

NBER WORKING PAPER SERIES

TARGETED SAVINGS AND LABOR SUPPLY

Louis Kaplow

Working Paper 15656

<http://www.nber.org/papers/w15656>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

January 2010

Harvard University and National Bureau of Economic Research. I am grateful to Alex Gelber and participants at the National Tax Association 2009 annual meeting for comments and the John M. Olin Center for Law, Economics, and Business at Harvard University for financial support. A prior version of this paper carried the title “Non-Optimizing Savings Behavior and Labor Supply,” and it grew out of work on Kaplow (2006) and chapters 9 and 11 of Kaplow (2008). The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 by Louis Kaplow. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Targeted Savings and Labor Supply
Louis Kaplow
NBER Working Paper No. 15656
January 2010
JEL No. D11,D91,H24,H31,H55,J22,J26

ABSTRACT

Substantial evidence suggests that savings behavior may depart from neoclassical optimization. This article examines the implications of raising the savings rate – whether through social security, retirement plans, or otherwise – for labor supply, where labor supply is determined by behavioral utility functions that reflect the non-neoclassical character of savings behavior. Under one formulation, raising the targeted savings rate has the same effect on labor supply as that of raising the labor income tax rate; under a second, raising the targeted savings rate has no effect on labor supply; and under a third, raising the targeted savings rate increases labor supply regardless of the slope of the labor supply curve. Effects on labor supply are particularly consequential because of the significant preexisting distortion due to labor income taxation.

Louis Kaplow
Harvard University
Hauser 322
Cambridge, MA 02138
and NBER
meskridge@law.harvard.edu

1. Introduction

Some research calls into question whether many individuals' savings behavior can be reconciled with idealized models of life-cycle maximization. See, for example, Laibson's (1996, 1997) and Thaler and Benartzi's (2004) examination of myopia and the surveys by Bernheim (2002) and Bernheim and Rangel (2007). Empirical work, reaching differing conclusions, has focused on such subjects as whether individuals' retirement savings are adequate and, relatedly, the existence of and explanation for a significant drop in consumption upon retirement.¹ Additional studies have shown that retirement savings decisions in 401(k) plans are heavily influenced by employer default rules. See, for example, Madrian and Shea (2001), Choi et al. (2004), and the survey by Beshears et al. (2006a). Relatedly, Beshears et al. (2006b) report that reducing the complexity of the retirement savings decision raises participation rates, and Carroll et al. (2009) show that requiring an active decision significantly raises enrollment. Furthermore, departures from neoclassical optimization may have normative implications. For example, Laibson (1996) estimates that correctives for inadequate savings due to myopia could raise individuals' lifetime welfare to an extent equivalent to almost an additional year's worth of income. Social insurance and tax provisions for retirement savings are often justified by concerns that individuals may inadequately provide for themselves.

Less attention has been devoted to the consequences of non-neoclassical savings behavior for labor supply.² This relationship is of intrinsic interest, and its practical importance is great in light of the significant preexisting labor supply distortion due to labor income taxation. If public and private actions that are designed to influence savings – ranging from social security to employer retirement programs – have even a small influence on labor supply, this effect could be of first-order importance in assessing such schemes. Moreover, because various savings policies are motivated in significant part by a belief that individuals' savings behavior may not be optimal, when considering the effect of such policies on labor supply it is appropriate to account for how such non-neoclassical savings behavior may influence labor supply decisions.

In order to address this subject, it is necessary to specify just how individuals do behave if, indeed, they do not fully adhere to the dictates of traditional models. Myopia has received some attention. In addition to the work of Laibson and others on myopia and savings, the implications of myopic savings for labor supply are addressed in Kaplow (2006, 2008). Models of myopia, however, typically assume that individuals are still optimizers; indeed, in some work, the optimization is quite complex, despite being infected by myopia.

The present investigation takes a different path, following the suggestion of Bernheim (1994), Diamond (2004), and others that retirement savings may be suboptimal because of the sheer complexity of retirement planning and individuals' inexperience. The motivation is that the pertinent dynamic optimization (requiring knowledge of and computations about asset

¹See, for example, Kotlikoff, Spivak, and Summers (1982), Banks, Blundell, and Tanner (1998), Engen, Gale, and Uccello (1999), Moore and Mitchell (2000), Bernheim, Skinner, and Weinberg (2001), Aguiar and Hurst (2005), Scholz, Seshadri, and Khittrakun (2006), and Smith (2006).

²See, for example, Diamond (2002) and the survey by Feldstein and Liebman (2002).

returns, earnings uncertainty, longevity, and other factors) is exceedingly complicated. This view is supported by experimental evidence offered by Johnson, Kotlikoff, and Samuelson (2001) and is reinforced by Gokhale, Kotlikoff, and Warshawsky's (2001) finding that common financial planning software is itself overly simple and, as a result, makes recommendations that are significantly erroneous. Thus, particularly with regard to savings decisions, it may well be that individuals do not engage in elaborate dynamic analysis and calculations but instead satisfice in some fashion.³

To pursue this line of inquiry, more must be said about just what individuals actually do. Here it is postulated that they follow a simple savings rule. One possibility is that individuals just do what they are ordered to do. A social insurance system may force them to save a certain portion of their earnings, which may constitute nearly all of their savings. Similarly, some employers' retirement plans may do the same, setting aside a fraction of gross wages to fund retirement benefits. Another possibility is that individuals derive their savings rules from cues in their environment. They might follow rules of thumb recommended by investment advisors or financial institutions, or they may imitate others' behavior; in addition, following such rules may become a habit, continuing to guide behavior without ongoing reconsideration of the underlying motivation. The aforementioned work suggesting that default specifications in employer retirement plans significantly influence many individuals' savings behavior is suggestive of this sort of phenomenon.

Supposing that individuals do follow some such savings rule, more must be said about how they view their savings when they make their labor supply decisions. It seems inappropriate to suppose that they are complete optimizers, taking fully into account the effect of current earnings on savings and how that influences utility in future periods, because such an approach is inconsistent with the view that they are targeted savers in the first instance. Nevertheless, it is necessary to say something about how they view their earnings when choosing labor supply. Because this question has not been studied in depth empirically or analyzed theoretically, the only available approach, taken here, is to examine the implications of a range of assumptions without making claims about their relative plausibility. Of course, understanding these implications can serve as a guide to empirical work aiming to test such behavioral assumptions.

In the first case considered, individuals are assumed simply to ignore that their savings will generate utility from future consumption. This is perhaps the simplest assumption, and it certainly requires no calculations regarding the future. When savings is via mandatory social insurance funded by payroll taxation imposed on employers or through retirement plans funded from wages, all individuals need to examine is their take-home pay, which already reflects withholding for taxes and savings. This assumption has some *prima facie* plausibility, but it does seem extreme, especially for individuals whose savings are voluntary.

Second, individuals are taken to treat their savings as if presently consumed; that is, their

³On a somewhat related subject, Liebman and Zeckhauser (2004) examine the possibility that individuals may not appreciate marginal effective tax rates and instead respond to average rates and also that individuals may ignore how present behavior affects future prices or tax rates, which they refer to collectively as "schmeduling."

savings are taken to generate the same utility as would be generated if instead they consumed the funds. In this case, individuals contemplate the consumption utility that would be generated by their after-income-tax earnings, viewing any savings as truly theirs but, being unable to assess its long-run effects on their utility, treat it as a shorthand as if it was like having more present consumption. This assumption is also fairly simple and perhaps extreme in a different way. It is, however, consistent with individuals' savings decisions being easily swayed by employers' defaults and others' advice, behavior that is suggestive of near-indifference regarding non-marginal changes in how much one saves.

Third, and perhaps most interesting, it is imagined that individuals treat their future consumption as if it has the same marginal utility as present consumption. The rationale for this variant is that individuals might suppose that those who set savings rules have themselves solved the consumption allocation problem and chosen their prescriptions accordingly. For example, designers of mandatory rules in social security or retirement plans or of voluntary targets (investment advisors or default rule architects) might be thought to have individuals' well-being in mind. Even if this view is not taken literally, individuals might believe it as a first approximation and, being unable to analyze the problem further, as a final guideline. Note that this assumption is consistent with individuals changing their savings behavior, even significantly, in response to changes in default rules, the receipt of advice, and so forth. Furthermore, specific findings in the 401(k) literature, in addition to those on direct responses to default rules, support the notion that defaults and other employer actions are treated as advice. For example, Madrian and Shea (2001) note that participants hired before automatic enrollment went into effect but who do not begin to participate until afterward are more likely to invest in the default fund, and Beshears et al. (2006a) report that when an employer match goes to employer stock, employees' contributions to employer stock are higher.

The body of the paper, section 2, presents a simple model and uses it to examine the effect of changing savings rates on labor supply under each of these three assumptions. In the first case (in which individuals ignore future consumption), raising the targeted savings rate has an effect qualitatively similar to that of raising the marginal income tax rate, so if labor supply is upward sloping, labor supply falls. In the second case (in which savings are treated as if consumed presently), raising the targeted savings rate has no effect on labor supply. In the third case (in which savings are treated as if they are optimal), raising the targeted savings rate raises labor supply, and this is so regardless of the slope of the labor supply curve.

Section 3 offers concluding remarks. It is emphasized that, although none of the assumptions examined here may have high, universal descriptive accuracy, perhaps there is some truth in each of them. Additionally, because they span a substantial range, significant illumination of potential labor supply effects is provided. In any event, the analysis suggests avenues for further analytical and, especially, empirical work.

2. Analysis

2.1. Model

Consider a two-period model in which individuals supply labor, l , in the first period at the wage w and allocate disposable income between consumption, c_1 and c_2 , in periods 1 and 2 respectively. To focus on the issues at hand in a clear and concrete fashion, suppose that individuals' actual utility – which is what they would maximize if they were neoclassical optimizers – takes the following simple form:

$$(1) \quad u(c_1, c_2, l) = \frac{c_1^{1-\rho}}{1-\rho} + \delta \frac{c_2^{1-\rho}}{1-\rho} - z(l),$$

where δ is the subjective discount factor and z measures the disutility of labor effort, with $z' > 0$ and $z'' > 0$. (Note that δ is taken here to be a real trait of individuals' utility, for purposes of assessing social welfare.) The positive constant ρ is the coefficient of relative risk aversion, utility from consumption in each period taking the constant-relative-risk-aversion form (where it is understood that, when $\rho = 1$, utility from consumption c_i is instead given by $\ln c_i$). Individuals are assumed to be subject to a linear income tax with marginal rate t , and savings earn interest at the rate r , so their budget constraint is⁴

$$(2) \quad c_1 + \frac{c_2}{1+r} = (1-t)wl.$$

Neoclassical optimizers' first-order condition for consumption allocation between the two periods (which can be determined by solving the budget constraint (2) for c_2 , substituting it into expression (1), and differentiating) is

$$(3) \quad \frac{\partial u}{\partial c_1} = c_1^{-\rho} - \delta(1+r)c_2^{-\rho} = 0 \quad \text{or} \quad \frac{c_1}{c_2} = \left(\frac{1}{\delta(1+r)} \right)^{1/\rho}.$$

Non-neoclassical savings behavior is described here by the assumption that individuals' savings is set at a specified proportion α of disposable income $y = (1-t)wl$. The motivations, discussed in the introduction, include the possibility that α represents the degree of forced savings through social security (let us suppose an actuarially fair scheme) or employers' retirement plans, or a target savings rate determined perhaps by a default rule in employers'

⁴It would be straightforward to allow a nonzero intercept in the tax schedule, corresponding to a uniform lump-sum tax or transfer or representing virtual income in a nonlinear scheme (although in the latter case some expressions below would have a further term for the change in the marginal tax rate with respect to income). There would be little qualitative effect on the results. For further analysis of how unearned income influences the values of ρ for which labor supply is upward sloping, see Chetty (2006).

retirement plans, the suggestion of investment advisors, or imitation of others' savings behavior.⁵ Accordingly, consumption allocations are $c_1 = (1-\alpha)(1-t)wl$ and $c_2 = \alpha(1-t)wl(1+r)$, and the consumption ratio is simply

$$(4) \frac{c_1}{c_2} = \frac{1-\alpha}{\alpha(1+r)}.$$

If income (labor supply) were taken as fixed, the optimal level of α can be determined by equating the ratio on the right side of expression (4) with that on the right side of the latter expression in (3), which yields

$$(5) \alpha^* = \frac{1}{1 + (1+r) \left(\frac{1}{\delta(1+r)} \right)^{1/\rho}}.$$

For example, in the case in which $\rho = 1$ (so that consumption utility is $\ln c_i$), $\alpha^* = \delta/(1+\delta)$, and the fraction of income consumed in the first period is $1-\alpha^* = 1/(1+\delta)$.

If α had no effect on labor supply, the optimal α would be given by α^* in expression (5). Raising α from 0 to α^* would raise individuals' utility, and further increases beyond α^* would be welfare reducing. The focus of the analysis to follow is on whether this characterization requires modification in light of the possible effect of α on labor supply, which as will be seen depends on how individuals' labor supply is determined. Because of the preexisting distortion due to labor income taxation, positive (negative) effects on labor supply will tend to be desirable (detrimental).⁶

2.2. Future consumption ignored

The first case to be considered is that in which our non-neoclassical individuals treat what they save as if it vanishes, which in many settings corresponds to focusing solely on their paycheck. In this case, the behavioral utility function – i.e., that which is relevant for determining individuals' labor supply but not their actual, experienced utility over time – is

⁵Social security – and all of the other methods of savings postulated here – are taken to be actuarially fair in order to focus on the effect of non-neoclassical savings behavior. Note that, in the presence of a tax on labor income, the assumption that social security is actuarially fair is primarily a matter of notational convenience, for any difference between taxes and expected benefits at a given income level could instead be understood as a component of the labor income tax. Nevertheless, actuarially unfair schemes raise the further question of individuals' perceptions regarding the subject and how these, in turn, should be formulated in a behavioral utility function.

⁶Ordinarily, marginal effects on labor supply have no first-order effect on individuals' utility (the envelope theorem) but do affect government revenue, which has a positive shadow value in the optimal income tax problem. In the present setting, there is the further subtlety that the envelope theorem is applicable to individuals' behavioral utility functions, which diverge from their normative utility functions. In any event, no attempt is made here to undertake a complete welfare analysis of any particular policy.

$$(6) \quad u(c_1, c_2, l) = \frac{c_1^{1-\rho}}{1-\rho} - z(l).$$

The first-order condition for labor supply in this case is obtained by differentiating expression (6) with respect to l , making use of the fact that $c_1 = (1-\alpha)(1-t)wl$:

$$(7) \quad \frac{du}{dl} = c_1^{-\rho} (1-\alpha)(1-t)w - z' = 0.$$

(Primes denote derivatives.) Differentiating this expression with respect to α , rearranging terms, and making appropriate substitutions yields

$$(8) \quad l_\alpha = \frac{(\rho-1)(1-t)wc_1^{-\rho}}{-d^2u/dl^2},$$

where l_α denotes $dl/d\alpha$.

To interpret expression (8), begin with the denominator. We know from individuals' second-order condition (and it can readily be confirmed directly) that $d^2u/dl^2 < 0$, so the denominator as a whole is positive. Accordingly, the sign of the labor supply effect is given by the sign of $\rho-1$.

It is useful to compare expression (8) to that for the effect of a change in the marginal tax rate on labor supply, which is

$$(9) \quad l_t = \frac{(\rho-1)(1-\alpha)wc_1^{-\rho}}{-d^2u/dl^2}.$$

These expressions are nearly identical, the only difference being whether $1-t$ or $1-\alpha$ appears in the numerator, a difference in scaling effects that is apparent from the expression $c_1 = (1-\alpha)(1-t)wl$.

The magnitude of the income effect relative to that of the substitution effect (which is "1" in the term $\rho-1$) is indicated by ρ . When the curvature of utility with respect to consumption (risk aversion) is low – below 1 – the substitution effect dominates, so raising α reduces labor supply, just as does raising t . When curvature is high – above 1 – the income effect dominates

and raising α or t raises labor supply.⁷ To summarize the present case, raising the savings rate α is qualitatively similar in its effect on labor supply to raising the preexisting income tax rate t .

Proposition 1: If individuals save the constant fraction α and, when determining their labor supply, ignore the effect of future consumption on utility, the effect of raising the savings rate α on labor supply is:

- (a) negative, if $\rho < 1$,
- (b) positive, if $\rho > 1$, and
- (c) qualitatively similar to that of raising the labor income tax rate t .

2.3. Future consumption as if it is consumed presently

The second case is that in which individuals treat what they save as if it is consumed presently. In this case, the behavioral utility function is

$$(10) \quad u(c_1, c_2, l) = \frac{y^{1-\rho}}{1-\rho} - z(l).$$

It is obvious from expression (10) that labor supply is independent of α . In any event, if one differentiates the first-order condition for labor supply l implied by expression (10) with respect to α , one can confirm that

$$(11) \quad l_\alpha = 0.$$

Proposition 2: If individuals save the constant fraction α and, when determining their labor supply, treat what they save as if it is presently consumed, the effect of raising the savings rate α on labor supply is:

- (a) zero (independent of the magnitude of ρ), and
- (b) qualitatively different from that of raising the labor income tax rate t .

2.4. Future consumption as if it has the same marginal utility as present consumption

The final and most interesting case to be considered is that in which individuals treat what they save as if it has the same marginal utility as present consumption. Recall that the motivation for this assumption is that individuals behave as though those who set the savings percentage – designers of social security or private retirement plans, investment advisors, others who are imitated – have solved individuals' consumption allocation problem optimally.

⁷See Chetty (2006). This point is also apparent from considering the labor supply curve directly. In this case,

$$l_w = \frac{(1-\rho)(1-\alpha)(1-t)c_1^{-\rho}}{-d^2 u / dl^2}.$$

In this case, it is supposed that individuals behave as if, at the savings rate α , the first-order condition (3) for the allocation of consumption over time holds. A feature of the optimum is that the individual would be indifferent between allocating an additional dollar to present versus future consumption. When the individual contemplates earning an additional dollar, therefore, the total marginal utility it produces when allocated $1-\alpha$ to present consumption and α to savings that fund future consumption – which in fact it will be – will be the same as the marginal utility that the dollar would produce if the individual were to allocate the entire dollar to present consumption. Accordingly, all that individuals need to know is their marginal utility of present consumption. The fraction $1-\alpha$ of their earnings will actually be devoted to present consumption, but they will gross up the resulting marginal utility to the level it would be if the full marginal dollar were consumed presently (at the actual current marginal utility). Thus, the individual behaves as if future consumption per se may be ignored while the marginal utility of present consumption is weighted by the factor $1/(1-\alpha)$. In choosing labor supply, therefore, we can treat individuals *as if* they maximize the behavioral utility function

$$(12) \quad u(c_1, c_2, l) = \frac{1}{1-\alpha} \frac{c_1^{1-\rho}}{1-\rho} - z(l).$$

Note that the derivative of expression (12) with respect to c_1 is $(1/(1-\alpha))c_1^{-\rho}$, consistent with the foregoing explanation.

The first-order condition for labor supply in this case is

$$(13) \quad \frac{du}{dl} = c_1^{-\rho} (1-t)w - z' = 0.$$

This expression may be compared with expression (7) for the first case, in which future consumption is ignored. That expression had the additional term $1-\alpha$, which is absent here because it is offset by the fraction $1/(1-\alpha)$. Differentiating expression (13) with respect to α , rearranging terms, and making appropriate substitutions yields

$$(14) \quad l_\alpha = \frac{\frac{\rho}{1-\alpha} (1-t)w c_1^{-\rho}}{-d^2u/dl^2}.$$

To interpret expression (14), it is useful to compare expression (8) for l_α in the case in which future consumption is ignored. There are two differences. Most obviously, the “-1” in the numerator of expression (8) – corresponding to the substitution effect – is absent from expression (14). Second, the magnitude of the income effect (ρ) in expression (14) is higher relative to that in expression (8) because here it is grossed up by the factor $1/(1-\alpha)$.

Both differences are in accord with the intuition behind the behavioral utility function (12). Because individuals in the present case behave as if α was selected in light of their

marginal utility of future consumption – taken to be equated with the marginal utility of present consumption – a higher α , which reduces present consumption and thus raises the marginal utility of present consumption, indicates a greater value to raising labor supply (a sort of income effect, and one that is applicable both to present consumption itself and, by implication, to what is saved). Furthermore, there is no analogue to the substitution effect because in the present formulation individuals do not perceive anything to be taken from them when savings are raised. Because the first effect is unambiguously positive and the second is zero, the net effect is positive as well.

Another way to see these points is to consider a blended case, intermediate between the first and third cases, in which $\theta \in [0, 1]$ indicates the weight on the present case. The resulting behavioral utility function may be written as

$$(15) \quad u(c_1, c_2, l) = \left(\frac{1 - (1 - \theta)\alpha}{1 - \alpha} \right) \frac{c_1^{1-\rho}}{1 - \rho} - z(l).$$

As before, one can differentiate the first-order condition for labor supply implied by this utility function with respect to α , rearrange terms, and make appropriate substitutions, which yields

$$(16) \quad l_\alpha = \frac{\left(\rho \left(1 + \theta \frac{\alpha}{1 - \alpha} \right) - (1 - \theta) \right) (1 - t) w c_1^{-\rho}}{-d^2 u / dl^2}.$$

Consistent with the previous suggestions, a higher level of θ – that is, more weight on the present case – implies that the income effect is stronger and the substitution effect is weaker relative to the first case, in which the targeted savings rate acts like a tax.⁸ Compare expression (8). At its highest value, $\theta = 1$, the income effect is weighted by $1/(1 - \alpha)$ and the substitution effect is zero, as in expression (14).

Proposition 3: If individuals save the constant fraction α and, when determining their labor supply, treat future consumption as if it has the same marginal utility as present consumption, the effect of raising the savings rate α on labor supply is:

- (a) positive (independent of the magnitude of ρ), and
- (b) qualitatively different from that of raising the labor income tax rate t .

3. Conclusion

The effect on labor supply of raising savings rates of non-neoclassical individuals

⁸Supposing that labor supply is upward sloping so that raising the savings rate reduces labor supply in the first case, it is still true that the income effect may substantially (although not completely) offset the substitution effect, so it may not take very much weight on the present case for the overall labor supply effect to be positive.

depends on how such individuals conceive of the savings that they undertake. When they ignore future consumption entirely, raising the targeted savings rate acts much like raising the labor income tax rate. When they instead treat their savings as if consumed presently, raising the targeted savings rate has no effect on labor supply. Finally, in the interesting case in which they treat their savings rule as implicitly reflecting an optimal allocation of consumption over time, raising the targeted savings rate increases labor supply without regard to the slope of the labor supply curve.

These results are important because changes in labor supply have significant, first-order effects on welfare due to the preexisting distortion on account of labor income taxation. For example, if the first model is correct, even actuarially fair social security may impose a high distortionary cost, whereas if the third model is correct, social security may, in addition to improving individuals' intertemporal consumption allocations, produce a substantial offset to the distortionary cost of the income tax.⁹

Important extensions of the present analysis would include the introduction of heterogeneity regarding both savings behavior and the behavioral utility function that determines labor supply. Additional aspects of complexity that motivate the satisficing approach, such as uncertainty about future earnings, health, and longevity, should also be considered. In addition, the problem could be extended to multiple periods and, in particular, to embrace the decision when to retire, perhaps supposing that, as retirement nears, individuals pay more explicit attention to savings role in supporting future consumption. Finally, note that the approach undertaken here with regard to paternalistic savings might also be employed to study paternalistic insurance, again whether mandatorily provided by the government or employers or voluntarily purchased on the basis of advice or imitation rather than explicit optimization. In an analogous model, it seems plausible that similar results would be obtained.¹⁰

As mentioned in the introduction, it is not claimed that any of the three cases examined here depict reality, or, if so, for what proportion of individuals. Perhaps other assumptions about labor supply are more apt, at least for some savers. Accordingly, future empirical work should not only continue to examine savings behavior per se but should also begin to explore how non-neoclassical savings behavior relates to labor supply decisions. An implication of the present analysis is that examination of the labor supply effects of social security, pensions, and other provisions that are aimed at savings may provide a basis for choosing among the sorts of models considered here.

⁹See also note 6 on the implicit assumption in the present analysis that the savings schemes examined are actuarially fair. Social security raises other issues, especially regarding individuals' information and beliefs about future benefits. For example, concerns about the misunderstanding of benefits motivate suggestions to provide clearer statements to individuals of how their future benefits accrue as they work and also proposals to pre-fund social security in the form of defined-contribution plans. Also relevant is political uncertainty, on which see Dominitz, Manski, and Heinz (2003).

¹⁰Just as optimal savings behavior equates marginal utilities across time, optimal insurance equates marginal utilities across states, so even the third, somewhat more refined, variation of the behavioral utility function for labor supply could be applied directly to insurance.

References

- Aguiar, Mark, and Hurst, Erik.** “Consumption versus Expenditure.” *Journal of Political Economy*, 2005, 113 no. 5: 919-948.
- Banks, James, Blundell, Richard, and Tanner, Sarah.** “Is There a Retirement-Savings Puzzle?” *American Economic Review*, 1998, 88, no. 4:769–88.
- Bernheim, B. Douglas.** “Personal Saving, Information, and Economic Literacy: New Directions for Public Policy.” In C.E. Walker, M. Bloomfield, and M. Thorning, eds., *Tax Policy for Economic Growth in the 1990s*. Washington, D.C.: American Council for Capital Formation, 1994: 53-78.
- Bernheim, B. Douglas.** “Taxation and Saving.” In Alan J. Auerbach and Martin Feldstein, eds., *Handbook of Public Economics*. Vol. 3. Amsterdam: North-Holland, 2002: 1173-1249.
- Bernheim, B. Douglas, and Rangel, Antonio.** “Behavioral Public Economics: Welfare and Policy Analysis with Non-Standard Decision Makers.” In P. Diamond and H. Vartiainen, eds., *Behavioral Economics and Its Applications*. Princeton: Princeton University Press, 2007: 7-77.
- Bernheim, B. Douglas, Skinner, Jonathan, and Weinberg, Steven.** “What Accounts for the Variation in Retirement Wealth among U.S. Households?” *American Economic Review*, 2001, 91 no. 4:832-857.
- Beshears, John, Choi, James J., Laibson, David, and Madrian, Brigitte C.** “The Importance of Default Options for Retirement Savings Outcomes: Evidence from the United States.” National Bureau of Economic Research, NBER Working Papers: No. 12009, 2006a.
- Beshears, John, Choi, James J., Laibson, David, and Madrian, Brigitte C.** “Simplification and Saving.” National Bureau of Economic Research, NBER Working Papers: No. 12659, 2006b.
- Carroll, Gabriel D., Choi, James J., Laibson, David, Madrian, Brigitte C., and Metrick, Andrew.** “Optimal Defaults and Active Decisions.” *Quarterly Journal of Economics*, 2009, 124 no. 4: 1639-1674.
- Chetty, Raj.** “A New Method of Estimating Risk Aversion.” *American Economic Review*, 2006, 96 no.5:1821-1834.
- Choi, James J., Laibson, David, Madrian, Brigitte C., and Metrick, Andrew.** “For Better or Worse: Default Effects and 401(k) Savings Behavior.” In David A. Wise, ed., *Perspectives on the Economics of Aging*. Chicago: University of Chicago Press, 2004. 81-121.
- Diamond, Peter A.** *Social Security Reform*. Oxford: Oxford University Press, 2002.

- Diamond, Peter A.** “Social Security.” *American Economic Review*, 2004, 94 no. 1:1-24.
- Dominitz, Jeff, Manski, Charles F., and Heinz, Jordan.** “Will Social Security Be There for You?": How Americans Perceive Their Benefits.” National Bureau of Economic Research, NBER Working Papers: No. 9798, 2003.
- Engen, Eric M., Gale, William G., and Uccello, Cori E.** “The Adequacy of Household Saving.” *Brookings Papers on Economic Activity*, 1999, no. 2:65-187.
- Feldstein, Martin, and Liebman, Jeffrey B.** “Social Security.” In Alan J. Auerbach and Martin Feldstein, eds., *Handbook of Public Economics*. Vol. 4. Amsterdam: North-Holland, 2002: 2245-2324.
- Gokhale, Jagadeesh, Kotlikoff, Laurence J., and Warshawsky, Mark J.** “Comparing the Economic and Conventional Approaches to Financial Planning.” In Laurence J. Kotlikoff, ed., *Essays on Saving, Bequests, Altruism, and Life-Cycle Planning*. Cambridge: MIT Press, 2001. 489-560.
- Johnson, Stephen, Kotlikoff, Laurence J., and Samuelson, William.** “Can People Compute? An Experimental Test of the Life Cycle Consumption Model.” In Laurence J. Kotlikoff, ed., *Essays on Saving, Bequests, Altruism, and Life-Cycle Planning*. Cambridge: MIT Press, 2001. 335-385.
- Kaplow, Louis.** “Myopia and the Effects of Social Security and Capital Taxation on Labor Supply.” National Bureau of Economic Research, NBER Working Papers: No. 12452, 2006.
- Kaplow, Louis.** *The Theory of Taxation and Public Economics*. Princeton: Princeton University Press, 2008.
- Kotlikoff, Laurence J., Spivak, Avia, and Summers, Lawrence H.** “The Adequacy of Savings.” *American Economic Review*, 1982, 72 no. 5:1056-1069.
- Laibson, David.** “Hyperbolic Discount Functions, Undersaving, and Savings Policy.” National Bureau of Economic Research, NBER Working Papers: No. 5635, 1996.
- Laibson, David.** “Golden Eggs and Hyperbolic Discounting.” *Quarterly Journal of Economics*, 1997, 112 no. 2:443-477.
- Liebman, Jeffrey B., and Zeckhauser, Richard J.** “Schmeduling.” Harvard University, Kennedy School of Government, 2004.
- Madrian, Brigitte C., and Shea, Dennis F.** “The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior.” *Quarterly Journal of Economics*, 2001, 116 no. 4:1149-1187.
- Moore, James, and Mitchell, Olivia S.** “Projected Retirement Wealth and Saving Adequacy.”

In Olivia S. Mitchell, P. Brett Hammond, and Anna M. Rappaport, eds., *Forecasting Retirement Needs and Retirement Wealth*. Philadelphia: University of Pennsylvania Press, 2000: 68-94.

Scholz, John Karl, Seshadri, Ananth, and Khitatrakun, Surachai. “Are Americans Saving ‘Optimally’ for Retirement?” *Journal of Political Economy*, 2006, 114 no. 4:617-643.

Smith, Sarah. “The Retirement–Consumption Puzzle and Involuntary Early Retirement: Evidence from the British Household Panel Survey.” *Economic Journal*, 2006, 116 no. 510:C130-C148.

Thaler, Richard H., and Benartzi, Shlomo. “Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving.” *Journal of Political Economy*, 2004, 112 no. 1, pt. 2:S164–S187.