Published in final edited form as:

J Popul Econ. 2018 January ; 31(1): 295–335. doi:10.1007/s00148-017-0665-9.

Optimal Population and Exhaustible Resource Constraints

Nicholas Lawson* and Dean Spears†

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

†Department of Economics. University of Texas at Austin. 2225 Speedway. Austin, TX, USA. 78712. +1-918-493-6406

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

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

Compliance with ethical standards

The authors declare that they have no conflict of interest.

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.

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.

We present this result in three variants of a general model of preferences over population and consumption. To highlight the basic intuition, we begin by demonstrating the result in a simple exogenous-utility model. Next, we provide a single-crossing proof in a more general setting with endogenous utility, and then we show that similar results can be obtained in a numerical simulation of a standard recursive dynasty model in the style of Becker and Barro (1988) and Barro and Becker (1989). 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 our main model, and derives analytical and numerical results for three specific versions of the model. Section 4 examines the climate change example, and section 5 concludes the paper. An appendix contains algebra and proofs, while a supplementary online appendix presents a number of extensions and sensitivity analyses.

2 Background and Literature

Dating back at least to Dasgupta (1969), there is a substantial literature on the optimal population problem, 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).² 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).

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? Theoretical advantages and disadvantages of average and total utilitarianism are both well-explored in the literature; in empirical surveys of social choice, both views receive support (Spears, 2017). 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. While this distinction between average and total utility makes no practical difference to welfare analysis if the total size of the population is constant across all possible policies,⁵ special attention needs to be paid to the way population enters the social welfare function when the population size is itself a variable, as different approaches can lead to very different conclusions.

²An important related paper in the population economics literature by Nerlove, Razin, and Sadka (1986a) studies a similar, but distinct question. They consider whether parents' procreative decisions would be inefficient in the face of two externalities: a fixed stock of land that makes a larger population less productive, and public goods that are cheaper for a larger population to support; they find that these externalities offset. An extension in online appendix B.4, which brings their paper and ours closer together, models the technological progress as a public good that, as in their model, is more easily achieved by a larger population.

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 .

⁵Parfit (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.

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.⁶ 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.⁷

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 (1986b), and Razin and Yuen (1995).⁸

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 policies that reduce the size of the near-term human population 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,⁹ 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

⁶An alternative approach to accounting for population is suggested by Blackorby, Bossert, and Donaldson (1995): critical-level utilitarianism, where adding a new person to the population increases welfare only if their utility is above the critical value η .

⁷We thank Marc Fleurbaey for suggesting an analysis of such a specification.

⁸However, 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.

⁹Recent surveys on the economics of climate change are provided by Pindyck (2013) and Stern (2013).

tax on children is also required to internalize the “population externality.”¹⁰ 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).¹¹

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.¹² 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.

3 Model of Optimal Population

We now present the main model of 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.¹³ In this way, we keep our model 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 use the standard social welfare function given in (1). While we could assume that this function represents the exogenous preferences of a social planner, we prefer the standard dynasty interpretation as in Jones and Schoonbroodt (2010) and Bohn and Stuart (2015): assume that each individual in generation t receives utility from their own consumption, and

¹⁰Having an additional child imposes externalities upon other people, because that child will create additional emissions that impose costs on other people, but an emissions tax is powerless on this margin as long as the proceeds of the tax are redistributed to society: 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.

¹¹In 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.

¹²They 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.

¹³However, we do present an extension of our recursive dynasty model that includes physical capital in online appendix B.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).

from the number and the welfare of their children. If β is the common discount factor, and if all individuals in generation t are identical, the utility of a representative individual in generation t is:

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

where c_t is consumption, $u'(c) > 0$ and $u''(c) < 0$, 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. Recursive substitution implies that this can be written as a discounted stream of consumption utilities, and we assume that $q(f_t) = f_t^\alpha$ as in Bohn and Stuart (2015). The

fertility rate f_t can be written as $\frac{n_{t+1}}{n_t}$, where n_t is the size of the population in period t . Therefore, W_t can be written as:

$$W_t = u(c_t) + \sum_{j=1}^{\infty} \beta^j \left(\frac{n_{t+j}}{n_t} \right)^\alpha u(c_{t+j}).$$

We start in period 0, where we assume that the starting population n_0 is exogenously given and can thus be normalized to one. The social planner – or the period-0 representative individual – chooses n_t from period 1 forward. The welfare function can then be written as:

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

which is exactly as written in (1).

The exponent on fertility in $q(f_t)$, α , is the welfare weight on population size, and $\alpha \in [0, 1]$, where $\alpha = 1$ corresponds to the standard Benthamite total-utility social welfare function, and $\alpha = 0$ to the Millian average-utility formulation. In most of our analysis, we will focus on derivatives with respect to α at interior values, to avoid degenerate solutions, but appendix B considers a special case in which we can compare results with $\alpha = 0$ and $\alpha = 1$; we show that similar results apply in that case. We assume that utility from consumption is always positive, $u(c_t) > 0$; if utility were negative, we would need to make α negative and regard a planner with a more negative α as more total-utilitarian (see Jones and Schoonbroodt (2010)).

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), Boucekine and Fabbri (2013), and Bohn and Stuart (2015). Some researchers argue on philosophical grounds that 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 (1986b), 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.

One way to understand the role of α is to note that both α and β (the time preference factor) influence the weight placed on the average welfare of future generations in the social welfare function. Where β governs the extent to which the social objective attends to the fact that a generation is temporally distant, α governs the extent to which it attends to the fact that a generation is large or small. Indeed, when $\alpha = 1$, constant population growth would weight future generations' average welfare exponentially, just as a constant exponential rate of time preference does.¹⁴

We will specifically model consumption in the subsections to come, but in each case we will assume that consumption must be produced from an exhaustible resource. Specifically, we assume that each individual's consumption requires one unit of production, which requires

the use of $\frac{1}{T_t}$ units of the exhaustible resource; T stands for technology, and this productivity at using the exhaustible resource increases exogenously over time at a rate γ : $T_{t+1} = \gamma T_t \forall t$. The resource is limited to a quantity of X starting in period 1, and so the resource constraint is $\sum_{t=1}^{\infty} \frac{n_t}{T_t} \leq X$.¹⁵ This exogenous technological progress is an important mechanism for our model: it is the key reason why a more total-utilitarian social planner might have an incentive to wait for a larger population later.¹⁶

We present results for three specific cases of this general model, starting with the most simple and specific model and proceeding gradually to the most general. First, in section 3.1 we consider a simple version of the model in which utility per person in each period is exogenous, and population is chosen freely by the social planner; a full analytical solution is available for this case. Section 3.2 focusses on a case in which utility depends on fertility costs, and section 3.3 allows for recursivity in utility; these increases in generality are offset by the complexity of the analysis, such that we are limited to a theoretical proof in the former case and numerical results in the latter case. Algebra and proofs are collected in appendix A.

3.1 Simple Exogenous-Utility Model

We begin with the simplest version of our general model, in which any person alive in a period receives an exogenous utility $u_t > 0$; consumption is assumed to be exogenously fixed at a value that we normalize to one.¹⁷ In this simple model, we abstract from fertility

¹⁴The role of population growth in weighting the importance of future periods under a TU SWF is discussed in the context of climate policy by (Budolfson, Dennig, Fleurbaey, Scovronick, Siebert, Spears, and Wagner, 2017).

¹⁵The supplementary online appendix analyzes the robustness of our results to the weakening of these assumptions. An extension in which resource use per person is a choice variable is presented in online appendix B.1, while online appendix B.3 presents an extension in which the resource is renewable, and online appendix B.4 contains one in which technology is endogenous to the size of the population.

¹⁶This role of technological progress is not a mere convenient modeling assumption: the interaction between environmental constraints and technological progress is central to the empirical understanding of 20th century population growth by demographers and economists. For an overview, see Lam's (2011) overview of how technological change, incentivized by prices in part as a result to population growth, shaped the use of limited natural resources by a growing population.

decisions, and assume that a social planner chooses the size of each generation subject to the resource constraint. The planner's problem can be written as:

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

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). To assist with the intuition, we also present a version of this model in appendix A.1 in which there are only two periods.

We can then solve for the optimal allocation by differentiating with respect to any n_t and n_{t+1} .¹⁸ The algebra can be found in appendix A.2, 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 appendix

A.2 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. The importance to our result of technological growth ($\gamma > 1$) is clear from this simple expression, and from the fact that time preference, β , is no greater than 1 (the social planner does not strictly prefer utility in the further future).

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. To show this, start with $n_{t+1} = gn_t$ and

recursively substitute to find $n_t = (\beta\gamma)^{\frac{t-1}{1-\alpha}} n_1$. In appendix A.2, we show that this can be combined with the resource constraint to give:

¹⁷A leading precursor to this model is presented by Cigno (1981), which gives the first characterization of a steady state with exhaustible resources and endogenous population assuming a constant saving ratio. Our present exogenous-utility framework abstracts away from savings; nevertheless, this model is an important precedent for moving away from Stiglitz's (1974) prior assumption of a constant population growth rate. Instead, Cigno investigates the consequences of making population growth a function of consumption and capita per capita. More generally, this early literature considered carefully how to model the interaction of technological progress and population and consumption growth – we, too, depend on technological progress in our model, as the source of an incentive for the social planner to permit a larger population later.

¹⁸It is simple to confirm that the second-order conditions are satisfied; they are also satisfied in the other models in the body of the paper, and in all models in the appendices with two exceptions. As discussed in online appendices B.2 and C, 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.

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

as long as $\beta\gamma^\alpha < 1$. Otherwise, if β is sufficiently close to one, the planner always prefers to delay population arbitrarily far into the future, and a degenerate solution arises in which optimal n_t is infinitesimally small at all finite times.

Further algebra in appendix A.2 leads to the following proposition:

Proposition 1—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 appendix A.2.

Therefore, as long as discounting is sufficiently small so that population growth is positive,¹⁹ 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.²⁰

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

¹⁹If $\beta\gamma < 1$, 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.

²⁰Note 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. We thank David de la Croix for the suggestion of analyzing the comparative statics with respect to γ .

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, at least over any time period of intertemporal constraint and sufficiently fast technological growth. 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 two subsections, we will extend our analysis to a case with endogenous (but non-recursive) utility, and to a standard recursive dynasty model as in Becker and Barro (1988).

3.2 Intermediate Model

The model in the previous subsection was very simple, allowing for a complete analytical solution. However, dynasty models typically feature an explicit fertility decision, in which having more children imposes some cost in terms of consumption. Therefore, in this subsection, we will introduce a version of our model in which consumption is a function of population – but we will keep the model non-recursive to simplify the analysis, and reserve the full recursive dynasty model for the next subsection.

Specifically, we assume that utility is a function of consumption, which is itself a function of population n_t , $c_t = c(n_t)$, where $c'(n_t) < 0$ and $c''(n_t) \leq 0$. Each additional person in generation t imposes a cost, in the form of lower consumption per capita for the same generation, due to fertility costs, division of capital over a larger workforce, etc. As a result, the planner's problem is now:

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

We start by deriving the first-order condition for n_t ; appendix A.3 shows that it is:

$$(\beta\gamma)^t n_t^{\alpha-1} \left[\alpha u(c(n_t)) + n_t u'(c(n_t)) c'(n_t) \right] = \frac{\lambda}{T_0}$$

and appendix A.3 also demonstrates that $\frac{dn_t}{dt} > 0$ if $\beta\gamma > 1$.

Without assuming simple functional forms for $u(c_t)$ and $c(n_t)$, we cannot solve for a closed-form solution for population over time. However, we can examine the derivative of the first-order condition with respect to α , and see if we can derive any patterns regarding the effect of α on n_t at different points in time. We are able to prove the following single-crossing

condition for $\frac{dn_t}{d\alpha}$:

Proposition 2—If $\beta\gamma > 1$ and the resource constraint is binding, 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 $t > t^*$.

Proof: See appendix A.3. $\square$

Thus, while we do not have a closed form solution for n_t as in the previous subsection, or a solution for the critical time t^* , we have the same basic results as in Proposition 1: a higher α , and thus a higher welfare weight on population size, translates into a reduction in optimal population in the near future, and an increase in optimal population in the more distant future. This applies as long as $\beta\gamma > 1$, so that population growth is positive; and it applies to any model in which consumption depends on fertility as specified above. However, we are still abstracting from recursivity of utility from fertility decisions. This further generalization of our model is considered in the following subsection, where it becomes more difficult to present algebraic proofs, necessitating a numerical approach.

3.3 Recursive Dynasty Model

In this subsection, we consider a standard recursive dynasty model.²¹ We use the welfare function from (1), but now we allow for recursivity in utility. Specifically, we assume that

consumption in period t is a function of fertility in that period, $f_t = \frac{n_{t+1}}{n_t}$: 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.²² Thus, 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.

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 online appendix B.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:

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

²¹We continue to abstract from capital – except in the extension in online appendix B.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.

²²In 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 online appendix B.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.

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; appendix A.4 shows that the result is:

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

(3)

A full analytical solution is not feasible, as a balanced-growth equilibrium is only optimal at one specific value of α ; nor is a theoretical proof of single-crossing feasible as in section 3.2. In appendix A.4, we show that a weaker proof is possible: if α is initially at the value that ensures a balanced growth path and κ is not too large (so that the resource constraint is binding), and if 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 $\beta\gamma > 1$. That is, 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 result is analogous to those from the previous cases of the model, and to obtain further illustration of our result and its generality, we 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 set $T_0 = 1$ and $X = 100$. 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 $\alpha = 0.6$.

The combined analytical and numerical evidence in this section demonstrates 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 and when technology is improving quickly. In appendix B, we show that an analogous result also applies to a comparison of $\alpha = 0$ and $\alpha = 1$, in a special case of our model in which those values do not lead to degenerate results.

In online appendix B, 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 exogenous-utility model; and we allow larger populations to induce innovations that raise future technological efficiency in a simple two-period exogenous-

utility model.²³ 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 model, because in each case, the new feature added to the model amounts to a weakening of the exhaustibility of the resource. Online appendix B 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.

4 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. Whether climate policy should include population policy – meaning, whether part of humanity’s climate mitigation efforts should include policies to reduce the size of the near-term human population – is a matter of considerable debate in the climate and environmental economics literatures (Gaffin and O’Neill, 1997; O’Neill, Dalton, Fuchs, Jiang, Pachauri, and Zigova, 2010; O’Neill, MacKellar, and Lutz, 2005). A larger future population will cause greater emissions under near-term technologies (O’Neill, Liddle, Jiang, Smith, Pachauri, Dalton, and Fuchs, 2012; Liddle, 2011), and implies that more future people will be vulnerable to climate damages. These facts may constitute a reason to promote policies that reduce the size of the near-term population, but whether such a policy is best depends on the magnitude of the achievable reduction in population size (Bradshaw and Brook, 2014) and its environmental consequences (Spears, 2015). The optimal policy also depends on the costs of reducing population, which we ignore: where Connelly (2008) emphasizes that some policies intended to reduce human fertility have historically had large social costs in practice, Abel, Barakat, Samir, and Lutz (2016) suggest that welfare-enhancing human development programs (such as female education) could also have the consequence of non-coercively promoting fertility reduction.

To this longstanding debate, our paper adds the following observation: evaluating whether a population-reduction policy is socially optimal further depends upon the social value of population size, such that a smaller near-term population could surprisingly be preferred by a social planner who values large population, if the interaction of environmental constraints and technological progress would therefore permit an even larger population in the further future.

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.²⁴ We also assume that, immediately after the year 2200, the discount factor exogenously drops to zero, so that

²³This is in the spirit of Nerlove, Razin, and Sadka’s (1986a) model of a larger population contributing to a public good more inexpensively – here, technological change.

²⁴This 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.

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).²⁵

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 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 t has a generation-specific emissions per capita x_t and technology level T_t , which is unchanged throughout that generation's life. T_t is now interpreted as a technological efficiency at producing consumption for a given level of emissions. Therefore, 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 .

This calibration of the model assumes an intertemporal emissions budget X . In other words, the human population has agreed to impose a limit of X on the intertemporal sum of carbon emissions. To link this budget to the size of the population, we need an estimate of the effect of population growth on emissions growth. Most 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). Therefore we follow this empirical literature in modeling carbon emissions in a period as proportional to the size of the population, abstracting away from any inequality across people in carbon emissions: total emissions in each period are equal to $x_t n_t$.²⁶

Under these assumptions, the emissions budget constraint X is isomorphic to an intertemporal population size constraint. We assume that global climate politics agrees upon a carbon budget that would permit a constant global population of 9 billion at a constant level of per-capita emissions x_b , or a constant size of 3 billion people per generation; we

normalize the constant x_t to one, so that the intertemporal constraint would be $\sum_{t=1}^8 n_t = X$, where we assume that $X = 24$. This constraint is imposed with certainty as just sufficient to avoid a disaster.²⁷

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

²⁵For 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.

²⁶Of course, actual emissions are highly unequally distributed among people living today. 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.

²⁷By 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.

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 billion of age 25–49, and 1.082 billion of age 50+. Therefore, we set $n_0 = 2.934$, 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 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 .

Therefore, depending on the society's valuation of population size, the constraints of climate change may offer an important practical reason to attempt to reduce the size of the near-term human population. Pursuing such a policy would require considerable care: our model has abstracted away from any utility costs of population policy itself, but some past coercive fertility policies have proven harmful (Connelly, 2008), and even the principle that one may not have a right to unlimited procreation has been called objectionable or even offensive (Conly, 2015). One approach may be to strengthen international and domestic policy commitments to human development policies that are known to increase average well-being while permitting or encouraging women to freely choose to reduce their own fertility (Abel, Barakat, Samir, and Lutz, 2016). Wheeler and Hammer (2010), for example, argue quantitatively that promoting female education in developing countries would be a climate mitigation policy strategy, due to its consequences for fertility, that would be more cost-effective than investments in low-carbon energy sources.

5 Conclusion

In this paper, we have studied optimal population levels and growth rates in the presence of an exhaustible resource constraint, and in the presence of technological change over time in the environmental costliness of population size. In this setting, we show that the standard result that a 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 a simple exogenous-utility model, an endogenous-utility model, and a standard recursive dynasty model, and it can imply significant negative impacts of the population welfare weight on current population in a simple model of climate change.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

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. We are especially grateful for the help and guidance provided by two anonymous referees of this journal and by the editor. Any errors or omissions are the responsibility of the authors.

This project received no funding. 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.

A Algebra and Proofs

This appendix presents the algebra and proofs from our main model in section 3 of the paper.

A.1 Simple Two-Period Exogenous-Utility Model

Consider first the simplest possible version of the model, with only two relevant periods: $t \in \{1, 2\}$ (the population and utility in period 0 is exogenously fixed). The planner's problem is:

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

Differentiating with respect to n_1 and n_2 , we find:

$$\alpha n_1^{\alpha-1} u_1 = \frac{\lambda}{T_1}$$

$$\beta \alpha n_2^{\alpha-1} u_2 = \frac{\lambda}{T_2}.$$

Combining these expressions, and assuming that $u_2 = u_1$, we find:

$$\frac{n_2}{n_1} = (\beta \gamma)^{\frac{1}{1-\alpha}}.$$

Defining the growth factor as $g \equiv (\beta \gamma)^{\frac{1}{1-\alpha}}$, and combining $n_2 = g n_1$ with the resource constraint:

$$n_1 = \frac{T_2 X}{g + \gamma} \quad (4)$$

$$n_2 = \frac{g T_2 X}{g + \gamma}. \quad (5)$$

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

Proposition 3

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

Proof—Differentiate (4) and (5) with respect to g :

$$\frac{dn_1}{dg} = \frac{-T_2 X}{(g+\gamma)^2} < 0$$

$$\frac{dn_2}{dg} = \frac{\gamma T_2 X}{(g+\gamma)^2} > 0.$$

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 $\beta\gamma > 1$. $\square$

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.²⁸ 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.

A.2 Simple Infinite-Horizon Exogenous-Utility Model

As noted in section 3.1, in the exogenous-utility model, the lagrangian for the planner's problem can be written as follows:

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

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

$$\beta^t \alpha n_t^{\alpha-1} u_t = \frac{\lambda}{T_t}$$

$$\beta^{t+1} \alpha n_{t+1}^{\alpha-1} u_{t+1} = \frac{\lambda}{T_{t+1}}.$$

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

²⁸In 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 (1986b) 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.

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

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

$$\frac{dg}{d\alpha} > 0.$$

Using the equation for population in any period $n_t = (\beta\gamma)^{\frac{t-1}{1-\alpha}} n_1$, and the resource constraint

$$\sum_{t=1}^{\infty} \frac{n_t}{\gamma^{t-1} T_1} = X, \text{ we can derive the starting value } n_1:$$

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

For this to give a positive solution for n_1 , it must be the case that $\beta\gamma^\alpha < 1$.²⁹ Assuming this to be true, the solution for n_1 is:

$$n_1 \left[1 - (\beta\gamma^\alpha)^{\frac{1}{1-\alpha}} \right] X T_1$$

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

$$n_t = (\beta\gamma)^{\frac{t-1}{1-\alpha}} \left[1 - (\beta\gamma^\alpha)^{\frac{1}{1-\alpha}} \right] X T_1.$$

Finally, we can differentiate n_t with respect to α :

$$\frac{dn_t}{d\alpha} = \frac{X T_1 \ln(\beta\gamma) (\beta\gamma)^{\frac{t-1}{1-\alpha}}}{(1-\alpha)^2} \left[(t-1) \left(1 - (\beta\gamma^\alpha)^{\frac{1}{1-\alpha}} \right) - (\beta\gamma^\alpha)^{\frac{1}{1-\alpha}} \right].$$

²⁹Otherwise, if β is sufficiently close to one, then the optimal n_1 is infinitesimally small, as the planner is willing to delay population indefinitely.

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 1.

A.3 Intermediate Model

We want to prove the single-crossing property in the intermediate model; that is, the

existence of 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 derivative of the planner's problem with respect to n_t is:

$$\frac{\partial \mathcal{L}}{\partial n_t} = \beta^t \alpha n_t^{\alpha-1} u(c(n_t)) + \beta^t n_t^\alpha u'(c(n_t)) c'(n_t) - \frac{\lambda}{T_t} = 0$$

which leads immediately to:

$$(\beta\gamma)^t n_t^{\alpha-1} [\alpha u(c(n_t)) + n_t u'(c(n_t)) c'(n_t)] = \frac{\lambda}{T_0}. \quad (6)$$

We first prove that $\frac{dn_t}{dt} > 0$, for which purpose we take the derivative with respect to t ; after a bit of simplification, we have:

$$\begin{aligned} & \ln(\beta\gamma) n_t [\alpha u(c(n_t)) + n_t u'(c(n_t)) c'(n_t)] \\ & + [(\alpha - 1) \alpha u(c(n_t)) + n_t (2\alpha u'(c(n_t)) c'(n_t) + n_t u''(c(n_t)) c'(n_t) c'(n_t) + n_t u'(c(n_t)) c''(n_t))] \frac{dn_t}{dt} = 0. \end{aligned}$$

All of the terms multiplying $\frac{dn_t}{dt}$ are at least weakly negative (and the first two are strictly negative), while if $\beta\gamma > 1$, the first term is positive. Therefore, the derivative must be

positive: $\frac{dn_t}{dt} > 0$ if and only if $\beta\gamma > 1$.

Next, we take the derivative of the first-order condition with regard to α ; whereas the multiplier λ is a constant with regard to t , it is not with regard to α : a change in α will

necessitate a different λ to ensure that the resource constraint is binding. Therefore, we have:

$$\begin{aligned} & (\beta\gamma)^t \left[(\alpha - 1) n_t^{\alpha-2} \left(\alpha u(c(n_t)) + n_t u'(c(n_t)) c'(n_t) \right) + n_t^{\alpha-1} \left(\alpha u'(c(n_t)) c'(n_t) + u'(c(n_t)) c'(n_t) + n_t u''(c(n_t)) c'(n_t) c' \right) \right. \\ & \quad \left. + (\beta\gamma)^t \left[\ln(n_t) n_t^{\alpha-1} \left(\alpha u(c(n_t)) + n_t u'(c(n_t)) c'(n_t) \right) + n_t^{\alpha-1} u(c(n_t)) \right] \right] \\ & = \frac{1}{T_0} \frac{d\lambda}{d\alpha}. \end{aligned}$$

Simplifying, the total derivative is:

$$\begin{aligned} & T_0(\beta\gamma)^t n_t^{\alpha-2} \left[(\alpha - 1) \alpha u(c_t) + n_t \left(2\alpha u'(c_t) c'(n_t) + n_t u''(c_t) c'(n_t)^2 + n_t u'(c_t) c''(n_t) \right) \right] \frac{dn_t}{d\alpha} \\ & \quad + \lambda \ln(n_t) \\ & \quad + T_0(\beta\gamma)^t n_t^{\alpha-1} u(c_t) = \frac{d\lambda}{d\alpha}. \end{aligned}$$

To further simplify notation, define $g(t)$ as the terms multiplying $\frac{dn_t}{d\alpha}$; as every term in $g(t)$ is at least weakly negative, and the first two are strictly negative, we can conclude that $g(t) < 0$

for all values of t . Similarly, define $f(t) \equiv \frac{d\lambda}{d\alpha} - \lambda \ln(n_t) - T_0(\beta\gamma)^t n_t^{\alpha-1} u(c_t)$, and then we have:

$$\frac{dn_t}{d\alpha} = \frac{f(t)}{g(t)}.$$

The effect of α on n_t is $f(t)$, and how it changes over time. Because $g(t)$ is always

negative, proving single-crossing of $f(t)$ would also prove single-crossing of $\frac{dn_t}{d\alpha}$. We start by defining $f(t)$ as a function of t only through n_t using (6) to substitute for $T_0(\beta\gamma)^t n_t^{\alpha-1}$:

$$\begin{aligned} f(t) &= \frac{d\lambda}{d\alpha} - \lambda \ln(n_t) \\ & \quad - \frac{\lambda u(c(n_t))}{\alpha u(c(n_t)) + n_t u'(c(n_t)) c'(n_t)} \\ &= \frac{d\lambda}{d\alpha} - \lambda \left[\ln(n_t) + \left(\alpha + \frac{n_t u'(c(n_t)) c'(n_t)}{u(c(n_t))} \right)^{-1} \right]. \end{aligned}$$

Defining the derivative with respect to t as $f'(t)$, we now have $f'(t) = \frac{\partial f(t)}{\partial n_t} \frac{dn_t}{dt}$, and we

have already shown that $\frac{dn_t}{dt} > 0$ if $\beta\gamma > 1$. Therefore, consider $\frac{\partial f(t)}{\partial n_t}$:

$$\frac{\partial f(t)}{\partial n_t} = -\lambda \left[\frac{1}{n_t} - \left(\alpha + \frac{n_t (c(n_t)) c'(n_t)}{u(c(n_t))} \right)^{-2} \times \left(\frac{(u'(c(n_t)) c'(n_t) + n_t u''(c(n_t)) c'(n_t)^2 + n_t u'(c(n_t)) c''(n_t)) u(c(n_t)) - n_t u'(c(n_t))}{u(c(n_t))^2} \right) \right]$$

Every term in the numerator of the second line is negative, which means that $\frac{\partial f(t)}{\partial n_t} < 0$. Therefore, $f'(t) < 0$ if and only if $\beta\gamma > 1$.

Next, notice that $\frac{dn_t}{d\alpha}$ cannot be either always positive or always negative, as long as the resource constraint is binding: if α increases, any increase in population at some point in time must be offset by a decrease at some other point in time. Therefore, $f(t)$ cannot be always positive or always negative; given that $f'(t) < 0$, $f(t)$ must initially be positive and then become negative beyond some critical time t^* . Finally, we know that $g(t)$ is always

negative, and therefore $\frac{dn_t}{d\alpha} = \frac{f(t)}{g(t)}$ must be negative for $t < t^*$ and positive for $t > t^*$, if and only if $\beta\gamma > 1$. This is the result presented in Proposition 2.

A.4 Recursive 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:

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

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

Our goal is to use this equation to derive a similar result to that in section 3.1 or 3.2. For this purpose, we begin by considering the possibility of a balanced growth path, along which $n_{t+1} = g n_t$. 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:

$$\sum_{t=0}^{\infty} \frac{n_t}{T_t} = X \rightarrow \sum_{t=0}^{\infty} \left(\frac{g}{\gamma} \right)^t = X T_0$$

$$g^* = \left(\frac{X T_0 - 1}{X T_0} \right) \gamma. \quad (7)$$

These two conditions can only be satisfied at one particular value of α , specifically

$\alpha^* = 1 - \frac{\ln(\beta\gamma)}{\ln(g^*)}$. This is the only α for which a balanced-growth equilibrium exists; at any other value, implementing the feasible growth path g^* from (7) 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. (3) holds regardless, but to ensure that the constraint is

binding, we need $\frac{\partial \mathcal{L}}{\partial n_t} > 0$ on the balanced growth path, 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:

$$\kappa < \frac{\alpha u(c)}{(\gamma - g) u'(c)} \quad (8)$$

where $c \equiv 1 - \kappa g$ is constant per-capita consumption. Greater intuition into condition (8) can be obtained if we assume that utility is linear in consumption, so that (8) simplifies to:

$$\kappa < \frac{\alpha}{\gamma - (1 - \alpha)g}. \quad (9)$$

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 (7) above, g^* approaches γ as XT_0 becomes large relative to one. If $g \simeq \gamma$, then

(9) simplifies to $\kappa < \frac{1}{g}$, which is exactly the condition required for consumption $1 - \kappa g$ to be positive. (9) 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, allowing for risk-aversion using a CRRA utility function further weakens this

condition by raising utility relative to marginal utility. Let $u(c) = \frac{c^{1-R}}{1-R}$ 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 (8) 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 this a weaker 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 (7) 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. We find the following condition for the effect of α on n_t and n_{t+1} :

Proposition 4

If α is initially at the value that ensures a balanced growth path and κ satisfies (8), 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—Totally differentiate (3) with respect to α , n_t and n_{t+1} :

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

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, 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. $\square$

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.³⁰

B Comparison of $\alpha = 0$ and $\alpha = 1$

In this appendix, we present a special case of our model in which we can compare results with $\alpha = 0$ and $\alpha = 1$. In each of the models presented so far, at least one of these values of α leads to degenerate results, so we now consider a setting in which utility from consumption is quadratic. Additionally, to simplify the analysis to come, which will rely on continuity of derivatives with respect to time, we focus on a continuous-time setting:

$$u(t) = 1 + \psi_1 n(t) - \psi_2 n(t)^2$$

where $\psi_1, \psi_2 > 0$.

In this case, the planner's problem is:

$$\mathcal{L} = \int_{t=0}^{\infty} e^{-rt} n(t)^\alpha \left(1 + \psi_1 n(t) - \psi_2 n(t)^2 \right) - \lambda \left(\int_{t=0}^{\infty} \frac{n(t)}{T(t)} - X \right)$$

where r is the instantaneous discount rate, and the first-order condition for $n(t)$ is:

$$e^{-rt} n(t)^{\alpha-1} \left[\alpha (1 + \psi_1 n(t) - \psi_2 n(t)^2) + n(t) (\psi_1 - 2\psi_2 n(t)) \right] = \frac{\lambda}{T(t)}.$$

Simplifying, normalizing $T(0) = 1$, and writing $T(t) = e^{pt}$ where p is the technological growth rate (to maintain consistent notation), we arrive at:

$$e^{(p-r)t} n(t)^{\alpha-1} \left[\alpha + (1+\alpha) \psi_1 n(t) - (2+\alpha) \psi_2 n(t)^2 \right] = \lambda.$$

We then differentiate with respect to t :

³⁰Indeed, 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.

$$(p-r)\lambda + e^{(p-r)t}n(t)^{\alpha-2} \left[(\alpha-1)\alpha + \alpha(1+\alpha)\psi_1 n(t) - (1+\alpha)(2+\alpha)\psi_2 n(t)^2 \right] \frac{dn(t)}{dt} = 0$$

and therefore we have the following results for $\alpha = 0$ and $\alpha = 1$:

$$\left. \frac{dn(t)}{dt} \right|_{\alpha=0} = \frac{\lambda_0 (p-r) e^{-(p-r)t}}{2\psi_2}$$

$$\left. \frac{dn(t)}{dt} \right|_{\alpha=1} = \frac{\lambda_1 (p-r) e^{-(p-r)t}}{6\psi_2 n(t) - 2\psi_1}$$

where λ_α denotes the equilibrium value of λ for a given value of α .

We want to prove that, at any value of t for which $n(t)$ is equal for $\alpha = 0$ and $\alpha = 1$,

$\left. \frac{dn(t)}{dt} \right|_{\alpha=1} > \left. \frac{dn(t)}{dt} \right|_{\alpha=0}$; if so, then given continuity, it is impossible for $n(t)|_{\alpha=1}$ to descend from above $n(t)|_{\alpha=0}$ to below $n(t)|_{\alpha=0}$ as t increases. If the resource constraint is binding in both cases – which we assume is the case – $n(t)|_{\alpha=1}$ cannot be always above or always below

$n(t)|_{\alpha=0}$. Therefore, if $\left. \frac{dn(t)}{dt} \right|_{\alpha=1} > \left. \frac{dn(t)}{dt} \right|_{\alpha=0}$ when $n(t)|_{\alpha=1} = n(t)|_{\alpha=0}$, we will have proved a single-crossing property, in which $n(t)|_{\alpha=1}$ starts out below $n(t)|_{\alpha=0}$ and then rises above $n(t)|_{\alpha=0}$ after some critical t^* .

To prove this result, we need:

$$\frac{\lambda_1 (p-r) e^{-(p-r)t}}{6\psi_2 n(t) - 2\psi_1} > \frac{\lambda_0 (p-r) e^{-(p-r)t}}{2\psi_2}.$$

Assuming that $p > r$ – which is analogous to our usual assumption that $\beta\gamma > 1$ – and substituting for λ_α , we need:

$$\frac{1 + 2\psi_1 n(t) - 3\psi_2 n(t)^2}{3\psi_2 n(t) - \psi_1} > \frac{\psi_1 - 2\psi_2 n(t)}{\psi_2}.$$

We further assume that $3\psi_2 n(t) > \psi_1$, which is implicitly a condition on the parameters of the utility function: the increase in utility with $n(t)$ cannot be too steep. Then the necessary condition is:

$$\psi_2 + 2\psi_1\psi_2 n(t) - 3\psi_2^2 n(t)^2 > 3\psi_1\psi_2 n(t) - \psi_1^2 - 6\psi_2^2 n(t)^2 + 2\psi_1\psi_2 n(t)$$

which simplifies to:

$$\psi_1^2 + \psi_2 > 3\psi_2 n(t) (\psi_1 - \psi_2 n(t)).$$

It is straightforward to prove that this final condition is always satisfied. The right-hand side is maximized at $n(t) = \frac{\psi_1}{2\psi_2}$, at which value the right-hand side is equal to $\frac{3\psi_1^2}{4}$, which is less than $\psi_1^2 + \psi_2$. Therefore, as long as the parameters are such that $3\psi_2 n(t) > \psi_1$ at the optimum, and as long as $p > r$, so that the optimal population is growing over time, there will be a critical time t^* , before which $n(t)$ is higher for $\alpha = 0$, and after which $n(t)$ is higher for $\alpha = 1$. When the planner is total-utilitarian rather than average-utilitarian, they will desire a population that starts out lower but rises higher in the long-run.

References

- Abel GJ, Barakat B, Samir K, Lutz W. Meeting the Sustainable Development Goals Leads to Lower World Population Growth. PNAS. 2016; 113(50):14294–14299. [PubMed: 27911797]
- Asheim GB, Buchholz W, Hartwick JM, Mitra T, Withagen C. Constant Savings Rates and Quasi-Arithmetic Population Growth under Exhaustible Resource Constraints. Journal of Environmental Economics and Management. 2007; 53(2):213–229.
- Barro RJ, Becker GS. Fertility Choice in a Model of Economic Growth. Econometrica. 1989; 57(2):481–501.
- Baudin T. Family Policies: What Does the Standard Endogenous Fertility Model Tell Us? Journal of Public Economic Theory. 2011; 13(4):555–593.
- Becker GS, Barro RJ. A Reformulation of the Economic Theory of Fertility. Quarterly Journal of Economics. 1988; 103(1):1–25. [PubMed: 11617986]
- Blackorby C, Bossert W, Donaldson D. Intertemporal Population Ethics: Critical-Level Utilitarian Principles. Econometrica. 1995; 63(6):1303–1320.
- Bohn H, Stuart C. Calculation of a Population Externality. American Economic Journal: Economic Policy. 2015; 7(2):61–87.
- Boucekkine R, Fabbri G. Assessing Parfit's *Repugnant Conclusion* Within a Canonical Endogenous Growth Set-Up. Journal of Population Economics. 2013; 26(2):751–767.
- Boucekkine R, Fabbri G, Gozzi F. Egalitarianism under Population Change: Age Structure Does Matter. Journal of Mathematical Economics. 2014; 55:86–100.
- Bradshaw CJ, Brook BW. Human Population Reduction is not a Quick Fix for Environmental Problems. PNAS. 2014; 111(46):16610–16615. [PubMed: 25349398]
- Budolfson M, Dennig F, Fleurbaey M, Scovronick N, Siebert A, Spears D, Wagner F. Optimal Climate Policy and the Future of World Economic Development. World Bank Economic Review. 2017 forthcoming.
- Cigno A. Growth with Exhaustible Resources and Endogenous Population. Review of Economic Studies. 1981; 48(2):281–287.
- Cole MA, Neumayer E. Examining the Impact of Demographic Factors on Air Pollution. Population & Environment. 2004; 26(1):5–21.
- Conde-Ruiz JI, Giménez EL, Pérez-Nievas M. Millian Efficiency with Endogenous Fertility. Review of Economic Studies. 2010; 77(1):154–187.
- Conly, S. One Child: Do We Have a Right to More?. Oxford University Press; New York: 2015.
- Connelly, M. Fatal Misconception: The Struggle to Control World Population. Harvard University Press; Cambridge, MA: 2008.
- Dasgupta P. On the Concept of Optimum Population. Review of Economic Studies. 1969; 36(3):295–318. [PubMed: 12275700]

- Dasgupta P. The Population Problem: Theory and Evidence. *Journal of Economic Literature*. 1995; 33(4):1879–1902.
- Dasgupta, S., Mitra, T. On Some Problems in the Formulation of Optimal Population Policies when Resources are Depletable. In: Eichhorn, W.Henn, R.Neumann, K., Shephard, R., editors. *Economic Theory of Natural Resources*. Springer-Verlag; 1982. p. 409–429.
- de la Croix D, Gosseries A. Population Policy Through Tradable Procreation Entitlements. *International Economic Review*. 2009; 50(2):507–542.
- Dietz T, Rosa EA. Effects of Population and Affluence on CO_2 Emissions. *PNAS*. 1997; 94(1):175–179. [PubMed: 8990181]
- Gaffin SR, O'Neill BC. Population and Global Warming With and Without CO_2 Targets. *Population and Environment*. 1997; 18(4):389–413.
- Golosov M, Jones LE, Tertilt M. Efficiency with Endogenous Population Growth. *Econometrica*. 2007; 75(4):1039–1071.
- Harford JD. Stock Pollution, Child-Bearing Externalities, and the Social Discount Rate. *Journal of Environmental Economics and Management*. 1997; 33(1):94–105.
- Harford JD. The Ultimate Externality. *American Economic Review*. 1998; 88(1):260–265.
- Jones LE, Schoonbroodt A. Complements Versus Substitutes and Trends in Fertility Choice in Dynastic Models. *International Economic Review*. 2010; 51(3):671–699.
- Kelly DL, Kolstad CD. Malthus and Climate Change: Betting on a Stable Population. *Journal of Environmental Economics and Management*. 2001; 41(2):135–161.
- Lam D. How the World Survived the Population Bomb: Lessons From 50 Years of Extraordinary Demographic History. *Demography*. 2011; 48(4):1231–1262. [PubMed: 22005884]
- Liddle B. Consumption-Driven Environmental Impact and Age Structure Change in OECD Countries: A Cointegration-STIRPAT Analysis. *Demographic Research*. 2011; 24:749–770.
- Mitra T. Limits on Population Growth under Exhaustible Resource Constraints. *International Economic Review*. 1983; 24(1):155–168.
- Nerlove M, Razin A, Sadka E. Population Size and the Social Welfare Functions of Bentham and Mill. *Economics Letters*. 1982; 10(1-2):61–64. [PubMed: 12313537]
- Nerlove M, Razin A, Sadka E. Population Size: Individual Choice and Social Optima. *Quarterly Journal of Economics*. 1985; 100(2):321–334.
- Nerlove M, Razin A, Sadka E. Endogenous Population with Public Goods and Malthusian Fixed Resources: Efficiency or Market Failure. *International Economic Review*. 1986a; 27(3):601–609.
- Nerlove M, Razin A, Sadka E. Some Welfare Theoretic Implications of Endogenous Fertility. *International Economic Review*. 1986b; 27(1):3–31.
- Nordhaus, WD. *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press; New Haven and London: 2008.
- O'Neill BC, Dalton M, Fuchs R, Jiang L, Pachauri S, Zigova K. Global Demographic Trends and Future Carbon Emissions. *PNAS*. 2010; 107(41):17521–17526. [PubMed: 20937861]
- O'Neill BC, Liddle B, Jiang L, Smith KR, Pachauri S, Dalton M, Fuchs R. Demographic Change and Carbon Dioxide Emissions. *Lancet*. 2012; 380(9837):157–164. [PubMed: 22784534]
- O'Neill, BC., MacKellar, FL., Lutz, W. *Population and Climate Change*. Cambridge University Press; New York: 2005.
- Palivos T, Yip CK. Optimal Population Size and Endogenous Growth. *Economics Letters*. 1993; 41(1):107–110. [PubMed: 12318664]
- Parfit, D. *Reasons and Persons*. Oxford University Press; Oxford: 1984.
- Pindyck RS. Climate Change Policy: What do the Models Tell Us? *Journal of Economic Literature*. 2013; 51(3):860–872.
- Razin A, Ben-Zion U. An Intergenerational Model of Population Growth. *American Economic Review*. 1975; 65(5):923–933.
- Razin A, Yuen CW. Utilitarian Tradeoff between Population Growth and Income Growth. *Journal of Population Economics*. 1995; 8(1):81–87. [PubMed: 12289751]
- Renström TI, Spataro L. The Optimum Growth Rate for Population under Critical-Level Utilitarianism. *Journal of Population Economics*. 2011; 24(3):1181–1201.

- Robinson, JA., Srinivasan, T. Long-Term Consequences of Population Growth: Technological Change, Natural Resources, and the Environment. In: Rosenzweig, M., Stark, O., editors. Handbook of Population and Family Economics. Vol. 1b. Elsevier; 1997. p. 1175-1298.
- Spears D. Smaller Human Population in 2100 Could Importantly Reduce the Risk of Climate Catastrophe. PNAS. 2015; 112(18):E2270. [PubMed: 25848063]
- Spears D. Making People Happy or Making Happy People? Questionnaire-Experimental Studies of Population Ethics and Policy. Social Choice and Welfare. 2017; 49(1):145–169.
- Stern, N. The Economics of Climate Change: The Stern Review. Cambridge University Press; Cambridge and New York: 2007.
- Stern N. 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. 2013; 51(3):838–859.
- Stiglitz JE. Growth with Exhaustible Natural Resources: The Competitive Economy. Review of Economic Studies. 1974; 41:139–152. Symposium on the Economics of Exhaustible Resources.
- Weitzman ML. On Modeling and Interpreting the Economics of Catastrophic Climate Change. Review of Economics and Statistics. 2009; 91(1):1–19.
- Wheeler D, Hammer D. The Economics of Population Policy for Carbon Emissions Reduction in Developing Countries. 2010 Working Paper 229, Center for Global Development

Figure 1. Optimal Population Paths in Simple Infinite Horizon Model

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 $\beta = \{0.963, 0.967, 0.971, 0.975\}$.

(a) Optimal Population Path

(b) Optimal Population Growth Rate

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.

Figure 3. Optimal Generation Sizes in Climate Change Model

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 $\hat{\beta} = \{0.985, 0.9875, 0.99, 0.9925\}$.

Figure 4. Optimal Population Sizes in Climate Change Model

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 $n_0 = 2.934$, $n_9 = \frac{n_8 + 4}{2}$, $\gamma = 1.381$, $X = 24$, $\kappa = 0.1811$, 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.