

The End of Economic Growth? Unintended Consequences of a Declining Population

Charles I. Jones*

Stanford GSB and NBER

January 6, 2020 — Version 1.0

Abstract

In many models, economic growth is driven by people discovering new ideas. These models typically assume either a constant or a growing population. However, in high income countries today, fertility is already below its replacement rate: women are having fewer than two children on average. It is a distinct possibility — highlighted in the recent book, *Empty Planet* — that global population will decline rather than stabilize in the long run. What happens to economic growth when population growth turns negative?

*I am grateful to Sebastian Di Tella, Jesus Fernandez-Villaverde, Jeremy Greenwood, Pete Klenow, Chris Tonetti, and David Weil for helpful comments.

1. Introduction

In many growth models based on the discovery of new ideas, the size of the population plays a crucial role. Other things equal, a larger population means more researchers which in turns leads to more new ideas and to higher living standards. This basic feature is shared by the original endogenous growth models of Romer (1990), Aghion and Howitt (1992), and Grossman and Helpman (1991) as well as by the semi-endogenous growth models of Jones (1995), Kortum (1997), and Segerstrom (1998) in which standard policies have level effects instead of growth effects. It is a feature of numerous other models.¹


In a recent book entitled *Empty Planet*, Bricker and Ibbitson (2019) make the case based on a rich body of demographic research that global population growth in the future may not only fall to zero but may actually be negative. For example, the natural rate of population growth (i.e. births minus deaths, ignoring immigration) is already negative in Japan and in many European countries such as Germany, Italy, and Spain (United Nations, 2019).

Figure 1 shows historical data on the total fertility rate for various regions. This measure is the average number of live births a cohort of women would have over their reproductive life if they were subject to the fertility rates of a given five-year period. To sustain a constant population requires a total fertility rate slightly greater than 2 in order to compensate for mortality. The graph shows that high income countries as a whole, as well as the U.S. and China individually, have been substantially below 2 in recent years. According to the U.N.'s *World Population Prospects 2019*, the total fertility rate in the most recent data is 1.8 for the United States, 1.7 for China and for High Income Countries on average, 1.6 for Germany, 1.4 for Japan, and 1.3 for Italy and Spain. In other words, fertility rates in the rich countries of the world are already consistent with negative long-run population growth: women are having fewer than two children throughout much of the developed world.

A sharp downward trend in India and for the world as a whole is also evident in the figure. As countries get richer, fertility rates appear to decline to levels consistent, not with a constant population, but actually with a declining population.

¹Examples include Kremer (1993), Acemoglu (1998), Ngai and Samaniego (2011), Doepke and Zilibotti (2014), Acemoglu and Restrepo (2018), Akcigit and Kerr (2018), Atkeson and Burstein (2019), and Buera and Oberfield (2019).

Figure 1: The Total Fertility Rate (Live Births per Woman)



Note: The total fertility rate is the average number of live births a hypothetical cohort of women would have over their reproductive life if they were subject during their whole lives to the fertility rates of a given period and if they were not subject to mortality. Each data point corresponds to a five-year period. Source: United Nations (2019).

Conventional wisdom holds that in the future, global population will stabilize at something like 8 or 10 billion people. But maybe that is not correct. We surely do not know for certain what will happen in the future. The fact that so many rich countries already have fertility below replacement indicates that a future with negative population growth is a possibility that deserves further consideration.

The models of economic growth cited above assume a constant or growing population, and for understanding economic growth historically, that is clearly the relevant case. The demographic evidence, however, suggests that this may not be the case in the future. Hence the focus of this paper: what happens to economic growth if population growth is negative?

We show below — first in models with exogenous population growth and then later in a model with endogenous fertility — that negative population growth can be particularly harmful. When population growth is negative, both endogenous and semi-endogenous growth models produce what we call an *Empty Planet* result: knowledge and living standards stagnate for a population that gradually vanishes. In a model with endogenous fertility, a surprising result emerges: even the social planner can get stuck in this trap if society delays implementing the optimal allocation and suffers from

inefficient negative population growth for a sufficiently long period. In contrast, if the economy switches to the optimal allocation soon enough, it can converge to a balanced growth path with sustained exponential growth: an ever-increasing population benefits from ever-rising living standards. Policies related to fertility may therefore determine whether we converge to an “empty planet” or to an “expanding cosmos”; they may be much more important than we have appreciated.

1.1 Literature Review

Many models feature endogenous fertility, modeled in a variety of ways. Becker and Barro (1988) and Barro and Becker (1989) take an altruistic approach in which the utility of children enters the utility function for parents, giving rise to a dynastic utility function. Papers that follow this approach include Doepke and Zilibotti (2005) and Manuelli and Seshadri (2009). Other papers emphasize a “warm-glow” effect in which parents care about the number of their offspring; for example, see De La Croix and Doepke (2003), Hock and Weil (2012), and Doepke and Tertilt (2016). Finally, many papers feature a quantity-quality tradeoff and assign a key role to education, often in the context of explaining the demographic transition and the emergence of modern economic growth. These include Becker, Murphy and Tamura (1990), Galor and Weil (1996, 2000), Greenwood and Seshadri (2002), and Kalemli-Ozcan (2002).

On the empirical side, Jones and Tertilt (2008) provide a detailed account of the decline in U.S. fertility using Census data, while Delventhal, Fernandez-Villaverde and Guner (2019) study the demographic transition using data from 186 countries and 250 years. Feyrer, Sacerdote and Stern (2008) highlight negative population growth in Japan and parts of Europe and raise the possibility that it could revert back to being positive as the status of women in society improves. Doepke and Tertilt (2016) and Greenwood, Guner and Vandenbroucke (2017) provide general surveys of family macroeconomics, including fertility. In general, this literature sometimes recognizes the possibility that population growth could ultimately be negative, but that is not its emphasis.

More generally, demographic forces are garnering broader attention in the macroeconomic literature. Several recent papers suggest that falling labor force growth may explain a substantial part of the decline in firm entry and dynamism in the U.S. economy, including Karahan, Pugsley and Sahin (2019), Hopenhayn, Neira and Singhanian

(2018), Engbom (2019), and Peters and Walsh (2019).

The literature that explicitly considers negative population growth in a growth context is much smaller. Manuelli and Seshadri (2009) explain the heterogeneity in international fertility rates by emphasizing that taxes and transfers in Europe may in part be responsible for low fertility. Sasaki and Hoshida (2017) study negative population growth in a semi-endogenous growth setting. They show that the rate of technological change falls to zero as people endogenously exit the research sector. More surprisingly, they provide a setting where negative population growth leads to *positive* steady-state growth in income per person because capital per person rises as the number of people declines. However, this result is incomplete in that they assume a zero depreciation rate for capital: if there is a fixed amount of capital but the population declines, then capital per person grows. One can easily generalize their result to positive depreciation rates using a Solow model. If the rate of population decline is η and capital depreciates at rate δ , then there are two possible regimes. If $\eta > \delta$, i.e. the rate of population decline is faster than the depreciation rate of capital, then K/L rises asymptotically along a balanced growth path. But when $\eta < \delta$, instead, you get the standard Solow result of constant K/L in steady state. Empirically, rates of population decline are perhaps 1% or smaller, whereas depreciation rates are 3% or 5% or more. The Sasaki and Hoshida (2017) case of exponential growth in capital per person from declining population therefore seems implausible as an empirical matter. Christiaans (2011) has results along these lines in a model with increasing returns that results from externalities to capital, showing the two possible regimes.

This motivates Sasaki (2019a) to consider a model with non-renewable resources, where a zero depreciation rate is more natural. In that case, though, one might wonder about elasticities of substitution: if a single Robinson Crusoe populated an earth full of land and natural resources, would her income be extremely high? Sasaki (2019b) considers a Solow model with CES production and finds that with an elasticity of substitution less than unity, the long-run growth rate is determined only by the rate of technological progress, with no contribution from the rising capital-labor ratio that results from negative population growth. Because capital is not essential, even an infinite capital-labor ratio gives finite output. These results suggests that capital and non-renewable resources can be omitted from the model without much loss in generality,

which is what we do below.

2. The Empty Planet Result

How do idea-based growth models behave when population declines? We begin by introducing exogenous, negative population growth into a simplified version of the Romer (1990), Aghion and Howitt (1992), and Grossman and Helpman (1991) endogenous growth models. This case turns out to be especially easy to analyze. Then we consider semi-endogenous growth models.

2.1 Fully Endogenous Growth as in Romer/AH/GH

Consider the following simplified version of idea-driven endogenous growth models:

$$Y_t = A_t^\sigma N_t \quad (1)$$

$$\frac{\dot{A}_t}{A_t} = \alpha N_t \quad (2)$$

$$N_t = N_0 e^{-\eta t}, \quad \eta > 0 \quad (3)$$

According to equation (1), a single consumption-output good is produced using people N_t and the stock of ideas (“knowledge”) A_t . Crucially, as in Romer (1990), there is constant returns in this production function to rival inputs — here just people — and therefore increasing returns to people and ideas together. The degree of increasing returns is parameterized by σ .

Equation (2) is the endogenous growth equation. It says that the growth rate of knowledge is proportional to the population. The literature often distinguishes between researchers and workers who produce the consumption good, but not always. Here, we make the simplifying assumption that is closer in spirit to learning by doing: people can work to make consumption goods and get new ideas at the same time. An exogenous split of people into workers and researchers would deliver identical results to those below. One can also allow this allocation to be determined endogenously.

Finally, equation (3) specifies that the population declines exogenously at the rate η . For example, $\eta = .005$ corresponds to a population that declines exponentially at a half a percent per year. We write the model here and throughout the paper so that all

parameter values (Greek letters) are positive.

Combining (2) and (3) gives the following differential equation, in which the growth rate of knowledge declines exponentially:

$$\frac{\dot{A}_t}{A_t} = \alpha N_0 e^{-\eta t}.$$

This differential equation is easy to solve, yielding the following result (detailed derivations for this and other results are available in Appendix A):

Result 1 (Romer/AH/GH with Negative Population Growth): *In the Romer/AH/GH model with negative population growth, the stock of knowledge A_t is given by*

$$\log A_t = \log A_0 + \frac{g_{A0}}{\eta} (1 - e^{-\eta t})$$

Both A_t and income per person $y_t \equiv Y_t/N_t$ converge to constant values A^ and y^* as t goes to infinity, where*

$$A^* = A_0 \exp\left(\frac{g_{A0}}{\eta}\right)$$

$$y^* = y_0 \exp\left(\frac{g_{y0}}{\eta}\right)$$

where variables indexed by 0 denote initial values.

We refer to this as the **Empty Planet result**. Economic growth stagnates as the stock of knowledge and living standards settle down to constant values. Meanwhile, the population itself falls at a constant rate, gradually emptying the planet of people. This outcome stands in stark contrast to the conventional result in growth models in which knowledge, living standards, and even population grow exponentially: not only do we get richer over time, but these higher living standards apply to an ever rising number of people.

The last equation in Result 1 is amenable to calibration. For example, if $g_{y0} = 2\%$ and $\eta = 1\%$, so that the population is falling at 1% per year, the long-run level of GDP per person will be $e^2 \approx 7.4$ times higher than current income. Obviously, slower declines in population would make this factor even higher.

In what follows, we explore the robustness of this finding. First, we see that it occurs in semi-endogenous growth models as well, and then we consider what happens when

the population growth rate itself is an endogenous outcome.

2.2 Semi-Endogenous Growth with Declining Population

With positive rates of population growth, semi-endogenous growth models in the tradition of Jones (1995), Kortum (1997), and Segerstrom (1998) give very different results from the fully endogenous growth models. We see next that with negative population growth, the *Empty Planet* result still emerges.

A simplified semi-endogenous growth model is obtained by changing the idea production function:

$$Y_t = A_t^\sigma N_t \quad (4)$$

$$\frac{\dot{A}_t}{A_t} = \alpha N_t^\lambda A_t^{-\beta} \quad (5)$$

$$N_t = N_0 e^{-\eta t}. \quad \eta > 0 \quad (6)$$

In particular, we introduce the parameter $\beta > 0$, capturing the extent to which new ideas (proportional improvements in productivity) are getting harder to find.²

Combining (5) and (6) gives the following differential equation:

$$\frac{\dot{A}_t}{A_t} = \alpha N_0^\lambda e^{-\lambda \eta t} A_t^{-\beta}.$$

Integrating this differential equation gives the next result.

Result 2 (Semi-Endogenous Growth with Negative Population Growth): *In the semi-endogenous growth model with negative population growth, the stock of knowledge A_t is given by*

$$A_t = A_0 \left(1 + \frac{\beta g A_0}{\lambda \eta} \left(1 - e^{-\lambda \eta t} \right) \right)^{1/\beta}.$$

Defining $\gamma \equiv \lambda \sigma / \beta$ to capture the overall degree of increasing returns to scale in this economy, both A_t and income per person $y_t \equiv Y_t / N_t$ converge to constant values A^ and*

²An alternative in the literature is to write the idea production function as $\dot{A}_t = \alpha N_t^\lambda A_t^\phi$ with $\phi < 1$. These are equivalent, with $\beta = 1 - \phi$.

y^* as t goes to infinity, where

$$A^* = A_0 \left(1 + \frac{\beta g_{A0}}{\lambda \eta} \right)^{1/\beta} .$$

$$y^* = y_0 \left(1 + \frac{g_{y0}}{\gamma \eta} \right)^{\gamma/\lambda} . \quad (7)$$

Along the transition path, the growth rate satisfies

$$\frac{\dot{y}_t}{y_t} = g_{y0} \cdot \left(\frac{y_t}{y_0} \right)^{-\frac{\lambda}{\gamma}} e^{-\lambda \eta t} = \frac{g_{y0} e^{-\lambda \eta t}}{1 + \frac{g_{y0}}{\gamma \eta} (1 - e^{-\lambda \eta t})}$$

In other words, the growth rate falls to zero slightly faster than $e^{-\lambda \eta t}$.

This result confirms that both endogenous growth and semi-endogenous growth lead to the *Empty Planet* outcome. Rather than sustained exponential growth in living standards and population, living standards stabilize for a vanishing number of people.

Quantitatively, however, the level at which lower living standards stagnate can be much lower with semi-endogenous growth. To illustrate, we need to calibrate one additional parameter relative to what we had before. Across a range of different case studies, Bloom, Jones, Van Reenen and Webb (2019) find estimates of $\beta \approx 3$ when $\sigma = 1$ (a normalization in the general case where we do not observe ideas directly) and $\lambda = 1$ (a standard value in the literature), which gives $\gamma = 1/3$. Plugging these values into equation (7), the long-run level of GDP per person would be $(1 + 3 \cdot 2)^{1/3} \approx 1.9$ times higher than current income. The difference versus the endogenous growth case is striking: with $\beta = 0$ (so that $\gamma = \infty$), long run income was 7.4 times higher than current income for the same parameter values.

3. Endogenous Fertility and the Equilibrium Allocation

We now endogenize the population growth rate itself. There are many related ways to accomplish this, and the literature has not converged on a single best practice; see the literature review at the start of this paper for references. Almost all approaches assume that having offspring is a time intensive activity, and this is at the center of the approach we take below.

In models of endogenous fertility, population growth in a decentralized equilibrium can be optimal or may be either above or below the optimal rate. In fact, because the number of people is endogenous, the definition of “optimal” is itself not obvious; for example, see Golosov, Jones and Tertilt (2007). The most natural case of interest here is one in which parents do not fully internalize the fact that their offspring create nonrival ideas that benefit the entire economy, so that equilibrium fertility is too low.

But there are also other possible nuances. For example, Farhi and Werning (2007) note that the social planner may care about future generations both because individuals care about their own children and because the social planner puts weight on each generation. This means that social welfare will generally put more weight on future generations than individuals do, also leading optimal fertility to be higher than equilibrium fertility. Externalities to human capital in models with a quality-quantity tradeoff can also give rise to inefficiently low fertility. Alternatively, one can construct idea-based models in which optimal fertility is below equilibrium fertility; see Jones (2003) and Futagami and Hori (2010) for some discussion. Here, we do not attempt to draw any firm conclusion about the range of possible externalities that may exist. Instead, we focus on one interesting case in which the equilibrium features negative population growth while the optimal allocation has positive population growth.

To simplify, we abstract from the demographic transition. That is, we are not focused on how fertility fell from 5 to 3 to 1.5 children per woman. Instead, the focus is on the stable fertility rate at the end of the demographic transition, and what happens if it implies negative population growth.

3.1 Environment

The economic environment for the setup with endogenous fertility is in Table 1. It builds on our earlier model, with one enhancement. There is now a single allocative decision that has to be made at each date: each person is endowed with one unit of time that can be used to produce either consumption or offspring. Devoting ℓ_t units of time to producing children leads to a fertility rate of $b(\ell_t) = \bar{b}\ell_t$. The linear function is convenient analytically but not essential. There is a constant death rate, δ , and the population growth rate is $n_t = \bar{b}\ell_t - \delta$. Thus if ℓ_t is sufficiently small, the population growth rate can be negative.

Table 1: Economic Environment: Endogenous Fertility Model

Final output	$Y_t = A_t^\sigma (1 - \ell_t) N_t$
Population growth	$\frac{\dot{N}_t}{N_t} = n_t = b(\ell_t) - \delta$
Fertility	$b(\ell_t) = \bar{b}\ell_t$
Ideas	$\frac{\dot{A}_t}{A_t} = N_t^\lambda A_t^{-\beta}$
Generation 0 utility	$U_0 = \int_0^\infty e^{-\rho t} u(c_t, \tilde{N}_t) dt, \quad \tilde{N}_t \equiv N_t/N_0, \rho \equiv \bar{\rho} + \delta$
Flow utility	$u(c_t, \tilde{N}_t) = \log c_t + \epsilon \log \tilde{N}_t$
Consumption	$c_t = Y_t/N_t$

This setup excludes many other considerations that would be interesting to explore in the future such as human capital, physical capital, and a quantity-quality tradeoff. We instead focus on the simplest model that allows us to make our points.

People obtain utility from consumption and from having descendents. The expected lifetime utility of a member of the generation born at date 0 is

$$U_0 = \int_0^\infty e^{-(\bar{\rho} + \delta)t} u(c_t, \tilde{N}_t) dt$$

where $\bar{\rho}$ is the pure rate of time preference, c_t is consumption, and $\tilde{N}_t \equiv N_t/N_0$ is the number of descendents of generation 0 at date t . Discounting also occurs because of the death rate, and we define $\rho \equiv \bar{\rho} + \delta$ as the overall discount rate. The flow of utility at date t is

$$u(c_t, \tilde{N}_t) = \log(\tilde{N}_t^\epsilon c_t) = \log c_t + \epsilon \log \tilde{N}_t.$$

This is a special case of the preferences considered by Becker and Barro (1988) as used by Barro and Sala-i-Martin (1995). The log functional forms are convenient in several ways. First, $\log c_t$ leads income and substitution effects to cancel, so that a constant fertility rate is consistent with balanced growth. The $\log \tilde{N}_t$ is helpful for time consistency and for simplifying the value function in solving for transition dynamics. This setup can be generalized in various ways and the qualitative results will be preserved.

But the analysis of the transition dynamics is sufficiently subtle that it is very helpful to illustrate them with the simplest possible setup.³

3.2 A Competitive Equilibrium with Externalities

As in Romer (1990), the nonrivalry of ideas leads to increasing returns. Some departure from pure perfect competition is necessary, and the equilibrium in general will not be efficient. We consider a simple equilibrium in which the production of ideas is purely external. Also, we start with the equilibrium allocation because it is designed to be simple. Section 4 below considers an optimal allocation.

Firms produce final output in perfectly competitive markets, taking the stock of ideas A_t as exogenous. Each person chooses time spent raising children versus working in the market sector, ℓ_t versus $1 - \ell_t$, in order to maximize utility, also taking the time path of A_t as exogenous. Hence ideas evolve according to the idea production function entirely as an externality: people do not recognize that by having children, their kids may produce new knowledge in the future that makes the entire economy more productive. Markets are perfectly competitive, subject to the idea externality, and the only price is the wage per unit of work, given by $w_t = A_t^\sigma$ in equilibrium.

Taking $\{w_t\}$ as given, people in each generation solve

$$\max_{\{\ell_t\}} \int_0^\infty e^{-\rho t} u(c_t, \tilde{N}_t) dt$$

subject to

$$\dot{N}_t = (b(\ell_t) - \delta)N_t$$

$$c_t = w_t(1 - \ell_t)$$

and given the function forms assumed in Table 1.

The Hamiltonian for this problem is

$$\mathcal{H} = u(c_t, \tilde{N}_t) + v_t(b(\ell_t) - \delta)N_t$$

³Here are some examples of how things can be generalized. We currently have all people generating ideas rather than only people who work in the market sector. Introducing a $1 - \ell_t$ into the idea production function just leads to a level effect and so should not change the results below. Allowing the flow of utility to be $N_t^\epsilon u(c_t)$ can also be permitted, as long as $u(c)$ is CRRA with curvature between 0 and 1.

where v_t is the shadow price (in utils) of another person.

The first-order condition for this problem with respect to ℓ_t is

$$\underbrace{v_t N_t b'(\ell_t)}_{\text{MU of time in fertility}} = \underbrace{u_c(c_t, \tilde{N}_t) w_t}_{\text{MU of time in making goods}} = \frac{1}{1 - \ell_t}. \quad (8)$$

On the left side, spending a little more time on fertility leads each of N_t people to have $b'(\ell_t)$ additional offspring, valued at shadow price v_t . Alternatively, the time could be spent working to earn the wage w_t , which is converted to utility units using the marginal utility of consumption. At the maximum, individuals are indifferent between these two options.

The first-order condition for N_t gives an arbitrage-like equation for the shadow price of people:

$$\rho = \frac{\dot{v}_t}{v_t} + \frac{1}{v_t} \left(\frac{\epsilon}{N_t} + v_t n_t \right). \quad (9)$$

To solve further, define $V_t \equiv v_t N_t$ to be the shadow value in utils of the entire population. Then (9) can be written as

$$\dot{V}_t = \rho V_t - \epsilon.$$

This differential equation obviously has a steady state at $V_{eq}^* = \epsilon/\rho$. Simple inspection of the differential equation reveals that this steady state is unstable: if V_t differs from V_{eq}^* , then either V_t speeds off to ∞ violating the transversality condition or the law of motion would require $\ell_t < 0$, which is not allowed. The solution is therefore

$$V_t = V_{eq}^* = \frac{\epsilon}{\rho}$$

so that V_t is constant at all points in time.

Substituting this into (8) and using our functional form assumptions gives the equilibrium amount of time spent on fertility. This is stated in our next result:

Result 3 (The Equilibrium with Endogenous Fertility): *In the equilibrium with endogenous fertility, the allocation of time devoted to offspring is constant and given by*

$$\ell_t = 1 - \frac{1}{bV_t} = 1 - \frac{1}{bV_{eq}^*} = 1 - \frac{\rho}{b\epsilon} \equiv \ell_{eq} \quad (10)$$

Equilibrium population growth is then also constant at the rate $n_{eq} = \bar{b} - \delta - \rho/\epsilon$. Depending on parameter values, equilibrium population growth can be positive or negative.

Fertility is constant at each point in time, even along the transition path. A higher rate of time preference (ρ) lowers fertility, while a higher preference for offspring (ϵ) or a better fertility technology (\bar{b}) raises fertility. This setup, therefore, does not explain the demographic transition or how we got to the situation we are in today; that's not its purpose.

Instead, it allows us to study the current situation: depending on parameter values, equilibrium population growth can be positive or negative. The negative case is the one of interest here. In that case, we have an equilibrium setup with endogenous fertility that feeds naturally into the results from Section 2. The negative population growth combined with the idea production function implies that the equilibrium with endogenous fertility features a growth rate that falls to zero so that output per person converges to a steady state, as in equation (7). Therefore, the *Empty Planet* result can be supported as an equilibrium outcome with endogenous fertility.

4. The Optimal Allocation

Now instead consider the optimal allocation in this economic environment. With endogenous fertility, there is no unique criterion for social welfare, as discussed by Golosov, Jones and Tertilt (2007). Instead, we consider the allocation that maximizes the utility of a representative generation. The key reason this differs from the equilibrium allocation considered above is that the optimal allocation takes into account the fact that a larger population generates more nonrival ideas, raising everyone's income. This will lead optimal fertility to be higher than its equilibrium rate.

Defined this way, the optimal allocation solves

$$\max_{\{\ell_t\}} \int_0^{\infty} e^{-\rho t} u(c_t, \tilde{N}_t) dt$$

subject to

$$\dot{N}_t = (b(\ell_t) - \delta)N_t$$

$$\frac{\dot{A}_t}{A_t} = N_t^\lambda A_t^{-\beta}$$

$$c_t = Y_t/N_t.$$

The Hamiltonian for the optimal allocation is

$$\mathcal{H} = u(c_t, \tilde{N}_t) + \mu_t N_t^\lambda A_t^{1-\beta} + v_t (b(\ell_t) - \delta) N_t$$

where μ_t is the shadow price of an idea and v_t is the shadow price of another person.

The first-order condition for this problem with respect to ℓ_t is

$$\underbrace{v_t N_t b'(\ell_t)}_{\text{MU of time in fertility}} = \underbrace{\frac{u_c(c_t, \tilde{N}_t) y_t}{1 - \ell_t}}_{\text{MU of time in making goods}} = \frac{1}{1 - \ell_t}. \quad (11)$$

Notice that this equation has the same form as the equilibrium first-order condition; however, the shadow value of people, v_t , will be different. (We abuse notation for now by not using a different letter for the equilibrium versus optimal v_t .)

The first-order condition with respect to A_t can be expressed as an arbitrage equation:

$$\rho = \frac{\dot{\mu}_t}{\mu_t} + \frac{1}{\mu_t} \left(u_c \sigma \frac{y_t}{A_t} + \mu_t (1 - \beta) \frac{\dot{A}_t}{A_t} \right). \quad (12)$$

The required rate of return is ρ , and the production of ideas yields both a capital gain and a dividend. Continuing this analogy, this equation can be solved to yield the shadow price of an idea along a balanced growth path as the initial dividend divided by “r-g”:

$$\mu_t^* = \frac{\frac{1}{c_t} \sigma \frac{y_t}{A_t}}{\rho - g_\mu - (1 - \beta) g_A} = \frac{\sigma / A_t^*}{\rho + \beta g_A} \quad (13)$$

where the second equality uses the fact that the growth rate of the dividend (and hence the growth rate of μ) equals $-g_A$ along a BGP.

Similarly, the first-order condition for N_t in arbitrage form is

$$\rho = \frac{\dot{v}_t}{v_t} + \frac{1}{v_t} \left(\frac{\epsilon}{N_t} + \mu_t \lambda \frac{\dot{A}_t}{N_t} + v_t n_t \right). \quad (14)$$

Comparing this equation to the equivalent condition in the equilibrium, equation (9),

reveals that they differ because of the second term in parentheses: the optimal allocation values people not only for the direct utility they provide, but also because of the additional ideas they produce.

4.1 Steady State Balanced Growth

Using (13), the social value of people along a balanced growth path is given by

$$V_{sp}^* = v_t^* N_t^* = \frac{\epsilon + \lambda z^*}{\rho} \quad (15)$$

where

$$z^* \equiv \mu_t^* \dot{A}_t^* = \frac{\sigma g_A^*}{\rho + \beta g_A^*}. \quad (16)$$

We use z_t to denote the social value of the new ideas produced at a point in time. Therefore, equation (15) makes clear that the social planner values people for both the utility they bring, via ϵ , as well as the additional ideas they produce, via z^* .

Finally, if the planner solution features positive population growth in the steady state, i.e. if $n_{sp}^* > 0$, then

$$g_A^* = \frac{\lambda n_{sp}^*}{\beta} \quad (17)$$


$$g_y^* = \gamma n_{sp}^*, \quad \text{where } \gamma \equiv \frac{\lambda \sigma}{\beta}. \quad (18)$$

As usual in semi-endogenous growth models, the long-run growth rate is the product of the overall degree of increasing returns to scale, $\gamma \equiv \lambda \sigma / \beta$, and the rate at which scale is growing, n_{sp}^* . Alternatively, if the planner solution features zero or negative population growth in the steady state, then $g_A^* = g_y^* = 0$.

These two growth regimes — the focus of the first half of the paper — are shown graphically in Figure 2. If steady-state population growth is positive, then steady-state knowledge growth is proportional to n . But if population growth is negative, then steady-state knowledge growth is zero. This kink — these two regimes — gives rise to the possibility of multiple steady states, which we will see shortly.

Equations (15) and (16) make clear the difference between the equilibrium and the optimal allocations in steady state. The equilibrium value of people is $V_{eq}^* = \frac{\epsilon}{\rho}$, while

Figure 2: Knowledge Growth and Population Growth in Steady State



Note: There is a fundamental “kink” in the technology for generating steady-state growth in the model. If population growth is positive, then steady-state knowledge growth is proportional to n . But if population growth is negative, then steady-state knowledge growth is zero. This kink — these two regimes — gives rise to the possibility of multiple steady states.


the social value includes an additional term reflecting the ideas that are produced:

$$V_{sp}^* = \frac{\epsilon + \lambda z^*}{\rho} \Rightarrow V(n) = \begin{cases} \frac{1}{\rho} \left(\epsilon + \frac{\gamma}{1 + \frac{\rho}{\lambda n}} \right) & \text{if } n > 0 \\ \frac{\epsilon}{\rho} & \text{if } n \leq 0 \end{cases} \quad (19)$$

Both allocations value an extra person because of the additional utility that descendants provide, via ϵ . Only the optimal allocation values the extra ideas produced at the margin by the extra person. However, if $n^* \leq 0$, then the growth rate g_A^* is also zero, implying from (16) that $z^* = 0$. That is, the value of the additional ideas produced by people falls to zero.⁴ Importantly, this implies that with zero growth in ideas, the optimal and equilibrium allocations value additional people only because of the utility benefit, so that $V_{sp}^* = V_{eq}^*$.

⁴In (16), there is a price effect and a quantity effect: the quantity of new ideas, \dot{A} , is zero in this case. The price of ideas, μ_t , is given by (13) and remains finite; intuitively, if income and consumption are constant and finite in some steady state, then the price of a new idea is as well. So if \dot{A}_t is zero, then so is the social value of new ideas.

Figure 3: A Unique Steady State for the Optimal Allocation when $n_{eq}^* > 0$



Note: When equilibrium fertility is positive, there is a unique solution to equations (19) and (20). That is, the optimal allocation features a unique steady state with $n_{sp}^* > n_{eq}^*$.

Solving (11) gives the optimal allocation of time along a BGP as well as the implied population growth rate:


$$\ell_{sp}^* = 1 - \frac{1}{bV_{sp}^*} \Rightarrow n(V) = \bar{b} - \delta - \frac{1}{V}, \quad (20)$$

where we've left implicit the constraint that $\ell \geq 0$ implying $n \geq -\delta$.

Equations (19) and (20) give two nonlinear equations in two unknowns, n and V that characterize the optimal allocation in steady state. These equations are characterized graphically in Figures 3 and 4.

Figure 3 considers the case in which equilibrium fertility is positive. In this case, there is a unique solution to equations (19) and (20). The optimal allocation features a unique steady state in which optimal population growth exceeds the equilibrium rate, i.e. $n_{sp}^* > n_{eq}^*$.

The case of interest in this paper, however, is when equilibrium population growth is negative, and this case gives rise to a rich set of outcomes, as suggested by Figure 4.

Figure 4: Multiple Steady States in the Optimal Allocation when $n_{eq}^* < 0$ 

Note: When equilibrium fertility is negative and there exists an “expanding cosmos” steady state for the planner problem, there are generically three candidate solutions to equations (19) and (20) that characterize the steady state. We will see later that the middle steady state is unstable and can be ruled out.

With $n_{eq}^* < 0$, the optimal allocation generically features *three* candidate steady states. A high “expanding cosmos” steady state has positive population growth. The low steady state has the same population growth rate as the equilibrium allocation and features negative population growth; this is the *Empty Planet* outcome. Finally, there is a middle steady state in between. We will see shortly that this steady state is unstable and would never be reached along the optimal path. But even the presence of two stable steady states suggests that the dynamics of the optimal allocation are rich.⁵

⁵The middle and high steady states can be characterized analytically as the solution to a quadratic equation. In particular, $V_{sp}^* = \frac{\rho + \bar{\alpha}\epsilon + (\bar{b} - \delta)\gamma \pm \sqrt{(\rho + \bar{\alpha}\epsilon + (\bar{b} - \delta)\gamma)^2 - 4\bar{\alpha}\rho(\epsilon + \gamma)}}{2\bar{\alpha}\rho}$, where $\bar{\alpha} \equiv \bar{b} - \delta$. For some parameter choices (e.g. for γ sufficiently small, which rotates the $V(n)$ schedule back to the left), the term inside the square root can be negative so that the only steady state is the equilibrium solution with negative population growth.

4.2 Stability and Transition Dynamics

The transition dynamics turn out to be economically interesting but must be solved for numerically. We begin by describing the baseline parameter values that we use in the analysis; the results are robust to a range of alternative values.

Parameter values. Table 2 summarizes our parameter choices for the numerical examples. They are chosen to be somewhat realistic, but the main point is to show qualitatively what the transition dynamics of the model look like.

Because we do not observe ideas directly, it is convenient to normalize $\sigma = 1$ so that A has the units of total factor productivity. The extensive evidence on idea production functions in Bloom, Jones, Van Reenen and Webb (2019) suggests that $\beta > 0$ so that the exponential growth of ideas is getting harder to achieve. We choose $\beta = 1.25$, a typical value. With little lost, we assume $\lambda = 1$ so that the overall implied degree of increasing returns to scale is $\gamma \equiv \lambda\sigma/\beta = 0.80$.

We assume a conventional rate of time preference of $\rho = 1\%$ and set $\delta = 1\%$, implying an expected lifetime of 100 years.

Motivated by the recent fertility experience in the OECD, Japan, and the United States, we assume $n_{eq}^* = -0.5\%$, so that in equilibrium, the population will decline at half a percent per year. Finally, we assume that the typical person spends about 1/8 of his or her time endowment producing and raising children.

Given these assumptions, the following four equations determine the values of \bar{b} , ϵ , n_{sp}^* , and ℓ_{sp}^* :

$$n_{eq}^* = \bar{b}\ell_{eq}^* - \delta = -0.5\% \quad (21)$$

$$\ell_{eq}^* = 1 - \frac{\rho}{\bar{b}\epsilon} = 1/8 \quad (22)$$

$$n_{sp}^* = \bar{b}\ell_{sp}^* - \delta \quad (23)$$

$$\ell_{sp}^* = 1 - \frac{\rho}{\bar{b}(\epsilon + \lambda z^*)} \quad (24)$$

where z^* is given by (16) and (17).

Implied Parameter Values and Steady-State Results. The implied parameter values and steady-state outcomes (corresponding to the “high” steady state of the planner

Table 2: Parameter Values and Steady-State Results for the Numerical Examples

Key Assumed Values as Inputs to Numerical Examples

Parameter/Moment	Value	Comment
σ	1	Normalization
λ	1	Duplication effect of ideas; 0.75 possible
β	1.25	Bloom, Jones, Van Reenen and Webb (2019)
ρ	.01	Standard value
δ	1%	Death rate
n_{eq}^*	-0.5%	Suggested by fertility rates in Europe, Japan, U.S.
ℓ_{eq}^*	1/8	Time spent raising children

Implied Parameter Values and “Expanding Cosmos” Steady-State Results

Result	Value	Comment
\bar{b}	.04	$n_{eq}^* = \bar{b}\ell_{eq}^* - \delta = -0.5\%$
ϵ	.286	From equation (22) for ℓ_{eq}^*
n_{sp}^*	1.74%	From equations (23) and (24) for ℓ_{sp}^* and n_{sp}^*
ℓ_{sp}^*	0.68	From equations (23) and (24) for ℓ_{sp}^* and n_{sp}^*
$g_y^{sp} = g_A^{sp}$	1.39%	Equals γn_{sp}^* with $\sigma = 1$
...	64.0%	Idea share of social value of people: $\lambda z^*/(\epsilon + \lambda z^*)$

Note: The first panel in the table shows key assumptions that are an input into the numerical examples. The second panel shows implied parameter values and steady-state results given these assumptions. The steady-state for the planner problem corresponds to the “high” steady state.

problem) are then shown in the bottom part of Table 2. It is worth noting that the optimal population growth rate given these values is substantially higher than the equilibrium rate: 1.74% versus -0.5%. Even with sharp dynamic diminishing returns in the idea production function ($\beta = 1.25$), there is a large positive externality to offspring in this calibration. The associated steady-state growth rate of income per person is 1.39%.


Equation (19) suggests an alternative way to make the point about the size of the population externality. The equilibrium social value of people is $V_{eq}^* = \epsilon/\rho$ while the optimal value in steady state is $(\epsilon + \lambda z^*)/\rho$. The last row of Table 2 reports the fraction of the optimal social value of people that is associated with ideas: $\lambda z^*/(\epsilon + \lambda z^*) \approx 64\%$. Of course, this is not a formal calibration that one should take too seriously, but the values give a sense of the magnitudes that may be relevant.

Stability and Dynamics. Given the presence of three candidate steady states in the optimal allocation, the transition dynamics are subtle. We derive the results in detail in Appendix B. Here, we summarize the key findings.

While there are two fundamental state variables in our environment (N_t and A_t), it is convenient for studying transition dynamics to consider the state variables as N_t and $x_t \equiv A_t^\beta/N_t^\lambda$. We will refer to x_t as “knowledge per person,” which is a slight abuse of language in that it ignores the exponents. As shown in Appendix B, the optimal allocation of labor can be expressed solely as a function of x_t , which means that population growth can as well. This allows us to study the transition dynamics in a simple two-dimensional plane. It is also useful to keep in mind that with this definition of the state, $\dot{A}_t/A_t = 1/x_t$. That is, the state-like variable x_t can also be interpreted as the inverse of the growth rate of knowledge. Along any balanced growth path, x_t will be constant. Finally, notice that a bigger x is not necessarily better: it is the ratio of two state variables that are each good for welfare, A and N .


Figure 5 shows the transition dynamics for the equilibrium allocation. The equilibrium features a constant negative rate of population growth, which causes x_t to increase over time: intuitively, A_t goes to a constant, while N_t is falling, which causes x_t to increase. These dynamics are very simple, but this is a good way to introduce the figure and it complements the more complicated dynamics of the optimal allocation, which we turn to now.

Figure 5: Transition Dynamics for the Equilibrium Allocation



Note: This figure shows the transition dynamics in the equilibrium allocation. The state variable on the horizontal axis is $x \equiv A^\beta / N^\lambda$, which we somewhat loosely refer to as “knowledge per person.” Arrows indicate transition dynamics. The equilibrium features a constant negative rate of population growth, which causes x_t to increase over time.

Figure 6: Transition Dynamics for the Optimal Allocation




Note: This figure shows the transition dynamics in the optimal allocation. The state variable on the horizontal axis is $x \equiv A^\beta / N^\lambda$, which we somewhat loosely refer to as “knowledge per person.” Arrows indicate transition dynamics. If the economy begins with a stock of knowledge per person that is not too high, it converges to the stable “high” steady state. Alternatively, if knowledge per person is sufficiently high, the economy converges to the low steady state with negative population growth, which equals the equilibrium rate. The candidate “middle steady state” is unstable with imaginary eigenvalues that generate a spiral path locally. The dynamics shown correspond to the path that maximizes welfare for each value of the state variable and correspond to the upper and lower arms of the spiral. See the appendix for details.

Figure 6 shows the more complicated dynamics of the optimal allocation. The high “expanding cosmos” steady state is saddle-path stable. There are a wide range of potential starting points for knowledge per person, x , such that the optimal allocation ultimately converges to the high steady state. This is what one would generally expect in a problem like this.


The middle steady state is unstable: linearizing the first-order conditions from the Hamiltonian of the optimal problem reveals that it is unstable with imaginary eigenvalues, generating spiral dynamics. In other words, for a given value of knowledge per person, x , close to the middle steady state, there are multiple candidate paths that satisfy the first-order conditions from the Hamiltonian. A stylized picture of the spiral dynamics around the middle steady state is shown in Figure 7.

Figure 7: A Stylized Depiction of the Spiral Dynamics for the Middle Steady State



Note: For our baseline parameter values, the linearized dynamics around the middle steady state features imaginary eigenvalues with positive real parts, generating an unstable spiral out away from the steady state.

Figure 8: Population Growth Near the Middle Steady State



Note: The graph zooms in on the spiral around the middle steady state. The blue line is the upper arm and the green line is the lower arm. Solid lines are optimal paths, while dashed lines are points that are dominated. The numbers in red report welfare for $N = 1$ (we subtract 500 from the welfare numbers to make the graph easier to read). The jump between the upper arm and the lower arm is the point where the welfares are equal.

For the parameter values in Table 2, the path that delivers the highest welfare is shown in Figure 6: it involves staying on the “upper” arm of the spiral for values of x slightly higher than x_{middle}^* and then a jump down to the “lower” outer arm of the spiral for even higher values of x ; see Figure 8, discussed next, for more details. This jump does not violate any first-order condition because the dynamics are unstable around the jump: the economy never jumps along an optimal path.

Figure 8 zooms in on the middle candidate steady state and gives the details underlying the jump. The blue line is the upper arm and the green line is the lower arm. Solid lines are optimal paths, while dashed lines are points that are dominated. The numbers in red report welfare for $N = 1$.⁶ The jump between the upper arm and the lower arm occurs where the welfares are equal. Additional spirals not visible to the eye exist near

⁶We subtract 500 from the welfare numbers to make the graph easier to read.

the middle steady state, but these too are dominated.⁷

If the starting value for x is below the jump point in Figure 6, the economy converges to the “expanding cosmos” steady state. Conversely, if the starting value for x is above the jump point, then the economy asymptotically converges to the “low” steady state, the *Empty Planet* outcome. This features negative population growth at a rate that asymptotically equals the equilibrium rate n_{eq}^* and occurs as knowledge per person, $x \equiv A^\beta/N^\lambda$, goes off to infinity. This happens because N_t falls due to the negative population growth while the stock of knowledge A_t stabilizes at some positive value, as in Result 2.

The intuition for the two steady states — and for the stability of both — goes as follows. An increase in knowledge per person x causes optimal fertility to decline because the extra ideas produced by offspring have a diminishing marginal benefit; this explains the negative slope of $n(x)$ in Figure 6. If equilibrium fertility is positive, then optimal fertility will also remain positive — the planner values people at least as much as the equilibrium. But if equilibrium fertility is negative, then for x high enough, optimal fertility becomes negative as well. This is because as x goes to infinity, the stock of knowledge divided by the number of people is so high that the “knowledge value” of additional offspring falls to zero. But once population growth is negative, x increases over time rather than decreasing since the denominator of $x \equiv A^\beta/N^\lambda$ is falling. That causes x to increase, reinforcing the change. That is the intuition for the bifurcation in Figure 6 and the perhaps surprising stability of the *Empty Planet* outcome.⁸

What do the transition dynamics look like for different parameter values? The intuition for the answer can be found by looking back at the $n(V)$ and $V(n)$ figure that characterized the multiple steady states, Figure 4. If we reduce the importance of ideas in the economy — i.e. if we reduce γ , say by making ideas even harder to find via a higher β — the $V(n)$ curve rotates back to the left. This pushes down the “expanding cosmos” steady state and raises the middle steady state. That is, the two points get closer together. In Figure 6, this shrinks the range of x values for which the transition

⁷This type of unstable middle steady state in economics is called a “Skiba point” after Skiba (1978). Theoretical analysis of one-state models can be found in Beyn, Pampel and Semmler (2001) and suggest that additional analytic progress could be made in this application. We leave this to future research.

⁸In fact, the intuition is slightly more subtle. For different parameter values, it is possible for the population growth rate on the lower arm after the jump to be, temporarily, positive. Nevertheless, x continues to increase because the rate of population growth is slow relative to the rate of knowledge growth. See footnote 9 for more details.

dynamics lead x_t to decline. If we continue to lower γ and reduce the importance of ideas, eventually the $n(V)$ and $V(n)$ curves are tangent — the “high” and “middle” steady states become the same point. In Figure 6, the transition dynamics then *always* involve x increasing, and the only steady state is the *Empty Planet* outcome. This same general result obtains if we reduce the importance of ideas even further, so that the $n(V)$ curve in Figure 4 lies entirely below the $V(n)$ schedule. Intuitively, if ideas are not very important, both the equilibrium and the optimal allocations will feature negative population growth.⁹

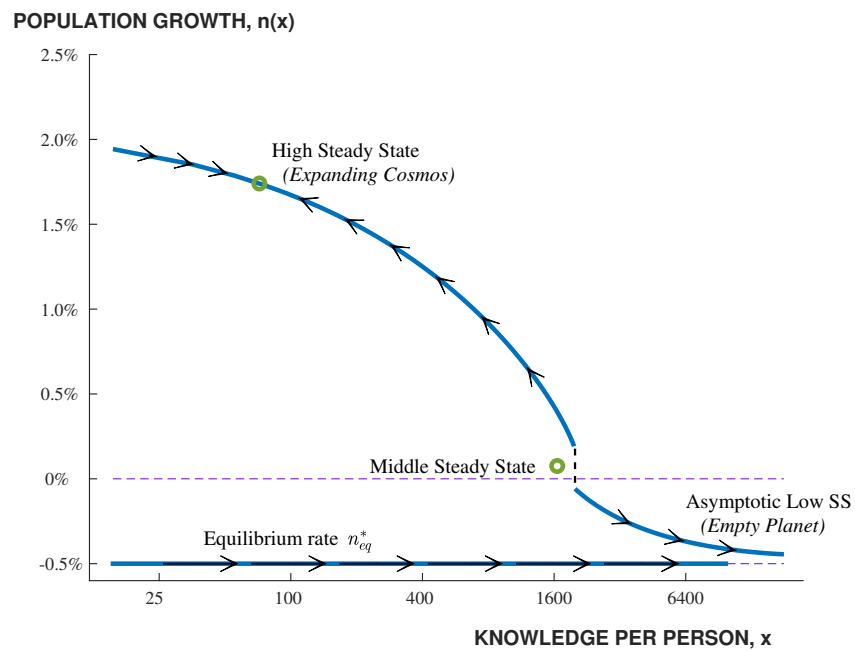
4.3 The Economics of the Transition Dynamics

The transition dynamics lead to an important economic point, summarized in Figure 9. Consider an economy that is governed by the equilibrium allocation. It features negative population growth at rate n_{eq}^* , and suppose the economy is initially endowed with a certain population and stock of knowledge such that knowledge per person, x_0 , equals 50. The social planner would like the economy to have a much higher fertility rate and converge to the “expanding cosmos” steady state, with positive population growth and positive economic growth: both the number of people and the amount of income per person would rise exponentially forever. In contrast, the equilibrium allocation will simply move the economy steadily to the right, to higher values of x , along the lower line: there will be a constant negative rate of population growth, so knowledge per person, x , will rise as the number of people declines.

At any point in time, society may adopt better institutions, such as a fertility subsidy, that move the economy to the optimal allocation. If this occurs at $x = 100$ or $x = 400$ or $x = 1600$, then the economy will eventually transition to the high steady state and exhibit exponential growth forever. However — and this is the surprising point — if the economy delays adopting good institutions for too long, eventually knowledge per person x will rise *above* the jump point (around 1990 in the figure). Once this happens, the

⁹A few other points are worth noting. First, as discussed in Appendix B, the law of motion for $x_t \equiv A_t^\beta / N_t^\lambda$ is $\dot{x}_t = \beta - \lambda n_t x_t$. Clearly if n_t is negative, then $\dot{x} > 0$, but this is sufficient rather than necessary: $\dot{x}_t > 0$ can occur with positive population growth as long as n_t is sufficiently small. This explains how the transition dynamics can feature rising x_t throughout for γ sufficiently small. Second, the imaginary eigenvalues that deliver the spiral dynamics around the middle steady state also disappear eventually as we reduce γ . In that case, the middle steady state is an unstable source with two exit paths that converge to the high and low steady states. This is intuitive if you imagine what the middle state looks like as the spiral gets smaller and smaller.

Figure 9: Transition Dynamics: Summary



Note: The bottom line in the figure shows the transition dynamics for the equilibrium allocation while the curved lines show the dynamics for the optimal allocation. An economy governed by the equilibrium dynamics can get trapped in the Empty Planet outcome if it waits too long to switch to the optimal allocation. In this example, that would occur if knowledge per person, x , rises above about 1990.

optimal regime changes. Adopting good institutions that deliver the optimal allocation will involve negative population growth, albeit at a higher rate than the equilibrium. This means that knowledge per person will continue to grow — not because of large increases in knowledge but rather because of declines in the number of people — and the economy will converge to the low steady state. Population will decline, knowledge will remain below an upper bound, and income per person will stagnate. This is the *Empty Planet* outcome. **The surprise is that if society waits too long to adopt good institutions, the optimal allocation switches from one of sustained exponential growth in population, knowledge, and living standards to one of stagnation and an empty planet.**

This is summarized in our last main result:

Result 4 (The Optimal Allocation with Endogenous Fertility): *The allocation that maximizes the welfare of each generation converges to one of two steady states. If the economy adopts the optimal allocation while knowledge per person, x , is sufficiently low, it leads to the Expanding Cosmos outcome of sustained exponential growth in population, knowledge, and living standards. Alternatively, if the economy waits too long to switch to the optimal path, it converges to the Empty Planet outcome: living standards stagnate as the population gradually declines toward zero.*

5. Conclusion

Historically, fertility rates in high-income countries have fallen from 5 children per women to 4, 3, 2, and now even fewer. From a family's standpoint, there is nothing special about "above two" versus "below two" and the demographic transition may lead families to settle on fewer than two children. The macroeconomics of the problem, however, make this distinction one of critical importance: it is the difference between an Expanding Cosmos of exponential growth in both population and living standards and an Empty Planet, in which incomes stagnate and the population vanishes.

Endogenizing fertility leads to an additional subtlety. When the equilibrium fertility rate is negative, the optimal allocation features two stable steady states. If the economy adopts the optimal allocation soon enough, it converges to the Expanding Cosmos. But

if the economy waits too long to switch, even the optimal allocation can be trapped by the Empty Planet outcome.

These results are not a forecast, of course — the paper is designed to suggest that a possibility we have not considered carefully until now deserves more attention. Still, there are many ways in which this model could be misleading. Automation could enhance our ability to produce ideas sufficiently that growth in living standards continues even with a declining population, for example. Or new discoveries could eventually reduce the mortality rate to zero, allowing the population to grow despite low fertility. Nevertheless, the emergence of negative population growth in many countries and the fact that it has profound implications for the future of economic growth make this a topic worthy of further exploration.

References

- Acemoglu, Daron, “Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality,” *Quarterly Journal of Economics*, 1998, 113, 1055–1089.
- and Pascual Restrepo, “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,” *American Economic Review*, June 2018, 108 (6), 1488–1542.
- Aghion, Philippe and Peter Howitt, “A Model of Growth through Creative Destruction,” *Econometrica*, March 1992, 60 (2), 323–351.
- Akcigit, Ufuk and William R. Kerr, “Growth through Heterogeneous Innovations,” *Journal of Political Economy*, 2018, 126 (4), 1374 – 1443.
- Atkeson, Andrew and Ariel Burstein, “Aggregate Implications of Innovation Policy,” *Journal of Political Economy*, 2019, *forthcoming*, xxx–xxx.
- Barro, Robert J. and Gary S. Becker, “Fertility Choice in a Model of Economic Growth,” *Econometrica*, March 1989, 57 (2), 481–501.
- and Xavier Sala-i-Martin, *Economic Growth*, New York: McGraw-Hill, 1995.
- Becker, Gary S. and Robert J. Barro, “A Reformulation of the Economic Theory of Fertility,” *Quarterly Journal of Economics*, February 1988, 108 (1), 1–25.
- , Kevin M. Murphy, and Robert Tamura, “Human Capital, Fertility and Economic Growth,” *Journal of Political Economy*, October 1990, 98, S12–S37.

- Beyn, Wolf-Jürgen, Thorsten Pampel, and Willi Semmler, “Dynamic optimization and Skiba sets in economic examples,” *Optimal Control Applications and Methods*, 2001, 22 (5-6), 251–280.
- Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, “Are Ideas Getting Harder to Find?,” 2019. Stanford University manuscript.
- Bricker, Darrell and John Ibbitson, *Empty Planet: The Shock of Global Population Decline*, Crown Publishing Group, 2019.
- Buera, Francisco J and Ezra Oberfield, “The Global Diffusion of Ideas,” April 2019. Princeton University manuscript.
- Christiaans, Thomas, “Semi-Endogenous Growth when Population is Decreasing,” *Economics Bulletin*, 2011, 31 (3), 2667–2673.
- Croix, David De La and Matthias Doepke, “Inequality and Growth: Why Differential Fertility Matters,” *American Economic Review*, 2003, 93 (4), 1091–1113.
- Delventhal, Matthew J., Jesus Fernandez-Villaverde, and Nezhir Guner, “Demographic Transitions Across Time and Space,” March 2019. UPenn unpublished paper.
- Doepke, M. and M. Tertilt, “Families in Macroeconomics,” in John B. Taylor and Harald Uhlig, eds., *Handbook of Macroeconomics*, Vol. 2, Chapter 23, Elsevier, 2016, pp. 1789 – 1891.
- Doepke, Matthias and Fabrizio Zilibotti, “The Macroeconomics of Child Labor Regulation,” *American Economic Review*, December 2005, 95 (5), 1492–1524.
- and — , “Culture, Entrepreneurship, and Growth,” in Philippe Aghion and Steven N. Durlauf, eds., *Handbook of Economic Growth*, Vol. 2, Chapter 1 of *Handbook of Economic Growth*, Elsevier, 2014, pp. 1 – 48.
- Engbom, Niklas, “Firm and Worker Dynamics in an Aging Labor Market,” 2019. NYU unpublished manuscript.
- Farhi, Emmanuel and Ivan Werning, “Inequality and Social Discounting,” *Journal of Political Economy*, 2007, 115 (3), 365–402.
- Feyrer, James, Bruce Sacerdote, and Ariel Dora Stern, “Will the Stork Return to Europe and Japan? Understanding Fertility within Developed Nations,” *Journal of Economic Perspectives*, September 2008, 22 (3), 3–22.
- Futagami, Koichi and Takeo Hori, “Technological Progress and Population Growth: Do We Have Too Few Children?,” *The Japanese Economic Review*, 2010, 61 (1), 64–84.

- Galor, Oded and David N. Weil, "The Gender Gap, Fertility, and Growth," *American Economic Review*, June 1996, 86 (3), 374–387.
- and David Weil, "Population, Technology, and Growth: From the Malthusian Regime to the Demographic Transition," *American Economic Review*, September 2000, 90, 806–828.
- Golosov, Mikhail, Larry E. Jones, and Michele Tertilt, "Efficiency with Endogenous Population Growth," *Econometrica*, 2007, 75 (4), 1039–1071.
- Greenwood, Jeremy and Ananth Seshadri, "The U.S. Demographic Transition," *American Economic Review*, May 2002, 92 (2), 153–159.
- , Nezhil Guner, and Guillaume Vandenbergue, "Family Economics Writ Large," *Journal of Economic Literature*, December 2017, 55 (4), 1346–1434.
- Grossman, Gene M. and Elhanan Helpman, *Innovation and Growth in the Global Economy*, Cambridge, MA: MIT Press, 1991.
- Hock, Heinrich and David N. Weil, "On the dynamics of the age structure, dependency, and consumption," *Journal of Population Economics*, 2012, 25 (3), 1019–1043.
- Hopenhayn, Hugo, Julian Neira, and Rish Singhania, "From Population Growth to Firm Demographics: Implications for Concentration, Entrepreneurship and the Labor Share," Working Paper 25382, National Bureau of Economic Research December 2018.
- Jones, Charles I., "R&D-Based Models of Economic Growth," *Journal of Political Economy*, August 1995, 103 (4), 759–784.
- , "Population and Ideas: A Theory of Endogenous Growth," in P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, eds., *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, Princeton University Press, 2003, pp. 498–521.
- Jones, Larry and Michele Tertilt, "An Economic History of Fertility in the U.S.: 1826-1960," in Peter Rupert, ed., *Frontiers of Family Economics*, Elsevier, 2008, pp. 165–230.
- Kalemli-Ozcan, Sebnem, "Does the Mortality Decline Promote Economic Growth?," *Journal of Economic Growth*, Dec 2002, 7 (4), 411–439.
- Karahan, Fatih, Benjamin Pugsley, and Aysegul Sahin, "Demographic Origins of the Startup Deficit," Working Paper 25874, National Bureau of Economic Research May 2019.
- Kortum, Samuel S., "Research, Patenting, and Technological Change," *Econometrica*, 1997, 65 (6), 1389–1419.

- Kremer, Michael, "Population Growth and Technological Change: One Million B.C. to 1990," *Quarterly Journal of Economics*, August 1993, 108 (4), 681–716.
- Manuelli, Rodolfo E. and Ananth Seshadri, "Explaining International Fertility Differences*," *The Quarterly Journal of Economics*, 05 2009, 124 (2), 771–807.
- Moll, Benjamin, "Lecture 3: Hamilton-Jacobi-Bellman Equations," Spring 2018. Unpublished lecture slides.
- Ngai, Rachel and Roberto Samaniego, "Accounting for Research and Productivity Growth Across Industries," *Review of Economic Dynamics*, July 2011, 14 (3), 475–495.
- Peters, Michael and Conor Walsh, "Declining Dynamism, Increasing Markups and Missing Growth: The Role of the Labor Force," 2019 Meeting Papers 658, Society for Economic Dynamics 2019.
- Romer, Paul M., "Endogenous Technological Change," *Journal of Political Economy*, October 1990, 98 (5), S71–S102.
- Sasaki, Hiroaki, "Non-Renewable Resources and the Possibility of Sustainable Economic Development in a Positive or Negative Population Growth Economy," 2019. MPRA Paper No. 92204.
- , "The Solow Growth Model with a CES Production Function and Declining Population," 2019. Kyoto University unpublished manuscript.
- and Keisuke Hoshida, "The Effects of Negative Population Growth: An Analysis using a Semi-endogenous R&D Growth Model," *Macroeconomic Dynamics*, 2017, 21 (7), 1545–1560.
- Segerstrom, Paul, "Endogenous Growth Without Scale Effects," *American Economic Review*, December 1998, 88 (5), 1290–1310.
- Skiba, A. K., "Optimal Growth with a Convex-Concave Production Function," *Econometrica*, 1978, 46 (3), 527–539.
- United Nations, "World Population Prospects 2019," Technical Report 2019.

Appendix to “The End of Economic Growth? Unintended Consequences of a Declining Population”

A. Derivation of Results

Derivation of Result 1. Romer/AH/GH with Negative Population Growth

Integrate the differential equation:

$$\int \frac{dA_t}{A_t} = \alpha N_0 \int e^{-\eta t} dt$$

which gives

$$\log A_t = C_0 - \frac{\alpha N_0}{\eta} e^{-\eta t}$$

Setting $t = 0$ to solve for the constant gives

$$C_0 = \log A_0 + \frac{\alpha N_0}{\eta}$$

Next, note that $g_{A0} = \alpha N_0$. Then the time path for the stock of ideas over time:

$$\log A_t = \log A_0 + \frac{g_{A0}}{\eta} (1 - e^{-\eta t})$$

So that as $t \rightarrow \infty$,

$$\log A_t \rightarrow \log A^* \equiv \log A_0 + \frac{g_{A0}}{\eta}$$

In other words, an exponentially declining growth rate leads to a steady state level of technology and income per person.

$$y_t \rightarrow y^* \equiv \left(A_0 e^{g_{A0}/\eta} \right)^\sigma$$

Finally, converting fully into output terms using $g_y = \sigma g_A$:

$$\frac{y^*}{y_0} = e^{g_{y0}/\eta} = \exp\left(\frac{g_{y0}}{\eta}\right).$$

Derivation of Result 2. Semi-Endogenous Growth with Negative Population Growth

Integrate the differential equation:

$$\int A_t^{\beta-1} dA_t = \alpha N_0^\lambda \int e^{-\lambda\eta t} dt$$

which gives

$$\frac{1}{\beta} A_t^\beta = C_0 - \frac{\alpha N_0^\lambda}{\lambda\eta} e^{-\lambda\eta t}$$

Setting $t = 0$ to solve for the constant gives

$$C_0 = \frac{1}{\beta} A_0^\beta + \frac{\alpha N_0^\lambda}{\lambda\eta}$$

Then the time path for the stock of ideas over time:

$$A_t^\beta = A_0^\beta + \frac{\beta\alpha N_0^\lambda}{\lambda\eta} (1 - e^{-\lambda\eta t})$$

Dividing by A_0^β and noting that $g_{A0} = \alpha N_0^\lambda A_0^{-\beta}$ gives

$$\frac{A_t}{A_0} = \left(1 + \frac{\beta g_{A0}}{\lambda\eta} (1 - e^{-\lambda\eta t}) \right)^{1/\beta}$$

Converting to output using $y = A^\sigma$ and defining $\gamma \equiv \lambda\sigma/\beta$ to measure the overall degree of increasing returns to scale:

$$\frac{y_t}{y_0} = \left(1 + \frac{g_{y0}}{\gamma\eta} (1 - e^{-\lambda\eta t}) \right)^{\gamma/\lambda} \quad (25)$$

Taking the limit as $t \rightarrow \infty$,

$$\frac{y^*}{y_0} = \left(1 + \frac{g_{y0}}{\gamma\eta} \right)^{\gamma/\lambda}$$

Taking logs and derivatives of equation (25) gives the growth rate over time:

$$\frac{\dot{y}_t}{y_t} = g_{y0} \cdot \left(\frac{y_t}{y_0} \right)^{-\frac{\lambda}{\gamma}} e^{-\lambda\eta t} = \frac{g_{y0} e^{-\lambda\eta t}}{1 + \frac{g_{y0}}{\gamma\eta} (1 - e^{-\lambda\eta t})}$$

Derivation of Result 3. The Equilibrium with Endogenous Fertility

These results are all derived in the main text.

Derivation of Result 4. The Optimal Allocation with Endogenous Fertility

These results are derived in the main text and in Appendix B on transition dynamics.

B. Solving Numerically for Transition Dynamics

The numerical solution of the transition dynamics for the optimal allocation proceeds in four broad steps, first listed here and then described in more detail below:

1. We linearize the transition dynamics from the FONC associated with Hamiltonian to see that the high steady state is stable while the middle steady state is unstable with spiral dynamics.
2. We state the optimal allocation problem in terms of a value function and a Hamilton-Jacobi-Bellman (HJB) equation, which can be specialized so that the policy function depends on a single state variable (x_t).
3. We solve the HJB equation numerically; see the Matlab program `HJB.m`.
4. We solve the original Hamiltonian system numerically; see the Matlab program `HamiltonianDynamics.m`.

Solving the transition dynamics in these two ways helps in understanding the dynamics (and allows for an independent check that the solution is correct).

B.1 Linearizing around the Candidate Steady States

Using the idea production function $\dot{A}_t/A_t = N_t^\lambda A_t^{-\beta}$ as well equations (12) and (14), the FOCN conditions from the Maximum Principle for the optimal allocation can be expressed in terms of three key variables, a state-like variable $x_t \equiv A_t^\beta/N_t^\lambda$ and the two

costate-like jump variables, $V_t \equiv v_t N_t$ and $z_t \equiv \mu_t \dot{A}_t$:

$$\begin{aligned}\dot{x}_t &= \beta - \lambda f(V_t)x_t \\ \dot{V}_t &= \rho V_t - (\epsilon + \lambda z_t) \\ \dot{z}_t &= (\rho + \lambda f(V_t))z_t - \frac{\sigma}{x_t}\end{aligned}$$

where $n_t = f(V_t) \equiv \bar{b} - \delta - 1/V_t$.

In steady state, each of these key variables is constant. Linearizing this differential system around the steady state gives

$$\begin{aligned}\dot{x}_t &= -\lambda n^*(x_t - x^*) - \lambda x^* f'(V^*)(V_t - V^*) \\ \dot{V}_t &= \rho(V_t - V^*) - \lambda(z_t - z^*) \\ \dot{z}_t &= \frac{\sigma}{(x^*)^2}(x_t - x^*) + \lambda z^* f'(V^*)(V_t - V^*) + (\rho + \lambda n^*)(z_t - z^*)\end{aligned}$$

where $f'(V^*) = (1/V^*)^2$. Expressing the linearized system in matrix form with $X \equiv [x \ V \ z]'$ allows us to write it as $\dot{X}_t = B(X_t - X^*)$ where B is the matrix of coefficients, which in turn depends on various steady-state values.

We can now evaluate this linearized system around the candidate steady states using the parameter values in Table 2. First, consider the “high” steady state. The matrix B_{high} has one negative eigenvalue and two positive eigenvalues, all real, indicating a saddle-path stable steady state. In contrast, the “middle” steady state is unstable: all its eigenvalues have positive real parts, and two are imaginary. These results are computed in the Matlab program `HamiltonianDynamics.m`.

The imaginary eigenvalues that deliver the spiral dynamics around the middle steady state can disappear for alternative parameter values, for example for γ sufficiently small. In that case, the middle steady state is an unstable source with two exit paths that converge to the high and low steady states. This is intuitive if you imagine what the middle state looks like as the spiral gets smaller and smaller.

This general characterization is broadly robust to the parameter values and shows our first point: the high steady state is stable, while the middle steady state is an unstable source.

B.2 The Value Function Approach

We begin by setting up the value function using the “natural” state variables, A and N . The Hamilton-Jacobi-Bellman equation for the optimal allocation is

$$\rho V(A, N) = \max_{\ell} \sigma \log A + \log(1 - \ell) + \epsilon \log N + V_A N^\lambda A^{1-\beta} + V_N N n(\ell) \quad (26)$$

Taking the FOC with respect to ℓ gives

$$\ell = 1 - \frac{1}{bV_N N}. \quad (27)$$

This expression for the planner takes the same form as the equilibrium solution in equation (10), where $V_N N$ replaces V_t : time invested in children depends on the value of people in both cases.

The Value Function Approach, Revisited. We now show how to redefine the state variables to simplify the problem. Define new state variables to be $p_t \equiv \log N_t$ and $x_t \equiv \frac{A_t^\beta}{N_t^\lambda}$. We will refer to x_t as “knowledge per person”; because of the exponents, this is not strictly accurate, but it is convenient and helps with intuition. Notice that in terms of these states,

$$\frac{\dot{A}_t}{A_t} = \frac{1}{x_t}$$

and

$$c_t = A_t^\sigma (1 - \ell_t) = x_t^{\frac{\sigma}{\lambda}} N_t^\gamma (1 - \ell_t), \quad \gamma \equiv \lambda\sigma/\beta$$

and therefore flow utility is

$$\begin{aligned} u &= \log c + \epsilon \log N \\ &= \frac{\gamma}{\lambda} \log x + (\epsilon + \gamma) \log N + \log(1 - \ell) \\ &= \frac{\gamma}{\lambda} \log x + (\epsilon + \gamma) p + \log(1 - \ell). \end{aligned}$$

Also, the law of motion for x is

$$\dot{x}_t = \beta - \lambda n(\ell) x_t. \quad (28)$$

With this setup, the Hamilton-Jacobi-Bellman equation for the optimal allocation is

$$\rho \tilde{W}(x, p) = \max_{\ell} (\epsilon + \gamma) p + \frac{\gamma}{\lambda} \log x + \log(1 - \ell) + \tilde{W}_p n(\ell) + \tilde{W}_x [\beta - \lambda n(\ell)x] \quad (29)$$

Key progress is made by guessing a partial form of this value function. In particular, we guess that

$$\tilde{W}(x, p) = \alpha p + W(x) \quad (30)$$

for some constant α . With this guess $\tilde{W}_p = \alpha$, $\tilde{W}_{pp} = 0$, and $\tilde{W}_{xp} = 0$. Computing these derivatives of the value function in equation (31) directly, using the envelope theorem, gives

$$\rho \tilde{W}_p = (\epsilon + \gamma) + \tilde{W}_{pp} n(\ell) + \tilde{W}_{xp} [\beta - \lambda n(\ell)x]$$

Applying our guess for the form of the value function means that

$$\alpha = \frac{1}{\rho} (\epsilon + \gamma),$$

and verifies that our guess is correct: the value function can be written as a separable function, linear in the log of population.

Because $\tilde{W}_x = W_x$, we can rewrite the key part of the HJB equation as

$$\rho W(x) = \max_{\ell} \frac{\gamma}{\lambda} \log x + \log(1 - \ell) + \alpha n(\ell) + W_x [\beta - \lambda n(\ell)x.] \quad (31)$$

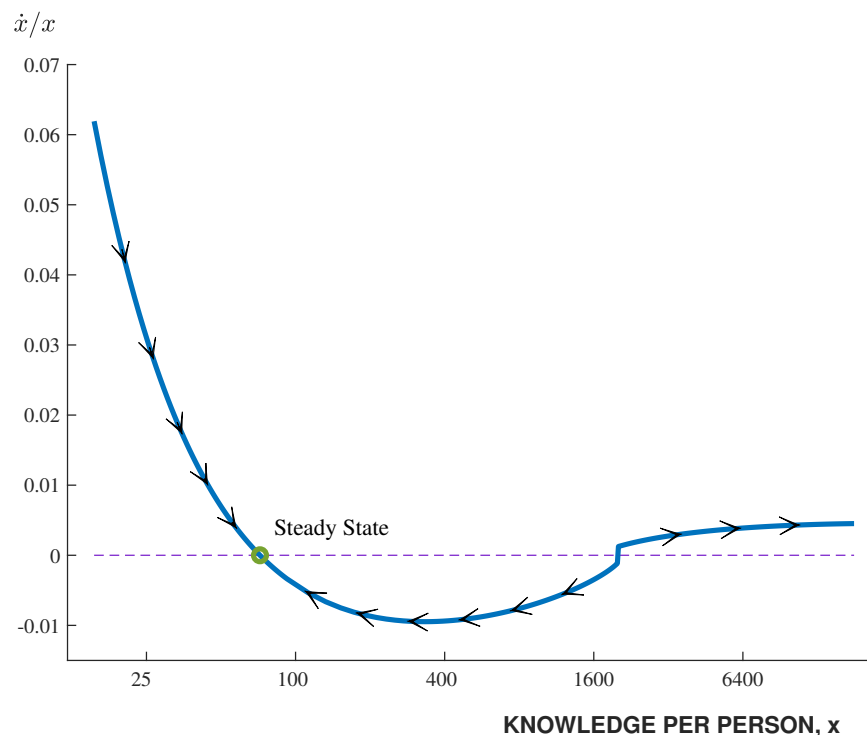
This reduces our two-state problem to a one-state problem that is much easier to analyze both graphically and numerically. Notice that under the initial condition $N_0 = 1$ (a free normalization given that we can choose the units of population), $W(x) = \tilde{W}(x, 0)$ and the value functions are equal; otherwise, welfare just shifts up based on the size of the population.

Taking the FOC with respect to ℓ gives

$$\ell = 1 - \frac{1}{\hat{b}\hat{V}} \quad (32)$$

where $\hat{V} \equiv \alpha - \lambda W_x x$. Comparing this to (27) earlier reveals that \hat{V} is the “social” value of people.

Figure 10: The Phase Diagram



Note: Numerical solution of the HJB equation using the Matlab program HJB.m.

B.3 Numerical Transition Dynamics using the HJB Value Function

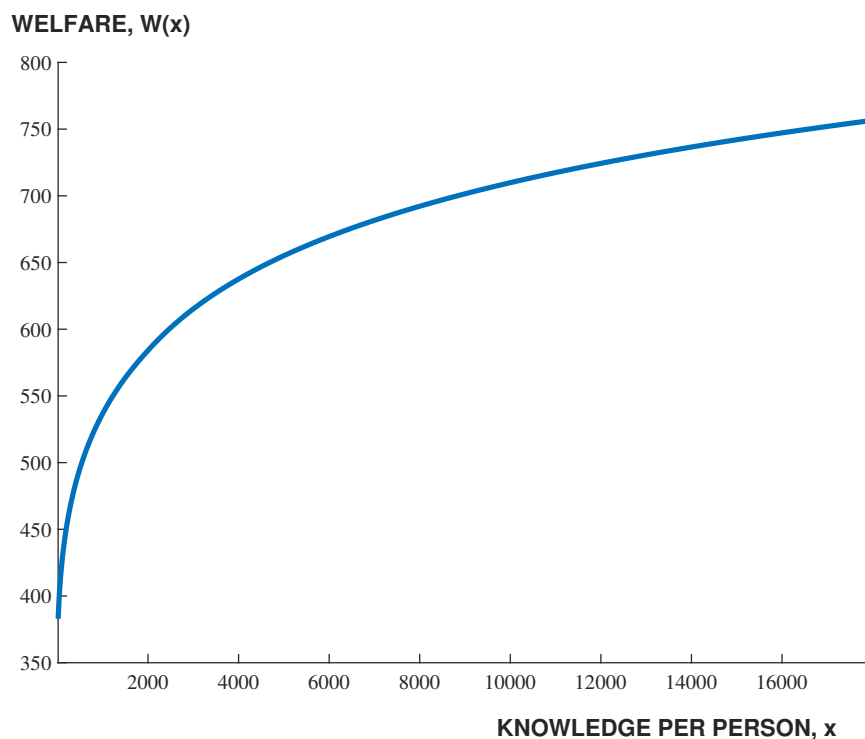
We solve for the value function numerically using the finite difference approach discussed by Moll (2018). Given that we have reduced the problem to a single state variable, the method discussed in those notes applies directly.

The only subtlety is the presence of the unstable middle candidate steady state. However, the numerical solution using the FONC from the Hamiltonian, discussed next in Section B.4 delivers values for the outer arms of the spirals, providing good guesses for the place where the transition path become vertical so that $\dot{x} = 0$. These points can be used as endpoints for the grid for x in the numerical solution.

Figure 10 shows the phase diagram in (x, \dot{x}) space. From equation (28), as $x \rightarrow \infty$, $\dot{x}/x \rightarrow -\lambda n_{eq}^* > 0$. The phase diagram makes the nature of the transition dynamics clear and justifies the dynamics shown in the main text in Figure 6.

Figure 11 shows the numerical solution of the value function, which is nicely con-

Figure 11: The Value Function



Note: From the numerical solution of the HJB equation using the Matlab program HJB.m.

cave and appears continuous despite the jump in the phase diagram. This is because the jump point occurs precisely where welfare is the same on either side of the jump.

B.4 Numerical Transition Dynamics using the Original Hamiltonian

Earlier in the appendix in Section B.1, we stated the differential system from the FONC for the Hamiltonian approach:

$$\begin{aligned}\dot{x}_t &= \beta - \lambda f(V_t)x_t \\ \dot{V}_t &= \rho V_t - (\epsilon + \lambda z_t) \\ \dot{z}_t &= (\rho + \lambda f(V_t))z_t - \frac{\sigma}{x_t}\end{aligned}$$

where $n_t = f(V_t) \equiv b(\ell(V_t)) - \delta$.

We solve this system numerically using a “reverse shooting” approach. That is, we

start from a candidate steady state, move a tiny amount away in the direction of the eigenvector of the linearized system, and solve the full nonlinear dynamics backwards to characterize the optimal path. This is especially useful in tracing out the spiral dynamics around the unstable middle steady state. For example, if we begin with $x_0 > x_{high}$, the reverse dynamics move to the right and eventually spiral into the middle steady state (the “upper arm” of the spiral). Alternatively, if we begin with $x_0 = 10,000$, i.e. close to the “low” steady state, and solve backwards, this takes us along the “lower arm” of the spiral back to the unstable middle steady state. Figure 8 in the main text shows the numerical solution of these dynamics.

The advantage of the Hamiltonian approach numerically is that it traces out the spirals “automatically,” so to speak. The HJB value function approach seeks instead to find the optimal allocation. In this application, both play important roles in understanding the transition dynamics. And of course solving the transition dynamics using two different approaches verifies that the solution is correct.