Our fragile future under the cumulative cultural evolution of two technologies

Charles Efferson[1,*], Peter J. Richerson[2], and Vanessa Weinberger[3]

[1]Faculty of Business and Economics, University of Lausanne, Switzerland

[2]Department of Environmental Science and Policy, UC Davis, USA

[3]GEMA Center for Genomics, Ecology & Environment, Universidad Mayor, Santiago, Chile

[*]Address correspondence to charles.efferson@unil.ch.

Key words: cultural evolution; human population; technological progress; sustainability

Abstract

We derive and analyse a model with unusual features characterising human activities over the long-run. First, human population dynamics draw heavily on consumer-resource modelling in ecology in that humans must consume biological resources to produce new humans. Second, the model also draws heavily from economic growth theory in that humans do not simply consume biological resources; they also produce the resources they consume. Finally, humans use two types of technology. Consumption technology affects the rate at which humans can extract resources. Production technology controls how effectively humans convert labour into new resources. The dynamics of both types of technology are subject to cumulative cultural evolutionary processes that allow both technological progress and regress. The resulting model exhibits a wide range of dynamical regimes. That said, the system is routinely sensitive to initial conditions, with wildly different outcomes given the same parameter values. Moreover, the system exhibits a basic fragility in the sense that human activities often lead to the endogenous extinction of the human species. This can happen gently, or it can follow periods of explosive human activity with super-exponential growth that ends in collapse.

1 Introduction

Imagine an evolutionary ecological system with a consumer species having the following characteristics. The consumers are smart, but not overly so [1]. They do, however, have an unusual talent for social learning [2]. They are also slowly developing cooperative breeders with cognitive and motivational psychologies that amplify the effects of being good at social learning. Specifically, they are highly cooperative, motivated to learn from and teach others [3], and prone to imitate even when they do not understand exactly how and why the behaviours they imitate are useful [4, 5]. Finally, cooperative breeding [6] helps support the costly drawn-out childhood consumers need to acquire the skills, norms, and principles they must learn from others in order to become productive adults.

Humans are, of course, a consumer species with exactly these characteristics. As a result, human populations function as collective systems for generating, disseminating, and storing knowledge. Each individual contributes a small addition to the accumulating whole, while aggregate knowledge can reach levels far beyond what any individual could produce or master on her own [7–11]. Social learning, in short, allows the population to retain the past insights of others while individuals build on these insights little by little. The result is cumulative cultural evolution. First come traditional varieties of corn, rice, wheat, and potatoes. Then come modern high-yield varieties built on decades of scientific research. First comes geometry, then topology. First Art Tatum, then Keith Jarrett.

A key question is, Does this process lead to a sustainable outcome? Humans show an extreme and undeniable ability to shape our environment and construct our ecological niches [12, 13]. But do we know where we are headed, and if so do we have the wherewithal to plan accordingly? With respect to the latter question, even in terms of basic psychology, the answer may be no. Humans have inconsistent preferences through time, and we struggle to control ourselves as a result [14, 15]. However cooperative we may be, we are also happy to make decisions that seem desirable now but will seem stupid when the future becomes the present. Today's choices readily become tomorrow's regrets. This problem can only get worse if we consider the costs we impose not on our future selves, but on future generations of other people and other species.

With respect to whether we know where we are headed, the answer may also be no. Human economies are embedded in evolutionary ecological systems, and such systems frequently generate complex dynamics [16, 17, Lenton and Scheffer, Lima *et al.* in this volume] even without the added complexity of a long-lived cooperatively breeding social learner. The ecological future is often hard to see. We can readily understand that a car is more convenient for getting around than a horse. Anticipating the many mechanisms feeding into anthropogenic climate change is not so easy [18–21]. We can readily comprehend that a modern variety of wheat yields more food than a traditional variety. Anticipating the loss of crop diversity as farmers abandon traditional varieties, with associated risks to our food supply, is not so easy [22].

Nonetheless, the past clearly suggests that we should not underestimate the power of cumulative cultural evolution. Aside from everything else, cultural evolution may yet help us predict, understand, and respond to where we are going. This is the point of departure for the present paper. As suggested above, we develop and analyse a model of consumer-resource dynamics in which the consumer has an unusual talent for social learning. Our task is to examine how this talent might shape the dynamics of human economies, including resource dynamics, the dynamics of different types of technology, and the dynamics of the human population itself.

The model rests on two core ideas that link human population dynamics, resource dynamics, and the intense productivity gains that follow from cumulative cultural evolution. First, humans do not simply consume the biological resources on which they depend; humans also produce the biological resources they consume. A golden field of wheat, ready for harvest, represents an important food source for humans. It also represents an environment that humans have heavily modified to the advantage of the wheat and to the disadvantage of many species that might otherwise live there. Humans are, in effect, resource-producing consumers. They produce wheat only to consume it. The same is true, of course, for Alpine pastures, the corn fields of Oaxaca, Cretan olive orchards, and the rice terraces of Luzon. Until the Holocene, our hunter-gatherer ancestors were pure consumers. During the Holocene, we also became producers. As domestication has proceeded apace, we have invested heavily in a few species while more or less ignoring many others we also depend on [23, Richerson *et al.* in this volume].

Second, the technologies humans use both to produce and consume resources are subject to cumulative cultural evolution, and associated cultural evolutionary processes can lead to either the loss or gain of technology [11]. In particular, a given level of technology requires a certain collective investment to maintain. If the total investment in a technology, aggregated over the entire population, is sufficiently large, the population maintains or even adds to the level of technological sophistication [24, 25]. Otherwise, the level of technology should decline, as seems to have been the case in Tasmania before European contact [24, 26].

We take both of these general ideas and adapt them to develop a model of consumer-resource dynamics in the tradition of theoretical ecology. Our model extends the classic predator-prey model of DeAngelis and coauthors [27] in numerous ways. In the classic model, the consumer allocates its time between various activities associated with acquiring prey and then consuming prey once acquired. We maintain this framework, but we extend the time allocation problem to include the production of the resource, the maintenance and development of technologies required to acquire and consume the resource, as well as the maintenance and development of technologies required to produce the resource (Supplementary Information, § 1).

2 Key ideas, core limitations

In 1971, Jay Forrester published *World Dynamics* [28]. The book included a model of the world that soon became the basis for one of the founding documents of environmentalism, namely *The Limits to Growth* [29]. Both books painted stark pictures of a future in which human population growth and consumption readily combine to reach and even temporarily surpass the unforgiving limits imposed by the global environment. The claim was that, without some fundamental change in the way humans reproduce and consume, this over-expansion happens all too soon, and collapse ensues.

Many people pointed out that the models used in both books lack technological progress [30, 31]. Indeed, the first publication of Robert Boyd, one of the founders of gene-culture coevolutionary theory, addressed exactly this point [32]. Boyd simply added technological progress to Forrester's model and arrived at relatively rosy predictions for the future. Boyd's point was not that technological progress will save us. Rather, he demonstrated the limits of such modelling efforts by showing that a given model of the world readily supports wildly different conclusions. Boyd distinguished between what he called the "Malthusian" perspective, which emphasises fixed environmental constraints, and the "technological-optimist" perspective, which emphasises technological progress and the impressive productivity gains that follow. He concluded that Forrester's model cannot resolve this difference in perspectives precisely because model results are sensitive to whether one does or does not assume technological progress.

Boyd's distinction represents two traditions that have persisted in the intervening decades. Researchers in the Malthusian tradition have continued to focus on the fundamental limits our environments impose [33–35], while researchers in the technological-optimist tradition have maintained faith in the seemingly limitless potential of our collective ingenuity [36, 37]. In the present paper, we lean towards the latter perspective with a model that allows technological progress and human productive activities to relax environmental constraints with no fundamental limit. In this sense, we consider a kind of best-case scenario for the future. Like Boyd, we do not mean to imply that technological progress will usher us gently towards such a scenario; we simply want to see what such progress might imply.

We do, however, insist on two additional ideas, both of which appear in the schematic of our model in Fig. 1. First, we insist that technological progress concerns not just production, but also consumption. We have gotten better at producing resources, but we have a much longer history of improving our technologies to acquire and consume resources. Second, as technologies advance, they require an increasingly large aggregate investment to maintain the knowledge behind the technology. Otherwise, the population loses technological sophistication via drift-like processes [24, 25]. Populations may lose technologies for other reasons, but we focus on the aggregate investment relative to what the current technology requires. As we will see, with these additions in place, even a best-case scenario is not especially optimistic.

Before we turn to the details of the model itself, we would like to discuss some of the core limitations of our approach. As mentioned above, the model includes two forms of technology, and each technology can either advance or decline based on the human population's collective investment in the technology. We model this investment by extending the time allocation problem central to many consumer-resource models in theoretical ecology [38]. Specifically, these models often start by assuming that individual consumers divide their time among multiple mutually exclusive activities like searching for resources, interacting with other consumers, and handling the resources they have already acquired [27]. We extend this approach by assuming that individual consumers can also devote time to maintaining and improving the consumption

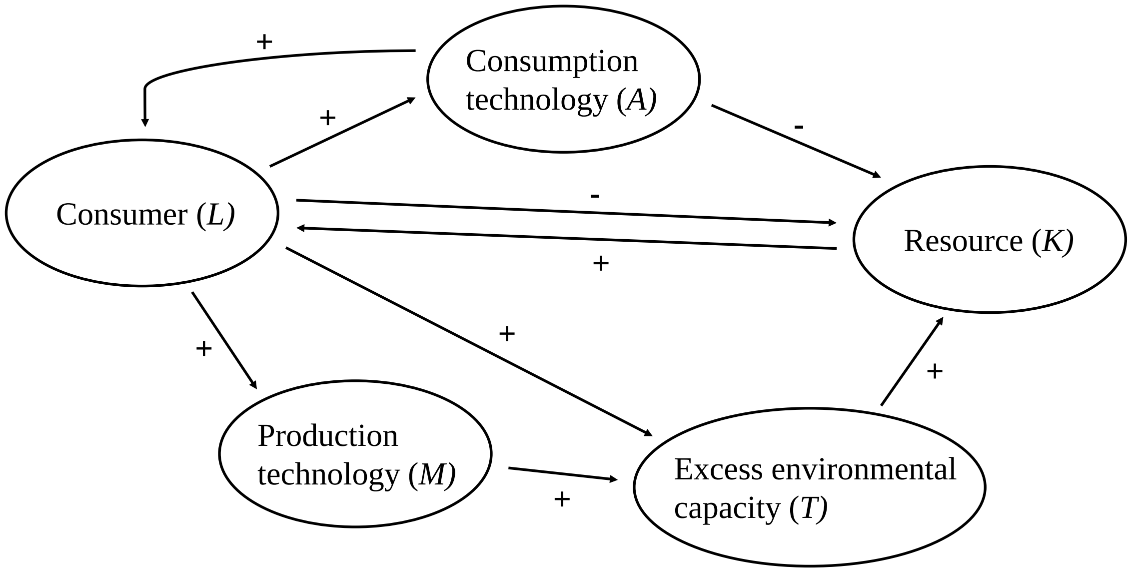

Figure 1: Schematic of variables in the model (see (2) below) and their effects on each other. K is the biological resource. L is the human population. T is the capacity of the environment to support the resource at levels in excess of a system without humans. Put differently, T measures environmental capacity that is specifically due to human productive activities. A is the technology humans use to convert their own labour into consumption of the resource. M is the technology humans use to convert their own labour into the environment's capacity (T) to support the resource. The arrows show all feedbacks between variables. Specifically, the sign shows the direction of the marginal effect of the variable at the root of an arrow on the growth of the variable at the end of the arrow under the assumption that all variables and parameters are positive. For example, the positive arrow from T to K indicates that $\partial \dot{K}/\partial T > 0$. Importantly, the feedbacks from L to the two technologies are positive. This means, conditional on allocating some time to each technology, a larger population represents a larger aggregate investment in each technology. Thus, $\partial \dot{A}/\partial L > 0$, and $\partial \dot{M}/\partial L > 0$. One or both of the technologies, however, may nonetheless regress ($\dot{A} < 0$ or $\dot{M} < 0$) if the aggregate investment is not sufficient to maintain the technology at its current level, a possibility that seems to be central to cumulative cultural evolutionary processes [24, 25].

technology or maintaining and improving the production technology.

With this approach, we do not explicitly model social learning and the cultural evolutionary dynamics of technologies that result. Our approach to cumulative cultural evolution of technologies is instead somewhat phenomenological; the only first principles involved are related to time allocation. People devote time to maintaining and developing a technology, and this affects whether the technology improves or worsens through time. Moreover, time allocation is parameterised. Consumers cannot reallocate their time, as an example, by reducing the time they spend on the production technology and increasing the time they spend consuming resources.

Time allocation is exogenous and thus fixed. Alternative models could allow time allocation to be continually optimised or to change through time because of social learning and attendant cultural evolutionary dynamics. Such models are easy to imagine but probably hard to analyse.

In addition, we treat the consumer population as homogeneous. For example, we can choose parameter values appropriate for an early Neolithic population. Most people engage in food production, and the average time allocated to resource production per consumer is thus relatively high. Alternatively, we can choose parameter values for a contemporary industrialised population. Most people do not produce food, and thus the average time allocated to resource production per consumer is extremely low. The model readily accepts parameter values consistent with the averages in both of these scenarios. It ignores, however, the underlying heterogeneity responsible for the low average in the second scenario. In contemporary industrial societies, the average time allocated to food production is not low because everyone spends 10 seconds a day farming. It is low, of course, because of the division of labour [Lenton and Scheffer in this volume]. Only a handful of farmers produce food for everyone, which in turn frees up capacity for other specialised forms of labour like those needed for armies, saxophone playing, and research in human evolutionary ecology. Because our model only works with time allocation averages, it is silent about the effects of such subtleties.

Similarly, the conversion of consumed resources into new humans is also parameterised. Accordingly, we can choose parameter values representing a society in which people tend to have many children and invest little in each child, or we can choose values for a society in which people invest heavily in few children. Our model cannot capture an endogenous transition between these two regimes [39, Lima *et al.* in this volume]. All in all, we cannot pinpoint the joint effect of our model's limitations, but identifying these limitations explicitly highlights our primary contribution. Namely, we examine a potentially unlimited system in which consumers use two distinct technologies to convert their labour into the production and consumption of resources, and both technologies are subject to cumulative cultural evolutionary processes. Again, even though the system has no fundamental limit, model results suggest that taking technological progress seriously does not necessarily lead one to an optimistic vision of the future.

3 Model overview and simulation methodology

To explain the key principles, we begin with a resource species, K_D, that consumers, L_D, both produce and consume (Supplementary Information, § 2). In the absence of consumption, resource dynamics are based on the logistic model. The carrying capacity, however, has both an exogenous component and an endogenous component that depends on human production of the resource. Specifically, the carrying capacity is $K_{\max} + T_D$. The quantity $K_{\max} \in \mathbb{R}_{++}$ (i.e. $\{K_{\max} \in \mathbb{R} \,|\, K_{\max} > 0\}$) is a parameter, and it specifies the carrying capacity of the environment for the resource species in the absence of production. T_D is a state variable that summarises the effects of production. It has the same units as K_D and $K_{\max}$, and we assume that $T_D \in \mathbb{R}_+$ (i.e. $\{T_D \in \mathbb{R} \,|\, T_D \geq 0\}$). Thus, we limit attention to a world in which human activities can potentially increase the capacity of the environment to sustain the resource, but they cannot reduce this capacity below some exogenous lower limit, namely $K_{\max}$. Moreover, because T_D is unbounded, the carrying capacity is variable and takes values in $[K_{\max}, \infty)$, and we can also think about cases in which $T_D \to \infty$. The lack of an upper bound, of course, is implausible, but this is what we mean when we say we consider a kind of best-case scenario in a system with no fundamental limit.

The dynamics of T_D depend on a production function with inputs consisting of both T_D and effective human labour. Effective human labour is a summary measure of the work, so to speak, that goes into humans producing the resource. Effective human labour depends on the size of the human population, L_D, the time allocated to production, and the technology, M_D, used to transform human effort into the environment's capacity to support the resource. $M_D \in \mathbb{R}_+$ is thus a state variable and one of two technologies in the model. As the technology that supports biological production, it spans the cultural evolutionary history ranging from our very first attempts to modify environments in favour of the species we value to all the technologically-intensive processes that constitute modern agriculture.

The other technology in the model is $A_D \in \mathbb{R}_+$, which is also a state variable. Like M_D, the variable A_D represents technology in the sense that it controls how human effort is converted into something else. Whereas M_D converts human effort into an environment that supports the growth of the resource, A_D converts human effort into acquiring resources for consumption. A_D, in short, captures the cultural evolutionary processes that range from simple wooden clubs to modern harvesting combines and the various industries that prepare and distribute our food around the world.

Finally, we complete the model by specifying technology dynamics. For each of the two technologies, dynamics unfold according to the basic idea in Henrich's discussion of the Tasmanian toolkit [24]. Namely, if human investment in technology is sufficiently high, technology progresses. If sufficiently low, technology regresses. We adapt this idea by incorporating it into the consumer's time allocation problem (Supplementary Information, § 2) at the centre of our model and other models in theoretical ecology [38]. To see the intuition, let x_1 be the total time actually invested in maintaining and developing the consumption technology, aggregated over all humans, divided by the total time required to just maintain this technology at its current level, A_D. With x_1 defined, the function $f_1 : \mathbb{R}_+ \longrightarrow \mathbb{R}$ maps x_1 to the change in A_D per unit of A_D per unit of time,

$$f_1(x_1) = \delta_1(-1 + x_1^{\gamma_1}). \tag{1}$$

If no one invests any time in the technology, then $x_1 = 0$, and A_D declines exponentially based on $\delta_1 \in \mathbb{R}_{++}$. If the total time invested is exactly what is required, then $x_1 = 1$, and the technology neither regresses nor progresses, $f_1(1) = 0$. More broadly, if $x_1 < 1$, technology regresses in the precise sense that the change in technology per unit technology per unit of time is negative. If $x_1 > 1$, technology progresses in the opposite sense. We are silent about exactly how $f_1(x_1)$ increases in x_1. Although $\gamma_1 \in (0, 1)$ would be a natural assumption, we do not limit attention to these values and instead simply assume $\gamma_1 > 0$.

To complete the model of A_D dynamics, we need to specify x_1. As explained, x_1 is the total time invested in the technology divided by the total time required to maintain the technology at its current level. The total time invested is derived from the time allocation problem for consumers (Supplementary Information, § 2). In other words, consumers divide their time between several different tasks, one of which is investing in the technology used to consume the resource.

For the time required to maintain the technology, we assume (Supplementary Information, § 2) this value is proportional to $A_D^{\lambda_1}$, where $\lambda_1 > 0$. If $\lambda_1 < 1$, the total time required to maintain A_D is a concave function, and thus the average time required per unit of technology declines as technology progresses. If $\lambda_1 > 1$, the total time required to maintain A_D is a convex function, and thus the average time required per unit of technology rises as technology progresses. Although we do not repeat the logic for M_D, the production technology, the logic behind M_D dynamics is exactly parallel. That said, the relevant parameters and functions that control M_D dynamics are independent of those that control A_D dynamics. The only necessary link between the two technologies is that time invested in A_D cannot be invested in M_D and vice versa (Supplementary Information, § 2).

The Supplementary Information (§ 2) presents the full derivation of the model. Following a dimensional analysis, the unit-free version of the model is the following, where we have dropped the "D" subscripts to indicate that all quantities are pure numbers without units (Fig. 1). For

Table 1: A summary of parameters in model (2). For each parameter, we provide an intuitive description of the parameter's role in the model. We use the qualifier "intuitive" here because all parameters listed are actually unit-free parameter combinations. Thus, their complete and technically precise definitions are much more elaborate than these intuitive definitions (Supplementary Information, § 2), but without much additional insight. All parameters are non-negative, and in addition $\beta \in [0, 1]$.

Unit-free parameter	Intuitive description
θ	Intrinsic growth of resource
σ	Converts T into carrying capacity
χ	Time handling acquired resources
η	Time managing interspecific interactions
μ	Converts acquired resources into new humans
β	When producing T, controls the returns to T vs. ML
ξ	Decay of T in absence of production
γ_1	Controls how time allocated to A contributes to A
λ_1	Controls the total time required to maintain A at current level
ψ	Decay of A in the absence of investments
γ_2	Controls how time allocated to M contributes to M
λ_2	Controls the total time required to maintain M at current level
ω	Decay of M in the absence of investments

the moment, we assume all parameters (Table 1) are strictly positive.

$$\begin{aligned}
\dot{K} &= \frac{\theta K \left(1 + \sigma T - K\right)}{1 + \sigma T} - \frac{AKL}{1 + \chi AK + \eta L} \\
\dot{L} &= \frac{\mu AKL}{1 + \chi AK + \eta L} - L \\
\dot{T} &= T^{\beta} \left(ML\right)^{1-\beta} - \xi T \\
\dot{A} &= A^{1-\lambda_1\gamma_1} L^{\gamma_1} - \psi A \\
\dot{M} &= M^{1-\lambda_2\gamma_2} L^{\gamma_2} - \omega M
\end{aligned} \tag{2}$$

The system has two steady states that readily admit local stability analyses. The first of these steady states is $\left(\hat{K}, \hat{L}, \hat{T}, \hat{A}, \hat{M}\right) = (0, 0, 0, 0, 0)$, which is locally unstable (Supplementary Information, § 2). The second is $\left(\hat{K}, \hat{L}, \hat{T}, \hat{A}, \hat{M}\right) = (1, 0, 0, 0, 0)$, which is locally stable (Supplementary Information, § 2). As explained below, the model often ends up converging to this latter state, after some period of human activity, and for this reason we will refer to this state below as the "post-human state". Additional steady states exist, but they are not analytically tractable.

We now explain the simulations we used to make further headway.

As a general strategy for managing further analyses of a quite complex model, we implemented the following protocol. We first defined a bounded region of parameter space (Supplementary Information, § 3) in which the bounds were minimally restrictive. We then randomly selected a point in this space and numerically estimated any associated steady states using the `nleqslv` package [40] in `R` [41]. For each steady state, we used the `rootSolve` package [42] to calculate the Jacobian matrix evaluated at the steady state, along with associated eigenvalues. Finally, we randomly selected five initial conditions distributed around the steady state and numerically estimated the dynamics of the system using the `deSolve` package [43]. We repeated this exercise until we had reached 100 steady states. Because some combinations of parameter values yielded more than one steady state, the number of unique points in parameter space considered was not 100, but rather 69. Of the 100 steady states considered, only four of them were numerically estimated to be locally stable in the sense that the dominant eigenvalue was negative. We implemented the protocol described here for each of the 100 steady states, regardless of whether or not we estimated the steady state to be locally stable.

We chose 100 steady states and five initial conditions per steady state in an effort to automate, at least to some extent, the simulation project given the practical challenges associated with numerically solving ordinary differential equations. In practical terms, programming a fully general routine that implements a differential equations solver across a wide range of conditions can be challenging. As parameter values and initial conditions vary, the numerical properties of the system vary, which means the challenges created for the solver also vary. As an intuitive illustration, imagine a system that grows exponentially versus one that takes a long time to converge to a steady state. If we naively ask the computer to simulate both systems under exactly the same conditions, we quickly discover one problem or another based on whether we choose a short time frame or a long time frame. Under a short time frame, the first system continues to take values small enough for the computer to handle, but the second system does not have time to converge. Under a long time frame, the second system has time to converge, but the first system eventually takes on values too large for the computer to handle. In our case, we found that 100 steady states and five initial conditions allowed us automate the process to some extent, although we still had to isolate and manually simulate the system for a few of our 100 steady states. We include our `R` code for this exercise as supplementary materials, which would allow the interested reader to repeat the entire exercise with a new set of 100 steady states based on random points in parameter space.

4 Results

The model supports five different types of steady state. In addition to the two discussed above, three other types of steady state exist. A simple inspection of system (2) shows that these additional steady states must involve positive values for L and A. Humans cannot exist without consumption technology, and consumption technology cannot persist without humans. Thus, a steady state with humans present implies that consumption technology is also present, and vice versa. The same is not true for production. Although production (T, M) cannot exist without humans, humans can exist without production precisely because the environment always has some exogenous potential to support the resource upon which humans depend.

Accordingly, the three additional types of steady state all involve $\hat{K} > 0$, $\hat{L} > 0$, and $\hat{A} > 0$. Given this condition, the three types of steady state include (i) $\hat{T} = 0$ and $\hat{M} = 0$, (ii) $\hat{T} = 0$ and $\hat{M} > 0$, and (iii) $\hat{T} > 0$ and $\hat{M} > 0$. We ignore type (ii) steady states because they imply that consumers continue to invest in maintaining a production technology (i.e. $\hat{M} > 0$) that they do not use (i.e. $\hat{T} = 0$). Because time allocation is parameterised in the model, this kind of steady state can arise under some combinations of parameter values, but it is not a realistic scenario. Only two of our 100 steady states based on randomly selected parameter values took this form.

Across the 100 steady states considered, an extremely common outcome is that the system eventually converges on $\left(\hat{K}, \hat{L}, \hat{T}, \hat{A}, \hat{M}\right) = (1, 0, 0, 0, 0)$. This is the post-human state, a state in which the resource is present at its exogenous carrying capacity, but the consumer and everything associated with the consumer is extinct. The path to the post-human state, however, often involves some period of growth in the human population, with associated growth in production and technology, followed by collapse. This cycle of growth and decline can be quite tame in the sense that the maximum size of the human population is roughly the same order of magnitude as the initial population size. Fig. 2a shows an example.

In contrast, the cycle of growth and decline can also be spectacular, with the human population exhibiting super-exponential growth to reach a maximum size far beyond its initial size, only to collapse and go extinct. To our knowledge, this is the only model of human population dynamics that exhibits this pattern of super-exponential growth followed by endogenous collapse. Even more surprisingly, this outcome does not necessarily occur just because parameters take certain values. Rather, the outcome hinges on initial conditions. Fig. 2b shows an example. Depending on initial conditions, the population either converges smoothly to the post-human state, or it exhibits a wild period in which human activity explodes super-exponentially before collapse and extinction.

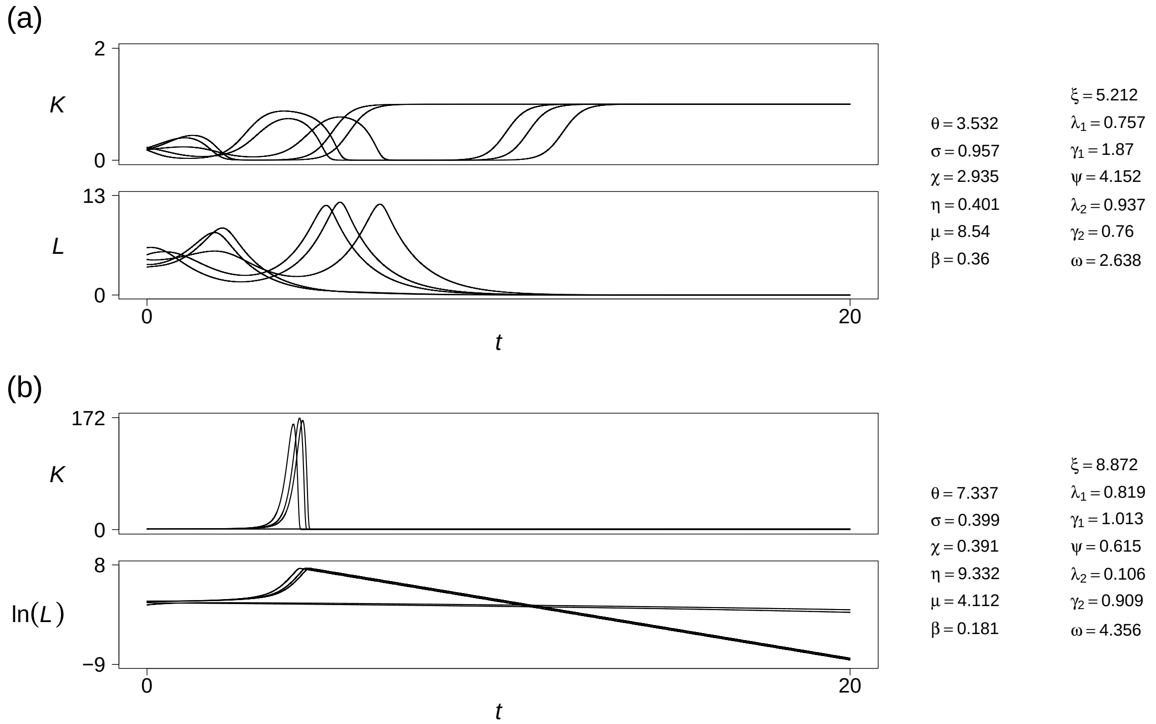

Figure 2: **(a)** Dynamics of model exhibiting relatively tame cycles before converging to $(\hat{K}, \hat{L}, \hat{T}, \hat{A}, \hat{M}) = (1, 0, 0, 0, 0)$, a state in which humans are extinct. **(b)** Dynamics of model exhibiting extreme boom-bust cycles dependent on initial conditions. For three initial conditions, the system grows in spectacular fashion, with human population growth (L) super-exponential (i.e. $\ln(L)$ convex), only to collapse and converge on human extinction. For these three populations, even the resource (K) goes extinct in the simulations, although analytical results show that this steady state would not be stable in the face of perturbations. For the remaining two initial conditions, the system converges in a relatively modest way to a state in which humans are extinct. Parameter values to the right. The graphs only show K and L for brevity. See Supplementary Information (§ 3) for more detail, including the dynamics of T, A, and M.

The model also produces persistent cycles in which all the state variables are positive, but again this outcome exhibits a clear dependence on initial conditions. To illustrate, Fig. 3 shows two sets of simulations under the same parameter values. Although the parameter values are the same, the two graphs differ in the sense that the initial conditions are distributed around two different steady states associated with these parameter values. For some initial conditions, the system converges on stable cycles in which all variables remain positive (Fig. 3a,b). For other initial conditions (Fig. 3b), in contrast, the system converges quickly and smoothly to the post-human state in which productive consumers and their technologies are all extinct.

Continuing in the same vein, the model also exhibits regimes of growth or decline that depend on initial conditions. Fig. 4a shows dynamics under a single set of parameter values, but with initial conditions having a profound effect. As the graph shows, one of two basic outcomes holds. Either the system quickly enters a state in which all state variables grow exponentially,

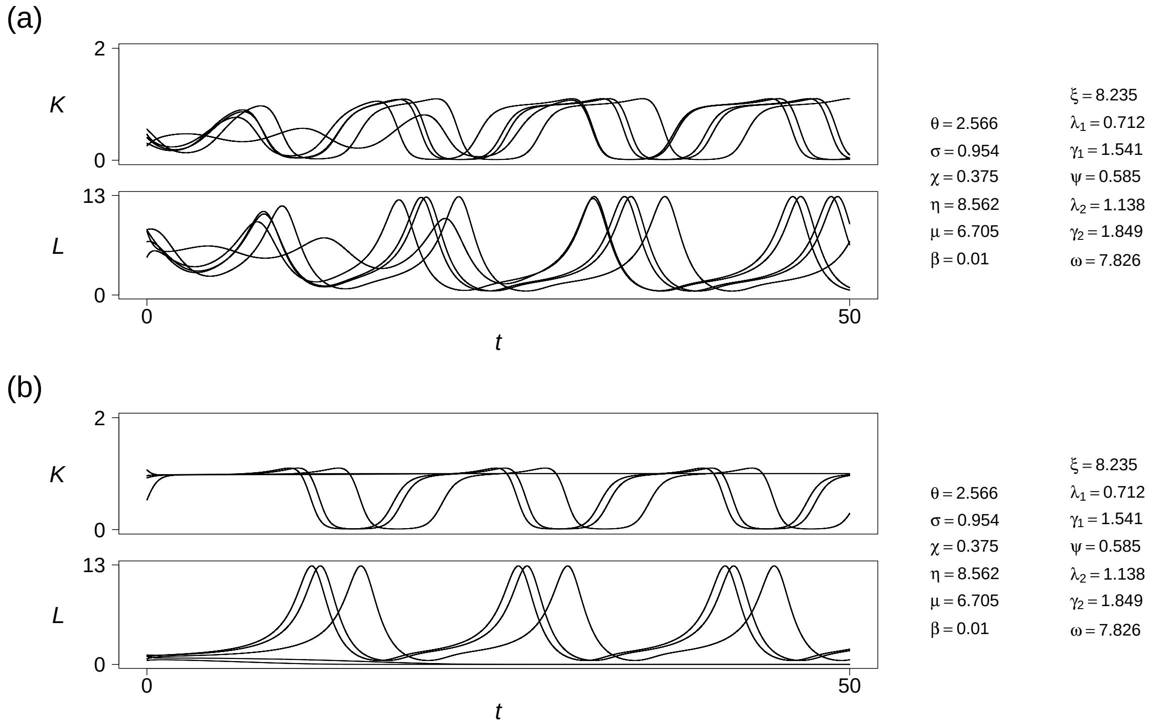

Figure 3: **(a)** Dynamics of model with persistent cycles in which all variables remain positive. **(b)** Dynamics of model with persistent cycles dependent on initial conditions. Simulated populations can exhibit persistent cycles, but for some initial conditions the system converges smoothly to a state in which humans are extinct. Parameter values to the right. The graphs only show K and L for brevity. See Supplementary Information (§ 3) for more detail, including the dynamics of T, A, and M.

or the system converges on the post-human state (Fig. 4a). The outcomes could hardly be more different, but the underlying mechanism is simply a matter of initial conditions.

Finally, the model also supports dynamical regimes in which all state variables converge on positive values, but again initial conditions play a decisive role. Fig. 4b shows the system either settling on positive values for all state variables (Supplementary Figure 12) or, with the same parameter values but different initial conditions, a state in which humans are extinct. When the system converges on a state in which humans and their technologies persist in a stable equilibrium, this outcome only occurs after the entire system completes a number of transient cycles. Moreover, to repeat the common refrain, human extinction remains a possible outcome depending on initial conditions (Fig. 4b).

The results above highlight two characteristics of model (2). First, the model supports many different dynamical regimes, including gentle convergence to a steady state, exponential growth, limit cycles, and transient cycles with super-exponential growth then collapse of the human population. Second, initial conditions routinely play a decisive role in terms of which dynamical regime obtains. The importance of initial conditions per se is perhaps not surprising given the

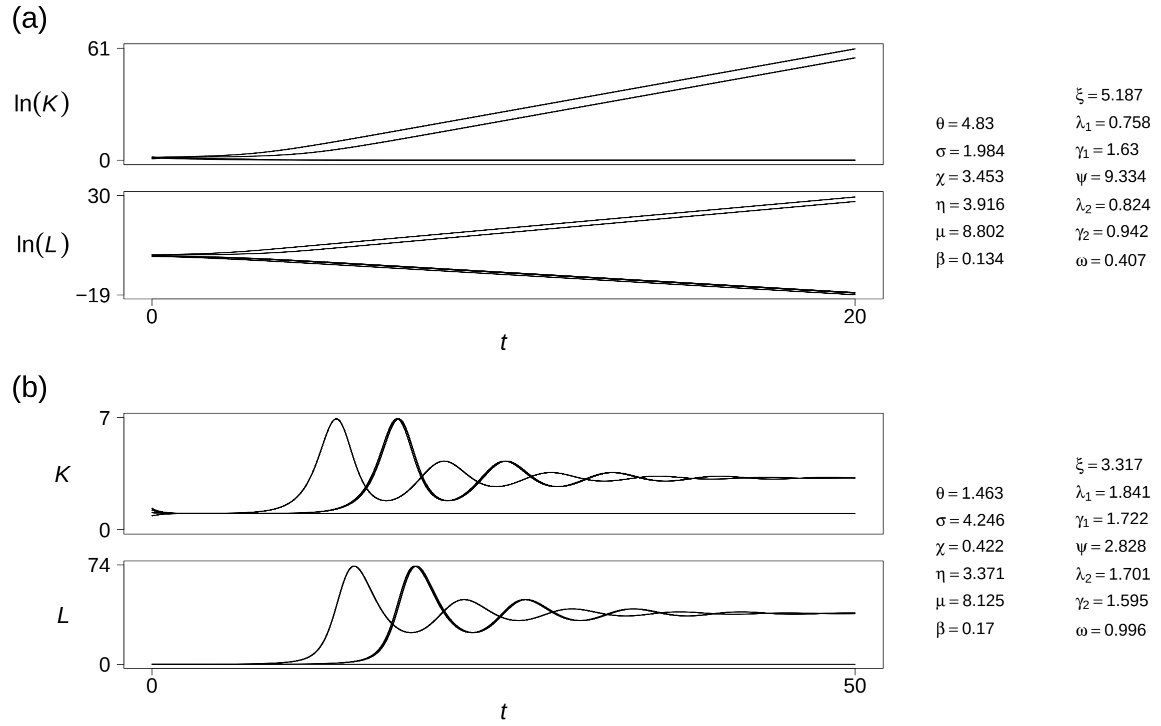

Figure 4: **(a)** Dynamics of model converging, based on initial conditions, on either exponential growth for all variables or the extinction of humans. **(b)** Model dynamics exhibiting, depending on initial conditions, one of two regimes, either transient cycles followed by positive values for all variables or human extinction. Parameter values to the right. The graphs only show K and L for brevity. See Supplementary Information (§ 3) for more detail, including the dynamics of T, A, and M.

assumptions that underlie our approach to technology evolution.

Specifically, our approach to technology evolution closely follows the spirit of Henrich's model of cumulative cultural evolution [24], even if the details are entirely different. The key idea is that a human population represents a distributed means of storing the technological knowledge accumulated in the population's cultural evolutionary past. By extension, the characteristics of the population must shape whether or not it can maintain and add to this knowledge. If it can, technology progresses, which may support human population growth and thus the potential to accumulate more technological knowledge. If it cannot, technology regresses, which presumably leads to a trajectory in the opposite direction. The dominance of initial conditions and associated path-dependence we routinely observe from our model is built into this key idea.

Although the importance of initial conditions in general may not be surprising, the fact that extreme differences in outcomes seem to follow routinely from differences in initial conditions is. The selected results discussed above show that single points in parameter space routinely support dramatically different regimes. In an attempt to identify additional patterns, we ran a cluster analysis on the 69 randomly selected points in parameter space using the `pvclust` package [44] in R. The analysis identified 13 non-degenerate clusters (supplementary materials), and we then inspected and characterised the dynamics across all simulations within each of these clusters.

The results parallel the discussion above closely. We did not identify any clear tendency for a given cluster to exhibit a certain kind of dynamics. Instead, parameter combinations within a cluster regularly support a wide array of dynamical regimes, with the regimes that occur frequently a matter of initial conditions. This finding is, of course, consistent with the tendency for single points in parameter space to support multiple dynamical regimes. That said, one common theme is evident. All 13 clusters seem to support dynamics leading to the post-human state in which only the resource remains. This implies a surprising fragility in the system. If we interpret parameter values as the fundamentals of the human society, so to speak, our model suggests that the potential for societal collapse is generally present in a way that is largely independent of the fundamentals. This is especially surprising given that we have ignored the possibility that humans can degrade the resource's environment via pollution or the exploitation of non-renewable resources.

5 Additional results for reduced model with Type I functional response

Finally, we consider parameter values that should be appropriate for societies with efficient technologies for handling resources and a highly institutionalised division of labour. In many contemporary industrial societies, for example, the proportion of the population actually producing and handling our food is small, largely because food production is an industrial process in which machines do much of the work. This means the average time per consumer spent handling resources approaches zero. Moreover, such societies often have highly developed markets and enforceable contracts to manage and regulate the distribution of our food with minimal conflict. This means the average time per consumer spent negotiating intraspecific interactions related to the distribution of food also approaches zero. As the average handling time and average time engaged in intraspecific interactions both approach zero, the functional response characteristic of the model by DeAngelis and colleagues [27], namely $AK/(1 + \chi AK + \eta L)$, takes a simplified form. Specifically, in the limit one can show that $\chi = 0$ and $\eta = 0$ (Supplementary Information, § 2). As a result, the functional response simply becomes AK, which is often called a "Type I" functional response. To reiterate the caveats discussed in § 2, we can parameterise the model in this way to match certain properties of contemporary industrial societies. In doing so, however, we work only with population averages, and we ignore the fact that $\chi, \eta \to 0$ specifically because of the highly developed division of labour that holds in these societies [Lenton and Scheffer in this volume]. Intuitively, such an approach captures the average effects of a division of labour, but it does not represent the division of labour itself.

To simplify model (2) further, we make three additional assumptions. First, we assume $\mu = 0.1$. This assumption reflects the fact that μ captures the conversion of resource biomass into consumer biomass, and the energy losses along the way ensure that this conversion should occur at a ratio much lower than one-to-one, especially in contemporary societies. Second, we assume that $\lambda_1 = \lambda_2 = 1$. This assumption simply means that the total time required to maintain a given level of technology grows in a linear fashion (Supplementary Information, § 2). Intuitively, if a bow and arrow and a kayak are equally complex pieces of technology, maintaining both technologies requires twice as many person-hours as maintaining each on its own. Third, we assume that $\gamma_1 = \gamma_2$ and $\psi = \omega$. This assumption simply means that, although A and M are technologies used for different activities, both are subject to the same kinds of cumulative cultural evolutionary processes. Thus, their dynamics work in the same way. Importantly, this assumption does not mean that the two types of technology must grow or decline together. Rather, it simply means the parameter values underlying cultural evolutionary dynamics are the same.

With these assumptions in place, a simplified version of the model follows for societies with industrial production and an institutionalised division of labour.

$$\begin{aligned}
\dot{K} &= \frac{\theta K\left(1+\sigma T-K\right)}{1+\sigma T} - AKL \\
\dot{L} &= (0.1)(AKL) - L \\
\dot{T} &= T^{\beta}\left(ML\right)^{1-\beta} - \xi T \\
\dot{A} &= A^{1-\gamma_1} L^{\gamma_1} - \psi A \\
\dot{M} &= M^{1-\gamma_2} L^{\gamma_2} - \omega M
\end{aligned} \tag{3}$$

Note that, although we retain γ_2 and ω in the notation for model (3), we are assuming $\gamma_1 = \gamma_2$ and $\psi = \omega$. For this model, we first examined all 144 parameter combinations based on $\theta \in \{0.1, 10\}$, $\sigma \in \{0.1, 10\}$, $\beta \in \{0.1, 0.5, 0.9\}$, $\xi \in \{0.1, 10\}$, $\gamma_1 \in \{0.1, 0.5, 0.9\}$, and $\psi \in \{0.1, 10\}$. Using `nleqslv`, we identified 94 steady states in this parameter space. For any given steady state, we used `deSolve` to simulate dynamics from various initial conditions randomly distributed around the steady state. Additionally, we also simulated dynamics for each of the

144 points in parameter space from random initial conditions without regard for steady states. We include the code for both approaches as supplementary material.

This exercise produced three conclusions. First, the model exhibits an extraordinary tendency to converge sooner or later to a state in which consumers do not exist. Sometimes the system moves smoothly toward such a state. Sometimes it cycles first, and sometimes it does so in spectacular fashion (e.g. Fig. 5a). Either way, most of the 144 parameter combinations we considered readily lead to the disappearance of humans.

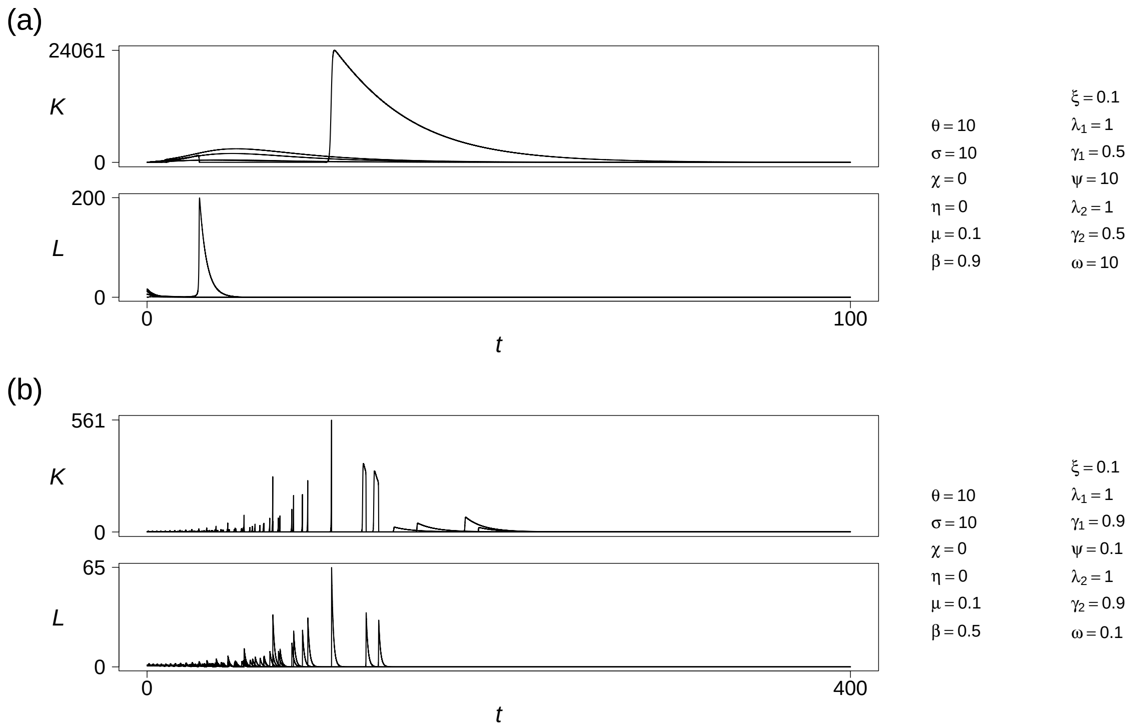

Figure 5: **(a)** An example of model (3) exhibiting a spectacular cycle before collapse. For four of the initial conditions shown here, the population converges smoothly on human extinction. For one initial condition, however, the system cycles first, with extreme changes in K and L in the process. **(b)** Model (3) exhibiting increasing cycles before collapse. For all five initial conditions, the system shows cycles of increasingly large amplitude until it collapses, with humans and their technologies extinct. Parameter values to the right. The graphs only show K and L for brevity. See Supplementary Information (§ 3) for more detail, including the dynamics of T, A, and M.

Second, however, humans do not always disappear, and the dynamical regimes in which humans persist share specific features. In particular, θ is always 10, never 0.1, which means that the intrinsic growth potential of the resource is not too small. In addition, γ_1 and γ_2 are always 0.5 or 0.9, never 0.1, which means that, in terms of maintaining and producing technology, the returns to human effort are also not too small. Finally, ψ and ω are always 0.1, never 10, which means that existing technological knowledge does not decay too fast. When these conditions hold, consumers may persist, but we find no evidence of a growth regime. Persistence consists of simple convergence towards a steady state or perhaps some gentle cycling that diminishes over time.

Finally, however, the region of parameter space in question by no means guarantees that consumers will persist. Fig. 5b provides a counterexample. Although the parameter values are consistent with regimes that sometimes lead to the persistence of humans, all initial conditions in Fig. 5b lead to moderate cycling at first. Cycles, however, increase in amplitude until the nadirs of the cycles start to flirt with extinctions in L. As cycles continue to grow, humans eventually go extinct.

6 Discussion

Our model attempts to bridge the strange and regrettable gap between theoretical ecology and economic growth theory, a gap that mirrors the equally strange and regrettable gap between population biology and economics more broadly. To try and bridge this gap, our model combines multiple ideas, some of which are common in one field but rare in the other. First, atypical in economic growth theory but typical in theoretical ecology, we tie the dynamics of the consumer population directly to resources consumed. Although this basic idea seems obvious from a biological perspective, economic growth theory has a long tradition of simply positing that human population growth is exogenous and thus exponential [45]. Perhaps this is simply a convenient simplification for modelling societies that have escaped the Malthusian trap by transitioning to low fertility. Nonetheless, the difficulty with models of exogenous growth is that, once they start to go wrong, they go really wrong really fast. In any case, more recent work in economic growth theory has abandoned the traditional approach. Economic growth theorists and human evolutionary ecologists alike have addressed the linked dynamics of humans and the resources upon which humans depend. Much of this work has focused on the all-important distinction between increasing resources that are transformed into an increasingly large human population versus increasing resources that are transformed into an increasingly wealthy human population [26, 46–51].

Second, atypical in theoretical ecology but typical in economic growth theory, we assume that the consumer produces resources with a technology that evolves endogenously [52]. Technology transforms inputs, one of which is human effort, into outputs. Technological progress means more output for the same input, and an important finding from growth theory is that sustained economic growth requires technological progress [53, 54]. This claim pertains to production technology (e.g. M), not consumption technology (e.g. A), and our results are broadly consistent with the same idea. Although our model does not seem to generate sustained growth readily, when it does, sustained progress in the production technology occurs in parallel (Supplementary Information, § 3). Interestingly, the consumption technology also progresses steadily in these cases. This means that, when sustained growth happens, the consumer is getting better at both producing and consuming the resource.

Third, atypical in economic growth theory but typical in a sense in theoretical ecology, we assume that the consumer consumes the resource with a technology that evolves endogenously. Interestingly, production technology gets all the attention in economic growth theory; technology transforms resources and labour into new resources via production. The field tends to assume, often tacitly, that consumption does not require technology, but of course this is incorrect. The evolution of consumption technology has a vastly longer history than the evolution of production technology. Humans have been developing technologies to transform resources and labour into new resources for maybe 10,000 years as domestication began to evolve in the early Holocene [55, 56]. In contrast, we have been developing technologies to transform labour into consuming resources, resources whose existence we did not contribute to directly, for much longer. The earliest stone tools date to three million years ago, and these tools were widespread in Eurasia by two million years ago (Richerson *et al.*, this volume). Population biology implicitly recognises the importance of consumption technology in the sense that the field has a long tradition of studying coevolution between predators and their prey [57–59]. Although we may not traditionally think of predator evolution in such systems as examples of technology evolution, they are. The use of the word "technology" may only seem unusual because we tend to reserve the word for humans.

What is unusual is that we combine the idea of a resource-producing consumer with the consumer's reliance on multiple forms of technology subject to cumulative cultural evolution. Specifically, we assume that any given technology is essentially distributed know-how based on the unusual propensity humans have for social learning. As a technology progresses, it consists of an increasing amount of knowledge that is more widely distributed across the human population. At some point, a technology becomes sufficiently advanced and complex that individuals must specialise; the aggregate far exceeds what any single individual can master. Technology in this case becomes a truly collective phenomenon, and even the most talented of individuals may know only a tiny fraction of what the collective knows.

Crucially, this kind of technology evolution involves some special properties. As a collective phenomenon, technology requires collective maintenance [24]. Moreover, as technology progresses, the collective requirements in terms of maintaining the technology increase in tandem. If the collective investment in maintenance is not sufficient, technology not only fails to progress;

it regresses. Consequently, for societies with advanced technologies, technology can regress even if the population is relatively large [25]. This idea represents a key distinction between our approach and approaches in economic growth theory. Although many models in economic growth theory also posit a positive relationship between technological growth and human population size [60, 61], they do not typically assume that technological regression can occur in a way that depends jointly on both the population size and the current degree of technological complexity (Fig. 1).

All in all, what does this interdisciplinary constellation of ideas suggest about cultural evolution and a sustainable human future? We would like to offer three qualitative findings. First, our model exhibits an exceedingly diverse range of dynamics that includes smooth convergence to steady states, transient cycles, apparent limit cycles, and stable growth regimes. Some of these dynamical regimes resemble the super-exponential growth that humans have exhibited so far in the Holocene and Anthropocene [62]. Second, dynamics are routinely path-dependent. For a given combination of parameter values, depending on initial conditions, the system can grow indefinitely or collapse. It can converge on the post-human state or cycle before converging on a steady state in which humans persist. It can converge on limit cycles or the post-human state. Routine path-dependence, of course, undercuts the notion of economic fundamentals and their supposed effects on growth and prosperity [63]. Moreover, humans with a capacity for accumulating advanced technologies have existed for eight glacial/inter-glacial cycles, each of roughly 100,000 years. Perhaps the Holocene's trademark shift to resource-producing consumption was possible in previous inter-glacials, but initial conditions were only appropriate in the Holocene.

Third, the system demonstrates a basic instability. It supports a wide range of dynamical regimes in which humans disappear sooner or later. Sometimes the disappearance of humans is a gentle process. Sometimes it is exactly the opposite, even to the extent that the human population can grow super-exponentially before collapse. To our knowledge, our model is the only one to date that supports such a regime. Other models that also consider technological progress support super-exponential growth in the human population [64–67]. Instead of collapse, however, the human population reaches infinity in finite time, an obviously impossible scenario. We would summarise this instability as follows. When resource-producing consumers excel at social learning and cumulative cultural evolution, the result is a kind of fragility in which the resource-producing consumer generates its own extinction under a wide range of conditions. This remains true even with endogenous technology dynamics and no fundamental limit to the system. Importantly, this does not mean that the resource-producing consumer does not enjoy periods of success. Quite the opposite, they can be wildly productive with explosive growth before collapse. Moreover, long-run growth regimes are possible. They seem, however, to occur under a relatively narrow range of conditions.

Importantly, any model of large-scale human systems in an exercise in presumption. All models leave out details, which is part of their appeal, but the problems this can create may be especially serious for models of the sort we present here. Tractability requires compromise, and the compromises are potentially extreme when thinking about human populations and economies over extended time scales. We would like to close by discussing some of our results in light of the compromises we made. First, as discussed, consumers in our model do not optimise, not even in some crude approximate sense. The time per capita invested in each activity related to the consumption and production of resources is fixed. Agents cannot adjust their behaviour in response to changes in the relative benefits of these activities. Our intuition is that a model that would allow such adjustments would attenuate the instabilities we find with our model. That said, empirical evidence shows that people are not especially good at making smart choices for the future [14, 15].

In any case, by fixing time allocation, we can consider different combinations of parameter values that we think represent different periods of human history, but we cannot capture the endogenous transition between major periods [Lima *et al.* in this volume]. In spite of this limitation, we can in principle ask the following question. Given present conditions as a starting point, can a strategically selected set of parameter values put us on a sustainable trajectory? The challenges with this question, however, are legion. The model leaves out much of what surely matters [Lenton and Scheffer in this volume], and evaluating the model is difficult against the single observation that human history represents. Moreover, human societies in the past have routinely suffered random shocks. When coupled with the sensitivity to initial conditions routinely on display here, the model's predictive value may be limited, especially over time scales shorter than those associated with the model's tendency to transition to the post-human state.

Finally, consumers in our model can improve the environment for the resource species without limit, but they cannot degrade this environment. Given that we ignore environmental degradation, the propensity for collapse our model exhibits is especially surprising. Our intuition is that including environmental degradation would exacerbate this propensity. The upshot is that, when consumption and production technologies are both subject to cumulative cultural evolution in a system with no fundamental limit, even our collective ingenuity can lead to a fragile future. Despite the model's limitations, this result effectively turns the technological-optimist's perspective on its head.

Acknowledgements and funding

We thank Matthias Schief for comments on an earlier version of the paper. C.E. acknowledges the support of the Swiss National Science Foundation (Nr. 100018_185417/1). V.P.W. acknowledges the support of a Postdoctoral Fellowship from Agencia Nacional de Investigación y Desarrollo de Chile (ANID 3190609).

Author contributions

C.E. developed and analysed the models. C.E. wrote the paper with input and insights from P.J.R. and V.W.

Data and code

Available as supplementary materials.

Competing interests

We have no competing interests.

References

[1] Joseph Henrich. *The Secret of Our Success: How Culture is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton University Press, 2017.

[2] Peter J. Richerson and Richerson Boyd. *Not By Genes Alone: How Culture Transformed the Evolutionary Process*. Chicago: University of Chicago Press, 2005.

[3] L. G. Dean, R. L. Kendal, S. J. Schapiro, B. Thierry, and K. N. Laland. Identification of the social and cognitive processes underlying human cumulative culture. *Science*, 335(6072):1114–1118, 2012. doi: 10.1126/science.1213969. URL `http://www.sciencemag.org/content/335/6072/1114.abstract`.

[4] Gyoergy Gergely, Harold Bekkering, and Ildikó Király. Rational imitation in preverbal infants. *Nature*, 415(14 February):755, 2002.

[5] Maxime Derex, Jean-François Bonnefon, Robert Boyd, and Alex Mesoudi. Causal understanding is not necessary for the improvement of culturally evolving technology. *Nature Human Behaviour*, 3(5):446–452, 2019.

[6] Sarah Blaffer Hrdy. *Mothers and Others: The Evolutionary Origins of Mutual Understanding*. Cambridge, MA: Harvard University Press, 2011.

[7] Magnus Enquist, Stefano Ghirlanda, Arne Jarrick, and C-A Wachtmeister. Why does human culture increase exponentially? *Theoretical Population Biology*, 74(1):46–55, 2008.

[8] Magnus Enquist, Stefano Ghirlanda, and Kimmo Eriksson. Modelling the evolution and diversity of cumulative culture. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1563):412–423, 2011.

[9] Lewis G Dean, Gill L Vale, Kevin N Laland, Emma Flynn, and Rachel L Kendal. Human cumulative culture: a comparative perspective. *Biological Reviews*, 89(2):284–301, 2014.

[10] Michael Muthukrishna, Ben W Shulman, Vlad Vasilescu, and Joseph Henrich. Sociality influences cultural complexity. *Proceedings of the Royal Society B: Biological Sciences*, 281 (1774):20132511, 2014.

[11] Oren Kolodny, Nicole Creanza, and Marcus W Feldman. Evolution in leaps: the punctuated accumulation and loss of cultural innovations. *Proceedings of the National Academy of Sciences*, 112(49):E6762–E6769, 2015.

[12] F. John Olding-Smee, Kevin N. Laland, and Marcus W. Feldman. *Niche Construction: The Neglected Process in Evolution*. Princeton: Princeton University Press, 2003.

[13] Kevin N. Laland, Kim Sterelny, John Odling-Smee, William Hoppitt, and Tobias Uller. Cause and effect in biology revisited: Is Mayr's proximate-ultimate dichotomy still useful? *Science*, 334(6062):1512–1516, 2011. doi: 10.1126/science.1210879. URL `http://www.sciencemag.org/content/334/6062/1512.abstract`.

[14] Shane Frederick, George Loewenstein, and Ted O'Donoghue. Time discounting and time preference: a critical review. *Journal of Economic Literature*, 40(2):351–401, 2002.

[15] Keith Marzilli Ericson and David Laibson. Intertemporal choice. In *Handbook of Behavioral Economics: Applications and Foundations 1*, volume 2, pages 1–67. Elsevier, 2019.

[16] Robert M. May. Simple mathematical models with very complicated dynamics. *Nature*, 261(10 June):459–467, 1976.

[17] J. M. Cushing, R. F. Costantino, Brian Dennis, Robert A. Desharnais, and Shandelle M. Henson. *Chaos in Ecology: Experimental Nonlinear Dynamics*. Academic Press, 2003.

[18] J Doyne Farmer, Cameron Hepburn, Matthew C Ives, T Hale, Thomas Wetzer, Penny Mealy, Ryan Rafaty, Sugandha Srivastav, and Rupert Way. Sensitive intervention points in the post-carbon transition. *Science*, 364(6436):132–134, 2019.

[19] Matto Mildenberger and Dustin Tingley. Beliefs about climate beliefs: the importance of second-order opinions for climate politics. *British Journal of Political Science*, 49(4): 1279–1307, 2019.

[20] Ilona M Otto, Jonathan F Donges, Roger Cremades, Avit Bhowmik, Richard J Hewitt, Wolfgang Lucht, Johan Rockström, Franziska Allerberger, Mark McCaffrey, Sylvanus SP Doe, et al. Social tipping dynamics for stabilizing earth's climate by 2050. *Proceedings of the National Academy of Sciences*, 117(5):2354–2365, 2020.

[21] Sara M Constantino, Gregg Sparkman, Gordon T Kraft-Todd, Cristina Bicchieri, Damon Centola, Bettina Shell-Duncan, Sonja Vogt, and Elke U Weber. Scaling up change: A critical review and practical guide to harnessing social norms for climate action. *Psychological Science in the Public Interest*, 23(2):50–97, 2022.

[22] Stephen B Brush. *Farmers? Bounty: Locating Crop Diversity in the Contemporary World.* Yale University Press, 2008.

[23] Peter J. Richerson, Robert Boyd, and Robert L. Bettinger. Was agriculture impossible during the pleistocene but mandatory during the holocene?: a climate change hypothesis. *American Antiquity*, 66(3):387–411, 2001.

[24] Joseph Henrich. Demography and cultural evolution: how adaptive cultural processes can produce maladaptive losses–the Tasmanian case. *American Antiquity*, 69(2):197–214, 2004.

[25] Maxime Derex, Marie-Pauline Beugin, Bernard Godelle, and Michel Raymond. Experimental evidence for the influence of group size on cultural complexity. *Nature*, 503(7476):389, 2013.

[26] Michael Kremer. Population growth and technological change: One million bc to 1990. *The Quarterly Journal of Economics*, 108(3):681–716, 1993.

[27] D. L. DeAngelis, R. A. Goldstein, and R. V. O'Neill. A model for trophic interaction. *Ecology*, 56(4):881–892, 1975.

[28] Jay W Forrester. *World Dynamics*. Wright Allen, 1971.

[29] Donella H. Meadows, Dennis L. Meadows, Jørgen Randers, and William W. III Behrens. *The Limits to Growth*. Universe Books, 1972.

[30] Mahbub ul Haq. " the limits to growth:" a critique. *Finance and Development*, 9(4):2, 1972.

[31] William Ophuls. Technological limits to growth revisited. *Alternatives: Perspectives on Society, Technology and Environment*, 4(2):4–11, 1975.

[32] Robert Boyd. World dynamics: a note. *Science*, 177(4048):516–519, 1972.

[33] Robert Costanza and Herman E Daly. Natural capital and sustainable development. *Conservation Biology*, 6(1):37–46, 1992.

[34] Johan Rockström, Will Steffen, Kevin Noone, Åsa Persson, F Stuart Chapin, Eric F Lambin, Timothy M Lenton, Marten Scheffer, Carl Folke, Hans Joachim Schellnhuber, et al. A safe operating space for humanity. *Nature*, 461(7263):472–475, 2009.

[35] Will Steffen, Katherine Richardson, Johan Rockström, Sarah E Cornell, Ingo Fetzer, Elena M Bennett, Reinette Biggs, Stephen R Carpenter, Wim De Vries, Cynthia A De Wit, et al. Planetary boundaries: Guiding human development on a changing planet. *Science*, 347(6223):1259855, 2015.

[36] Ester Boserup. *Population and Technological Change: A Study of Long-Term Trends*. University of Chicago Press, 1981.

[37] Julian Lincoln Simon. *The Ultimate Resource 2*. Princeton University Press, 1996.

[38] Ted J. Case. *An Illustrated Guide to Theoretical Ecology*. Oxford: Oxford University Press, 2000.

[39] Oded Galor and David N. Weil. From Malthusian stagnation to modern growth. *The American Economic Review*, 89(2):150–154, 1999.

[40] Berend Hasselman. *nleqslv: Solve Systems of Nonlinear Equations*, 2022. URL `https://CRAN.R-project.org/package=nleqslv`. R package version 3.3.3.

[41] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2021. URL `https://www.R-project.org/`.

[42] Karline Soetaert. *rootSolve: Nonlinear root finding, equilibrium and steady-state analysis of ordinary differential equations*, 2009. R package 1.6.

[43] Karline Soetaert, Thomas Petzoldt, and R. Woodrow Setzer. Solving differential equations in R: Package deSolve. *Journal of Statistical Software*, 33(9):1–25, 2010. doi: 10.18637/jss.v033.i09.

[44] Ryota Suzuki and Hidetoshi Shimodaira. Pvclust: an r package for assessing the uncertainty in hierarchical clustering. *Bioinformatics*, 22(12):1540–1542, 2006.

[45] Robert J. Barro and Xavier Sala-I-Martin. *Economic Growth*. Cambridge: The MIT Press, 2nd edition, 2004.

[46] Monique Borgerhoff Mulder. The demographic transition: are we any closer to an evolutionary explanation? *Trends in Ecology and Evolution*, 13(7):266–277, 1998.

[47] Hillard Kaplan, Jane Lancaster, and Arthur Robson. Embodied capital and the evolutionary economics of the human life span. *Population and Development Review*, 29:152–182, 2003.

[48] Oded Galor. From stagnation to growth: unified growth theory. In Philippe Aghion and Steven N. Durlauf, editors, *Handbook of Economic Growth*, pages 171–294. Amsterdam: Elsevier, 2005.

[49] Lesley Newson and Peter J Richerson. Why do people become modern? a darwinian explanation. *Population and Development Review*, 35(1):117–158, 2009.

[50] Quamrul Ashraf and Oded Galor. Dynamics and stagnation in the malthusian epoch. *American Economic Review*, 101(5):2003–41, 2011.

[51] Nico Voigtländer and Hans-Joachim Voth. How the west “invented” fertility restriction. *American Economic Review*, 103(6):2227–64, 2013.

[52] VP Weinberger, C Quiñinao, and PA Marquet. Innovation and the growth of human population. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372 (1735):20160415, 2017.

[53] Robert M. Solow. A contribution to the theory of economic growth. *The Quarterly Journal of Economics*, 70(1):65–94, 1956.

[54] Trevor W Swan. Economic growth and capital accumulation. *Economic Record*, 32(2): 334–361, 1956.

[55] Dorian Q Fuller, Tim Denham, Manuel Arroyo-Kalin, Leilani Lucas, Chris J Stevens, Ling Qin, Robin G Allaby, and Michael D Purugganan. Convergent evolution and parallelism in plant domestication revealed by an expanding archaeological record. *Proceedings of the National Academy of Sciences*, 111(17):6147–6152, 2014.

[56] Greger Larson, Dorian Q Fuller, et al. The evolution of animal domestication. *Annual Review of Ecology, Evolution, and Systematics*, 45(1):115–136, 2014.

[57] Peter A Abrams and HIROYUKI Matsuda. Fitness minimization and dynamic instability as a consequence of predator–prey coevolution. *Evolutionary Ecology*, 11(1):1–20, 1997.

[58] Michael H Cortez and Joshua S Weitz. Coevolution can reverse predator–prey cycles. *Proceedings of the National Academy of Sciences*, 111(20):7486–7491, 2014.

[59] Ramith R Nair, Marie Vasse, Sébastien Wielgoss, Lei Sun, Yuen-Tsu N Yu, and Gregory J Velicer. Bacterial predator-prey coevolution accelerates genome evolution and selects on virulence-associated prey defences. *Nature Communications*, 10(1):1–10, 2019.

[60] Philippe Aghion and Peter Howitt. A model of growth through creative destruction, 1990.

[61] Paul M Romer. Endogenous technological change. *Journal of Political Economy*, 98(5, Part 2):S71–S102, 1990.

[62] Andrej V Korotaev. *Introduction to social macrodynamics: compact macromodels of the world system growth*. Editorial URSS, 2006.

[63] Samuel Bowles. *Microeconomics: Behavior, Institutions, and Evolution*. New York: Russell Sage, 2004.

[64] Joel E Cohen. Population growth and earth's human carrying capacity. *Science*, 269(5222): 341–346, 1995.

[65] Anders Johansen and Didier Sornette. Finite-time singularity in the dynamics of the world population, economic and financial indices. *Physica A: Statistical Mechanics and its Applications*, 294(3-4):465–502, 2001.

[66] Luís MA Bettencourt, José Lobo, Dirk Helbing, Christian Kühnert, and Geoffrey B West. Growth, innovation, scaling, and the pace of life in cities. *Proceedings of the National Academy of Sciences*, 104(17):7301–7306, 2007.

[67] Vyacheslav I Yukalov, Elizaveta P Yukalova, and Didier Sornette. Extreme events in population dynamics with functional carrying capacity. *The European Physical Journal Special Topics*, 205(1):313–354, 2012.