Population Growth, Ideas and the Speed of History*

Kevin Kuruc†

February 1, 2024

Abstract

Leading models imply a positive relationship between population growth and per capita economic growth—a larger labor force produces more non-rival ideas, which propel per capita productivity improvements. This paper demonstrates that this relationship does *not* imply that faster population growth leaves individuals better off. The reason is straightforward: Faster population growth increases the arrival rate of both people and ideas, leaving each life unchanged. Rather than influencing the well-being of individual lives, the population size governs the *speed of history*. A larger population would merely give the illusion of increasing individual living standards, because we are used to measuring how these things advance per unit of time, not per human life. This reframing leads to a novel conclusion: in even ideas-based growth models with increasing returns to scale, the social value of population growth must come through increases in the quantity, not the quality, of lives.

Keywords: Population growth, economic growth, scale effects, aggregate welfare.
JEL Codes: J11; I31; O41.

*I would like to thank Gustav Alexandrie, Maya Eden, Mike Geruso, John Firth, Loren Fryxell, Tom Houlden, Antony Millner, Luis Mota, Toby Ord, Dean Spears, Christian Tarsney, Phil Trammel, Dietrich Vollrath, David Weil, Anson Zhou, and participants at Oxford University's Global Priorities Workshops for comments that have greatly improved the paper. All remaining errors are my own.

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

1 Introduction

The human population is set to enter an unprecedented phase of decline (Basten et al., 2013; United Nations, 2022). From the vantage of leading theories of economic growth, this is troubling. The ideas and knowledge that propel human progress are created by people. Larger populations, other things equal, create more of this non-rival knowledge, which increases per capita living standards. Jones (2022) formalizes that an unintended consequence of declining populations may be the end of economic growth, full stop. This line of thought has lead many to the normative conclusion that it would be preferable to avoid population decline.

This paper demonstrates that this normative conclusion may be mistaken, even under a standard positive population-economic growth relationship. The core insight comes from focusing on living standards for people rather than time periods. To see how the argument works, imagine that we can identify individuals by their birth order (i.e., $i = 100$ is the 100th person ever born). By the time the 100 millionth person lives, by definition, there will have been 100 million potential innovators who lived before her. If cumulative human research effort prior to her existence is what matters for her quality of life, her life will not have been improved by large historical per-period populations and fast economic growth. What would change is merely the date of her birth. In short, I show that population size should be thought to govern when existences and innovations happen—*the speed of history*—not living standards.

I begin by drawing out the core mechanism in the simplest possible setting. The key assumption on economic growth is that the level of innovation in period t is determined by the cumulative number of people-years prior to period t. This is not a novel assumption: it is equivalent to assuming a standard parameterization of standard endogenous growth models. However, this reframing makes immediately clear the normative point: By the time the ith person lives, the amount of knowledge that they have access to is invariant to prior per-period population sizes. A larger population size brings forward the date at which certain ideas are discovered, but it also brings forward every individual's existence. Therefore, individuals' living standards are invariant to flow (i.e., per-period) population sizes, even if the macroeconomy exhibits this well-known channel of increasing returns to scale.

This does not imply that social welfare is independent of the flow population size—a larger flow population may allow more individuals to exist. Moreover, I show that these marginal existences are the highest quality lives. Any reasonable social welfare function will be increasing in the number of the very-high-quality lives lived. So, flow population sizes can plausibly increase

aggregate social welfare by making possible these marginal existences.

However, whether larger flow populations lead to a larger timeless (i.e., cumulative, historical) popualation depends on facts about human extinction. Consider the simplest case where our civilizational lifespan is limited by an exogenous event, such as an existential asteroid, or some other cosmic event. Then, a larger flow population corresponds to a larger timeless population. There are more people per period, for a fixed number of periods. On the other hand, suppose that extinction is endogenous. That is, suppose we invent some technology that causes our end (e.g., a weapon of mass destruction). Here, the timeless population may be independent of the flow population because our technologically-driven extinction is plausibly brought forward just as fast as the arrival rate of people and technologies. Taken together, the social value of additional people *today* depends on facts about what will cause our extinction and what the quality of life will be just prior to this event.

The simple model understands the implications of endogenous growth in a novel and intuitive way: larger populations speed up human history. When seen through this lens it is clear why facts about our extinction will be relevant. We experience *more of our potential history* before an exogenous extinction event if human history is moving faster. This will be valuable to the extent that the lives occuring at that time are good. If extinction is endogenous, speeding things up merely compresses the entire human story into a shorter window. In either case, no one's life is made better via the faster economic growth, so larger populations can only be valuable insofar as they allow more people to ever exist.

I then move to a more general version of the simple model. This allows me to reinterpret the implications of existing models (e.g., Jones, 2022) as well as explore a range of common parameterizations. Specifically and least obviously, non-linearities in the pass-through from per period populations to per period economic growth do not change the main takeaways. Likewise, continuous and stochastic hazard functions for humanity's survival preserve the points from the simple model. The unintuitive normative implications of this class of growth models are not sensitive to the simplifications I use to highlight them.

This paper contributes to literatures in both economic growth theory and long-run welfare economics. Since Romer (1986, 1990) increasing returns to scale has been at the center of long-run studies of economic growth. Jones (1995) built on this insight to highlight the importance of population growth (see also Jones, 2003, 2005, 2022). To summarize this trailblazing line of work: exponential growth in aggregate research inputs can be maintained only if the population grows exponentially, so the rate of population growth is the *only* variable relevant for long-run economic

growth.

This stronger claim has been met with some skepticism, and alternative models have been built whereby slower population growth can drive faster economic growth. Specifically, if human capital investments per-child are higher in smaller families this could offset the relative decline in population sizes (Galor and Weil, 2000). However, it is in general extremely difficult to avoid the implication that a larger population level would have higher living standards if any knowledge that is produced is non-rival (Jones, 1999). This weaker claim is what is necessary for the arguments in this paper.

This weaker claim about population size and living standards has recent empirical support. Peters (2022) studies exogenous increases in population sizes coming from refugee assignments after WWII. He finds that localities with assigned more people have higher long-run incomes. Bloom et al. (2020) and Ekerdt and Wu (2023) study whether increases in aggregate research effort contribute to industry-wide (and economy-wide) productivity increases—finding that they do. It is a short jump from that finding to the result that a larger population, other things equal, would generate faster research progress. Finally, Kremer (1993) includes a case-study of ancient population groups that were geographically dislocated following the melting of polar ice caps (which submerged land crossings). The groups that were larger in size post-separation had much higher rates of long-run technological progress.[1,2]

By asking about the implications for long-run welfare of these models—that is, the quantity and quality of human lives lived—this paper also intersects with a literature in applied welfare economics. Most relevant is Klenow et al. (2023), which accounts for social welfare gains coming from population growth versus income growth. They find that increases in the quantity of lives may matter more than the increase in living standards over this period. Separately, there is a recent literature interested in risks to humanity's survival that directly deals with the question of how many people will ever live. Greaves (2019) studies the case of climate change, where each human life has some non-zero carbon footprint in the long run. She shows that this implies that a fixed number of people can ever live: either the planet is warmed beyond a level that we could survive, or we run out of whatever the necessary dirty input is. The present study works with a premise

[1]The largest was the Eurasian continent, which had the most technological advancement at the time of connection being re-established. The smallest groups in fact died out, as they had trouble even *maintaining* technologies that were discovered and shared pre-separation.

[2]Interestingly, anthropologists have separately come to believe there must be a positive relationship between population size and living standards, coining the term "the collective brain" to mean the aggregate mental capacity of a society, which is increasing in the number of individual brains (Henrich, 2004).

that is more relevant for economists (and more realistic), but generates analogous takeaways using a similar style of reasoning. Ord (2024) develops a very general framework for assessing the long-run welfare implications of changes to the trajectory and survival probability of humanity. The lessons of the current paper could be ported into Ord's framework, which has proven useful for verifying and better understanding the mechanisms and generality of the specific population size channel of interest here. Finally, Ashenbrennar and Trammell (2024) study an endogenous growth model with existential risk that can be endogenously mitigated. Much of their inquiry is distinct, but they experiment with population changes as a mechanism for speeding up growth and endogenously increasing our survival probability.

2 Ideas-based growth models do not imply a positive relationship between population and individual living standards

This section presents the simplest model that illustrates the main point of the paper. The key assumption is that the level of technology in a given period is a function of cumulative people-years lived prior to that year. I start by showing this assumption is equivalent to a special case of endogenous growth models, and that it is not unreasonable given the historical record. I then show that it implies that each individual's income is invariant to historical flow population sizes, even though larger flow population sizes increase the rate of productivity growth. Finally, I show that whether an increase in the size of the flow population is welfare increasing depends on views about existence-value and facts about how the model is closed (i.e., how humanity is assumed to end).

2.1 Cumulative innovation is a function of cumulative people-years

Equation 1 is a production function of productivity improvements (i.e., *ideas*) that is standard in the semi-endogenous growth literature (e.g., Jones, 2022).

$$\frac{\dot{A}(t)}{A(t)} = \alpha(t)(s(t)N(t))^{\lambda}A(t)^{-\beta} \tag{1}$$

Time, t, is continuous; $\dot{A}$ is the instantaneous change in productivity, A. This depends on three factors. First, the number of invididuals engaged in research. In (1) this is decomposed as the share of the population in research, s, multiplied by the population size, N. Second, the productivity of

these researchers, α. Third, the accumulated knowledge stock $A(t)$, which allows for non-linearity in ideas production. This could speed up productivity improvements ($\beta < 1$) if we use past inventions (e.g., the computer) to help generate new inventions, or it could represent a slowing down ($\beta > 1$) if earlier discoveries are systematically easier to make. In practice, $\beta > 1$ matches the historical data better, both in aggregate and at the industry level (Bloom et al., 2020).

Notice that the number of researchers is raised to some power λ that determines whether there are diminishing returns to researchers within a period. There is considerable uncertainty about what value λ should take. One reason to believe that $\lambda < 1$ is *duplication*; people alive at the same time might solve the same problems, resulting in wasted effort. On the other hand, *collaboration* might be a reason that $\lambda > 1$; the ability to communicate with other researchers might result in faster progress than a counterfactual where those same researchers have non-overlapping lives.[3]

The high-level intuition behind this equation is simple: if ideas are produced by people, then a larger population ought to produce more of them. Productivity increases should be faster, other things equal, in periods with larger populations. Romer (1990) recognized that non-rival ideas would give rise to increasing returns to scale. Work since then has confirmed just how difficult it is to neutralize this scale effect in innovation-based growth models (Jones, 1999).

For expositional simplicity, let $\lambda = 1$ and $s(t), \alpha(t)$ be some constants $\bar{s}, \bar{\alpha}$. The assumption of constant α, s, is employed because the objective is to isolate the effects that population size changes have. Of course it is true that increasing α or s would also increase economic growth rates, but that is not the focus of this paper. The assumption of $\lambda = 1$ will be more substantive. It imposes a constant marginal effect of population within a period (i.e., if the world population were twice as large, exactly twice as much progress on productivity improvements would occur). In Section 4 I will explore how the implications of the model change for other values of λ.

With these assumptions, we can rewrite (1) as follows, where $\theta = \bar{\alpha}\bar{s}$.

$$\frac{\dot{A}(t)}{A(t)} = \theta N(t) A(t)^{-\beta} \tag{2}$$

Integrating both sides with respect to t delivers the following expression for the level of A.

$$A(t) = \left(\beta\theta \int_0^t N(\tau)d\tau + A_0^{\beta}\right)^{\frac{1}{\beta}} \tag{3}$$

[3] In reality, this relationship likely has an 'S'-shape, where collaboration benefits dominate at small populations, but duplication issues dominate for large populations. It is unclear which dominates for current population sizes.

Equation 3 captures a core idea leveraged throughout the remainder of the paper. It indicates that the level of $A(t)$ is determined by $\int_0^t N(\tau)d\tau$. That is, cumulative innovation by period t is pinned down by cumulative people-years lived by period t. This is not how this class of models is typically described, but the intuition is straightforward. If innovations and ideas contribute to a non-depreciating, non-rival stock of knowledge that is built over time by people, the stock of available ideas will be larger after more people have ever lived. Notice that this does not need to be a linear relationship. In the case where $\beta > 1$, a person-year lived when the knowledge stock is large will generate a smaller increase in A.

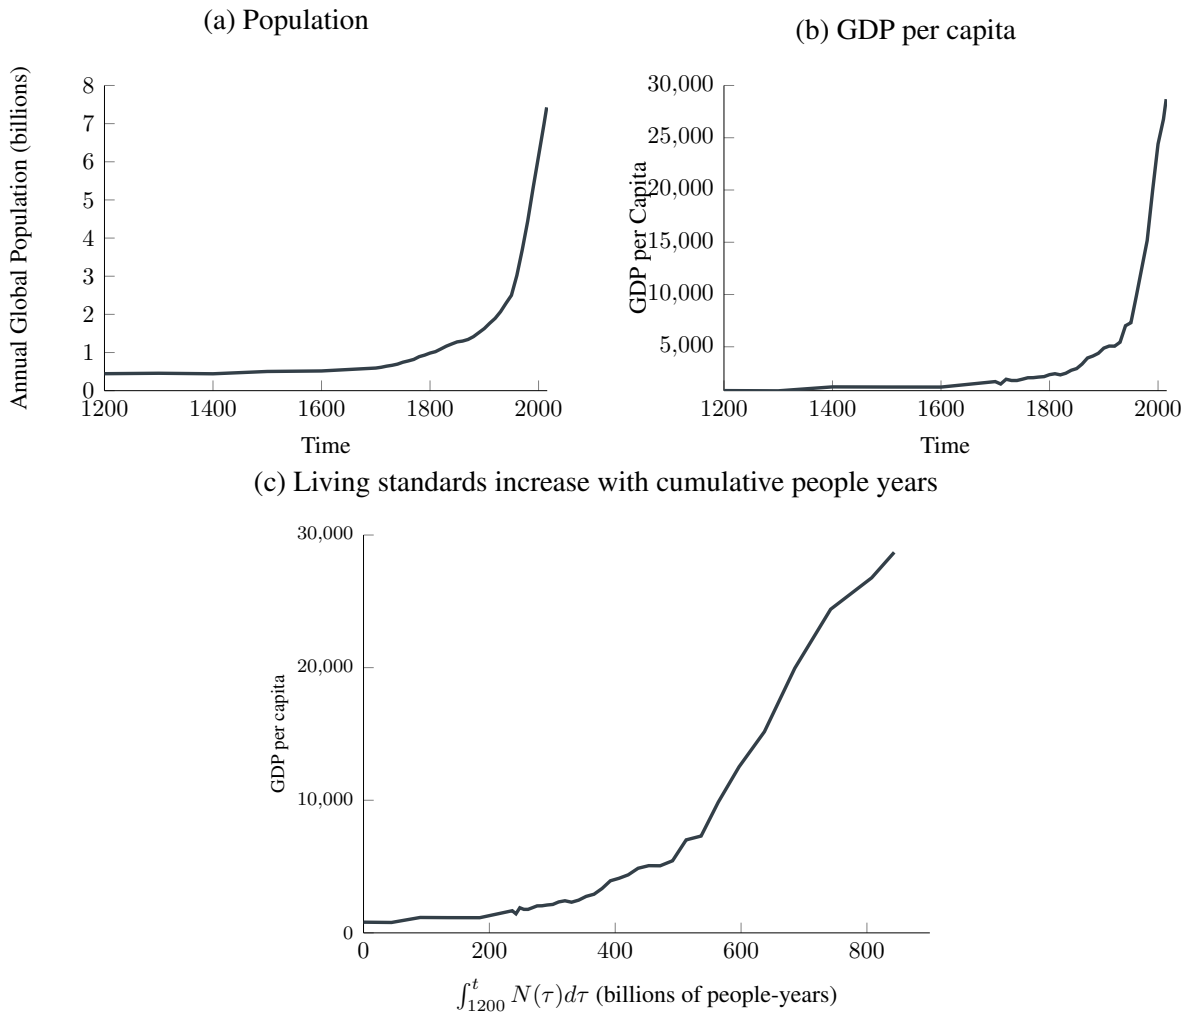
Substantively, what Equation 3 does is eliminate t as an independent variable. It is cumulative human effort, not time, that drives innovation. Time is correlated with innovation *because* it is correlated with cumulative human effort. So, we can reinterpret this equation to eliminate time without losing anything of conceptual importance.

Figure 1 demonstrates this elimination of variables in practice. Panels (a) and (b) plot these variables in the standard way: populations and living standards change as time passes. Panel (c) graphs cumulative people years against living standards (which I take as a proxy for cumulative productivity gains). The pattern in (c) is not itself decisive evidence in favor of Equation 3; other models could generate this broad pattern. However, as argued in an influential paper by Michael Kremer (1993), this class of models does an excellent job of matching the broad historical fact that for most of human history populations were small and progress was very slow, until recently when populations and productivity growth simultaneously exploded. Likewise, increasing returns via the non-rivalry of knowledge are at the center of more sophisticated theoretical efforts to endogenize the industrialization of humanity (Galor and Weil, 2000; Jones, 2001, etc.).[4]

To reiterate before moving on to the normative implications of these models: the results and analysis that follow will not rely on the stronger claim that population growth is the *only* source by which we could sustain long-run economic growth Jones (1995). I am interested in whether a positive relationship between population size and economic growth gives us reason to prefer faster population growth, or larger long-run populations. I find it does not, even when focusing on a model in which population increases are the only thing that can generate long-run growth. Therefore, deviations that allow for non-population-driven growth could only strengthen the conclusion that this relationship is normatively irrelevant.

[4]Unified Growth Theory (Galor and Weil, 2000) emphasizes the role of human capital as well, but nonetheless has at its core an equation whereby a larger population creates more non-rival ideas that propel progress. An open question is whether population *growth* increases growth rates, because changing the growth rate (rather than the level) may influence per capita human capital investments.

Figure 1: Income growth tracks cumulative people-years better than time

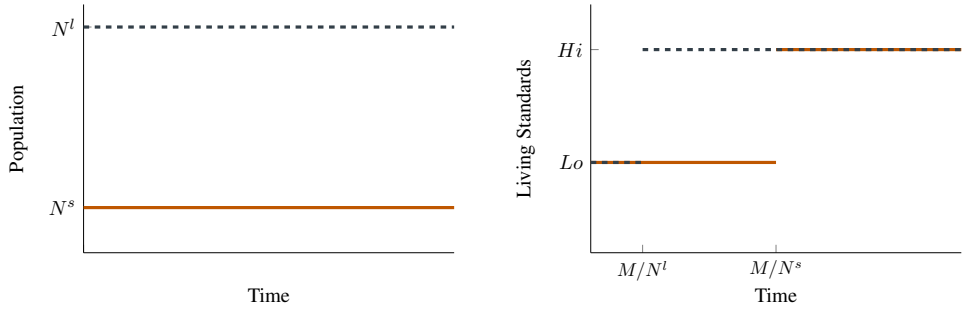


2.2 Population sizes govern the *speed of history*

This section will work with the most transparant version of a model that builds from the result of the prior sub-section: cumulative people years pins down cumulative innovation. Time, t and the flow population size, $N(t)$, are continuous. There are two possible technology levels, $A = \{Lo, Hi\}$, representing periods with low or high per capita living standards. In accordance with the results from the prior section, A increases from $Lo \rightarrow Hi$ after some fixed number of people-periods, M, have been lived. For ease of exposition I will assume that each individual lifetime is only one period, so that M is a measure of lifetimes that need to be lived before humanity's accumulated knowledge is enough to transition from $Lo \rightarrow Hi$.

Figure 2 depicts this simple economic environment for two hypothetical population sizes. For

Figure 2: Larger populations speed up economic growth



simplicity I will depict these graphs for some fixed $N^s(t) = N^s < N^l = N^l(t)$. In the history with a large population, the jump from $Lo \to Hi$ happens earlier. Equivalently, economic growth is faster in the larger population history.

Now, imagine that we can identify each person i in history by the order of their birth. In other words, we are counting the number of people who have ever lived as they are born. So far there have been about 100 billion people who have ever been born, so each of us alive now would be i's in the 100 billions. Our grandparents were the lower is; our grandchilren will be the higher is. Consider the living standard of any i in the high and low population histories of Figure 2. A larger population history brings forward the is just as fast as it brings forward the jump to Hi. What is different is merely the date at which an individual exists. Figure 3 depicts this case by graphing living standards for each i, rather than for each period t. From this vantage, the two histories are exactly identical, aside from the speed at which history has progressed.

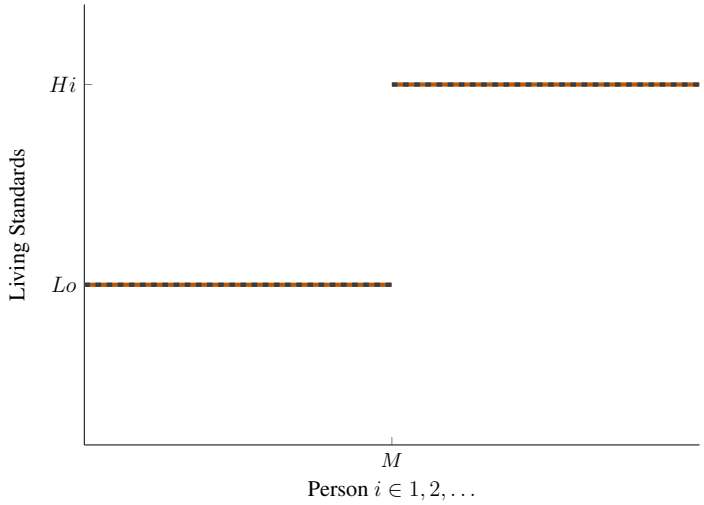
Let's refer to this as Theorem 1, where I will denote person i's per capita income during their lifetime as $y(i|t) \propto A(t)$ (where this value takes 0 or is undefined for periods when the individual does not exist, and is proportional to A in the period the individual lives).

Theorem 1. *Assume that productivity improvements scale linearly (within a period) with the size of the population, N. Then, an increase in N increases the rate of economic growth, but is independent of $y(i)$ $\forall$ i.*

Theorem 1 immediately delivers a corollary that will be the starting point for doubting the normative relevance of the population-economic growth channel.

Corollary 1. *Under these same assumptions, if $\sum_t N(t)$ equals some fixed $\tilde{N}$, then (undiscounted) aggregate welfare, $\sum_{t=0}^{\infty} \sum_i y(i,t)$ is invariant to $N(t)$ $\forall$ t.*

Figure 3: No individual has a better life because of the faster growth



Corollary 1 states that if we have some fixed timeless population—that is, a fixed number of individuals who will ever be born—then the sum of welfare over time is invariant to the number of people alive in any given period. In other words, it would not have decreased aggregate timeless welfare if the 1900s had half as many people, and the rate of technological progress had been halved over this period.

This corollary is more striking when viewed through a slightly different lens. *If* aggregate welfare is unaffected by flow population sizes, conditional on the timeless population, *then* any welfare effects must come via changes in the timeless population. This is a novel understanding of the implications of the non-rivalary of ideas that are endogenously produced: a larger population merely gives the illusion of increasing individual living standards, because we are used to measuring how these things change over time. If a larger population is to have aggregate welfare effects, it is going to be because it causes more good lives to be lived.

2.3 Discounting, Extinction, and the Value of Existing

Corollary 1 leaves open two avenues for population sizes to influence aggregate welfare. First, and most familiarly, we might discount the future. Effects closer to us in time may matter more. If so, *speeding up history* would be socially valuable. It is easy to verify that bringing forward the jump from $Lo \to Hi$ will increase discounted aggregate welfare.

There are three main reasons standard models discount future consumption gains. First, given economic growth, future people will be richer than us. So it is not a time effect, so much as

a diminishing marginal utility effect. But here I am considering utilities directly, so this is not relevant. Second, there is an expected value argument. If there is some chance humanity goes extinct, for example, then any benefit we attempt to confer on our future has some probability of being irrelevant. Here I will discuss extinction directly, so I do not need the discount rate as a reduced-form stand-in for extinction risk.[5] Finally, there is the empirical fact that people are impatient, so perhaps the social welfare function should inherit this impatience. There is broad consensus that the first two reasons are legitimate reasons to discount, but much less over the third. Imagine that we knew for certain that 500 billion people would ever live, and that we knew exactly the quality of their lives. Could it possibly matter—from an impartial planner's perspective—whether those lives happen earlier or later in chronological time? That is the style of discounting relevant for this paper. I follow a tradition in public and welfare economics that rules out this form of discounting.[6]

Setting aside pure time discounting amounts to focusing only on the quantity and quality of lives that are ever lived, in expectation. Theorem 1 makes clear that, for the set of people who come into existence with certainty, life is equally good regardless of the flow population size. However, it also seems empirically straightforward that more people will ever exist if the flow population size is larger. To make this conjecture precise, we need to take a stance on what constrains existences—i.e., we have to make assumptions about human extinction.

2.3.1 Exogenous Extinction

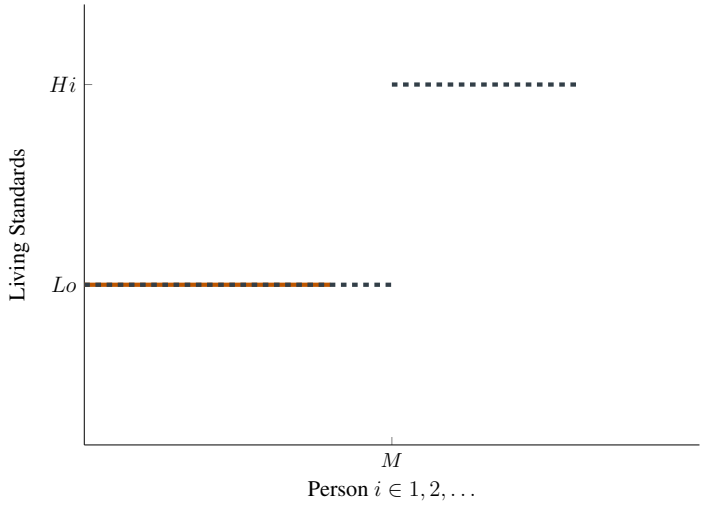
The simplest case is exogenous extintion. Whether this is deterministic or stochastic makes no difference in this simple setting, so I will start with deterministic extinction in period T^X. Recall the *speeding up history* framing: larger flow populations increase the arrival rate of ideas and people, and thus compress more of an equivalent history into less time. The result with exogenous extinction is that we get more of our potential history by T^X when flow populations are larger.

To see this graphically, consider Figure 4. Here I've again assumed the two population sizes under consideration are some fixed N^s, N^l. Furthermore, imagine that T^X is such that $N^s \times T^X < M < N^l \times T^X$; the smaller population does not generate enough cumulative people years prior to extinction to progress to $A = Hi$. The large population does.

[5] I will also assume away uncertainty about our effects on the future for expositional simplicity.

[6] Conversely, if one did want to employ a positive pure rate of time discounting, larger populations in this setting are something of an aggregate welfare loophole: without making any life better, aggregate welfare can be increased by moving up the date that each of these lives is lived.

Figure 4: With exogenous extinction, larger populations get through more of history



These curves overlap for all $\int N(\tau)d\tau \leq N^s \times T^X$. At this point, the small population has gone extinct. The large population proceeds on the same path until $N^l \times T^X$. Humanity progresses through more of its potential history before the period of extinction.

Corollary 2. *The effect of a marginal increase to the population—in any or all periods—is to increase the number of lives lived just prior to the time of extiction.*

Consider a some small ε increase in the population size in this setting. For concreteness, focus on N^l increasing to some $N^L = N^l + \varepsilon$. In our usual way of looking at things, it would appear as if we are increasing the size of the population in each period holding fixed per capita well-being. But consider what this increase looks like using the logic of Figure 4. Humanity would get through just a bit more of the potential history prior to T^X. This implies that the marginal existinces from flow population changes are the "best" lives (assuming A maps to well-being).

This will be generally true for a one-time pulse to the size of the population in any year. The x-axis of these figures—which drives progress in A—is cumulative people years, independent of year lived. Increasing the population size by one today speeds everything up by one person, allowing one more person to live prior to T^X. So, an additional person today has the effect of one additional person living at the back-end of history. As strange as it may seem, the social value of an additional child today depends on what life will look like for humanity at the time of exogenous extinction.

2.3.2 How does aggregate social welfare respond to additional people?

Whether population size has *any* normative effects will depend on how we value additional existences. There is an entire field devoted to understanding the public good value of existences that I will draw on here. One natural option is to take a firm stance that existence cannot be compared to non-existence. We might want our social welfare function to only compare a function of the well-being of people that exist in both options under consideration.

This approach runs into a few problems, making it an undesirable choice. Most significantly, it is not clear how to select individuals from the larger group who count as "existing" regardless of the option chosen. Any option that makes the population size slightly larger in any given period is going to change everyone's identity. Consider the earlier example where the world population was half as large in the 20th century. The result would not be that half of us would have been conceived and birthed—an entirely different set of people (of about 4 billion) would inhabit the planet. So, we might compare the first $N^s \times T^X$ people in the large population, but that would be a bit arbitrary. Plus, in these cases the social welfare function would not increase if the "later" people had flourishing existences rather than awful existnces.

In any case, if the welfare function was indifferent to these additional lives—regardless of their level—then it is straightfoward that there is no social value from larger flow population sizes. It is worth noting that this differs from the usual assumed implication of a positive relationship between population sizes and economic growth. The standard interpretation is that these non-rival goods improve everyone's well-being and so larger populations would be preferred regardless of your stance on existence value (e.g., Budolfson et al, 2022).

Because most of the literature rejects this view of existences, debates center around *Average Value* and *Total Value*. *Averagism* takes social welfare to be the average value among everyone who ever exists. Again using $y(i)$ to represent well-being of person i, this would take the form:

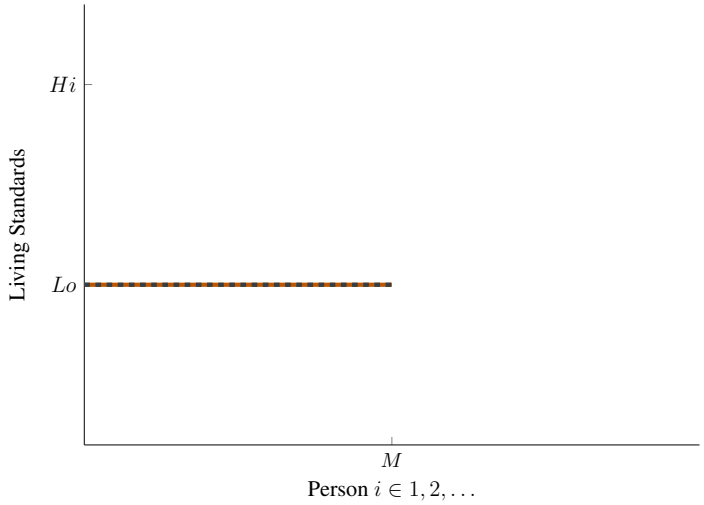
$$V^{avg}(\{y(i)\}) = \frac{1}{N}\sum_{i=1}^{N} y(i). \tag{4}$$

Totalism merely sums the well-being of anyone who ever lives[7]:

$$V^{avg}(\{y(i)\}) = \sum_{i=1}^{N} y(i). \tag{5}$$

[7]Note that this implicitly normalizes the value of non-existence to $y(j) = 0$. If that is the level of well-being you experience, it is equivalent to not existing, according to this additive welfare function

Figure 5: With endogenous extinction, total existences are invariant to per-period population sizes



Corollary 2 tells us that functionally, any marginal person comes at the end of history, when $y(i)$ is at its peak. Therefore, under either *Averagism* or *Totalism* an increase in the size of the population increases social welfare. Human history includes slightly more of the very best lives that we will ever achieve, which will be valuable on any theory of aggregate social welfare.

So, as long as the effect of a larger flow population is to increase the number of people who ever live, larger flow populations will have social value. Our understanding of this social value is dramatically revised—it is because humanity lives more of its potential history—but it is nonetheless valuable. What could break this relationship is a breaking of the link between flow populations and total existences.

2.3.3 Endogenous Extinction

Leading scholars of extinction risk believe that nearly all of the risk that we face is *endogenous* (Rees, 2003; Ord, 2020). That is, human extinction will be caused by some technology being used by some human beings. For example, suppose that humans continue to use emit some fixed, non-zero, amount of CO_2 each period. There is some amount of warming that would drive us extinct, and it would be reached after some fixed numbers of humans have lived according to this model (see Greaves, 2019). Alternatively, we have weapons of mass destruction that threaten to destory us if individuals set enough of them off. Some observers worry that the invention of sufficiently powerful artificial intelligence will lead to our doom (see e.g. Jones, 2023).

Put simply, these models, to a first-order, break the link between flow populations and total

populations. To see this, consider again the framing where population sizes govern the speed of human history. If extinction is endogenous to the progress of human history, then its arrival will be brought forward at the same rate that existences and innovations are. This is easiest to see in a case like artificial intelligence: imagine in our simple model that once we reach $A = Hi$, humanity ends quickly thereafter. Then, M people will live, regardless of how large or small populations are. We are either spread out over many years, or packed in over few, but the same number of lives happen.

Figure 5 graphs this case. Both populations get just to M people-lives, and then go extinct. This picture is so simple precisely because time is not graphed. These histories play out on different timelines, but under the assumption that every individual—conditional on A—advances progress by the same ex-ante amount, we will reach advanced artificial intelligence after the same number of existences. Given the intentional simplicity of this model, I will not formalize endogenous extinction in this setting more than that. In the section that follows the main lessons will be shown to generalize.

To recap what these lessons are:

(1.) In standard endogenous growth models, cumulative people-years lived prior to an individual's existence determines the level of technology available at their birth $\Rightarrow$

(2.) The size of the flow population changes how fast we get to technological milestones and cumulative existences, but not how good each of those existences is $\Rightarrow$

(3.) Any potential benefits coming from increases in flow populations sizes must come through increases in the size of the timeless population of individuals who ever exist.

(4.) Moreover, whether (and by how much) we should expect larger flow populations to pass through to larger timeless populations depends crucially on the form that existential risks to humanity take.

3 Extensions to continuous models of A and existential risk probabilities

This section demonstrates that the four main lessons of the paper generalize to more realistic settings of continuous technological progress, non-constant population sizes and stochastic existential

risk. Additionally, it revists the conclusion of Jones (2022), wherein living standards stagnate under a regime of population decline but grow towards infinity otherwise.

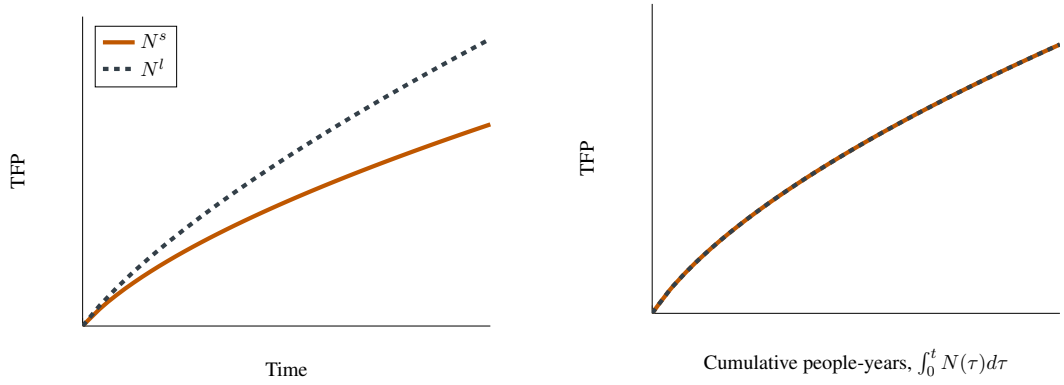
3.1 Continuous versions of the simple discrete model

Return again to the result in Equation 3 describing the continuous version of $A(t)$. This is the format most growth economist are familiar with.

$$A(t) = \left(\beta\theta \int_0^t N(\tau)d\tau + A_0^{\beta}\right)^{\frac{1}{\beta}}$$

It should not be surprising that the main results pass through to the continuous version of the model. The results rest on the fact that cumulative people-years pin down technological progress in a period. Taking again two constant populations, N^s, N^l we get the following two analogous figures plotted in two panels of Figure 6.

Figure 6: Lessons extend to continuous setting



The left panel plots how living standards would diverge over time.[8] The right panel shows, unsurprisingly, that each individual has the same living standard.

It is less obvious how this extends to more realistic cases of extinction. Our two extreme cases retain their original interpretations for straightforward reasons. First, if extinction happens exogenously with some probability each period, then each population history has the same expected number of periods to live. It will remain true that the larger population traverses more of this potential history. On the other extreme, if there is some level of A such that we go extinct once

[8]Notice that growth is not exponential here. That is because I am assuming a constant population with $\beta > 1$, which implies a constant population discovers fewer and fewer ideas as time goes on (ideas are getting "harder to find").

we reach it—e.g., if we invent a hyper-intelligent AI that wipes us out—this will happen after the same number of cumulative people years regardless of population sizes.

There are other ways we might imagine endogenous extinction occurring, however. I will consider two here. In one specification A is inherently dangerous. As we invent new technologies, there are more ways for one of us to kill everyone. Conversely, A might be inherently safe going forward. That would correspond to a case where our saefty-promoting technologies progress faster than our dangerous technologies (e.g., if vaccine technology progresses faster than our ability to design dangerous pathogens). In each case I assume that our risk per period scales with the population. Other things equal, a world population of 1 billion is much less likely to have a rogue actor that attempts to kill everyone, or a conflict between two or more parties that spirals into an existential conflict, than a world population of 10 billion.

$$P(survive(t)|alive) = \frac{1}{1 + \theta N(t) \times A(t)^{\phi}} \tag{6}$$

$$P(survive(t)|alive) = \frac{1}{1 + \theta N(t) \times e^{-\phi A(t)}} \tag{7}$$

Equation 6 is the case where A increases make life more risky; Equation 7 is the case where our probability of extinction falls as A increases.

Rather than provide analytical solutions for the surivial probabilities per period, or the survival probability that a given i faces, I provide computational solutions to simulations of these histories. It turns out that under these forms of existential risk, the exact same number of people live, in expectation, regardless of the flow population size.

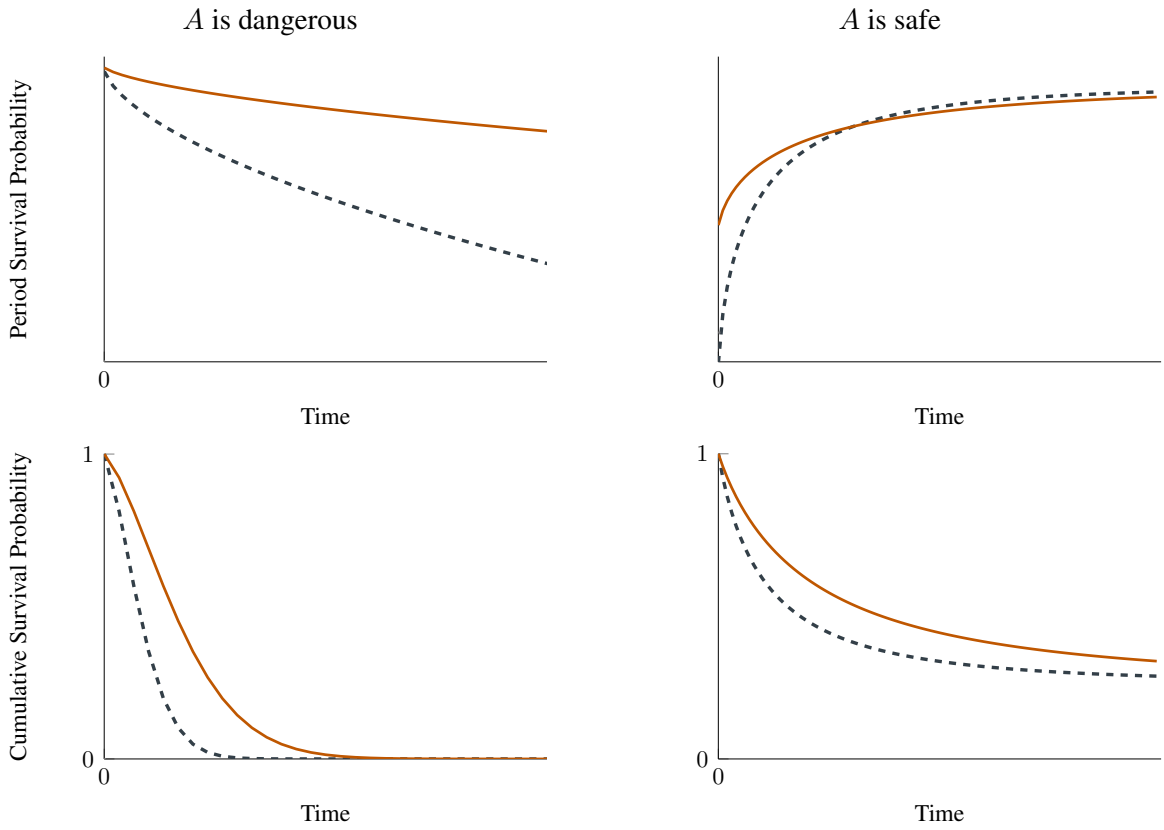
3.2 Does economic growth end if populations decline?

This paper is in large part motivated by the striking findings of Jones (2022). Jones demonstrates that the future path of living standards is qualitatively different under different regimes of population growth. Namely, declining populations have stagnant long-run incomes with an upper-bound, whereas constant or growing populations have incomes that continue to increase indefinitely.

That appears inconsistent with the findings here. I claim that different population sizes merely change the time at which progress occurs and people exist. Here, I show the claims in each of these papers are in fact consistent with one another.

The key is to recognize that there is an implicit extinction assumption in Jones (2022) that can account for the qualitative differences. The declining population is one with constant negative

Figure 7: If populations are smaller, we survive longer in expectation



population growth. This implies exponential declines in population until there are essentially zero people on the planet.[9] If the population exponentially decays towards zero people, cumulative people-years will be finite even over an indefinitely long future.

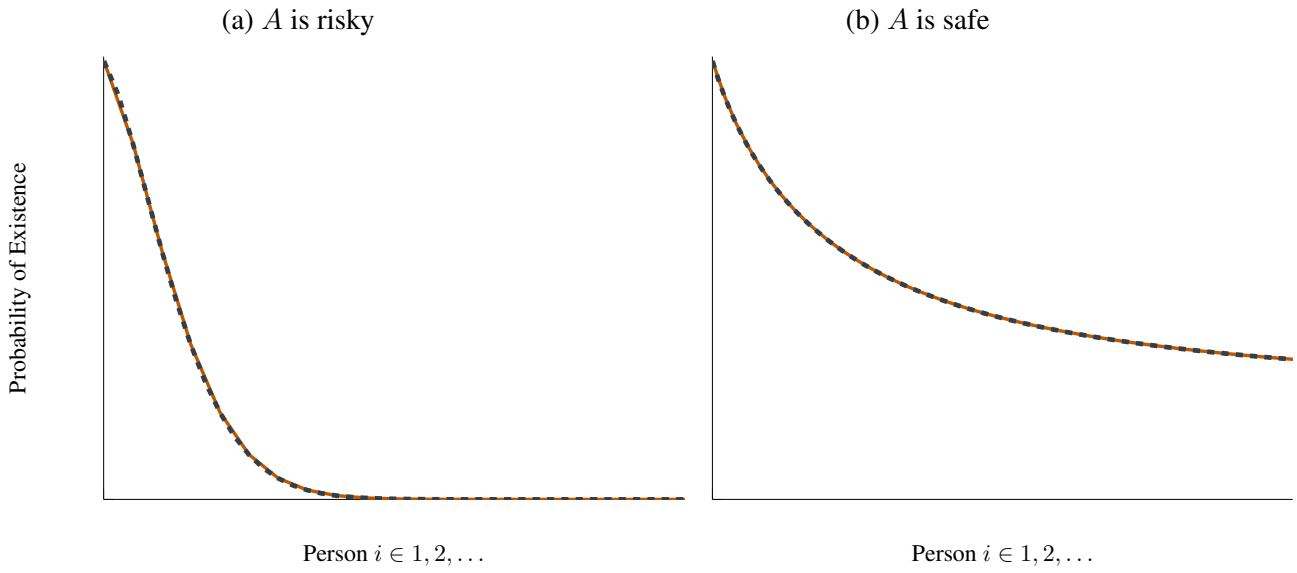
When seen through this lens the result becomes notably more believable: if there are zero people on the planet in the limit, of course there will be no technological progress in the limit.[10] Where else could the progress come from? On the other hand, it leaves one to wonder if economic growth is really the thing to worry about—as opposed to our extinction—regarding the possibility constant below-replacement fertility rates. The "end of people" seems more troubling than the end of economic growth.

To put this in the context of the now-familiar graph with cumulative people-years on the x-axis, Figure 9 plots Jones' *Expanding Cosmos* outcome (with positive population growth) against

[9] Jones understands and does not hide this fact: he denotes that future the *Empty Planet* equilibrium.

[10] Technically, the population could go to zero slowly enough that A never reaches an upper-bound for some function forms, so this is not universally true. Nonetheless, this is what happens for population decline that results in a finite integral over the infinite future.

Figure 8: Timeless population size invariant to flow population sizes with endogenous extinction



his *Empty Planet* outcome (with negative population growth). The difference is that the *Expanding Cosmos* outcome traverses much more of our potential history.

4 Changes to intra-period returns to population size, λ

This section demonstrates that the assumption of within-period linearity between population size and TFP growth does not drive any of the qualitative takeaways thus far. This simplification may overstate the quantitative degree of some of the claims, but not the directional takeaways that are the contribution of this paper.

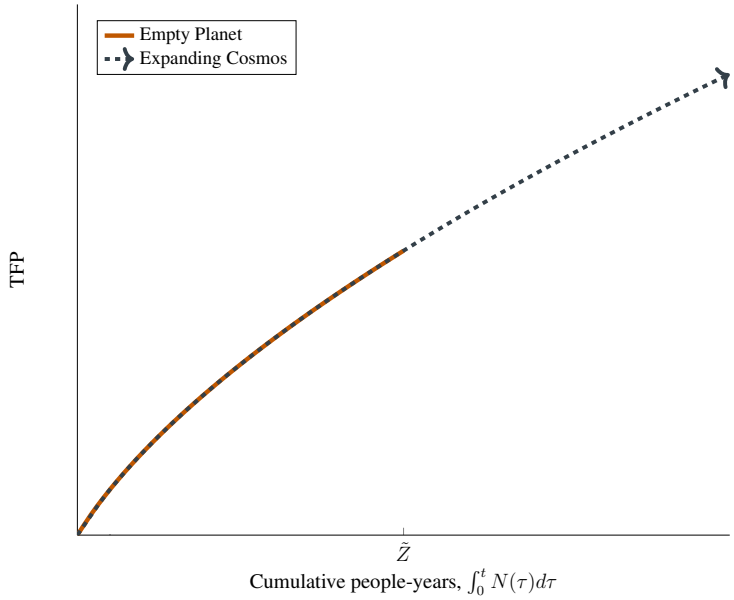
Below I reprint for convenience the standard endogenous growth equation that I began with, including the simplification that $\alpha(t), s(t)$ are fixed constants.

$$\frac{\dot{A}(t)}{A(t)} = \theta N(t)^{\lambda} A(t)^{-\beta}$$

The assumption that $\lambda = 1$ is the linearity assumption that I have worked with for simplicity. However, this a knife's edge condition where agglomeration effects exactly offset duplication conerns.[11] Here I will consider cases where $\lambda \neq 1$.

[11] Recall that duplication concerns are that if more people are researching in any given period, the probability of overlapping findings is higher, so it is not a one-to-one pass through from population to new non-rival knowledge.

Figure 9: The End of Economic Growth, or The End of People?



It will be easiest to see the effect that λ has by first considering the problem of allocating a fixed number of existences across time. Suppose that for some reason there will only ever be 100 lives lived. If $\lambda > 1$, agglomeration effects dominate and per capita living standards are maximized by having all 100 lives happen in the same period. More generally, on any margin, it will be true that a 1% increases in per period population sizes will increase per capita living standards for those 100 individuals. Conversely, if $0 < \lambda < 1$, long-run technological progress is maximized by spreading those 100 individuals out in non-overlapping lives. What is interesting here is that if $\lambda < 1$, which is by no means a controversial parameterization, and the timeless population is fixed, then the Romer-channel in fact prescribes *smaller* populations to maximize per capita living standards.

However, this takeaway is not robust to more realistic cases where flow populations influence the size of the timeless population. This will be easiest to see in constant population cases, where I will consider marginal long-run increases. First, consider the case of exogenous extinction with an asteroid striking in period T^X. Imagine some baseline N compared against some small increase $N + \varepsilon$. It will be true that for each $t \in [0, T^X]$ living standards are slightly higher in each period for any $\lambda > 0$. Each period has slightly more people, which translates into slightly more R&D output.[12] So, if each period has higher living standards and the population is uniformally distributed

[12]The dynamics are non-obvious because of the $A^{-\beta}$ on the right hand side. But it is well known that a 1% increase in long-run research intensity passes through to long-run productivity at a rate of $\frac{\lambda}{\beta}$%. So it will in general be true that in each period the knowledge stock is slightly larger for a slightly larger long-run population.

across time in each case, it must be true that timeless per capita living standards are higher in the larger population world under exogenous extinction.

Under endogenous extinction things become very counterintuitive. Suppose we are dealing with a simple case where A^x is a level of A such that we go extinct once it is reached. If $\lambda < 1$ the planner may prefer large per period population sizes precisely because they make less progress than if that same number of lives were spread out over time. So, the number of lives lived are maximized by putting people in the same period *because* their effect on technology is smaller in those periods. I do not know how seriously to take this possibility. If nothing else, it calls for more research on λ and the sources of exinction risk that humans face.

5 Discussion and Conclusion

The size of the human population is almost certainly going to begin shrinking later this century. Whether humanity stabilizes at 8 billion, rather than 4 billion, for example, will likely influence rates of economic growth. At first glance, the economic growth alone appears to be reason enough to prefer stabilizing at larger population levels.

This paper shows that this is not in fact the case in standard models that generate a positive population-productivity growth relationship. These models imply that, because the arrival rate of ideas depends on the arrival rate of individual lives, a larger population leaves no one better off. Improving living standards via additional economic growth is not a reason to prefer larger populations.

However, a larger population could improve social welfare if it leads to additional existences. This is a more contentious claim, but it is difficult to construct a coherent social welfare function that does not respond positively to additional high-quality lives. Therefore, the key question is whether a larger per-period population leads to a larger timeless population. Under most, but notably not all, models of existential risks that we face, more people in any period increases the (ex-ante) timeless population. So, while larger populations may not be valuable for the reason of additional economic growth, they are almost certainly beneficial insofar as a larger population increases the number of lives ever lived.

References

Basten, S., Lutz, W. and Scherbov, S. (2013), 'Very long range global population scenarios to 2300 and the implications of sustained low fertility', *Demographic Research* **28**, 1145–1166.

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.

Ekerdt, L. and Wu, K.-J. (2023), 'Self-selection and the diminishing returns to research', *Working Paper* . Available at https://sites.google.com/view/lorenz-ekerdt/research.

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

Henrich, J. (2004), 'Demography and cultural evolution: how adaptive cultural processes can produce maladaptive losses—the tasmanian case', *American Antiquity* **69**(2), 197–214.

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

Jones, C. I. (1999), 'Growth: With or without scale effects?', *American Economic Review* **89**(2), 139–144.

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. (2003), Popuation and ideas: A theory of endogenous growth, *in* P. Aghion, R. Frydman, J. Stiglitz and M. Woodford, eds, 'Knowledge, Information, and Expectations in Macroeconomics (in honor of Edmund S. Phelps)', Princeton University Press, pp. 498–521.

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. (2023), 'The a.i. dilemma: Growth versus existential risk'. Available at: https://web.stanford.edu/~chadj/existentialrisk.pdf.

Klenow, P., Jones, C. I., Bils, M. and Adhami, M. (2023), 'Population and welfare: The greatest good for the greatest number'. Available at: https://amirrorclear.net/files/shaping-humanity's-longterm-trajectory.pdf.

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

Ord, T. (2020), *The Precipice*, Hatchette Books.

Ord, T. (2024), 'Shaping humanity's longterm trajectory'. Available at: https://web.stanford.edu/~chadj/slides-popwelfare.pdf.

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

Rees, M. (2003), *Our Final Hour*, Hatchette Books.

Romer, P. M. (1986), 'Increasing returns and long-run growth', *Journal of Political Economy* **94**(5), 1002–1037.

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

United Nations (2022), 'World population prospects 2022', *New York, NY: United Nations Department for Economic and Social Affairs* .