

ABSTRACT

Title of Dissertation: ENSEMBLE KALMAN INVERSE PARAMETER ESTIMATION FOR HUMAN AND NATURE DYNAMICS TWO

Maia Karpovich, Doctor of Philosophy in Applied Mathematics & Statistics, and Scientific Computation, 2023

Dissertation Directed by: Eugenia Kalnay
Department of Atmospheric and Oceanic Science and Institute for Physical Science and Technology

Safa Mote
Department of Atmospheric and Oceanic Science and Institute for Physical Science and Technology

Since the widespread development of agriculture 10,000 years ago and particularly since the Industrial Revolution beginning in the 18th century, the coupled Earth and Human systems have experienced transformative change. The world's population and Gross Domestic Product have each increased by factors of at least eight in the last two centuries, powered by the intensive use of fossil energy and fossil water. This has had dramatic repercussions on the Earth system's stability, threatened by Human system activities such as habitat destruction, global warming, and depletion of Regenerating and Nonrenewable energy resources that increasingly alter environmental feedbacks.

To analyze these changes, we have developed the second generation of the Human and

Nature Dynamics model, HANDY2. HANDY2 is designed to simulate the dynamics of energy resources and population over the Industrial era from 1700 to 2200, flexibly incorporating real-world observations of population and energy consumption in an expanded suite of mechanisms that track capital investment, labor force allocation, class mobility, and extraction and production technologies. The use of automated Ensemble Kalman Inversion (EnKI) estimation for HANDY2's parameters allows us to accurately capture the rapid 20th-century rise in the use of phytomass and fossil fuels, as well as the global enrichment of Elites that puts pressure on natural resources and Commoners. EnKI-derived HANDY2 ensembles project that current world policies may lead to a collapse in the world's population by 2200 caused by rapid depletion of resources. However, this collapse can be prevented by a combination of actions taken to support voluntary family planning, lower economic inequality, and most importantly, invest in the rapid expansion of Renewable energy extraction.

ENSEMBLE KALMAN INVERSE PARAMETER ESTIMATION FOR
HUMAN AND NATURE DYNAMICS TWO

by

Maia Karpovich

Dissertation submitted to the faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2023

Dissertation Committee:

Distinguished University Professor Eugenia Kalnay, Chair
Assistant Professor Safa Mote, Co-Chair
Professor William Dorland
Professor Brian Hunt
Professor Ning Zeng

© Copyright by
Maia Karpovich
2023

Dedication

This dissertation is dedicated to Anna Harren, whose memory lives on in all of us.

The guardian stands resolute.

Acknowledgments

I thank my advisors, Professors Eugenia Kalnay and Safa Mote, and my external co-advisor, Jorge Rivas, for creating the original HANDY model and the original HANDY2, and for constant mentorship and feedback in its development. Their generous support was instrumental in bringing the project to completion.

I thank Professors Brian Hunt, William Dorland, Konstantina Trivisa, and Ning Zeng for their advice and suggestions about the direction of my dissertation research, and for graciously agreeing to serve on my dissertation committee.

I thank Dr. Eviatar Bach for suggesting the use of EnKI for parameter estimation in HANDY2.

I thank Professor Amy Sapkota and Dr. Rianna Teresa Murray for their work and teaching in the UMD Global STEWARDS program, and for providing a platform to discuss HANDY2 in the context of the food-energy-water (FEW) nexus.

This research was supported in part by NRT-INFEWS: UMD Global STEWARDS (STEM Training at the Nexus of Energy, Water Reuse and Food Systems) that was awarded to the University of Maryland School of Public Health by the National Science Foundation National Research Traineeship Program, Grant number 1828910.

This research was additionally supported by the Department of Mathematics' Hauptman Summer Fellowship, 2023-2024, which was generously funded by Carol Fullerton and Distin-

guished University Research Professor James A. Yorke.

I additionally thank Professor James Yorke for providing feedback on an earlier version of HANDY2 during a critical period in its development.

Table of Contents

Preface	ii
Dedication	ii
Acknowledgements	iii
Table of Contents	v
List of Tables	vii
List of Figures	viii
List of Abbreviations	x
Chapter 1: Introduction	1
1.1 Why HANDY?	1
1.2 Outline of Dissertation	4
Chapter 2: Reproduction of HANDY1	7
2.1 Model Structure	7
2.2 Replication of HANDY1 Results	9
Chapter 3: Structure and Results of HANDY2	14
3.1 Overview of HANDY2 Structure	14
3.2 HANDY2 Key Concepts and Sectors	16
3.2.1 Types of Energy	16
3.2.2 Extraction, Production, Efficiency, and Use Technologies	17
3.2.3 Labor Force and Investment	19
3.2.4 Births and Deaths	20
3.2.5 Class Mobility and Inequality	23
3.2.6 Burnable Carbon, Reserves, and Resources	25
3.3 HANDY2 Equations and Parameters	27
3.3.1 Variables and Parameters	27
3.3.2 Equations	31
3.4 HANDY2 Data and Estimates	39
3.4.1 Estimates of Regenerating and Nonrenewable Stocks, and Renewable Flows	39
3.4.2 Population Estimates	44

3.5	Results	46
3.5.1	Scenario 1: Baseline	47
3.5.2	Scenarios 2 and 2-1: No Nonrenewables	54
3.5.3	Scenario 3: No Renewables	58
3.5.4	Scenario 4: Inequality Reduction	60
3.5.5	Scenario 5: Higher Renewables Use	62
3.5.6	Scenario 6: Higher Renewables Use and Lower Birth Rates	67
3.6	Conclusion	69
Chapter 4: Overview of EnKI and Parameter Estimation		72
4.1	EnKI Mathematics	73
4.1.1	The Algorithm	73
4.1.2	Properties of EnKI	76
4.2	EnKI for Lorenz96	77
Chapter 5: HANDY1 Parameter Estimation		83
5.1	The EnKI Implementation for HANDY1: Building on Lorenz96	84
5.2	Parameter Estimation Results	88
5.2.1	Egalitarian Scenarios	88
5.2.2	Equitable Scenarios	92
5.2.3	Unequal Scenarios	95
5.2.4	Soft Landings in an Unequal Society	100
5.3	Conclusion	102
Chapter 6: HANDY2 Parameter Estimation		104
6.1	Overview	104
6.2	Adapting EnKI for HANDY2	105
6.2.1	Learned Parameters and Starting Priors	109
6.3	Results	113
6.3.1	Convergence Dynamics	113
6.3.2	Examining Sample Runs	116
6.3.3	Parameter Estimates	118
6.4	Ensemble HANDY2 Runs	128
6.4.1	Scenario 1: Baseline	129
6.4.2	Scenarios 2 and 2-1: No Nonrenewables	129
6.4.3	Scenario 3: No Renewables	130
6.4.4	Scenario 4: Inequality Reduction	131
6.4.5	Scenario 5: Higher Renewables Use and Inequality Reduction	131
6.4.6	Scenario 6: Higher Renewables Use, Inequality Reduction, and Lower Birth Rates	132
6.5	Conclusion	133
Bibliography		144

List of Tables

3.1	HANDY2 variables and parameters	28
3.2	HANDY2 baseline scenario parameter values	48
5.1	HANDY1 Egalitarian parameter estimation values	92
5.2	HANDY1 type-N collapse parameter estimation values	97
5.3	HANDY1 type-L collapse parameter estimation values	99
6.1	HANDY2 EnKI priors	112
6.2	HANDY2 EnKI ensemble parameter values	122

List of Figures

2.1	HANDY1 Egalitarian runs	10
2.2	HANDY1 Equitable runs	11
2.3	HANDY1 Unequal runs	12
3.1	Relative vs absolute population growth	22
3.2	Extraction efficiency ξ	35
3.3	Gross energy use	40
3.4	Population increase over time	45
3.5	Baseline scenario plots	52
3.6	Regenerating-only scenario plots	55
3.7	Regenerating-only scenario plots, with flat ψ_R	57
3.8	No-Renewables scenario plots	59
3.9	Inequality reduction in HANDY2	61
3.10	Inequality reduction scenario plots	63
3.11	Higher Renewables use scenario plots	64
3.12	Higher Renewables use scenario population change	66
3.13	Higher Renewables and lower birth rates scenario plots	68
3.14	Higher Renewables and lower birth rates scenario, population change	70
4.1	Lorenz96 parameter estimation, iteration 1	81
4.2	Lorenz96 parameter estimation, iteration 5	82
5.1	Single-parameter estimation of δ , Egalitarian HANDY1	90
5.2	Multi-parameter estimation, Egalitarian HANDY1	91
5.3	Multi-parameter estimation, Equitable HANDY1	94
5.4	Single-parameter estimation of δ , Unequal HANDY1	96
5.5	Multi-parameter estimation, Unequal type-N collapse	98
5.6	Multi-parameter estimation, Unequal type-L collapse	99
5.7	Estimation of β , HANDY1 Unequal soft-landing scenario	101
6.1	EnKI error during convergence of 10 HANDY2 estimation runs	114
6.2	HANDY2 parameter estimation, seed 1	119
6.3	HANDY2 parameter estimation, seed 2	120
6.4	HANDY2 parameter estimation, seed 3	121

6.5	Baseline scenario EnKI ensemble	136
6.6	EnKI ensemble for Regenerating-only scenario	137
6.7	EnKI ensemble for Regenerating-only scenario with flat ψ_R	138
6.8	EnKI ensemble for no-Renewables scenario	139
6.9	EnKI ensemble for inequality reduction scenario	140
6.10	Inequality reduction in the EnKI ensemble	141
6.11	EnKI ensemble for higher Renewables use scenario	142
6.12	EnKI ensemble for higher Renewables and lower birth rates scenario	143

List of Abbreviations

C	Carbon
CAFE	Corporate Average Fuel Economy
CCS	Carbon capture and storage
CE	Common Era
CLiMA	Climate Modeling Alliance
CSV	Comma-separated value
EJ	Exajoule (10^{18} joule)
EKS	Ensemble Kalman Sampler
EnKF	Ensemble Kalman Filter
EnKI	Ensemble Kalman Inversion
EU	European Union
GDP	Gross Domestic Product
HANDY	Human and Nature Dynamics (all versions)
HANDY1	Human and Nature Dynamics, Generation 1
HANDY2	Human and Nature Dynamics, Generation 2
HYDE	History of the Global Environment Database
IPCC	Intergovernmental Panel on Climate Change
LED	Light-emitting diode
RCP	Representative Concentration Pathway
RK4	Runge-Kutta method
RMS	Root mean square
RNG	Random number generator
TFR	Total fertility rate
TJ	Terajoule (10^{12} joule)
TW	Terawatt (10^{12} watt)
UN	United Nations
US	United States
ZJ	Zettajoule (10^{21} joule)

Chapter 1: Introduction

1.1 Why HANDY?

The Earth and Human systems, when considered together, have experienced rapid and accelerating change since the start of widespread agriculture about 10000 years before present. Over that time period, the world population has grown from several million to eight billion [Klein Goldewijk et al., 2017], and so has the world's wealth, driven by an ever-increasing allocation of Earth's food, energy, and water resources. At first, that was done on a small scale, with local communities burning primarily woody phytomass for cooking food and heating dwellings [Smil, 2017]. But with improving technology, and particularly with the discovery of fossil fuels and fossil water in the Industrial Revolution, no part of the globe remains untouched. Today's rich society has dramatically altered the balance of the carbon, nitrogen, and phosphorus cycles [Galloway et al., 2004, Le Quéré et al., 2015, Steffen et al., 2015, Zeng et al., 2014], depleted aquifers throughout the globe, engaged in widespread habitat destruction in a bid to gain more arable land [Sodhi and Ehrlich, 2010], emitted large amounts of trace metals into the atmosphere [Pacyna and Pacyna, 2001], and has caused an unprecedented climate crisis through anthropogenic global warming [Hansen et al., 2013, McGlade and Ekins, 2015, Steffen et al., 2015]. These trends show no signs of abating, as the Human system continues to grow with respect to the Earth system [Motesharrei et al., 2016, Steffen et al., 2015].

Despite this overall growth, individual human existence has been extremely precarious for most people at most times in history. Subsistence farmers, 80% of the population in Rome, often only produced a bare surplus of food for themselves and their families; in bad years, widespread famine was common [Erdkamp, 2005]. The agricultural surplus supported a small elite class engaged in many other forms of economic activity, that could be organized more effectively with increasing state complexity, or could seize more land or resources with military power [Erdkamp, 2005, Rosenstein, 2004]. Yearly and century-scale variations in harvests or resource accessibility often led to ‘boom and bust’ cycles in overall population, seen in places as far afield as South America or ancient China [Chu and Lee, 1994, Goldberg et al., 2016].

Yet sufficient shocks to food production, state capacity, or environmental conditions often led to full-scale collapses in the population and wealth of complex societies. This occurred in the western Roman Empire beginning in the third century CE. As climate change, plague, and military weakness disrupted Mediterranean trade networks, estimated per capita wealth dropped substantially from three times subsistence to near the subsistence level. Elites attempted to maintain their relative status in the disruption by imposing serfdom on Commoners and devaluing citizenship with the *Constitutio Antoniniana* [Jongman, 2007]. Similar collapses took place in other ancient societies such as those of Easter Island [Roman et al., 2017] or the Khmer Empire at Angkor [Evans et al., 2023, Motesharrei et al., 2014]. Collapse was almost never complete or total. It instead manifested in transitions from high-equilibrium to low-equilibrium states across time in response to changes in either the Earth or Human systems [Evans et al., 2023, Jongman, 2007].

In light of this history, the world order originating in the post-1750 Industrial Revolution and particularly the post-1950 Green Revolution has been remarkably stable, as the consistent

and rapid growth in population and GDP per capita has shown no significant reversals [Motesharrei et al., 2016, United Nations Population Division, 2023]. Yet this may not last. The ongoing disruption to the Earth System is likely to lead to shortages of critical resources and disruptions of ecosystem services [Raudsepp-Hearne et al., 2013] as society's ecological footprint transgresses several 'ecological boundaries' at once [Steffen et al., 2015]. Unexpected weather events resulting from anthropogenic global warming will add additional strain to systems built on the assumption of climate stability [Masson-Delmotte et al., 2021, Pörtner et al., 2022]. These types of shocks have led to recent Earth subsystem models showing a high risk of collapse on a decadal timescale [Willcock et al., 2023], indicating that our present global society, with a high equilibrium powered by 200 EJ/yr of usable fossil fuel energy [Energy Institute, 2023], faces significant existential risk.

In response to these challenges, I and my advisors have developed the second generation of the Human and Nature Dynamics model, or HANDY2. HANDY2 is meant as a simple partially-coupled model of the Earth and Human systems that captures the main dynamics of world population and resource use both before and after the Industrial Revolution. It follows the seminal first version of HANDY [Motesharrei et al., 2014], which investigated the causes of simple pre-Industrial societal collapses, and showed how Elite wealth concentration can induce or worsen them. We have made numerous improvements and additions to HANDY2, which can now simulate multiple possibilities of collapse in the next two centuries by separately tracking three types of energy, incorporating technological advancements, and allowing for changing societal policies. The structure and functionality of HANDY2 has been enhanced via the use of Ensemble Kalman Inversion (EnKI) [Iglesias et al., 2013] to automatically determine the values of model parameters, improving accuracy and error estimates in our scenarios of the medium-term future

to 2200.

1.2 Outline of Dissertation

The following chapters of this dissertation will document the structure and development of HANDY2, show results we have obtained from the model, and discuss its implications for our knowledge of past and future resource collapses.

In Chapter 2, we present a brief overview of the original Human and Nature Dynamics model (HANDY1) and the results in [Motesharrei et al. \[2014\]](#) obtained from it, discussing its reimplementations in the form of a Julia notebook. We show that each scenario in HANDY1 can be replicated and is robust to changes in the code used to run the model.

Chapter 3 covers the development of HANDY2 Industrial and the results obtained from it. Its first sections document the many new additions to the model structure not present in HANDY1, including the distinction between Regenerating, Nonrenewable, and Renewable energy resources, as well as variables that track production and efficiency technologies, labor force allocation, and government policy favoring Renewables investment. We also explain how we derived real-world historical estimates of population, energy stocks, and energy flows, using data from the UN Population Division, HYDE, and the Energy Institute’s (formerly BP) Statistical Review of World Energy [[Energy Institute, 2023](#), [Klein Goldewijk et al., 2017](#), [United Nations Population Division, 2023](#)]. In the last section, we demonstrate in a ‘business as usual’ run of HANDY2 that continuing current global policies poses a great risk of society undergoing resource collapse within the next 100 to 200 years. But we show in several other runs how a combination of three key policy changes—reducing global birth rates, leveling economic inequality, and in-

creasing investment into Renewable energy—averts collapse and enables a stable and prosperous future for humanity.

Chapter 4 gives an overview of the Ensemble Kalman Inversion (EnKI) [Iglesias et al., 2013] technique we use to more accurately estimate parameters in HANDY2. In its basic form, EnKI attempts to find an unknown set of model parameters u from a known set of observations $y = G(u) + \eta$, model structure G , and a Gaussian error η with fixed covariance Γ . Unlike in standard Kalman inversion, this is done by searching for solutions within a space A generated by an ensemble of particles z_0^j , and minimizing the L_2 error of the output y . This is done with the aid of existing Julia libraries provided by CLiMA [Dunbar et al., 2022]. EnKI is particularly useful for models like HANDY1 or HANDY2 which have many free parameters, as it does not directly require the model inverse or derivative, and deals well with local minima that can make other forms of error minimization difficult.

Chapter 5 documents how we apply EnKI to the HANDY1 model with the aid of the CLiMA libraries. We demonstrate recovery of HANDY1’s free parameters from noisy pseudo-observations of a specified Egalitarian, Equitable, or Unequal scenario, in situations with soft landings to equilibrium and in full type-N (for Nature) or type-L (for Labor) collapses. We also document the choices made in applying EnKI to HANDY family models. For these experiments in particular, we minimize using a loss function based on the mean-square error of Nature, Wealth, and Commoner population over the length of the entire model run, add independent multiplicative noise η to create each pseudo-observation y^n , and constrain parameter priors to avoid unrealistic collapses that take place with unphysical parameter values.

Finally, Chapter 6 presents the results of EnKI parameter estimation for the HANDY2 model. Unlike with HANDY1, we perform loss minimization based on real-world observations

of population as well as Regenerating, Nonrenewable, and Renewable energy use, and optimize 22 model parameters at once. This allows us to expand on the findings of Chapter 3, creating scenarios that track the historical record far more closely than could be done manually. We elucidate how certain HANDY2 mechanisms and dynamics affect model outcomes by comparing mean estimates of parameter values in a super-ensemble of 50 EnKI runs, made possible by the stochastic nature of EnKI. Changing policy parameters in the super-ensemble enables us to test the robustness of HANDY2's projections of resource collapse by 2200, assuming certain boundary conditions. In particular, we confirm that the promotion of higher Renewables use is by far the most important policy mechanism that increases Commoner prosperity and prevents collapse. We also discuss the implications of HANDY2 for the hypothesis of [Pomeranz \[2000\]](#) that a 'land shortage' affecting phytomass availability was the main impetus behind the use of Nonrenewable coal in Britain and Europe, at the start of the Industrial Revolution.

Chapter 2: Reproduction of HANDY1

2.1 Model Structure

The original version of HANDY, HANDY1 [[Motesharrei et al., 2014](#)], is a minimal model that simulates 500 to 1000 years of pre-Industrial human activity with only four basic state variables. Within HANDY1, stocks of Wealth w are produced from Nature y , both represented in units of ecodollars, by Commoners x_C . The Wealth is then consumed by both Commoner and Elite populations x_E . Various scenarios presented in the HANDY paper explore the cycles of growth and collapse that can occur on these multi-century timescales, resulting from the depletion and replenishment of Nature. HANDY1 uses Regenerating biomass energy as a proxy for Wealth, consistent with the historical evidence for forests (in North America and Europe) and crop residues (in East and South Asia) being the main energy source for productive societal activities before the Industrial Revolution [[Pomeranz, 2000](#), [Smil, 2015](#), [2017](#)].

We begin by showing the governing equations for the state variables of HANDY1, which are as follows. ([Motesharrei et al. \[2014\]](#), eq. 3)

$$\dot{x}_C = \beta_C x_C - \alpha_C x_C \quad (2.1a)$$

$$\dot{x}_E = \beta_E x_E - \alpha_E x_E \quad (2.1b)$$

$$\dot{w} = \delta y x_C - C_C - C_E \quad (2.1c)$$

$$\dot{y} = \gamma y (\lambda - y) - \delta x_C y. \quad (2.1d)$$

Specifically, Commoners and Elites in HANDY1 increase in population at a per capita birth rate given by a constant β_C or β_E respectively. They die at a variable death rate α_C and α_E , with the overall population growth rate being the difference of the two. Extraction of Nature to produce wealth is given by the product of Commoner population x_C , Nature y , and a depletion constant δ , and decreases by the sum of Commoner and Elite consumption C_C and C_E . Finally, Nature regenerates to a maximum of λ with a logistic function that has a rate proportional to γ , both constant, and is drawn down by the yearly production of Wealth given above.

The non-constants $C_{C,E}$ and $\alpha_{C,E}$ in the above formula are given below, in forms equivalent to those of [Motesharrei et al. \[2014\]](#) (eqs. 4-6).

$$\omega = w / (\rho(x_C + \kappa x_E)) \quad (2.2a)$$

$$C_C = \min(1, \omega) s x_C \quad (2.2b)$$

$$C_E = \min(1, \omega) \kappa s x_E \quad (2.2c)$$

$$\alpha_C = \alpha_{min} + \max(0, 1 - \omega) (\alpha_{max} - \alpha_{min}) \quad (2.2d)$$

$$\alpha_E = \alpha_{min} + \max(0, 1 - \kappa \omega) (\alpha_{max} - \alpha_{min}). \quad (2.2e)$$

Commoner consumption C_C is given as the product of the per capita consumption rate ω (capped at 1), population x_C , and the subsistence salary s . Likewise, the commoner death rate α_C is also determined by ω and the constants of maximum death rate α_{max} and minimum α_{min} , reaching the minimum when $\omega = 1$. The formulae for Elites are similar with the addition of an inequality factor κ , as Elites will continue to consume a disproportionate amount of Wealth, κ times as much as the Commoners, even as the Commoners starve. Finally, ω itself is determined by the total wealth w divided by the minimum subsistence wealth ρ times the sum of Commoners and equivalent Elites.

The above formulae determine a consistent theoretical model with twelve free parameters and four state variables. While the original authors implemented it in the systems-modeling program Vensim, it is straightforward to translate to other programming languages including Python and Julia. The supplementary data files provide an example of a Jupyter notebook (`handy1_plot_code.ipynb`) which replicates the HANDY1 model and plots the results.

2.2 Replication of HANDY1 Results

By using the same parameters as found in [Motesharrei et al. \[2014\]](#), we replicate the scenarios and results in the original HANDY paper. Figure 2.1 shows what happens in an Egalitarian society (as in section 5.1 of [Motesharrei et al. \[2014\]](#)), with no Elites and with various levels of depletion δ . In the first scenario with the lowest level of depletion (2.1a), the Human system reaches a point of stability, or *soft landing*, within about 600 years from the scenario's start. As δ increases (2.1a, 2.1c), the model variables begin to oscillate around a steady state. At the highest levels of depletion (2.1d), a full collapse occurs as the stock of Nature is completely depleted.

All state variables, including Commoner and Elite population, eventually become zero.

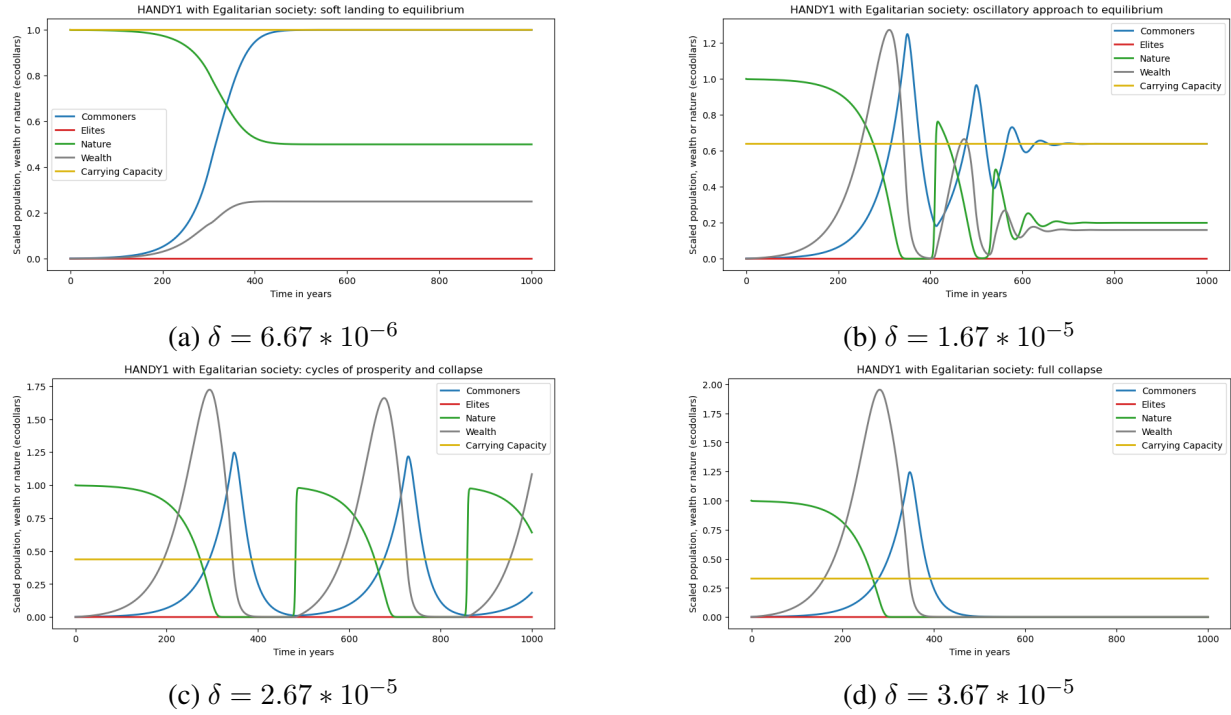


Figure 2.1: A graph of state variables from Egalitarian runs of HANDY1 with increasing depletion δ . The color scheme and units are as in [Motesharrei et al. \[2014\]](#), but vertical scale is different.

Similar dynamics occur in an Equitable society (section 5.2 of [Motesharrei et al. \[2014\]](#))—one where there is a small pool of non-workers that consume resources at a rate $\kappa = 1$ identical to that of the Commoners, but do not extract Nature. Figure 2.2 once again shows a progression from stable equilibrium (2.2a), to oscillatory collapse (2.2b, 2.2c), and then to full collapse (2.2d), with slightly higher values of depletion δ than in the Egalitarian case.

Conversely, the Human system becomes much more unstable in an Unequal society, where there are non-working Elites that consume $\kappa = 10$ or $\kappa = 100$ times more than the Commoners. Figure 2.3a replicates the type-L (for Labor) collapse scenario from section 5.3.1 of [Motesharrei et al. \[2014\]](#), where a small amount $x_E = 10^{-3}$ of Elites with very high levels of inequality are

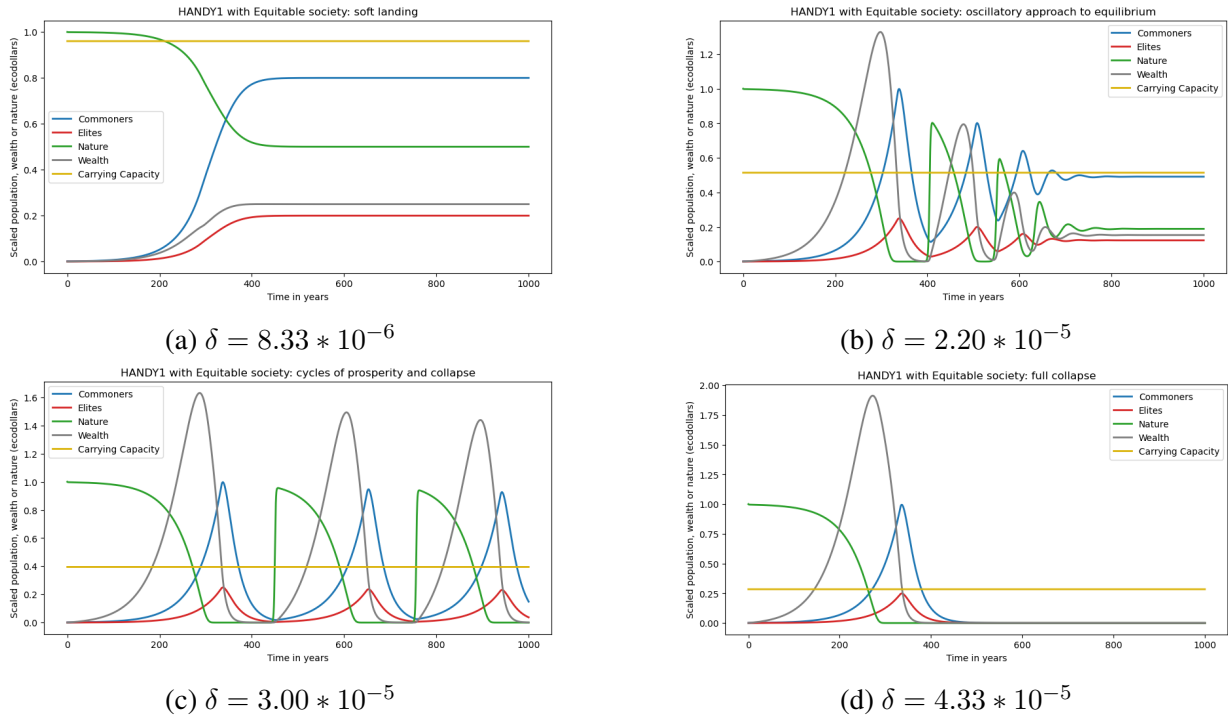


Figure 2.2: A graph of state variables from Equitable runs of HANDY1 with increasing depletion. The color scheme and units are as in [Motesharrei et al., 2014], but vertical scale is different.

present at the start of the run. Within 700 years, the Elites consume so much wealth that they prevent the Commoner population from being able to sustain itself, and both groups ultimately face total societal collapse by 1000 years from run start.

Figure 2.3b shows a type-N (for Nature) collapse scenario, caused both by overdepletion and inequality (section 5.3.2 of Motesharrei et al. [2014]). The much higher resource consumption of both Commoners and Elites ensures Nature is consumed much faster than it can regenerate, resulting in a swift collapse lasting only 500 years. Elites collapse after the Commoners as they can access much more of the shrinking stock of Wealth. With these runs, we verify the main result of HANDY1—that overconsumption and economic inequality can quickly cause the destabilization and downfall of a theoretical human society powered by Regenerating energy.

Finally, by reducing both inequality and depletion, and by lowering the Elite birth rate, a stable equilibrium is possible even in an Unequal society. This is shown in Figures 2.3c and 2.3d. The first scenario (5.3.3) with low δ shows a soft landing to equilibrium, while the second (5.3.4) shows oscillatory damping with lower overall carrying capacity. This shows how a few minor policy changes can entirely prevent collapse in HANDY1.

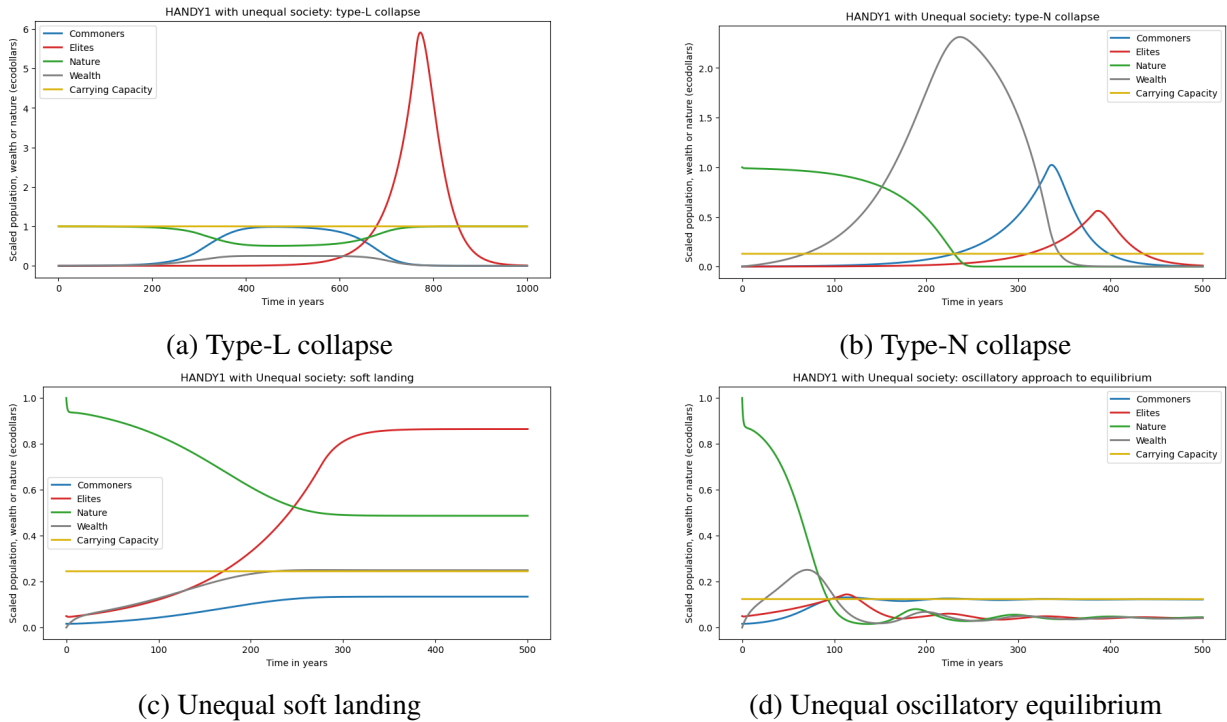


Figure 2.3: A graph of state variables from Unequal runs of HANDY1 with type-L and type-N collapses (top), as well as runs with lower inequality κ and stable or oscillatory approaches to equilibrium (bottom). The color scheme and units are as in [Motesharrei et al., 2014], but vertical scale is different.

The graphs produced in the Jupyter notebook nearly precisely match the Vensim version of HANDY1 originally shown in Motesharrei et al. [2014], with a few minor technical exceptions. Most importantly, the notebook uses Euler integration of the state variables with a time step of 0.2 years, which may have slightly different error properties. The notebook also uses double-

precision variables (Julia type `Float64`) compared to the single precision of Vensim, which in certain edge cases can cause a nearly infinitesimal amount of Nature to persist after a resource collapse and regenerate after a long period. This has no practical effect on the scenarios presented in the HANDY paper, as they each terminate after 500-1000 years. If warranted, code may be added to set the state variables of HANDY1 to zero if they cross below a certain threshold. We make this change in HANDY2 (see Chapter 3). Notwithstanding these minor discrepancies, we can independently confirm that the implemented results of HANDY1 reflect the theoretical model dynamics as presented in [Motesharrei et al. \[2014\]](#) for all scenarios in the paper, and that the model is robust to differences in programming and parameterization. We therefore use this version of HANDY1 as a foundation for implementing its more complex second generation—HANDY2.

Chapter 3: Structure and Results of HANDY2

3.1 Overview of HANDY2 Structure

The original formulation of HANDY1 was tailored to describe one cycle of growth and collapse in a pre-Industrial society. But this is not the world of today. Instead, the vast increases in the use of fossil fuels and water since 1800 have greatly transformed the human way of life, as these have become the primary drivers of productive human activity. Whereas the use of Regenerating biomass energy peaked at about 20 exajoules (EJ) per year before the introduction of Nonrenewables [Smil, 2017], the modern use of coal, oil, and gas is now about 200 EJ/yr and continues to grow rapidly [Energy Institute, 2023]. This trend has only accelerated since the Green Revolution, when improved crop varieties are now grown rapidly with the aid of artificial fertilizers—themselves a product of Nonrenewables-intensive mechanisms such as the Haber-Bosch process [Wang et al., 2021] and mining of phosphate rock throughout the globe [Geissler et al., 2018]. This has had the twin effects of greatly altering global element cycles [Galloway et al., 2004, Le Quéré et al., 2015, Steffen et al., 2015, Zeng et al., 2014] and increasing humanity's use of Regenerating energy to 50 EJ/yr [Smil, 2017], as it remains to this day an important resource input in large parts of the developing world [Smil, 2015, 2017].

These transformative changes can be summarized by the idea (Mote et al. [2020]; Figure 5) that we live in a 'Full World' that contrasts with the 'Empty World' captured in HANDY1.

No part of the Earth's environment remains unaffected by human activity. In addition to the changes discussed above to the global carbon and nitrogen cycles, Earth is now experiencing unprecedented decreases in the reserves of fossil water [Steffen et al., 2015], increases in the concentration of greenhouse gases such as carbon dioxide, methane, and chlorofluorocarbons [Masson-Delmotte et al., 2021, Mote et al., 2020], rapid species extinctions, urbanization and sprawl disrupting natural habitats, and a global climate crisis that is causing manifold disruptions to human and animal life [Hansen et al., 2013, Masson-Delmotte et al., 2021, Sodhi and Ehrlich, 2010, Steffen et al., 2015]. This has happened concurrent with an eightfold increase in both human population and GDP per capita since 1800 [Mote et al., 2020, Motesharrei et al., 2016]) reflecting a greatly improved standard of living for many people even in the face of all of these challenges. Humans are, more than ever before, an integral component of the Earth System whose choices and behaviors must be included in our model of today's environment.

For this reason, we have created the second generation of HANDY, HANDY2. HANDY2 is a model specifically designed to simulate population and energy dynamics over the 500-year period between 1700 and 2200, with a greater emphasis on Nonrenewable and Renewable resources, and an expanded suite of mechanisms relating to population, capital investment, inequality, and technology specifically tailored to the Full World. The model is designed to flexibly incorporate real observations of population and energy consumption in setting its parameters and starting conditions, allowing it to accurately track the effects of the Industrial and Green Revolutions, with the aid of automated EnKI parameter estimation. In several near-future scenarios, HANDY2 shows that current global policies are likely to cause a type-N resource collapse and a sharp decline in the world's population. Yet this can be averted by a combination of changes to decrease birth rates, lower economic inequality, and rapidly adopt Renewable sources of energy.

3.2 HANDY2 Key Concepts and Sectors

3.2.1 Types of Energy

HANDY2 Industrial introduces a number of key concepts, distinctions, and mechanisms not included in HANDY1. As HANDY is primarily an *energy* model, the most important of these is the difference between three types of energy—Regenerating, Nonrenewable, and Renewable Flows. Regenerating energy (subscript *R*) represents stocks of accessible energy that replenish themselves and can be depleted on human timescales. Most importantly, this includes both traditional biomass (e.g. crop residues, firewood, animal dung) [Smil, 2015, 2017], and industrial-scale ‘modern biomass’ energy (e.g. trash incineration, corn and cellulosic ethanol) [Energy Institute, 2023]. This energy type is included in HANDY1.

The other two types of energy are new. Nonrenewables (subscript *N*) include those stocks of energy that are not replenished on a human timescale. This is dominated by the burning of coal, oil, and gas, but also includes nuclear power in the absence of large-scale fuel reprocessing. In HANDY2, we use the sum of coal, oil, and gas energy outputs to represent the yearly Nonrenewables use. We also subtract these energy outputs from our estimate of the total stock of ‘burnable Nonrenewables’, following the numbers given in McGlade and Ekins [2015]’s paper on the amount of fossil fuels that must be left in the ground to prevent anthropogenic global warming from likely exceeding 2.5 to 3 degrees C. This figure tracks the RCP4.5 Representative Concentration Pathway identified in the IPCC’s 5th and 6th Assessment Reports [Masson-Delmotte et al., 2021, van Vuuren et al., 2011], and represents about 60 ZJ of fossil fuels. More details on the calculation are included below.

Finally, Renewable Flows (subscript F) are defined as those energy sources which do not have a stock that appreciably accumulates; i.e. no reasonable amount of human activity can deplete the instantaneously accessible amount. This most notably includes solar and wind power, which are rapidly expanding their contribution to the world's electricity generation. Both have an upper practical extraction limit that is several times greater than the current use of Nonrenewables [Li et al., 2018, Lu et al., 2009]. Renewable flows also include hydroelectricity, the use of which has plateaued since 2010 [Energy Institute, 2023]. In many areas, hydroelectric power plants have limited expansion potential due to their capital-intensive nature [Bogoviz et al., 2020] and due to habitat and biodiversity loss from the construction of large reservoirs [Palmeirim et al., 2014].

Regenerating stocks are often confused with Renewable flows in both the public imagination and policymaking, and promoted as a 'greener' alternative to Nonrenewables in the same way as Renewables are, which has led to many egregious policy failures with negative impacts on the environment [Raven et al., 2021]. High-profile examples of such failures include Northern Ireland's Renewable Heat Incentive scandal [Coghlin et al., 2020], and EU demand for wood pellets causing large-scale deforestation in North Carolina [Grunwald, 2021].

3.2.2 Extraction, Production, Efficiency, and Use Technologies

Other important concepts are those related to technological change and resource efficiency. It has long been confirmed by many studies that increases in resource use efficiency counter-intuitively *increase* total resource utilization. This is referred to as the Jevons Paradox [York and McGee, 2016], and appears in both Industrial and pre-Industrial human societies, and in

the HANDY2 model as well. We distinguish four types of technological and efficiency factors, which arise from different sources and have different effects on production. Only the first two, *extraction* and *production* technology, are explicitly included in HANDY2.

Extraction technology (ξ) represents the development of better methods to access more and different kinds of resources. Examples include hydraulic fracturing (fracking) of tight oil [Heinberg, 2013], artificial fertilizers used in the Green Revolution that encourage higher plant growth rates [Smil, 2017], and the buildout of large offshore wind farms from the 1990s onwards [Díaz and Guedes Soares, 2020]. This technology directly encourages greater depletion by making it less costly, and leads to a increase in wealth in real life as in HANDY2.

Production technology (ψ) instead represents the ability to gain more resources from an *existing* stock, or otherwise reduce production losses. Production technologies include those as simple as burning wood in stoves instead of open fires, raising energy efficiency from less than 5% to around 15% [Smil, 2015, 2017], or as complex as combined-cycle natural gas turbines that now displace coal in electricity grids and regularly reach thermal efficiencies of almost 65% [Energy Information Administration, General Electric, Kwon et al., 2019]. In HANDY2, increased production technology increases wealth per unit depletion, while indirectly stimulating overall resource depletion by encouraging increased labor force allocation to the resource.

Next are two types that are not currently in HANDY2: consumption technology (ζ) and use efficiency (η). Increased consumption technology entails a greater ability for a consumption sector to use more resources; an example is found in the increasing energy needs of larger houses and vehicles in the US [Covington, 2023, NBC News, 2022]. Conversely, increased use efficiency includes what are classically thought of as efficiency technologies meant to use fewer resources, like low-flow shower heads, LED lighting, fuel-efficient cars, or better-insulated water heaters.

While those efforts are laudable, use efficiency is subject to the Jevons paradox as it drives the adoption of consumption technology. This effect has already negated the effectiveness of the US's Corporate Average Fuel Economy (CAFE) standards [Whitefoot and Skerlos, 2012]. Because these two factors cancel each other out, we do not include them in HANDY2.

3.2.3 Labor Force and Investment

HANDY2 includes a mechanism that shifts the use and development of energy resources between the three types by means of tracking the labor force allocation and accumulated investment in each resource. At the start of the model run, each resource type has a starting investment A . The labor force allocation L of each resource type at each timestep is determined by its proportion of the overall investment; this allocation in turn affects the production rate. Then, the *production* $\pi_{R,N,F}$ of each resource type as a proportion of the total production π in turn determines the allocation of the investment at each time step, out of a total investment pool I whose size is governed by wealth ω . Said investment is then added to the accumulated total A , subject to a depreciation rate σ .

This mechanism creates a positive investment feedback loop that is a primary driver of change over time in HANDY2. If a particular resource has a production rate that is growing slowly, then it will receive more investment over time. This directly increases the labor force allocation as the investment accumulates, and indirectly increases the production and extraction efficiency factors, which in turn increases production. The process continues until the resource production reaches a new equilibrium with the depletion of stocks, in the case of Regenerating and Nonrenewable energy, or approaches the maximum F of accessible Renewable flows. As we

will see in section 3.3, the dynamics and timing of this positive feedback are strongly affected by investment-related parameters in HANDY2; small changes in the investment threshold A_{Nth} can have large effects on when (or even if) Nonrenewable use grows exponentially within the model.

Governments and policymakers have recently begun attempts to deliberately incentivize the use of Renewables, in the form of government subsidies for construction of Renewable infrastructure [Bistline et al., 2023], taxation of carbon emissions [Hájek et al., 2019, Metcalf, 2021] or even investment into basic research [Trancik, 2014]. This is represented in HANDY2 by means of an I_{fp} Renewables subsidy. After a certain starting year (usually 1990 CE), representing the recent advent of clean-energy policies, some amount of Nonrenewables investment I_N is artificially transferred to the Renewables investment I_R , keeping the total I unchanged. The rate of transfer depends on the difference in labor force allocation between Nonrenewables and Renewables. Once Renewable flows become established enough, the transfer stops. We thus obtain a semi-endogenous mechanism to trigger the growth in Renewable flows late in the history of the model run, which makes sense within the context and structure of HANDY and is in accordance with the historical record.

3.2.4 Births and Deaths

One of the biggest changes in population dynamics over the period covered by HANDY2 has been the decrease in birth rates across large portions of the world. This trend is one that began in the rich world in the early 20th century and has since spread to developing countries by the 21st century. Sweden's total fertility rate (TFR) stood at 4 children per woman in 1910 before dropping dramatically to the replacement level of 2 by 1930 [Dribe and Scalone, 2014].

The same decrease occurred in China by 1990, spurred by the One-Child Policy [Yang et al., 2022], while India has only recently reached a TFR of just above 2 children per women as of 2020 [Pearce, 2021]. Demographers at the UN Population Division [United Nations Population Division, 2023] thus project much lower fertility rates in the future even where they remain high currently, such as in sub-Saharan Africa.

However, demographic predictions have been notoriously poor for many years, based on unfounded assumptions that tend to underestimate overall population growth [Mote et al., 2020]. In addition, while declining world TFR has meant that the *relative* population growth has peaked at 2 percent per year by 1960, the *absolute* population growth remains relatively steady near its peak of around 80 million per year to the present day (United Nations Population Division [2023]; see Figure 3.1).

We thus state the question of predicting birth rates to be out of scope for HANDY2 and instead externally impose a reduction in birth rates based on a set of *birth rate policy* parameters. In our baseline model scenario, these parameters are set to cause the birth rates $\beta_{C,E}$, constant in HANDY1, to decline in such a way as to roughly match the historical estimates and projections of world population. Estimates are derived from HYDE [Klein Goldewijk et al., 2017] and the [United Nations Population Division, 2023], and projections to 2100 are sourced from the UN Population Division. In scenarios where we impose a greater population reduction, the policy parameters are changed further to cause the world population to peak below the UN projection and slowly shrink afterwards. They are also changed in scenarios where the estimated peak in world population is higher and later. We consider this a good compromise to avoid excessively increasing model complexity, while still taking into account population factors important for projecting the medium-term future.




Figure 3.1: Relative (dotted) and absolute (solid) population growth per year from 1700 to 2022, derived from population estimates from HYDE [Klein Goldewijk et al., 2017] and the UN Population Division [United Nations Population Division, 2023]. Arrows mark the turning points of the first Industrial Revolution in 1750 and the Green Revolution in 1950.

The death rates within the HANDY2 population model are endogenous. As in HANDY1, death rates are determined by the Commoner wealth ω . However, the minimum death rate at high ω is now imposed gradually, to eliminate the discontinuities present in HANDY1. This is discussed further in the next section.

3.2.5 Class Mobility and Inequality

The dynamics of class inequality between Elites and Commoners take on a new and more important role in HANDY2. In HANDY1, Elites do not perform any production or generate wealth, but merely consume resources at a constant rate κ greater than that of the Commoners. HANDY2 introduces the concept of *class mobility*. For most of world history, while Elites controlled a vastly greater share of a society's resources than their small numbers would indicate, membership in the Elite class was always contingent. Losing a war could decimate the land holdings, Regenerating resources, and thus wealth any given Elite could extract [Gat, 2008]. High birth rates meant wealth either had to be shared between new Elites in the case of partible inheritance, reducing their economic power over the course of several generations, or transferred to one successor in the case of primogeniture, creating a profusion of landless younger sons who were at great risk of falling into Commoner status [Christelow, 1996, Kokkonen and Sundell, 2014].

After the Industrial Revolution, a similar process has continued to take place to this day. While wealth is now largely derived from capital, itself largely a product of Nonrenewable energy inputs, many people born into wealth fail to maintain it, either through a failure to inherit enough to compete with others, poor business investments, disputes between Elites, or any number of other reasons [Harrington, 2017, Hughes, 2004].

The reverse move—from Commoner to Elite—is far less common. Before the Industrial Revolution, when the majority of Commoners were subsistence farmers, it was nearly impossible. A small plot of land simply could not produce enough agricultural wealth, even in a good year, to directly attain Elite status. Nor could that farmer accumulate wealth as they had little secure access to credit or banking [Erdkamp, 2005]. Today, becoming an Elite is merely rare. About 70% of billionaires on the Forbes 400 list were born in upper-middle class or (in rare cases) lower class conditions [Kaplan and Rauh, 2013], leaving only 30% born into Elite status. Yet this represents only a small handful of people. For the vast majority of the world population, there is no reasonable chance of rising from Commoner status as the gap between the wealth of Elites and the rest of society continues to grow at a rapid pace in much of the developed world. This is particularly true in the United States [Mitchell, 2020, Saez and Zucman, 2016].

In HANDY2, we thus model class mobility as a restricted two-way process. Below a threshold of wealth ω_μ , a fraction of Elites become Commoners each year at a rate governed by the parameter μ_M and the wealth of the Commoners. Above the critical value of ω_μ , a (small) fraction of Commoners become Elites at a rate given by $1/\kappa$ times the downward mobility. This represents the greater opportunities for Commoner advancement present in a richer society, but it may be considerably too generous in practice—even low amounts of upward mobility can cause sharp Elite growth due to the Elites' much smaller population.

Of note also is that the downward elite mobility rate μ has a very strong influence on the distribution of a society's resource consumption between classes, Commoner population growth, and overall HANDY2 model results and predictions, which we will discuss in section 3.5. This sensitivity makes μ_M a critical parameter to estimate automatically via EnKI, which is covered at length in Chapter 6.

Finally, there is also a mechanism to impose reductions in the inequality κ . We use this to simulate scenarios in which policies are enacted to reduce wealth accumulation by the Elites, thereby distributing a larger share of the resource pie to Commoners and increasing their relative wealth. This is shown in Scenarios 4-6 of HANDY2 discussed in section 3.5; we leave it to policymakers to determine the specific details of this inequality reduction.

3.2.6 Burnable Carbon, Reserves, and Resources

Nonrenewables are by far the most important resource in HANDY2, but determining their actual stock size is very difficult. Ongoing fossil fuel exploration continues to find new stocks of coal, oil, and gas whose sizes may be poorly or not at all known, and existing stocks that were previously too unprofitable to extract may become economically useful with higher energy prices or better production technology. For this reason, it is common to categorize fossil fuel stocks into *reserves* and *resources*.

Reserves, more fully *proved reserves*, “are generally taken to be those quantities [of Nonrenewables] that geological and engineering information indicates with reasonable certainty can be recovered in the future from known deposits under existing economic and operating conditions” [Energy Institute, 2023]. Reserves have a generally well-defined amount and may increase in quantity over time due to exploration of new deposits or changes in economic conditions. Conversely, *resources* are defined as “The quantity of oil, gas, or coal remaining that is recoverable over all time with both current and future technology, irrespective of current economic conditions” [McGlade and Ekins, 2015]. This is a poorly defined quantity, and not reported in such data sources as the Energy Institute (formerly maintained by BP)’s *Statistical Review of World*

Energy [Energy Institute, 2023]. The size of resources may even decrease with better knowledge of fossil fuel distribution or different assumptions about hypothetical future technology. We thus avoid using this estimate in HANDY2 entirely.

Instead, we focus on a third quantity, that of *burnable carbon* as defined by McGlade and Ekins [2015]. This is the quantity of fossil fuels that can be used by humans without exceeding some climate-related threshold, most commonly a certain global temperature increase above the pre-Industrial average. Burnable carbon estimates are necessarily an approximation; assumptions about production efficiency, climate forcing sensitivity, or future technologies such as carbon capture and storage (CCS) may greatly affect this value. In their paper, McGlade and Ekins [2015] assume a maximum emissions estimate in line with the Paris conference target of 2 degrees C above pre-Industrial temperatures. However, we think this value is too small, as global Nonrenewables use and fossil fuel emissions continue to rise steadily [Energy Institute, 2023]. With current levels of climate action, emissions are likely to approach the RCP4.5 pathway defined in the IPCC 5th Assessment Report [van Vuuren et al., 2011], leading to a temperature rise by 2100 of 2.5 to 3 degrees C above pre-Industrial [Masson-Delmotte et al., 2021]. Higher values are less likely, as coal and oil resources may be overestimated and unable to support continued expansion of Nonrenewables output until the end of the century [Höök et al., 2010, Laherrère, 2019].

The actual calculation of burnable carbon amounts are straightforward, converting the estimates of carbon emissions into the atmosphere from any given RCP [van Vuuren et al., 2011] into a gigatons of carbon (Gt C) equivalent, and then to energy equivalents (in EJ) following industry-standard conversion factors [Engineering Toolbox]. The result we obtain is close to the current estimated value of fossil fuel reserves (see section 3.4 for details). We thus have more confidence in using this shared value to set the starting stock of Nonrenewables y_N in HANDY2.

3.3 HANDY2 Equations and Parameters

3.3.1 Variables and Parameters

HANDY2 has many more variables and parameters that can be changed between model runs than HANDY1. They can be roughly divided into four broad categories.

The first are state variables, representing the persistent quantities of the model that receive increments at each timestep. There are eight in HANDY2, four of which are inherited from HANDY1: the Commoner and Elite populations x_C , x_E , the accumulated wealth w , and the Regenerating stock y_R . Four more are new: the Nonrenewables stock y_N , and accumulated resource investment for the three resource types A_R , A_N , and A_F . A special case is the Renewable flows parameter F , which is a *pseudo-stock*, governing an important sector of the model's structure as a production limit on one of three energy types, but has a dimension of power rather than energy. As such, we do not call it y_F .

The second group consists of eleven parameters estimated at least partially from the empirical record and existing literature; the estimates are described more fully in the following section 3.4. Four are the demographic parameters for birth and death rates α_{min} , α_{max} , β_C , and β_E , also used in HANDY1. One is the starting inequality κ_{start} , and three are natural resource-related: nature regeneration γ , nature capacity λ_{start} , and maximum accessible Renewable flows F_{base} . Finally, each resource type has a production technology factor (ψ_R, ψ_N, ψ_F) that couples depletion to production. A special consideration is that these technology parameters are functions that can vary with time. ψ_R in particular has ancillary variables $\psi_{R,start}$, $\psi_{R,end}$, $\psi_{R,offset}$, and $\psi_{R,scale}$ that govern the shape of the logistic function.

Seventeen more parameters can be tuned as part of the model validation, and have no existing external constraints. These include the depletion δ_R , δ_N , and δ_F for each resource type, the Regenerating capacity growth rate λ_{growth} , and two demographic parameters: s , the subsistence salary, and ρ , the famine wealth threshold. There is also the class mobility parameter ω_μ that determines the wealth at the onset of mobility, and the maximum mobility rate μ_M . Nine more parameters govern the economic model of HANDY2: the thresholds for cumulative accessibility investment A_{Rth} , A_{Nth} , A_{Fth} , minimum extraction efficiencies $\xi_{m,R}$, $\xi_{m,N}$, $\xi_{m,F}$, the wealth at investment onset ω_i , investment rate J , and investment depreciation rate σ . All except A_{Rth} can be tuned with EnKI parameter estimation, as covered in Chapter 6.

Finally, five parameters are *policy knobs* whose values can change in different scenarios to reflect deliberate societal choices. The birth rate policy β_{final} and inequality reduction κ_{final} represent policies that may be enacted by a society to lower birth rates $\beta_{C,E}$ or reduce inequality κ . To govern the shape of the logistic curves of both policies, there exist ancillary parameters β_{offset} , β_{scale} , κ_{offset} , and κ_{scale} . Next, there are the Renewable flows subsidy parameters $I_{Fp,start}$ and $I_{Fp,base}$. Finally, there is F_{policy} , a function that determines the increase over time in the maximum amount of accessible Renewable flows.

Table 3.1 summarizes each of the parameters listed above.

Table 3.1: List of variables and parameters used in HANDY2, with units.

Variable	Name in code	Short description
State variables		
x_C	xc, xcts	Commoner population (millions)
x_E	xe, xets	Elite population (millions)
w	w, wts	Accumulated wealth (EJ)
y_R	yr, yrts	Regenerating resource stock (EJ)

y_N	yn, ynts	Nonrenewable resource stock (EJ)
A_R	Ar, Arts	Regenerating accumulated investment (EJ)
A_N	An, Ants	Nonrenewable accumulated investment (EJ)
A_F	Af, Afts	Renewable accumulated investment (EJ)
<hr/>		
Empirically-constrained parameters		
α_{min}	alpha_min	Minimum death rate (1/yr)
α_{max}	alpha_max	Maximum death rate (1/yr)
β_C	beta_c	Starting Commoner birth rate (1/yr)
β_E	beta_e	Starting Elite birth rate (1/yr)
κ_{start}	kappa_start	Starting inequality
γ	gamma	Nature regeneration rate (1/yr)
λ_{start}	lambda_start	Nature regeneration capacity (EJ)
F_{start}	F_base	Maximum accessible Renewables at run start (EJ/yr)
$\psi_{R,start}$	psi_r_start	Starting Regenerating production efficiency
$\psi_{R,end}$	psi_r_end	Ending Regenerating production efficiency
$\psi_{R,offset}$	psi_r_offset	Offset of Regen prod. efficiency increase (yr CE)
$\psi_{R,scale}$	psi_r_scale	Timescale of Regen prod. efficiency increase (yr)
ψ_N	psi_n	Nonrenewables production efficiency
ψ_F	psi_f	Renewables production efficiency
<hr/>		
Other tunable parameters		
δ_R	delta_r	Regenerating depletion rate (1/(10 ⁶ persons-yr))
δ_N	delta_n	Nonrenewables depletion rate (1/(10 ⁶ persons-yr))
δ_F	delta_f	Renewables depletion rate (1/(10 ⁶ persons))
λ_{growth}	lambda_growth	Nature regeneration capacity growth rate (1/yr)
s	s	Subsistence salary (TJ/person-yr)
ρ	rho	Famine wealth threshold (TJ/person)
ω_μ	omega_mu	Class mobility threshold consumption (TJ/person-yr)
μ_M	mu_M	Maximum downward class mobility (1/yr)
A_{Rth}	Arth	Regenerating investment threshold (EJ)
A_{Nth}	Anth	Nonrenewable investment threshold (EJ)
A_{Fth}	Afth	Renewable investment threshold (EJ)
$\xi_{m,R}$	xim_r	Regenerating minimum extraction efficiency
$\xi_{m,N}$	xim_n	Nonrenewable minimum extraction efficiency
$\xi_{m,F}$	xim_f	Renewable minimum extraction efficiency
ω_i	omega_i	Consumption at investment onset (TJ/person-yr)
J	J	Investment scaling factor
σ	sigma	Investment depreciation rate (1/yr)
<hr/>		
Policy knobs and ancillary parameters		
β_{final}	final_br	Final birth rate (1/yr)
β_{offset}	bra_offset	Time offset of birth rate reduction (years CE)

β_{scale}	bra_scale	Timescale of birth rate reduction (yr)
κ_{final}	kappa_final	Final inequality
κ_{offset}	kappa_offset	Time offset of inequality reduction (years CE)
κ_{scale}	kappa_scale	Timescale of inequality reduction (yr)
$I_{Fp,start}$	ifpstartyear	Starting year of Renewables subsidy (years CE)
$I_{Fp,base}$	ifp_base	Base rate of Renewables subsidy
F_{policy}	F_policy	Renewables growth policy factor

Model settings

t	fyear	Model current time (years CE)
t_{start}	startyear	Model start year (CE)
	nyear	Model run time (yr)
	timestep	Model timestep (yr)

Internal model quantities

π_R	prodr, prodrts	Regenerating production (EJ/yr)
π_F	prodn, prodnts	Nonrenewables production (EJ/yr)
π_F	prodf, prodfts	Renewables production (EJ/yr)
ξ_R	xir	Regenerating extraction efficiency
ξ_N	xin	Nonrenewable extraction efficiency
ξ_F	xif	Renewable extraction efficiency
ψ_R	psi_r	Regenerating production efficiency
L_R	Lr	Regenerating labor force allocation
L_N	Ln	Nonrenewable labor force allocation
L_F	Lf	Renewable labor force allocation
I	I	Total investment (EJ/yr)
I_r	I_R	Regenerating investment (EJ/yr)
I_n	I_N	Nonrenewable investment (EJ/yr)
I_f	I_F	Renewable investment (EJ/yr)
I_{fp}	ifp, ifpts	Renewables subsidy
ω	omega	Commoner consumption per capita (TJ/person-yr)
C_C	Cc	Commoner resource consumption (EJ/yr)
C_E	Ce	Elite resource consumption (EJ/yr)
α_C	alpha_c	Commoner death rate (1/yr)
α_E	alpha_e	Elite death rate (1/yr)
β_{RC}	bra_c	Commoner birth rate reduction
β_{RE}	bra_e	Elite birth rate reduction
μ	mu_downward, mu_upward	Class mobility rate (1/yr)
κ	kappa	Current inequality
λ	lambda	Nature regeneration capacity (EJ)
F	F	Maximum accessible Renewables (EJ/yr)

3.3.2 Equations

The formulae governing the HANDY2 state variables are as follows.

$$\dot{x}_C = \beta_C \beta_{RC} x_C - \alpha_C x_C + \mu x_E \quad (3.1a)$$

$$\dot{x}_E = \beta_E \beta_{RE} x_E - \alpha_E x_E - \mu x_E \quad (3.1b)$$

$$\dot{y}_R = \gamma y_R \left(1 - \frac{y_R}{\lambda}\right) - \frac{\pi_R}{\psi_R} \quad (3.1c)$$

$$\dot{y}_N = -\frac{\pi_N}{\psi_N} \quad (3.1d)$$

$$\dot{w} = \pi_R + \pi_N + \pi_F - C_C - C_E - I_R - I_N - I_F \quad (3.1e)$$

$$\dot{A}_{R,N,F} = I_{R,N,F} - \sigma A_{R,N,F}. \quad (3.1f)$$

These equations are broadly similar to HANDY1, with the new additions described previously. Commoner and Elite populations are now affected by the birth rate policy parameters β_{RC} , β_{RE} , and class mobility μ . Regenerating energy y_R replenishes at the rate γ as before (though γ has slightly different units) and is reduced at the rate of raw Regenerating production. Non-renewables y_N are similarly depleted but do not regenerate. Wealth w increases as the sum of all production quantities, and is reduced by both consumption and the yearly investment I . Finally, the accumulated investment rates $A_{R,N,F}$ grow with yearly investment and decrease with depreciation σ .

The final (as opposed to raw) production rates π are:

$$\pi_R = \delta_R L_R x_C \psi_R \xi_R y_R \quad (3.2a)$$

$$\pi_N = \delta_N L_N x_C \psi_N \xi_N y_N \quad (3.2b)$$

$$\pi_F = \tanh(\delta_F L_F x_C) \psi_F \xi_F F. \quad (3.2c)$$

As in HANDY1, these are the product of the depletion rate δ , the commoner population x_C , and the resource stock/pseudo-stock y_R , y_N , or F , but now include the efficiency technology parameters ψ and ξ as well as the labor force allocation L . Note the presence of the \tanh term in the Renewables production—this serves to cap production at the maximum extraction rate of flows given by F .

We now examine the investment sector. The total investment I is determined by:

$$I = J x_C \max(0, \omega - \omega_i). \quad (3.3)$$

It is proportional to the population x_C , the base investment rate J , and the per capita commoner consumption ω over some small threshold ω_i . That is usually set as $1/\kappa$ as it is the Elites who drive most investment into new production with their control of capital—and they will continue doing so even as Commoners starve.

Investment rates are divided amongst the resource types in proportion to their production π :

$$I_R = \frac{\pi_R}{\pi_R + \pi_N + \pi_F} I \quad (3.4a)$$

$$I_N = \frac{\pi_N - \pi_N I_{Fp}}{\pi_R + \pi_N + \pi_F} I \quad (3.4b)$$

$$I_F = \frac{\pi_F + \pi_N I_{Fp}}{\pi_R + \pi_N + \pi_F} I. \quad (3.4c)$$

As mentioned previously, a proportion of Nonrenewables investment equal to I_{Fp} of its total is given to Renewables to represent government policy promoting Renewable flows. I_{Fp} begins at zero, and its yearly change is determined as:

$$\dot{I}_{Fp} = I_{Fp,base}(L_N - 2L_F). \quad (3.5)$$

After the start year given by the parameter $I_{fp,start}$, ordinarily 1990, the Renewables subsidy grows for each year that the Nonrenewables labor force is more than twice the Renewables labor force. Once those proportions reverse, the subsidy shrinks again to zero. I_{Fp} is not allowed to become negative.

Similarly, the labor force allocation L is divided amongst the resource types in proportion to their total investment $A_{R,N,F}$. The sum of allocations is equal to one.

$$L_R = \frac{A_R}{A_R + A_N + A_F} \quad (3.6a)$$

$$L_N = \frac{A_N}{A_R + A_N + A_F} \quad (3.6b)$$

$$L_F = \frac{A_F}{A_R + A_N + A_F}. \quad (3.6c)$$

The extraction technology functions are coupled directly to the investment as well. These have a complicated form but are very similar to each other: ξ_R , ξ_N , and ξ_F are determined by a variant of a logistic function with a fixed cap at 1 and a starting value at zero investment of ξ_m . The slope is non-zero at zero investment, which represents the fact that even a small investment with no technology base can increase production a fair amount. Figure 3.2 shows the shape of the general ξ function, normalized to different A_{th} values.

$$\xi_r = 1 / \left(1 + \left(\frac{1}{\xi_{m,R}} - 1 \right) \exp \left(- \frac{A_R}{A_{Rth}} \right) \right) \quad (3.7a)$$

$$\xi_n = 1 / \left(1 + \left(\frac{1}{\xi_{m,N}} - 1 \right) \exp \left(- \frac{A_N}{A_{Nth}} \right) \right) \quad (3.7b)$$

$$\xi_f = 1 / \left(1 + \left(\frac{1}{\xi_{m,F}} - 1 \right) \exp \left(- \frac{A_F}{A_{Fth}} \right) \right). \quad (3.7c)$$

But the production technology functions determining ψ_R , ψ_N , and ψ_F in HANDY2 differ from one another. ψ_F is simply a constant. Conversely, ψ_R is a logistic function similar in form to those determining birth rates $\beta_{C,E}$ and inequality κ . The starting value is asymptotically equal to $\psi_{R,start}$, the end value approaches $\psi_{R,end}$, while $\psi_{R,offset}$ determines the midpoint year and $\psi_{R,scale}$ the slope of the change in Regenerating production technology.

$$\psi_R = \psi_{R,end} - (\psi_{R,end} - \psi_{R,start}) / \left(1 + \exp \left((t - \psi_{R,offset}) / \psi_{R,scale} \right) \right). \quad (3.8)$$

The formula for ψ_N is taken from the estimates for yearly thermal power efficiency provided by the Statistical Review of World Energy [Energy Institute, 2023]. It is piecewise linear.




Figure 3.2: Extraction efficiency ξ as a function of normalized investment $\frac{A}{A_{th}}$. The efficiency approaches a maximum of 1 as investment becomes large.

$$\psi_N = \begin{cases} 0.36 & \text{if } t \leq 2000 \\ 0.36 + \frac{0.04}{17(t-2000)} & \text{if } 2000 < t \leq 2017 \\ 0.40 + \frac{0.05}{33(t-2017)} & \text{if } 2017 < t \leq 2050 \\ 0.45 & \text{if } t > 2050 \end{cases} \quad (3.9)$$

Next, we describe the population sector. Consumption $C_{C,E}$ and wealth ω are the same as in HANDY1, with the major exception that there is no cap on consumption per capita—in line with what is observed in 20th century data [[Our World in Data, World Bank, 2023](#)].

$$\omega = w / (\rho(x_C + \kappa x_E)) \quad (3.10a)$$

$$C_C = \omega s x_C \quad (3.10b)$$

$$C_E = \omega \kappa s x_E. \quad (3.10c)$$

But the formula for death rates α_{min} and α_{max} is substantially different. The death rate no longer decreases as a piecewise function of ω , reaching a minimum for Commoners at $\omega = 1$. Instead, we use a hyperbolic tangent-based function as with many other new portions of HANDY2. Death rates decrease steeply when the consumption is small, but more gradually when it gets large, tending toward an asymptote of α_{min} in the limiting case of a rich society ($\omega \rightarrow \infty$).

$$\alpha_C = \alpha_{max} - (\alpha_{max} - \alpha_{min}) \tanh(\omega) \quad (3.11a)$$

$$\alpha_E = \alpha_{max} - (\alpha_{max} - \alpha_{min}) \tanh(\omega\kappa). \quad (3.11b)$$

As discussed in the previous section, the birth rate policy β_{RC} is a parameter used in certain scenarios to externally impose the decreases in fertility observed throughout the world and projected by the [United Nations Population Division \[2023\]](#). The ancillary parameters govern the start time of the β_{RC} logistic function (β_{offset}), its duration (β_{scale}), and its end value (β_{final}). In the Julia code it can also be turned off by the parameter `bra_switch`.

$$\beta_{RC} = (\beta_{final}/\beta_C) + ((\beta_C - \beta_{final})/\beta_C) * 1 / \left(1 + \exp((t - \beta_{offset})/\beta_{timescale})\right) \quad (3.12a)$$

$$\beta_{RE} = (\beta_{final}/\beta_E) + ((\beta_E - \beta_{final})/\beta_E) * 1 / \left(1 + \exp((t - \beta_{offset})/\beta_{timescale})\right). \quad (3.12b)$$

Several equations govern the inequality sector of HANDY2. Inequality κ is regulated by the policy parameter κ_{final} ; the decrease from its starting value κ_{start} is determined by a log-logistic function that has similar ancillary parameters, κ_{scale} , κ_{offset} to the birth rate policy function. It takes on a slightly different form to β_{RC} to account for the fact that κ 's effects scale logarithmically; hence $\log(\kappa)$ is itself logistic. The inequality reduction may be disabled with the code parameter `kappa_switch`.

$$\kappa = \kappa_{start} * \exp\left(\log\left(\frac{\kappa_{final}}{\kappa_{start}}\right) / \left(1 + \exp((-t + \kappa_{offset})/\kappa_{scale})\right)\right). \quad (3.13)$$

Next, we have a pair of piecewise linear equations that govern the class mobility. The rationale behind them is extensively discussed in the previous section. Downward mobility of elites happens at a rate μ dependent on normalized wealth $1 - \frac{\omega}{\omega_\mu}$. If that quantity would be negative, upward mobility occurs instead but with a rate $1/\kappa$ as large.

$$\mu = \begin{cases} \mu_M * \left(1 - \frac{\omega}{\omega_\mu}\right) & \text{if } \frac{\omega}{\omega_\mu} \leq 1 \\ \frac{\mu_M}{\kappa} * \left(1 - \frac{\omega}{\omega_\mu}\right) & \text{if } \frac{\omega}{\omega_\mu} > 1. \end{cases} \quad (3.14)$$

Finally, two special equations change the maximum size of Regenerating stocks and Renewable flows over time. The maximum stock size of Nature, λ , grows exponentially with a (usually small) rate constant of λ_{growth} .

$$\lambda = \lambda_{start} * \exp\left((t - t_{start}) * \lambda_{growth}\right) \quad (3.15)$$

Conversely, the maximum accessible Renewable flows can grow over time at a rate set by F_{policy} , starting from a value of F_{base} . The exact formula is somewhat complex but ensures a smooth transition from a constant starting value to linear growth around the year 2000 CE, when real-world governments started heavy investment into Renewable flows.

$$F = F_{base} \left(1 + F_{policy} * \left(\frac{t - 2060}{80} + \sqrt{\frac{1}{8} + \frac{t - 2060^2}{80}}\right)\right). \quad (3.16)$$

In practical implementations of HANDY2, there are also a number of other quantities that are calculated to aid the creation of graphs and other forms of data visualization, or as practical

conveniences to determine other model variables. Some of these are returned as outputs in our Julia implementation of the model. They are not however part of its theoretical core and not necessary for evaluating HANDY2, so will not be listed in this section. The reader may refer to our reference Julia implementation (in the supplementary data files) to see them.

3.4 HANDY2 Data and Estimates

As with any model, HANDY2 is only as capable of making projections as the data that determines its inputs. Unfortunately, parameters in several sectors of HANDY2, including the labor force, investment, and inequality mechanisms, do not have reliable data for the entire model time period of 1700 to 2020, or are difficult to adequately measure (e.g. who counts as an Elite?). Nevertheless, there do exist reliable records of energy stocks and flows, as well as world population. We use these records to set the starting conditions and parameters in HANDY2. Resource and population figures at later times also allow us to constrain the unobserved parameters and keep the model scenarios in line with the historical record—particularly with the aid of EnKI parameter estimation, as we will cover in Chapter 6.

3.4.1 Estimates of Regenerating and Nonrenewable Stocks, and Renewable Flows

The three varieties of energy in HANDY2 have reliable historical data with coverage in most of the model prediction domain from 1800 to 2022, drawn mostly from the two sources of the [Energy Institute \[2023\]](#) and [Smil \[2017\]](#). Figure 3.3 shows the estimates of gross energy use we have obtained, broken down by energy type.

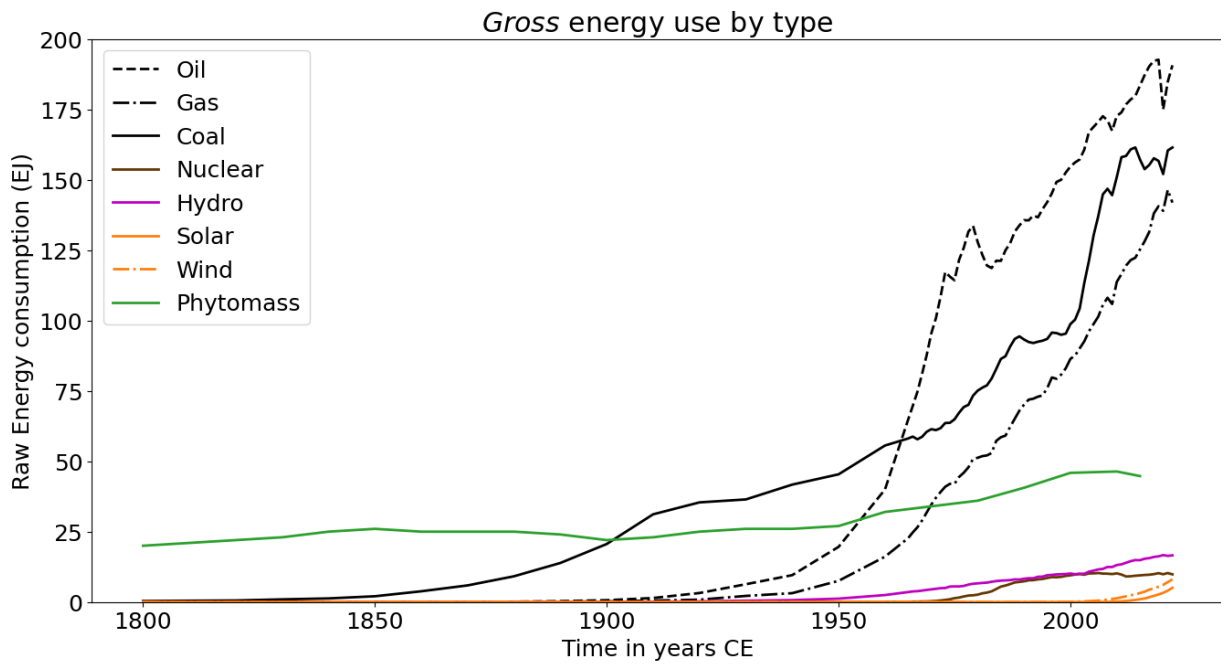


Figure 3.3: Plot of gross energy use by resource type from 1800 to 2022. The green line represents Regenerating energy (mostly biomass). Black lines show Nonrenewable fossil fuels, brown nuclear power, purple hydropower, while orange lines show Renewable solar and wind.

For Regenerating energy, yearly production can be treated as almost entirely composed of woody phytomass. Other types of Regenerating energy, such as animals being used for food or labor, are negligible over this time period. Modern biomass energy and biofuels also have negligible production before 2000 (below 1 EJ/year), and even afterwards account for at most 5 EJ/yr by 2020 [Energy Institute, 2023]. We thus use the figures in Smil [2017]’s appendix to estimate traditional biomass, adding them to the estimates of modern biomass from the Energy Institute [2023] to determine Regenerating production. These values, given every 10 years and in 2015, start at approximately 20 EJ/yr in 1800 and grow to 45 EJ/yr after the Green Revolution by 2000. The table in Smil [2017] stops in 2015; we assume consumption remains constant after that time. To this day, there are many citizens of less developed nations that use traditional biomass for energy in their daily life [Smil, 2015, 2017].

Slightly more difficult is the task of estimating the maximum regenerating capacity λ and the starting amount of Regenerating energy y_R . For this, we turn to Smil [2015], whose Table 12.1 estimates that the global woody phytomass stock of 0 CE was about 1000 gigatons of carbon (Gt C), decreasing to 600 Gt C by 1900 CE. We use a value slightly lower than the average of these two figures, 750 Gt C, as a starting point for λ . Per Smil’s estimates, trees generally contain about 50% carbon and on average have a dry energy density of 18 MJ/kg when burned [Smil, 2015], which equates to a phytomass stock in energy terms of 27 zettajoules (ZJ). We thus set the starting λ and y_R to this value.

Next, we look at the use of Nonrenewables. Energy Institute [2023] has excellent breakdowns of energy production and consumption by resource type from 1965 onwards. We simply add up their figures for gross coal, oil, and gas production to arrive at the total yearly production of fossil fuels, which we use as the HANDY2 Nonrenewables production. As of 2022, this is

494.05 EJ/yr. Before 1965, [Smil \[2017\]](#) has figures for coal, oil, and gas production every 10 years from 1800 to 2015. These are given on an net-energy basis, so we divide by the production efficiency $\psi_N = 0.36$ to obtain the gross production. When their data overlaps, from 1970 to 2015, the two sources are in close agreement to within 1%. Nuclear energy, while also Nonrenewable, is excluded from the total because it is neither a growing source of energy production today, nor subject to the same production technology ψ_N constraints or fossil fuel burnable stock concerns as coal, oil, or gas. We split it out as a separate entry in our supplementary energy data file.

Calculating the Nonrenewables stock size y_N following the burnable carbon methodology takes a bit more creativity, as this value is not directly listed in our sources. Beginning with the RCP4.5 pathway [[van Vuuren et al., 2011](#)] we use to estimate future carbon use (see discussion in section 3.2), [Rogelj et al. \[2014\]](#) calculate that it accounts for the use of 720 Pg = Gt C from 2015 to 2100. The conversion between carbon content and energy output of a fossil fuel varies by type, with natural gas being less carbon-intensive than coal and oil. But given an equal mix of the three fossil fuels, the average energy output is about 56.5 MJ/kg C [[Engineering Toolbox](#)], resulting in a total energy production over this time period of 40.7 ZJ. We add this to the 19.6 ZJ burned until 2015, most of it since 1965, [[Energy Institute, 2023](#), [Smil, 2017](#)] to obtain a total burnable Nonrenewables stock of 60.3 ZJ. Depending on the fossil fuel mix or climate sensitivity to carbon forcing, this value can have a margin of error of 20%; we thus set the value of y_N in HANDY2 to a slightly higher value of 62 ZJ.

It is informative to compare this value to the proved reserves of fossil fuels. As of 2020, the [Energy Institute \[2023\]](#) gives proved reserve values of 1074 billion tons of coal, 188.1 trillion cubic meters of natural gas, and 244 billion tons of oil. A conversion to energy equivalents gives

a total of 49.8 ZJ [Energy Institute, 2023, Engineering Toolbox]. 22.0 ZJ have been burned from 1800 to 2020, for a total reserves-based y_N of 71.8 ZJ. As [McGlade and Ekins \[2015\]](#) noted, this is higher than the Burnable carbon total, but not by much; i.e. HANDY2 assumes that a small fraction of proved reserves will be left unburnt to meet modest global warming targets.

Finally, we look at the estimates of Renewables production. As with Nonrenewables, the [Energy Institute \[2023\]](#) has data for hydroelectric, wind, and solar generation, while geothermal energy in that source is conflated with Regenerating biomass. However, the world’s geothermal production amounts to less than 1 EJ/yr and can be ignored [US Energy Information Administration]. These figures are given on an ‘input-equivalent’ basis, assuming incorrectly they have the same thermal efficiency as Nonrenewables; we reverse this by dividing by ψ_N . These data are combined with the table in [Smil \[2017\]](#) to give the total production per year, linearly interpolated when necessary, for each Renewables type. We then group wind and solar together to estimate current π_F , as their production is rapidly expanding as of 2022. Conversely, hydropower has reached a peak production of 18 EJ/yr by 2015. Like nuclear power, it is graphed separately [Energy Institute, 2023].

Determining the maximum accessible Renewable flows F_{start} requires some level of speculation. We begin by restricting the calculation to wind and solar power. While hydroelectricity has historically been more widely used, it has recently become increasingly controversial in large parts of the world due to its potential for environmental damage and biodiversity loss when built in areas such as the Amazon rainforest [Palmeirim et al., 2014] that do not already have dams. Hydropower has also faced difficulties with large capital costs for construction, and may be undercut in the future by cheaper solar and wind energy [Bogoviz et al., 2020, Smil, 2017]. We thus ignore it in calculations of F_{start} . In particular, wind and solar power are the two Renew-

able flows with the greatest long-term growth potential, enough to displace Nonrenewables. [Lu et al. \[2009\]](#) estimate Earth’s wind energy accessible by ground-based turbines as 78 TW, or 2450 EJ/yr. If only 10% of this is accessible, this would by itself be enough to displace half of current fossil fuel production.

Solar power has even greater potential. A quick calculation gives the incoming solar energy on Earth’s accessible land area as 15000 TW (4500 ZJ/yr)—which far surpasses all conceivable future needs [[Smil, 2017](#)]. Yet only a small amount of this power is feasibly accessible today. [Li et al. \[2018\]](#) investigate the climatic effects of a plan to cover 20% of the Sahara in solar panels, with a total collection area of $1.96 * 10^{12} \text{ m}^2$. Assuming a more modest initial buildout of about 1% of the desert (10^{11} m^2), and an average near-future generation capacity of 60 W/m^2 , this gives a power production of 220 EJ/yr, within an order of magnitude of the wind estimate above. We thus set F_{start} in HANDY2 to 80 EJ/yr, with F_{policy} in some scenarios driving an increase to over 350 EJ/yr. Such a scenario seems reasonable in light of the rapid expansion and decreasing cost of solar power installations [[Lal et al., 2017](#), [Louwen et al., 2016](#), [Swanson, 2007](#)].

3.4.2 Population Estimates

As with energy, we use two sources to estimate population in HANDY2. Between 1700 and 2020, the HYDE3.2 dataset maintained by Utrecht University provides population data [[Klein Goldewijk et al., 2017](#)] every 10 years. Interpolation is not necessary for our use case, as we employ only the available data in our plots and in EnKI parameter estimation. From 1950 onward, we use the [United Nations Population Division \[2023\]](#)’s yearly July 1 estimates of world population, currently up to date through 2022. The two data sets are in close agreement when

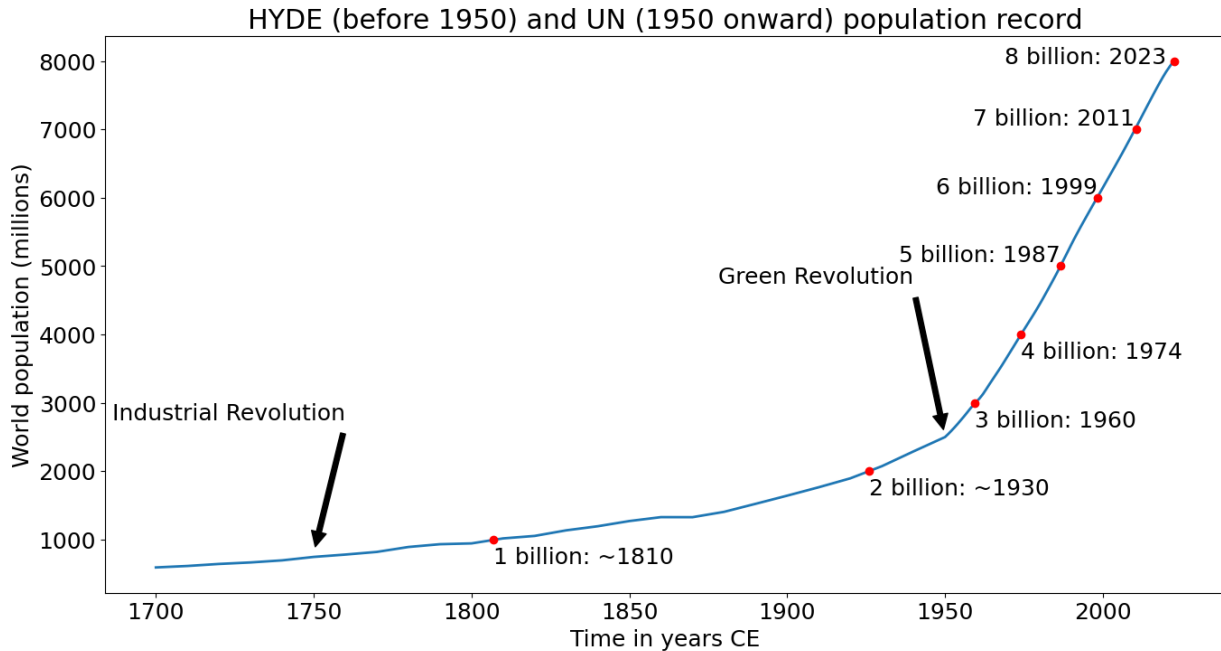


Figure 3.4: Population increase over time from 1700 to 2020, using data from HYDE3.2 [Klein Goldewijk et al., 2017] and the UN Population Division [United Nations Population Division, 2023]. Dots mark increments of 1 billion humans on Earth. Arrows point to the Industrial Revolution and Green Revolution, two events that greatly increased the rate of population growth.

they overlap, in the decades from 1950 onwards, which allows us to form a consistent record. From the model start date in 1650, the reconstructed population increases from just below 500 million to almost 8 billion in 2022, with an especially rapid increase from 1950 onward. We therefore set the starting population $x_C = 500$ in HANDY2; one unit of population represents one million people. Figure 3.4 shows the overall trend.

We use this HYDE and UN-derived record of world population to make reasonable estimates for HANDY2's starting parameter values. In the baseline scenario, we keep the starting birth rates $\beta_{C,E} = 0.03$ per year the same as in HANDY1 for the sake of consistency. This value is slightly lower than the crude birth rates of 0.04 to 0.05 per year observed in much of the least

developed world around the year 1950 [United Nations Population Division, 2023], but it works well with the retained HANDY1 maximum death rate $\alpha_{max} = 0.07$, and new HANDY2 minimum death rate $\alpha_{min} = 0.005$, the approximate lower limit seen in the developed world during the 20th century [United Nations Population Division, 2023]. We also use the UN’s moderate population projections to 2100 to guide our initial guess at the birth rate policy parameters. Values for β_{scale} of 23 years, β_{offset} of 2021 CE, and β_{final} of 0.01 per year match the projections, and are also displayed on the model results plots. We once again caution that these projections assume a decline in global birth rates to a below-equilibrium level by 2100, and may be substantially too low as a consequence [Mote et al., 2020].

3.5 Results

We now present the results for the HANDY2 model. Similar to HANDY1 [Motesharrei et al., 2014], the model parameters can be adjusted to investigate the effect of various policies on the overall trajectory of population, wealth, and resource use in the model time period. As HANDY2 is meant to cover industrial society, we focus on the five centuries from 1700 to 2200, starting the model in 1650 to allow for spinup of state variables.

In the first part of most HANDY2 scenarios, the evolution to 2022 is strongly constrained by the historical record of resource use [Energy Institute, 2023, Smil, 2017] and population [Klein Goldewijk et al., 2017, United Nations Population Division, 2023] as described in the previous section. But we do not have real data in the future. Thus, HANDY2 uses a set of scenarios to show what *might* happen in a certain situation, given a particular course of action by global policymakers or the public, or a shift in the global economy to favor one type of energy source

over another. In this section, we present seven main scenarios in total—including a 'business as usual' baseline scenario, three scenarios excluding Nonrenewables and Renewables, and three scenarios with some combination of a greatly increased use of Renewables, decreased population growth, and decreased economic inequality compared to the baseline. The first four scenarios show a high potential for global economic and social collapse by 2200, as Nonrenewables resources are rapidly depleted and not replaced. The last three avoid this collapse, and the last Scenario 6 with all three changed policy factors shows an outcome with a very high standard of living among Commoners.

3.5.1 Scenario 1: Baseline

The baseline scenario represents our best guess estimation of the moderate-term future and its social and economic evolution, assuming the policies and dynamics that are in place today continue until at least 2100. In particular, we keep inequality κ at its HANDY1 level of 100 over the entire model run, and assume global birth rates decline by 2100 to $\beta_C = 0.008$, just below the replacement level. This results in a global peak population of 10.5 billion, equal to the moderate projection from the [United Nations Population Division \[2023\]](#). Nonrenewables use peaks at near its 2022 level before declining rapidly as burnable carbon is depleted, while Renewables use is capped at 80 EJ/yr after growing at its present-day rapid pace, as easily accessible sites for solar and wind energy become more scarce over time after a rapid initial buildout. Parameters in this baseline scenario were tuned by hand to allow it to display these key features, and to match the historical record of observations to the year 2022. A full list of the parameters used is given in [Table 3.2](#).

Table 3.2: Parameters used in the baseline scenario (Scenario 1) of HANDY2. Definitions are as in Table 3.1.

Variable	Name in code	Short description	Value
Starting values of state variables			
x_C	startcommoners	Commoner population (millions)	500
x_E	startelites	Elite population (millions)	5
w	startwealth	Accumulated wealth (EJ)	0
y_R	startyrts	Regenerating resource stock (EJ)	$2.7 \cdot 10^4$
y_N	startynts	Nonrenewable resource stock (EJ)	$6.2 \cdot 10^4$
A_R	startAr	Regenerating accumulated investment (EJ)	0.4
A_N	startAn	Nonrenewable accumulated investment (EJ)	0.4
A_F	startAf	Renewable accumulated investment (EJ)	0.04
Empirically-constrained parameters			
α_{min}	alpha_min	Minimum death rate (1/yr)	0.005
α_{max}	alpha_max	Maximum death rate (1/yr)	0.07
β_C	beta_C	Starting Commoner birth rate (1/yr)	0.03
β_E	beta_E	Starting Elite birth rate (1/yr)	0.03
κ_{start}	kappa_start	Starting inequality	100
γ	gamma	Nature regeneration rate (1/yr)	0.04
λ_{start}	lambda_start	Nature regeneration capacity (EJ)	$2.7 \cdot 10^4$
F_{start}	F_start	Maximum accessible Renewables at run start (EJ/yr)	80
$\psi_{R,start}$	psi_r_start	Starting Regenerating production efficiency	0.05
$\psi_{R,end}$	psi_r_end	Ending Regenerating production efficiency	0.15
$\psi_{R,offset}$	psi_r_offset	Time offset of Regen efficiency increase (years CE)	2000
$\psi_{R,yscale}$	psi_r_scale	Timescale of Regen efficiency increase (yr)	40
ψ_N	psi_n	Nonrenewables production efficiency	see 3.9
ψ_F	psi_f	Renewables production efficiency	1

Other tunable parameters

δ_R	delta_r	Regenerating depletion rate (1/(10 ⁶ person-yr))	1.3*10 ⁻⁴
δ_N	delta_n	Nonrenewables depletion rate (1/(10 ⁶ person-yr))	2.3*10 ⁻⁶
δ_F	delta_f	Renewables depletion rate (1/(10 ⁶ persons))	7*10 ⁻⁴
λ_{growth}	lambda_growth	Nature regeneration capacity growth rate (1/yr)	0.04
s	s	Subsistence salary (TJ/person-yr)	0.005
ρ	rho	Famine wealth threshold (TJ/person)	0.05
ω_μ	omega_mu	Class mobility threshold consumption (TJ/person-yr)	1.5
μ_M	mu_M	Maximum downward class mobility (1/yr)	0.06
A_{Rth}	Arth	Regenerating accessibility investment threshold (EJ)	10 ⁷
A_{Nth}	Anth	Nonrenewable accessibility investment threshold (EJ)	315
A_{Fth}	Afth	Renewable accessibility investment threshold (EJ)	12
$\xi_{m,R}$	xim_r	Regenerating minimum extraction efficiency	0.1
$\xi_{m,N}$	xim_n	Nonrenewable minimum extraction efficiency	0.25
$\xi_{m,F}$	xim_f	Renewable minimum extraction efficiency	0.05
ω_i	omega_i	Consumption at investment onset (TJ/person-yr)	0.01
J	J	Investment scaling factor	0.01
σ	sigma	Investment depreciation rate (1/yr)	0.01

Policy knobs and ancillary parameters

β_{final}	final_br	Final birth rate (1/yr)	0.008
β_{offset}	bra_offset	Time offset of birth rate reduction (years CE)	2030
β_{scale}	bra_scale	Timescale of birth rate reduction (yr)	30
κ_{final}	kappa_final	Final inequality	100
κ_{offset}	kappa_offset	Time offset of inequality reduction (years CE)	none
κ_{scale}	kappa_scale	Timescale of inequality reduction (yr)	none
$I_{Fp,start}$	ifpstartyear	Starting year of Renewables subsidy (years CE)	1990
$I_{Fp,base}$	ifp_base	Base rate of Renewables subsidy	3*10 ⁻³
F_{policy}	F_policy	Renewables growth policy parameter	0


Model settings

t_{start}	startyear	Model start year (CE)	1650
	nyear	Model run time (yr)	550
	timestep	Model timestep (yr)	0.5

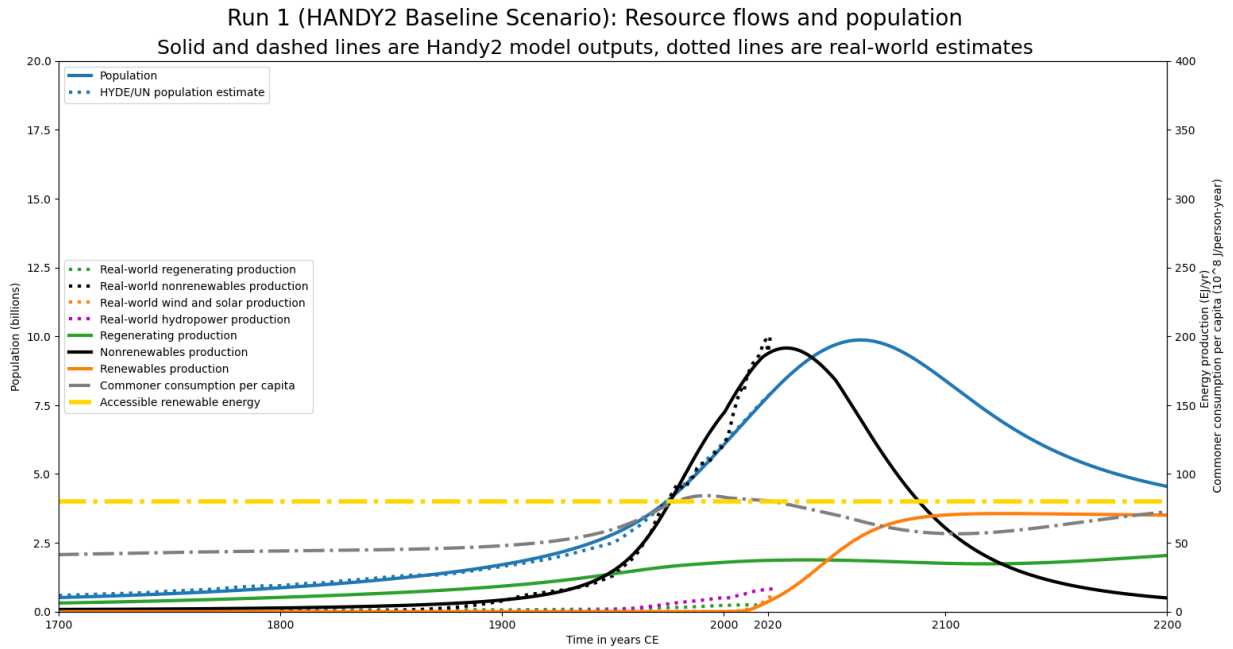
The results of the baseline scenario are shown in Figure 3.5 below, displaying both a *stocks plot* and *flows plot* for the scenario. In the stocks plot, the state variables of HANDY2 are represented in solid color, including Commoner population (blue), *equivalent* Elite population (red), Regenerating Nature stocks (green), Nonrenewables stocks (black), and accumulated wealth (gray). As in HANDY1, the equivalent Elite population is defined as the Elite population x_E times the inequality κ , representing their total consumption and impact on society. Also included is the commoner consumption per capita C_C/x_C (gray dash-dot line). Observed variables derived from historical data are represented with dotted lines. These include the HYDE/UN observed population (blue dotted line), and the estimate of Nonrenewables stock derived from the [Energy Institute \[2023\]](#) (black dotted line). The population scale is plotted in billions on the left. Energy and wealth is shown on the right in units of 10^{21} or 10^{19} J respectively; C_C/x_C is normalized to units of 10^8 J/person-yr.

In the flows plot, we show the HANDY2 estimated production of each type of resource with solid lines. Regenerating energy is in green, Nonrenewables are in black, and Renewables are in orange. The solid blue line displays the sum of Commoner and Elite populations. Dotted lines represent the real-world estimates of population (blue), Nonrenewables production (black), Regenerating production (green), solar and wind Renewables production (orange), and hydropower (purple), all scaled by usable final energy output. The gray dash-dot line is once again C_C/x_C , while the yellow dotted line is the accessible renewable energy F . The left population scale is in billions while the right energy scale is in EJ/yr, or 10^8 J/person-yr for the per capita commoner consumption.

The baseline scenario shows a striking example of a type-N collapse from insufficient resource use. The parameters chosen in the historical part of the model run allow a very close



(a) Baseline scenario stocks plot



(b) Baseline scenario flows plot

Figure 3.5: Stocks and flows plots for the HANDY2 baseline scenario. In the stocks plot, solid lines represent state variables of HANDY2, and dotted lines represent real-world population and resource stocks. In the flows plot, solid lines represent yearly energy production by resource type, dotted lines show real-world energy production, and the yellow dash-dot line shows the accessible renewable energy F . Both plots show commoner consumption per capita C_C/x_C as the grey dash-dot line. An overall collapse in the population takes place from 2080.

match to the observed history of Nonrenewables use and population. By carefully setting the Nonrenewables investment threshold A_{Nth} , the positive feedback in investment, production, and efficiency technologies for fossil fuel use activates around 1950, exactly when it historically did, causing it to grow linearly to the present. The population parameters are tuned with similar effectiveness. Slow growth causes the world population to approach 2.5 billion by the time of the Green Revolution, which then increases sharply as a massive influx in wealth from greater Nonrenewables use causes the consumption per capita C_C/x_C to rise enough to make death rates small. The Commoner population then peaks at 10.5 billion, following the [United Nations Population Division \[2023\]](#) projections and confirming our choice of β_{offset} and β_{scale} . The Elite population stays relatively low for most of the run due to the overall high maximum level of class mobility μ_M , before greatly expanding as the consumption of the Commoners begins to exceed ω_μ , and they move up in social status through the 21st century. We also note that the increased depletion of Nonrenewables and greater production technology ψ_R puts pressure on Regenerating stocks, which reach a minimum of about 10 ZJ in the mid-21st century.

The Renewable production in the baseline scenario (Figure 3.5b) grows from an average of about 13 EJ/yr today to near the limit given by F , 80 EJ/yr, in 2100. However, the Nonrenewables production peaks in 2025 and drops steeply after that time, as burnable carbon begins to run out and labor force allocation and investment begins to shift away from fossil fuels toward Renewables. By 2200, only about 54 ZJ of the starting fossil fuel stock of 62 ZJ is burned—but this includes most of the easily accessible proven reserves of carbon, leaving only harder-to-access resources of dubious economic feasibility for future generations.

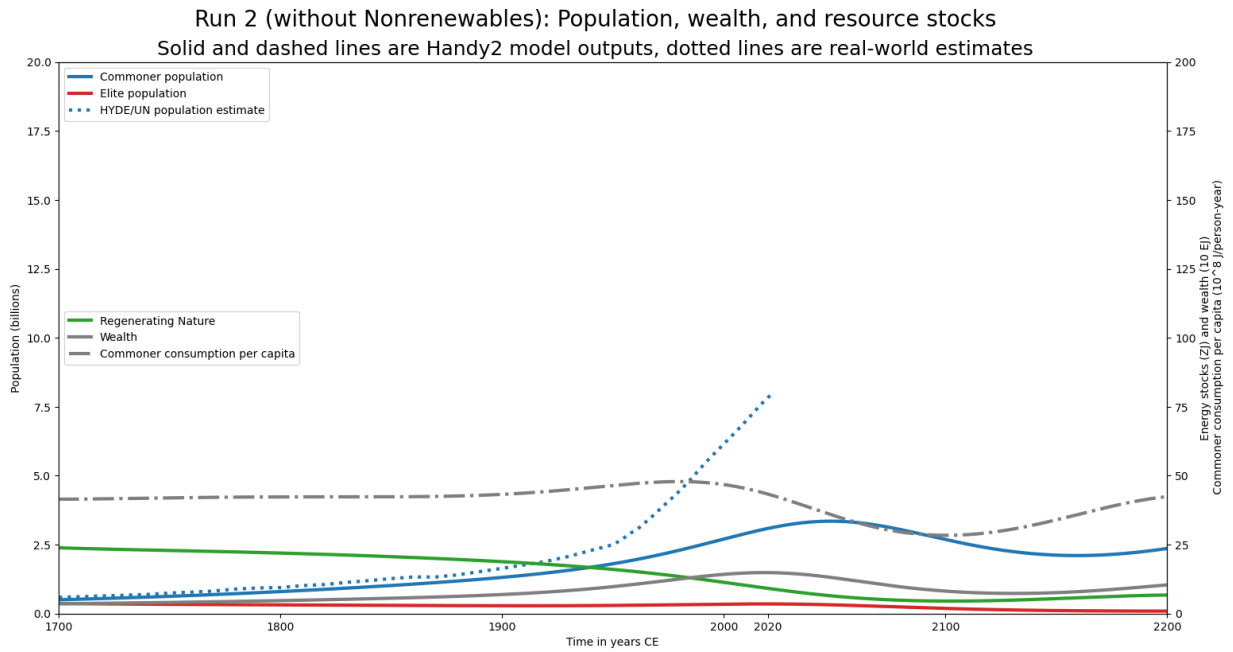
80 EJ/yr of Renewables energy is simply not enough to replace the 200 EJ/yr of fossil fuels π_N maintaining today's living standards, so C_C/x_C starts to greatly decline after 2020, reaching

a trough by 2100. This modest decrease is enough to sharply elevate Commoner death rates α_C to over 0.015. Combined with lower birth rates, the world population is reduced by over half its peak level to 4.5 billion in 2200. The situation then begins to stabilize, as the low use of Renewables restores a moderate standard of living to the survivors. Yet this society has a much smaller population and industrial base than today, and a very large and rich pool of Elites. In effect, we observe a modest Type-N collapse, and a return to a moderate-equilibrium society from today's high-equilibrium state—showing how long-run economic growth is impossible without an expanding pool of resources. This outcome is not desirable, and would be difficult to escape from due to cheap fossil fuels being largely depleted. But it is not a complete collapse of human society. From examining this baseline scenario of HANDY2 and others later in this chapter, a complete collapse is unlikely as long as some Renewables exist to maintain energy flows.

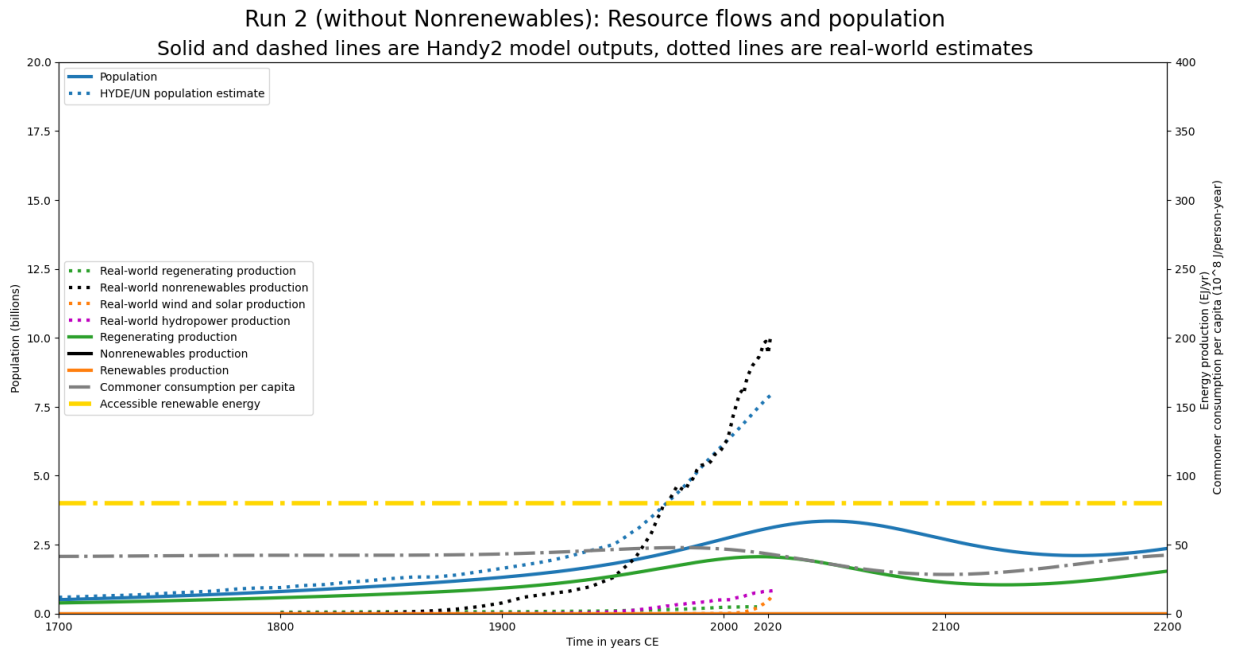
3.5.2 Scenarios 2 and 2-1: No Nonrenewables

To validate HANDY2, it is instructive to look at what might have happened in a scenario where fossil fuels were neither discovered nor exploited. This mirrors the pre-Industrial era of human society over a short time period, which will be explored further at a later date in its own branch of HANDY. We implement this Scenario 2 by setting the starting Nonrenewables stock y_N to a value of zero, and the birth rate policy toggle `bra_switch` to `false`. The latter change is equivalent to setting $\beta_{final} = \beta_C = \beta_E$. The results of this are shown in a stocks plot and flows plot with the same scale as in the baseline scenario.

In this scenario, we see the drastic effects of fossil fuels on the growth of human society and the economy. Without them, the Commoner and Elite populations grow much more slowly



(a) Regenerating-only scenario stocks plot




(b) Regenerating-only scenario flows plot

Figure 3.6: Stocks and flows plots for HAN DY2 scenario 2, with no use of Non-renewables. These plots use the same color scheme and scales as in the baseline scenario, Figure 3.5.

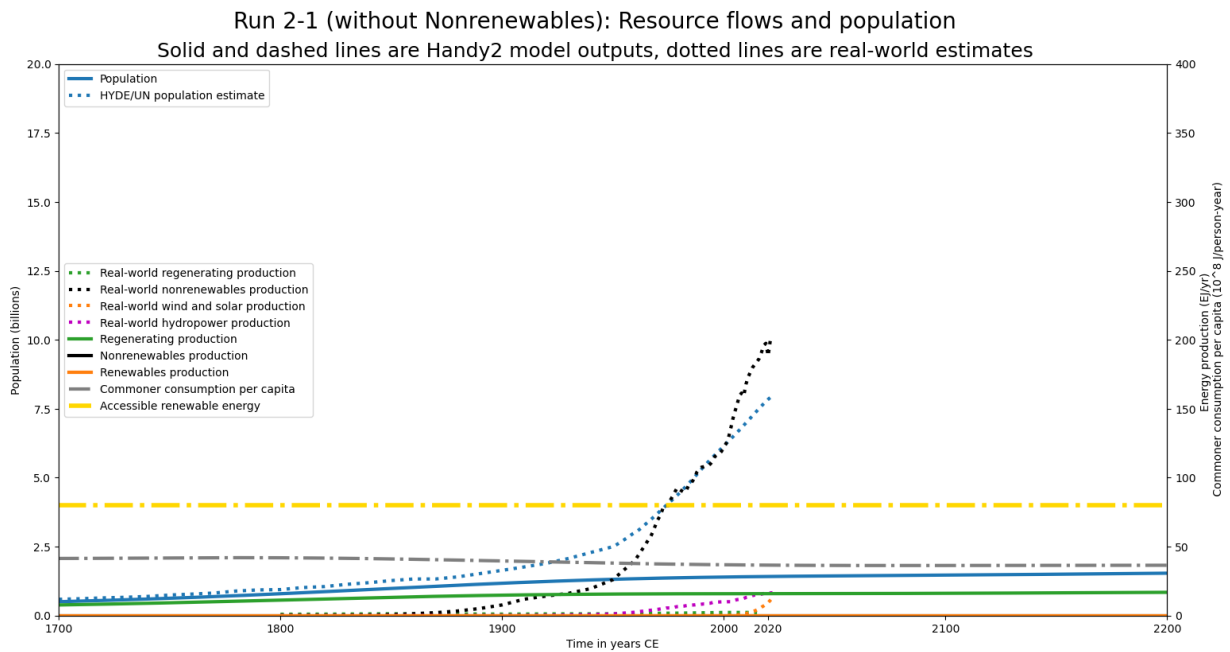
than they did in Scenario 1, peaking at 3 billion in 2050 due to the increased drawdown of Nature and the prescribed increase in Regenerating production efficiency. However, the Commoner consumption per capita C_C/x_C never exceeds pre-Industrial levels, and this causes the population to drop later in the model run and stabilize at about 2.5 billion. Nature clearly does not recover by the end of the scenario. Without fossil fuels, the limited regeneration rate of woody phytomass keeps human society in a low-equilibrium state, and this is true even with the technological possibilities of increased Regenerating energy use made possible in modern times.

We note that the increases of ψ_R production technology in HANDY2 represent the Green Revolution of the later 20th century, enabled by massive inputs of Nonrenewable energy inputs for fertilizers, specialized plant breeding, and the modern supply chain [Smil, 2017]. It is therefore instructive to examine a version of Scenario 2 without the increase in ψ_R , which we prescribe by setting $\psi_{R,final}$ equal to $\psi_{R,start}$. Figure 3.7 shows the results of this modified Scenario 2-1 below.

In this scenario, we see a slow approach to equilibrium as the Commoner population slightly grows over time, approaching a maximum value of 1.5 billion with a small production of Regenerating phytomass. Commoner wealth ω remains nearly constant, and Nature does not become significantly depleted as it does in the baseline scenario or Scenario 2, without the modeled Green Revolution putting pressure on Regenerating stocks. Maximum Regenerating flows remain near the pre-Industrial value of 20 EJ/yr [Smil, 2017]. This version of Scenario 2-1 seems to indicate that the pre-Industrial Human system was stable, but this may not be true on a local level, and is certainly less true with lower values of the class mobility factor μ_M . At lower values near $\mu_M = 0.04$, Elites grow over time, resulting in a slow type-N collapse over the HANDY2 time domain. Such an observation seems to bolster the argument made in Pomeranz [2000] that



(a) Regenerating-only scenario with flat ψ_R , stocks plot



(b) Regenerating-only scenario with flat ψ_R , flows plot


Figure 3.7: Stocks and flows plots for HANDY2 scenario 2-1, with no use of Nonrenewables and no ψ_R increase. These plots use the same color scheme and scales as in the baseline scenario, Figure 3.5.

scarcity of land and therefore woody phytomass forced people in increasingly economically productive areas of the world, such as Great Britain, to seek Nonrenewable coal as to satisfy their energy needs—which was enough to start the Industrial Revolution. We will revisit this statement in Chapter 6.

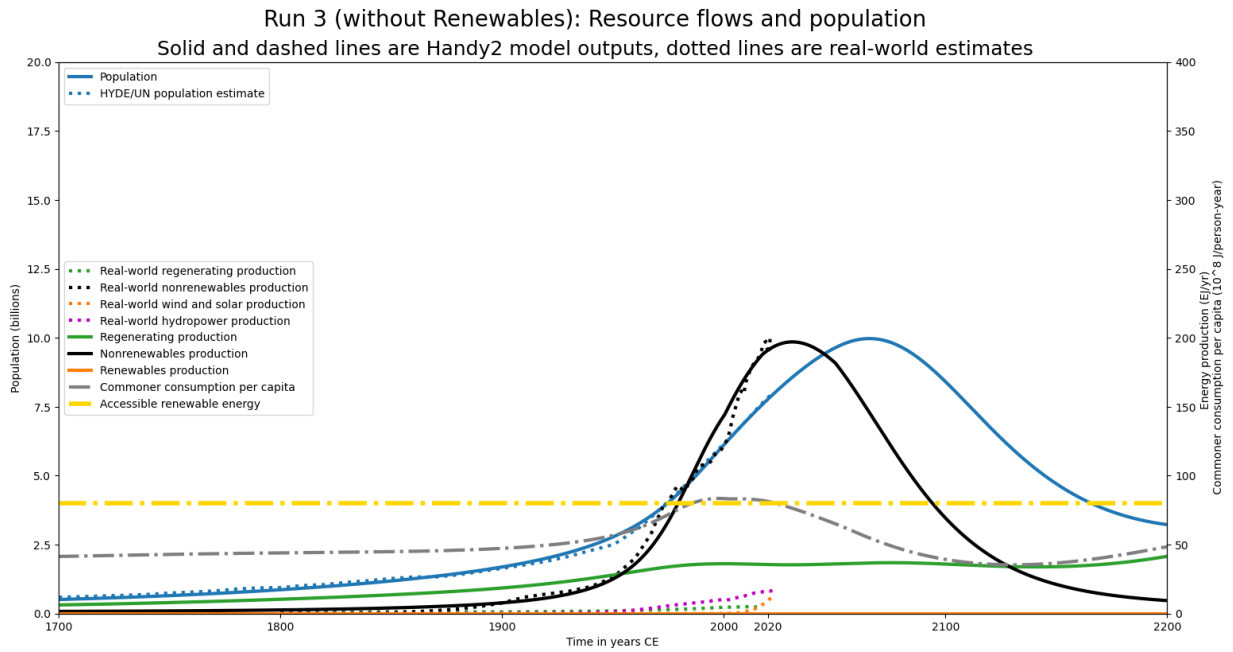
3.5.3 Scenario 3: No Renewables

In this scenario, we revert to the conditions seen in the baseline scenario with its full complement of Nonrenewable resources, but this time excluding all Renewable flows. This is done by setting δ_F to 0. We also modify the birth rate policy parameters to ensure a present-day rate of population growth until 2200—setting β_{offset} to 2000 CE, β_{final} to 0.018 per year, and β_{scale} to 11 years. We also increase A_{Nth} to 328 EJ, to ensure a better match of the population and resource curves to the historical record. Doing so allows us to see a worst-case collapse scenario for present-day industrial society, where population growth continues indefinitely into the future and Renewables are never used. We display this in Figure 3.8, with the same stocks and flows plots as in the baseline scenario.

The worst-case scenario is quite dire for the Commoners. With nothing to supplant a shrinking reserve of fossil fuels, C_C/x_C drops more rapidly than in the baseline scenario from 2020 onwards, reaching a minimum by 2120 at near pre-Industrial levels. The population grows rapidly before the mid-21st century but peaks in 2070 at the same value of 10.5 billion as in Scenario 1, due to the sharp rise in death rates and increased global poverty. At this point, the Human system is entirely unsustainable, and undergoes a dramatic type-N/L mixed collapse by the end of the model run, returning to the low-equilibrium state of Scenario 2. Even more so than in the baseline



(a) No-Renewables scenario stocks plot



(b) No-Renewables scenario flows plot

Figure 3.8: Stocks and flows plots for HANDY2 scenario 3, with no use of Renewables. These plots use the same color scheme and scales as in the baseline scenario, Figure 3.5.

scenario, the remaining 3.5 billion people of Earth have no real ability to access Nonrenewables or Renewables, making a return to present-day standards of living nearly impossible. Such an outcome is thankfully unlikely, as today’s solar and wind production is instead growing rapidly at the levels assumed in the baseline scenario.

3.5.4 Scenario 4: Inequality Reduction

We now introduce a scenario with a more positive outcome, which may be obtained by enacting policies that reduce the consumption of Elites. In the stocks plot of the baseline scenario displayed previously, and in HANDY1 [Motesharrei et al., 2014], one can see that Elite consumption in a modeled society is similar in scale to the Commoner consumption, even though the Elite population is only $1/(1 + \kappa)$ of the total—at most 1%. In many cases, this uncontrolled growth of the Elites can lead to a type-L collapse (see Chapter 2). But collapse can be avoided by limiting Elite wealth accumulation. In this scenario, we start with the same parameters as in Scenario 1 but set `kappa_switch` to `true`, and thus allow the inequality reduction policy knobs $\kappa_{offset} = 2100$ CE, $\kappa_{scale} = 20$ yr and $\kappa_{final} = 10$ to become active. The final wealth inequality $\kappa = 10$ is modest by historical standards, yet still not fully egalitarian. Inequality reduction is imposed mostly in the latter half of the 21st century, as shown in Figure 3.9. All other parameters are the same as in the baseline scenario.

The results of HANDY2 Scenario 4, as shown in Figure 3.10, are mostly the same as in the baseline scenario, but Commoners lose less wealth by 2200 in a slower resource collapse. The final Commoner population at the end of the integration is 6 billion, about 20% higher than in Scenario 1. Commoner consumption C_C/x_C is also slightly higher, staying near the model

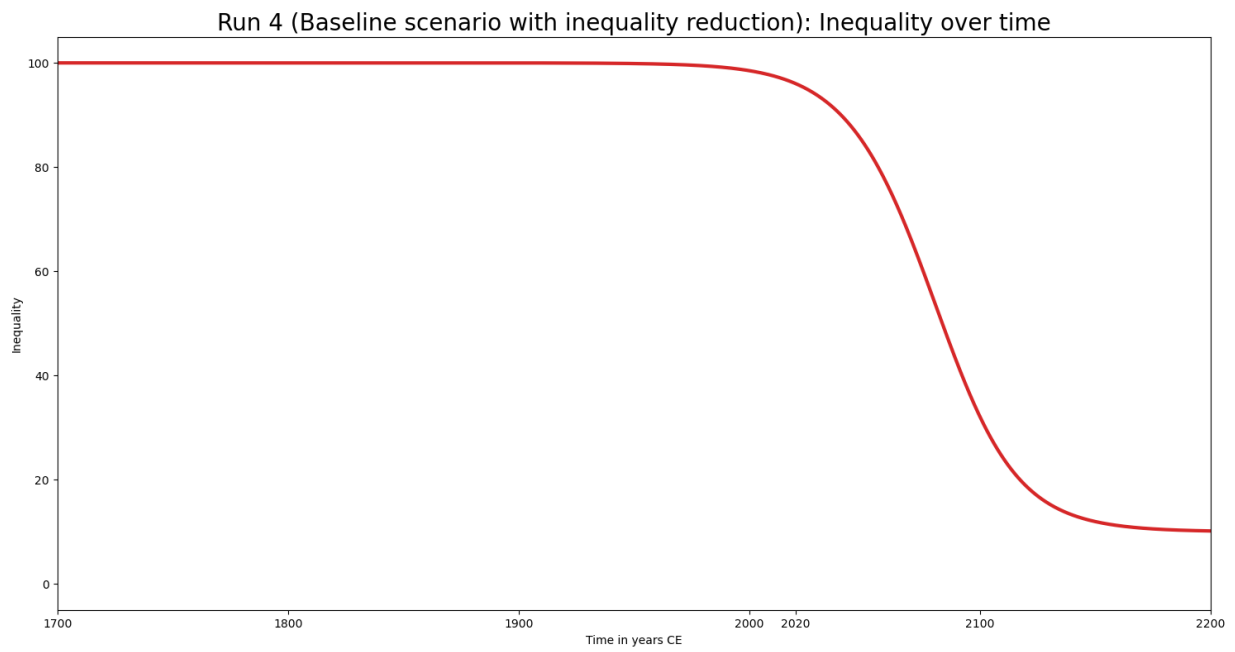


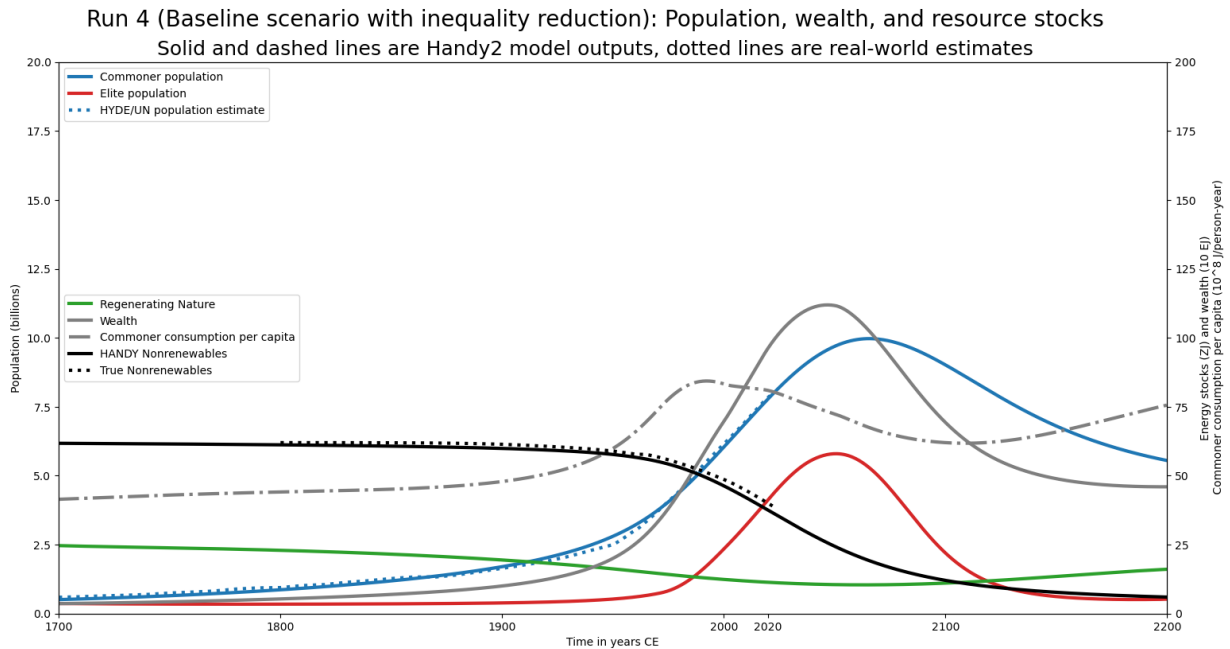
Figure 3.9: The level of inequality κ in HANDY2 Scenario 4 over time. It decreases from a starting value of $\kappa = 100$ in a log-logistic manner, reaching a floor of $\kappa = 10$ by 2150.

peak level of 6-7.5 GJ/person-yr throughout the run. Elite consumption is much lower after 2100 due to the smaller value of κ —even though Elites are still richer than the Commoners, they are not so rich as to consume more than about 10% of society’s energy production. However, the overall shortage of energy from 2050 onwards does not change. Regenerating and Nonrenewable resources are still depleted at approximately the same rate as in the baseline scenario, and the cap of 80 EJ/yr of Renewables is only barely enough to restore a moderate equilibrium by 2200. To do better, industrial society must access more resources.


3.5.5 Scenario 5: Higher Renewables Use

In Scenario 5, we greatly increase the amount of Renewables that are accessible to society over time compared to the baseline scenario, by setting F_{policy} to 1. This allows the maximum use of Renewables F to grow from about 80 EJ/yr in 2000 CE to 360 EJ/yr in 2200 CE. The linear increase in F is postulated to arise as another result of Renewables accessibility technology, with particularly improved access to solar and wind power. We also impose the inequality reduction seen in Scenario 4, with the same parameters, i.e. `kappa_switch` remains `true`. All other parameters are kept the same, and we plot the results in Figure 3.11.

This small change completely overhauls the dynamics of HANDY2, entirely preventing the resource collapse that takes place in the preceding Scenarios 1, 3, and 4. The much larger amount of accessible Renewable flows is almost fully utilized beyond 2070, as it is now a sector large enough to attract its own labor force L_F and completely replace Nonrenewables without any great decrease in per capita energy use and thus Commoner wealth. C_C/x_C remains nearly constant, with a slight dip near 2070, at about 10 GJ/person-yr. As a result, the world population remains

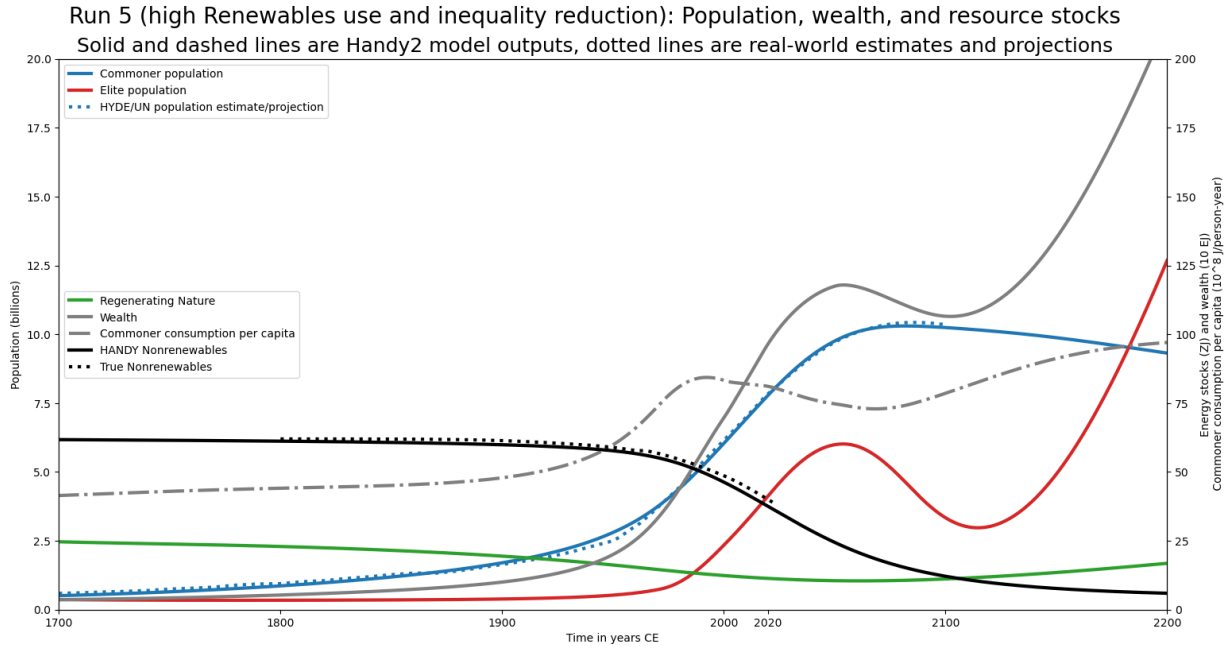


(a) Inequality reduction scenario stocks plot

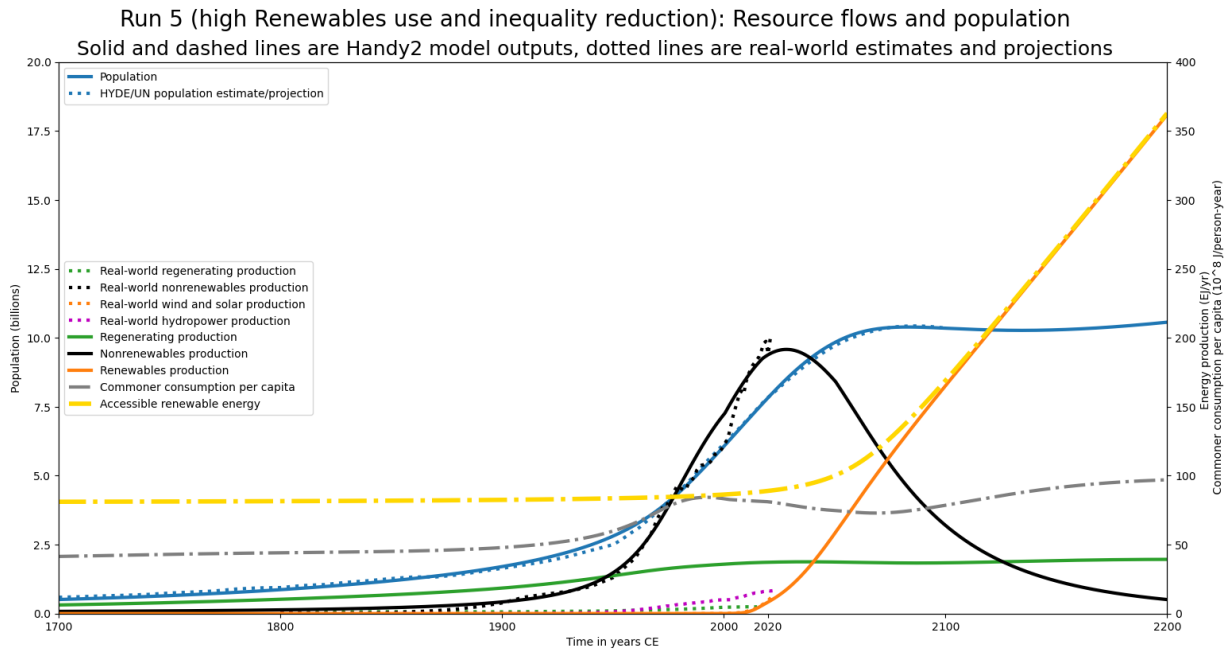


(b) Inequality reduction scenario flows plot

Figure 3.10: Stocks and flows plots for HAN DY2 scenario 4, imposing a reduction in Elite inequality over time to $\kappa = 10$. These plots use the same color scheme and scales as in the baseline scenario, Figure 3.5.



(a) Higher Renewables use scenario stocks plot



(b) Higher Renewables use scenario flows plot

Figure 3.11: Stocks and flows plots for HANDY2 scenario 5, with F_{policy} set to 1 to allow increased use of Renewables and inequality lowered over time to $\kappa = 10$. These plots use the same color scheme and scales as in the baseline scenario, Figure 3.5.

steady at its peak of 10.5 billion from 2070 onward. While the birth rate policy β_{RC} is enough to stop the population from growing by this time, in line with the UN projections to 2100 shown on the stocks and flows plots, the Commoner birth and death rates remain close to equal with each other at about 0.009 per year. Both rates slowly decrease after that time (Figure 3.12). Another dramatic change in Scenario 5 is the large growth in the number of Elites. Initially, their equivalent population declines from its present value due to the reduction in κ . But afterwards, the Commoners become rich enough, and inequality low enough, to support a very high rate of upward class mobility, almost quadrupling the Elite population from 2100 to 2200. This can be seen by the difference between the population curves in the stocks and flows plots. When Commoner population declines slowly in Scenario 5, Commoners are simply becoming Elites rather than dying.

In Scenario 5, fossil fuels are phased out at a similar rate as in the baseline scenario, not exhausting the burnable reserves of carbon despite a higher population and energy intensity. We also see Regenerating consumption remain constant, as the higher Renewables use displaces Regenerating labor force allocation L_R , reducing the extent to which biomass is currently used in the developing world [Smil, 2017]. Commoner energy needs are instead met by Renewables, with an investment rate high enough that the production π_F utilizes all of the available capacity F shown by the yellow dash-dot line. We thus see that a larger accessible pool of Renewables is the single most important change that avoids resource collapse in the 21st century—one that policymakers can easily take action on by investing in rapid construction of solar and wind energy infrastructure, large amounts of which are needed to support today's energy-intensive lifestyle. In Chapter 6, we note this finding is robust to changes in HANDY2 parameter values that we obtain from EnKI parameter estimation.

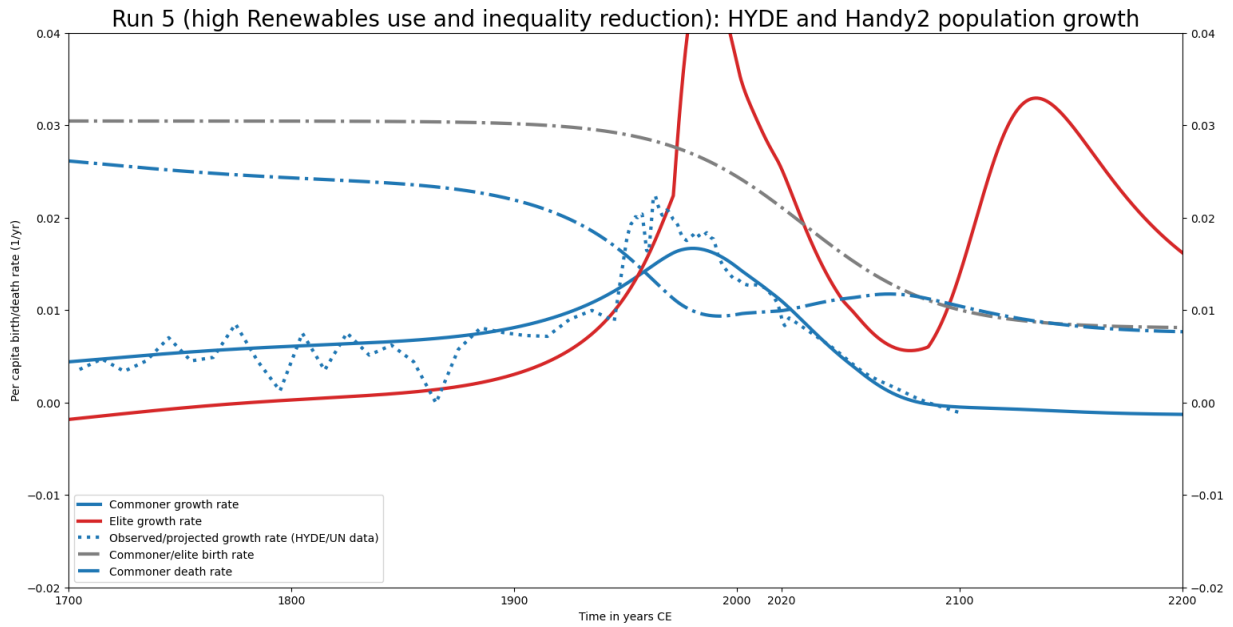


Figure 3.12: Population growth rate plot for HANDY2 scenario 5. The blue and red solid lines show the Commoner and Elite growth rate per year, while the blue dotted line shows the real-world yearly growth rate of population, using HYDE [Klein Goldewijk et al., 2017] estimates to 1950, United Nations Population Division [2023] observations to 2022, and United Nations Population Division [2023] projections afterwards. The dash-dot blue line represents the Commoner death rate in HANDY2, while the grey line is the Commoner birth rate. The area between them shows the natural growth of the Commoner population.

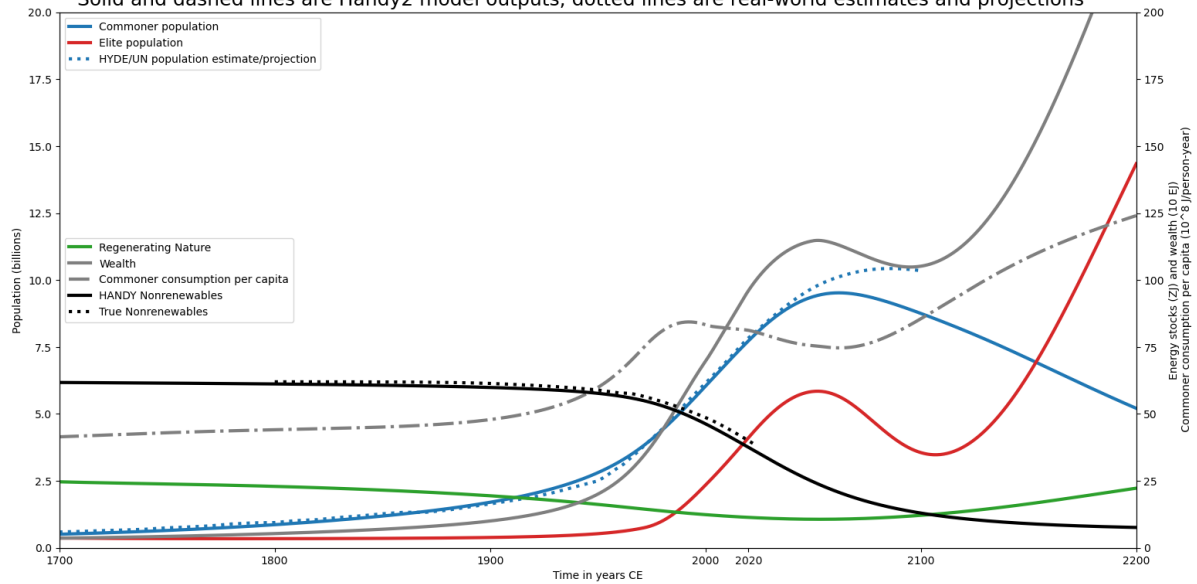
3.5.6 Scenario 6: Higher Renewables Use and Lower Birth Rates

We can make one further change in Scenario 6, compared to Scenario 5, to increase societal prosperity even further. This is done by instituting a birth rate policy that causes birth rates to decline below replacement levels. In particular, we set $\beta_{final} = 0.003$ and $\beta_{offset} = 2037$ CE, while β_{scale} remains at 30 years. This change may seem small, but it aims to represent global action by policymakers to reduce fertility more than in the UN's moderate projection, by increased promotion of family planning and education over the next 50 years. By 2070, we estimate these HANDY2 parameters equate to a total fertility rate (TFR) of about 1.7 children per woman. This is moderately low, yet still substantially higher than fertility rates in many rich world countries as of 2022 [[United Nations Population Division, 2023](#)]. The stocks and flows plots of Scenario 6 are shown as Figure 3.13.

As with Scenario 5, the rapid takeoff in Renewables output π_F in Scenario 6 entirely prevents societal collapse. Renewables grow at a linear rate for almost all of the scenario past 2020, remaining extremely close to F from 2100 onwards. Nonrenewables production peaks just after 2025 and is eventually almost entirely replaced by Renewables production, while the Regenerating output even begins to drop after 2150 due to the decrease in total population. The overall production of energy is very similar between Scenarios 5 and 6.

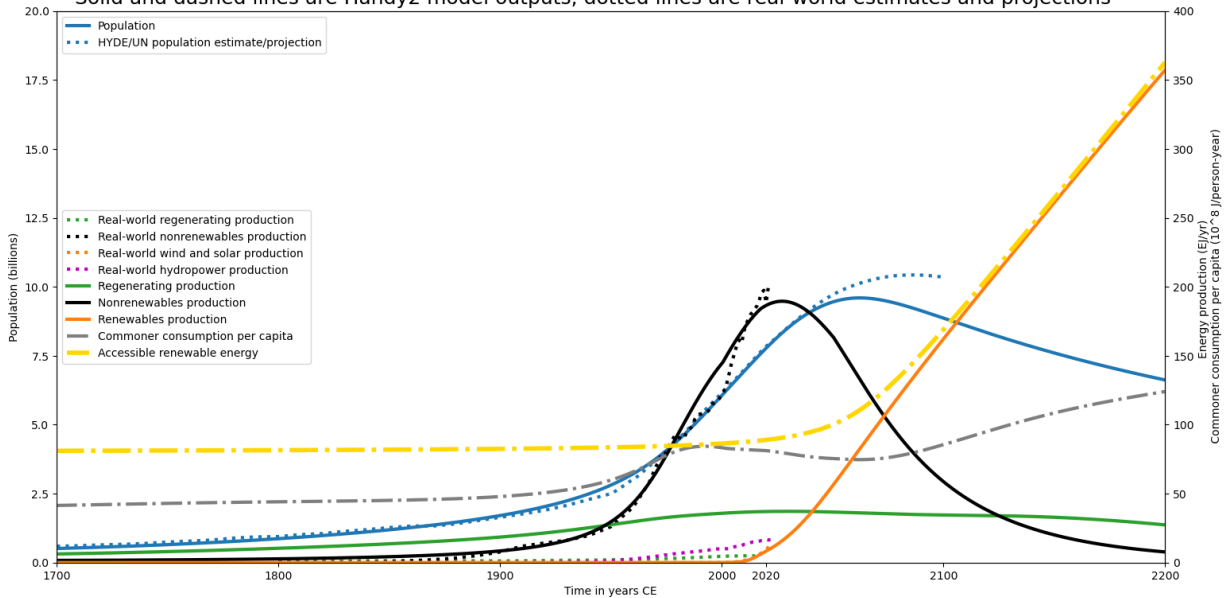
It is in the population domain where the lower birth rates demonstrate their effect of increasing Commoner wealth, as fewer people share the same resources. In Scenario 6, the total world population peaks at only 9.5 billion by 2060, and begins declining slowly to a value of about 6.5 billion in 2200. Unlike in the baseline scenario, this decrease is *entirely* due to a lower birth rate. This can be seen in the population growth plot, Figure 3.14, as death rates remain at

Run 6 (very high Renewables use, low birth rates, inequality reduction): Population, wealth, and resource stocks
 Solid and dashed lines are Handy2 model outputs, dotted lines are real-world estimates and projections



(a) Higher Renewables use scenario stocks plot

Run 6 (very high Renewables use, low birth rates, inequality reduction): Resource flows and population
 Solid and dashed lines are Handy2 model outputs, dotted lines are real-world estimates and projections



(b) Higher Renewables use scenario flows plot

Figure 3.13: Stocks and flows plots for HANDY2 scenario 6, setting F_{policy} to 1 to allow increased use of Renewables, adopting a stronger policy to lower birth rates β_C , and reducing inequality over time to $\kappa = 10$. These plots use the same color scheme and scales as in the baseline scenario, Figure 3.5.

present-day levels past the year 2000. Commoner consumption C_C/x_C reaches a value of 12.5 GJ/person-yr by 2200, far higher than in any other scenario and nearly 1.5 times the baseline scenario's value in 2020. Even the total wealth w grows faster than in Scenario 5 due to the smaller population consuming less. Elite growth also follows a similar pattern as in Scenario 5, greatly increasing after 2100 due to upward class mobility becoming high at larger values of ω . This only has a small impact on the overall standard of living due to the low income inequality. By the end of Scenario 6, Elites make up about 20% of the world population, showing how the wealth in this rich world is spread far more broadly than in the baseline scenario. Scenario 6 shows the best outcome of any of the HANDY2 scenarios, reaching a high-equilibrium state where everyone can enjoy a prosperous life. This is made possible by a combination of three policy changes—increasing the accessibility of Renewable flows, reducing population growth, and limiting economic inequality.

3.6 Conclusion

HANDY2 successfully extends the domain of applicability of the original HANDY1 model to the radically different world after the Industrial Revolution. We have designed several mechanisms that improve the predictive accuracy of HANDY2 by capturing the fundamental features of modern society. In this model, the looming threat of Nonrenewable resource depletion interacts with labor force and investment dynamics, population growth projections, and class mobility, among other factors, to shape seven scenarios of the near-term future of humanity. The predictive value of HANDY2 has been buttressed by extensive evidence from demography, history, economics, and the biosciences that explain the choices made in its construction and the parameters

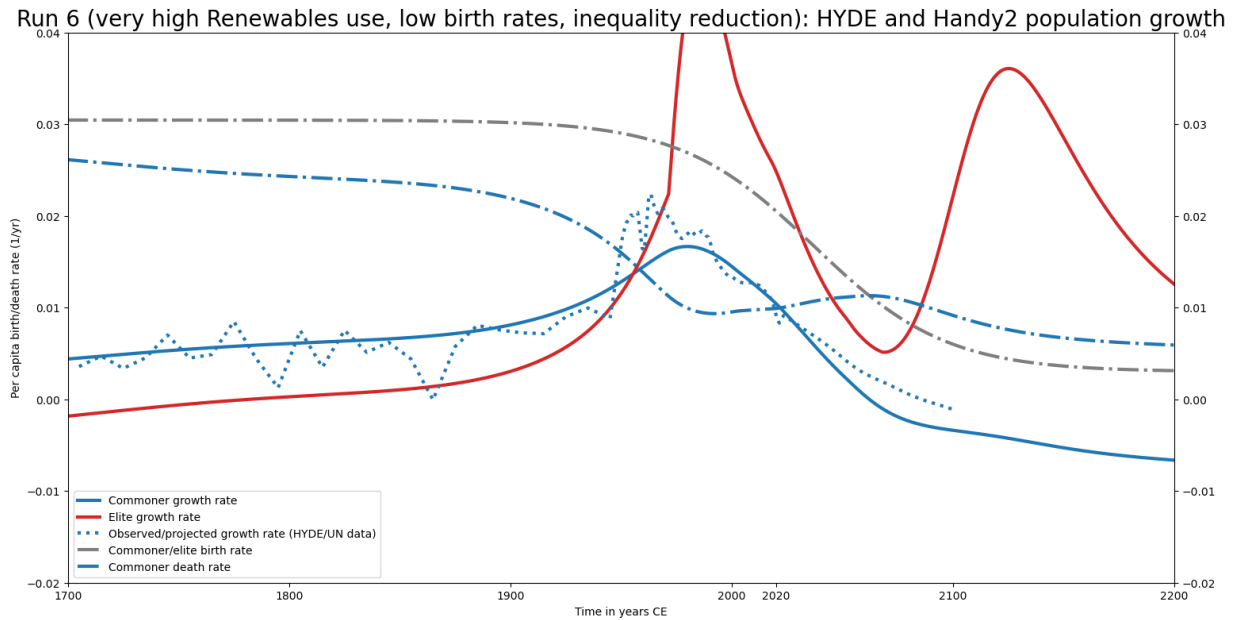


Figure 3.14: Population growth rate plot for HANDY2 scenario 6. The lines and scale are the same as in Figure 3.12.

we run it with. Those constraints, alongside plenty of experimentation, have allowed the model to accurately replicate the growth in world population, fossil fuel consumption, and phytomass use from 1800 to the present. However, we postulate that all else being equal, these current conditions will likely lead to a type-N ecosystem and Human system collapse by the end of the 21st century, greatly disrupting the global standard of living.

Collapse can be averted by making three changes—reducing wealth inequality, lowering global birth rates, and increasing the accessible amount of renewable energy. For instance, wealth and consumption inequality may be reduced through a global tax on capital above a certain level of income [Nguyen and Khieu, 2020, Piketty et al., 2022]. Birth rates across the developing world have already sharply declined, and may be further lowered through efforts by world policymakers to encourage family planning, bolster economic development, and promote the status and ability

of women to make their own reproductive decisions [de Silva and Tenreyro, 2017, Piepmeier and Adkins, 1973]. Large-scale investment in renewable energy has already taken place across the rich world [Bistline et al., 2023], and even more so in China, where a rapid expansion of solar and wind capacity has been seen in just the five years preceding 2022 [Energy Institute, 2023]. It is clear from experience that none of these three policy measures are infeasible to implement. But they will require significant political will and large-scale coordination with the aim of preventing resource collapse and ensuring a prosperous future for humanity.

There are some caveats to the current version of HANDY2. It is clear from our experimentation that even small changes to the model's baseline scenario have significant effects on the dynamics and overall trajectory of the society it simulates. With certain parameter values examined in variants of the baseline scenario, the type-N collapse may be mild, or particularly severe, or even take on aspects of a mixed type-L collapse [Motesharrei et al., 2014] from extreme growth in Elite population. These variations of course greatly affect the interpretation of all HANDY2 scenarios, and the extent of the action required to avoid bad outcomes for the Human system. In the following chapter, we will explain one method to estimate hidden true parameters in complex models, Ensemble Kalman Inversion (EnKI), and we will demonstrate its applicability to both HANDY1 and HANDY2 through further examination and optimization of both model generations.

Chapter 4: Overview of EnKI and Parameter Estimation

Throughout the process of developing both HANDY1 and HANDY2, most parameters in the model were set and fine-tuned by hand, in an attempt to match the observed trends of population, wealth, and resource use in the Human system. This process was to some degree successful, but it shows multiple drawbacks. Testing subtly different values of model parameters is extremely time-consuming and cumbersome, and there is very little possibility of sampling most of a high-dimensional parameter space in just a few, or even a few hundred runs of HANDY2. By changing one parameter at a time, we may end up optimizing the model toward a local minimum which is unrepresentative of its dynamics in the general case—thereby potentially drawing flawed or erroneous conclusions from it. Manual optimization also gives us no good estimate of the model error or its sensitivity to changes in particular parameters, which is extremely important to know if one is to have confidence in the results and future projections of HANDY2.

For these reasons, we implement the *Ensemble Kalman Inversion* (EnKI) method of parameter estimation, first described in [Iglesias et al. \[2013\]](#), to automatically determine the most likely parameter values of HANDY2 which best fit the historical observations of Commoner and Elite population, and Renewable, Nonrenewable and Regenerating energy use. EnKI is based on the ensemble Kalman filter (EnKF), which sees widespread use in meteorological and oceanographic data assimilation to determine a *future* model state from present observations [[Evensen and van](#)

Leeuwen, 1996, Houtekamer and Mitchell, 2001]. The innovation of Iglesias et al. [2013] is to instead use this method to solve the inverse problem of determining *past* states and observations. In their paper, the authors construct the EnKI method by iteratively applying the ensemble Kalman filter, estimating a model state u from observations y , and show that the estimate lies within the span of the initial ensemble A . They also note the EnKI process maintains similar accuracy to least-squares approximations over A while not requiring the derivative of the model G . Calculating the derivative would be impractical for HANDY and HANDY2. The authors then give examples of EnKI parameter estimation for various simple models. This is extended further by the Climate Modeling Alliance (CLiMA) [Dunbar et al., 2022], who provide code and documentation for estimating state within the Lorenz96 [Lorenz, 2006] model, and a 'Cloudy' cloud-aerosol model. Other authors have applied EnKF-based parameter estimation to models in the Earth sciences with similar complexity to HANDY2 [Schöniger et al., 2012]. However, applying EnKI to population and resource dynamics models like HANDY is to our knowledge novel, and is demonstrated in Chapters 5 and 6.

4.1 EnKI Mathematics

4.1.1 The Algorithm

In its most basic form, EnKI attempts to find the solution to model inverse problems of the following type.

$$y = G(u) + \eta. \tag{4.1}$$

G in this formalism is the forward model whose parameters are to be estimated. Most commonly, it is realized by a system of differential equations. u is taken as an initial set of unknowns in some Hilbert space X , while the known set of observations y is in another Hilbert space Y ; there is also a Gaussian error η with known covariance Γ . In practice, attempting to find the most accurate realization of u on a high-dimensional space is very difficult. Hence we attempt to find the solution in the constrained, finite-dimensional space A , incorporating prior knowledge of the solution in the form of the starting ensemble. The abstract model G is thus realized in the form of a dynamic model $\Xi : Z \rightarrow Z$, where $Z = X \times Y$, the product of the Hilbert spaces, $\Xi(z) = \begin{pmatrix} u \\ G(u) \end{pmatrix}$, and artificial dynamics $z_{n+1} = \Xi(z_n)$ arise from repeated applications of Ξ . In this setting, [Iglesias et al. \[2013\]](#) define observed data as (eq. 3):

$$y_{n+1} = H(z_{n+1}) + \eta_{n+1}, \quad (4.2)$$

where H is the projection operator taking Z to Y , and η is a sequence of independent, identically distributed normal deviates in Y with covariance Γ .

To initialize the ensemble Kalman inversion, one begins by taking an ensemble $\psi_{j=1}^{(j)}$ of observations in X and defining the initial state in Z to be ([Iglesias et al. \[2013\]](#), pg. 4):

$$z_0^{(j)} = \begin{pmatrix} \psi^{(j)} \\ G(\psi^{(j)}) \end{pmatrix}. \quad (4.3)$$

The finite-dimensional space A is then defined as the span of the set of observations $z_0^{(j)}$, chosen deliberately to ensure an optimal estimate. The first guess of the model truth is simply the mean of the starting ensemble. If no prior probability distribution is known, the ensemble becomes a

truncated basis of X ; otherwise with a specified prior, we obtain a Karhunen-Loève (KL) basis [Iglesias et al., 2013].

In the EnKI prediction step, the artificial dynamics are first used to evolve the ensemble and introduce the forward model (Iglesias et al. [2013], eq. 9).

$$\hat{z}_{n+1}^{(j)} = \Xi(z_n^{(j)}). \quad (4.4)$$

This evolved state is used to obtain the ensemble mean and covariance (Iglesias et al. [2013], eq. 10,11):

$$\bar{z}_{n+1} = \frac{1}{J} \sum_{j=1}^J \hat{z}_{n+1}^{(j)} \quad (4.5)$$

$$C_{n+1} = \frac{1}{J} \sum_{j=1}^J \hat{z}_{n+1}^{(j)} (\hat{z}_{n+1}^{(j)})^T - \bar{z}_{n+1} \bar{z}_{n+1}^T. \quad (4.6)$$

Next comes the analysis step, where the particles in the ensemble are adjusted to better fit the data. The Kalman gain K_n is defined as (Iglesias et al. [2013], eq. 12):

$$K_{n+1} = C_{n+1} H^T (H C_{n+1} H^T + \Gamma)^{-1}. \quad (4.7)$$

The ensemble is then updated as follows, where $\eta_{n+1}^{(j)}$ are as before independent normal error vectors with covariance Γ (Iglesias et al. [2013], eq. 14, 15). The right term introduces the observation innovation at each timestep.

$$z_{n+1}^{(j)} = (I - K_{n+1} H) \hat{z}_{n+1}^{(j)} + K_{n+1} (y + \eta_{n+1}^{(j)}). \quad (4.8)$$

Finally the mean of the ensemble update, u_{n+1} , is calculated. If the ensemble has not yet converged, the update steps are repeated until it does. Convergence can in general be determined according to the discrepancy principle—EnKI iteration stops when the norm of $y - G(u_n)$ in covariance space Γ falls below some constant at least equal to the overall norm of the noise η^* . This noise is also beneficial to the stability of the error minimization. Adding noise allows the particles to herd less within A and so gives a better final approximation of u [Iglesias et al., 2013].

4.1.2 Properties of EnKI

Most importantly, the EnKI iterative process always keeps the particles $z_{n+1}^{(j)}$ in the same space A as they were initialized. This follows from the analysis step and Kalman gain being linearly conservative in any EnKF method, and the EnKI update functions being defined to be linearly conservative. Iglesias et al. [2013] prove this statement by induction on the time step n (their Theorem 2.1). This allows those authors to make the following two observations. First, as the starting ensemble forms a basis of A , any solution $p_{n+1}^{(j)} \in A$ may be expressed as a linear combination of the ensemble. Second, the minimum EnKI error at any iteration step n cannot be less than the infimum of the error in A ; i.e.:

$$\|u_n - u^*\| \geq \inf_{v \in A} \|v - u^*\|, \quad (4.9)$$

where u^* is the model truth.

Third, the EnKI solution is usually not much worse than the least-squares minimization of the error in A , $u_{BA} = \min_{u \in A} \|u - u^*\|_{\Gamma}^2$. This is confirmed in the numerical experiments conducted by Iglesias et al. [2013]. In the particularly simple case where the forward model G is

linear, as are the operators C and Γ used to calculate the Tikhonov-Phillips regularization of the least squares solution, the authors show (eq. 41) that in the limit of more ensemble members and iterations, the EnKI solution asymptotically approaches the least-squares solution.

$$u_1 \rightarrow u_{TP} = \bar{u} + CG^*(GCG^* + \Gamma)^{-1}(y - G\bar{u}) \quad (4.10)$$

This is highly advantageous in the case of simple models G , as calculating the least squares regression requires that the derivative of the forward model exist and can be run in a reasonable time, while EnKI does not. For more complex cases where the model is non-linear, this equivalence is not necessarily true. But the two solutions are quite close to each other in practice [Iglesias et al., 2013], with EnKI much faster to iterate. The number of data points usually assimilated with EnKI at each iteration is generally small, so the runtime is dominated by the number of particles N_{ens} times the number of iterations N_{iter} . These tend to both be much smaller than the number of dimensions of many practical models G [Iglesias et al., 2013]. Estimating the model gradient, as needed in any variational method, can itself be a significant mathematical and programming task which is likely infeasible in the case of HANDY2. We gladly take the opportunity to avoid it by using EnKI.

4.2 EnKI for Lorenz96

Iglesias et al. [2013] and their subsequent work done as part of CLiMA [Dunbar et al., 2022] provide several examples of EnKI used in practical Earth system models. They also maintain Julia-language libraries that implement EnKI and related algorithms. One particularly instructive example is given in `lorenz_example.jl`, whose tools we borrow for HANDY2

EnKI. This code estimates two parameters in the Lorenz96 [Lorenz, 2006] model: the mean F and the amplitude of transient oscillations A , given a starting prior estimate of F and A , and observations of the wind speeds y_0 and y_1 throughout the Lorenz96 runs. The code then generates plots of the estimated parameters in both the unconstrained and constrained parameter spaces (see below).

`lorenz_example.jl` begins by calling the `EnsembleKalmanProcesses` (EKP) library, specifying where the output figures are to be saved, and using a fixed random (RNG) seed for reproducibility. The model settings are then specified, including the run duration and characteristic timescale τ_c of Lorenz96. There is also an option to select which output statistics y^* to use: a mean y^* over the entire run duration, or a linear approximation over time. Then the model truth is specified, including the hidden parameters F and A , and the transient oscillation frequency ω . The overt parameters are also provided, including zonal model size N , starting time t_{start} , timestep dt , batch length t_s , integration length t_{fit} , and initial perturbation F_p .

The next step is to construct the prior parameter distribution for the ensemble. This is done by invoking `constrained_gaussian`, which takes as input the prior mean, standard deviation, and the upper and lower bounds of a specified interval. The function then returns a prior distribution in *constrained space*, and a mapping $T : U \rightarrow \Phi$ that transforms the initial specified normal distribution into a distribution whose support is on that interval. This prevents unphysical parameters from breaking the model if introduced into the ensemble. A full explanation and guidance for using constrained distributions is provided at [Climate Modeling Alliance \[a\]](#).

A Boolean flag exists in the code to select whether to prescribe the EnKI covariance Γ or to learn it from the forward model. In the latter case, Lorenz96 is run once with the true model parameters, generating a *truth run*. The average of y_0 and y_1 within a data collection period

of length t_{fit} is determined for $n_{samples}$ windows. These observations are used to construct the 4×4 covariance matrix Γ . Otherwise, the mean wind observations over the entire truth run are adjusted with multiplicative noise of magnitude given by `noise_level`. Γ is computed from $n_{samples}$ such noisy samples. In EnKI for HANDY1 and 2, we use somewhat different strategies to estimate starting covariance, which are discussed further in Chapters 5 and 6.

Code objects are constructed to initialize the EnKI iteration process. An `Observation` object is first used to contain the true observations and covariance Γ . Then `construct_initial_ensemble` constructs the initial ensemble of model parameters u_0 from the prior distributions, the number of ensemble members N_{ens} , and the random seed. An EnKI object `ekiobj` is then made using `EnsembleKalmanProcess`, containing the initial ensemble, the mean of the observations y , the covariance Γ , and the `Inversion()` object, with the option to specify additional parameters such as failure handlers or learning rate schedulers that are not used for `Lorenz96`.

At each of N_{iter} iteration steps, `get_phi_mean_final` obtains the ensemble parameters u from the `ekiobj`. The `Lorenz96` model is then run with those parameters, producing a set `g_ens` of N_{ens} outputs. `update_ensemble!` then performs the EnKI update algorithm, modifying the ensemble in-place, which is then used to return the mean square error estimate between the ensemble mean u_n and the true parameters u^* . The EnKI error is displayed in terminal output. At the end of the iteration loop, the mean of the ensemble is displayed as the posterior estimate for the `Lorenz96` parameters u^* and saved to disk.

Finally, two animated plots are produced. The first plot shows the convergence of the ensemble toward the `Lorenz96` model truth in unconstrained parameter space, using `transform_constrained_to_unconstrained` to recover the normal distribution of the priors.

The second plot is the same graph in constrained space. We show the first and final frames of the second plot in Figures 4.1 and 4.2.

From the plots, it is clear that EnKI works quite well for parameter estimation in Lorenz96. Within 5 iterations, the ensemble of parameter estimates quickly converges from a wide spread of initial guesses for F and A to being tightly clustered around the true values of both parameters. The ensemble mean u_n is off from the truth by less than 1%. Dunbar et al. [2022] demonstrate their code works well to recover hidden variables in simple situations where the truth run for a model is known. We therefore adapt some of this code in Chapter 5 to estimate the key parameters in HANDY1, after making significant changes to account for our model's increased complexity. HANDY2, due to its even more complex structure, lack of known model truth u^* , and existence of good prior constraints on some of its parameters but not others, is handled separately in Chapter 6.




Figure 4.1: A plot of the EnKI parameter estimation ensemble in constrained space after 1 iteration, produced by the `lorenz_example.jl` code from [Dunbar et al. \[2022\]](#). Blue dots show F (horizontal) and A (vertical) parameters for each ensemble particle. The intersection of the red lines shows the model truth.

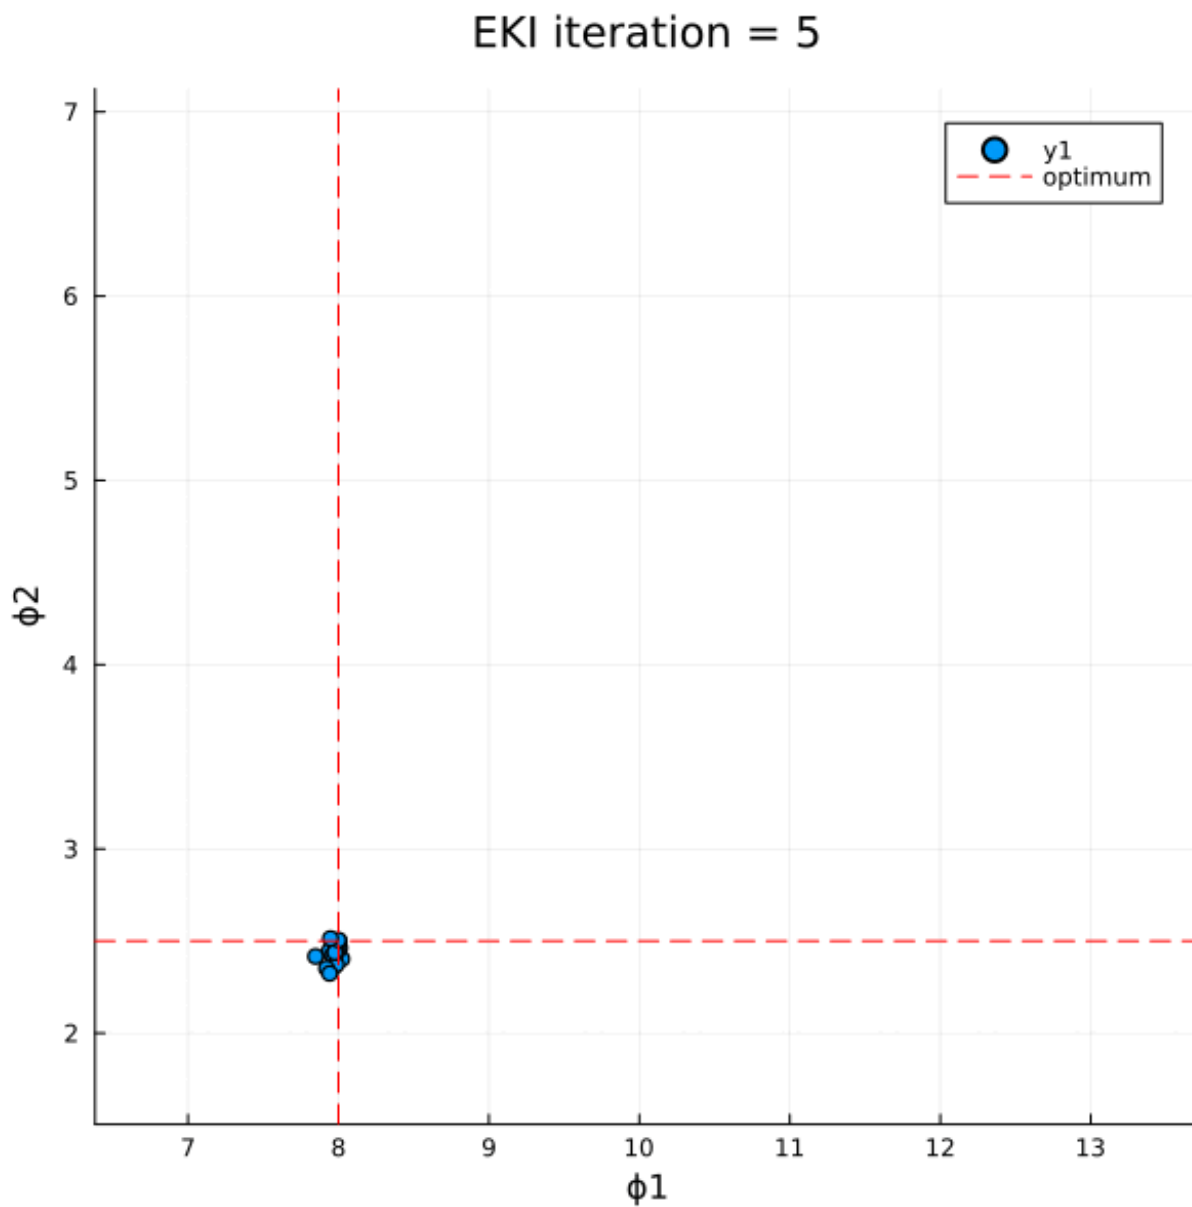


Figure 4.2: A plot of the final EnKI ensemble after 5 iterations, produced by the `lorenz_example.jl` code from [Dunbar et al. \[2022\]](#). The plot elements are the same as in Figure 4.1.

Chapter 5: HANDY1 Parameter Estimation

We now turn to implementing EnKI parameter estimation for the HANDY family of models. In Chapter 2, we examined the structure of HANDY1 as published by [Motesharrei et al. \[2014\]](#), that covers one growth and collapse cycle in a hypothetical pre-Industrial human society. It can be seen from their findings that even in a model as simple as HANDY1 with only 12 parameters, small changes in initial and boundary conditions can guide the development of very different societies. As we further expanded upon in HANDY2 (Chapter 3), low rates of depletion and inequality lead to stable societies that approach their ecological carrying capacity, while high rates of either can quickly lead to resource collapse.

With the tools of EnKI, we can ask the reverse question: given the observations of population, wealth, and resource use of a society in a HANDY1 scenario, can we recover the model parameters that gave rise to that scenario? The answer is yes. Subject to certain constraints on the starting ensemble, we use the parameter estimation algorithm and implementations outlined in Chapter 4 to create code (`handy1_parameter_estimation.jl` in our supplementary materials) that finds the parameter values in HANDY1 scenarios to great accuracy. We successfully recover the parameters in Egalitarian approaches to equilibrium, as well as Equitable and Unequal scenarios with type-L and type-N collapses, given pseudo-observations with noise of the model's state variables x_C , x_E , y , and w . We also establish that the EnKI process is more

fragile for HANDY models than for Lorenz96. There are state transitions in HANDY1 beyond which any ensemble particle will fail to converge to the model truth. We give examples of those inversion failures and discuss their overall effect on parameter estimation.

As with Lorenz96, HANDY1 parameter estimation approximates a *known* model truth that may not correspond to conditions in the real world. We demonstrate and validate this process as a stepping stone toward parameter estimation in HANDY2, which approximates an *unknown* model truth in a higher-dimensional search space. That process is covered in Chapter 6.

5.1 The EnKI Implementation for HANDY1: Building on Lorenz96

The `handy1_parameter_estimation.jl` code we write to implement EnKI [Iglesias et al., 2013] for HANDY1 works very similarly to CLiMA’s Lorenz96 example, but is extended to account for the differences in our model. Most importantly, because HANDY1 is not periodic, we cannot simply average the observation signal over time in order to produce usable data points for inversion. We are forced to take a different approach by producing a mean squared error, or *loss* vector, between the ensemble output and modeled truth for each parameter i ($\|y_i^* - y_{n,i}\|_2^2$) and using that in the `EnsembleKalmanProcesses` library [Dunbar et al., 2022] by optimizing the loss vector to a target of zero. We also determine an assumed covariance Γ based on variability *within* the truth (or nature) run, perturbed slightly by an estimate of multiplicative noise.

The code begins by invoking the Julia version of the HANDY2 model. This is identical to the Jupyter notebook version used in Chapter 2, and returns an array `H1Output` of the state variables of the model at each timestep. Then, the parameters and settings for the model truth

scenario are defined, used in a first pass of HANDY1 which computes the 'true' observations y^* . These values can be changed to test parameter estimation efficiency for the various scenarios defined in [Motesharrei et al. \[2014\]](#) or elsewhere. The random number generator (RNG) seed is also set to a specific integer to ensure reproducibility of the results. Here, we use the value 42.

Next, the list of parameters to estimate with EnKI, contained in fields of the `H1Params` struct, is defined in `param_names`. The other parameters are kept at the true values set in a particular HANDY1 scenario. For each of these test parameters, a prior distribution needs to be set using `combine_distributions` and `constrained_gaussian`. As in [Lorenz96](#), these functions take a prior mean, standard deviation, upper bound, and lower bound; the output distributions are then packaged into a convenient `priors` object. Some thought has to be provided when selecting the upper and lower bounds in order to prevent unphysical model results. In particular, HANDY1 breaks if the birth rates are not within the range of the minimum and maximum death rates; i.e. $\alpha_{min} < \beta_C, \beta_E < \alpha_{max}$ must be satisfied. If this constraint does not hold, the population necessarily goes to zero or infinity in the model run and parameter inversion is not successful. Similarly, extreme values of other parameters can cause particle failure and consequent loss of dimensionality of the inversion ensemble.

A list of output parameters used to constrain the inversion is then specified. In most cases, we use observations of all four HANDY1 state variables: Commoner population x_C , Elite population x_E , wealth w , and Regenerating nature y . These are each vectors y_i^* representing the entire history of the HANDY1 scenario at each time step. To each observation, we then add a multiplicative independent noise η of mean 0 and standard deviation given by `noise_level`, usually 0.02, that creates a set of pseudo-observations. That is, for each timestep t and parameter

i :

$$\tilde{y}_{i,t} = (1 + \eta)y_{i,t}^*. \quad (5.1)$$

The perturbed output pseudo-observations often change greatly over time during the spinup, growth, and collapse phases of HANDY1, even in runs that converge to an equilibrium. Therefore, we can use them to estimate the true covariance matrix Γ_u , with elements $\Gamma_u(i, j) = \text{Cov}(\tilde{y}_i, \tilde{y}_j)$, providing better results than *a priori* estimation. This strategy is a continuous version of the one used in `lorenz_example.jl`, which breaks the truth run up into $n_{\text{samples}} = 20$ periods before taking the parameter mean in each of them and using that to calculate covariance.

The standard `EnsembleKalmanProcess` objects are now instantiated with the initial ensemble constructed from the ensemble priors, noisy 'model truth' pseudo-observations, the covariance matrix Γ_u , the `SampleSuccGauss` failure handler, the `EKSStableScheduler`, and the RNG seed. To use the loss-based inversion within the code, we set `truth_sample` to a zero vector. `update_ensemble!` accepts only a single observation vector per ensemble member and inversion iteration, so it must be the error estimate instead of a list of observations y_t .

As of the latest version, `EnsembleKalmanProcess` gives an option to use one of several learning rate schedulers [[Climate Modeling Alliance, b](#)]. In most cases, we use the `EKSStableScheduler` [[Kovachki and Stuart, 2019](#)]. This stable scheduler takes two parameters: the first is the *numerator* α , which determines the learning rate far from the area of convergence, while the second parameter, the *kernel* ϵ , determines the learning rate Δt close to the area of convergence. Large values of ϵ cause the ensemble to converge more slowly. Specifically, this learning rate is given as:

$$\Delta t = \frac{\alpha}{\|U\| + \epsilon} \quad (5.2a)$$

$$U = (G(u) - \bar{G}(u))^T \Gamma_u^{-1} (G(u) - y) \quad (5.2b)$$

where G is the forward HANDY model, and Γ_u is the covariance matrix at each iteration. For HANDY1, we choose $\alpha = 10$ and $\epsilon = 0.01$.

Other schedulers exist that can change learning rate over time, such the `DataMisfitController` [Iglesias and Yang, 2021], but we have not tested them with HANDY. In addition, we invoke the `SampleSuccGauss` failure handler to replace failed ensemble particles with the average of the remaining particles. While this reduces the dimension of the search space and does not resolve underlying model stability issues, it can be useful to prevent inversion runs from failing when only a few particles diverge.

EnKI iteration begins by creating a `handylensparams` object containing ensemble parameters for HANDY1, generally different from the truth-scenario parameters. For each iteration, `get_phi_final` obtains the list u_n of ensemble test parameters to be estimated and `setfield!` updates each particle’s copy of `handylensparams` with the test parameters. HANDY1 is then run for each particle, producing the output y_t (`ensembleoutputs`). Some code optimization took place to ensure this can be done quickly. The mean square loss vector `g_ens` is then calculated for each particle. This is equal to $\|\tilde{y}_i - y_{t,i}\|_2^2$ for each HANDY1 state variable i under observation. The ensemble is updated with `update_ensemble!`, and the ensemble EnKI error over all output variables and particles $\|\tilde{y} - y_t\|_2$ is then calculated. Parameter estimation is continued for N_{iter} iterations; it can also end early if the relative change in the global EnKI error

between iterations is less than a certain threshold (usually 10^{-3}).

We then produce plots from the ensemble mean output identical to those in Chapter 2. The posterior parameter estimates are obtained by `get_φ_mean_final` before being used in a final pass of HANDY1. From the plots, we can see at a glance if the EnKI process produces a scenario identical to the model truth. If that is the case, we deem the parameter estimation a success. We also use the plots to analyze what goes wrong in the case of inversion failure.

5.2 Parameter Estimation Results

5.2.1 Egalitarian Scenarios

5.2.1.1 Estimating δ

We begin by estimating parameters for an Egalitarian scenario of HANDY1 (Figure 2.1a and sections 5.1.1-5.1.3 in Motesharrei et al. [2014]). In this scenario, we set the starting Elite population $x_{E,start}$ to zero and the baseline depletion δ to $6.67 * 10^{-6}$ (see Table 5.1). Since there are no Elites in an Egalitarian society, we delete x_E from our list of output observations in the parameter estimation code for all Egalitarian runs. All scenarios are run for 600 years with a timestep of 0.5 years, except when otherwise specified.

EnKI estimation can be done most easily for one parameter at a time. The most obvious candidate is δ , the depletion rate, as it varies greatly between runs and gives rise to very distinct dynamics in HANDY1 depending on its value, with stable societies, oscillating equilibria, and full collapses all possible. We note that δ tends to vary across several orders of magnitude. Thus, a reasonable prior for estimating δ is the log-normal distribution given by a lower parameter

limit of 0, an upper limit of ∞ , a particular prior mean μ_p , and a standard deviation σ_p . We set $\sigma_p = 0.2\mu_p$ to make the starting particle density lower near the model truth.

Estimating the single parameter δ in an Egalitarian society works quite well. We quickly attain convergence to the scenario truth u^* across a wide range of starting prior means μ_p , within just 10-20 iterations and a wall-clock time between 0.5 and 1 seconds on an M1 MacBook Air. The radius of convergence extends at least one or two orders of magnitude, from $0.06u^* \leq \mu_p \leq 7u^*$ in the soft-landing scenario (Figure 5.1), and this works regardless of the true value of δ in soft-landing runs and full collapses alike. Convergence stability however degrades outside these bounds. While the particles do not diverge to infinity at very high prior mean depletion δ , they are instead attracted to a spurious local minimum. It is still evident that HANDY1 is very stable when it comes to perturbations in δ .

5.2.1.2 Estimating Many Parameters

To further examine the convergence of parameter estimation in the Egalitarian HANDY1 scenarios, we estimate eight of the model's twelve parameters simultaneously: α_{min} , α_{max} , γ , λ , $x_{C,start}$, δ , s , and ρ . The Elite-related parameters $x_{E,start}$, β_E , and κ are excluded as they are not relevant for the Egalitarian run. β_C also remains constant since it is difficult to keep between α_{min} and α_{max} as required for model stability. To simplify the estimation, we let all the prior distributions be log-normal, specifying the parameter bounds as $[0, \infty)$. The prior standard deviation σ_p is equal to $0.2\mu_p$ in all runs, and μ_p is allowed to vary.

The estimation of many parameters is somewhat troublesome, but still works fairly well with the addition of noise to the observations. When the prior mean μ_p equals the scenario

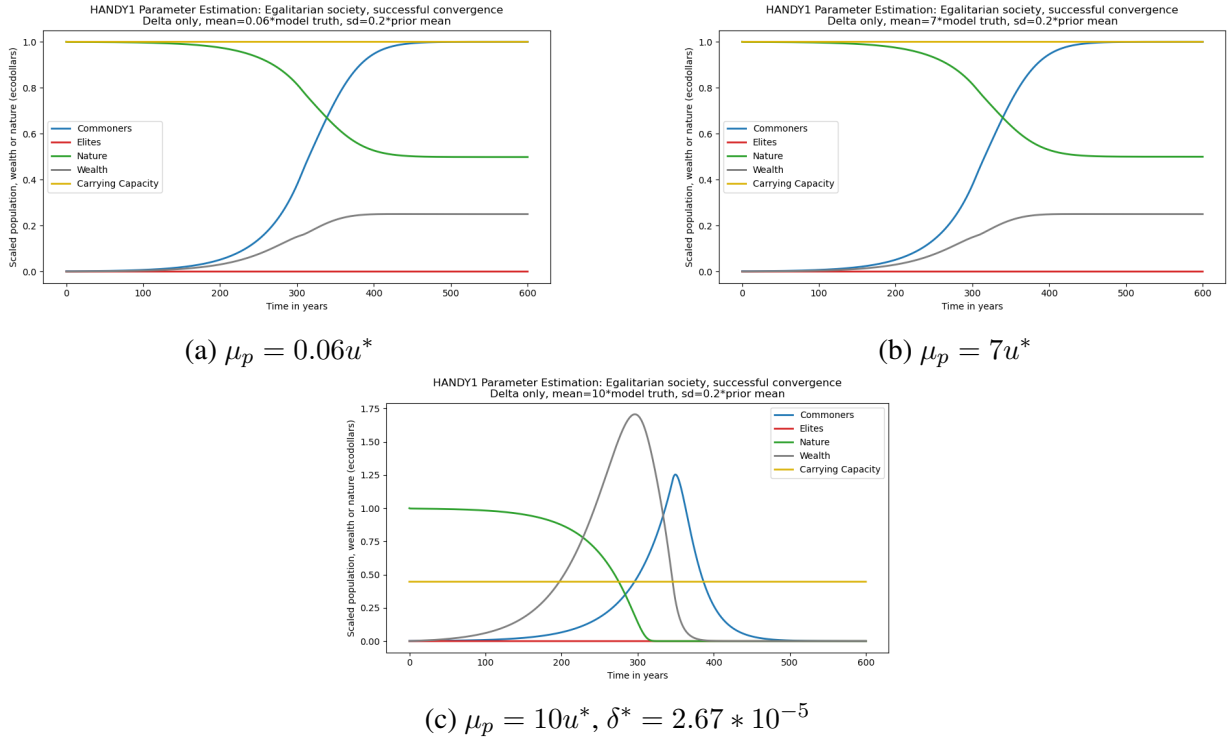


Figure 5.1: Parameter estimation of δ in the HANDY1 Egalitarian scenario. The left two figures show estimation for true $\delta = 6.67 * 10^{-6}$, $\sigma_p = 0.2\mu_p$, and labeled μ_p . The right figure shows estimation for an oscillatory Rgalitarian run with true $\delta = 2.67 * 10^{-5}$ and $\sigma_p = 0.2\mu_p$. The plot legend is as in Figure 2.1.

truth u^* , convergence to the true parameter values occurs after about 60-80 iterations. The wall clock time is about 0.7 seconds. Parameter values approach the truth only approximately and not exactly—the final EnKI error is within an order of magnitude of 10^{11} . Some variables are learned more effectively than others, since the final α_{min} is within 1% of the true scenario value u^* , while the starting x_c is 25% different from the truth. The convergence result is shown in Figure 5.2a below, with learned parameters in Table 5.1.

Changing the prior mean somewhat reduces the effectiveness of convergence. With a prior mean of $2.3u^*$, or any value above $1.5u^*$, the estimation only partially converges after 100 iterations, with some model parameters such as δ being estimated at about $1.47 * 10^{-5}$, a value about

2.5 times the model truth. Other parameters like α_{min} continue to be learned accurately. This partial success produces a kink in the wealth curve seen in Figure 5.2b.

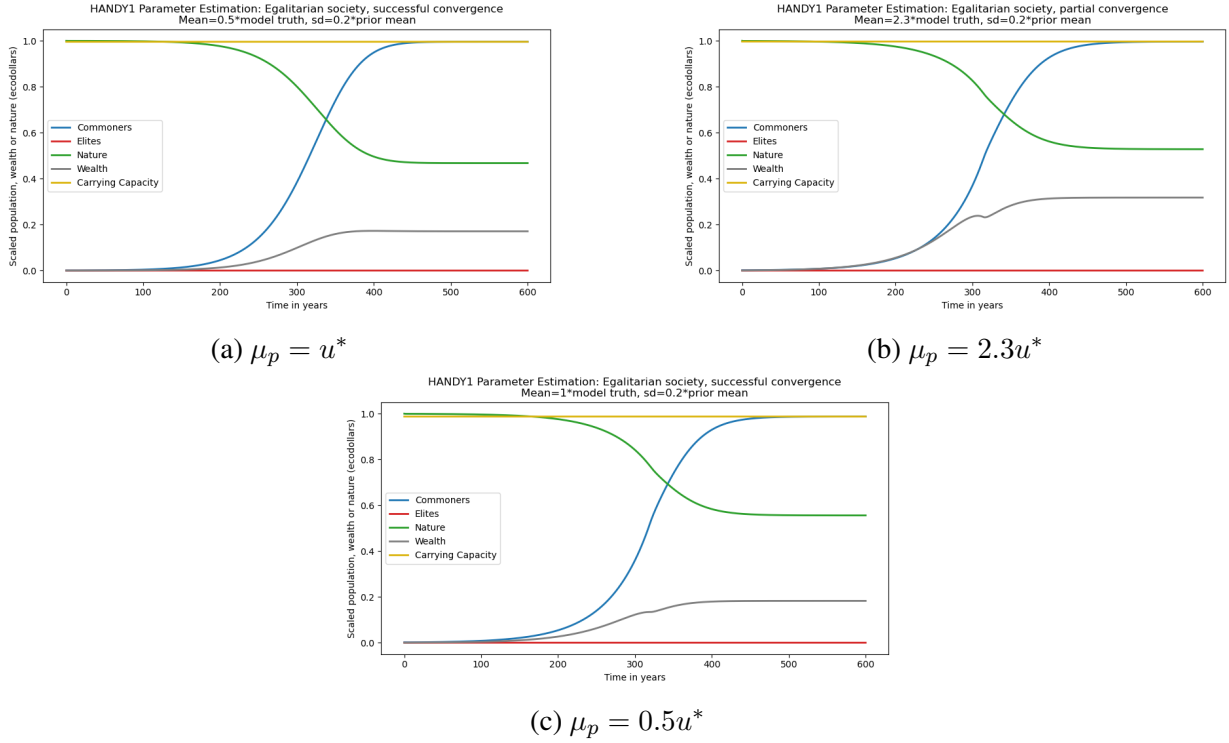


Figure 5.2: Multiple parameter estimation performance in the HANDY1 Egalitarian scenario. The left figure shows estimation for true $\delta = 6.67 \times 10^{-6}$, $\sigma_p = 0.2\mu_p$, and $\mu_p = u^*$. The right figure shows a partial estimation success with $\mu_p = 2.3u^*$. The bottom figure shows a successful estimation with $\mu_p = 0.5u^*$. The legend is as in Figure 2.1.

Convergence is also somewhat less good when the prior mean is reduced. With $\mu_p = 0.5u^*$, the posterior particle mean for most parameters is somewhat lower than u^* (see Table 5.1). In this case, α_{min} is learned less effectively with a final value of only about 0.005 per year, and the starting value of x_C is estimated at only 37.7. The EnKI error stabilizes at about 4×10^{11} . Nevertheless, the scenario still shows soft-landing to equilibrium, just as in the original HANDY1 run (Figure 2.1a, and Motesharrei et al. [2014]). This can be seen in Figure 5.2c below; it appears

as if many parameter values u_n can give rise to similar model dynamics.

The HANDY1 model gradient is well behaved for parameter values close to the model truth, but poorly behaved for extreme values. This is particularly true in the case of the maximum death rate parameters α_{max} and the depletion δ , which have a pathological effect on model state if they get too small or large. As mentioned before, the death rates must satisfy $\alpha_{min} < \beta_C, \beta_E < \alpha_{max}$ in order for HANDY1 to be stable. Larger values of α_{min} cause immediate population collapse, while smaller values of α_{max} lead to unstoppable exponential population growth. These parameter values cause HANDY1 to break down and produce extremely large, zero, or even spurious negative results for other state variables, leading to particle convergence failure. However, using noise in the EnKI pseudo-observations makes this issue less severe than it would be otherwise.

Table 5.1: List of model truth values u^* and posterior mean parameter estimates for HANDY1 Egalitarian runs, with varying prior means μ_p . All values to 3 significant figures.

Variable	Name in code	u^*	$\mu_p = 0.5u^*$	$\mu_p = 1u^*$	$\mu_p = 2.3u^*$
α_{min}	alpha_min	0.01	0.00511	0.0110	0.0106
α_{max}	alpha_max	0.07	0.0643	0.0778	0.0674
γ	gamma	0.01	0.00665	0.00976	0.0239
λ	lambda	100	105	103	98.9
$x_{C,start}$	startcommoners	100	37.7	124	102
δ	delta	$6.67 * 10^{-6}$	$4.96 * 10^{-6}$	$5.91 * 10^{-6}$	$1.47 * 10^{-5}$
s	s	0.0005	0.000421	0.000472	0.00117
ρ	rho	0.005	0.00413	0.00348	0.00630

5.2.2 Equitable Scenarios

In the equitable scenarios of HANDY1 (Figure 2.2, and section 5.2 in [Motesharrei et al. \[2014\]](#)), there is a pool of non-working 'Elites' x_E with inequality $\kappa = 1$ who take up resources but otherwise do not contribute to the society. With the exception of that population increasing the

equilibrium depletion δ , these runs are otherwise identical to the Egalitarian scenarios. It would thus not be surprising if the results of Equitable scenario parameter estimation were largely the same as in the Egalitarian case.

However, it becomes clear from our results that multi-parameter estimation functions more effectively in HANDY1 runs with a single cycle of growth and collapse than in oscillatory runs. In the latter case, the very sharp rises in y and w at multiple times degrade the EnKI stability of HANDY1 and lead to poorer convergence. We thus estimate parameters specifically for the full-collapse case (Figure 2.2d and section 5.2.4 in Motesharrei et al. [2014]) with $\delta = 4.33 * 10^{-5}$ and $x_{E,start} = 25$. We keep the other parameter priors the same as in the Egalitarian scenario, and run estimations for δ in isolation and for the same eight parameters as before. Starting x_E cannot be reliably learned in an Equitable run, as elite population x_E is always a constant multiple of commoner population x_C . Including x_E in the list of outputs y^* would lead to EnKI failure due to the ensemble covariance becoming singular.

When estimating δ only, the domain of convergence for μ_p is much broader than in the Egalitarian scenario. δ can be reliably estimated after 10-20 iterations, in 0.1 to 1 seconds of wall-clock time, with starting values of the prior mean ranging from $0.01u^*$ to over $400u^*$.

In the estimation of several parameters, the results are slightly better than in the Egalitarian scenario. With $\mu_p = u^*$ and $\sigma_p = 0.2\mu_p$, parameter convergence is slow but reliable, halting after about 40 iterations and 0.8 seconds. The EnKI error is about $1.4 * 10^8$ after 100 iterations, and all eight parameters are estimated to within an accuracy of better than 10%. As previously, the best convergence results were found for α_{min} , accurate to within 1%, and the worst for $x_{C,start}$ and δ . The radius of convergence is also improved—convergence occurs with $0.2u^* \leq \mu_p \leq 1.7u^*$, and the final EnKI error estimate is consistent within that region. For parameter priors $\mu_p > 1.7u^*$, the

EnKI process fails to converge entirely as α_{min} begins to exceed β_C . We additionally hypothesize convergence fails due to the highly skewed tails of the log-normal prior distributions. Figure 5.3 shows the final HANDY1 model plots at various μ_p .




Figure 5.3: Multiple parameter estimation successes in the HANDY1 Equitable scenario. All figures show estimation for true $\delta = 4.33 * 10^{-5}$ and $\sigma_p = 0.2\mu_p$, with μ_p as shown. The legend is as in 2.1.

Removing the α_{min} and α_{max} parameters from the estimation slightly improves the domain of convergence for μ_p , as it becomes approximately $0.15u^* \leq \mu_p \leq 2.5u^*$. This is broadly consistent with our hypothesis above and observations from tests done in the Egalitarian scenario.

Overall, Equitable HANDY1 runs with type-N collapses are handled more gracefully by our parameter estimation code than the soft-landing Egalitarian runs tested in Subsection 5.2.1. There are several possible reasons why this may be the case. In particular, a small change in model variables greatly affects the timing of population decline in type-N collapses in a way that does not occur with other scenario dynamics, spreading the estimation gradient across a wider area of parameter space. This is beneficial for the use of EnKI in HANDY2 (see Chapter 6), as its runs are largely also type-N collapses with one peak in population, wealth, and energy use in a run lasting several centuries.

5.2.3 Unequal Scenarios

5.2.3.1 Estimating δ

In this section, we repeat the single-parameter estimation of δ for the Unequal scenarios of HANDY, that contain Elites with an inequality of $\kappa = 100$. One examined scenario shows a type-N collapse (Figure 2.3b and section 5.3.2 in Motesharrei et al. [2014]), with high $\delta = 10^{-4}$ and $x_{E,start} = 0.2$. Another shows a type-L collapse (Figure 2.3a and Section 5.3.1) with $\delta = 6.67 * 10^{-6}$ and $x_{E,start} = 10^{-3}$. The type-L collapse scenario is uniquely run for 1000 years to capture the full collapse cycle. As we now introduce true Elites to the society whose population is no longer a fixed multiple of the Commoner population, we restore x_E to the output observations used for EnKI. The other parameter priors are the same as in the Egalitarian run.

The addition of Elites does not appreciably affect the wide domain of convergence for the type-N collapse, but makes the rate of convergence somewhat slower. EnKI takes about 40-60 iterations to reach a value of δ at most 10% from the model truth, with an error estimate near $1.5 * 10^8$. The final value of u^* does not depend on the starting μ_p , which shows convergence on a domain of $0.14u^* \leq \mu_p \leq 200u^*$ (Figure 5.4a).

For the type-L collapse scenario, which has a much lower true value of δ and a much higher peak of Elites late in the run, the parameter estimation is far less stable. While the EnKI ensemble takes only 20 iterations to converge to u^* for a wide range of input parameter values, the domain of convergence is substantially smaller at $0.08u^* \leq \mu_p \leq 8u^*$. Of note is that high values of starting δ , above $8u^* = 5.33 * 10^{-5}$, fail to converge at all (Figure 5.4b). This is peculiar, since convergence does occur for the type-N collapse scenario with a starting δ two orders of magnitude

greater at $2 * 10^{-2}$. It is likely that the longer integration period and flat dynamics in the type-L scenario do not provide enough feedback for the parameter estimation to work well with extreme values of δ . This emphasizes the importance of choosing the right priors, observations, and time domain for HANDY in EnKI experiments.




Figure 5.4: Parameter estimation of δ in HANDY1 Unequal scenarios. The left figure shows successful estimation for a type-N collapse scenario with $\mu_p = 10u^*$. The right figure shows unsuccessful estimation for a type-L collapse scenario with $\mu_p = 10u^*$. The legend is as in Figure 2.1.

5.2.3.2 Multi-parameter Estimation

The estimation of multiple parameters in Unequal HANDY1 societies is done along similar lines to the Egalitarian and Equitable scenarios. As we now have an estimation input x_E derived from the Elite population, we introduce two new parameters to determine: the inequality κ , and the starting number of Elites $x_{E,start}$. This gives a total of ten parameters, still excluding the birth rates β_C and β_E . As with the other parameter priors, we specify a support on the upper half-plane $[0, \infty)$. σ_p remains at $0.2\mu_p$ for this series of runs.

For the type-N collapse scenario, convergence is successful only in a small range of prior means satisfying $0.6u^* \leq \mu_p \leq 1.5u^*$. In this range, each parameter is estimated to within 5% of

the model truth u^* except for κ , which often deviates by up to 50% from u^* (Table 5.2; Figures 5.5a and 5.5b). For higher values of μ_p up to $1.9u^*$, the parameter estimation converges only partially to the truth, with over 100% estimation errors in many parameters and an EnKI error of $2.07 * 10^{10}$. Such a run can be seen in Figure 5.5c. We hypothesize that κ tends to be poorly estimated once it exceeds a value of 10, since it does not greatly affect the value of *raw* Elite population x_E . A similar proportion of wealth w is taken up by Elites at $\kappa = 100$ as at $\kappa = 10$ or even $\kappa = 200$, so any minor change in w does not affect the timing of Elite or Commoner collapse and therefore does not cause the higher EnKI error to disturb parameter estimation convergence.

Removing the estimation of κ predictably reduces the EnKI error for marginal values of μ_p , extending the radius of convergence at the low end to about $\mu_p = 0.3u^*$. However, this change has little to no effect at high values of μ_p and is not pursued further.

Table 5.2: List of model truth values u^* and posterior mean parameter estimates for HANDY1 Unequal type-N collapses, with varying prior means μ_p . All values are given to 3 significant figures.

Variable	Name in code	u^*	$\mu_p = 0.6u^*$	$\mu_p = 1.2u^*$	$\mu_p = 1.9u^*$
α_{min}	alpha_min	0.01	0.0100	0.0100	0.0128
α_{max}	alpha_max	0.07	0.0702	0.0707	0.0721
γ	gamma	0.01	0.0107	0.0101	0.0181
λ	lambda	100	96.6	98.3	77.4
$x_{C,start}$	startcommoners	100	102	100.2	241
δ	delta	$1 * 10^{-4}$	$9.47 * 10^{-5}$	$1.09 * 10^{-4}$	$5.18 * 10^{-5}$
s	s	0.0005	0.000447	0.000451	0.000633
ρ	rho	0.005	0.0469	0.00477	0.00694
$x_{E,start}$	startelites	0.2	0.208	0.195	0.276
κ	kappa	100	161	147	50.3

For multiple parameter estimation in the type-L collapse scenario, we adjust the run length to 1000 years as in the single-parameter case. Surprisingly, the resulting model inversion is

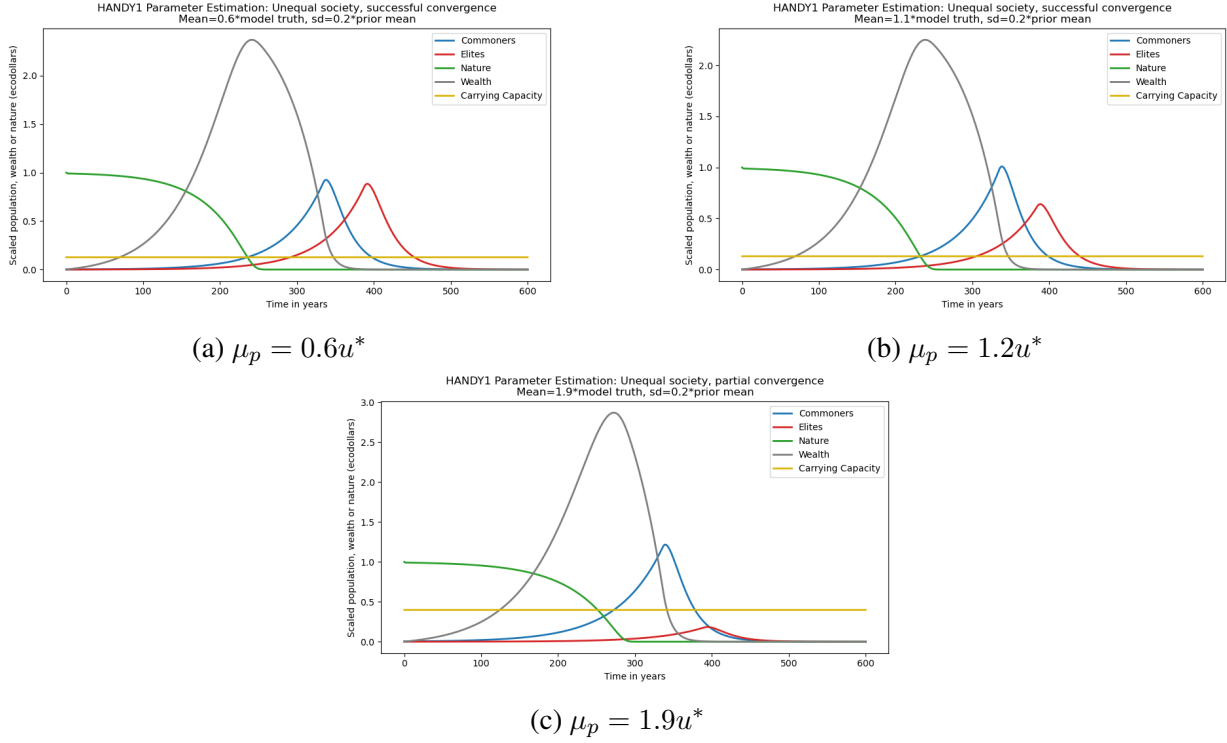


Figure 5.5: Multiple parameter estimation successes and failures for a type-N collapse scenario in HANDY1 with inequality. All figures show estimation for true $\delta = 10^{-4}$. The top two figures show mostly successful estimations for low μ_p as shown. The bottom figure shows the effects of a partial estimation failure at $\mu_p = 1.9u^*$. The legend is as in Figure 2.1.

slightly *less* fragile, but also less effective than in the type-N collapse. The radius of convergence for prior means is now $0.4u^* \leq \mu_p \leq 3.5u^*$, and almost all of the values of μ_p within that interval display only partial convergence. EnKI error remains 10^{11} or above in all runs, and the convergence process for even $\mu_p = u^*$ takes the entirety of the $N_{ens} = 100$ allowed estimation iterations and over 1 second of real time. Some of the parameters, including α_{max} , are estimated with reasonable accuracy (Table 5.3), while others like δ or γ retain errors of over 100% at the end of all examined runs (Figure 5.6). In particular, the estimates of κ in the type-L collapse scenario tend to be more accurate than in the type-N collapses, with inversion errors generally

less than 20% for most runs. This is likely because small variations in κ can advance or delay the timing of the sharp peak in Elite consumption, altering the estimation fit to the nature run greatly in a linearizable manner.

Table 5.3: List of model truth values u^* and posterior mean parameter estimates for HANDY1 Unequal type-L collapses, with varying prior means μ_p . All values are given to 3 significant figures.

Variable	Name in code	u^*	$\mu_p = 0.4u^*$	$\mu_p = 1u^*$	$\mu_p = 3.5u^*$
α_{min}	alpha_min	0.01	0.00844	0.00987	0.0121
α_{max}	alpha_max	0.07	0.0679	0.0670	0.0716
γ	gamma	0.01	0.00660	0.00947	0.0260
λ	lambda	100	101	101	98.1
$x_{C,start}$	startcommoners	100	72.8	87.6	163
δ	delta	$6.67 * 10^{-6}$	$4.63 * 10^{-6}$	$6.25 * 10^{-6}$	$1.46 * 10^{-5}$
s	s	0.0005	0.000354	0.000494	0.00115
ρ	rho	0.005	0.00397	0.00353	0.00594
$x_{E,start}$	startelites	0.001	0.000351	0.00119	0.00497
κ	kappa	100	115	80.8	66.3

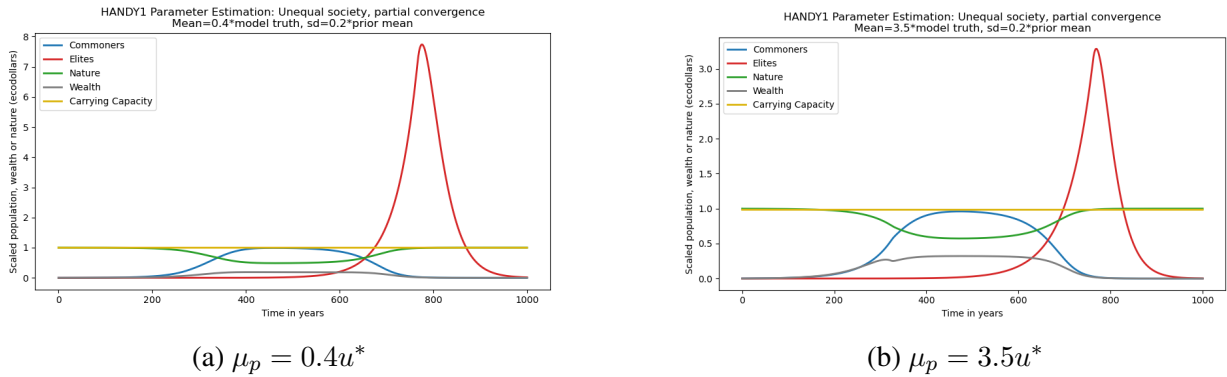


Figure 5.6: Multiple parameter estimation for a type-L collapse scenario in HANDY1 with inequality. All figures show estimation for true $\delta = 6.67 * 10^{-6}$ and $x_{E,start} = 0.001$. Convergence was partially successful at $\mu_p = 0.4u^*$ (left) and $\mu_p = 3.5u^*$ (right), and at all values in between, with an EnKI error of over 10^{11} in both runs. The legend is as in Figure 2.1.

In both single-parameter and multi-parameter estimation, the EnKI process seems to experience difficulty processing the long run time and high spike of Elite population seen in the reference type-L collapse. But the strong signal of Elite collapse timing better conditions the

multiple parameter estimation, as critical collapse parameters like κ , $x_{E,start}$, and α_{max} are allowed to vary among the ensemble particles. We mercifully encounter less estimation fragility with the predominantly type-N and mixed type-N/L collapses seen in HANDY2 parameter estimation.


5.2.4 Soft Landings in an Unequal Society

For a final experiment, we estimate the parameters of the soft-landing Unequal scenario of HANDY1 (Figure 2.3c; section 5.3.3 in Motesharrei et al. [2014]), starting with a prior mean equal to the type-N Unequal scenario truth. This soft-landing run differs significantly in its starting conditions and evolution from the other scenarios we have examined above. In particular, the birth rates β_C and β_E differ from each other and from their usual values of 0.03 per year, and the starting populations of Elites and Commoners are much larger. For this reason, we add β_C and β_E to the parameter estimation, using a log-normal prior for each with support on the non-negative reals $[0, \infty)$. We keep $\sigma_p = 0.2\mu_p$. To speed convergence, we reduce the `noise_level` parameter η to 0.002, set the `EKSStableScheduler` parameters to $\alpha = 20$ and $\epsilon = 10^{-3}$, and extend the run time to $N_{iter} = 500$ iterations. The other parameter priors are identical to those used for the type-N collapse scenario.

The estimation run with all 12 parameters of HANDY1 is largely successful. The posterior estimate slowly converges to the soft-landing scenario by 500 iterations and 3 seconds of clock time, with a final EnKI error of $6.62 * 10^{11}$ and an output model that reaches a stable equilibrium as in the nature run (Figure 5.7a). But the posterior parameter estimates differ somewhat from the truth. Parameters like λ which start close to u^* remain close with up to 10% error, while

other parameters like the birth rate are estimated as $\beta_C = 0.0553$ and $\beta_E = 0.0365$ per year. These values are somewhat divergent from the respective true values of 0.065 and 0.02 per year. Convergence is sometimes impeded by a numerical instability present in the Nature stock y which disrupts the learning process of EnKI. At other times, this has no effect. Unlike in HANDY2, reducing the time step does not solve the issue.

Restricting parameter estimation to subsets of the HANDY1 parameters is even more effective. Estimating β_C and β_E alone is completely successful when setting a prior for each bounded on $[0, 0.07 = \alpha_{max}]$, and ineffective otherwise. Both parameters are learned within 50 EnKI iterations to within 0.01% of their true values (Figure 5.7b). Estimating the ten parameters which are *not* the birth rates is also somewhat effective, but convergence often depends on the random initial ensemble state and only partially approaches the scenario truth.



(a) All parameters

(b) β_C, β_E only

Figure 5.7: Successful estimation of birth rates β_C and β_E in the HANDY2 Unequal soft-landing scenario. The left subfigure shows EnKI estimation with all 12 HANDY1 parameters, starting from a prior mean μ_p with true parameter values for a type-N collapse run. The right subfigure shows estimation restricted to β_C and β_E alone, with other parameters fixed at their true values. The legend is as in Figure 2.1.

The above results show that reliable parameter estimation of birth rates is possible but difficult in HANDY1. In particular, issues exist for scenarios with unusual values of β_C and β_E .

In order to improve convergence, stability-improving modifications to the model structure and better extraction of observations may be needed. The difficulty of estimating birth rates is one of several reasons why we did not change β_C or β_E in the HANDY2 parameter estimation, instead treating it as a policy measure. We instead opted to use other mechanisms, such as class mobility, to constrain Commoner and Elite population growth.

5.3 Conclusion

The EnKI parameter estimation method described in [Iglesias et al. \[2013\]](#) and [Dunbar et al. \[2022\]](#) has been adapted successfully for HANDY1. Estimation of certain parameters that can vary greatly between scenarios and have well-defined behavior, like δ , is fast and effective across a wide range of initial conditions and prior means. Yet the estimation of most or all parameters within a scenario is somewhat slower and more fragile. Certain dynamics in HANDY1 give rise to unstable behavior when parameters are set to unusual values, easily causing particle divergence to zero-population attractors or causing Regenerating nature y_N to briefly go negative. This generally occurs for prior means μ_p differing by a factor of more than about two from the model truth u^* . Some parameters like the birth rates β_C, β_E are difficult to estimate, as each particle must satisfy $\alpha_{min} < \beta < \alpha_{max}$ to maintain stability. But these can still be determined effectively in certain circumstances, as with the Unequal soft-landing scenario shown in [Figure 2.3c](#). Issues with the stability of EnKI are also worse for certain types of oscillatory Equitable scenarios or type-L Elite-driven collapses. But the method remains largely robust in most cases, and is even more stable for the most common HANDY1 scenarios that lead to either equilibrium soft landings or type-N collapse.

The occasional instability of parameter estimation in HANDY1 raises questions about its effectiveness in the more complex HANDY2 model. However, we believe these problems do not impact the use case of HANDY2. In particular, good starting estimates for the prior means μ_p of HANDY2 parameter values already exist in the form of the Chapter 3 manually-optimized scenarios. These help us confirm that the EnKI inversion results do not generally differ much from our expectations. We also perform HANDY2 parameter estimation on scenarios with a dramatic upward trend in population and resource use from 1800 to 2020—a pure type-N collapse situation where EnKI performs best. And the more noisy nature of real-world observations compared to artificial pseudo-observations adds robustness by preventing the inversion from becoming easily trapped by model instabilities [Iglesias et al., 2013]. Finally, we will perform multiple estimation runs with different pseudorandom seeds, giving a broader picture of the HANDY2 parameter space and preventing a single ensemble failure thwarting the entire process. With these thoughts in mind, Chapter 6 describes the EnKI implementation for HANDY2.

Chapter 6: HANDY2 Parameter Estimation

6.1 Overview

In chapter 3, we presented several scenarios of HANDY2 which successfully captured the growth in population, energy use, and GDP per capita from 1800 to 2020, and outlined several scenarios for the future direction of human society to 2200 that depend on the actions taken by global policymakers. However, HANDY2 has quite a few free parameters. Tuning their values by hand does not make it clear which model configuration is best supported by the historical record. The wide range of model outcomes also makes it difficult to know the overall probability of worrisome outcomes for humanity in the baseline or increased-renewables scenarios. It is clear that a better method of exploring the large parameter space would be beneficial.

We therefore introduce EnKI parameter estimation to HANDY2 to solve these problems. Unlike in HANDY1, we estimate parameters based on the distance to real-world observations of Renewable, Nonrenewable, and Regenerating energy use, and population, from 1800 to 2020. By starting with parameter priors from existing scenarios already close to true observations, we can track their convergence to develop better estimates for tunable parameters we have had difficulty estimating in practice, such as μ_M , s , or ρ . We evaluate the impact of these choices on the overall evolution of the HANDY2 scenarios to 2200, and discuss them in the context of economics and history. We also examine the limitations of the parameter estimation approach for HANDY2,

including the lack of accurate observations and concerns with convergence instability, and show where the EnKI process may work better with externally determined constraints.

Finally, we run the baseline HANDY2 scenario many times to build an ensemble of 50 different posterior estimates of the critical model variables. We run this ensemble to 2200, and change its policy parameters analogously to the other HANDY2 scenarios discussed in section 3.5, in order to more robustly examine the possibility for socially optimal outcomes given the three key policy measures of encouraging Renewable energy, decreasing economic inequality and lowering birth rates discussed there. The results from this superensemble confirm the analysis of the manually-tuned scenarios—combining all three changes almost always prevents a type-N resource collapse.

6.2 Adapting EnKI for HANDY2

To perform parameter estimation on all or most parameters of the HANDY2 model, we use the `handy2_parameter_estimation.jl` code adapted from the HANDY1 EnKI estimation in Chapter 5. The preamble of the code begins by calling the authoritative Julia copy of HANDY2, contained in a function that takes in the parameters and variables used in the model (Table 6.1) and returns a `HANDY2Output` structure with the state variables of the model at each timestep.

The EnKI estimation process functions similarly to the one presented in section 5.2, and we present a summary below highlighting key differences where applicable. First, the initial-guess scenario is defined, used for display purposes and to state the values of fixed HANDY2 parameters. Presently, it is identical to the baseline Scenario 1 defined in subsection 3.5.1, whose

parameter values are given in Table 3.2. Various parameter estimation settings are also defined here, including the RNG seed for the model run, the names of the output variables y^* to be learned, the number of ensemble particles N_{ens} and iterations N_{iter} to use (usually 50 and 500 for this application), and the choice of learning rate scheduler.

We then specify each variable to be learned and construct prior distributions for those variables. Ideally, we would seek to learn every variable in HANDY2 that may be tunable or weakly empirically constrained, while excluding those variables that represent deliberate policy changes, often altered in the non-baseline HANDY2 scenarios. We are able to use most such parameters in the estimation. However, there are some we must exclude because they cause the estimation to become unstable or lead to particle divergence. Specific details are in subsection 6.2.1 below.

We next extract the true values of output state variables y^* from the real-world datasets of population and energy use detailed in Chapter 3.4. These data from the [United Nations Population Division \[2023\]](#), [Klein Goldewijk et al. \[2017\]](#), [Smil \[2017\]](#), and the [Energy Institute \[2023\]](#) are read from the same CSV file used to produce the values in HANDY2 figures, and provide values for the five inversion output parameters x_C , x_E , y_R , y_N , and y_F . Unlike in the HANDY1 parameter estimation, there is no external estimate of wealth w that can be used to estimate the HANDY2 output. Some post-processing is done on the observations to make them comparable to the HANDY variables. x_C is divided by 10^4 to make the baseline error for population comparable in size to the other errors, and is thus expressed in units of 10^7 people. y_N is multiplied by ψ_N to obtain the extractable fossil fuel consumption, for comparison with HANDY2, and y_F is multiplied by ψ_F . This is not done for ψ_R , so as to maintain consistency with the manual estimation used for the baseline scenario. Finally, x_E is an artificial output variable constructed by multiplying the commoner population x_C by $1/\kappa$, in order to better estimate Elite-related

parameters; we use that value to approximate the historical number of very rich Elites being about 1% of the total world population. (See discussion in subsection 3.2.5.)

The outputs y_{init} from the initial guess scenario, with population x_C, x_E divided by 10 for error normalization, are used to construct the baseline loss $\|y_{init} - y^*\|_2$, individually computed for each output variable. This is used for display purposes only and as a diagnostic tool to compare the baseline scenario estimation error to results of the EnKI estimation process.

In the HANDY1 estimation, a covariance matrix Γ was constructed from a nature run. We do not have model truth when estimating variables from real-world data. We therefore instead use the initial ensemble loss to create the covariance. After constructing N_{ens} ensemble particles from the already-defined parameter priors using `construct_initial_ensemble` and `transform_unconstrained_to_constrained` to sample the prior space, we simply run HANDY2 once for each particle to get outputs y_0 , and calculate the error $\|y_0 - y^*\|_2$ for each particle and parameter. Γ is then calculated straightforwardly from this list of errors.

As before, the EnKI iteration begins by constructing the `EnsembleKalmanProcess` object with the initial ensemble, the zero 'truth vector' used to do error minimization, the covariance Γ , `Inversion()`, the RNG seed, the learning rate scheduler, and the `SampleSuccessGauss` failure handler. As in HANDY1 EnKI, the `EKSStableScheduler` [Kovachki and Stuart, 2019] is used to dynamically adjust the learning rate throughout the run. This is extremely beneficial for maintaining particle stability throughout the inversion process, by keeping a low learning rate until the parameter estimation gets reasonably close to the observations. Significant experimentation was required to determine the optimal values of both stable scheduler parameters; we determined them to be 100 for the numerator α and 0.1 for the kernel ϵ .

The EnKI iterations largely proceed the same way as for HANDY1. A number `N_ens`

of test particles $u_{z,i}$ are drawn from the parameter distribution at each iteration, then are run through the HANDY2 model to produce an output mean square loss $(g_{mean})^2 = \|y_{z,i} - y^*\|_2^2$ for each output variable and an EnKI error $\|y_z - y^*\|_2$ over all output variables. The ensemble is then updated using `update_ensemble!` with the error vector. In addition, several summary statistics are printed. We display the EnKI error of the ensemble at each iteration as well as the RMS error g_{mean} , which can be compared to the baseline scenario error estimate to track the progress of the parameter estimation. The EnKI updates stop after `N_iter` iterations are complete, or when the EnKI error stops changing. This is defined as a difference in relative error in successive iterations below the threshold value (here 10^{-3}), with an absolute error below another user-defined value (here $2 * 10^3$). As convergence tends to be quite slow for HANDY2, we usually set N_{iter} to 500 iterations.

At the end of the EnKI process, we obtain the posterior mean estimate of the parameters u_i with `get_phi_mean_final`, then save them to an output file alongside the initial estimate u_{init} . These parameters are then used to run a final pass of HANDY2 that produces the posterior mean output y_i , used in the HANDY2 plotting code to create the usual stocks, flows, and birth rate plots as shown in Section 3.5.

To examine a broader parameter space during the estimation process, we run the EnKI code with many different random seeds, producing different posterior estimates and HANDY2 scenarios on each run. The estimation attempts that fail to converge to the observed truth, around 25-30% of all starting conditions, are discarded, and the remaining output scenarios obtained from the lowest positive integer seeds are collected in a model ensemble of 50 members. We use this model ensemble to rerun the six HANDY2 scenarios from section 3.5, modifying the policy parameters altered in each scenario in order to examine their robustness under various

initial conditions.

The output of the HANDY2 parameter estimation process, discussion about the parameter values obtained from it, and the ensemble runs for the Chapter 3 scenarios are presented in the next section of this chapter.

6.2.1 Learned Parameters and Starting Priors

Here, we list the 22 parameters learned in HANDY2 parameter estimation, and their prior distributions as implemented in the code, going in the order presented in Table 3.2. Of the starting state variables we estimate only x_E . This is the one starting value which is neither strongly constrained by observations, nor irrelevant to the later model evolution. The Elite population is only weakly constrained in 1650 as we have few good records or definitions of how many people were part of the world's rich at this time. x_E also strongly affects the later number of Elites throughout the run. x_C , y_R , and y_N are well-characterized; we extensively discuss the reasons these parameters are set at their values in Section 3.4. Starting w , A_R , A_N , and A_F do not affect the model. The starting accumulated investment values are just placeholders used to set the initial labor force allocation, and are not used elsewhere. The initial wealth w is set at zero to prevent early sharp spikes in wealth-related parameters that can destabilize HANDY2. In both cases, parameter estimation would not be particularly useful.

Of the empirically-constrained parameters, we estimate α_{min} , α_{max} , κ_{start} , γ , and λ_{start} . These are all parameters which have at best weak empirical constraints and which would be quite useful to characterize with EnKI. Those variables that are excluded are β_C , β_E , F_{start} , the production efficiencies ψ_N , ψ_F , and parameters related to ψ_R . The birth rates suffer the same

estimation limitations as with HANDY1. It is difficult to constrain them within the minimum and maximum death rates while learning both at the same time and retaining particle stability. We also use the birth rates in HANDY2 as part of a birth rate policy knob. It is thus clearer to maintain both values at 0.03 per year to ensure that birth rate policies make sense in the EnKI ensembles. The latter justification also applies to F_{start} , which couples closely to the Renewable policy knobs. The production efficiencies ψ_N and ψ_F are externally specified to agree with the values from [Energy Institute \[2023\]](#). Finally, changing the values of the ψ_R parameters causes particle stability issues; they are also kept outside the estimation to be consistent with the other ψ values.

Of the remaining tunable parameters, all are included in the estimation with the exception of A_{Rth} . These are the values that have no direct external constraints, and where parameter estimation has the highest benefit. In many cases, the values used in the baseline scenario are based on those inherited from HANDY1, educated guesses, or manually tuned—the automated EnKI process makes this work much more efficient. A_{Rth} was excluded because of its extremely wide range and high baseline value of 10^7 ; when introduced to the estimation, it often quickly diverges to extreme values because of its negligible effect on the HANDY2 model output. For this reason it can safely be excluded.

The model settings are kept the same as in Chapter 3 except for the timestep, which is reduced to 0.1 years. This helps eliminate a numerical instability present when integrating HANDY with certain unusual parameter combinations. No HANDY2 settings are learned in the estimation.

Of the 22 parameters, the prior means μ_p are set identically to the values of the baseline scenario, to ensure the particles are close to a 'known good' estimate of the parameter truth iden-

tified by hand, and to avoid the issues with particle collapse outside of small parameter domains identified in Chapter 5. Likewise, the prior standard deviations are set to values approximately one third to one half of the parameter means. Some σ_p values for parameters with small numerical values are on the high end of this range, which avoids a bug in the CLiMA libraries that prevents constrained parameter distributions from being initialized. These values represent a good compromise between model novelty and adherence to the baseline scenario, and are more stable than those used in Chapter 5.

The lower and upper limits for each parameter prior are also set by hand. Lower limits for most priors are zero, as all parameters are assumed to be non-negative real numbers in the design of HANDY2; this helps prevent particle collapse. The exception is α_{max} whose lower limit is 0.03 to prevent it from ever being lower than the birth rates β_C, β_E . Upper limits are set at reasonably high values that prevent the underlying particle distributions from undergoing excessive divergence. In most cases where the parameter priors can be assumed to be approximately normal, the upper limits are the common-sense bounds on the range: 1 in the case of variables that must lie on $[0, 1]$ such as ξ , 0.3 or 0.5 in the case of other small variables, and values about 5-10 times the parameter mean for larger-valued parameters. α_{min} has an upper bound of $\beta_{C,E} = 0.03$ to retain model stability. Variables such as $\delta_{R,N,F}$ or κ that are approximately lognormal have ∞ as an upper bound, which creates a log-normal prior distribution.

The full list of prior means, standard deviations, lower bounds, and upper bounds for all estimated variables is given in Table 6.1.

Table 6.1: List of parameters and constrained prior distributions used in HANDY2 parameter estimation. Definitions are as in Table 3.2.

Variable	Name in code	Prior mean	μ_p	Standard deviation	σ_p	Lower bound	Upper bound
x_E	startelites	5		2		0	50
α_m	alpha_min	0.005		0.003		0	0.03
α_M	alpha_max	0.07		0.02		0.03	0.3
κ_{start}	kappa_start	100		33		0	∞
γ	gamma	0.04		0.015		0	1
λ_{start}	lambda_start	$2.7 \cdot 10^4$		10^4		0	$6 \cdot 10^4$
δ_R	deltar	$1.3 \cdot 10^{-4}$		$4.5 \cdot 10^{-5}$		0	∞
δ_N	deltan	$2.3 \cdot 10^{-6}$		$8 \cdot 10^{-7}$		0	∞
δ_F	deltaf	$7 \cdot 10^{-4}$		$2.5 \cdot 10^{-4}$		0	∞
λ_{growth}	lambda_growth	$4 \cdot 10^{-4}$		$1.5 \cdot 10^{-4}$		0	0.05
s	s	0.005		0.002		0	0.2
ρ	rho	0.05		0.02		0	1
ω_μ	omega_mu	1.5		0.5		0	30
μ_M	mu_M	0.06		0.02		0	1
A_{Nth}	Anth	315		110		0	1000
A_{Fth}	Afth	12		4		0	100
$\xi_{m,R}$	xim_r	0.1		0.04		0	1
$\xi_{m,N}$	xim_n	0.25		0.1		0	1
$\xi_{m,F}$	xim_f	0.05		0.02		0	1
ω_i	omega_i	0.01		0.004		0	0.5
J	J	0.01		0.004		0	0.5
σ	sigma	0.01		0.004		0	0.5

6.3 Results

6.3.1 Convergence Dynamics

EnKI parameter estimation works well to recover plausible values for the HANDY2 parameters that fit the observed data of population and energy use. When we ran the code for small positive RNG seeds between 1 and 69, we found that 50 of 69 runs, 72.5% of the total, display satisfactory convergence to the true observations u^* , as measured by the EnKI error being below $5 * 10^3$ after 500 iterations. The convergence is rather slow at first. Even the fastest-converging runs do not approach the truth until after 50-100 iterations, and some of the slowest that did successfully converge (e.g. seed 3) found the truth region less than 50 iterations from the end of the run. As shown in Figure 6.1, most particle states begin outside of the convergence region with an EnKI error of about 10^8 , slowly declining to 10^7 often with some irregularity. But once the truth is approached, the convergence is rapid. Error quickly drops to 10^3 within an order of magnitude, although it is often very erratic at first before the posterior estimate reaches a stable local minimum.

There is some variability among the ensembles that do not converge to the true observations. Most commonly, non-converging ensembles have EnKI error that stays near 10^7 to 10^8 , and simply fail to follow the weak model gradient toward the region of convergence. More rarely, a few ensembles have their parameter values diverge, causing the mean error to increase rapidly until total particle collapse. Another type of convergence instability happens in certain cases when the ensemble finds the convergence region—the high sensitivity of the estimation loss to timing in Nonrenewables and population growth can cause the covariance matrix update to be-




Figure 6.1: A plot of the EnKI error of 10 converging HANDY2 parameter estimation runs, with iteration number on the x-axis and EnKI error on the y-axis. Each color represents a different starting random seed. Convergence is stopped early if the error is below $2 * 10^3$ and does not relatively change by more than 0.1% per iteration.

come singular, or the ensemble update to overshoot into the high or moderate error region with a much lower learning rate. In one singular instance, convergence failed as the ensemble found a numerical instability in the HANDY2 Euler integration. This is much less likely with a low timestep of 0.1 yr than at the higher timestep used in the baseline scenario.

The greater than 70% success rate of convergence is remarkable considering the evident instability of HANDY1 parameter estimation in unusual starting conditions. This success is for two reasons. The introduction of real-world observations with more complex noise definitionally produces more realistic output dynamics, and a more robust EnKI update across multiple starting particle states. This reduces the risk of collapse in particle dimensionality that often happened with estimates in HANDY1. Additionally, we found that adjusting the learning rate over time with the `EKSStableScheduler` [Kovachki and Stuart, 2019] is even more effective in HANDY2 than in HANDY1. Without it, most particles end up diverging quickly from the observation truth due to excessive learning rate at low errors, leading to ensemble collapse in well over half of runs. While originally implemented for the related Ensemble Kalman Sampler (EKS) method [Garbuno-Inigo et al., 2020], dynamically changing the learning rate is extremely useful in many practical EnKI applications such as HANDY2 due to the dramatically different environment close and far from the true observations. It is even possible that the learning rate is not reduced enough near the observation truth, due to the evident inversion instability upon initial particle convergence. Further experiments will be required to establish if this is the case.

6.3.2 Examining Sample Runs

Posterior estimates from the successful parameter estimation runs produce HANDY2 scenarios that confirm many of the findings we have seen in Chapter 3. Figures 6.2, 6.3, and 6.4 show the stocks and flows plots respectively for runs with seed values 1, 2, and 3. One can immediately determine that the curves of Nonrenewable production and population agree far more closely with the real-world data than in the manually estimated baseline scenario (subsection 3.5.1; Figure 3.5). The Commoner population in each run follows a smooth upward trend, peaking in 2060 to 2080 at a value of about 11-12 billion. This is a notably higher value than in the [United Nations Population Division \[2023\]](#) moderate scenario, an indication that their projections may need to be revised upwards [[Motesharrei et al., 2016](#)]. Resource collapse still happens as in the baseline scenario, with the Commoner population dropping by half to 6.5 billion, and the Elite population following a very similar growth curve to the Commoners in accordance with the x_E parameter truth imposed on the run evolution. Note that the proportion of equivalent Elites varies greatly between runs, since κ_{start} is poorly constrained by the observations. The raw number of Elites x_E is correctly estimated at close to 0.01 times that of the Commoners x_C . Unlike in the baseline scenario, higher values of ω_μ generally prevent a large-scale rise in Elite populations during the peak of Commoner consumption C_C/x_C . This is likely more realistic in the far future.

In all three sample runs, the maximum of Nonrenewables use occurs in about 2030-2040, more accurate than in the baseline scenario which predicts a peak in 2020. The results still necessarily align with the rapid fossil fuel use reductions necessary to keep within the RCP4.5 concentration pathway [[Masson-Delmotte et al., 2021](#), [van Vuuren et al., 2011](#)]. Because of the later decline in Nonrenewables, these runs in general use more of the 62 ZJ of burnable carbon

y_N , leaving far less usable for future humanity. Renewable energy hits the cap of $F = 80$ EJ/yr at about the same rate in the EnKI runs as in the baseline, almost exactly in accordance with the measured values of energy use through 2022. In some runs, this occurs implausibly quickly in combination with much higher δ_F , leading to minor destabilization of the Commoner population, Elite population, and wealth curves in the time period of this jump. A few runs (not shown) demonstrate a much more gradual uptake of Renewable flows, and even a decline past 2100 as Commoner population starts to collapse.

These resource changes do not broadly affect the overall trend of the Commoner per capita consumption. The general trend is for C_C/x_C to rise from its pre-industrial value to reach a peak at or just after the year 2020, when the energy inputs from Nonrenewables and Renewables are greatest. A dip follows, causing the population collapse observed in most runs, and a slight recovery occurs after 2100 as population drops. In all, this is in line with the baseline scenario. However, the absolute value of C_C/x_C changes dramatically from run to run, governed by the values of s which are poorly constrained by the estimation output—some scenarios show an average consumption of 5-10 GJ/person-yr, while others may show values as low as 2 or as high as 15 GJ/person-yr. This seems to be mainly an issue of scaling rather than structural dynamics, as the effect of consumption on birth rates and elite attrition remains more or less the same between scenarios. External observations of per capita energy consumption will likely be necessary to further constrain s .

Finally, consumption of Regenerating energy in each run follows the estimates of [Smil \[2017\]](#) closely, with a slow ramp in use after the Industrial Revolution, a peak in the 2000-2020 period, another drop as the global economy and population partially collapses to 2100, followed by a slight increase at the end of the century. Nature stocks correspondingly decline, reaching


a minimum by 2050 to 2100 before recovering slightly by the end of the run. This is in line with historical estimates of global phytomass from Smil [2015], and constitutes a very plausible projection for the medium-term future.

Overall, the parameter estimation runs successfully capture the present-day state of the global population and economy in HANDY2, to a much better degree than the manual estimations done in Chapter 3. Their findings are still in accordance with our previous evaluations of the baseline scenario to 2200 and conventional wisdom. However, some constraints or additional data may need to be added to the EnKI y^* to account for the more unusual features of these runs.

6.3.3 Parameter Estimates


Table 6.2 shows the values of the estimated parameters in the three sample runs shown previously, as well as the arithmetic or geometric (for δ_R , δ_N , δ_F , and κ_{start}) means, and standard deviations of the parameters across all 50 estimation runs. What is notable is that each of the parameters often converge to very different values in each run. Some of the parameters, such as the starting wealth inequality κ_{start} , or the threshold for investment ω_i , are only weakly constrained by the observations and show a broad range of possible values, including the manually-estimated value used in HANDY2. Others such as λ_{start} have relatively strong constraints, and converge across scenarios to values similar to those we identified in Chapter 3, raising our confidence that they make sense in the context of the model dynamics. Still others such as δ_N converge to values consistently different to those of the base run, and it seems in those cases that the parameter-estimated values are more appropriate for HANDY2. Below we will discuss each of the parameters in detail.

HANDY2 Parameter Estimation Run: Population, wealth, and resource stocks



(a) Stocks plot


HANDY2 Parameter Estimation Run: Resource flows and population



(b) Flows plot

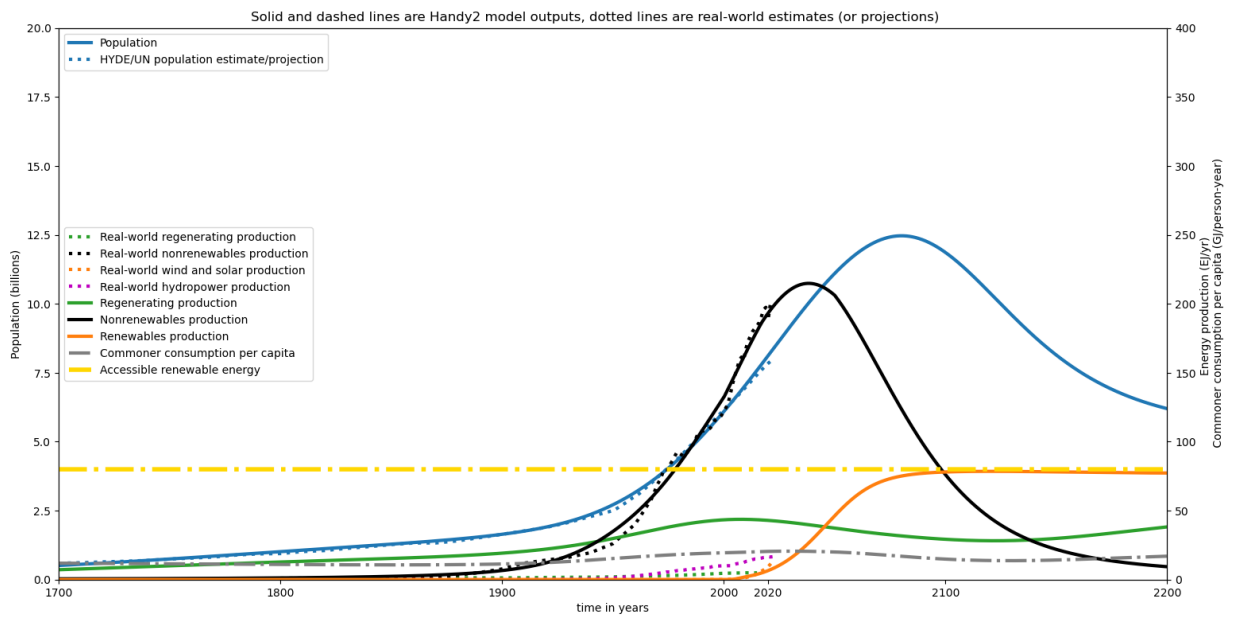
Figure 6.2: Plots of the results from a HANDY2 scenario obtained via EnKI parameter estimation, with random seed 1. The stocks plot is on the left, and the flows plot is on the right. Legends and labels are as in Figure 3.5.

HANDY2 Parameter Estimation Run: Population, wealth, and resource stocks



(a) Stocks plot

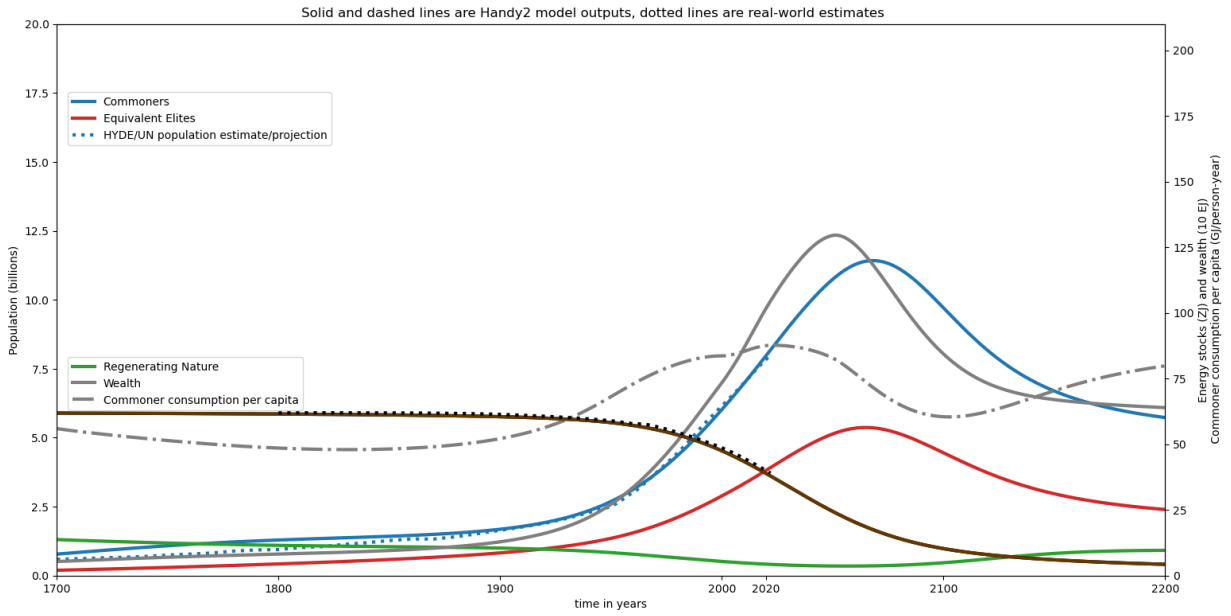
HANDY2 Parameter Estimation Run: Resource flows and population



(b) Flows plot


Figure 6.3: As in Figure 6.2 but with random seed 2.

HANDY2 Parameter Estimation Run: Population, wealth, and resource stocks



(a) Stocks plot

HANDY2 Parameter Estimation Run: Resource flows and population



(b) Flows plot

Figure 6.4: As in Figure 6.2 but with random seed 3.

Table 6.2: List of estimated parameter values for HANDY2 EnKI estimation runs 1-3, as well as mean values and standard deviations across all 50 successful runs. All values are to 3 significant figures. Definitions of parameters are in Table 3.1.

Variable	Name in code	Run 1 value	Run 2 value	Run 3 value	Mean value	Standard deviation
x_E	startelites	17.2	2.73	2.29	16.3	15.5
α_m	alpha_min	0.00316	0.00299	0.00232	0.00285	0.00173
α_M	alpha_max	0.115	0.0862	0.109	0.113	0.0435
κ_{start}	kappa_start	16.7	183	50.0	32.6	$9.57 * 10^4$
γ	gamma	0.0540	0.0991	0.0862	0.136	
λ_{start}	lambda_start	$2.95 * 10^4$	$2.57 * 10^4$	$1.64 * 10^4$	$2.72 * 10^4$	$1.08 * 10^4$
δ_R	deltar	$1.23 * 10^{-4}$	$5.29 * 10^{-5}$	$2.87 * 10^{-4}$	$1.48 * 10^{-4}$	$5.08 * 10^{-4}$
δ_N	deltan	$2.37 * 10^{-6}$	$2.80 * 10^{-6}$	$2.64 * 10^{-6}$	$2.57 * 10^{-6}$	$2.20 * 10^{-7}$
δ_F	deltaf	$8.04 * 10^{-4}$	$1.04 * 10^{-3}$	$1.09 * 10^{-3}$	$8.53 * 10^{-4}$	$1.04 * 10^{-3}$
λ_{growth}	lambda_growth	$9.05 * 10^{-6}$	$5.38 * 10^{-9}$	$4.17 * 10^{-8}$	$2.10 * 10^{-4}$	$5.60 * 10^{-4}$
s	s	0.00349	0.00116	0.00485	0.00511	0.00265
ρ	rho	0.105	0.0651	0.0478	0.0801	0.0551
ω_μ	omega_mu	1.89	1.96	2.08	2.13	0.829
μ_M	mu_M	0.0647	0.0378	0.0406	0.0591	0.0385
A_{Nth}	Anth	31.5	151	145	207	235
A_{Fth}	Afth	17.7	2.83	2.20	15.3	12.7
$\xi_{m,R}$	xim_r	0.132	0.252	0.0894	0.197	0.195
$\xi_{m,N}$	xim_n	0.204	0.167	0.247	0.252	0.0913
$\xi_{m,F}$	xim_f	0.0898	0.0106	$4.57 * 10^{-4}$	0.0863	0.125
ω_i	omega_i	0.00304	0.00651	0.00278	0.0264	0.0411
J	J	0.0102	0.0110	0.00878	0.00929	0.00409
σ	sigma	0.00853	$5.40 * 10^{-4}$	0.00477	0.0170	0.0289

6.3.3.1 Population Parameters

The mean estimates for the death rates α_{min} and α_{max} are notably lower and higher respectively than the values of 0.005 and 0.07 per year in the baseline scenario. It seems as if the model functions better with a wider range of death rates across Commoner wealth ω , with the usual values of α_{max} in particular around 0.11 per year giving rise to a steeper drop in model population during model collapses. This may seem high compared to what is historically observed, but is in line with the peaks in total death rate seen during famine events, at a rate of about 0.07 to 0.08 per year according to the [United Nations Population Division \[2023\]](#). The minimum death rate being lower at around 0.003 per year makes sense as well. In HANDY2, this value is an asymptotic lower bound, so the minimum death rates actually observed in the model are about 0.005 to 0.01 per year, in line with the lowest crude death rates observed in the world [[United Nations Population Division, 2023](#)] and the lower limit used in HANDY1.

The Elite population parameters κ_{start} , $x_{E,start}$, ω_μ , and μ_M are also interesting. The first two are very poorly constrained, and the baseline values of 100 and 5 (million people) lie squarely in the ensemble's distribution of values. For κ_{start} , with values ranging from about 1.58 to $6.77 * 10^5$, this spread seems to be because of not observing the *equivalent* Elite population, as discussed above. In particular, using the equivalent Elites κx_e as a parameter output would fix this issue. The starting Elite population x_E is also highly variable, but is inversely related to the values of μ_M ; higher values of the class mobility require a higher starting number of Elites to match the estimation target of 1% of the Commoner population.

Conversely, ω_μ and μ_M are strongly constrained and tightly linked to each other, addressing one of the major problems of hand-tuned HANDY2—providing a robust justification for the class

mobility parameters. In nearly all of the runs, ω_μ exceeds the 1.5 value in the baseline scenario but is below 2.5. We are reasonably confident that values near the mean of 2.12 are the correct setting that keeps the Elite growth rate and population curve at a sensible level throughout HANDY2 scenarios, while preventing excessive rates of upward mobility being introduced to the model. The values of μ_M cluster tightly near the average of 0.06 per year, with a few upper outliers around 0.20 per year. They are also reasonably correlated to the output x_E , and this greatly raises our confidence in the baseline value of 0.06 per year derived from extensive investigations by hand. We determine that these two parameter values are not at all arbitrary, but arise from the structure of Elite population movements across a wide range of initial model conditions, staying relatively constant even as other model quantities can differ greatly between scenarios.

6.3.3.2 Regenerating Quantities

Parameters in the model related to Nature are also some of the most tightly constrained by the EnKI process and are quite consistent with the baseline estimates. The value of the regeneration rate γ is substantially higher on average, about 0.086 per year, compared to the baseline 0.04, but many scenarios show lower values that are equally consistent with the observations of global phytomass stocks and consumption. The higher values do contribute to the greater recovery of phytomass seen in the preceding estimation scenarios. In most runs, λ_{start} is very close to our derived estimate of 27 ZJ for baseline HANDY2, discussed in section 3.4. We therefore validate the data collated by [Smil \[2015\]](#) in the context of our model. λ_{growth} is however largely irrelevant to the estimation. Most runs show values near or equal to 0, and only a few runs show values even near our manual estimate of $4 * 10^{-4}$ per year. Growth in Nature capacity over time seems

to be largely irrelevant in industrial society.

Values of the Regenerating depletion rate δ_R vary substantially between runs. This parameter's geometric mean is about $1.5 * 10^{-4}$ with the standard deviation being much higher, overlapping our baseline estimate. It is genuinely interesting that the production of phytomass is only weakly constrained by observations. This may be due to the low use of Regenerating resources compared to Nonrenewables, making them less important as a factor constraining the population, or more likely due to variations in δ_R being compensated by changes in the extraction efficiency $\xi_{m,R}$ to give a relatively constant production rate. There also seems to be little correlation between δ_R and λ .

6.3.3.3 Nonrenewable and Renewable quantities

The average value of δ_N in the estimation scenarios is very consistently higher than in the baseline scenarios, averaging $2.57 * 10^{-6}$ per million persons per year, with a standard deviation one tenth the size at $2.20 * 10^{-7}$. This strongly suggests that our initial value of $2.30 * 10^{-6}$ is too low, causing a premature decline in fossil fuel production not seen in reality and avoided in the previously shown estimation scenarios. Higher values were initially not used in our manual tuning because of excessive Nonrenewables use in the early part of the run around 1800-1900. Yet this was not at all an issue with EnKI—the simultaneous optimization of many variables allows for other parameters such as the starting extraction technology $\xi_{m,N}$ or investment threshold A_{Nth} to be changed in the context of δ_N , showing the power of EnKI in validating results that exogenously capture the historical growth of fossil fuel use.

A_{Nth} varies greatly between runs, and is often a bit lower than our baseline estimate of 315

EJ. But this is expected in light of it being the main parameter affecting the timing of Nonrenewables growth. Its value is thus tightly regulated by the other estimated parameters.

For the Renewables side of HANDY2, the mean δ_F of $8.5 * 10^{-4}$ per million persons per year is consistent with our observations but highly variable between runs. This is due to the limited history of Renewable flows use providing only weak constraints on their overall growth rate. It is as yet unclear what the speed of solar and wind buildout will be by 2040, as it strongly depends on external policy measures that are outside the scope of the model's endogenous projections. It does still appear that the higher values in a few parameter estimation runs lead to a rate of growth likely too fast to be realistic. Similarly, A_{Fth} is largely a timing variable that can take on very different estimated values depending on other model features. The baseline value of 12 EJ is correct in the context of other baseline scenario parameters.

6.3.3.4 Wealth-related Parameters: s and ρ

The values of the subsistence salary s and the famine wealth threshold ρ take on a wide range of values throughout the 50 estimation scenarios, greatly affecting the dynamics of the model runs as well as the absolute magnitude of parameters like Commoner and Elite consumption C_C, C_E . The mean value of s , 0.00511 TJ/person-yr, is close to the baseline scenario value but with significant variation. The mean value of ρ is conversely higher than the baseline value at 0.08 TJ/person. Both values strongly affect the population growth rate in all stages of the model, so they must either compensate or be compensated by other population growth parameters, particularly the Commoner death rate α_C and to a lesser extent the various investment thresholds. It appears as if the estimates in HANDY1 and manually-tuned HANDY2 are just one of many that

are reasonable in the context of the model structure, and can produce consistent model output.

What is also interesting is that the mean ratio of ρ to s is about 21.1 years, though with significant variance. This ratio represents the average wealth accumulation by a Commoner with respect to their yearly wage. HANDY1 assumes a ratio of 10, which is in line with the minimal ability of peasants to store wealth in a preindustrial society [Erdkamp, 2005]. In the industrial world with easy access to banks, investment, and stores of wealth and finance, a higher value seems to be called for, and the parameter estimation results seem to confirm this hypothesis.

6.3.3.5 Investment Mechanisms

The parameter estimates of extraction technology $\xi_{m,R}$, $\xi_{m,N}$, and $\xi_{m,F}$ are each highly variable between model runs, seemingly without any particular consistency and weakly constrained due to the opacity of the investment mechanism. Of these, the mean value of $\xi_{m,N}$ is the best constrained and closest to the baseline estimate at 0.251, which thus seems to be a validated initial guess. The original sketch of HANDY2 used 0.1, which is lower than any value observed in the 50 scenarios. $\xi_{m,R}$ conversely has a mean of 0.196 that is about twice the initial estimate. This is intriguing but not surprising—especially when considering that it is effectively equal to ξ_R due to the very high A_{Rth} value not being included in the parameter estimation. It does however indicate that the extraction of Regenerating energy in HANDY2 may be higher than previously anticipated. $\xi_{m,F}$ has a mean value of 0.086, above the baseline estimate of 0.05 but still lower than for the other two extraction technology parameters. The baseline value is still reasonable in light of this parameter’s high variability.

Of the final three parameters, ω_i shows a mean of 0.026 but with many runs having values

extremely near zero. It seems that any small value of the starting investment threshold, roughly near $1/\kappa$, suffices to construct a plausible HANDY2 scenario. J is more tightly constrained, with a mean of 0.0093 and standard deviation of 0.004; the global investment rate is the primary control of the speed at which extraction technology develops, and it is mostly insensitive to other model dynamics. Our hand-tuned guess of 0.01 is consistent with the estimation. The accumulated investment attrition rate σ shows a scenario mean of 0.017 per year, though with a distribution very skewed toward higher values. Investment attrition is not a major mechanism in HANDY2, so the exact value does not particularly matter.

6.4 Ensemble HANDY2 Runs

The ensemble of 50 parameter estimation runs discussed above was trained with policy parameters representing the baseline scenario of HANDY2. In each of the three runs shown so far, we have confirmed the basic finding of the manually-tuned model presented in Chapter 3—that our current projections of Nonrenewable and Renewable resource depletion will cause a moderate type-N collapse by 2200. But to more reliably test the model, we can create a super-ensemble. We run HANDY2 with each of the 50 sets of parameters found via estimation, while changing policy parameters in line with the six other scenarios of Chapter 3. This serves as an independent test of validity of these scenarios, and is also instructive to see how changes in societal policies, such as a decrease in birth rates, greater use of Nonrenewables, or reduction in inequality, affect HANDY2 when there is a wider array of unusual initial conditions. It also allows us to qualitatively estimate the likelihood of societal collapse or survival through 2200.

6.4.1 Scenario 1: Baseline

To this end, we begin by presenting a baseline stocks plot, Figure 6.5a and flows plot, Figure 6.5b, for the super-ensemble subset containing 48 of the 50 EnKI estimation runs. Runs 4 and 46 were excluded because of numerical instability in certain later scenarios and extreme Regenerating outlier status respectively. The results confirm what was seen in the first three scenarios of manually-estimated HANDY2. A type-N collapse occurs in most of the ensemble by 2200, following the depletion of Nonrenewables and the drawdown of Nature. Commoner population drops by at least 50% from its peak value with little variance between runs, and Nature in all but one run is severely depleted by 2100. The uptake of Renewable energy is fairly rapid in almost all of the ensemble, but is not enough to prevent per capita commoner consumption C_C/x_C from falling well below its 2020 peak value.

6.4.2 Scenarios 2 and 2-1: No Nonrenewables

In scenario 2 without Nonrenewables but with an increase in ψ_R following the Green Revolution (Figures 6.6a and 6.6b), we again confirm our findings from Chapter 3. Regenerating energy alone, even at modern day consumption levels, cannot sustain today's rates of population growth. In all but one run, we see a global population peak of about 2.5-3 billion people, which slowly decreases to 2200. When they do exist in large numbers, Elites gradually increase their wealth at the expense of the Commoners, preventing Commoner consumption from ever going above subsistence levels. Even though Nature is substantially drawn down, there is enough Regenerating energy to maintain a low-equilibrium steady state similar to what was observed in the centuries preceding the widespread use of fossil fuels.

We note once again that the Green Revolution and consequent increase in ψ_R was made possible through fossil-fuel inputs [Smil, 2017]. Thus the scenario 2-1 ensemble (Figures 6.7a, 6.7b) omits the increase in ψ_R . Unlike in the manual estimation scenario, the average of runs shows Commoner population reaching a low equilibrium of 1-1.5 billion people, with a steady consumption of Regenerating energy near 20 EJ/yr. These are quite close to the values of population and phytomass consumption seen at the eve of the Industrial Revolution—and shows even more how fossil fuels are necessary to support a large Human system. Interestingly, the stocks plot does not show an overall depletion of y_R with time; i.e. our simulations indicate that the ‘land shortage’ of Pomeranz [2000] was not global, but limited instead to a few high-development regions like Great Britain that intensively used their phytomass stocks.

6.4.3 Scenario 3: No Renewables

In this scenario (Figures 6.8a, 6.8b), we apply modifications to the ensemble that set δ_R to 0 and change the birth rate policy to a projection above the replacement level, but unlike in Section 3.5, we keep A_{Nth} at its estimated value. Almost all runs show a full and dramatic collapse in the Commoner and Elite populations alike, with x_C dropping 80% from its peak value of 11-12 billion by the end of the run, and x_E shrinking similarly due to strong Elite downward mobility. Commoner consumption C_C/x_C in the 22nd century is below even its pre-industrial level, indicating population decline due to a mass starvation event just as in the manually-tuned scenario. Regenerating Nature recovers somewhat from its very low levels in some runs by 2200, but this is small comfort for a society forced to return to a low-equilibrium state ravaged by global heating and without access to easily-obtained Nonrenewables. This situation is one thankfully

rendered unlikely due to the rapid buildout of solar and wind power in many areas of the world.

6.4.4 Scenario 4: Inequality Reduction

In this scenario (Figures 6.9a, 6.9b), we turn `kappa_switch` on and thereby impose a reduction (or slight increase in some runs) of the wealth inequality κ to 10 by 2200. This reduction is plotted as Figure 6.10. The scenario dynamics are otherwise very similar to the baseline: world population peaks at a value of 11-12 billion by 2100, before declining by about half in almost all runs due to resource scarcity by 2200. Commoner consumption is not significantly higher than it is for Scenario 1; in most cases, it matches the 2020-era peak by 2200 instead of exceeding it. This is likely because κ is already significantly smaller than 100 in many parameter-estimated runs. There is simply not enough Elite wealth accumulation to transfer to the Commoners in these instances. In addition, the slightly higher wealth means that the EnKI runs on average have slightly less regeneration of Nature y_R even after the partial type-N collapse, posing long-run sustainability issues. It is clear that for this ensemble, mitigating inequality alone will not sufficiently address the issue of resource exhaustion, while it may do so for ensembles with higher average κ . More societal change is still necessary.

6.4.5 Scenario 5: Higher Renewables Use and Inequality Reduction

As with the hand-tuned Scenario 5, increasing the use of Renewables is the single most important policy change that prevents resource collapse (Figures 6.11a, 6.11b). Setting F_{policy} to 1 ensures a steadily growing rate of Renewables use in every ensemble run, and almost all maintain a steady global population from about 2080 onwards, now controlled largely by the birth rate

policy parameters rather than by insufficient income to survive. Following a brief dip in wealth ω during the transition from fossil fuels to Renewables, per-capita Commoner consumption rises steadily in most runs to end up slightly above the peak levels of the 21st century. In many runs, there is even enough wealth to allow Commoner upward mobility—something that does not happen in the super-ensemble’s Scenario 1. We can thus say with confidence that large-scale investment in solar and wind power will likely ensure a prosperous future for humanity.

The depletion of Regenerating resources and slow rise by the end of the model integration is similar to that of the manually-tuned scenario, and slightly slower than in the super-ensemble baseline. But a minority of ensemble runs show a more rapid decrease in the Nonrenewables consumption past 2030, reaching a value of less than 25 EJ/yr by 2100. This is enough to limit the total burned carbon in HANDY2 to less than 50 ZJ, leaving a small stock of fossil fuel reserves in the ground. While climate change is not simulated in HANDY2, these outcomes would likely limit global warming to somewhat less than 2.5 degrees C above pre-industrial levels [Masson-Delmotte et al., 2021, van Vuuren et al., 2011], further reducing the overall disruption incurred by the Human System.

6.4.6 Scenario 6: Higher Renewables Use, Inequality Reduction, and Lower Birth Rates

In the final set of parameter estimation runs (Figures 6.12a, 6.12b), we combine all three proposed policy changes—a decrease in societal inequality, greater use of Renewables, and lower overall birth rates—to produce the most socially optimal outcome. The results are in line with what was seen in Chapter 3. Per-capita commoner consumption increases steadily past 2100

CE and reaches values that allow very rapid upward class mobility by 2200. But the decline in Commoner population in Scenario 6 is far less than in the manually tuned run, with a final population mean of 8 billion people. Only one ensemble run undergoes a total collapse, and all the others show a population in 2200 of well over 5 billion, with some particles even maintaining a steady state population of 12 billion Commoners. This is likely due to the lower minimum death rates α_{min} obtained via the EnKI process. Consequently, the society is not quite as well-off as it may otherwise be, and there is room for improvement in overall wealth if birth rates are reduced further. While less effective with this particular super-ensemble, birth rate policy is still an effective tool to ensure global prosperity.

As in Scenario 5, about 20% of the ensemble runs burn less than 50 ZJ of their starting Nonrenewables stock y_N , potentially mitigating climate risk. Nature y_R recovers slightly more by the end of the run in Scenario 6, but still less so than in the manually tuned scenario.

6.5 Conclusion

The EnKI parameter estimation process [Dunbar et al., 2022, Iglesias et al., 2013] is an effective means of determining the hidden internal state of many types of models, and we have successfully applied it to HANDY2. By numerically inverting observations of Commoner and Elite population and energy consumption to constrain the value of 22 parameters, the high dimensionality of HANDY2 is no longer an obstacle to a better understanding of its phase space. Careful timestepping techniques like the `EKSStableScheduler` [Kovachki and Stuart, 2019] ensure that over 70% of EnKI runs converge to the the observation truth within 500 iterations, with far more accuracy than the hand-tuned baseline scenario in Chapter 3, capturing key de-


tails such as the later peak date of fossil fuel use and the overall proportion of Elites in society. Running the estimation process multiple times creates an ensemble of 50 runs with many distinct boundary conditions, which we use to analyze where our initial estimates of model parameters were correct (as with s and λ_{start}), and where they were incorrect and a different value may have been appropriate (as with δ_N , α_{max} , or the class mobility wealth threshold ω_μ).

The resulting HANDY2 ensemble was also used successfully to simulate the other scenarios of Chapter 3, and despite never being trained on them, changes in several different policy parameters did not destabilize more than two runs in total. The results from the EnKI superensemble largely confirm results from the manually-tuned scenarios, with some minor differences. For instance, global population in the hypothetical Regenerating-only Scenario 2 was found to be stable over time with moderate depletion levels, while reducing the inequality κ in Scenario 4 does not cause the Commoners to become much richer if κ is low to begin with. We also see some intriguing emergent properties in certain unusual runs, such as a steeper drop in fossil fuel consumption that occurs in Scenarios 5 and 6 if Nonrenewables are replaced more rapidly by greater amounts of Renewables. We now have greater confidence that the present-day use of Renewable energy is not enough to prevent a severe type-N resource collapse within the next two centuries. But this can be almost certainly averted with even a linear increase in Renewables use at the present-day rate. Societal affluence can be improved even further by reducing overall birth rates through family planning, and by decreasing high levels of inequality.

EnKI parameter estimation is not perfect, but has several minor issues likely resolvable through more careful selection of parameter priors or output variables for the inversion. For instance, the wealth inequality κ and subsistence salary s is poorly learned in many ensemble runs. In the first case, this may be fixed by adding pseudo-observations of *equivalent* Elites


instead of *raw* Elites, and in the second case, by adding a historical estimate of global per capita energy consumption for years when that data is available. The EnKI process was also unable to learn the starting values of birth rates β_C , β_E at the same time as the death rates α_{min} and α_{max} because of model instability. One potential fix is to transform β into a parameter on the domain $[0, 1]$, corresponding to $[\alpha_{min}, \alpha_{max}]$ on the unconstrained space. This will require some minor changes to the HANDY2 code. The estimation convergence rate of only about 70-75% is also a persistent weakness of our EnKI implementation. More testing, debugging, and careful selection of inversion hyperparameters may be required to improve this figure, and we may wish to replace HANDY2's Euler integrator with one such as the Runge-Kutta method (RK4) that is more numerically stable. Further work along these lines is planned to refine and improve HANDY2, making it an even more effective basic model of the Earth and Human systems, with deep implications for the medium-term future outlook of global civilization.

HANDY2 ensemble, Baseline Scenario 1: Population, wealth, and resource stocks
 Solid and dashed lines are Handy2 model outputs, dotted lines are real-world estimates



(a) Stocks plot


HANDY2 ensemble, Baseline Scenario 1: Resource flows and population
 Solid and dashed lines are Handy2 model outputs, dotted lines are real-world estimates



(b) Flows plot


Figure 6.5: Stocks and flows plots for a super-ensemble of 48 randomly-seeded members in the HANDY2 baseline scenario (Scenario 1), obtained via EnKI parameter estimation. Each solid and dash-dot line represents the super-ensemble mean values for HANDY2 output variables, with the legends and labels for each line the same as in Figure 3.5. Each vertically shaded interval extends one sample standard deviation from the ensemble mean. The yellow dash-dot line is the accessible renewable energy F .

HANDY2 ensemble, Scenario 2 with no Nonrenewables: Population, wealth, and resource stocks
 Solid and dashed lines are Handy2 model outputs, dotted lines are real-world estimates



(a) Stocks plot

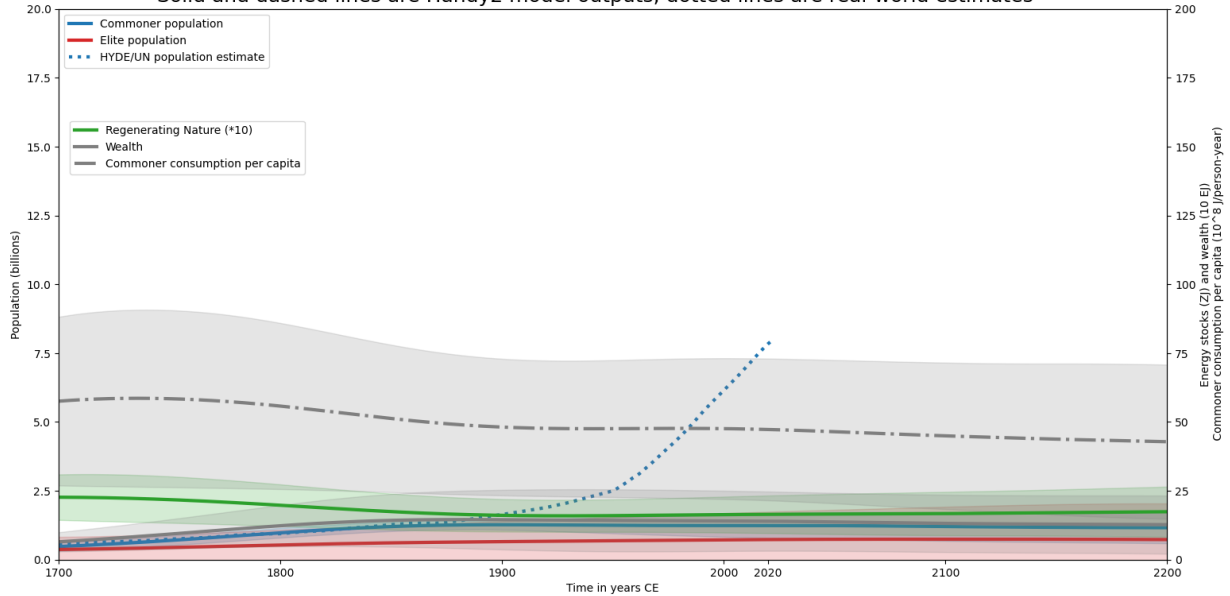
HANDY2 ensemble, Scenario 2 with no Nonrenewables: Resource flows and population
 Solid and dashed lines are Handy2 model outputs, dotted lines are real-world estimates



(b) Flows plot

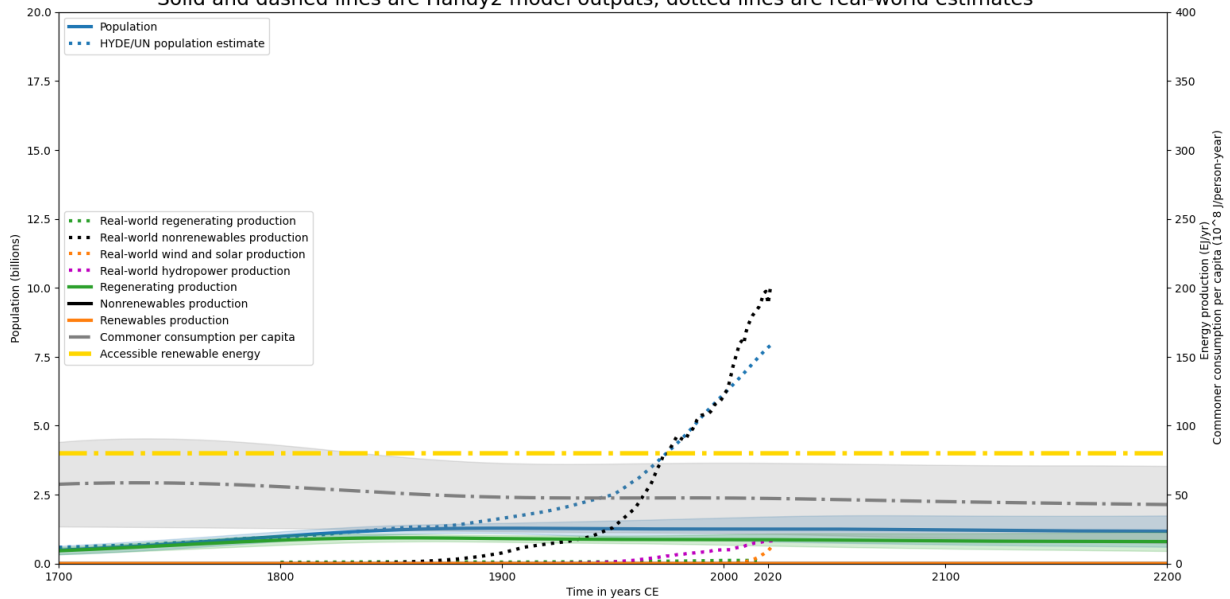
Figure 6.6: Stocks and flows plots for a super-ensemble of 48 randomly-seeded members in the HANDY2 scenario without Nonrenewables (Scenario 2), obtained via EnKI parameter estimation. This plot uses the same color scheme and scale as Figure 6.5.

HANDY2 ensemble, Scenario 2-1 with no Nonrenewables: Population, wealth, and resource stocks
Solid and dashed lines are Handy2 model outputs, dotted lines are real-world estimates




(a) Stocks plot

HANDY2 ensemble, Scenario 2-1 with no Nonrenewables: Resource flows and population
Solid and dashed lines are Handy2 model outputs, dotted lines are real-world estimates




(b) Flows plot

Figure 6.7: Stocks and flows plots for a super-ensemble of 48 randomly-seeded members in the HANDY2 scenario without Nonrenewables or increasing ψ_R (Scenario 2-1), obtained via EnKI parameter estimation. This plot uses the same color scheme and scale as Figure 6.5.




(a) Stocks plot




(b) Flows plot

Figure 6.8: Stocks and flows plots for a super-ensemble of 48 randomly-seeded members in the HANDY2 scenario without Renewables (Scenario 3), obtained via EnKI parameter estimation. This plot uses the same color scheme and scale as Figure 6.5.



(a) Stocks plot



(b) Flows plot

Figure 6.9: Stocks and flows plots for a super-ensemble of 48 randomly-seeded members in the HANDY2 scenario with reduction of inequality κ (Scenario 4), obtained via EnKI parameter estimation. This plot uses the same color scheme and scale as Figure 6.5.






Figure 6.10: The level of inequality κ over time for EnKI parameter estimation ensembles with inequality reduction (Scenario 4). Values of κ converge in a log-logistic manner to a value of $\kappa = 10$ by 2150.

HANDY2 ensemble, Scenario 5 with high Renewables use: Population, wealth, and resource stocks
Solid and dashed lines are Handy2 model outputs, dotted lines are real-world estimates



(a) Stocks plot


HANDY2 ensemble, Scenario 5 with high Renewables use: Resource flows and population
Solid and dashed lines are Handy2 model outputs, dotted lines are real-world estimates



(b) Flows plot

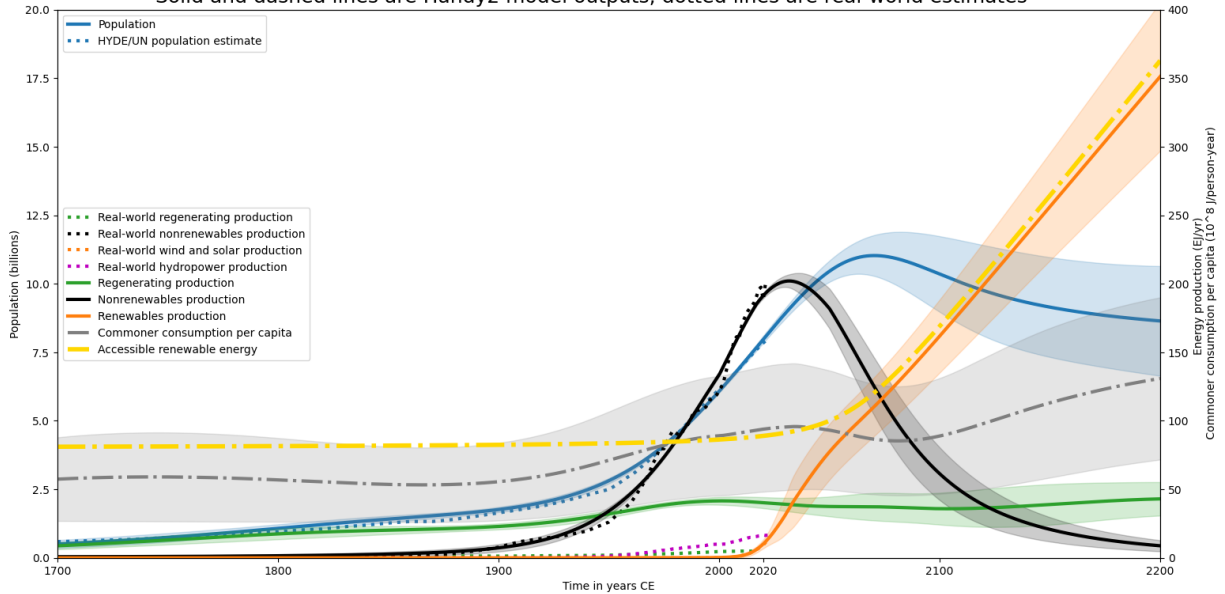
Figure 6.11: Stocks and flows plots for a super-ensemble of 48 randomly-seeded members in the HANDY2 scenario, with higher use of Renewable flows ($F_{policy} = 1$) and reduction of inequality κ (Scenario 5), obtained via EnKI parameter estimation. This plot uses the same color scheme and scale as Figure 6.5.

HANDY2 ensemble, Scenario 6 with high Renewables use, lower birth rates: Population, wealth, and resource stocks
Solid and dashed lines are Handy2 model outputs, dotted lines are real-world estimates



(a) Stocks plot

HANDY2 ensemble, Scenario 6 with high Renewables use, lower birth rates: Resource flows and population
Solid and dashed lines are Handy2 model outputs, dotted lines are real-world estimates



(b) Flows plot

Figure 6.12: Stocks and flows plots for a super-ensemble of 48 randomly-seeded members in the HANDY2 scenario, with aggressive reduction in birth rates β_C , higher use of Renewable flows ($F_{policy} = 1$), and reduction of inequality κ (Scenario 6), obtained via EnKI parameter estimation. This plot uses the same color scheme and scale as Figure 6.5.

Bibliography

- J. Bistline, N. Mehrotra, and C. Wolfram. Economic Implications of the Climate Provisions of the Inflation Reduction Act, May 2023. URL <https://www.nber.org/papers/w31267>.
- A. V. Bogoviz, S. V. Lobo, and A. N. Alekseev. CURRENT STATE AND FUTURE PROSPECTS OF HYDRO ENERGY IN RUSSIA. *International Journal of Energy Economics and Policy*, 10(3):482–488, Mar. 2020. ISSN 21464553. doi: 10.32479/ijeep.8968. URL <https://www.econjournals.com/index.php/ijeep/article/view/8968>.
- S. M. Christelow. The Division of Inheritance and the Provision of Non-Inheriting Offspring among the Anglo-Norman Elite. *Medieval Prosopography*, 17(2):3–44, 1996. ISSN 0198-9405. URL <https://www.jstor.org/stable/44946233>. Publisher: Board of Trustees of Western Michigan University through its Medieval Institute Publications.
- C. Y. C. Chu and R. D. Lee. Famine, revolt, and the dynastic cycle. *Journal of Population Economics*, 7(4):351–378, Nov. 1994. ISSN 1432-1475. doi: 10.1007/BF00161472. URL <https://doi.org/10.1007/BF00161472>.
- Climate Modeling Alliance. Prior distributions · EnsembleKalmanProcesses.jl, a. URL https://clima.github.io/EnsembleKalmanProcesses.jl/dev/parameter_distributions/.
- Climate Modeling Alliance. Learning rate schedulers, b. URL https://clima.github.io/EnsembleKalmanProcesses.jl/dev/learning_rate_scheduler/.
- P. Coghlin, U. O’Brien, and K. MacLean. *The Report of the Independent Public Inquiry into the Non-Domestic Renewable Heat Incentive (RHI) Scheme, Volume 1 - Chapters 1-19*. 2020.
- T. Covington. [STUDY] Supersized: Americans Are Living in Bigger Houses With Fewer People | The Zebra, 2023. URL <https://www.thezebra.com/resources/home/median-home-size-in-us/>.
- T. de Silva and S. Tenreyro. Population Control Policies and Fertility Convergence. *Journal of Economic Perspectives*, 31(4):205–228, Nov. 2017. ISSN 0895-3309. doi: 10.1257/jep.31.4.205. URL <https://www.aeaweb.org/articles?id=10.1257/jep.31.4.205>.

- M. Dribe and F. Scalone. Social Class and Net Fertility Before, During and After the Demographic Transition: A Micro-Level Analysis of Sweden 1880–1970. *Demographic Research*, 30:429–464, Feb. 2014. doi: 10.4054/DemRes.2014.30.15.
- O. R. Dunbar, I. Lopez-Gomez, A. Garbuno-Iñigo, D. Z. Huang, E. Bach, and J. long Wu. Ensemblekalmanprocesses.jl: Derivative-free ensemble-based model calibration. *Journal of Open Source Software*, 7(80):4869, 2022. doi: 10.21105/joss.04869. URL <https://doi.org/10.21105/joss.04869>.
- H. Díaz and C. Guedes Soares. Review of the current status, technology and future trends of offshore wind farms. *Ocean Engineering*, 209:107381, Aug. 2020. ISSN 0029-8018. doi: 10.1016/j.oceaneng.2020.107381. URL <https://www.sciencedirect.com/science/article/pii/S002980182030411X>.
- Energy Information Administration. Increased U.S. renewable and natural gas generation likely to reduce summer coal demand. URL <https://www.eia.gov/todayinenergy/detail.php?id=56760>.
- Energy Institute. Resources and data downloads, 2023. URL <https://www.energyinst.org/statistical-review/resources-and-data-downloads>.
- Engineering Toolbox. Combustion of Fuels - Carbon Dioxide Emission. URL https://www.engineeringtoolbox.com/co2-emission-fuels-d_1085.html.
- P. Erdkamp. *The Grain Market in the Roman Empire: A Social, Political and Economic Study*. Cambridge University Press, Nov. 2005. ISBN 978-1-139-44768-3. Google-Books-ID: Iij9uvGtJFEC.
- D. Evans, M. Polkinghorne, R. Fletcher, D. Brotherson, T. Hall, S. Klassen, and P. Wijker. Perspectives on the ‘Collapse’ of Angkor and the Khmer Empire. In *The Angkorian World*, pages 541–553. Routledge, London, 1 edition, Apr. 2023. ISBN 978-1-351-12894-0. doi: 10.4324/9781351128940-37. URL <https://www.taylorfrancis.com/books/9781351128940/chapters/10.4324/9781351128940-37>.
- G. Evensen and P. J. van Leeuwen. Assimilation of Geosat Altimeter Data for the Agulhas Current Using the Ensemble Kalman Filter with a Quasigeostrophic Model. *Monthly Weather Review*, 124(1):85–96, Jan. 1996. ISSN 1520-0493, 0027-0644. doi: 10.1175/1520-0493(1996)124<0085:AOGADF>2.0.CO;2. URL https://journals.ametsoc.org/view/journals/mwre/124/1/1520-0493_1996_124_0085_aogadf_2_0_co_2.xml. Publisher: American Meteorological Society Section: Monthly Weather Review.
- J. Galloway, F. Dentener, E. Boyer, R. Howarth, S. Seitzinger, G. Asner, C. Cleveland, P. Green, E. Holland, D. Karl, A. Michaels, J. Porter, A. Townsend, and C. Vöosmarty. Nitrogen Cycles: Past, Present, and Future. *Biogeochemistry*, 70:153–226, Jan. 2004. doi: 10.1007/s10533-004-0370-0.

- A. Garbuno-Inigo, F. Hoffmann, W. Li, and A. M. Stuart. Interacting langevin diffusions: Gradient structure and ensemble kalman sampler. *SIAM Journal on Applied Dynamical Systems*, 19(1):412–441, 2020. doi: 10.1137/19M1251655. URL <https://doi.org/10.1137/19M1251655>.
- A. Gat. *War in Human Civilization*. Oxford University Press, Oxford, illustrated edition edition, Apr. 2008. ISBN 978-0-19-923663-3.
- B. Geissler, L. Hermann, M. C. Mew, and G. Steiner. Striving Toward a Circular Economy for Phosphorus: The Role of Phosphate Rock Mining. *Minerals*, 8(9):395, Sept. 2018. ISSN 2075-163X. doi: 10.3390/min8090395. URL <https://www.mdpi.com/2075-163X/8/9/395>. Number: 9 Publisher: Multidisciplinary Digital Publishing Institute.
- General Electric. Aeroderivative and Heavy-Duty Gas Turbines | GE Gas Power. URL <https://www.ge.com/gas-power/products/gas-turbines>.
- A. Goldberg, A. M. Mychajliw, and E. A. Hadly. Post-invasion demography of prehistoric humans in South America. *Nature*, 532(7598):232–235, Apr. 2016. ISSN 1476-4687. doi: 10.1038/nature17176. URL <https://www.nature.com/articles/nature17176>. Number: 7598 Publisher: Nature Publishing Group.
- M. Grunwald. The ‘Green Energy’ That Might Be Ruining the Planet, Mar. 2021. URL <https://www.politico.com/news/magazine/2021/03/26/biomass-carbon-climate-politics-477620>.
- J. Hansen, P. Kharecha, M. Sato, V. Masson-Delmotte, F. Ackerman, D. J. Beerling, P. J. Hearty, O. Hoegh-Guldberg, S.-L. Hsu, C. Parmesan, J. Rockstrom, E. J. Rohling, J. Sachs, P. Smith, K. Steffen, L. V. Susteren, K. v. Schuckmann, and J. C. Zacher. Assessing “Dangerous Climate Change”: Required Reduction of Carbon Emissions to Protect Young People, Future Generations and Nature. *PLOS ONE*, 8(12):e81648, Dec. 2013. ISSN 1932-6203. doi: 10.1371/journal.pone.0081648. URL <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0081648>. Publisher: Public Library of Science.
- B. Harrington. Capital without Borders: Wealth Managers and the One Percent. In *Capital without Borders*. Harvard University Press, Jan. 2017. ISBN 978-0-674-97361-9. doi: 10.4159/9780674973619. URL <https://www.degruyter.com/document/doi/10.4159/9780674973619/html>.
- R. Heinberg. Technology to the Rescue, Oct. 2013. URL <https://www.resilience.org/stories/2013-10-02/snake-oil-chapter-2-technology-to-the-rescue/>.
- P. L. Houtekamer and H. L. Mitchell. A Sequential Ensemble Kalman Filter for Atmospheric Data Assimilation. *Monthly Weather Review*, 129(1):123–137, Jan. 2001. ISSN 1520-0493, 0027-0644. doi: 10.1175/1520-0493(2001)129<0123:ASEKFF>2.0.CO;2. URL https://journals.ametsoc.org/view/journals/mwre/129/1/1520-0493_2001_129_0123_asekff_2.0.co_2.xml. Publisher: American Meteorological Society Section: Monthly Weather Review.

- J. E. Hughes. *Family Wealth: Keeping It in the Family—How Family Members and Their Advisers Preserve Human, Intellectual, and Financial Assets for Generations*. John Wiley & Sons, June 2004. ISBN 978-1-57660-151-8. Google-Books-ID: Sg_9CgAAQBAJ.
- M. Hájek, J. Zimmermannová, K. Helman, and L. Roženský. Analysis of carbon tax efficiency in energy industries of selected EU countries. *Energy Policy*, 134:110955, Nov. 2019. ISSN 0301-4215. doi: 10.1016/j.enpol.2019.110955. URL <https://www.sciencedirect.com/science/article/pii/S0301421519305427>.
- M. Höök, A. Sivertsson, and K. Aleklett. Validity of the Fossil Fuel Production Outlooks in the IPCC Emission Scenarios. *Natural Resources Research*, 19:63–81, June 2010. doi: 10.1007/s11053-010-9113-1.
- M. Iglesias and Y. Yang. Adaptive regularisation for ensemble Kalman inversion. *Inverse Problems*, 37(2):025008, Feb. 2021. ISSN 0266-5611, 1361-6420. doi: 10.1088/1361-6420/abd29b. URL <https://iopscience.iop.org/article/10.1088/1361-6420/abd29b>.
- M. A. Iglesias, K. J. H. Law, and A. M. Stuart. Ensemble Kalman methods for inverse problems. *Inverse Problems*, 29(4):045001, Mar. 2013. ISSN 0266-5611. doi: 10.1088/0266-5611/29/4/045001. URL <https://dx.doi.org/10.1088/0266-5611/29/4/045001>. Publisher: IOP Publishing.
- W. Jongman. Gibbon was right: The decline and fall of the roman economy. In *Crises and the Roman Empire*, pages 183–199. Brill, Jan. 2007. ISBN 978-90-474-2090-3. URL https://brill.com/display/book/edcoll/9789047420903/Bej.9789004160507.i-448_015.xml. Section: Crises and the Roman Empire.
- S. N. Kaplan and J. D. Rauh. Family, Education, and Sources of Wealth among the Richest Americans, 1982-2012. *American Economic Review*, 103(3):158–162, May 2013. ISSN 0002-8282. doi: 10.1257/aer.103.3.158. URL <https://www.aeaweb.org/articles?id=10.1257/aer.103.3.158>.
- K. Klein Goldewijk, A. Beusen, J. Doelman, and E. Stehfest. Anthropogenic land use estimates for the Holocene – HYDE 3.2. *Earth System Science Data*, 9(2):927–953, Dec. 2017. ISSN 1866-3508. doi: 10.5194/essd-9-927-2017. URL <https://essd.copernicus.org/articles/9/927/2017/essd-9-927-2017.html>. Publisher: Copernicus GmbH.
- A. Kokkonen and A. Sundell. Delivering Stability—Primogeniture and Autocratic Survival in European Monarchies 1000–1800. *American Political Science Review*, 108(2): 438–453, May 2014. ISSN 0003-0554, 1537-5943. doi: 10.1017/S000305541400015X. URL <https://www.cambridge.org/core/journals/american-political-science-review/article/abs/delivering-stabilityprimogeniture-and-autocratic-survival-in-european-monarchies-10001800/2399079C174599A840E5230E8827609C#access-block>. Publisher: Cambridge University Press.

- N. B. Kovachki and A. M. Stuart. Ensemble Kalman inversion: a derivative-free technique for machine learning tasks. *Inverse Problems*, 35(9):095005, Sept. 2019. ISSN 0266-5611, 1361-6420. doi: 10.1088/1361-6420/ab1c3a. URL <https://iopscience.iop.org/article/10.1088/1361-6420/ab1c3a>.
- H. M. Kwon, S. W. Moon, T. S. Kim, D. W. Kang, J. L. Sohn, and J. Lee. A study on 65 % potential efficiency of the gas turbine combined cycle. *Journal of Mechanical Science and Technology*, 33(9):4535–4543, Sept. 2019. ISSN 1976-3824. doi: 10.1007/s12206-019-0850-8. URL <https://doi.org/10.1007/s12206-019-0850-8>.
- J. Laherrère. Are there enough fossil fuels to generate the IPCC CO2 baseline scenario?, Aug. 2019. URL <https://aspofrance.org/2019/08/30/are-there-enough-fossil-fuels-to-generate-the-ipcc-co2-baseline-scenario/>.
- N. N. Lal, Y. Dkhissi, W. Li, Q. Hou, Y.-B. Cheng, and U. Bach. Perovskite Tandem Solar Cells. *Advanced Energy Materials*, 7(18):1602761, 2017. ISSN 1614-6840. doi: 10.1002/aenm.201602761. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/aenm.201602761>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/aenm.201602761>.
- C. Le Quéré, R. Moriarty, R. M. Andrew, J. G. Canadell, S. Sitch, J. I. Korsbakken, P. Friedlingstein, G. P. Peters, R. J. Andres, T. A. Boden, R. A. Houghton, J. I. House, R. F. Keeling, P. Tans, A. Arneeth, D. C. E. Bakker, L. Barbero, L. Bopp, J. Chang, F. Chevallier, L. P. Chini, P. Ciais, M. Fader, R. A. Feely, T. Gkritzalis, I. Harris, J. Hauck, T. Ilyina, A. K. Jain, E. Kato, V. Kitidis, K. Klein Goldewijk, C. Koven, P. Landschützer, S. K. Lauvset, N. Lefèvre, A. Lenton, I. D. Lima, N. Metz, F. Millero, D. R. Munro, A. Murata, J. E. M. S. Nabel, S. Nakaoka, Y. Nojiri, K. O’Brien, A. Olsen, T. Ono, F. F. Pérez, B. Pfeil, D. Pierrot, B. Poulter, G. Rehder, C. Rödenbeck, S. Saito, U. Schuster, J. Schwinger, R. Séférian, T. Steinhoff, B. D. Stocker, A. J. Sutton, T. Takahashi, B. Tilbrook, I. T. van der Laan-Luijkx, G. R. van der Werf, S. van Heuven, D. Vandemark, N. Viovy, A. Wiltshire, S. Zaehle, and N. Zeng. Global Carbon Budget 2015. *Earth System Science Data*, 7(2): 349–396, Dec. 2015. ISSN 1866-3508. doi: 10.5194/essd-7-349-2015. URL <https://essd.copernicus.org/articles/7/349/2015/>. Publisher: Copernicus GmbH.
- Y. Li, E. Kalnay, S. Motesharrei, J. Rivas, F. Kucharski, D. Kirk-Davidoff, E. Bach, and N. Zeng. Climate model shows large-scale wind and solar farms in the Sahara increase rain and vegetation. *Science*, 361(6406):1019–1022, Sept. 2018. doi: 10.1126/science.aar5629. URL <https://www.science.org/doi/10.1126/science.aar5629>. Publisher: American Association for the Advancement of Science.
- E. N. Lorenz. Predictability – a problem partly solved. In R. Hagedorn and T. Palmer, editors, *Predictability of Weather and Climate*, pages 40–58. Cambridge University Press, Cambridge, 2006. ISBN 978-0-521-84882-4. doi: 10.1017/CBO9780511617652.004. URL <https://www.cambridge.org/core/books/predictability-of-weather-and-climate/predictability-a-problem-partly-solved/3221BDE379DEB669BA52C66263AF3206>.

- A. Louwen, W. van Sark, R. Schropp, and A. Faaij. A cost roadmap for silicon heterojunction solar cells. *Solar Energy Materials and Solar Cells*, 147:295–314, Apr. 2016. ISSN 0927-0248. doi: 10.1016/j.solmat.2015.12.026. URL <https://www.sciencedirect.com/science/article/pii/S0927024815006741>.
- X. Lu, M. B. McElroy, and J. Kiviluoma. Global potential for wind-generated electricity. *Proceedings of the National Academy of Sciences of the United States of America*, 106(27):10933–10938, July 2009. ISSN 0027-8424. doi: 10.1073/pnas.0904101106. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2700152/>.
- V. Masson-Delmotte, P. Zhai, A. Pirani, S. L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M. I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J. B. R. Matthews, T. K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou, editors. *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2021. doi: 10.1017/9781009157896.
- C. McGlade and P. Ekins. The geographical distribution of fossil fuels unused when limiting global warming to 2 °C. *Nature*, 517(7533):187–190, Jan. 2015. ISSN 1476-4687. doi: 10.1038/nature14016.
- G. E. Metcalf. Carbon taxes in theory and practice. *Annual Review of Resource Economics*, 13(1): 245–265, 2021. doi: 10.1146/annurev-resource-102519-113630. URL <https://doi.org/10.1146/annurev-resource-102519-113630>.
- T. Mitchell. 1. Trends in income and wealth inequality, Jan. 2020. URL <https://www.pewresearch.org/social-trends/2020/01/09/trends-in-income-and-wealth-inequality/>.
- S. Mote, J. Rivas, and E. Kalnay. A Novel Approach to Carrying Capacity: From a priori Prescription to a posteriori Derivation Based on Underlying Mechanisms and Dynamics. *Annual Review of Earth and Planetary Sciences*, 48(1):657–683, 2020. doi: 10.1146/annurev-earth-053018-060428. URL <https://doi.org/10.1146/annurev-earth-053018-060428>. eprint: <https://doi.org/10.1146/annurev-earth-053018-060428>.
- S. Motesharrei, J. Rivas, and E. Kalnay. Human and nature dynamics (HANDY): Modeling inequality and use of resources in the collapse or sustainability of societies. *Ecological Economics*, 101:90–102, May 2014. ISSN 0921-8009. doi: 10.1016/j.ecolecon.2014.02.014. URL <https://www.sciencedirect.com/science/article/pii/S0921800914000615>.
- S. Motesharrei, J. Rivas, E. Kalnay, G. R. Asrar, A. J. Busalacchi, R. F. Cahalan, M. A. Cane, R. R. Colwell, K. Feng, R. S. Franklin, K. Hubacek, F. Miralles-Wilhelm, T. Miyoshi, M. Ruth, R. Sagdeev, A. Shirmohammadi, J. Shukla, J. Srebric, V. M. Yakovenko, and N. Zeng. Modeling sustainability: population, inequality, consumption, and bidirectional coupling of the Earth and Human Systems. *National Science Review*, 3(4):470–494, Dec. 2016. ISSN 2095-5138. doi: 10.1093/nsr/nww081. URL <https://doi.org/10.1093/nsr/nww081>.

- NBC News. America's cars and trucks are getting bigger, and so are their blind zones. Kids are paying the price., Oct. 2022. URL <https://www.nbcnews.com/news/us-news/americas-cars-trucks-are-getting-bigger-are-front-blind-zones-children-rcna52109>.
- T. V. Nguyen and H. Khieu. Does a global wealth tax reduce inequality? When Piketty meets Mankiw. *Research in Economics*, 74(2):119–130, June 2020. ISSN 1090-9443. doi: 10.1016/j.rie.2020.02.004. URL <https://www.sciencedirect.com/science/article/pii/S1090944320300387>.
- Our World in Data. GDP per capita vs. energy use. URL <https://ourworldindata.org/grapher/energy-use-per-capita-vs-gdp-per-capita>.
- J. M. Pacyna and E. G. Pacyna. An assessment of global and regional emissions of trace metals to the atmosphere from anthropogenic sources worldwide. *Environmental Reviews*, 9(4):269–298, Dec. 2001. ISSN 1181-8700. doi: 10.1139/a01-012. URL <https://cdnsiencepub.com/doi/10.1139/a01-012>. Publisher: NRC Research Press.
- A. F. Palmeirim, C. A. Peres, and F. C. W. Rosas. Giant otter population responses to habitat expansion and degradation induced by a mega hydroelectric dam. *Biological Conservation*, 174:30–38, June 2014. ISSN 0006-3207. doi: 10.1016/j.biocon.2014.03.015. URL <https://www.sciencedirect.com/science/article/pii/S0006320714001232>.
- F. Pearce. India defuses its population bomb: Fertility falls to two children per woman, 2021. URL <https://www.science.org/content/article/india-defuses-its-population-bomb-fertility-falls-two-children-woman>.
- K. B. Piepmeyer and T. S. Adkins. The status of women and fertility. *Journal of Biosocial Science*, 5(4):507–520, Oct. 1973. ISSN 1469-7599, 0021-9320. doi: 10.1017/S0021932000009378. URL <https://www.cambridge.org/core/journals/journal-of-biosocial-science/article/abs/status-of-women-and-fertility/B33ACA3D308565D4FC0D449F6CB0913A>. Publisher: Cambridge University Press.
- T. Piketty, E. Saez, and G. Zucman. Rethinking Capital and Wealth Taxation. working paper or preprint, Nov. 2022. URL <https://shs.hal.science/halshs-04104410>.
- K. Pomeranz. *The Great Divergence: China, Europe, and the Making of the Modern World Economy*. Princeton University Press, 2000. ISBN 978-0-691-09010-8. doi: 10.2307/j.ctt7sv80. URL <https://www.jstor.org/stable/j.ctt7sv80>.
- H.-O. Pörtner, D. C. Roberts, M. M. B. Tignor, E. S. Poloczanska, K. Mintenbeck, A. Alegría, M. Craig, S. Langsdorf, S. Löschke, V. Möller, A. Okem, and B. Rama, editors. *Climate Change 2022: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. 2022.

- C. Raudsepp-Hearne, G. Peterson, M. Tengö, E. Bennett, T. Holland, K. Benessaiah, G. MacDonald, and L. Pfeifer. Untangling the Environmentalist's Paradox: Why Is Human Well-being Increasing as Ecosystem Services Degrade? *BioScience*, 60:576–589, June 2013. doi: 10.1525/bio.2010.60.8.4.
- P. Raven et al. Scientist Letter to Biden van der Leyden Michel Suga Moon (February 11 2021) - DocumentCloud, 2021. URL <https://www.documentcloud.org/documents/20482842-scientist-leter-to-biden-van-der-leyden-michel-suga-moon-february-11-2021>.
- J. Rogelj, M. Meinshausen, J. Sedláček, and R. Knutti. Implications of potentially lower climate sensitivity on climate projections and policy. *Environmental Research Letters*, 9 (3):031003, Mar. 2014. ISSN 1748-9326. doi: 10.1088/1748-9326/9/3/031003. URL <https://dx.doi.org/10.1088/1748-9326/9/3/031003>. Publisher: IOP Publishing.
- S. Roman, S. Bullock, and M. Brede. Coupled Societies are More Robust Against Collapse: A Hypothetical Look at Easter Island. *Ecological Economics*, 132:264–278, Feb. 2017. ISSN 0921-8009. doi: 10.1016/j.ecolecon.2016.11.003. URL <https://www.sciencedirect.com/science/article/pii/S0921800916307509>.
- N. S. Rosenstein. *Rome at War: Farms, Families, and Death in the Middle Republic*. Univ of North Carolina Press, 2004. ISBN 978-0-8078-2839-7. Google-Books-ID: CGgwN9ZLaPYC.
- E. Saez and G. Zucman. Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data *. *The Quarterly Journal of Economics*, 131(2):519–578, May 2016. ISSN 0033-5533. doi: 10.1093/qje/qjw004. URL <https://doi.org/10.1093/qje/qjw004>.
- A. Schöniger, W. Nowak, and H.-J. Hendricks Franssen. Parameter estimation by ensemble Kalman filters with transformed data: Approach and application to hydraulic tomography. *Water Resources Research*, 48(4), 2012. ISSN 1944-7973. doi: 10.1029/2011WR010462. URL <https://onlinelibrary.wiley.com/doi/abs/10.1029/2011WR010462>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2011WR010462>.
- V. Smil. *Harvesting the Biosphere: What We Have Taken from Nature*. MIT Press, Aug. 2015. ISBN 978-0-262-52827-6. Google-Books-ID: DFT5DwAAQBAJ.
- V. Smil. *Energy Transitions: Global and National Perspectives*. ABC-CLIO, LLC, 2017. ISBN 978-1-4408-5324-1. Google-Books-ID: loKnDAEACAAJ.
- N. S. Sodhi and P. R. Ehrlich. *Conservation Biology for All*. OUP Oxford, Jan. 2010. ISBN 978-0-19-157425-2. Google-Books-ID: bCiQDwAAQBAJ.
- W. Steffen, W. Broadgate, L. Deutsch, O. Gaffney, and C. Ludwig. The trajectory of the Anthropocene: The Great Acceleration. *The Anthropocene Review*, 2(1):81–98, Apr. 2015. ISSN 2053-0196. doi: 10.1177/2053019614564785. URL <https://doi.org/10.1177/2053019614564785>. Publisher: SAGE Publications.

- R. M. Swanson. Developments in Silicon Solar Cells. In *2007 IEEE International Electron Devices Meeting*, pages 359–362, Dec. 2007. doi: 10.1109/IEDM.2007.4418946. ISSN: 2156-017X.
- J. E. Trancik. Renewable energy: Back the renewables boom. *Nature*, 507(7492):300–302, Mar. 2014. ISSN 1476-4687. doi: 10.1038/507300a. URL <https://www.nature.com/articles/507300a>. Number: 7492 Publisher: Nature Publishing Group.
- United Nations Population Division. World population prospects, 2023. URL <https://population.un.org/wpp/>.
- US Energy Information Administration. Use of geothermal energy. URL <https://www.eia.gov/energyexplained/geothermal/use-of-geothermal-energy.php>.
- D. P. van Vuuren, J. Edmonds, M. Kainuma, K. Riahi, A. Thomson, K. Hibbard, G. C. Hurtt, T. Kram, V. Krey, J.-F. Lamarque, T. Masui, M. Meinshausen, N. Nakicenovic, S. J. Smith, and S. K. Rose. The representative concentration pathways: an overview. *Climatic Change*, 109(1):5, Aug. 2011. ISSN 1573-1480. doi: 10.1007/s10584-011-0148-z. URL <https://doi.org/10.1007/s10584-011-0148-z>.
- M. Wang, M. A. Khan, I. Mohsin, J. Wicks, A. H. Ip, K. Z. Sumon, C.-T. Dinh, E. H. Sargent, I. D. Gates, and M. Golam Kibria. Can sustainable ammonia synthesis pathways compete with fossil-fuel based Haber–Bosch processes? *Energy & Environmental Science*, 14(5):2535–2548, 2021. doi: 10.1039/D0EE03808C. URL <https://pubs.rsc.org/en/content/articlelanding/2021/ee/d0ee03808c>. Publisher: Royal Society of Chemistry.
- K. S. Whitefoot and S. J. Skerlos. Design incentives to increase vehicle size created from the U.S. footprint-based fuel economy standards. *Energy Policy*, 41:402–411, Feb. 2012. ISSN 03014215. doi: 10.1016/j.enpol.2011.10.062. URL <https://linkinghub.elsevier.com/retrieve/pii/S0301421511008779>.
- S. Willcock, G. S. Cooper, J. Addy, and J. A. Dearing. Earlier collapse of Anthropocene ecosystems driven by multiple faster and noisier drivers. *Nature Sustainability*, pages 1–12, June 2023. ISSN 2398-9629. doi: 10.1038/s41893-023-01157-x. URL <https://www.nature.com/articles/s41893-023-01157-x>. Publisher: Nature Publishing Group.
- World Bank. World Development Indicators | Data Catalog, 2023. URL <https://datacatalog.worldbank.org/search/dataset/0037712/World-Development-Indicators>.
- S. Yang, Q. Jiang, and J. J. Sánchez-Barricarte. China’s fertility change: an analysis with multiple measures. *Population Health Metrics*, 20(1):12, Mar. 2022. ISSN 1478-7954. doi: 10.1186/s12963-022-00290-7. URL <https://doi.org/10.1186/s12963-022-00290-7>.
- R. York and J. A. McGee. Understanding the Jevons paradox. *Environmental Sociology*, 2(1):77–87, Jan. 2016. ISSN null. doi: 10.1080/23251042.2015.1106060. URL

<https://doi.org/10.1080/23251042.2015.1106060>. Publisher: Routledge -
eprint: <https://doi.org/10.1080/23251042.2015.1106060>.

N. Zeng, F. Zhao, G. J. Collatz, E. Kalnay, R. J. Salawitch, T. O. West, and L. Guanter. Agricultural Green Revolution as a driver of increasing atmospheric CO₂ seasonal amplitude. *Nature*, 515(7527):394–397, Nov. 2014. ISSN 1476-4687. doi: 10.1038/nature13893. URL <https://www.nature.com/articles/nature13893>. Number: 7527 Publisher: Nature Publishing Group.