax and benefit reform, and full tax reform. Our calculations of welfare losses take into account the full distribution of α , rather than just the three realization shown here, but the figure would be too messy to show all realizations over α .

Figure 5 illustrates that the individual will reduce consumption if the shock turns out to be benefit reform, whereas consumption actually jumps up a bit if the shock takes the form of a tax increase. This is because the individual was rationally planning for the contingency that benefits would be dramatically cut, and hence the news that reform takes the structure of a tax increase

is a positive shock to expected wealth because there are only a few years of tax payments left and now benefit are certain. To make the intuition clear, suppose at the very moment the individual retires the government announces that the Social Security program will reform its budget solely through increased taxation. This is a strictly positive wealth shock for this individual because he/she doesn't bear any of the burden of the tax increase, and now Social Security benefits are no longer in jeopardy. But good news should not be confused with welfare, because the uncertainty about the timing and structure of reform distort consumption and saving allocations in the years leading up to reform, even if reform just so happens to never actually end up costing the individual anything.

Figure 6 plots the exact same information as Figure 5, but for the case of a later realization of the reform shock ($t_1 = 0.6$ instead of 0.4). Taken together, the two figures show similar qualitative responses to the resolution of timing and structural uncertainty. However, the individual's responses are more exaggerated in the case of the later shock, at least up to a certain age, after which there is eventually so little remaining Social Security benefits to collect that uncertainty about the timing and structure of reform becomes irrelevant.

Figure 7 shows a little more detail. We add to Figures 5 and 6 the optimal consumption responses to a wider variety of realizations of the timing of reform, $t_1 \in \{0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$, while still plotting just three particular realizations of the structure of reform $\alpha \in \{0, 0.5, 1\}$. This figure helps to illustrate, at a glance, a few of the many contingent consumption plans formulated by the individual, and it helps to illustrate the breath of the distortion to consumption that is caused by the presence of double uncertainty about Social Security reform. For instance, at this particular parameterization, the variation in consumption during the middle part of the retirement period is especially severe. The maximum possible consumption levels over retirement result from the cases in which full tax reform (with no benefit cut) strikes during retirement and the individual therefore realized that he is totally off the hook, while minimum possible consumption levels during the retirement period result from cases in which full benefit reform strikes near the transition from work to retirement.

Table 2. Welfare Loss from Uncertainty about Timing & Structure of Reform:

Panel A. Baseline Density Function with Constant Hazard Rate of Reform

Income Level	Repl. Rate	Full Benefit Reform (21% benefit cut)	Full Tax Reform (3.1 ppt increase)	Double Uncertainty
very low	67.5%	2.39	1.70	3.09
low	49.0%	1.51	1.91	1.96
average	36.4%	1.00*	2.07	1.45
high	30.1%	0.78	2.16	1.27
max	24.0%	0.60	2.25	1.19

Panel B. Alternative Density Function with Mode at 2033

Income Level	Repl. Rate	Full Benefit Reform (21% benefit cut)	Full Tax Reform (3.1 ppt increase)	Double Uncertainty
very low	67.5%	1.02	1.31	3.50
low	49.0%	0.73	1.49	2.12
average	36.4%	0.57	1.62	1.44
high	30.1%	0.50	1.69	1.19
max	24.0%	0.45	1.77	1.03

*All other welfare costs are relative to this case.

7. Robustness

In this section we show how our results change when we alter certain assumptions. First, we consider different assumptions about preference parameters (the coefficient of risk aversion and the discount rate). Second, we show how much larger the welfare costs of uncertainty can be as we increase the variance of timing uncertainty and as we increase the variance of structural uncertainty. These are important extensions because we do not have good data to guide our baseline calibration of timing and structural uncertainty. Third, we show that the poor are still hit much harder than the rich by reform uncertainty even if we consider differential mortality across income groups, which is commonly believed to unwind at least some of the progressivity of Social Security (e.g., Coronado et al. (1999)).

7.1. Preference Parameters

As we increase the coefficient of relative risk aversion, σ , from the baseline value of 3 up to 5, the magnitude of the welfare cost of reform uncertainty increases dramatically. For instance, for the baseline case of exponentially distributed uncertainty about the timing of a full benefit cut, the cost to an average earner increases by more than 50%. The same is also true for the case of exponentially distributed uncertainty about the timing of a full tax increase.

Our baseline assumption is that the individual only discounts for mortality and not for pure time preference ($\rho = 0$). If we make a slightly more standard assumption that the discount rate equals the market rate of interest ($\rho = r = 2.9\%$), then the welfare effects of uncertainty about reform grow stronger. For instance, for the baseline case of exponentially distributed uncertainty about the timing of a full benefit cut, the cost of uncertainty for an average earner increases by 20%, while uncertainty about the timing of a full tax increase becomes 10% more costly to an average earner.

7.2. Uniform Reform Dates

A plausible assumption is that the random reform date is distributed uniformly over the individual's maximum lifespan, $\phi(t_1) = 1$ for all $t_1 \in [0, 1]$. Such an assumption spreads out the distribution of reform shocks more than in the baseline densities. The welfare results are reported

in Table 3, holding everything else at baseline parameters and assuming the structure of reform (α) continues to be distributed uniformly as well. Uncertainty about the timing of full benefit reform is now much more costly than in the baseline cases, while the cost of uncertainty about the timing of full tax reform is on the same order of magnitude as the baseline cases. Double uncertainty now hits the poor almost 4 times as hard as it hits the rich.

Table 3. Welfare Loss from Uncertainty about Timing & Structure of Reform:
Uniform Distribution of Reform Dates over Max Lifespan

Income Level	Repl. Rate	Full Benefit Reform (21% benefit cut)	Full Tax Reform (3.1 ppt increase)	Double Uncertainty
very low	67.5%	10.38	1.61	7.05
low	49.0%	6.17	1.79	4.41
average	36.4%	3.74	1.92	2.94
high	30.1%	2.70	2.00	2.33
max	24.0%	1.83	2.08	1.84

All welfare costs are relative to the numeraire case in Panel A of Table 2.

7.3. Extreme Political Risk

We have thus far used a uniform distribution over α (the structural uncertainty parameter) in all of our calculations for the double uncertainty case. This assumption implies that any particular convex combination (compromise) between a full tax increase and a full benefit cut is just as likely as any other combination. This may be a reasonable baseline assumption, but it could be the case that the structure of reform will ultimately look more like the outcome of a tug-of-war contest between two political parties, with one side winning completely and the other side losing, rather than a compromise. In this subsection we assume that the structure of reform is uncertain and will be either all on the tax side or all on the benefit side, with no probability of a convex combination. We assume that these two possibilities are equally likely, and we leave everything else at the baseline parameterization.

We obtain two key results. First, extreme political risk causes the welfare costs of reform uncertainty to be significantly larger than in the baseline calculations, for all income groups. Second, extreme political risk widens the inequality in welfare costs among the rich and poor. In the baseline calculations, reform uncertainty hits the poor about 3 times harder than the rich, but in the current case it is closer to a factor of 4.

7.4. Differential Mortality

It is well known that low-income individuals suffer from lower survival probabilities at all ages relative to high-income individuals. We have recomputed our welfare analysis for the case of income-specific survival functions based on Social Security administrative data provided in Table 1 of Waldron (2013). Based on this data, we assume that, throughout their lives, maximum earners face age-specific hazard rates of dying that are 0.5 times those of average earners, high earners face age-specific hazard rates that are 0.8 times those of average earners, low earners face age-specific hazard rates that are 1.4 times those of average earners, and very low earners face age-specific hazard rates that are 2.5 times those of average earners. (If scaling by these factors causes an age-specific hazard rate to exceed 1, that hazard rate is set to 1). Allowing hazard rates to differ by income group causes almost no change to our calculations of the welfare costs of uncertainty about the timing of reform.

Why doesn't differential mortality undo our results? One may think that because our results are driven by the progressivity of the Social Security system, including differential mortality would undo our results since this would unwind (or even reverse) the progressivity of the system. While it is true that differential mortality unwinds the progressivity of Social Security in a cross-sectional sense, it does not unwind the progressivity in a longitudinal sense. At a moment in time, the ratio of aggregate benefits collected by survivors to aggregate taxes paid by workers would tend to be low for segments of the population with lower survival probabilities. But this has nothing to do with how a given worker treats Social Security taxes and benefits in an expected utility (longitudinal) model. In such a model, the progressivity of Social Security is only related to the benefit-earning rule and not to mortality risk, because the latter enters the model only through the discount factor in the utility function. Regardless of their survival type, expected utility maximizers make an optimal consumption-saving plan that accounts for the contingency that

they survive until the maximum possible date. In other words, income flows (like Social Security) in the budget constraint of an expected utility maximizer are not discounted for survival risk.

8. Conclusion

In this paper we attempt to understand the welfare consequences of uncertainty over the timing and structure of social security reform. Household decision-making in the face of reform uncertainty can be modeled as a stochastic two-stage optimal control problem. We have formalized the tools that are required for studying this issue in a continuous-time setting.

We have paid special attention to how these welfare costs are distributed across income groups. While the precise magnitude of the welfare costs depend on a variety of factors, a consistent theme throughout our paper is that reform uncertainty may strike low-income groups especially hard. While the need for reform itself is driven by unavoidable demographic forces, uncertainty about the timing and structure of reform are avoidable. All of the costs that we report in this paper go away if policymakers simply announce when and how Social Security will be reformed.

References

- Abel, Andrew B. and Janice C. Eberly**, “Investment, Valuation, and Growth Options,” *Quarterly Journal of Finance*, 2012, 2 (1).
- Amit, Raphael**, “Petroleum Reservoir Exploitation: Switching from Primary to Secondary Recovery,” *Operations Research*, 1986, 34 (2), 534–549.
- Belo, Frederico, Vito D. Gala, and Jun Li**, “Government Spending, Political Cycles, and the Cross-Section of Stock Returns,” Working Paper October 2011.
- Benítez-Silva, Hugo, Debra S. Dwyer, Frank Heiland, and Warren C. Sanderson**, “Retirement and Social Security Reform Expectations: A Solution to the New Early Retirement Puzzle,” Working Paper, Stony Brook 2007.
- Bernanke, Ben S.**, “Irreversibility, Uncertainty, and Cyclical Investment,” *Quarterly Journal of Economics*, 1983, 98 (1), 85–106.
- Blundell, Richard and Thomas M. Stoker**, “Consumption and the Timing of Income Risk,” *European Economic Review*, 1999, 43, 475–507.
- Boucekkine, Raouf, Aude Pommeret, and Fabien Prieur**, “Optimal Regime Switching and Threshold Effects: Theory and Application to a Resource Extraction Problem under Irreversibility,” Working Papers, LAMETA, Universtiy of Montpellier 2012.
- , – , **and** – , “On the timing and optimality of capital controls: Public expenditures, debt dynamics and welfare,” *International Journal of Economic Theory*, 2013, 9 (1), 101–112.
- , – , **and** – , “Technological vs. Ecological Switch and the Environmental Kuznets Curve,” *American Journal of Agricultural Economics*, 2013, 95 (2), 252–260.
- , **Cagri Saglam, and Thomas Vallee**, “Technology Adoption Under Embodiment: A Two-Stage Optimal Control Approach,” *Macroeconomic Dynamics*, April 2004, 8 (02), 250–271.
- Boutchkova, Maria, Hitesh Doshi, Art Durnev, and Alexander Molchanov**, “Precarious Politics and Return Volatility,” *Review of Financial Studies*, 2012, 25 (4), 1111–1154.

- Bütler, Monika**, “Anticipation Effects of Looming Public-Pension Reforms,” *Carnegie-Rochester Conference Series on Public Policy*, 1999, 50 (1), 119–159.
- Clarke, Harry R. and William J. Reed**, “Consumption/Pollution Tradeoffs in an Environment Vulnerable to Pollution-Related Catastrophic Collapse,” *Journal of Economic Dynamics and Control*, 1994, 18, 991–1010.
- Coronado, Julia Lynn, Don Fullerton, and Thomas Glass**, “Distributional Impacts of Proposed Changes to the Social Security System,” NBER Working Papers, National Bureau of Economic Research January 1999.
- Croce, M. Max, Howard Kung, Thien T. Nguyen, and Lukas Schmid**, “Fiscal Policies and Asset Prices,” *Review of Financial Studies*, 2012, 25 (9), 2635–2672.
- Dasgupta, Partha and Geoffrey Heal**, “The Optimal Depletion of Exhaustible Resources,” *Review of Economic Studies*, 1974, 41, 3–28.
- Dogan, Erol, Cuong Le Van, and Cagri Saglam**, “Optimal Timing of Regime Switching in Optimal Growth Models: A Sobolev Space Approach,” *Mathematical Social Sciences*, 2011, 61, 97–103.
- Dominitz, Jeff, Charles F. Manski, and Jordan Heinz**, “Will Social Security be There for You?: How Americans Perceive their Benefits,” NBER Working Papers 9798, National Bureau of Economic Research, Inc 2003.
- Durnev, Art**, “The Real Effects of Political Uncertainty: Elections and Investment Sensitivity to Stock Prices,” Working Paper September 2010.
- Eeckhoudt, Louis, Christian Gollier, and Nicolas Treich**, “Optimal Consumption and the Timing of the Resolution of Uncertainty,” *European Economic Review*, 2005, 49, 761–773.
- Epstein, Larry G., Emmanuel Farhi, and Tomasz Strzalecki**, “How Much Would You Pay to Resolve Long-Run Risk,” *American Economic Review*, 2014, *forthcoming*.

- Evans, Richard W., Laurence J. Kotlikoff, and Kerk L. Phillips**, “Game Over: Simulating Unsustainable Fiscal Policy,” NBER Working Papers 17917, National Bureau of Economic Research, Inc 2012.
- Fernández-Villaverde, Jesús, Pablo Guerrón-Quintana, and Keith Kuester**, “Fiscal Volatility Shocks and Economic Activity,” Working Paper April 2011.
- Gokhale, Jagadeesh**, “Social Security Reform: Does Privatization Still Make Sense?,” *Harvard Journal on Legislation*, 2013, 50, 169–207.
- Gomes, Francisco J., Laurence J. Kotlikoff, and Luis M. Viceira**, “The Excess Burden of Government Indecision,” NBER Working Papers 12859, National Bureau of Economic Research, Inc 2007.
- Gourinchas, Pierre-Olivier and Jonathan A. Parker**, “Consumption Over the Life Cycle,” *Econometrica*, January 2002, 70 (1), 47–89.
- Hassett, Kevin A. and Gilbert E. Metcalf**, “Investment with Uncertain Tax Policy: Does Random Tax Policy Discourage Investment?,” *The Economic Journal*, 1999, 109 (457), 372–393.
- Hoel, Michael**, “Resource Extraction when a Future Substitute has an Uncertain Cost,” *Review of Economic Studies*, 1978, 45 (3), 637–644.
- Hugonnier, Julien, Florian Pelgrin, and Aude Pommeret**, “Technology Adoption under Uncertainty in General Equilibrium,” Working Paper 2006.
- Julio, Brandon and Youngsuk Yook**, “Political Uncertainty and Corporate Investment Cycles,” Working Paper 2012.
- Kamien, Morton I. and Nancy L. Schwartz**, “Optimal Maintenance and Sale Age for a Machine Subject to Failure,” *Management Science*, 1971, 17 (8), 495–504.
- Kemp, Murray C. and Ngo Van Long**, “Optimal Control Problems with Integrands Discontinuous with Respect to Time,” *The Economic Record*, June-Sept 1977, 53 (142&143), 405–20.
- Kitao, Sagiri**, “Sustainable Social Security: Four Options,” *Review of Economic Dynamics*, 2014, 17 (4), 756–779.

- Liebman, Jeffrey B. and Erzo F.P. Luttmer**, “Would People Behave Differently If They Better Understood Social Security? Evidence From a Field Experiment,” *American Economic Journal: Economic Policy*, 2014, *forthcoming*.
- Luttmer, Erzo F. P. and Andrew A. Samwick**, “The Welfare Cost of Perceived Policy Uncertainty: Evidence from Social Security,” Working Paper, Dartmouth College 2012.
- Makris, Miltiadis**, “Necessary conditions for infinite-horizon discounted two-stage optimal control problems,” *Journal of Economic Dynamics and Control*, December 2001, *25* (12), 1935–1950.
- Mangasarian, Olvi L.**, “Sufficient Conditions for the Optimal Control of Nonlinear Systems,” *SIAM Journal on Control*, 1966, *4* (1), 139–152.
- Manski, Charles F.**, “Measuring Expectations,” *Econometrica*, 2004, *72* (5), 1329–1376.
- McGrattan, Ellen R. and Edward C. Prescott**, “On Financing Retirement with an Aging Population,” Working Paper, University of Minnesota 2014.
- Pantzalis, Christos, David A. Strangeland, and Harry J. Turtle**, “Political elections and the resolution of uncertainty: The international evidence,” *Journal of Banking and Finance*, 2000, *24* (10), 1575–1604.
- Pastor, Lubos and Pietro Veronesi**, “Political Uncertainty and Risk Premia,” Working Paper September 2011.
- Pommeret, Aude and Katheline Schubert**, “Abatement Technology Adoption Under Uncertainty,” *Macroeconomic Dynamics*, 2009, *13* (04), 493–522.
- Rodrik, Dani**, “Policy Uncertainty and Private Investment in Developing Countries,” NBER Working Papers 2999, National Bureau of Economic Research June 1989.
- Saglam, Cagri**, “Optimal Pattern of Technology Adoptions under Embodiment: A Multi-Stage Optimal Control Approach,” *Optimal Control Applications and Methods*, 2011, *32*, 574–586.
- Sargent, Thomas**, “Ambiguity in American Monetary and Fiscal Policy,” Working Paper, New York University May 2005.

- Sialm, Clemens**, “Stochastic taxation and asset pricing in dynamic general equilibrium,” *Journal of Economic Dynamics and Control*, 2006, 30 (3), 511–540.
- Tahvonen, Olli and Cees Withagen**, “Optimality of Irreversible Pollution Accumulation,” *Journal of Economic Dynamics and Control*, 1996, 20, 1775–1795.
- Tomiyama, Ken**, “Two-stage optimal control problems and optimality conditions,” *Journal of Economic Dynamics and Control*, 1985, 9 (3), 317–337.
- Ulrich, Maxim**, “How Does the Bond Market Perceive Government Interventions?,” Working Paper May 2012.
- van der Wiel, Karen**, “Preparing for Policy Changes: Social Security Expectations and Pension Scheme Participation,” IZA DP No. 3623, Institute for the Study of Labor 2008.
- Waldron, Hilary**, “Mortality Differentials by Lifetime Earnings Decile: Implications for Evaluations of Proposed Social Security Law Changes,” *Social Security Bulletin*, 2013, 73 (1).
- Wright, Ian, Nick Bloom, and Jose Barrero**, “Investment and the Term Structure of Uncertainty,” Working Paper, Stanford University 2014.

Appendices

Appendix A provides a detailed proof of Theorem 1. Appendix B provides a full derivation of the solution to the household's optimization problem.

Appendix A: Proof of Theorem 1

Step 1. This step of the backward induction procedure is a completely standard Pontryagin problem and needs very little justification. The optimal control and state paths after the switch is realized must solve a standard deterministic control problem. We denote the solution to this subproblem $(u_2^*(t|t_1, x(t_1)), x_2^*(t|t_1, x(t_1)))_{t \in [t_1, T]}$, where the extra notation is meant to convey the dependence of the solution on the switch date t_1 and on the state variable at that date $x(t_1)$. The last part of Step 1 is strictly for convenience: we take this solution and change the time dummy t to z , and change the switch point t_1 to t and write $(u_2^*(z|t, x(t)), x_2^*(z|t, x(t)))_{z \in [t, T]}$. Thus we have the optimal control and state paths for all points in time z greater than switch point t . This change of dummies is innocuous but proves to be very helpful below in the process of linking the subproblems together through the continuation functional in the next step.

Step 2. This step requires a little more explanation. The purpose of this step is to find the optimal paths for the control and state before the realization of the switch. Hence, using the solution from Step 1, the objective functional is

$$\max_{u(t)_{t \in [0, T]}} : J_1 = \mathbb{E} \left[\int_0^{t_1} f_1(t, u(t), x(t)) dt + \int_{t_1}^T f_2(t, u_2^*(t|t_1, x(t_1)), x_2^*(t|t_1, x(t_1))) dt \right] \quad (\text{A1})$$

or

$$\begin{aligned} \max_{u(t)_{t \in [0, T]}} : J_1 = & \int_0^T \int_0^{t_1} \phi(t_1) f_1(t, u(t), x(t)) dt dt_1 + \int_0^T \phi(t_1) \mathcal{S}(t_1, x(t_1)) dt_1 \\ & + \int_T^\infty \int_0^T \phi(t_1) f_1(t, u(t), x(t)) dt dt_1, \end{aligned} \quad (\text{A2})$$

where $\mathcal{S}(t_1, x(t_1))$ is the continuation value or continuation function,

$$\mathcal{S}(t_1, x(t_1)) = \int_{t_1}^T f_2(t, u_2^*(t|t_1, x(t_1)), x_2^*(t|t_1, x(t_1))) dt, \quad (\text{A3})$$

and likewise $\int_0^T \phi(t_1)\mathcal{S}(t_1, x(t_1))dt_1$ is the continuation **functional**. The constraints are

$$\frac{dx(t)}{dt} = g_1(t, u(t), x(t)), \text{ for } t \in [0, T], \quad (\text{A4})$$

$$x(0) = x_0, x(T) = x_T. \quad (\text{A5})$$

Note that the control and state variables are defined over the entire planning interval because the switch could happen at any time on this interval, and hence the pre-switch problem amounts to choosing a path for these variables over the full interval. The marginal valuation of the state variable reflects the fact that the switch could happen at any moment in time, and hence the entire timepath of the state variable is relevant in determining the continuation value during the second stage. This makes matters more complicated than in the standard two-stage model with a deterministic switch, because there the costate variables from the two stages are simply equated at the switch point to ensure that the value of finishing the first stage with an extra unit of the state variable is equal to the value of starting the second stage with another unit. Whereas here the two stages are linked together through a continuation functional. The solution to the above problem is the one that will be followed up to the random switch point.

This appears to be a non-standard control problem, but with some algebra we can convert it into a standard one for which the standard Maximum Principle applies. To make progress, let's change the dummy of integration in the second integral in J_1 from t_1 to t , and also change the dummy in the continuation functional from t to z . Now we restate the objective functional as

$$\begin{aligned} \max_{u(t)_{t \in [0, T]}} : J_1 = & \int_0^T \int_0^{t_1} \phi(t_1) f_1(t, u(t), x(t)) dt dt_1 + \int_0^T \phi(t) \mathcal{S}(t, x(t)) dt \\ & + \int_T^\infty \int_0^T \phi(t_1) f_1(t, u(t), x(t)) dt dt_1, \end{aligned} \quad (\text{A6})$$

where

$$\mathcal{S}(t, x(t)) = \int_t^T f_2(z, u_2^*(z|t, x(t)), x_2^*(z|t, x(t))) dz. \quad (\text{A7})$$

This innocuous change of variables is helpful because it allows us to write the continuation function \mathcal{S} as a function of $x(t)$. In doing so, $u_2^*(z|t, x(t))$ and $x_2^*(z|t, x(t))$ now take on the interpretation of the optimal control and state variables for all points in time z that are beyond the switch date

t and conditional on $x(t)$.

Changing the order of integration in the first term in J_1 , along with applying Fubini's Theorem to the third term in J_1 , causes a more manageable Pontryagin problem to emerge. Let us now restate our problem one last time

$$\max_{u(t)_{t \in [0, T]}} : J_1 = \int_0^T \left\{ \left[\int_t^\infty \phi(t_1) dt_1 \right] f_1(t, u(t), x(t)) + \phi(t) \mathcal{S}(t, x(t)) \right\} dt, \quad (\text{A8})$$

subject to

$$\mathcal{S}(t, x(t)) = \int_t^T f_2(z, u_2^*(z|t, x(t)), x_2^*(z|t, x(t))) dz, \quad (\text{A9})$$

$$\frac{dx(t)}{dt} = g_1(t, u(t), x(t)), \text{ for } t \in [0, T], \quad (\text{A10})$$

$$x(0) = x_0, \quad x(T) = x_T. \quad (\text{A11})$$

This reformulated objective functional has an intuitive interpretation. The first term in the integrand gives the payoff of $u(t)$ and $x(t)$ through the function $f_1(t, u(t), x(t))$, weighted by the probability that the random switch will occur sometime after t . The second term gives the payoff of holding $x(t)$ through the continuation value $\mathcal{S}(t, x(t))$, weighted by the density function $\phi(t)$. Thus, the first payoff term f_1 is weighted by one minus the c.d.f. because this payoff is relevant as long as the switch comes later than t , whereas the second payoff term \mathcal{S} is weighted by the p.d.f. because this payoff is relevant only at the switch point. The standard Maximum Principle can be applied to this reformulated subproblem. We denote the solution to this subproblem $(u_1^*(t), x_1^*(t))_{t \in [0, T]}$. This is the solution path for all t before the realization of the random switch. Note that the switch could happen at any point in time. This is why this subproblem is defined over the entire planning interval.

Appendix B. Derivation of Solution to Household Problem

Using Theorem 1 as our guide, we can solve the stochastic two-stage household problem recursively by breaking it into two Pontryagin subproblems as follows.

Step 1. Solve the post-reform ($t = t_1$) subproblem:

We first solve the Pontryagin subproblem corresponding to the moment that reform occurs.

This is a deterministic, fixed endpoint control problem

$$\max_{c(t)_{t \in [t_1, T]}} : J_2 = \int_{t_1}^T e^{-\rho t} \Psi(t) \frac{c(t)^{1-\sigma}}{1-\sigma} dt, \quad (\text{B1})$$

subject to

$$\frac{dk(t)}{dt} = rk(t) + y_2(t) - c(t), \text{ for } t \in [t_1, T], \quad (\text{B2})$$

$$t_1 \text{ given, } k(t_1) \text{ given, } k(T) = 0. \quad (\text{B3})$$

Form the Hamiltonian \mathcal{H}_2 with multiplier $\lambda_2(t)_{t \in [t_1, T]}$ and compute the necessary conditions

$$\mathcal{H}_2 = e^{-\rho t} \Psi(t) \frac{c(t)^{1-\sigma}}{1-\sigma} + \lambda_2(t) [rk(t) + y_2(t) - c(t)], \quad (\text{B4})$$

$$\frac{\partial \mathcal{H}_2}{\partial c(t)} = e^{-\rho t} \Psi(t) c(t)^{-\sigma} - \lambda_2(t) = 0, \text{ for } t \in [t_1, T], \quad (\text{B5})$$

$$\frac{d\lambda_2(t)}{dt} = -\frac{\partial \mathcal{H}_2}{\partial k(t)} = -r\lambda_2(t), \text{ for } t \in [t_1, T]. \quad (\text{B6})$$

Rewrite the costate equation

$$\lambda_2(t) = \lambda_2(t_1) e^{-r(t-t_1)}, \quad (\text{B7})$$

and collapse the necessary conditions into a single equation

$$e^{-\rho t} \Psi(t) c(t)^{-\sigma} = \lambda_2(t_1) e^{-r(t-t_1)}. \quad (\text{B8})$$

Solve for $c(t)$

$$c(t) = \lambda_2(t_1)^{-1/\sigma} e^{[(r-\rho)t - rt_1]/\sigma} \Psi(t)^{1/\sigma}. \quad (\text{B9})$$

Solve differential equation (B2) using the boundary conditions in (B3)

$$k(t_1) + \int_{t_1}^T y_2(v) e^{-r(v-t_1)} dv = \int_{t_1}^T c(v) e^{-r(v-t_1)} dv. \quad (\text{B10})$$

Insert (B9) into (B10)

$$k(t_1) + \int_{t_1}^T y_2(v) e^{-r(v-t_1)} dv = \int_{t_1}^T \lambda_2(t_1)^{-1/\sigma} e^{-r(v-t_1) + [(r-\rho)v - vt_1]/\sigma} \Psi(v)^{1/\sigma} dv, \quad (\text{B11})$$

and solve for the constant

$$\lambda_2(t_1)^{-1/\sigma} = \frac{k(t_1) + \int_{t_1}^T y_2(v) e^{-r(v-t_1)} dv}{\int_{t_1}^T e^{-r(v-t_1) + [(r-\rho)v - rt_1]/\sigma} \Psi(v)^{1/\sigma} dv}. \quad (\text{B12})$$

Insert this into (B9) to obtain the solution consumption path

$$c_2^*(t|t_1, k(t_1)) = \frac{k(t_1) + \int_{t_1}^T e^{-r(v-t_1)} y_2(v) dv}{\int_{t_1}^T e^{-r(v-t_1) + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} e^{(r-\rho)t/\sigma} \Psi(t)^{1/\sigma}, \text{ for } t \in [t_1, T]. \quad (\text{B13})$$

This is the optimal consumption path after the reform shock has hit.

Anticipating the method for solving the pre-reform subproblem in the next step of Theorem 1, we will need to change the time dummies: now think of t as the reform date and z as any time after the reform date. Thus, rewrite the solution as

$$c_2^*(z|t, k(t)) = \frac{k(t) + \int_t^T e^{-r(v-t)} y_2(v) dv}{\int_t^T e^{-r(v-t) + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} e^{(r-\rho)z/\sigma} \Psi(z)^{1/\sigma}, \text{ for } z \in [t, T]. \quad (\text{B14})$$

Step 2. Solve the pre-reform ($t = 0$) subproblem:

Next we solve the $t = 0$ subproblem as in Theorem 1,

$$\max_{c(t)_{t \in [0, T]}} : J_1 = \int_0^T \left\{ \left[\int_t^\infty \phi(t_1) dt_1 \right] e^{-\rho t} \Psi(t) \frac{c(t)^{1-\sigma}}{1-\sigma} + \phi(t) \mathcal{S}(t, k(t)) \right\} dt, \quad (\text{B15})$$

subject to

$$\begin{aligned} \mathcal{S}(t, k(t)) &= \int_t^T e^{-\rho z} \Psi(z) \frac{c_2^*(z|t, k(t))^{1-\sigma}}{1-\sigma} dz \\ &= \frac{1}{1-\sigma} \left[\frac{k(t) + \int_t^T e^{-r(v-t)} y_2(v) dv}{\int_t^T e^{-r(v-t) + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} \right]^{1-\sigma} \int_t^T e^{[r(1-\sigma) - \rho]z/\sigma} \Psi(z)^{1/\sigma} dz, \end{aligned} \quad (\text{B16})$$

$$\frac{dk(t)}{dt} = rk(t) + y_1(t) - c(t), \text{ for } t \in [0, T], \quad (\text{B17})$$

$$k(0) = 0, \quad k(T) = 0. \quad (\text{B18})$$

Form the Hamiltonian \mathcal{H}_1 with multiplier $\lambda_1(t)_{t \in [0, T]}$ and compute the necessary conditions²⁶

$$\mathcal{H}_1 = \left[\int_t^\infty \phi(t_1) dt_1 \right] e^{-\rho t} \Psi(t) \frac{c(t)^{1-\sigma}}{1-\sigma} + \phi(t) \mathcal{S}(t, k(t)) + \lambda_1(t) [rk(t) + y_1(t) - c(t)], \quad (\text{B19})$$

$$\frac{\partial \mathcal{H}_1}{\partial c(t)} = \left[\int_t^\infty \phi(t_1) dt_1 \right] e^{-\rho t} \Psi(t) c(t)^{-\sigma} - \lambda_1(t) = 0, \text{ for } t \in [0, T], \quad (\text{B20})$$

$$\frac{d\lambda_1(t)}{dt} = -\frac{\partial \mathcal{H}_1}{\partial k(t)} = -\phi(t) \left[\frac{k(t) + \int_t^T e^{-r(v-t)} y_2(v) dv}{\int_t^T e^{-r(v-t) + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} \right]^{-\sigma} e^{-rt} - r\lambda_1(t), \text{ for } t \in [0, T]. \quad (\text{B21})$$

Differentiate (B20) with respect to t

$$\begin{aligned} 0 &= -\phi(t) e^{-\rho t} \Psi(t) c(t)^{-\sigma} + \left[\int_t^\infty \phi(t_1) dt_1 \right] \left[\frac{d\Psi(t)}{dt} e^{-\rho t} - \rho \Psi(t) e^{-\rho t} \right] c(t)^{-\sigma} \\ &\quad - \left[\int_t^\infty \phi(t_1) dt_1 \right] \sigma e^{-\rho t} \Psi(t) c(t)^{-\sigma-1} \frac{dc(t)}{dt} - \frac{d\lambda_1(t)}{dt}. \end{aligned} \quad (\text{B22})$$

Insert (B20) into (B21)

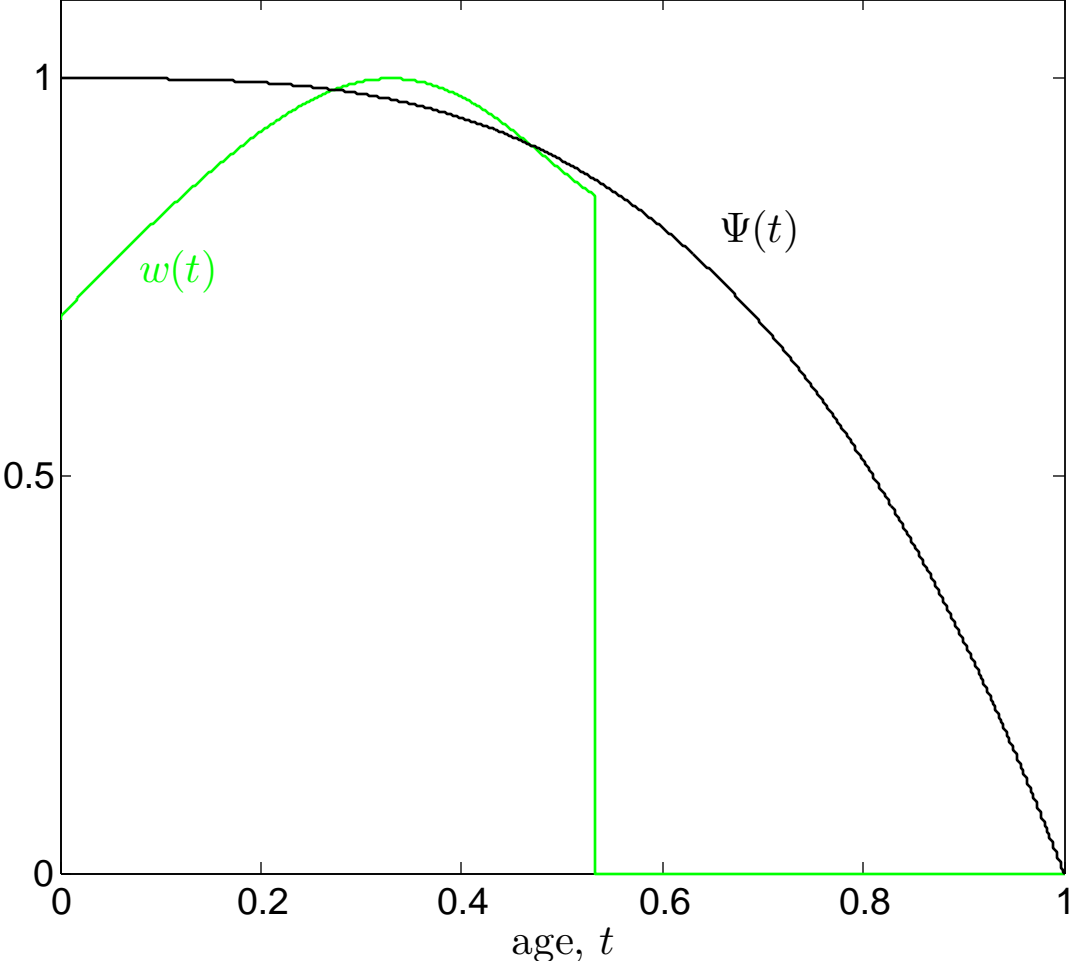
$$\frac{d\lambda_1(t)}{dt} = -\phi(t) \left[\frac{k(t) + \int_t^T e^{-r(v-t)} y_2(v) dv}{\int_t^T e^{-r(v-t) + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} \right]^{-\sigma} e^{-rt} - r \left[\int_t^\infty \phi(t_1) dt_1 \right] e^{-\rho t} \Psi(t) c(t)^{-\sigma}, \quad (\text{B23})$$

and then insert (B23) into (B22) to obtain the Euler equation

$$\begin{aligned} \frac{dc(t)}{dt} &= \left(\frac{c(t)^{\sigma+1}}{\Psi(t)} \left[\frac{k(t) + \int_t^T e^{-r(v-t)} y_2(v) dv}{\int_t^T e^{-r(v-t) + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} \right]^{-\sigma} e^{(\rho-r)t} - c(t) \right) \times \left[\frac{\sigma}{\phi(t)} \int_t^\infty \phi(t_1) dt_1 \right]^{-1} \\ &\quad + \left[\frac{d\Psi(t)}{dt} \frac{1}{\Psi(t)} + r - \rho \right] \frac{c(t)}{\sigma}. \end{aligned} \quad (\text{B24})$$

²⁶The necessary conditions are also sufficient because the integrand of J_1 is concave in $c(t)$ and $k(t)$ (see Corollary 1).

Figure 1. Baseline Wage Profile and Survival Uncertainty



Survival data: SSA 1990 cohort. Wage data: Gourinchas and Parker (2002).

Figure 2. Weibull Density Functions over Random Timing of Reform

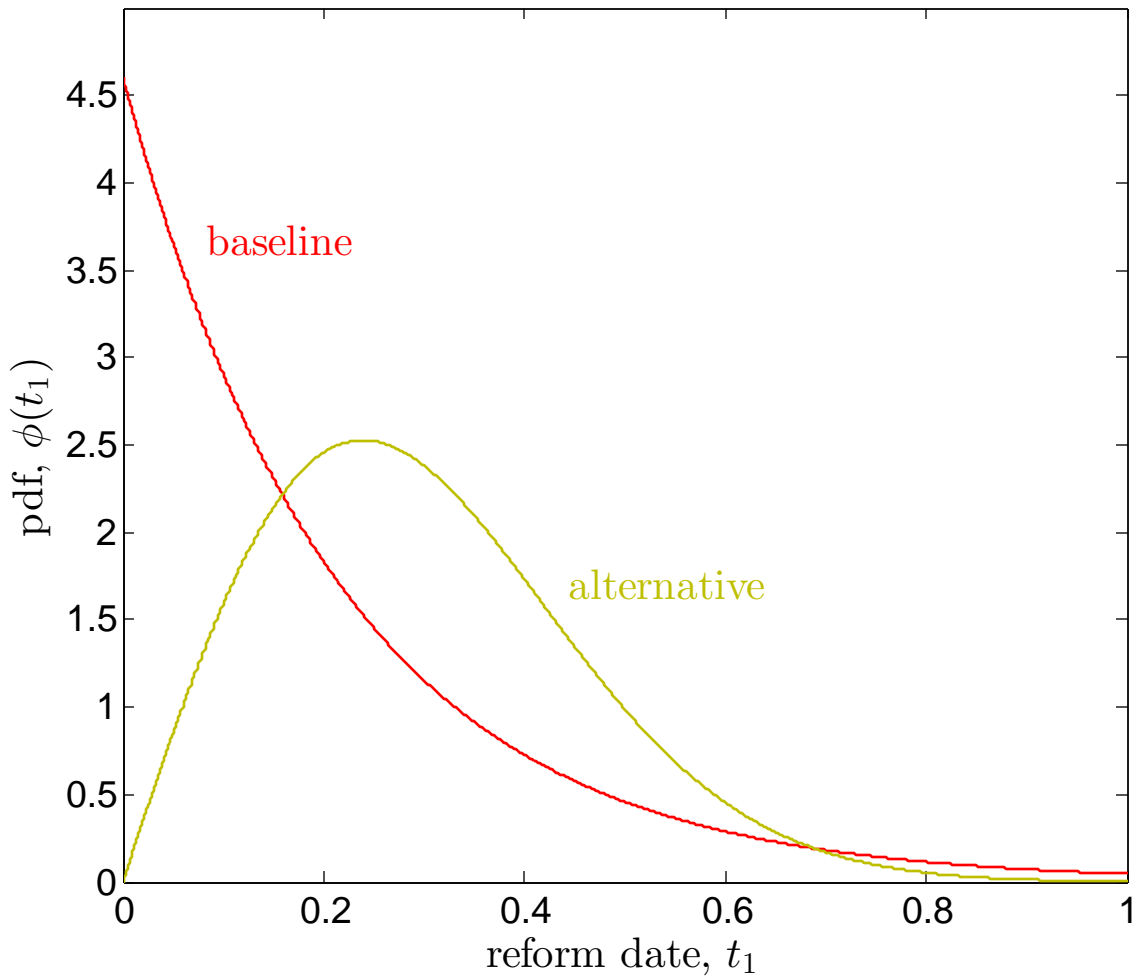
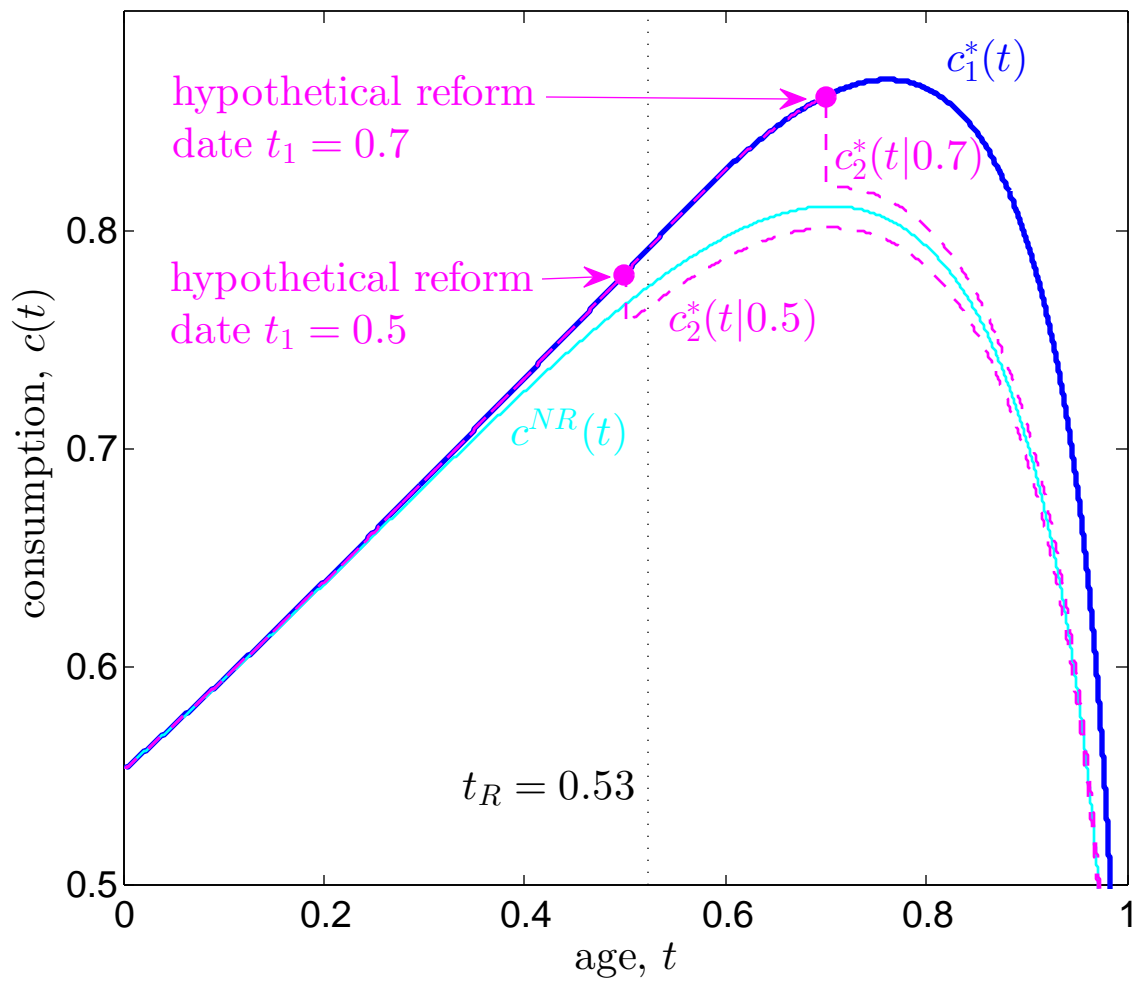
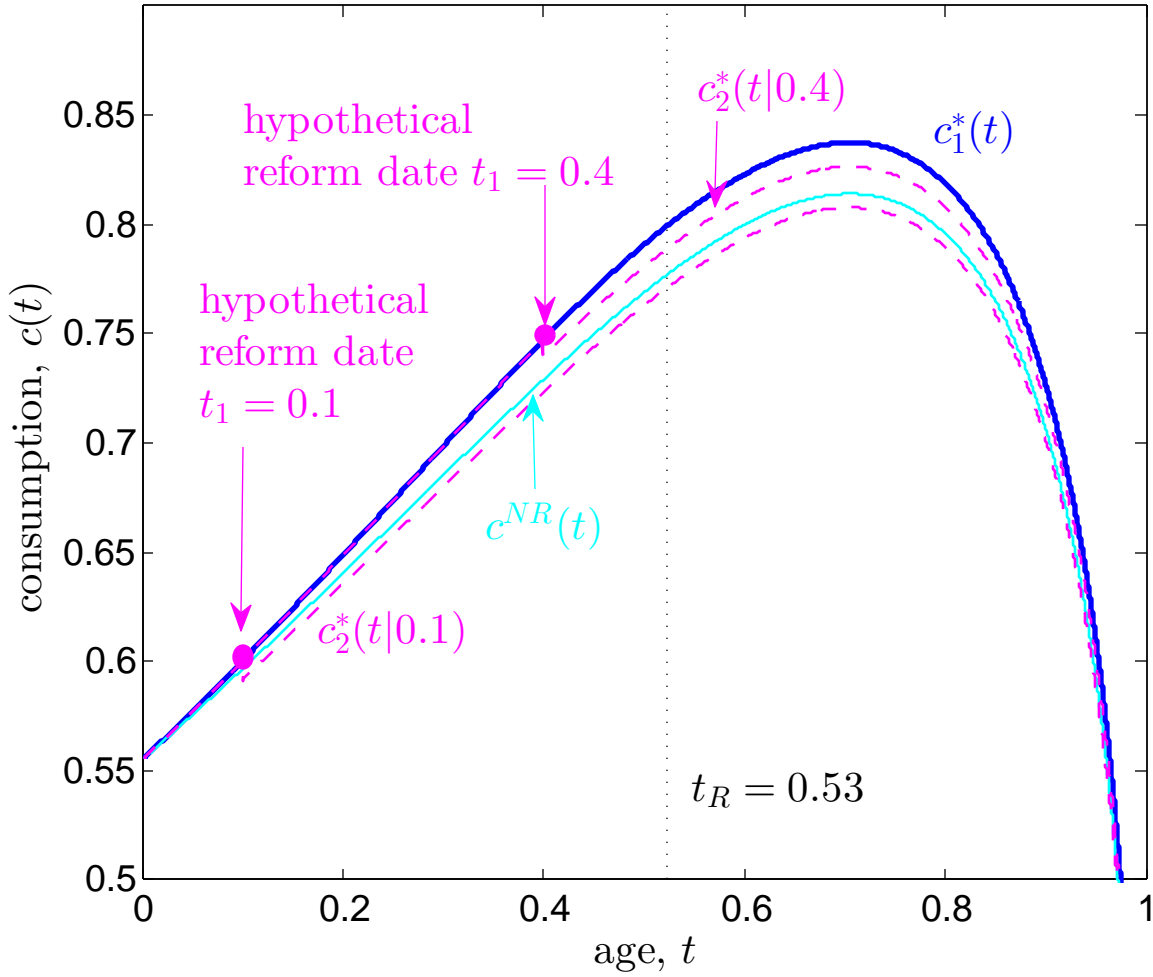


Figure 3. The Case of **Benefit Reform** with Stochastic Reform Date



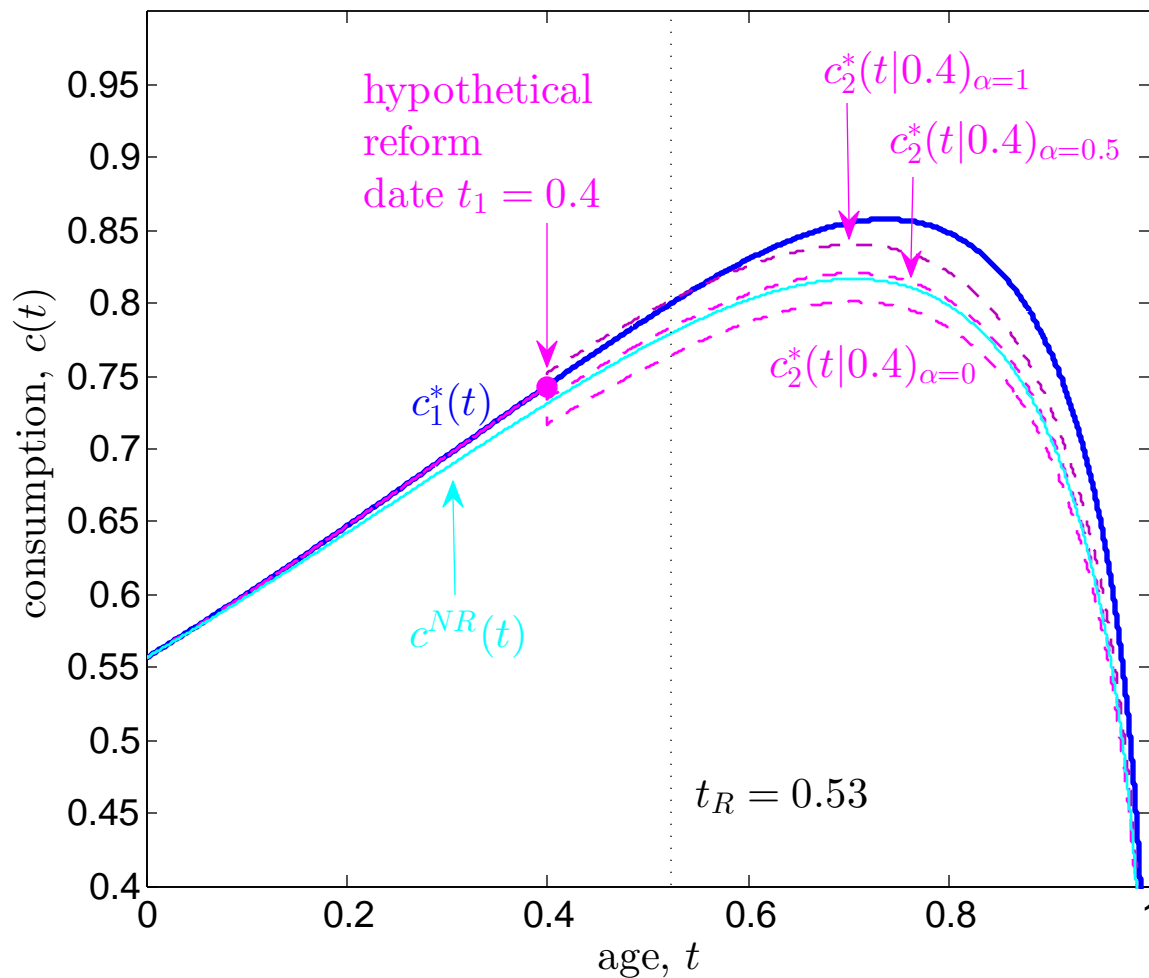
Social security parameters: $\tau_1 = 0.106$, $\tau_2 = 0.106$, $b_1 = 0.322$, $b_2 = 0.254$.

Figure 4. The Case of **Tax Reform** with Stochastic Reform Date



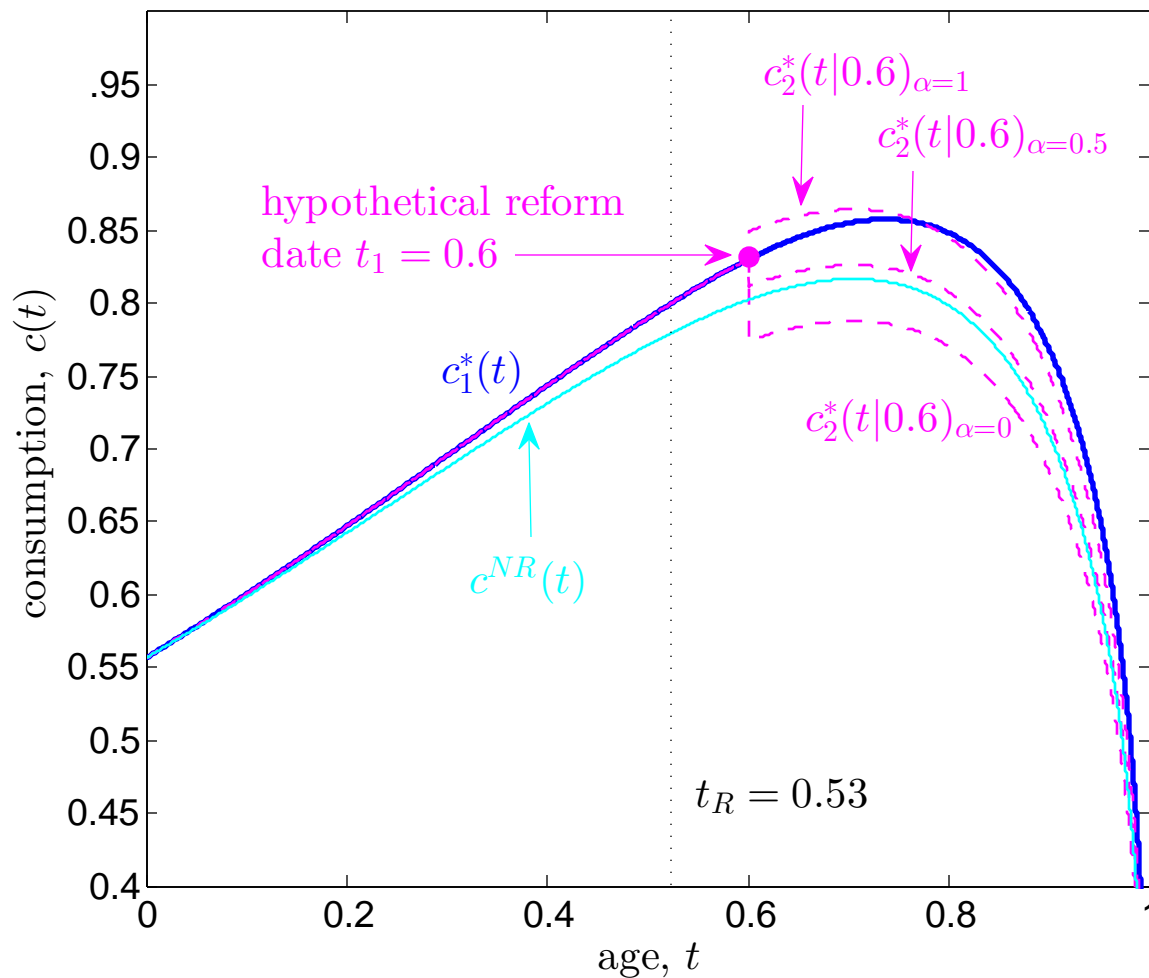
Social security parameters: $\tau_1 = 0.106$, $\tau_2 = 0.137$, $b_1 = 0.322$, $b_2 = 0.322$.

Figure 5. **Double Uncertainty: Timing and Structural Uncertainty**



The share of the budget crisis resolved through extra taxation is α , and the share resolved through benefit adjustments is $1 - \alpha$.

Figure 6. **Double Uncertainty: Timing and Structural Uncertainty**



The share of the budget crisis resolved through extra taxation is α , and the share resolved through benefit adjustments is $1 - \alpha$.

Figure 7. **Double Uncertainty:** Timing and Structural Uncertainty

