Macroeconomic Dynamics, **16**, 2012, 230–239. Printed in the United States of America.
doi:10.1017/S1365100510000064

A CLOSED-FORM SOLUTION TO A MODEL OF TWO-SIDED, PARTIAL ALTRUISM

ANA FERNANDES
University of Bern

This paper presents a closed-form characterization of the allocation of resources in an overlapping generations model of two-sided, partial altruism. Three assumptions are made: (i) parents and children play Markov strategies, (ii) utility takes the CRRA form, and (iii) the income of children is stochastic but proportional to the saving of parents. In families where children are rich relative to their parents, saving rates—measured as a function of the family's total resources—are higher than when children are poor relative to their parents. Income redistribution from the old to the young, therefore, leads to an increase in aggregate saving.

Keywords: Closed Form, Value Function, Partial Altruism, Markov Strategies

1. INTRODUCTION

The ubiquitous presence of the family in many aspects of decision-making—from human capital investments by the young [Becker (1993)] to the role of financial transfers among kin in the distribution of wealth and persistence of inequality [Becker and Tomes (1979), Loury (1981)]—has long been recognized in the literature. Models of the family are especially important in analyzing topics that involve resource redistribution across generations, such as the implementation of social security programs; altruism has been a recurrent paradigm in this context [see, e.g., Barro (1974), Becker (1974), Laitner (1979, 1988)].

Dynamic models of the altruistic family typically involve generalizations of the overlapping-generations framework.[1] The utility of old agents at a given moment in time no longer depends solely on their own consumption path; it is affected by the well-being of their descendants. The models used range from assuming very strong family ties—people value the utility of their descendants as much as their own—to considering less extreme views of affection across kinship, whereby the utility of a family member is somewhat discounted relative to one's own.[2] In this literature, altruism is generally considered to be one-sided, with parents

This paper draws on chapter 7 of my thesis, Fernandes (1999). I thank Robert E. Lucas, Jr., Fernando Alvarez, Gary Becker and Sherwin Rosen for guidance. Financial support from *Banco de Portugal*, *Fundação Para a Ciência e a Tecnologia* and *Fundación Ramón Areces* is gratefully acknowledged. Address correspondence to: Ana Fernandes, Department of Economics, University of Bern, Schanzeneckstrasse 1, 3001 Bern, Switzerland; e-mail: ana.fernandes@vwi.unibe.ch.

© 2012 Cambridge University Press 1365-1005/12

https://doi.org/10.1017/S1365100510000064 Published online by Cambridge University Press

caring about children and children caring about their children—but not about their parents. One important reason for this one-sidedness has been the analytical difficulty of obtaining tractable characterizations of the family's choices otherwise. The roots of the problem are sketched below.[3]

Altruistic family members will donate income to less wealthy relatives to equalize the marginal utility from consumption. This equalization will not be full in the case of partial altruism, as the utility of relatives is less valued than one's own. In a static environment of two-sided altruism, this implies that transfers will flow from parents to children when parents are wealthy relative to their children, from children to parents otherwise, and that there will be an intermediate range where no transfers take place, where parent and child have similar incomes and where each would like to receive transfers from the other family member. Consider now the utility the parent derives from consumption. Let us fix the child's income and examine what happens as the parent's income increases. For low income values, the parent receives a transfer from the child. An additional unit of income is "taxed" by the child in the form of a transfer reduction. As parental income rises, transfers decrease until they become exactly zero. At this point, additional parental income reverts fully to parental consumption, as the child cannot impose a negative transfer on the parent. The marginal utility from consumption therefore increases locally in a discrete fashion (seen from the point of view of the parental utility from consumption, the utility function has an upward kink at the point where transfers from the child are exactly zero).[4] This causes concavity of the value function to break down and prevents standard theorems—to demonstrate existence of a value function summarizing parental future utility—from being used. Instead, people have resorted either to full altruism (no kink exists, because transfers are always positive and fully equalize consumption across family members) or to one-sided altruism (there is a kink at the point where the parent's income is high enough to trigger transfers to a poorer child but the value function remains concave).[5]

This paper presents a closed-form characterization of the allocation of resources (consumption and saving) in an overlapping generations (OLG) model of two-sided, partial altruism. Under the assumptions that (i) parents and children play Markov strategies, (ii) utility takes the CRRA form, and (iii) the income of children is stochastic but proportional to the saving of parents, it is possible to solve explicitly for optimal consumption, transfers (from parents to children or from children to parents), and saving, as well as for the value function summarizing the utility parents derive from the future. As in the static setup, partial altruism causes transfers to flow from wealthy relatives to relatively poorer family members, but not to take place at all whenever incomes are sufficiently close. It is shown that the saving decisions of young dynasty members are a function of the total income of the family. In fact, when interfamily transfers flow from the old to the young (bequests), the family saves a smaller share of its total resources compared to the case where the young provide the older cohort with transfers (gifts). The consumption of the old cohort is also a higher fraction of the family's total income when bequests are positive compared to the old's consumption share when the

https://doi.org/10.1017/S1365100510000064 Published online by Cambridge University Press

young provide positive gifts. When intergenerational transfers are zero, family members act independently and kinship is an irrelevant consideration for economic decisions such as saving or consumption.

The paper is organized as follows. Section 2 describes the environment. Section 3 provides a parametric example. Section 4 concludes.

2. AN OLG MODEL WITH TWO-SIDED, PARTIAL ALTRUISM

2.1. Preferences

The environment is that of a standard OLG model with partial altruism [see Laitner (1988)]. Each family member lives for two periods. When young, the child coexists with the parent. In the second period, she raises her own child. In the following period, that child will raise her grandchild. The extended family is therefore infinitely lived.

Parent and child value each other's utility but by less than they value their own. Consider a parent–child pair. Let V_p denote the parent's forward-looking utility associated with decisions concerning the present date and the entire future. Similarly, let V_c denote the child's lifetime utility, and c_p and c_c stand for parent's and child's consumption, respectively. For numbers $\beta, \lambda_H, \lambda_L \in (0, 1)$ and $\lambda_H > \lambda_L$, consider the following system:

$$V_p = \lambda_H u(c_p) + \lambda_L V_c, \tag{1}$$

$$V_c = \lambda_H [u\,(c_c) + \beta V_{p'}] + \lambda_L V_p. \tag{2}$$

The function $V_{p'}$ represents the utility accruing to the young agents in the second period of their lives, because they will be old then and therefore assume the role of parents. The assumption $\lambda_H > \lambda_L$ reflects the partial nature of altruism: individuals value their direct utility more than they value that of their relatives.

Solving the system (1) and (2), we get

$$V_p = \frac{\lambda_H}{1-\lambda_L^2} u(c_p) + \frac{\lambda_H \lambda_L}{1-\lambda_L^2} [u\,(c_c) + \beta V_{p'}], \tag{3}$$

$$V_c = \frac{\lambda_H \lambda_L}{1-\lambda_L^2} u(c_p) + \frac{\lambda_H}{1-\lambda_L^2} [u\,(c_c) + \beta V_{p'}]. \tag{4}$$

We may define $\tilde{\lambda}_H \equiv \lambda_H/(1-\lambda_L^2)$ and $\tilde{\lambda}_L \equiv \lambda_H\lambda_L/(1-\lambda_L^2)$. There is no guarantee that $\tilde{\lambda}_H$ and $\tilde{\lambda}_L$ are in the interval $(0, 1)$ and, in particular, it is also possible that the new discount factor in V_c, $\beta\lambda_H/(1-\lambda_L^2)$, is not in this interval. It may be convenient, therefore, to impose restrictions on λ_H and λ_L so that the discount factor lies strictly in the unit interval. But the general idea that each agent weighs more the utility from own consumption than the utility of other family members goes through, in (3) and (4).

Also, parents discount the future more than their children, as the term multiplying β in V_p, $\lambda_H\lambda_L/(1-\lambda_L^2)$, is strictly smaller than the corresponding term in

https://doi.org/10.1017/S1365100510000064 Published online by Cambridge University Press

V_c, $\lambda_H/(1-\lambda_L^2)$. The reason for this has to do with the fact that they value the future indirectly, through their child's future consumption, which they value less than their own consumption.

In what follows, I will use the preference format of (3) and (4) to motivate the functional forms

$$V_p = \lambda u(c_p) + (1-\lambda)\,[u\,(c_c) + \beta V_{p'}], \tag{5}$$

$$V_c = (1-\lambda)\,u(c_p) + \lambda[u\,(c_c) + \beta V_{p'}], \tag{6}$$

where $\lambda \in (0.5, 1)$.

2.2. A Dynasty

Equations (5) and (6) characterize the preferences of old and young agents, respectively. Let us now turn to other aspects of the extended family. Nature acts at the beginning of every period of this infinite-horizon game by drawing the income of young agents, z_t. Parental income is assumed nonstochastic, for simplicity. Old agents in period t have income $a_t R$, where a_t represents assets accumulated from period $t-1$, which yields the exogenous return R.[6]

Each agent moves in two adjacent periods. When he is born, he observes his income and his parent's resources and decides on a nonnegative transfer to the parent, g_t, and on the amount of savings to accumulate for the following period, a_{t+1}. At time t, the parent (who was born in period $t-1$) observes z_t and decides how to split the resources $a_t R$ into own consumption and a nonnegative transfer to the child, b_t. Both parent and child know the entire past history of endowment realizations, transfers, and asset accumulation decisions made by other dynasty members. Their choices of gifts, savings, and bequest—the agents' strategies—map the entire set of past decisions and income realizations, as well as the income realizations of the period in which they move, into the nonnegative real numbers. In addition, strategies are additionally constrained to verify $b_t, g_t \geq 0$, $g_t + a_{t+1} \leq z_t + b_t$, and $b_t \leq a_t R + g_t$.[7]

At time t, for fixed interest rate R, given resources a_t, z_t and decisions b_t, g_t, a_{t+1}, consumption is given by

$$c_{pt} = a_t R - b_t + g_t, \tag{7}$$

$$c_{ct} = z_t + b_t - g_t - a_{t+1.} \tag{8}$$

2.3. Markov Perfect Equilibria

To characterize subgame-perfect equilibria, one could use the ideas of self-generation and factorization, from Abreu et al. (1986, 1990), and determine the set of lifetime discounted utilities associated with subgame-perfect strategies of this infinite game. A more restrictive procedure, which also confines attention to

https://doi.org/10.1017/S1365100510000064 Published online by Cambridge University Press

Markov perfect strategies, would be to find a solution to a functional equation, as follows.

Consider a space of functions with domain in $\mathbf{R}^2$, $F(\mathbf{R}^2)$. Suppose that $V \in F(\mathbf{R}^2)$ satisfies the functional equation

$$V(a, z) = \max_b \left\{ \lambda u\,(aR - b + g) + (1 - \lambda)\left[u(z + b - g - a') + \beta \int V(a', z')\mu(dz') \right] \right\} \tag{9}$$

subject to

$$g, a' \in \arg\max \left\{ (1 - \lambda)\, u\,(aR - b + g) + \lambda \left[u(z + b - g - a') + \beta \int V(a', z')\mu(dz') \right] \right\}. \tag{10}$$

Equations (9) and (10) characterize *one* Markov perfect equilibrium. The gain from using this very restrictive approach will be shown in the next section: the parametric example solved below enables a complete characterization of the function $V(\cdot)$ as well as of the optimal choices (b, g, a') as functions of the state (a, z).

3. A PARAMETRIC EXAMPLE

In this section, I assume that the income of the child is stochastic but proportional to her parent's saving. Say the parent saves a, which earns the deterministic return R. The child's resources are the product of the stochastic shock z and the parent's saving a. Therefore, changes in a affect the income of both parent and child in the same direction, something that would tend to reduce potential disincentives to save; these could arise from the fact that, from altruism, children know they will not appropriate the totality of their saved income if their own children turn out to be poor, and, by saving more, the current generation also reduces the likelihood that their children will transfers resources to them in the future. Together with CRRA utility, the fact that changes in a do not affect the ratio of parent's to child's resources suggests that decision rules will be proportional to a, therefore eliminating the potential lack of global concavity of $V(\cdot)$ with respect to a. This will indeed be the case, as shown below.

Is the assumption of stochastic proportionality of incomes plausible? Although its analytical convenience is undeniable, it is not as restrictive as it may at first appear. In fact, the overlapping-generations model as it now stands could be calibrated to reproduce a certain degree of mean reversion around a by allowing the distribution of the next period's shocks, z', to depend on today's shock realization, z, as in a Markov process.[8] As long as the distribution of future shocks does not depend on current or past values of a, the analysis goes through.

https://doi.org/10.1017/S1365100510000064 Published online by Cambridge University Press

TABLE 1. Optimal consumption and savings

	Low z $b > 0$	Medium z $b = g = 0$		High z $g > 0$
		z_l	z_h	
c_p	$\lambda[1-\beta(1-\lambda)]a(z+R)$	Ra		$\frac{(1-\lambda)[1-\beta(1-\lambda)]}{1-\beta+2\beta\lambda}a(R+z)$
c_c	$(1-\lambda)(1-\beta(1-\lambda))a(z+R)$	$\frac{1-\beta(1-\lambda)}{1+\beta\lambda}az$		$\frac{\lambda[1-\beta(1-\lambda)]}{1-\beta+2\beta\lambda}a(z+R)$
a'	$\beta(1-\lambda)a(z+R)$	$\frac{\beta}{1+\beta\lambda}az$		$\frac{\beta\lambda}{1-\beta+2\beta\lambda}a(z+R)$

Given resources a and the realization of z, once transfers b, g and current savings a' are determined, consumption of family members is now given by

$$c_p = aR - b + g \quad \text{and} \quad c_c = az + b - g - a'.$$

Within CRRA, I choose logarithmic utility, as this particular choice provides more insightful results. With the choice $u(c) = \log(c)$, I guess that $V(a, z) = f(z, R) + B\log(a)$, a strictly concave function in a. The guess is verified for $B = 1/[1-\beta(1-\lambda)]$.[9] With the properties of V given above, it is easy to see that, for small values of z (relative to R), transfers flow from parents to children; the converse happens if z is sufficiently high, whereas intermediate realizations of this random variable are associated with no transfers taking place at all. The results are presented in Table 1.[10]

The values of z that divide the transfer regimes are given by

$$z_l = \frac{(1-\lambda)}{\lambda}\left[\frac{1+\beta\lambda}{1-\beta(1-\lambda)}\right]R \quad \text{and} \quad z_h = \frac{\lambda}{1-\lambda}\left[\frac{1+\beta\lambda}{1-\beta(1-\lambda)}\right]R.$$

Several things are worth noting, concerning the results in Table 1. First, if we consider the extreme case $\lambda = 0.5$ (full, rather than partial, altruism), the table collapses into a single column: transfers will always be positive and completely equalize consumption across family members for all values of z and R. The fact that altruism is partial, $\lambda \in (0.5, 1)$, ensures that $z_l < z_h$ and gives rise to the three different transfer regimes.

As in a static context, transfers flow from wealthy family members to less well-off relatives. When $z < z_l$, the parent provides positive transfers to the child. For $z > z_h$, the child provides gifts to the parent. For intermediate realizations of z, $z_l < z < z_h$, no transfers take place.

Perhaps the most interesting aspect of this problem is the fact that optimal consumption and savings decisions are generally functions of the family's total resources, $(R + z)a$. In fact, we have $c_j = c_j(a, R, z)$, for $j = p, c$. Similarly, $a' = a'(a, R, z)$. It is only when $z \in (z_l, z_h)$ that individual consumption is a

https://doi.org/10.1017/S1365100510000064 Published online by Cambridge University Press

function of that family member's income, and savings depends on the child's income alone. When transfers are positive (either low or high realizations of the shock z), optimal consumption and savings are expressed as a constant fraction of the family's total resources, $a(R + z)$.

In regimes with either $b > 0$ or $g > 0$, the family member providing the transfers is solving a Pareto problem of maximizing a weighted sum of utilities subject to the total resources of the family. Comparing the low-z with the high-z regimes, we see that the parent's marginal propensity to consume out of total resources is highest when the parent himself is providing the transfers. In fact, parental consumption equals the fraction $\lambda[1 - \beta(1 - \lambda)]$ of total family income when parental transfers are positive and only $(1-\lambda)[1-\beta(1-\lambda)]/(1-\beta+2\beta\lambda)$ when the child provides positive transfers to the parent. This is not a surprising outcome: the transfer giver places the highest weight on his own direct well-being. Likewise, the child's marginal propensity to consume and save out of the family's total resources is highest when the child provides positive transfers to the parent. When gifts are positive, the saving rate is $\beta\lambda/(1 - \beta + 2\beta\lambda)$, compared to only $\beta(1 - \lambda)$, when parents transfer resources to children.

Consider now the Ricardian equivalence experiment of *lump-sum* income redistribution across generations. Say that the government manages to transfer income from the parent and give it to the child, while keeping the family's total resources $a(R + z)$ constant. If this transfer causes a change in transfer regime to take place (either by moving the family from a positive bequest region to a no-transfer or to a positive gift regime), then it will lead to an increase in saving and in the child's consumption at the expense of a reduction in the consumption of the old generation. This nonneutrality is also a property of the static models of partial altruism. If the income transfer leaves the family in the same positive transfer regime (either bequests or gifts), then nothing changes: consumption and saving remain unchanged and Ricardian equivalence prevails.

Although it is not obvious from the table, saving is a constant fraction of the saver's after-transfer income. In the same spirit of *lump-sum* redistribution as before, consider the comparison across different families whose total income is the same, $a(R + z)$, but where ownership of resources differs. Rank families in increasing order of the child's income, az. When bequests are provided, the child's after-transfer income, $az + b$, is less than the corresponding amount she will get when gifts are given, $az - g$. As a consequence, total saving will also be lower in the positive bequest regime. Despite the constant saving rate measured as a function of the after-transfer income, the examination of individual saving data would suggest that the saving rate of the child—measured in terms of her income az—is monotonically decreasing in her income [except in the range (az_l, az_h), where it would stay constant].

How does saving vary with the intensity of altruism—i.e., as λ decreases? In the positive bequest region, lower λ makes the old care more about the young and raises the amount of the bequest. The young receive a larger transfer and they raise saving. In this regime, altruism raises the saving rate. In the positive gift

https://doi.org/10.1017/S1365100510000064 Published online by Cambridge University Press

TABLE 2. Saving rates and limits of transfer intervals as functions of λ

λ	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95
s_h/s_l	1	1.12	1.26	1.45	1.69	2.03	2.55	3.4	5.11	10.24
z_h/z_l	1	1.49	2.25	3.45	5.44	9	16	32.11	81	361

region, lower λ indicates that the young value the well-being of the old more and, consequently, provide a greater gift to them. Saving is therefore reduced. In the no-transfer regime, saving rates are decreasing functions of λ, as this parameter also measures the intensity of altruism regarding the children of the current savers.

These changes in saving and consumption choices across transfer regimes introduce an element of heterogeneity across families with the same total income but where the distribution of resources varies. One consequence of this analysis is, therefore, the breakdown of the representative family or household paradigm. Another implication of the analysis is the separation between income inequality and the distribution of well-being. The fact that transfers partly equalize differences in living standards implies that the distribution of consumption across individuals will have lower variance than the distribution of income; in fact, transfers will transform pairs of extreme income observations (very high and very low) into pairs of more balanced consumption outcomes.

Let $s_h = s_h(\lambda)$ denote the saving rate out of total familial resources when $z > z_h$. Define $s_l = s_l(\lambda)$ to be the saving rate for realizations of z below z_l. Table 2 computes the ratio s_h/s_l for varying values of λ in the interval (0.5, 1). It also shows how the ratio z_h/z_l varies with λ.

From Table 2, we see that small departures from the full altruism benchmark of $\lambda = 0.5$ have important consequences for the saving rate as well as for the income differential needed to trigger transfers. With $\lambda = 0.65$, for example, the saving rate associated with positive gifts is roughly one and one-half times as high as the one prevailing under positive bequests; also, a child providing to her progeny earns now approximately three and one-half times as much as a child receiving parental transfers. These magnitudes naturally increase with λ, as we move further away from full altruism.

4. CONCLUSION

This paper provided a closed-form solution to a model of two-sided, partial altruism. The most important assumption in overcoming the typical breakdown of value-function concavity—a consequence of partial altruism—was the proportionality of parent's and child's income. As argued, this seemingly restrictive assumption is compatible with conventional income processes and well suited to replicate, for example, mean reversion in intergenerational income.

https://doi.org/10.1017/S1365100510000064 Published online by Cambridge University Press

The analysis draws attention to an interesting tension between intergenerational mobility and income inequality at a point in time. Partial altruism implies that families pooling resources are those in which income is relatively dissimilar across their members. Those families whose differences in income are largest will reallocate resources in order to reduce differences in wellbeing. Consequently, partial altruism suggests that accurate measurement of wealth inequality must take into account the degree of persistence of income across generations. In fact, individual well-being depends not only on individual wealth but also on how the wealth of a particular individual ranks among that of his family members. Quantifying the importance of altruism in the reduction of income inequality is the object of ongoing research.

NOTES

1. Samuelson (1958).

2. See Laitner (1997) for a survey.

3. As is common in the altruism literature, attention is restricted to the case where the utility of an agent at a moment in time is a weighted sum of his own direct utility from consumption and that of his relatives (parent and/or child). See Laitner (1997).

4. This is trivially true of the parent's own marginal utility from consumption. More importantly, it is also true if one considers the parent's total marginal utility from consumption, a weighted sum of his direct marginal utility from consumption and that of his child, as long as the weight placed on his own direct utility exceeds the one placed on the child's.

5. Laitner (1988) considers two-sided partial altruism; there, parents are allowed to invest in lotteries and this restores concavity to the value function.

6. All the results in this paper are robust to having parental income be stochastic, instead, as long as the ratio of parental to child's income remains independent of the parent's actions.

7. The action space for the parent—$b_t \in [0, a_t R + g_t]$—for example, depends on thetransfer received from the child. Also, the set of feasible values for the child's choices of gift and savings depend on the parent's action b_t. The utility of all family members is still a well-defined object, for each feasible action b_t, g_t, and a_{t+1}, in all time periods t. In fact, it is not possible to write down action sets that are independent of the players' actions, in the more general formulation of this game, without imposing artificial limits on the transfer and savings choices. In the subset of Markov perfect equilibria, discussed below, we could refine the limits on transfers as follows. Optimal parental transfers, b_t, will be contained in $[0, a_t R]$, and the child's optimal choice of gift, g_t, will only take values in the interval $[0, z_t]$. However, the child's choice of savings will depend on the parental bequest, even in the subset of Markov perfect strategies, because positive parental transfers will induce higher savings. One could, as in Laitner (1988), impose that savings are chosen *after* transfer decisions are implemented. Given the sequential nature of choices, savings would then be a function of the child's total income after transfers: $z_t + b_t$. This modification would have no implications for Markov perfect strategies or the results of the parametric example of Section 3, and the simultaneous choice of transfers and savings is therefore maintained here.

8. See Mulligan (1999) for a comparison of different theories of intergenerational income dynamics.

9. For general CRRA utility, with $u(c) = c^{1-\sigma}/(1-\sigma)$, the guess $V(a, z) = f(z) + Ba^{1-\sigma}/(1-\sigma)$ is verified for the same value of B as above.

10. The function $f(z, R)$ falls in one of three ranges, corresponding to the possibilities $z \leq z_l$, $z \in (z_l, z_h)$, and $z > z_h$. Say that $f^1(z, R)$ corresponds to the functional form associated with the lowest interval of z values, and that $f^2(\cdot)$ and $f^3(\cdot)$ are those associated with the other intervals. Simple but tedious algebra shows that $f^1(z_l, R) = f^2(z_l, R)$, and $f^2(z_h, R) = f^3(z_h, R)$.

https://doi.org/10.1017/S1365100510000064 Published online by Cambridge University Press

REFERENCES

Abreu, Dilip, David Pearce, and Ennio Stacchetti (1986) Optimal cartel equilibria with imperfect monitoring. *Journal of Economic Theory* 39(1), 251–269.

Abreu, Dilip, David Pearce, and Ennio Stacchetti (1990) Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica* 58(5), 1041–1063.

Barro, Robert J. (1974) Are government bonds net wealth? *Journal of Political Economy* 82(6), 1095–1117.

Becker, Gary S. (1974) A theory of social interactions. *Journal of Political Economy* 82(6), 1063–1093.

Becker, Gary S. (1993) *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, 3d ed. Chicago: University of Chicago Press.

Becker, Gary S. and Nigel Tomes (1979) An equilibrium theory of the distribution of income and intergenerational mobility. *Journal of Political Economy* 87(6), 1153–1189.

Fernandes, Ana (1999) *Familial Preferences and Economic Choices: Does Distribution Matter?* Ph.D. Dissertation, University of Chicago.

Laitner, John (1979) Bequests, golden-age capital accumulation and government debt. *Economica* 46, 403–414.

Laitner, John (1988) Bequests, gifts and social security. *Review of Economic Studies* 55(2), 275–299.

Laitner, John (1997) Intergenerational and interhousehold economic links. In M.R. Rosenzweig and O. Stark (eds.), *Handbook of Population and Family Economics*, Vol. 1A, Chap. 5. North-Holland.

Loury, Glenn C. (1981) Intergenerational transfers and the distribution of earnings. *Econometrica* 49(4), 843–867.

Mulligan, Casey B. (1999) Galton versus the human capital approach to inheritance. *Journal of Political Economy* 107(6), S184–224.

Samuelson, Paul A. (1958) An exact consumption-loan model of interest with or without the social contrivance of money. *Journal of Political Economy* 66(6), 467–482.

https://doi.org/10.1017/S1365100510000064 Published online by Cambridge University Press