Journal of Public Economics 256 (2026) 105601

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/pubec

On the optimality of deferred public annuities[☆]

Liran Einav[a], Amy Finkelstein[b,*]

[a] Stanford University and NBER, Stanford, USA
[b] MIT and NBER, Cambridge, USA

ARTICLE INFO

Keywords:
Social security
Annuities
Deferred annuities

ABSTRACT

What is the optimal path of Social Security benefits for an individual who has retired with a stock of wealth, faces stochastic mortality, and has no access to annuities and no preferences for bequests? It is a deferred annuity in which the government annuity pays out zero for some periods and a constant amount after that. The optimal length of the deferral period increases with the retiree's initial wealth and their survival probability.

1. Introduction

Consider the following, simple, two-period problem for optimal Social Security benefit design: A retired individual arrives in period 1 with sufficient initial wealth, which they can save. The retiree may die before period 2 with probability $p < 1$, so they cannot borrow. They have no access to annuities, no preferences for bequests, and there is no interest rate or discount factor. The government needs to decide its Social Security policy, in which it disburses S dollars per retiree in period 1 (s_1) and period 2 (s_2), subject to the government budget constraint that $S = s_1 + ps_2$ (as in a standard overlapping-generations setting) and the behavioral response of the retiree to the government's choice of annuity payment profile. What is the optimal Social Security benefit structure?

The answer to the simple problem – as well as to the natural extension to any finite time horizon – is that the optimal government social security policy is to defer the entire Social Security benefit payment to the last period. Intuitively, the government wants to save as much as possible on behalf of the individual, as it has access to a better savings technology: it can transfer money between periods at a price of $p < 1$, whereas the individual – because they have no access to annuity markets – can only do so at a price of 1.

With an infinite horizon and continuous time, the solution generalizes to an optimal government Social Security benefit structure of a deferred annuity, in which the government annuity pays out nothing early on and at a constant rate thereafter. The optimal length of the deferral period increases with initial wealth and survival probability. If individuals start with no initial wealth, then a constant annuity per period becomes optimal. Thus, the optimal deferral length would vary across heterogeneous individuals.

Section 2 derives these results. In Section 3 we discuss some implications for Social Security design, as well as the connection between our paper and a number of other highly related papers. In particular, our result is not "new" in the sense that there are papers that have our result as their (perhaps obvious) implication, but we had not seen it spelled out this way before.

2. Theory

We consider a setting in which a representative retiree – with a standard per-period utility from consumption $u(c)$[1] – reaches retirement age with accumulated personal savings (or initial wealth) w, facing stochastic mortality. The government is endowed with the retirees' accumulated Social Security contributions S (per retiree). The government problem is to choose its optimal Social Security payout schedule given its total budget of S, taking into account the retirees' optimal consumption-saving decisions and their survival probabilities.

For simplicity, we assume that the timing of retirement and the initial retirement wealth are exogenous, and that the individual cannot borrow against the future, does not discount the future or receive an interest rate on savings, has no bequest motive, and has no access to annuities. The no-borrowing constraint is standard, imposed to eliminate the possibility that the individual may die in debt. The assumption of no discounting

☆ We are grateful to the Editor, three anonymous referees, Doug Bernheim, Jim Poterba, Pietro Tebaldi, Chris Tonetti, and Ivan Werning for helpful comments, and to Carine You and especially Jack Kelly for excellent research assistance.

* Corresponding author.
Email addresses: leinav@stanford.edu (L. Einav), afink@mit.edu (A. Finkelstein).

[1] By "standard" we mean that we assume throughout that $u(c)$ is strictly increasing, strictly concave, twice continuously differentiable, and that it satisfies the Inada conditions ($\lim_{c\to 0} u'(c) = \infty$ and $\lim_{c\to\infty} u'(c) = 0$).

https://doi.org/10.1016/j.jpubeco.2026.105601
Received 12 October 2025; Received in revised form 23 January 2026; Accepted 13 February 2026

or interest rate on savings is inconsequential, and is merely done for simplicity.

The assumption that the individual has no access to annuities is more extreme than necessary; many frictions in private annuity markets would create the same qualitative result that there is scope for potentially welfare-improving government policy in providing an annuity through Social Security. The existence of frictions in private annuity markets is motivated by the so-called "annuity puzzle": empirically, very few individuals voluntarily purchase annuities, despite the well-known theoretical result that the optimal choice for a retired individual facing an unknown date of death and actuarially fair annuity markets is to fully annuitize their retirement wealth (Yaari, 1965; Davidoff et al., 2005). A variety of explanations have been proposed to address this puzzle, including non-actuarially fair pricing of annuities (Mitchell et al., 1999), bequest motives (Lockwood, 2018), late-life medical expenditure risk (Reichling and Smetters, 2015), and various behavioral biases (Brown, 2009). Rather than micro-found the lack of voluntary annuitization, we study the optimal government policy under the assumption that individuals have no access to private annuities.

2.1. Finite horizon

$T = 2$. The key insight is most easily illustrated using the two-period setting we alluded to in the first paragraph of the Introduction. The individual retires at the beginning of period 1 with financial wealth w. They live in period 1 for sure, and survive to period 2 with probability p, after which they die for sure. The government's Social Security benefit policy is defined by the amount it pays the retiree in period 1 (s_1) and in period 2 (s_2). The government's problem is to choose the payment profile (s_1, s_2) that maximizes the retiree's expected utility, subject to the government's budget constraint for an actuarially fair Social Security benefit structure ($s_1 + p \cdot s_2 \leq S$) and the behavioral response of the retiree to the government's choice of payment profile.

It is instructive to start with the first-best benchmark in which the government can fully tax away the individual's starting wealth w, and hence assign the individual's per-period consumption directly subject to their budget constraint, $s_1 + p \cdot s_2 \leq w + S$. Equivalently, we can derive the first best benchmark by assuming that the individual's starting wealth is $w + S$ and that they have access to actuarially fair annuities in which they could transfer a dollar in period 1 to $1/p$ dollars in period 2. In either scenario, the government (or the individual) would optimally choose equal consumption across the two periods ($c_1^{FB} = c_2^{FB}$) so that the individual's marginal utility of consumption is constant across periods ($u'(c_1^{FB})/u'(c_2^{FB}) = 1$). This is simply the classic result that in the presence of actuarially fair insurance and the absence of moral hazard, full insurance is optimal (Arrow, 1963), and more specifically, full annuitization is optimal (Yaari, 1965).

If, however, the retiree retains w and has no ability to annuitize it, the first best outcome of equal consumption across periods generally cannot be achieved. Rather, they choose consumption in periods 1 and 2 to maximize expected utility subject to their no-borrowing constraint and their lifetime budget constraint:

$$\max_{c_1,c_2} \left[u(c_1) + p \cdot u(c_2)\right] \tag{1}$$

subject to

$$c_1 \leq w + s_1, \tag{2}$$

and

$$c_2 \leq w + s_1 + s_2 - c_1. \tag{3}$$

If the individual's no-borrowing constraint (Eq. (2)) is not binding, they will optimally choose to front-load their consumption toward period 1 (i.e., $c_1^* > c_2^*$, such that $u'(c_1^*)/u'(c_2^*) = p$) because every unit of consumption in period 2 costs them a unit of consumption in period 1 (see the retiree's budget constraint in Eq. (3)), but they only receive the period 2 consumption with some probability $p < 1$.[2]

The government therefore chooses the annuity benefits s_1 and s_2 to maximize the retiree's expected utility subject to the government budget constraint and the retiree's optimal choices of consumption (denoted by c_1^* and c_2^*):

$$\max_{s_1,s_2} \left[u(c_1^*) + p \cdot u(c_2^*)\right] \tag{4}$$

subject to

$$s_1 + p \cdot s_2 \leq S, \tag{5}$$

The government's optimal policy is to defer more of the annuity payment to the second period, until either the retiree's consumption is equalized across periods or all of the government's payout is concentrated in the second period, whichever occurs first. Specifically, for $w \geq S/p$, the government's optimal schedule is to defer the entire Social Security payment to period 2, so that $s_1 = 0$ and $s_2 = S/p$; as a result, the retiree will choose $c_1^* \geq c_2^*$, with the equality holding (and the first best outcome obtained) at $w = S/p$. For $w < S/p$, the government's optimal policy is $s_1 = \frac{S-pw}{1+p}$, which achieves the first best outcome of the retiree equating their consumption across periods ($c_1 = c_2 = \frac{S+w}{1+p}$).

Crucially, the government has access to annuity markets, while (by assumption) the retiree does not. This enables the government to transfer dollars across periods at a lower cost than the retiree can. Specifically, transferring an additional dollar in s_2 costs the government only $p < 1$ in s_1 (see the government's budget constraint in Eq. (5) compared to the retiree's in Eq. (3)). If the retiree were given the choice of whether and how much of S to annuitize at actuarially fair rates, they would choose the same solution as the government. In other words, given our other assumptions, there is no wedge between the privately optimal and the socially optimal allocation.

$T > 2$. Consider now $T > 2$ periods, with p_t being the (strictly positive and strictly decreasing in t) probability of surviving to period $t \leq T$ (and death is certain after period T). The back-loaded nature of the two-period solution generalizes. However, with a large enough horizon T (and fixed initial wealth w) the consumer's no-borrowing constraint will eventually be binding, so paying everything in the last period will generally not be optimal. Instead, the optimal Social Security schedule is a deferred annuity, in which there are initially no payments until the individual exhausts w, at which point there is a flat payout schedule for the remainder of the periods; during this flat payout period the first-best consumption path obtains.[3] The full characterization of this solution is in Appendix C. As T gets large, the solution approaches the infinite horizon case, which we turn to next.

2.2. Infinite horizon

We now turn to our main deferred-annuity result, which is obtained in the context of an infinite-horizon version of the problem. It will be tractable to do this in continuous time.[4] Specifically, the individual now

[2] A non-binding no-borrowing constraint is what we were assuming in the opening paragraph when we said (much less precisely) that the individual has "sufficient initial wealth."

[3] The discrete nature of the problem means that it is common (depending on the parameters) to have a single "top up" period, just before the flat schedule begins, in which the optimal Social Security payment is positive but lower than the flat per-period amount.

[4] As shown in Appendix C, the discrete-time version of this problem converges (as the time increment between periods goes to zero) to the continuous-time result.

has to solve the following problem:

$$\max_{\{c_t\}} \int_0^\infty p_t u(c_t) dt, \tag{6}$$

subject to the no-borrowing constraint

$$\int_0^t c_\tau d\tau \le w + \int_0^t s_\tau d\tau \quad \forall t, \tag{7}$$

where $\{s_t\}$ is the government's social security payout schedule and p_t is the probability of survival to date t.[5] The government's problem is then to choose the Social Security payout schedule $\{s_t\}$ to maximize the retiree's utility subject to the individual making optimal consumption choices in response to the government's choice and subject to the Social Security schedule satisfying the government's budget constraint:

$$\int_0^\infty p_t s_t dt = S. \tag{8}$$

As shown in the following theorem, the solution to the government's problem is a deferred annuity that pays nothing until some point, after which it pays a constant amount thereafter:

Theorem 1. *The optimal government transfer schedule takes the form of a deferred annuity:*

$$s_t = \begin{cases} 0 & t < t^* \\ \left(\int_{t^*}^\infty p_t dt\right)^{-1} S & t \ge t^*. \end{cases} \tag{9}$$

The formal proof is in Appendix A, but here we try to provide intuition for the key steps. The first thing to note is that the objective function of the government and the individual is the same, so we can rewrite the problem: instead of the government optimally choosing $\{s_t\}$ subject to the individual making optimal consumption choices $\{c_t\}$ (given the government's choice $\{s_t\}$), we can solve the equivalent problem in which the government chooses both $\{s_t\}$ (subject to its budget constraint) and $\{c_t\}$ (where its choice of $\{c_t\}$ is constrained by the same no-borrowing condition as the retiree in Eq. (7)).

The second step of the proof is to note that it is suboptimal to have positive Social Security payments ($s_t > 0$) when the retiree's assets are positive ($w_t > 0$). This is because, as in the two-period problem in Section 2.1, the government has a more efficient way to transfer assets to the future than the retiree. Therefore, if the retiree were going to save at time t and the Social Security payment rate were positive, the government could increase the retiree's expected utility by slightly reducing the Social Security payment rate, so that the government – rather than the retiree – is the one that does more of the saving. The converse is also true: because of the Inada conditions on the retiree's utility function, if the retiree has no assets at time t, then any optimal solution must have a positive Social Security payment rate to guarantee a positive consumption rate.

This in turn implies that the optimal solution will have a two-phase structure. In the first phase, saving rates are positive and Social Security payments are zero. In the second phase Social Security payments are positive and the individual has no other assets and does not save, so they consume their Social Security payments. We denote by t^* the time at which the first phase ends and the second begins.

Once the two-phase solution is established, we can solve it backwards, conditional on t^*. In the second phase, the government is effectively facing a standard consumption-saving problem (on behalf of the individual) but has access to the equivalent of a frictionless annuity market. The optimal solution is thus the first best, in which consumption is equalized across periods. The government's budget constraint (Eq. (8)) dictates the level of this constant consumption path.

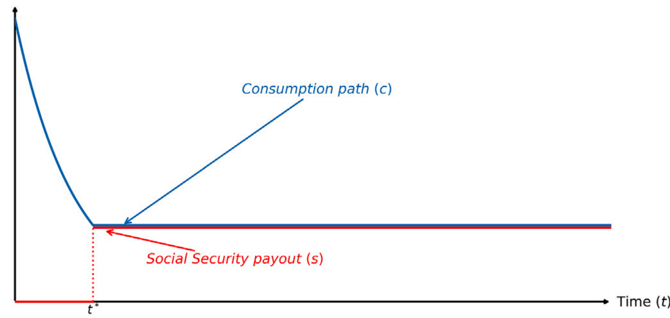

Fig. 1. Illustration of Theorem 1. **Notes**: The figure illustrates the optimal annuity payout rate and the associated optimal consumption path in the infinite-horizon continuous-time setting.

Prior to t^*, the individual consumes out of their own wealth. Just like the finite-horizon model presented in Section 2.1, this means that – absent access to private annuities – the individual will front-load their consumption because of their positive mortality risk; as a result, they will choose an optimal c_t that is a declining function of t.

Finally, to determine the value of t^* at the optimum, it is instructive to realize that c_t would optimally be continuous at t^* (if it is not, one could "locally" smooth out consumption around t^* and increase utility). Then one can notice that in the second phase, the constant rate of consumption (which is equal to the constant Social Security payout rate) is increasing in t^* (see Eq. (9)), while the consumption rate just before t^* is decreasing in t^* (as the initial assets get spread over a longer phase-one period), so there is a unique value of t^* at which c_t is continuous at t^*.

Fig. 1 illustrates the optimal solution, showing the qualitative nature of the optimal consumption rate and optimal Social Security payments. The individual will optimally front-load their consumption out of their initial wealth, and will gradually lower their consumption rate until it runs out of their own wealth and hits the Social Security payment rate at t^*, at which point they are on a constant consumption rate and consume their Social Security payment.

Comparative statics. Thus far we have considered a representative retiree. In Appendix B we allow for heterogeneity and show that the optimal value of t^* satisfies intuitive comparative statics properties. Specifically, we show that the optimal deferral period t^* is longer for richer individuals (higher w), individuals with higher survival rates, and individuals with greater absolute risk aversion (or, similarly, with lower elasticity of intertemporal substitution).

The key intuition behind these comparative statics is that – given the government's budget constraint – a longer deferral period means a higher rate of consumption after t^*. Individuals with greater wealth can consume longer out of their own wealth, which allows the government to delay their annuity payments longer and reduce the difference in consumption rates between the two phases. Likewise, individuals with greater survival rates will choose a less front-loaded consumption path prior to t^*, making it more beneficial to extend the period during which they consume out of their own wealth. Finally, more risk averse individuals or individuals who are associated with lower elasticity of intertemporal substitution have a stronger preference for consumption smoothing and thus for reducing the difference in consumption rates across the two phases, so delaying the start of Social Security payments allows for a higher level of consumption after t^*.

3. Discussion

One of the primary functions of Social Security programs is to ensure that retirees have annuity income (Feldstein and Liebman, 2002; Diamond, 2004). This stems from the welfare-enhancing role

[5] More specifically, we assume that $p_0 = 1$, that p_t is strictly decreasing in t, and that $\int_0^\infty p_t < \infty$ (which implies that $\lim_{t\to\infty} p_t = 0$).

of annuities (Yaari, 1965; Davidoff et al., 2005), combined with the extremely limited voluntary purchase of annuities; as noted, the limited voluntary annuity market in turn may reflect some combination of market failures and/or optimization frictions.

Given the potential rationale for publicly-mandated annuities, we explored the optimal form of the publicly-mandated annuity stream. We conclude that it takes the form of a deferred annuity, in which, after the individual is retired, the mandated annuity pays out zero for some period, and a constant amount thereafter. The optimal deferral length increases with the individual's retirement wealth and decreases with their remaining expected lifetime.

In practice, of course, no country to our knowledge provides its Social Security benefit in the form of a deferred annuity. If the individual has reached the minimum eligibility age for Social Security, payments can typically start immediately upon retirement. This presumably reflects the fact that the optimal deferral length can vary dramatically with the individual's wealth at retirement, creating both practical difficulties (in that wealth is typically private information) and potentially political ones in trying to tailor the deferral period to the individual's wealth.

There are, however, countries where retirees can take their pension as a deferred annuity, and choose the length of the deferral. Specifically, Chile has a mandatory national savings program for workers, which they can receive at retirement either as a phased withdrawal or as an annuity. Retirees who choose an annuity have a choice over various annuity features, including the length of any deferral (Illanes and Padi, 2019). In our stylized framework, the retiree, like the government, would choose a deferred annuity, with the length of the deferral increasing in their wealth and in their survival probability.

Naturally our stylized setting abstracts from a number of potentially important channels that may influence the optimal form of the publicly-mandated annuity stream. These include, but are not limited to, our assumption of no bequest motives, exogenous retirement, and exogenous accumulation of savings prior to retirement. Our objective is not policy prescription per se, but rather to highlight a particular economic force that is at play.

This economic force can also be found in prior academic research. Feldstein (1987) appears to be the most closely related to our theoretical results. He considers the optimal age profile of retirement benefits for rational consumers with positive wealth in the absence of annuity markets. While the focus is different from ours,[6] the paper includes the result that in a two-period model the annuity should pay out only in the second period.[7] Scott et al. (2007) also presents highly-related theoretical results. They show that once there are costs to private annuitization such that full annuitization of retirement wealth is no longer optimal, the optimal partial annuity contract involves a deferred annuity. Moreover, similarly to our two-phase consumption path result, they show that optimal consumption in their setting is first drawn from the individual's non-annuitized assets and then (once the non-annuitized assets are exhausted) from the annuitized assets.

[6] Feldstein (1987) focuses on the tradeoff between the gains from risk protection from backloading annuities relative to the loss of the higher returns from private assets if the annuity is paid out sooner.

[7] Similarly, Bernheim (1987) derives the consumer's optimal consumption choices as a function of their initial wealth and the annuity payment profile, while Bernheim (1987) similarly solves for the individual's optimal wealth drawdown as a function of the annuity payment profile. Both highlight the distinction between the individual's and the planner's effective discount factors.

In addition, versions of our theoretical results can be found in results from various calibrated models. For example, Horneff et al. (2023) analyze retirees' optimal choices for payouts from their defined contribution plans and find that choosing a deferred annuity can be beneficial; Horneff et al. (2025) simulate the expected utility gains from defaulting some portion of retirees' 401(k) balances into immediate or deferred annuities, and find that the gain from a deferred annuity (relative to a non-deferred one) increases with the individual's socioeconomic status; Braun et al. (2017) simulate the welfare gains from means-tested social insurance programs and find that scaling up the program would benefit both the rich and the poor; and Coile et al. (2002) consider the decision of individuals to delay claiming their Social Security benefits and find via simulation that the optimal length of the delay in claiming benefits increases with wealth. Our theoretical note helps clarify and conceptualize the key economic intuition behind these and related results.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary data for this article can be found online at doi:10.1016/j.jpubeco.2026.105601.

Data availability

No data was used for the research described in the article.

References

Arrow, K.J., 1963. Uncertainty and the welfare economics of medical care. Am. Econ. Rev. 53 (5), 941–973.
Bernheim, B.D., 1987a. Dissaving after retirement: testing the pure life cycle hypothesis. In: Shoven Zvi Bodie, J.B. and Wise, D.A. (Eds.), Chap. 9 in *Issues in Pension Economics*. University of Chicago Press, pp. 237–280.
Bernheim, B.D., 1987b. The economic effects of social security: toward a reconciliation of theory and measurement. J. Public Econ. 33 (3), 273–304.
Braun, R.A., Kopecky, K.A., Koreshkova, T., 2017. Old, sick, alone, and poor: a welfare analysis of old-age social insurance programmes. Rev. Econ. Stud. 84 (2), 580–612.
Brown, J.R., 2009. Understanding the role of annuities in retirement planning. In: Lusardi, A. (Eds.), Overcoming the Savings Slump: How to Increase the Effectiveness of Financial Education and Saving Programs. University of Chicago Press, pp. 178–206.
Coile, C., Diamond, P., Gruber, J., Jousten, A., 2002. Delays in claiming social security benefits. J. Public Econ. 84 (3), 357–385.
Davidoff, T., Brown, J.R., Diamond, P.A., 2005. Annuities and individual welfare. Am. Econ. Rev. 95 (5), 1573–1590.
Diamond, P., 2004. Social security. Am. Econ. Rev. 94 (1), 1–24.
Feldstein, M., Liebman, J.B., 2002. Social security. Handb. Public Econ. 4, 2245–2324.
Feldstein, M.S., 1987. Should Social Security Benefits Increase with Age? NBER Working Paper No. 2200.
Horneff, V., Maurer, R., Mitchell, O.S., 2023. Fixed and variable longevity income annuities in defined contribution plans: optimal retirement portfolios taking social security into account. J. Risk Insur. 90 (4), 831–860.
Horneff, V., Maurer, R., Mitchell, O.S., 2025. Defaulting 401 (K) Assets into Payout Annuities for "Pretty Good" Lifetime Incomes. NBER Working Paper No. 33915.
Illanes, G., Padi, M., 2019. Retirement Policy and Annuity Market Equilibria: Evidence from Chile. NBER Working Paper No. 26285.
Lockwood, L.M., 2018. Incidental bequests and the choice to self-insure late-life risks. Am. Econ. Rev. 108 (9), 2513–2550.
Mitchell, O.S., Poterba, J.M., Warshawsky, M.J., Brown, J.R., 1999. New evidence on the money's worth of individual annuities. Am. Econ. Rev. 89 (5), 1299–1318.
Reichling, F., Smetters, K., 2015. Optimal annuitization with stochastic mortality and correlated medical costs. Am. Econ. Rev. 105 (11), 3273–3320.
Scott, J.S., Watson, J.G., Wei-Yin, H., 2007. Efficient Annuitization: Optimal Strategies for Hedging Mortality Risk. Pension Research Council Working Paper No. 2007-09.
Yaari, M.E., 1965. Uncertain lifetime, life insurance, and the theory of the consumer. Rev. Econ. Stud. 32 (2), 137–150.