Seesaws and Social Security Benefits Indexing

Citation

Weinzierl, Matthew. "Seesaws and Social Security Benefits Indexing." Harvard Business School Working Paper, No. 15-039, November 2014. (NBER Working Paper Series, No. 20671.)

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:13479078>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Seesaws and Social Security Benefits Indexing

Matthew Weinzierl

Working Paper

15-039

November 12, 2014

Copyright © 2014 by Matthew Weinzierl

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Seesaws and Social Security Benefits Indexing

Matthew Weinzierl*

November 12, 2014

Abstract

The price indexation of Social Security benefit payments has emerged in recent years as a flashpoint of debate in the United States. I characterize the direct effects that changes in that price index would have on retirees who differ in their initial wealth at retirement and mortality rates after retirement. I propose a simple but flexible theoretical framework that converts benefits reform first into changes to retirees' consumption paths and then into a net effect on social welfare. I calibrate that framework using recently-produced data on Social Security beneficiaries by lifetime income decile and both existing and new survey evidence on the normative priorities Americans have for Social Security. The results suggest that the value retirees place on protection against longevity risk is an important caveat to the widespread enthusiasm for a switch to a slower-growing price index such as the chained CPI-U.

Introduction

The indexation of Social Security benefit payments may seem like an issue about which only an economist could get excited, but it has emerged in recent years as a flashpoint of debate in the United States. In his 2014 budget, President Obama proposed changing the price index with which retiree benefits are adjusted for inflation. In brief, the change was expected to lower the growth rate of benefits for all retirees, though at advanced ages that change would have been offset by progressive "benefit enhancements." Because it was not tied to an increase in the starting level of those benefits, the President's proposal was expected to reduce the total present value of benefits. The President's proposal was explicitly intended to appeal to Congressional Republicans eager to reduce future spending on Social Security, but it was deeply unpopular with many of his fellow Democrats¹. When negotiations on more general fiscal policy challenges yielded little progress over the subsequent year, the President removed the proposal from his 2015 budget. His spokesperson

*This paper was prepared for the Fall 2014 BPEA conference. Thanks to the editors, David Romer and Justin Wolfers, as well as to my discussants, Martin Feldstein and Aleh Tsyvinski, and the many participants in that conference for helpful comments and discussions. Thanks to Darren A. Rippy at the BLS for sharing the data on CPI-E. The author can be contacted at 277 Morgan Hall, Harvard Business School; mweinzierl@hbs.edu.

¹The indexing change was intended to cover tax brackets, as well, and the President's proposal was thereby intended to be a net positive contributor to the government budget.

made clear, however, that changes to indexation were still on the table if included in broader budget deals.

While the overall fiscal implications of benefits-indexing reform have been widely discussed, this paper’s contribution is to explore both the positive and normative aspects of its distributional consequences across the population of retirees.² In particular, I study the direct effects that changes in benefits-indexing have on retirees who differ in two important ways: initial wealth at retirement and mortality rates after retirement. I propose a simple but flexible theoretical framework that converts benefits reform first into changes to retirees’ consumption paths and then into a net effect on social welfare. I use recently-produced data on the net worth, benefit levels, and mortality risks of Social Security beneficiaries by lifetime income decile to provide quantitative results. Finally, I introduce survey evidence on the priorities Americans have for Social Security, a first step in pinning down the normative implications of these effects of indexing reform. The specific questions I use in the survey take a novel form that may be useful for estimating normative preferences across a wide range of policy issues.

In brief, I find that a useful metaphor for thinking about the direct effects of indexation on heterogeneous retirees is a playground seesaw, where two facts about retired households in the United States push down on opposite ends.

Pushing down on the left end of the seesaw (i.e., toward an increasing path of real benefits) is the large majority of Social Security beneficiaries who worry about outliving their private assets and having to rely nearly exclusively on those benefits to fund expenditure late in life. In fact, a core purpose—and achievement—of Social Security is to prevent the elderly from falling into poverty as they age (see Gary V. Englehardt and Jonathan Gruber, 2004). As has long been understood, for instance in Feldstein (1987), benefits that rise in real terms over retirement and are therefore backloaded later in life will provide valuable protection against longevity risk for retirees with positive private wealth (throughout this paper, I assume that private annuitization of wealth outside of DB pensions is unavailable). A faster-growing price index will therefore generate welfare gains through its effects on these households.

Pushing down on the other end of the metaphorical seesaw (i.e., toward a decreasing path of real benefits) sit the poorest retiree households, who also have the highest mortality rates (see James E. Duggan, Robert Gillingham, and John S. Greenlees, 2006). Most directly, the poorest retirees sit on this end because, with little wealth at the start of retirement, they benefit less from Social Security’s effective annuitization. A more subtle reason is that a faster-growing price index that backloads the present value of benefits has the effect of redistributing, through an actuarially unfair adjustment, some of the total value of benefits away from poorer retirees when mortality is inversely related to income. To the extent that these retirees are the ones who most need support from Social Security, a faster-growing price index thereby generates welfare losses.

²The effects of moving to a slower-growing price index, such as the chained CPI-U, on the paths of benefits across retirees have been analyzed by a number of researchers. See CRFB (2013) and Anya Olsen (2008), for example. These prior analyses did not translate the effects on benefits into implications for consumption or welfare, and they did not compare alternative indexing schemes—the two main contributions of this paper.

In other words, heterogeneity across retiree households means that any given reform to benefits-indexing generates effects with exactly opposite welfare implications. In this way, benefits-indexing policy inevitably has distributional consequences and, as we will see when considering the proposal by President Obama, may even be used to pursue distributional goals.

It is important to emphasize from the start that this paper focuses on the direct effects of benefits-indexing reform on retirees, abstracting from a number of general equilibrium effects and other factors that matter for the optimal path of benefits and that, therefore, ought to be part of a comprehensive evaluation of indexing reform.³ Most prominently, changing the path of benefits may affect individuals' labor effort and saving decisions during their working lives, but my calculations hold fixed households' behavior prior to retirement. Related, I do not consider the implications of benefits-indexing reform for the accumulation of the economy's capital stock, and I abstract from the controversial possibility that benefits paid earlier will yield gains to households who can achieve a higher rate of return in the private investment market than they obtain from the natural rate of return of a pay-as-you-go Social Security system (see Feldstein 1987, 1990). Finally, technological change, especially in the context of medical care for the elderly, may affect the optimal response of policy to an increase in real benefits and therefore matter for the choice of indexing.⁴ This paper's omission of these factors is not meant to imply that they can be ignored. Instead, I omit them to better focus on one piece of that broader question.

This paper also abstracts from several complications specific to the Social Security system that may matter for the results but that would make the analysis and intuition for the results substantially less straightforward. In particular, I do not model spouses' joint decisions about benefits or surviving spouses' decisions about benefits options, instead treating the household as the unit of analysis; I do not allow for early or late retirement, instead having all households retire at the same age; and I do not include the disability benefits portion of Social Security in the analysis. Microsimulation models that capture much of the complexity of the actual Social Security system, for instance the MINT model as described in Smith and Favreault (2013), may be useful for including these features in future analyses.

Which side of the seesaw carries more weight? I show that the answer depends on both positive factors about which we have some good existing evidence and normative factors about which we have very little. For the positive factors, I show that a large majority of retirees are likely to sit on the left end of the seesaw, that is, favor a steeper path of benefits that effectively annuitizes more of a given retiree's total wealth. Moreover, the simulations below suggest that the gains to the poor from frontloading benefits are much smaller, in consumption or individual utility terms, than

³In principle, as suggested to me by Martin Feldstein, the design of the optimal path of benefits and the identification of an ideal price index are separate tasks. If we believe the path of real benefits has been chosen optimally in current policy, such a separation is natural. In this paper, I explore the question of how proposed price indexes affect the path of real benefits and, therefore, retirees and social welfare.

⁴Suppose that advances in medical care for the elderly allowed them to purchase a higher quality of life at a lower real cost, such as through the introduction of a new product. Their real benefits would rise in this case, but so too would their ability to generate extra welfare with additional resources. In that case, it may be important to target benefits toward households with high marginal utilities of consumption, not low values of real benefits.

the gains to the majority of retirees from backloading benefits. These positive results suggest that the direct effects on retirees of frontloading benefits, as in a switch to the chained CPI, are likely to generate a net loss of welfare unless society puts a strong normative priority on the poorest, shortest-lived retirees relative to the rest of the population and, in particular, relative to poor and middle-class retirees who outlive their life expectancy at retirement. In other words, the results of this paper suggest that the value retirees place on protection against longevity risk is an important caveat to the widespread enthusiasm for a switch to a slower-growing price index such as the chained CPI-U.

To explore the normative aspects of this question, I consider two classic normative criteria and generate novel opinion survey evidence on the relevant preferences of Americans. The two classic criteria would endorse opposite reform proposals as simulated here: that is, the utilitarian criterion would endorse backloading while the Rawlsian would endorse frontloading. Survey respondents put equal value on increasing benefits to poor retirees who die young and poor retirees who outlive their life expectancies, and they put substantially less value—perhaps even negligible value—on increasing the benefits of average retirees. These results are inconsistent with either a utilitarian or Rawlsian criterion on its own, but applying them to the simulated reform results suggests that backloading of benefits is likely to generate net welfare gains, at least in terms of its direct effects on retirees, as is the case under the standard utilitarian criterion.

From a policymaking perspective, the net positive welfare implications of the direct effects on retirees of moving to a faster-growing price index might be expected to translate into political support for such a reform, but two political realities make that support less likely: namely, public opposition to benefit reductions, and pressure from some policymakers to lower total Social Security spending.⁵ To see why, note that such a reform automatically means a decrease in initial benefits for retirees (if total spending is held fixed) or an increase in total spending (if initial benefits are held fixed). Taking those political realities into account, the results of this paper shed some light on the specific reform President Obama proposed in his 2014 budget. That proposal, which was designed to reduce total spending by maintaining initial benefit levels but slowing their growth rate, used "benefit enhancements" at advanced ages to protect some of the effective annuitization that frontloading would otherwise have sacrificed. As I show below, the progressive design of those benefit enhancements meant that they would provide this protection largely to lower-income households. The President's proposed reform would thus simultaneously achieve the positive effects of frontloading on the poorest, shortest-lived retirees and the positive effects of backloading on the poorest, longest-lived retirees, and it would bring a substantial net welfare gain under the utilitarian, Rawlsian, or survey-based normative criterion.⁶ Of course, that reform would generate losses as well, reducing the well-being of the higher-income half of the retiree population and—therefore—potentially having disincentive effects that would reduce its appeal in a more

⁵Of course, the indirect effects of reform not included in this paper's analysis, such as the effects on private saving and capital accumulation, may also explain resistance to reform.

⁶Note that the frontloaded element of the proposal adds to its appeal under the Rawlsian criterion, but not under the utilitarian or survey-based criteria.

comprehensive analysis. The President's proposal thus illustrates the inherent connection between benefits-indexing policy and the redistributive role of Social Security.

The paper proceeds as follows. Section 1 reviews how Social Security uses indexing today, lays out the seesaw metaphor described above, and briefly summarizes the empirical literature behind the factors at each end. Section 2 presents a simple model that allows us to analyze these direct effects on retiree households with a small set of positive and normative parameters. Section 3 simulates a version of that model using U.S. data, and it considers three prominent indexing reform proposals: the chained CPI-U, the CPI-E (an experimental series calculated by the BLS "using households whose reference person or spouse is 62 years of age or older"), and the chained CPI-U augmented with late-in-life "benefit enhancements" as proposed by President Obama in 2014. Section 4 presents novel, but far from definitive, opinion survey evidence on the normative components of the model and uses that evidence as well as conventional normative criteria to provide suggestive welfare evaluations of the direct effects of the three policy options. Section 5 extends the analysis to include several aspects omitted from the baseline case, and Section 6 concludes.

1 Background and Key Considerations

The current Social Security system uses indexing—that is, adjusting nominal values over time—in three ways. First, it scales the income earned during a beneficiary's working years into current dollars when calculating the value at retirement of his total accumulated Social Security earnings. Second, it indexes the bracket points of the progressive function that converts that scaled lifetime earnings into a monthly benefit. Third, it indexes benefits upon retirement. For the first two instances of indexing, the current system uses a wage index; for the third it uses the CPI-W, the consumer price index for urban wage earners.

These three instances of indexing can be seen as serving different purposes. The first, which I will call earnings-indexing, is most naturally seen as trying to capture the natural rate of return of the pay-as-you-go (or "unfunded") Social Security system, which is closely related to the growth rate of nominal wages.⁷ The second, which I will call brackets-indexing, tries to preserve the desired progressivity (across lifetime earnings levels) of the system despite changes in the wage distribution and nominal wages. The third, which I will call benefits-indexing, tries to protect the real value of retirees' benefits over time, though as emphasized throughout this paper it also has implications for the effective progressivity of the system due to differences in mortality by lifetime income levels. It is benefits-indexing that has been the focus of public debate, and it will be my focus in this paper

⁷Earnings-indexing could, in principle, serve many purposes. Because the lifecycle path of earnings varies systematically with the value of lifetime earnings, the choice of indexing will tend to favor some earners over others. One could try to use that choice, therefore, as a new optimal tax instrument that would relax the classic efficiency-equality tradeoff. Similarly, one might try to take advantage of the effect that expected earnings-indexing has on the extent to which workers view the payroll tax as a tax, rather than as a form of saving. These are purposes that can be more directly pursued by adjusting the history-dependent redistributive elements of Social Security, such as the replacement rates in each income bracket.

as well.⁸

1.1 Budget-neutral benefits-indexing

This paper's baseline analysis focuses on budget-neutral benefits-indexing reforms. By "budget-neutral" I mean that the expected present value of benefits (across all individuals in an age cohort) is unaffected by the way in which benefits are indexed. Therefore, in the analysis below in which I consider a shift to an index that causes a steeper rise in benefits over time, I adjust (down) the starting value of benefits for all beneficiaries in the cohort by the factor required to keep the expected present value of total benefits the same. The assumption of budget neutrality is not necessary, but it allows us to focus on the direct effects of the time path of benefits rather than their level. As I show in Section 5, reforms that are not budget neutral can be analyzed using this paper's approach as well, and the main lessons are unaffected.

A simple but useful observation about changes to budget-neutral benefits-indexing is that their effects on benefit levels are highly concentrated toward the beginning and end of retirement, as illustrated in the following figure.

⁸There appears to be little interest in reform to the other two uses of indexing. Even the Bowles-Simpson proposal (White House 2010), which suggests changing the bracket points to increase progressivity, does not change the methods of earnings-indexing or brackets-indexing. This is somewhat unfortunate, in that changes to the method of earnings-indexing hold substantial promise for more closely aligning the current system with its true "natural" rate of return. In particular, earnings-indexing could include changes to projected beneficiary-worker (dependency) ratios and aggregate life expectancies. Bracket-indexing would be a simple way of implementing a limited version of the inequality adjustments suggested by, for example, Robert Shiller (2003).

Figure 1: Benefits received (in \$2005) under three paths with the same expected present value for a male, 65-year-old retiree with median household earnings and average mortality rates.

To produce Figure 1, I start with the initial annual Social Security benefit for the median retired household from the fifth decile of household lifetime earnings, as calculated in John Karl Scholz, Ananth Seshadri, and Surachai Khitatrakun (2006) and shown in real (2005) dollars. I assume the "Status Quo" policy would provide a constant stream of real benefits at this level (in Section 5 I show that alternatives to this assumption do not change the lessons of the baseline analysis). As alternatives to a constant real benefit, I consider two paths : a "Frontloaded" path in which real benefits grow at an annual rate of -0.27 percentage points, and a "Backloaded" path in which they grow at +0.37 percentage points. These alternatives correspond to two prominent proposals for benefits-indexing reform: namely the use of the chained CPI-U and the CPI-E indexes calculated by the BLS. Using average mortality rates from the Social Security Administration and an annual discount factor of 0.96, I adjust initial benefits under these two alternatives to ensure that the expected present value total cost of each path is the same as for the Status Quo.

Figure 1 makes clear the roots of the seesaw metaphor: the sizeable differences in benefits early and late in retirement across benefit paths. Retirees who value Social Security's insurance against longevity risk, and especially those who come to rely on Social Security benefits because they outlive their private savings, will prefer the backloaded benefits path. Retirees with little private wealth or high mortality risks, and especially those who do not survive to advanced ages, will prefer the frontloaded path. I now turn to a brief discussion of the existing literature on these two competing

features in the retired population in the United States.

1.2 Evidence on variation in private retirement savings

An extensive literature has examined whether retirees enter retirement with sufficient assets to sustain their economic well-being. In general, the results have drawn a qualitative distinction between the status of the large majority of the Social Security population and approximately the bottom quintile of retirees. Reassuringly, most retirees appear to reach retirement with sufficient assets (both Social Security and non-Social Security) to smooth shocks, supplement Social Security benefits, and maintain what a rational life-cycle consumer with their lifetime earnings history would plan as an optimal path of expenditures. The bottom quintile, in contrast, enter retirement (or soon find themselves) almost entirely dependent upon Social Security benefits and other transfers. A number of studies have obtained such a result, and some have drawn policy implications from it. I briefly review several of them here, though of course the literature is too large to do it full justice in such a short discussion.

Rudolph G. Penner and Karen E. Smith (2010) summarize the findings of Smith, Mauricio Soto, and Penner (2009), who use Health and Retirement Survey (HRS) data from 1998 through 2006 to study the assets held by retirees, including the expected present value of their pension and Social Security benefits. They conclude: "The net worth of those in the top quintile of the income distribution increased until age 85 (figure 1). For those in the three middle quintiles, net worth began declining after age 70, but only very slowly. Evidently, the vast majority in this portion of the income distribution will die with a significant amount of assets. Few older households, including those with little income, used home equity to finance retirement consumption. The bottom income quintile never accumulated much wealth and spent their assets quickly, leaving them dependent on Social Security and whatever DB pensions they had earned...Our results are reassuring for households in the top 80 percent of the income distribution, but the data indicate that the lowest income quintile quickly becomes almost wholly dependent on Social Security after retirement....Reformers must be sensitive to the heavy dependence on Social Security in the lowest part of the income distribution."

Consistent with these findings, David A. Love, Michael G. Palumbo, and Paul A. Smith (2009) note that "It is reasonably well known that retirees in the bottom quintile of the income distribution (conditional on their age and marital status) rely almost exclusively on DB pension benefits, Social Security benefits, and other government transfers to finance spending." Barbara A. Butrica, Joshua H. Goldwyn, and Richard W. Johnson (2005) also use HRS data and find that "Individuals in the lowest income quintile consume between 99 and 107 percent of their after-tax income plus annuitized assets." Michael D. Hurd and Susann Rohwedder (2008) focus on preparedness for retirement across levels of educational attainment and estimate that 17 percent of married people and 36 percent of singles are not adequately prepared for retirement (meaning that they are likely to exhaust their wealth before death).

Finally, Scholz, Seshadri, and Khitatrakun (2006), hereafter SSK, use HRS data and a dynamic

lifecycle optimization model and show that only a small minority of individuals are failing to save adequately to sustain desired consumption paths. While doing so, they find that Social Security wealth dominates for at least the bottom lifetime-income decile, arguably the bottom three deciles, of retirees.⁹ As discussed below, I will rely on these authors' research for estimates of retiree wealth and Social Security benefits by income group.

1.3 Evidence on the relation between income and mortality

It has become a staple of commentary on the fiscal health of Social Security that mortality differences across income groups matter for the true impact of a range of proposed reforms, such as to the full retirement age (see Paul Krugman 2012, for example). While the literature quantifying these mortality differences is less developed than that on wealth, I draw heavily on a few relatively recent contributions.

Gopi Shah Goda, John B. Shoven, and Sita Nataraj Slavov (2009) use mortality estimates for the top and bottom halves of the earnings distribution to show the dramatic result that "Under the assumption of constant mortality across lifetime income subgroups, the Social Security system is progressive regardless of the measure shown. However, a good deal of the progressivity is undone or even reversed when differential mortality is taken into account. The results are similar for both stylized earners at different points of the earnings distribution and actual workers' earnings histories." They also offer this speculation: "Rather than analyzing the mortality differences between those in the top and bottom halves of the lifetime earnings distributions, we would have liked to have the information by lifetime income decile so that we could examine the mortality experience of the genuinely poor vs. those at other parts of the distribution. It seems likely that the extent of mortality inequality is even greater than reflected in the top half/bottom half analysis." In fact, it appears that variation in mortality does widen at the extremes of the income distribution. Duggan et al. (2007) use administrative Social Security data to show a consistently positive relationship between average age of death and lifetime earnings deciles.

Related to Goda et al.'s suggestion that more disaggregated estimates would yield additional insights, Hilary Waldron (2007) uses Social Security Administration data to characterize life expectancy for men by income quartile at 5-year increments from age 60 to age 85. A complementary data source is Bosworth and Burke (2014), who use the HRS to calculate life expectancy at age 55 for men and women as well as relative mortality rates for men and women aged 50-74 and 75+ by income quintile. Both of these sources show that retirees in approximately the bottom quarter (for example) of the lifetime earnings distribution have life expectancies 15 to 20 percent less than those in the top quarter prior to retirement. These gaps are larger than those between the second and third quarters of the income distribution, and Waldron's estimates suggest they are not narrowing over time. I explain in Section 3 how I use these results.

⁹SSK find that the underaccumulation of wealth is driven not by lifetime income per se but by being single rather than married, as single retirees have systematically lower incomes. I do not distinguish between single and married households in this analysis.

2 A Partial Reform approach to optimal benefits-indexing

In this section I lay out a simple formal structure through which to model how benefits-indexing reforms turn into changes in the consumption paths of retirees and how these consumption changes can be aggregated into a measure of social welfare.¹⁰ As noted in the Introduction, this analysis focuses on the direct effects of benefits-indexing reform on retirees in the context of heterogeneity in initial wealth and mortality risks, setting aside a number of other factors that matter for a more general approach to the topic of optimal benefits-indexing. In particular, I abstract from any distortionary effects that changes to the method of Social Security benefits-indexing may have on labor supply or saving decisions of households during their working lives.

In the model, there are I types of Social Security beneficiaries, indexed by $i \in \{1, 2, \dots, I\}$ and equally prevalent at the time of retirement, $t = 1$. Type indicates the level of lifetime income y_i , the level of non-annuitized wealth $A_{i,1}$ available at $t = 1$, and age-specific mortality risks $m_{i,t}$. A more general model would not impose a one-to-one link between net wealth, mortality risk, and lifetime income, but the theoretical and (especially) empirical challenges to the analysis are substantially reduced with this assumption. The probability of individual i being alive at age t is $\prod_{\tau=1}^t (1 - m_{i,\tau})$. As the use of private annuities in the United States is quite limited (see Brown et al, 2001), I assume annuitization of A_i is unavailable or unappealingly costly.

Once reaching retirement, each beneficiary receives streams of real-valued Social Security benefits denoted $\{B_{i,t}\}_{i,t}$ and (possibly zero) DB pension benefits denoted $\{P_{i,t}\}$ for type i at age t (note that all quantities in this paper's analysis are real, not nominal, unless otherwise noted). In the Status Quo policy, we will assume that this stream is constant in real terms, so that $B_{i,s}^{SQ} = B_{i,t}^{SQ}$ for all ages s, t . A reform to the method of benefits-indexing generates a stream of changes in benefits that I will denote $\{dB_{i,t}\}_{i,t}$. Note that I treat $B_{i,t}$ as an after-tax benefit, implicitly assuming that reform to benefits-indexing does not change the tax rates on retiree benefits.

Though in principle a reform could take a wide range of forms, in this paper we are especially interested in one class:

$$B_{i,t} = \lambda B_i^{SQ} (1 + \pi)^t \text{ for } t \in \{1, 2, \dots, T\}, \quad (1)$$

such that:

$$\sum_i \sum_t \beta^t \left(\prod (1 - m_{i,t}) \right) (B_{i,t} - B_i^{SQ}) = 0, \quad (2)$$

where β is the uniform discount factor in the economy (I consider heterogeneity in β in Section 5). This class of reforms scales the initial benefit level by the factor $\lambda \geq 0$ and grows that scaled benefit by the rate π each year, such that the total present value cost of benefit payments is the same in the Status Quo and reform policies. For example, a reform that increased the initial benefit level and then reduced the rate of growth in real benefits would have $\lambda > 1$ and $\pi < 0$.

Individuals solve a standard utility-maximization problem once they reach retirement. They

¹⁰I consider a relatively narrow set of reforms that deviate only slightly from the status quo policy, so it is natural to use this so-called "partial reform" approach of Guesnerie (1977), Feldstein (1976), and more recently Saez and Stantcheva (2014).

use their accumulated assets and their streams of Social Security and DB pension benefits to fund consumption in each period they are alive, and they obtain time-separable utility from that consumption. Note that there is no uncertainty in the utility they obtain from spending in the future.¹¹ Utility is zero when not alive, and there is no bequest motive (in Section 5 I show the results are robust to adding a bequest motive). They are subject to the (real-world) constraint that they cannot borrow against future Social Security or DB pension benefits. Formally, individual i solves:

$$\max_{\{c_{i,t}\}_t} E[U_i] = \sum_t \beta^t \left(\prod (1 - m_{i,t}) \right) u(c_{i,t})$$

subject to a constraint that (non-Social Security) net worth must be non-negative at all points during retirement:

$$A_{i,t} \geq 0, \text{ for all } t \in \{1, 2, \dots, T\},$$

where $A_{i,1}$ is given and

$$A_{i,t+1} = (A_{i,t} + P_{i,t} + B_{i,t} - c_{i,t})(1 + r).$$

where $(1 + r) = \beta^{-1}$ is the annual return (net of taxes) that an individual may earn on net wealth.

Note that in this model, were households able to fully annuitize their wealth, they would choose a constant consumption level throughout retirement. Without such full annuitization, mortality risk will cause the household's optimal consumption path to decline throughout retirement until reaching the level of annuity benefits (provided by Social Security and DB pensions, if applicable). After that point, the household will be dependent on these benefits to fund consumption.

The expected change in social welfare from reform $\{dB_{i,t}\}_{i,t}$ is evaluated as the weighted sum of the welfare values of the consumption changes it causes. In particular, social welfare is denoted W , so the change in social welfare from the stream of changes in benefits $\{dB_{i,t}\}_{i,t}$ is:

$$dW = \sum_i \sum_t \left(\prod (1 - m_{i,t}) \right) \beta^t \frac{dc_{i,t}}{\{dB_{i,t}\}_{i,t}} g_{i,t} \quad (3)$$

where $\beta^t dc_{i,t} / \{dB_{i,t}\}_{i,t}$ denotes the present value of the change in consumption by type i in year t in response to the change in policy, and $g_{i,t}$ is the marginal social welfare value of a present value unit of consumption for a beneficiary of type i in year t .

It is important to note that these $g_{i,t}$ parameters can take essentially any values, though Pareto efficiency would require them to be nonnegative. This flexibility enables us to use a wide variety of welfare criteria, including those inferred from public opinion, to evaluate policy reforms. An alternative formal approach would locate the welfare costs from the low lifetime utility of the shortest-lived, poorest retirees in their own utility functions, perhaps by having their utility be a highly concave function of total consumption in retirement or some other version of time non-

¹¹Medical expenditure shocks have been shown by many previous researchers to be important for retirees' decisions and welfare. Though not included in this paper, upward shocks to the marginal utility of spending at later ages would likely increase the appeal of the backloaded benefits streams.

separability.¹² Then, helping those retirees would be a matter of insurance, not redistribution (this logic is related to the justification Rawls offers for the maximin priority). It is far from clear that individuals have such preferences, however, so I take the approach that granting large weights to those worst-case outcomes is a normative decision by society, not a feature of individual preferences.

3 Simulated effects of benefits-indexing reform proposals

To simulate the effects of benefits-indexing reform, we need to specify functional forms and parameter values for the preceding section’s model, determine the values of the model’s key empirical inputs, and choose candidate reform policies.

3.1 Functional forms and parameter values

The per-period utility function takes the familiar constant relative risk aversion form:

$$u(c_{i,1}) = \frac{1}{1-\gamma} \left((c_{i,1})^{1-\gamma} - 1 \right), \quad (4)$$

where $\gamma = 3$ following SSK.¹³ I follow SSK in setting the annual discount factor $\beta = 0.96$ as well, and I assume that the return to saving $(1+r) = \beta^{-1}$.

3.2 Data on initial wealth, benefit levels, and mortality

To determine the key empirical inputs to the model, I use estimates drawn from the existing literature on Social Security. I divide the population of retiree households into deciles by lifetime income, so $I = 10$ and each type $i=\{1, 2, \dots, 10\}$ corresponds to a lifetime income decile. The use of ten types is made possible by recent empirical work estimating household wealth, benefits, and mortality data at that level of disaggregation. Some of that data is not available by gender, so I treat households as the unit of analysis throughout.

For the initial wealth and benefit levels of retirees I rely on SSK, which is a careful and uniquely detailed source of these data, in that no other source of which I am aware provides both median overall (non-Social Security) net worth and median (present value) Social Security and DB pension wealth data by lifetime income decile. This level of detail is especially important for capturing the right end of the seesaw: for example, data that divides the population into quintiles, or that groups households by point-in-time income rather than lifetime income, can obscure the difficult position in which the lowest decile of retirees appear to find themselves. To infer annual benefit amounts for both Social Security and DB pensions, I use average mortality rates (for men) in the

¹²I thank Bob Hall for suggesting this discussion.

¹³This value for γ is toward the upper end of typical ranges for this parameter, which measures the degree of risk-aversion of the individual. Though a high value for γ may be appropriate if retirees are generally more risk-averse than the average person, I have also run the analysis assuming $\gamma = 1.5$. All qualitative results described in the baseline case hold there as well, though the size of the welfare gains generated by the Hybrid Progressive Reform are smaller. Intuitively, with less concave utility from consumption that policy’s redistribution is less valuable in terms of social welfare.

United States and the same real interest rate r as in SSK to calculate the constant real benefit amounts that yields SSK's reported wealth figures by lifetime earnings decile (in their Table 2). Of course, SSK's data are not perfectly designed for my purposes. Most obviously, the average age of their sample is 55.7 years, several years prior to typical retirement age. Ideally, we would have data at age 62 or 65. While it is possible that the last few years prior to retirement differentially affect retiree households by income decile, a comparison of the SSK data with calculations by Love, Palumbo, and Smith for (point-in-time) income quintiles suggests that this is not likely to be a serious concern. A different concern is that the SSK data are relatively old, focused on the 1992 HRS wave. Gale, Scholz, and Seshadri (2009) attempt to address this concern and show that their core findings are largely unaffected by considering later cohorts (though they do not reproduce the estimates needed for this paper for later waves). Finally, recent work on the progressivity of the overall OASDI program has noted that, in the words of CBO (2006), "the progressivity of Social Security is driven mainly by disabled-worker and auxiliary [survivor] benefits.". While this paper focuses on the retirement portion of benefits, for which SSK's estimates are well-suited, we may be interested in the implications of indexing reform that applies to disability benefits as well. (Note that SSK implicitly includes disability benefits after retirement age has been reached, as disability benefits are automatically converted to retirement benefits at that point).

Table 1 shows the median (non-Social Security) net worth A_i , annualized DB pension benefit $P_{i,t}$, annual benefit level under the Status Quo Social Security system B_i^{SQ} , and present value Social Security wealth, all in 2005 dollars and by household lifetime earnings decile. To be clear, all of these estimates are from SSK other than B_i^{SQ} , which is inferred from SSK's Social Security wealth estimates.

Table 1: Net wealth and initial benefits (\$2005)				
Lifetime earnings decile, type i	A_i	$P_{i,t}$	B_i^{SQ}	SS wealth
1 (lowest)	6,938	0	3,086	36,357
2	35,383	0	4,965	58,499
3	60,336	0	6,813	80,263
4	104,069	1,259	9,122	107,471
5	124,883	2,504	11,180	131,722
6	172,543	3,872	14,016	165,138
7	178,416	4,756	15,717	185,175
8	231,727	6,058	17,831	210,076
9	313,594	8,883	19,273	227,063
10 (highest)	545,321	10,779	23,273	281,206

Standing out from Table 1 are the small initial net wealth holdings A_i of the lowest deciles of the lifetime-earnings distribution.¹⁴

¹⁴Though not included in the baseline analysis, in Section 5 I show the robustness of the main results to including in the simulations a means-tested (by income) transfer payable to all individuals, modeled on the Supplemental Security

To estimate mortality rates by income decile, I rely on recent work by Bosworth and Burke (2014), who calculate relative mortality rates by lifetime-income quintiles and gender in the HRS for the age range 50-74. From the Social Security Administration’s current period life table, I have average mortality rates by age and gender. Combining these data sources, I adjust the SSA’s overall average mortality rates at age 65 by a vector of scalars to approximately match BB’s mortality patterns by income quintiles. BB also report relative mortality rates for the age range 75+, indicating some convergence of mortality rates across income quintiles as retirees age. To roughly match this convergence, I calculate mortality rates after age 65 so that each decile’s mortality rate approaches linearly the average gender-specific mortality rate by age 119, the Social Security Administration life table’s terminal age (the results change very little if I assume no such convergence in relative mortality). The following table shows the resulting one-year mortality rates for each decile, by gender, at 10-year increments from age 65 through 95.

Table 2: One-year calculated mortality rates (in percent)								
Lifetime earnings decile, type i	Men’s age				Women’s age			
	65	75	85	95	65	75	85	95
1 (lowest)	2.2	4.9	12.3	30.4	1.5	3.6	9.6	25.4
2	2.0	4.6	11.8	29.4	1.4	3.4	9.2	24.6
3	1.9	4.3	11.2	28.4	1.2	3.0	8.4	23.0
4	1.9	4.3	11.2	28.4	1.1	2.9	8.0	22.2
5	1.9	4.3	11.2	28.4	1.1	2.8	7.8	21.8
6	1.6	3.8	10.2	26.4	1.0	2.7	7.6	21.4
7	1.5	3.6	9.6	25.4	1.0	2.7	7.6	21.4
8	1.3	3.3	9.1	24.4	1.0	2.7	7.6	21.4
9	1.2	3.0	8.5	23.5	0.9	2.3	6.8	19.8
10 (highest)	1.1	2.8	8.0	22.5	0.5	1.6	5.2	16.6
Social Security Administration average	1.6	3.8	10.2	26.4	1.0	2.7	7.6	21.4

The calculated mortality rates in Table 2 show the dramatic negative relationship between lifetime earnings and mortality rates, especially early in retirement. These rates roughly match existing related estimates along a number of dimensions. For example, shown at the bottom of Table 2 are the Social Security Administration’s official average mortality rates for each gender at each age; these rates match the calculated values for the sixth decile in all cases.

Income program of the United States.

3.3 Reform proposals

I consider three reform proposals, two informed by recent experience with chained CPI-U and the experimental CPI-E series and a third based on President Obama's proposal that is a hybrid of the first two. Figure 2 shows historical data for the December values of three price indices, the chained CPI-U, CPI-E, and the currently-used CPI-W, with each index set equal to 100 in December, 1990. The chained CPI-U was first reported by the BLS in 2000, while the BLS has backcast the experimental CPI-E through 1983.

Figure 2: Historical data on three price indices.

The figure makes clear that chained-CPI-U has risen more slowly than CPI-W over the 2000-2013 period, though year-to-year changes are not always smaller. By the end of the fourteen available years, chained CPI-U is approximately 4 percent lower than CPI-W, for an average annual gap of -0.27 percentage points.¹⁵ Over the same period, CPI-E has been nearly identical to CPI-W. That stands in stark contrast to its more rapid growth from 1983 through 2000, when it exceeded CPI-W by an average annual rate of 0.37 percentage points.

The first reform policy—"Backloaded Reform"—is designed to mimic the CPI-E's behavior in the 1983-2000 era, having benefits grow at a faster rate than the Status Quo. Specifically, I set $\pi = 0.0037$ for Backloaded Reform, implying a steeper path of benefits and a smaller initial benefit than in the Status Quo.

The second reform policy—"Frontloaded Reform"—is designed to mimic the chained CPI-U's be-

¹⁵Thanks to Alan Viard for noting an error in my calculation of this value for π in an earlier draft of the paper.

havior since its origination, having benefits grow at a slower rate than the Status Quo. Specifically, I set $\pi = -0.0027$ for Frontloaded Reform, implying a flatter path of benefits and a larger initial benefit than in the Status Quo. This reform has received much attention in public debates, as the chained price index is generally seen to address upward bias in the traditional CPI. It may be of interest to note that The National Commission on Fiscal Responsibility and Reform (White House, 2010), commonly known as the Bowles-Simpson commission, recommended a shift to chained CPI-U for Social Security benefits-indexing.

The third reform policy—"Hybrid Progressive Reform"—is designed to match the proposal made by President Obama. In that proposal, the chained CPI-U would be used to index benefits, but so-called "benefit enhancements" would phase in at age 75 and 95, each eventually raising benefits by five percent of the average retiree's benefit over a ten-year phase-in period. This reform combines features of the two others, but it also includes a substantial increase in the progressivity of Social Security benefits. The source of this increase is the use of the average retiree's benefit, rather than each individual retiree's benefit, in the calculation of the benefit enhancement. As the average benefit is approximately four times greater than the lowest-decile's benefit and half as large as the top decile's (see Table 2), the first ten-year benefit enhancement would effectively raise benefits by 20 percent for the lowest decile retiree and 2.5 percent for the highest decile retiree. The President's proposal thereby illustrates how the debate over indexing is closely linked to the broader debate over progressivity. Related, note that this proposal's redistributive impacts make it more likely to affect labor supply during households' working lives—effects from which this paper abstracts throughout (see Jeffrey B. Liebman, Erzo F.P. Luttmer, and David G. Seif, 2009 for evidence on how labor supply responds to Social Security benefits changes).

Table 3 summarizes these proposals and shows the equilibrium value of λ that satisfies the government budget constraint when we simulate the economy's response to each policy. Technically, to obtain these values we set π for each reform policy and have the simulation guess a value for λ . All individuals maximize their utilities given these parameters and the data on benefits, net wealth, and mortality. The simulation searches for a value of λ that satisfies the government's budget constraint as shown in (2).

Table 3: Reform policy parameters

Policy	π	λ
Status Quo	0.0000	1.000
Backloaded Reform (CPI-E)	0.0037	0.970
Frontloaded Reform (chained CPI-U)	-0.0027	1.022
Hybrid Progressive Reform	$\left\{ \begin{array}{l} -0.0027 + \text{benefit enhancements} \\ \text{for ages 76-85 and 95+} \end{array} \right.$	1.011

Note how the Hybrid Progressive Reform leaves the starting value of benefits closest to, and above, the Status Quo value, a feature that may be relevant for political feasibility.

3.4 Simulated effects of reform

Now we turn to the effects of these reform policies.

I begin by showing the policies' effects on real benefit payments in Figure 3. The four subplots of Figure 3 show results for the lowest and second-lowest income deciles, the fifth income decile, and the top income decile. In each subplot, I show the benefit paths under the Status Quo and three reform policies at each age.

Figure 3: Annual benefits paths under Status Quo (solid line), Backloaded Reform (dashed line), Frontloaded Reform (dotted line), and Hybrid Progressive Reform (dash-dot line), in thousands of real 2005 dollars, for four of ten household lifetime-income deciles.

The subplots in Figure 3 for the Frontloaded Reform (dotted lines) and Backloaded Reform (dashed lines) closely resemble Figure 1, of course. In fact, because we apply the same λ, π pair to all benefits paths in each reform, the figure shows that the effects of reform are quite similar across income deciles. It is also apparent how the Hybrid Progressive Reform (dash-dotted lines) takes on a zigzag shape that combines the two other reforms—providing higher benefits at the start of retirement than the backloaded reform and higher benefits at the end of retirement than frontloaded reform. In

fact, compared to the Status Quo (solid lines) policy, it achieves both higher initial benefits and higher final benefits for low income retirees, reflecting its substantial redistribution of benefits from higher to lower deciles.

The effects on consumption paths chosen by retirees in the model are much more variable across income deciles. In the next figure, we plot consumption under the Status Quo and these reforms following the same structure as above:

Figure 4: Annual consumption paths under Status Quo (solid line), Backloaded Reform (dashed line), Frontloaded Reform (dotted line), and Hybrid Progressive Reform (dash-dot line), in thousands of real 2005 dollars, for four of ten household lifetime-income deciles.

Figure 4 shows the pattern of declining consumption until private assets are exhausted, as discussed in Section 1, which obtains due to the lack of private annuitization. A similar pattern is found in, for example, Hurd and Rohwedder (2008).

To interpret Figure 4, it may be helpful to focus first on the comparison of Backloaded Reform to Frontloaded Reform, postponing a consideration of the Hybrid Progressive Reform until later. For these two "simple" reforms, two prominent features jump out from Figure 4. First, while a household's chosen paths of consumption are hardly distinguishable across benefits-indexing meth-

ods early in retirement, there are sharp divergences when they exhaust their non-Social Security wealth. Remarkably, all deciles experience a substantially higher median path of consumption in these later years under the Backloaded Reform than under the Frontloaded Reform, despite the latter's inability to generate substantial increases in consumption earlier in retirement. Second, the ages at which households exhaust their private assets and become dependent on Social Security benefits rise substantially with lifetime income.

One potentially puzzling nuance related to the first of these features is that for all but the bottom decile consumption is in fact slightly greater at all ages under the Backloaded reform than under the Frontloaded reform. The key intuition for this result is that the Frontloaded Reform provides less insurance against longevity risk than the Backloaded Reform. Therefore, households choose to consume less of their private assets in order to self-insure against longevity risk, offsetting the mechanical increase in benefits at early ages that the Frontloaded Reform provides.

For the lowest decile households, however, consumption is greater at early ages under the Frontloaded Reform than under the Backloaded Reform. Two factors explain this exception. First, these households have little wealth and high mortality rates. Thus, the effective annuitization provided by the backloaded path enables only small increases in consumption out of their private assets early in retirement, in contrast to higher-decile retirees. A second, more subtle reason is that the frontloading that comes from using a slower-growing price index is not actuarially fair. To see why, note that the Frontloaded Reform allocates that total value of benefits through a *uniform* proportional adjustment to Status Quo benefits. Thus, it causes a redistribution of resources from low-mortality to high-mortality retirees, increasing the consumption of lower-income retirees. Note that this factor provides a second reason why consumption paths do not rise for higher-decile retirees under the Frontloaded Reform.

As for the second prominent feature of Figure 4, consistent with prior research we find that most retirees exhaust their private assets only late into retirement, while a substantial share of lower-income retirees depend on Social Security benefits throughout much of retirement. Overall, only 18 percent of individuals exhaust their non-Social Security, non-DB assets in this simulation, also consistent with prior research. For example, Love, Palumbo, and Smith (2008) calculate "annualized comprehensive wealth," which is the value of a retiree's total resources divided by his or her remaining life expectancy at any given age. In their research they find that "in (real) dollar terms, the median household's...real annualized wealth actually tends to rise with age over retirement." In our simulations, we find consistent patterns, with annualized wealth calculated in this way greater fifteen years into retirement than at the start and positive until at least age 90 for retirees in the third income decile or higher. At the same time, lower income decile retirees exhaust their non-Social Security wealth much earlier. For the lowest decile, in these simulations non-Social Security wealth is nearly exhausted fifteen years into retirement and is less than the level of annual benefits only eight years in.

Figure 4 appears to make a strong case in favor of Backloaded Reform relative to Frontloaded Reform, and that case looks all the stronger if we convert these results on consumption paths

into changes to expected utility during retirement. All deciles—even the lowest—prefer Backloaded Reform to the Status Quo and prefer the Status Quo to Frontloaded Reform in expected utility terms at retirement.

The seeming dominance of Backloaded Reform is not airtight, however, because it generates losses for the poorest, shortest-lived retirees relative to Frontloaded Reform or the Status Quo. To examine this feature of the reforms, we calculate each individual's change in "realized retirement utility": the change in total utility during retirement for an individual from decile i who lives t years under each policy. We then convert these changes, which are in units of utility, into consumption equivalents by calculating the percentage change in the total present value of consumption during retirement that, when multiplied by the marginal utility of consumption in the last year of life for a retiree, yields the given change in realized retirement utility. Figure 5 shows the results.

Figure 5: Differences in realized retirement utility levels, converted into consumption equivalents, for four household lifetime-earnings deciles under three reforms from the Status Quo: Backloaded Reform (dashed line), Frontloaded Reform (dotted line), and Hybrid Progressive Reform (dash-dot line). Realized retirement utility for type i and age t is the total utility obtained in retirement for a retired household of income decile i who lives to the age t .

Figure 5 shows the extent to which the seesaw apparent in benefit paths translates into a similar shape in realized retirement utilities.

For all but the lowest decile of households, the Backloaded Reform generates higher realized utility than the Frontloaded Reform no matter the age of death, but especially at later ages when its ability to provide longevity insurance has its greatest value. The same preference holds for the poorest households who live beyond approximately age 82. In other words, most retirees sit squarely on the left end of the seesaw when it comes to these direct effects of benefits-indexing reform, preferring a steeper path of benefits with a lower starting point. However, the poorest households who die earlier in retirement prefer Frontloaded Reform over Backloaded Reform, as shown in the top left subplot of Figure 5. That is, they sit on the right end of the seesaw and prefer a flatter benefit profile.

Now, we turn to a consideration of the Hybrid Progressive Reform in Figures 4 and 5. The Hybrid Progressive Reform generates very different consumption effects across deciles: while its path lies below the Status Quo at all ages for the top decile retiree, it exceeds all other paths at all ages for the bottom decile retiree. These differences reflect both its combination of the two other reforms and its extensive redistribution of benefits, as it can achieve wide-ranging improvements for low-income retirees at the cost of a general decrease in consumption for higher-income retirees. As would be expected, these implications for consumption translate into gains in realized retirement utility for every retiree in the bottom two (three, in fact) lifetime income deciles relative to the Status Quo and losses for every retiree in the top five deciles.

In the next section, I explore how we might convert these heterogeneous results across retiree households into net welfare implications.

4 Welfare criteria and net welfare implications of the direct effects of reform

As summarized formally in expression (3), in this paper I calculate the net welfare effects of reform by multiplying discounted changes in consumption from the Status Quo by two things: the population proportion of individuals who survive to enjoy that consumption, and a non-negative welfare weight $g_{i,t}$. The weight $g_{i,t}$ measures the value society puts on a marginal increase in consumption for a household of type i at age t , relative to all other households. Because each reform has some retirees who gain and some who lose, their net welfare implications depend on how those welfare weights vary across the population of retirees.

4.1 Two familiar principles

The conventional approach to normative evaluation in much of applied public economics research is to rely on well-known principles with roots in political philosophy, the two most commonly-used being the simple-sum utilitarian criterion and the so-called "Rawlsian," or maximin criterion.¹⁶

¹⁶See Weinzierl (2014) for a critique of this conventional choice in the optimal tax literature.

These two criteria have especially clear implications for the welfare weights $g_{i,t}$, and it turns out that the choice between them illustrates well the policymaking challenge that benefits-indexing reform presents.

Under the simple-sum utilitarian criterion, the change in welfare dW is the sum of the experienced annual utility changes across all individuals. In the language of the general formula (3) above, this option sets $g_{i,t} = u'(c_{i,t})$, so that society puts greater weight on the annual consumption changes of individuals with lower consumption levels (and thus higher marginal utilities of consumption). Figure 6 shows the $g_{i,t}$ for the same four deciles as in previous figures, given the consumption levels in the simulated Status Quo economy from the previous section and scaled so that the maximum $g_{i,t}$ equals one.

Figure 6: Marginal welfare weights $g_{i,t}$ for four of the ten deciles of retirees under the Utilitarian criterion, where $g_{i,t}$ equals the marginal utility from consumption at age t for a retiree of type i .

The vertical axes in Figure 6 all have the same scale, making it clear that the utilitarian criterion puts much greater weight on consumption changes for lower-income retirees than others, and in particular on consumption changes at advanced ages for those households, when their consumption

levels fall to the level of their Social Security benefits. Note that these differences are especially large given our assumption (following SSK) that $\gamma = 3$, a value that is toward the upper end of conventional ranges for that parameter. If we use $\gamma = 1.5$, the marginal weights on the fifth decile rise to around 0.20.

The Rawlsian criterion prioritizes the well-being of the worst-off member of society. It therefore sets $g_{i,t} = 1.00$ on the consumption change of the household in the lowest income decile who lives only one year in retirement—i.e., the household with the lowest overall utility in retirement—and $g_{i,t} = 0.00$ on all other consumption changes.¹⁷

The net welfare implications of each benefits-indexing alternative under the Rawlsian criterion are immediately apparent from examining Figure 5 for $i = 1$, $t = 1$, where we find the effects of each path on the total utility in retirement for the household from decile 1 who dies in the first year of retirement. From that figure, it is clear that the Rawlsian criterion would endorse the Frontloaded Reform over the Status Quo and both over the Backloaded Reform.

The net welfare implications under the utilitarian case are not so immediately clear, as that criterion puts substantial weight not only on the same worst-off household that drove the Rawlsian results but also on poor retirees who live long into retirement and spend down their private assets. To calculate the change in social welfare under the utilitarian criterion, I multiply the relevant g_{it} values by discounted consumption changes adjusted for survivorship and take the sum, as in expression (3). The utilitarian criterion turns out to endorse the Backloaded Reform over the Status Quo and both over the Frontloaded Reform: in other words, exactly the opposite order as under the Rawlsian criterion.

As this result implies, the extent of backloading most preferred under the utilitarian criterion may be substantially larger than that implied by a switch to the CPI-E. If we solve for the utilitarian-optimal π and λ (i.e., for the class of reforms formalized in expression 1), we find that $\pi = 0.012$ (about three times the rate increase from the switch to the CPI-E) and $\lambda = 0.91$ maximize total expected utility of all retirees at retirement (we cannot use the marginal welfare weights approach in this case because the changes are too large).

The contrast between the Rawlsian and utilitarian rankings suggests that there may exist a mixture of the two that would be consistent with the Status Quo policy being chosen as optimal. In fact, if we put a weight of 0.91 on the Rawlsian weights and 0.09 on the utilitarian weights, the planner prefers the Status Quo policy to both the backloaded and frontloaded reforms.¹⁸ The large implied weight on the Rawlsian weights in the Status Quo makes sense in light of the finding that backloading is preferred by most agents. That is, for the Status Quo policy to be optimal the

¹⁷The Rawlsian priority as modeled here is an extreme case of a social objective in which weights on individuals decrease in their lifetime, rather than annual, utility—see Pestieau and Ponthiere (2012) and the comments on this paper by Aleh Tsyvinski for discussions. A related pattern for MSWWs (in this paper’s framework) weights consumption changes by the retiree’s total utility in retirement raised to a negative exponent (e.g., as if we were taking the marginal utility of total consumption in retirement). Such weights can generate a preference for frontloading if the curvature over total retirement utility is steep enough, because the weights in that case approach Rawlsian weights. For less steep curvature, backloading is still preferred.

¹⁸Thanks to Aleh Tsyvinski for suggesting this analysis.

planner must have a large weight on the worst-off retiree.

The Hybrid Progressive Reform, however, is the most-preferred of these policies under both the utilitarian and Rawlsian criteria. By combining the two other reforms' positive implications for the poorest retirees, the Hybrid Progressive Reform outperforms them both. That is, the frontloading in the early years of the proposal benefits the worst-off retirees, increasing its appeal under the Rawlsian criterion, while the backloading through the benefit enhancements brings utilitarian gains. Both of these benefits are substantially augmented by the redistribution pursued under this reform, while the corresponding negative effects on the top half of retirees are given, under both criteria, very little weight. We can quantify the potential gain from this reform under the utilitarian criterion by calculating the uniform proportional increase in consumption across all retiree types and ages that would generate the same increase in social welfare as does this reform over the Status Quo. That "consumption-equivalent" gain is 0.75 percent of consumption for retirees in the case of the Hybrid Progressive Reform. For comparison, the Backloaded Reform generates a gain of 0.12 percent of consumption for retirees, and the Frontloaded Reform generates a slightly smaller size loss. Recall that all of these calculations abstract from a number of indirect effects of benefits-indexing reform on households, including on their labor supply and saving decisions during their working lives.

Of course, these conventional criteria may not match true social preferences, and to explore this possibility we now turn to an attempt to empirically study society's normative priorities for Social Security.

4.2 Evidence on prevailing normative priorities for Social Security

This section presents some novel survey evidence on the American public's priorities for Social Security generated through Amazon's Mechanical Turk (M-Turk) interface. The way in which this survey elicits marginal social welfare weights $g_{i,t}$ may prove useful to other researchers interested in using a positive approach to normative questions.¹⁹

The survey was completed in August, 2014 by 150 members of the Amazon Mechanical Turk worker population from the United States who demonstrated good past performance on tasks. Respondents had up to 15 minutes to complete the survey, and they were asked to enter their M-Turk identification number as well as a completion code at the end of the survey for verification purposes. The respondents completed the survey in less than seven minutes on average. They were paid \$2.50 for the task, for an average hourly rate of \$23.00.

Mechanical Turk is of course an imperfect tool: e.g., it is not a representative sample of Americans. That said, it has proven to be a popular alternative to surveys costing orders of magnitude more (with their own problems with representativeness), and analysis by subgroup can provide some reassurance on the robustness of the results to sample composition. Horton, Rand, and Zeckhauser (2010) study the use of Mechanical Turk, and find: "Online experiments, we show, can be

¹⁹A growing literature in public economic theory has considered using positive evidence on prevailing normative priorities, rather than exogenously specified normative criteria, to inform evaluations of policy. See Gaertner and Schokkaert (2012) for an overview of "empirical social choice" research. Weinzierl (2014) and Saez and Stantcheva (2014) are recent examples applied to tax policy.

just as valid— both internally and externally— as laboratory and field experiments, while often requiring far less money and time to design and conduct."

The survey has three parts. The first part tests whether respondents understand and can perform simple calculations related to the concepts of percentages, averages, and life expectancy. The third part of the survey asks respondents to self-report their political views and demographic traits (age, gender, education, and economic status).

I gather data on normative priorities for Social Security in the second part of the survey. Respondents are given a one-sentence (official) description of Social Security, told that policymakers must decide (among other things) how much in benefits to pay out to different retirees, and then told they will be asked a couple of questions to get their "opinions on how policymakers should make this choice." They are then shown the following screen:

Please consider the following situation.

Suppose the Social Security system has raised some extra tax revenue that must be allocated among the three retirees described below. Please assume that these retirees worked equally hard during their working years and saved equally well for retirement.

Please rank these retirees in terms of who ought to receive an increase in his Social Security benefits, where #1 is the retiree you think ought to be the first to receive an increase and #3 is the retiree you think ought to be the last. (Drag the retiree descriptions to change their ranks).

- John is 65 years old. He has just retired and is expected to live to age 70. While he was working, his income was in the bottom 10% of incomes, and he currently has about \$10,000 per year to spend.
- Robert is 75 years old. He retired at age 65, when he was expected to live to age 83. While he was working, his income was in the middle 10% of incomes, and he currently has about \$25,000 per year to spend.
- William is 90 years old. He retired at age 65, when he was expected to live to age 81. While he was working, his income was in the bottom 10% of incomes, and he currently has about \$9,000 per year to spend.

The three retirees in this first question represent three important points in the joint age-income distribution.²⁰ In particular, John represents a very low income individual with a short life expectancy, the point given particular priority by the Rawlsian criterion. William is also very low income but has lived a long life, giving him a greater overall utility level than John but leaving him with a smaller current (according to the survey) level of consumption. Thus a Utilitarian would allocate more to William, while a Rawlsian would allocate more to John. Finally, Robert is a middle-income individual approaching his expected lifespan. He is much better off than either of the other retirees and provides a simple way for us to gauge how quickly marginal welfare weights decline with well-being.²¹

This first question is largely intended to get respondents to engage with the descriptions of

²⁰These names were the most popular names, according to the Social Security Administration's names database, for boys born in 1949 and 1924.

²¹The consumption levels indicated in the survey implicitly include other transfers for John and William, such as SSI and SNAP. See Section 5 for a discussion of how these programs relate to this paper's analysis of Social Security.

these retirees. Nevertheless, the responses may be of interest. William is rated first by 62% the respondents, John by 29%, and Robert by 9%. The preference for William directly casts some doubt on the possibility that a Rawlsian criterion will emerge from the survey evidence.

The key questions for this paper's purposes comes next, when respondents are shown a series of screens starting with one like the following, tailored according to which retiree the respondent ranked last in the previous question. In this screen, the respondent ranked Robert last in that question.

On the next several pages, we'll ask for your opinion on some specific options for changing these retirees' benefits.

When making your choices, please imagine that you are a policymaker trying to choose what is best. Ignore any effects these options might have on the rest of the economy, and focus on the effect each option has on the retirees.

For your reference, we'll copy the descriptions of the retirees on each page.

- John is 65 years old. He has just retired and is expected to live to age 70. While he was working, his income was in the bottom 10% of incomes, and he currently has about \$10,000 per year to spend.
- Robert is 75 years old. He retired at age 65, when he was expected to live to age 83. When he was working, his income was in the middle 10% of incomes, and he currently has about \$25,000 per year to spend.
- William is 90 years old. He retired at age 65, when he was expected to live to age 81. While he was working, his income was in the bottom 10% of incomes, and he currently has about \$9,000 per year to spend.

Which of the following would you prefer?

- ☐ Increasing Robert's benefit by \$100
- ☐ Increasing John's benefit by \$100

If the respondent chooses Robert over John in this question, he or she is reminded (by the computer) that Robert was ranked last in the earlier question, and he or she is asked to make the choices consistent. Then, the following choice appears:

Which of the following would you prefer?

- ☐ Increasing Robert's benefit by \$100
- ☐ Increasing John's benefit by \$75

If the respondent chooses John over Robert, he or she then faces the same choice but with the increase for John at \$50 and then \$25. After that, or whenever the respondent chooses Robert over John, he or she then faces a similar set of choices between benefits increases for Robert and William.

These series of questions are designed to allow the direct inference of marginal welfare weights.²² To see how, suppose a respondent ranks Robert last and (implicitly) assigns marginal value g_{Robert}

²² An earlier version of the survey used sliders to elicit the same information. Though preferable in many ways, the slider interface appeared to confuse respondents, who often implicitly assigned lower weights to retirees they preferred.

to Robert's consumption. Then, an \$100 increase in Robert's benefits provides a benefit to the respondent (acting as a policymaker) of $100g_{Robert}$. The respondent is then asked to choose between this gain and alternative gains. Suppose the respondent chooses the \$50 increase for William (but not the \$25 increase) over the \$100 increase for Robert. Then, we can infer that $100g_{Robert} > 25g_{William}$ and $100g_{Robert} < 50g_{William}$ implying that $g_{Robert}/g_{William} \in [0.25, 0.5]$. Similarly, we can calculate a range for g_{Robert}/g_{John} for each respondent, indicating the profile of relative welfare weights across these retirees.

Components of these questions are designed to counteract some potential confounding influences on the respondents. I ask respondents to "ignore any effects these options might have on the rest of the economy, and focus on the effect each option has on the corresponding retiree." This request is intended in particular to minimize the extent to which respondents consider the efficiency costs of raising different amounts of extra revenue for the benefits increases. I also ask them to "imagine that you are a policymaker" in the hopes that it will cause the respondents to take a considered, objective perspective. Inconsistent answers across the ranking question and the series of choices cause error messages to appear, preventing the respondent from making errors in interpreting the questions. Finally, the wide range of potential relative valuations implied by the choices (from 1.3 to 4.0 in each case) is intended to reduce concerns that respondents would default to equality and thereby imply smaller differences between g weights than is accurate, as might be natural in other designs (such as splitting an amount between the retirees).

A number of potential risks remain with survey evidence of this kind. One risk is that respondents may not be accustomed to thinking about these policy choices in terms of indifference points, which seem natural to most economists but which reverse the intuitive idea that the respondent would like to grant his or her preferred retirees greater increases, not smaller ones, than his or her least preferred retiree. Of course, more general concerns about how the questions are framed and whether the survey primes respondents toward any particular outcome also apply to this specific survey.

The results of the survey for the relative weight on Robert versus John and William are consistent with those weights implied by the utilitarian criterion. The median choices across all respondents imply a range of values for both $g_{Robert}/g_{William}$ and g_{Robert}/g_{John} of $[0.00, 0.25]$. That is, these median responses indicate a very small value for the welfare weight on Robert relative to both John and William, consistent with Figure 6 that shows a negligible weight on Robert under the utilitarian criterion. Of course, it is possible that the true relative weight put on Robert lies closer to 0.25 than to zero (which the survey cannot pin down). It is also possible that respondents systematically took the mental shortcut provided by choosing that John or William receive an increase—no matter how small—rather than Robert, which would bias our estimate of the relative value of g_{Robert} toward zero. Of the 116 respondents who ranked Robert last in the first question, approximately three-quarters (85) chose the benefit increase for both John and William in all cases.

At the same time, the results of the survey for the relative weights on John and William are not consistent with those implied by either the utilitarian criterion or the Rawlsian criterion on their

own. Specifically, the median choice across all respondents implies that $g_{John}/g_{William} = 1.00$, so that respondents put similar value on benefits increases for John and William, contrary to the utilitarian preference for William (which Figure 6 suggests would approximately set $g_{John}/g_{William} = 0.5$) and the Rawlsian extreme preference for John. In fact, the mean choice among those who ranked Robert last, which is in general too sensitive to outliers to be a useful measure of preferences in this survey, implies that $g_{John}/g_{William} = 1.05$ with a standard error of 0.05 (the means are very large for those who ranked William last and very small for those who ranked John last).

Both of these sets of results hold across virtually all subgroups. They hold for respondents aged 18-25, 26-40, 41-64, and 65+; for both male and female respondents, and for respondents who place themselves on the political left, middle, and right. The only exceptions are across race and household incomes, where for (the small number of) black and high-income respondents the relative weight on Robert is in the range $[0.25, 0.5]$.

One possibility suggested by these results is that respondents' moral reasoning reflects a mixture of these two standard criteria. Such a mix can easily generate $g_{i,t}$ values for $t = 1$ and $t = 35$ for the lowest income decile $i = 1$ that are very similar, in keeping with the survey evidence on John and William. The same mix yields extremely small values for the welfare weight on the "average" retiree (i.e., $i = 5$ and $T = 10$), in keeping with the survey evidence on Robert.

When we apply these weights to the reform options, the rankings and consumption-equivalent welfare gains and losses are the same for all reform proposals as under the utilitarian criterion. The costs of the Backloaded Reform for the worst-off retirees are not enough to offset the gains it generates for the poor retirees who outlive their private assets, so backloaded benefits as under a switch to the CPI-E are preferred to frontloaded benefits as under a switch to the chained CPI-U. The Hybrid Progressive Reform, by combining the Backloaded Reform's appeal to long-lived poor retirees with the Frontloaded Reform's appeal to the short-lived poor retirees, dominates the policy ranking under this criterion, reflecting survey respondents' low concern for consumption decreases among better-off retirees.

5 Extensions to the baseline analysis

In this section, I extend the analysis above along a number of dimensions. Though each extension modifies the baseline results somewhat, the basic seesaw metaphor, the tradeoff between the effects on the vast majority of retiree households and the worst-off, and the likely net welfare impacts of the direct effects of benefits-indexing reform all continue to apply. To simplify the discussion, I focus on the effects of these extensions on the Backloaded and Frontloaded Reform policies.

5.1 Myopic households

The fully rational, foresighted utility maximizing household modeled above may not represent all, or even most, retirees' consumption and saving behavior. In particular, though the evidence reviewed in Section 1 suggests that myopia is not an issue for most retirees outside of the lowest income

deciles, we might be interested in the effects of assuming that some share of retirees have difficulty delaying consumption early in retirement. Feldstein (1985, 1987) made clear the importance of myopic households to determining the optimal path of benefits.

To gauge the effects of this myopia, I consider a model in which retirees from the bottom through sixth income deciles choose consumption at each age using a lower discount factor $\hat{\beta}$ than the true discount factor β upon which their utility depends, as in Table 6:

Table 6: Impatience by income decile										
	Bottom	2	3	4	5	6	7	8	9	Top
$\hat{\beta}$	0.90	0.91	0.92	0.93	0.94	0.95	0.96	0.96	0.96	0.96
β	0.96	0.96	0.96	0.96	0.96	0.96	0.96	0.96	0.96	0.96

That is, individuals use $\hat{\beta}$ to choose their consumption paths, but the utilitarian evaluation of their overall retirement utility uses β (as would, it is assumed, the individuals if they were able to adopt a disinterested perspective). Of course this is only a crude version of this extension to the baseline model, for example a more sophisticated model would have heterogeneity in impatience within deciles.

The results are similar to the baseline results, but more extreme. That is, the gains from Backloaded Reform are larger for the majority of households that value its insurance against longevity risk—a feature even more beneficial in a setting where households have difficulty saving. For example, households in the second decile ($i = 2$) who survive thirty years into retirement see a 30 percent larger gain (in utility terms) from Backloaded Reform in this setting than in the baseline. At the same time, the gains from frontloading are even higher for those retirees with short ex post lives and few initial assets. For example, the shortest-lived household in the bottom decile has more than twice the gain from frontloading in this setting as in the baseline. Moreover, the shorter-lived half of households in the second income decile now gain from Frontloaded Reform (whereas they lost in the baseline case), as their impatience causes them to benefit more from the higher initial benefits and their limited assets make the appeal of backloading small. Their (impatient) consumption of an even higher share of the frontloaded benefits means that, when they (ex post) do not survive later into retirement, their realized utility during retirement was even higher than in the patient case. On balance, under the utilitarian criterion the increase in the gains to the majority of households outweighs the increase in the losses to a few, such that the net welfare impacts of the direct effects on retirees are more positive for backloading and more negative for frontloading in this model than in the baseline case. Similarly, the difference between the two policies under the Rawlsian criterion, which ranks the frontloading policy ahead of the backloading policy, also grows.

5.2 Budget non-neutrality

Thus far I have imposed budget neutrality to disentangle the effects of changing the shape of the time-path of benefits from the effects of changing the expected present-value of those benefits. Of

course, much of the energy in the policy debate over benefits-indexing is due to the likelihood that choosing a more slowly-growing price index, such as the chained CPI-U, would generate savings for the Social Security program.

The approach taken above can readily include a requirement that reform lower the expected present-value cost of benefits. To illustrate this, I reduce all Status Quo benefits by 10 percent and impose the same restriction on reform policies as before, namely that they have the same expected present value total cost of benefits. The baseline results change very little with this variation, with the same households lining up on either end of the seesaw as in the baseline case and the same net welfare implications obtaining. The intuition for these results is that the relative effects of the reforms are largely unaffected by the shift in their total value. Once all are adjusted to provide 10 percent smaller total benefits, the backloaded reform continues to provide better longevity risk protection than the modified Status Quo or frontloaded reform, while the poorest, shortest-lived retirees continue to prefer the frontloaded reform that provides greater benefits early on. These results support the argument that analyses of the level and shape of benefits may be done separately.

5.3 Non-constant Status Quo benefits

Second, thus far I have assumed that Status Quo benefits are constant in real terms. In reality, there is considerable debate and uncertainty over whether they are increasing or decreasing in real terms. Goda et al. (2007) argue that current benefits-indexing, and even the faster-growing CPI-E, fail to provide enough protection against the rising costs of medical expenditure among retirees as they age and over time. Specifically, they calculate the real Social Security benefit net of medical expenses and show that it grew more slowly from 1983 to 2007 than did a price index of non-medical goods and services (so the real non-medical purchasing power of Social Security beneficiaries declined). On the other hand, fixed-basket price indices such as the CPI-W are susceptible to the well-known problem that they overestimate the inflation faced by individuals due to quality changes and substitution away from expensive goods and services (see Boskin et al, 1996).

To test the sensitivity of our baseline results to this assumption, I consider two alternatives.

First, to study the possibility that the CPI-W underestimates the inflation faced by retirees, I assume that the CPI-E is, in fact, the correct price index for retirees. This means that Backloaded Reform now has $\pi = 0.00$, and I set its $\lambda = 1.00$ to impose budget neutrality as in the baseline case. The Status Quo now has $\pi = -0.0037$ and $\lambda = 1.031$, while Frontloaded Reform has $\pi = -0.0064$

and $\lambda = 1.054$. Visually, the benefit paths are as follows:

Figure 7: Annual benefits paths under Status Quo (solid line), Backloaded Reform (dashed line), Frontloaded Reform (dotted line), and V-shaped Reform (dash-dot line), in thousands of real 2005 dollars, for four of ten household lifetime-income deciles, when the Backloaded Reform is the one that provides a constant real benefit.

Though the benefits paths in Figure 7 look quite different from those in Figure 3, the relative effects of reform on households are remarkably similar in this variation on the baseline analysis. Essentially the same households benefit from Backloaded Reform and Frontloaded Reform, and to very similar degrees. The results on the net welfare implications of reform are very similar, as well, under either the utilitarian or Rawlsian criterion.

Second, to study the possibility that the CPI-W overestimates the inflation faced by retirees, I assume that the chained CPI-U is, in fact, the correct price index for retirees. This means that Frontloaded Reform now has $\pi = 0.00$, and I set its $\lambda = 1.00$, the Status Quo now has $\pi = +0.0027$ and $\lambda = 0.978$, while Backloaded Reform has $\pi = +0.0064$ and $\lambda = 0.947$. The baseline results are robust to this variation, as well.

5.4 Bequest motive

The retirees in the baseline model have no reason to retain wealth other than longevity risk. In much of the existing literature explaining retiree wealth dynamics, a bequest motive is used as an ingredient to explain the retention of substantial assets late into retirement. As noted earlier, the simulations in this paper generate paths for what Love et al. (2008) call "annualized comprehensive wealth," that fit well with what appears in the data. Nevertheless, it may be valuable to understand the robustness of our results to the existence of a bequest motive, given its prominence in previous, more sophisticated, simulations of retiree behavior.

To test this, I have households value any assets left at death as if those assets were consumed by their heirs in the next period, multiplied by a factor scaling the strength of their bequest motive, using the same utility function specification (4) as for the household while it was alive. That parameter, κ is:

Table 7: Bequest motive by income decile										
	Bottom	2	3	4	5	6	7	8	9	Top
κ	1	1	1	1	1	10	20	30	40	50

such that the top deciles have a relatively strong bequest motive. The resulting simulation has, as would be expected, no households die with zero assets (whereas in the baseline simulation 19 percent of households die with zero assets). This likely overstates the degree of bequeathing done at the bottom of the income distribution and understates it at the top. Nevertheless, the results are informative in that the changes from the baseline simulation are quite minor outside of the pattern of asset holdings. That is, the top decile retirees now die with assets equal to about three times their consumption late in retirement. The sets of households who gain and lose from each reform are largely the same as in the baseline analysis, as is the ranking of policies under either the utilitarian or Rawlsian criteria. One minor but interesting difference from the baseline is that high-decile retirees raise their consumption less later in life under the backloading reform than in the baseline case. Intuitively, with a bequest motive these retirees save more of their private assets until death.

Of course, the non-surviving household leaves assets unspent in the Frontloaded policy, and if those assets were reclaimed by the government, the difference between the policies would diminish. In reality, the U.S. government raises very little revenue from the taxation of bequests, and none from households for which Social Security benefits materially change their accumulation of assets, so I assume that the direct fiscal costs of benefits is not offset by any posthumous taxation. Similarly, I do not consider the value inheritors place on bequests, a topic analyzed in Feldstein (1990), to retain this paper's focus on the direct effects of reform on retirees.

5.5 Additional transfers

Throughout the analysis I abstract from additional transfers made to poor retirees. In reality, the very poor elderly receive support from the Supplementary Security Income program as well as more

targeted programs such as in-kind "food stamps" through the SNAP program. The baseline SSI transfer was approximately \$10,000 per year for an elderly couple and \$7,000 for an individual in 2005, according to the Social Security Administration. SSI benefits are displaced dollar-for-dollar by almost any income source, including Social Security benefits.

While including these transfers in the simulations above is technically straightforward, there are conceptual complications. In particular, in reality these transfers are indexed for inflation just as are Social Security benefits. If we include these transfers without adjusting their indexing approach—something the President’s 2014 budget proposal suggested—changes to the Social Security benefits of the lowest-decile retirees are entirely canceled out by changes to their SSI benefits (though the government would save some money that could be allocated to other retirees). This mechanically neutralizes the benefits of Frontloaded reform. In fact, in simulations of the baseline model modified to include a guaranteed minimum benefit of \$9,000 that falls dollar-for-dollar with Social Security benefit increases, the lowest-decile households are unaffected by any indexing reform, most other households lose from the Frontloaded Reform, and all other households gain from Backloaded Reform. That is, the Backloaded Reform can produce a Pareto improvement in this case relative to the Status Quo. Another, simpler scenario is that such transfers would also be adjusted in any reform to benefits-indexing, so that the net effect on beneficiaries of a reform to Social Security benefits-indexing may be only partially offset, not offset at all, or even magnified. Because of this ambiguity, and the likelihood that including such transfer programs in the analysis would strengthen the results of the baseline, I chose to omit them from the main paper. Of course a more comprehensive analysis that included a range of potential changes to these transfer programs would be valuable.

Finally, note that our omission of these transfers causes the marginal utility values of consumption for individuals in the lowest income decile to be larger than if these transfers were included. This factor will cause the baseline analysis to overestimate the appeal of frontloaded reform and underestimate the appeal of backloaded reform.

6 Discussion and Conclusion

The choice of a price index for Social Security benefits may seem to have, for most purposes, small stakes. One exception, however, is its implications for retirees who rely on Social Security benefits to fund their consumption either because their own resources are limited or they outlive their expected lifespan. For these retirees, half of a percentage point faster growth in benefits—approximately the difference between two of the most prominent proposals for indexing reform—turns into a 20 percent increase in benefits if they outlive their private savings. On the other hand, assuming budget-neutral reform, it also can mean benefits that are 7 percent lower at the start of retirement, when they are sure to be alive to receive them.

In this paper, I outline a flexible and relatively simple formal structure for modeling this tradeoff in the direct effects of benefits-indexing reform on a population of heterogeneous retiree households.

I bring to that model evidence from recent empirical work on Social Security, quantifying the effects of three prominent policy proposals. I gather some new evidence on the priorities Americans appear to have for Social Security benefits, using a methodology that may prove useful more broadly. Finally, using that evidence, as well as conventional normative criteria, I provide suggestive estimates of those proposals' effects in terms of social welfare.

The results of this analysis suggest that reform to a backloaded benefits-indexing approach, such as the CPI-E, has substantial appeal, at least in terms of its direct effects on retirees. Note that this is, of course, the opposite proposal to the one that has generated the most enthusiasm in Washington: namely, a switch to the slower-growing chained CPI-U. A backloaded approach's ability to concentrate resources at later ages, when retirees face longevity risk and have exhausted their own resources, makes it the preferred approach for most retirees. While a normative criterion that concentrates priority on the worst-off retirees would therefore endorse a frontloaded reform, the standard utilitarian criterion and the criterion implied by the survey evidence in this paper prefer to backload the path of benefits.

Political considerations make the case for backloaded benefits-indexing reform extremely difficult, however. Such a reform would require a reduction in initial benefits to retain budget neutrality or an increase in total spending on benefits to retain initial benefit levels. Both requirements are likely to be deal-killers in Washington.

In this context, the appeal of President Obama's 2014 budget proposal for a benefits-indexing reform that combines a shift to the chained CPI-U with benefit enhancements at advanced ages becomes clear. Such a proposal can capture the best parts of both of the simpler reforms—protecting both the poorest, shortest-lived retirees who would prefer frontloading and the large majority of retirees, especially those who live to advanced ages, who prefer backloading. It is important to note that the President's specific proposal combined this hybrid of frontloading and backloading with an increase in progressivity, which might be achieved through other means, as the benefit enhancements at advanced ages were to be uniform across the lifetime-income distribution. In the simulations above it causes the top five income deciles of retirees to prefer the Status Quo to this reform, and the potential disincentive effects from which this paper abstracts may therefore reduce this proposal's appeal. Nevertheless, if those disincentive effects are limited and the normative preferences of Americans resemble those of either the conventional utilitarian criterion or those implied by the survey results in this paper, the Hybrid Progressive Reform is likely to generate—in terms of the direct effects on retirees—a sizeable net welfare gain.

As this result and the rest of the analysis in this paper has demonstrated, benefits-indexing reform is more than just a fiscal issue; its distributional implications and its possible role as a vehicle for redistribution make it a flexible and potentially powerful policy tool. Given that, it is important to reiterate that this paper uses a simplified model that abstracts from a number of effects of shifting the time-path of benefits on household behavior and the general economic environment, as well as from complexities of the Social Security system and retiree household structure. My hope is that it puts that simplicity to good use, clarifying a piece of the tradeoffs involved in choosing a

method of benefits-indexing, and that further analyses will refine our understanding of the lessons learned here.

References

- [1] Boskin, Michael J., E. Dulberger, R. Gordon, Z. Griliches, and D. Jorgenson (1996). "Toward a More Accurate Measure of the Cost of Living," Final Report to the Senate Finance Committee, December 4
- [2] Bosworth, Barry and Kathleen Burke (2013). "Differential Mortality and Retirement Benefits in the Health and Retirement Study," Brookings.
- [3] Burdick, Clark and Lynn Fisher (2007), "Social Security Cost-of-Living Adjustments and the Consumer Price Index," Social Security Bulletin, 67(3).
- [4] Butrica, Barbara A., Joshua H. Goldwyn, and Richard W. Johnson (2005). "Understanding Expenditure Patterns in Retirement," Center for Retirement Research at Boston College WP 2005-03, January.
- [5] Committee for a Responsible Federal Budget (2013). "The President's Proposed Chained CPI Protections." Accessed at <http://crfb.org/blogs/presidents-proposed-chained-cpi-protections>, August 2014.
- [6] Congressional Budget Office (2006). "Is Social Security Progressive?" December 15, 2006.
- [7] Duggan, James E., Robert Gillingham, and John Greenlees (2007), "Mortality and Lifetime Income: Evidence from Social Security Records," Working paper.
- [8] Englehardt, Gary V. and Jonathan Gruber (2004), "Social Security and the Evolution of Elderly Poverty," NBER WP 10466.
- [9] Feldstein, Martin (1976), "On the Theory of Tax Reform," Journal of Public Economics, 6, 77-104.
- [10] Feldstein, Martin (1985), "The Optimal Level of Social Security Benefits," Quarterly Journal of Economics, 50(2).
- [11] Feldstein, Martin (1987), "Should Social Security Benefits Increase with Age?" NBER Working Paper 2200.
- [12] Feldstein, Martin (1990), "Imperfect Annuity Markets, Unintended Bequests, and the Optimal Age Structure of Social Security Benefits." Journal of Public Economics 41(4), pp. 31-43.
- [13] Gaertner, Wulf and Erik Schokkaert (2012). Empirical Social Choice: Questionnaire-Experimental Studies on Distributive Justice. Cambridge.
- [14] Gale, William G., John Karl Scholz, and Ananth Seshadri (2009). "Are All Americans Saving 'Optimally' for Retirement?"

- [15] Goda, Gopi Shah, John B. Shoven, and Sita Natari Slavov (2009). "Differential Mortality by Income and Social Security Progressivity," Working Paper.
- [16] Guesnerie, R (1977), "On the Direction of Tax Reform," *Journal of Public Economics*, 7, 179-202.
- [17] Horton, John J., David G. Rand, and Richard J. Zeckhauser, "The Online Laboratory: Conducting Experiments in a Real Labor Market," *Experimental Economics*.
- [18] Hurd, Michael D. and Susann Rohwedder, (2008), "The Adequacy of Economic Resources in Retirement," Michigan WP 2008-184.
- [19] Krugman, Paul (2012). "Life, Death, and Deficits," *New York Times*, November 15.
- [20] Liebman, Jeffrey B.; Erzo F.P. Luttmer; and David G. Seif, (2009). "Labor supply responses to marginal Social Security benefits : Evidence from discontinuities." *Journal of Public Economics* 93, pp. 1208-1223.
- [21] Love, David A., Michael G. Palumbo, and Paul A. Smith, (2009) "The trajectory of wealth in retirement," *Journal of Public Economics*.
- [22] Olsen, Anya (2008). "Distributional Effects of Reducing the Cost-of-Living Adjustments," Social Security Administration Policy Brief No 2008-03, November.
- [23] Penner, Rudolph G. and Karen E. Smith, (2010). "How Quickly Do Older Adults Spend Their Wealth? Urban Institute.
- [24] Pestieau, Pierre and Gregory Ponthiere, (2012), "The Public Economics of Increasing Longevity," *Hacienda Public Espanola/Revista de Economia Publica* 200, pp. 35-68.
- [25] Saez, Emmanuel and Stefanie Stantcheva (2014). "Generalized Social Marginal Welfare Weights for Optimal Tax Theory." Available as NBER working paper 18835.
- [26] Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun, (2006), "Are Americans Saving "Optimally" for Retirement?" *Journal of Political Economy*
- [27] Shiller, Robert J. (2003). *The New Financial Order*. Princeton.
- [28] Smith, Karen E. and Melissa M. Favreault (2013), "A Primer on Modeling Income in the Near Term, Version 7 (MINT7)," Urban Institute, Washington D.C., October.
- [29] Smith, Karen, Mauricio Soto, and Rudolph G. Penner (2009). "How Seniors change their Asset Holdings during Retirement," Center for Retirement Research at Boston College WP 2009-31 Released: December 2009.
- [30] Waldron, Hilary, (2007), "Trends in Mortality Differentials and Life Expectancy for Male Social Security-Covered Workers, by Socioeconomic Status," *Social Security Bulletin*, 67(3).

- [31] Weinzierl, Matthew (2014). "The Promise of Positive Optimal Taxation," *Journal of Public Economics*.
- [32] White House (2010). *The National Commission on Fiscal Responsibility and Reform*, December.