Households, Consumption, and Energy Use:
Population Age Structure and Future Carbon Emissions for the United States

March 4, 2005
Prepared for Population Association of America Conference
Philadelphia PA, April 2005

Michael Dalton,[†•] Brian O'Neill,[‡] Alexia Prskawetz,[*] Jiang Leiwen,[ç] John Pitkin[è]

[†] California State University Monterey Bay, CA, USA.
[‡] International Institute for Applied Systems Analysis, Laxenburg, Austria.
[*] Vienna Institute for Demography, Vienna, Austria.
[ç] Brown University, Providence, RI, USA.
[è] Analysis and Forecasting, Inc., Cambridge, MA, USA.

Abstract

Expected changes in the age composition of U.S. households over the next 25-100 years could affect demand for energy and emissions of carbon dioxide, the most important greenhouse gas. We use the Population-Environment-Technology model, household projections from the ProFamy model, and data from the U.S. Consumer Expenditure Survey, to assess the effects of population aging on labor supply, demand for different consumer goods, energy used to produce these goods, and carbon emissions. Our results show that demographic factors, combined with economic heterogeneity across households, can have substantial effects on energy use and carbon emissions.

1. Introduction

The links between population, consumption, and energy use are key components of the relationship between population and environment (Curran and de Sherbinin, 2004). Population can affect energy use in two major ways. First, demographic factors can affect the scale of the economy, either directly through population size, or indirectly through the effects of aging, urbanization, and other changes in economic growth (Birdsall et al., 2001). Second, compositional change in the population across various demographic categories could affect aggregate consumption patterns, and therefore total energy use. For example, over the past century living arrangements in the United States (U.S.) have shifted toward smaller households with older household members, and this shift is expected to continue in the future, driven by population aging and changes in behavior. The energy studies literature has identified these household characteristics as key determinants of direct residential energy demand (Schipper, 1996), and changes in the composition of U.S. households could have substantial effects on national energy demand (O'Neill and Chen, 2002). A small set of studies have included household level detail in projections of future energy demand, but these

[•] Corresponding author: Michael_Dalton@csumb.edu. We thank Warren Sanderson, Ross Guest, and other participants at the Symposium on Population Ageing and Economic Productivity, Vienna Institute for Demography, December 2, 2004 for helpful comments and discussions. Although the research described in this article has been funded wholly or in part by the United State Environmental Protection Agency through grant/cooperative agreement # R-82980101 to Brown University, it has not been subjected to the Agency's required peer and policy review and therefore does not necessarily reflect the views of the Agency and no official endorsement should be inferred.

have been limited to short time horizons, and simple household projections (Lareau and Darmstadter, 1983; Weber and Perrels, 2000).

Energy-economic growth models are popular tools for the analysis of climate change policies (Weyant and Hill, 1999). These models have a dynamic structure that is often based on the infinitely lived agent (ILA) of the Ramsey (1928) savings model. The energy-economic models are part of much larger family of dynamic computable general equilibrium (CGE) models that are used for analysis on a range of topics including, increasingly, policy issues related to population aging and other demographic events.

While these models have many similarities, they also exhibit important differences. Many models adopt a recursive, or backwards-looking, formulation of investment decisions, and are based on a variation of the Solow (1956) growth model that assumes some type of fixed savings rule, usually a constant fraction of income in each period. Fixed savings rules are a simplification that avoids solving a dynamic optimization problem. Nonetheless, models with fixed savings rules often compensate for this simplification with detailed energy sectors, and other realistic features such as land-use, and demographic change (e.g. MacCracken, et al., 1999).

Other models in the energy economics literature adopt a forward-looking approach to capital accumulation that assumes perfect-foresight about the future productivity of capital, prices, and other variables (e.g. Goulder, 1995). The properties of a dynamic competitive equilibrium with forward-looking behavior are substantially different from models based on fixed savings rules. In fact, a dynamic equilibrium with fixed savings rules is not an authentic competitive equilibrium because households are not, strictly speaking, utility maximizers. While the assumption of perfect-foresight is not realistic, it does incorporate information about the future into current decisions, and is thus an improvement over fixed-savings rules from the point of view of economic theory. Moreover, perfect-foresight can be interpreted as a first-order approximation to rational expectations (Fair and Taylor, 1983). Some economic growth models mix different types of savings behavior by assuming a proportion of the population solves a dynamic optimization problem, while others follow a fixed savings rule (McKibbin and Vines, 2000).

Properties of the equilibrium concept used in different economic growth models are further obfuscated by the choice of solution method. Most models rely on theorems from welfare economics that show, under certain conditions, a competitive equilibrium is the solution to a social planner's problem, and conversely, that solutions to a planner's problem can be supported with market prices to form a competitive equilibrium. While this approach is popular for computing and analyzing a competitive equilibrium, it is an indirect approach that is generally applicable only under special conditions such as perfect capital markets, and no externalities. While this technique is often applied in decentralized settings, and thus a tool for positive or descriptive economic analysis, solutions to a social planner's problem have been used extensively in the economics literature on climate change to evaluate the costs and benefits of emissions reduction. Results from this literature have been influential in policy discussions about the economically efficient level of emissions abatement. Therefore, the structure of the ILA models is open to scrutiny.

According to Manne (1999), ILA models are the typical approach for comparing costs and benefits of alternative emissions abatement strategies. He cites ILA models of Cline (1992), Peck and Teisberg (1992), Nordhaus (1991, 1994), and Manne, Mendelsohn, and Richels (1995). Nordhaus and Yang (1996) follow a similar approach, with multiple global regions. The economy in each of these models is analyzed as though there were a benevolent planner acting as a trustee on behalf of both present and future generations. This assumption is obviously heroic. Nonetheless, ILA models have been developed with detailed production sectors for energy and other intermediate goods, have a transparent dynamic structure to describe capital accumulation, and can be calibrated to historical data. In other words, ILA models are broadly consistent with economic theory, and currently provide the most detailed empirical tools for evaluating the costs, and perhaps benefits, of controlling greenhouse gas emissions. However, Schelling (1995) and others (e.g. Azar and Sterner, 1996) have criticized the strong welfare assumptions implicit in the representative agent, planner-based ILA approach.

Overlapping generations (OLG) models provide an alternative to ILA models for dealing explicitly with sustainability and other intergenerational welfare issues (Howarth and Norgaard, 1992; Farmer and Randall, 1997). The OLG models have an explicit demographic structure to describe key life-cycle stages. Recently, several OLG models have been used to re-examine the policy implications derived from the planner-based ILA models cited above. In some cases, OLG models yield results that are similar, or complement, corresponding ILA models (Stephen, et al., 1997; Manne, 1999).

However, other studies find substantial differences between results with OLG and ILA models. Howarth (1996, 1998) matches a two-period OLG model to assumptions in Nordhaus (1994), and finds that modest to aggressive reductions in greenhouse gas emissions are justifiable in terms of economic efficiency. Howarth shows that, in general, ILA models can be represented as reduced-form OLG models without qualitatively important demographic features. He concludes that Nordhaus' (1994) model, in particular, is strongly sensitive to changes in the intergenerational weights used in the social welfare function.

Gerlagh and van der Zwaan (2000, 2001) reach similar conclusions, or even stronger, and they question whether ILA models are appropriate for analysis of climate change policies. Differences in their results from other OLG models, notably Stephen et al. (1997) and Manne (1999), are attributed to an explicit representation of longer life expectancy and population aging in their three-period OLG model. However these features, while innovative, may undermine the force of the argument in favor of OLG models over an ILA approach, while supporting the importance of population and demographic factors in the economics of climate change.

Like their ILA counterparts, OLG models come with a variety of structural assumptions and solution techniques. In general, OLG models have dynamic properties that are different from ILA models (Auerbach and Kotlikoff, 1987; Geanakoplos and Polemarchakis, 1991; Kehoe, 1991). However these differences depend critically on the assumption that savers in OLG models plan only for their own retirement, and do not care about future generations. For example if parents care about the welfare of their children, a bequest motive exists that influences savings behavior, and leads to an OLG model that is similar to ILA models in terms of discounting (Barro, 1974). The Blanchard-Yaari-Weil model of perpetual youth provides a set of conditions under which OLG and ILA approaches are equivalent (Blanchard, 1985, Blanchard and Fischer, 1987). Marini and Scaramozzino (1995) use a version of this model to show that solving a social planner's problem with overlapping generations collapses to the representative agent framework as a special case only when there is an absence of heterogeneity among generations. In other words, questioning the suitability of the planner-based ILA approach to environmental policy analysis reduces to an empirical issue of whether there is significant heterogeneity in the savings and consumption decisions of different generations. Analyzing this issue is the subject of this paper.

In this paper, we develop an energy-economic growth model that shares features of ILA and OLG approaches. While the methods used in the literature described above are relevant to our approach, we do not evaluate costs or benefits of emissions abatement in this paper. Instead, the objective of this paper is to develop a modeling framework capable of explicitly accounting for the effects of demographic heterogeneity and compositional change over time on future economic growth, energy use, and CO2 emissions. We focus on a specific case study: the potential effects of changes in age structure on energy use and emissions in the U.S. Specifically, we use recently developed household projections for the U.S., household level economic data, and an economic growth model to develop energy and emissions scenarios that test whether explicitly accounting for demographic heterogeneity can substantially effect outcomes. Our approach differs in two key ways from existing energy and emissions projections: it is based on households, rather than individuals, as the demographic unit of analysis, and it incorporates demographic heterogeneity by grouping households into separate dynasties. Each dynasty has its own capital stock, labor supply, demand for consumption goods, and other household level variables. These multiple dynasties are introduced into the Population-Environment-Technology (PET) model (Dalton and Goulder, 2001). The model has seventeen consumer goods, including energy intensive goods like utilities and fuels, and less intensive goods such as education or health (Goulder, 1995). To calibrate the model, estimates of consumption expenditures, savings, asset accumulation, labor supply, and other variables for households in each household age group were derived from the U.S. Consumer Expenditure Survey (CES).

Analysis with the PET model is based on a set of household projections for alternative scenarios of future living arrangements that are intended to span a wide range of plausible outcomes in term of the future distribution of the population across households of various sizes and age compositions (Jiang and O'Neill, 2005). These household projections are carried out using the ProFamy model (Zeng et al., 1998, 2005). The ProFamy model projects individuals and households simultaneously and represents a substantial improvement over previous household projection models, which have relied on simple headship rate methods that do not give projections for some household types (Jiang and O'Neill, 2004).

We introduce demographic dynamics into the PET model by using ProFamy projection results to construct "cohorts" of households, where household age is defined by the age of the household head (Deaton, 1997). Household cohorts are then grouped into three infinitely-lived dynasties. Each dynasty contains households separated in age by the average length of a generation, taken to be 30 years. For example, eighty-year-old, fifty-year-old, and 20-year-old households are grouped in a single dynasty, based on the assumption that the younger households are, on average, descendents of the older households. Note that by increasing the length of a generation, our approach converges to the simplest OLG framework, with each dynasty representing only one cohort, excluding any altruistic behavior. Conversely, a shorter generational length is closer to a typical ILA framework. Therefore, heterogeneity in dynasties increases with generational length.

Households in different age groups are associated with distinct income and consumption levels, based on the CES data. Differences among age groups imply that each dynasty is associated with a specific pattern of income and consumption, based on its age distribution at each point in time. These differences have implications for energy demand, both directly and indirectly. In our results, the most important effects are caused by differentials in labor income across age groups that create complex dynamics for consumption and savings. These dynamics, and other relationships implied by the household projections and CES data, create interacting effects that feedback, and forward, to influence each dynasty's current and future consumption and savings decisions. A dynamic general equilibrium model is required to analyze these interacting effects on behavior, including how price changes for individual consumer goods affects tradeoffs between consumption and savings at the household level.

Using the PET model, we are able to decompose, and analyze, these general equilibrium effects. In particular, we use the PET model to analyze how household level variables respond to plausible changes in the age composition of U.S. households over the next 25-100 years. We also use the model to estimate how changes in household level variables affect the whole economy, and whether projected changes in the age composition of U.S. households could have a substantial influence on total energy demand, and CO2 emissions. Our results show that combining ILA and OLG approaches creates complicated dynamics for the age structure of each dynasty, which cause cycles in labor income that affect savings and consumption directly, and that have also indirect effects on energy demand. In general, we find that including heterogeneity among U.S. households reduces emissions, by almost 40% in our low population scenario. Effects of heterogeneity are less in other scenarios, and our results show that emissions are around 15% lower. Sec. 2 of the paper describes the PET model and household data. The household populations are described in Sec. 3, and results are given in Sec. 4. We conclude with a discussion of our analysis, results, and directions for future research in Sec. 5.

2. Population-Environment-Technology Model

The PET model is a global-scale dynamic computable general equilibrium model designed to analyze economic trade-offs associated with production and use of fossil fuels, and carbon dioxide emissions. A separate document, available from the authors, gives mathematical descriptions and data sources of the PET model (Dalton and Goulder, 2001). An overview is given here. The production component of the PET model has industries with many perfectly competitive firms that produce intermediate goods, including energy and materials, and final goods. Consumption and investment are final goods, and a government sector produces a final good. Production functions for each industry in the model have a capital-labor-energy-materials (KLEM) structure, with a nested constant elasticity of substitution form. There is a separate nest for energy inputs with oil and gas, coal, electricity, and refined petroleum. Other intermediate goods are aggregated, and produced by a single materials industry.

Each production function in the PET model has a substitution parameter for energy inputs that is assumed greater than the substitution parameter for KLEM-inputs, implying that energy inputs are more substitutable in production with one another, than energy is with other inputs. Estimating or assigning appropriate values for substitution parameters is an important topic in applied general equilibrium analysis, and has been the subject of past work with the PET model. We assign values here based on what we judge to be a standard configuration of the model, with the substitution elasticity for energy inputs set equal to 2.0 for all industries, implying modest substitutability of energy inputs, and an elasticity for KLEM inputs of 0.4, so that demand for these inputs is relatively inelastic. A wide range of assumptions regarding production functions, and substitution elasticities, are used in the energy-economic growth models that appear in the energy and climate change literature (e.g. Weyant and Hill, 1999). The substitution elasticities given above are consistent with this literature. Because oil and gas, and coal industries produce primary energy from fossil fuels, outputs of these industries account for CO2 emissions in the model.

The consumption component of the PET model is based on a population with many households that take prices as given. Each consumer good in the model is produced by a different industry, and one industry produces investment goods. Households demand consumer goods, and receive income by supplying capital and labor to producers. Households save by purchasing investment goods, and in the model, savings behavior is determined by solving an infinite horizon dynamic optimization problem for the dynasty to which the household belongs. Consumption and savings are described in more detail below.

Fig. 1: Overview of the PET Model. Households demand consumption and investment goods (C and I), and supply capital and labor (K and L). Final good producers supply C, I, and a government good (G). Intermediate goods producers supply energy and materials (E and M). The primary energy producers, which are coal, oil and gas industries, create CO2 emissions.

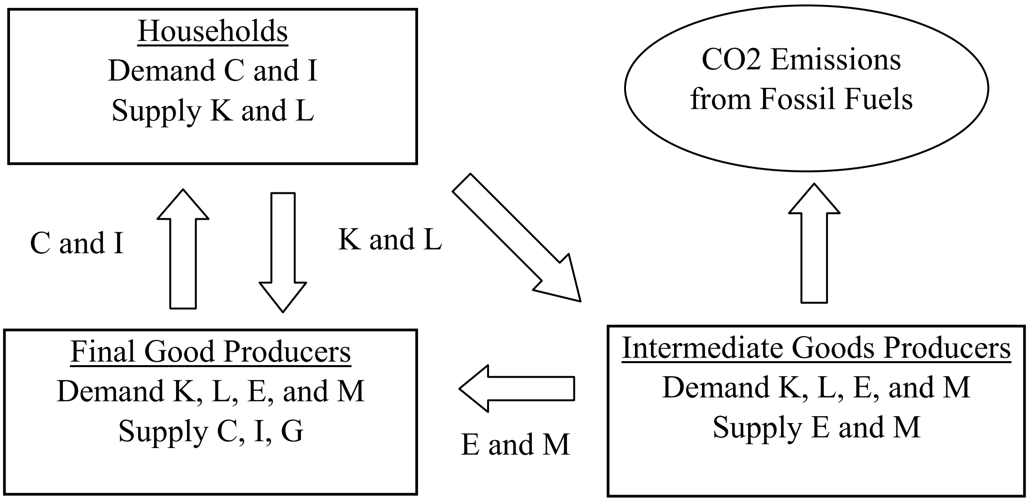

The following sections present parts of the PET model related to household consumption and savings. These parts of the model are central to our general equilibrium analysis of demographic factors that affect energy use and CO2 emissions. The PET model includes international trade, and can analyze different countries and world regions, but so far we have developed household data and population projections for the U.S. only. In the this paper, we are primarily interested in interactions between household consumption and factor supply within the U.S. economy. Therefore, we have omitted trade here to simplify the model, and isolate effects of demographic factors. We recognize that results are likely to be affected by this omission, but an initial assessment without effects of trade provides a useful benchmark against which further work can be compared, and still allows an informative comparison of results with and without demographic heterogeneity.

Household Consumption and Savings

Using age of the household head, we classify individuals in the population into three separate dynasties, indexed by i. Each dynasty consists of a large number of identical households, extending a standard assumption in neoclassical growth models that the population consists of a large number of identical households. Our extension to multiple household types, which result in separate dynasties, is consistent with neoclassical growth theory, and from the point of view of general equilibrium analysis, is more natural and interesting than assuming all households are the same.

Let n_{it} denote the total number of people living in each household type at time $t \geq 0$. Each household is endowed with labor l_{it}, and an initial stock of assets $\bar{k}_i$, which are expressed in per capita terms. Likewise, other variables are expressed in per capita terms, except where noted. Capital owned by different household is homogeneous, and perfectly substitutable in production. Households save by purchasing investment goods x_{it}, at price q_t. Investment is added to a stock of household assets, or capital k_{it}, which depreciates at rate $\delta > 0$ that is the same for all households, according to the law-of-motion

$$k_{it+1} = (1-\delta)k_{it} + x_{it}. \tag{1}$$

Household capital income is determined by the rental rate of capital, r_t, which is the same for all households. Labor's wage rate, w_t, is also assumed to be equal across households, so that differences in labor income are from variations in per capita labor supply or productivity. Labor is assumed, without loss of generality, to be the numeraire good in our analysis, and $w_t = 1$ for all t.

The model has 17 consumer goods, indexed by j. Per capita consumption for a household of type i, of good j, at date t is denoted by c_{ijt}. The price of each consumer good is denoted by p_{jt}. Households have a common discount factor $0 < \beta < 1$, and intertemporal substitution parameter $-\infty < \rho < 1$. Preferences for different consumer goods are characterized by a substitution parameter $-\infty < \sigma < 1$ that is also assumed to be the same for all households. The expenditure share parameters μ_{ijt} are differentiated for households, and can vary over time.

This paper evaluates the importance of demographic factors during a transition period, and does not to address possible effects on the long run equilibrium. Therefore, we assume that households are identical in the long run. The rationale for this assumption is to compare results in cases with and without demographic heterogeneity. In cases with demographic heterogeneity, values for per capita labor supply, l_{it}, and expenditure shares, μ_{ijt}, tend over time to equal values for all i. These long run conditions imply the terminal or long run balanced growth path equilibrium with demographic heterogeneity is the same as the reference case with representative households.

The transition period in the model includes time span of the demographic projections described below, and continues for another century, so that demographic heterogeneity accumulated through 2100 gradually disappears by 2200. Even without the long run restrictions on l_{it} and μ_{ijt}, if capital income tax rates ϕ_{it} are the same for each i, then other assumptions in the model, described below, imply that asset stocks of each dynasty, k_{it}, expressed in per capita terms, converge endogenously to equal values. In other words, per capita asset holdings are the same across dynasties in the long run, even if labor income or consumption patterns are different. This result depends on the tax rates for capital income being the same for each dynasty, but is not directly affected by the tax rate on labor income θ_{it}.

In the model, households receive per capita lump-sum transfers from the government, g_{it}, which is a net value so that negative values represent net payments. Private transfers, among households, are represented in the model, but are not distinguished here to save notation. The budget constraint for a household in dynasty i at date t is

$$\sum_{j=1}^{17} p_{jt}c_{ijt} + q_t x_{it} = (1-\theta_{it})w_t l_{it} + (1-\phi_{it})r_t k_{it} + g_{it}. \tag{2}$$

Demand for consumption goods is influenced by tradeoffs across goods at each t, and by dynamic factors related to savings and investment. Households take prices as given, are rational with forward-looking behavior, and in particular have perfect foresight of future values for all variables that affect their investment decisions. These variables include relevant prices, such as q_t and r_t, and future asset holdings by other households. Forward-looking behavior implies that equilibrium conditions in the model are dynamically consistent. Although the assumption of perfect foresight is restrictive in terms of the information structure of the model, this approach is preferable to an even more restrictive information structure, such as ignoring the value of future information altogether, which is implicit in models that use fixed savings rules to drive investment. Perfect foresight may be justified either by appealing to some type of certainty equivalence, or as the first step in an algorithm that converges to a rational expectations equilibrium (Fair and Taylor, 1983).

Tradeoffs across goods are described with a constant elasticity of substitution expenditure function, and over time by a constant elasticity of substitution intertemporal utility function. The PET model does not include leisure in household utility functions. Therefore, labor supply is inelastic, and determined by each household's labor endowment, l_{it}. Given prices, and subject to constraints (1) and (2), each household of type i chooses sequences of consumption $\{c^*_{ijt}\}$, for all j, and investment $\{x^*_{it}\}$, to maximize

$$\frac{1}{\rho}\sum_{t=1}^{\infty} \beta^t n_{it}\left(\sum_{j=1}^{17} \mu_{ijt} c_{ijt}^{\sigma}\right)^{\frac{\rho}{\sigma}}. \tag{3}$$

We describe two steps in the solution algorithm for each household's optimization problem to aid explanation of results below. Other parts of the dynamic algorithm are described in detail in the PET model's technical document (Dalton and Goulder, 2001). In the first step, demand for each consumer good is determined from prevailing prices by minimizing total expenditures, subject to a given level of utility, at each date t. A dual price index is used to calculate the marginal cost of consumption for each household, which varies across household types because of heterogeneity in expenditure shares. The price index dual to the expenditure function in (3) has a closed-form expression for each household type,

$$\overline{p}_{it} = \left(\sum_{j=1}^{17} \mu_{ijt}^{\frac{1}{1-\sigma}} p_{jt}^{\frac{\sigma}{\sigma-1}}\right)^{\frac{\sigma-1}{\sigma}}. \tag{4}$$

Each price index includes a weighted sum that depends on expenditure shares for each household type, and the prices of consumer goods faced by all households. In the general equilibrium PET model, prices of consumer goods are influenced in complex ways by changes in factor supply, including effects on labor with an aging population. The dual price index summarizes price changes across goods to indicate overall effects on the marginal cost of consumption for each household type. The marginal cost of consumption is compared to the price of investment goods to determine optimizing tradeoffs for households between consumption and savings at each t.

The second step in each household's problem is solving for paths of consumption expenditures and investment, for all t, that maximize (3). While price changes for consumer goods have static effects on the pattern of consumption, the tradeoff between consumption and savings affects model dynamics. The model's solution algorithm uses first-order conditions, or Euler equations, from maximizing (3), subject to (1) and (2), which imply

$$\frac{q_t}{\bar{p}_{it}}\left(\sum_{j=1}^{17}\mu_{ijt}c_{ijt}^{\sigma}\right)^{\frac{\sigma-1}{\sigma}} = \beta\left(\frac{r_{t+1}+(1-\delta)q_{t+1}}{\bar{p}_{it+1}}\right)\left(\sum_{j=1}^{17}\mu_{ijt+1}c_{ijt+1}^{\sigma}\right)^{\frac{\sigma-1}{\sigma}}. \qquad (5)$$

A solution to the Euler equations (5), capital law-of-motion (1), and budget constraint (2), which also satisfies a set of transversality conditions, is sufficient to maximize (3). Moreover, the solution is unique (Stokey and Lucas, 1989). The transversality conditions use the shadow value of capital for each household λ_{it}, and require that

$$\lim_{t\to\infty}\lambda_{it}k_{it} = 0. \qquad (6)$$

The transversality conditions ensure that each household's sequence of capital stocks is bounded. We use this fact to compute a steady-state level of the capital stock that is the same for all households, k^*, which satisfies conditions assumed above.

The PET model allows labor augmenting and other types of technical change. Let γ denote the long run rate of labor augmenting technical change. The long run condition used to compute the steady-state level of the capital stock is given by the steady state, or balanced growth path, ratio of the return on capital to the price of investment goods

$$(1-\phi_{it})\frac{r_t}{q_t} = \frac{1}{\beta}(1+\gamma)^{1-\rho} - (1-\delta). \qquad (7)$$

By assumptions above, parameters on the right-hand side of (7) do not depend on time, and are the same across household types. Because households face the same prices on capital and investment, if capital income tax rates are the same across households, then per capita asset accumulation is equal in the long run, which was mentioned above in the description of long run conditions. The PET model uses the Euler equations, and a variation of the Gauss-Seidel algorithm, to compute the dynamic transition from $\bar{k}_i$ to k^* for each household (Fair and Taylor, 1983).

Production, Consumption, and Income Data

The pattern of expenditure shares on energy and other inputs varies across industries. Benchmark input-output (IO) data for the model are from MacCracken, et al. (1999). These data are used to calibrate the PET model's production functions, and are derived from U.S. National Income and Product Accounts (NIPA), and other sources. To calibrate the model's household demand system, we use data from the U.S. Consumer Expenditure Survey (CES). The CES is a nationally representative survey composed of two parts: an Interview survey, and a Diary survey. In some cases, CES survey results are different from NIPA data. To resolve differences in the consumption and production data, we use CES data to determine aggregate expenditure shares of each consumer good at the economy-wide level, and apply these economy-wide shares to total consumption expenditures in order to determine the output of each consumer good industry. Conditional on the CES-determined output levels, demands for energy and other inputs of each industry are determined using input-output ratios derived from NIPA data. Additional details on the calibration procedure are available in the model's technical appendix.

The CES Interview survey has a sample size of approximately 5,500 households and is based on recall of expenditures over the past three months and income over the past year. It is aimed at capturing relatively large expenditures and those that occur on a regular basis. The Interview survey has a rotating panel design: each panel is interviewed for five consecutive calendar quarters and then dropped from the survey. A new panel is then introduced. Therefore about 20% of the addresses are new to the survey each quarter. The Diary survey is based on a written account of expenditures over the past two weeks, and is aimed at better capturing small, frequent purchases.

The CES data are used for economic analyses of consumption (e.g., Schmitt, 2004; Paulin, 2000). Details of our work with the CES data are described in a separate document (O'Neill, in preparation). In brief, data are integrated by choosing the most reliable source (Interview or Diary) for each consumption category, according to BLS. The CES categories are then aggregated into the 17 consumer good categories used in the PET model (Goulder, 1995). Mean annual per capita expenditures for these goods are calculated by household type. Household types are defined by characteristics of the "reference person" in the household, defined in the CES data as the first member mentioned by the respondent when asked to "Start with the name of the person or one of the persons who owns or rents the home." We use the reference person as our "householder" or "household head".

Values in Tab. 1 show how consumption of the 17 consumer goods varies across age groups using expenditure shares, or fraction of total expenditures, for each good. We use these expenditure shares as benchmark data for the PET model, which are converted to share parameters μ_{ijt} that calibrate the model's household demand system.

Tab.1: Expenditure Shares for Different Age Groups (%). Source: Consumer Expenditure Survey, 1998.

		Age of Household Head							
Good	Mean	15-25	25-35	35-45	45-55	55-65	65-75	75-85	85-95
1. Food	15.29	15.41	14.71	15.55	15.29	15.31	15.55	16.43	12.43
2. Alcohol	1.02	1.69	1.22	0.96	0.84	1.02	0.87	0.99	0.24
3. Tobacco	0.85	0.93	0.83	0.89	0.86	0.98	0.76	0.43	0.37
4. Utilities	4.22	2.90	3.74	4.01	3.98	4.69	5.53	6.71	6.07
5. Housing Services	20.50	21.54	23.80	21.69	18.82	17.80	16.19	17.63	33.63
6. Furnishings	4.48	3.76	4.29	4.35	4.84	5.07	4.66	4.16	1.21
7. Appliances	1.35	1.65	1.25	1.41	1.33	1.49	1.21	1.19	0.87
8. Clothing	4.93	5.35	5.31	5.28	5.40	4.07	4.00	2.85	1.59
9. Transportation	8.25	7.71	8.33	7.99	8.68	8.90	8.25	6.78	4.70
10. Motor Vehicles	12.01	14.47	13.06	12.65	12.57	11.20	9.42	5.08	5.12
11. Services	7.22	5.48	6.25	6.53	7.31	8.35	9.53	10.04	9.19
12. Financial Service:	2.99	1.93	2.95	3.20	2.80	3.55	2.88	3.26	1.58
13. Recreation	3.75	3.38	3.67	3.65	4.02	3.70	3.99	3.88	2.07
14. Nondurables	1.98	2.12	2.16	2.09	2.07	1.76	1.74	1.06	0.70
15. Fuels	3.40	3.50	3.29	3.40	3.50	3.59	3.42	3.02	2.25
16. Education	1.76	5.50	1.29	1.75	2.41	1.14	0.50	0.19	0.37
17. Health	5.99	2.69	3.84	4.60	5.28	7.39	11.51	16.30	17.62

To summarize key differences in expenditure patterns, we distinguish between younger versus older households. As discussed below, the household projections show that future compositional changes are driven by shares of the population at opposite ends of the age range in Tab. 1. As seen in the table, older households spend a substantially larger share of income than younger households on utilities, services, fuels, and health care, and a substantially smaller share on housing services, clothing, motor vehicles, non-durables, and education.

Since the most energy intensive goods are utilities and fuels, expenditure patterns in Tab. 1 imply that aggregated consumption in older households is more energy intensive than consumption in younger households. The utilities category is about two-thirds electricity, with the remaining third split between natural gas, and payments for water and sewer services. Electricity demand is driven principally by appliance use, and natural gas consumption by space conditioning (EIA, 2004). Although older households spend a larger fraction of income on utilities, absolute levels of expenditures are roughly the same across the two household types when income differences are taken into account, which is consistent with previous work on patterns in residential energy use (Bin and Dowlatabadi, 2005). The fuels category is 80-90% gasoline, and is therefore influenced mainly by car use. The remainder is split primarily between fuel oil and natural gas. While old households spend a larger share of per capita income on fuels than young households, income differences imply the absolute level of fuel use is substantially smaller, which is consistent with other work (O'Neill and Chen, 2002).

Tab.2: Total Consumption Expenditures, Savings, Income, Government (Gov.) and Household (HH) Transfers, and Income Tax Rates for Different Age Groups (per capita values in 1998 dollars). Source: Consumer Expenditure Survey, 1998.

		Age of Household Head							
	Mean	15-25	25-35	35-45	45-55	55-65	65-75	75-85	85-95
Consumption	13,214	11,355	11,824	12,175	15,987	15,336	14,156	12,555	12,084
Savings	3,316	1,080	2,253	3,442	4,674	5,020	2,299	3,036	6,808
Labor income	14,198	9,659	14,753	15,278	21,583	14,440	4,014	1,324	1,325
Capital income	2,020	192	769	1,336	2,081	4,115	4,998	5,019	3,777
Capital	33,377	3,076	5,894	17,040	43,867	66,295	95,910	87,351	83,277
Gov. transfers	371	-440	-882	-811	-1,066	1,270	6,098	7,957	7,384
HH transfers	48	342	210	32	7	65	-244	-364	-474
Capital tax rate	0.23	0.39	0.34	0.31	0.30	0.17	0.16	0.15	0.17
Labor tax rate	0.09	0.06	0.08	0.08	0.10	0.10	0.18	0.26	0.18

Government transfers in Tab. 2 include social security, workers compensation, unemployment benefits, and other kinds of public assistance, and these favor older households in per capita terms by a wide margin. Savings includes retirement contributions, down payments on purchases of property, mortgage payments, capital improvements, and investments in own businesses or farms. Assets include the value of financial accounts and securities plus the equity share of property.

3. Household Projections and Dynasties

In Tab. 3, we present household projections from the ProFamy model for three scenarios. The ProFamy projections run from 2000 to 2100; for simplicity population is assumed to stay constant after 2100 in our analysis. Values in the table give total population in each year of the projection, followed by percentage shares of the population living in households of different ages, in order to clearly distinguish differences in both scale and composition across scenarios. Work with the ProFamy model, and methods for developing the U.S. household projections, are described in a separate paper (Jiang and O'Neill, 2005), and an overview is given here.

The scenarios we use are based on a set of plausible demographic assumptions for fertility, mortality, migration, and union formation and dissolution rates that span a wide range of outcomes in terms of population size, age structure, and household size. Assumptions for demographic rates, and how to combine them in each scenario, were chosen in order to produce one scenario with relatively small, old households (our low scenario), one scenario with relatively large, young households (our high scenario), and one scenario with moderate outcomes (our medium scenario). Between the three scenarios, population size varies by more than a factor of four. An important property of the projections is that the age composition of households in the low scenario is remarkably different from the pattern in high and medium scenarios, with people living in older households making up a much greater percentage of the population under conditions of low fertility and mortality.

Tab.3: U.S. Population (millions) and Shares (%) Living in Households of Different Ages in High, Medium, and Low Population Scenarios. Source: ProFamy model.

		Population Shares (%) by Age of Household Head								
Year	Population	15-25	25-35	35-45	45-55	55-65	65-75	75-85	85-95	95+
High Population Scenario										
2000	281.4	6.5	23.0	30.5	19.7	9.4	6.5	3.6	0.9	0.1
2010	316.6	7.3	21.7	25.0	20.7	13.6	7.0	3.6	1.0	0.1
2020	361.2	6.2	21.8	24.5	17.5	14.6	10.1	4.2	1.0	0.1
2030	414.3	6.4	19.9	24.9	17.4	12.5	11.2	6.3	1.3	0.2
2040	475.0	6.7	20.5	23.3	17.7	12.4	9.6	7.3	2.3	0.2
2050	546.3	6.9	20.8	23.9	16.5	12.5	9.5	6.5	3.0	0.4
2060	630.2	6.9	20.6	24.0	16.9	11.7	9.6	6.6	3.0	0.6
2070	728.3	7.0	20.5	23.7	17.0	12.0	9.1	6.7	3.3	0.7
2080	841.5	6.9	20.4	23.5	16.8	12.1	9.3	6.5	3.6	0.9
2090	970.4	6.9	20.2	23.3	16.7	12.0	9.4	6.8	3.6	1.1
2100	1117.0	6.8	20.1	23.1	16.6	11.9	9.4	7.0	3.8	1.2
Medium Population Scenario										
2000	281.4	6.5	23.0	30.5	19.7	9.4	6.5	3.6	0.9	0.1
2010	307.8	6.7	21.0	25.2	21.3	13.9	7.1	3.6	1.1	0.1
2020	333.8	5.8	20.6	23.9	18.0	15.4	10.6	4.4	1.1	0.2
2030	360.6	5.8	18.9	23.9	17.4	13.2	12.2	6.9	1.4	0.2
2040	387.8	5.6	19.2	22.5	17.7	12.9	10.7	8.5	2.6	0.3
2050	414.5	5.4	19.0	22.9	16.8	13.2	10.7	7.8	3.7	0.5
2060	442.3	5.3	18.6	22.7	17.2	12.6	11.1	7.9	3.9	0.7
2070	472.3	5.2	18.4	22.2	17.0	13.0	10.7	8.4	4.3	0.8
2080	504.9	5.0	18.1	22.1	16.8	12.9	11.0	8.2	4.8	1.1
2090	538.3	4.9	17.7	21.8	16.8	12.8	11.1	8.7	5.0	1.4
2100	573.0	4.7	17.4	21.5	16.6	12.8	11.1	8.9	5.4	1.7
Low Population Scenario										
2000	281.4	6.5	23.0	30.5	19.7	9.4	6.5	3.6	0.9	0.1
2010	303.7	6.8	21.0	24.9	21.1	14.0	7.2	3.7	1.1	0.1
2020	321.2	5.3	20.3	23.6	17.9	15.7	11.1	4.7	1.2	0.2
2030	331.4	4.5	17.6	23.6	17.5	13.7	13.1	7.8	1.8	0.3
2040	334.1	3.9	16.4	21.8	18.0	13.8	12.0	10.0	3.7	0.4
2050	328.5	3.3	14.9	21.0	17.2	14.8	12.6	9.7	5.6	0.9
2060	317.9	2.9	13.4	19.8	17.0	14.6	14.1	10.7	6.0	1.6
2070	305.0	2.5	12.0	18.4	16.5	15.0	14.2	12.3	7.0	2.0
2080	287.7	2.3	10.7	16.9	15.7	15.0	15.1	12.9	8.5	2.9
2090	269.9	2.1	10.1	15.5	14.8	14.7	15.5	14.0	9.4	3.9
2100	250.5	2.0	9.5	14.9	13.7	14.1	15.5	14.8	10.7	4.8

We use the population distribution by household age to construct dynasties that consist of a series of cohorts of households of different ages at each point in time. The procedure for constructing cohorts and dynasties from the ProFamy projections is outlined in Fig. 2.

Figure 2: Cohort Structure of Population in the PET Model. Each color shows cohorts that are members of the same dynasty (i.e. household type). Thus, the first household type is dynasty 1, consisting of cohorts 1a-f (boxes). The second household type is dynasty 2, consisting of cohorts 2a-f (circles). The third household type is dynasty 3, consisting of cohorts 3a-e (triangles).

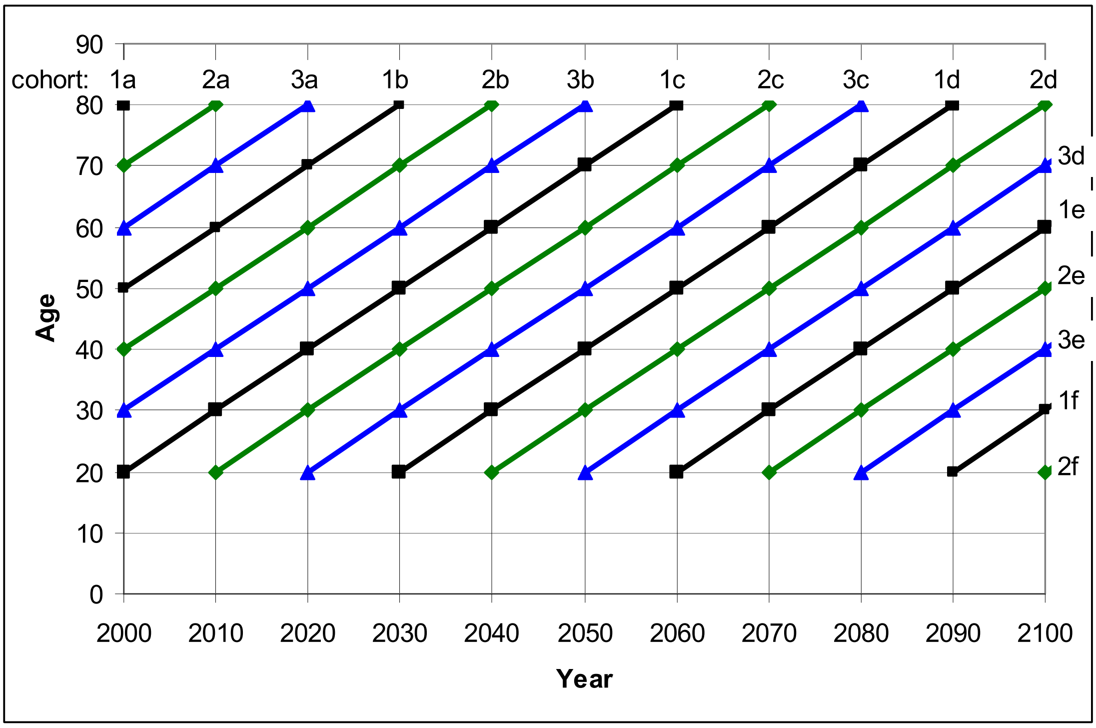

This procedure implies that each dynasty has a specific household age distribution at each point in time, based on the population size of each cohort. We use benchmark data from the Consumer Expenditure Survey for households of different ages to derive weighted-mean per capita labor supply and expenditure shares for consumer goods for each dynasty over time, with weights given by the proportion of the population living in households of different ages. Total labor supply for each dynasty is then calculated as the product of per capita labor supply and population size. Expenditure shares are translated into share parameters for the PET model's demand system at the time of model calibration. In this way, the ProFamy projections are used to determine the changing composition of the population across household types within each dynasty; the CES data are then used to calculate mean characteristics of households within each dynasty that change over time to reflect the changing demographic composition.

4. Results

The primary objective of this paper is to evaluate the extent to which population aging in the United States could affect long-term energy use, and CO2 emissions projections over the next 50-100 years. For this objective, we conducted forty-eight simulations with the PET model to analyze the effects on emissions of population aging in the United States over the next hundred years. To isolate the effects of population aging, technical change is not included in these simulations, nor is international trade. Obviously, both are important and the subject of ongoing work that is discussed below. In the simulations, households may exhibit heterogeneity, or be influenced by other dynamic factors, through 2100. After 2100, differences between households are assumed to gradually disappear until a stationary equilibrium with identical households is reached at 2200, the simulation horizon used for results below. Results in the first 50-100 years of simulations are of interest. The transition period after 2100 is used to stabilize the approach to the long run equilibrium, and results from this period are not presented. Households anticipate convergence to the long run equilibrium, and assumptions about the long run conditions can affect results, even in the first decades of a simulation, through forward-looking savings behavior.

The simulations are divided into four groups that include varying degrees of demographic heterogeneity. At one extreme, the group of simulations is based on mean values that are neutral with respect to population aging. At the other, the group of simulations includes three dynasties with age-specific demographic heterogeneity in both consumption patterns and labor supply. In between these extremes, a group of simulations includes heterogeneity in consumption patterns, but not labor supply, and another group includes only effects of aging on labor supply. The twelve simulations in each group are based on the low, medium, and high household projections described above, and are stratified by four combinations of household substitution parameters for sensitivity analysis. To evaluate various effects of population aging, comparisons are made between the four groups. In particular, results compare simulations that include effects of population aging to the group with identical or representative households, which do not include effects of aging

Our primary comparison is for simulations demographic heterogeneity, and representative households without aging. This first comparison brackets the range of results with demographic heterogeneity for the three household projections, and is based on a set of reference values for the household substitution parameters. Other results are refinements of the first comparison that either distinguish between different sources of demographic heterogeneity, or test sensitivity to alternative values of the substitution parameters.

The second comparison considers heterogeneity in age-specific consumption patterns, with labor supply that is equal to the level of representative households without aging. Results turn out to be minor for this comparison because differences in expenditure shares for each good are generally similar among age groups, and the consumption pattern assigned to each dynasty is a weighted average that tends to smooth these differences even more.

The third comparison addresses the important issue of whether separate dynasties are needed to account for the effects of population aging on energy use and CO2 emissions. The main methodological contribution of this paper is the development of an energy-economic growth model with multiple dynasties to explicitly incorporate the general equilibrium effects of demographic events, such as population aging. However, other energy-economic growth models have identical households and a single dynasty structure. The third comparison addresses whether these other models can incorporate the main effects of population aging simply by scaling the labor supply of representative households.

Fig. 3: Per Capita Dynamics for Labor Income (top) and Capital Stock (bottom) in Thousands of Year 2000 Dollars for the 3 Dynasties in the Low (hatched), Medium (light solid), and High (dark solid) Population Scenarios. Note that dynasty 1 has the largest labor income in 2000, dynasty 2 has the lowest, and dynasty 3 is in the middle. The order for dynasties is the same in 2000 for capital stocks.

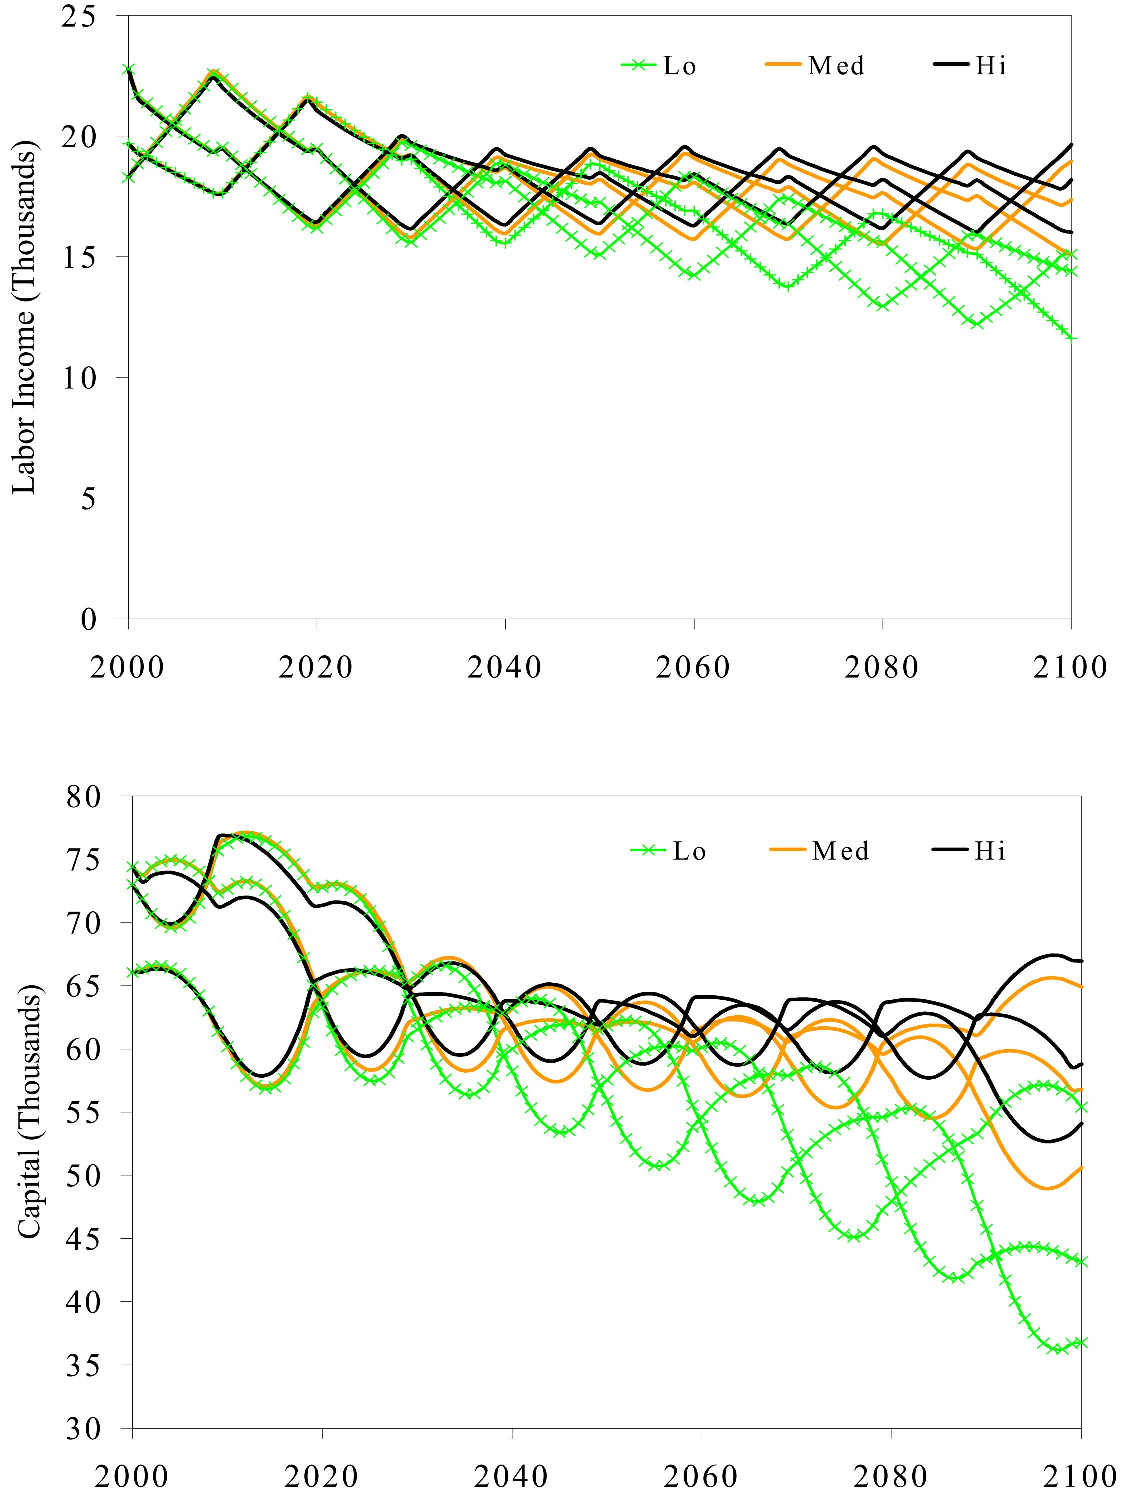

Heterogeneous versus Representative Households

The model configuration with heterogeneous households has three dynasties that follow the dynamics in Fig. 2. For each dynasty, age-specific weights for consumption expenditures are derived from values in Tab. 1, and weights for labor supply are derived from Tab. 2. Initial capital for each dynasty is also derived from Tab. 2. Benchmark values in the PET model for transfers and income tax rates are set to zero in the simulations for this paper to simplify the interpretation of results.

Values in Tab. 2, the procedure in Fig. 2, and the household projections in Tab. 3 have interesting implications for the dynamics of labor income, which cycle and exhibit a downward trend due to the effects of aging. The first graph in Fig. 3 shows these results using per capita labor income for the three dynasties, in the three population scenarios. Each curve represents a single dynasty's per capita values over time, in thousands of year 2000 dollars. Note that per capita labor income in the representative household configuration for each scenario would appear as a flat line at $20,000 per year. The dynasties can be identified from their supply of labor in 2000. For example, dynasty 1 has a cohort in the 45-55 age group in 2000, and this group has the largest per capita labor income. Thus, dynasty 1 has the greatest labor income in 2000, which is $22,785 per person. The labor income in 2000 of dynasty 2 is $18,286 per person, and dynasty 3 has a slightly larger value, $19,684 per person. The patterns of labor income in Fig. 3 directly affect the dynamics of savings and capital accumulation for each dynasty, which are presented in the second graph. The capital stock for each dynasty follows the same order as labor income in 2000. Dynasty 1 has the largest value, $74,407 per person in 2000, dynasty 3 is next with $73,004 per person, and dynasty 2 has the lowest value, $66,044 per person. Although there are minor variations at the beginning of simulations, capital in the representative household configuration would appear as a flat line, at $70,000 per person.

The model configuration for representative households without aging has three identical dynasties to be consistent with the structure of the model used for other simulations. Each dynasty has one-third of the total population at each point in time, and per capita expenditure shares are equal to the mean values in Tab. 1, which are estimates for the U.S. population. Labor supply, consumption expenditures, and other variables are equal in per capita terms, and are derived from mean values in Tab. 2. These conditions are equivalent to scaling a representative household to the level of an exogenous population projection.

The graphs in Fig. 4 compare results for total CO_2 emissions, and per capita CO_2 emissions, over time for heterogeneous and representative households. Total emissions with heterogeneous households are driven by changes in both the size and age composition of the population. Results show that total emissions with heterogeneous households range from 0.9 to 5.1 billion metric tons per year by 2100. For representative households, changes in emissions over time are due to changes in the size of the population, and emissions range from 1.4 to 5.9 billion metric tons per year by 2100 in the three population scenarios. Because of aging, heterogeneity leads to lower emissions in each population scenario. Differences between emissions in simulations with heterogeneous and representative households are a combination of direct effects from changes in labor supply due to aging, and indirect or general equilibrium effects from changes in capital accumulation, prices, or other factors. Aging implies fewer young workers, whose per capita labor contribution tends to be greater than the population mean. Hence, aging implies a reduction in aggregate labor supply for a given population size.

Fig. 4: Range of CO2 emissions and Per Capita CO2 Emissions for Heterogeneous (Het) and Representative (Rep) Households in Low (Lo), Medium (Med), and High (Hi) Population Scenarios.

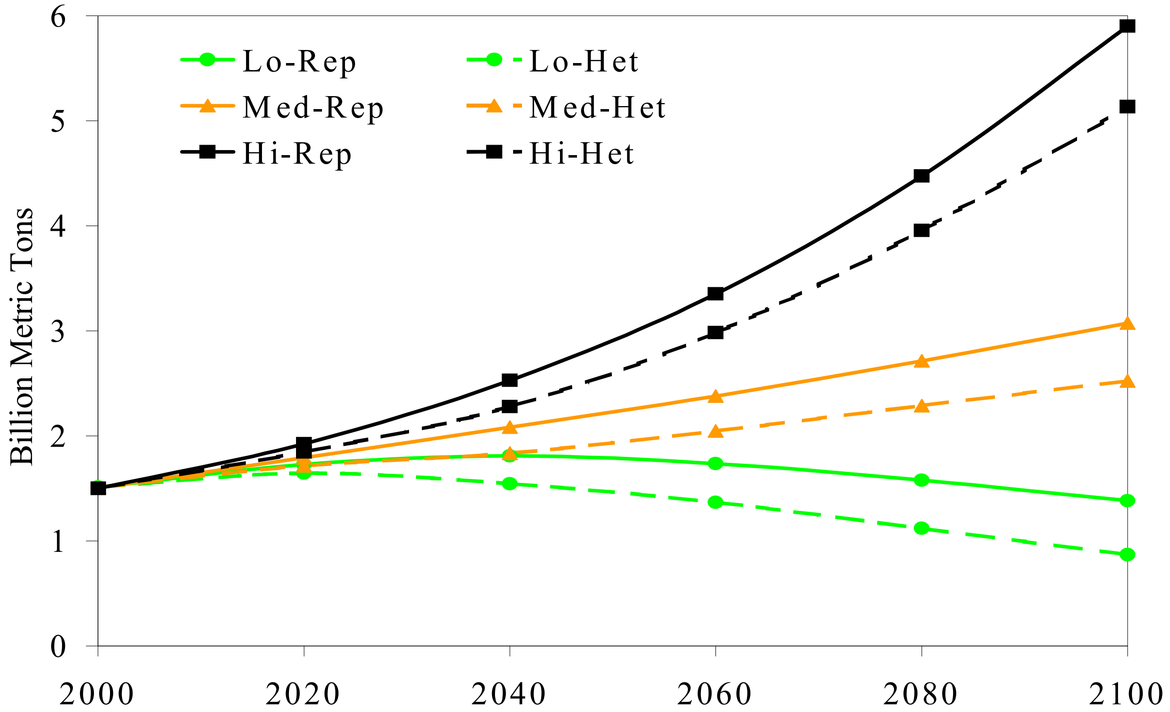

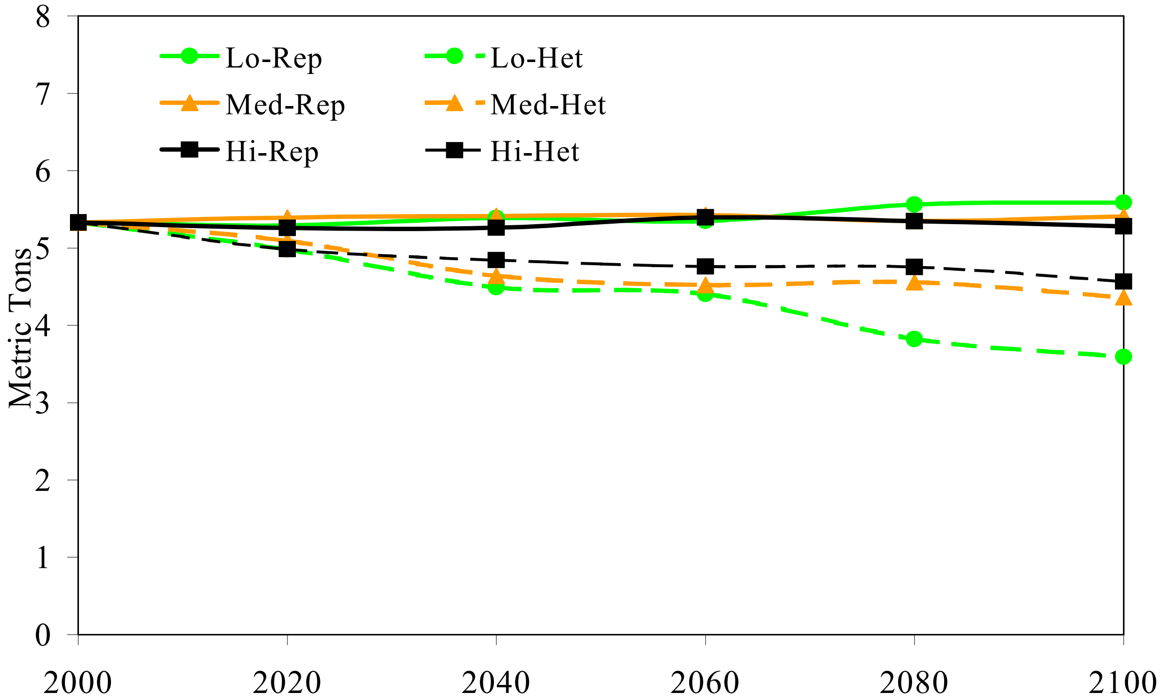

The second graph in Fig. 4 shows per capita emissions with heterogeneous and representative households in each population scenario. Because total population within each scenario is the same, differences in per capita emissions in Fig. 4 are caused exclusively by changes in total emissions. Per capita growth in output, measured by gross domestic product (GDP) per person, is essentially zero with representative households, and changes in carbon intensity, represented by CO2 emissions per dollar of GDP, are also minor. Consequently, per capita emissions with representative households are essentially constant over time and across population scenarios, around 5.3 tons per person. The second graph in Fig. 4 shows that demographic heterogeneity in the low population scenario reduces per capita emissions by about 2.0 metric tons per person by 2100. Per capita labor supply, which is a weighted average over different age groups, are similar in medium and high population scenarios, which is why per capita emissions are relatively close. The scarcity of young workers drives results in the low population scenario, which has substantial effects on per capita emissions. The range in per capita emissions between low and high population scenarios is about one ton per person by 2100, but because of lags in the onset of population aging, these effects are not apparent until 2060.

Population Aging and Representative Households

A second model configuration with identical households is used to evaluate whether the main effects of population aging can be incorporated in the model simply by scaling the labor supply of representative households. The representative household configuration with aging has the same level of aggregate labor as the model with heterogeneous households. For the reference values of our household substitution parameters, total emissions for the representative household configuration with aging are about 2-10% greater than with heterogeneous households, in the high and low population scenarios respectively. However for other values of the substitution parameters, total emissions are about 2-8% lower. The exact values are presented with other results of the sensitivity analysis.

Sensitivity Analysis for Household Substitution Parameters

The substitution parameters ρ and σ in each household's utility function from (3) directly affect results. Our reference value for households' intertemporal substitution parameter is $\rho = 0.5$, or an elasticity of $1/(1-\rho) = 2.0$. This value is taken from Goulder (1995), who reports it is in the range of estimates obtained by Hall (1988), and Lawrance (1991). Our reference value for the substitution elasticity of consumer goods is also 2.0, or $\sigma = 0.5$. We conduct a sensitivity analysis to examine how results with less elastic values for ρ and σ differ.

Values for the intertemporal substitution elasticity are important in macroeconomic models (Guvenen, 2003), and obtaining reliable and consistent estimates has been a problem. Beudry and van Wincoop (1996) use panel data for U.S. states, and report estimates close to a value of one, and significantly different from zero. Note that an elasticity of one implies a ρ of zero, which is equivalent in the limit to the natural log utility function. An elasticity of zero implies ρ approaches minus infinity, or the Leontief case of perfect complements. A recent study, using a new econometric approach, estimates intertemporal substitution elasticities less than one, but not significantly different from zero (Yogo, 2004). Therefore, negative values for ρ seem plausible. Inelastic values for σ may also be plausible

Values in Tab. 3 summarize effects of population aging on U.S. CO2 emissions, relative to the first representative household configuration in the three population scenarios, and for alternative values of the substitution parameters. The values in the table are intended to include a plausible range with substitutes and complements in consumption. Our inelastic value for the consumption substitution parameter is $\sigma = -3.0$, or an elasticity of 0.25. We also consider an inelastic value for the intertemporal substitution parameter of $\rho = -3.0$.

Tab.3: Percentage Differences in U.S. CO2 Emissions with Population Aging Compared to the First Representative Household Configuration in Low (L), Medium (M), and High (H) Population Scenarios, and for Alternative Values of the Intertemporal (ρ) and Consumption (σ) Substitution Parameters.

	Rep. W/Aging			Heterogeneous			Rep. W/Aging			Heterogeneous		
Year	L	M	H	L	M	H	L	M	H	L	M	H
	$\rho = 0.5, \sigma = 0.5$						$\rho = 0.5, \sigma = -3.0$					
2000	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.2	0.1	0.1
2020	-4.9	-4.4	-4.1	-4.8	-4.2	-4.1	-4.9	-4.4	-4.1	-5.1	-4.5	-4.3
2040	-12.7	-10.1	-8.3	-14.6	-11.9	-9.8	-12.7	-10.2	-8.3	-14.8	-12.0	-9.9
2060	-18.3	-11.8	-9.2	-21.2	-14.0	-11.0	-18.2	-11.8	-9.2	-21.4	-14.0	-11.0
2080	-25.0	-13.2	-9.7	-29.0	-15.7	-11.5	-24.9	-13.2	-9.7	-29.3	-15.7	-11.5
2100	-31.5	-14.9	-10.8	-37.2	-17.9	-13.0	-31.6	-14.9	-10.8	-37.4	-17.9	-13.0
	$\rho = -3.0, \sigma = 0.5$						$\rho = -3.0, \sigma = -3.0$					
2000	0.2	0.2	0.2	-0.1	0.1	0.0	0.3	0.2	0.2	0.1	0.1	0.1
2020	-5.5	-4.8	-4.3	-1.0	-0.8	-1.0	-5.3	-4.7	-4.3	-2.6	-2.2	-2.2
2040	-12.6	-9.8	-8.1	-8.6	-7.3	-6.3	-12.6	-9.9	-8.1	-10.5	-8.7	-7.3
2060	-18.4	-11.7	-9.1	-13.7	-10.0	-8.0	-18.3	-11.7	-9.1	-16.2	-11.2	-8.9
2080	-25.1	-13.3	-9.8	-19.0	-11.3	-8.4	-25.1	-13.2	-9.7	-22.3	-12.6	-9.3
2100	-31.1	-14.8	-10.7	-25.3	-13.0	-9.5	-31.3	-14.8	-10.8	-29.0	-14.4	-10.5

Values in Tab. 3 provide information to summarize comparisons among the model configurations, substitution parameters, and population scenarios. Our primary comparison, is between the two model configurations that consider population aging. The second representative household configuration includes effects of aging by scaling aggregate labor supply so that it is equal to aggregate labor supply for the simulations with heterogeneous households. Values in the table for the case with $\rho = 0.5$ and $\sigma = 0.5$ relate directly to results in Fig. 4. In this case, for the low population scenario, emissions are about 37% less in 2100 with heterogeneous households relative to the representative household configuration without aging. Most of this difference, about 85%, is due directly to changes in labor supply associated with population aging, since emissions in the second representative household configuration are more than 31% less than in the first household configuration. The remaining difference occurs through capital dynamics and general equilibrium effects. The effects of population aging on emissions are smaller for medium and high population scenarios, about 18% and 13% respectively, because the effects of population aging are reduced in these scenarios.

5. Discussion

While demographic events are usually treated implicitly in energy-economic growth models, we directly link population age structure to consumption and energy use in a dynamic general equilibrium model to estimate effects on future CO2 emissions. First, we analyzed CES data to develop age profiles of expenditure patterns, labor income, and other variables. Second, we developed population projections for each age group using the ProFamy model. Third, we incorporated the population projections, and a household demand framework calibrated with the CES data, into the PET model.

For results in this paper, we ran four sets of simulations with the PET model, based on representative and heterogeneous household configurations of the model in low, medium, and high population scenarios. Each set of simulations corresponds to a particular set of assumptions about the substitutability of consumption over goods, and over time. Our results compare two types of heterogeneous households to representative households. In the first type, expenditure shares for different consumer goods depend on age of the household head. The second type has heterogeneity in expenditure shares, and in sources of household income, including capital and labor. The first type of heterogeneity affects the composition of demand, but our results show these effects are negligible.

In contrast, effects associated with age-specific heterogeneity in labor income have strong negative effects on CO2 emissions. The range of results across population scenarios is a 10-37% decrease in emissions, per year, by 2100. These effects are accompanied by changes in many other variables, including per capita emissions, GDP, demand for different types of energy, and demand for individual consumer goods. Therefore, results in this paper support further consideration of demographic factors in emissions projections, and suggest these factors may be critical to the development of new emissions scenarios, particularly those based on low population projections for the U.S., because effects of aging are most important in this scenario.

However, our model and current approach are based on simplifying assumptions that ignore feedbacks, which could dampen, or ameliorate, economic effects of an aging population. For example, this paper considers population age structure, but changes in household size, immigration, or other demographic factors may also be important. Resolving these issues is beyond the scope of this paper, the aim of which is to isolate effects of population heterogeneity for age, the most widely recognized demographic factor, in a dynamic general equilibrium setting, and establish an initial set of empirical bounds on these effects. This initial assessment provides an informative comparison of results with and without demographic heterogeneity, in the absence of potentially confounding factors such as technical change or international trade, and thus providing a useful benchmark against which further work can be compared. Results in this paper suggest that demographic factors have the potential to substantially affect long-term emissions for the U.S., and motivate further study of relationships between demographic change, economic growth, and energy use.

In future work, we plan to relax many of the limiting assumptions used in this paper. First, the analysis in this paper does not include effects of technical change, even though it is obviously important. The PET model can handle many types of technical change, including Harrod, Hicks, and Solow neutral types. A straightforward extension of results in this paper is to include technical change in the analysis, including sources of technical change in energy inputs that imply empirically realistic declines over time in the carbon intensity of output. Gains in labor productivity are well known to be a main driver of economic growth. In standard economic growth models, these gains imply changes in scale, often with no compositional effects. In this case, potential losses in labor productivity due to population aging, as demonstrated by results in this paper, may be partially or fully offset by labor augmenting technical change. However, carbon intensity is decreasing over time, which implies other types of technical change besides the labor augmenting variety are present. The net effect of general equilibrium interactions between changes in labor productivity due to aging, and these other types of technical change, are impossible to predict without further analysis, which we will conduct in future work.

We do not have good empirical estimates of several key elasticity values used in our analysis. These values are associated with the substitutability of consumption over time, and across different goods, including energy intensive goods like utilities and fuels, and less intensive goods such as education or health. Our results are sensitive to the value of the intertemporal substitution elasticity, and in some cases, the substitution elasticity for consumer goods. We plan to use data from the U.S. Consumer Expenditure Survey (CES) to estimate substitution elasticities for consumer goods, and test hypotheses about whether these vary among age groups. We will use these estimates in future work.

An important limitation of our approach, currently, is that labor supply for each age group is inelastic, and does not respond to changes in real wages or other variables. Clearly, increasing labor supply is a plausible response by older age groups to changes in real wages, policy, life expectancy, or other factors that provide an incentive to delay retirement, or otherwise continue working. The standard way to endogenize labor supply in general equilibrium models is to introduce a term for leisure in households' utility functions. The PET model could be modified in this way by introducing another elasticity parameter into our analysis, which may not provide an improvement over the current approach. Before labor supply is endogenous in the PET model, we will conduct a thorough analysis of CES data to infer a reasonable range of alternatives for age profiles of labor supply.

Finally, another important restriction in this paper is that results are for the U.S. only, under assumptions of a closed economy. The PET model has the structure to include multiple countries or regions, and international trade, but we do not currently have demographic data for other countries to support the type of analysis in this paper. We are currently developing demographic data for China, India, and other countries. The data demands of future work are expected to be intensive, including household projections, household economic data, and production data for different consumer good industries. Results with international trade are difficult to predict *a priori*, and will depend on the countries being compared. Countries that differ in age distribution will gain from trade, since labor intensive goods can be exported by the country with the younger population. So, we might expect international trade to diminish the effects of aging on energy use and CO2 emissions, relative to an autarky situation without trade. However, population aging is a global event (O'Neill, MacKellar, and Lutz, 2001). Extrapolating results in this paper suggests that we might find a general upward bias in global emissions projections, which further motivates the need for research in this area.

6. References

Arthur, W. and G. McNicoll. 1977. Optimal time paths with age dependence, *Review of Economic Studies*, 44, 111-123.

Auerbach, A., and L. Kotlikoff. 1987. *Dynamic Fiscal Policy*. Cambridge, MA: Cambridge University Press.

Azar C. and T. Sterner. 1996. Discounting and distributional considerations in the context of global warming, *Ecological Economics*, **19**, 169-184.

Barro, R.1974. Are government bonds net wealth? *J. Political Economy*, **82**, 1095-1117.

Beaudry, P and E. van Wincoop. 1996. The intertemporal elasticity of substitution: An exploration using a U.S. panel of state data, *Economica*, New Series, 63 (251), 495-512.

Bin, S. and Dowlatabadi, H. 2005. Consumer lifestyle approach to US energy use and the related CO_2 emissions, *Energy Policy* 33, 197-208.

Birdsall, N., A. Kelley, and S. Sinding. 2001. Population Matters: Demographic Change, Economic Growth, and Poverty in the Developing World. Oxford University Press, Oxford, UK.

Blanchard, O. 1985. Debt, deficits, and finite horizons, *J. Political Economy*, **93**, 223-247.

Blanchard, O. and S. Fischer. 1987. Lectures on Macroeconomics. Cambridge: MIT Press.
Bryant, R. and W. McKibbin. 1998. Issues in modeling the global dimensions of demographic change, Brookings Discussion Papers in International Economics, 141, December.
Cline, W. 1992. Global Warming: The Economic Stakes. Washington D.C.: Institute for International Economics.
Curran, S. and A. de Sherbinin. 2004. Completing the picture: The challenges of bringing "consumption" in the population-environment equation. *Population and Environment* 26(2), 107-131.
Cutler, D., J. Poterba, L. Sheiner, L. Summers. 1990. *An Aging Society: Opportunity or Challenge?* Brookings Papers on Economic Activity, 1/1990.
Dalton, M. and L. Goulder. 2001. An intertemporal general equilibrium model for analyzing global interactions between population, the environment and technology: PET model structure and data. Unpublished document, California State University Monterey Bay. (science.csumb.edu/~mdalton/EPA/pet.pdf)
Deaton, A. 1997. *The Analysis of Household Surveys*. The World Bank, Washington, DC.
Energy Information Administration (EIA), 2004. *Annual Energy Review 2003*. Energy Information Administration, U.S. Department of Energy, Washington, DC.
Fair, R. and J. Taylor. 1983. Solution and maximum likelihood estimation of dynamic nonlinear rational expectations models, *Econometrica*, 51(4), 1169-1185.
Farmer, M., A. Randall. 1997. Policies for sustainability: Lessons from an overlapping generations model, *Land Economics*, **73**(4).
Geanakoplos, J., and H. Polemarchakis. 1991. Overlapping generations. Ch. 35 in Handbook of Mathematical Economics IV, W. Hildenbrand and H. Sonnenschein (eds.), Amsterdam: North-Holland.
Gerlagh, R. and B. van der Zwaan. 2000. Overlapping generations versus infinitely-lived agent the case of global warming, *in* The Long-Term Economics of Climate Change, Vol. 3, 301-327, R. Howarth and D. Hall (eds.), Stamford CT: JAI Press.
Gerlagh, R. and B. van der Zwaan. 2001. The effects of ageing and an environmental trust fund in an overlapping generations model on carbon emissions reduction, *Ecological Economics*, **36** (2), 311-326.
Goulder, L. 1995. Effects of carbon taxes in an economy with prior tax distortions: an intertemporal general equilibrium analysis, *J. Env. Econ. Management*, 29, 271-297.
Guest, R. and I. McDonald. 2001. Ageing, immigration and optimal national saving in australia, *Economic Record* 77, 237, 117-134.
Guvenen, F. 2003. Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective. RCER Working Paper No. 491, University of Rochester, January.
Hall, R. 1988. Intertemporal substitution in consumption, *J. Political Economy*, **96** (2), 339-357.
Howarth, R. 1996. Climate change and overlapping generations, *Contemporary Economic Policy*, **14**(4), 100-111.
Howarth, R. 1998. An overlapping generations model of climate-economy interactions, *Scand. J. Economics*, **100**(3), 575-591.
Howarth, R. and R. Norgaard. 1992. Environmental valuation under sustainable development, *American Economic Review*, **82**(2), 473-477.
Jiang, L., B. O'Neill. 2004. Towards a new model for probabilistic household forecasts, *International Statistical Review*, **74**(1), 51-64.
Jiang, L., B. O'Neill. 2005. Impact of demographic events on household formation and dissolution in the United States, working paper, Brown University.

Kehoe, T. 1991. Computation and multiplicity of equilibria. Ch. 38 in Handbook of Mathematical Economics IV, W. Hildenbrand and H. Sonnenschein (eds.), Amsterdam: North-Holland.

Lareau, T. and J. Darmstadter. 1983. Energy and Household Expenditure Patterns. Resources for the Future, Washington, DC.

Lawrance, E. 1991. Poverty and the rate of time preference: Evidence from panel data, *J. Political Economy*, **99** (1), 54-77.

MacCracken, C., J. Edmonds, S. Kim, and R. Sands. 1999. The economics of the Kyoto Protocol, *Energy J.*, Kyoto Special Issue, vii-xiv.

Manne, A. 1999. Equity, efficiency, and discounting, Ch. 12 *in* Discounting and Intergenerational Equity, P. Portney and J. Weyant (eds.), Washington D.C.: Resources for the Future.

Manne, A., R. Mendelsohn, and R. Richels. 1995. MERGE: A model for evaluating regional and global effects of GHG reduction policies, *Energy Policy*, **23**(1), 17-34.

Marini, G. and P. Scaramozzino. 1995. Overlapping generations and environmental control, *J. Environmental Economics and Management*, **29**(1), 64-77.

McKibbin, R. and D. Vines. 2000. Modelling reality: The need for both inter-temporal optimization and stickiness in models for policy-making, *Oxford Review of Economic Policy*, **16**(4), 106-137.

Nordhaus, W. 1991. To slow or not to slow: The economics of the greenhouse effect, *Economic J.*, **101**, 920-937.

Nordhaus, W. 1994. Managing the Global Commons: The Economics of Climate Change. Cambridge: MIT Press.

Nordhaus, W. and Z. Yang. 1996. A regional dynamic general equilibrium model of alternative climate change strategies, *American Economic Review*, **86**(4), 741-765.

O'Neill, B. and B. Chen. 2002. Demographic determinants of household energy use in the United States. In Methods of Population-Environment Analysis, A Supplement to *Population and Development Review*, **28**, 53-88.

O'Neill, B., L. MacKellar, and W. Lutz. 2001. Population and Climate Change. Cambridge, UK: Cambridge University Press.

O'Neill, B.C. Household expenditure patterns by demographic characteristics based on the 1998 U.S. Consumer Exependiture Survey data. Working Paper, in preparation.

Paulin, G.D. 2000. Expenditure patterns of older Americans, 1984-1997. *Monthly Labor Revi*ew, May 2000, 3-28.

Peck, S. and T. Teisberg. 1992. CETA: A model for carbon emissions trajectory assessment, *Energy J.*, **13**, 55-77.

Ramsey, F. 1928. A mathematical theory of saving, *Economic J.*, **38**(4), 543-559.

Samuelson, P. 1958. An exact consumption-loan model of interest with or without social contrivance of money. *J. Political Economy*, 66, 467–482.

Schelling, T. 1995. Intergenerational discounting, *Energy Policy*, **23**(4), 395-401.

Schipper, L. 1996. Lifestyles and the environment: The case of energy, *Daedalus* 125, 113-138.

Schmitt, J. 2004. Estimating household consumption expenditures in the United States using the Interview and Diary portions of the 1980, 1990, and 1997 Consumer Expenditure Surveys, DEMPATEM Working Paper No. I, available at http://www.uva-aias.net/lower.asp?id=186.

Solow, R. 1956. A contribution to the theory of economic growth, *Quarterly J. Economics* 70, 65-94.

Stephan, G., G. Muller-Furstenberger, and P. Previdoli. 1997. Overlapping generations or infinitely-lived agents, *Environmental and Resource Economics*, **10**, 27-40.

Stokey, N., and R. Lucas. 1989. Recursive Methods in Economic Dynamics, Cambridge MA: Harvard University Press.
Weber, C. and A. Perrels. 2000. Modelling lifestyle effects on energy demand and related emission, *Energy Policy*, **28**, 549-566.
Weyant, J., and J. Hill. 1999. Introduction and overview, costs of the Kyoto Protocol: a multimodel evaluation, *Energy Journal*, Kyoto Special Issue, vii-xiv.
Yogo, M. 2004. Estimating the Elasticity of Intertemporal Substitution when instruments are weak, *Review of Economics and Statistics*, **86** (3), 797-810.
Zeng, Y., J. Vaupel, and W. Zhenglian. 1998. Household projection using conventional demographic data. In W. Lutz, J. Vaupel, and D. Ahlburg (eds.), Frontiers of Population Forecasting, supplement to *Population and Development Review*, **24**, 59-87.