

Optimal Population and Exhaustible Resource Constraints

Nicholas Lawson*

Université du Québec à Montréal

Dean Spears†

Indian Statistical Institute - Delhi Centre and Princeton University

August 23, 2016

Abstract

A large literature considers the optimal size and growth rate of the human population, trading off the utility value of additional people with the costs of a larger population. In this literature, an important parameter is the social weight placed on population size; a standard result is that a planner with a larger weight on population chooses larger population levels and growth rates. We demonstrate that this result is conditionally overturned when an exhaustible resource constraint is introduced: if the discount rate is small enough, the optimal population today decreases with the welfare weight on population size. That is, a more total-utilitarian social planner could prefer a smaller population today than a more average-utilitarian social planner. We also present a numerical illustration applied to the case of climate change, where we show that under plausible real-world parameter values, our result matters for the direction and magnitude of optimal population policy.

Keywords: optimal population, climate change, social choice and welfare, exhaustible resources, population ethics and policy, utilitarianism

JEL Codes: J10, J19, I31

*lawson.nicholas@uqam.ca. Université du Québec à Montréal, Département des sciences économiques, 320 rue Ste-Catherine Est, Bureau DS-5819, Montréal, Québec, H2X 1L7, Canada.

†dspears@princeton.edu. Wallace Hall, Woodrow Wilson School, Princeton University, Princeton, NJ. +1-918-493-6406. We would like to thank Raouf Boucekkine, David de la Croix, Giorgio Fabbri, Marc Fleurbaey, and seminar participants at the University of Copenhagen and the Indian Statistical Institute - Delhi for their comments. Any errors or omissions are the responsibility of the authors.

1 Introduction

A large and active literature considers the optimal size and growth rate of the human population, trading off the utility value of additional people with the need to spread capital or other resources across a larger population. In this literature, a critically important parameter is the weight placed on population size in the social welfare function.¹ Numerous papers study the effect of changing this weight, or of assuming extreme values that correspond to total utilitarianism (sometimes called Benthamite, where welfare equals the sum of utilities across the population) or average utilitarianism (sometimes called Millian, where welfare equals the average utility across the population). All previous studies have found that a planner with a larger weight on population – a more “total-utilitarian” planner – would choose larger population levels and growth rates.

In this paper, we introduce an exhaustible resource constraint into optimal population analysis, and show that this consensus result can be conditionally overturned: if the discount rate is small enough, the optimal population today instead *decreases* with the welfare weight on population size. In the future, the effect of this weight on optimal population must eventually become positive, but the critical date can be pushed arbitrarily far into the future as the discount rate decreases; thus, optimal population may well be lower for more total-utilitarian planners for the immediate and possibly quite distant future.

This result is based on a simple intuition: more total-utilitarian planners place a higher weight on periods in which there are more people, and if parameters are such that population is growing over time, this implies that more total-utilitarian planners act as if they have a lower discount rate over the number of people. Given that an exhaustible resource constraint places a limit on population over time, planners with a larger weight on population are more willing to delay that population into the future and have fewer people today.

To highlight the basic intuition, we begin by demonstrating this result in a simple exogenous-utility model. We then switch to a standard dynasty model in the style of Becker and Barro (1988) and Barro and Becker (1989), and show both analytically and numerically

¹As we will discuss later, an alternative approach that we do not study in this paper is to consider Pareto efficiency rather than social welfare maximization; Golosov, Jones, and Tertilt (2007) present a generalization of Pareto efficiency to a case with variable population.

that similar results can be obtained in that setting. Finally, we present a calibrated model in which the resource constraint is reinterpreted as a constraint on the amount of greenhouse gases that can be emitted, once again demonstrating that raising the welfare weight on population can generate significant reductions in current optimal population if the discount rate is low enough.

The rest of the paper proceeds as follows. Section 2 provides a brief overview of relevant literatures and the optimal population problem. Section 3 then presents the simple exogenous-utility model and derives results, while section 4 does the same for the dynasty model. Section 5 examines the climate change example, and section 6 concludes the paper. An appendix presents a number of extensions and sensitivity analyses

2 Background and Literature

The literature studying the optimal population problem is an established one, dating back at least to Dasgupta (1969). It is now a substantial literature, with numerous branches; many papers perform a descriptive analysis, seeking to understand the implications of the fertility decision, as in Becker and Barro (1988) and Barro and Becker (1989). A portion of the literature focusses on Pareto efficiency as an objective rather than the maximization of a social welfare function, with a recent influential example being Golosov, Jones, and Tertilt (2007), who demonstrate how Pareto efficiency can be generalized to a setting with variable population. Meanwhile, a considerable literature studies prescriptive questions, including optimal policy research such as the recent examination of child taxes in Bohn and Stuart (2015) and tradable procreation entitlements in de la Croix and Gosseries (2009). While the early literature focussed largely on models in which population is chosen exogenously, as in Dasgupta (1969), the modelling of the reproduction process has become increasingly sophisticated, using models of dynastic altruism; early examples include Razin and Ben-Zion (1975) and Nerlove, Razin, and Sadka (1985). Becker and Barro (1988) and Barro and Becker (1989) reinvigorated a literature that continues today with examples such as Jones and Schoonbroodt (2010) and Conde-Ruiz, Giménez, and Pérez-Nievas (2010).²

²Although the latter is not a true altruism model, as parents receive utility from their number of children but not from the utility of the latter.

A parameter of critical importance in optimal population analysis is the weight placed on population size in the social welfare function: to what extent is the social planner like a *total utilitarian*, who adds utilities over persons and therefore values a larger population, or like an *average utilitarian*, who averages utilities over persons and therefore is insensitive to population size? To be concrete, consider the following discrete-time intertemporal social welfare function:

$$W = \sum_{t=0}^{\infty} \beta^t n_t^\alpha u(c_t) \quad (1)$$

where β is the discount factor, $u(c_t)$ is utility per person from consumption in period t ³, and n_t is the size of the population at time t . α is the key parameter of which we have been speaking: if $u(c_t) > 0$,⁴ a higher value of α represents a higher weight placed upon population size and, all else equal, tends to lead the planner to choose a larger population. At the extremes, $\alpha = 1$ corresponds to total (or Benthamite) utilitarianism, while $\alpha = 0$ represents average (or Millian) utilitarianism.

There is some disagreement in the economics and philosophy literatures about the appropriate way to model average utilitarianism, between the alternatives of averaging by period or generation, or averaging over the entire human history. In our context, this corresponds to debating where to put the α exponent: on n_t or on some intertemporal aggregate of population. The former has become the standard in the economics literature, particularly in dynasty models, and thus (1) is the specification we will use. However, in online appendix C, we show that the same basic result can be obtained with an alternative specification in which average utilitarianism is interpreted as the average utility over all individuals in human history.⁵

Our distinction between average and total utility may seem somewhat odd to anyone who is not familiar with the literature, as in many economic analyses there is no practical difference. If one is interested, for example, in the optimal value of some policy parameter,

³We abstract from issues of inequality within generation by assuming that utility depends on consumption per person.

⁴If utility is negative, as would be the case if utility takes the CRRA form with a risk-aversion coefficient greater than 1, then Jones and Schoonbroodt (2010) show that α must be negative in any sensible social welfare function; a more total-utilitarian planner would then correspond to a smaller (more negative) value of α .

⁵We thank Marc Fleurbaey for suggesting an analysis of such a specification.

and the total size of the population is constant across all possible policies, then evaluating the welfare impact of the policy using the average utility of individuals will lead to identical conclusions to those reached by using the total utility; they will differ only through multiplication by a constant.⁶ However, when the population size is itself a variable, special attention needs to be paid to the way population enters the social welfare function, as different approaches can lead to very different conclusions.

An alternative approach to accounting for population is suggested by Blackorby, Bossert, and Donaldson (1995): critical-level utilitarianism, in which α is implicitly set to 1 but in which some value, which we will denote as η , is subtracted from consumption utility in each period:

$$W = \sum_{t=0}^{\infty} \beta^t n_t [u(c_t) - \eta].$$

In this situation, adding a new person to the population increases welfare only if their utility is above the critical value η . This specification can be generalized to include both the η and an α exponent on n_t , and α and η have different interpretations and philosophical justifications.⁷ However, their effects on optimal population are qualitatively similar: raising the critical level η and lowering the exponent α both tend to raise the proportionate importance of a policy's impact on per-person utility relative to the size of the population. Additionally, in the standard specification with $\alpha = 1$ and in our simple setting, critical-level utilitarianism leads to knife's-edge results, in which extreme values of population are optimal unless the critical values in particular periods satisfy a precise condition. It is implausible that an actual policy-maker could know the critical value with such precision. For both of these reasons, we do not focus on critical-level utilitarianism in this paper.

Numerous papers study the optimal population size or growth rate as a function of the weight placed on population in the social welfare function. The consensus is that a more total-utilitarian planner – or a planner with a higher α – will choose a larger population and/or faster growth rate. This result has been found by Nerlove, Razin, and Sadka (1982), Nerlove, Razin, and Sadka (1985), Nerlove, Razin, and Sadka (1986), and Razin and Yuen

⁶Parfit (1984) notes that almost any important policy change will change the *set* of persons who are born, especially in future generations, but to a social planner that values utilities anonymously, for many policy changes the effect on the *number* of people may be small.

⁷We thank Giorgio Fabbri for pointing this out to us.

(1995). However, interest in the problem was not extinguished by this finding. Palivos and Yip (1993) claim to have found the opposite result: with CRRA utility and a coefficient of relative risk-aversion greater than one, they find that a larger exponent on population in the social welfare function leads to a smaller optimal population size; Boucekkine and Fabbri (2013) find a similar result. However, this does not imply that a more total-utilitarian planner would prefer a smaller population. As noted by Boucekkine and Fabbri (2013) and Boucekkine, Fabbri, and Gozzi (2014), a risk-aversion coefficient greater than one implies negative utility, which – if the utility of non-existence is assumed to be zero – implies that even the best possible life is worse in utility terms than not existing; if the exponent on population size is positive, the existence of persons is assumed to make society directly worse off. Jones and Schoonbroodt (2010) demonstrate that, in this case, an economically sensible social welfare function must feature a *negative* α , and it is easy to show that a higher weight on population then corresponds to a *smaller* (more negative) α . Thus, the finding of Palivos and Yip (1993) is useful because it confirms that this case is in fact entirely consistent with the standard result: a Benthamite planner, who places a higher weight on population, would choose a larger population.

A focus on optimal population in the presence of exhaustible resource constraints, however, has been relatively limited. Dasgupta and Mitra (1982) and Mitra (1983) both incorporate such a constraint into their analysis, but they do not consider the total-vs-average margin, focussing only on total utilitarianism, while Asheim, Buchholz, Hartwick, Mitra, and Withagen (2007) consider only classical utilitarianism and maximin. Surveys of the effects of population growth on the environmental resource base are provided by Dasgupta (1995) and Robinson and Srinivasan (1997). Our analysis, therefore, fills an important gap in considering the interaction between the social welfare function and empirically relevant exhaustible resource constraints: we show that exhaustible resources can fundamentally alter the optimal population problem, and that it may be that population control should actually be favoured by those who have the *strongest* preference for a large population.

The final section of our paper also makes a contribution to the literature on the effect of population on climate change. While most of the focus on potential policy responses

to climate change has been on instruments targeting emissions,⁸ population can also be an important margin: Harford (1997) and Harford (1998) demonstrate that, in a context with variable population, an emissions tax is not sufficient to achieve the optimal allocation – a tax on children is also required to internalize the “population externality.”⁹ Therefore, reduction of population could, in principle, be an important component of the policy response to climate change. Indeed, one contribution of this paper is to bring a more sophisticated economic welfare analysis to a point actively debated in the contemporary scientific literature: see Spears (2015) in response to a pessimistic article by Bradshaw and Brook (2014).¹⁰

Bohn and Stuart (2015) provide an important contribution by studying the optimal population path in this setting; however, they simplify the problem by ignoring the intertemporal nature of the resource constraint, assuming a fixed amount of allowable emissions every period, and they do not focus on the average-vs-total utility margin that is central to our contribution. Kelly and Kolstad (2001) also study the role of population in climate change, arguing that the latter is a problem largely “caused” by population and productivity growth.¹¹ Our paper represents the first attempt to consider the role of total-vs-average utilitarianism for optimal population in the context of climate change, as well as more generally in the context of exhaustible resources.

⁸Recent surveys on the economics of climate change are provided by Pindyck (2013) and Stern (2013).

⁹A standard result in environmental economics is that an instrument should be targeted as directly as possible at the relevant margin, which suggests that an emissions tax would be the optimal policy. However, having an additional child imposes externalities upon other people, because that child will create additional emissions that impose costs on other people – and an emissions tax is powerless on this margin as long as the proceeds of the tax are in some way redistributed to society. In the simplest case, if each new child pays an expected amount of emissions taxes but receives an equal amount of lump-sum reimbursements, the tax on creating new emissions through fertility is zero. Thus, Harford (1997) finds that, whatever the level of the emissions tax, the externality generated by adding a person is the amount of the resources they and all of their descendants use up, including their greenhouse gas emissions.

¹⁰In the non-economic literature, Bradshaw and Brook (2014) argue that a reduction in the human population is not a quick fix for environmental problems, because reducing population would take too long; however, they do not compute or report any effects of population size on climate outcomes or social welfare. They appear to be considering massive reductions in population, as they speculate (but do not compute) that the optimal human population may be about 1 to 2 billion. Spears (2015) points out that these findings do not rule out that feasible reductions in population size or growth could importantly affect the risk of climate catastrophe. Neither of these papers compute an optimal population under any well-defined social welfare function, or even explicitly consider the social welfare implications of population size.

¹¹They evaluate optimal policies for both average and total utilitarianism, and find that total utilitarianism leads to higher taxes on both children and emissions, and more emissions abatement, but they do not study the impact on the optimal population size.

3 Simple Model

In our analysis throughout the paper, we focus on a setting in which each generation is active for only one period, as in Golosov, Jones, and Tertilt (2007) and Renström and Spataro (2011); thus, generations do not overlap, meaning that we abstract from intergenerational transfers and support of the elderly as a motivation for population growth. As in Bohn and Stuart (2015), we also abstract from capital (aside from that embodied in exhaustible resources), to abstract from capital-dilution as a cost of population growth.¹² In this way, we keep our models closely focussed on our central tradeoff: the direct positive effect on social welfare of adding more people versus the cost of using up exhaustible resources. A similar basic result would be obtained in a more complicated model, but the intuition would be less clear.

We begin our analysis with a simple toy model to build intuition, in which utility per person in each period is exogenous, and population is chosen freely by the social planner; endogenous and costly fertility is introduced in section 4. Each person necessarily uses up some amount of the exhaustible resource base, and we assume for now that this amount is exogenous but decreasing over time as the technology to turn resources into consumption improves.¹³ There is only one policy variable per period: how many people do we want, knowing that the resource constraint implies a constraint on the total number of people over time (discounted according to technological improvement)? We start with a general setting, and then study special cases of two-period and infinite-horizon versions of the model; algebra and proofs are collected in online appendix B.

As noted in the previous section, we model utilitarianism as a function of population in each period; thus, average utilitarianism ($\alpha = 0$) corresponds to a generational average. This is standard in the recent economics literature, as in Becker and Barro (1988), Barro and Becker (1989), Palivos and Yip (1993), Razin and Yuen (1995), Boucekkine and Fabbri (2013), and Bohn and Stuart (2015). Some researchers argue on philosophical grounds that

¹²However, we do present an extension of our dynasty model that includes physical capital in appendix A.2, to show that similar results apply in that case. We also abstract from the quantity-quality tradeoff examined by papers such as Baudin (2011).

¹³An extension in which the resource is renewable is presented in appendix A.3, whereas one in which technology is endogenous to the size of the population can be found in appendix A.4.

it is more appropriate to consider the overall average across time, as in Nerlove, Razin, and Sadka (1982), Nerlove, Razin, and Sadka (1985), and Nerlove, Razin, and Sadka (1986), and we show in online appendix C that we obtain the same basic results when we use that specification. The important point that we emphasize is that a planner that places greater weight on per-person utilities obtained in periods with more people will tend to act as if they have a lower discount rate when those periods take place in the future.

3.1 Analysis of Model

We assume that time consists of T periods, which is allowed to be either finite or $T \rightarrow \infty$. In each period there is a generation of size n_t , where each generation lives for only one period. Any person alive in a period requires a fixed amount of consumption c_t and receives an exogenous utility $u_t > 0$,¹⁴ but their consumption must be produced from an exhaustible resource x . Each unit of x used up in period t produces T_t units of consumption (where T stands for technology), and we assume that required consumption is set to unity, so that the amount of resources required per person is $x_t = \frac{1}{T_t}$. Meanwhile, productivity T at using the exhaustible resource increases over time at a rate γ : $T_{t+1} = \gamma T_t \forall t$. The resource is limited to a quantity of X , and so the resource constraint is $\sum_{t=1}^T \frac{n_t}{T_t} \leq X$.

In this simple model, we abstract from fertility decisions, and assume that a social planner chooses the size of each generation subject to the resource constraint. Here and throughout the paper, the planner weights population according to n_t^α , where $\alpha \in (0, 1)$; $\alpha = 1$ corresponds to the standard Benthamite total-utility social welfare function, and $\alpha = 0$ to the Millian average-utility formulation, but we consider strictly interior values of α to avoid degenerate solutions. The planner also has a discount factor of $\beta \leq 1$, and thus a social welfare function as in (1), but starting from period 1; therefore, the planner's problem can be written as:

$$\mathcal{L} = \sum_{t=1}^T \beta^{t-1} n_t^\alpha u_t - \lambda \left(\sum_{t=1}^T \frac{n_t}{T_t} - X \right)$$

where λ , here and throughout the paper, is a lagrange multiplier on the resource constraint (which can be written as an equality as there is no reason not to use up all the resources).

¹⁴If utility were negative, we would need to make α negative and regard a planner with a more negative α as more total-utilitarian.

We can then solve for the optimal allocation by differentiating with respect to any n_t and n_{t+1} .¹⁵ The algebra can be found in online appendix B, and combining the derivatives for n_t and n_{t+1} and assuming that utility u_t is constant across periods, we find:

$$\frac{n_{t+1}}{n_t} = (\beta\gamma)^{\frac{1}{1-\alpha}}. \quad (2)$$

Thus, the optimum features a balanced growth path, with a constant growth factor $g \equiv (\beta\gamma)^{\frac{1}{1-\alpha}}$. When technology is improving faster, or when discounting is smaller, the optimal population growth rate will tend to be larger. The most interesting effect for our purposes is that of α , the degree of total-utilitarian-ness, and a simple derivative in online appendix B shows that $\frac{dg}{d\alpha} > 0$ if and only if $\beta\gamma > 1$. That is, as long as optimal population growth is positive, a more total-utilitarian planner would want even higher population growth.

This is a common finding, but what is new is the resource constraint, which implies that it is not possible to simply increase population in all periods; to achieve higher population growth, the initial population must be smaller. The implications of this for the optimal population path can be shown most simply in two versions of the model: one with two periods, and one with an infinite horizon.

3.1.1 Results with Two Periods

Consider first the simplest possible version of the model, with only two periods: $t \in \{1, 2\}$. We know that $n_2 = gn_1$, and combining this with the resource constraint, we find:

$$n_1 = \frac{T_2 X}{g + \gamma}$$

$$n_2 = \frac{gT_2 X}{g + \gamma}.$$

These expressions depend on α only through g , leading to the following simple result:

Proposition 1. *n_2 is increasing and n_1 is decreasing with α if and only if $\beta\gamma > 1$.*

¹⁵It is simple to confirm that the second-order conditions are satisfied; they are also satisfied in the dynasty and climate change models, and in all models in the appendices with two exceptions. As discussed in online appendix C and appendix A.2, there is the possibility of non-concavity in the case of our alternative specification for average utilitarianism, and in the model with capital. We simply assume the existence of an interior optimum in the former case, which is supported by simulation results with reasonable parameters, and we confirm numerically that the second-order condition is satisfied in our simulations for the latter case.

Proof. See online appendix B. □

Thus, the optimal population today decreases with α , and the population tomorrow increases, whenever the optimal population growth is positive. When facing a resource constraint with technological improvement, total utilitarians prefer to have fewer people today so as to defer population to the future, when it will be less costly in terms of resource use and thus a larger population can be supported.¹⁶ In a sense, having a higher α is like having a lower discount rate over the number of people: as long as population is increasing over time, the future is valued more highly when α is larger and the planner wants more people to be alive in that future.

3.1.2 Results with Infinite Horizon

In a two-period model, the result is quite simple: if population goes up in one period, it must necessarily go down in the other period. A somewhat stronger result applies in an infinite horizon setting, where $t \in \{1, 2, \dots\}$. We know that $n_{t+1} = gn_t$, and so we can recursively substitute to find $n_t = (\beta\gamma)^{\frac{t-1}{1-\alpha}} n_1$. In online appendix B, we show that this can be combined with the resource constraint to give:

$$n_t = (\beta\gamma)^{\frac{t-1}{1-\alpha}} \left[1 - (\beta\gamma^\alpha)^{\frac{1}{1-\alpha}} \right] XT_1.$$

as long as $\beta\gamma^\alpha < 1$; otherwise, if β is sufficiently close to one, then optimal n_t is infinitesimally small as the planner is willing to delay population indefinitely.

Further algebra in online appendix B leads to the following proposition:

Proposition 2. *If $\beta\gamma > 1$ (so that optimal population growth is positive) and $\beta\gamma^\alpha < 1$ (so that population is not delayed indefinitely), there is a critical time t^* such that $\frac{dn_t}{d\alpha} < 0$ if $t < t^*$ and $\frac{dn_t}{d\alpha} > 0$ if $t > t^*$. The critical value is given by:*

$$t^* = \frac{1}{1 - (\beta\gamma^\alpha)^{\frac{1}{1-\alpha}}}.$$

Proof. See online appendix B. □

¹⁶In a somewhat different two-period setting where consumption is drawn from a fixed quantity of capital, Nerlove, Razin, and Sadka (1985) and Nerlove, Razin, and Sadka (1986) show that a more total-utilitarian planner prefers higher population growth, but in that setting the starting population is fixed, and a larger future population is made possible by reducing consumption per capita.

Therefore, as long as discounting is sufficiently small so that population growth is positive,¹⁷ but not so small that the planner prefers to delay population growth indefinitely, the starting optimal population n_1 decreases with α . Meanwhile, $\frac{dn_t}{d\alpha}$ is increasing with t , so that raising α has an increasingly positive effect on optimal population at later dates, with a single crossing point at t^* . However, the point in time at which a more total-utilitarian planner would prefer a larger population could be quite far into the future. Indeed, as $\beta\gamma^\alpha$ approaches one – which requires a small but positive discount rate – the critical value t^* recedes towards infinity; it is possible that for hundreds or thousands of generations, a planner that places a higher weight on population would choose a smaller population.¹⁸

A numerical illustration is provided in Figure 1, where we display four different cases in which $X = 100$, $\gamma = 1.04$, $T_1 = 1$, and $\alpha = \{0.4, 0.6\}$. In this setting, the lower bound on β such that optimal population growth is positive is 0.9615, whereas the upper bound for finite population is 0.9767 with $\alpha = 0.6$ (and higher for $\alpha = 0.4$), so we show four intermediate cases. We can see that in every case, the starting value of n_t is lower for $\alpha = 0.6$, but the growth rate is higher; most interestingly, the crossing point moves later and later as β increases, eventually occurring beyond the 100 periods presented in the figure.


This finding – that being more total-utilitarian can lead one to choose a lower population now and perhaps for a considerable amount of time in the future – is unique in the literature to the best of our knowledge. Given the existence of resource constraints in the real world, it is a finding of considerable empirical and policy relevance. However, the analysis so far has been based on a very simple model, with exogenous utility and with a complete abstraction from fertility. In the next section, we will extend our analysis to a standard dynasty model as in Becker and Barro (1988).

¹⁷If $\beta\gamma < 1$, the optimal population declines asymptotically to zero over time. In this case, a higher α leads to a higher starting population and a faster decline; the effect on current population is reversed, but the intuition is the same, as a more total-utilitarian planner wants to increase population more in the periods when it is already large.


¹⁸Note that faster technological improvement γ raises the optimal growth rate and lowers n_1 ; if people in the future will use less resources, the optimal population path allocates more resources to later periods when especially large populations can be supported. This effect is increasing in α : a more total-utilitarian planner reacts by lowering initial population and raising the growth rate by a greater amount. Therefore, $\frac{dt^*}{d\gamma} > 0$: the crossing point of population paths moves later when technology improves faster. This result also holds numerically in our analysis of the dynasty and climate change models in the following sections. We thank David de la Croix for the suggestion of analyzing the comparative statics with respect to γ .

Figure 1: Optimal Population Paths in Simple Infinite Horizon Model


(a) $\beta = 0.963$




(b) $\beta = 0.967$



(c) $\beta = 0.971$



(d) $\beta = 0.975$



Notes: All panels present the first 100 periods of the optimal population paths for $\alpha = \{0.4, 0.6\}$, from simulations in which $\gamma = 1.04$, $X = 100$, $T_1 = 1$, and $\beta = \{0.963, 0.967, 0.971, 0.975\}$.

4 Dynasty Model with Endogenous Fertility

In this section, we consider a standard dynasty model.¹⁹ Assume that any individual in generation t receives utility from their own consumption, from the number of children they have, and from the welfare of those children. Using the standard functional form as in Jones

¹⁹We continue to abstract from capital – except in the extension in appendix A.2 – and to assume that generations do not overlap as in Golosov, Jones, and Tertilt (2007). Note however that our results hold with overlapping generations if only one generation is working at a time, or if production by each generation uses a unique technology.

and Schoonbroodt (2010) and Bohn and Stuart (2015):

$$W_t = u(c_t) + \beta q(f_t) W_{t+1}$$

where c_t is consumption, $u'(c) > 0$ and $u''(c) \leq 0$, f_t is the fertility rate, $q(f_t)$ is a weight which depends on fertility, and W_{t+1} is the total utility of an individual in the next generation. Notice that recursive substitution implies that this can be written as a discounted stream of consumption utilities, and assuming that $q(f_t) = f_t^\alpha$ as in Bohn and Stuart (2015), it can be written as:

$$W_0 = \sum_{t=0}^{\infty} \beta^t n_t^\alpha u(c_t) \quad (3)$$

where $n_t \equiv \prod_{s=0}^{t-1} f_s$, and where we will normalize the starting population n_0 to one.

The standard practice in a dynasty model is to assume that the planner acts from the perspective of the current generation; that is, (3) is in fact the social welfare function, as in the simple model. However, there are two differences from the simple model: first, fertility is now a choice variable of each individual, and thus must come at some cost. Second, there is a fixed starting population n_0 , so n_1 cannot be chosen freely and without cost – having many people at time 1 requires many costly births at time 0, whereas reducing the time-1 population reduces the altruistic portion of welfare for the time-0 planner.

We assume that each child (or unit of child, as non-integer values are allowed to make the problem tractable) costs κ units of the parent's income.²⁰ Output takes the same form as before: each individual receives an income of 1 and uses up $\frac{1}{T_t}$ units of the resource, where T_t increases exponentially at a rate γ . A more complex version of the model in which per-person resource use in each period is also a choice variable can be found in appendix A.1, where we show that it leads to the same first-order condition for optimal population (except that the choice of resource use now explicitly enters the expression) and similar simulation results. Therefore, consumption is $c_t = 1 - \kappa f_t$, and the planner's problem is defined by:

$$\mathcal{L} = \sum_{t=0}^{\infty} \beta^t n_t^\alpha u \left(1 - \kappa \frac{n_{t+1}}{n_t} \right) - \lambda \left(\sum_{t=0}^{\infty} \frac{n_t}{T_t} - X \right)$$

²⁰In the current specification, since income is constant, this could instead be interpreted as a time cost of κ percent per child; replacement fertility reduces net income to $(1 - \kappa)$ percent of the maximum value. The distinction between time and monetary costs of fertility is beyond the scope of our analysis, but in our analysis of the variable resource use model in appendix A.1, we have confirmed that an altered specification of the model in which fertility costs are proportional to output does not qualitatively change the results from simulation. We thank David de la Croix for mentioning this issue.

where fertility f_t is replaced by $\frac{n_{t+1}}{n_t}$. The planner chooses n_t for all $t \geq 1$, and we take derivatives with respect to n_t and n_{t+1} and combine; online appendix D shows that the result is:

$$\begin{aligned} & \beta\alpha n_t^{\alpha-1}u(c_t) + \beta n_t^{\alpha-2}u'(c_t)\kappa n_{t+1} - n_{t-1}^{\alpha-1}u'(c_{t-1})\kappa \\ & = \beta\gamma \left(\beta\alpha n_{t+1}^{\alpha-1}u(c_{t+1}) + \beta n_{t+1}^{\alpha-2}u'(c_{t+1})\kappa n_{t+2} - n_t^{\alpha-1}u'(c_t)\kappa \right). \end{aligned} \quad (4)$$

Our goal is to use (4) to derive a similar result to that in section 3. For this purpose, we begin by considering the possibility of a balanced growth path, along which $n_{t+1} = gn_t$ for some constant g . Given a fixed starting value of $n_0 = 1$, there is only one possible constant growth rate g^* that would exactly satisfy the resource constraint:

$$\begin{aligned} \sum_{t=0}^{\infty} \frac{n_t}{T_t} = X & \rightarrow \sum_{t=0}^{\infty} \left(\frac{g}{\gamma} \right)^t = XT_0 \\ g^* & = \left(\frac{XT_0 - 1}{XT_0} \right) \gamma. \end{aligned} \quad (5)$$

Meanwhile, equation (4) implies that if a balanced-growth equilibrium exists, the population growth rate must be equal to $(\beta\gamma)^{\frac{1}{1-\alpha}}$, just as in the previous section. These two conditions can only be satisfied at one particular value of α , specifically $\alpha^* = 1 - \frac{\ln(\beta\gamma)}{\ln(g^*)}$. At any other value of α , implementing the feasible growth path g^* from (5) would make λ too small or too large to satisfy the first-order conditions.

Given the existence of the child-bearing cost κ , it is possible that having children could be sufficiently costly that desired population growth would be low and that the resource constraint would not be binding. (4) holds regardless, but to ensure that the constraint is binding, we need $\frac{\partial \mathcal{L}}{\partial n_t} > 0$ at $\lambda = 0$ on the balanced growth path, and online appendix D demonstrates that this requires:

$$\kappa < \frac{\alpha^* u(c)}{(\gamma - g^*) u'(c)} \quad (6)$$

where $c \equiv 1 - \kappa g^*$ is constant per-capita consumption. Greater intuition into condition (6) can be obtained if we assume that utility is linear in consumption, so that (6) simplifies to:

$$\kappa < \frac{\alpha^*}{\gamma - (1 - \alpha^*)g^*}. \quad (7)$$

This condition simply requires that the cost of raising children κ not be too large, and further note that it is quite a weak condition. γ is necessarily larger than $(1 - \alpha^*)g^*$ due to the

assumption that $\beta\gamma^\alpha < 1$, and in most cases g^* will be of the same order of magnitude as γ ; in fact, from (5) above, g^* approaches γ as XT_0 becomes large relative to one. If $g^* \simeq \gamma$, then (7) simplifies to $\kappa < \frac{1}{g^*}$, which is exactly the condition required for consumption $1 - \kappa g^*$ to be positive. (7) could only be significantly restrictive in a situation in which either β or XT_0 are very small, so that γ could be quite large and yet g^* could be close to 1. Additionally, online appendix D demonstrates that allowing for risk-aversion, using a CRRA utility function in (6), further weakens this condition by raising utility relative to marginal utility.

Now, suppose that α takes the value α^* defined above, the one that ensures that the balanced growth path described in (5) is optimal, and further assume that the value of κ is small enough that the resource constraint is binding. Consider the effect of a marginal increase in α . We will examine a case in which only two population values, n_t and n_{t+1} for some $t \geq 1$, are allowed to change from their balanced-growth-path values. Online appendix D contains the algebra, and we find the following condition for the effect of α on n_t and n_{t+1} :

Proposition 3. *If α is initially at the value that ensures a balanced growth path and κ satisfies (6), and only n_t and n_{t+1} are allowed to vary in response to a marginal change in α , $\frac{dn_t}{d\alpha} < 0$ and $\frac{dn_{t+1}}{d\alpha} > 0$ if and only if $\beta\gamma > 1$.*

Proof. See online appendix D. □

Therefore, starting from a resource-constrained balanced growth path featuring positive growth, an increase in α will lead to a convexification of the population path: n_t will decrease and n_{t+1} will increase. This analytical result demonstrates that the basic result from the previous section can apply in a more realistic model as well: a higher weight on population size in the social welfare function can lead to lower optimal population today, to save space for a larger population in the future.

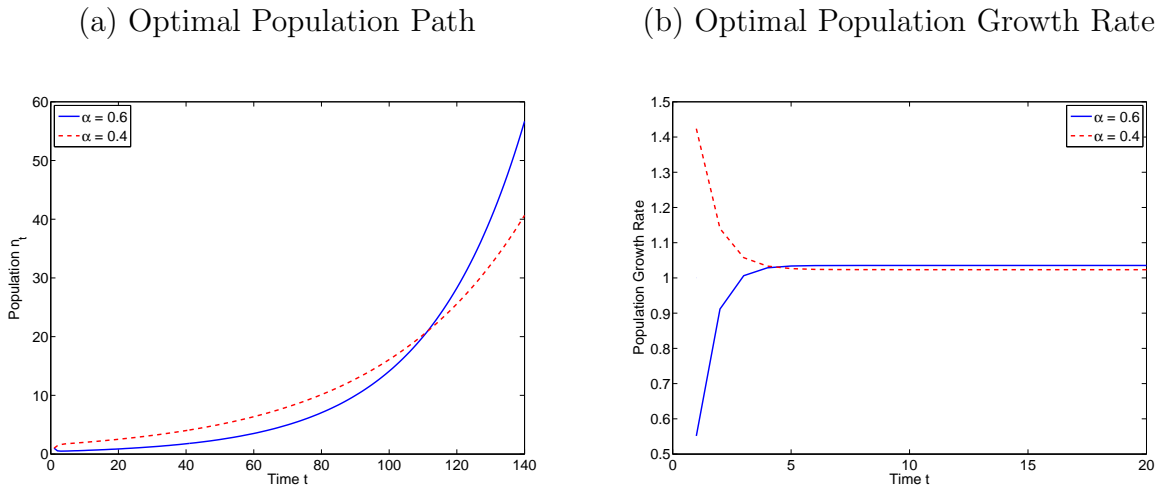
If all the n_t are allowed to vary with α , the problem becomes much more complicated; it is intuitive that the same result should typically hold, but since the overall pattern of effects could be quite complex, we are unable to extend our theoretical result to a more general case; similarly, we cannot extend the result analytically beyond the setting of a balanced growth path.²¹ Therefore, to obtain further illustration of our result and its generality, we

²¹Indeed, in our simulations, we have found numerically extreme cases in which a higher α causes optimal population to rise slightly at first, then decline for many periods, and then rise again.

now present the results of simulations of the dynasty model.

We assume values of $\beta = 0.975$ and $\gamma = 1.04$, and utility $u(c) = c$ that is linear over consumption, and we normalize the starting population to $n_0 = 1$ and set $X = 100$. Finally, we follow Bohn and Stuart (2015) in assuming that each child costs 13.8% of parental consumption, and we assume that having a child also consumes 5% of the parents' working time; then the replacement cost is $\kappa = 1 - 0.95 \times 0.862 = 0.1811$. We can then solve for the optimal population path, which is plotted for the first 140 generations in the left panel of Figure 2. As in the simple infinite horizon model, the optimal population path is initially lower when α is higher, and that state of affairs persists for over 100 generations. The right panel shows that while the long-run growth rate is increasing in α , the initial periods feature considerably higher population growth when α is smaller, since $n_0 = 1$ is smaller than optimal for $\alpha = 0.4$ and larger than optimal for $\alpha = 0.6$.

Figure 2: Optimal Population Paths in Dynasty Model



Notes: All panels present results for $\alpha = \{0.4, 0.6\}$, from simulations in which $\beta = 0.975$, $\gamma = 1.04$, $T_0 = 1$, $X = 100$, and $\kappa = 0.1811$. Panel (a) presents results from the first 140 periods, while (b) only presents the first 20 periods, as population growth rates remain flat after that point.

The combined analytical and numerical evidence in this section demonstrate that our result is robust and highly relevant: a stronger welfare weight on population size is likely to lead to a lower optimal population in the near future, and perhaps even the medium-term future, when population is chosen subject to an exhaustible resource constraint.

In appendix A, we present extensions of the models presented so far, in which we add

additional factors to test the sensitivity of our main result. We allow for choice of emissions in the dynasty model; we add physical capital to the dynasty model; we make the resource renewable in the infinite-horizon exogenous-utility model; and we allow larger populations to induce innovations that raise future technological efficiency. In each case, our result is somewhat weakened, and can even be overturned in the final two cases if we deviate sufficiently far from our baseline models, because in each case, the new feature added to the model amounts to a weakening of the exhaustibility of the resource. Appendix A thus demonstrates that the essential feature required for our result is the exhaustibility of the resource and the inability of society to find satisfactory substitutes.

5 Numerical Example with Climate Change

In this final section of the paper, we provide a simple calibrated numerical example to demonstrate the empirical relevance of our result. Climate change is perhaps the most discussed environmental issue in the world today, and with a few tweaks our model can easily fit this setting. We use our dynasty model, calibrated to the real world but in a simplified 8-period way, to provide an illustration of our principle in action.

We assume that each period lasts 25 years and represents a generation, so that the model covers the next 200 years; after year 2200, we assume the climate change problem is exogenously solved, and thus we treat the total sum of allowable emissions over the next 200 years as an exhaustible resource.²² We also assume that, immediately after the year 2200, the discount factor exogenously drops to zero, so that the planner only cares about the next 200 years (allowing us to abstract from impacts of the population path on outcomes in the distant future).²³

Individuals live for 75 years - 3 periods - but we denote them as individuals of generation t , where t is their first period of life and n_t is the size of the generation, and we assume that at

²²This is comparable to the time periods studied in standard economic climate models such as Nordhaus' DICE and RICE models, and the assumption that climate change eventually will be exogenously solved is similar to Nordhaus' "backstop" technology.

²³For example, if higher- α planners want a larger long-run population, or if technology continues to increase exponentially, expanding the population quickly after 2200 may be desirable, which would make steep population growth in the final decades of the 2100s beneficial; and yet it seems unrealistic to make our results depend on such unknowable very-long-term parameters.

the end of period t , young individuals choose the number of children that will be born in the following period, and pay the associated fertility costs. We assume that each generation is subject to its own emissions per capita and technology level, which is unchanged throughout that generation's life. Therefore, redefining the technology parameter T as the growth of output given fixed emissions, per-capita consumption of generation t can thus be written as $c_t = T_t \left(x_t - \kappa \frac{n_{t+1}}{n_t} \right)$, where x_t and T_t are the emissions per capita and technology level for generation t , and where child-raising costs also increase with T_t .

Although greenhouse gas emissions have been growing in recent decades, going forward we assume for simplicity that per-capita emissions are fixed at $x_t = 1$; thus we assume that emissions policy is successful in stopping the ongoing global increase in emissions per person, but to keep the analysis simple we assume that this value does not decline either.²⁴ Then the resource constraint is simply $\sum_{t=1}^8 n_t = X$,²⁵ and we assume that $X = 24$, implying that a constant global population of 9 billion (3 billion per generation) over the next 2 centuries would be just sufficient to avoid a disaster.²⁶

In period 8, the population size in period 9 enters into utility in the fertility cost term, so we must choose some plausible value to impose for n_9 ; on the assumption that the steady-state optimal population in years beyond 2200 is 12 billion, we assume that individuals will subsequently asymptotically approach a generation size of 4 billion, moving 50% of the way to the steady-state in each period, and thus $n_9 = \frac{n_8+4}{2}$. We assume linear utility as in the rest of our simulations, and then the planner's problem can be written as:

$$\mathcal{L} = \sum_{t=0}^8 \beta^t n_t^\alpha T_t \left(1 - \kappa \frac{n_{t+1}}{n_t} \right) - \lambda \left(\sum_{t=1}^8 n_t - X \right).$$

To calibrate the model, we use the UN's population estimates from 2000, in which a global population of 6.128 billion was divided into 2.934 billion people of age 0-24, 2.112

²⁴Of course, actual emissions are highly unequally distributed among people living in 2015. Although population-reduction policies are typically discussed in the context of poorer countries, emissions per person are much higher in richer countries. Therefore, any actual policy implications of our results may be more relevant to richer countries where emissions are greater.

²⁵Most of the empirical evidence suggests that population has a roughly proportional effect on greenhouse gas emissions; see Dietz and Rosa (1997), Cole and Neumayer (2004), and the survey in O'Neill, Liddle, Jiang, Smith, Pachauri, Dalton, and Fuchs (2012).

²⁶By ignoring uncertainty, we do not have to model the outcome in the case of catastrophic impacts as in Weitzman (2009); in that sense, our analysis is very like Bohn and Stuart (2015) in assuming a level of emissions that must not be exceeded, though here that limit is intertemporal rather than period-by-period.

billion of age 25-49, and 1.082 billion of age 50+. Therefore, the fixed $n_0 = 2.934$, and the planner needs to choose the sizes of each of the next 8 generations subject to the emissions constraint, where $n_9 = \frac{n_8+4}{2}$ as stated above.

Using values from the World Resources Institute, Bohn and Stuart (2015) calibrate their model to a 1.4% annual population growth rate, per-capita output growth of 1.7%, and aggregate greenhouse gas emissions growth of 1.8%; we use these numbers to calibrate the technological improvement parameter T_t in our model. Per-capita emissions growth is $\frac{1.018}{1.014} - 1 = 0.3945\%$, and thus the component of per-capita output growth that can be assigned to technology and not emissions growth is $\frac{1.017}{1.003945} = 1.013$. We continue to assume a constant growth rate of technology γ , which in this case is equal to $1.013^{25} = 1.381$. Finally, we set $\kappa = 0.1811$ as before.



With the model calibrated in this simple way, we can evaluate how the path of n_t changes with α . As suggested by the results earlier in the paper, the result depends strongly on the discount rate. Since we use β as the discount factor per 25-year period, let us denote $\beta \equiv \hat{\beta}^{25}$, so that $\hat{\beta}$ represents the annual discount factor. Figure 3 presents the optimal values of n_t for four different values of α , for each of $\hat{\beta} = \{0.985, 0.9875, 0.99, 0.9925\}$, while Figure 4 aggregates the three overlapping generations to produce population totals in each period (equal to $n_t + n_{t-1} + n_{t-2}$). When $\hat{\beta} = 0.985$, the planner wishes to front-load population, and thus a higher- α planner wants even more population now. However, when $\hat{\beta} = 0.9875$ or higher, we find the result from earlier in the paper: current optimal population decreases with α . This threshold in social discount rates turns out to correspond with the major divide in the literature on the economics of climate policy: Stern (2007) assumes a social discount rate below 1.5%, which our computations indicate could recommend reducing population now, and especially for more total-utilitarian planners; Nordhaus (2008) assumes a much larger social discount rate, which would have the opposite implications.

Quantitatively, changes in α can have substantial impacts on the optimal population levels; in the final $\hat{\beta} = 0.9925$ case, increasing α from 0.4 to 0.6 causes the optimal population in 2100 to drop from 6.45 billion to 4.84 billion, a drop of about 1.6 billion; we note that this is the approximate magnitude of change that Bradshaw and Brook (2014) dismiss as feasible but irrelevant. The final, thicker line in Figure 4 is the medium-fertility population

Figure 3: Optimal Generation Sizes in Climate Change Model



(a) $\hat{\beta} = 0.985$

(b) $\hat{\beta} = 0.9875$



(c) $\hat{\beta} = 0.99$

(d) $\hat{\beta} = 0.9925$



Notes: All panels present optimal generation sizes n_t for 4 different values of $\alpha = \{0.2, 0.4, 0.6, 0.8\}$, from simulations in which $n_0 = 2.934$, $n_9 = \frac{n_8+4}{2}$, $\gamma = 1.381$, $X = 24$, $\kappa = 0.1811$, and $\hat{\beta} = \{0.985, 0.9875, 0.99, 0.9925\}$.

projection by the UN over the next 100 years, and it is interesting to note that, in all but the $\hat{\beta} = 0.985$ case, the optimal population path calls for substantial reductions in fertility from the projection over the next century, followed in the higher- β cases by renewed growth later on; however, this is obviously dependent on our level of permissible emissions X .



6 Conclusion

In this paper, we have studied optimal population levels and growth rates in the presence of an exhaustible resource constraint. In this setting, we show that the standard result that a

Figure 4: Optimal Population Sizes in Climate Change Model

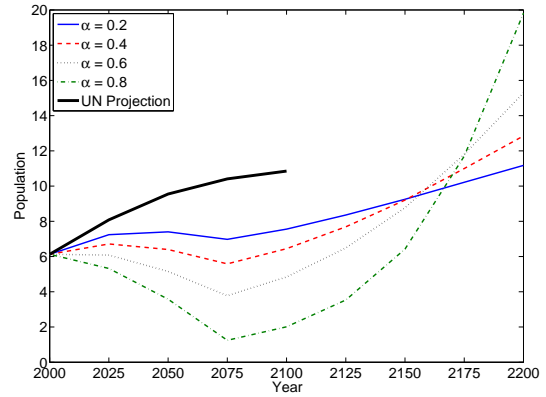

(a) $\hat{\beta} = 0.985$

(b) $\hat{\beta} = 0.9875$



(c) $\hat{\beta} = 0.99$

(d) $\hat{\beta} = 0.9925$



Notes: All panels present optimal populations $n_t + n_{t-1} + n_{t-2}$ for 4 different values of $\alpha = \{0.2, 0.4, 0.6, 0.8\}$, from simulations in which $n_0 = 2.934$, $n_9 = \frac{n_s+4}{2}$, $\gamma = 1.381$, $X = 24$, $\kappa = 0.1811$, and $\hat{\beta} = \{0.985, 0.9875, 0.99, 0.9925\}$. The final, thick black line in each figure is the UN medium-fertility projection up to 2100.

more total-utilitarian social planner will always choose a larger population no longer holds; a planner with a larger weight on population is actually likely to choose a smaller population now and perhaps well into the future. A larger weight on population, when parameters are such that population is growing over time, means that the planner acts as if they have a lower discount rate over the number of people, and thus they are more willing to delay large population sizes into the future. This result holds in both a simple exogenous-utility model and a standard dynasty model, and it can imply significant negative impacts of the population welfare weight on current population in a simple model of climate change.

A Extensions

This appendix contains four extensions of the main models in the paper, adding additional factors to the model to test the sensitivity of our main result. First, we consider a case where the planner has a choice over the optimal level of emissions; we show that the optimal population equation takes the same form, and that the simulation results are very similar to those from the dynasty model in the main text. Second, we add capital to the dynasty model with variable resource use, and show that increasing substitutability between capital and exhaustible resources in production weakens our result, by moving to an earlier date the point at which the effect of α on population becomes positive, but does not overturn the result in the cases we consider. We then examine optimal population when the resource is renewable; we find that making the resource renew faster weakens our result and eventually overturns it. Finally, we show that if larger populations induce innovations that improve technology in the future, our result can again be overturned.

In each case, the new feature added to the model amounts to a weakening of the exhaustibility of the resource, either directly or indirectly. In this way, this appendix shows the limit of our result, and demonstrates that the essential feature required is the exhaustibility of the resource and the inability of society to find satisfactory substitutes.

A.1 Variable Resource Use in Dynasty Model

In this appendix, the dynasty model from section 4 is extended to include a choice over the amount of resource used up in each period. To be precise, the planner chooses e_t , the amount of resource income obtained by each individual at time t , and the output function is linear in population size, so that per-capita income is $f(e_t)$, where $f'(e) > 0$ and $f''(e) < 0$. Thus, the amount of X used up by each generation is $\frac{n_t e_t}{T_t}$. The planner's problem is defined by:

$$\mathcal{L} = \sum_{t=0}^{\infty} \beta^t n_t^\alpha u \left(f(e_t) - \kappa \frac{n_{t+1}}{n_t} \right) - \lambda \left(\sum_{t=0}^{\infty} \frac{n_t e_t}{T_t} - X \right).$$

where we continue to assume a constant growth rate γ for T . As noted in footnote 20, our numerical results are very similar if consumption is instead specified as $f(e_t) \left(1 - \kappa \frac{n_{t+1}}{n_t} \right)$, where fertility costs are consistent with a time cost interpretation; the crossing point t^* moves slightly earlier, but the qualitative results are the same.

The planner chooses e_t and n_t for all $t \geq 1$, and we assume that both are fixed and normalized to one in period 0; differentiating with respect to e_t and n_t , and then using the derivatives for e_{t+1} and n_{t+1} to cancel out the λ , we find:

$$\begin{aligned} & \frac{e_{t+1}}{e_t} \left(\beta \alpha n_t^{\alpha-1} u(c_t) + \beta n_t^{\alpha-2} u'(c_t) \kappa n_{t+1} - n_{t-1}^{\alpha-1} u'(c_{t-1}) \kappa \right) \\ & = \beta \gamma \left(\beta \alpha n_{t+1}^{\alpha-1} u(c_{t+1}) + \beta n_{t+1}^{\alpha-2} u'(c_{t+1}) \kappa n_{t+2} - n_t^{\alpha-1} u'(c_t) \kappa \right) \\ & \quad n_t^{\alpha-1} u'(c_t) f'(e_t) = \beta \gamma n_{t+1}^{\alpha-1} u'(c_{t+1}) f'(e_{t+1}). \end{aligned}$$


Here, the problem is even more complicated than in section 4; a balanced growth path in which $n_{t+1} = g n_t$ and $e_{t+1} = e_t$ is feasible, and now for any value of α there is a constant e that satisfies the resource constraint, but there is no reason to think that any given combination of α and e would actually be optimal. Meanwhile, the analytical exercises of section 4 are complicated by the fact that the optimal path of e_t is not necessarily fixed.

Thus, we focus on simulations, using linear utility and the same parameter values as those presented in Figure 2, but with e_t chosen as well subject to a production function $f(e_t) = 2e_t - e_t^2$. The latter implies that $e_t = 1$, the implied value in section 4, is the maximum feasible value, and thus that the choice of e_t will be in $[0, 1]$. Figure 5 presents the optimal population and growth paths, exactly as in Figure 2, and the result is also qualitatively identical: a planner with a higher α - a more total-utilitarian planner - would prefer a lower population for approximately the first 100 generations. The population growth rate is once again significantly higher in the first few periods with $\alpha = 0.4$, as the population level is corrected to a higher path, and then growth settles down to a value slightly lower than that with $\alpha = 0.6$.

Figure 5: Optimal Population Paths in Dynasty Model with Variable Resource Use

(a) Optimal Population Path

(b) Optimal Population Growth Rate



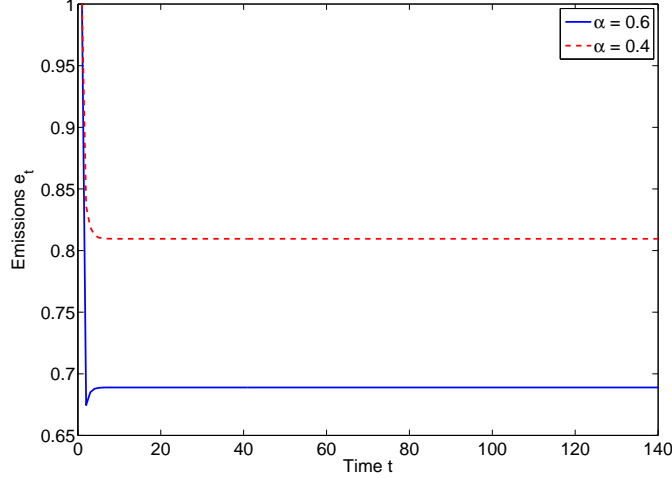
Notes: All panels present results for $\alpha = \{0.4, 0.6\}$, from simulations in which $\beta = 0.975$, $\gamma = 1.04$, $T_0 = 1$, $X = 100$, $\kappa = 0.1811$, and $f(e_t) = 2e_t - e_t^2$. Panel (a) presents results from the first 140 periods, while (b) only presents the first 20 periods, as population growth rates remain flat after that point.

Figure 6, meanwhile, presents the optimal paths for e_t , which settle down very quickly, with a high- α planner preferring a lower resource use per period. This is intuitive, as a planner who places more value on population size relative to individual utility would prefer to restrict consumption in order to make a larger population possible. However, even with this effect, the main result of the paper holds: it is quite possible, even likely, that a planner with a higher welfare weight on population size would prefer a smaller population now and long into the future.

A.2 Capital and Exhaustible Resources

We now add physical capital to the dynasty model, and allow the planner to choose both the level of resources used per period, e_t , and the level of capital investment. If m_t is the level of capital per person in period t and δ is the depreciation rate, investment per person in period t is $m_{t+1} - (1 - \delta)m_t$, which we assume can be made at a cost $w(m_{t+1} - (1 - \delta)m_t)^2$. Then, specify output per person as $f(e_t, m_t, n_t) = (2z_t(e_t, m_t) - z_t(e_t, m_t)^2)n_t^{\phi-1}$, where $z_t(e_t, m_t) = \left(ae_t^{\frac{\sigma-1}{\sigma}} + (1-a)m_t^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}$ is a CES aggregate of resource and capital with elasticity of substitution σ ; we assume diminishing marginal returns to population by assuming that $\phi < 1$ (this simply helps with ensuring convergence of the solution algorithm). Since $2z_t - z_t^2$ takes a maximum at 1, this prevents the planner from

Figure 6: Optimal Resource Income Paths in Dynasty Model



Notes: Results are presented for the first 140 periods for $\alpha = \{0.4, 0.6\}$, from simulations in which $\beta = 0.975$, $\gamma = 1.04$, $T_0 = 1$, $X = 100$, $\kappa = 0.1811$, and $f(e_t) = 2e_t - e_t^2$.

accumulating ever-increasing amounts of economic resources per person. However, planners with different tastes may prefer different divisions of z_t into e_t and m_t , and the more substitutable capital and the exhaustible resource are - that is, the larger is σ - the easier it will be to cut back on resource use and use capital instead, thus weakening the exhaustible resource constraint.

The planner's problem is summarized in the following lagrangian:

$$\mathcal{L} = \sum_{t=0}^{\infty} \beta^t n_t^\alpha u \left(f(e_t, m_t, n_t) - \kappa \frac{n_{t+1}}{n_t} - w(m_{t+1} - (1 - \delta)m_t) \right) - \lambda \left(\sum_{t=0}^{\infty} \frac{n_t e_t}{T_t} - X \right)$$

and we simplify the analysis by assuming linear utility, as in the simulations throughout the paper; then the first-order conditions are:

$$\frac{\partial \mathcal{L}}{\partial n_t} = \beta^t n_t^{\alpha-2} (\alpha n_t c_t - (1 - \delta) f(e_t, m_t, n_t) n_t + \kappa n_{t+1}) - \beta^{t-1} n_{t-1}^{\alpha-1} \kappa - \lambda \frac{e_t}{T_t} = 0$$

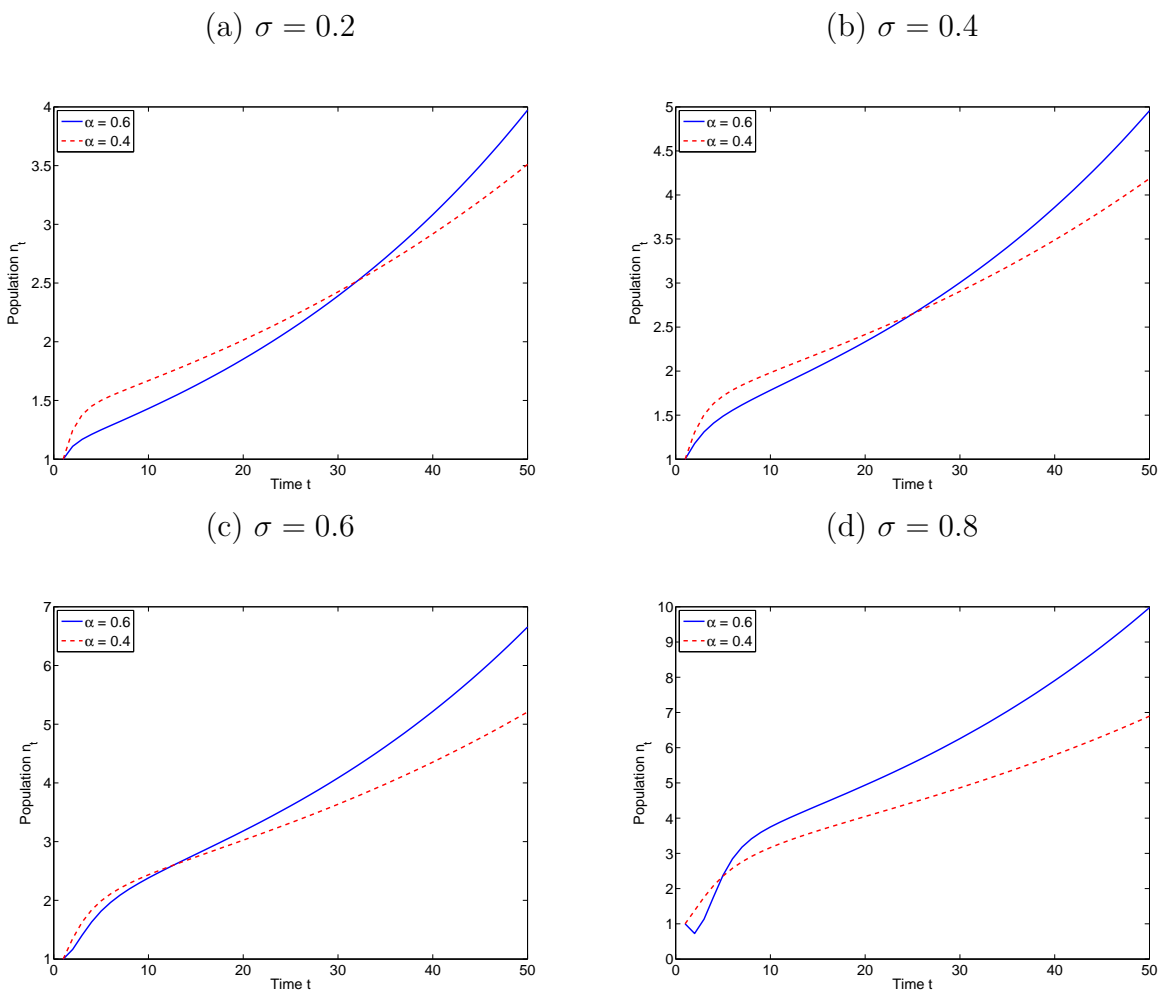
$$\frac{\partial \mathcal{L}}{\partial e_t} = \beta^t n_t^\alpha f_{e_t} - \lambda \frac{n_t}{T_t} = 0$$

$$\frac{\partial \mathcal{L}}{\partial m_t} = \beta^t n_t^\alpha (f_{m_t} + 2(1 - \delta)w(m_{t+1} - (1 - \delta)m_t)) - \beta^{t-1} n_{t-1}^\alpha 2w(m_t - (1 - \delta)m_{t-1}) = 0$$

where f_{e_t} and f_{m_t} are partial derivatives of output with respect to e_t and m_t . If $\phi < 1$, the second-order conditions are not guaranteed to be satisfied, but we have confirmed that they are satisfied in each of our numerical analyses to follow.

The first-order conditions are analytically quite complicated, but they can be simulated to illustrate their meaning. We use the standard calibration from before, except that $X = 50$ (and the simulations only run for 200 periods rather than 1000), and that we assume $a = 0.5$, $w = 1$, $\delta = 0.1$, and $\phi = 0.9$. Figure 7 presents optimal population paths for $\alpha = 0.4$ and $\alpha = 0.6$, as usual, but for four different values of σ , from 0.2 to 0.8.

Figure 7: Optimal Population Sizes in Capital & Resource Model



Notes: All panels present results for the first 50 periods for $\alpha = \{0.4, 0.6\}$, from simulations in which $\beta = 0.975$, $\gamma = 1.04$, $T_0 = 1$, $X = 50$, $\kappa = 0.1811$, $a = 0.5$, $w = 1$, $\delta = 0.1$, and $\phi = 0.9$.

We can see that, as σ increases, optimal populations generally increase, because easier substitution to capital means that the exhaustible resource constraint is less binding. More importantly for our purposes, however, the time at which the effect of α on optimal population becomes positive moves earlier. Higher- α planners have a greater incentive to switch from resource e_t to capital m_t , to loosen the exhaustible resource constraint and allow larger populations; we find that e_t drops faster and to a lower level when $\alpha = 0.6$, while m_t increases faster and to a higher level with $\alpha = 0.6$ in all but the $\sigma = 0.2$ case.

However, in the cases we study, our basic result continues to hold: the initial optimal population is lower for $\alpha = 0.6$. The reason is that, in order to build up a large capital stock, the $\alpha = 0.6$ planner must spend a lot on capital investment initially, which leaves few resources for fertility; as σ gets large, the optimal population actually drops initially to enable the planner to accumulate capital. At levels of σ above 0.8, we have had severe difficulties in finding a solution, as the planner wants to lower n_1 to a very low level - and at that level the cost of reproduction for period 2 is very high, making consumption in period 1 low and thus reducing the desire to have people in period

1. Thus, the problem appears to be unstable when substitution is sufficiently easy, but the reason is that high- α planners want to lower initial populations too much, so it is unlikely that a high- α planner will ever always want a larger initial population than a lower- α planner.

Therefore, we conclude that allowing for substitution between capital and exhaustible resources can weaken our result, in that the crossing point at which $\frac{dn_t}{d\alpha}$ becomes positive moves earlier. It is possible that sufficiently strong substitution could, under certain parameter configurations, completely overturn our result, but this does not appear to happen in the case we study.

A.3 Renewable Resource

Next, we consider a case in which the resource is not exhaustible, but rather renews at a slow rate. Denoting the remaining stock of the resource at time t as X_t , we assume that:

$$X_{t+1} = X_t - \frac{n_t}{T_t} + \eta(\bar{X} - X_t)$$

where $T_1 = 1$ and $T_t = \gamma^{t-1}$, and where we now use η to denote the rate of resource renewal, and where \bar{X} is the starting value of the resource stock. The resource renews itself in each period based on the amount existing at the beginning of the period, and grows at a faster rate when the resource is more depleted; one could easily think of alternative specifications for resource renewal, but this specification gives us a single parameter η that describes how renewable the resource is, which we can vary in the subsequent analysis.

We pursue our analysis in the infinite-horizon version of the simple exogenous-utility model. The planner's problem is most easily specified as follows:

$$\max_{n_t, X_t} W = \sum_{t=1}^{\infty} \beta^{t-1} n_t^\alpha u_t \quad \text{s.t.} \quad X_{t+1} = X_t - \frac{n_t}{T_t} + \eta(\bar{X} - X_t) \quad \& \quad \frac{n_t}{T_t} \leq X_t, \forall t.$$

Thus, we think about the planner's problem as choosing both n_t and X_t subject to the resource renewal equation and the feasibility constraint that states that no more resource can be used up in a period than exists at the start of the period. The lagrangian expression is:

$$\mathcal{L} = \sum_{t=1}^{\infty} \left[\beta^{t-1} n_t^\alpha u_t - \lambda_t \left(X_{t+1} - X_t + \frac{n_t}{T_t} - \eta(\bar{X} - X_t) \right) - \mu_t \left(\frac{n_t}{T_t} - X_t \right) \right]$$

and taking the partial derivatives while assuming $u_t = 1$ for all t :

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial n_t} &= \beta^{t-1} \alpha n_t^{\alpha-1} - \frac{\lambda_t}{T_t} - \frac{\mu_t}{T_t} = 0 \\ \frac{\partial \mathcal{L}}{\partial X_t} &= \lambda_t(1 - \eta) - \lambda_{t-1} + \mu_t = 0. \end{aligned}$$

Initially, the feasibility constraint will not be binding, unless the planner wishes to use up all of the available resource in the first period; we can check that the latter condition does not hold, and so we assume that $\mu_t = 0$ up to some time t^* at which the constraint begins to bind. This time t^* could be at infinity, or it could be a finite time. However, if t^* occurs in finite time, then once the feasibility constraint begins to bind, it must bind forever, and the resource stock must stay at the steady-state value of $\theta \equiv \frac{\eta \bar{X}}{1+\eta}$, which is used up and renewed in every subsequent period.

To understand this, consider that allowing X_t to increase from θ and converge to some higher value is clearly dominated by remaining at $X_t = \theta$: the allowable resource usage is lower in every

period in the former case, and thus population and welfare are also lower. Meanwhile, an oscillatory deviation from $X_t = \theta$ is also welfare-reducing under the standard assumption that $\beta\gamma^\alpha < 1$. To prove this, suppose that we have reached a resource stock of θ at time t , and have to decide whether or not to remain at θ in period $t + 1$ and beyond. The baseline case is that we stay at $X_{t+1} = X_{t+2} = \theta$, and so forth to infinity. Against this, consider the welfare obtained from an increase in X_{t+1} to $\theta + \epsilon$, followed by a return to θ for X_{t+2} and thereafter. In the former case, the population sizes are $n_t = \theta\gamma^{t-1}$ and $n_{t+1} = \theta\gamma^t$, and the welfare over those two periods is:

$$W_b = \beta^{t-1}(\theta\gamma^{t-1})^\alpha + \beta^t(\theta\gamma^t)^\alpha.$$

Meanwhile, in the oscillatory deviation, the population sizes are $n_t = (\theta - \epsilon)\gamma^{t-1}$ and $n_{t+1} = (\theta + (1 - \eta)\epsilon)\gamma^t$, and the welfare over those two periods is:

$$W_d = \beta^{t-1}((\theta - \epsilon)\gamma^{t-1})^\alpha + \beta^t((\theta + (1 - \eta)\epsilon)\gamma^t)^\alpha.$$

Therefore, the net welfare gain from the oscillation is:

$$\Delta W \equiv W_d - W_b = (\beta\gamma^\alpha)^{t-1} [(\theta - \epsilon)^\alpha - \theta^\alpha] + (\beta\gamma^\alpha)^t [(\theta + (1 - \eta)\epsilon)^\alpha - \theta^\alpha].$$

Now, assume that $\beta\gamma^\alpha < 1$ as before; then a necessary (but not sufficient) condition for the welfare gain to be positive is:

$$(\theta - \epsilon)^\alpha + (\theta + (1 - \eta)\epsilon)^\alpha > 2\theta^\alpha.$$

However, Jensen's inequality tells us that, since $f(x) = x^\alpha$ is an increasing and concave function:

$$(\theta - \epsilon)^\alpha + (\theta + (1 - \eta)\epsilon)^\alpha < (\theta - \epsilon)^\alpha + (\theta + \epsilon)^\alpha < 2\theta^\alpha.$$

Therefore, the oscillation must reduce welfare, and the only possible optimum is one in which X_t remains at θ after t^* if t^* is reached in finite time.

Therefore, for the time up to t^* , we can solve the first-order conditions setting $\mu_t = 0$, which leads us to the following condition:

$$\frac{n_{t+1}}{n_t} = ((1 - \eta)\beta\gamma)^{\frac{1}{1-\alpha}}.$$

If $\eta = 0$, then of course we are back in the standard exhaustible-resource world of section 3, and the population growth rate is $(\beta\gamma)^{\frac{1}{1-\alpha}}$. However, faster resource renewal reduces the population growth rate, at least until the feasibility constraint binds; this may seem counter-intuitive, but the reason is that larger populations in early periods, by depleting the resource, now lead to more resource renewal in future periods, so there is an added benefit of a larger population at the beginning. Thus, the starting population level will increase with η , while the growth rate declines, so that later population values may increase or decrease.



The effect of α on the optimal population levels and growth rate is the same as before up to t^* , except that $(1 - \eta)$ is added; if $(1 - \eta)\beta\gamma > 1$, so that population growth is positive, an increase in α raises the optimal population growth rate, and thus lowers the optimal starting population. However, this condition is less likely to be satisfied when η increases, so that an increase in resource renewability will eventually overturn our result: with sufficiently high η , population growth will be negative (though from a high value) until t^* is reached, and higher α will make that growth rate more negative, with a larger starting population.

This is illustrated in Figure 8, which presents optimal population paths for the standard parameter values, with $\eta = \{0.001, 0.004, 0.008, 0.015\}$. In each case, t^* is reached in finite time, as can be seen in the final three figures. In the first three figures, the population growth rate increases with α , at least up to the point at which the constraint begins to bind, and the initial optimal population is lower when $\alpha = 0.6$. When $\eta = 0.015$, population starts from a high level and then descends slowly, but faster and from a higher level when $\alpha = 0.6$, exactly as predicted above.

Figure 8: Optimal Population Sizes with Renewable Resource



(a) $\eta = 0.001$

(b) $\eta = 0.004$



(c) $\eta = 0.008$

(d) $\eta = 0.015$



Notes: All panels present optimal population paths for $\alpha = \{0.4, 0.6\}$, from simulations in which $\beta = 0.975$, $\gamma = 1.04$, $T_1 = 1$, $\bar{X} = 100$, and $\eta = \{0.001, 0.004, 0.008, 0.015\}$. Panels (a) through (c) present the first 100 periods, while (d) presents the first 50 periods.

A.4 Population-Induced Innovation

This final appendix considers the case in which population at time t determines the technology level at time $t + 1$ (and perhaps beyond). This captures, in a reduced-form way, the idea that a larger population could provide more opportunities or impetus for innovation, and is one more way in which the exhaustibility of our resource could be weakened.

Making T_t dependent on past population values adds considerably to the complexity of our model, so to illustrate the impact we focus on the simple 2-period exogenous-utility model. We assume that $T_1 = 1$, but that T_2 is an increasing function of n_1 : $T_2 = \gamma A n_1^\delta$, where $\delta > 0$, and where A is a constant with respect to α that we will vary with δ (as explained later) to prevent changes in δ from directly affecting the relative technological efficiencies.

We can then set up the maximization problem:

$$\mathcal{L} = n_1^\alpha u_1 + \beta n_2^\alpha u_2 - \lambda \left(n_1 + \frac{n_2}{\gamma A n_1^\delta} - X \right)$$

and if $u_1 = u_2 = 1$, the first-order conditions are:

$$\alpha n_1^{\alpha-1} = \lambda \left(1 - \frac{\delta n_2}{\gamma A n_1^{\delta+1}} \right)$$

$$\beta \alpha n_2^{\alpha-1} = \frac{\lambda}{\gamma A n_1^\delta}.$$

Substituting for λ , we have:

$$n_1^{\alpha-1} = \beta \gamma A n_1^\delta n_2^{\alpha-1} \left(1 - \frac{\delta n_2}{\gamma A n_1^{\delta+1}} \right)$$


along with the resource constraint, $n_1 + \frac{n_2}{\gamma A n_1^\delta} = X$.

Even in a two-period setting, the resulting algebra is rather complicated, so as before we use numerical simulations to explore the effect of δ on our results. In Figure 9, we present simulations for $\beta = 0.975$, $\gamma = 1.04$, and $X = 1$; for any value of α , we also set $A = (n_1^0)^{-\delta}$, where n_1^0 is the optimal n_1 when $\delta = 0$ for that value of α . In this way, changing the value of δ will affect the technological parameter only by shifting n_1 ; there will be no direct effect whereby raising δ directly raises (if $n_1 > 1$) or lowers (if $n_1 < 1$) the value of γn_1^δ .

Figure 9: Optimal Population Sizes and Growth Rate with Induced Innovation

(a) Population Sizes

(b) Population Growth Rates



Notes: Panel (a) presents the optimal population levels n_1 and n_2 , while panel (b) presents the optimal growth rate $\frac{n_2}{n_1}$, for $\alpha = \{0.4, 0.6\}$, from simulations in which $\beta = 0.975$, $\gamma = 1.04$, $X = 1$, and $\delta = [0, 0.05]$.

At low values of δ , of course, our standard result holds: population growth is positive, and a higher α raises population growth further while lowering n_1 . However, as δ increases, the same result obtains as in the previous appendix: population growth eventually falls below zero, and our result is flipped: a higher α means a larger initial population and a lower growth rate. As

with renewable resources, induced innovation raises the value of large populations at early times, as they help to make the resource constraint less binding in the future. Therefore, a more total-utilitarian planner who wants to increase population when it is already large may choose to expand the population more in earlier periods.

Compliance with ethical standards. This project received no funding. The authors declare that they have no conflict of interest. Although unrelated to this project, Dean Spears, in his capacity of Executive Director of r.i.c.e. (a 501(c)3 public charity), has received grants from the Bill and Melinda Gates Foundation, the NIH, USAID, and the IGC, and has been a paid consultant for the World Bank.

References

- ASHEIM, G. B., W. BUCHHOLZ, J. M. HARTWICK, T. MITRA, AND C. WITHAGEN (2007): “Constant Savings Rates and Quasi-Arithmetic Population Growth under Exhaustible Resource Constraints,” *Journal of Environmental Economics and Management*, 53(2), 213–229.
- BARRO, R. J., AND G. S. BECKER (1989): “Fertility Choice in a Model of Economic Growth,” *Econometrica*, 57(2), 481–501.
- BAUDIN, T. (2011): “Family Policies: What Does the Standard Endogenous Fertility Model Tell Us?,” *Journal of Public Economic Theory*, 13(4), 555–593.
- BECKER, G. S., AND R. J. BARRO (1988): “A Reformulation of the Economic Theory of Fertility,” *Quarterly Journal of Economics*, 103(1), 1–25.
- BLACKORBY, C., W. BOSSERT, AND D. DONALDSON (1995): “Intertemporal Population Ethics: Critical-Level Utilitarian Principles,” *Econometrica*, 63(6), 1303–1320.
- BOHN, H., AND C. STUART (2015): “Calculation of a Population Externality,” *American Economic Journal: Economic Policy*, 7(2), 61–87.
- BOUCEKKINE, R., AND G. FABBRI (2013): “Assessing Parfit’s *Repugnant Conclusion* Within a Canonical Endogenous Growth Set-Up,” *Journal of Population Economics*, 26(2), 751–767.
- BOUCEKKINE, R., G. FABBRI, AND F. GOZZI (2014): “Egalitarianism under Population Change: Age Structure Does Matter,” *Journal of Mathematical Economics*, 55, 86–100.
- BRADSHAW, C. J., AND B. W. BROOK (2014): “Human Population Reduction is not a Quick Fix for Environmental Problems,” *PNAS*, 111(46), 16610–16615.
- COLE, M. A., AND E. NEUMAYER (2004): “Examining the Impact of Demographic Factors on Air Pollution,” *Population & Environment*, 26(1), 5–21.
- CONDE-RUIZ, J. I., E. L. GIMÉNEZ, AND M. PÉREZ-NIEVAS (2010): “Millian Efficiency with Endogenous Fertility,” *Review of Economic Studies*, 77(1), 154–187.
- DASGUPTA, P. (1969): “On the Concept of Optimum Population,” *Review of Economic Studies*, 36(3), 295–318.

- (1995): “The Population Problem: Theory and Evidence,” *Journal of Economic Literature*, 33(4), 1879–1902.
- DASGUPTA, S., AND T. MITRA (1982): “On Some Problems in the Formulation of Optimal Population Policies when Resources are Depletable,” in *Economic Theory of Natural Resources*, ed. by W. Eichhorn, R. Henn, K. Neumann, and R. Shephard, pp. 409–429. Springer-Verlag.
- DE LA CROIX, D., AND A. GOSSERIES (2009): “Population Policy Through Tradable Procreation Entitlements,” *International Economic Review*, 50(2), 507–542.
- DIETZ, T., AND E. A. ROSA (1997): “Effects of Population and Affluence on CO_2 Emissions,” *PNAS*, 94(1), 175–179.
- GOLOSOV, M., L. E. JONES, AND M. TERTILT (2007): “Efficiency with Endogenous Population Growth,” *Econometrica*, 75(4), 1039–1071.
- HARFORD, J. D. (1997): “Stock Pollution, Child-Bearing Externalities, and the Social Discount Rate,” *Journal of Environmental Economics and Management*, 33(1), 94–105.
- (1998): “The Ultimate Externality,” *American Economic Review*, 88(1), 260–265.
- JONES, L. E., AND A. SCHOONBROODT (2010): “Complements Versus Substitutes and Trends in Fertility Choice in Dynastic Models,” *International Economic Review*, 51(3), 671–699.
- KELLY, D. L., AND C. D. KOLSTAD (2001): “Malthus and Climate Change: Betting on a Stable Population,” *Journal of Environmental Economics and Management*, 41(2), 135–161.
- MITRA, T. (1983): “Limits on Population Growth under Exhaustible Resource Constraints,” *International Economic Review*, 24(1), 155–168.
- NERLOVE, M., A. RAZIN, AND E. SADKA (1982): “Population Size and the Social Welfare Functions of Bentham and Mill,” *Economics Letters*, 10(1-2), 61–64.
- (1985): “Population Size: Individual Choice and Social Optima,” *Quarterly Journal of Economics*, 100(2), 321–334.
- (1986): “Some Welfare Theoretic Implications of Endogenous Fertility,” *International Economic Review*, 27(1), 3–31.
- NORDHAUS, W. D. (2008): *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press, New Haven and London.
- O’NEILL, B. C., B. LIDDLE, L. JIANG, K. R. SMITH, S. PACHAURI, M. DALTON, AND R. FUCHS (2012): “Demographic Change and Carbon Dioxide Emissions,” *Lancet*, 380(9837), 157–164.
- PALIVOS, T., AND C. K. YIP (1993): “Optimal Population Size and Endogenous Growth,” *Economics Letters*, 41(1), 107–110.
- PARFIT, D. (1984): *Reasons and Persons*. Oxford University Press, Oxford.

- PINDYCK, R. S. (2013): “Climate Change Policy: What do the Models Tell Us?,” *Journal of Economic Literature*, 51(3), 860–872.
- RAZIN, A., AND U. BEN-ZION (1975): “An Intergenerational Model of Population Growth,” *American Economic Review*, 65(5), 923–933.
- RAZIN, A., AND C.-W. YUEN (1995): “Utilitarian Tradeoff between Population Growth and Income Growth,” *Journal of Population Economics*, 8(1), 81–87.
- RENSTRÖM, T. I., AND L. SPATARO (2011): “The Optimum Growth Rate for Population under Critical-Level Utilitarianism,” *Journal of Population Economics*, 24(3), 1181–1201.
- ROBINSON, J. A., AND T. SRINIVASAN (1997): “Long-Term Consequences of Population Growth: Technological Change, Natural Resources, and the Environment,” in *Handbook of Population and Family Economics*, ed. by M. Rosenzweig, and O. Stark, vol. 1b, Elsevier, pp. 1175–1298.
- SPEARS, D. (2015): “Smaller Human Population in 2100 Could Importantly Reduce the Risk of Climate Catastrophe,” *PNAS*, 112(18), E2270.
- STERN, N. (2007): *The Economics of Climate Change: The Stern Review*. Cambridge University Press, Cambridge and New York.
- (2013): “The Structure of Economic Modeling of the Potential Impacts of Climate Change: Grafting Gross Underestimation of Risk onto Already Narrow Science Models,” *Journal of Economic Literature*, 51(3), 838–859.
- WEITZMAN, M. L. (2009): “On Modeling and Interpreting the Economics of Catastrophic Climate Change,” *Review of Economics and Statistics*, 91(1), 1–19.

Supplementary Appendix for: Optimal Population and Exhaustible Resource Constraints

Nicholas Lawson*

Université du Québec à Montréal

Dean Spears†

Princeton University and Indian Statistical Institute - Delhi Centre

August 23, 2016

A Introduction

In this supplementary appendix, we present additional analysis and proofs from our models. References to numbered equations, tables, etc. refer to the numbers from the main paper.

We begin with a presentation of the algebra and proofs from our simple model from section 3 of the paper in section B. Section C then demonstrates that similar results follow in a model featuring an alternative specification of average utilitarianism. Finally, section D presents algebra and proofs from the dynasty model.

B Algebra and Proofs from Simple Model

As noted in section 3, in the simple model, the lagrangian for the planner's problem can be written as follows:

$$\mathcal{L} = \sum_{t=1}^T \beta^{t-1} n_t^\alpha u_t - \lambda \left(\sum_{t=1}^T \frac{n_t}{T} - X \right)$$

*lawson.nicholas@uqam.ca.

†dspears@princeton.edu.

and differentiating with respect to n_t and n_{t+1} gives:

$$\begin{aligned}\beta^{t-1}\alpha n_t^{\alpha-1}u_t &= \frac{\lambda}{T_t} \\ \beta^t\alpha n_{t+1}^{\alpha-1}u_{t+1} &= \frac{\lambda}{T_{t+1}}.\end{aligned}$$

Combining these two expressions leads to an equation for optimal population growth:

$$\frac{n_{t+1}}{n_t} = \left(\frac{\beta\gamma u_{t+1}}{u_t} \right)^{\frac{1}{1-\alpha}}.$$

which, if we assume u_t to be constant across time, simplifies to:

$$\frac{n_{t+1}}{n_t} = (\beta\gamma)^{\frac{1}{1-\alpha}} \equiv g.$$

Differentiating g with respect to α , we find:

$$\frac{dg}{d\alpha} = \ln(\beta\gamma) (\beta\gamma)^{\frac{1}{1-\alpha}} \left(\frac{1}{(1-\alpha)^2} \right)$$

and the only term that has an uncertain sign is the first one: if $\beta\gamma > 1$, and thus $g > 1$, then $\frac{dg}{d\alpha} > 0$.

In the two-period version of the model, it is straightforward to solve for n_1 and n_2 as a function of g as in the text:

$$\begin{aligned}n_1 &= \frac{T_2 X}{g + \gamma} \\ n_2 &= \frac{g T_2 X}{g + \gamma}\end{aligned}$$

and then we can differentiate with respect to g :

$$\begin{aligned}\frac{dn_1}{dg} &= \frac{-T_2 X}{(g + \gamma)^2} < 0 \\ \frac{dn_2}{dg} &= \frac{\gamma T_2 X}{(g + \gamma)^2} > 0.\end{aligned}$$

Since n_1 and n_2 depend on α only through g , and $\frac{dg}{d\alpha} > 0$ if and only if $\beta\gamma > 1$, then $\frac{dn_1}{d\alpha} < 0$ and $\frac{dn_2}{d\alpha} > 0$ if and only if $\beta\gamma > 1$, as stated in Proposition 1.

In the infinite horizon version of the model, using the equation for population in any period $n_t = (\beta\gamma)^{\frac{t-1}{1-\alpha}} n_1$, and the resource constraint $\sum_{t=1}^{\infty} \frac{n_t}{\gamma^{t-1} T_1} = X$, we can derive the starting value n_1 :

$$n_1 \sum_{t=1}^{\infty} (\beta\gamma)^{\frac{t-1}{1-\alpha}} = X T_1.$$

For this to give a positive solution for n_1 , it must be the case that $\beta\gamma^\alpha < 1$ (otherwise, if β is sufficiently close to one, then the optimal n_1 is infinitesimally small, as the planner is willing to delay population indefinitely). Assuming this to be true, the solution for n_1 is:

$$n_1 = \left[1 - (\beta\gamma^\alpha)^{\frac{1}{1-\alpha}}\right] XT_1$$

and this can be used to give the value for any n_t :

$$n_t = (\beta\gamma)^{\frac{t-1}{1-\alpha}} \left[1 - (\beta\gamma^\alpha)^{\frac{1}{1-\alpha}}\right] XT_1.$$

Finally, we can differentiate n_t with respect to α :

$$\frac{dn_t}{d\alpha} = \frac{XT_1 \ln(\beta\gamma) (\beta\gamma)^{\frac{t-1}{1-\alpha}}}{(1-\alpha)^2} \left[(t-1) \left(1 - (\beta\gamma^\alpha)^{\frac{1}{1-\alpha}}\right) - (\beta\gamma^\alpha)^{\frac{1}{1-\alpha}} \right].$$

If we set this to zero, we find that the critical value of t is:

$$t^* = \frac{1}{1 - (\beta\gamma^\alpha)^{\frac{1}{1-\alpha}}}.$$

Notice that the first part of $\frac{dn_t}{d\alpha}$ – the fraction prior to the square brackets – is always positive given the assumption that $\beta\gamma > 1$. The part within the square brackets is clearly increasing with t given that $\beta\gamma^\alpha < 1$, and thus we have that $\frac{dn_t}{d\alpha} < 0$ for $t < t^*$ and $\frac{dn_t}{d\alpha} > 0$ if $t > t^*$, exactly as stated in Proposition 2.

C An Alternative Specification for Average Utilitarianism

In our main analysis, we treat average utilitarianism as a generational average, applying our α parameter to the size of each generation. While we describe in section 3 that this is the standard specification in the recent literature on the economics of population, it is not the universal standard: Nerlove, Razin, and Sadka (1982), Nerlove, Razin, and Sadka (1985), and Nerlove, Razin, and Sadka (1986) all find that total utilitarians prefer larger populations in a setting in which average utilitarianism evaluates the average utility across all individuals over time. Furthermore, some philosophical objections have been raised against the generational

average specification as a way of modelling average utilitarianism; for example, Dasgupta (1998) argues that such a specification is ad hoc and lacks philosophical foundations.¹

Therefore, in this appendix, we demonstrate that similar results hold when we model average utilitarianism as an average across all time, rather than across a generation; we thank Marc Fleurbaey for suggesting this analysis. We will specify our general social welfare function as:

$$W = \frac{\sum_{t=0}^{\infty} \beta^t n_t u_t}{(\sum_{t=0}^{\infty} \beta^t n_t)^{1-\alpha}}$$

where α is once again the weight placed on population; if $\alpha = 0$, the function expresses pure average utilitarianism, while if $\alpha = 1$, the term on the bottom is equal to one and the social welfare function corresponds to total utilitarianism.² We will evaluate the effect of marginal changes to α on our results, but in each case the same conclusion would hold if considering the extreme cases of average ($\alpha = 0$) and total ($\alpha = 1$) utilitarianism.

To keep the algebra simple, we will only present here the two-period exogenous utility case; exactly analogous results follow in the infinite-horizon and dynasty models when considering changes in n_t and n_{t+1} , and are available upon request. The planner's problem is:

$$\mathcal{L} = \frac{n_1 u_1 + \beta n_2 u_2}{(n_1 + \beta n_2)^{1-\alpha}} - \lambda \left(\frac{n_1}{T_1} + \frac{n_2}{T_2} - X \right).$$

However, we need to make one minor modification to the model to prevent knife's-edge solutions. If u_t is simply a constant (or if $u_t = 1 - \kappa \frac{n_{t+1}}{n_t}$ in the dynasty model), then an interior solution can only exist if $\beta\gamma = 1$, and then any interior solution is valid. Instead, we assume that output features diminishing marginal returns with respect to population size, so that total output is equal to n_t^δ for $\delta < 1$, and we assume linear utility over per-capita consumption so that $u_t = n_t^{\delta-1}$. Then the planner's problem becomes:

$$\mathcal{L} = \frac{n_1^\delta + \beta n_2^\delta}{(n_1 + \beta n_2)^{1-\alpha}} - \lambda \left(\frac{n_1}{T_1} + \frac{n_2}{T_2} - X \right).$$

¹However, Dasgupta (1998) also objects to the overall average utility specification, arguing that it is not time-consistent across generations. Thus, he argues that average utilitarianism as a philosophical concept is fundamentally flawed, an issue that we do not intend to get into; we simply assume that different planners could hold social welfare functions that differ according to the weight placed on population, whatever the philosophical issues underlying them.

²When dividing utility by population, we must use the discount rate β , as otherwise the term on the bottom is likely to be infinite; additionally, discounting on the top but not on the bottom would be inconsistent, and would imply that population should likely be made as large as possible in the near future, declining in the long run.

Differentiating with respect to n_1 and n_2 , we find:

$$\frac{\partial \mathcal{L}}{\partial n_1} = \frac{\delta n_1^{\delta-1} (n_1 + \beta n_2) - (1 - \alpha) (n_1^\delta + \beta n_2^\delta)}{(n_1 + \beta n_2)^{2-\alpha}} - \frac{\lambda}{T_1}$$

$$\frac{\partial \mathcal{L}}{\partial n_2} = \frac{\beta \delta n_2^{\delta-1} (n_1 + \beta n_2) - \beta(1 - \alpha) (n_1^\delta + \beta n_2^\delta)}{(n_1 + \beta n_2)^{2-\alpha}} - \frac{\lambda}{T_2}.$$

At the optimum, of course, both of these derivatives are equal to zero. Concavity cannot be guaranteed for all parameter values, but it holds in simulations with values similar to those used in section 3; therefore, we assume concavity of the value function with respect to n_1 and n_2 , and that thus an interior optimum exists.

However, the expressions are sufficiently complicated that we are not going to try to solve for n_1 and n_2 ; as it happens, we do not need to in order to sign the effect of α on n_1 and n_2 . Suppose that we start from some $\hat{\alpha}$; then we could solve the equations above for the optimal $n_1(\hat{\alpha})$ and $n_2(\hat{\alpha})$, and the associated $\lambda(\hat{\alpha})$. Now assume that α increases marginally to $\tilde{\alpha} = \hat{\alpha} + \varepsilon$, where ε is vanishingly small. Further assume that we hold n_1 and n_2 fixed for now, and that λ adjusts with α to hold $\frac{\partial \mathcal{L}}{\partial n_2} = 0$, so that λ takes the value $\lambda_2(\tilde{\alpha})$ that would make the planner willing to hold n_2 unchanged. This, however, will not generally be the value of λ that will make $\frac{\partial \mathcal{L}}{\partial n_1} = 0$, so we can evaluate whether the planner would like to increase or decrease n_1 . If $\frac{\partial \mathcal{L}}{\partial n_2} = 0$, then we know that:

$$\lambda_2(\tilde{\alpha}) = T_2 \frac{\beta \delta n_2^{\delta-1} (n_1 + \beta n_2) - \beta(1 - \tilde{\alpha}) (n_1^\delta + \beta n_2^\delta)}{(n_1 + \beta n_2)^{2-\tilde{\alpha}}}$$

and assuming that $T_2 = \gamma$ and $T_1 = 1$, this means:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial n_1} &= \frac{\delta n_1^{\delta-1} (n_1 + \beta n_2) - (1 - \tilde{\alpha}) (n_1^\delta + \beta n_2^\delta)}{(n_1 + \beta n_2)^{2-\tilde{\alpha}}} - \gamma \frac{\beta \delta n_2^{\delta-1} (n_1 + \beta n_2) - \beta(1 - \tilde{\alpha}) (n_1^\delta + \beta n_2^\delta)}{(n_1 + \beta n_2)^{2-\tilde{\alpha}}} \\ &= \frac{\delta (n_1 + \beta n_2) (n_1^{\delta-1} - \beta \gamma n_2^{\delta-1}) + (1 - \tilde{\alpha}) (n_1^\delta + \beta n_2^\delta) (\beta \gamma - 1)}{(n_1 + \beta n_2)^{2-\tilde{\alpha}}}. \end{aligned}$$

As stated, at $\hat{\alpha}$, we know that $\frac{\partial \mathcal{L}}{\partial n_1} = 0$ by definition, and thus the numerator must be equal to zero at $\hat{\alpha}$. Increasing α affects the bottom of the fraction, but without changing the sign, whereas it also reduces the top as long as $\beta \gamma > 1$. Therefore, if $\beta \gamma > 1$, $\frac{\partial \mathcal{L}}{\partial n_1}(\tilde{\alpha}) < 0$. If the equilibrium λ with $\tilde{\alpha}$ was such that both $\frac{\partial \mathcal{L}}{\partial n_1}$ and $\frac{\partial \mathcal{L}}{\partial n_2}$ took the same sign at $n_1(\hat{\alpha})$ and $n_2(\hat{\alpha})$, there would be no allocation consistent with constrained optimization, as concavity

would imply that both n_1 and n_2 should be moved in the same direction, which is not feasible given the resource constraint. Therefore, the optimal λ must be such that $\frac{\partial \mathcal{L}}{\partial n_1}$ and $\frac{\partial \mathcal{L}}{\partial n_2}$ take opposite signs at $n_1(\hat{\alpha})$ and $n_2(\hat{\alpha})$ given $\tilde{\alpha}$, and given that both are decreasing in λ , this is only possible at a $\lambda^* < \lambda_2(\tilde{\alpha})$ at which $\frac{\partial \mathcal{L}}{\partial n_1} < 0$ and $\frac{\partial \mathcal{L}}{\partial n_2} > 0$. Then, by concavity, it must be that n_1 decreases with α and n_2 increases.

Thus, we again find the standard result from the paper: if $\beta\gamma > 1$, then a planner with a larger value of α – a more total-utilitarian planner – will want to reduce population today to allow for a larger population tomorrow. This proves that our result holds in this alternative specification for average utilitarianism, in the two-period exogenous-utility case; as stated earlier, exactly analogous results follow from the infinite-horizon and dynasty models when we consider changes in n_t and n_{t+1} , holding all other n fixed. In this specification, an increase in α means that increases in total discounted population are weighted more strongly, as they do not count as heavily in the denominator; therefore, the planner would like to increase population in the most efficient way, and if $\beta\gamma > 1$, technology is improving faster than discounting and it is best to increase the future population at the expense of today.

D Algebra and Proofs from Dynasty Model

Solving the planner's problem with respect to n_t , we find

$$\beta^t \alpha n_t^{\alpha-1} u(c_t) + \beta^t n_t^{\alpha-2} u'(c_t) \kappa n_{t+1} - \beta^{t-1} n_{t-1}^{\alpha-1} u'(c_{t-1}) \kappa = \frac{\lambda}{T_t}.$$

Using the first-order condition for n_{t+1} to cancel out the λ , we then find:

$$\begin{aligned} & \beta \alpha n_t^{\alpha-1} u(c_t) + \beta n_t^{\alpha-2} u'(c_t) \kappa n_{t+1} - n_{t-1}^{\alpha-1} u'(c_{t-1}) \kappa \\ & = \beta \gamma (\beta \alpha n_{t+1}^{\alpha-1} u(c_{t+1}) + \beta n_{t+1}^{\alpha-2} u'(c_{t+1}) \kappa n_{t+2} - n_t^{\alpha-1} u'(c_t) \kappa). \end{aligned}$$

and it is easy to rearrange this expression to find that, if a balanced growth path is optimal, g^* must be equal to $(\beta\gamma)^{\frac{1}{1-\alpha}}$.

To ensure that the resource constraint is binding on the balanced growth path, we need $\frac{\partial \mathcal{L}}{\partial n_t} > 0$ at $\lambda = 0$, or:

$$\beta \alpha u(c) + \beta u'(c) \kappa g - g^{1-\alpha} u'(c) \kappa > 0.$$

Using the fact that $g^{1-\alpha} = \beta\gamma$, this can be simplified to:

$$\kappa < \frac{\alpha u(c)}{(\gamma - g)u'(c)}$$

which is the result in equation (6). Suppose for the moment that utility is CRRA, or $u(c) = \frac{c^{1-R}}{1-R}$ where $R < 1$ is the coefficient of relative risk-aversion; we restrict R to be less than one to ensure that utility remains positive. Then the condition becomes:

$$\kappa < \frac{\alpha c}{(\gamma - g)(1 - R)}$$

and it is clear that the right-hand side is increasing in R , so that risk-aversion makes (6) a weaker condition by raising utility relative to marginal utility.

Returning to the first-order conditions, we totally differentiate (4) with respect to α , n_t and n_{t+1} , and we find:

$$\begin{aligned} & [\beta\alpha(\alpha - 1)n_t^{\alpha-2}u(c_t) + 2\beta(\alpha - 1)n_t^{\alpha-3}u'(c_t)\kappa n_{t+1} + \beta n_t^{\alpha-4}u''(c_t)\kappa^2 n_{t+1}^2 + n_{t-1}^{\alpha-2}u''(c_{t-1})\kappa^2] dn_t \\ & - [\beta\alpha n_t^{\alpha-2}u'(c_t)\kappa + \beta n_t^{\alpha-3}u''(c_t)\kappa^2 n_{t+1} - \beta n_t^{\alpha-2}u'(c_t)\kappa] dn_{t+1} \\ & + [\beta n_t^{\alpha-1}u(c_t)(1 + \alpha \ln(n_t)) + \beta n_t^{\alpha-2}u'(c_t)\kappa n_{t+1} \ln(n_t) - n_{t-1}^{\alpha-1}u'(c_{t-1})\kappa \ln(n_{t-1})] d\alpha \\ & = \beta\gamma [\beta\alpha(\alpha - 1)n_{t+1}^{\alpha-2}u(c_{t+1}) + 2\beta(\alpha - 1)n_{t+1}^{\alpha-3}u'(c_{t+1})\kappa n_{t+2} \\ & \quad + \beta n_{t+1}^{\alpha-4}u''(c_{t+1})\kappa^2 n_{t+2}^2 + n_t^{\alpha-2}u''(c_t)\kappa^2] dn_{t+1} \\ & \quad - \beta\gamma [(\alpha - 1)n_t^{\alpha-2}u'(c_t)\kappa + n_t^{\alpha-3}u''(c_t)\kappa^2 n_{t+1}] dn_t \\ & + \beta\gamma [\beta n_{t+1}^{\alpha-1}u(c_{t+1})(1 + \alpha \ln(n_{t+1})) + \beta n_{t+1}^{\alpha-2}u'(c_{t+1})\kappa n_{t+2} \ln(n_{t+1}) - n_t^{\alpha-1}u'(c_t)\kappa \ln(n_t)] d\alpha. \end{aligned}$$

The resource constraint requires that $dn_{t+1} = -\gamma dn_t$; we also impose a balanced growth path so that $n_{t+1} = gn_t$, where $g = (\beta\gamma)^{\frac{1}{1-\alpha}}$, and $c_t = c \equiv 1 - \kappa g$. Then we rearrange to find:

$$\begin{aligned} & n_t^{\alpha-2} \left[\beta\alpha(\alpha - 1)u(c) \left(1 + \frac{\gamma}{g} \right) + 2\beta(\alpha - 1)u'(c)\kappa (2\gamma + g) + \beta u''(c)\kappa^2 (g^2 + \gamma^2 + 4\gamma g) \right] dn_t \\ & = \beta n_t^{\alpha-1} \ln(g) [\alpha u(c) + (g - \gamma)\kappa u'(c)] d\alpha. \end{aligned}$$

Each of the terms in the square brackets on the left-hand side contains either $(\alpha - 1)$ or $u''(c)$, both of which are negative, and all other terms in those brackets are positive, so it

is immediately apparent that the left-hand side is some negative number multiplied by dn_t . Therefore, $\frac{dn_t}{d\alpha} < 0$ if and only if the square bracket on the right-hand side is positive; on the assumption that $\beta\gamma > 1$ so that $g > 1$, we need:

$$\kappa > \frac{\alpha u(c)}{(g - \gamma)u'(c)}$$

which is exactly the condition for the resource constraint to be binding, as given by (6). Therefore, we have proved the result of Proposition 3: if α is initially at the balanced-growth-path value and κ satisfies (6), $\frac{dn_t}{d\alpha} < 0$ and $\frac{dn_{t+1}}{d\alpha} > 0$ if and only if $\beta\gamma > 1$.

References

- DASGUPTA, P. (1998): "Population, Consumption and Resources: Ethical Issues," *Ecological Economics*, 24(2-3), 139–152.
- NERLOVE, M., A. RAZIN, AND E. SADKA (1982): "Population Size and the Social Welfare Functions of Bentham and Mill," *Economics Letters*, 10(1-2), 61–64.
- (1985): "Population Size: Individual Choice and Social Optima," *Quarterly Journal of Economics*, 100(2), 321–334.
- (1986): "Some Welfare Theoretic Implications of Endogenous Fertility," *International Economic Review*, 27(1), 3–31.