Population Wellbeing Initiative at the UT-Austin PRC

Working Paper 2305

Marginal Benefits of Population: Evidence from a Malthusian Semi-Endogenous Growth Model

The relationship between human population sizes and economic well-being has been long debated. We revisit this question in a modern, forward-looking setting using a semi-endogenous growth model with natural resources: people innovate, but dilute fixed resources. Conditional on a stable population, the model produces an analytical steady-state with a parsimonious sufficient statistic for the (local) relationship between long-run per capita income and population sizes. Calibration suggests a positive relationship between steady state population and GDP per capita at current levels. Furthermore, trends in natural resource shares provide evidence that humanity is not yet near the income-maximizing population level.

Authors

Maya Eden

Kevin Kuruc

The University of Texas at Austin
Population Wellbeing Initiative

Marginal Benefits of Population: Evidence from a Malthusian Semi-Endogenous Growth Model*

Maya Eden[†] Kevin Kuruc[‡]

January 30, 2023

Abstract

The relationship between human population sizes and economic well-being has been long debated. We revisit this question in a modern, forward-looking setting using a semi-endogenous growth model with natural resources: people innovate, but dilute fixed resources. Conditional on a stable population, the model produces an analytical steady-state with a parsimonious sufficient statistic for the (local) relationship between long-run per capita income and population sizes. Calibration suggests a positive relationship between steady state population and GDP per capita at current levels. Furthermore, trends in natural resource shares provide evidence that humanity is not yet near the income-maximizing population level.

Keywords: Optimal population, innovation, endogenous growth, natural resources, Malthusian constraints

*We would like to thank Oded Galor, Mike Geruso, Chad Jones, Marta Prato, Dean Spears, David Weil, Mu-Jeung Yang, Anson Zhou and participants at the Global Priorities Institute's Workshops on Global Priorities Research, 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. All remaining errors are our own.

†Economics Department, Brandeis University; CEPR; Global Priorities Institute, University of Oxford

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

1 Introduction

Dating back to Malthus' (1798) *An Essay on the Principle of Population* scholars have speculated about the relationship between population sizes and per capita economic well-being. The classic Malthusian concern is that, in the presence of a fixed factor of production, the economy exhibits decreasing returns to scale. Each additional worker gets less of this factor which decreases average labor productivity (and, hence, living standards). More recently, Paul Romer's Nobel Prize winning work sparked a literature instead premised on the non-rivalry of ideas and how this implies *increasing* returns to scale (Romer, 1990). More people implies more ideas, which in turn raise everyone's productivity. While there exists work considering both of these forces (e.g., Boserup, 1965; Galor and Weil, 2000; Jones, 2001), it has mostly focused on explaining the transition to industrialization. Much less has been forward-looking, resulting in considerable uncertainty regarding the overall effects of population on living standards in the present and future.

Here we ask what leading models from these respective sub-disciplines can teach us when integrated. The core of the model has natural resource constraints—à la Dasgupta (2021)—and a semi-endogenous growth (SEG) component—à la Jones (1995, 2022). In the presence of a fixed population the model gives rise to a parsimonious (zero-growth) steady state. Furthermore, this steady state delivers a sufficient statistic that governs the (local) relationship between long-run per capita income and long-run populations. Calibrating this sufficient statistic to recently estimated moments implies that, at current population levels, the elasticity between population sizes and per capita income is positive. Determining the globally optimal population requires more parametric assumptions, but it seems plausibly quite large in this framework.

The first ingredient of the model is a renewable resource. In a recent comprehensive report on the economics of biodiversity, Dasgupta (2021) argues that at our level of abstraction, environmental constraints can be conceptualized as a renewable resource model with a tipping point. As long as the health of ecosystems remains above the tipping point, nature (partially) regenerates. The steady-state solution to such a problem is to draw a constant flow of environmental services that is independent of population size. Large and small populations share the intergenerational challenge of drawing only the highest *sustainable* level. The solution to this independent sub-problem can then be treated isomorphically to a fixed-factor.

The other key ingredient is the process of knowledge accumulation. We model this process as increasing with population, as in the (semi-)endogenous growth literature (Jones, 1995). A standard law of motion for TFP is employed with one exception. We relax the standard assumption of non-zero knowledge depreciation. This is a natural generalization—some of our intellectual effort (i.e., librarians) is certainly spent on knowledge upkeep, implying some depreciating momentum—and does not change the balanced growth properties of these models. In the fixed population case, however, non-zero depreciation gives rise to a steady-state that can be easily analyzed.[1]

The relationship between long-run population levels and per-capita income then depends directly on the two aforementioned ingredients: the fixed environmental service needs to be divided amongst more people, but TFP increases with population sizes. How these interact and compete against one another at different population sizes will be governed by the specific functional form and parameterization of the aggregate production function; we consider the generalized case of constant returns to scale in rival inputs. The elasticity of per capita income with respect to population can be shown to equal the elasticity of population on TFP—which is determined by the degree of diminishing returns in the innovation equation—minus the Malthusian elasticity—which is exactly equal to the income-share of natural resources.

This analytical elasticity has a few important implications. First, and the main result of the paper: the income share of natural resources is small enough that it is likely below the elasticity of productivity with respect to population, resulting in a global economy with all-things-considered (local) increasing returns to scale. Specifically, estimates of the endogenous growth channel (Bloom et al., 2020; Peters, 2022) imply that the elasticity of TFP with respect to populations is between 0.2-0.6; natural resources are estimated to command between 0.02-0.1 of total income. The central estimates for these parameters put the knowledge externality between 5-10 times the environmental externality, giving ample room for imperfections in either measurement. This implies that long-run incomes are increasing in long-run populations at our current size.

Imposing more structure, further statements can be made regarding the income-maximizing population level. In standard semi-endogenous growth formulations, the elasticity of TFP with respect to population is constant. Therefore, the all-things-considered population-income elasticity will be zero—i.e., a local maximum—only when the income share of natural resources increases to that level. The historical data suggests that natural resources' income share is stable, or even falling, with population sizes. From the view of the model, this supports a Cobb-Douglas specification between man-made inputs and natural resources; a specification that does not lend itself to a bounded optimal population.[2] Therefore, to reach such a local max, there must be a structural

[1] In a standard SEG model, zero population growth also implies a balanced growth path with zero TFP growth in the limit, but TFP growth converges to zero slowly enough that the level of TFP $\to \infty$ and thus the model does not admit a steady state representation.

[2] Alternatively, it could be that these inputs are complementary, but that technological progress is consistently

change in the pass-through from natural resources to GDP. Where such a structural change will occur is necessarily difficult to speculate about and leads us to the tentative qualitative conclusion that humanity is not yet near an income-maximizing population size.

1.1 Relationship to prior literature

This paper sits most closely to the literature studying the causes of the shift from a Malthusian growth regime—with stagnant incomes and relatively small populations—to a modern growth regime with unprecedented growth in both incomes and populations (e.g., Kremer, 1993; Galor and Weil, 2000; Jones, 2001; Galor, 2011). These models similarly rely on a fixed factor in production and TFP that increases in the size of the population. The key difference in this paper is our forward looking focus: models built to explain why the world experienced a phase change to sustained growth may be of limited relevance for understanding whether the modern economy would be more productive with 7 or 8 billion people.

This forward looking focus leads us to take a few noteworthy modeling departures from the existing literature. Most importantly, we study different long-run population *levels*. Long-run growth analyses almost always analyze balanced growth paths calibrated to the last few centuries of positive population growth. However, the rate of global population growth is expected to fall to or below zero this century. A more relevant question than "is indefinite 2% population growth better for long-run living standards than 1% population growth?" is "is an indefinite population of 7 billion preferable to 8 billion?" This is true too for related work arguing that perpetual population growth is preferable to perpetual population decline (Jones, 2022). While an illuminating and compelling strand of literature, it is hard to know what a world of trillions—or millions—looks like for living standards, and therefore how to apply balanced growth lessons to questions of population policy.

A second important difference is that we avoid modeling fertility decisions, or how these populations attain their eventual size. The objective is not to predict where world populations will settle—if indeed they settle anywhere—but to ask whether an increase in populations would increase or decrease per capita incomes. The benefits of this focus are to (i) greatly simplify the exposition and (ii) sidestep unresolved debates regarding the importance of human capital accumulation for growth under different family sizes. To see why little is lost when we abstract from

natural-resource augmenting. That leads to equally optimistic takeaways regarding the size of the population the earth can support.

fertility decisions, notice that if populations stabilize *anywhere*, fertility decisions must converge to about 2 children per family, regardless of the population level. Per capita child-rearing costs and human capital investments in the long-run should then not depend on population levels.[3][4]

There are other benefits to studying a steady state induced by a stable population as opposed to a balanced growth path (BGP). Foremost among these is our ability to generalize the production environment to any constant returns to scale (in rival inputs) function. To generate a non-zero BGP with a fixed factor it is necessary to have a Cobb-Douglas production structure. Cobb-Douglas imposes elasticities of substitution that are inconsistent with many observers view of environmental inputs to production (i.e., that they are necessary, hard to substitute from, and may bind more intensely as populations grow) and are known to have important implications for the effects of population sizes in this class of models (Wilde, 2012; Leukhina and Turnovsky, 2016). Our sufficient statistic allows for arbitrary and non-constant degrees of substitutability, helping to bring nuanced concerns of population pressures into a transparent macroeconomic representation.

A few recent and closely related papers share our forward-looking focus as well as an interest in long-run population levels (as opposed to growth rates). Córdoba et al. (2022) study an optimal population with a finite resource. However, Córdoba et al. (2022) has no positive feedback loop between population and technology. Instead, the model assumes a Malthusian environment with decreasing returns, but the planner cares about actual and potential people. Larger populations have intrinsic value in their framework because of the additional existences, but these existences drive down average living standards, generating the trade off studied. Similarly, Pindyck (2022) characterizes sustainable consumption paths conditional on a given relationship between per capita incomes and population. Our objective is complementary to these papers, studying whether (and at what population sizes) a trade-off exists between population size and per capita consumption. Peretto and Valente (2015) is even more closely related, merging a Schumpeterian model of growth into a model with a finite resource and studying long-runs with stable popultions. While sharing much of our motivation regarding non-perpetual population growth, they instead focus on modeling fertility with the explicit objective of highlighting under what conditions a stable solution exists

[3]This is in contrast to a BGP with slower or faster population growth: on a BGP with faster population growth the parent:child relationship is permanently worse than one with slower population growth, perhaps resulting in lower per capita human capital investment. This mechanism is at the core of Unified Growth Theory (Galor and Weil, 2000; Galor, 2011) and can lead to different takeaways than standard semi-engodegenous models without this explicit human capital decision.

[4]An admitted limitation is that we miss potentially important transition dynamics with this focus. Following a long tradition in macroeconomics, we recognize that steady-state effects are informative about the all-things-considered effects and especially easy to analyze, serving as a natural starting point.

(as opposed to endogenous explosion/collapse of populations). They say less about the question of whether per capita incomes would be higher or lower for larger or smaller long-run populations, a question that will be of relevance should governments want to arrest an endogenously occurring dynamic of persistent population decline.

Finally, there is a literature too large to adequately cover on the negative relationship between populations and living standards. Recent work by economists includes—but is certainly not limited to—Dasgupta (2021), Dasgupta et al. (2021), Henderson et al. (2022) and Johnson and Vollrath (2020). This literature focuses on the negative effects of resource scarcity—which our model shares—without the potential benefits of population size. These papers advance an important line of research on potential environmental problems. Indeed, as we argue later, our model shares the first-order concern of Dasgupta (2021) regarding sustainable resource use (unconditional of population size): we cannot easily conceive of a long-run model where humanity sub-optimally tip ecosystems into oblivion. Our focus is merely one level down. Solving the sustainability problem necessarily requires restricting aggregate per period resource use, which gives rise to the fixed factor representation we employ to ask about populations conditional on such a solution. More positively, Kuruc, Vyas, Budolfson, Geruso and Spears (2022) show that population sizes change too slowly relative to the urgency of climate change for contemporaneous fertility rates to matter for this challenge. Similarly, if society determined that larger long-run populations were beneficial conditional on environmental management, there would be many decades to develop the policies and institutions to support this goal before populations were meaningfully different.

2 The Model

This section presents the simplest model that illustrates the main point of the paper using two key equations.

$$Y = AF(N, \bar{E}) \tag{1}$$

$$\frac{\dot{A}}{A} = \alpha N^{\lambda} A^{-\beta} - \delta_A \tag{2}$$

A renewable resource problem gives rise to a Malthusian aggregate production function that has constant returns to scale in rival inputs and includes a fixed factor, $\bar{E}$. Total factor productivity, A, is Hicks-neutral and increases in population sizes, as in the semi-endogenous growth literature. N

represents the human population, but can generally be seen to incorporate factors that scale with it, such as physical and human capital. These model components are detailed in turn.

2.1 The Malthusian Component

In traditional Malthusian models designed to describe the pre-industrial world it is standard to include a fixed factor in production representing arable land (Kremer, 1993; Galor and Weil, 2000; Jones, 2001; Córdoba and Liu, 2022). The modern world is more complex. Concerns over non-renewable resources, climate change, loss of biodiversity, etc., have replaced concerns about the quantity of productive soil. However, in a recent treatise on the economics of biodiversity, Dasgupta (2021) argues that management and withdrawal from the biosphere can be roughly conceptualized as a renewable resource problem. If nature is undisturbed 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 exhibits. When withdrawal is exactly equal to regeneration, the stock remains constant (see Appendix A for details).[5] Notice that this description does not depend on the population size. Population has no independent effect on regeneration rates nor does its size influence whether it would be inefficient to not manage the biosphere at the level that maximizes what can be withdrawn. Both large and small populations face the challenge that their contemporaneous well-being is increasing in withdrawal, but their dynamic well-being requires preservation of resources.

One caveat of this representation is that we do not model the ways in which population sizes affect the management of natural resources. This is motivated by two observations. First, unsustainable paths that result in the stock of available ecosystem services going to zero are mathematically trivial, but sobering: long-run incomes go to zero in this model.[6] Regardless of population size it is imperative to avoid eroding ecosystems to sufficiently low levels. Second, it is ex-ante ambiguous whether larger populations are more likely to pursue an unsustainable path. On one hand, there are more people which makes coordination more difficult, other things equal (and perhaps a higher marginal utility per person of overdrawing). But there will also be more ideas for institutions and/or technologies that may increase the probability of overcoming these coordination problems. For example, global agricultural land use has fallen since 2000 in spite of population growth (FAO,

[5]Note that there are a continuum of steady-state solutions, each mapping into a steady-state stock of nature (broadly defined). The only condition is that the withdrawal is equal to whatever regeneration some stock of nature produces.

[6]That is, unless they are perfectly substitutable with man-made inputs. This seems unlikely and is not a view that is defended.

2022), counterintuitively suggesting we are operating more sustainably (in that sector, at least) relative to when the world population was about 20% smaller.[7] The cause of this seems like it must be improved agricultural TFP. The simplification here reflects our uncertainty over the long-run relationship between population, technology, wealth, and sustainability. As such, we presuppose a sustainable steady-state solution to resource extraction that is independent of N, resulting in the fixed-factor representation of Equation (1).

Transitory dynamics deserve mention. The first is climate change, perhaps the most serious environmental challenge of modern times. Kuruc, Vyas, Budolfson, Geruso and Spears (2022) address this quantitatively using William Nordhaus' DICE model and find that population effects on global warming are small for any realistic population change. This is because population is a stock. It takes many decades to influence population sizes through realistic changes in fertility—by the time differences can reasonably emerge, nearly all of the global warming story will be written. The second are non-renewable resources, which necessarily imply that humanity cannot reach a steady-state (as any limiting solution requires zero use of something non-accumulable). In Section 3.3.1 we show how this problem looks with a non-renewable resource that gets used up over time—the key results are strikingly similar. More substantively, the reason we ignore exhaustible resources is that they are mainly fossil fuels, of which there appear to exist renewable substitutes (wind, nuclear, geothermal, solar, ethanol, etc). Minerals like lithium are discussed as "non-renewable" in popular writing, but in fact they are *fully renewable* in this framework and fit precisely into a fixed factor representation. The important fact is that they are non-*exhaustible*. As with land, there exists some stock of minerals that are shared more widely as populations grow. But these minerals are never used up like fossil fuels are.

In summary then, the renewable resource framing is both realistic and analytically convenient. The environmental constraint is that as populations grow each worker has less E at her disposal, not that the environment is further eroded. This generates lower per worker productivity, holding fixed A, implying lower per capita incomes. Under the constant returns to scale assumption:

$$\frac{Y}{N} = y = AF\Big(1, \frac{\bar{E}}{N}\Big). \tag{3}$$

[7] Global agricultural land use available at: https://www.fao.org/faostat/en/#data.

2.2 The Semi-Endogenous Growth Component

Theories of endogenous growth—supported by empirical evidence—give strong reasons to believe that A is not invariant to population sizes and the level of economic activity (Romer, 1990; Kremer, 1993; Galor and Weil, 2000; Jones, 2005; Jones and Romer, 2010; Peters, 2022). Larger populations generate more ideas. Formally, the law of motion for A takes the form in Equation 2 (reproduced here), closely following the semi-endogenous growth literature.

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

Population enters the production function of knowledge directly. This can be thought of as a learning-by-doing process that is commonly employed (e.g., Jones, 2022) or as a scenario where a constant share of human capital is employed in R&D. The scalar on population, α, mediates the effect of people on knowledge accumulation; this term roughly captures the productivity of people as well as the share of time spent in research-focused activity. The exponent λ allows for intratemporal diminishing returns to research effort; β allows for intertemporal diminishing returns to research effort (i.e., if ideas get harder to find after picking the lowest hanging fruit). The final term, δ_A, represents depreciation of the knowledge stock.

This relaxation to non-zero depreciation is non-standard in the endogenous growth literature, but is both realistic and gives rise to an analytical steady state.[8] On the realism of this assumption, the micro literature estimating returns to R&D at the firm-level finds substantial knowledge depreciation (Hall et al., 2010). This is more difficult to estimate at the macro level. Diamond (1993) documents that as hunter-gather societies were cut off from one another—resulting in a reduced N for the respective closed economies—important technologies were lost. Even since the advent of writing as a preservation technology it seems we have lost technology such as Roman Concrete. Kremer (1993) notes these anthropological findings and suggests they are evidence for depreciation to be included in the law of motion for knowledge.

Independent, contemporaneous evidence in favor of a depreciation-like force is the effort spent preserving and organizing knowledge. Librarians, authors of review articles and textbooks, Wikipedia contributors, etc., are working to facilitate the use of existing knowledge, rather than generate new knowledge. Presumably these efforts are increasing in the level of the knowledge stock. Furthermore, its existence implies some counterfactual reduction of A in the absence of such work, where

[8] An exception is Dietz and Stern (2015) which includes knowledge depreciation in a climate-economy model.

A is better thought of as employable knowledge. Admittedly it is difficult to disentangle—both conceptually and empirically—knowledge being *lost* as opposed knowledge being unused. What matters here is only that the amount of knowledge that can be effectively used in any period has an endogenous upper bound that is increasing in the size of the population. This force is at present underexplored, so we follow the capital accumulation convention and use a parsimonious linear representation. Importantly, we show below that (i) any $\delta_A > 0$ generates a steady-state and (ii) the sufficient statistic governing the relationship between population and long-run incomes is invariant to the magnitude of δ_A. The qualitative existence of increasing upkeep/organization effort is crucial for this analysis, its quantitative magnitude is not.

To conclude this section, the simplest model has two key ingredients, respectively capturing the Malthusian concern that large populations dilute natural resources per person and the endogenous growth insight that an economy with more resources should produce more ideas. These components sit within a general CRS production function, allowing them to interact in the usual Cobb-Douglas setting as well as more generalized functions. How these forces push against one another, and on what parameters this depends, is discussed next.

3 Steady State Results

This section presents the steady state implications of this model. An analytical relationship between long-run per capita income, $\bar{y}$, and population levels, $\bar{N}$, arises. The elasticity of per capita income with respect to population can be calibrated to two empirical moments. The ratio between intra- and inter-period diminishing returns knowledge production, $\frac{\lambda}{\beta}$, governs the positive effect of population increases. The income-share of the fixed-factor determines the negative Malthusian effect. Comparing recent estimates of these terms implies that marginal increases to current population sizes increases long-run per capita incomes. Globally, the existence of an income-maximizing population depends on when the income share of natural resources gets sufficiently large. Historical trends suggest this income share has not increased as populations have grown, seeming to imply that we are not yet near such an income-maximizing population.

3.1 Main Results: The marginal effect of population on per capita income

The only endogenous dynamic variable in this setting is A; $\bar{E}, \bar{N}$ are exogenous. Restricting our focus to a stable population, Equation 2 implies a steady-state for A.[9]

$$\bar{A} = \left(\frac{\alpha \bar{N}^{\lambda}}{\delta_A}\right)^{\frac{1}{\beta}} \tag{4}$$

Conceptually, this steady state arises as the knowledge stock becomes so large that to even maintain, organize, and employ it takes all people-hours in this sector.[10] As the stock of knowledge gets unwieldy, the challenge is making use of existing knowledge, not generating new ideas.

The math here is simple for anyone familiar with the standard Solow model. But the implications are striking. Economists have long known that capital accumulation cannot lead to sustained growth because of increasing aggregate depreciation costs. This same issue plausibly exists one level up, where conventional wisdom currently holds that sustained exponential growth can come from knowledge accumulation. And indeed, with an exponentially growing population, exponential growth in A remains possible (see Appendix B). The likely reason this term has been thus far underexplored is because nearly all balanced growth analyses of semi-endogenous growth models assume a positive rate of population growth, where this term affects long-run levels but not growth rates.[11]

The goal of the current exercise is to determine the sign of the elasticity of y with respect to N. It will help to conceptualize the per-worker production function in the following way.

$$\bar{y}(\bar{N}) = \underbrace{\bar{A}(\bar{N})}_{\text{Eqn. 4}} \times \underbrace{F\left(1, \frac{\bar{E}}{\bar{N}}\right)}_{\text{Eqn. 3}} \tag{5}$$

Then, the elasticity is

$$\begin{aligned}\frac{\partial ln(\bar{y})}{\partial ln(\bar{N})} &= \frac{\partial ln(\bar{A})}{\partial ln(\bar{N})} + \frac{\partial ln(F)}{\partial ln(\bar{N})} \\ &= \frac{\lambda}{\beta} - \frac{\partial F/\partial E \times \bar{E}/\bar{N}}{F}.\end{aligned} \tag{6}$$

[9] Set $\dot{A} = 0$, and solve for A.

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

[11] One exception is Jones (2022) on depopulation, where this issue is touched on in the introduction but not explored analytically.

The positive term represents how much more knowledge can be accumulated and productively maintained in the steady-state of this economy from a 1% increase in $\bar{N}$. This is governed by the ratio of the intra-period diminishing returns to research effort, λ, and the degree to which knowledge becomes more difficult to accumulate/organize as A increases, β. Recent work by Bloom et al. (2020) directly targets this parameter and estimates that this ratio lies between (0.2,0.5) for the aggregate economy. In a different setting with plausibly exogenous long-run population changes, Peters (2022) estimates a closely related term to be about 0.5.

The negative term represents—holding A fixed—by how much per-worker productivity falls from a 1% increase in $\bar{N}$. This is the Malthusian channel. This sub-elasticity is exactly equal to the income share of natural resources in competitive markets, a result anticipated in a constant elasticity of substitution framework by Weil and Wilde (2009). We denote this term ϕ_E.

Existing data implies that the global share of natural resources have fluctuated between 2.5% and 6% over the past half century. Figure 1 plots data taken from the World Bank and United States Department of Agriculture on the rents earned by natural resources.[12] We follow the World Bank's *Changing Wealth of Nations* classification and include: subsoil energy and minerals, timber resources and all agricultural land.[13] Estimates of the factor share of agricultural land (in agricultural production) are between 20-30%. This is consistent with historical approximations of about $\frac{1}{3}$ for land's factor share in pre-industrial agriculturally based societies. The small estimates in Figure 1 come from the fact that agriculture itself is now a relatively small share of global production, at less than 10%. This implies that agricultural land's share of total global GDP is less than 2%.

Another reason for our low estimate is that we omit urban land from land's share of GDP. Urban land values are clearly tied to man-made structures and people living on or near it. Put differently, humanity could choose to make more urban land, so it is fundamentally not a fixed resource.[14] It would be a mistake to use the high value commanded by urban land as a reflection of natural resources becoming scarce. In fact, the high value of urban land seems, if anything, to be evidence for a desire (directly or indirectly) to have *more* nearby people.

Returning to Equation 6, estimates for the respective terms are then (0.2,0.5) for the knowledge-elasticity and (-0.1,-0.02) for the Malthusian-elasticity, suggesting an elasticity between (0.1,0.5)

[12]The USDA produces these estimates from the United Nation's Food and Agricultural Organization, so it is global in scope.

[13]This classification misses water resources (and fish, due to the international nature of fish stocks), but these are not plausibly large enough to increase the 2.5-6% resource rents to the 20+% that would be necessary for these resource rents to match the knowledge-elasticity.

[14]If the resources necessary to build cities were becoming scarce, that would be a different story. But that is already reflected in timber and mineral rents that are captured in our methodology.

for the all-things-considered elasticity between population size and per capita incomes. There are reasons both of these estimates should come with uncertainty. Natural resources are particularly difficult to measure. But for this quantity to exceed the knowledge elasticity, it would need to be the case that global data understates their importance between 3-10-fold, depending on the true value of $\frac{\lambda}{\beta}$.

More important than the quantitative takeaway is the qualitative finding that the relationship we study very likely has a (locally) positive relationship if the forces we capture are the dominant considerations. What is less clear is whether we should expect this relationship to hold if populations grow and natural resources become a more binding constraint on growth. We turn to this question next.

Figure 1: Natural resource shares are small and non-increasing

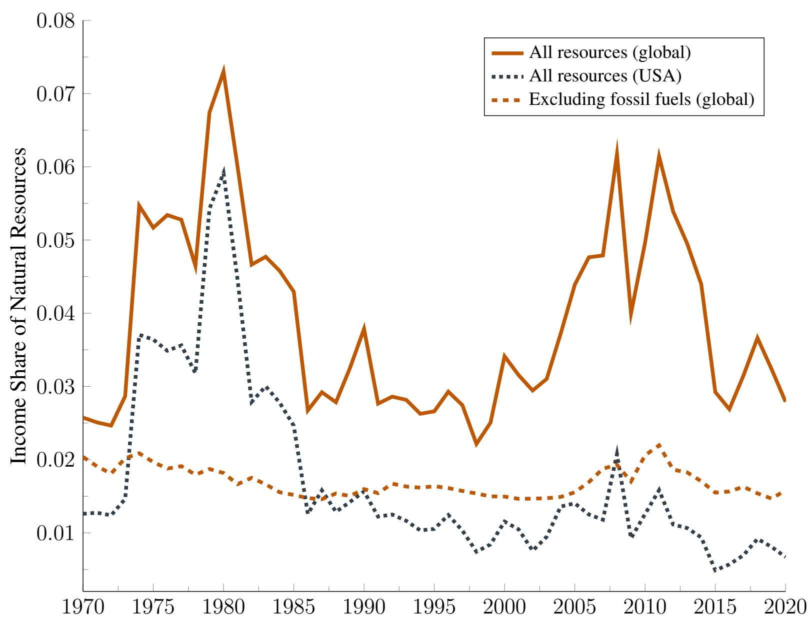

Notes: Share of income paid to natural resources over time. (Solid) Global income share paid to all natural resources, following World Bank classification to include: (a) subsoil energy and minerals; (b) timber resources; (c) crop land; (d) pasture land. Details of each category are contained in the Appendix. (Dashed) This same income share, but using only US data to ensure trend or level not driven by unreliable global data. (Dashed) Excludes fossil fuels which drive the level and volatility of this series, but not its flat trend.

3.2 Income-Maximizing Population

Beyond the local relationship, we can use the steady-state structure in this paper to examine the global relationship between populations and per capita incomes. In particular, we may be able to say something about how close we are to a local max of the relationship $\bar{y}(\bar{N})$, or if one even exists. Greaves (2022), writing on exactly this topic in the philosophical literature, stipulates that there must be an inverted-U relationship between populations and per capita well-being. At small populations, the gains from non-rival benefits we endow on one another exceed any crowding effects, but at sufficiently large populations the opposite must be true. This is a reasonable *a priori* view and one that is consistent with environmental worries.

Figure 2: Income and population plausibly exhibit a non-monotonic relationship

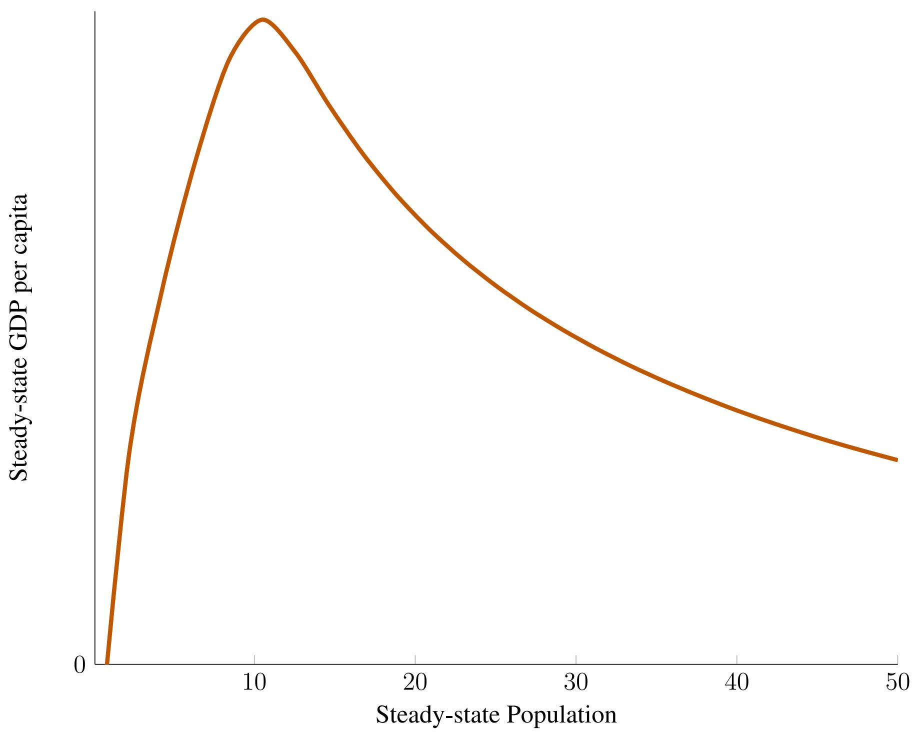

Specifically, consider an $F(.,.)$ that has a constant elasticity of substitution between natural and man-made inputs as in Equations 7, 8. Furthermore, if natural resources are necessary and become a more significant drag on productivity as population sizes increase, it would be natural to impose that these inputs are gross complements, or that $\rho < 1$. Figure 2 traces out what such a relationship looks like, generally.

$$Y = A\Big[aN^{\rho} + (1-a)E^{\rho}\Big]^{\frac{1}{\rho}} \tag{7}$$

This delivers the following function form for $\bar{y}(\bar{N})$:

$$\bar{y}(\bar{N}) = \Big(\frac{\alpha \bar{N}^{\lambda}}{\delta_A}\Big)^{\frac{1}{\beta}}\Big[a + (1-a)\Big(\frac{E}{N}\Big)^{\frac{1}{\rho}}\Big]^{\rho}. \tag{8}$$

The results in Section 3.1 at a minimum inform us that we would be on the left-hand side of such a hump-shaped relationship, if the empirical counterparts to the theory are reliable. Determining when the peak would be reached requires estimating when the income-share of natural resources grows to about 0.3, our central estimate for the elasticity between TFP and population sizes $\frac{\lambda}{\beta}$. However, notice that Figure 1 does not have a clear trend. Since 1970 the population has grown from around 3 to 8 billion, and yet the natural resource share of income has, if anything, fallen. This implies that something close to a Cobb-Douglas specification is reasonable. The Cobb-Douglas production function gives rise to the simple expression in 9—a function that is monotonic in N. If the idea-externality exceeds the income share of natural resources at current and past population levels, it will *always* exceed the income share of natural resources.

$$\bar{y}(\bar{N}) = \Omega \bar{N}^{\frac{\lambda}{\beta}-(1-a)} \tag{9}$$

While somewhat instructive, the simple time-series relationship is not powered enough to rule out other possibilities. We note a few additional pieces of evidence and/or structural possibilities that leave us confident in the qualitative statement that we are not at present near an income-maximizing population. The first is cross-sectional evidence. Weil and Wilde (2009) study the Malthusian channel at the country level, which likewise requires a stance on the elasticity of substitution. They observe that countries with more physical and human capital—what our framework and theirs capture as scaling with N—pay less of GDP to natural resources. They take that as evidence that man-made inputs must be substitutable for E. This case is qualitatively similar to the Cobb-Douglas case. Under gross substitutability, as N becomes a more dominant share of economic activity (as it would with larger populations) its factor share must grow. Accordingly, the factor share of natural resources must fall, implying that E becomes an even smaller proportional drag on incomes as populations grow.

A different concern is that our Hicks-neutral assumption on TFP growth is omitting the possibility of E-augmenting technological change. Then it would be possible for the effective ratio of $\frac{E}{N}$ to remain unchanged over the last half century even as populations grew. If that is true, there is no variation in the relative inputs, and hence no evidence about the degree of substitutability between them. It is correct that we cannot rule that out. However, this possibility merely relocates the trend question to a different variable. If there has been consistent and steady progress in E-augmenting technology in response to larger populations—as, for example, Boserupian theories would predict—it is still true that a structural break will be necessary to reach an income-maximizing population. Rather than the production function eventually exhibiting not-yet-observed gross complementary, it would instead be that the type of technological progress we have so far observed would need to relent at some larger population or future date.

A final possibility is that E withdrawals are not in practice fixed and have instead grown proportionately with populations. Even if E were truly complementary, if its growth has kept pace with population growth (unsustainably) then the factor share could remain fixed. However, based on data that can be observed for major categories of renewable natural resources, the growth in their use has been negligible relative to population growth. For example, in the appendix we show that agricultural land use and timber resources have grown only 30% while populations have grown nearly 300%. If the increase in human and physical capital that has accompanied the growth in population sizes were accounted for, growth rate of man-made inputs would be larger still. Under any plausible assumptions, growth in the ratio of man-made inputs to natural inputs has grown substantially, without a corresponding rise in the factor share of natural resources.

In conclusion, if the conjecture in Figure 2 is correct, the available trends suggest that the relationship must be roughly positive at 8 billion and not particularly close to flattening out. We of course cannot rule out an eventual flattening at population sizes outside of the historical experience, but calibrating a production function (or factor-augmenting technological change) to past data does not suggest we will reach such a flattening (let alone a negative relationship) at any nearby population level.

3.3 Robustness to alternative models of environmental inputs

There are two dimensions along which we can relax the previous assumptions for how E enters the production function. These alternative formulations provide the same takeaways.

3.3.1 Exhaustible resources

Some observers are worried about the stock of exhaustible resources, fossil fuels being foremost among these. Exhaustible resources can be modeled as a classic cake-eating problem with their initial stock denoted $\mathbb{E}$. For analytical simplicity, and consistent with the flat trend in factor shares, we assume a Cobb-Douglas specification for this exercise (and again abstract from capital accumulation).

$$\max \sum_{t=0}^{\infty} \delta^t \frac{1}{1-\sigma} y_t^{1-\sigma} \tag{10}$$

$$\text{with } \; y_t = A_t \left(\frac{E_t}{\bar{N}}\right)^{1-a} \tag{11}$$

$$\text{and } \; \sum_{t=0}^{\infty} E_t = \mathbb{E} \tag{12}$$

Here the discount rate is denoted δ. We have already established that $A \to \left(\frac{\alpha \bar{N}^{\lambda}}{\delta_A}\right)^{\frac{1}{\beta}}$ independent of the environmental side of the model. For simplicity then assume A reaches this level reasonably quickly and is therefore fixed at this level for the long-run study of this model. With a fixed A, N it can be shown that the solution is characterizied by $\frac{E_t}{E_{t+1}} = \delta^{\frac{1}{1+\sigma a - a}}$.

The key takeaway is that the solution is independent of the population size, just as before. Therefore, it remains true that the trade-off of a larger population is a $\frac{\lambda}{\beta}\%$ increase in A at the expense of a 1% lower $\frac{E_t}{N}$. The difference is that the trade-off is period-by-period and the use of E falls each period. However, in this Cobb-Douglas specification a 1% decline in $\frac{E_t}{N}$ reduces y_t by $(1-a)\%$ regardless of the level of E being drawn. The difference in this setting—that $\frac{E}{N}$ is declining each period—turns out to be unimportant for comparing a 1% loss in $\frac{E}{N}$ to a $\frac{\lambda}{\beta}\%$ increase in A.

3.3.2 Non-rival benefits from natural resources

Alongside the rival ecosystem services that earn profits—and are conceptually captured in our 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).[15]

$$Y = AB^{\xi} F(N, \bar{E}) \tag{13}$$

[15] See Dasgupta (2021) Chapter 4*.

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_20 throughout the water cycle—and helps promote innovation and learning—such as plants in the Amazon that provide the ideas for new pharamcueticals. If you take a broader view of Y beyond measured GDP, you might also consider B to contribute to Y via non-use values; we all derive some intellectual value from the mere existence of bio- and scenic-diversity.

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); this introduces a channel by which the stock itself is directly beneficial. This indirect channel changes the optimal level of E. Previously, it would have been optimal to erode the biosphere until it reached its peak growth rate.[16] When B enters directly we have a competing incentive to keep B larger than where it reaches its regenerative peak.

We reiterate here that this coordination problem is difficult to solve, regardless of population sizes, and is not the problem that we focus on. If it is true that our ecosystem withdrawals are not optimal (or even efficient) that problem almost certainly needs to be solved by means other than population reduction. Even smaller populations than our own will need to figure out how to coordinate away from over use.

4 Welfare Implications

Before concluding, we discuss two conceptual dimensions along which our analysis currently falls short of producing policy recommendations: transition dynamics and the (potential) benefits of existence. Nonetheless, the long-run differences in living standards this paper highlights will clearly be an input to an all-things-considered analysis of the betterness of population sizes.

4.1 Dynamic trade-offs and social impatience

Along a transition path there exist non-trivial dynamic trade-offs. This is true for reasons both within and outside of the scope of the current model. Within the bounds of the model, a larger population immediately reduces $\frac{E}{N}$, but only increases cumulative knowledge over time. For this

[16]For example, in an ocean with fish populations at their carrying capacity, there is no net-regeneration since populations cannot grow. Doing some fishing that reduces the stock in fact increases the rate of growth, and hence what can be drawn in a steady state.

reason, a sufficiently impatient planner, even if given free choice over the number of people, will not always choose to increase the level of long-run GDP per capita.

Specifically, consider a planner facing a decision between long-run stable population N_0 and some larger population N_1 such that $ln(N_1) = ln(N_0) + \varepsilon$ (i.e., a marginal percent increase). From the steady-state analysis above we know the long-run (log) increase in y will be $\frac{\lambda}{\beta} - \phi_E$, where ϕ_E is the term corresponding to the income share of natural resources. We also know that population influences the growth rate of A, and so on impact this increase in population does not change A. The on-impact effect will be $-\phi_E$; between $t = (0, \infty)$ the effect tends towards its long-run outcome. Figure 3 demonstrates. If the planner is impatient and is attempting to maximize the

Figure 3: An Instantaneous Population Increase Implies Dynamic Trade-offs

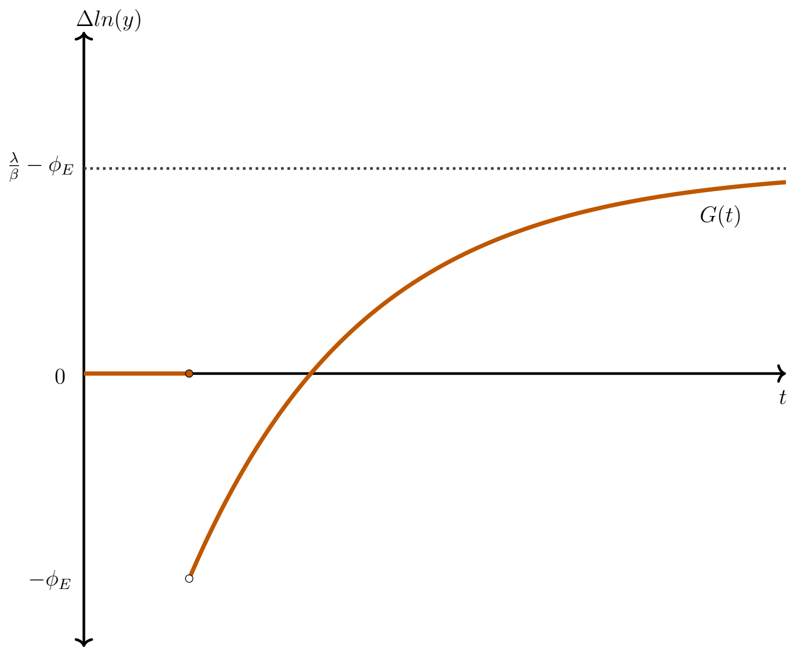

Notes: Response of per capita income to an instantaneous increase in N by 1%. y falls on impact as natural resources are split amongst more workers. Over time as A accumulates net benefits occur. For simplicity depiction is a transition from one steady-state towards a new steady-state. For a sufficiently impatient planner, the initial decline may out weigh the sustained benefit in an optimal policy setting.

discounted flow of per capita incomes, she may optimally forgoe this long-run benefit to avoid the short-run cost. Formally, there exists some pure rate of time preference $\delta > 0$, such that

$$W = \int_0^\infty e^{-\delta t} G(t) dt < 0,$$

where $G(t)$ are the per capita income gains depicted in Figure 3.

Formally solving for an optimal population path with social discounting is beyond the scope of this paper. This is because it would also require formalizing the costs of having children, and people's preferences for doing so. An advantage of our approach is that by focusing on a long-run with a stable population, we've omitted discussion on the costs and benefits of parenting: in our solution family units in any population size converge on the same fertility decisions. An adult population of N needs to have N children per generation in a steady-state; so the per capita costs and benefits are identical. By highlighting that the long-run population effects are positive, this analysis demonstrates that the marginal costs of getting to a larger population would need to be sufficiently large and long-lived to generate a negative (discounted) effect.

4.2 Social value of existences

Further complicating the welfare considerations is the fact that additional people exist in larger populations. There is a rich theoretical literature on how to make social welfare comparisons when population sizes are non-constant; i.e., how (if at all) to value additional existences.[17] A near consensus is that the standard practice to compare average well-being (proxied here by GDP per capita) is deficient.[18]

A simple alternative with more support is the *Total View*, wherein it is the sum of utility that is welfare relevant. The social welfare function would then take the form:

$$SWF(t) = N(t) \times \bar{u}(t), \tag{14}$$

where $\bar{u}(t)$ is the average utility in period t. Indeed, there are recent examples of climate and growth economists revisiting long-run analyses taking this conception of welfare seriously (e.g., Scovronick et al., 2017; Klenow et al., 2022). A world with more people living good lives may be, *ceteris paribus*, a better world.

Putting this into practice poses serious conceptual challenges. Notice that a 1% change in y is *not* a 1% change in u. Aside from diminishing returns in the utility function, there's the deeper problem that u needs to be cardinal and a y that generates 0 utility needs to be defined (normalized here to the point of indifference between existence and non-existence).

[17] See Greaves (2017) or Kuruc, Budolfson and Spears (2022) for overviews of this literature.

[18] To see why: ask yourself whether a world with one additional person who has a good—but below average—life is *strictly* worse than the status quo. Few, on reflection, find this a plausible property of a population-sensitive social welfare function.

We avoid taking quantitative stances on these parameters here and instead note that if any weight is given to additional existences, the marginal social welfare benefits of population will be larger than the per capita effects documented. As our estimates already indicate a locally positive relationship between populations and per capita incomes and a large income-maximizing population, additionally considering the value of existence only serves to strengthen such a result.[19] Interestingly, unlike most theoretical discussions of population ethics which pit per capita welfare against total welfare—and hence generate disagreement that depends on population ethics views (e.g., Córdoba et al., 2022)—the model here suggests that these disagreements may not be critical on the margin. Under any plausible population-welfare function, smaller populations with lower average living standards represents an important loss of welfare.

5 Conclusion

Whether a smaller world population would be better or worse for individual living standards is an unsettled and important question. This paper takes a first step towards answering this by combining the most notable benefits and costs into an analytically tractable framework. The empirical counterparts to the resulting sufficient statistic imply that the gradient between per capita incomes and population sizes is locally positive. Furthermore, the income-maximizing population is perhaps much larger than observed populations. While future research will certainly have more to say about the details of this problem, the transparent methodology and answer provided here serve as a much needed starting point to these debates.

References

Bloom, N., Jones, C. I., Van Reenen, J. and Webb, M. (2020), 'Are ideas getting harder to find?', *American Economic Review* **110**(4), 1104–44.

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

Córdoba, J. C. and Liu, X. (2022), 'Malthusian stagnation is efficient', *Theoretical Economics* **17**, 415–460.

[19] An underlying assumption is that the marginal existences are weakly better than non-existence. This strikes us as uncontroversial given current and expected global average living standards.

Córdoba, J. C., Liu, X. and Ripoll, M. (2022), 'Optimal population in a finite planet', *Working Paper* .

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

Dasgupta, P., Dasgupta, A. and Barrett, S. (2021), 'Population, ecological footprint and the sustainable development goals', *Environmental and Resource Economics* pp. 1–17.

Diamond, J. (1993), 'Ten thousand years of solitude', *Discover* .

Dietz, S. and Stern, N. (2015), 'Endogenous growth, convexity of damage and climate risk: How nordhaus' framework supports deep cuts in carbon emissions', *The Economic Journal* **125**(583), 574–620.

FAO (2022), Faostat, Technical report, United Nations. Retreived from: https://www.fao.org/faostat/en/#data on Dec. 14, 2022.

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

Galor, O. and Weil, D. N. (2000), 'Population, technology, and growth: From malthusian stagnation to the demographic transition and beyond', *American Economic Review* **90**(4), 806–828.

Greaves, H. (2017), 'Population axiology', *Philosophy Compass* **12**(11), e12442.

Greaves, H. (2022), Optimum population size, *in* G. Arrenhius, K. Bykvist and T. Campbell, eds, 'Oxford Handbook of Population Ethics', Oxford University Press.

Hall, B. H., Mairesse, J. and Mohnen, P. (2010), Measuring the returns to r&d, *in* B. H. Hall and N. Rosenberg, eds, 'Handbook of the Economics of Innovation', Vol. 2, Elsevier, pp. 1033–1082.

Henderson, J. V., Jang, B. Y., Storeygard, A. and Weil, D. N. (2022), 'Climate change, population growth, and population pressure', *Working Paper* . Available at: https://cpb-us-w2.wpmucdn.com/sites.brown.edu/dist/1/24/files/2022/11/Climate-Change-Population-Growth-and-Population-Pressure-September-20-200 pdf.

Johnson, R. T. and Vollrath, D. (2020), 'The role of land in temperate and tropical agriculture', *Economica* **87**(348), 901–937.

Jones, C. I. (1995), 'R&d-based models of economics growth', *Journal of Political Economy* **103**(4).

Jones, C. I. (2001), 'Was an industrial revolution inevitable? economic growth over the very long run', *The B.E. Journal of Macroeconomics* **1**(2).

Jones, C. I. (2005), Growth and ideas, *in* P. Aghion and S. N. Durlauf, eds, 'Handbook of Economic Growth', Vol. 1B, Elsevier, pp. 1063–1111.

Jones, C. I. (2022), 'The end of economic growth? Unintended consequences of a declining population', *American Economic Review* **112**(11), 3489–3527.

Jones, C. I. and Romer, P. M. (2010), 'The new kaldor facts: Ideas, institutions, population, and human capital', *American Economic Journal: Macroeconomics* **2**(1), 224–45.

Klenow, P., Jones, C. I., Bils, M. and Adhami, M. (2022), 'Population and welfare: The greatest good for the greatest number'. Available at: https://web.stanford.edu/~chadj/slides-popwelfare.pdf.

Kremer, M. (1993), 'Population growth and technological change: One million b.c. to 1990', *The quarterly journal of economics* **108**(3), 681–716.

Kuruc, K., Budolfson, M. and Spears, D. (2022), Population issues in welfare economics, ethics, and policy evaluation, *in* 'Oxford Research Encyclopedia of Economics and Finance'.

Kuruc, K., Vyas, S., Budolfson, M., Geruso, M. and Spears, D. (2022), 'A larger population yields long-run benefits net of climate damages', *Population Wellbeing Initiative Working Paper* .

Leukhina, O. M. and Turnovsky, S. J. (2016), 'Population size effects in the structural development of england', *American Economic Journal: Macroeconomics* **8**(3), 195–229.

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

Peters, M. (2022), 'Market size and spatial growth – evidence from germany's post-war population expulsions', *Econometrica* **90**(5), 2357–2396.

Pindyck, R. S. (2022), 'Population, productivity, and sustainable consumption', *NBER Working Paper* . No. 30019.

Romer, P. M. (1990), 'Endogenous technological change', *Journal of Political Economy* **98**(5, Part 2), S71–S102.

Scovronick, N., Budolfson, M. B., Dennig, F., Fleurbaey, M., Siebert, A., Socolow, R. H., Spears, D. and Wagner, F. (2017), 'Impact of population growth and population ethics on climate change mitigation policy', *Proceedings of the National Academy of Sciences* **114**(46), 12338–12343.

Weil, D. N. and Wilde, J. (2009), 'How relevant is malthus for economic development today?', *American Economic Review: Papers & Proceedings* **99**(2), 255–60.

Wilde, J. (2012), 'How substitutable are fixed factors in production? evidence from pre-industrial england', *MPRA* .

Figure 4: Different countries, similar long-run trend

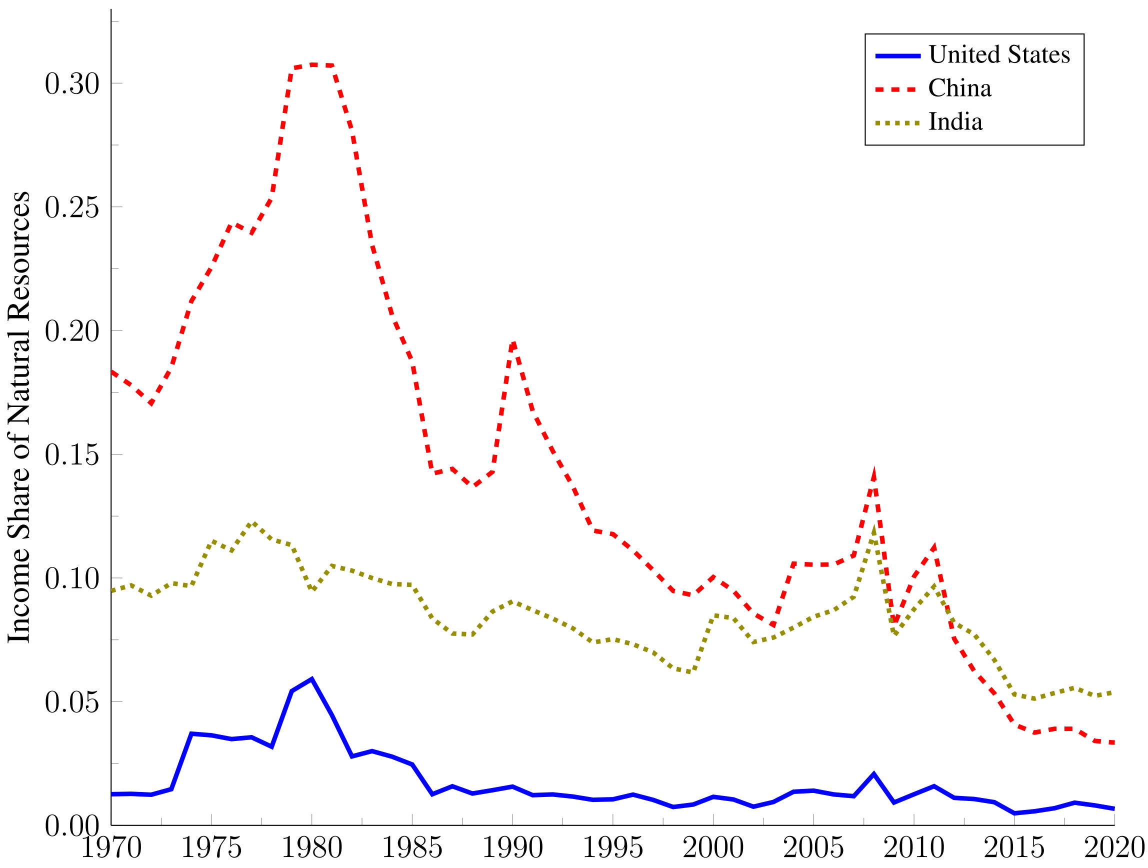

Notes: Income share of resources by country.

A 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 15, taken from Dasgupta (2021).

$$R(B) = rB\left[1 - \frac{B}{K}\right]\left[\frac{B - T}{K}\right] \tag{15}$$

B is the stock of biosphere as it relates to human production/consumption and $R(B)$ is the amount of regeneration. K is the carrying capacity on 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.

$$\dot{B}(B) = R(B) - E \tag{16}$$

In a steady-state solution $\bar{E} = R(\bar{B})$. Notice that any level of $B > T$ can be a steady state for some $\bar{E} > 0$; $\bar{E}$ must just be set to equal $R(B)$. Income is increasing in E, so the optimal solution is to manage B at the point which maximizes $R(B)$, and therefore maximizes $\bar{E}$.

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, $\bar{E}$, to this independent subproblem into the aggregate production function.

Figure 5: Regeneration rate with tipping point

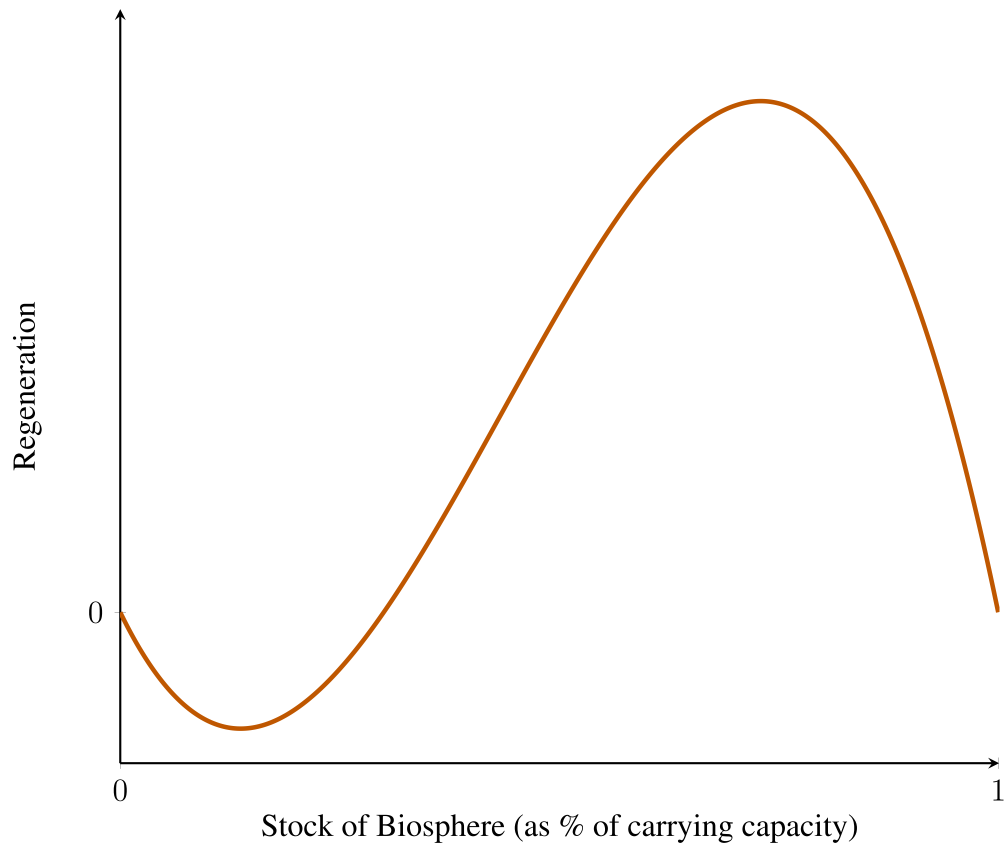

Note:.

B Balanced Growth Case

It is common when analyzing long-run trajectories to study balanced growth paths (BGP). Indeed, much of the existing literature that combines the elements we use here to study the historical escape from Malthusian stagnation does just this (e.g., Galor and Weil, 2000; Jones, 2001). Here we make explicit the relationship between existing work and our own, and confirm that the BGP properties of the model are consistent with the takeaways of the zero-growth steady state.

Jones (2001) is the model most closely related to our own. Jones uses a similar specification for knowledge accumulation and likewise employs a fixed factor as is standard in studying Malthusian stagnation. The differences are (i) his restriction to Cobb-Douglas hides meaningful insights gleaned from the generalized CES we employ, (ii) his focus is on *why* an economic and population boom took place (wherein he includes a structural model of fertility), and (iii) his historical focus makes it unclear whether the parametric or structural assumptions are relevant for understanding the present and future. This latter point is not a criticism of this work, indeed Jones himself notes the exercises forward-looking implications are of questionable value. Nonetheless, the takeaways from such a model are predictably similar to our own which we highlight here.

Consider the same three-equation model with a fixed positive rate of population growth, n.

$$
\begin{aligned}
y =& A\left(\frac{E}{N}\right)^{1-a} \\
\frac{\dot{A}}{A} =& \theta N^{\lambda} A^{-\beta} - \delta_A \\
\frac{\dot{N}}{N} =& n > 0
\end{aligned}
$$

Here, as in Jones (2001), we are restricted to a Cobb-Douglas specification as a CES does not admit a BGP.[20] The novel inclusion of depreciation in our setting does not change the BGP properties of a standard semi-endogenous formulation: for $\frac{\dot{A}}{A}$ to be constant, it is still the case that $N^{\lambda} \times A^{-\beta}$ needs to be constant. This implies that the growth in A must be proportional to n. Specifically,

$$\frac{\dot{A}}{A} = \frac{\lambda}{\beta} n. \tag{17}$$

For curious readers: the depreciation rate lowers the level of TFP along the BGP, but not the growth properties of the BGP. As before, the growth rate of y will be the sum of the growth rate of A and the growth rate of the Cobb-Douglas piece. The time derivative of the logged Cobb-Douglas function will be the elasticity with respect to population times the growth in population: $-(1-a)n$. In sum, on the BGP,

$$\frac{\dot{y}}{y} = \frac{\lambda}{\beta} n - (1-a)n = \Big(\frac{\lambda}{\beta} - (1-a)\Big) n. \tag{18}$$

For constant positive population growth with a Cobb-Douglas production structure the same qualitative solution arises: faster population growth leads to faster per capita income growth if and only if $\frac{\lambda}{\beta} > (1-a)$. Despite endogenonizing fertility—thus having a different set of differential equations—this inequality arises in Jones (2001). Rather than turn to data, Jones studies the case where this holds by assumption because his goal is to generate an industrial revolution where population and living standards grow together. Our contribution is to extend this result into a more

[20] As populations grow, A grows at a constant rate, but the percent change in the CES function is non-constant.

general, forward-looking setting where empirical moments can be brought to bear on whether this relationship is likely to hold moving forward in time.