Adaptive Learning, Social Security Reform, and Policy Uncertainty

Erin Cottle Hunt*
Department of Economics
Lafayette College

September 18, 2020

Abstract

I develop an adaptive learning model to study the welfare effects of Social Security policy uncertainty in an aging economy. Agents combine *full knowledge* of the political process (which Social Security reforms are possible and when they could occur) with *limited knowledge* about the structure of the economy. The adaptive learning amplifies cyclical dynamics along the transition path to the new steady state. This magnifies the welfare effects of policy uncertainty, compared to a standard rational expectations model. The ex-ante consumption equivalent variation that equates the expected utility of consumption (with policy uncertainty) to the utility of expected consumption (across the possible policies without uncertainty) ranges between -0.29% and 0.21% of lifetime consumption in the adaptive learning model compared to -0.012% to -0.018% in the standard model. The welfare cost to future generations is also larger in the adaptive learning model compared to the rational model.

JEL Codes: D83, D84, D91, E71, H55, H62
Key words: Learning; Bounded Rationality; Policy Uncertainty; Social Security; Fiscal Sustainability

*cottlee@lafayette.edu

1 Introduction

The aging of the US population puts upward pressure on Social Security expenditures. There is a limit to how much the government can debt finance benefits, as rapid debt growth could strain the economy (see, for example, Chalk (2000) and Davig et al. (2010)). It is likely that Social Security benefits will be reduced or Social Security taxes will be raised (or both) at some point in the future. Thus, Social Security reform is likely, but uncertain. This uncertainty may influence how households save and consume today. In turn, those choices may influence the accumulation of government debt.

I build a general equilibrium life-cycle model that incorporates demographic change, the potential for government debt accumulation, and uncertain Social Security reform.[1] I use the model to explore household saving and consumption. I consider two variations of the model. As a baseline, I use a standard full information or rational expectations model. In addition to the baseline, I use a limited cognition or adaptive learning model. Agents in the adaptive learning model combine *full knowledge* of the political process (which Social Security reforms are possible and when they could occur) with *limited knowledge* about the structure of the economy. Agents do not know how demographic and policy changes influence endogenous variables such as the interest rate and wages. Instead, agents use an adaptive rule to form their forecasts of these endogenous variables. This asymmetry reflects the assumption that agents have reliable information about potential Social Security reforms that they incorporate into their saving-consumption decision.[2] At the same time, agents lack full structural knowledge of the general equilibrium evolution of key endogenous variables which they forecast using a simple statistical technique as in the learning literature.

I use the model to calculate life-cycle consumption for cohorts of agents alive before, during, and after uncertain Social Security reform following a demographic transition. The model is parameterized to the US economy. Four reforms are possible: either benefits are cut or taxes are raised in 2025, or benefits are cut or taxes are raised by a larger amount in 2035. I allow the government to accumulate debt along the transition path; the reforms are

[1]See Nelson and Phillips (2019) and the references therein for a discussion of the literature dedicated to analyzing Social Security reforms using OLG models.

[2]In 2019 Pew Research found a large majority (83%) of non-retired Americans doubt that Social Security will provide benefits at current levels when they eventually retire (Pew Research Center, March 2019, "Looking to the Future, Public Sees an America in Decline on Many Fronts"), and the American Life Panel in 2011 found that 56 percent of respondents are "not too confident" or "not confident at all" that the Social Security system will be able to provide them the level of benefits currently promised (from the 2011 survey *Well-being ms179*, alpdata.rand.org).

large enough to both finance benefits indefinitely and stabilize any accumulated debt.

I calculate the ex-ante welfare effect of facing policy uncertainty in both the rational model and the adaptive learning framework. In each case, I use a consumption equivalent welfare metric comparing expected utility of consumption (with policy uncertainty) to utility of expected consumption (across the policy paths without uncertainty). The four possible policy reforms have different income and wealth effects. The welfare metric nets out those effects to measure the cost to a household of facing policy uncertainty. The adaptive learning mechanism introduces cycles of over-optimism and over-pessimism that amplify the effects of policy uncertainty as compared to the rational model. The welfare effects of policy uncertainty are small and always negative in the standard rational model. The welfare effects are larger and sometimes positive in the adaptive learning model (somewhat similar to Brunnermeier and Parker (2005)). This is because households in the adaptive learning model do not understand how the demographic transition or policy change(s) will impact future wages, interest rates, and bonds. By assumption, they believe these variables in the future will be similar to what they observed in the past. Thus, although they understand how the potential Social Security reforms will impact their personal taxes and benefits, they do not fully understand how the reforms will impact the aggregate economy.

I also use the model to construct a welfare metric for future generations (born after a reform) that compares utility of consumption in a world with policy uncertainty to utility of consumption in a world with certain policy. This metric includes both the welfare effect of policy uncertainty and also the wealth effects of a particular reform. This analysis is conducted in both the rational model and the learning model. The welfare effects are larger in the learning model. The learning dynamics create cycles in the capital stock which raise and lower wages and interest rates along the transition path to the terminal steady state. Generations are benefited (or harmed) by these dynamics because they enter the model in a state of the world with higher (or lower) capital as a result of the over- (or under-) saving of the generations who faced uncertainty.

The analysis in this paper contributes to a growing literature that evaluates the effects of Social Security policy uncertainty. The consumption equivalent variation I use to measure the effect of policy uncertainty is based on the welfare criteria developed in Caliendo, Gorry, and Slavov (2019). They build a continuous-time life-cycle model with rich Social Security policy uncertainty (reform can occur at any date along a continuum and can be a benefit cut, a tax increase, or any mixture of the two) and calculate the welfare cost of policy uncertainty.

In their baseline model, the cost is relatively small, as agents facing policy uncertainty are able to partially self-ensure by saving. They also calculate the welfare effect for hand-to-mouth consumers who cannot save and find the welfare cost is much larger (close to 1% of lifetime consumption). My analysis confirms the intuition of Caliendo, Gorry, and Slavov's work; the welfare effects of facing policy uncertainty are larger in my boundedly-rational model as compared to the standard baseline.

This paper is also closely related to Kitao (2018) who builds a general equilibrium overlapping generations model calibrated to the Japanese economy. Kitao models the aging of the Japanese economy and includes policy uncertainty (in the main analysis, reform can take place at three different dates). She quantifies the welfare trade-off across generations by delaying reform and increasing its scope. The welfare comparison for future generations in my paper is similar to the welfare comparison used by Kitao. The results in my baseline (fully rational) model are similar to the results in Kitao's paper: delaying the reform reduces capital and output which harms generations born after the reform. I find that the welfare effects, both positive and negative, are larger in magnitude in the adaptive learning framework.

Several other papers have explored the welfare effects of Social Security uncertainty, including Bütler (1999) who studies the Swiss economy, Gomes et al. (2007) who calculate how much consumption an agent would be willing to give up to have policy uncertainty resolved earlier in her life, and Nelson (2017) who explores how an agent's belief about the probability of a reform influences her lifetime consumption and saving.[3]

This paper complements the existing literature by extending the analysis of Social Security policy uncertainty to a boundedly-rational model. I show that the effects of policy uncertainty can be much larger in a boundedly-rational model where agents combine complete knowledge of the policy process with limited knowledge about the structure of the economy.

The paper also contributes to a growing literature that examines the response to anticipated fiscal policy in models of adaptive learning. My analysis embeds the main modeling assumptions of Evans, Honkapohja, and Mitra (2009) in an overlapping generations framework.[4] The methodology used is particularly suited for the case in which there is a one-time

[3]In a related empirical study, Luttmer and Samwick (2018) conduct a large survey and find that respondents are willing to forgo 6 percent of their promised Social Security benefits to avoid policy uncertainty.

[4]See also Mitra, Evans, and Honkapohja (2013), Mitra, Evans, and Honkapohja (2019), Gasteiger and Zhang (2014), and Caprioli (2015). The adaptive learning model developed in this paper is a special case of Finite Horizon Learning in an overlapping generations model as developed in Cottle Hunt (2019), based on Branch

novel policy change, as with the Social Security reform under examination. Agents have specific information about the ways the reform could be implemented, which they build into their decision making, but no data from the past that is specifically relevant concerning its general equilibrium effects over time on wages and interest rates. For the latter they need to rely on adaptive forecasts.[5]

The baseline model is developed in Section 2, followed by the adaptive learning model in Section 3. Model parameterization is in Section 4. Section 5 considers an announced change to Social Security. Section 6 conducts welfare analysis for uncertain Social Security reform. Robustness analysis is in Section 7 and Section 8 concludes.

2 Theory: Baseline Model

I begin by solving the dynamic problem of a household facing possible demographic and policy changes. The baseline full information rational expectations model is presented first. The alternative adaptive learning framework is presented in Section 3.

2.1 Households

Households live for J periods, choose asset allocation (a^j for $j = 1, \cdots J-1$) in the first $J-1$ periods of life, and consumption (c^j for $j = 1, \cdots, J$) in all J periods of life, to maximize utility, taking prices and government policy (Social Security tax rate τ, benefit z, and additional tax λ) as given. Agents receive wage w_t for labor provided in period t, and retire exogenously at date $T \leq J$. The gross real return on savings in period t is given by R_{t+1}. Superscripts on variables indicate life-cycle stage (i.e., age), and subscripts indicate time period.[6]

Households maximize discounted expected lifetime utility (1) subject to period budget constraints (2). Here $E_t^*(x)$ indicates the time t expectation of x. The star indicates that

et al. (2013). For an additional example of adaptive learning in an overlapping generations model see Bullard and Duffy (2001).

[5]Evans, Honkapohja, and Mitra (2009) point out that if the same policy change were to be repeated many times then agents could learn how to forecast wages and interest rates in a more sophisticated way, but in practice repeated experiments of exceptional policies do not happen.

[6]As an example, a_{t+1}^2 is age 2 savings (savings in the second stage of life), in time period $t+1$.

the expectations need not be rational.[7]

$$\max_{a^j_{t+j-1}} E^*_t \sum_{j=1}^{J} \beta^{j-1} u(c^j_{t+j-1}) \tag{1}$$

$$c^j_{t+j-1} + a^j_{t+j-1} \leq R_{t+j-1} a^{j-1}_{t+j-2} + y^j_{t+j-1} \qquad \text{for } j = 1, \cdots, J \tag{2}$$

Here y^j indicates net period labor income during working life and Social Security income after retirement:

$$y^j_{t+j-1} = (1 - \tau_{t+j-1} - \lambda_{t+j-1}) w_{t+j-1} \quad \text{for } j < T$$
$$y^j_{t+j-1} = z^j_{t+j-1} \quad \text{for } j \geq T.$$

The lifespan is certain and agents do not have a bequest motive, so they exhaust all of their resources in the final period of life

$$a^J_{t+J-1} = 0.$$

Labor is supplied inelastically and preferences are given as the standard constant elasticity of substitution function:

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma} \quad \text{if } \sigma \neq 1$$
$$u(c) = \ln(c) \quad \text{if } \sigma = 1.$$

2.2 Demographics

To start off the economy, I assume that in period zero, there are J cohorts who enter the economy with given asset holdings according to their age. I assume that the initial young enter the economy with zero assets, and all other cohorts enter with a^j_{-1} for $j = 2, \cdots, J$. Successive cohorts enter the model with zero assets when they are young.

N_t indicates generation t and is given by number of young (i.e., the generation born) at time t. The population grows at rate n_t such that

$$N_t = (1 + n_t) N_{t-1}.$$

[7] In Section 3 expectations are formed adaptively.

The population at any time t is the sum of all living cohorts. Labor is supplied inelastically, thus the labor force (denoted H_t) is simply the working age population:

$$H_t = \sum_{j=1}^{T-1} N_{t+1-j}.$$

2.3 Production

The consumption good in the economy (Y_t) is produced by a single firm (or equivalently many small firms) using a constant elasticity of substitution technology that takes aggregate capital (K_t) and labor (H_t) as inputs and produces the consumption good according to:

$$Y_t = F(K_t, H_t) = AK_t^{\alpha} H_t^{1-\alpha}.$$

The parameter α measures the intensity of use of capital in production, and A represents technology or total factor productivity.[8] Factor markets are competitive and capital and labor are paid their marginal products. The gross real interest rate R_t is given by:

$$R_t = F_K(K_t, H_t) + 1 - \delta \tag{3}$$

where δ is the rate of depreciation. The wage rate w_t is given by

$$w_t = F_H(K_t, H_t). \tag{4}$$

2.4 Government

The government runs a modified pay-as-you-go Social Security system. The government pays retirement benefits to the retired generations by taxing the working generations. Benefits are not required to be the same as tax revenue, as in the current U.S. system. Any deficit is financed by issuing debt (and surpluses are saved and earn the market rate of return).

The payroll tax rate is τ_t. Social Security benefits z_t^j are paid according to a benefit earning rule:

$$z_t^j = \phi_t AIPE_t^j$$

[8]I abstract from technological growth, although it is straightforward to incorporate this into the model. If technological process were included in the model, aggregate variables would grow at the rate $(1+n)(1+g)$ along the balanced growth path (if g indicates the growth rate of technology).

where z_t^j represents the benefit paid to a retiree of age j in time t who initially retired at age T. The parameter ϕ_t is the replacement rate and shows how much of a worker's earnings are replaced by Social Security benefits. The term $AIPE_t^j$ represents the average indexed period earnings of an agent who is currently age j in time t (who retired exogenously at age T). This term is defined as:

$$AIPE_t^j = \sum_{i=1}^{T} \frac{I_{i,T} w_{t-T+i}}{T} \quad \text{for } j \geq T$$

where $I_{i,T}$ indicates the wage index for time period i relative to retirement date T (the ratio of average wages in year T to average wages in year i). In the United States, Social Security benefits are a function of average indexed monthly earnings (or AIME) from the highest 35 years of wage earnings for a given individual.[9] I simplify slightly, and assume that benefits depend on average indexed earnings from all working years. I make one additional simplification by assuming the benefit earning rule is simply a fraction (ϕ) of average earnings rather than a piece-wise linear progressive function of earnings. This simplification is trivial since there is no within-cohort heterogeneity in this model. Note finally that benefits are constant during the retirement phase of life for any particular cohort.[10]

To slow down the accumulation of debt along a transition path in which tax revenue is less than benefits paid, I also introduce an additional tax that is proportional to outstanding government debt. This tax doesn't directly correspond to anything in the U.S. tax system, but can be interpreted as legislative unease with rapidly growing debt. The tax does not prevent government debt from growing infinitely, but it slows the growth of debt. The tax is similar in spirit to a tax in Leeper (1991) (as modeled more recently by Davig, Leeper, and Walker (2010)) and thus will be called a Leeper tax. The Leeper tax is modeled as an additional payroll tax:

$$\lambda_t(B_t/H_t)w_t$$

where λ_t is the incremental tax paid (surplus received) when government debt is positive (negative) and B_t/H_t is government debt per laborer.[11]

[9] For more information on how Social Security benefits are calculated see https://www.ssa.gov/oact/cola/Benefits.html.

[10] In the United States, Social Security benefits grow according to a cost of living adjustment, which ensures that retirees' purchasing power is not eroded by inflation or wage growth. There is no inflation or wage growth (i.e. technology or productivity growth) in this model, so constant benefits are consistent with the U.S. cost of living adjustment.

[11] The Leeper tax leaves open the possibility of a total tax rate greater than one hundred percent if bond

The Leeper tax is useful for three reasons. First, the tax allows me to calibrate a dynamically efficient steady state with government debt (without a Leeper tax, a dynamically efficient steady state would require government saving). Second, the Leeper tax increases the speed of convergence along a transition path to a new steady state. The Leeper tax reduces the number of periods it takes for the state variables to converge to the new steady state because it slows the growth of government debt (or asset accumulation) and reduces the magnitude of any swings in debt. Third, the tax is also helpful computationally, as it increases the basin of attraction for the stable steady state. The tax increases the region of values for the state-variables (capital and bonds) that converge to the steady state. It is worth emphasizing that the Leeper tax does not drive the *qualitative* results of the paper; the steady states analyzed in the model can be achieved without a Leeper tax, with an appropriate increase in the payroll tax.

In the model, the government is not required to balance the Social Security budget in any particular period. If Social Security taxes are less than (more than) Social Security benefits, the Social Security program runs a deficit (surplus). In the United States, Social Security is "off-budget" for accounting purposes which means the Social Security budget is separate from the general operating budget of the federal government. The SSA has had operating surpluses in several years since 1982. Those surpluses, often called a "trust fund", were lent to the Treasury. This allowed the federal government to borrow less from other lenders. The interest income from the trust fund allows the SSA to pay benefits that are greater than the revenue collected in taxes. The difference between "on-budget" or "off-budget" is only an accounting difference; the total borrowings of the government depend on the SSA deficit (or surplus) and the deficit (or surplus) from the rest of the government. For the model, I abstract away from the trust fund and include Social Security "on-budget." Each of the reforms considered is sufficient to stabilize government debt in the long run.

The government issuance of bonds is equal to the gross interest on outstanding debt plus the Social Security deficit from the previous period less any revenue generated from the Leeper tax:

$$B_{t+1} = R_t B_t + \sum_{j=T}^{J} N_{t+1-j} \phi_t AIPE_t^j - H_t \tau_t w_t - H_t \lambda_t (B_t/H_t) w_t. \tag{5}$$

levels are high or if τ or λ is large (i.e., $\tau_t + \lambda_t(B_t/H_t) > 1$, for large B_t/H_t and/or large τ_t or λ_t). In the policy experiments that follow, that never happens.

2.5 Markets

Prices adjust in equilibrium to clear all markets: labor, assets (capital and bonds), and goods.

Equilibrium in the labor market requires the labor exogenously supplied by households to equal the labor used by the representative firm:

$$\sum_{j=1}^{T-1} N_{t+1-j} = H_t. \tag{6}$$

The asset market clears when aggregate capital and bonds are equal to the total savings of each cohort from the previous period:

$$K_{t+1} + B_{t+1} = \sum_{j=1}^{J} N_{t+1-j} a_t^j. \tag{7}$$

Households are indifferent between holding capital or bonds, since they earn the same real rate of return R_t. Capital and bonds pay the same rate of return through a no-arbitrage condition because there is no production risk in the economy. The stock of bonds in any period is pinned down by the government's flow budget constraint equation (5).

The goods market clears when output is equal to consumption plus investment:

$$F(K_t, L_t) = \sum_{j=1}^{J} N_{t+1-j} c_t^j + K_{t+1} + (1-\delta)K_t. \tag{8}$$

A full definition of the equilibrium and the steady state are in the on-line Appendix A.

3 Theory: Adaptive Learning

There are no stochastic shocks in the model; in the absence of policy uncertainty, agents in the baseline fully rational model of the previous section anticipate changes in endogenous variables with perfect foresight. When there is policy uncertainty, the rational agents can no longer see the future with perfect foresight. Rather, they form their expectations of endogenous variables using rational expectations. The expected value of a variable is equal to the mathematical expected value. I will refer to the standard baseline model as a "rational expectations" model throughout the rest of the paper.

In the limited cognition or adaptive learning version of the model, agents do not know the future value of endogenous variables. Agents are assumed to have limited knowledge about the structure of the economy; although agents solve a dynamic life-cycle optimization problem, they do not know the value of future wages, interest rates, or government debt (which depends on wages and influences take-home pay via the Leeper tax). Following the learning literature (Evans and Honkapohja (2001)), agents are assumed to forecast these endogenous variables using statistical techniques. Learning papers often assume agents run recursive least-squares to form their expectations. This is not feasible in a non-stochastic model. Following Evans et al. (2009), who also model a non-stochastic economy, I assume that agents use a simple adaptive learning rule to form their expectations of the future value of endogenous variables. This type of learning is often referred to as "steady-state" learning, because agents are estimating the mean, or steady state, of a variable, rather than a complex time series model. Steady-state learning is particularly suited for the case in which there is a one-time novel policy change, as with the Social Security reform under examination. Agents have some specific information about the ways the reform could be implemented, which they build into their decision making, but no data from the past that is specifically relevant concerning its general equilibrium effects over time on wages and interest rates. For the latter they need to rely on adaptive forecasts. I assume homogenous expectations across agents, which is equivalent to assuming agents inherit expectations from the previous generation.[12]

In the absence of policy uncertainty, agents in the adaptive learning model form expectations according to an adaptive rule, which will described below. Agents choose saving and consumption according to their first order conditions, given their adaptive forecasts of the interest rate, wage, and bonds. In the presence of policy uncertainty, agents are still assumed to use an adaptive rule to forecast endogenous variables. The agents have complete knowledge of the policy process, and so they understand which policies are possible, when those changes could occur, and with what probability. The agents make their saving-consumption decisions based on their first order conditions which equate the utility of consumption in the current period with the discounted expected utility of consumption in future periods, where the expectation is over the possible realizations of policy, given their adaptive forecasts of prices and bonds. Thus, the agents are *fully rational* with regard to policy uncertainty, but behaviorally limited in that they do not know the future value of endogenous variables and

[12]Steady-state learning is analogous to constant gain least-squares learning in a model with random shocks (like a productivity shock).

use an adaptive rule to form their forecasts.

Agents in the adaptive learning model forecast wages, interest rates, and bonds using the following rule:

$$w^e_{t+1} = \gamma w_t + (1-\gamma) w^e_t \tag{9}$$

$$R^e_{t+1} = \gamma R_t + (1-\gamma) R^e_t \tag{10}$$

$$b^e_{t+1} = \gamma b_t + (1-\gamma) b^e_t \tag{11}$$

with $\gamma \in (0,1)$. Here w^e indicates expected wage, R^e indicates expected interest rate, and b^e indicated expected government debt per worker. Agents also form expectations at time t of the wage, interest rate, and bonds in period $t+j$ for $j > 1$:

$$w^e_{t+j} = w^e_{t+1} \tag{12}$$

$$R^e_{t+j} = R^e_{t+1} \tag{13}$$

$$b^e_{t+j} = b^e_{t+1} \quad \text{for } j > 1. \tag{14}$$

In other words, agents are forecasting the steady state value of each variable.

Agents inherit expectations from the previous generation. That is, agents enter the model with knowledge of the previous period's prices, bonds, and expectations. Thus, at any moment in time, all agents have the same expectation for future prices and bonds.

A young agent in the adaptive learning model chooses first period savings and consumption and plans future savings and consumption to satisfy her $J-1$ first order equations and her lifetime budget constraint. Her time $t+1$ plan does not have to be consistent with her time t plan. She can update her savings decision based on the new information she receives in period $t+1$. In the absence of policy uncertainty, the young agent in time t solves:

$$(R_{t+j-1} a^{j-1}_{t+j-2} + y^j_{t+j-1} - a^j_{t+j-1})^{-\sigma} = R^e_{t+j} \beta (R^e_{t+j} a^j_{t+j-1} + y^{e,j+1}_{t+j} - a^{j+1}_{t,t+j})^{-\sigma} \quad \text{for } j = 1, \cdots, J-1 \tag{15}$$

with $y^{e,j+1}$ indicating expected net period labor income or Social Security income. That is:

$$y^{e,j+1}_{t+j} = (1 - \tau_{t+j} - \lambda_{t+j}) w^e_{t+j} \quad \text{for } j < T$$

$$y^{e,j+1}_{t+j} = \phi_{t+j} w^e_{t+j+T-j} \quad \text{for } j \geq T.$$

Similarly, an agent of age two solves $J-2$ first order equations, for the remaining $J-2$ periods of her life cycle. In total, the J cohorts alive in any period solve $\sum_{j=1}^{J-1} j = \frac{(J-1)J}{2}$ first order conditions. Together, the decisions of households of all ages (15), market clearing (5), (6), (7), and (8), and the expectation equations (9) through (14), create a recursive system that governs the dynamics of the economy. The stability of the learning model is discussed in more detail in on-line Appendix B.

In the presence of policy uncertainty, a young agent solves a similar problem to (15), except that the right-hand side is expected utility of consumption, given the adaptive forecasts. That is, the right-hand side is a weighted average of the utility of consumption across the possible states of the world that correspond to each of the possible policy reforms. In each possible state, the expected wage, interest rate, and government bonds are determined from the adaptive process outlined above. The agents combine full knowledge of the policy process including the probability and timing of each reform with limited structural knowledge about the macroeconomy.

In that regard, adaptive learning can be viewed as a small deviation from the baseline rational expectations model in the sense that agents are still forward looking (they solve an optimization problem) and only use a different rule to forecast future aggregates (the forecasting rule is backward looking). Although the agents in the learning model are using a simple adaptive rule to forecast, adaptive expectations are plausible in a world in which the true data generating process is complex. Adaptive expectations are optimal if agents think that wages, interest rates, and bonds follow IMA(1,1) processes. That is, expectations of the form given in Equation (9) are *rational* if the change in the variable of interest has the following form $\Delta x_t = \epsilon_t + \theta\epsilon_{t-1}$ for variable $x \in \{w, R, b\}$, shocks ϵ, and parameter $\theta \in (0, 1)$. The use of adaptive expectations is equivalent to agents believing that the variable they are forecasting has a mixture of permanent and transitory shocks (Muth (1960)). The gain parameter, γ, has a natural interpretation in this context. If all shocks are transitory, the optimal forecasting rule is constant, i.e., γ=0. In contrast, if all shocks are permanent, then the best forecasting rule is a random walk, i.e., $\gamma = 1$. In this particular model, all shocks are permanent, but agents are not endowed with this knowledge. I choose the gain parameter used in simulations to minimize the welfare cost of agents inaccurately forecasting along the transition path.[13] I explore alternative gain parameters as robustness checks in Section 7.1.

[13]Evans and Ramey show that by appropriately tuning the free parameters of the forecast rule agents can obtain the best forecast rule within a given class of under-parameterized learning rules (Evans and Ramey (2006)).

4 Parameterization

The parameter values are listed in Table 1. Agents enter the model at age 25 and live for six periods ($J = 6$). Each period is 10 years, and agents die with certainty at age 85. The capital share of income, discount factor, inverse elasticity of substitution, and depreciation rate are all set to standard parameter values.[14] The total factor productivity parameter A is chosen to ensure that two steady states exist for reasonable Social Security parameters.

Parameter		Value	Target
α	Capital share of income	$\frac{1}{3}$	from literature
β	Discount factor *annual* rate	0.995	"
σ	Inverse elasticity of substitution	1	"
δ	Depreciation *annual* rate	0.1	"
initial n	Population growth rate	0.1802	beneficiaries/workers 0.3
new n	Population growth rate	0.0164	beneficiaries/workers 0.44
A	TFP factor	10	two steady states with soc sec
τ	Pre-reform Social Security tax rate	0.106	OASI portion of payroll tax
ϕ	Pre-reform soc sec replacement rate	0.37	average replacement rate
τ^{early}	Early tax increase, tax rate	0.1378	SSA report
ϕ^{early}	Early benefit cut, replacement rate	0.2960	SSA report
τ^{late}	Early tax increase, tax rate	0.1425	SSA report
ϕ^{late}	Early benefit cut, replacement rate	0.2849	SSA report

Table 1: Parameterization of the model

The ratio of Social Security beneficiaries to workers was around 0.3 from 1975-2005. This ratio is expected to increase to 0.48 by 2095 under the Social Security Administration's intermediate cost assumptions (SSA (2019)). The increase in the ratio of beneficiaries to workers is driven by both increasing lifespans for retirees, and declining birthrates for working generations. I abstract from these details in the model, and capture all population changes using the growth rate n. For the exercises in this paper, I will choose n such that the ratio of retirees to workers is 0.3 and then increases gradually to 0.48, as illustrated in Figure 1.

In the policy simulations that follow, the initial Social Security parameters are set to match the current US system. The payroll tax rate is set to match the current old-age survivor insurance (OASI) portion of the Social Security tax, $\tau = 0.106$. The initial replacement rate is set to $\phi = 0.37$, which is the average replacement rate for a worker with medium earnings for cohorts born between 1953 and the present.[15] I will consider both "early" and

[14]See Attanasio (1999) for a discussion of the inverse elasticity of substitution being close to unity.

[15]For taxes, see "Social Security Taxes" `https://www.ssa.gov/OACT/ProgData/oasdiRates.html`, for ben-

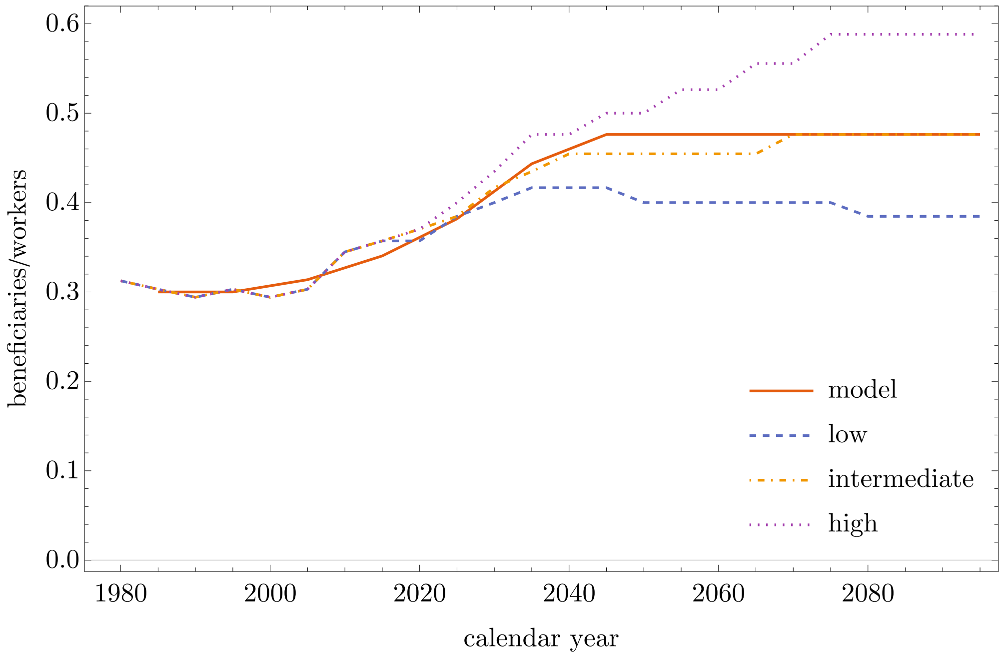

Figure 1: Ratio of beneficiaries to retirees. The solid red line in this graph illustrates the ratio of beneficiaries to retirees in the calibrated, six period model. The demographic change is modeled as a one-time reduction in the population growth rate n from 0.1802 to 0.01635. The dashed blue, and dash-dotted yellow, and dotted purple lines show the low, intermediate, and high cost projections of the ratio from the 2019 SSA Trustees' Report. The model calibration closely tracks the intermediate assumptions from the SSA.

"late" reforms (in 2025 and 2035, respectively). Late reforms will correspond to estimates from the Social Security Administration for the size of reform needed for long-run solvency if reform is delayed until the trust fund is depleted in the year 2034. Specifically, taxes would need to be increased by 3.65 points, or benefits would need to be cut by 23 percent (SSA (2019)). The early reforms I consider are the midpoint between the SSA's estimate of reforms that would be needed today, and the reforms that would be needed in 2034. The early and late reforms are sufficiently large to ensure that the economy converges to a steady state following reform and bonds do not grow infinitely. In the absence of reform, bonds would grow infinitely.

In the learning model, I set the gain parameter $\gamma = 0.93$. This gain parameter minimizes the maximum welfare cost to an agent of using adaptive forecasts in the learning model along a transition path that includes a demographic shock and a change to Social Security. It is also reasonable to use a large gain parameter because the length of a period is so long (10 years). The learning literature generally uses infinite horizon models where each period

efits, see Social Security Administration Actuarial Note 2019.9, `https://www.ssa.gov/oact/NOTES/ran9/an2019-9.pdf`.

is assumed to be short, perhaps one quarter (3 months). Gain parameters in constant-gain learning papers are usually around 0.05, which corresponds to 0.87 for a 10-year period.[16] The simulation results are sensitive to the gain parameter. Gain parameter selection and sensitivity are discussed in Section 7.1.

5 Announced Social Security Reform

Before examining policy uncertainty, it is useful to review the transition dynamics following an announced (or certain) Social Security reform. The example that follows is depicted in Figure 2. The model begins with an initial population of all age groups in 1935. The growth rate of the population slows in 1985, resulting in a gradual increase in the ratio of Social Security beneficiaries to workers. In order to continue financing benefits, the government increases the Social Security payroll tax in 2025. The demographic change and tax change are known from the beginning of the simulation (1935). The tax increase is parameterized to correspond to projections from the Social Security Administration, as described in Section 4. Figure 2 plots the path of capital and bonds for this experiment in the baseline full information rational expectations model (solid red line), and in the adaptive learning model (dashed blue line). Each will be described briefly.

There is no uncertainty in this example, and so in the baseline full information rational expectations model, this is equivalent to saying the demographic change and tax increase take place with perfect foresight. Agents in the model fully anticipate how changes in the demographic structure and Social Security system will impact their own consumption and saving, as well as the path of endogenous variables, such as the capital stock (and thus wages and interest rates). The stock of capital per worker k begins to fall before the demographic change as young agents in the model anticipate how the change will affect them in old age. The transition path to the new steady state is non-monotonic, with capital per worker increasing following the demographic change, falling, and then increasing again following the tax increase. Government bonds increase following the demographic change, and fall only after taxes are increased, converging to a new, higher steady state.

In the adaptive learning model, agents know the demographic change is coming, but they don't know how that will impact endogenous variables because they are assumed to have limited information about the structure of the economy. Therefore, they do not respond to

[16] Note: $1 - (1 - 0.05)^{40} = 0.8714$. See Mitra et al. (2013) and references therein for a discussion of gain parameter size.

the demographic change until after it happens and prices change. Following the decrease in the population growth rate, the adaptive learning agents update their expectations and extrapolate the change they have observed forward. This leads the capital stock to fall below the new steady state and then converge with an oscillatory pattern, consistent with the related literature (e.g. Mitra et al. (2013) and Gasteiger and Zhang (2014)). These fluctuations reflect periods of optimism or pessimism followed by subsequent corrections.

The expectations of the adaptive learning model are shown in Figure 3. This figure plots agents' time t expectation of interest rates (R^e_{t+1}), wages (w^e_{t+1}), and bonds (b^e_{t+1}), against the realized interest rates (R_{t+1}), wages (w_{t+1}), and bonds (b_{t+1}). In the initial steady state, agents' expectations match the data. Agents in the learning model overestimate the interest rate and underestimate the wage in the first period following the change in the population growth rate (which moves the economy away from the initial steady state). They are overly optimistic about the interest rate and overly pessimistic about wages. Because agents are expecting a higher interest rate and a lower wage, they save (relatively) more. This increases the capital stock relative to the rational expectations model. The increased capital stock decreases the interest rate and increases wages, which causes agents to save less, which eventually drives down the capital stock. This oscillatory pattern continues along the convergence path to the new steady state. The gain parameter is large for this example $\gamma = 0.93$, and the expectations track the data from the model closely. If the gain parameter were smaller, it would take longer for the expectations of the households to catch up with the data from the model.

6 Policy Uncertainty

The policy change of the previous section is non-stochastic and announced. In this section, I explore the path of the economy when Social Security policy is uncertain, or risky.[17]

I assume that policy uncertainty takes the following form: reform is possible in either 2025, or one model-period later in 2035. Within each period, two reforms are possible, either a benefit cut or a tax increase. Thus, there are four possible paths for the economy. The probability of each path is $p = 0.25$.[18] It is as if Congress flips a coin to decide if they

[17] I follow Caliendo et al. (2019) and use the words "uncertainty" and "risk" interchangeably in this paper. All of the examples I will consider have known probabilities, and thus might be called "risky."

[18] Nelson (2017) uses a similar stochastic process, and presents results for different probabilities of a given reform. Changing the probability of a given reform can qualitatively change the transition dynamics leading up to and following the reform. For example, if the probability of a benefit cut is higher, this

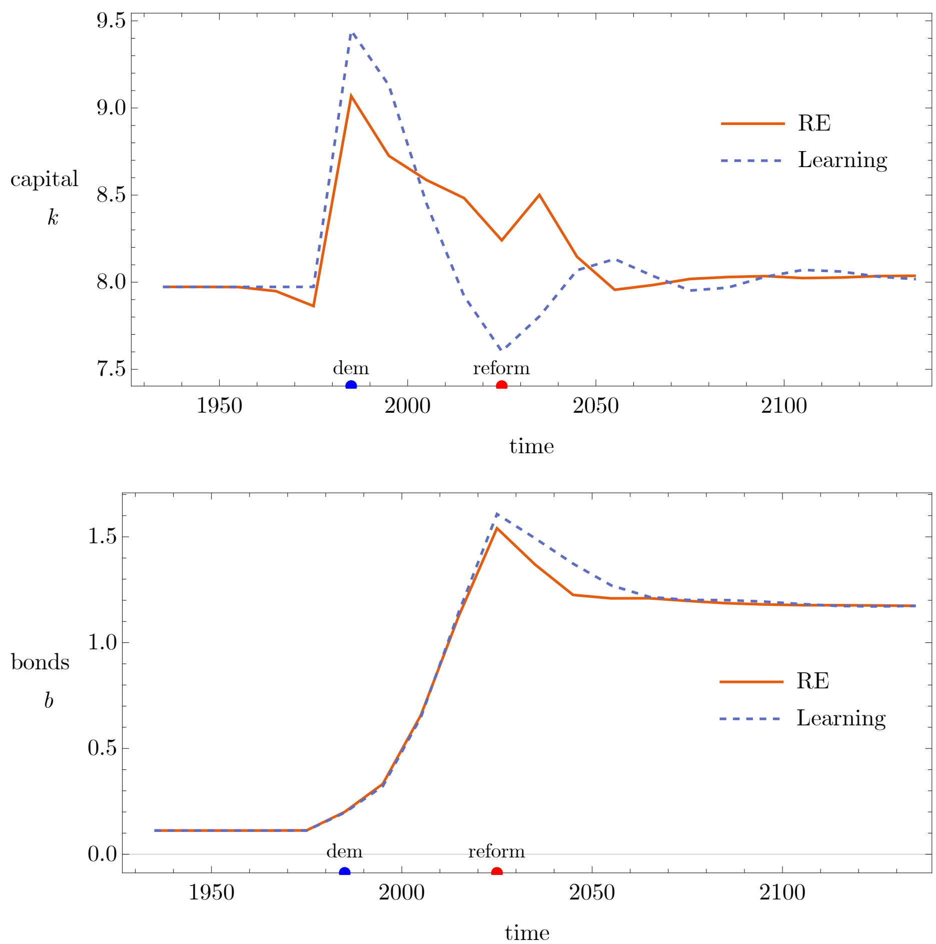

Figure 2: Paths of capital and bonds following a demographic change in 1985 and a Social Security tax increase in 2025. The path of the economy under rational expectations is the solid red line, the path with adaptive learning is the dashed blue line.

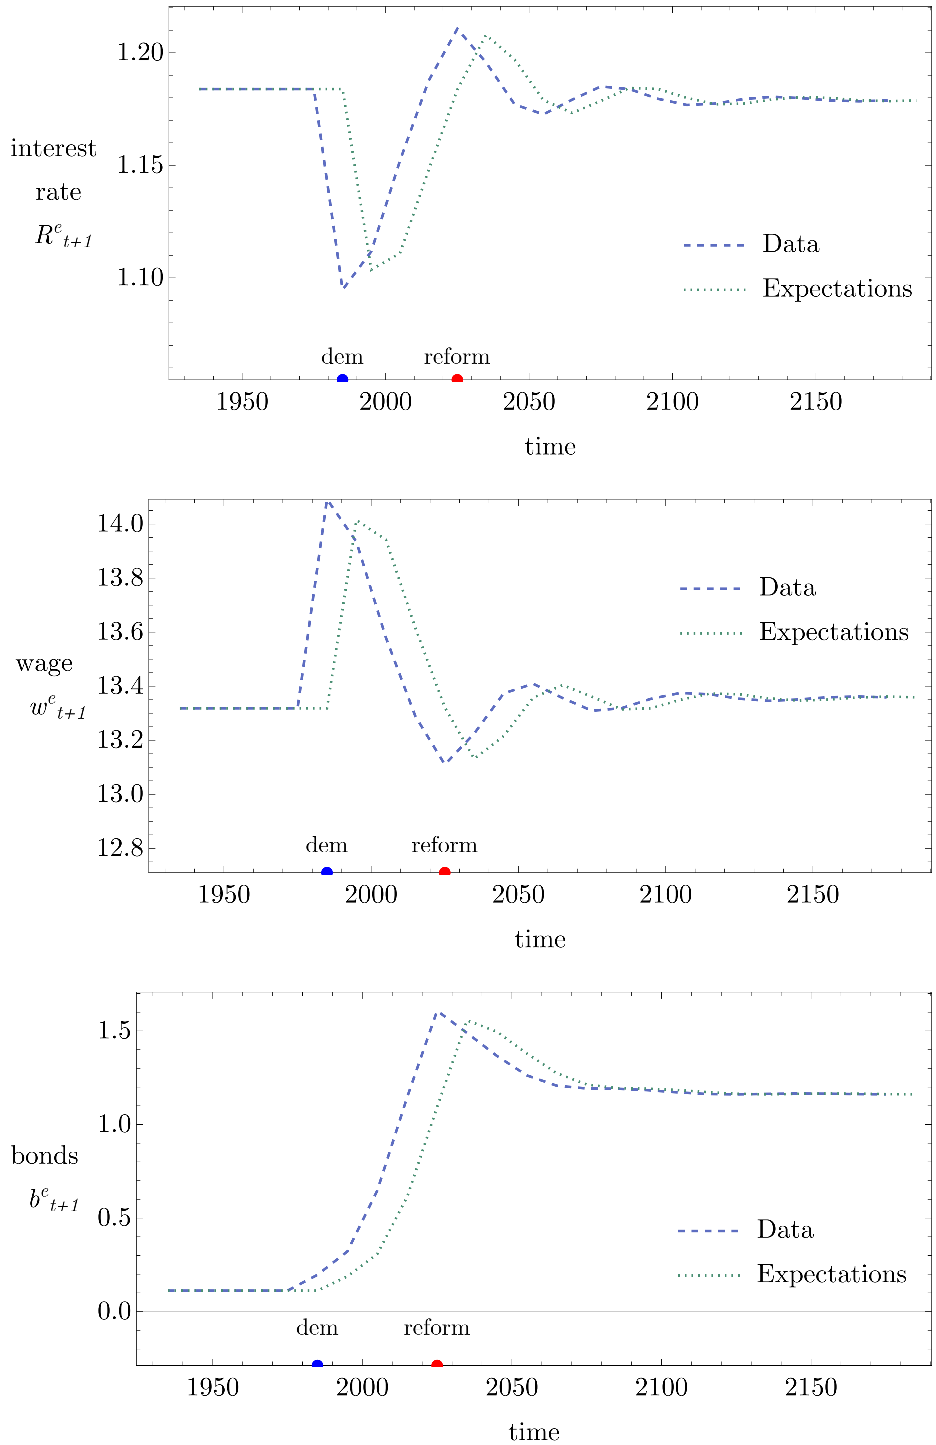

Figure 3: Time paths for expected interest rates, wages, and government debt levels (dotted teal lines), and realized interest rates, wages, and debt (dashed blue lines) in the learning model following a demographic change in 1985 and a Social Security tax increase in 2025.

should take action, and then they flip a coin again to see what type of action they should take. This is a simplification compared to the actual uncertainty Americans face regarding the future of Social Security; reform could take place in more dates, and could be a variety of reforms not considered in this paper.[19]

All agents in the economy know the four possible reforms and the relative probability of each. This is true in both models. The left side of Figure 4 shows the transition dynamics for capital and bonds of the four possible equilibrium paths in the rational expectations model. The left side of Figure 5 shows the path of age-specific asset holdings. The four possible paths for asset holdings are identical up to 2015. Agents in the initial years face the same uncertainty about future Social Security policy. This uncertainty is either fully or partially resolved in 2025. Along the first two paths of the economy, agents observe that policy was reformed in 2025 (either as a tax increase or benefit cut). Uncertainty is resolved for these two paths. Along the other two paths, agents observe that policy was not reformed in 2025, so they know reform will take place in 2035, but they do not know which reform until the policy is realized. In 2035, all uncertainty is resolved, and all four paths are different. There are four steady states for this example, one for each of the reforms.

Capital and bonds are both predetermined variables; the saving in period t corresponds to capital and bonds in period $t+1$. Thus, the four possible paths for capital and and bonds in Figure 4 are identical until 2025, because the asset choices of individuals were the same along all four paths until 2015.

The right side of Figures 4 and 5 plots the paths of capital and bonds and age-specific assets (saving) from the adaptive learning model. The agents in the adaptive learning model combine *complete* information about the policy process (the possible reforms and the likelihood of each) with *incomplete* information about the structure of the economy (market clearing and pricing equations). The adaptive learning agents know that Social Security policy change will influence their own ability to save and consume, but they are unsure how the policy changes will impact the aggregate economy. They are also unsure how changing demographics will influence the economy. The agents form adaptive forecasts of the endogenous variables they need in order to make their saving consumption decision. Each period, they update their forecasts as they observe new information. Prior to the change in the pop-

increases the precautionary saving of young households.

[19] Caliendo et al. (2019) use a different modeling framework and are able to explore a much richer distribution of policy uncertainty: in their model reform can take place at any moment in the life cycle and can be any combination of benefit cut and/or tax increase. For computational reasons, I did not explore such complex uncertainty in this general equilibrium overlapping generations model.

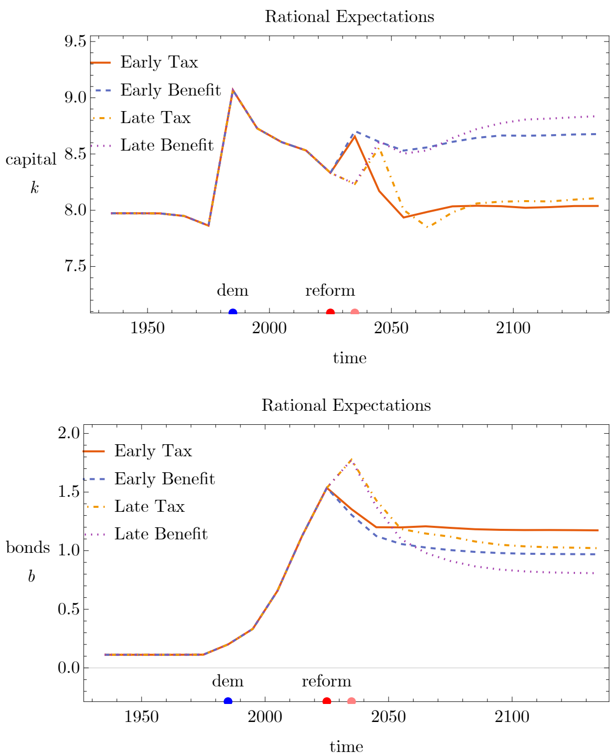

Figure 4: Time paths for capital and bonds for the four possible paths under policy uncertainty. The left graphs are for the rational expectations model. The graphs on the right are for the adaptive learning model. The solid red line is a tax increase in 2025, the dashed blue line is a benefit reduction in 2025, the yellow dot-dashed line is a (larger) tax increase in 2035, and the purple dotted line is a (larger) benefit cut in 2035.

ulation growth rate, the adaptive learning agents over-save compared to the rational agents because they are overestimating the interest rate. This drives up capital in the learning model compared to the baseline model. Following a policy reform, capital falls below the steady state and then converges.

I consider a relatively simple type of policy uncertainty in which there are only four total possible reform options and only two possible reform dates. If I were to include more policy options, this would impact the saving choices of households and the dynamics of the economy. For example, if a third policy option were possible that raised taxes a little and cut benefits a little, this could either increase or decrease the incentive for cohorts to save. If the third option were more likely than the full benefit cut option, this would reduce the incentive to save for cohorts who might experience the smaller benefit cut. Similarly, if more reform dates were possible, policy uncertainty would impact a greater number of generations, and would impact the generations differently than the simple uncertainty considered here. For example, if a particular cohort faced the possibility of a benefit reduction in three periods with equal probability (middle age, early retirement, and late retirement), rather than two periods with

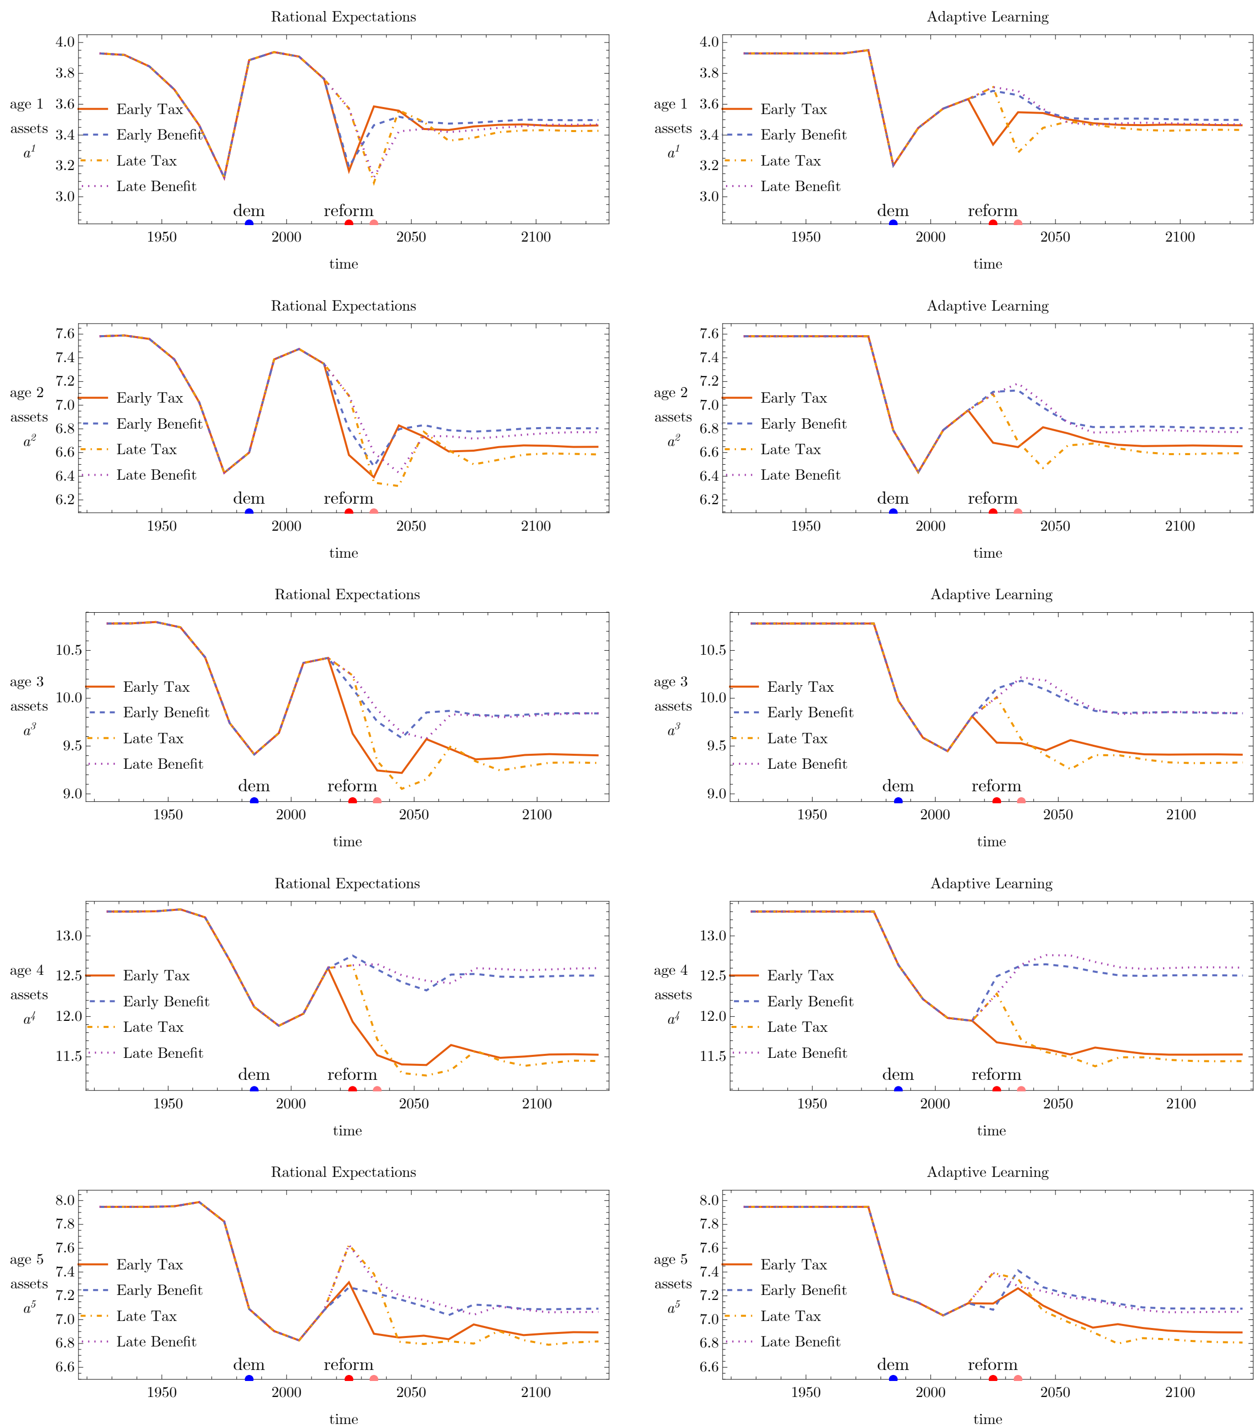

Figure 5: Time paths for age-specific assets (savings) for the four possible paths under policy uncertainty. The left graphs are for the rational expectations model. The graphs on the right are for the adaptive learning model. The solid red line is a tax increase in 2025, the dashed blue line is a benefit reduction in 2025, the yellow dot-dashed line is a (larger) tax increase in 2035, and the purple dotted line is a (larger) benefit cut in 2035.

equal probability (early retirement and late retirement), they would save relatively more, since the chance that their benefits would be cut for all of their retirement years is greater. Caliendo et al. (2019) consider a richer set of reform dates and a continuous set of reform options, ranging from a benefit cut only to a tax increase only.

6.1 Ex-ante Welfare Effect of Policy Uncertainty

I consider the ex-ante welfare effect of policy uncertainty for generations alive before or during the possible Social Security reform. I calculate this effect using a consumption equivalent variation (CEV) that equates the expected utility of consumption (in the world *with* policy uncertainty) to the utility of expected consumption plus a fraction Δ (i.e., the expected value of consumption over the possible policies without uncertainty, plus Δ). The ex-ante CEV for policy uncertainty is calculated for each cohort alive before or during the reform as:

$$E[U\left(c(uncertain)\right)] = U\left(E[c(certain)(1+\Delta)]\right) \tag{16}$$

where U indicates discounted lifetime utility, $U = \sum_{j=1}^{J} \beta^{j-1} u(c^j_{t+j-1})$, and expectations E are taken with regard to the possible policy paths. That is, letting $i = 1, 2, 3, 4$ indicate the four possible policy paths (early tax increase, early benefit cut, late tax increase, late benefit cut) and p^i indicate the probability of a given path, Equation 16 can be written as:

$$\sum_{i}^{4} p^i U\left(c(i, uncertain)\right) = U\left(\sum_{i}^{4} p^i c(i, certain)(1+\Delta)\right). \tag{17}$$

Here $c(i, uncertain)$ is the consumption of an individual in the world with policy uncertainty along the path where policy i is implemented. Note that $c(i, uncertain)$ is the same along all four paths until 2025 (when the early reforms take place). On the right hand side of the equation, $c(i, certain)$ is the consumption of an individual where policy i happens with certainty–that is, the policy change is announced from the beginning of the model.

The consumption equivalent variation represents the fraction of lifetime consumption that a young individual with the expected value of consumption over all possible realizations of the reform would need to be given such that her utility is the same, ex-ante, as her expected utility of consumption with policy uncertainty. A negative consumption equivalent variation indicates that policy uncertainty harms the individual and she would be willing to pay to avoid uncertainty. The welfare effect is *always* negative in the standard, fully rational, model. This reflects risk aversion. Indeed, this consumption equivalent variation measures the individual's willingness to pay to insure herself against the policy risk.

This ex-ante consumption equivalent variation is similar in spirit to the welfare cost calculated in Caliendo et al. (2019). In both cases, the welfare metric captures the ex-ante effect of facing policy uncertainty, taking into consideration how the policies change the consumption opportunities of the individual.

The ex-ante welfare effects for both the rational and adaptive models are presented in Figure 6. The ex-ante welfare effect is small and negative in the rational model. The effect is not the same for each generation, since different cohorts face the possibility of reform at different stages in their life cycle. For example, the cohort that enters the model in 1995 (at age 25) faces reform in their prime working years or retirement when they are either 55 (early reform) or 65 (late reform). This cohort would be willing to pay the most to avoid policy uncertainty. The consumption equivalent variation is -0.018% for this cohort. That is, an individual in the 1995 cohort would be willing to give up 0.018% of her lifetime consumption to live in a world without policy uncertainty where her consumption is equal to the expected value of consumption across the policy reforms. In contrast, the cohort that enters the model earlier in 1975 is only impacted by reform if it takes place in 2025 (when they are 75). Thus, the welfare cost is smaller (-0.0012%) for this older cohort. All of the cohorts entering the model between 1975 and 2015 face uncertain reform during their life cycle; the consumption equivalent variations for each of these cohorts is between -0.012% and -0.018%. This small negative welfare effect is similar to the welfare cost of facing Social Security policy uncertainty calculated by Caliendo et al. (2019). They find that an individual with average wage earnings would be willing to give up 0.02% of life-cycle consumption to live in a world with no uncertainty, endowed with expected income across the policy options. Caliendo et al. (2019) model wage heterogeneity in a single cohort, while my paper models overlapping cohorts with a representative agent in each cohort.

The ex-ante welfare effect of policy uncertainty in the adaptive learning model differs from the welfare effect in the rational model in three main ways. First, the welfare effect in the learning model is not always negative; second, the welfare effect is often larger in magnitude; and third, the uncertainty impacts a larger number of generations before a reform. All three differences are visible in Figure 6.

In the rational model, the welfare effect of policy uncertainty is always negative. This is because of risk aversion in the utility function; an individual is necessarily worse off ex-ante facing a risky consumption stream compared to certain consumption equal to the expected value of the risky consumption. In the adaptive learning model, in contrast, the welfare effect is sometimes positive. An individual in the adaptive learning model is still risk averse (the same utility function is used in both models) and does not enjoy uncertainty for the sake of uncertainty. However, because the policy uncertainty indirectly changes the agents' forecasts of future endogenous variables, it can lead to self-fulfilling cycles of optimism or pessimism. The policy uncertainty amplifies the cycles by changing how much agents in the model save. This saving increases the capital stock, which increases wages and reduces interest rates. The adaptive agents forecast these changes into the future and reduce their saving, which lowers the capital stock and reduces wages and increases interest rate. Then the agents save more, which increases the capital stock, and so on and so on. If a particular cohort is young during a period of high wages, they benefit from these endogenous cycles. Similarly, if a cohort is older when interest rates are high, they can benefit. Thus, the feedback effects of policy uncertainty can be positive for some cohorts because they experience higher consumption in the world with policy uncertainty than they would have if they consumed the expected value of consumption across the possible policy paths.

The fact that uncertainty sometimes increases welfare for adaptive learning agents is related to several papers that explore optimism, irrationality, and saving externalities. Feigenbaum et al. (2011) explore the pecuniary externality of saving in a general equilibrium model. They show agents in an overlapping generations economy can exploit this pecuniary externality by coordinating behavior across generations. The welfare gains to some generations in my model come from the same underlying feedback loop between saving and capital. In a somewhat similar vein, Caprioli (2015) explores the trade-off between distortions from taxes and distortions from expectations. Caprioli shows that if agents form expectations using a learning algorithm, a benevolent planner could choose policy to manipulate expectations, reducing taxes and issuing debt at times of pessimism and doing the opposite at times of optimism. Finally, the potentially positive welfare effect of policy uncertainty is related to Brunnermeier and Parker (2005). They show that forward looking agents can experience higher current utility if they are optimistic. The optimism leads to worse decision making, but potentially higher anticipatory utility.

The second main difference between the fully rational model and the adaptive learning model is the size of the welfare effect of policy uncertainty. The ex-ante welfare effect of policy uncertainty in the learning model ranges between -0.29% and 0.21%. This is more than an order of magnitude larger than in the rational model, where the effect ranges between -0.012% and -0.018%. The consumption equivalent variation compares expected utility of consumption with policy uncertainty to expected consumption from the possible reforms, as described in equations 16 and 17. In the rational model, the consumption on the left-hand side of the equation (used to calculate expected utility of consumption) is very similar to the consumption on the right-hand side of the equation (expected consumption across the four possible policy paths). The welfare cost comes from the curvature of the utility function. In contrast, in the adaptive learning model, the consumption on the left- and right-hand sides of equation 16 are less similar. The difference between these consumption streams contributes to the welfare cost in addition to the risk aversion in the utility function. Individuals in the adaptive learning model do not understand how policy changes or policy uncertainty will impact wages, interest rates, and government debt. So, although they respond to potential policy changes, the difference between their consumption (facing uncertainty) and the expected consumption averaged across all possible paths is larger than in the rational model. Therefore, the difference between expected utility of consumption (with policy uncertainty) and utility of expected consumption (averaged across the possible policy paths) is larger in the adaptive learning model compared to the rational model.

The result that the welfare effects of policy uncertainty are larger in the behavioral model relative to the fully rational model is consistent with evidence from Caliendo et al. (2019). They calculate the welfare cost of facing policy uncertainty for a hand-to-mouth consumer. They find that an individual with average wage earnings who cannot save would be willing to give up 1.23% of life-cycle consumption to live in a world with no uncertainty, endowed with expected income across the policy options, compared to 0.02% for a fully rational consumer.[20] By construction, this individual is not able to save to self-insure against policy risk. Thus, facing the possibility that Social Security benefits (the only source of retirement income for a non-saver) might be reduced is costly to the individual. In contrast, in the adaptive learning model in this paper, the individual saves, but she does not know the rate of return or future wages, and so she might save too much or too little. In both cases, the individual is limited relative to the fully rational model and is not able to self-insure as effectively.

The third main difference is which cohorts are impacted by the policy uncertainty. In the rational model, the ex-ante welfare effect of facing policy uncertainty only impacts individuals who are alive during a possible reform (the effects for older generations are very close to zero).

[20]This welfare effect is for an individual facing policy option *and* policy timing uncertainty. The effect is smaller for an individual who knows reform will be a tax increase, but does not know the timing of reform.

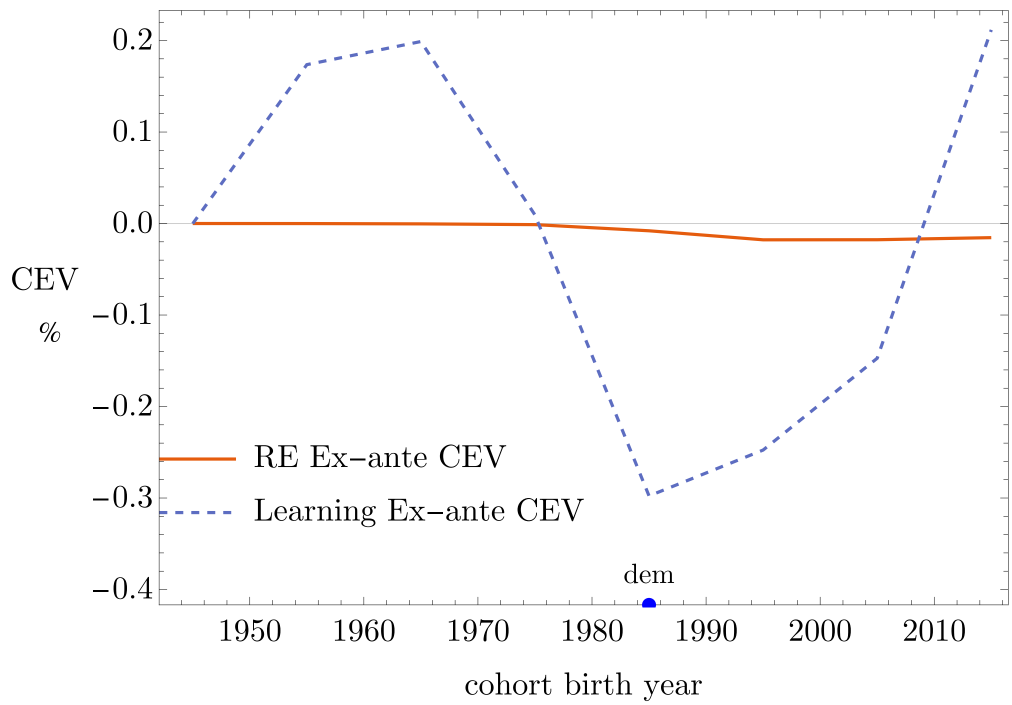

Figure 6: Ex-ante welfare effect of policy uncertainty. The solid red line corresponds to the baseline full-information rational expectations (RE) model. The dashed blue line corresponds with the adaptive learning model. The CEV compares expected utility of consumption to the utility of expected consumption. A negative number indicates the household would be willing to give up consumption to avoid policy uncertainty.

Given the parameterization of the model, only the cohorts born between 1975 and 2025 face policy uncertainty. In contrast, in the adaptive learning model, cohorts of agents alive *before* a reform who are not alive *during* the reform can still be impacted by policy uncertainty. If an older generation overlaps with younger generations who face policy uncertainty, the older generation can be impacted by the general equilibrium feedback of the decisions of the younger generations who face uncertainty. The adaptive learning agents do not understand how policy changes or policy uncertainty will influence wages or interest rates and so they save and consume differently than they would if they had full information. This over- or under-consumption changes the capital stock which changes prices which in turn, impacts the welfare of all cohorts who are alive at that moment, even if some of the cohorts do not face the policy changes or uncertainty in their lifespan.

Policymakers may be interested in the average ex-ante welfare effect of policy uncertainty. To measure this, I calculate a "social" consumption equivalent variation (CEV) that equates the average expected utility of consumption (in the world with policy uncertainty) to the average utility of expected consumption plus the CEV. This consumption equivalent variation represents the average fraction of lifetime consumption that young individuals with the expected value of consumption over all possible realizations of the reform would need to be given such that their average utility is the same, ex-ante, as the average expected utility of consumption with policy uncertainty. This calculation considers the discounted lifetime utility of cohorts of agents born from the beginning of the simulation (in 1935) to the time of late reforms (2035). I consider two different welfare weights when constructing the averages. In the first, each individual receives equal weight; in the second, each cohort receives equal weight (and thus individuals in younger cohorts are discounted because their cohorts are larger). In the rational expectations model, the ex-ante social welfare effect of policy uncertainty is -0.009% when all individuals receive equal weight, and -0.008% when each cohort receives equal weight. This suggests that on average, agents born before or during a reform would be willing to give up less than a basis point of lifetime consumption to avoid policy uncertainty. In the adaptive learning model, the ex-ante social welfare effect of policy uncertainty is -0.020% when all individuals receive equal weight, and -0.008% when each cohort receives equal weight. Although some cohorts are ex-ante better off with policy uncertainty, on average, the agents born before or during an uncertain reform would be willing to pay a few basis points of lifetime consumption to avoid policy uncertainty.

6.2 Welfare Effect of Policy Uncertainty: Future Generations

The effects of policy uncertainty are not isolated to cohorts of agents born prior to a reform. Agents entering the model after a reform is enacted are also impacted by the macroeconomic effects of the policy uncertainty from the preceding periods. For example, the capital stock could be higher because of the precautionary saving of previous cohorts. One way to examine this effect is to construct an ex-post consumption equivalent variation.[21] Specifically, I calculate how much consumption an agent born after a reform took place in the world without uncertainty (the policy was announced at the beginning of the model in 1935) would need to be given to live in an alternate universe where the same policy was enacted as the result of the uncertain process. For example, suppose taxes are increased in 2035 (the late tax reform). The consumption equivalent variation is the fraction, Γ, that equates the discounted lifetime utility of consumption plus Γ of a particular cohort in a world where the 2035 tax increase was announced at the beginning of the model to the discounted lifetime utility of consumption of the same cohort in the world where taxes were increased in 2035 as

[21] I will use the phrase ex-post to describe this welfare effect; however, since all uncertainty has been resolved and there are no other shocks in the model, ex-post and ex-ante comparisons are the same for cohorts born after a reform is implemented.

the result of the uncertain process. If the consumption equivalent variation is negative, that means utility is lower in the world with policy uncertainty compared to announced policy.

Figure 7 plots the consumption equivalent variation for future cohorts, both in the standard rational expectations model, and in the adaptive learning model. This consumption equivalent variation is not directly comparable to the previous section. The consumption equivalent variation in this section includes the wealth effects of a particular reform. This consumption equivalent variation can be positive or negative in both the rational and adaptive models. For example, in the rational model, the consumption equivalent variation is positive for the first few cohorts following a tax increase reform. Agents who live in a world where taxes were increased as the result of the uncertain process have higher utility than agents who experienced the same tax increase in a world without policy uncertainty. This is because the previous generations saved a little more in the world with uncertainty (to offset the chance that their benefits would be cut), which raised the capital stock and increased wages. Similarly, the consumption equivalent variation is negative in the rational model for a few generations born after a benefit reduction.

The future-generation consumption equivalent variation in the rational model ranges between -0.32% and 0.32%. These welfare effects are much larger than the previous section because they include the wealth effect of a particular reform. The future-generation consumption equivalent variations are even larger in the adaptive learning model, ranging between -1.9% and 1.2%. The welfare effects are more than an order of magnitude larger in the adaptive learning model. The welfare effects of uncertainty are larger in the learning model compared to the rational expectations model because agents in the learning model do not anticipate how Social Security policy changes will impact wages, interest rates, and bonds. Even after a policy has been implemented, it takes the agents several periods to learn the new prices and bond levels. The welfare effects are most negative in the adaptive learning model for cohorts born three periods after one of the early reforms. The agents born after an early reform experience lower wages and thus lower consumption because their predecessors did not save as much as they would have if they had known the reform was coming in 2025 rather than possibly occurring in 2035.

The analysis of this section compares the realization of an uncertain policy to the same policy as if it were announced. An alternative approach is to compare the realization of each uncertain policy to the same announced baseline reform, as in Kitao (2018). This is presented in the on-line Appendix C.

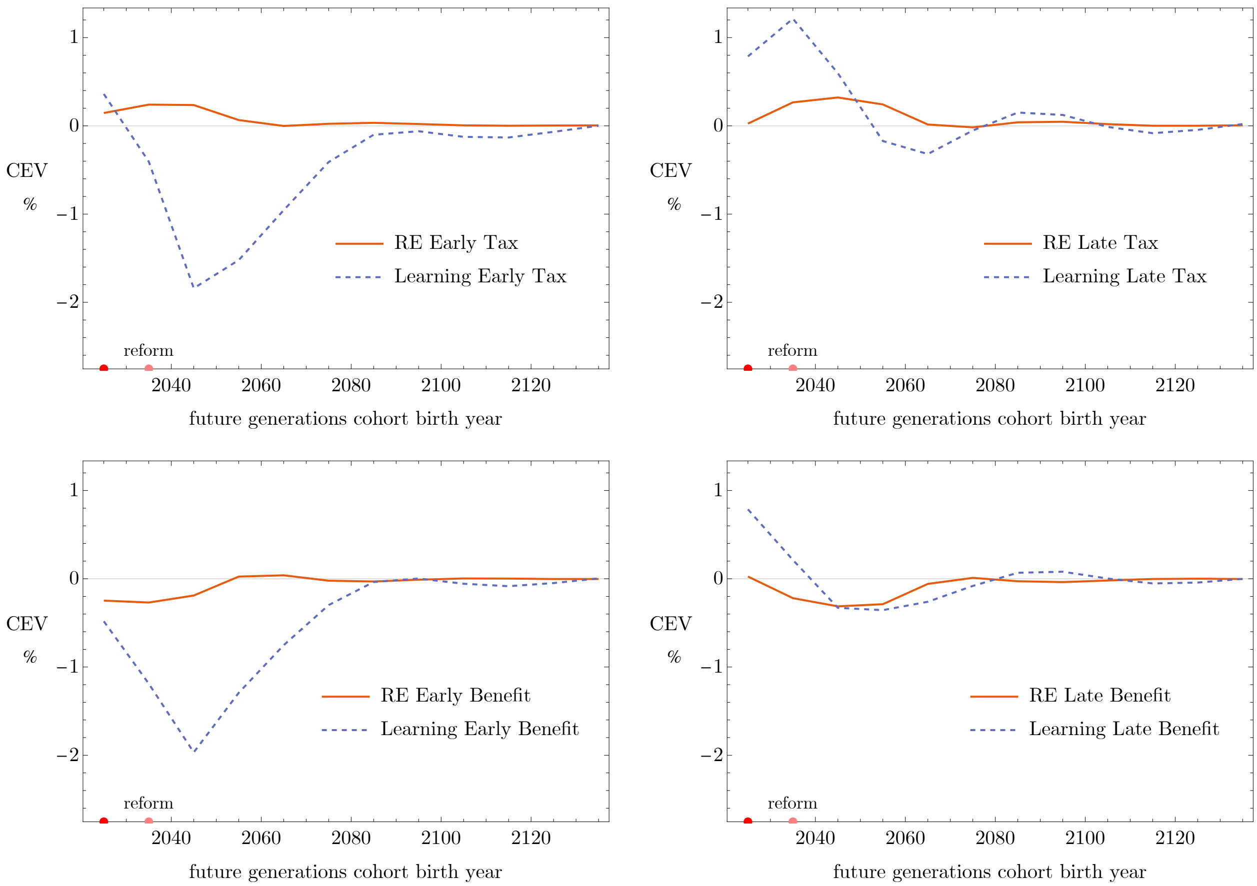

Figure 7: Future generations: welfare effect of policy uncertainty. The solid lines correspond to the fully rational model. The dashed lines correspond with the learning model. The CEV compares utility of consumption of certain (announced) policy to the utility of consumption where the same policy occurs as the result of the uncertain process. A negative number indicates the household is harmed by the policy uncertainty.

6.3 Alternative Welfare Cost: Initial Steady State

An additional ex-post welfare approach to measure the impact of demographic change and a particular policy reform is to compare the utility of consumption of each cohort along a particular transition path to the utility of consumption in the initial steady state. This comparison is depicted in Figure 8. This comparison calculates the fraction of lifetime consumption, η, such that a particular cohort would be indifferent between their own consumption and the consumption they would have experienced in the initial steady state plus η. A negative number indicates a cohort would have been better off in the initial steady state. Because the demographic change necessitates a reduction in the generosity of Social Security, all cohorts in the standard rational model are worse off along the transition path compared to the initial steady state. Similarly, most cohorts are worse off in the learning model. Only a small number of cohorts who are old when taxes are increased (either as the result of policy uncertainty or because policy was announced) are better off than the initial steady state. The most harmed cohorts are those in the adaptive learning model born during or after a reform. The welfare cost is sometimes 2 percentage points lower in the adaptive learning model compared to the rational model.

The consumption equivalent variations in the rational model with certain reform and with uncertain reform are close in size for all cohorts across each of the four possible reforms. In contrast, sometimes there is a larger difference between the consumption equivalent variations in the adaptive model when there is policy uncertainty compared to announced policy. This is particularly visible when the policy enacted is either the early tax increase or the early benefit cut. The consumption equivalent variations for cohorts entering the model right after the early reform takes place are worse off when the policy was the result of an uncertain process rather than announced policy. This is consistent with Section 6.2, which showed the lowest consumption equivalent variations in the adaptive learning model for the early reforms.

6.4 Timing or Policy Option Uncertainty Only

The welfare analysis of sections 6.1, 6.2 and 6.3 focused on simultaneous policy option and timing uncertainty. Policymakers may want to know which dimension of policy uncertainty (timing or option) is driving the welfare effects. In this section, I consider policy option uncertainty and policy timing uncertainty separately.

Figure 9 shows the ex-ante consumption equivalent variation for four different types of

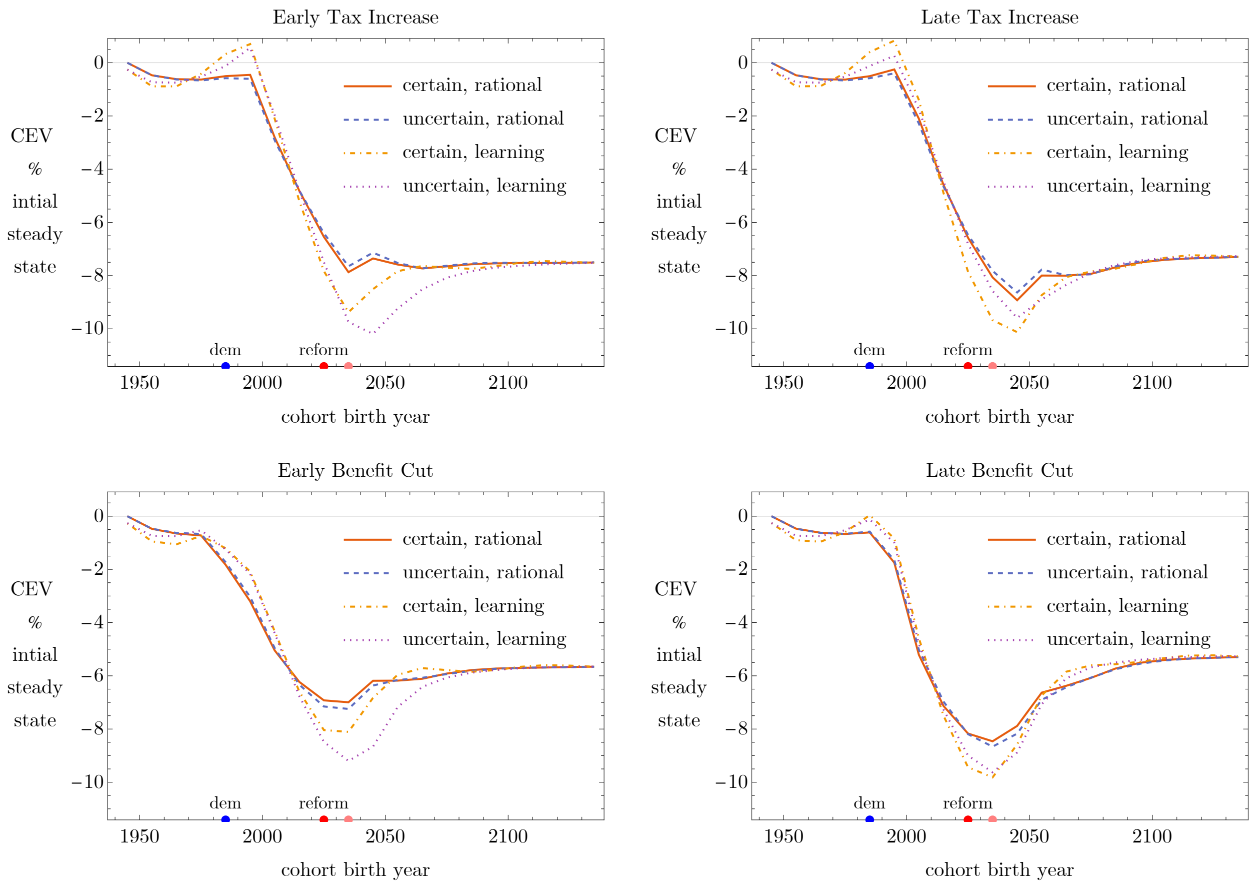

Figure 8: Comparison to initial steady state. The CEV compares utility of consumption of a particular path to the utility of consumption in the initial steady state. A negative number indicates that a particular cohort is worse off compared to the initial steady state. The top left shows the early tax increase, top right the late tax increase, bottom left the early benefit cut, bottom right the late benefit cut. The solid red line is the rational model with certain (announced) policy, the dashed blue line is the rational model with policy uncertainty, the yellow dash-dot line is the adaptive learning model with certain policy, and the dotted purple line is the adaptive learning model with policy uncertainty.

policy uncertainty: (i) policy timing uncertainty where taxes are increased in 2025 or increased by a larger amount in 2035, (ii) policy timing uncertainty where benefits are reduced in 2025 or reduced by a larger amount in 2035, (iii) policy option uncertainty where either taxes are increased or benefits are reduced in 2025, and (iv) policy option uncertainty where either taxes are increased or benefits are reduced by larger amounts in 2035. In all four scenarios the two reforms are equally likely. The consumption equivalent variation is defined as in Section 6.1 and represents the fraction of lifetime consumption that a young individual with the expected value of consumption over both possible realizations of the reform would need to be given such that her utility is the same, ex-ante, as the expected utility of consumption with policy uncertainty. A negative number indicates that policy uncertainty harms the individual.

In the fully rational model, policy option uncertainty is more harmful to households than policy timing uncertainty. This is shown on the left panel of Figure 9. The welfare effect is largest in magnitude for the cohort entering the model in 2005 facing 2035 policy option uncertainty. These agents face the possibility of taxes being increased or benefits being cut in the last stage of the working life right before they retire. Similarly, when reform takes place in 2025 but the policy option is uncertain, the most harmed cohort is the cohort that would experience the policy change in their last working years. In contrast, if individuals know that benefits will be cut (or that taxes will be increased), but they don't know exactly when, they are better able to prepare for the reduction in retirement income (or take-home pay). The welfare effects are smallest if households face timing uncertainty with tax increases.[22]

The adaptive learning model is shown in the right panel of Figure 9. The timing of reforms seems to matter more than the particular type of uncertainty. The welfare cost is smallest (the welfare benefit is largest) for the 1985-2015 cohorts if reforms are guaranteed to take place in 2025. The welfare cost is largest (the welfare benefit is smallest) for these cohorts when reforms are guaranteed to take place in 2035. The resolution of uncertainty helps agents' forecasts catch up to the data and reduces the welfare cost for cohorts who are alive after the reform.

[22]Caliendo et al. (2019) observe a similar pattern for low-income households in their analysis.

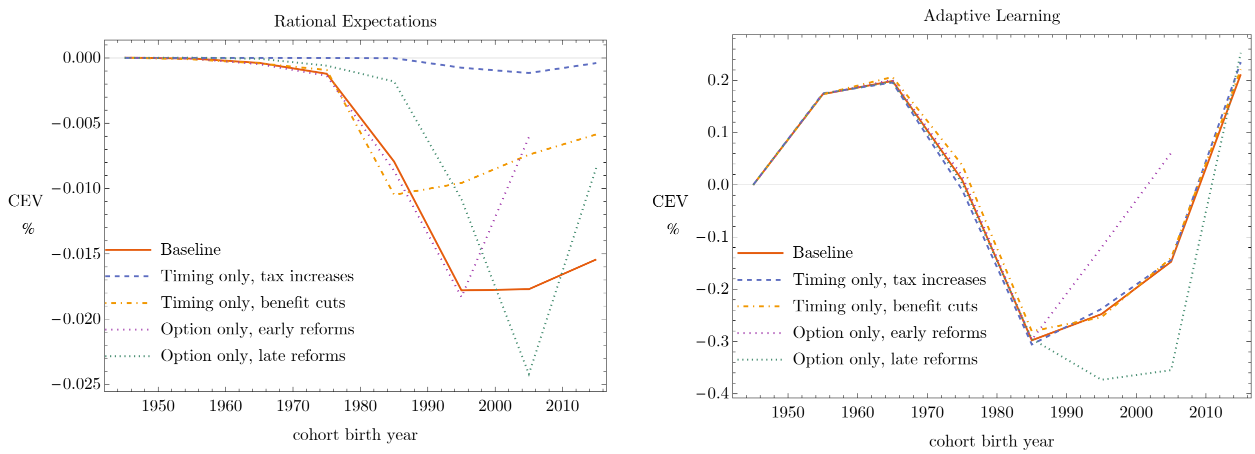

Figure 9: Ex-ante welfare effect of policy uncertainty in the rational model (left) and adaptive learning model (right). The solid red line is the baseline parameterization with timing and option uncertainty. The dashed blue line is timing uncertainty only (both reforms are tax increases), the dot-dashed yellow line is timing uncertainty only (benefit cuts only), the dotted purple line is reform option uncertainty only (both reforms in 2025), and the tiny green dotted line is option uncertainty only (both reforms in 2035). The CEV compares expected utility of consumption to the utility of expected consumption. A negative number indicates the household would be willing to give up consumption to avoid policy uncertainty.

7 Robustness

7.1 Gain Parameter Selection and Sensitivity

The gain parameter affects the transition paths for the learning economy. The use of adaptive expectations is equivalent to agents believing that the variables they are forecasting have a mixture of permanent and transitory shocks. If the agents knew that all of the shocks to the economy were permanent, it would be optimal to use a gain parameter equal to one. However, the agents are assumed to have incomplete information about the economy, and so they do not realize that all shocks are permanent. They update their forecasts with a gain parameter γ that is between zero and one.[23]

In the main sections of the paper, I use a gain parameter of $\gamma = 0.93$. This gain parameter minimizes the maximum welfare cost to an agent of using adaptive forecasts in the learning model along the transition path that includes a demographic shock and a change to Social Security. To measure this welfare cost, I construct a consumption equivalent variation for each cohort that compares the utility of consumption in the learning model to the utility of consumption for an "infinitesimally rational" agent. An infinitesimally rational agent is able to predict future prices with perfect foresight in the learning world, but is such a small part of

[23]For the particular policy changes considered, debt only converges for gain parameters $\gamma \geq 0.25$. A smaller gain parameter is possible with smaller demographic changes and/or different policy parameters.

the market that she does not change prices. I compute the consumption equivalent variation that makes an infinitesimally rational agent indifferent between her own consumption and the consumption of the learners. I chose the gain parameter γ to minimize the maximum welfare cost. It is unsurprising that the gain parameter is close to $\gamma = 1$, since all shocks in the model are permanent.

Figure 10 plots the path of capital and bonds with different gain parameters for the announced tax increase (the same exercise as in Section 5, Figure 2). As the size of the gain parameter falls, the learning agents make larger forecast errors (since they place more weight on their previous expectations and less weight on the new information they're observing), which generates larger swings in the endogenous variables, such as the capital stock. The endogenous cycles are larger in amplitude and magnitude when the gain parameter is smaller. A similar pattern is observed in Gasteiger and Zhang (2014), Mitra et al. (2013), and Evans et al. (2019) which are focused on permanent changes to fiscal policy. In contrast, in much of the learning literature the changes in models are transitory shocks, and so a smaller gain parameter leads to smaller swings in aggregate variables.

The ex-ante welfare analysis of Section 6.1 is conducted with different gain parameters below in Figure 11 and Table 2. As the gain parameter falls, the endogenous cycles are large in magnitude. It takes more periods for the economy to adjust because agents place less weight on new information. Thus, the most harmed cohort is a later cohort when the gain parameter is smaller. In Figure 11, the most harmed cohort is born in 1985 when the gain parameter is 0.93. The most harmed cohort is born one period later, in 1995, if the gain parameter is 0.6, and in 2015 if the gain parameter is 0.3. The same information is displayed in Table 2. Note that the gain parameter $\gamma = 0.93$ was chosen to minimize the welfare cost of making forecast errors using adaptive expectations instead of rational expectations in a world with certain policy. Thus the gain parameter does not minimize the welfare cost of experiencing policy uncertainty.

7.2 Robustness: Balanced Budget Reforms

In the baseline parameterization of the paper, the Social Security system runs a deficit before and after reform. The reforms are large enough to prevent debt from growing infinitely, but the reforms do not completely eliminate the Social Security deficit. An alternative parameterization is presented in this section. The Social Security system is assumed to run a balanced budget before the demographic shift, and the reforms enacted are such that

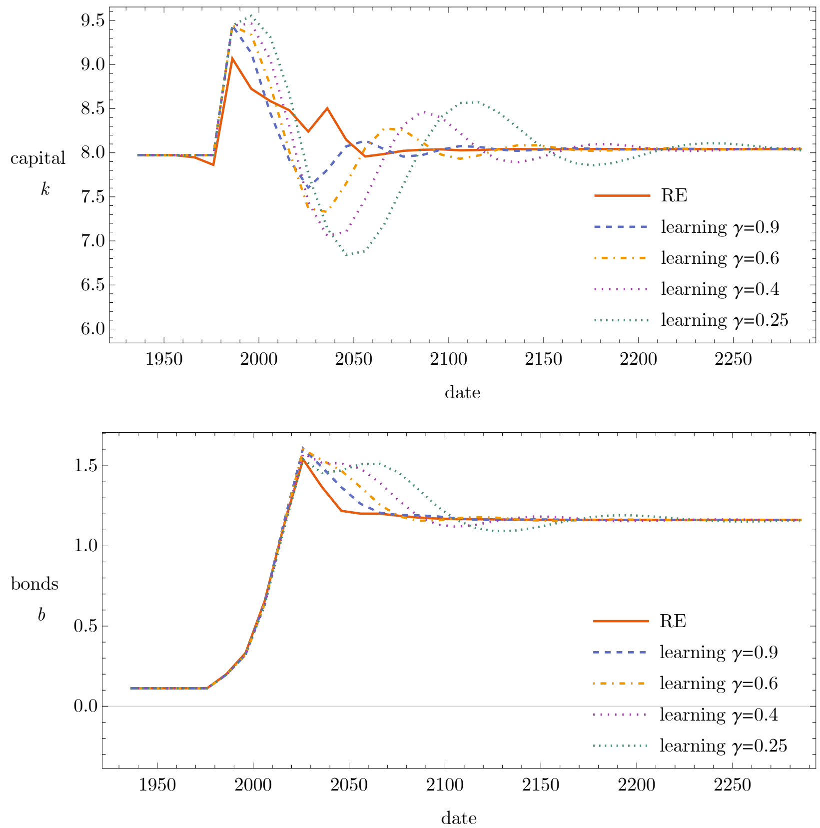

Figure 10: Paths of capital and bonds for different gain parameters γ following a demographic change and a Social Security tax increase in 2025. The path of the economy under rational expectations is the solid red line, learning with a gain parameter of $\gamma = 0.93$ is the dashed blue line, $\gamma = 0.6$ is the dot-dashed yellow line, $\gamma = 0.4$ is the dotted purple line, and $\gamma = 0.25$ is the tiny dotted green line.

	Ex-ante CEV				Social ex-ante CEV	
Gain γ	CEV min (%)	cohort most harmed	CEV max (%)	cohort most helped	individuals weighted equally	cohorts weighted equally
1	-0.273	1985	0.290	2015	-0.009	0.002
0.93	-0.294	1985	0.296	2015	-0.020	-0.008
0.9	-0.298	1985	0.212	2015	-0.025	-0.013
0.8	-0.295	1985	0.211	1965	-0.047	-0.032
0.7	-0.318	1995	0.213	1965	-0.068	-0.051
0.6	-0.317	1995	0.206	1965	-0.084	-0.065
0.5	-0.327	2005	0.189	1965	-0.086	-0.067
0.4	-0.371	2015	0.200	1975	-0.067	-0.050
0.3	-0.350	2015	0.189	1975	-0.020	-0.010
0.25	-0.285	2015	0.173	1975	0.013	0.018

Table 2: Minimum, maximum, and average ex-ante CEV and the birth year of the most harmed and most helped cohort for policy uncertainty for different gain parameters γ in the adaptive learning model.

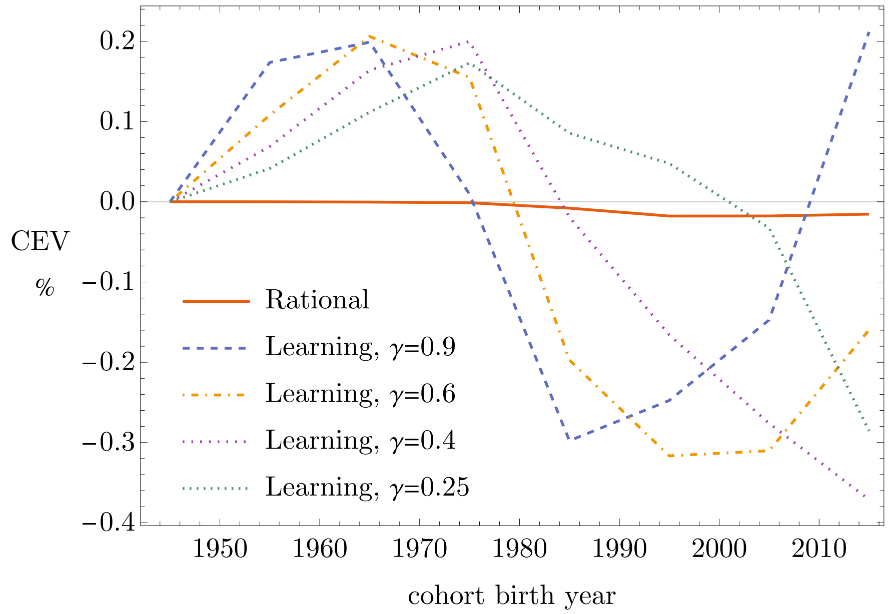

Figure 11: Ex-ante Consumption Equivalent Variation with different gain parameters. The solid red line corresponds to the baseline full information rational expectations (RE) model. The dashed blue line corresponds with the adaptive learning model with $\gamma = 0.93$, the dot-dashed yellow line is $\gamma = 0.6$, the dotted purple line is $\gamma = 0.4$, and the tiny dotted green line is $\gamma = 0.25$. The CEV compares expected utility of consumption to the utility of expected consumption. A negative number indicates the household would be willing to give up consumption to avoid policy uncertainty.

the Social Security budget is balanced in the terminal steady states (Social Security taxes collected equal benefits paid in the steady state).

The path of the economy for this new parameterization is presented in Figure 12. The initial Social Security tax rate is still set to the OASI portion of payroll taxes: $\tau = 0.106$. The initial benefit replacement rate is set to exactly balance the budget $\phi = 0.353$ (compared to 0.37 in the baseline parameterization). The two possible reforms that balance the budget following the demographic change are (i) increase the payroll tax to $\tau = 0.168$ (compared to 0.138 or 0.1425 in the baseline parameterization) and (ii) reduce the benefit replacement rate to $\phi = 0.223$ (compared to 0.296 or 0.285 in the baseline parameterization). Note that there are only two reforms, rather than four, since the reforms have been chosen to balance the budget in the steady state after the demographic change. The transition path to a particular steady state differs if the reform is enacted in 2025 or 2035, but the steady state is the same. Figure 12 depicts the possible paths for the economy when there is policy uncertainty and all four reforms are equally likely. The time paths for the economy are qualitatively similar in the baseline parameterization and in the balanced budget parameterization. Bonds increase as the population ages and converge to a steady state following reform. The swings in the

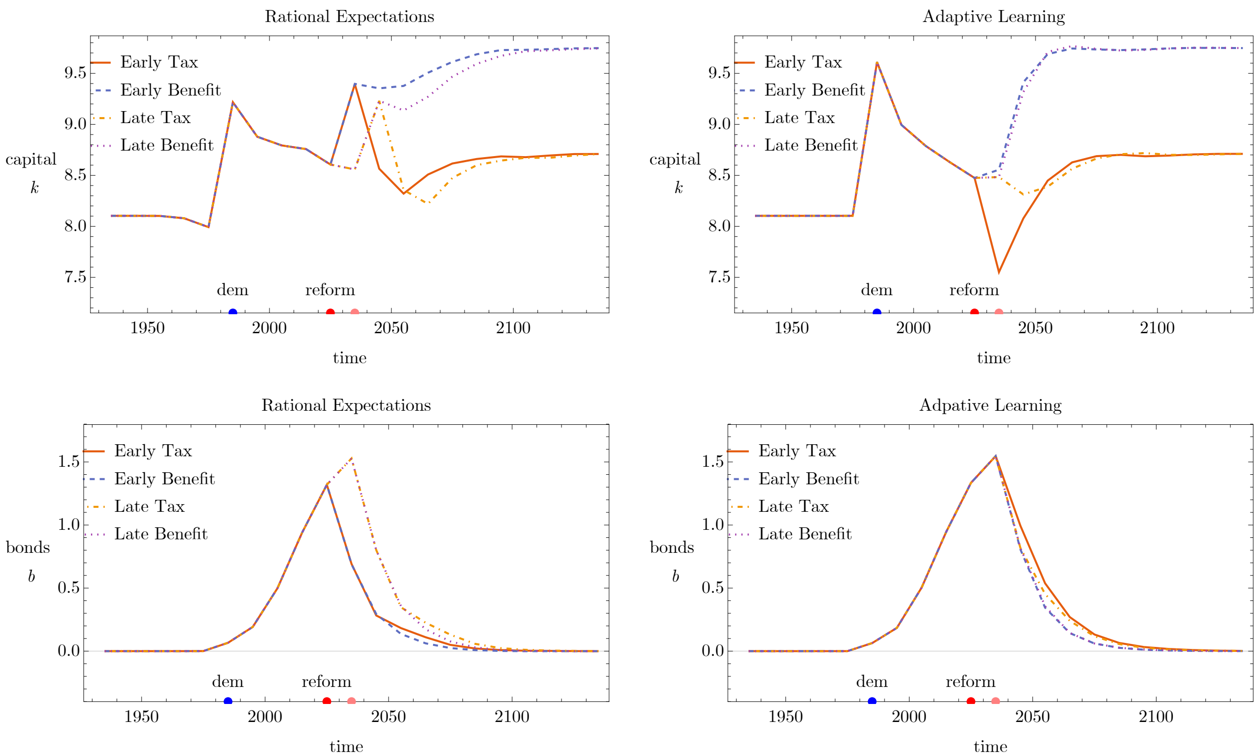

Figure 12: Time paths for capital and bonds for the balanced budget parameterization under policy uncertainty. Time paths for capital and bonds for the four possible paths under policy uncertainty. The left graphs are for the rational expectations model. The graphs on the right are for the adaptive learning model. The solid red line is a tax increase in 2025, the dashed blue line is a benefit reduction in 2025, the yellow dot-dashed line is the same tax increase in 2035, and the purple dotted line is the same benefit cut in 2035. The tax increase and benefit cut balance the Social Security budget in the terminal steady state.

capital stock are larger in the adaptive learning model compared to the fully rational model. The capital stock is lowest along the path with an early tax increase in the adaptive learning model.

The balanced budget reforms are larger in magnitude than the reforms in the baseline parameterization. As such, the welfare effects of facing policy uncertainty are also larger in magnitude. The ex-ante consumption equivalent variation for the balanced budget parameterization is presented in Figure 13. As in the baseline case, the welfare effect is always negative in the rational model. Recall, this consumption equivalent variation measures the fraction of lifetime consumption that a young individual with the expected value of consumption over all possible realizations of the reform would be need to be given such that her utility is the same, ex-ante, as the expected utility of consumption with policy uncertainty. The most harmed cohort in the rational model is still the cohort who enters the model in 1995 and experiences reform either when they are in their late working years or starting retirement. In the balanced budget parameterization, this cohort would be willing to pay

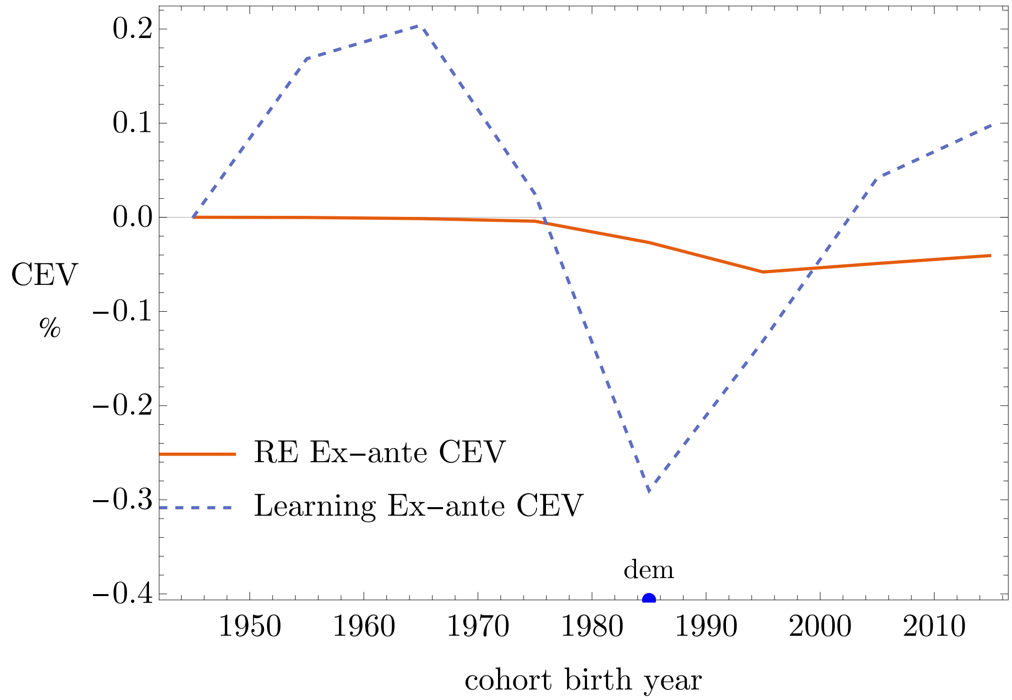

Figure 13: Ex-ante welfare effect of policy uncertainty, with the balanced budget parameterization. The solid red line corresponds to the baseline full-information rational expectations (RE) model. The dashed blue line corresponds with the adaptive learning model. The CEV compares expected utility of consumption to the utility of expected consumption. A negative number indicates the household would be willing to give up consumption to avoid policy uncertainty.

0.05% of lifetime consumption to avoid policy uncertainty. The welfare effects are larger in magnitude and sometime positive in the adaptive learning model. The most harmed cohort would be willing to give up 0.29% of lifetime consumption to avoid policy uncertainty. The cohort who benefits the most is better off by 0.20% of lifetime consumption.

With the balanced budget parameterization, the average ex-ante welfare effect of policy uncertainty in the fully rational model is -0.026% (-0.022%) with all individuals (cohorts) weighted equally. This is larger than the effect with the baseline parameterization, but has the same sign. The average ex-ante welfare effect of policy uncertainty in the adaptive learning model is 0.005% (0.014%) with all individuals (cohorts) weighted equally. This is different than the baseline parameterization where the average effect was negative. The average effect is positive in the adaptive learning model with the balanced budget reforms because only twp cohorts are harmed by the policy uncertainty (the 1985 and 1995 cohorts), while the others benefit ex-ante. In the baseline parameterization three cohorts are harmed by the policy uncertainty.

The welfare effects of policy uncertainty on future generations are similar in the balanced budget parameterization compared to the baseline parameterization. The welfare effects are a little larger in magnitude since the reforms are larger, but are otherwise qualitatively similar. The same comparison is true when examining policy option and policy type uncertainty separately; the results are similar in the balanced budget parameterization compared to the baseline parameterization.

8 Discussion and Conclusion

The analysis in the paper focuses on policy uncertainty in an adaptive learning model and a rational expectations model. Both models include a feedback loop between the expectations of households and the evolution of the aggregate economy. In both models, households make decisions based on their expectations of future factor prices and government debt levels. The central distinction between the models is how that feedback takes place. In the standard, rational expectations model, households form their expectations using the mathematical conditional expectation of the relevant variables using all available information, including knowledge of all equilibrium equations. In contrast, in the adaptive learning model, households combine *full knowledge* of the political process (which Social Security reforms are possible and when they could occur) with *limited knowledge* about the structure of the economy. Agents forecast endogenous variables using an adaptive rule, based on their observation of the economy from the previous period.

The ex-ante welfare effect represents the fraction of lifetime consumption that a young individual with the expected value of consumption over all possible realizations of the reform would be need to be given such that her utility is the same, ex-ante, as the expected utility of consumption with policy uncertainty. This welfare effect is equivalent to 0.029% of period consumption for the cohort of agents most harmed in the learning model compared to 0.018% percent in a rational expectations framework. Similarly, the social (or average) ex-ante welfare cost of policy uncertainty is equivalent to 0.02% of period consumption in the learning model compared to 0.009% percent in a rational expectations framework. The welfare implications for future generations are also larger in the adaptive learning model compared to the rational model.

Under rational expectations the welfare implications of policy uncertainty are relatively modest, but this hinges in large part on the assumption that agents can perfectly anticipate the general equilibrium effects of policy changes on the path of wages and interest rates. This paper highlights that the inability to forecast these general equilibrium effects significantly increases the welfare effects of policy uncertainty in the context of Social Security reform, both ex-ante and for future generations. The welfare effect of policy uncertainty is sometimes positive in the adaptive learning model. This is true from an ex-ante perspective for cohorts born before a reform takes place, and ex-post for cohorts who enter the model after a reform has been implemented. The positive effects are driven by the general equilibrium feedback of other cohorts who changed their saving because of the policy uncertainty. If policymakers wanted to increase saving, it would probably be better to try to change saving directly, rather than through policy uncertainty, although a direct comparison is beyond the scope of this paper. Although the welfare effects are positive for some cohorts in the learning model, on average, the welfare effects are negative, both ex-ante and for future generations.[24] Thus, policymakers could improve welfare by reducing uncertainty about the future of Social Security policy.

This paper focuses on the differences between an adaptive learning model and a rational expectations model with a single type of uncertainty: namely Social Security policy uncertainty. A large literature has examined the implications of boundedly-rational adaptive learning in models with other types of uncertainty, primarily focusing on the enhanced ability of bounded rationality models to fit the data.[25] In order to isolate this effect (and also for reasons of computational tractability) this paper has taken the simplest and most direct route of eliminating all stochastic shocks other than those connected with policy uncertainty. Future work could examine the impact on key results of extending the model to include productivity shocks and other sources of uncertainty.

[24]The social welfare effect is negative for gain parameters $\gamma \in [0.28 - 0.98]$.

[25]See, for example, Eusepi and Preston (2011), who examine the implications for sample moments and impulse response functions of productivity shocks within an RBC model; and Giusto (2014), who studies the implications of adaptive learning for income and wealth distribution within a Krusell and Smith (1998) incomplete-markets model.

References

Attanasio, O. (1999). Consumption. In Taylor, J. B. and Woodford, M., editors, *Handbook of Macroeconomics*, volume 1, Part B, chapter 11, pages 741–812. Elsevier, 1 edition.

Branch, W., Evans, G. W., and McGough, B. (2013). Finite horizon learning. In Sargent, T. and Vilmunen, J., editors, *Macroeconomics at the Service of Public Policy*, chapter 8, pages 141–163. Oxford University Press, Oxford.

Brunnermeier, M. K. and Parker, J. A. (2005). Optimal expectations. *The American Economic Review*, 95(4):1092–1118.

Bullard, J. and Duffy, J. (2001). Learning and excess volatility. *Macroeconomic Dynamics*, 5(2):272–302.

Bütler, M. (1999). Anticipation effects of looming public-pension reforms. *Carnegie-Rochester Conference Series on Public Policy*, 50:119 – 159.

Caliendo, F. N., Gorry, A., and Slavov, S. (2019). The cost of uncertainty about the timing of social security reform. *European Economic Review*, 118:101 – 125.

Caprioli, F. (2015). Optimal fiscal policy under learning. *Journal of Economic Dynamics and Control*, 58(Supplement C):101 – 124.

Chalk, N. A. (2000). The sustainability of bond-financed deficits: An overlapping generations approach. *Journal of Monetary Economics*, 45(2):293–328.

Cottle Hunt, E. (2019). Finite Horizon Life-cycle Learning. working paper.

Davig, T., Leeper, E. M., and Walker, T. B. (2010). "unfunded liabilities" and uncertain fiscal financing. *Journal of Monetary Economics*, 57(5):600–619.

Eusepi, S. and Preston, B. (2011). Expectations, learning, and business cycle fluctuations. *The American Economic Review*, 101(6):2844–2872.

Evans, G. and Ramey, G. (2006). Adaptive expectations, underparameterization and the lucas critique. *Journal of Monetary Economics*, 53(2):249–264.

Evans, G. W., Guesnerie, R., and McGough, B. (2019). Eductive stability in real business cycle models. *The Economic Journal*, n/a(n/a).

Evans, G. W. and Honkapohja, S. (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton.

Evans, G. W., Honkapohja, S., and Mitra, K. (2009). Anticipated fiscal policy and adaptive learning. *Journal of Monetary Economics*, 56(7):930 – 953.

Feigenbaum, J., Caliendo, F. N., and Gahramanov, E. (2011). Optimal irrational behavior. *Journal of Economic Behavior and Organization*, 77(3):285 – 303.

Gasteiger, E. and Zhang, S. (2014). Anticipation, learning and welfare: the case of distortionary taxation. *Journal of Economic Dynamics and Control*, 39(Supplement C):113 – 126.

Giusto, A. (2014). Adaptive learning and distributional dynamics in an incomplete markets model. *Journal of Economic Dynamics and Control*, 40(C):317–333.

Gomes, F., Kotlikoff, L., and Viceira, L. (2007). The excess burden of government indecision. NBER Working Papers 12859, National Bureau of Economic Research, Inc.

Kitao, S. (2018). Policy uncertainty and cost of delaying reform: The case of aging japan. *Review of Economic Dynamics*, 27:81 – 100.

Krusell, P. and Smith, Jr., A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896.

Leeper, E. M. (1991). Equilibria under 'active' and 'passive' monetary and fiscal policies. *Journal of Monetary Economics*, 27(1):129 – 147.

Luttmer, E. F. P. and Samwick, A. A. (2018). The welfare cost of perceived policy uncertainty: Evidence from social security. *American Economic Review*, 108(2):275–307.

Mitra, K., Evans, G. W., and Honkapohja, S. (2013). Policy change and learning in the rbc model. *Journal of Economic Dynamics and Control*, 37(10):1947 – 1971.

Mitra, K., Evans, G. W., and Honkapohja, S. (2019). Fiscal policy multipliers in an rbc model with learning. *Macroeconomic Dynamics*, 23(1):240–283.

Muth, J. F. (1960). Optimal properties of exponentially weighted forecasts. *Journal of the American Statistical Association*, 55(290):299–306.

Nelson, J. (2017). Welfare implications of uncertain social security reform. working paper.

Nelson, J. and Phillips, K. (2019). Macroeconomic effects of reducing oasi benefits: A comparison of seven overlapping-generations models. *NATIONAL TAX JOURNAL*, 72(4):671 – 692.

SSA (2019). The 2019 annual report of the board of trustees of the federal old-age and survivors insurance and federal disability insurance trust funds. Technical report, Social Security Administration.