Front. Econ. China 2014, 9(1): 6–24
DOI 10.3868/s060-003-014-0002-4

RESEARCH ARTICLE

Larry Karp

Overlapping Generations and Environmental Policy: An Introduction

Abstract A small but growing body of literature uses overlapping generations (OLG) models to study environmental policy for long-lived problems such as climate change. An OLG model, unlike the infinitely lived representative agent model, distinguishes between impatience with respect to one's own future utility, and attitudes toward successors' utility. I discuss the problem of time inconsistency, the role of Markov perfection, and show that a class of OLG models can be studied using methods developed to analyze models of non-constant discounting. An example illustrates the techniques and determines the conditions under which, in equilibrium, there is under-investment or over-investment in natural capital.

Keywords overlapping generations, climate policy, time consistency, Markov perfection, under-investment

JEL Classification C73, D62, D63, D64, H41, Q54

1 Introduction

Most dynamic analyses of environmental problems assume that the decision maker is an infinitely lived representative agent (ILRA). A small body of work uses overlapping generations (OLG) models to study environmental problems. OLG models recognize that at any point in time individuals of different ages are alive; over time some of these people die and new people are born. Decisions taken today affect not only the future utility flows of people currently alive, but also the utility flows of those who have not yet been born. OLG models, unlike ILRA models, recognize the difference between an individual at two points in time, and two individuals at these two points in time. This difference is particularly important if people making current decisions

Received March 13, 2013

Larry Karp (✉)
Agricultural and Resource Economics Department, University of California, Berkeley. Berkeley, CA 94720-3310, USA
E-mail: karp@berkeley.edu

have different preferences with respect to their own future utility versus the utility of their successors. An impatient person might, for example, prefer to receive 0.9 units of utility today rather than one unit of utility ten years from now. There is no reason to think that this person would also be willing to take 1 unit of utility from her successor born 110 years from now in order to give an extra 0.9 units of utility to someone born 100 years from now. Our impatience with respect to our own utility might bear little relation to our willingness to make trade-offs between our successors, who are different people. OLG models can recognize the difference between our views about our own future utility, and about our successors' utility. This distinction, which is absent in ILRA models, may be particularly important for long-lived environmental problems, e.g., those related to climate change, where decisions taken today may affect people's welfare centuries from now.

This paper is intended as an introduction to OLG models used to study long-lived dynamic environmental problems. A simple model, involving only two parameters, provides a way to distinguish between impatience with respect to our own future utility, from our willingness to make transfers across different people. When individuals make this distinction, there are different ways to think about the social planner (or the succession of social planners) who represents those individuals. One alternative yields time consistent preferences, and the second yields time inconsistent preferences. I explain the logic behind these two alternatives, together with the meaning of time inconsistency, and the approaches taken to study problems that exhibit it. I then discuss the relationship between a particular type of OLG model and a model of an infinitely lived representative agent with non-constant discounting. Some OLG models can be studied using tools developed to study models of non-constant discounting. A three period model shows how to conduct analysis under non-constant discounting and illustrates the type of strategic incentives that arise in this setting. In the interest of brevity, this paper does not provide a review of the literature. The interested reader should consult the papers cited, which contain such reviews.

2 Time Consistency

A simple model of preferences in an OLG setting uses two parameters to measure attitudes toward utility trade-offs over time and across individuals. The first parameter, known as the pure rate of time preference, and denoted $r \geqslant 0$, measures a person's willingness to sacrifice their own future utility in order to increase their own current utility. A larger value of r corresponds to greater impatience or a lower valuation of their own future utility. People care about the utility stream of their descendants to a greater or lesser extent depending on their degree of intergenerational altruism. In order to distinguish between these different components of a person's welfare, I label the welfare arising from the person's own utility flow as the "selfish" component of their welfare. For example, if a person born at time t lives for T years, has a pure rate of time preference r, and experiences a constant utility flow u during their lifetime, the selfish component of their welfare equals

$$W^{\text{selfish}}(t) = u\int_t^{T+t} e^{-rs}ds = \frac{1-e^{-rT}}{r}u.$$

Here, t merely records the person's date of birth.

A second discounting parameter, denoted $\lambda \geqslant 0$, provides an inverse measure of altruism with respect to unborn future generations. For a person alive at time 0, the altruistic component of their welfare, associated with the agent born at time t, is $e^{-\lambda t}W^{\text{selfish}}(t)$; the person alive at time 0 values the welfare of all agents born in the future ($t > 0$) in this manner. Perhaps individuals feel altruism with respect to future generations, and a social planner adopts those preferences, e.g. by means of a political process. Alternatively, perhaps a planner somehow exists outside the society, e.g., the planner is appointed by a deity. This distinction becomes important later, but for the time being I merely speak of a planner who has altruism. The planner discounts the selfish welfare of agents born in the future at rate λ.

A larger value of r means that individuals are more impatient as regards their own utility, and a larger value of λ means that the social planner is more willing to require sacrifices of a distant successor in order to benefit a nearer successor. The parameter λ is *inversely* related to altruism. For very large values of λ the planner is willing to take a great deal away from someone born in the future, in order to achieve a small benefit for someone born only slightly earlier. Large values of λ imply that the planner does not care about people born in the distant future, and in that sense is not altruistic.

This model helps to illustrate and explain the problem of time consistency. At time 0 a planner evaluates a policy that increases Leticia's time t utility but reduces Jianfeng's time t utility. Jianfeng is born at time 0 and Leticia is born at time $0 < s < t$ and they will both be alive at time t. Both people have the pure rate of time preference r, and the planner discounts the selfish welfare of agents born in the future at rate λ. The present value, to Jianfeng, at the time of his birth, of a unit of utility at time t is e^{-rt}. The planner at time 0 applies no discount rate to Jianfeng's future utility because he has just been born. The present value, to Leticia, discounted back to the time of her birth, s, of a unit of utility at time t, equals $e^{-r(t-s)}$. The planner discounts Leticia's time s welfare using the factor $e^{-\lambda s}$, so the planner attaches the weight $e^{-\lambda s}e^{-r(t-s)}$ to a unit of Leticia's time t utility. The ratio of these weights is

$$\Phi(0, s, t) \equiv \frac{e^{-\lambda s}\left[e^{-r(t-s)}\right]}{e^{-rt}}.$$

The first argument of the function Φ gives Jianfeng's age (0) when the planner is considering this transfer; the second argument gives Leticia's date of birth (s), and the third argument gives the time at which the policy will cause a reduction in Jianfeng's utility and an increase in Leticia's. If the policy increases Leticia's utility by one unit, then the planner at time 0 wants to adopt the policy if and only if it reduces Jianfeng's utility by no more than $\Phi(0, s, t)$. Thus, $\Phi(0, s, t)$ is a threshold. Utility losses for Jianfeng below this threshold make the policy acceptable to the planner at time 0, whereas utility losses for Jianfeng above the threshold make the policy unacceptable.

After $a < s$ years have elapsed, Jianfeng is a years old, and Leticia has still not been born. A planner at time a is taking a second look at the policy that transfers utility from Jianfeng to Leticia at time t. The present value to Leticia of a unit of utility at t, discounted back to her birth, is still $e^{-r(t-s)}$, and the planner now attaches the weight $e^{-\lambda(s-a)}$ to this quantity, because at time a Leticia will be born $s - a$ years in the future. Suppose that the planner attaches the total weight $d(a, t)$ to Jianfeng's utility, where "total" means that the weight includes both the planner's and Jianfeng's preferences. The ratio of the weights for Leticia and Jianfeng is now

$$\Phi(a, s, t) = \frac{e^{-\lambda(s-a)}\left[e^{-r(t-s)}\right]}{d(a, t)}.$$

Again, this ratio is a threshold. The planner at a will not accept losses to Jianfeng at t, in exchange for a one unit gain in utility to Leticia at t, in excess of this ratio.

Preferences are said to be time consistent if and only if, in the absence of new information, the planner at time a makes the same decision as the planner at the earlier time, 0. The two planners might be the same agent, appointed by the deity, or they might be different agents, representing the individuals alive at time 0 and at time a respectively For the time being, I ignore this distinction, speaking merely of the two planners. Comparison of the ratios shows that they are the same, and thus preferences are time consistent, if and only if

$$d(a, t) = e^{a(\lambda-r)}e^{-r(t-a)}.$$

The quantity $e^{-r(t-a)}$ is the present value that Jianfeng attributes to a unit of his own utility $t - a$ years in the future, so $e^{a(\lambda-r)}$ is the value that the planner at a attaches to Jianfeng's selfish welfare.

To determine whether this discounting formula is reasonable, it helps to consider matters more generally, by introducing the variable $\tau = t - a$, the number of years in the future (from the current time) at which the transfer between Jianfeng and Leticia will occur. As before, a denotes the age of the person currently alive (here Jianfeng) whose utility flow the planner is evaluating. With this definition, we see that the time consistency of plans requires the planner, at any time before Leticia is born, to attach the weight $d(a, \tau) = e^{a(\lambda-r)}e^{-r\tau}$ to the utility τ years in the future of a person currently a years old (Jianfeng). The quantity $e^{-r\tau}$ equals the weight that Jianfeng attaches to his own future utility flow (because he has the pure rate of time preference r), so $e^{a(\lambda-r)}$ is the weight that the planner attaches to Jianfeng's selfish welfare when Jianfeng is a years old.

In the special case where $\lambda = r$, the planner attaches the same weight to everyone, regardless of their age. The equality $\lambda = r$ makes the modeler's life simple, but is not especially compelling. The parameters r and λ measure fundamentally different concepts, the willingness to make intertemporal trades for the same person, and the willingness to make intertemporal trades for two different people. There is no reason to think that the two parameters "should" have the same value. For the rest of this section, I take $\lambda < r$; the results change in an obvious manner if the inequality is reversed.

For $\lambda \neq r$, the modeler has a choice. A modeler who requires that preferences be time consistent, sets $d(a, \tau) = e^{a(\lambda - r)} e^{-r\tau}$. In order to appreciate the significance of this formulation, consider Jianfeng and Bob, who are both currently alive; Bob is older, e.g. Bob is a' years old and Jianfeng is a years old with $a' > a$. The planner attaches more weight to Jianfeng than to Bob, because the assumptions $r > \lambda$ and $a' > a$ imply $e^{a(\lambda - r)} > e^{a'(\lambda - r)}$. Whether this weighting scheme is reasonable depends on how we think of the planner. A planner appointed by a deity in some sense "stands outside of time"; for such a planner, there is no special significance to the current time or to the ordering of time. This planner might reasonably discount every person's welfare back to their date of birth, not to the current time. A planner appointed by a deity should have the obvious attributes of a deity, including time consistency.

Matters are different if we think of this planner, or sequence of planners, as arising from a political process that somehow reflects the preferences of those currently alive when a decision is made. We can imagine many such political processes, but it is not likely that any would assign influence to individuals that monotonically decreases with a person's age.

A neutral assumption is that a person's influence is independent of their age. With that assumption, Jianfeng and Bob have the same weight in the social planner's objective function, and we replace $d(a, \tau) = e^{a(\lambda - r)} e^{-r\tau}$ with $\hat{d}(\tau) = e^{-r\tau}$. Although this weighting scheme gives the same weight to the future utility of two currently living agents, it gives different weight (for $\lambda \neq r$) to the future utility of a currently living agent, and one who is born in the future. For example, at time $t = a$, Bob and Jianfeng receive the weight $e^{-r\tau}$ for a unit of utility τ years in the future, but Leticia, who is born $s < \tau$ years in the future receives the weight $e^{-\lambda s} e^{-r(\tau - s)}$. She receives a higher weight than Bob and Jianfeng if $\lambda < r$ and a lower weight if $\lambda > r$. The rationale for this difference is that at the current time Bob and Jianfeng are both alive to influence the decision maker. Leticia has not yet been born, so her interests are represented only to the extent that currently living agents care about her.

The remarks above establish that with the weighting scheme that is independent of a person's age, the planner's preferences are time inconsistent. To illustrate this inconsistency, suppose that the policy in question increases Leticia's time t utility by 1 unit and decreases Jianfeng's time t utility by $\Phi(0, s, t) - \varepsilon$, with $\varepsilon > 0$. The planner at time 0 approves of this policy, because, from this planner's standpoint, the gain to Leticia is worth the cost to Jianfeng. In line with the remarks above, suppose that planners ignore a living person's age in evaluating their payoff (assuming of course that the planner is not worried about the person in question dieing before the transfer takes place). After $a < s$ years have elapsed, the net gain to the planner, of adopting the policy, equals the gain that the planner attributes to the benefit that Leticia enjoys, minus the cost that the planner attributes to Jianfeng's loss:

$$\begin{aligned} &e^{-\lambda(s-a)}\left[e^{-r(t-s)}\right](1) - e^{-r(t-a)}\left[\Phi\left(0,s,t\right)-\varepsilon\right] \\ &= e^{r(s-t)}\left(e^{\lambda(a-s)} - e^{-s\lambda}e^{ar}\right) + e^{-r(t-a)}\varepsilon. \end{aligned}$$

As $a \rightarrow s$ the expression on the second line approaches

$$e^{r(s-t)}\left(1 - e^{s(r-\lambda)} + \varepsilon\right),$$

which (in view of $r - \lambda > 0$) is negative for $\varepsilon < e^{s(r-\lambda)} - 1$. If the policy is extremely attractive for the initial planner ($\varepsilon > e^{s(r-\lambda)} - 1$) then it remains attractive for all subsequent planners, and there is no preference reversal. However, if the policy is only moderately attractive for the initial planner ($0 < \varepsilon < e^{s(r-\lambda)} - 1$), then it becomes unattractive for some subsequent planner, and there is a preference reversal.

Modelers who want to think of planners as emerging from a political process, rather than being appointed by a deity, may prefer the weighting scheme $\hat{d}\left(\tau\right) = e^{-r\tau}$ rather than $d\left(a,\tau\right) = e^{a(\lambda-r)}e^{-r\tau}$. These modelers must confront the problem of time inconsistency, or preference reversals. If the planner at a point in time can make commitments about actions taken in the future, the possibility that preferences can be reversed is irrelevant: what the current boss (planner) says, goes. This formulation is not reasonable where we are interested in sequences of decisions over many decades: policymakers today cannot perfectly commit their distant successors to a particular course of action, although they can influence their successors.

There have been two principal responses to the recognition of time inconsistency. One alternative assumes that the decision-maker is naive, and makes plans about the future under the mistaken belief that in the current period she can choose policies that will be implemented in the future. This naive planner does not bother to try to influence her successors, because she acts as if she can dictate their actions to them. Why assume this particular type of irrationality, or wishful thinking, on the part of agents? Without the constraint of (some degree of) rationality, a modeler has too much freedom to bake the results into the cake. I therefore do not consider the naive planner further. The second alternative assumes that today's planner is sophisticated. The planner can choose an action today, and understands that her successors choose actions during the period when they are in power.

Consider the case where the current environment can be described by a "state variable," e.g. the stock of atmospheric greenhouse gas (GHG). This state variable is said to be "directly payoff-relevant" because it has a physical, as distinct from psychological, effect on agents' payoffs. Policies that are conditioned on only the directly payoff-relevant state variable, are said to be Markovian. In the game amongst the sequence of planners, each indexed by the time at which they choose an action, a subgame perfect Markov equilibrium is called Markov Perfect. In such an equilibrium, each of the sequence of planners takes an action conditioned on the state variable that they inherit, and anticipates that her successors will do the same. Given beliefs about future policy functions (a mapping from the future value of the state variable to the future action), today's planner can influence her successors by influencing the state variable that the successors inherit. For example, a change in current GHG emissions alters the trajectory of the stock of GHGs, thereby altering the emissions that future policy makers select.

3 The Relationship between the Planner and Society

This section shows that certain OLG models are equivalent to models with an infinitely lived representative agent with non-constant discount rate. In these cases, methods developed to study models of non-constant discounting can be used to analyze the OLG model. Suppose that all agents have the same utility function, which is independent of their age, and that there is a mechanism for transferring utility among people currently alive, so that all agents alive at a point in time have the same utility flow; those flows may change over time. A policy that equalizes utility flows across agents at a point in time is unrealistic, but it makes the model tractable while also emphasizing the effect of transferring utility across points in time. With these strong assumptions, it is possible to aggregate the preferences, with respect to future utility flows, of the agents currently alive, using a discount factor with non-constant discount rate. Let $\{u_\tau\}_{\tau=t}^{\infty}$ be a trajectory of utility from the current time, t, to ∞. Because I assume that all agents alive at a point in time have the same utility flow, it is unnecessary to keep track of different individuals at a point in time.

The objective is to find a discount factor, a function $D(\tau)$ that aggregates the trajectory of utility, producing a welfare criterion $\int_0^\infty D(\tau)\, u_{t+\tau} d\tau$. In constructing the discount function $D(\tau)$, I retain the assumption that agents currently alive discount their own future utility at rate r and discount the welfare of their successors at rate λ. I also assume that the population is constant; relaxing that assumption adds only minor complications. Most significantly, I assume that all agents alive at a point in time have the same weight in the welfare aggregation. Their weight is not linked to their age. From the previous section, we know that this assumption implies that preferences are time inconsistent

Before considering this model in more detail, I contrast it with Schneider et al.'s (2012) OLG model in which agents accumulate private capital. A person's equilibrium capital stock depends on their age; therefore, two people of different age, alive at a point in time, have different capital stocks and different utility flows. This distinction among currently living agents greatly complicates the model. Schneider's et al. also impose time consistency of preferences. As noted above, this requirement implies that older agents receive less weight in the social planner's welfare criterion if $\lambda < r$. Time consistency is a reasonable property if we think of the social planner as standing outside time. It also provides an appropriate comparison with the standard model where the social planner is identified with an infinitely lived representative agent (ILRA). Schneider's et al. describe the problems that arise in using preference parameters of agents in the OLG setting to calibrate preference parameters of the ILRA. Most integrated assessment models used to evaluate climate policy are based on the ILRA. The OLG setting distinguishes between individual's discounting of their own future utility, and their (or a social planner's) discounting of the utility of unborn future generations. This distinction is usually neglected in ILRA models, which therefore cannot fully represent preferences of agents in an OLG setting, even in the case where both are time consistent.

I now return to the model where two agents currently living have the same utility flow, and the same weight in the social planner's welfare criterion. In this case, the planner's (or the sequence of planners') preferences are time inconsistent. This model is appropriate if we think of decisions being made through a political process in which an agent's age does not determine their influence.

Completing the model requires an assumption on agents' lifetime. There are two obvious candidates. Under "exponential lifetimes," an agent's lifetime is exponentially distributed, with mortality rate (the hazard rate) θ. Under "finite lifetimes" each agent has a known, finite lifetime, T. The two models are comparable if the expected exponential lifetime equals the finite lifetime, which requires $\theta = \frac{1}{T}$. The exponential distribution is "memoryless," which means that the expected additional lifetime of an agent currently alive is independent of their current age. In this setting, any two agents currently alive are indistinguishable from each other: they both have the same distribution of remaining lifetime, and by assumption their utility flows are equal. In this case, there is a representative agent in the usual sense. Any individual currently alive can be taken as the social planner at that point in time.

Under finite lifetimes, an older agent dies sooner than a currently living younger agent. The older agent therefore has less time in which to obtain selfish welfare from his own future utility flows, compared to the younger agent. By assumption, both agents evaluate welfare of unborn future generations in the same way. In this setting, the two agents currently alive are different, even though (by assumption) they obtain the same utility flows while alive. With finite lifetimes, I assume that a utilitarian social planner aggregates the preferences of those currently alive, assigning equal weight to their preferences.

Ekeland and Lazrak (2010) obtain the formula for the function D under exponential lifetime, and Karp (2013) obtains the formula under finite lifetime. Under exponential lifetime, D is a weighted sum of exponentials, and under finite lifetime, D has a somewhat more complicated form. Rather than reproduce these formula, I use Fig. 1, taken from Karp (2013), which shows the discount rates associated with the discount factors, under exponentially distributed and finite lifetimes, for parameter values $r = 0.02 = \theta = \frac{1}{T}$, and for $\lambda \in \{0.01, 0.06\}$. For $\lambda = 0.01$ (where $\lambda < r$), agents value the future utility flows of unborn generations more highly than their own future utility flows, and the reverse holds for $\lambda = 0.06$. The discount rates are constant for $\lambda = r$ and they fall for $\lambda < r$, a case often called hyperbolic discounting; the discount rates rise if $\lambda > r$. For $\lambda < r$, the discount rates under exponentially lived and finitely lived agents are similar, but the two discount rates are quite different when $\lambda >> r$, at least for distant times (large t).

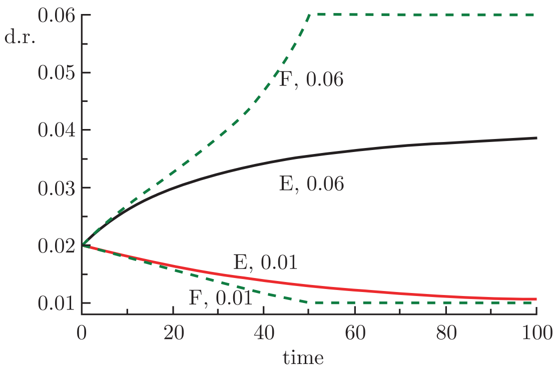

Fig. 1 Discount Rates (d.r.) for $\theta = 0.02 = r = \frac{1}{T}$

Note: Solid curves (labelled E) correspond to exponentially distributed lifetime, and dashed curves (labelled F) correspond to fixed lifetime. Numerical values in label show value of λ.

A discrete time example helps to explain the logic behind these discount factors and rates. Suppose that agents live for two periods: $T = 2$. The population is constant, with measure normalized to 1, so the measure $\frac{1}{2}$ of new agents are born in each period. The agents alive in period 0 discount future generations' utility at $e^{-\lambda}$ per period, and they discount their own utility at e^{-r} per period. Consider matters from the standpoint of agents at time 0. In the next period, half of those agents are still alive, so the population-weighted present value of a unit of their utility is $\frac{1}{2}e^{-r}$; half of the population has just been born, so the population-weighted present value of their utility (as valued by agents at time 0) is $\frac{1}{2}e^{-\lambda}$. The total weight applied to a unit of utility in the next period is therefore $D(1) = \frac{1}{2}\left(e^{-r} + e^{-\lambda}\right)$. At time $t > 1$ all of the agents alive at time 0 have died. Half of the population at t was born at $t-1$. Their present value, at time of birth, of a unit of time t utility is e^{-r}, and agents at time 0 value that utility at $e^{\lambda(t-1)}e^{-r}$. The time 0 agents value a unit of utility for those born at time t at $e^{-\lambda t}$. Thus, the discount factor that those alive at time 0 apply to a unit of utility at time t is $D(t) = \frac{1}{2}\left(e^{\lambda(t-1)}e^{-r} + e^{-\lambda t}\right)$. Defining $\beta = \frac{e^{-r} + e^{-\lambda}}{2}e^{\lambda}$, $\delta = e^{-\lambda}$, this formula simplifies to $D(1) = \beta\delta$ and $D(t) = \beta\delta^t$ for $t > 1$. This particular form of discounting is known as β, δ or quasi-hyperbolic discounting (Laibson 97). The case usually emphasized is $\lambda < r$, where $\beta < 1$.

Karp (2005), and Gerlagh and Liski (2012) study climate models with β, δ discounting. Both of these papers describe the model as the representation of an infinitely lived agent who happens to have non-constant discounting. There are various motivations for non-constant discounting in the climate context, including the idea that a person is impatient with respect to their own future utility, and is better able to distinguish between utility flows at two nearby times, say time t and $t+1$, compared to the utility flows between two distant times, say t' and $t'+1$, where $t' \gg t$ (Karp and Tsur, 2011). This section shows that these β, δ models can also be interpreted as an OLG model in which agents live for two periods and discount the utility of future generations at a rate different from the discount rate applied to their own future utility. More generally, a class of OLG models can be studied by examining the problem of an infinitely lived agent with non-constant discounting.

4 A Three-Period Example

Here I use a three period model, adapted from Gerlagh and Liski (2012), to illustrate the strategic incentives that arise in an OLG setting—equivalently, under non-constant discounting.[1] I find the conditions that determine whether, in equilibrium, the investment in environmental capital relative to man-made capital is greater or less than the efficient ratio of investment in the two stocks.

The agent in the first period has two ways to transfer utility to the future. By leaving an additional unit of man-made capital, the first agent directly benefits his immediate successor, the second period agent; this bequest also indirectly benefits the third period agent, to the extent that the second period agent invests some of that additional capital

[1] The previous section shows that a particular OLG model in which an infinite sequence of agents each live for two periods is equivalent to a model of an infinitely lived agent with β, δ discounting. This equivalence is not exact in the three stage model considered here. But the motivation for this three period model is that it provides a simple way to think about an infinite horizon setting.

for the third period. The agent in the first period can also reduce its GHG emissions, thus leaving a larger stock of environmental capital, i.e., a better climate system. The assumption here, supported by climate science, is that emissions have a long term effect, represented by the third period, but not an effect in the middle term, represented by the second period. An additional assumption is that there are no GHG emissions in the second and third periods; e.g., society has discovered an alternative to fossil fuels by that time.

The first period agent discounts second period utility by $\beta\delta$ and discounts third period utility by $\beta\delta^2$, with $\beta < 1$ and $\delta < 1$. The second period agent discounts the third period utility by $\beta\delta$. Thus, the first agent is willing to take δ units of utility from the second period in order to give the third period one additional unit of utility; but the second period is willing to surrender only $\beta\delta < \delta$ units of utility in order to give the third period one additional unit of utility. The second period savings decision is therefore suboptimal from the standpoint of the agent in the first period.

The first period agent has one unit of a composite commodity, "capital," which can be either consumed or invested. The agent decides how much to consume, c_1, yielding utility $u(c_1)$, and how much to lock up as natural capital, an "environmental gift," E, for the third period. He invests the residual, $1 - c_1 - E$, in a project that earns exogenous gross return R. The second period agent has wealth

$$w = (1 - E - c_1) R. \tag{1}$$

Emissions in the first period are an input into production; equivalently, it is costly to reduce emissions. At a given level of first period consumption, a reduction in emissions creates a better third period environment (higher E), but lower second period wealth, as in equation (1). The second period agent decides how much to consume, c_2, yielding utility $u(c_2)$, and invests the residual in a project that earns exogenous gross return R for the third agent. The agent in the third period has the endowment E and $c_3 = [(1 - E - c_1) R - c_2] R$, and obtains utility $V(c_3, E)$. The functions u and V are increasing and concave.[2]

The first-period agent has two ways to transfer utility to the third agent: by means of the environmental stock, or by means of a higher endowment of capital (wealth) left to the second-period agent, who will pass some of that on to the third period agent. The second period agent cannot alter the level of the environmental stock, E, left by the first agent. In that sense, the environmental stock is "protected" against the potential depredations of the second period agent. If the first period agent leaves her successor

[2] The rate of return in Gerlagh and Liski (2012) is endogenous, while I take it as exogenous (R). However, I use more general utility functions, so neither model is a special case of the other. Their setting gives rise to the comparative statics reported in equation (3). However, in their model, there is always over-investment in the environmental stock, whereas in my setting there may be either over- or under-investment in the environmental stock.

more capital, the second-period agent consumes some of it rather than investing all of it for the third period agent. For this reason, there is an advantage, to the first period agent, of transferring utility to the third period agent by means of the cleaner environment, rather than by means of capital. That advantage tends to increase the equilibrium value of $\frac{E}{c_3}$.

In a first best setting, efficiency requires that the third-period marginal rate of substitution between capital and the environment, denoted MRS^3_{E,c_3}, equal the compounded return on investment over two periods, i.e.,

$$MRS^3_{E,c_3} \equiv \frac{\frac{\partial V}{\partial E}}{\frac{\partial V}{\partial c_3}} = R^2. \tag{2}$$

To interpret this equation, use

$$dV = \frac{\partial V}{\partial E}dE + \frac{\partial V}{\partial c_3}dc_3 = 0 \Longrightarrow MRS^3_{E,c_3} = -\frac{dc_3}{dE}_{|\text{utility constant}}.$$

The function MRS^3_{E,c_3} equals the number of units of capital that the agent in the third period would require in order to be willing to sacrifice one unit of the environmental stock. If the first period agent were able to reduce the environmental stock by one unit, and lock up that unit in an investment for two periods, the additional third period capital would be R^2. Efficiency requires that the third period agent's willingness to trade between capital and the environmental stock equals the ability to transform the environmental stock to capital by means of investment. I use the following

Definition 1. *The equilibrium involves over-investment in environmental stocks (in equilibrium, $\frac{E}{c_3}$ is larger than the efficient ratio) if $MRS^3_{E,c_3} < R^2$ and the equilibrium involves under-investment in environmental stocks (in equilibrium $\frac{E}{c_3}$ is smaller than the efficient ratio) if $MRS^3_{E,c_3} > R^2$.*

Recall that the second period under-invests, from the perspective of the first period. In contrast, Definition 1 determines over- and under-investment from the perspective of efficiency, not the perspective of the first agent, and in addition it concerns the ratio of stocks, $\frac{E}{c_3}$, not their levels. The environmental stock is "locked up," whereas capital must pass through the hands of the second period agent. This difference creates a strategic incentive for the first period agent to prefer transferring utility to the third period by means of the environmental stock rather than capital: this strategic incentive tends to result in a larger ratio $\frac{E}{c_3}$ than is consistent with equation (2), creating pressure for over-investment in the environmental stock (where $MRS^3_{E,c_3} < R^2$).

The second period agent takes E and w as given. Its optimal bequest is

$$c_3 = c_3\left(w, E\right) = \arg\max_{c_3} \left[u\left(w - \frac{c_3}{R}\right) + \beta\delta V(c_3, E)\right].$$

The comparative statics of the policy function are

$$R > \frac{\partial c_3}{\partial w} > 0 \text{ and } \operatorname{sign}\left(\frac{\partial c_3}{\partial E}\right) = \operatorname{sign}\left(\frac{\partial^2 V}{\partial c_3 \partial E}\right). \tag{3}$$

(An appendix contains additional calculations and proofs.) System (3) shows how the first-period agent is able to manipulate the second-period decisions. The first pair of inequalities confirms a previous assertion: if the first agent leaves his successor an additional unit of wealth, the successor consumes some but not all of it, investing the remainder for the third period agent.The second equation shows that the effect of a cleaner third period environment (higher E), on second period investment, depends on the cross partial, $\frac{\partial^2 V}{\partial c_3 \partial E}$. I adopt

Definition 2. *Capital and the environmental stocks are third period complements in utility if an increase in one increases the marginal value of the other* $\left(\frac{\partial^2 V}{\partial c_3 \partial E} > 0\right)$ *and they are third period substitutes in utility if an increase in one decreases the marginal value of the other* $\left(\frac{\partial^2 V}{\partial c_3 \partial E} < 0\right)$.

To obtain an economic interpretation of this definition, consider the case where third period utility depends only on third period consumption, C, which depends on the capital and environmental stocks: $C = h\left(c_3, E\right)$, where, h is a concave production function, increasing in both arguments. With these assumptions, $V\left(c_3, E\right) = U\left(h\left(c_3, E\right)\right)$, where the utility function U is concave. In this case,

$$\frac{\partial^2 V}{\partial c_3 \partial E} = U'\left(C\right)\left[h_{c_3,E} + \frac{U''}{U'} h_{c_3} h_E\right],$$

where subscripts denote partial derivatives. A sufficient condition for $\frac{\partial^2 V}{\partial c_3 \partial E} < 0$ (so that $\frac{\partial c_3}{\partial E} < 0$) is $h_{c_3,E} < 0$, and a necessary condition for $\frac{\partial^2 V}{\partial c_3 \partial E} > 0$ (so that $\frac{\partial c_3}{\partial E} > 0$) is $h_{c_3,E} > 0$. DICE and other integrated assessment models assume that a cleaner environment increases the marginal productivity of capital, so that $h_{c_3,E} > 0$; in those models, capital and the environment are complements in production, but need not be complements in third period utility. (Note the distinction between "complements in utility" and "complements in production".) With $h_{c_3,E} > 0$, a larger value of E increases the productivity of the second period bequest, c_3, and tends to increase the equilibrium c_3. However, if the utility function U is "very concave," the higher third period income associated with higher E reduces the marginal utility of the extra income generated by a higher bequest, a fact that causes larger E to reduce the equilibrium c_3. The net effect on c_3 of larger E is ambiguous when $h_{c_3,E} > 0$. If U is nearly linear, and $h_{c_3,E} > 0$, a larger E increases the equilibrium c_3. In this case, the first period agent has a strategic incentive to increase E in order to increase second period investment. This incentive promotes over-investment in the environmental stock.

The two inputs, E, c_3, might be substitutes in production, in which case they are also substitutes in utility. For example, suppose that a worse environment (lower E) requires increased expenditures on non-discretionary adaptive expenditures such as seawalls or on the prevention of climate-related epidemics. These expenditures cannot be used for production of the consumption good. Denote non-discretionary adaptive expenditures as $A(E)$, a decreasing convex function of environmental quality. The amount of capital remaining for production, after expenditures on adaptation, is $c_3 - A(E)$, and the production function for the consumption good is $h(c_3 - A(E))$, a concave increasing function. With this model, $h_{c_3,E} = -h''(c_3 - A(E))A' < 0$. Using these facts and equation (3), I have:

Lemma 1. *If c_3 and E are inputs into third period production, and third period utility is a function of third period output, then $V(c_3, E) = U(h(c_3, E))$. In this case, a necessary condition for a cleaner environment to induce higher second period investment $\left(\frac{\partial c_3}{\partial E} > 0\right)$ is that the inputs c_3, E are complements in production; if, in addition, the third period utility function is nearly linear in consumption, then $\frac{\partial c_3}{\partial E} > 0$. A sufficient condition for a cleaner environment to induce lower second period investment $\left(\frac{\partial c_3}{\partial E} < 0\right)$ is that the inputs c_3, E are substitutes in production.*

I noted above that the first period agent has an incentive over-invest in the environmental good, relative to capital, because doing so provides a way to transfer utility to the third period without the risk that the second period agent will consume some of the transfer. Lemma 1 shows that an offsetting strategic incentive arises if the third period inputs are substitutes in production, or if the third period utility function is very concave. In either of these cases, higher environmental quality implies lower second period investment, a fact that tends to promote under-investment in the environmental stock.

By studying the optimization problem of the first period agent, I obtain

Proposition 1. *The equilibrium involves under-investment in the environmental stock ($MRS^3_{E,c_3} > R^2$) if and only if in equilibrium*

$$\frac{\partial^2 V}{\partial c_3 \partial E} < R^2 \frac{\partial^2 V}{\partial c_3^2}. \tag{4}$$

and the equilibrium involves overinvestment in environmental stocks if and only if inequality (4) is reversed.

The right side of inequality (4) is negative by concavity of V, so a necessary (but not sufficient) condition for the inequality is that $\frac{\partial^2 V}{\partial c_3 \partial E} < 0$, which by equation (3) implies $\frac{\partial c_3}{\partial E} < 0$. Lemma 1 provides an example showing that $\frac{\partial^2 V}{\partial c_3 \partial E}$ might be either positive or negative.

The proposition states that whether under- or over-investment in the environmental stock actually occurs in equilibrium depends on a ratio of second derivatives of the third period utility function. In contrast, the *definition* of under- and over-investment depends on the ratio of first derivatives. To demonstrate that inequality (4) is consistent with concavity of V, consider the linear-quadratic example

$$V = a_1 c_3 + a_2 E - \frac{1}{2}\left(c_3^2 + fE^2 + 2gc_3E\right).$$

Concavity of V requires $f > 0$ and $f > g^2$, or $g < +\sqrt{f}$, and inequality (4) requires $g > R^2$. Both inequalities are satisfied if $+\sqrt{f} > g > R^2$.

This model illustrates the strategic incentives arising under non-constant discounting (or an OLG model). The second period agent under-invests, from the standpoint of the first period agent. The agent in the first period has two ways to transfer utility to the third period agent: by increasing investment in either man-made or natural capital. The model is consistent with the science of climate change, which suggests that the full effect of current emissions is felt only many decades in the future. This feature, together with the assumptions that GHG emissions occur only in the first period, implies that a first period investment in natural capital cannot be altered by the second period agent. That agent would, however, consume some of the capital that the first period agent would like to bequeath the third period. The fact that investment in natural capital is protected from the depredations of the second period agent, whereas investments in man-made capital are vulnerable to those depredations, promotes over-investment in natural capital in equilibrium. However, a lower stock of natural capital increases second period equilibrium investment in man-made capital, if man-made and natural capital are substitutes in third period utility. When these stocks are substitutes, the fact that second period investment is too low, from the perspective of the first agent, gives the first period agent an incentive to under-invest in natural capital. Thus, when the two stocks are substitutes in third period utility, the equilibrium might result in over-investment of natural capital or in under-investment of natural capital.

5 Discussion

Overlapping generations models, unlike the infinitely lived representative agent model with constant pure rate of time preference, distinguish between attitudes toward our own future utility, and our successors' utility. OLG models offer new insights to questions regarding the investment in long-lived public goods, such as the climate system. In an effort to reduce the entry cost to this literature, I have reviewed and illustrated some significant features of these models.

Perhaps the most basic question is whether to require a planner's preferences to be time-consistent. In a two-parameter model, time consistency requires giving less weight to the future utility of people who are currently living as they grow older. Time consistency is a reasonable requirement if we think of plans being imposed by a god-like agent who stands outside of time; the resulting model also provides a useful comparison to the ILRA model, which (typically) is time consistent. However, if we think of plans (e.g., for climate protection) as being determined by a political system, there is no reason for old agents to be less influential. In that case, the succession of planners who represent the interests of those alive at the time plans are chosen and implemented, are likely to have time-inconsistent preferences. With this model, sophisticated agents (plausibly) condition their actions on directly payoff-relevant state variables.

For a class of OLG models in which all agents alive at a point in time have the same utility flow, the preferences of agents currently alive can be aggregated using a discount factor with non-constant discount rate. This equivalence makes it possible to study some OLG models using tools developed to study models of non-constant discounting. It also permits models of non-constant discounting, recently employed to study climate policy, to be interpreted as OLG models.

I illustrated the methods used to study finite horizon versions of these models using a three period model. The simplicity of this model makes it possible to understand the relation between primitives, such as the utility functions, and strategic incentives that agents face when they cannot commit their successors to specific actions. These strategic incentives might lead to either over- or under-investment of natural capital, relative to man-made capital.

Acknowledgements This paper was prepared for the workshop "Global Environmental Challenges: The Role of China," at the Shanghai University of Finance and Economics, December 13–14, 2012. I thank the workshop participants for their comments.

References

Ekeland I, Lazrak A (2010). The golden rule when preferences are time inconsistent. Mathematical and Financial Economics, 4(1): 29–55

Gerlagh R, Liski M (2012). Carbon prices for the next thousand years. Unpublished working paper

Karp L (2005). Global warming and hyperbolic discounting. Journal of Public Economics, 89: 261–282

Karp L (2013). Provision of a public good with altruistic overlapping generations and many tribes. Giannini Foundation working paper

Karp L, and Tsur Y (2011). Time perspective and climate change policy. Journal of Environmental Economics and Management, 62: 1–14

Laibson D (1997). Golden eggs and hyperbolic discounting. Quarterly Journal of Economics, 62: 443–78

Schneider M, Traeger C, Winkler R (2012). Trading off generations: Equity, efficiency and climate change. European Economic Review, forthcoming

Appendix: Technical Material

Derivation of equation (3). To ease notation, I define

$$\nu \equiv \frac{\partial^2 V}{\partial c_3^2} < 0 \text{ and } \theta \equiv \frac{\partial^2 V}{\partial c_3 \partial E}.$$

Concavity of V implies $\nu < 0$, but θ could be positive or negative. Using $c_3 = (w - c_2) R$, the second period bequest solves

$$\max_{c_3} \left[u\left(w - \frac{c_3}{R} \right) + \beta\delta V(c_3, E) \right] \Longrightarrow$$

$$f(w, E, c_3) \equiv -u'(c_2) + \beta\delta R \frac{\partial V(c_3, E)}{\partial c_3} = 0 \tag{5}$$

$$\frac{\partial f}{\partial c_3} = \frac{1}{R}\left(u''(c_2) + \beta\delta R^2 \nu \right) < 0.$$

The second line of this system, the first order condition, implicitly defines the second-period policy function, $c_3 = c_3(w, E)$; the third line is the second order condition. Totally differentiating the first order condition gives

$$\frac{\partial c_3}{\partial w} = \frac{1}{\frac{\partial f}{\partial c_3}} u''(c_2) > 0 \text{ and } \frac{\partial c_3}{\partial E} = \left(\frac{-\beta\delta R}{\frac{\partial f}{\partial c_3}} \right) \theta \lesseqgtr 0.$$

Using $\nu < 0$, the first equality implies

$$\frac{\partial c_3}{\partial w} = \frac{R u''(c_2)}{u''(c_2) + \beta\delta R^2 \frac{\partial V^2(c_3, E)}{\partial c_3^2}} < R.$$

Proof. Proposition (1)

Using the second-period policy function and equation (1), the first-period optimization problem is

$$\max_{w,E}\left\{u\left(1-E-\frac{w}{R}\right)+\beta\delta\left[u\left(w-\frac{c_3}{R}\right)+\delta V\left(c_3\left(w,E\right),E\right)\right]\right\}.$$

The first order condition for w is

$$-u'\left(c_1\right)\frac{1}{R}+\beta\delta\left[A\frac{dc_3}{dw}+u'\left(c_2\right)\right]=0$$

$$\text{with } A\equiv\frac{1}{R}\left(-u'\left(c_2\right)+\delta R\frac{\partial V}{\partial c_3}\right)=\delta\frac{\partial V}{\partial c_3}\left(1-\beta\right),$$

where the last equality in the definition of A uses the second line of equation (5). Using the definition of A, I rewrite the first order condition with respect to w as

$$-u'\left(c_1\right)+\beta\delta R\left[\delta\frac{\partial V}{\partial c_3}\left(1-\beta\right)\frac{dc_3}{dw}+u'\left(c_2\right)\right]=0.$$

Using similar manipulations, the first order condition with respect to E is

$$-u'\left(c_1\right)+\beta\delta\left[\delta\frac{\partial V}{\partial c_3}\left(1-\beta\right)\frac{dc_3}{dE}+\delta\frac{\partial V}{\partial E}\right]=0.$$

Setting these two first order conditions equal to each other and simplifying gives

$$\delta\left[\frac{\partial V}{\partial c_3}\left(1-\beta\right)\left(\frac{dc_3}{dE}-R\frac{dc_3}{dw}\right)+\frac{\partial V}{\partial E}\right]=Ru'\left(c_2\right).$$

Now using the second line of equation (5) to eliminate $u'\left(c_2\right)$ gives

$$\delta\left[\frac{\partial V}{\partial c_3}\left(1-\beta\right)\left(\frac{dc_3}{dE}-R\frac{dc_3}{dw}\right)+\frac{\partial V}{\partial E}\right]=\beta\delta R^2\frac{\partial V}{\partial c_3}\Longrightarrow$$

$$\frac{\partial V}{\partial E}=\frac{\partial V}{\partial c_3}R^2\left[\beta+\frac{\left(1-\beta\right)}{R^2}\left(R\frac{dc_3}{dw}-\frac{dc_3}{dE}\right)\right]\Longrightarrow$$

$$MRS^3_{E,c_3}\equiv\frac{\frac{\partial V}{\partial E}}{\frac{\partial V}{\partial c_3}}=R^2\left(\beta+\left(1-\beta\right)\Omega\right), \tag{6}$$

with

$$\Omega\equiv\frac{1}{R^2}\left(R\frac{dc_3}{dw}-\frac{dc_3}{dE}\right)=\frac{1}{R\frac{\partial f}{\partial c_3}}\left(u''\left(c_2\right)+\beta\delta\theta\right)=\frac{u''\left(c_2\right)+\beta\delta\theta}{u''\left(c_2\right)+\beta\delta R^2\nu},$$

where the last equality uses the third line of equation (5) to eliminate $\frac{\partial f}{\partial c_3}$. Equation (6) implies

$$MRS^3_{E,c_3} \left\{ \begin{array}{c} > R^2 \\ < R^2 \end{array} \right\} \text{if} \left\{ \begin{array}{c} \Omega > 1 \\ \Omega < 1 \end{array} \right\}.$$

Thus, there is overinvestment in the environment if and only if $\Omega < 1$ and underinvestment if and only if $\Omega > 1$. Using the definition of Ω, the last relation is equivalent to

$$MRS^3_{E,c_3} \left\{ \begin{array}{c} > R^2 \\ < R^2 \end{array} \right\} \text{if} \left\{ \begin{array}{c} \theta < R^2\nu \\ \theta > R^2\nu \end{array} \right\}. \quad \blacksquare$$