Working Paper Series
Congressional Budget Office
Washington, DC

BEHAVIORAL EFFECTS OF SOCIAL SECURITY REFORM IN A DYNAMIC MICRO-SIMULATION WITH LIFE-CYCLE AGENTS

Amy Rehder Harris (e-mail: amy.harris@cbo.gov)
John Sabelhaus (e-mail: john.sabelhaus@cbo.gov)
Almudena Sevilla-Sanz (e-mail: almudena.sevilla-sanz@cbo.gov)
Congressional Budget Office
Washington, DC

May 2005
2005-6

Abstract

The Congressional Budget Office Long-Term (CBOLT) model uses dynamic micro-simulation to analyze Social Security policy. The version of CBOLT currently being used to analyze policy for the Congress incorporates micro behavioral effects insofar as agents alter their timing of initial claiming of Old Age Insurance (OAI) worker benefits when benefits change, and that has a direct impact on government outlays and a feedback on the macro economy through changes in labor supply. However, the change in benefit claim age is only one of three behavioral responses that could be considered in Social Security analysis–the other two are labor supply (before or after claim age) and saving behavior. This paper develops a structural life-cycle model in which agents make choices over all three margins. Because the structural model is developed using the same stochastic processes in CBOLT's micro-simulation, the state-dependent behavioral "rules" obtained from solving the life-cycle model can be used to determine behavior in a CBOLT baseline or reform-analysis simulation.

1. Introduction[1]

Social Security taxes and benefits play an important role in most people's financial planning, so changing Social Security rules has the potential for significantly affecting saving and labor supply behavior. Those changes in behavior could either reinforce or offset the first-order budgetary and distributional effects of the policy change, depending on the nature of the responses. The Congressional Budget Office Long-Term (CBOLT) micro-simulation model was developed to evaluate Social Security proposals using a representative, longitudinal sample of the population, which makes it well-suited for analyzing the extent to which behavioral responses might alter conclusions about policy effects.[2] This paper describes a new set of modules being developed for potential use in the CBOLT project: simulating labor supply and saving behavior using a structural life-cycle model.

The analysis here combines the existing large-scale micro-simulation approach in CBOLT with a structural model of life-cycle behavior. In CBOLT, demographic, economic, and program outcomes are simulated at the micro level for a one in one-thousand representative sample of the population (about 400,000 observations per year) which is large enough to capture interactions between population heterogeneity and complex program rules. Those micro outcomes are then aggregated, which allows coordinated analysis of distributional and macro results. The amount of detail tracked at the micro level in CBOLT was initially specified to be

---

[1]The authors would like to thank Julia Coronado who provided useful comments on an earlier version of the paper.

[2]For an overview of the CBOLT project see O'Harra, Sabelhaus, and Simpson (2004). The CBOLT micro-simulation model has been used as part of the analysis of the long-term outlook for Social Security and the Kolbe-Stenholm (HR.3821), Presidents's Commission to Strengthen Social Security (PCSSS) Plan 2, and Diamond-Orszag reform proposals. See CBO (2004-A, B, C, D).

just sufficient for simulating Social Security taxes, benefits, and (under proposed reforms) individual accounts. The required detail for projecting Social Security outcomes includes variables like micro-level births, deaths, marital status transitions, assignment of spouses, labor supply, earnings, and benefit claim status.[3] In addition to the equations for simulating that basic detail, the CBOLT micro model also has limited behavioral response, insofar as agents react by altering initial benefit claim-age decisions when Social Security benefits change. That change in benefit claiming feeds back on the macro economy through increased labor supply.

However, the claim-age response currently in CBOLT is one part of the overall story about potential behavioral responses. Agents could potentially react to changes in Social Security along at least three dimensions. In addition to changing benefit claim age, they could also respond by changing saving or labor supply before or after initial benefit claiming. When considering how one might evaluate any given behavioral response, it might be important to simultaneously consider all of the potential responses, because they are not likely to be independent.

These criteria for introducing behavioral responses are the basis for the strategy suggested here–adopting structural life-cycle decision rules for labor supply, consumption, and benefit claiming behavior. Those decision rules are generated by solving a parsimonious life-cycle model using standard dynamic programming techniques. The key is to specify a life-cycle model that has the same basic stochastic structure as in CBOLT and enough detail about Social Security and other programs to capture incentives and behavioral effects. The solution to the life-cycle problem is a set of consumption, labor supply, and claiming rules in terms of a handful

[3] There are several technical papers on the CBO website describing various CBOLT micro modules. See, in particular, Harris and Sabelhaus (2003), O'Harra and Sabelhaus (2002), and Perese (2002).

of state variables that will show up in both models–age, wealth, health status, idiosyncratic earnings differentials, and average lifetime earnings.

The derived life-cycle decision rules then become inputs to CBOLT simulations. The current-law solution to the dynamic programming problem is the input for baseline micro saving and labor supply projections, but the approach can also be used to directly evaluate reforms by re-solving the dynamic program under alternative policies and using the new decision rules as inputs to the CBOLT reform simulation. From a computational perspective, it is critical that this approach does not require forward-looking calculations by agents in the actual course of a CBOLT simulation.[4]

The specific goals of this paper are to characterize the existing CBOLT claiming responses and to provide a progress report on the overall behavioral-response agenda. The next section describes the CBOLT micro model and shows how the simple interconnected benefit claim-age and labor supply responses in the current model together affect conclusions about policy effects. The policy experiment shown is a fairly dramatic and immediate across-the-board ten percent cut in benefits, which comes close to eliminating the present-value funding gap in Social Security over a 75 year horizon. Agents in CBOLT are assumed to delay claiming benefits when the policy change is enacted, and that feeds back on both system finances and the overall economy. The behavioral effect of the policy change is both higher growth and tax receipts, though the effects are best described as modest, and timing issues are important.

The third section presents the structural life-cycle model of consumption, benefit claiming, and labor supply behavior being developed for use in CBOLT. The model is specified

[4]A standard CBOLT 100 year simulation runs in about 20 minutes on a high-end desktop computer, which makes (overnight) Monte Carlo simulation feasible.

using the same stochastic earnings process in CBOLT, and incorporates the crucial determinants of real-world consumption and labor supply outcomes which previous research has focused on: uncertain lifetimes, a realistic Social Security system, liquidity constraints, consumption floor, uncertainty about health status and medical expenditures, and a utility function that includes consumption, leisure, and health status. The life-cycle model is solved and the behavioral rules are used in micro-simulation of both cross-section and longitudinal outcomes.

The version of the life-cycle model used for the prototype baseline and reform simulations in the last section includes endogenous consumption and benefit claiming behavior, but with labor supply fixed at working full time before benefit claiming and not working thereafter. Even though the model results shown here do not incorporate the full range of potential labor supply responses, the consumption and wealth accumulation patterns generated in the micro-simulation are consistent with what is known about actual behavior from various micro data sources. In particular, the model generates realistic income, consumption, and wealth outcomes across two important dimensions: trajectories by age within lifetime income groups, and cross-section snapshots of a representative population. The effects of the 10 percent benefit cut experiment are a little mixed, and sensitive to what level of consumption floor is specified.

2. Reduced-Form Labor Supply Effects in the CBOLT Micro Model

Micro-simulation is a powerful tool for analyzing public policy issues when the interaction of complex program rules and population heterogeneity is likely to be of first-order importance. In addition to capturing the aggregate and distributional effects of policy rules when micro behavior is exogenous, it is possible to consider how incentives built into the rules might affect individual labor supply and saving behavior. There are two important criteria that arise when thinking about how best to introduce behavioral responses: although the primary goal is to develop realistic responses, the complexity of the response mechanism has to be weighed against computational constraints and model manageability. The current version of CBOLT reflects the existing resolution of that tradeoff. The model has most of the micro/macro linkages in place that would allow a comprehensive analysis of behavioral responses, but the margins of response are currently limited to retired worker benefit claiming probabilities and the impact of beneficiary status on labor supply.

The first principles pursued in the development of the micro-simulation were generating realistic demographic, economic, and policy outcomes for a large representative sample, and then applying the complex Social Security program rules to determine budgetary and distributional outcomes. The CBOLT micro-simulation operates on the basic processes (birth, education, labor supply, earnings, first marriage, divorce, remarriage, mate matching, benefit claiming, benefit awards, and ultimately death) needed to calculate Social Security taxes and benefits, and integrates the micro outcomes with a macro growth model and unified budget framework. Most of the micro processes in CBOLT were kept parsimonious in the initial specifications in order to get a working version of the model and thus distill the first-order impact of using micro-simulation as a basis for Social Security baseline and reform analysis.[5]

Even though the claiming behavior in CBOLT's micro-simulation is limited, the architecture is largely in place for analyzing the aggregate and distributional effects of Social Security. For example, the macro growth model framework in CBOLT employs a standard Cobb-Douglas production technology, where aggregate labor input is the sum of hours worked in the micro model.[6] Thus, any policy-induced changes in micro labor supply will appropriately feed back on aggregate output and system finances.

There are a number of ways one might introduce a reduced-form approach to simulating OAI worker benefit claiming behavior. For example, Coile and Gruber (2003) consider a forward-looking model in which agents, as of the earliest eligibility age, compare the benefit level that would result by claiming at any given age to the benefit if they choose to claim at the optimal or "peak-value" age. Using this approach with the Health and Retirement Study, their analysis captures some important aspects of claim-age responses that are also reflected in the CBOLT model. For example, because actuarial adjustments applied to benefits neutralize most of the impact of claiming behavior changes on outlays, the first-order budgetary impact of delays in benefit claiming are fairly modest. However, the same observation suggests a more important

---

[5]Even the limited-behavior version of the micro-simulation based-approach leads to important insights about Social Security that do not come through in other analyses. First, all else equal, the micro-simulation generates projected benefit awards for male OAI workers below those based on standard actuarial techniques, because CBOLT properly captures observed shifts in the historical relative earnings profiles (CBO, 2004-A). Second, direct analysis of the micro-level outcomes suggests there are serious problems with using hypothetical "example" workers to analyze the impact of proposed reforms (CBO, 2004-B, C, D).

[6]In general the macro growth model framework is closely related to the approach in Bosworth and Burtless (2002, 2004) with two important differences: aggregate labor input is summed from the micro model in CBOLT, and aggregate private saving adjusts to target a stable long-run capital output ratio. That saving assumption effectively neutralizes the impact of assumptions about other components of the Federal budget and creates a stable baseline similar to the one used by the Social Security Trustees when they analyze system finances. One of the primary goals of the life-cycle approach described in the next section is to replace the simple aggregate private saving rule by summing over individual saving, as is currently the case for labor supply.

route by which policy might generate feedback effects: if agents delay claiming when benefits change, they are likely to keep working, raising production and tax receipts.

The CBOLT micro claiming and labor supply modules capture these basic principles as well. Benefit claiming rates vary by cohort and sex based on differences in average benefit replacement rates–the ratio of average benefits to the after-tax average wage. If replacement rates fall, agents delay claiming to offset some of the policy change through the actuarial adjustments built into the system. The claiming probability model is a simple logistic formulation which naturally constrains claiming probabilities to the zero/one range. In addition to the benefit replacement rate determining claiming, there are also exogenous adjustments at three crucial ages: the Early Eligibility Age (EEA), the Normal Retirement Age (NRA), and the Medicare eligibility age. These adjustments capture the actual behavior of claimants that a simple model which relates claiming to benefit levels by age cannot.[7] The CBOLT labor force participation equation and hours worked decision include beneficiary status as an independent variable, so the delay in claiming has the expected positive impact on labor supply.[8]

The simple claiming behavior and subsequent labor supply responses in CBOLT are intended to capture expectations about future retiree behavior under current law and to generate reasonable responses to various types of benefit changes. For example, scheduled increases in the Social Security NRA are expected to indirectly induce some claiming delay because of the actuarial benefit reductions at ages below the NRA, and directly delay claiming because of the exogenous NRA effect. Interestingly, the baseline CBOLT projections still suggest some

---

[7]This apparent exogeneity of benefit claiming at particular ages arises in the context of specifying the life-cycle model in the next section.

[8]See Harris and Sabelhaus (2003) for further details.

bunching of claiming at the old NRA (age 65) because the Medicare eligibility age (65) is not scheduled to change. In general, the extent of response is controlled by (ad hoc) parameters chosen to generate outcomes suggested by previous research and, to some extent, intuition.[9]

Figure 1 and Table 1 show the result of a simple experiment designed to give a sense of how different benefit claiming responses affect policy conclusions in CBOLT. The experiment is an immediate and permanent 10% reduction in benefits for all new beneficiaries, implemented by reducing the 90%, 32%, and 15% benefit formula replacement factors to 81%, 28.8%, and 13.5%. Figure 1 shows that, depending on the specified value for the claim response, the impact on claiming probabilities could be quite large, especially at the earliest claim ages. That effect is mitigated at older ages because the benefit claiming module adjusts as the pool of potential claimants changes, preserving the asymptotic property that everyone eligible claims by age 70 (there is no reason not to in the U.S.).

Using different assumptions about claiming responses affects conclusions about the impact on measures of interest like beneficiary counts, total benefits, labor supply, and GDP (Table 1). The years 2014 and 2050 are chosen as reference points to highlight an important aspect of the effects; in the short run, simulating delays in claiming while ignoring the phase-in effect could lead one to overestimate the budgetary impact. The "no-behavior" effect of the policy change is a roughly 5.5% drop in both average and total OAI worker benefits by 2014, because roughly half of benefits are still being paid to people who claimed before the experiment was initialized in 2004. However, with the baseline claiming response turned on, the count of

[9] The parameters controlling responses are the coefficient on benefit replacement rates and the exogenous ERA, NRA, and Medicare effects. Given these values, age-specific intercepts are adjusted so the model's predicted claiming behavior matches the actual claiming behavior of the most recent cohort of beneficiaries in the U.S.

beneficiaries is nearly 3% lower by 2014, and total worker benefits have fallen 7.2%. The short-run effects are even larger with the higher claiming response.

In the long run, the claiming effect on total outlays is largely washed out because those beneficiaries who delay claiming receive significantly higher average benefits through the actuarial adjustments. Indeed, the reduction in long-run outlays is fairly insensitive to how claiming is modeled, measuring very close to the ten percent reduction specified in the experiment across all three levels of responses.[10] However, even these simple claiming responses do show a potential for affecting outcomes through the impact on labor supply and GDP. As claiming is delayed, total hours worked rises about one percent in the long run, which seems significant given the size of the elderly labor force.

Although this limited claiming response reflects a distinct improvement over assuming no behavior, it leaves much to be desired. As a matter of principle, a more realistic micro-simulation could project individual saving and wealth accumulation in both baseline and policy reform simulations. In addition, since it is possible that agents will respond to benefit changes by working more when young, the micro-simulation could include labor supply effects before claim age and the corresponding interactions with individual saving. This result leads naturally to the life-cycle approach for modeling behavioral responses.

[10]One oft-used measure of the budgetary effects–the change in the 75 year summary actuarial balance–is effectively the same under the various claiming specifications ranging from 1.13 to 1.15 percent of taxable payroll.

3. A Life-Cycle Model of Consumption, Labor Supply, and Benefit Claiming

This section presents a structural life-cycle model that captures important aspects of saving and labor supply behavior in a way that is consistent with simulating outcomes for a large representative micro sample, as in CBOLT. Previous research has shown that stochastic versions of the life-cycle model are capable of explaining consumption behavior and wealth distribution, benefit claiming behavior, and labor supply across full and part time opportunities at different points in the life-cycle. The goal in specifying the model below is to capture these desired properties while keeping the dynamic programming problem parsimonious and consistent with CBOLT micro-processes, so that establishing a link between CBOLT simulations and the life-cycle solution is feasible.[11]

The first principle built into the specification below is forward-looking consumption and saving behavior, using the buffer-stock formulation (Zeldes (1989), Deaton (1991), Carroll (1992, 1997)) as a starting point. The buffer-stock model can generally be described as the solution to a life-cycle consumption problem where labor supply is exogenous and individual earnings evolve over time because of shocks to transitory and permanent stochastic components. The buffer-stock model has proved useful in explaining why (for example) young agents don't save as much as certainty-equivalent models suggest but don't demand extensive borrowing

[11]The approach here is not a general equilibrium solution to the life-cycle problem, as in papers like Castenada, Diaz-Gimenez, and Rios-Rull (2002), Huggett (1996), Huggett and Ventura (1999, 2000), and Nishiyama (2004). One practical limiting factor is computational: the CBOLT micro sample is quite large, averaging about 400,000 active observations per year in a standard 100 year projection, so solving forward with the traditional consistent expectations approach assumed in general equilibrium models is infeasible. Indeed, almost any conceivable algorithm that involves looping over agents' expectations is computationally infeasible, which is the basis for the recursive approach suggested here. Looking ahead, one can imagine how the approach suggested here could someday be modified for expectations: one could work with a set of consumption and labor supply decision rules that vary by policy and expected factor prices, and then allow changes in actual factor prices to feed back on expectations over time.

either. Basically, the effect of expected income growth dominates patience early in the life-cycle, but uncertainty about future income precludes imprudent borrowing. This general class of forward-looking consumption model has also been used to explain the accumulation of wealth across the income distribution. Particularly, Hubbard, Skinner, and Zeldes (1994, 1995) use a forward-looking model with asset-based, means-tested social insurance program to explain why low-income households are more likely than high income households to hold little or no wealth.

The second important principle built into the specification below is focused on explaining life-cycle labor supply and benefit claiming behavior. One aspect of behavior not completely explained by a simple model is the bunching of benefit claiming (retirement from full-time jobs) at particular ages, notably 62 and 65, which are the early (ERA) and normal (NRA) retirement ages under Social Security. This bunching is somewhat of an anomaly because the U.S. Social Security system is close to actuarially fair; that is, delaying claiming by one year at age 62 does not cause a big change in the present value of lifetime net transfers. Therefore, one answer to the puzzle is that retirement and claiming are based on some sort of "reference point" behavior. However, Rust and Phelan (1997) propose a solution to this anomaly that involves idiosyncratic health shocks and the availability of Medicare (which begins at age 65) and employer-provided health insurance.[12] Gustman and Steinmeier (1986, 2002, 2003, 2004) also develop an explanation for differences in claiming behavior, where the key determinants include (for example) the availability of part-time work after retirement and heterogeneity in the rate of time preference across agents.

In these various approaches to predicting labor supply and benefit claiming in a life-cycle

[12] This approach is further developed, with an application to the Disability Insurance (DI) program, in Rust, Buchinsky, and Benitez-Silva (2001, 2003).

context, it is important to consider leisure and health status along with consumption in the utility formulation. For example, the value of leisure is often assumed to rise as agents age, which makes them gradually more likely to retire and start claiming benefits. The actual dates of claiming are then determined by the interaction (or tension) between program rules (ERA and NRA, Medicare eligibility age) and preferences. Thus, although behavior may seem somewhat exogenous when the independent variable used to predict claiming is just the (actuarial) value of benefits at various ages, a comprehensive life-cycle formulation is able to explain a lot of the variation in behavior.[13]

These insights from the literature on life-cycle consumption and labor supply are reflected in the specification below. The following notation is used throughout this section and in the Appendix where the solution and simulation strategy is described in more detail:

U_t	expected present value of utility in period t
c_t	consumption in period t
c^{min}	consumption floor, guaranteed by government transfers
l_t	normalized leisure in period t
h_t	health status in period t
R	age at which Social Security (old-age) benefit claiming begins
T	maximum age
β	single period discount factor
$\pi_{t+s/t}$	probability of surviving to period t+s given survival through period t
a_t	wealth, beginning of period t
r	single period (certain) net return on wealth
e_t	earnings in period t
${\tau^{ss}}_t$	Social Security taxes paid in period t
τ^{ss}	Social Security tax rate
e^{max}	Social Security taxable maximum
b_t	Social Security benefit received in period t

[13]Improvements in computing power and more wide-spread use of life-cycle modeling have led several authors to combine insights and make significant advances that build on these pioneering works. These types of papers focus on the joint consumption and labor supply decision, as in the model here. Notable among these are Woolley (2004), Blau (2004), French et al. (2003), and Van der Klauw and Wolpin (2003).

m_t out of pocket medical expenditures in period t
μ_t medical expenditures shock (deviation from expected) in period t
σ_μ standard deviation of medical expenditure shocks
$\bar{e}_t$ average indexed taxable earnings through period t
a^h_t cash on hand in period t; sum of beginning wealth and realized net cash inflows
λ_t exogenous (age-related) component of earnings in period t
δ_t idiosyncratic permanent earnings differential at time t
σ_η standard deviation of shocks to idiosyncratic permanent earnings differential
ϵ_t transitory earnings shock at time t
σ_ϵ standard deviation of transitory earning shocks
α utility function parameter, determines intratemporal substitution
γ utility function parameter, determines intertemporal elasticity of substitution
Γ_t age-specific utility function parameter, affects disutility of working by age

The life-cycle problem facing agents is to choose a sequence of consumption and leisure values

($\tilde{c},\tilde{l}$) and benefit claiming age (R) to maximize expected utility,

$$\max_{\tilde{c},\tilde{l},R} U_t = \sum_{s=0}^{T-t} \beta^{-s} U(c_{t+s,} l_{t+s}, h_{t+s}) \pi_{t+s/t}$$

The evolution of beginning of period wealth is given by,

$$a_{t+1} = a_t (1 + r) + e_t - \tau^{ss}_t + b_t - m_t - c_t$$

and there are two other constraints; $a_t \geq 0$ and $c_t \geq c^{min}\ \forall$ t. Those conditions are guaranteed by assuming government transfers that cover the gap between cash on hand ($a^h_t = a_t (1 + r) + e_t - \tau^{ss}_t + b_t - m_t$) and the consumption floor (c^{min}).

Potential earnings (full-time equivalent) is modeled in logs as the sum of the fixed (age-specific) component (λ_t), the idiosyncratic permanent differential (δ_t), and the transitory shock

(ϵ_t). Potential earnings are adjusted for leisure to determine actual earnings. That is,

$$e_t = \exp[\lambda_t + \delta_t + \epsilon_t]\, g(l_t)$$

where $g(l_t)$ represents normalized labor supply as a function of normalized leisure. In particular, normalized leisure in the model takes on the value of 1 for no work, 0.6 for full-time, and 0.8 for part-time. Then, the mapping from leisure to labor supply is g(1)=0 (no work), g(0.6)=1 (full-time), and g(0.8) = .425 (part-time).[14]

Both the permanent and transitory components of earnings are stochastic. The permanent differential evolves over time according to $\delta_t = \delta_{t-1} + \eta_t$, and both shocks ($\eta_t$ and ϵ_t) are assumed normally distributed with mean zero and fixed standard deviations ($\sigma_\eta, \sigma_\epsilon$). Given outcomes for earnings, the value of average indexed taxable earnings through age t ($\overline{e}_t$) is a weighted average of last period's average ($\overline{e}_{t-1}$) and current earnings (e_t).[15]

Social Security taxes are computed by applying the tax rate to actual earnings but only up to the taxable maximum. That is, $\tau^{ss}_t = \tau^{ss} \min(e_t, e^{max})$. Social Security benefits are somewhat more complicated. If age (t) is greater than claim age (R) then benefits are computed initially as a function of average taxable earnings through claim age minus one and claim age itself. That

[14] This specification for stochastic earnings is identical to the one used in the CBOLT micro-simulation model. The value for part-time work corresponds to 50% hours worked, but is adjusted down from g(0.8)=.5 to g(0.8)=.425 because CBOLT uses a 15% adjustment to salaries for part-time workers.

[15] In the model, the age-specific component of earnings (λ_t)includes an economy-wide average wage index (AWI) variable (currently normalized to grow at 1% per year, but adjustable) that factors in when computing average indexed earnings. Without indexing, average taxable earnings evolve using $\overline{e}_t = \overline{e}_{t-1}((t-1)/t) + \min(e_t, e^{max})(1/t)$. With indexing, the formula is $\overline{e}_t = \overline{e}_{t-1}(AWI_t/AWI_{t-1})((t-1)/t) + \min(e_t, e^{max})(1/t)$

is, $b_t = b(\bar{e}_{R-1}, R, e_t, t)$.[16] The basic benefit (Primary Insurance Amount, or PIA) is a piece-wise linear function of average taxable earnings (Average Indexed Monthly Earnings, or AIME, denoted by $\bar{e}_{R-1}$), with higher replacement factors at low average earnings. Claim age plays a role in benefit determination because the basic benefit (PIA) is also actuarially adjusted; agents get exactly the PIA at Normal Retirement Age (NRA), but less if they start claiming earlier, and more if they start claiming later. The actual Social Security benefit paid is still potentially affected by current earnings, however, because some beneficiaries with current earnings above the "earnings test" are subject to having benefits reduced during that time period.[17]

Health status evolves using a simple, age-dependent Markov transition process denoted $h(t, h_{t-1})$, and out of pocket medical spending depends on health status and a stochastic element, $m_t = m(h_t) + \mu_t$, where the error term has mean zero and standard deviation σ_μ. In the current model, there are two health states, good and bad (although, given stochastic mortality, there is implicitly a third health state, which is dead), with values of h=1 in the good health state and h=.75 in the bad health state. The process determining the health status transition is very simple: agents are more likely to stay healthy when young, and not very likely to regain good health

[16]The average earnings measure used in the AIME formula is actually only the average of the highest 35 years of earning. Because it is infeasible to track that many state variables, this feature of the actual rules is simplified in the dynamic programming algorithm.

[17]Individuals with earnings in years prior to their NRA face a reduction in monthly benefits of $1 for every $2 of monthly earnings above the monthly earnings test amount. Individuals with earnings in the year of attaining the NRA face a reduction in monthly benefits of $1 for every $3 of monthly earnings above the monthly earnings test amount; however, not all earnings are considered in applying this test, only the share of earnings that are attributed to the months prior to attaining the NRA (someone with an NRA of 65 years and 6 months would only face the earnings test on half of their annual earnings). For 2004, the threshold for those below the NRA is $11,640, for those in the year of the NRA, $31,080.

once in the bad health state.[18] Out of pocket medical spending is also very simple in the current model. Total outlays are a stochastic variable, with different means and variances for the good and bad health states.[19] The only institutional detail in the model is the effect of Medicare: at age 65, out of pocket health spending is always zero, because Medicare is assumed to cover all bills.

The period utility function is a standard CRRA intertemporal formulation with Cobb-Douglas substitution between consumption and leisure. In particular,

$$U(c_t, l_t, h_t) = \frac{(c_t^{\alpha}(l_t h_t \Gamma_t)^{1-\alpha})^{1-\gamma}}{1-\gamma}$$

In this formulation age and health status operate directly on the value of leisure to affect overall utility. When h_t is less than one (as in the poor health state) it is as though the agent simply has less normalized leisure; so, for example, a person in good health (h_t=1) working full-time (l_t=.6) has roughly the same utility as a person in poor health (h_t=.75) not working (l_t=1).

The dynamic programming approach to solving this life-cycle problem in terms of state variables involves working backward from the last period of life, recursively solving the series of two-period utility maximization problems (Bellman's equation) at each step. The model above has, in principle, five state variables in terms of which to solve the problem: two are discrete (age and health status) while the other three are continuous (wealth, average earnings through the current age, and the permanent earnings differential). However, the actual solution

[18] Everyone starts at age 21 in good health. There is a 2% chance each year of moving to the bad health state until age 51, at which point there is a 5% chance of moving to the bad health state. There is always a 1% chance of moving back from the bad health state to the good health state.

[19] The mean expenditures for good health agents is $1,000, with a $500 standard deviation. The mean expenditures for poor health agents is $5,000, with a $2,500 standard deviation.

developed for the problem above (described in the Appendix) only seems tractable when the state space is expanded to explicitly reflect the fact that labor supply decisions are conditional on earnings shocks and expected benefit claim age, and consumption decisions are in turn conditional on labor supply. The proposed solution is feasible, but currently not computationally cost-effective.

Given the computational problems currently associated with solving the full model, the discussion below focuses on a version of the model in which the labor supply choices are restricted to working full-time before benefit claiming, and not working once claiming has begun. This specification effectively makes the problem one of choosing consumption at each age and choosing the value for the benefit claim age. The simplest approach to solving this problem (see Rust, Buchinsky, and Benitez-Silva (2001, 2003)) is to make benefit claim age (R) an additional discrete state variable along with age and health status (the continuous states are still wealth, average earnings through the current age, and the permanent earnings differential). The dynamic program is then solved recursively for consumption in terms of all six states, and the choice of claim age is actually made only in a simulation context. At each point in the simulation, agents (conditional on the other five states) choose the claim age with the highest expected utility, and then choose the value of consumption consistent with all six states.

4. Life-Cycle Behavior in a Dynamic Micro-Simulation

The structural life-cycle model developed in the last section is designed to be used in a two-step procedure for analyzing policy. The first step is solving the structural model using standard dynamic programming techniques, thereby generating decision rules in terms of state variables. The second step is using those derived decision rules in the context of a dynamic micro-simulation, integrating the behavioral predictions with the other stochastic processes operating on agents at the micro level. This section describes the results of prototype simulations for baseline and policy reform applications of this two step procedure.

The first step is to show that the specified model meets the basic goal of generating realistic baseline outcomes. In order to keep the experiments initially manageable, the simulations here use a limited subset of the CBOLT micro-simulation machinery. Unisex agents are born, begin working at age 21, work full-time until benefit claim age, and then fully retire. They face mortality hazards derived from Social Security Administration actuarial projections, as in the main CBOLT model. There is no sex, marriage, education, or other demographic detail, so the age-earnings equation in the simulator is a collapsed version from the main CBOLT model that averages coefficients over demographic groups. The model is also simulated using only the "going-forward" (not historical) version of CBOLT–that is, all agents are simulated from birth.

There are two distinct ways to simulate the life-cycle model, and both are used here. The first is to focus on a single cohort, and track outcomes over their life histories. In the results presented below (labeled "longitudinal") that are based on a single cohort, the sample size is always set to 100,000. The second approach is to generate a representative sample for some point in time. In these simulations, a certain number of agents per cohort (for all sample-year ages 21 to 100) are simulated, where the exact count is determined by the population growth rate. In the results presented below (labeled "cross-section") the oldest cohort sample size is set to 2,000, and with the population growth rate at 1% per year, the youngest cohort (age 21 in the sample year) has about 4,000 observations. Given simulated deaths, the cross-section sample year observation count is about 185,000 agents, the same order of magnitude as in CBOLT.

In both the longitudinal and cross-section simulations, there is a key feature associated with real-world timing: the earnings process in the life-cycle model (adopted from CBOLT) generates relative earnings at each point in time, and the actual level of earnings is determined by the interaction of the overall average wage index and the idiosyncratic relative earnings component. The level of real wage growth in the simulations is set to 1%, and all of the Social Security parameters are adjusted (in real terms) for the change in average wage growth (as in the real world). Thus, the longitudinal earnings trajectories, cross-section age-earnings profiles, and benefit outcomes are consistent with an economy that has steady-state real wage growth (independent of population weighting effects) equal to 1%.

The baseline and reform simulations generally share a common set of parameter setting and specifications for stochastic processes, summarized in Table 2. The first set of permutations shown in the various tables involve changing the minimum consumption floor maintained by government transfers (c^{min}) with values of $1, $5,000, and $10,000. The second set of permutations in the baseline involve turning the medical expenditure uncertainty on and off. These two sets of assumptions are shown to play an important role in how well the model replicates certain real-world observations.

The first set of results to focus on is from the cross-section simulations, shown in Table

3. The table compares simulated percentiles of the cross-section wealth distribution with values from the 2001 Survey of Consumer Finances (SCF). In both the actual and simulated data sets, the wealth percentiles are divided by overall average income to make the two comparable. The SCF values are shown for two measures: total net worth, and net worth excluding housing net worth. The simulated data are shown for three values of the consumption floor with and without the medical expenditure shocks turned on.

Two conclusions can be drawn from Table 3. First, the structural model is able to replicate several important features of the actual wealth distribution. Second, conclusion number one is very sensitive to the specification of consumption floor and uncertainty about medical expenditures. When medical expense shocks are turned off the gradual (and ultimately complete) resolution of income uncertainty leads to very low saving. With medical expense shocks turned on, if the consumption floor is set too low the model generates a lot more normalized wealth than is observed in the actual data. Although it is not clear exactly which SCF wealth concept should be used for the comparison, the model with a $5,000 consumption does a good job of replicating the wealth distribution all the way up to the highest percentiles. It is particularly noteworthy that the model generates negligible wealth for a large share of the population, which is clearly consistent with the data.

The next set of results focus on longitudinal outcomes for two stylized agents, shown in Figures 2 through 5. The first stylized agent earns exactly the average (age-specific) earnings, claims benefits at age 62, and remains in good health through age 100. The second stylized agent is identical, but earns exactly half the average at each age. Figures 2 and 3 show consumption and income trajectories, and Figures 4 and 5 focus on wealth to income ratios by age.

The consumption and income trajectories in Figures 2 and 3 are based on consumption floors of $1 and $5,000 respectively. Even with a $1 consumption floor in Figure 2, the stylized agents reveal an important feature of the model which is consistent with real-world behavior–consumption tracks income for most of a person's work life and only begins to adjust for life-cycle smoothing reasons as they approach retirement. That effect is even stronger when the consumption floor is raised to $5,000. Indeed, comparing Figures 4 and 5 makes this distinction even clearer: the low earning agent saves very little when the floor is $1, but nothing at all (until right before retirement) when the floor is $5,000. This occurs even though the agent is earning well above the actual consumption floor of $5,000–the principle that even being near the floor has an impact on behavior comes through very clearly.

Results for the stylized agents show that the model is also able to capture other important features of behavior, though there are also areas where model refinement is suggested. One interesting observation is the drop in consumption that occurs at benefit claiming age, 62. This type of drop is also observed in actual data, and is occurring in this model because of the non-separable consumption and leisure in the utility function. Note that the drop is more noticeable for the average income agent who is in a standard life-cycle situation (well above the consumption floor) and saving for retirement. The figures also show at least one area where the model can use some refinement–the rate of wealth depletion after retirement (Figures 4 and 5) is not consistent with actual patterns of behavior.

The last set of baseline results to focus on is also longitudinal but based on the entire simulated sample. Table 4 shows the ratios of wealth to average lifetime income as of age 62 for every agent in the simulation who lives to age 62. The first column shows the overall average for the cohort, and the other columns split agents into lifetime earnings quintiles based on their average earnings from age 21 to age 61–these are the lifetime rich versus lifetime poor. (Note that this table is designed to ultimately be compared to a data set like the Health and Retirement Study with its longitudinal earnings histories). The differences in results across medical expense and consumption floor assumptions is consistent with the findings in Table 3, and the reasons are the same. Table 4 really drives this point home: given realistic earnings shocks, the structural model still suggests that if there is a meaningful consumption floor then a large number of agents will enter retirement with little or no accumulated wealth.

Although benefit claiming behavior is turned on in the model results shown here, the effect is still mixed. About 10% of agents start claiming at age 62 (fewer than in the base case) and it is an interesting u-shaped mix of people whose (actuarially reduced) Social Security benefits are near or below the consumption floor and the very wealthy for whom the size of the Social Security benefit is irrelevant. These results are really only suggestive, however, because the consumption/leisure tradeoff parameters in the utility function need to be adjusted.

The model results are very rough at this stage, but it is feasible to implement a policy change and look at the results. Table 5 shows the impact of a ten percent benefit reduction on the longitudinal wealth accumulation results. As expected, the consumption floor effect dominates for the lowest earning agents, and the life-cycle effect dominates for the high earners. Although significantly more analysis is in order before strong conclusions are drawn, it is clear that the interaction across government programs may be crucial for understanding the saving effects of Social Security reform.

5. Extensions

The structural life-cycle model developed here offers significant promise for accomplishing the goal of introducing additional behavior into the CBOLT dynamic micro-simulation. Several extensions should be implemented to reach that goal. Some of those are fairly modest changes, and can be implemented with little effect on computational burden. Others are more fundamental and will require deciding whether to work around those issues using post-model adjustments or try to incorporate the extensions directly into the model.

The first possible step is to adopt a new algorithm for solving the dynamic program that will work efficiently on the full model, and thus enable completely endogenous labor supply, before and after retirement. The second area for improvement is to adopt more realistic health states, health transitions, and health expenditures, and make mortality a function of health status. Also, the model is also basically set up to include disability insurance (DI) beneficiary status as a state variable, because the discrete claim age state (R) can be extended to include one more dimension, as in Rust, et. al (2001, 2003). Adding DI status as a choice variable will involve an application and acceptance module, but that can be randomized. It is also straightforward (and important) to add an income tax system to the model, because all of the relevant income variables are tracked in levels, and it is just a question of parameterizing the (basic) income tax code. Finally, introducing bequest motives and bequest receipts are fairly modest changes.

The more elaborate extensions to the life-cycle model that one would like to ultimately incorporate include adding features like heterogeneous preferences, family structure, different types of pension coverage, and portfolio choices. The problem with these types of innovations is that they involve (sometimes significant) expansion of the state space–for example, allowing heterogeneity in rates of time preference involves solving the model for each possible value of time preference to consider, and that effectively makes the rate of time preference a (continuos) state variable. Some of these effects are too important to ignore, though, and a good deal of research is warranted to decide exactly how to incorporate those . One possibility is to solve the life-cycle model using a limited set of states (as in this paper) but then adjust predicted behavior systematically using characteristics in the micro-simulation that are not in the state space. For example, the model might be solved generally for couples and singles, but then predicted consumption is adjusted ex post based on number of children to reflect the fact that people with children spend more (all else constant) than people without children.

References

Blau, David M. (2004). "Retirement and Consumption in a Life-Cycle Model," Mimeo. University of North Carolina at Chapel Hill.

Bosworth, Barry, and Gary Burtless (2004). "Supply-Side Consequences of Social Security Reform: Impacts on Saving and Employment," CRR Working paper 2004-1. Center for Retirement Research at Boston College (January).

Bosworth, Barry, and Gary Burtless (2002). "The Effects of Social Security Reform on Saving, Investment, and the Level and Distribution of Economic Well-Being," CRR Working paper 2000-2. Center for Retirement Research at Boston College (January).

Coile, Courtney, and Jonathan Gruber (2003). "Fiscal Effects of Social Security Reform in the United States," CRR Working paper 2003-5. Center for Retirement Research at Boston College (April).

Congressional Budget Office (2004-A). *The Outlook for Social Security*. Washington, D.C. (June).

Congressional Budget Office (2004-B). *Long-Term Analysis of H.R. 3821, the Bipartisan Retirement Security Act of 2004*. Washington, D.C. (July).

Congressional Budget Office (2004-C). *Long-Term Analysis of Plan 2 of the President's Commission to Strengthen Social Security*. Washington, D.C. (July).

Congressional Budget Office (2004-D). *Long-Term Analysis of Diamond-Orszag Social Security Plan* . Washington, D.C. (December).

Carroll, Christopher D. (1997). "Buffer Stock Saving and the Life-Cycle/Permanent Income Hypothesis," *The Quarterly Journal of Economics,* CXII(1), 1-56.

Carroll, Christopher D. (1992). "The Buffer Stock Theory of Saving: Some Macroeconomic Evidence." *Brookings Papers on Economic Activity,* 2: 61-135.

Castenada, Ana, Diaz-Gimenez, Javier, and Rios-Rull, Jose-Victor (2002) "Accounting for the U.S. Earnings and Wealth Inequality." Mimeo, Universidad Carlos III de Madrid and University of Pennsylvania.

Deaton, Angus (1991) "Saving and Liquidity Constraints." *Econometrica*, 59:5 1221-1248.

French, Eric (2003). "The Effects of Helath, Wealth, and Wages on Labor Supply and Retirement Behavior", Working Paper, Federal Reserve Bank of Chicago.

Gustman, Alan L., and Thomas L. Steinmeier (2004). "The Social Security Retirement Earnings Test, Retirement and Benefit Claiming," NBER Working Paper 10905. National Bureau of Economic Research (November).

Gustman, Alan L., and Thomas L. Steinmeier (2003). "Retirement Effects of Proposals by the President's Commission to Strengthen Social Security," NBER Working Paper 10030. National Bureau of Economic Research (October).

Gustman, Alan L., and Thomas L. Steinmeier (2002). "The Social Security Early Entitlement Age in a Structural Model of Retirement and Wealth," NBER Working Paper 9183. National Bureau of Economic Research (September).

Gustman, Alan L., and Thomas L. Steinmeier (1986). "A Structural Retirement Model," *Econometrica*, 54:3 555-584.

Harris, Amy Rehder, and John Sabelhaus (2003) "Projecting Longitudinal Earnings for Long-Term Policy Analysis," Technical Paper 2003-2. Congressional Budget Office (April).

Hubbard, R. Glenn, Skinner, Jonathan, and Zeldes, Stephen P. (1995) "Precautionary Saving and Social Insurance." *Journal of Political Economy*, 103: 360-399.

Hubbard, R. Glenn, Skinner, Jonathan, and Zeldes, Stephen P. (1994) "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving." *Carnegie-Rochester Series on Public Policy*, 40: 59-126

Huggett, Mark, and Gustavo Ventura (2000) "Understanding Why High Income Households Save More Than Low Income Households." *Journal of Monetary Economics*, 45:2, 361-97 (April).

Huggett, Mark, and Gustavo Ventura (1999) "On the Distributional Effects of Social Security Reform." *Review of Economic Dynamics*, 2:498-531.

Huggett, Mark (1996) "Wealth Distribution in Life-Cycle Economies." *Journal of Monetary Economics*, 38, 469-494.

Nishiyama, Shinishi (2004) "Analyzing and Aging Population–A Dynamic General Equilibrium Approach," Technical Paper 2004-3. Congressional Budget Office (February).

O'Harra, Josh, and John Sabelhaus (2002) "Projecting Longitudinal Marriage Patterns for Long-Term Policy Analysis," Technical Paper 2002-2. Congressional Budget Office (October).

O'Harra, Josh, John Sabelhaus, and Michael Simpson (2004) "Overview of the Congressional Budget Office Long-Term (CBOLT) Policy Simulation Model," Technical Paper 2004-1. Congressional Budget Office (January).

Perese, Kevin (2002) "Mate Matching for Micro-Simulation Models," Technical Paper 2002-3, Congressional Budget Office (November).

Rust, John, Moshe Buchinsky, and Hugo Benitez-Silva (2003). "Using a Life-Cycle Model to Predict Induce Entry Effects of a $1 for $2 Benefit Offset in the SSDI Program." Mimeo, University of Maryland (November).

Rust, John, Moshe Buchinsky, and Hugo Benitez-Silva (2001). "An Empirical Model of Social Insurance at the End of the Life-Cycle." Mimeo, University of Maryland (August).

Rust, John, and Christopher Phelan (1997) "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets." *Econometrica*, 65:4 781-831.

President's Commission to Strengthen Social Security (PCSSS) (2001). "Strengthening Social Security and Creating Wealth for All Americans," Report of the President's Commission, (December).

Social Security Administration, Office of the Chief Actuary (2002). "Estimates of Financial Effects for Three Models Developed by the President's Commission to Strengthen Social Security," Memorandum to Daniel Patrick Moynihan and Richard D. Parsons from Stephen C. Goss and Alice H. Wade (January 31).

Van der Klaauw, Wilbert and Kenneth I. Wolpin (2003). "Social Security, Pensions, and the Savings and Retirement Bahavior of Households", manuscript, University of North Carolina at Chapel Hill.

Woolley, Mathew (2004) "Optimal Consumption and Investment With Labor Income Uncertainty and Endogenous Retirement," Mimeo, University of North Carolina (September).

Zeldes, Stephen P. (1989) "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence." *Quarterly Journal of Economics*, 104: 275-298.

Appendix: Life-Cycle Model Solution and Simulation Strategies

This appendix describes the simulation and solution strategies for the life-cycle model developed in the text. The first section describes the algorithms for solving and simulating the version of the model used for the baseline and policy experiments in the paper, which has labor supply conditional on benefit claim age. The second section describes the algorithm developed for working with the full model.

<u>A1. Solving and Simulating the Model with Endogenous Consumption and Benefit Claiming</u>

This first section describes the algorithm is the restricted version of the full model in which labor supply is restricted to equal one before claim age, and zero thereafter. Otherwise the model is completely as described in the text. The description below is actually in reverse order; the timing of events is sketched out in terms of how the model is simulated in CBOLT, and following that is a discussion of how the model is solved.

Agents begin each period with some initial wealth (a_t) and immediately get realizations for earnings shocks: the permanent shock η_t which updates the permanent differential term δ_t, and the transitory shock, ϵ_t. They also get an immediate realization for health status (h_t) and their out of pocket medical expenses (m_t). The value of average earnings $\overline{e}_t$ is known from the previous period. Then, conditional on the value of benefit claiming age (R), the agent knows cash on hand ($a^h{}_t$) because labor supply is solely a function of age relative to claim age. Given cash on hand, the agent chooses a value for consumption based on the derived decision rule $c_t = c(\delta_t, \overline{e}_t, a^h{}_t, R, h_t, t)$. The decision rule can only be used to solve for consumption directly if all the state variables are exactly at grid points, which is never the case for the continuous variables. Therefore, linear interpolation (in three dimensions) is used to solve for consumption. The actual state space for the discrete variables is straightforward (t has 80 values from 21 to 100, R has 9 values from 62 to 70, and h has two values for good and bad health) but the dimensioning of the three continuous variables reflects a tradeoff between precision and computational constraints. The results in this paper use a wealth grid of 25 points (in logs), and average earnings/permanent differential grids of 10 points each.

Because the consumption decisions are conditional on claim age R, the complication when simulating the model is that (prior to an actual R being realized) the agent must consider every possible value for claim age, and choose the one with the highest expected utility. To accomplish this, expected utility is tracked in the recursion using the six state variables $EU_t(R) = EU(\delta_t, \overline{e}_t, a^h{}_t, R, h_t, t)$. There are nine discreet values for benefit claim age (62 through 70; everyone claims by age 70, because there is no reason not to in the U.S.) therefore choosing an optimal path (before R is realized) involves computing separate outcomes for each of those nine possible values, each time solving for cash on hand and then consumption. The claim age R with the highest expected utility then becomes the reference point value *as of the current age*; that is, the agent behaves as though they will claim at age R, even though next period the various shocks may change that decision.

The sequencing of events when simulating life-cycle outcomes is structured to be consistent with the approach used to solve the dynamic program. There are different ways to solve these types of problems, all starting with the basic principle of recursively solving Bellman's Equation across discreet grids. The starting point for that process is the observation that in the last period of life (maximum age T equals100 in the current version of the model) rational agents will consume all available cash on hand regardless of the values for any other state variables and leisure, thus $c_T = c(\delta_T, \bar{e}_T, a^h{}_T, R, h_T, T) = a^h{}_T$. Then, the measure of expected utility in period T can be expressed in terms of the state variables $EU_T = EU(\delta_T, \bar{e}_T, a_T, R, h_T, T)$.

The problem for consumption involves a tradeoff in periods before T, because now the consumer must decide (given the state) how to allocate cash on hand between current consumption and next period wealth. Given leisure in period T-1 (which depends only on R in the exogenous labor supply version of the model) the consumer solves for a decision rule $c_{T-1} = c(\delta_{T-1}, \bar{e}_{T-1}, a^h{}_{T-1}, R, h_{T-1}, T-1)$ by maximizing the expression $U(c_{T-1}, l_{T-1}, h_{T-1}) + \beta\pi_{T/T-1} E[U(c_T, l_T, h_T)]$. Note that expected value of period T utility is a function of the T-1 consumption choice and probable realizations for the period T shocks, which depends on the set of probability states.

There are differences in approaches to dealing with probability states when evaluating and maximizing expected utility. The approach used here is to discretize the distribution of permanent and transitory earnings and medical out of pocket expenditure shocks, and directly compute expected utility by evaluating outcomes at discrete combinations of earnings and health shock outcomes. This process is very time consuming because there are (arbitrarily) 11 discrete outcomes for each of the two earnings shocks, three for the medical expenses, and two outcomes for the health status transitions, so every guess at consumption require 11*11*3*2 = 726 valuations of cash on hand. Developing an alternative approach to computing expected utility (perhaps using Monte Carlo integration) is a first-order priority for future work.

A2. Solving the Model with Endogenous Consumption, Benefit Claiming, and Labor Supply

The complete model with endogenous labor supply is incrementally more complex. In principle, if consumption, leisure, and claiming decisions are made simultaneously, the problem in the text can be solved in terms of five state variables: age, wealth, health status, average earnings through the current age, and the permanent earnings differential. However, to make the dynamic programming feasible (or at least more tractable) the consumption and labor supply problems are, as in the simpler problem, initially solved for every possible value of benefit claim age (R), which effectively brings the count of state variables for the consumption and labor supply problems to six. Also, logic and programming considerations suggest making the leisure and consumption decisions sequential over time.

Agents begin each period with some initial wealth level (a_t) and immediately get realizations for earnings shocks: the permanent shock η_t, which updates the permanent differential term δ_t, and the transitory shock, ϵ_t. They also get an immediate realization for health status (h_t) and their out of pocket medical expenses (m_t). Denote the pair of transitory earnings and out of pocket medical shocks using Ω_t. The value of average earnings $\bar{e}_t$ is known from the previous period.

Then, conditional on the value of benefit claiming age (R), the agent chooses a value for leisure in period t based on the derived decision rule $l_t = l(\Omega_t, \delta_t, \bar{e}_t, a_t, R, h_t, t)$. Conditional on the value of leisure (which determines earnings (e_t), Social Security taxes (τ^{ss}_t) and Social Security benefit (b_t)) the agent then solves for cash on hand (a^h_t). Finally–still conditional on the value of R, but now also conditional on the value for leisure–the agent chooses optimal consumption based on the derived decision rule $c_t = c(l_t, \delta_t, \bar{e}_t, a^h_t, R, h_t, t)$. Because the leisure and consumption decisions are conditional on claim age R, the complication when simulating the model is that (prior to an actual R being realized) the agent must consider every possible value for claim age, and choose the one with the highest expected utility. To accomplish this, expected utility is tracked in the recursion using the subset of state variables $EU_t(R) = EU(\delta_t, \bar{e}_t, a_t, R, h_t, t)$. There are nine discreet values for benefit claim age (62 through 70; everyone claims by age 70, because there is no reason not to in the U.S.) therefore choosing an optimal path (before R is realized) involves computing separate outcomes for each of those nine possible values, each time solving for leisure, cash on hand, then consumption. The claim age R with the highest expected utility then becomes the reference value *as of the current age*; that is, the agent behaves as though they will claim at age R, even though next period the various shocks may change that decision. Given consumption, the last steps needed to end the simulation for period t are updating average earnings and incrementing next period beginning wealth using $a_{t+1} = a^h_t - c_t$.

As above, the starting point for the solution process is the observation that in the last period of life (maximum age T equals100 in the current version of the model) rational agents will consume all available cash on hand regardless of the values for any other state variables and leisure, thus $c_T = c(l_T, \delta_T, \bar{e}_T, a^h_T, R, h_T, T) = a^h_T$. Choosing a leisure decision rule in the last period of life is also straightforward. The idea is to maximize utility in the last period of life, which is a function of leisure and consumption. Thus, the leisure decision rule involves choosing between the three discrete values (no work, part-time, and full-time) and evaluating which leads to higher expected utility. Given a choice of leisure, the outcome for last-period cash on hand (and hence last period consumption) depends only on the common state variables (permanent earnings differential, average earnings, health status, benefit claim age), beginning of period wealth, and the value of the transitory earnings and out of pocket medical shocks (Ω_t). Thus, the solution to the leisure decision can be expressed notionally as $l_T = l(\Omega_T, \delta_T, \bar{e}_T, a_T, R, h_T, T)$ being equal to the value for leisure that maximizes $U(c_T, l_T, h_T)$ conditioned on Ω_T and the other states. Summing probability-weighted outcomes across values for ϵ_t and μ_t then allows one to construct a measure of expected utility in period T, which can be expressed (as above) in terms of the subset of state variables $EU_T = EU(\delta_T, \bar{e}_T, a_T, R, h_T, T)$.

The problem for consumption involves a tradeoff in periods before T, because now the consumer must decide (given the state) how to allocate cash on hand between current consumption and next period wealth. This is where the sequential solution proves most useful, because the structure is such that the consumer chooses an optimal consumption conditional on each possible value of leisure in the current period. For example, given leisure in period T-1, the consumer solves for a decision rule $c_{T-1} = c(l_{T-1}, \delta_{T-1}, \bar{e}_{T-1}, a^h_{T-1}, R, h_{T-1}, T-1)$ by maximizing the expression $U(c_{T-1}, l_{T-1}, h_{T-1}) + \beta\pi_{T/T-1} E[U(c_T, l_T, h_T)]$. Note that everything needed to calculate the expected value of period T-1 utility is known when the consumption decision is being made. Given the states, a value for consumption determines $a_{T,}$ and the values (from discretized probability distributions) for the various shocks are then sufficient to solve for l_T in any state of the world that could exist in period T. Given l_T and the probability state, the solution for c_T is a straightforward application of the period T decision rule $c_T = a^h{}_T$. Given a period T-1 consumption rule, the period T-1 leisure rule is derived exactly the same way the period T leisure rule was derived. That is, for each discretized combination of the transitory and out of pocket medical shocks, solve for the value of leisure that maximizes utility conditional on those shocks. The period T-1 solution ends with the calculation of expected utility, which, as in period T, is solved for by summing over the probability-weighted values for ϵ_t and $\mu_{t.}$

Figure 1. Male Old Age Insurance (OAI) Worker Cumulative Claiming Rates in the CBOLT Model (1960 Cohort, For Whom Normal Retirement Age (NRA) is Fully Phased-In to Age 67)

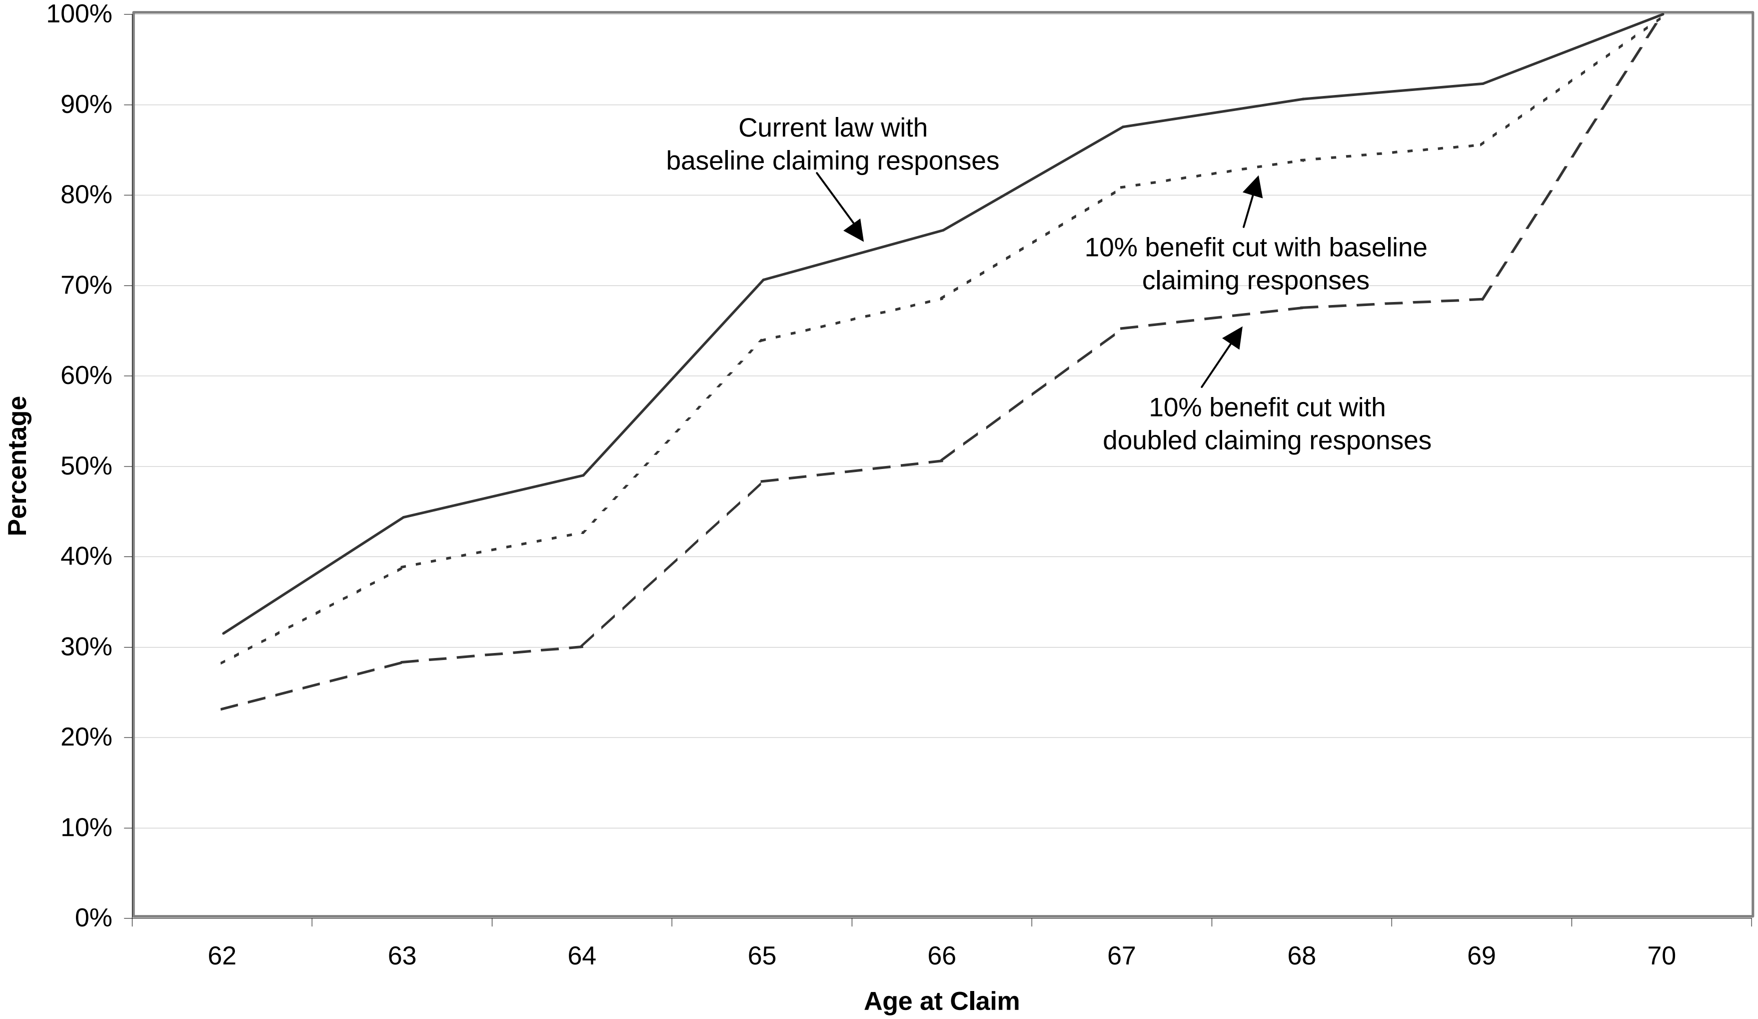

Table 1. Percent Change in CBOLT Outcomes Under 10% Benefit Cut Policy Experiment

Year	Claiming Assumption		
2014	**No claiming responses**	**Baseline claiming responses**	**Doubled claiming responses**
OAI Worker Beneficiaries	0.0%	-2.7%	-5.1%
OAI Worker Average Benefits	-5.5%	-4.7%	-4.0%
OAI Worker Total Benefits	-5.6%	-7.2%	-9.0%
Total Hours Worked	0.0%	0.3%	0.6%
Real GDP	0.0%	0.2%	0.4%
Payroll and Income Taxes	0.0%	0.2%	0.3%
2050	**No claiming responses**	**Baseline claiming responses**	**Doubled claiming responses**
OAI Worker Beneficiaries	-0.5%	-1.6%	-4.9%
OAI Worker Average Benefits	-10.0%	-7.8%	-4.4%
OAI Worker Total Benefits	-10.5%	-9.3%	-9.1%
Total Hours Worked	0.1%	0.3%	0.9%
Real GDP	0.1%	0.3%	0.9%
Payroll and Income Taxes	0.0%	0.3%	0.9%

Table 2. Summary of Parameter Settings and Stochastic Process Assumptions

Basic Parameters

- Maximum life span (T) = 100
- Single period discount factor (β) = 0.95
- Utility function parameter, determines intratemporal substitution (α) = .75
- Utility function parameter, determines intertemporal elasticity of substitution (γ)=3
- Age-specific utility function parameter, affects disutility of working (Γ_t) = 1.0
- Consumption floor (c^{min}) varies from $1 to $10,000 across simulations
- Normalized leisure (l_t) = .6 if working full-time, 1.0 if not working

Earnings Process

- Exogenous earnings component (λ_t) set so age 60 average earnings = $30,000
- Standard deviation of shocks to permanent earnings differential (σ_η) = .05
- Standard deviation of transitory earning shocks (σ_ϵ) = .375
- Initial (age 21) dispersion of permanent differentials (δ_t) =.08

Health States and Health Shock Transitions

- Normalized health status in period t (h_t) = 1.0 in good health, 0.75 in poor health
- Health status transition ($h(t,h_{t-1})$) probability of moving from good to poor health = .02 before age 50, and = .05 after age 50. Probability of moving from poor health to good health always = .01.

Medical Expenditure Shocks

- Expected value of medical expenses ($m(h_t)$) in good health state = $1,000, poor health state = $5,000.
- Standard deviation of medical expenses (σ_μ) in good health state = $500, in poor health state = $2,5000.

Macroeconomic Assumptions

- Population growth rate = 1% (for Cross-Section Simulations)
- Growth in real Average Wage Index (AWI) = 1%
- Single period (certain) net return on wealth (r) = .05

Table 3. Comparison of Cross-Section Simulated and Survey of Consumer Finances (SCF) Wealth Distributions

	Normalized Wealth Percentiles (Ratio of Each Percentile to Overall Average Income)										
Wealth Percentile	**10th**	**20th**	**30th**	**40th**	**50th**	**60th**	**70th**	**80th**	**90th**	**95th**	**99th**
Actual Percentiles, 2001 SCF											
Total Networth	0.0	0.1	0.3	0.7	1.3	2.0	3.3	5.5	10.8	19.2	86.3
Non-Housing Networth	0.0	0.0	0.1	0.3	0.5	0.9	1.7	3.2	7.7	14.6	72.8
Simulated, Medical Expenditure Shocks On											
Consumption Floor = $1	0.2	0.7	1.1	1.5	1.9	2.3	2.9	3.8	5.6	8.0	33.3
Consumption Floor = $5,000	0.0	0.1	0.4	0.7	1.2	1.8	2.6	3.7	5.7	8.1	33.9
Consumption Floor = $10,000	0.0	0.0	0.1	0.3	0.6	1.1	1.8	2.9	5.0	7.6	32.6
Simulated, Medical Expenditure Shocks Off											
Consumption Floor = $1	0.0	0.0	0.0	0.0	0.0	0.1	0.3	0.9	2.3	4.0	11.6
Consumption Floor = $5,000	0.0	0.0	0.0	0.0	0.0	0.0	0.2	0.8	2.2	3.9	11.2
Consumption Floor = $10,000	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.7	2.1	3.9	11.6

Figure 2. Consumption and Income Levels by Age
(Consumption Floor=$1, Medical Expenditure Shocks On, Health=Good, Claim Age=62)

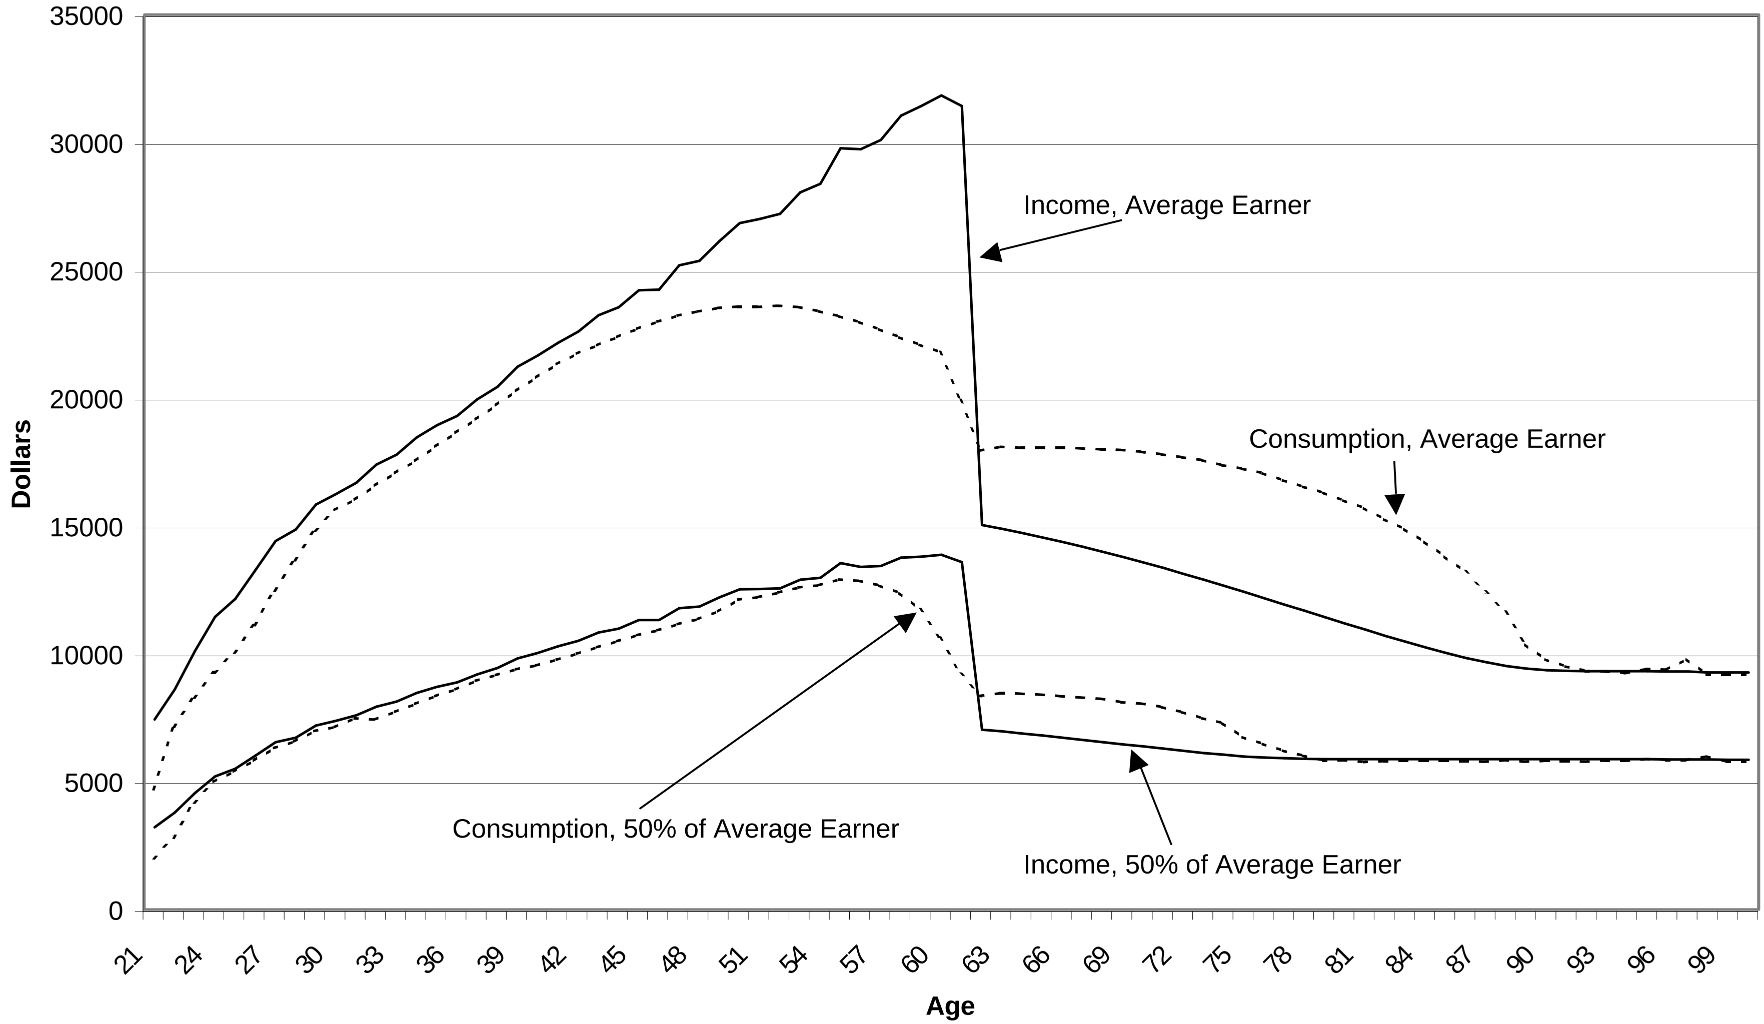

Figure 3. Consumption and Income Levels by Age
(Consumption Floor=$5,000, Medical Expenditure Shocks On, Health=Good, Claim Age=62)

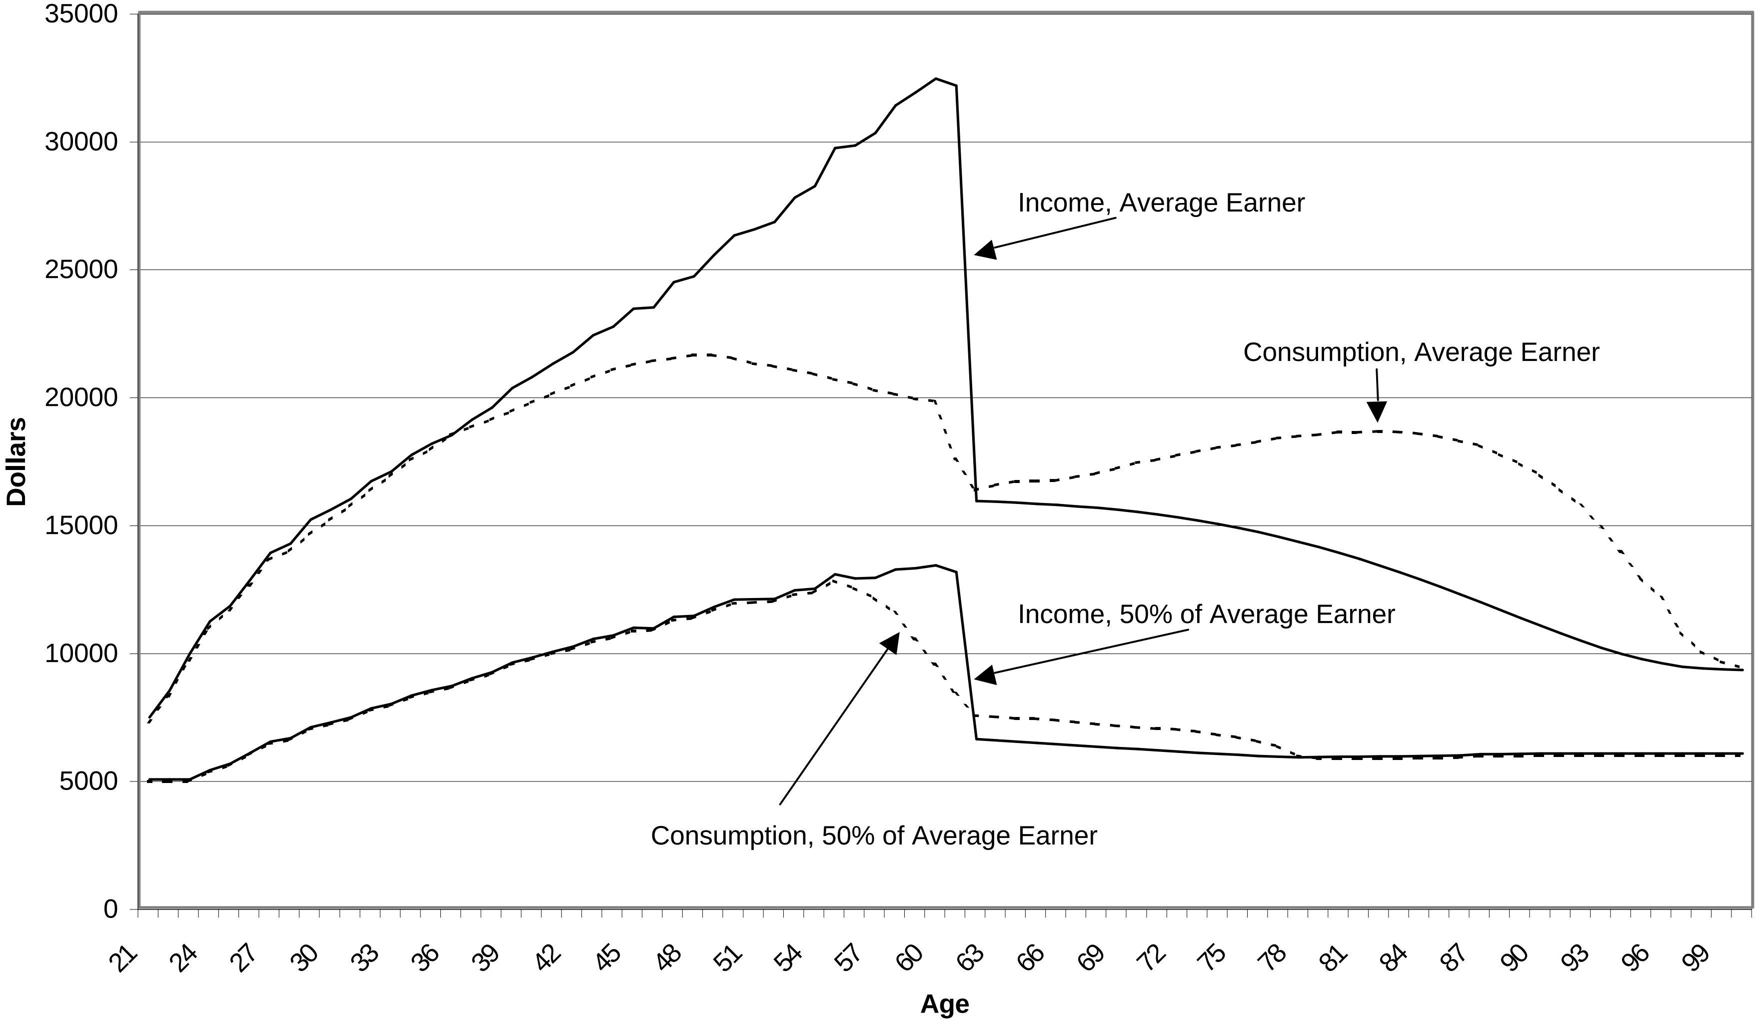

Figure 4. Ratio of Wealth to Income by Age
(Consumption Floor=$1, Medical Expenditure Shocks On, Health=Good, Claim Age=62)

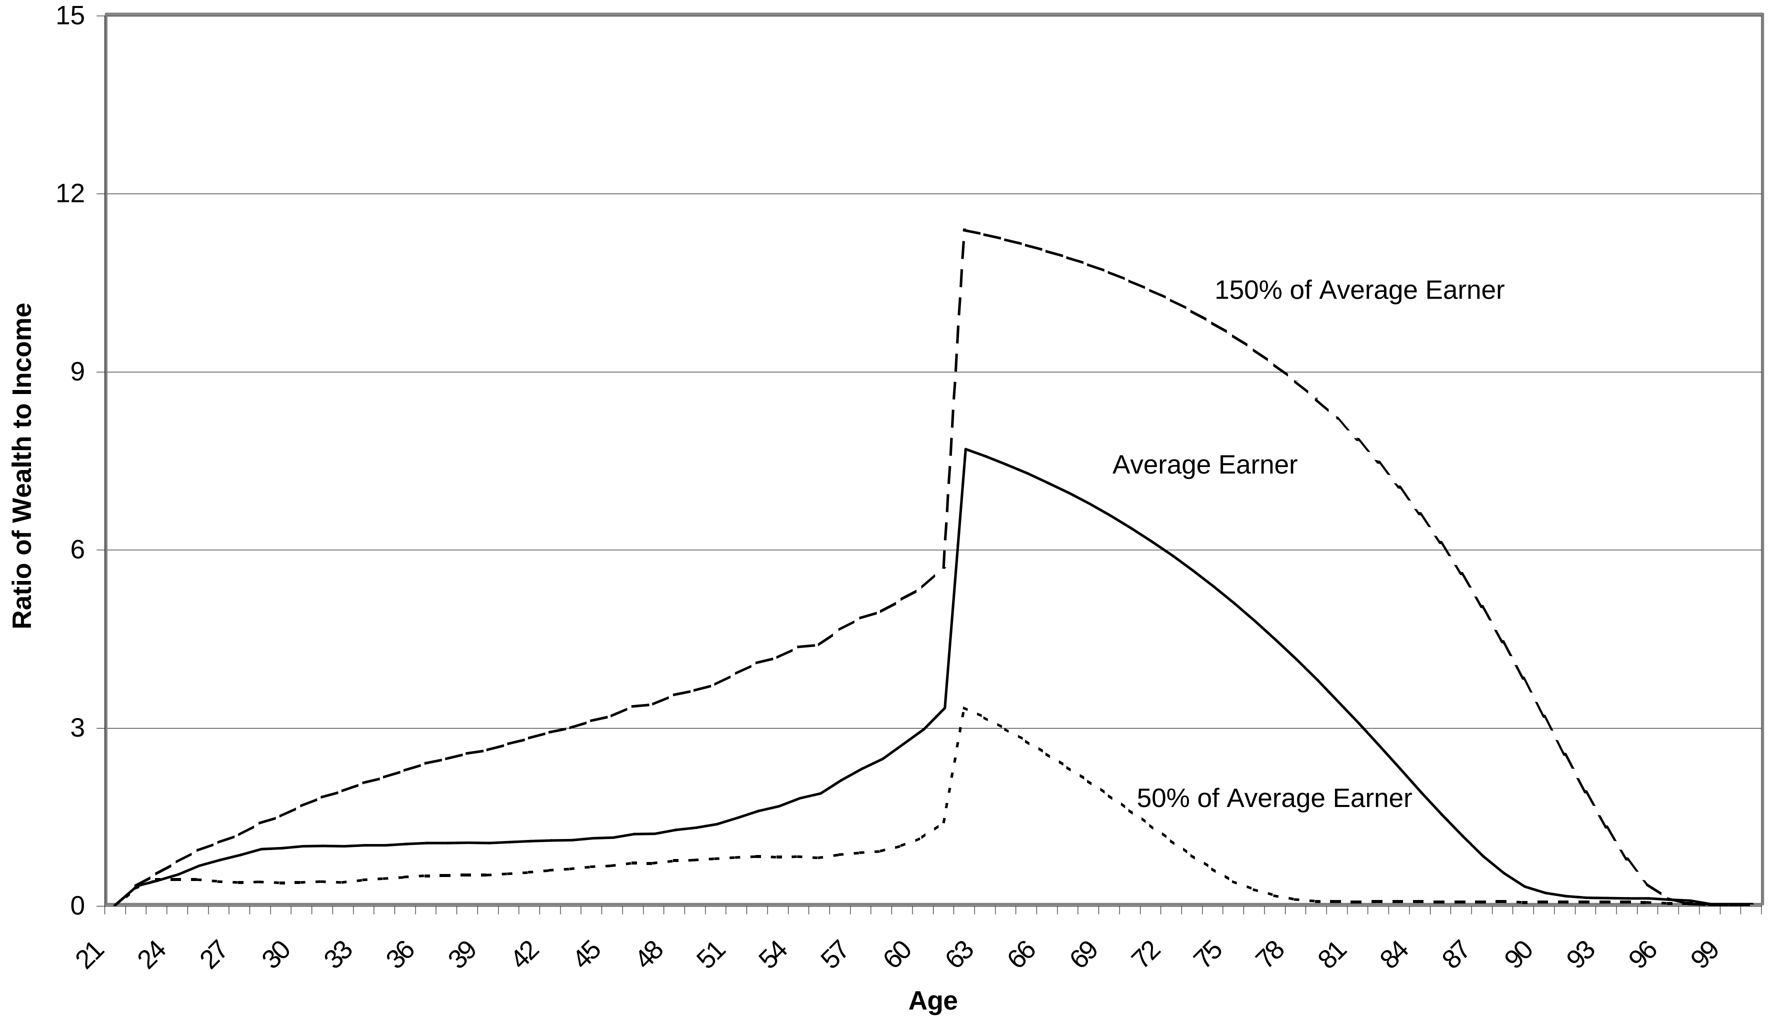

Figure 5. Ratio of Wealth to Income by Age
(Consumption Floor=$5,000, Medical Expenditure Shocks On, Health=Good, Claim Age=62)

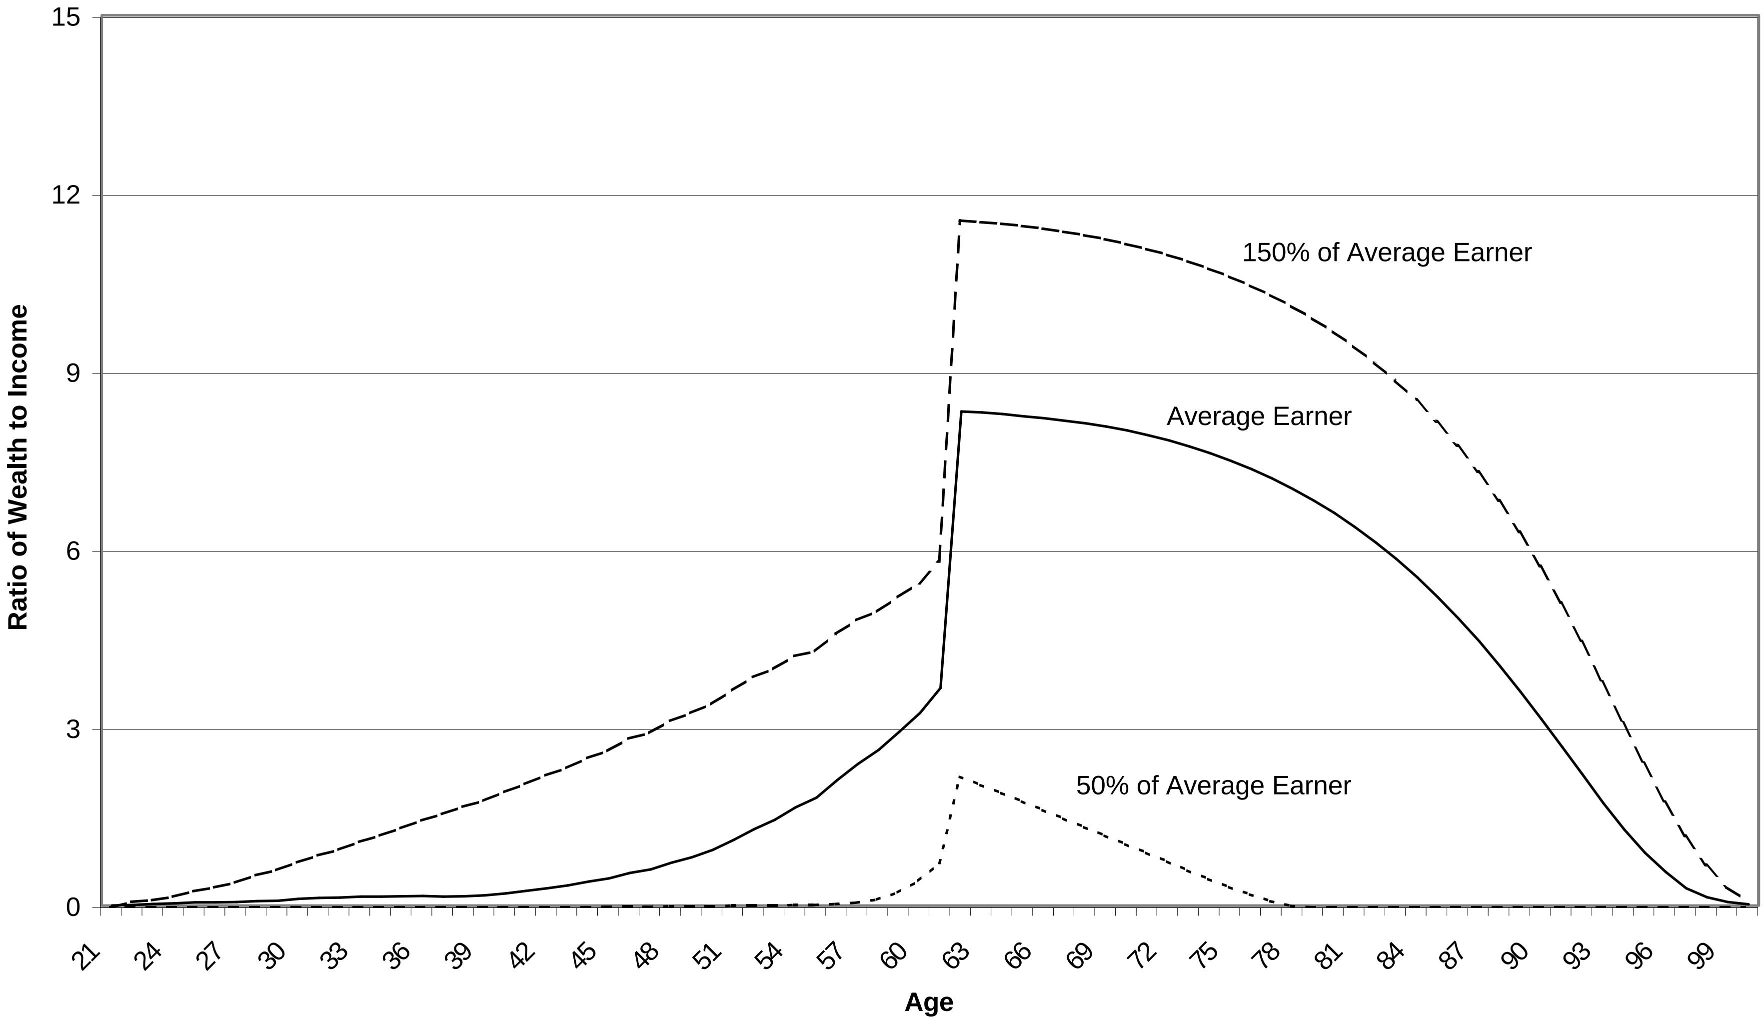

Table 4. Simulated Longitudinal Wealth Accumulation Patterns

Lifetime Average Earnings Quintile	Ratio of Wealth at Age 62 to Average Lifetime Income					
	All	Lowest	Second	Middle	Fourth	Highest
Medical Expenditure Shocks On						
Consumption Floor = $1	5.3	5.1	4.2	3.8	3.5	6.7
Consumption Floor = $5,000	5.0	1.3	3.4	3.9	3.6	6.5
Consumption Floor = $10,000	4.4	0.1	1.3	2.5	3.0	6.5
Medical Expenditure Shocks Off						
Consumption Floor = $1	1.7	0.9	1.1	1.2	1.3	2.3
Consumption Floor = $5,000	1.7	0.6	1.0	1.1	1.2	2.3
Consumption Floor = $10,000	1.6	0.2	0.7	1.0	1.1	2.3

Table 5. Effect of Ten Percent Benefit Cut on Simulated Longitudinal Wealth Accumulation Patterns (Medical Shocks on For All Simulations)

Lifetime Average Earnings Quintile	Ratio of Wealth at Age 62 to Average Lifetime Income					
	All	Lowest	Second	Middle	Fourth	Highest
Consumption Floor = $5,000						
Baseline	5.0	1.3	3.4	3.9	3.6	6.5
Ten Percent Benefit Cut	5.1	1.1	2.6	3.3	3.4	7.3
Consumption Floor = $10,000						
Baseline	4.4	0.1	1.3	2.5	3.0	6.5
Ten Percent Benefit Cut	4.8	0.1	1.4	2.7	3.1	7.1