Would a larger population increase long-run incomes?*

Maya Eden† Kevin Kuruc‡

January 2025

Abstract

The relationship between the human population size and per capita incomes has long been debated. Two competing forces feature prominently in these discussions. On the one hand, a larger population means that limited natural resources must be shared among more people. On the other hand, more people means more innovation and faster technological progress, other things equal. We study a model that features both of these channels, and which gives rise to a simple sufficient statistic governing the sign of the effect of population size on steady-state incomes. A calibration suggests that, in the long run, (marginal) increases in population would likely lead to (marginal) increases in per capita incomes.

*We would like to thank Gustav Alexandrie, John Firth, Oded Galor, Mike Geruso, Chad Jones, Luis Mota Freitas, Marta Prato, Bernardo Ribeiro, Dean Spears, Phil Trammell, David Weil, Mu-Jeung Yang, Anson Zhou and participants at CEPR's workshop on Sustainability and Public Policy, the Global Priorities Institute's workshops on Global Priorities Research, UC Davis' ARE seminar, the University of Texas Macroeconomics Workshop, and the Population Wellbeing Initiative's Innovation and Scale Workshop for conversations and comments that have greatly improved the paper. KK's research is supported by grant P2CHD042849 awarded to the Population Research Center at UT Austin by the Eunice Kennedy Shriver National Institute of Child Health and Human Development and by the Musk Foundation. Gage Weston provided excellent research assistance. All remaining errors are our own.

†University of Zurich; CEPR; Global Priorities Institute, University of Oxford. Email: maya.eden@econ.uzh.ch

‡Population Wellbeing Initiative, University of Texas at Austin; Global Priorities Institute, University of Oxford. Email: kevinkuruc@utexas.edu.

1 Introduction

Later this century the global population is expected to peak and then begin declining. This prompts one of the oldest open questions in economics: should we expect smaller populations to generate better per capita outcomes?

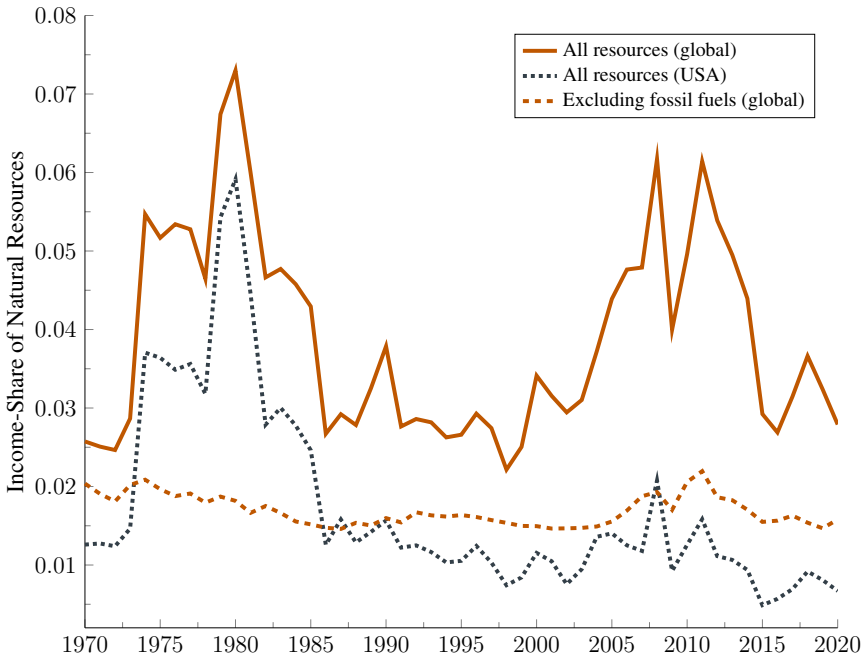
There are, of course, many dimensions to this question. Here, we focus on one: the long-run relationship between population levels and per capita incomes. The classic Malthusian concern is that larger populations strain our productive natural resources and spread the benefits of ecosystem services more widely. This suggests a negative relationship between population sizes and per capita incomes. More recently, economists have come to formalize important channels by which a larger population could have competing benefits. In particular, people produce infinitely shareable knowledge that increases everyone's productivity. If the number of ideas is increasing in the number of people, then larger populations will contribute to a more productive economy (Romer, 1990; Jones, 2005).

This paper compares the quantitative importance of these two channels using leading models from these respective sub-disciplines. In our framework, the Malthusian concern is captured by a production function that utilizes natural resources that are in fixed supply. The innovation channel is captured by a positive relationship between population size and productivity growth. One difficulty in comparing these two forces is that diluting environmental services across people has a persistent level effect on income, while innovation gains have growth effects. We overcome this with a tweak to the standard semi-endogenous growth framework. The result is a steady-state relationship between population size and the level of productivity. When combined with the Malthusian channel, this delivers a characterization of steady state per capita incomes that depends directly on population sizes.

Using this solution we can calibrate the long-run relationship between population levels and per capita income based on previously-estimated parameters. The innovation process depends on

a ratio of parameters that has been recently studied in detail (Bloom et al., 2020; Peters, 2022; Ekerdt and Wu, 2023; Terry et al., 2024). The effect of diluting the fixed factor across more individuals is directly captured by the income-share of natural resources, a feature anticipated by Weil and Wilde (2009). The overall effect of population size on long-run income is determined by whether the income-share of natural resources is larger or smaller than the relevant parameters in the innovation equation.

Figure 1: The share of income going to natural resources is small and non-increasing



Notes: Share of income paid to natural resources over time. (Solid) Global income-share paid to all natural resources (following World Bank classification, see Appendix E). Included are: (a) subsoil energy and minerals; (b) timber resources; (c) crop land; (d) pasture land. (Dotted) This same income-share, but using only US data to ensure that the level or trend is not driven by less reliable data sources (see also Figure A2). (Dashed) Same income-share, at the global level, excluding fossil fuels because they disproportionately contribute to both the level and volatility, but do not drive the flat trend.

The main finding is that the innovation effect dominates. Underlying this finding is the observation that the share of global income accruing to natural resources is small, at less than 5%, and has remained trendless over the past half century, as documented in Figure 1. Over the same period there has been a doubling of the global population, indicating a substantial change in factor input-

shares. These observations—coupled with the fact that long-run populations are unlikely to be far outside of the historical experience—constitute strong evidence that the long-run natural resource income share will not grow substantially over time. We show that this implies the Malthusian channel is likely to be smaller than estimates of the benefits of long-run productivity improvements: The population-income relationship is therefore positive under our baseline calibration.

The baseline model we use to generate this result is intentionally simple. The point of the main exercise is to draw out the implications of the most straightforward combination of these two forces once they have been modified to fit together and calibrated to existing data. This, of course, leaves other open questions: What if technological progress is biased against natural resources in the long run? What if there are non-rival benefits of ecosystems, such as those provided by lower global temperatures? What if there is a negative, rather than null, relationship between the population size and the aggregate supply of natural resources? We extend the model along these dimensions to ask whether the initial finding can be overturned. These extensions do not lead to a qualitatively different comparison, leaving the main takeaway unchanged.

This paper contributes to a longstanding and still active literature on the effects of population sizes on per capita outcomes in two ways. Theoretically, we propose a novel, parsimonious framework that can compare growth effects of additional innovation with the level effects of natural resource dilution. Quantitatively, we leverage the time-series of natural resource income-shares to sign the model-implied population-income relationship. We are not the first to recognize that natural resource income-shares are relevant for quantifying the Malthusian channel (see e.g., Weil and Wilde, 2009), but we contribute a method of transparently comparing this series to the innovation effects of population increases.

By theoretically merging and quantifying competing channels, we advance multiple active lines of research. First, work such as Peters and Walsh (2021), Hopenhayn et al. (2022), Jones (2022), Karahan et al. (Forthcoming) and others highlight reasons to expect a positive effect of long-run population growth on long-run economic growth. Recent empirical work has documented support,

and provided parameter estimates, for the specific semi-endogenous growth channel that we employ (Bloom et al., 2020; Peters, 2022; Kruse-Andersen, 2023; Terry et al., 2024). However, these papers are, for the most part, not performing evaluations that also consider the drawbacks of population growth or size. The micro-evidence in Peters (2022) and Terry et al. (2024)—that localities receiving larger population in-flows end up with higher incomes—is evidence in this direction, but these analyses leave open whether global environmental constraints would begin to bite if *all* localities got larger. We complement these existing efforts by advancing the theoretical literature in a way that demonstrates what must be true for a positive relationship to arise when these global environmental concerns are accounted for.

There is likewise a robust literature studying the *negative* effects of population growth and/or size. Weil and Wilde (2009), Ashraf et al. (2013) and Karra et al. (2017) are focused on the case of population growth in developing countries, where agricultural resources are a larger share of inputs and a quality-quantity trade-off in child raising may be more relevant. Recently, Henderson et al. (2022) studies local population growth and how the sharing of resources compares to the magnitude of damages from climate change (with each channel sharing the common property of reducing per capita natural resource availability). Dasgupta et al. (2021) uses a simple global model like our own to argue that elevating per capita incomes to an adequate level requires large population reductions. However, any channel by which population size or growth can improve incomes is absent from their analysis. We take the relevant findings from this literature for our evaluation of the net-effects of population size.

In merging these considerations, we are building from and extending an earlier literature focused on the transition from a Malthusian growth regime to a modern growth regime (e.g., Kremer, 1993; Galor and Weil, 2000; Jones, 2001). Rather than seeking to explain how and why growth increased in the past, we look forward. Therefore, we study the effect of different long-run population *levels*. The era of positive human population growth appears to be coming to an end, making the question about size rather than growth rates the more relevant one. This has theoretical and

empirical advantages that allows us to study this relationship in a novel way, which we describe in detail in Section 1.1.

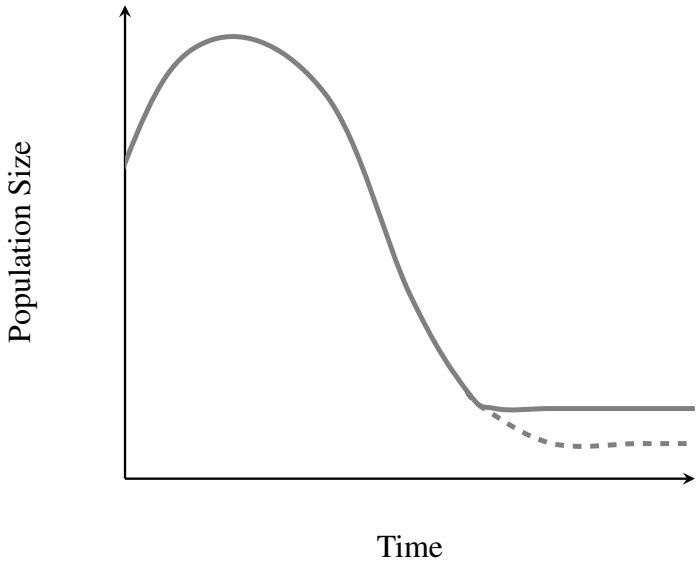
More recently, Peretto and Valente (2015), Bretschger (2020), Budolfson et al. (2023), Gerlagh (2023) and Kruse-Andersen (2023) have, like us, merged models of endogenous growth with natural resource constraints to look forward. Their respective focuses are distinct from ours. Peretto and Valente (2015) seeks to understand which parameter conditions lead endogenously determined populations to explode, collapse, or stabilize. We assume that populations will stabilize *somewhere*—perhaps via social or economic policy that induces replacement fertility rates—and ask whether incomes would be higher if populations stabilize at higher levels. Bretschger (2020), Budolfson et al. (2023), Gerlagh (2023) and Kruse-Andersen (2023) are interested in fossil fuel use and climate change under different population growth regimes, an important but distinct issue that we discuss in more detail in Section 2.3.

1.1 Focusing on levels, rather than growth rates, is analytically simpler and more policy relevant

An important feature of our study relative to past work is our focus on levels, rather than growth rates. This focus is a primary innovation of this paper—which informs the formal modeling choices—so we discuss it prior to presentation of model equations.

To fix ideas, Figure 2 illustrates a case of what we have in mind. In line with leading demographic projections, the global population peaks and then begins declining. At some future date, we assume that a zero-growth steady state is reached; for example, policy is enacted that raises fertility rates, or social norms change after a phase of global population decline. (If zero growth is never reached humanity dwindles until extinction, a distinct case that we ignore.) A marginal population increase is as simple as imagining that societies arrest the decline earlier than otherwise, but have common long-run policy and preferences.

Figure 2: A marginal long-run population increase as a different stabilization date



Notes: In line with Spears et al. (2024), we take population decline to be the default long-run outcome until some point of stabilization.

The first thing to notice is the conceptual departure from balanced growth analyses, which are commonly used to assess the population-productivity relationship. On a BGP all variables grow at some constant exponential rate. If population growth is positive, the limiting outcome is an infinitely large population that has infinitesimally few natural resources per person, but is infinitely productive. A comparative static with respect to population growth asks whether it would be better to race towards that future a bit faster.[1] While analyses of these sorts have been informative for understanding the historical record, it remains unclear what to make of their extrapolations.

Our question about levels is less ambitious—we do not need to take a stance on what happens at much larger population levels—so it requires fewer structural assumptions. For example, only a small class of production functions deliver BGPs, and these restrictions omit important complementarities that are prominent in discussions about environmental inputs.[2] A question about

[1]Conversely, a comparative static relative to a negative population growth regime asks if it would be better to race towards a zero-population outcome a bit slower.

[2]This is because a BGP is necessarily one where a constant rate of population growth leads to a constant rate of

steady-state relationships does not need these restrictions. Moreover, observable empirical moments may be more reliable for our marginal analysis because we are focused on population sizes near or within the historical experience.

This focus also simplifies the analysis by isolating scale effects from other sources by which population growth could influence economic growth; for example, larger average family sizes (e.g., human capital investment per child) or economy-wide aging (e.g., pension program financing). In fact, the existence of these forces is an important reason why a model is necessary at all. The last century of per capita income growth has come alongside both growth in the size of the population and declining rates of fertility, making it difficult to infer from aggregate trends alone the role that population size has played. But stabilizing at any aggregate size requires fertility rates to converge to one child per adult-life, which implies that the ratios between children, working adults and the elderly will be independent of the population size. So, assessing the effects of stabilizing at different sizes requires calibrating only the productivity-based scale effects, allowing us to abstract from the complicated and competing effects of population growth.

Despite being less ambitious and side-stepping important questions related to growth rate differences, our local-levels question is likely the more relevant one for forward-looking analyses. It is no longer the case that we should anticipate exponential growth in populations; the prevailing equilibrium appears to be one of zero or negative population growth. Unless humanity is to endogenously dwindle to zero people, there must eventually be something approximating stabilization. A relevant input to understanding whether it would be better to stabilize sooner or later is the relationship between per capita incomes and long-run population sizes. It happens to be that this question also (i) makes the analysis significantly simpler, (ii) allows us to relax standard assumptions and (iii) avoids open debates about human capital accumulation and other implications of changes to age pyramids and fertility rates.

per capita income growth—it cannot be the case that environmental inputs become more binding as population sizes grow.

2 Model

Time is continuous and indexed $t \in \mathbb{R}_{++}$. In each period aggregate GDP is given by

$$Y(t) = A(t)F(N(t), K(t), E(t)) \tag{1}$$

where $A(t) > 0$ is total factor productivity, $N(t)$ is the global labor force, $K(t)$ is physical capital and $E(t)$ is the natural resources used in production.

We assume that the production function, F, has constant returns to scale. Holding fixed productivity, doubling the amount of labor, capital, and natural resources would double aggregate output. In practice, natural resources do not increase, creating scope for a negative relationship between per capita GDP and population: if we double population (and even capital) without increasing natural resources, then GDP will less-than-double and GDP per capita will decline.

The possibility of a positive relationship between population and per capita GDP is introduced through the process of technological progress. The technology parameter, $A(t)$, consists of ideas, which are produced by people and can be used freely and indefinitely to improve the production process. Because more people generate more ideas, productivity is increasing in the number of people. Our investigation concerns the relationship between long-run income per capita, $\bar{y}$, and the long-run population level, $\bar{N}$ (where $\bar{x}$ indicates a long-run, steady-state value of x).

2.1 Natural resources

The $E(t)$ in Equation 1 is conceptually an aggregator of arbitrarily many environmental or agricultural inputs.

$$E(t) = H(e_1(t), e_2(t), e_3(t), \ldots, e_Z(t)) \tag{2}$$

Some of these inputs will be minerals (e.g., lithium), some will be renewable resources (e.g., timber), some will represent fixed agricultural land. What makes the subcomponents of $E(t)$

unique is that we assume they cannot be created by people. So, their availability and use does not scale with population sizes. This is in contrast to the subcomponents of $K(t)$, which can be accumulated through savings and investments. Formally, we assume that

$$\frac{\partial \bar{E}}{\partial \bar{N}} \leq 0. \tag{3}$$

It can be temporarily true that natural resource use increases with population size, but in the steady state that we study, the long-run use of E cannot respond positively to N. The baseline analysis will assume this derivative is precisely zero—that is, natural resources are fixed. In robustness exercises we relax this to explore cases where the availability of aggregate natural resource stocks is reduced by a larger population.

It is important to notice that we have not imposed any restrictions here on the functional form of H. These natural resources can interact in any arbitrary way with one another. Just as it is standard to assume that K aggregates a wide variety of individual capital inputs—computers, buildings, railways, etc.—this E is aggregating a wide variety of individual natural resources. In Appendix A we show that the main results pass through to the case where these individual e_is enter the aggregate production function directly and can each likewise interact with capital and labor in any arbitrary way.

To understand the properties of $E(t)$ if/when it converges to a long-run steady state, it will help to consider three categories of natural resources: (i) exhaustible, non-renewable; (ii) exhaustible, renewable; (iii) non-exhausted. We discuss these in turn.

Exhaustible, non-renewable resources (e.g., fossil fuels). To start, note that, in the long run, exhaustible and non-renewable resources will be exhausted (and not renewed), and therefore, for any e_i in this category

$$\lim_{t \to \infty} e_i(t) = \bar{e}_i = 0.$$

Setting these e_is to zero need not, and will not, imply zero output. This is because we take them to have perfect substitutes, albeit with different productivity levels. In particular, fossil fuels make up a large majority of what we consider in this category. It seems clear that the total elimination of fossil fuels would not end all, or even most, economic production.[3]

Exhaustible, renewable resources (e.g., timber, fish). The second category—exhaustible but renewable natural resources—is more complicated. It is possible to generate a long-run cycle in which resources are periodically exhausted, regenerated, and exhausted again; in this case, there will be no steady state. However, for tractability, we will restrict our focus to steady state solutions where the level of withdrawal of these resources is constant. Appendix B details the problem of renewable resource management in long-run steady states. The relevant and intuitive takeaway is that the level drawn each period depends only on the regeneration rate of the natural resources, and is therefore independent of population size. The number of fish we can sustainably draw from the ocean is co-determined with the size of the long-run fish population, not the long-run human population.

Non-exhausted resources (e.g., agricultural land, solar energy). The third category consists of natural resources that are not exhausted when used. Because there is no dynamic cost to using these resources, eventually, all economically viable inputs of this type will be in use each period.[4] Thus, there exists some $\bar{e}_i$ such that $e_i(t) = \bar{e}_i$ for non-exhausted resources.

The most trivial example of this is solar energy. There is some fixed amount of solar energy that will be viable to harness each period (conditional on a long-run level of A), and the amount available in the following period is independent of how much is used today. Minerals like lithium

[3]Furthermore, Appendix C discusses how the population level affects the rate at which non-renewable natural resources are exhausted, which will be important for transition dynamics. It turns out that, in standard settings, it does not (Bretschger, 2020).

[4]*Economically viable* here meaning that the marginal benefit of its use exceeds the marginal cost of accessing/capturing it.

likely fall into this category as well. Eventually all economically viable lithium will be embedded in products because it cannot be created or destroyed by standard human activity.[5] And, of course, agricultural land is a canonical example of this category of resources.[6]

In summary, the substantive assumption is that the aggregate production function relies on inputs that cannot grow indefinitely over the long-run. The constraint on how much of each e_i can be used in steady-state is the regeneration of that resource each period. For exhaustible, non-renewable resources this is zero; for renewable resources, it will be non-zero but less than the entire stock; for non-exhaustible resources the entire accessible stock can be used each period. Furthermore, though not modeled explicitly, these constraints may interact with one another. If, for example, farmland crowds out forests, the amount of farmland in use will partially determine how many trees can be harvested for timber. The steady state vector $\mathbf{e} = (\bar{e}_1, \bar{e}_2, \ldots)$ should be thought of as lying on some production-possibilities-like surface. That detail is not consequential for our purposes since we will assume the vector $\mathbf{e}$ is fixed with respect to small changes in population size in our main analysis.

The straightforward implication of this conceptualization is that a larger population necessarily has fewer of these inputs per capita. Holding fixed A, this would imply a negative relationship between population and per capita income:

$$\bar{y} = \frac{\bar{Y}}{\bar{N}} = \bar{A}F\Big(1, \frac{\bar{K}}{\bar{N}}, \frac{\bar{E}}{\bar{N}}\Big). \tag{4}$$

(where the first equality is a definition and the second equality follows from the assumption that F has constant returns to scale). This is what we refer to as the Malthusian channel.

[5] Perfect, zero-depreciation, recycling between products is admittedly idealistic, but if the rate of loss between uses is small enough these resources are approximately non-exhausted. If the depreciation is significant, these minerals conceptually belong conceptually in the exhaustible, non-renewable category.

[6] Technically, the nutrients in soil may be better categorized as "exhaustible, renewable". This will not matter for our purposes.

2.2 Technology

The literature on endogenous economic growth highlights a positive relationship between population size and productivity (Jones and Romer, 2010). Larger populations generate more goods, some of which are non-rival. An increase in aggregate non-rival goods increases per capita variables, since they can be used/consumed by everyone, regardless of the population size. Drawing on the semi-endogenous growth literature (Jones, 1995, 2022), we focus on the non-rival good that is knowledge and assume that the law of motion for total factor productivity is

$$\frac{\dot{A}}{A} = \alpha N^{\lambda} A^{-\beta} - \delta_A \tag{5}$$

where $\alpha, \lambda, \beta, \delta_A > 0$. To interpret this expression, consider first the case in which $\lambda = 1$ and $\beta = 0$. In this case, the accumulation of knowledge is proportional to the size of the population. This captures a simplistic model in which each person adds a constant amount to the stock of knowledge. The parameter δ_A governs the rate at which the stock of knowledge depreciates: unless people invest in knowledge preservation, ideas get forgotten or go unused.

When $\beta > 0$, we have a negative relationship between the rate of productivity growth and the level of productivity. If the stock of knowledge is already large, new ideas become harder to find, leading to ever-declining productivity growth in the absence of increases in research inputs. A larger β implies more decreasing returns to innovation due to this channel.

The parameter λ governs the extent to which population size affects the amount of innovation in each period. Specifying $\lambda > 1$ captures a situation in which innovation benefits from collaboration among more people. In contrast, $\lambda < 1$ captures a situation in which there are diminishing returns to R&D efforts. For example, if innovation happens through sequential discoveries, there may be diminishing returns to having more people work on discovering the same thing in the same period.

When population is constant at $\bar{N}$, Equation 5 implies a steady state for A.[7]

$$\bar{A}(\bar{N}) = \left(\frac{\alpha \bar{N}^{\lambda}}{\delta_A}\right)^{\frac{1}{\beta}} \tag{6}$$

Conceptually, this steady state might be thought of as the stock of knowledge that is sufficiently large that to even maintain, organize, and employ it commands all people-hours in this sector.[8] As the stock of knowledge gets unwieldy, the challenge is making use of existing knowledge, not generating new ideas.

This interpretation is not crucial: any arbitrarily small δ_A generates this expression.[9] We do not take a stance on whether this is a quantitatively significant force. Of course, the value of δ_A matters for the level of long-run per capita incomes. But it can be easily seen that the relevant elasticity between $\bar{N}$ and $\bar{A}$ relies only on the ratio of λ to β. The introduction of an arbitrarily small δ_A is the minimal modification we make to the standard innovation equation that generates the analytical solution for the long-run per capita incomes that we study.

These two model components—endogenous TFP; fixed natural inputs—make up the core of our steady-state analysis. As we will demonstrate below, very few additional assumptions are necessary to derive an expression for the sign of the long-run population-income relationship.

2.3 Model Omissions: Unsustainable Resource Use and Climate Change

This long-run framing omits two related challenges at the forefront of environmental debates: (1) achieving long-run sustainable resource use, and (2) mitigating climate change. Despite widely held beliefs that population reduction could contribute to addressing these challenges, it is not obvious that this is the case. As these are critically important issues, we detail these omissions

[7] Simply set $\dot{A} = 0$ and solve for A.

[8] This could be due to the sheer breadth of knowledge society acquires or the increased domain-expertise necessary to contribute to organizing, using and preserving knowledge.

[9] In fact, the main relationship we derive holds even for $\delta_A = 0$ (see Appendix D).

before proceeding to the results of the model.

Long-run Sustainability. The model we study is one where humanity reaches a long-run regime of sustainable natural resource use: each e_i is drawn at a level that preserves a long-run stock that can support those withdrawals. We implicitly assume that whether such a balance is achieved is independent of population sizes, and therefore unimportant for the analysis.

We assume this because it is ambiguous how population size influences the probability of long-run sustainable resource use; sustainability must be achieved through institutions that protect the commons. With low enough extraction costs, even small populations will use up their resources in the absence of proper disincentives.[10] The relevant (open) question is whether a larger or smaller population makes it more likely that these disincentives are designed and enforced. On the one hand, the costs of sustainability will be larger for a larger population, pushing against efforts to control resource use: larger populations will need to constrain themselves to lower levels of per capita use, which raises the per capita opportunity costs of not defecting. On the other hand, some of the additional non-rival ideas that larger populations produce will almost certainly relate to the design of optimal institutions. Moreover, if a larger population is wealthier over the medium term (as our later calibration of this model suggests might be the case), this wealthier population may be willing to spend more of its resources promoting sustainability. This is a relationship that arises for most environmental issues—development comes with more successful protection of the commons (Dasgupta et al., 2002; Ritchie, 2024)—and in models studying general longevity trade-offs against consumption gains (Hall and Jones, 2007; Jones, 2016; Aschenbrenner and Trammell, 2024).[11]

[10] In a striking example, very small bands of early humans appear to have hunted more than half of the world's (non-African) large mammals to extinction (Ritchie, 2022). (Interestingly, the likely reason that African mammals are an exception is that they co-evolved alongside humans.)

[11] This is likewise why we do not expect increased transient, local pollution to be a consequential direct effect of population size: local air and water quality appear more sensitive to per capita income levels than population size, so the overall effect of population size on pollution levels seems likely to be determined by whether a larger population size is wealthier per capita rather than an additional independent effect of population.

In sum, we assume long-run sustainability not only because it is analytically convenient, but also because it is not clear whether marginal population changes are an important factor in whether sustainability is achieved.

Climate change. We omit climate change from the model for two related reasons. First, it must be the case for our long-run survival that per capita emissions are eventually negligible. This renders greenhouse gases conceptually unimportant for the marginal analysis we conduct about changes to long-run population sizes.

We acknowledge, however, that this framing may miss important transition dynamics that occur prior to full decarbonization. A larger medium-term population may emit more greenhouse gases for the next century or more, impacting long-run living standards. If the population changes that we study occur over the next decades, this could potentially be a factor that our analysis misses.

Quantitatively, research on this question indicates this is unlikely to be an important omission. This literature has emphasized that, because population is a stock, changes to its size happen too slowly to influence near- and medium-term environmental challenges, such as limiting greenhouse gas emissions (Bradshaw and Brook, 2014; Budolfson et al., 2023; Lutz, 2023). Budolfson et al. (2023) demonstrate that in realistic models of population dynamics, even immediate, large changes in fertility rates have an extremely small effect on long-run levels of global warming: by the time significant changes in the size of the population occur the long-run stock of atmospheric greenhouse gases is almost entirely determined.[12,13] This line of reasoning has led us to focus on long-term environmental constraints as the first-order issue, though we direct readers to this related literature for a more thorough discussion of the specific population-climate relationship.

[12]Specifically, an instantaneous change in fertility rates large enough to nearly double the population by 2200 only increases global temperatures from 4.17C to 4.22C on a pessimistic trajectory of emissions policy where decarbonization is not achieved until 2125.

[13]Moreover, this is arguably conservative. Bretschger (2020) shows that the optimal rate of aggregate emissions is independent of the population size, so we might expect policy makers to endogenously hasten the decline in per capita emissions if population growth rates increased.

3 Main result

We begin this section with an analytical solution for the relationship of interest: the elasticity of per capita income with respect to population size. We then discuss our process for determining the sign of this elasticity by calibrating the long-run values for the key parameters of this equation. Together, this provides strong evidence for a positive long-run population-income relationship.

3.1 Analytical Elasticity

The goal of this exercise is to determine the sign of the elasticity of GDP per capita, $\bar{y}$, with respect to population, $\bar{N}$. Note that GDP per capita is given by

$$\bar{y}(\bar{N}) = \frac{\bar{Y}(\bar{N})}{\bar{N}} = \frac{\bar{A}(\bar{N})F(\bar{N}, \bar{K}, \bar{E})}{\bar{N}}. \tag{7}$$

We can derive the elasticity in a few simple steps.

$$\begin{aligned}
\frac{\partial \ln(\bar{y})}{\partial \ln(\bar{N})} &= \frac{\partial \ln(\bar{A})}{\partial \ln(\bar{N})} + \frac{\partial \ln(F(\bar{N}, \bar{K}, \bar{E}))}{\partial \ln(\bar{N})} + \frac{\partial \ln(F(\bar{N}, \bar{K}, \bar{E}))}{\partial \ln(\bar{K})} \frac{\partial \ln(\bar{K})}{\partial \ln(\bar{N})} - \frac{\partial \ln(\bar{N})}{\partial \ln(\bar{N})} \\
&= \frac{\lambda}{\beta} + \frac{\frac{\partial F(\bar{N}, \bar{K}, \bar{E})}{\partial \bar{N}} \bar{N}}{F(\bar{N}, \bar{K}, \bar{E})} + \frac{\frac{\partial F(\bar{N}, \bar{K}, \bar{E})}{\partial \bar{K}} \bar{K}}{F(\bar{N}, \bar{K}, \bar{E})} \frac{\partial \ln(\bar{K})}{\partial \ln(\bar{N})} - 1 \\
&= \frac{\lambda}{\beta} + \phi_N + \phi_K \frac{\partial \ln(\bar{K})}{\partial \ln(\bar{N})} - 1
\end{aligned}$$

The first equality merely takes advantage of the additive nature of elasticities between the multiplicative A, F, and N (in the denominator). There is no term corresponding to a natural resource elasticity because the derivative of $\bar{E}$ with respect to $\bar{N}$ is zero in steady state. The second equality subs in what we know about how $\bar{A}$ responds to $\bar{N}$ from Equation 6 and replaces the other log-derivatives with more useful terms, which in the third equality are replaced by ϕs. These ϕs have an economically meaningful interpretation; they represent the share of income paid to the respective factors in competitive markets. To see this, note that $\bar{A}\frac{\partial F(\bar{N}, \bar{K}, \bar{E})}{\partial \bar{N}}$ is the marginal

product of labor. Thus, the term $\bar{A}\frac{\partial F(\bar{N},\bar{K},\bar{E})}{\partial \bar{N}}\bar{N}$ is the aggregate payments to labor, and the ratio $(\frac{\partial F(\bar{N},\bar{K},\bar{E})}{\partial \bar{N}}\bar{N})/F(\bar{N},\bar{K},\bar{E}) = (\bar{A}\frac{\partial F(\bar{N},\bar{K},\bar{E})}{\partial \bar{N}}\bar{N})/(\bar{A}F(\bar{N},\bar{K},\bar{E}))$ is the share of these payments in output.

Now, it will help to re-write $\frac{\partial \ln(\bar{K})}{\partial \ln(\bar{N})} = 1 + \frac{\partial \ln(\bar{k})}{\partial \ln(\bar{N})}$, where k is the capital-to-worker ratio. This allows us to separate a direct effect on per capita productivity from an indirect effect on per capita capital stocks. Additionally, we leverage the fact that with our CRS function $1 - \phi_N - \phi_K = \phi_E$. The share of income going to E is the share of income not going to capital and labor.[14] These two re-definitions deliver a straightforward analytical result for the elasticity of interest.

$$\begin{aligned}
\frac{\partial \ln \bar{y}}{\partial \ln \bar{N}} &= \frac{\lambda}{\beta} + \phi_N + \phi_K\Big[\frac{\partial \ln \bar{k}}{\partial \ln \bar{N}} + 1\Big] - 1 \\
&= \frac{\lambda}{\beta} - (1 - \phi_N - \phi_K) + \phi_K \frac{\partial \ln \bar{k}}{\partial \ln \bar{N}} \\
&= \underbrace{\frac{\lambda}{\beta} - \phi_E}_{\text{direct effect}} + \underbrace{\phi_K \frac{\partial \ln \bar{k}}{\partial \ln \bar{N}}}_{\text{indirect}}
\end{aligned} \tag{8}$$

Conceptually, the direct effect captures how productivity gains from TFP improvements compare against productivity losses from reduced per capita natural resources, holding fixed capital per worker. The indirect effect captures the fact that capital per worker will, in general, endogenously adjust to these changes in productivity via the effect on savings behavior and the investment production function. Below we argue that it is very unlikely that the indirect effect could overturn the sign of the direct effect—if anything, the indirect effect should amplify the direct effect; more productive economies tend to accumulate more capital per worker—so we will focus first and foremost on signing the direct effect.

[14] We recognize that under increasing returns, all factors cannot in fact be paid their marginal product. In spite of this issue, it is common to retain this assumption when studying factor shares in other contexts. We follow that convention. Formally, our model employs a learning-by-doing assumption rather than an endogenous choice to develop a new idea, so it is consistent with our framework to assign ideas zero income despite a positive marginal product.

3.2 Calibration of key terms suggests that incomes increase in N

The first term in the direct effect represents how much more knowledge can be accumulated and productively used in the steady state of this economy from a 1% increase in N. This is governed by the ratio of the intra-period returns to research effort, λ, and the degree to which knowledge becomes more difficult to accumulate as A increases, β. These are critical parameters in leading growth models, and a host of recent papers have made progress in estimating these values or closely related elasticities. Perhaps most influentially and directly relevant, Bloom et al. (2020) targets the ratio we are interested in and estimates that $\frac{\lambda}{\beta} \in (0.2, 0.5)$ for the U.S. economy.[15] Ekerdt and Wu (2023) build on Bloom et al. (2020) by generalizing the analytical assumptions employed and find lower values of β, and hence *larger* values for $\frac{\lambda}{\beta}$. Peters (2022) and Terry et al. (2024) take an entirely different approach, leveraging micro-data on historical population flows, and each estimate a closely related term to also be significantly larger than what is implied by Bloom et al. The totality of the evidence suggests that the elasticity between population size and long-run productivity is at least in the (0.2,0.5) range and could perhaps be even closer to one.

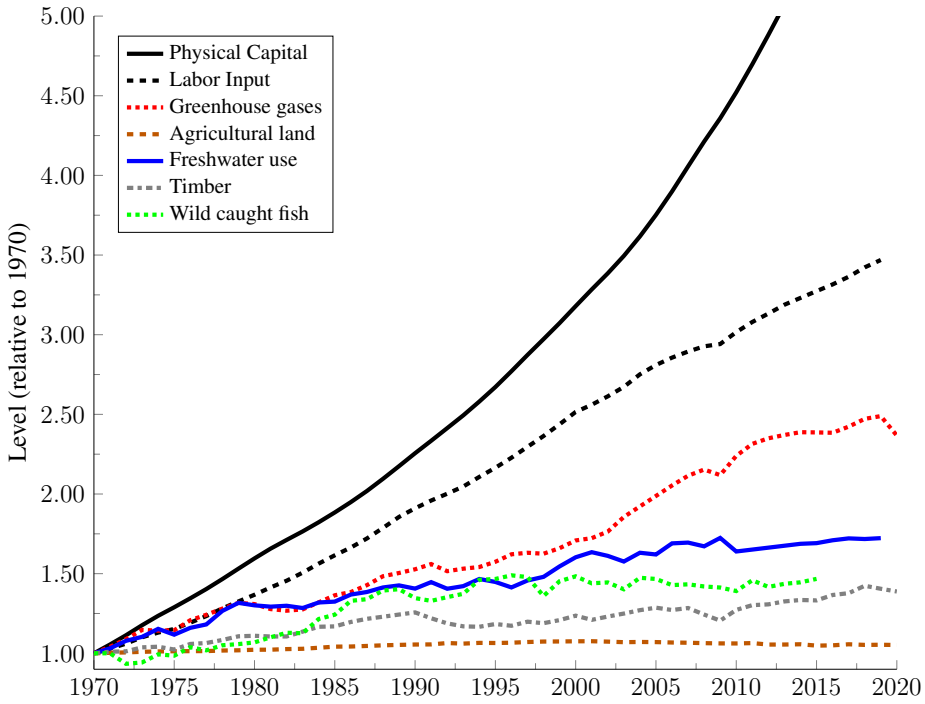
Turning to ϕ_E, consider again Figure 1 which plots the time-series for this value. Our construction of ϕ_E is detailed in Appendix E, but we note here that we follow Caselli and Feyrer (2007) and Monge-Naranjo et al. (2019) who themselves follow the World Bank's *Changing Wealth of Nations* reports. Subsoil energy and minerals, timber resources and all agricultural land are included as the economically relevant natural resources earning non-trivial rents. The levels of ϕ_E in Figure 1 correspond closely to the values others in this literature report and use (see also e.g., Weil and Wilde, 2009; Hassler et al., 2021).

The value of ϕ_E that is relevant for Equation 8 is its long-run value. Historical values are

[15] Even if this elasticity is only an estimate for innovation in frontier economies, it is still the relevant global elasticity as long as the TFP of non-frontier countries scales with innovations in frontier countries (which seems likely given historical knowledge spillovers and observed catch-up dynamics). A 1% uniform increase in the global population is a 1% increase in the population on the research frontier, which is a λ/β% increase in both frontier and non-frontier long-run TFP.

only partially informative about where this term will converge. The special case when observed values of ϕ_E are directly relevant without further assumptions is when the input-share of E in the long run is similar to the input-shares observed during the sample period. Because we have assumed that A is Hicks-neutral, if each of N, K, and E remain at their 2019 levels indefinitely, then the 2019 value of ϕ_E will be its long-run value. In cases like this, the main result follows immediately: observed levels of ϕ_E (about 0.03) are well below estimates of λ/β (about 0.3), so the direct productivity effect of interest will be positive.

Figure 3: Natural resource use does not match human or physical capital growth



Notes: Levels of resource inputs since 1970. Physical capital and the labor input are taken from the Penn World Tables 10.01 where labor input is an aggregate of hours provided and a human capital index (Feenstra et al., 2023). Aggregate natural resources are proxied by a range of resources with well-documented withdrawals. Greenhouse gases come from total fossil fuel use; agricultural land is the sum of crop and grazing land; freshwater measures agricultural, industrial and domestic use; timber is measured as roundwood (the pre-production measure of wood retrieved from forests); wild caught fish measures all fish production not from aquaculture.

For ϕ_E to grow significantly, (1) the long-run input-share of E must be outside of its observed range and (2) the aggregate production function must be such that this *input*-share deviation causes

a large deviation in the *income*-share of E. Figures 1 and 3, taken together, provide reason to believe that this is unlikely for population sizes near historical levels. Figure 3 plots the growth of N, K and various components of E since 1970. Both N and K have grown significantly faster than any of the E components plotted. This implies that the input-share of E must have fallen during this period. Nonetheless, Figure 1 documents that the income-share of natural resources has not systematically evolved over this period. This suggests that the input-shares would need to change dramatically to generate a ten-fold increase in ϕ_E.

To say something more formal about these long-run values, more structure needs to be imposed on the production function. We will consider an F that is a constant-elasticity of substitution production function of the following form:

$$Y = A\Big[ag(N,K)^{\rho} + (1-a)E^{\rho}\Big]^{\frac{1}{\rho}}. \tag{9}$$

Here, $g(N,K)$ is a CRS production function that combines capital and labor; $a \in (0,1)$ is a parameter that governs the "importance" of capital and labor relative to natural resources; and ρ governs the elasticity of substitution between $g(N,K)$ and natural resources. We follow Hassler et al. (2021) by separating out E to focus on the elasticity of substitution for this input of interest. The case $\rho \to 0$ corresponds to the Cobb-Douglas production function, with $(1-a)$ being the income-share of natural resources.

Given this functional form, it holds that

$$\frac{\phi_E}{1-\phi_E} = \frac{\frac{\partial Y}{\partial E}E}{\frac{\partial Y}{\partial g(N,K)}g(N,K)} = \frac{(1-a)E^{\rho}}{a(g(N,K))^{\rho}} = \frac{(1-a)}{a}\left(\frac{E}{g(N,K)}\right)^{\rho}$$

Taking logs, we have that

$$\ln\left(\frac{\phi_E}{1-\phi_E}\right) = \ln\left(\frac{1-a}{a}\right) + \rho\ln\left(\frac{E}{g(N,K)}\right)$$

Using Δ to denote a change over time ($\Delta x = x_t - x_0$ for some variable x), it must hold that

$$\Delta \ln\left(\frac{\phi_E}{1-\phi_E}\right) = \rho \Delta \ln\left(\frac{E}{g(N,K)}\right) \tag{10}$$

This equation allows us to easily summarize nine relevant cases, depicted in Table 1. The middle row and middle column represent the cases where the observed values for ϕ_E are a good approximation for its long-run value. These are cases where the input-shares remain roughly unchanged (middle row) or the function is roughly Cobb-Douglas in E (middle column).

Table 1: Sign of direct productivity effect as a function of parameters

	$\rho < 0$	$\rho \approx 0$	$\rho > 0$
$\bar{E}/g(\bar{N},\bar{K})$ smaller than observed	$-$	$+$	$+$
$\bar{E}/g(\bar{N},\bar{K})$ in observed range	$+$	$+$	$+$
$\bar{E}/g(\bar{N},\bar{K})$ larger than observed	$+$	$+$	$-$

Notes: General cases for the long-run relationship between population and per capita income in a CES production function where the relative inputs evolve over time. Row definitions are imprecise: "smaller" and "larger" are implicitly defined as large enough divergences from observed ratios to plausibly flip the sign of interest (which will depend on a combination of the growth in $\bar{E}/g(\bar{N},\bar{K})$ and the size of ρ). If, for example, ρ is just less than 0, and the ratios are just outside observed ranges, the sign would not necessarily be negative (this is formalized in Figure 4).

Outside of those cases, there are two (top right, bottom left) where the income-share of natural resources falls in the long run and two (top left, bottom right) where the income-share grows. The income-share of natural resources is already small, so the qualitative takeaways will be unchanged if it shrinks further. The more interesting cases are the top left—where inputs are complementary and the input-share of E shrinks in the long run—and the bottom right—where inputs are substitutable and the input-share of E grows in the long run.[16] These are cases where the current low values of ϕ_E could be potentially misleading; we therefore need to determine which of these nine cases appears most likely based on the evolution of these respective terms over the medium term.

There is little we can say with confidence about the future of long-run input-shares. Many

[16]That ρ is important is also highlighted by Wilde (2017).

believe that climate change or medium-term unsustainable withdrawal of resources implies that long-run levels of E will decline (see e.g., Dasgupta, 2021; Henderson et al., 2022). Additionally, continued growth in the capital stock would decrease the input-share of E. On the other hand, low fertility may imply that future population sizes are much smaller than today (see e.g., Jones, 2022; Geruso and Spears, 2023). This would directly reduce N and indirectly reduce K. Which of these forces will dominate is a question of forecasts that are beyond the scope of this paper.

We can instead make progress by calibrating ρ. Equation 10 implies a value for ρ once the growth in income-shares is compared to growth in factor inputs. Figure 1 displays no pervasive trend in the income-share of natural resources. Based on this, we set

$$\Delta \ln \left(\frac{\phi_E}{1-\phi_E} \right) = 0. \tag{11}$$

To calibrate ρ based on (10), it is also necessary that we have established that factor input-shares have changed, i.e., that $\Delta \ln \left(\frac{E}{g(N,K)} \right) \neq 0$. Otherwise, the parameter ρ is unidentified—in any constant returns to scale production function, when inputs increase by the same proportion then income-shares remain unchanged. Formally, Figure 3 documents that:

$$\Delta \ln(E) < \Delta \ln(N) < \Delta \ln(K).$$

As g is a constant-returns to scale production function, it holds that

$$\Delta \ln(E) < \Delta \ln(N) \leq \Delta \ln(g(N,K)) \leq \Delta \ln(K)$$

and hence

$$\Delta \ln \left(\frac{E}{g(N,K)} \right) = \Delta \ln(E) - \Delta \ln(g(N,K)) \neq 0. \tag{12}$$

By (10), (11) and (12) it follows that $\rho = 0$, and hence

$$Y = (g(N, K))^a E^{1-a} \text{ and } \phi_E = 1 - a.$$

In particular, the income-share ϕ_E is independent of factor inputs in this calibration. We can thus reasonably take the long-run value of ϕ_E to be the current income-share of natural resources, at least for long-run population sizes near or within the historical experience.[17]

Returning to Equation 8 we have the following approximate values for the direct effect:

$$\frac{\partial \ln \bar{y}}{\partial \ln \bar{N}} = \underbrace{\frac{\lambda}{\beta}}_{\approx 0.3} - \underbrace{\phi_E}_{\approx 0.03} + \phi_K \frac{\partial \ln \bar{k}}{\partial \ln \bar{N}}.$$

Given that ϕ_K—capital's share of income—is on the order of 0.4, $\frac{\partial \ln \bar{k}}{\partial \ln \bar{N}}$ would need to be approximately -0.66 to nullify the direct effect. That is, capital per worker would need to fall in response to this positive productivity effect. Equivalently, $\frac{\partial \ln \bar{K}}{\partial \ln \bar{N}}$ would need to be around 0.33. Such a small elasticity of aggregate capital with respect to population size, and the direct productivity benefits it brings, is unlikely. In a theoretical benchmark with fixed savings rates (e.g., a Cobb-Douglas environment), capital per worker would rise with productivity. The same share of a larger per capita income is saved. Likewise, the aggregate data displays roughly constant capital-output ratios, which imply an increase in capital per worker as productivity rises (Kaldor, 1961). Indeed, our own Figure 3 indicates that, historically, capital-worker ratios have grown as populations and productivity have increased. Therefore, for the remainder of the paper, we make the simple and conservative assumption that $\frac{\partial \ln \bar{K}}{\partial \ln \bar{N}} = 1$, so that the indirect channel is zero.

[17] We include this caveat because careful readers will recognize that if the production function is truly Cobb-Douglas over the entire domain, then the income-share of natural resources will *never* rise. We have little confidence in statements about what happens at much larger population sizes—if populations grow substantially and the income-share begins rising, this calibration can be re-examined—so we do not wish to push our results far outside of historical observations.

Combinations of off-the-shelf model ingredients and existing empirical estimates suggest that the elasticity of long-run income per capita with respect to population is positive.

4 Extensions and Robustness

Here we consider extensions and generalizations to the model that could plausibly strengthen the Malthusian channel. First, we examine how a range of non-Cobb-Douglas elasticity of substitutions map long-run uses of E into long-run values of ϕ_E, to determine whether our takeaway is likely to hold if the Cobb-Douglas calibration is mistaken or misleading. Second, we explore the possibility of factor-augmenting technological change—in particular, technological progress biased away from natural resources. Third, we formalize the inclusion of non-rival *benefits* of ecosystem services, such that there is an additional channel through which the environment contributes to income. Finally, we consider a case where humans reduce the availability of aggregate natural resources, rather than E being independent of the population size. None of these model relaxations appear significant enough to plausibly overturn the main result.

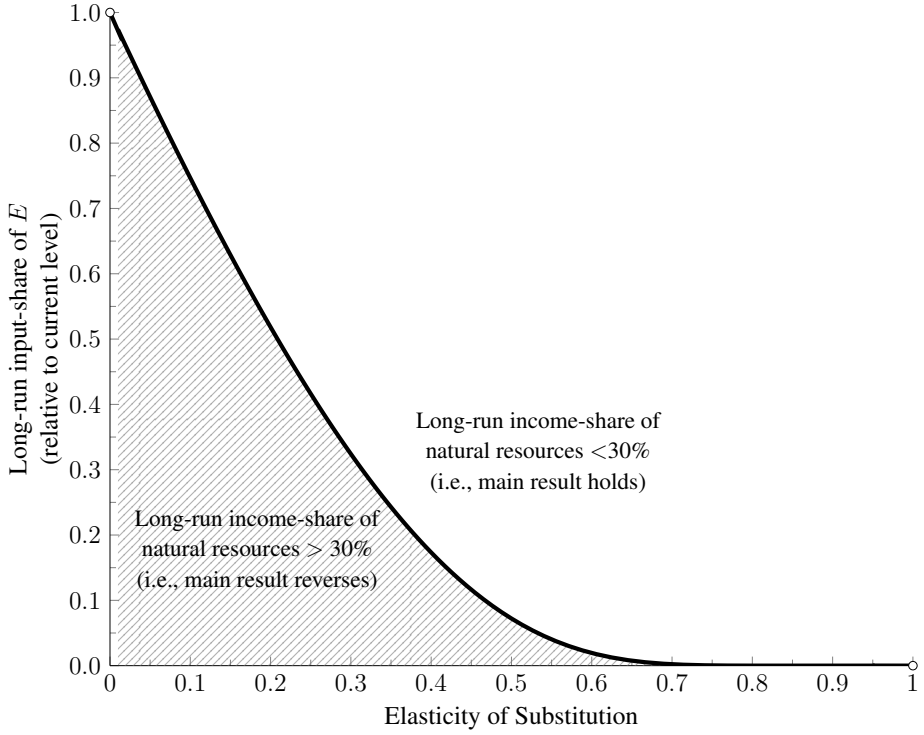
4.1 Long-run complementarities between E and human-provided inputs

While the patterns in Figures 1 and 3 provide strong evidence for a baseline case where environmental inputs do not have important complementarities with human-provided inputs, we can relax this assumption to study the case where such complementarities exist. This would imply that if the input-share of E fell—e.g., if natural resources are currently used at an unsustainable level—then the income-share of E would rise. Accordingly, current levels of ϕ_E would underestimate the long-run value relevant for our calibration.

Recall that our main result indicates that ϕ_E would need to rise by roughly ten times to be competitive with the elasticity of TFP with respect to population (from about 3% to about 30%). Figure 4 maps the combinations of production function parameterizations and long-run levels of E

that do and do not cause a ten-fold increase in ϕ_E. The x-axis is the elasticity of substitution, $\sigma = \frac{1}{1-\rho}$, from the earlier CES production function. The y-axis is the ratio of $\frac{\bar{E}}{g(\bar{N},\bar{K})}$ to $\frac{E_{2024}}{g(N_{2024},K_{2024})}$; that is, it is the future input-share of E as a fraction of its current input-share. The shaded region below the curve is the region of the parameter space that would cause the long-run income-share of E to rise near or above 30% of GDP, potentially challenging our main takeaways.

Figure 4: Elasticity of substitution must be very far from Cobb-Douglas for $\phi_E \to 30\%$



Notes: Combinations of elasticity of substitutions, $\sigma = \frac{1}{1-\rho}$, and changes in the input-share of E that could plausibly reverse our result. For example, if the long-run use of E is reduced such that its input-share falls by 50% relative to today (i.e., y-axis value of 0.5), then the EoS (x-axis value) would need to less than or equal to roughly 0.21 for this decline in E to increase ϕ_E to 30% or larger. The (0,1) point reflects that under a Leontief production function *any* reduction in E raises its income-share above 30%; (1,0) reflects that for Cobb-Douglas *no* change in E raises ϕ_E.

The two extreme points are intuitive. The lowest the elasticity of substitution can be is zero—the case of a Leontief production function. With a Leontief production function, the only way that each factor could have a positive marginal product is if they are supplied in their exact Leontief proportions. Any decline in the relative input-share of E would then raise its income-share to

100% of GDP. This is why the curve is arbitrarily close to one (the point where factor shares remain exactly at their current level) as the elasticity of substitution goes to zero. Conversely, if $\sigma \to 1$, the required decline in factor shares needs to be arbitrarily large to generate a ten-fold increase in ϕ_E. This is the Cobb-Douglas case where E's income-share is unchanging, so that no decline is large enough to increase its income-share to ten times its current level.

Between these endpoints the function is convex. For example, suppose our current use of resources is twice what would be sustainable.[18] That would imply—holding fixed $g(N, K)$—that the input-share of E will need to fall to about half its current level in a sustainable steady state. For this decline to raise the income-share of E to 30% or more, σ would need to be less than 0.25. Needless to say, this is far from the Cobb-Douglas case we believe is best supported by the last half-century of aggregate data. More generally, for the main conclusion to be reversed, either the input-share of E must fall dramatically (e.g., to less than 10% of its current level) or the elasticity of substitution must be *much* smaller than is implied by historical trends.

4.2 Factor-augmenting technical change

A more realistic version of this model might have factor-augmenting technological change, which could strengthen or weaken the results depending on the direction of augmentation. For example, endogenous directed technological change (e.g., Boserup, 1965; Acemoglu et al., 2012) may be an important mechanism for pinning down the quantitative elasticities. In a series of recent papers, Hassler et al. (2021, 2023) argue that this is a necessary mechanism to explain patterns related to fossil fuel use. In particular, Hassler et al. (2021) demonstrate that high degrees of complementarity over short time horizons can be consistent with flat long-run factor shares if technological change is E-augmenting (whether exogenously or endogenously). Likewise, Kruse-Andersen (2023) in-

[18] For reference, this is a more conservative assumption than employed by Dasgupta et al. (2021) who cite ecologists claims' that we are using the resources of 1.7 Earths (or that our long-run use needs to fall to about 60% of its current level).

cludes this as a possible channel by which population increases may not have any environmental related drawbacks in the case of climate change.

Our omission of this channel has been for conservatism: the main result demonstrates that even without directing technological change towards increasingly tight resource constraints, the marginal relationship between population and per capita income is estimated to be positive. However, it is straightforward to generalize the main finding to an arbitrary degree of factor-augmenting technological improvements, including cases where non-directed change only improves the efficiency of human-provided inputs (N, K).

To show this, we must generalize the original production function one degree further. Let A be the level of TFP that is relevant for human-provided inputs N and K; A_E is the level of TFP relevant for environmental inputs.

$$
\begin{aligned}
&Y = F(AN, AK, A_E E) \Rightarrow \\
&y = \frac{AF\Big(N, K, a_E E\Big)}{N}, \quad \text{where } \; a_e \equiv \frac{A_E}{A}
\end{aligned}
$$

The corresponding elasticity (focusing only on the direct effect) between population and per capita income becomes the following, where again ϕ_E is the share of income going to the fixed factors.

$$
\frac{\partial \ln(\bar{y})}{\partial \ln(\bar{N})} = \frac{\partial \ln(\bar{A})}{\partial \ln(\bar{N})} - \Big(1 - \frac{\partial \ln(\bar{a_E})}{\partial \ln(\bar{N})}\Big)\phi_E \tag{13}
$$

The relationship is intuitive and informative. If the elasticity of $\bar{a}_E$ with respect to $\bar{N}$ is one, then it is as if there has been no (relative) change in the fixed factor and the Malthusian channel drops from the expression: increasing $\bar{N}$ by 1% decreases $\frac{\bar{E}}{\bar{N}}$ by 1% (by construction), but increases $\bar{a}_E$ by 1% (by assumption), leaving the product of these terms unchanged. In the original case of Hicks-neutral technological change, $\bar{a}_E$ is unchanging, the elasticity equals zero, and this collapses to the baseline case in Equation 8.

In the case where technological progress operates disproportionately on N and K, $\bar{a}_E$ will decrease. This compounds the losses from $\frac{E}{N}$ decreasing. However, there is a limit on how severe the log-decline in $\bar{a}_E$ can be. This fraction can only decrease in proportion to the increase in $\bar{A}$ (since this is the denominator of $\bar{a}_E$). Formally, $\frac{\partial \ln(\bar{a}_E)}{\partial \ln(\bar{N})} \geq -\frac{\partial \ln(\bar{A})}{\partial \ln(\bar{N})}$. In this extreme case where *all* technological progress is N-,K-augmenting, we can write Equation (13) as:

$$(1 - \phi_E)\frac{\partial \ln(\bar{A})}{\partial \ln(\bar{N})} - \phi_E. \tag{14}$$

The TFP-elasticity is now mitigated by $\phi_E\%$ before being compared to ϕ_E. Under the baseline values for these parameters this will not be quantitatively meaningful. Rather than comparing roughly 0.3 (λ/β) to 0.03 (ϕ_E), the comparison would be 0.29 to 0.03. In short, variants of factor-augmenting technical change can not themselves overturn the results, even under the most pessimistic version of this assumption.

4.3 Non-rival benefits from natural resources

Alongside the rival ecosystem services that earn profits—and are conceptually captured in the income-share terms—there may also be non-rival benefits of nature that increase with its stock. Consider a generalized production function building from Dasgupta (2021).

$$Y = AB^{\xi}F(N, K, E) \tag{15}$$

Here B captures the general health (or stock) of the nature, what Dasgupta (2021) calls the biosphere. It performs regenerative services—such as filtering H_2O throughout the water cycle—and helps promote innovation and learning—such as plants in the Amazon that provide the ideas for new pharmaceuticals. If a broader view of Y beyond measured GDP is taken, B might be considered to contribute to Y via non-use values.

What we have called E throughout the paper is the flow of ecosystem services being drawn from B. Previously the stock of B was only indirectly useful (in that it determines how much E can be drawn), so was left undiscussed. These non-rival benefits introduce a channel by which the stock itself is directly beneficial. This indirect channel changes the optimal level of E, as the flow of E will determine the level of B (see Appendix B). Previously, it would have been optimal to erode the biosphere until it reached its peak growth rate.[19] When B enters directly we have a competing incentive to keep B larger than where it reaches its regenerative peak.

However, our sustainable solution assumes that the withdrawal rate in steady state is independent of the population level (recall the discussion in Section 2.3). In this case, the value of B is independent of population, as long as the rate of E is unaffected by population.[20] Our analysis therefore is unaffected by the inclusion of a non-rival benefit of ecosystem services.

4.4 Human activity erodes E

Here we relax the assumption that E is fixed with respect to population. For example, if land that produces important ecosystem services is converted into urban land, that would reduce the aggregate availability of E per period. This would serve as a channel that compounds the initial decline in per capita natural resources. Previously, the per capita decline from a 1% increase in population size was a 1%. The denominator of E/N increased by 1%; the numerator was fixed. In this conceptual extension, the numerator declines as well. Formally, the elasticity of the direct productivity benefits becomes:

$$\frac{\partial \ln \bar{y}}{\partial \ln \bar{N}} = \lambda/\beta - \Big(1 - \frac{\partial \ln \bar{E}}{\partial \ln \bar{N}}\Big)\phi_E. \tag{16}$$

[19] For example, in an ocean with fish populations, to keep these populations at their carrying capacity in steady state implies zero withdrawal, because there is zero-growth when they are at their carrying capacity.

[20] In general, it could be the case that a larger population finds it optimal to move closer to the maximum sustainable withdrawal rate—e.g., if the marginal utility of E increases with population size—but this would be an even better outcome for the larger population than our assumed solution where $\bar{E}$ is fixed.

In the baseline case where these is no relationship between $\bar{E}$ and $\bar{N}$, this new parenthetical term reduces to one. However, suppose this elasticity is itself negative one (i.e., a 1% increase in $\bar{N}$ reduces available aggregate resource withdrawals by 1%). Then, the negative Malthusian channel is doubled. Each 1% increase in population now reduces per capita natural resources by 2%; 1% from $\bar{N}$ increasing, 1% from $\bar{E}$ decreasing.

It would need to be the case that a 1% increase in population reduces $\bar{E}$ by nearly 10% to increase its magnitude ten-fold. This is not impossible given the unintuitive dynamics renewable resource problems can exhibit, but it seems unlikely unless the system is near a tipping point. We know of no current evidence that would indicate this should be the baseline assumption, so we do not take this channel to be one that reverses our main conclusion at present. However, this result indicates that further research on the elasticity between human population size and the availability of natural resources will be useful for numerically calibrating the size of this relationship in future work.

5 Conclusion

The human population is projected to peak—and likely begin shrinking—during the lifetimes of people alive today. Whether it would be better to stabilize at higher population sizes depends on many things: transition costs related to medium-term growth or decline (Galor, 2011), the social value of additional existences (Klenow et al., 2023), etc. This analysis focuses on one important aspect of this question: the effect on long-run per capita incomes.

Using a model that captures the most frequently discussed competing forces—that (i) nature imposes aggregate constraints and (ii) larger populations produce more non-rival goods—a simple analytical relationship between long-run population sizes and per capita incomes arises. Calibrating the relevant parameters suggests this is relationship is positive. Insofar as non-rival goods and fixed natural resources are the primary channels by which populations influence per capita incomes,

this result provides reason to believe that a future that stabilizes at a larger global population will be richer per capita than a world that stabilizes at a smaller population.

References

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous (2012) "The environment and directed technical change," *American Economic Review*, 102 (1), 131–166.

Aschenbrenner, Leopold and Philip Trammell (2024) "Existential Risk and Growth," *Global Priorities Institute Working Paper*, Available at: https://philiptrammell.com/static/Existential_Risk_and_Growth.pdf.

Ashraf, Quamrul H, David N Weil, and Joshua Wilde (2013) "The effect of fertility reduction on economic growth," *Population and Development Review*, 39 (1), 97–130.

Bloom, Nicholas, Charles I Jones, John Van Reenen, and Michael Webb (2020) "Are Ideas Getting Harder to Find?" *American Economic Review*, 110 (4), 1104–44.

Boserup, Esther (1965) *The conditions of agricultural growth: the economics of agrarian change under population pressure*: London: Allen & Unwin.

Bradshaw, Corey JA and Barry W Brook (2014) "Human population reduction is not a quick fix for environmental problems," *Proceedings of the National Academy of Sciences*, 111 (46), 16610–16615.

Bretschger, Lucas (2020) "Malthus in the light of climate change," *European Economic Review*, 127, 103477.

Budolfson, Mark, Mike Geruso, Kevin Kuruc, Dean Spears, and Sangita Vyas (2023) "Is Less Really More? Comparing the climate and productivity effects of a shrinking population," *Population Wellbeing Initiative Working Paper*, Available at: https://nbviewer.org/github/kevinkuruc/Papers_for_Download/blob/main/Stabilization.pdf.

Caselli, Francesco and James Feyrer (2007) "The Marginal Product of Capital," *The Quarterly Journal of Economics*, 122 (2), 535–568.

Dasgupta, Partha (2021) *The Economics of Biodiversity: The Dasgupta Review*: London: HM Treasury.

Dasgupta, Partha, Aisha Dasgupta, and Scott Barrett (2021) "Population, Ecological Footprint and the Sustainable Development Goals," *Environmental and Resource Economics*, 1–17.

Dasgupta, Susmita, Benoit Laplante, Hua Wang, and David Wheeler (2002) "Confronting the environmental Kuznets curve," *Journal of economic perspectives*, 16 (1), 147–168.

Ekerdt, Lorenz and Kai-Jie Wu (2023) "Self-Selection and the Diminishing Returns to Research," *Working Paper*, Available at https://sites.google.com/view/lorenz-ekerdt/research.

Feenstra, Robert C, Robert Inklaar, and Marcel P Timmer (2015) "The next generation of the Penn World Table," *American Economic Review*, 105 (10), 3150–3182.

——— (2023) "Penn World Tables (10.01)," Accessed on 7/15/23. Available at: https://www.rug.nl/ggdc/productivity/pwt/?lang=en.

Galor, Oded (2011) *Unified Growth Theory*: Princeton University Press.

Galor, Oded and David N Weil (2000) "Population, Technology, and Growth: From Malthusian Stagnation to the Demographic Transition and Beyond," *American Economic Review*, 90 (4), 806–828.

Gerlagh, Reyer (2023) "Climate, technology, family size; on the crossroad between two ultimate externalities," *European Economic Review*, 152, 104376.

Geruso, Mike and Dean Spears (2023) "With a whimper: Depopulation and Longtermism," in Barrett, Jacob and Hilary Greaves eds. *Essays on Longtermism*: Oxford University Press, Available at: https://sites.utexas.edu/pwi/files/2023/01/Geruso_Spears_Longtermism_and_Population_OUP.pdf.

Greaves, Hilary (2019) "Climate change and optimum population," *The Monist*, 102 (1), 42–65.

Hall, Robert E and Charles I Jones (2007) "The value of life and the rise in health spending," *The Quarterly Journal of Economics*, 122 (1), 39–72.

Hassler, John, Per Krusell, and Conny Olovsson (2021) "Directed Technical Change as a Response to Natural Resource Scarcity," *Journal of Political Economy*, 129 (11), 3039–3072.

——— (2023) "Finite Resources and the World Economy," *Journal of International Economics*.

Henderson, J Vernon, Bo Yeon Jang, Adam Storeygard, and David N. Weil (2022) "Climate Change, Population Growth, and Population Pressure," *Working Paper*.

Hopenhayn, Hugo, Julian Neira, and Rish Singhania (2022) "From population growth to firm demographics: Implications for concentration, entrepreneurship and the labor share," *Econometrica*, 90 (4), 1879–1914.

Jones, Charles I (1995) "R&D-Based Models of Economics Growth," *Journal of Political Economy*, 103 (4).

——— (2001) "Was an Industrial Revolution Inevitable? Economic Growth Over the Very Long Run," *The B.E. Journal of Macroeconomics*, 1 (2).

——— (2005) "Growth and Ideas," in Aghion, Philippe and Steven N Durlauf eds. *Handbook of Economic Growth*, 1B, 1063–1111: Elsevier.

——— (2016) "Life and growth," *Journal of Political Economy*, 124 (2), 539–578.

——— (2022) "The End of Economic Growth? Unintended Consequences of a Declining Population," *American Economic Review*, 112 (11), 3489–3527.

Jones, Charles I and Paul M Romer (2010) "The New Kaldor Facts: Ideas, Institutions, Population, and Human Capital," *American Economic Journal: Macroeconomics*, 2 (1), 224–45.

Kaldor, Nicholas (1961) "Capital Accumulation and Economic Growth," in *The Theory of Capital*, 177–222, The International Economic Association.

Karahan, Fatih, Benjamin Pugsley, and Ayşegül Şahin (Forthcoming) "Demographic origins of the startup deficit," *American Economic Review*.

Karra, Mahesh, David Canning, and Joshua Wilde (2017) "The effect of fertility decline on economic growth in Africa: A macrosimulation model," *Population and Development Review*, 43, 237–263.

Klenow, Peter, Charles I Jones, Mark Bils, and Mohamad Adhami (2023) "Population and Welfare: The Greatest Good for the Greatest Number," Available at: http://web.stanford.edu/~chadj/popwelfare.pdf.

Kremer, Michael (1993) "Population Growth and Technological Change: One Million B.C. to 1990," *The Quarterly Journal of Economics*, 108 (3), 681–716.

Kruse-Andersen, Peter (2023) "Directed technical change, environmental sustainability, and population growth," *Journal of Environmental Economics and Management*, 122.

Kruse-Andersen, Peter K (2023) "Testing R&D-Based Endogenous Growth Models," *Oxford Bulletin of Economics and Statistics*, 85 (5), 1083–1110.

Lawson, Nicholas and Dean Spears (2018) "Optimal population and exhaustible resource constraints," *Journal of Population Economics*, 31 (1), 295–335.

Lutz, Wolfgang (2023) "Population decline will likely become a global trend and benefit long-term human wellbeing," *Vienna Yearbook of Population Research*, 21.

Monge-Naranjo, Alexander, Juan M Sánchez, and Raül Santaeulalia-Llopis (2019) "Natural Resources and Global Misallocation," *American Economic Journal: Macroeconomics*, 11 (2), 79–126.

Peretto, Pietro F and Simone Valente (2015) "Growth on a finite planet: resources, technology and population in the long run," *Journal of Economic Growth*, 20 (3), 305–331.

Peters, Michael (2022) "Market Size and Spatial Growth – Evidence from Germany's Post-War Population Expulsions," *Econometrica*, 90 (5), 2357–2396.

Peters, Michael and Conor Walsh (2021) "Population Growth and Firm Dynamics," *National Bureau of Economic Research Working Paper*, No. 29424.

Ritchie, Hannah (2022) "Did humans cause the Quaternary megafauna extinction?" *Our World in Data*, https://ourworldindata.org/quaternary-megafauna-extinction.

——— (2024) *Not the end of the world: How we can be the first generation to build a sustainable planet*: Hachette UK.

Romer, Paul M (1990) "Endogenous Technological Change," *Journal of Political Economy*, 98 (5, Part 2), S71–S102.

Spears, Dean, Sangita Vyas, Gage Weston, and Michael Geruso (2024) "Long-term population projections: Scenarios of low or rebounding fertility," *Plos one*, 19 (4), e0298190.

Terry, Stephen, Thomas Chaney, Konrad Burchardi, Lisa Tarquinio, and Tarek Hassan (2024) "Immigration, Innovation, and Growth," *NBER Working Paper*, Available at: https://websites.umich.edu/~sjterry/PaperIIG.pdf.

Weil, David N and Joshua Wilde (2009) "How Relevant is Malthus for Economic Development Today?" *American Economic Review: Papers & Proceedings*, 99 (2), 255–60.

Wilde, Joshua (2017) "How substitutable are fixed factors in production? Evidence from pre-industrial England," *Working Paper*, Available at: https://www.usf.edu/arts-sciences/departments/economics/documents/substitutability_1.pdf.

Appendix

A Putting sub-components of E directly into F does not change the analysis

In this appendix we demonstrate that where in the process we aggregate is inconsequential. To see this, rewrite the aggregate production function as:

$$\tilde{F}(N, K, e_1, e_2, e_3, \ldots, e_Z) \tag{A1}$$

These es are fixed, just as E was fixed, in the main analysis. So, the log-derivative of this production function with respect to N continues to omit any of the es. In fact, it is exactly unchanged:

$$\frac{\partial \ln(\bar{y})}{\partial \ln(\bar{N})} = \frac{\partial \ln(\bar{A})}{\partial \ln(\bar{N})} + \frac{\partial \ln(\tilde{F}(\bar{N}, \bar{K}, \bar{e_1}, \ldots, \bar{e_Z}))}{\partial \ln(\bar{N})} + \frac{\partial \ln(\tilde{F}(\bar{N}, \bar{K}, \bar{e_1}, \ldots, \bar{e_Z}))}{\partial \ln(\bar{K})} \frac{\partial \ln(\bar{K})}{\partial \ln(\bar{N})} - \frac{\partial \ln(\bar{N})}{\partial \ln(\bar{N})}.$$

Recall that ϕ_E entered because it was the payments *not* going to capital or labor: we used the fact that $\phi_E = 1 - \phi_N - \phi_K$ in a CRS production function. Here, there are many e_is that share the income not going to capital or labor. But we still know that the sum of payments to these resources must be what is not paid to capital and labor; i.e., $\sum \phi_{e_i} = 1 - \phi_N - \phi_K$.

This implies that our sufficient statistic can simply be rewritten as:

$$\frac{\frac{\lambda}{\beta} - \sum \phi_{e_i}}{1 - \phi_K}.$$

This changes nothing—the empirical counterpart to the model already treated ϕ_E as being $\sum \phi_{e_i}$.

The reason we can introduce arbitrary degrees of complementarities between any natural resources with capital or labor without affecting the analysis is that we are studying a local derivative, which can be approximated by a separable, linear function. So, it may well be the case that at pop-

ulation sizes much larger than our own, a complementarity starts to bite. This would be reflected in ϕ_E rising, and our sufficient statistic would (correctly) respond to this change. Of course, this makes our methodology less powerful for out-of-sample conjectures. Nonetheless, our results are useful for comparing population sizes within the historical experience. Given that population growth appears to be ending this is useful for the questions societies will grapple with over the medium-term.

B Details of renewable resource problem

As noted in the main text, Dasgupta (2021) argues that the entire biosphere, B, can be roughly conceptualized as renewable resource problem. If nature is undisturbed by human activity, most resources will regenerate.

Steady state solutions to such problems are characterized by withdrawals of constant ecosystem services, $\bar{E}$, exactly equal to the amount of regeneration the renewable resource produces (which will be a function of its stock). Formally, assume that the biosphere has a regeneration function, $R(B)$, as in Equation A2, taken from Dasgupta (2021).

$$R(B) = rB\left[1 - \frac{B}{Z}\right]\left[\frac{B-T}{Z}\right] \tag{A2}$$

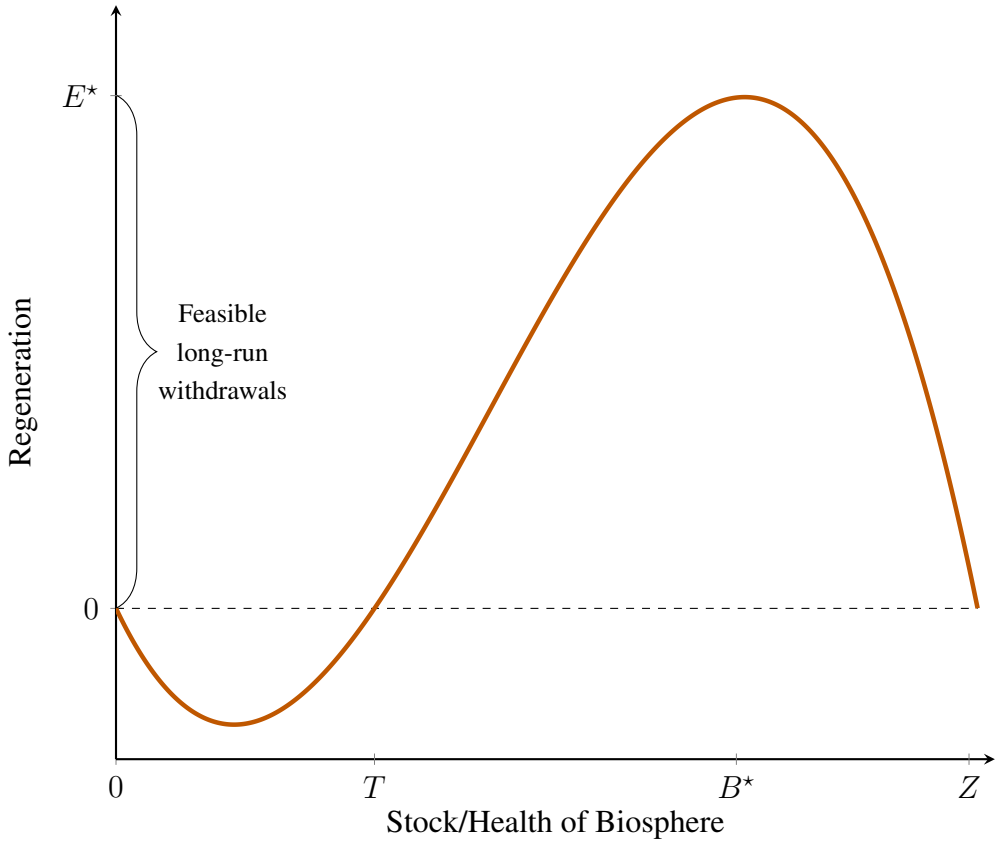
B is the stock of biosphere as it relates to human production/consumption and $R(B)$ is the amount of regeneration. Z is the carrying capacity of this renewable resource—where the natural world would converge with minimal human interference. T is a "tipping point"—should we degrade the environment below T, regeneration becomes negative and the system collapses towards $B = 0$. r is the rate of regeneration in the absence of a tipping point or carrying capacity. The law of motion on B is then governed by the difference between regeneration, $R(B)$, and the amount of ecosystem services drawn for human production/consumption, E.

$$\dot{B}(B) = R(B) - E \tag{A3}$$

In a steady state solution $\bar{E} = R(\bar{B})$. Notice that any level of $B \in [T, Z]$ can be consistent with a steady state outcome. According to Equation A3, the steady state level of B pins down the steady state level of E. In this formulation, there is a unique $B^\star$ that optimizes the regeneration rate and hence the value of $\bar{E}$.[21] Income is increasing in E, so the optimal solution in the baseline model is

[21] Notice that it is *not* the maximal B. Once B is at it's carrying capacity there is not net growth to drawn down, by

Figure A1: Regeneration rate with tipping point



to manage B such that the peak of regeneration is reached. In the model of Section 4.3, there is an additional benefit from a higher level of B. Generally, this will increase $B^\star$ beyond the level that maximizes only $\bar{E}$.

Nothing about this solution depends on the population size. The population size has no effect on regeneration rates, conditional on E, and it would be similarly inefficient for any population size to not manage B at the level that maximizes $R(B)$. Large and small populations alike face the challenge of intertemporal externalities. We abstract from the underlying details of this resource management problem and import a steady state solution, $\bar{E}$, of this independent subproblem into the aggregate production function.

definition.

C Exhaustible resources and population size

Some observers are worried about the stock of exhaustible resources rather than the long-run withdrawal of renewable resources. Fossil fuels are foremost among these concerns. Exhaustible resources amount to a classic "cake-eating" problem where we can denote the initial stock as $\mathbb{E}$. For analytical simplicity, and consistent with the flat trend in factor shares here and in Hassler et al. (2021), assume a Cobb-Douglas specification for this exercise. A planner maximizes per capita income that is subject to diminishing returns through a CRRA parameter σ.

$$\max \sum_{t=0}^{\infty} \delta^t \frac{1}{1-\sigma} y_t^{1-\sigma} \tag{A4}$$

$$\text{with } \; y_t = A_t \Big(\frac{E_t}{N}\Big)^{1-a} \tag{A5}$$

$$\text{and } \; \sum_{t=0}^{\infty} E_t = \mathbb{E} \tag{A6}$$

Here the discount rate is denoted δ, which now matters as we are focused on environments without a steady state level of consumption to analyze. We have already established that $A \rightarrow \left(\frac{\alpha N^{\lambda}}{\delta_A}\right)^{\frac{1}{\beta}}$ independent of the environmental side of the model. For simplicity then assume A reaches this level reasonably quickly and is therefore fixed for the long-run study of this model. With a fixed A, N it can be easily shown that the solution is characterized by $\frac{E_t}{E_{t+1}} = \delta^{\frac{1}{1+\sigma a - a}}$. The level of environmental withdrawals remains independent of the population size as in Bretschger (2020).

Therefore, since withdrawal levels are independent of N, each period the effect of a 1% increase in N is a 1% decline in $\frac{E}{N}$ and a λ/β% increase in A. This is the same as in the long-run steady state with a constant E, despite this being a case where long-run per capita incomes converge to zero (as A hits an upper limit, but E converges to zero). In other words, our results are not contingent on innovation growth ensuring that per capita incomes stay indefinitely high, even as resource use necessarily trends to zero. In this particular case, incomes fall regardless of the population size, but less quickly with a larger population.

If instead one conceptualizes non-renewable resources being used at a constant per capita rate, but also being in fixed aggregate supply, what is mechanically constrained is the number of people who can ever exist. It is true that humanity can persist for longer (temporally) at a smaller size, but the normative relevance of this is not obvious if the total number of individual existences is fixed (see Lawson and Spears (2018); Greaves (2019) for further discussion).

D Zero depreciation of knowledge retains main result

Suppose we took δ_A, the depreciation rate on knowledge, to be zero. Our law of motion then reduces to the standard form in the semi-endogenous growth literature.

$$\frac{\dot{A}}{A} = \alpha N^{\lambda} A^{-\beta}$$

It is well-known that the long-run level of productivity approaches infinity, regardless of population size, under this functional form (Jones, 1995). So, per capita incomes go to infinity as well, making the analysis of how these incomes differ across populations less straightforward.

However, it is possible to find the difference in incomes in each finite time-period and take the limit of that term, for an analogous expression.[22] It turns out that the marginal analysis is unchanged. So, our result holds not only for any finite δ_A as we demonstrate in the main text, but also for $\delta_A = 0$, as is more commonly assumed.

What we need to retain is that the elasticity of A with respect to N approaches λ/β with $\delta_A = 0$ (because the natural resource negative elasticity is unchanged). Integrating the law of motion for ideas gives the following expression for A_t:

$$A_t = \left(\alpha\beta \int_0^t N(\tau)^{\lambda} d\tau + A_0^{\beta} \right)^{\frac{1}{\beta}}. \tag{A7}$$

In the case of simplified cases of constant populations that we study, $N(\tau) = \bar{N} \forall\ \tau$. This then collapses to:

$$A_t = \left(\alpha\beta \bar{N}^{\lambda} t + A_0^{\beta} \right)^{\frac{1}{\beta}}. \tag{A8}$$

[22] We thank Bernardo Ribeiro for sending us a note pointing this out with a corresponding derivation.

This implies that the elasticity is:

$$\begin{aligned}\frac{\partial \ln A_t}{\partial \ln \bar{N}} =& \frac{1}{\frac{\beta}{\lambda} + \frac{A_0^{\beta}}{\alpha \lambda N^{\lambda} t}} \Rightarrow \\ \lim_{t \to \infty} \frac{\partial \ln A_t}{\partial \ln \bar{N}} =& \frac{\lambda}{\beta}.\end{aligned}$$

E Data Appendix

E.1 Resource rents and useage

There exists public estimates on natural resource shares from the World Bank and the US Department of Agriculture. This leaves us only with the task of aggregating existing estimates. Figures on rents as a percent of GDP for timber/forests, minerals, coal, oil, and natural gas come from the "Adjusted Net Savings" dataset (updated 9/23/2022) from the World Bank's 2021 "Changing Wealth of Nations" report. These figures are provided on an annual basis and beginning in 1970. The dataset has missing years at the country-level but makes estimates at the global level in each year.

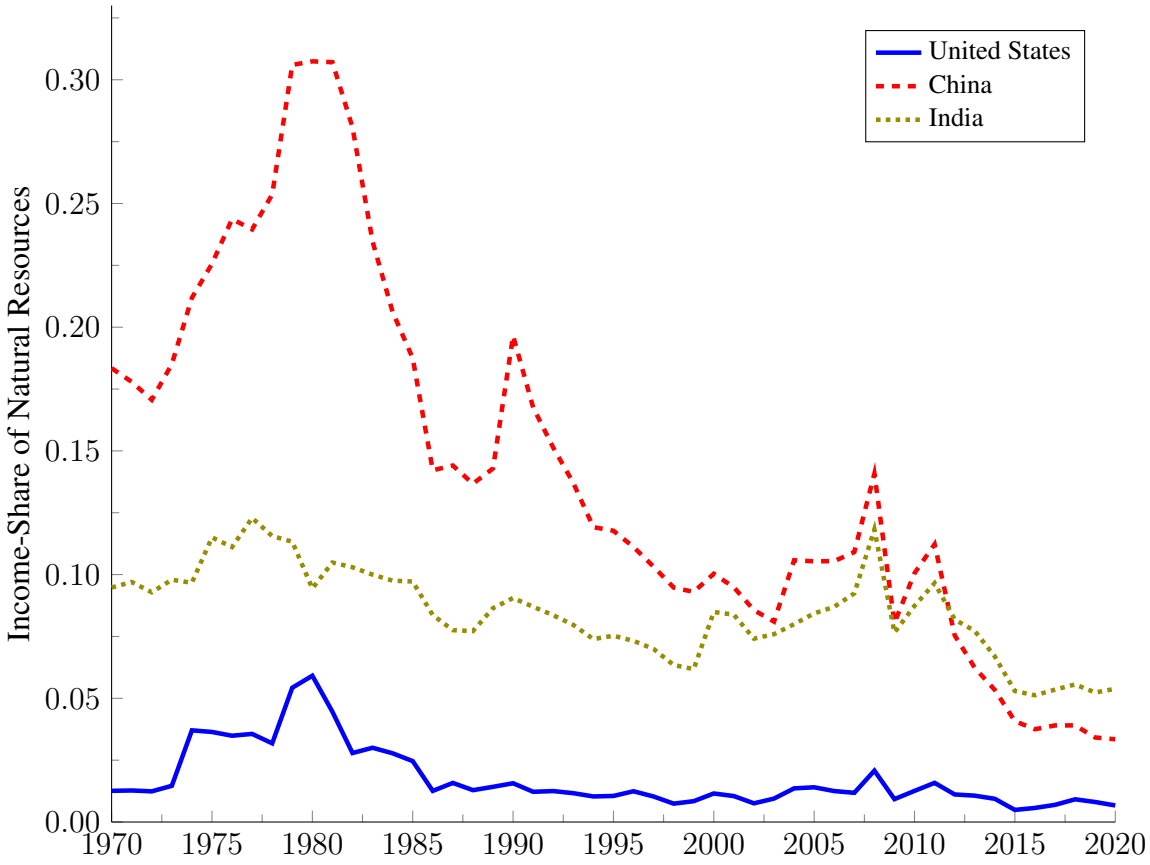
To determine the rents paid to agricultural land—inclusive of both crop and pasture applications—three data series are employed. First, FAOSTAT's "Value of Agricultural Production" (updated 11/15/2022) provides the total value of agricultural output in units of "Gross Production Value (constant 2014-2016 thousand US$)". We drop all animal products from the dataset except for cattle, sheep and goat products because the cost share due to land rents for animal products other than these 3 are small enough to be justifiably ignored. We then sum the total value of agricultural output for each country and at the global level.

Second, we need the share of total agricultural cost that is paid to agricultural land in each country and at the global level. The USDA's "International Agricultural Productivity" (updated 10/7/2022) can be leveraged here. Estimates of factor shares within agricultural production are provided for every decade from 1961-2020. Total agricultural revenues equal total agricultural costs inclusive of implicit land and capital rents. Therefore we can multiply the decadal factor shares for land by total annual agricultural revenue to get the total rents paid to agricultural land in each location-year.

Finally, we use the World Bank's GDP dataset (updated in constant 2015 US$ to be in the same units as agricultural land rents). We divide agricultural land rents by GDP to get the percent

of GDP paid to all agricultural land. We then simply combine the World Bank rent estimates for timber and subsoil minerals with the agricultural land rents to get the total percent of GDP paid to the recorded natural resources.

Figure A2: Different countries have similar long-run trend



Notes: Income-share of resources by country. Both across and within countries levels of economic development (e.g., human and physical capital accumulation) the share of income going to natural resources shrinks. This suggests natural resources have been complements over the domain of economic development observed through 2020.

Estimates of the factor share of agricultural land (in agricultural production) are between 20-30%, which are consistent with estimates exceeding 20% for land's factor share in agriculturally based societies (Weil and Wilde, 2009). Agricultural output is now a small share of global production, partially giving rise to the small estimates in Figure 1.

Another reason for these low values is that we omit urban land from "land's" share of GDP.

Urban land values are clearly tied to man-made structures and the people living on or near it. Put differently, humanity could choose to make more urban land, so it is fundamentally not a fixed factor.[23] It would be a mistake to use the high value commanded by urban land as a reflection of natural resources becoming scarce. If anything, this seems to be evidence for a desire (directly or indirectly) to have *more* nearby people.

In terms of resource use in Figure 3 we take data from various sources. Timber use is measured in pre-processed ("roundwood") and comes from the FAO's Forestry Production and Trade Database. The rest of the variables come from ourworldindata.org, which aggregates and hosts data from other sources.

E.2 Labor input and physical capital definitions

This data is taken with minimal modifications directly from the Penn World Tables v.10.01 (Feenstra et al., 2015, 2023). Physical capital is the variable *rnna* which is the real (in 2017 USD) national accounts recordings of physical capital (aggregated to the global level, of course). The labor input is the product of number of people employed (*emp*) $\times$ hours worked per person (*avh*) $\times$ a human capital index constructed from data on schooling (*hc*).

Not all countries have these data available for all years. Therefore, before aggregating across countries we restrict the sample to countries with non-missing data from all variables between 1970-2019. This ensures that the growth over the period in our global approximation is a (weighted) average growth rate across countries with consistent data, rather than a feature of more countries having the relevant data over time.

[23] If the resources necessary to build cities were becoming scarce that would of course matter. But this is already reflected in non-urban land prices, mineral rents, and to some extent timber values that are captured in our methodology.