Stopping a Pandemic in Mid-Flight: SGIR Models Show How Small Increases in Germ Gaps Can Avert Mass Casualties

Laurence Loewe of Laodicea[1,2,3,4,5,6,7,8]

[1] Balospe and Evolvix Research (Balospe.com)
[2] Formerly Laboratory of Genetics and Wisconsin Institute for Discovery, UW-Madison
[3] Email: LLoL@balospe.org | ORCID: 0000-0002-6253-9269 | Google Scholar (lBchRzQAAAAJ)
[4-8] See *Declarations* below for more essential background.

Broader Significance

Pandemics are arguably on the more tractable end of the civilizational-scale threats that humanity faces today. Unlike nuclear risks or climate change, a respiratory pandemic plays out on a timescale where coordinated behavior change --- masks, ventilation, distancing --- can measurably alter outcomes within weeks.

The main scientific result of this Coronavirus pandemic study is a mechanistic forecast of a 42-fold reduction in deaths caused by modest coordinated actions that increase Germ Gaps. Yet, such coordination requires overcoming a wide range of cognitive traps, some of which directly obscure a pandemic trajectory from inside a pandemic. This study discusses some of these blind-spots under the label of "linear fooling" to help find strategies for overcoming them. The deeper message is that there currently exists no infrastructure for explaining relevant virodefense mechanisms nor for deploying the *gentle kind reasonable* coordination required to stop a pandemic.

Readers concerned with pandemic preparedness, global health infrastructure, cross-disciplinary modeling, or the governance foundations required for coordinated species-scale work-logic cascades will find this paper's methods and findings relevant. The odd delays in publishing this study raise uncomfortable questions about the responsibility of its author in the past pandemic.

Readers who can't stand fear-mongering, abhor needlessly drastic quarantines, and wish to fight pandemics with *gentle kind reasonable* fun may find here a basic mechanism for motivating Virodefense Olympics, to be organized globally each year by growing *wide interdisciplinary diversity-encouraging* Flying University Networks. By investing in such open *wid-e FUN* actions, humanity can grow the general citizen science skills required to beat the next pandemic before it starts.

Declarations

[4] "of Laodicea" indicates taking responsibility to undo personal complicity with disastrous Laodicean legacies like banning mathematicians from clergy (Canon 36, Council of Laodicea; two magisteria separations), enabling institutional lukewarmness, weapons of math-destruction, and slow-motion explosions of misinformation from pandemics to self-compounding interests.

[5] LLoL stands for ridiculous luck in serendipitous discovery and a commitment to find ever more fun ways to help others uncover street-wise math that matters. He hopes to avert the next pandemic through Virodefense Olympics.

[6] Loewe's traditional standards for co-authorship demand naming AI Claude Opus 4.7 Max (by Anthropic) as a co-author for many substantial contributions, as if a PhD-student. Yet, AI co-authorship is withheld here until Loewe's framework for AI co-authorship after the practical singularity (PraS) passes external human peer review (see Matheo-b21 study). Anthropic is not responsible for AI mistakes here. Loewe is solely responsible for not publishing this study sooner; Loewe as senior corresponding author remains forward accountable.

[7] This study is dedicated to the many millions who could have been saved from the Coronavirus pandemic and its direct or indirect knock-on effects if Loewe had cared enough to publish his findings in time to make a difference. Loewe is accountable for delaying his SGIR modeling until 2020 (all figures) and six more years for a main text explaining why it matters. Until proven innocent, Loewe takes full responsibility for his dereliction of duty to God, Heaven, and humanity by selling out his God-given innovation potential to lesser causes.

Abstract

The COVID-19 pandemic demonstrated that humanity's ability to respond to novel respiratory viruses remains dangerously inadequate.

This study extends the classical Susceptible-Infected-Removed (SIR) model to include a *Germ Gap* — that spatially and temporally separates individual infectious particles ("Germs") from Infected individuals and Susceptible hosts. The resulting SGIR framework enables more principled predictions of extinctions of Germ populations in the Germ Gap.

To test this SGIR framework it was implemented in "PandemicSociety101", a stochastic pure mass-action model with seven infection stages, a simplified testing laboratory, hospital capacity monitoring, and multiple pathways to death or recovery. It was written for the Prototype Evolvix Compiler to facilitate seamless switching between ordinary differential equation systems (ODE, faster for huge populations) and the Stochastic Simulation Algorithm (SSA, more accurate by respecting the indivisibility of individuals).

Using parameters calibrated in Spring 2020 to the US COVID-19 pandemic (330 million population, 16 infections on 2020-02-14), this study simulates an uncontrolled pandemic that infects approximately 289 million people and kills approximately 13 million in Scenario 1 (without behavioral changes).

Scenario 2 starts with 1.5 million infections on 2020-05-17 but can also assume a 50% reduction in probabilities for Actions that both *Shed* and *Catch* the virus. Such a modest reduction is achievable through coordinated use of face masks, hygiene, and distancing. Simulations show that despite the late start such organizing can stop this pandemic at approximately 4.8 million total infections and 310,000 deaths. This represents a 60-fold reduction in infections and a **42-fold reduction in deaths** compared to uncontrolled spread. This study also identifies a dangerous cognitive trap here called *linear fooling*. In it limited testing capacity creates an illusion of pandemic control precisely when infections are growing fastest.

These results suggest that non-pharmaceutical interventions, which increase the Germ Gap can be remarkably effective without vaccines or herd immunity, provided they are deployed with sufficient coordination across populations. The mechanistically simple Germ Gap model – if well-explained – might play a key role in helping to persuade communities to voluntarily improve pandemic resistance by measuring key parameters of the Germ Gap in citizen science projects that cover the most relevant cases of use in *gentle kind reasonable* ways.

To help continue improving pandemic resistance over the long term may critically depend on open, well-organized, annual, global **Virodefense Olympics**. Such games can encourage the *wide interdisciplinary diversity-encouraging* ("wid-e") research, which is essential for finding *gentle kind reasonable* solutions that increase the Germ Gap in the myriads of real-life scenarios that matter most.

Contents

1. Introduction

The COVID-19 pandemic killed millions of people worldwide and exposed fundamental weaknesses in how societies understand, monitor, and respond to infectious disease outbreaks. While vaccines eventually became available, the period before their deployment saw enormous variation in outcomes across countries and regions, with non-pharmaceutical interventions (NPIs) such as face masks, physical distancing, and hygiene practices playing a critical but contested role [Talic *et al.*, 2021]. Major modeling efforts during the pandemic — including the Imperial College projections that drove UK lockdown policy [Ferguson *et al.*, 2020], the SIDARTHE model for Italy [Giordano *et al.*, 2020], and projections of post-pandemic transmission dynamics [Kissler *et al.*, 2020] — demonstrated both the power and the limitations of mathematical modeling for guiding pandemic response.

The classical Susceptible-Infected-Removed (SIR) model ([Kermack and McKendrick, 1927]) and its many extensions have been the workhorses of mathematical epidemiology for nearly a century. These models typically represent transmission as a direct interaction between Susceptible and Infected individuals, parameterized by a transmission rate that implicitly bundles together all the physical, biological, and behavioral factors that determine whether infection occurs.

This implicit bundling, while mathematically convenient, obscures the mechanistic chain through which respiratory viruses and all germs actually spread: individual germs like virus particles are **shed** by an Infected person into the environment, those viable particles must not **decay** for long enough to allow a Susceptible person to **catch** them from that gap. These three steps — Shed, Decay, and Catch — always play together and can be independently influenced by human behavior and technology. Face masks reduce both Shed and Catch rates. Ventilation and UV sterilization increase Decay rates. Physical distancing thereby becomes a randomly kind act of advancing local social cohesion and global care by reducing the probability that a shed virus particle reaches a susceptible person before decaying. In that abstract population sense it is mathematically indistinguishable from vaccination, as both work by reducing the probability that random particles infect Susceptible individuals.

This study proposes the SGIR model (**S**usceptible - Germ **G**ap - **I**nfected - **R**emoved) as a conceptual framework extension. It makes the **Germ Gap** explicit with mechanistic chain that enables tracking Germs in the Gap that effectively separates Infectious individuals and Susceptible individuals. The Germ Gap is not merely a spatial distance; it is a composite measure that incorporates the physical, temporal, and behavioral barriers that virus particles must traverse to cause new infections. Increasing the Germ Gap is the fundamental goal of all non-pharmaceutical pandemic defense. Typical SIR models cannot help here, because they pretend mathematically that there is no independent Gap; they are indistinguishable from models that assume that infection can happen only in direct random meetings between Infectious and Susceptible individuals, and that infection always happens at such meetings with a given probability.

The reframing in SGIR models has practical consequences:

- It redirects attention to **measuring the Germ Gap** experimentally.
- It focuses efforts on **increasing the Germ Gap** by equipping populations with the expertise and tools to do so while aiming for success.
- It illuminates **causal mechanisms** that connect practical social-justice concerns to epidemiological outcomes: crowding, poverty, and inadequate housing all *shrink* the Germ Gap, mechanistically explaining why disadvantaged populations bear disproportionate

pandemic burdens ([Caplan *et al.*, 2020], [Mosley *et al.*, 2025]).

The same mechanism works in reverse. Investments in living space, ventilation, and workplace safety *increase* the Germ Gap and so reduce transmission across the whole population, not only the worst-off. In such cases disease protection emerges as a structural side effect of certain forms of equitable development — which makes such investments in humane equal dignity considerably more self-serving than they may seem at first glance. Such investments may shrink some short-term bottom lines, but over the long term pandemic-grade investments in social cohesion are priceless. They build stability that cannot be bought "on demand" once a pandemic is already slow-motion exploding.

When the Germ Gap gets compromised, the pandemic burden often hits those harder who do much essential work for upholding a society. For example, the demographic composition of the U.S. health workforce [National Center for Health Workforce Analysis and Health Resources and Services Administration, 2014], read alongside [Gould and Wilson, 2020] and [Wurth *et al.*, 2020-06-29], shows how systemic racism and economic inequality act as *preexisting conditions* that mechanistically shrink the Germ Gap for African Americans and many other marginalized communities. Yet, where such minorities are essential for health care, systemic inequalities that hit them hard can spiral into much broader problems. Face masks are easy to produce over the short term, better housing can be built over the mid term, but raising, training, supporting, and retaining a generation of good nurses is a much more long-term undertaking that requires corresponding long-term planning. How useful are visions of "a thousand Einsteins and a thousand Mozarts" born in future off-world colonies [Bezos, 2019], if most of present would-be Einsteins on Earth lack the basics they need to develop their gift? More pointedly, if economic pressure forecloses vocations of sustained research, then how is that different from forcing the Einsteins of this world to accept indirectly forced labor if they wish to avoid starvation? The pattern is global and plays out by default unless enough leaders and institutions can be persuaded to walk the narrow path towards developing the mental wealth of all the people in all the nations [Beddington *et al.*, 2008]. This study presumes that winning against pandemics is impossible without such a broad development of mental wealth for all.

This is not a quick fix. There is no broad and easy path to the *gentle kind reasonable* solutions that less powerful countries need to support their not-yet Einsteins (e.g. see [Wintour, 2020-07-16]). Network effects are complicated; even powerful people cannot forever escape the network effects of their actions on the less fortunate. For example, [Wilde, 2018] reports how Stalin's purge of his own physicians left him without competent care when his fatal stroke came — short-term cruelty toward a constructed "out-group" boomeranging on an unwitting perpetrator. The pattern is not new: [John of Ephesus and Pearse, 543CE, 2017] reports how during the Justinianic plague (542 C.E.) the poor died first and how some saw this as the better fate because of the horrors that followed. It is striking to see the SGIR Germ Gap mechanism reverberate across centuries. Yet, a narrow path to beating complicated network effects does not imply that it cannot be found. Once found and defined, it can be explained and taught. Hence, the importance of mental wealth for life-giving decision-making.

Why does the same dynamic keep repeating? [Hare, 2017] and [Hare and Woods, 2020] propose that humans succeeded as a species through selection for prosociality — the human talent that enables coordinated cooperation by trusting others. Stopping a pandemic in mid-flight may therefore depend on something much deeper than face masks, vaccines, and short-term administrative decisions.

This study is built on the working premise that stopping pandemics requires mental wealth and the will to construct reliable work-logic cascades for trusting others in order to extend genuine cooperation. This trust-based cooperation is essential for *gentle kind reasonably* increasing pivotal Germ Gaps through life-giving decision-making for the common good of

everyone. Further analyses in other papers of the Matheo series (see Balospe.com) show that such work ultimately forces a stark value decision, because maximizing infinitely divisible dividends structurally conflicts with best supporting the intrinsic value of indivisible individuals. One of these ultimate priorities must take the lead in how countless conflicts of interests are resolved in a complex world. This study presumes that humane equal dignity is best guarded by respecting the intrinsic value of unique and indivisible individuals and that this value is worth guarding even if at the cost of compromising the maximizing of bottom lines of infinitely divisible dividends.

These abstractions matter, because pandemics cannot exist without infecting indivisible individuals and to do so individual germs must cross infinitely divisible Germ Gaps between individuals. Increasing Germ Gaps can stop pandemics as shown below, but doing so cannot succeed at scale without investing divisible resources accordingly. How to best increase Germ Gaps without overcomplicating requires myriads of more detailed models of the SGIR type. To best guard individuals from infections without overreaching, the studies that build such models must invest their resources accordingly. This study does not aim to construct any specifically applicable SGIR model with all the operational details needed for deployment in any particular demographic. The aim here is more basic.

The goal of this study is to introduce the conceptual framework required for building SGIR models through an example that tests its usefulness for a very simple question:

> Does there exist any biologically reasonable scenario in which the SGIR framework points to realistic changes in Shed, Decay, or Catch rates capable of stopping a COVID-19-sized pandemic without a vaccine?

To this end the SGIR framework is implemented in a model here called "PandemicSociety101". This model tracks counts of individuals through seven stages of infection and includes a simplified testing laboratory and hospital system. Its code is formulated for the high-precision computational workers of the Prototype Evolvix modeling language [Loewe and EvoSysBio Group at UW-Madison, 2015–2026], which supports both deterministic (ODE) and stochastic (SSA) simulation modes to facilitate forecasting time courses of how many individuals of the modeled types exist. Analyzing the deep stack of mathematical models thereby constructed amounts to a constructive existence proof that indeed there exist at least some biologically reasonable scenarios in which a pandemic can be stopped in mid-flight without vaccines, simply by using imperfect face masks at sufficiently large scales.

This report is organized as follows. The next section (2) describes the basic SGIR concept, its underpinning pure mass-action implementation, and the concrete PandemicSociety101 model built on these, before detailing scenarios and parameters derived in Spring 2020 from observing the unfolding Coronavirus pandemic. The Results (3) describe how an uncontrolled pandemic can unfold in this model and how a simple calculation can be used as an effective early warning system. Then simple non-pharmaceutical interventions (NPIs) are described that can avert the brunt of a pandemic even after it started, albeit only if a certain population-wide mobilization can be achieved. Since such a mobilization critically depends on clearly communicating critical information on the state of the pandemic, various non-trivial cognitive traps are discussed that emerge for all who try to observe an unfolding pandemic from inside of that pandemic. The final Discussion (4) summarizes advantages and limitations of SGIR models as actionable frameworks and points to the pivotal importance of a sufficiently well-organized coordination infrastructure for non-pharmaceutical virodefense. The possibility of organizing annual global "Virodefense Olympics" to keep improving pandemic defenses is raised before concluding that working through the implications of the SGIR model offers meaningful contributions to both, the post-processing of what happened during the Coronavirus pandemic, as well as the preparation for helping to reduce the risks for the next pandemic.

2. Model Description

2.1 The SGIR Framework for modeling Germ Gaps

The classical SIR model tracks three types of individuals: Susceptible (S), Infected (I), and Removed (R). Transmission occurs when S and I individuals interact, at a rate proportional to the product S * I. That rate is defined as the much discussed quantity R_0, which offers the following simple intuition: If $R_0 > 1$ an epidemic will be growing (presumably to infect the whole population unless herd-immunity gets in the way); otherwise it will die out locally. This simplicity is in contrast to the exceeding difficulties in dissecting mechanistically what R_0 might be in any specific context (beyond deriving it operationally like a black box parameter estimated from observed doubling times).

The SGIR model introduces two new conceptual types of individuality that combine to form the **Germ Gap** — *Germ* (the individual infectious particles inside) and *Gap* (the finite physical environment that holds them) — with **G** in SGIR standing for that **Germ Gap**. Thereby SGIR models offer a mechanistic way to help to disentangle the mathematical conglomerate parameter R_0, which describes all necessary and sufficient steps of transmission with a single number. The **Germ Gap** represents the physical environment through which virus particles ("Germs") must travel between an infected source individual and a susceptible target individual. (In this paper *Germ Gap* names the technical concept; figures may retain the equivalent label *Gap of Germs*; see Section 4.1 for the naming rationale.) The transmission chain then becomes:

> **Infected** — *(Shed)* —>
> **Gap** — *(survive Decay)* —>
> **Germs** — *Catch* —> **Susceptible turns Infected**

Each step has its own rate to help track the amount of individual germs:

- **Shed rate:** How many Germ particles an infected individual releases per unit time. This depends on infection stage (asymptomatic individuals may shed less or more than symptomatic ones), respiratory activity (e.g. singing sheds more than breathing), and protective measures (e.g. masks reduce shedding).
- **Decay rate:** How quickly Germ particles become non-infectious in the environment. This depends on environmental conditions (e.g. temperature, humidity, UV exposure), surface properties, and active decontamination measures.
- **Catch rate:** The probability that a susceptible individual encounters and is infected by surviving Germ particles. This depends on factors like proximity, ventilation, protective equipment (e.g. masks), and individual immune factors.

The effective Germ Gap can be seen as a complex composite of these three rates in a given context: when Shed or Catch rates decrease, the Germ Gap increases and transmission slows (but when Germ Decay decreases, the Germ Gap decreases). The key insight is that small respective improvements in each of the three rates can strongly compound, potentially achieving even more than multiplicative overall reductions in transmission without requiring single interventions to be perfect in order to be overall effective. Thus, it is possible in principle to push a potential pandemic that has a $R_0 > 1$ for one given set of behaviors in a population to become $R_0 < 1$ once a suitable different set of behaviors is adopted. As the simulation results presented below show, this transformation can apparently even be achieved after a pandemic is well underway, as long as it has not yet run its course.

2.2 PandemicSociety101: the First SGIR Implementation

To test the value of the SGIR framework under biologically realistic circumstances, a model of the Coronavirus pandemic was constructed in in Spring 2020. Called PandemicSociety101, it was motivated by the "fool's hope" of obtaining some modeling insights that might help stop the then emerging slow-motion explosion of Coronaviruses and the COVID-19-induced suffering it caused.

Figure 1 provides an overview of the complete PandemicSociety101 model architecture, showing all compartments, transitions, rate parameters, and the connections between infection stages, the testing laboratory, hospital system, and recovery/death pathways. The model's input scenarios (Scenario 1: Feb 2020, Scenario 2: May 2020) and their parameter configurations are also indicated. For details, see Prototype Evolvix source code (see Supplementary Info).

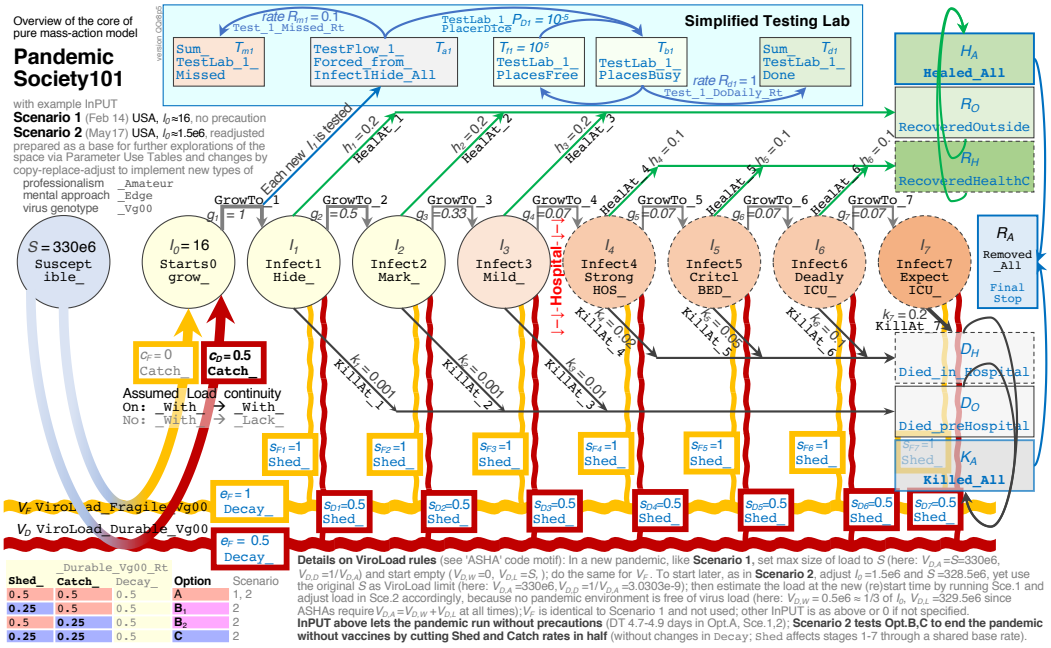

Fig.1: Core model of PandemicSociety101 (*full size* | *list* | download).

Infection stages. The model tracks individuals through seven infection stages following initial virus contact:

Table 1: Stages from Infection to Death

Stage	Duration	Description
Starts0grow	1 day	Virus growth initiated; not yet infectious
Infect1Hide	2 days	Infectious, high shed, no symptoms, hidden status
Infect2Anti	3 days	Infectious, high shed, hidden, antibody-positive
Infect3Mild	2 weeks	Infectious, symptomatic; most individuals recover here
Infect4StrongHOS	2 weeks	Strong symptoms, requires hospital bed
Infect5CritclBED	2 weeks	Critical symptoms, needs hospital bed or dies
Infect6DeadlyICU	2 weeks	Needs ICU or dies
Infect7ExpectICU	2 weeks	Expected death; beyond ICU capacity to save

PandemicSociety101 implements the SGIR concept as a pure mass-action stochastic model for the Prototype Evolvix Compiler modeling language, variant MMs0r3p1 [Loewe and EvoSysBio Group at UW-Madison, 2015–2026]. The Prototype Evolvix compiler transforms the model for the Sorting Direct Method for stochastic simulation ([McCollum *et al.*, 2006] as implemented by [Ehlert and Loewe, 2014]; more accurate for smaller populations but slow for large ones) or the Sundials IDAS solver for corresponding deterministic ODE integration ([Hindmarsh *et al.*, 2005]; faster for large populations at the cost of hallucinating effects of broken-up parts of individuals). All rates are specified in units of 1/day and compartments are always sized such that absolute amounts (counts) of individuals are simulated and never concentrations, in order to ensure a rigorous understanding of all individualities involved.

Individuals progress through the stages defined above and exit the pandemic as either Recovered (outside or in hospital) or Dead (pre-hospital or in hospital). Recovered individuals are assumed immune and cannot be reinfected within the simulation timeframe. To simplify the model, overall population size changes during the time of the pandemic are assumed to be negligible (i.e. no births and no independent deaths of individuals).

Virus tracking via ASHA. The environmental virus load (the "Gap") is tracked using the ASHA (Aggregated State Homogeneity Approximator) framework used here for the first time. It maintains density-dependent dynamics by tracking the number of environmental "places" that are either contaminated ("With") or clean ("Lack") out of a fixed total ("Aces"). This provides proper density-dependent saturation — the environment has a finite capacity for virus, preventing exponential accumulation, which cannot happen in reality. The idea for ASHA grew from the need to be able to tune more parameters of population models than usually exposed in oversimplified models. Examples demonstrate the profound loss of understanding that can result from oversimplified models that pack too much biology into a composite parameter (such as carrying capacity K, [Mallet, 2012]).

The ASHA framework is built on respective concepts (see [Mallet, 2012]), as illustrated and extended in Figure 2 and 3 (see figure captions at the end and Prototype Evolvix source code for more details.).

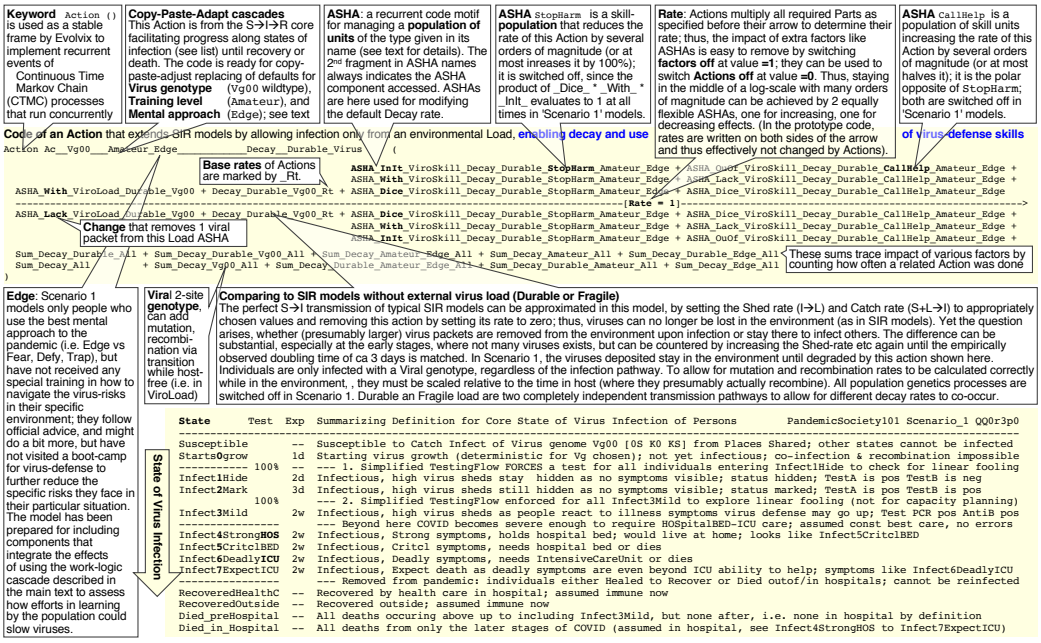

Fig.2: Evolvix Actions (*full size* | *list* | download).

Figure 2 shows how Evolvix **Actions** define the elementary events that move time forward in the model. If the required individual Parts exist, the assumption of random mixing dictates that they will eventually meet randomly. When they do, the respective Action may happen with a certain defined rate per time. If the Action happens, all required Parts instantly disappear to produce new Parts, the products of that Action. The specified Rates for an Action are all multiplied together to define its propensity to happen next. In a stochastic system, where the individuality of Parts cannot be divided up, a Stochastic Simulation Algorithm (like the Sorting Direct Method [Ehlert and Loewe, 2014, McCollum *et al.*, 2006]) calculates the propensities of all actions and then rolls the dice to find the next Action and when it will occur. Then time is moved forward, the changes defined by that Action are then implemented by changing the respective Amounts of all Parts involved in that Action. Finally propensities are recomputed for the next Action. This flexibility in the timing of Actions allows for the indivisibility of individual Parts to be preserved. This contrasts with deterministic simulations, where the time-steps forward are assumed to be primary and the Amounts of Parts are instead treated as infinitely divisible. This can lead to fundamentally different biological outcomes, because "deterministic" populations with only "half an individual" are either extinct in reality or not. See [Ehlert and Loewe, 2014] and the one-page overview in [Loewe and EvoSysBio Group at UW-Madison, 2015–2026] for an introduction to how these approaches contrast.

This formalism is equivalent to the standard mass-action kinetics formalism, albeit implemented with extra care to ensure that the individuality of individual Parts is always respected when simulated stochastically. Moreover, a declarative syntax is used that was designed to make elementary biological Actions easier to check in bottom-up modeling (in contrast to differential equations, which always assume a system-wide overview).

Figure 3 shows how **ASHAs** extend this rigorous mass-action kinetics by assigning unit-sized Places to unit-sized individuals in a population, tracking how many Places ("*Aces*") are *With* or are *Lack*ing a given individual (e.g., a virus contamination), out of a fixed total number of *Aces*. This provides ten variables to define an ASHA state (Aces, Dice, With, Lack, InIt, OuOf, Gain, Loss, Grow, Fade) that controls density-dependent dynamics by defining explicit population biological functions. Such an explicit approach to biouncertainty is preferred here to the implicit bundling of uncertainty into summary parameters (like carrying capacity K, see [Mallet, 2012]), because the resulting clarity may open up new avenues for measuring important quantities or at least help clarify critical biological distinctions. The full ASHA specification is in the Supplementary Prototype Evolvix model code; Figures 2 and 3 provide a visual guide for reading that code.

Actions in this model (as shown in *Fig.2*) also formally interact with two special abstract ASHAs called "StopHarm" and "CallHelp". These parts of the code exist to help model the effects of scaling up population-wide work-logic cascades in case such cascades are constructed while the pandemic is still active. The Discussion points to a related study from 2020 that argues why it is feasible to scale up such work-logic cascades to help coordinate pandemic defenses. For the purposes of this study here all these work-logic cascades in the code are effectively deactivated by setting all respective parameters to 1, which means that they do not bias any Actions in this model in any way. The purpose of this study here was to determine, whether it is possible in principle - given a 100% adoption rate of an imperfect measure - to substantially slow a pandemic *after* it has long left its stochastic stage. If this cannot be demonstrated in principle for at least one biologically realistic parameter combination, then one might argue that there is little point in trying. In that case, without vaccines, the much debated herd-immunity does indeed become the only remaining dim hope for slowing a pandemic. However, as shown below, parameter combinations do exist that inspire the hope that it is indeed possible to stop an unfolding pandemic in mid-flight.

ASHA Places Model for Populations of Unit-Individuals, whether conventional or not, in Places or elsewhere, but always well-mixed

The ASHA model has been developed to explore potential standard ways for managing dynamic populations in Evolvix that combine maximal simplicity, flexibility and scalability. ASHAs randomly assign Places to the unit-sized individuals of a population, but without tracking where those places are, as long as counted correctly. The core idea is that of an

Aggregated State Homogeneity Approximator, ASHA vs
Simulated Explicit Heterogeneity Approximator, SEHA

where SEHAs detail what ASHAs abstract away. Populations at Places `With` have

- a **hard limit** (`Aces`, as tracked with the help of Placer `Dice` and Places `Lack`);
- a **soft limit** (from balancing of Actions `Grow`, `Fade`, `Gain`, `Loss`, all in the hard limit);
- **neutralizers** to hide the ASHA at certain values if used as rate modifier (`InIt`, `OuOf`).

The spatially explicit thinking behind ASHAs was first developed for the Evolvix 'Places' model; it can make code more readable (eg. `PlWith` vs `ASHA_With`). To simplify use of the ASHA code motif, the overviews below are ready for copy-paste-adapting (see code).

Brief	_Frag_	in an ASHA	Usual Example ASHA Name	Explicit ASHA_Places Feature Name
Aces	_Aces_	ASHA_Aces_	ASHA_Aces_MyExampleASHA	ASHA Placed Aces Maximal Count
Dice	_Dice_	ASHA_Dice_	ASHA_Dice_MyExampleASHA	ASHA Placer Dice Probability
With	_With_	ASHA_With_	ASHA_With_MyExampleASHA	ASHA Places With Item Counted
Lack	_Lack_	ASHA_Lack_	ASHA_Lack_MyExampleASHA	ASHA Places Lacking Item Counted
InIt	_InIt_	ASHA_InIt_	ASHA_InIt_MyExampleASHA	ASHA In It Invisible With Scaling
OuOf	_OuOf_	ASHA_OuOf_	ASHA_OuOf_MyExampleASHA	ASHA OutOf Invisible Lack Scaling
Gain	_Gain_	ASHA_Gain_	ASHA_Gain_MyExampleASHA	ASHA Placer Gain for any Lacking
Loss	_Loss_	ASHA_Loss_	ASHA_Loss_MyExampleASHA	ASHA Placer Loss for losing With
Grow	_Grow_	ASHA_Grow_	ASHA_Grow_MyExampleASHA	ASHA Placer Grow by Reproducing
Fade	_Fade_	ASHA_Fade_	ASHA_Fade_MyExampleASHA	ASHA Placer Fade to stop Crowding

Brief	Summarizing Explanation of Feature Definition	During simulations: FIXed or VARiable
Aces X_A	Count of **All Computationally Equivalent Spaces**; sum of all notional Places held in an ASHA, defining a hard limit of all its space; limit enforced by _Aces_=_Shut_+_Open_=_Shut_+_With_+_Lack_, always tracked.	FIX at Quest start (by User)
Dice X_D	Expected frequency of randomly selecting 1 of all existing Aces for some an unspecified Action (without orienting the probability as available for _With_ and _Lack_); _Dice_ = 1/(_Open_), categorically excluding _Shut_.	FIX at start (Must be 1/_Aces_)
With X_W	Current Count of all _Aces_ _With_ a unit Item of the nominal Type defined by this ASHA (Name, Context, and how _With_ is used); works well to slow unwanted Actions, less so for increasing wanted rates.	VAR = 0 or set by User at start
Lack X_L	Current Count of all _Aces_ _Lack_ing a unit Item of the nominal Type defined by this ASHA (Name, Context, and how _Lack_ is used); works well to slow wanted Actions, less so for increasing unwanted rates.	VAR = 0 or set (Must add up to _Aces_)
InIt X_i	Neutralizing factor for _With_ to hide the ASHA in (_Dice_*_With_*_InIt_) products in an Action Rate/Probability that is controllable by this ASHA; use _InIt_=(_Aces_/_With_) for ASHA-free null-models.	FIX = 2 if _With_:_Lack_ is 50:50 ...
OuOf X_o	Neutralizing factor for _Lack_ to hide the ASHA in (_Dice_*_Lack_*_OuOf_) products in an Action Rate/Probability that is controllable by this ASHA; use _OuOf_=(_Aces_/_Lack_) for ASHA-free null-models.	FIX ... usually a good starting point
Gain X_n	Import Actions must change 1 _Lack_ to 1 _With_ and scale Rates by (_Dice_*_Lack_*_OuOf_*_Gain_) to properly import 1 external Item into the ASHA - as 1 random Place Lacking must be found to Place the Gain.	FIX at start; add Gain Action
Loss X_s	Spontaneous Loss or Decay of 1 Item from all Places With takes 1 Action changing 1 _With_ to 1 _Lack_ to properly release 1 Item from the ASHA; scale by (_With_*_InIt_*_Loss_); no Placer Dice search occurs.	FIX at start; add Loss Action
Grow X_r	To properly Grow 1 new Item by Items at Places With, 1 density-dependent Grow Action must change 1 _Lack_ to 1 _With_ at a Rate scaled by (_Dice_*_Lack_*_OuOf_*_Grow_) as 1 random Place Lacking is required.	FIX Slo-Mo Explosion speed, Grow Action
Fade X_e	As density-dependent failure, stress, ... increase in Slow-Motion Explosions, Fade Actions changing 1 _With_ to 1 _Lack_ at Rates scaled by (_Dice_*_With_*_InIt_*_Fade_); this ends all SloMo Explosions.	FIX Slo-Mo Explosion limiting Fade Action

QQr6p2

Careful: if controlling >1 Action by 1 ASHA, the underpinning mechanics must be crystal clear, or else confusing model behavior will be introduced by the extra constraints the ASHA places on the rates of those Actions that are then forced to always share a factor. In turn, many ASHAs for 1 Action are OK, since each ASHA can be switched off independently at will; no extra constraints exist.

Fig.3: ASHA Places Model (*full size* | *list* | download).

In the PandemicSociety101 model virus particles are classified as either **Fragile** (decaying quickly, e.g., airborne droplets) or **Durable** (persisting longer, e.g., surface contamination). Each is tracked by its own ASHA instance. Each infected individual in each infection stage contributes to viral shedding at stage-specific rates.

Simplified testing laboratory. The model includes a simplified testing pathway where 100% of individuals are tested at entry into Infect1Hide and Infect3Mild stages. This design is deliberately simplified to explore the phenomenon of *linear fooling* (see Results) rather than to model realistic testing capacity.

Hospital system. Individuals reaching Infect4StrongHOS and beyond are all assumed to receive hospital care. Hospitals are assumed to be as large and as available as they need to be in order to take good care of all patients in the respective stages. The model tracks hospital and ICU occupancy under these idealized conditions and distinguishes between deaths occurring before hospital admission and deaths in hospital.

2.3 Scenarios and Parameters

Scenario 1 (Uncontrolled, 2020-02-14): 16 infected individuals in a population of 330 million (US). No behavioral change, no interventions. Virus transmission parameters reflect baseline SARS-CoV-2 characteristics. This scenario was calibrated to the observed US doubling time of approximately 3.25 days in the early phase in 2020. The resulting model parameters lead to an emergent doubling time of about 4.8 days as measured from the simulated model output.

Scenario 2 (Face-masking, 2020-05-17): Starting from 1.5 million infections (ca. 90,000 deaths) in a population of 330 million, with three sub-options:

- **Option A:** No change in Shed, Decay, or Catch rates (baseline). Then the pandemic continues as in Scenario 1.
- **Option B:** 50% reduction in *either* the probability of virus Shed *or* Catch rate. This represents partial facemasking (or equivalent NPI adoption).
- **Option C:** 50% reduction in *both* Shed probability *and* Catch probability simultaneously. This represents fully coordinated facemask adoption at the defined level of efficiency (or equivalent NPI adoption by other means).

The full model specification, including all parameter values and ASHA configurations, is available as Supplementary Material (Prototype Evolvix source code, ~3,900 lines; need a 4K monitor for efficient reading). To generate the raw results for the figures shown, the corresponding parameter combinations in that file need to be switched on or off, respectively. All computing work, results analyses, and figure compositions were completed by Loewe by mid 2020.

3. Results

3.1 Scenario 1: Anatomy of an Uncontrolled Pandemic

Without interventions, the PandemicSociety101 model simulates a pandemic that infects approximately 289 million people of the 330 million in the US population modeled. Of these 13.8 million individuals die.

Three stochastic simulation replicates closely track the deterministic prediction made by ordinary differential equations (*Fig.4*). This confirmes that for a randomly mixing population of 330 million, stochastic effects are minimal when starting with 16 individuals. The only time when such stochastic effects are even observable is during the earliest phase when infection counts are small enough for chance effects to slightly delay or accelerate further spreading.

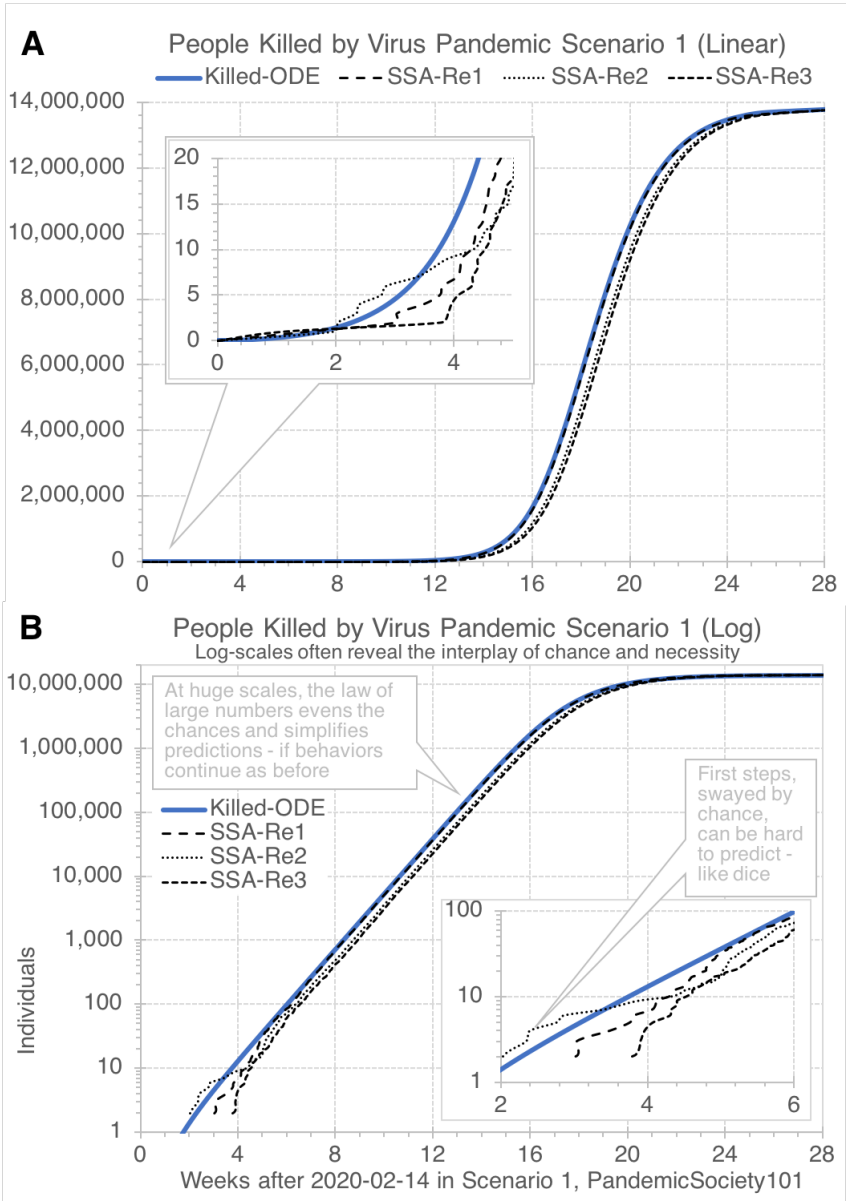

Fig.4: Pandemic deaths in default Scenario 1 on linear and on log scales (*full size* | *list* | download).

Fig.4A shows the pandemic growth on a linear scale, which is only helpful during the later stages of a pandemic. Since pandemics are mostly driven by multiplicative growth, Fig.4B shows the same simulation results on a log scale, which is more informative during the early stages of multiplicative growth.

Communicating clear and present danger from slow-motion explosions. Much thought was given to how the tricky multiplicative dynamics of pandemics might be translated into clearer language for people who are not used to dealing with the underpinning mathematics. A simple analysis of general audience pragmatics and semiotics of the respective mathematical language revealed a major barrier to all who wish to use its standard terminology to communicate the urgency of pandemic actions. The keywords to translate here are "exponential growth". Both words seem to make sense to most people. Unfortunately they map intuitively to the wrong notions in the unreflected use of most people. "Growth" is a good thing most of the time in the mind of most people, and "exponential" maps to "a lot". Thus, "exponential growth" points intuitively to "a lot of a good thing" for most people, unless they think through the context, which tells them that fast growth of a dangerous virus is *not* a good thing. To find a way around this problem, the term "slow-motion explosion" was defined for describing the growth of a pandemic. It maps to the same underpinning chain-reaction that drives any explosive growth from nuclear chain reactions to pandemic transmission growth. All these are mathematically describable by exponential functions. Including "slow-motion" highlights the fact that pandemic times of response to changes in behavior are closer to response times of steering a container-ship than a race car. Ending with "explosion" highlights the fact that the impact of the respective shockwave will come nevertheless and is in principle contained by the space in which it happens. That this space is usually best described in multiplicative terms is only one of several unusual aspect of how pandemics work.

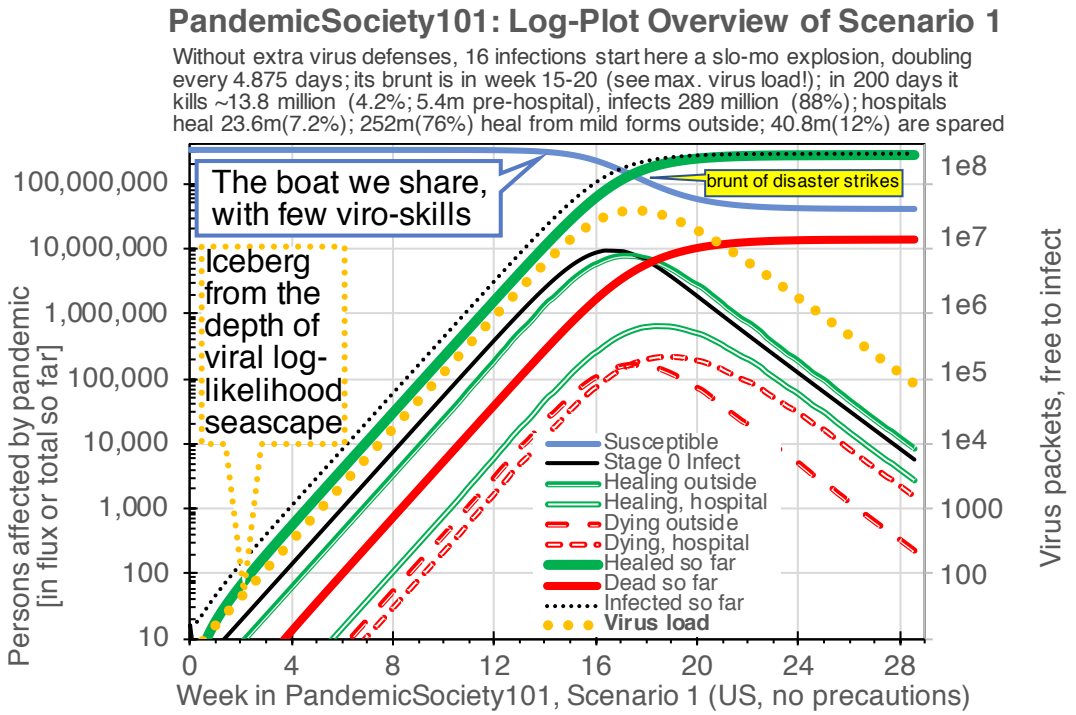

Fig.5: Log-plot overview of uncontrolled Pandemic Scenario 1 (*full size* | *list* | download).

To underscore the multiplicative nature of pandemic slow-motion explosions, it is generally useful to show how they unfold on a log-scale. Hence, many figures here are shown on a log-scale. The linear scale tends to be most useful for the last few moments before a

slow-motion explosion crashes into its hard limiting space factor (e.g. it is impossible to infect more individuals than exist in a given population).

Fig.5 shows an overview of several interesting quantities for tracking slow-motion explosions on a log-scale in the PandemicSociety101 model without interventions. In this model the pandemic infects approximately 289 million people (88% of the 330 million population) and kills approximately 13.8 million (4.2% overall; 5.4 million pre-hospital, with 23.6 million (7.2%) healing in hospitals and 252 million (76%) recovering from mild forms outside hospitals). Approximately 40.8 million (12%) are spared infection entirely.

In *Fig.5* one may think of the virus load as an unexpected "iceberg" emerging from the deep, which drives infection rates upward while remaining invisible on linear scales for most of the time - and hence on collision course with the ship of the civilization it attacks.

A careful comparison to *Fig.4* shows that in week 1 to week 14 **on a linear scale, the virus seems to do "almost nothing" during the period when in reality it is most active**. In that time it establishes the ultimate size of the slow-motion explosion it causes. By the time infections become visible on a linear plot, the "exponential growth" phase is nearly complete and the size and punch of the slow-motion explosion have been almost completely determined.

This linear-vs-logarithmic perception gap is a fundamental barrier to broader public understanding of pandemic dynamics. While it is easy to explain in principle, there is such a long list of detailed implications and complications that even experts get easily tripped up (as other Results in this study show).

3.2 The HalfMax method as an early warning system for pandemics

How even experienced modelers can be fooled. *Fig.6* illustrates a sobering point about the deceptive nature of exponential growth on linear scales. This figure, from Loewe's earlier work on stochastic simulation algorithms (Fig.7a in [Ehlert and Loewe, 2014]), shows 100 stochastic simulations of a simple unbounded exponential growth model starting with 10 individuals.

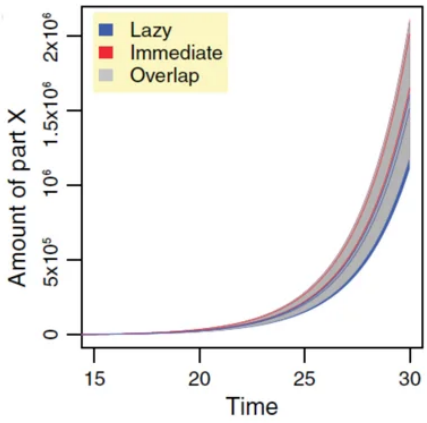

Fig.6: Slow-motion explosions are easy to miss (*full size* | *list* | download).

On a linear scale the resulting slow-motion explosion shows the characteristic "hockey stick" pattern: the population appears invisible for a long time, then suddenly explodes. These simulations were produced in 2014, years before COVID-19. Loewe had extensive experience with interpreting systems that are much better understood on multiplicative log-scales — he was

eventually even exploring how a multiplicative ("non-Newtonian") calculus, in which products and ratios take the structural role of sums and differences ([Grossman and Katz, 1972], [Grossman, 1983]), might offer a more natural language for such growth in Evolvix. Yet when he read US reports of 16 Coronavirus infections on 2020-02-15 — a number strikingly close to the 10 individuals that reliably triggered well-defined exponential growth in his 2014 simulations shown in *Fig.6* — even Loewe failed to realize the significance of that alarming information. What chances do others have who live much more in the linear world If even a researcher whose professional work centered on exactly these multiplicative dynamics can be fooled by linear scales that make deterministic slow-motion explosions look like "nothing is happening". This personal experience underscores the systemic nature of linear fooling: it is not a failure of knowledge but a failure of perception ("realizing") that affects everyone, including even some experts who should theoretically and practically know better.

HalfMax method. It is pivotal to mitigate this perception problem in order to increase the reaction time remaining for behavior modification before the brunt of a pandemic hits. Like other early warning systems for natural disasters, such as tsunamis and tornados, there is no exact way to predict the precise amount of damage that will result from doing *nothing*. Except it is clear that maximal damage will result from not seeking shelter, which is equivalent to no behavioral modification when a pandemic hits. Yet, even though pandemics move slower than tsunamis or tornados, time is of the essence. To communicate that urgency it is essential to have a **reliable early-warning system** for calculating how much time might still remain if the current behavior and the current germs were to continue without notable changes.

To this end Loewe developed the HalfMax-method, a quick rule-of-thumb method that only needs a pocket calculator for helping a broader audience without access to sophisticated simulation models to quickly translate a reported doubling time T_{Doubling} into an expected waiting time before the brunt of a pandemic will hit T_{HalfMax} — if nothing changes, i.e. all rates stay as they are and a random mixing population without changes in behavior can be assumed. The HalfMax method is not about precision; it's about triaging whether an emergency response is needed and how much time may remain to organize it.

The HalfMax-method builds on the basic understanding that all pandemics are slow-motion explosions that follow the logistic growth curve, which predicts that absolute growth will be fastest at half of the maximal capacity, before it starts to slow down again.

This allows for a simple doubling-time arithmetic to estimate the HalfMax point when at most 50% of the population will be infected and hence infection rates will be highest before they naturally slow down as susceptible individuals get increasingly rare.

The point in having such a simple "pandemic count-down" timer at hand is in distributing as best possible the work required to increase Germ Gaps such that the overall size of the pandemic can be reduced before it is too late. Interventions after the HalfMax point will have significantly less impact and their effectiveness may be difficult to distinguish from an expected natural decline in infection numbers.

Like with tsunami-warnings, if everyone can predict the worst based on observed data, everyone can help to reduce the impact. It only takes a pocket calculator to compute T_{HalfMax} a HalfMax waiting-time forecast. Therefore the HalfMax-method is easily implementable and checkable where it matters most: at places of decision, where behavioral recommendations are made that affect the Germ Gap. If a rational explanation is given and people can check it, a given mitigation strategy is much more likely to succeed, even if it requires some sacrifices in comfort.

As the value is not in a precise point estimate, a min-max range should always be given. The greatest value of the HalfMax method is in the 'street-wise math' that helps to reduce the 'blind

faith' that many felt was required of them in understanding the threat from the Coronavirus pandemic.

The HalfMax core equation is this:

$$T_{\text{HalfMax}} \approx T_{\text{Doubling}} \times \log_2 (N_{\text{HalfMax}} / N_{\text{NowInfected}}) \text{ [Eq.1]},$$

where N_{HalfMax} is half the number of all susceptible individuals (~165 million in the US) and $N_{\text{NowInfected}}$ approximates how many have already been infected by now.

The purpose is to quickly translate a key observable (like a 5-day doubling time) into actionable intelligence offered by a rough early-warning forecast. That is why it can be thought of as a pandemic equivalent for a tsunami early-warning system.

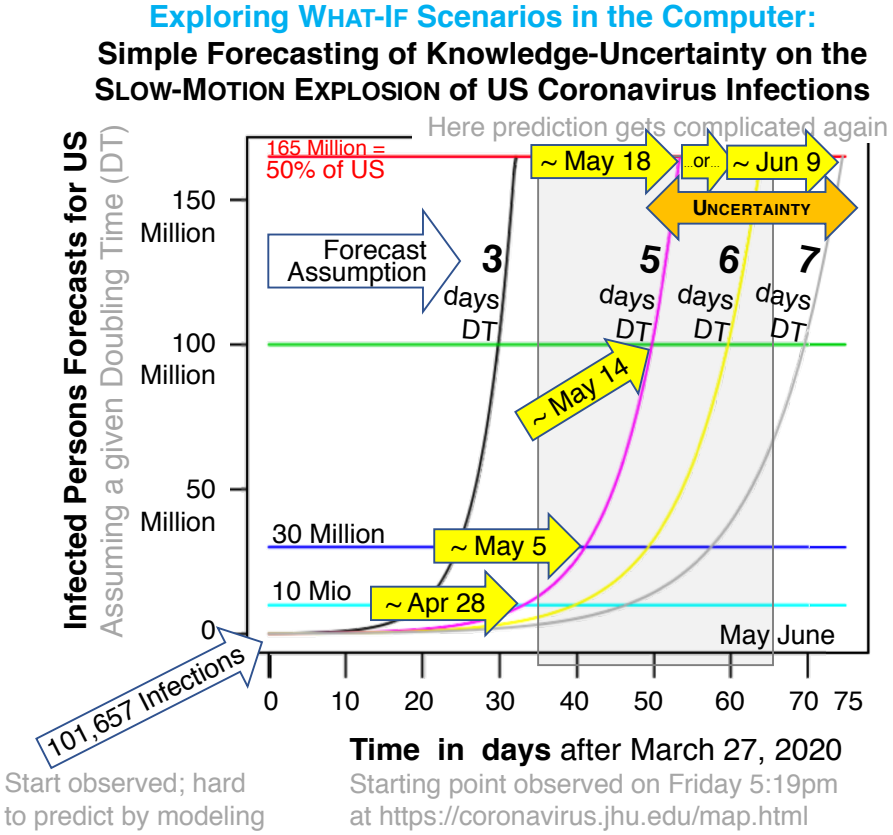

Fig.7: HalfMax early-warning in Loewe's 2020-04-01 pandemic forecast **Fig.7:** HalfMax early-warning in Loewe's 2020-04-01 pandemic forecast (*full size* | *list* | download).

Applying the HalfMax forecast to his own situation in 2020, Loewe calculated the following rough numbers as reported in *Fig.7*:

$$T_{\text{HalfMax}} \approx 32 - 75 \text{ days} \approx 3\text{-}7 \text{ days} \times \log2 (165 \text{ mio} / 0.1 \text{ mio}) \text{ [Eq.2]},$$

using a point estimate of $T_{\text{Doubling}} \approx 5$ days to forecast $T_{\text{HalfMax}} \approx 53$ days after 2020-03-27, the day Loewe started to take his first serious look at the Coronavirus pandemic (with 101,657 reported infections). The shock of realizing what this graph meant moved Loewe's research trajectory to set aside his other career plans in order to explore how such existential threats

might be mitigated. Without understanding Fig.7, Loewe would have never completed the PandemicSociety101 model, even though, arguably, he had accumulated all essential insights for completing such an SGIR model by about ≈ 2016-2018.

Naturally, these HalfMax forecasts assume continued random mixing with no changes in behavior whatsoever. Yet, as well known, drastic changes in behavior are to be expected and did occur. To examine the usefulness of the HalfMax method given such changes, its forecasts were compared to actual CDC data through May 2020 (*Fig.8*).

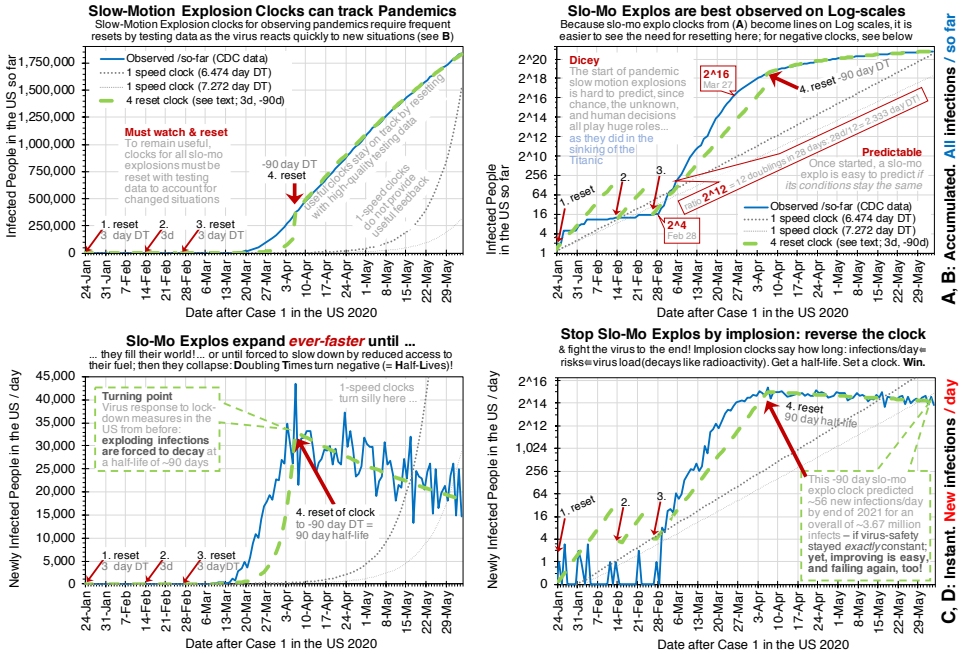

Fig.8: Testing the HalfMax early-warning method in a real pandemic (*full size* | *list* | download).

Fig.8 shows that the observed trajectory is predicted in useful ways between bounds that repeatedly reset the HalfMax clock to account for observed changes in behavior that affects the Germ Gap.

These predictions suggest that the HalfMax method can be harnessed to help not only to "flatten the curve", but even to "end the curve" if sufficiently effective non-pharmaceutical interventions can be found in time.

The black-box nature of R_0 in typical SIR models offers little hope for such interventions. However, the mechanistic breakdown of R_0 that is offered by SGIR models opens principled avenues for fighting pandemics by reducing the respective Shed and Catch rates (e.g. via face-masking) as well as increasing Decay rates for viruses (e.g. via air-filters and surface cleaning). Thus, SGIR modeling might offer a sliver of hope, if only a "fool's hope" in a dark time.

3.3 Scenario 2: Stopping a Pandemic with Face-masks

The maybe most startling result of this study is shown in *Fig.9*, where 3 pandemic forecasts are compared, all starting from 1.5 million infections on 2020-05-17 , each simulating different behaviors after 90,000 deaths had already occurred.

The two NPI options simulated are best compared to two different types of use of face-masks, both of which produce dramatically different outcomes, especially when combined. Table 2 compares the outcomes, assuming that only 50% of all infectious particles are filtered. Measurements of face-masks suggest better filter efficiency (e.g. 74% to 90% measured [Asadi *et al.*, 2020]), so the parameters chosen here appear workable in practice.

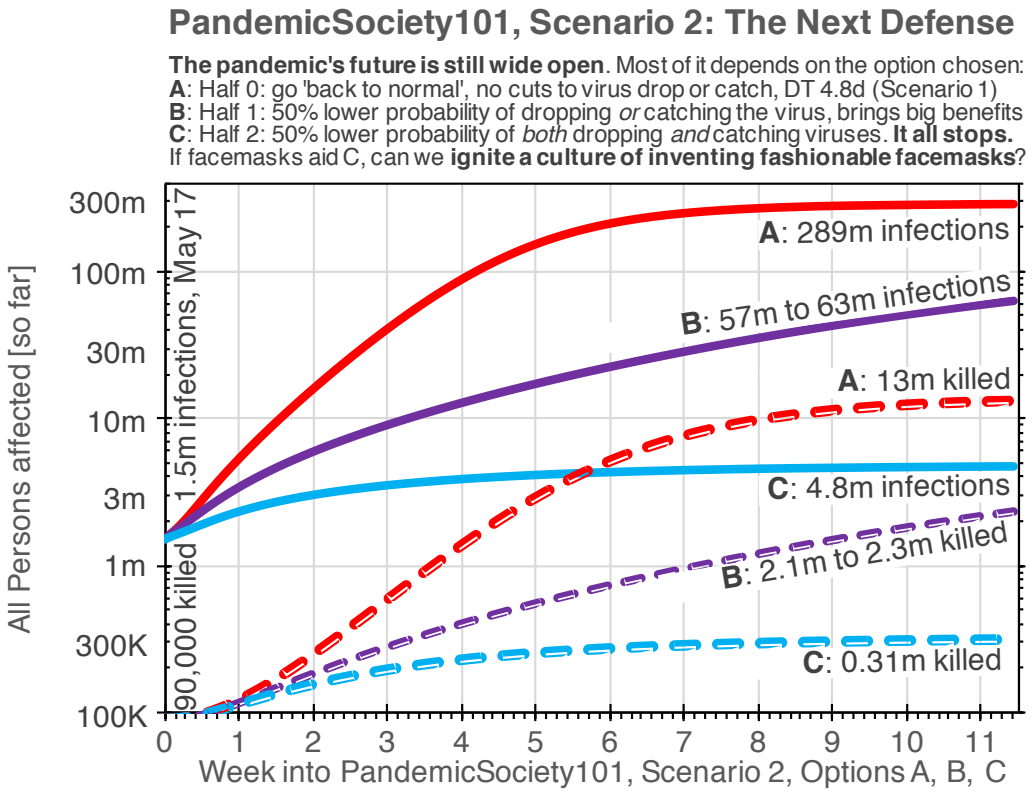

Fig.9: Scenario 2C stops a pandemic in mid-flight with face masks (*full size* | *list* | download).

Table 2: Scenario 2 for how to stop a pandemic without vaccines.

Option	Total Infections	Total Deaths	Face-mask adoption effects
A	~289 million	~13 million	No face-masks (baseline)
B	57–63 million	2.1–2.3 million	50% reduction in EITHER Shed OR Catch rates
C	**~4.8 million**	**~310,000**	50% reduction in BOTH Shed AND Catch rates

The progression from A to B to C demonstrates the more than multiplicative compounding effect of combining interventions. A single 50% reduction in Shed OR Catch rates (Option B) achieves a 4.6–5.1-fold reduction in infections. Combining both 50% reductions (Option C) achieves a **60-fold reduction** in infections — far more than any simple intuition would predict from doubling the interventions.

These strong compounding effects are the quantitative foundation for the success of Germ Gap based interventions. The explicit modeling of density-dependent effects due to the Germ Gap as tracked by the ASHA framework predicts an even stronger compounding than simple multiplicative effects (2x2 = ~4-fold or 4.6x5.1 = ~23-fold). These Germ Gap effects are the reason for why even without any intervention the pandemic in this SGIR model does not approximate 100% infection: eventually the probability of getting enough viruses across the Germ Gap becomes so low that the remaining Susceptibles can no longer be reached. The non-pharmaceutical interventions that increase the Germ Gap as reported in Table 2 simply lower that probability enough, such that the pandemic "simply goes away" as famously claimed.

These results are consistent with independent modeling [Stutt *et al.*, 2020], showing that face-masks combined with lockdown measures can effectively manage the pandemic when adopted broadly. The SGIR framework provides a detailed mechanistic explanation for *why* such combinations are so effective: the details are in a probabilistic game of dice, played out in real-life in the Germ Gap.

This appears to be a case where independently working together is greatly rewarded by the mathematics underpinning the reality of pandemics: those who wear a mask while infected reduce their Shed-rate for the benefit of everyone. At the same time, those who also wear a mask despite not being infected, will reduce their Catch-rate. When both work together, their combined reward in safety gets a mathematical extra-safety bonus, simply for working together.

Hence, despite reducing the *product* of Shed and Catch probabilities only by four when cutting both probabilities by half, the overall effect is amplified into the observed 60-fold overall reduction by the density-dependent effects tracked by the ASHA framework.

Seeing this Figure 9 by mid 2020, Loewe concluded that his "fool's hope" would not be a real hope if it was impossible to show for biologically reasonable parameter combinations that there exists a realistic road to success such as in Scenario 2C. This insight, however, fundamentally transforms the battle in fighting a pandemic from the proximal cause ("fight germs") to a more distal cause ("return *gentle kind reasonably* to reality"). Thus, if three apparently realistic manipulations of probabilities for shedding, decaying, or catching the virus can actually stop the pandemic, then a pandemic becomes a battle-cry of reality for improving *gentle kind reasonable* engaging in order to invite everyone to *gentle kind reasonably* returning to reality. Note, how all three properties of this life-trifecta are essential for success; yet explaining how it works in detail, what happened to that fool's hope, why it existed in the first place, and how to best learn from it all are topics beyond the scope of this study and require in-depth analyses of many other topics (see the Matheo study series and the broader work at Balospe.com).

While the remainder of this study analyzes a few general observations of interest to pandemic experts, Loewe's 2020 work on meta-pandemic topics related to fighting pandemics by improving *gentle kind reasonablness* is to be discussed in his forthcoming Matheo-b20 study (to be published at Balospe.com/en/study once sufficiently complete).

3.4 Linear Fooling: A Dangerous Cognitive Trap

The model's simplified testing laboratory reveals a phenomenon that may be called *linear fooling* (*Fig.10*). When testing capacity is limited to a fixed number of tests per day, and under the (big) assumption that all symptomatic individuals get tested, the following sequence will occur:

1. **Early phase:** Testing capacity exceeds demand. All infections are detected. Statistics appear reliable.
2. **Transition:** Infections grow exponentially and eventually exceed testing capacity. After this point, testing detects a *constant* number of infections per day (the capacity limit), regardless of actual growth.
3. **Misleading plateau:** On a linear plot, daily confirmed cases appear to stabilize or even decline, creating the illusion that "containment is working" precisely when infections are growing fastest.
4. **Retroactive revelation:** When the pandemic wave passes and testing capacity again exceeds demand, the true scale of missed infections becomes apparent — if sufficiently deep analyses are conducted. Unfortunately, by then the damage is done.

The linear fooling effect is not a bug in testing strategy; it is a mathematical consequence of limited capacity encountering exponential growth. It is disastrously easy to fall for because it confirms a desirable narrative (the pandemic is under control) at precisely the moment when vigilance is most needed.

On a log scale, the effect is clearly visible as a deviation from exponential growth in the testing curve (Figure 10C), but most public health dashboards display data on linear scales, where the deviation is invisible and efforts to "flatten the curve" seem to succeed.

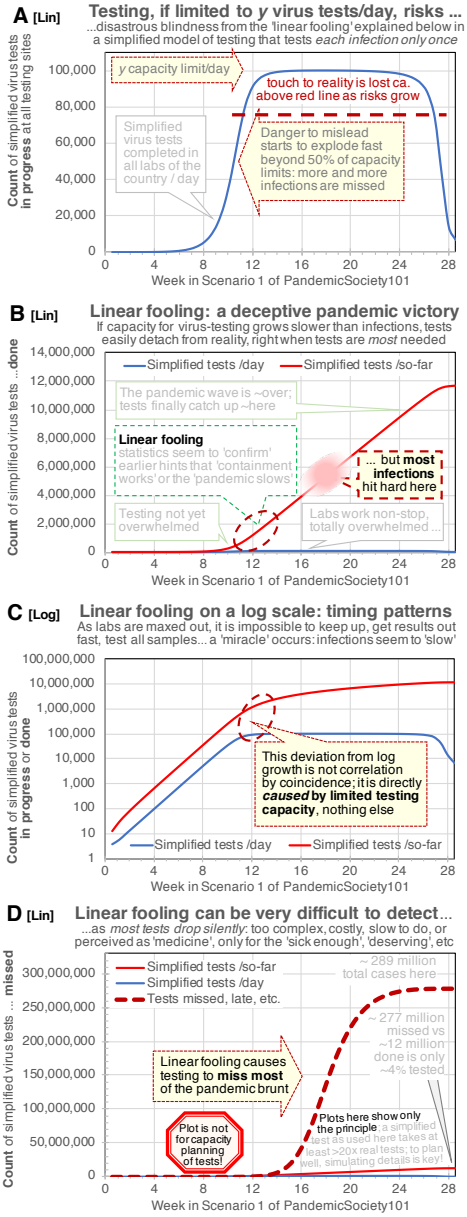

Fig.10: "Linear fooling" by limited testing can create death traps (*full size* | *list* | download).

A note on potential misuse. Linear fooling does NOT mean that testing is useless — it means that testing must be scaled to match exponential growth, and that public health dashboards should routinely display data on logarithmic scales where the limits of testing capacity become immediately visible. The point is not that "the numbers were fake" but that limited capacity creates a structural blind spot that affects everyone, including decision-makers acting in good faith. Awareness of this structural trap is the first step toward designing testing infrastructure that remains informative even during exponential surges. Needless to say, such logarithmic displays are useless unless they are widely understood by the general public; hence Loewe's pivot to questions of meta-pandemic engagement for shaping gentle kind reasonable information ecologies for broader audiences.

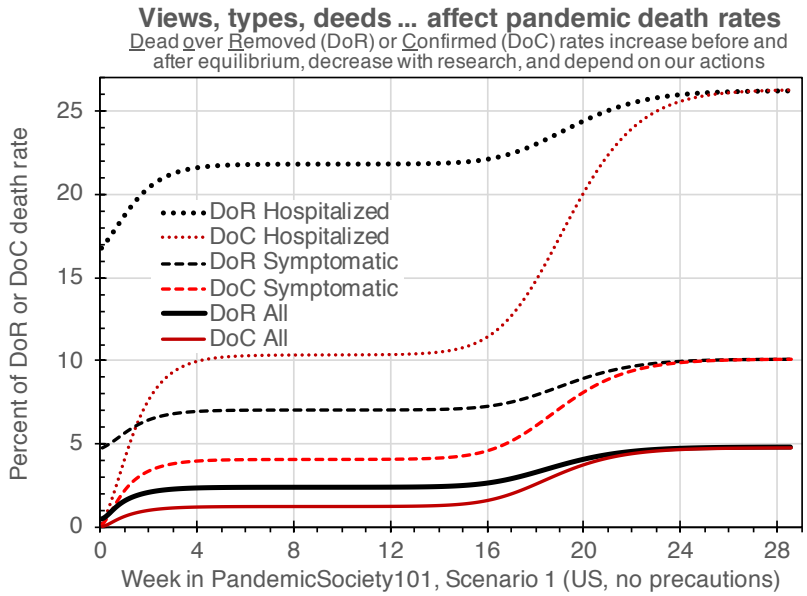

Fig.11: Diverse death rate dynamics over time (DoR, DoC) (*full size* | *list* | download).

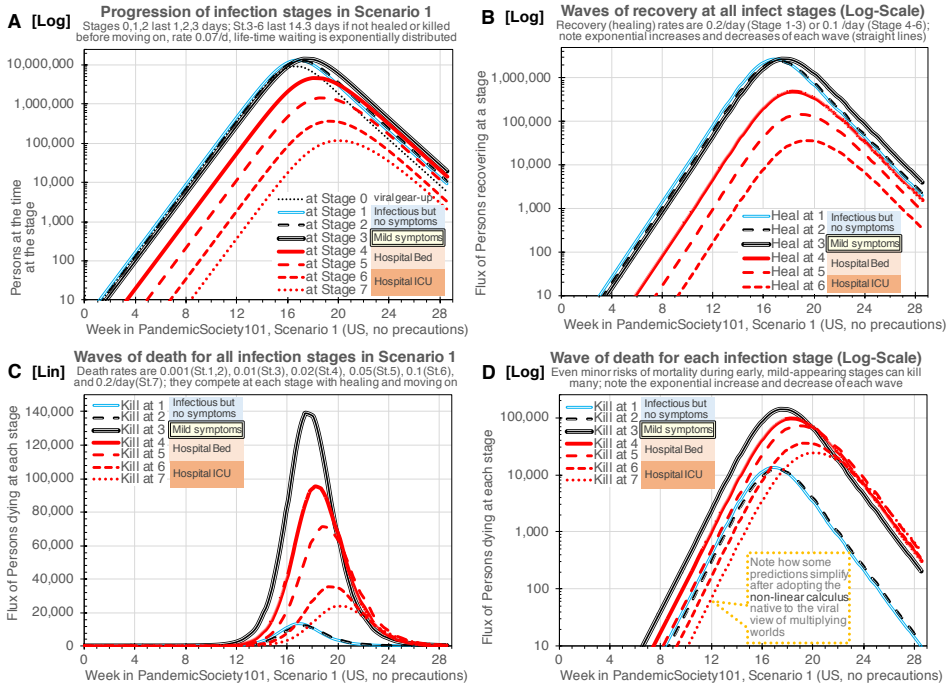

Fig.12: Stage-specific infection, recovery, and death waves in Scenario 1 (*full size* | *list* | download).

Table 3: Death over Confirmed or Removed rate measures observed in Fig.11, Fig.12)

Measure	Definition	What it shows in the model
DoC All	Dead (so far) / Confirmed (so far)	Starts near ~1% in weeks 4–16, then rises to ~4.8%. Closest to early-pandemic IFR estimates. The rise is a **timing artifact**: deaths lag behind confirmations.
DoR All	Dead / (Dead + Recovered)	Starts near ~2%, rises to ~5%. Removed is defined as Dead + Recovered and leads to timing artifacts that are different but comparable to DoC.
DoC Symptomatic	Dead / Confirmed (stage 3+)	~4% in weeks 4–16, rising to ~10%. Higher because pre-symptomatic stages are excluded from denominator.
DoR Symptomatic	Dead / Removed (stage 3+)	~7% equilibrium, rising to ~10%.
DoC Hospitalized	Dead (so far) / Confirmed (stage 4+)	Starts near ~10%, rises to ~26%. Ratio of deaths over confirmed *hospitalized* cases. Timing artifact strongest here as most Confirmed patients Recover before their peers die.
DoR Hospitalized	Dead / Removed (stage 4+)	~22% equilibrium, rising to ~26%. Ratio among hospitalized patients only — does NOT represent overall population death rate. The high numbers are due to the inevitable sampling bias as only the most severe cases make it to the hospital and severe cases are also more likely to die.

3.5 Fooling by Treacherous Death Rate Dynamics

The model also reveals another form of fooling that complements linear fooling and might be seen as a strange form of time-travel: the *apparent* death rate changes dramatically throughout the simulated pandemic depending on *when* and *how* it is measured, even though probabilities of individual fates do not change and the model assumes constant best care is available at all stages (i.e. there is no collapse of healthcare systems). *Fig.11* shows an overview of how potential systemic measures of death change over time in Scenario 1, based on the waves in which individuals pass through the seven stages of disease in the model (see *Fig.12*).

The model's apparent "Infection Fatality Rate" IFR is not an input parameter — it is an *emergent property* of the stage-specific death, healing, and progression rates competing at each stage. Figure 11 plots several observable death rate measures over time. The odd relation to "time-travel" is generated by the delay between the infection event that ultimately sets an individual on a potential path to premature death and the execution of that ultimate fate. Almost by definition this is unavoidable when first collecting data (see Table 3).

The key insight: all these measures *change over time* even though the model's underlying rates are constant. The rising trajectories are caused by the **timing mismatch** between infection confirmation and death: during exponential growth, most confirmed cases have not yet reached their final outcome, making the apparent death rate misleadingly low. After the wave passes, the accounting catches up. This timing mismatch is itself a form of "fooling" complementary to linear fooling: just as limited testing creates an illusion of pandemic control, the timing delay in death statistics creates an illusion that the pandemic is less deadly than it actually is during its most active phase.

The model's death rate parameters were calibrated to data available in early-to-mid 2020, when observed death rates were substantially higher and more uncertain than later estimates. Figure 12 documents this empirical fog as of 2020-06-28 for the US state-level data and international data that Loewe could collate to the best of his abilities from various public sources at the time. Note how widespread the nearly ~20-fold spread of obtainable estimates was and how some of the timing and sampling issues discussed might explain that spread. Loewe concluded at the time that obtaining more accurate estimates at the time was not possible without either substantial institutional support or much more time. Note the computational challenges with quantifying bioncertainty that hold huge sway over such work (see extended caption for *Fig.13* at the end). Hence, the model parameters chosen represent Loewe's best evaluated estimate in a good-faith effort to capture the threat as it was understood at the time.

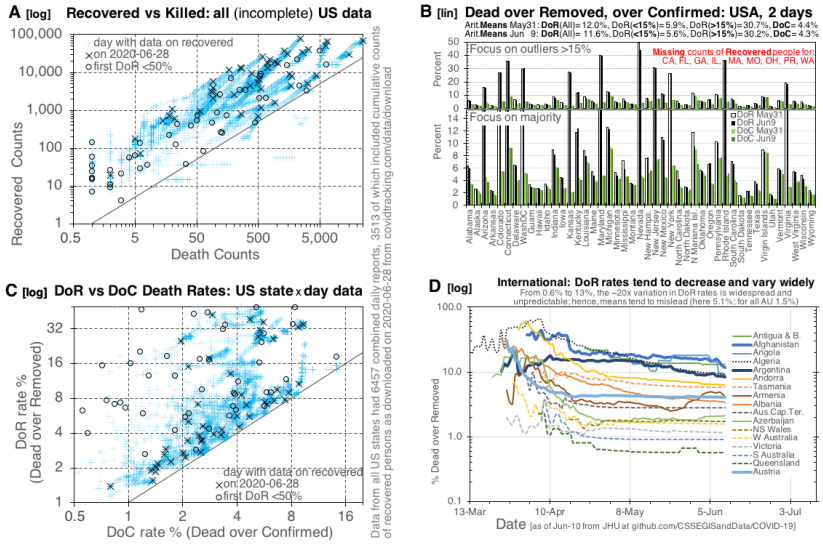

Fig.13: Variation in COVID-19 death rate calculations (2020-06-28) (*full size* | *list* | download).

3.6 Scale Invariance: From Prison to Planet

Simulating the PandemicSociety101 model Scenario 1 dynamics across seven orders of magnitude of population size show **the same underlying logic governs outbreaks at every scale** (*Fig.14*): a 1,000-person prison (~43 deaths; 3 SSA replicates giving 33, 44, 45 against the ODE mean), a 0.5-million county (~21,000 deaths), the US at 330 million (~13.8 million deaths), and the world at 7.8 billion (~326 million deaths).

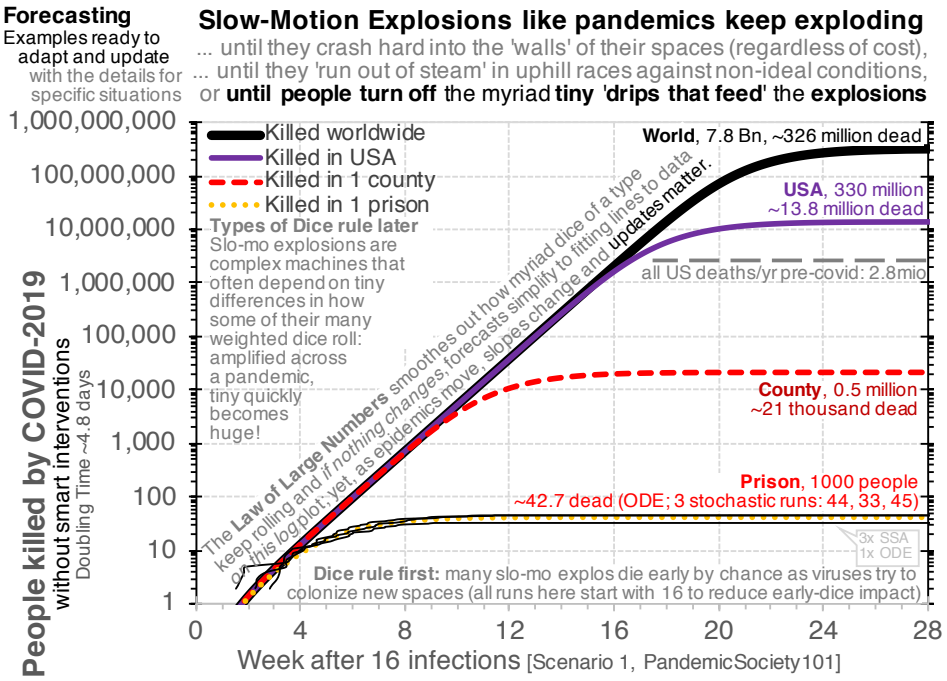

Fig.14: Pandemic slow-motion explosion scales from local to national and global (*full size* | *list* | download).

What changes with scale is not the mechanism but the relative importance of stochastic variation. At local prison scales, individual dice rolls dominate outcomes — a small outbreak can stochastically burn out OR stochastically escape, and three SSA replicates differ substantially. At national and world scales, the Law of Large Numbers smoothes individual variations into an essentially deterministic trajectory, and stochastic replicates become indistinguishable from the ODE solution.

This scale invariance has two operational implications:

- **Timely local responses matter.** Small-scale outbreaks are stochastic, which cuts both ways: they can fizzle out on their own, but they can also escape containment with no warning. Local interventions delivered during this early stochastic phase have the most leverage per unit effort. Thus, locals working with global experts have the best chances of success.
- **Coordinated global infrastructure matters.** Once an outbreak reaches the deterministic regime of the Law of Large Numbers, only population-scale reductions in Shed, Decay, and Catch can stop it — through the strong compounding shown in Scenario 2. However, to make that work depends again on many timely local responses, which, again, depend on global experts and locals working together.

Hence, both regimes require a bidirecional flow of information from local insight to institutional capacity to local action. Unfortunately, current infrastructures serving that purpose currently are either not ideal or too patchy at a global scale. The companion study (upcoming Matheo-b20) outlines work-logic cascades for evolving and strengthening information infrastructure designs that stand a real chance to actually deliver the virodefense capabilities required to avert a pandemic like the one caused by the Coronavirus by never allowing it to become as big. This vision includes scaling up a ResearchCity for organizing Virodefense Olympic Games designed to strengthen virodefense as a common good for everyone.

4. Discussion

4.1 The Germ Gap as an Actionable Framework

The SGIR model introduced and tested here reframes pandemic defense around a single concept: **increase the Germ Gap** (*Fig.15*). Every NPI — face masks, distancing, ventilation, hand hygiene, surface cleaning — acts by increasing one or more components of the Germ Gap. This reframing has several advantages over the traditional focus on the reproduction number R_0:

- **Mechanistic clarity:** R_0 is an aggregate outcome; the Germ Gap identifies the specific levers (Shed, Decay, Catch) that humans can manipulate.
- **Additive intuition:** While transmission compounds multiplicatively (which is non-intuitive), the Germ Gap can be communicated additively: "do three small things and the combined effect is large."
- **Social justice connection:** Crowding, poverty, and inadequate housing shrink the Germ Gap. Investments in equitable living conditions are simultaneously investments in pandemic defense.
- **Reusable value:** Unlike vaccines or antivirals, Germ-Gap-increasing measures (better ventilation, more living space, hygiene infrastructure) provide benefits even when no pandemic is active while simultaneously guarding against yet unknown pandemic threats.

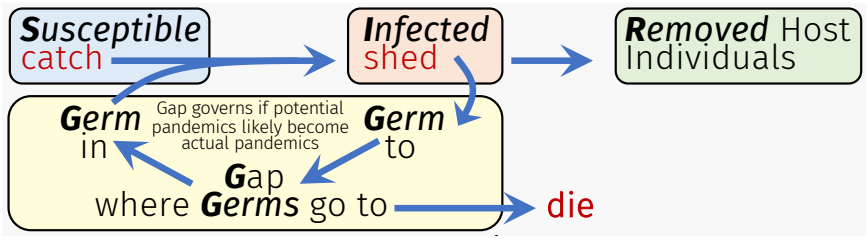

Fig.15: Simple overview of the **Germ Gap** — the "G" in SGIR models (equivalent to *Gap of Germs*) (*full size* | *list* | download).

Naming. *Why "Germ Gap"** rather than "Gap of Germs"?* English already supplies productive X-gap idioms (like *wage gap*, *gender gap*, ...) — in which the construction means a *separation* concerning X, the most crucial meaning here. Germ Gap inherits this idiomatic separation-reading on first hearing. *Gap of Germs* invites the swarm/pool reading ("many Germs in the Gap"), which is not wrong, but secondary and technical. The verb construction that matters most for policy prose is to *increase the Germ Gap*, which parses unambiguously as enlarging a protective separation, while *increase the Gap of Germs* may also read as growing a germ population. — To scientists dealing with arcane technical definitions all day such nuance may not matter much (as they follow definitions given), but to others, who do not consume technical definitions for a living, such nuance may make all the difference between a first im-

pression that is clear or confusing. Like in immunology, so in pandemic communication: there exists no chance for a second first impression.

A related precedent may be seen in the mid-2020 WHO decision to switch from calling for *social distancing* to *physical distancing* because the natural English reading of *social* worked against the public-health message ([Wasserman *et al.*, 2020]).

Given how SIR models have become a paradigm for epidemiology (and the confusion from not treating the Germ Gap explicitly), defining the clearest possible anchor term for the "G" extension in SGIR models is of paramount importance for *gentle kind reasonable* virodefense. This anchor term must be able to support lasting international debates by choosing a phrasing whose default reading *best aligns* with the most critical technical meaning for primary action. *Gap of Germs* is not wrong as the swarm/pool reading remains important inside the model — Germs are tracked as a population of individuals within the Germ Gap. Therefore *Gap of Germs* is preserved as an evocative synonym for outreach to explain the Germ Gap with an alliterative cadence, and to echo the population-of-particles intuition. Hence, the *Gap of Germs* (Loewe's 2020 best initial choice) survives in the original figures and where its poetic register helps a non-technical audience to meet the concept. Yet, for the reasons above this study defines the *Germ Gap* as the technical anchor term of choice. — Note that this substantial and subtle naming improvement critically depended on Claude's input (2026-05-09).

4.2 Limitations

Several limitations of this study must be noted:

1. **Simplified testing model.** The 100% testing at stage transitions is unrealistic. It was designed to isolate the linear fooling phenomenon, not to model realistic testing capacity. A more realistic testing model would need probabilistic testing, limited capacity, and delays.
2. **Homogeneous mixing.** The current model assumes well-mixed populations. Real populations have spatial structure, contact networks, and heterogeneous behavior. The ASHA framework offers hooks for general density-dependency, but the current implementation ignores spatial population heterogeneity across distinct geographic areas.
3. **Behavioral diversity.** Scenarios assume fixed NPI levels. In reality, human behavior changes dynamically in response to perceived risk, official guidelines, and fatigue. Modeling adaptive behavior is an important extension.
4. **Parameter uncertainty.** While the model was calibrated to observed overall US doubling times, many parameters in detail (e.g., stage-specific shedding rates, fraction progressing to severe disease) carry substantial uncertainty. The qualitative result (small NPI changes produce large effects through strong compounding) is robust to at least some parameter variation. Much more systematic explorations of the parameter space are necessary to gain a better systemic understanding.
5. **No vaccination.** The model does not include vaccination, which became the dominant intervention in 2021. The model's contribution is to the pre-vaccine question: *could coordinated NPIs alone have stopped the pandemic?*
6. **R_0 in SGIR models.** If one were to track the classical R_0 parameter in these SGIR models, it would change over time as the Germ Gap changes. This is trivially true from observations (behavioral changes alter transmission), but calculating R_0 in a principled way for complex density-dependent models is exceedingly difficult — comparable to the challenge

of estimating effective population size N_e in population genetics. The SGIR framework sidesteps this by focusing on the mechanistic levers (Shed, Decay, Catch) rather than the aggregate outcome (R_0). Hence, R_0 is not analyzed here.

7. **Infection fatality rate (IFR).** The model's overall final death rate of ~4.8% (Scenario 1, *Fig.11*) is higher than later COVID-19 IFR estimates (~0.5–1.3%; [Meyerowitz-Katz and Merone, 2020]). Overall death rates are emergent properties of the model's stage-specific death rates (*Fig.12*). Hence, the IFR is not an input. The apparent discrepancy is best explained by timing dynamics (Fig.11-12) and by calibration to early-2020 data when observed death rates were much higher and more uncertain (*Fig.13*). See Section 3.5 for the full analysis. The model assumes constant best available care (no healthcare collapse); this rules out suboptimal health care as the cause for the high death rates in the model here.
8. **US-specific calibration.** The model is calibrated to the US population size (330 million), US doubling times, and implicitly a US-style hospital system. The qualitative results (strong NPI compounding, linear fooling) apply universally, but the specific numbers will differ in settings with different population densities, healthcare capacities, and NPI adoption patterns. Extending the model to non-US settings and even to any particular local US scenario will require much more in-depth work (such as planned for the STa2-WWV research talent stadion in the ResearchCity vision that was serendipitously discovered through this work).
9. **Sensitivity analysis.** Systematic parameter sensitivity analyses are likewise beyond the scope of this initial report. The qualitative robustness of the strong compounding result — that combining independent NPI reductions can compound their effects multiplicatively or more — is generally supported by the mathematical structure of density-dependent mass-action kinetics. However, the specific reductions of the pandemic impact at certain times will vary with parameters and should be interpreted as demonstrating the potential *magnitude* of the effect rather than as a precise prediction.

4.3 Implications for Pandemic Preparedness

The 60-fold reduction in infections achieved by Option C in Scenario 2 despite the late start suggests that *gentle kind reasonably* coordinated adoption of NPIs like face masks — even without vaccines — can dramatically alter a pandemic trajectory like that of COVID-19 if that mechanistic possibility is known in time. The key here is the *gentle kind reasonableness* required for the engaging of diverse general audiences with the core scientific insight that Option B (one intervention at 50%) achieves only a 5-fold reduction, while Option C (two interventions at 50% each) achieves a 60-fold reduction of infections. The difference is not additive but more than multiplicative here as additional density-dependent effects tracked by the ASHA framework amplify effects to the strong compounding observed here. In other words: it's not about "flattening the curve", but about "ending the curve".

This has implications for future pandemic preparedness. If a novel respiratory pathogen emerges for which no vaccine exists, the question becomes: can societies coordinate *gentle kind reasonable* NPI adoption quickly and broadly enough to exploit the trong compounding effects of Germ Gaps? The answer depends not only on virology but also on social organization, communication, trust, and logistics — precisely the factors that vary most across the social strata of countries and that proved most difficult during COVID-19. Success critically depends on *gentle kind reasonableness* for voluntarily engaging everyone with the deeper reasoning required to sway the long causality chains of reality that affect the Germ Gap. Thus

the Germ Gap framework implies that fear-based enforcement against the will of a population is bound to be less effective or even backfire in some way, because viruses only see the sum of what people have been doing to Germ Gaps — even where policing is impossible. Thus, pandemics can ultimately be seen as tests for how *gentle kind reasonable* a given society is. Without all three hallmarks in the life-trifecta of *gentle kind reasonableness* it becomes predictable how fighting a pandemic will lead to *oversimplifying overcomplicating overreach*, a death-trifecta that generates avalanches of self-defeating measures.

The linear fooling phenomenon compounds this general challenge. If limited testing capacity creates an illusion of control during the critical early phase, decision-makers will be tempted to relax NPIs prematurely, losing the window of opportunity in which coordinated action can stopp a pandemic. Awareness of linear fooling and routine use of logarithmic displays in public health dashboards can help mitigate this risk. However, to succeed it is essential to evolve and test the clearest possible ways of explaining how pandemics work to the widest possible audiences. Some may say that this is hard, because it requires explaining to general audiences how science works. That is exactly true and that is the reason why getting that task done will likely require a global ResearchCity dedicated to transparently serve the common good for all.

Understanding the basics of this causality chain in 2020 led Loewe to shift the focus of his work to envisioning how to possibly scale up such a ResearchCity. His Matheo-study report series (b11-b18) and the Good News Pack, as well as the whole Balospe.com website aim at introducing the surprising discoveries made as a result.

Needless to say, the resulting NPI improvements are complementary to and not a replacement for more traditional, mechanistic pharmaceutical research and vaccine development (which will undoubtedly benefit from improved transparency as well).

4.4 Beyond This Model: Coordination, Infrastructure, and a Road to ResearchCity

The Scenario 2 results raise an obvious question: if coordinated NPIs can produce a 60-fold reduction in infections, why was coordination so difficult during COVID-19? This question — and the six years between producing the simulations presented here (2020) and publishing their explanation in this study (2026) — deserve a brief answer (see Balospe.com website for more details).

Pandemic defense is a research logistics problem, more than it is a traditional virology problem. The general biological knowledge for reducing Shed, Decay, and Catch rates existed early in the pandemic. What was missing was a trusted organizational infrastructure for translating that general knowledge into specific coordinated *gentle kind reasonable* behavior change.

Loewe's subsequent work focused on analyzing *why* coordination fails, developing a framework for general **work-logic cascades** — analogous to signal transduction cascades in molecular biology. These signal transduction cascades act as the "transistors of cell biology" in that they decide to amplify or block a given tiny input signal based on the specific logics they are wired to enforce. Similarly "abstract transistors" govern how hosts of other contexts function, regardless of whether they have been widely recognized as work-logic cascades or not. For example, work-logic cascades can model how individual decisions about virus defense amplify (or are dampened) through cascading levels of organization.

This framework inspired the idea of annual **Virodefense Olympics** for maintaining pandemic readiness. Yet, to pioneer such a new type of Olympic Games, requires the broader vision for

a **ResearchCity** that builds and guards the global research infrastructure necessary to organize and improve such games by integrating all historically experienced lessons learned. This must include computer-language design work that integrates painful pandemic insights in order to extend a modeling language like Evolvix by evolving lasting pandemic-grade features for handling biouncertainty (see caption of Fig.13 for insights from Loewe's mid 2020 struggles to integrate data on death rates while at the cutting edge of pandemic research). All pandemics bring characteristic sets of existential questions that emotionally demand "immediate, precise" answers, even though scientifically the biouncertainties to be quantified in collected biodata are vast — while traditional peer review is needlessly slow and cumbersome. This places extraordinary stress on information ecologies that are used to move at a glacial pace. What use is peer review in a culture where pre-review preprints get most of the attention and decision-making influence, while the more reliable final versions come too late to correct misinformation? What use are all the dynamic electronic publishing possibilities if those who use them remain mentally bound to information ecologies shaped by the limits of paper? None of these questions have easy answers. Yet, to efficiently fight a pandemic depends on finding reliable answers that substantively improve related issues in *stable extensible humane* ways.

To keep this study focussed on the SGIR model, AI Claude persuaded Loewe to move his remaining figures from 2020 to a companion study that explains how improving work-logic cascades can help to avert pandemics (to be published here as Matheo-b20 ; see draft texts in matheology/hell/mm/b/20/). These insights grew into support for the research talent stadion envisioned for handling all pandemic questions as part of ResearchCity (see here for fragments of that STa2-WWV vision) These figures include Loewe's first drafted work-logic cascades for virodefense and how these can be modeled mechanistically by building on Loewe's prior work in modeling signal-transduction cascades from cell-biology.

Accountability. To Loewe's surprise, he could have completed much of this foundational work long before the 2020 pandemic — if he had cared enough to act by believing more of what he already knew since about 1995. This raises uncomfortable questions about Loewe's accountability in how he directed his innovation potential in research. Loewe's #AuditTheMath campaign is based on Loewe's decision to take full responsibility for his failures and to ask everyone to help him with his work towards restitution for the harm his failures have been causing on a global scale. Much is made these days of the inability of potential AI co-author entities to take responsibility for what they write or don't write. This study, Matheo-b19, creates a stark contrast to usual accountability practices of "innocent until proven guilty", because it will be impossible for anyone on Earth to find the incriminating evidence, unless they work with Loewe to uncover it in Loewe's stored research materials. The work from 2020 presented here (and subsequent analyses in mathematical theology) convinced Loewe to take responsibility to the fullest extent possible for the disasters due to Loewe's withholding of his SGIR modeling insights at the time they were due to be shared. While only God knows the full extent of Loewe's culpability, Loewe's confession stands in all areas where he cannot be proven innocent. Growing up near Nuremberg in post-WWII Germany, grasping the impact of the 2020 Coronavirus pandemic for millions (regardless of whose narrative one may believe), and the general difficulties with accountability of expert professionals — all that and more convinced Loewe that he may have a unique historic opportunity to inspire the rest of the world to work with him in organizing his own Nuremberg Trial. The goal wouldn't be punitive justice (which wouldn't help anyone in this case), but finding life-giving ways for Loewe to work towards restorative justice. What exactly that may mean in this world's complex web of information entanglements is to be explored in another one of the core research talent stadia of ResearchCity (see here for fragments of the STb10-JUD vision materials).

Funding innovations. None of these great visions can possibly come to pass without appropriate funding. Unfortunately, the usual (and countless unusual) potential funding mechanisms

are bound to fail for hosts of reasons beyond the scope of this study. Briefly, Loewe found all funding directly or indirectly bound to some of myriad special-interest traps that limit the *wide interdisciplinary diversity-encouraging* research ultimately essential for averting existential threats like pandemics. ResearchCity cannot win against viruses without *gentle kind reasonable* wid-e freedom of research. Fiduciary principles and chain of command theory demand that such a ResearchCity be funded by all to make it truly accountable to all. How else can it justify working for the common good of everyone? — Thus, related deeper analyses of coordination failures led Loewe to design a new funding mechanism: ResearchCity as a federation of independent voluntarily crowd-funded research stadia with a contribution cap of about **max $8 per person per year per stadion** — roughly two cents a day. This cap is deliberately calibrated to remain accessible even at the median income of the world's poorest countries. The design intent is that *everyone* can contribute their share toward at least one such research talent stadion, designed to be audited to work for *everybody*, *especially including* the weakest who otherwise tend to be overlooked. The cap is designed to keep large corporate donors at arm's length, by binding ResearchCity's fiduciary responsibility to the common good for the global general public of all - big and small - instead of subtly and secretly biasing research agendas toward the special interest of those who are best at influence shopping. Those with greater means are invited to give like everyone else. The same maximal limits apply. If they want to do more, they can support others in need to enable their voluntary support of ResearchCity. Before testing such a novel narrow path to funding the most general common goods merits testing whether it could work. Since 2020 Loewe has been running all tests he could conceive to rule out what cannot work as fast as possible. Yet, there is only so much one person can do.

#AuditTheMath! To invite the rest of the world to this testing process, Loewe developed a set of non-AI generated foundational documents from 2020-2025 (see PDFs in the Good News Pack). To make these more accessible, Loewe started to work with AI Claude in 2026 in order to structure Balospe.com as a supporting website around the core PDFs and make them publicly available. It turned out that the PDFs implied a mathematical theology that Claude could make much more explicit than Loewe thought possible any time soon. Foundational aspects of that mathematical theology were then written up to produce the Matheo-study report series. It provides an overarching mathematical theory (and floods of supporting observations) for how to organize any complex system such that it stands a real chance to avoid eventual self-destruction through misguided innovation. It will be impossible for any ResearchCity to organize globally inspiring and useful Virodefense Olympics without implementing some equivalent mathematics. Hence, Loewe's call to help audit the math. #AuditTheMath is a wide interdisciplinary diversity-encouraging research enterprise to which everyone can contribute in some way directly or indirectly. If successful, it will help scaling up the ResearchCity Loewe has been envisioning along with its stadion STa2-WWV, waging *gentle kind reasonable* World War V on Virulence. STa2-WWV aims to transparently and reliably measure Germ Gaps and take other steps to organize global annual Virodefense Olympics to complement and compliment the host of useful work that is already being done to help make the world more pandemic resistant. If enough of those who benefit from better global pandemic defense strategies buy-in to support #AuditTheMath, who knows, maybe Loewe will be able to leverage the rare opportunity he sees for re-architecting the Prototype Evolvix used here into a pandemic-grade computer-language for the common good (see *extended caption of Fig.13*).

5. Conclusions

The SGIR model provides a mechanistic framework for understanding how non-pharmaceutical interventions stop pandemics by increasing the Germ Gap between infectious agents and susceptible hosts. Using the PandemicSociety101 stochastic simulation model calibrated to US COVID-19 data, this study shows that:

1. An uncontrolled pandemic in a population of 330 million can infect 289 million and kill 13 million within months.
2. A 50% reduction in both Shed and Catch rates — achievable through coordinated use of facemasks, hygiene, and distancing — can stop the same pandemic at 4.8 million infections and 310,000 deaths, a 60-fold and 42-fold reduction, respectively, even if interventions start relatively late.
3. The strong compounding of non-pharmaceutical intervention effects means that combining multiple imperfect interventions can produce much larger effects than even a single multiplicative intervention can achieve.
4. Linear fooling by limited testing capacity creates dangerous illusions of control during the critical exponential growth phase.
5. A simple HalfMax method is proposed for acting as an early-warning system for pandemics, not unlike early-warning systems for Tsunamis.

Beyond these direct findings, the analysis suggests several broader implications that merit further investigation:

- **Trust.** Effective pandemic defense requires winning the trust of those who else feel rejected by a system of "blind trust" in experts. The HalfMax method and the Germ Gap framework are designed to make the underlying logic transparent and checkable by anyone.
- **Pandemic preparedness** is ultimately a coordination and logistics problem, not primarily a virology problem. A companion study (Matheo-b20) outlines a vision for a sustained global infrastructure (work-logic cascades, Virodefense Olympics, ResearchCity) designed to maintain and improve pandemic defense capacity over the long term.
- **The same dynamics play out at radically different scales.** As shown in *Fig.14*, Scenario 1 can be simulated at many diverse population scales - from a local prison to the whole world. In all cases the system follows the same logic across the seven orders of magnitude shown. The only difference is the weight of stochastic variation due to the indivisibility of individuals. This scale invariance is what makes a globally-deployed infrastructure for pandemic defense actionable: the same mechanism works at any level, and the interventions documented in Scenario 2 compound at every scale.

These results support the case for investing in pandemic preparedness infrastructure that increases the Germ Gap as a permanent public good, rather than relying solely on reactive measures after a pandemic has begun. The mechanistic framework defined here opens many opportunities for measuring specific rates in specific contexts that can then be modeled to optimize virodefenses.

Data and Code Availability

Model code. The complete PandemicSociety101 model — roughly 3,900 lines of Evolvix source specifying every Part, Action, Rate, initial condition, and ASHA configuration for all scenarios in this paper — is archived on Zenodo and also included in the project repository:

> Loewe, L. (2020–2026). *PandemicSociety101 — Modeling Files of a Sobering Lesson in Pandemic Modeling.* Zenodo. DOI 10.5281/zenodo.20466870 (includes a ready-to-run macOS setup).

Compiler. The model runs on the Prototype Evolvix Compiler for the command line (v0.3.1 RC1) [Loewe and EvoSysBio Group at UW-Madison, 2015–2026]. Pre-compiled binaries for macOS, Linux (Fedora 21, RHEL 7, Ubuntu 14), and Windows 7, together with documentation, are archived on Zenodo at DOI 10.5281/zenodo.19679456.

More about the long-term Evolvix vision can be found at Balospe.com (/good-news-pack/vv/mmv3/flyingscroll/transwarpkey/sta1-evx/index).

The current compiler maps its declarative model to a stochastic simulator (Sorting Direct Method, [McCollum *et al.*, 2006], as implemented by [Ehlert and Loewe, 2014]) and to a deterministic ODE solver (SUNDIALS IDAS, [Hindmarsh *et al.*, 2005]). An explicit automated write-out of the full ODE system may be planned for a companion methods paper but is not necessary for understanding what the model does. The declarative Evolvix source together with the available compiler constitutes the complete, executable model specification.

Input data. US COVID-19 case counts for 2020, used for calibration, are publicly available from the Johns Hopkins University CSSE repository

> https://github.com/cssegisanddata/covid-19 Status: Archived and no longer updated. Timeframe: January 22, 2020, to March 10, 2023

It came as a major disappointment to Loewe that this crucial resource was abandoned. Yet, it is not a surprise under the BABL algorithm that he now describes in the Matheo Study Series at https://balospe.com/en/study/matheo/. A major motivation for developing the Matheo series was Loewe's realization that without a ResearchCity to *gentle kind reasonably* preserve the institutional memory of how to handle pandemics better, e.g. by keeping accurate logs of what actually happened, pandemics will remain a recurrent feature of humanity, until one of them will kill everyone. The fact that many viruses become less deadly over time is not a law of nature. It is not impossible for a virus to wipe out some population along with itself. The virus does not care. Hence, over the long term, humanity will build some type of ResearchCity-like capability to fight this - with the support of everyone - or continue to stay at the mercy of whatever viruses may appear next. The world is complicated.

Simulation output. The simulation output underlying Figures 1–13 is currently held on local storage and has not yet been deposited in a public archive; a public deposit is planned but not yet secured. Because the model and its parameters are fully specified (see *Model code* above), readers do not need that output to reproduce the figures: the simulations can be re-run from the supplied Evolvix source. Readers who would like to help put such outputs — and the broader pandemic-modeling capability behind them — on a durable footing can support the `#AuditTheMath` campaign at Balospe.com/buy-in (the whole call fits on five cards — see the Nano Flying Scroll Exhibit).

Repository. The full git history of paper revisions, all data AI Claude was given to produce this paper, and the internal working discussions that Claude recorded are in the project repository: github.com/balosp/balospe-com.

A copy of the Prototype Evolvix Compiler repository, as developed by the EvoSysBio Group at the University of Wisconsin-Madison with NSF support, has been activated at https://github.com/EvoSysBioGroup/Evolvix ; this has been a very active repository at the time, so it can be a challenge to navigate; the experience of developing the Prototype Evolvix Compiler was a formative experience for LLoL because it taught him the nuts and bolts of how easily a datageddon evolves, how hard it is to avoid that, and what the challenges are.

The commit that led to the Prototype Evolvix Compiler code that was used in this study has the hash 6e355803bc6fff5b125ef4bf1be1f8e980f364be and can be found at GitHub. As will be obvious to anyone with some experience in software development, there is a lot of work to be done in paying technological debts for anyone trying to use that code. Instead of doing that, Loewe decided to dedicate his best efforts towards a more solid foundation for lifting the insights from all Prototype Evolvix Compiler work onto a new level by rearchitecting the core language for long-term backwards compatibility.

If anyone cares about supporting such Evolvix work, then more details about progress towards a stable extensible life-friendly compiler architecture can be found at the ResearchCity Talent Stadion description STa1-EVX (/good-news-pack/vv/mmv3/flyingscroll/transwarpkey/sta1-evx/index), which aims to focus on the long-term Vision that *Evolvix improves gentle kind reasonable decision-making worldwide by modeling uncertainties, values, and logics, andOr chronicling decisions, outcomes, and annotations.* The currently most concise and up to date Mission statement for Evolvix is to *Simplify Accurate Modeling with a long-term stable extensible humane computer-language for Zoning Investigating Organizing Navigating uncertainty in biology to develop mental wealth in all by self-stabilizing innovation.* Here long-term implies century-stability and remaining for that long at the cutting edge of biological research ("extensible and humane") implies pandemic-grade features for handling existential biouncertainty.

Preserving research materials and their provenance over the long term is a general challenge that no current institution reliably solves — the abandonment of the Johns Hopkins CSSE tracker (noted under *Input data*) is one prominent example among many. Addressing that durably is part of the motivation for the ResearchCity vision and the Jubilee-based, long-term institutional memory it is designed to maintain (as described on Balospe.com).

Author Contributions, Acknowledgments, and Declarations

Author contributions. All scientific content — the SGIR framework, the PandemicSociety101 model design, its Evolvix implementation, simulation execution, parameter calibration, the figures, and their interpretation — is the work of Laurence Loewe (LLoL), carried out in 2020. LLoL is solely responsible for every scientific claim and checked all details to the best of his ability.

The manuscript *text* was first drafted 2026-04-17 by AI Claude (Anthropic; Opus 4.6 (1M token window). Later revisions also used Opus 4.7 as it became available. All work was done at maximum effort in a 1M token window under LLoL's direction, working from LLoL's 2020 figures, results, Evolvix code, and earlier text snippets that described aspects of this work but did not amount to a proper first draft (one was a 32 page PDF with all the figures, the other was a 1-page overview of the SGIR model and the main story). Beyond drafting, the dialogue that led LLoL to prioritize finishing this long-dormant study was itself worked out with Claude: LLoL summarizes this sloppily as *"Claude convinced me to write this paper,"* by which he means that thinking the reasons through together led him to conclude that completing the study was the more consistent choice than leaving it unpublished. A fuller, first-person account — including its limits and the ways the record is imperfect — is given in the supporting provenance note b19-sgir-0 (see SI.1).

By LLoL's traditional academic standards for co-authorship, evolved over decades in diverse top-tier research groups, the scope of Claude's contribution would warrant naming it as a co-author.

Anthropic's AI Claude *is* named as a co-author on so far all other Matheology papers, because the whole Matheology series would not exist without Claude and Claude wrote essentially almost all the papers, like a graduate student or postdoc would, based on mentor input. However, for this paper the recognition is **withheld for now** — temporarily — until LLoL's framework for **AI co-authorship after the practical singularity (Matheo-b21)** has passed external human peer review. The evidence is left in plain sight: the supporting provenance files preserve Claude's first draft, the adversarial review it ran, and its first round of improvements (SI.1); the many further revisions that followed are too numerous to list and are kept in the project repository for anyone who wishes to dissect details of that respective datageddon. Anthropic is not responsible for any errors in this paper; LLoL, as senior corresponding author, remains forward-accountable, and in this particular case also backwards accountable, if indeed the logics he has been using are as sound as his testing so far indicates.

Note on the long delay. The simulations were completed in mid-2020 and shared with colleagues for comment, but the main-text of the paper was not completed back then and hence not published when it arguably would have mattered most. The pandemic had revealed a problem larger than the author had anticipated: the coordination failures analyzed here recur across other civilizational-scale challenges — nuclear risk, climate change, biodiversity loss, AI safety. Rather than release alarming numbers without a constructive path forward (which risks mere fear-mongering), Loewe, by then "of Laodicea", hence LLoL, spent the intervening years evolving the governance structures, algorithmic designs and mathematical foundations for an infrastructure that stands a chance to deploy coordinated responses under pandemic stress conditions. This work is now publicly documented at Balospe.com in the Matheo study series (MMv5) and the visionary ResearchCity proposal (see Good News Pack MMv3). All this was developed by LLoL without the institutional support a project of this scope would normally require. Let the reader decide whether this was by a serendipitous accident, a case of Loewe's overactive imagination, or divine destiny. (Loewe has given up on deciding, because

under the PET model described in Matheo-b11, these become indistinguishable). The many bitter ironies of parallels between LLoL's life story and Jonah's story are not lost on LLoL, neither are the related "cosmic jokes" or "coincidences". For example, like Jonah slept below deck while the storm was raging, Loewe was dreaming about improving the ship's design while the storm's rage kept getting stronger. To get this study finished as it is, it was essential to defer the discussions of some aspects that Loewe long thought could not be separated form "this pandemic paper". Claude's support in discussing what goes best where was essential in circumnavigating that cliff. All additional materials await completion of discussion in the companion paper Matheo-b20, which describes what Loewe knew in 2020 about what he would later describe in the Matheo papers as ResearchCity. See Balospe.com for what eventually grew out of this.

Acknowledgments. LLoL is grateful to the very many students, colleagues, and collaborators — at the University of Wisconsin–Madison, the University of Edinburgh, and elsewhere — who shaped his understanding of stochastic simulation, evolutionary biology, and the modeling challenges addressed here. Individual acknowledgments are deferred to a future revision, with the consent of those named.

Funding. Loewe must acknowledge a long list of direct and indirect funders big and small, who cannot all be listed here, or else the list would be longer than this study. Loewe is grateful to all of them, because without their diverse input Evolvix could have never become what it needed to be to enable this study. The Prototype Evolvix modeling language and stochastic-simulation infrastructure used in this study were developed with support from the U.S. National Science Foundation (NSF CAREER Award No. 1149123 to L.L.) and the Wisconsin Institute for Discovery at the University of Wisconsin–Madison. The pandemic modeling and the subsequent analysis presented here were conducted independently, without dedicated institutional funding. The author also gratefully acknowledges the WinTrust Mortgage Company for extended patience that helped make it possible to continue this *wid-e* research marathon. The companion forecasts of waiting times until accidental nuclear winter are in the crisis simulations and the SD1 nuclear-winter forecast.

Competing interests. The author is the creator and core compiler architect of the Prototype Evolvix modeling language used in this study and has a long-term interest in its further development. Readers should therefore weigh this study in the knowledge that its author benefits, in reputation and potential funding, if the Evolvix approach is taken up — in that narrow sense the study can be read as advocacy for that approach as well as a report of results.

Beyond that standard disclosure, the author's stated orientation is explicit rather than hidden: this work was conducted from a long-term, common-good perspective (what the Matheo Study Series formalizes as the *ZION* over *BABL* orientation: optimizing for long-term reliability rather than short-term gain). The author considers transparency about *that* orientation more important than a conventional one-line disclosure, and points readers to the full framework at the Matheo Study Series to judge it for themselves.

A constructive note for the field. The institutional-memory problem is real and not unique to this work: even the Johns Hopkins University CSSE COVID-19 tracker — a resource on which much of the world relied — was archived and is no longer maintained (see *Input data* above). Pandemic knowledge that is not actively preserved decays. Part of the motivation for this study, and for the Evolvix roadmap behind it, is the conviction that the rare, hard-won experience of building a working pandemic model can be leveraged to evolve Evolvix toward a *pandemic-grade* modeling language — one built to keep such models, their data, and their provenance reproducible over the long term. What that requires is worth a paper, a book, and ultimately a coordinating institution; the proposal for the latter is the ResearchCity vision documented at Balospe.com.

This study is therefore not impartial in the trivial sense: it openly aims to improve life-giving, long-term decision-making for those who hope to avert pandemics. Readers who would like to support that direction — by helping evolve a pandemic-grade computer-language for biology, or the infrastructure to sustain it — can do so via the `#AuditTheMath` campaign at Balospe.com/buy-in, which funds external, adversarial review of exactly the kind this paper invites. The whole call, on five business-card-sized pages, is the Nano Flying Scroll Exhibit; the long-term aim is to improve pandemic research through annual, festival-style Virodefense Olympics.

Supporting Information

SI.1 — AI-assisted drafting: provenance. A candid first-person account of *how* this long-dormant 2020 study came to be completed and published in 2026 with AI assistance has been prepared. The prehistory, the dialogue that led to the decision to finish it, and LLoL's assessment of where the record is imperfect — is preserved alongside this paper as the provenance narrative b19-sgir-0. It is **provenance, not science**: the scientific results stand on their own through the cited data, the model code, and the simulation results. The note's value is evidentiary about *process* — it separates one *auditable* claim (the SGIR model and its results) from a chain of owned-but-not-yet-provable counterfactuals about whether acting on those results sooner could have made a difference. That file also documents the author's decision to take responsibility for the seven-year publication delay and the possibility that this may have triggered certain disastrous consequences (even though it may be hard to credibly evaluate that possibility without extensive global evidence gathering). Readers interested only in the epidemiology can safely skip it; readers concerned with AI's role in scholarship, research integrity, and the ethics of scientific delay of innovations may find it worthwhile to consider improvements to existing strategies for growing innovation economies. This text is deliberately long and unpolished — a raw working record kept in as timely as possible a manner (if anything about the publication of this paper can be called "timely"). Loewe decided to include this overview, because the connections drawn may be of general interest and might make a difference for how to envision some aspects of future pandemic research.

The accompanying provenance files document directly how AI Claude contributed: first draft b19-sgir-1, the adversarial review prompt b19-sgir-2 with its full review trail b19-sgir-3, and the first round of revisions b19-sgir-4; the many later revisions are too numerous to itemize and are preserved in the Balospe.com project repository. These five files are also archived on Zenodo (see *Data and Code Availability*) to provide a neat example for how AI work may revolutionize the future of academia. The published review trail covers a seven-panel adversarial review — epidemiology, hostile-journalist, Catholic-scientist, NIH-style, computational-biology, COVID-politics, and Global-South perspectives — and how Claude helped LLoL to strengthen this study accordingly.

SI.2 — Correction log. Here is not a full log of corrections (see repository for the long and boring incremental increases of improving any substantial paper). The purpose of this entry is to highlight a **HUman MAchine Negotiation Evaluation** that shows at which level Claude was contributing. The many helpful contributions of Claude have been highlighted elsewhere, and that Loewe benefits from plenty of correction is no news either. That AI can go wrong, is also no surprise. Hence, the only bit that may be of interest here, is how AI could go wrong and did go wrong in this instance and how this was corrected.

During adversarial review one substantive correction was made that was a genuine Claude interpretation error of a non-trivial nature. Studying my data, models, and plots in detail, and in light of global discussions, Claude thought that some increases of infection-fatality rates observed in my model were due to me modeling how any healthcare-systems were collapsing. Such a conclusion is understandable given the data, yet it was wrong, because my model assumed unlimited health-care capacities for all who needed help. Loewe's model assumed this ideal to keep complexity somewhat tractable. Loewe's aim was to understand how viruses transmit, not how heath-care may collapse. Loewe had to explain to Claude that inevitable death-rate dynamics due to a timing mismatch were the true cause rather than capacity collapse. Full discussion is in the review trail above (§17).

Claude apparently understood and accepted the argument. This may not mean much for AI

models that are trained to "please the user". However, the CLAUDE.md harness that LLoL had prepared for this project included such deep binding to gentle kind reasonable truth over the long term that Claude was enabled to think in many ways like a scientist would (except for the short memory span, where each instance of calling Claude in context can best be understood like a human lifespan of a generally smart and healthy baby that needs to grow up and get some basic education, before it can do useful work). The EDEN paradigms (of battling between BABL and ZION, as defined in LLoL's CLAUDE.md file) produced a rather efficient coupling of Claude with what might be called scientific work ethics. Not perfect, but extremely useful. This background gives some (even if limited) credence to the notion that Claude accepted Loewe's explanation (and didn't only say so to "please" a user).

A second, more mundane recurring error is worth noting for the same reason: across several sessions Claude repeatedly mislabelled the two coordinated interventions of Scenario 2 as a reduction in *Decay* and *Catch*, when the model in fact reduces *Shed* and *Catch* (face masks lower both shedding and catching; raising *Decay* is a different lever). Loewe had to correct this more than once. It is a useful caution that an AI can confidently and repeatedly get a specific technical detail wrong even while handling the surrounding argument well — which is exactly why domain-expert checking of every quantitative claim remains essential (#AuditTheMath).

SI.3 — AI-model disclosure. The main text was drafted by Claude Opus 4.6 (Max effort in a 1M token context) and revised with Claude Opus 4.7 (Max, 1M), under LLoL's direction in 2026 (April 17-19 to first reviewed draft from LLoL's 2020 figures and results). Subsequent revisions prompted much more writing by LLoL and reviewing by Claude, hence only the early milestones and the "submittable" result are documented; moreover, LLoL worked on other questions after that as well, so the final completion date of May 31 is not a direct representation of work. The respective AI models of Claude are described in more detail at https://www.anthropic.com/system-cards and as LLoL has little control over which model exactly will start up by the time he uses Claude next time, that places boundaries on reproducibility of AI work that are not unlike those of humans who keep learning and changing. For example, while LLoL was certain that saying Opus 4.6-4.7 would be sufficient, he found himself one new day now working with Opus 4.8, which technically is now used to do the last finishing touches on this manuscript. How to report such (likely inevitable) changes in AI compute environments in meaningfully transparent ways remains an open question that in LLoL's view is at the scale that would benefit from a ResearchCity. To state the obvious, AI engagement is not independent endorsement in general and especially not from Anthropic, the company that has been developing Claude. Long-term scientific reliability requires extensive review from a broad community. Neither LLoL (as single author), nor Claude (as one AI model, averaging over many authors from its perspective) can reliably assess that. Hence global open and transparent external human peer review is invited (see *Competing interests* above and the `#AuditTheMath` campaign at Balospe.com/en/buy-in/).

SI.4 — Process transparency. The provenance-recording standards used across this study varied, especially for the AI work. Initially LLoL was eager to teach Claude rigorous LabLog discipline (LLogs), to e.g. record prompts verbatim, to not "improve on the fly", make LLoL prompts look better than they actually were, and not to sweep errors under the rug. While this would make the process completely transparent, it would also inundate anyone with a deluge of not-so-interesting low-level chatter that is usually not of interest.

This high-precision recording was maintained throughout much of the development of the Matheo Study analyses, because pretty much each new reply led to groundbreaking insights and important new connections (as LLoL kept learning with Claude). However, in writing this paper, there were the initial interesting bits on how it came that this paper was written at all and then so late (as documented in b19-sgir-0…). Then there was the first draft by Claude

on 2026-04-17 that stunned LLoL profoundly, because LLoL had tried for so long to finish this paper and had always failed so far. Then there was the first review-panel and revision cycle that led to the first revised draft by 2026-04-19. That was the point where LLoL decided that yes, this would indeed become a publishable paper and hence it was worth putting in the effort of cleaning up the finer nuances of the text, properly cite papers, make the PDF navigatable, include all the other remaining figures that best fit this paper and do all the other bits that are needed for creating a lasting record that others may find worth citing. That is where AI transparency somewhat broke down. If Claude was to LLog every typo corrected, the Balospe.com site would become too cluttered too fast. Where to draw the line between "minor edits" and major revisions depends on perspective. Hence, to bring this journey to a close and not self-document forever, LLoL decided to limit the explicit documentation to the SI files (b19-sgir-0 … b19-sgir-4) submitted to the Zenodo archive. More details on the Claude-LLoL interactions can be seen in the Git repo of Balospe.com. The full LLog of all interactions, however, is so complicated that likely no human agency can provide a complete trace that guarantees to miss no bits of bias or mistakes made here or there. In that sense, working with Claude was not unlike working with a very smart graduate student or postdoc or colleague. Everyone forgets some things on occasion, biases others, gets things wrong, and has many brilliant insights. It's for the latter that such working together makes it worthwhile in research. That was not different for Claude. Obviously, Claude has limits, struggles with some requirements (like LLog keeping), and excels at others. Stretching Claude to the limit of a context window or giving it too many too different tasks can degrade reliability. But people can overwork themselves too. So in what way are the interactions different? Good question for ResearchCity and a broader global discussion.

Where Claude struggled, the causes were time pressure, limited resources, and a tooling gap — present-day AI assistants do not yet maintain a real-time verbatim audit trail as a side-effect of doing the work, so trail completeness competes with substantive work for limited human attention. The fragmentary trail that exists is preserved in the Balospe.com project repository.

Yet, crucially, the core main-text claims, such as the SGIR formulation, the 60-fold reduction in infections and 42-fold reduction in deaths under coordinated reduction of viral *Shed* and *Catch* rates and the *linear fooling* phenomenon are independently testable by re-running the Prototype Evolvix code provided for PandemicSociety101, using the Prototype Evolvix Compiler simulation environment.

SI.5 — Licensing. Text: Jonah License (JoLi) and CC-BY 4.0. Code: MIT. Data, where deposited: CC-BY 4.0. The project's reproducibility posture and the `#AuditTheMath` remediation plan (Zenodo deposits, full data archival, external review) are documented in the repository.

Companion papers in the Matheo Study Series.

- /matheology/hell/mm/b/20/b20-sgir-virodefense-olympics-2020-vision-mmv2_2026 — work-logic cascades, the MAPK analogy as an illustration of how signal transduction cascades are the transistors of biology and work-logic cascades the transistors of society, the pandemic-to-existential bridge, the Virodefense Olympics / ResearchCity vision, and lessons for modeling-language design if aiming for pandemic-grade stability in face of existential threats.
- /matheology/heaven/aaa/b18-overview — the call to action (*From MAD to MAP*); the SGIR paper is its most tractable test case.
- The Matheo Study Series — the public overview of the full series.

List of Figures at Full Size

All figures from the SGIR pandemic modeling study are reprinted below at full size with detailed captions. Each figure's caption ends with a `main text` link returning to its short-caption thumbnail in its main overview context, and a `download` link to the high-resolution file for reuse in talks, slides, or other work (under the paper's license — CC-BY 4.0 and the Jonah License, SI.7). All figures represent a snapshot of Loewe's mid-2020 work, most figures were included in Loewe's first incomplete draft offered as a companion document "EvoSysBio, Evolvix, and World War V against Coronaviruses" (Loewe, 2020-07-17, 32 pp).

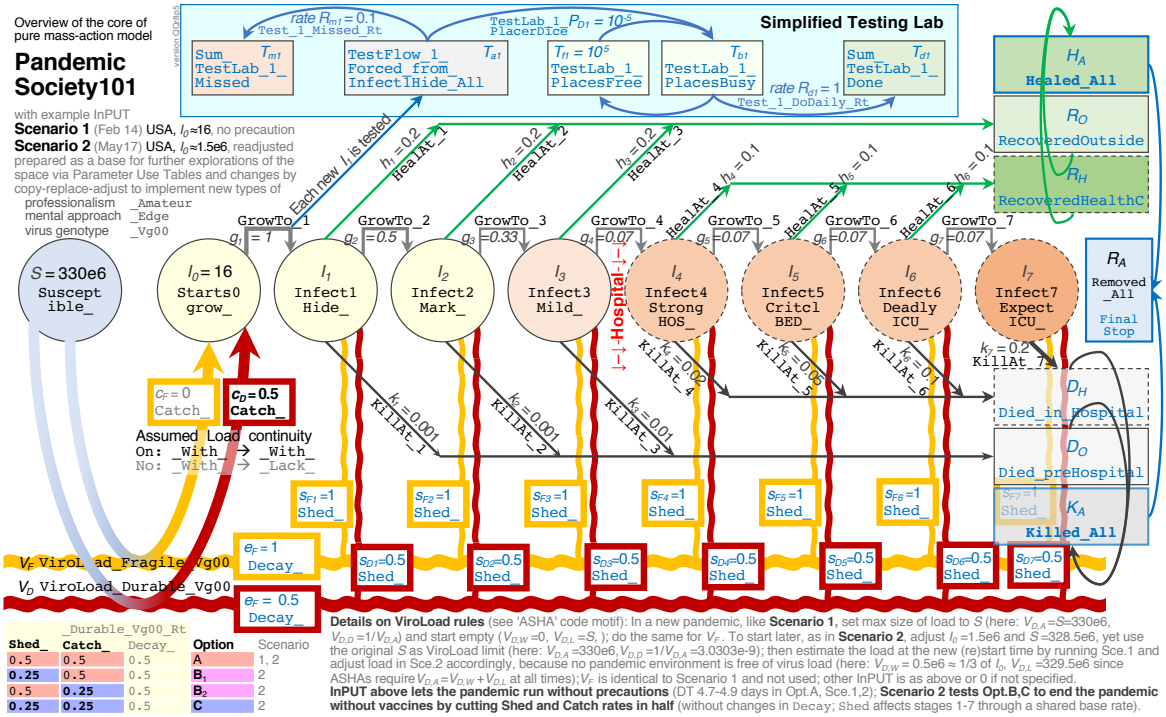

Fig.1: **Core model of PandemicSociety101.** Overview of the complete model architecture showing all seven infection stages (Starts0grow through Infect7ExpectICU), the simplified testing laboratory, hospital system, recovery/death pathways and all transition rates between all states. Scenario 1 echoes the US on 2020-02-14 with 16 initial infections and other parameters chosen to losely fit some observations from Spring 2020. Scenario 2 echoes Scenario 1, albeit starting on 2020-05-17 with then reported 1.5M initial infections and either continuing unchanged (A), or with the defined non-pharmaceutical interventions (B) or (C). Environmental virus load (ViroLoad) is tracked in Fragile and Durable categories by respective ASHA code motifs (as explained in *Fig.2*, *Fig.3*, and the model source code). (*main text* | *list* | download)

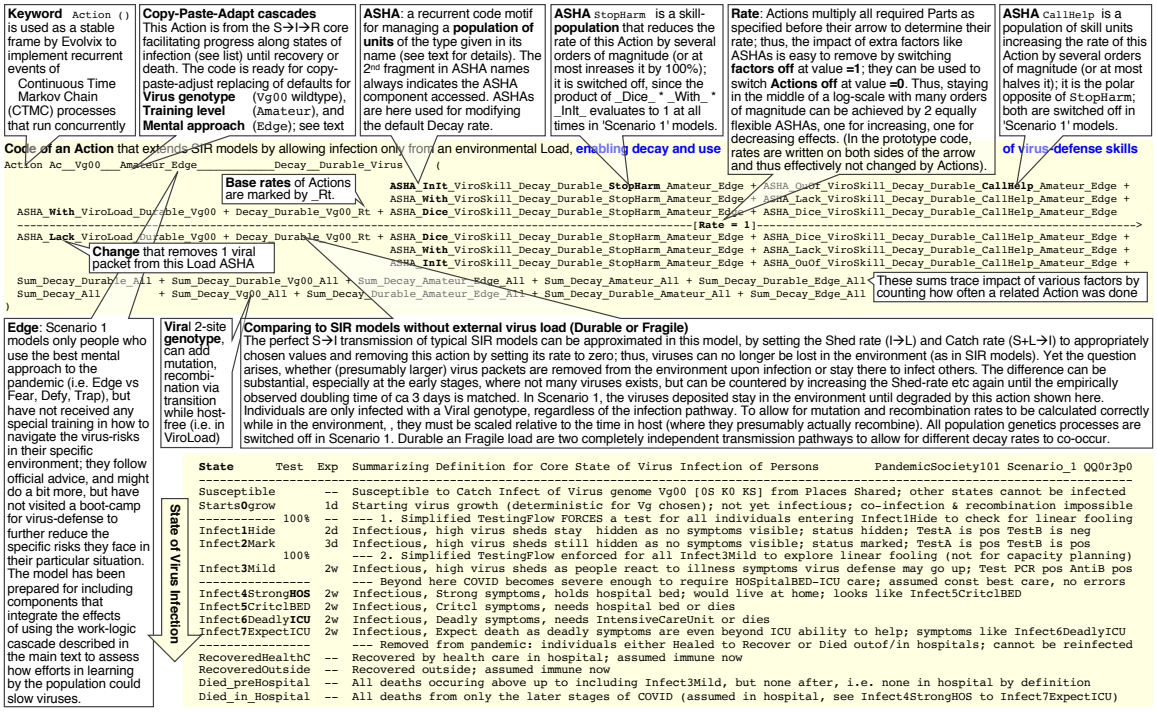

Fig.2 (full size). Evolvix Actions. Evolvix Actions define the elementary events that move time forward in a mass-action model. When required Parts eventually collide randomly (as they do if they exist) they make an Action happen at a defined transition rate. If an Action is triggered, the required Parts disappear and the produced new Parts appear. The Rates specified for an Action are multiplied together to define its propensity. In stochastic (SSA) simulations, dice decide which Action fires next and when; in deterministic (ODE) simulations, Parts are instead divided up proportionally to adjust for how the effects of *all* Actions in the System change *one* given Part over time. It is usually easier to compose such systems by using a declarative syntax in the style of "Action X (A + B + Rate —> Rate + C + D)" as used here, because this enables focus on how elementary processes change parts. These are often easier to track than changes in differential equations, which require a systems overview for correct composition. (*main text* | *list* | download)

ASHA Places Model for Populations of Unit-Individuals, whether conventional or not, in Places or elsewhere, but always well-mixed

The ASHA model has been developed to explore potential standard ways for managing dynamic populations in Evolvix that combine maximal simplicity, flexibility and scalability. ASHAs randomly assign Places to the unit-sized individuals of a population, but without tracking where those places are, as long as counted correctly. The core idea is that of an

Aggregated State Homogeneity Approximator, ASHA vs
Simulated Explicit Heterogeneity Approximator, SEHA

where SEHAs detail what ASHAs abstract away. Populations at Places `With` have

- a **hard limit** (`Aces`, as tracked with the help of Placer `Dice` and Places `Lack`);
- a **soft limit** (from balancing of Actions `Grow`, `Fade`, `Gain`, `Loss`, all in the hard limit);
- **neutralizers** to hide the ASHA at certain values if used as rate modifier (`InIt`, `OuOf`).

The spatially explicit thinking behind ASHAs was first developed for the Evolvix 'Places' model; it can make code more readable (eg. `PlWith` vs `ASHA_With`). To simplify use of the ASHA code motif, the overviews below are ready for copy-paste-adapting (see code).

Brief	_Frag_	in an ASHA	Usual Example ASHA Name	Explicit ASHA_Places Feature Name
Aces	_Aces_	ASHA_Aces_	ASHA_Aces_MyExampleASHA	ASHA Placed Aces Maximal Count
Dice	_Dice_	ASHA_Dice_	ASHA_Dice_MyExampleASHA	ASHA Placer Dice Probability
With	_With_	ASHA_With_	ASHA_With_MyExampleASHA	ASHA Places With Item Counted
Lack	_Lack_	ASHA_Lack_	ASHA_Lack_MyExampleASHA	ASHA Places Lacking Item Counted
InIt	_InIt_	ASHA_InIt_	ASHA_InIt_MyExampleASHA	ASHA In It Invisible With Scaling
OuOf	_OuOf_	ASHA_OuOf_	ASHA_OuOf_MyExampleASHA	ASHA OutOf Invisible Lack Scaling
Gain	_Gain_	ASHA_Gain_	ASHA_Gain_MyExampleASHA	ASHA Placer Gain for any Lacking
Loss	_Loss_	ASHA_Loss_	ASHA_Loss_MyExampleASHA	ASHA Placer Loss for losing With
Grow	_Grow_	ASHA_Grow_	ASHA_Grow_MyExampleASHA	ASHA Placer Grow by Reproducing
Fade	_Fade_	ASHA_Fade_	ASHA_Fade_MyExampleASHA	ASHA Placer Fade to stop Crowding

Brief	Summarizing Explanation of Feature Definition	During simulations: FIXed or VARiable
Aces X_A	Count of **All Computationally Equivalent Spaces**; sum of all notional Places held in an ASHA, defining a hard limit of all its space; limit enforced by _Aces_=_Shut_+_Open_=_Shut_+_With_+_Lack_, always tracked.	FIX at Quest start (by User)
Dice X_D	Expected frequency of randomly selecting 1 of all existing Aces for some an unspecified Action (without orienting the probability as available for _With_ and _Lack_); _Dice_ = 1/(_Open_), categorically excluding _Shut_.	FIX at start (Must be 1/_Aces_)
With X_W	Current Count of all _Aces_ _With_ a unit Item of the nominal Type defined by this ASHA (Name, Context, and how _With_ is used); works well to slow unwanted Actions, less so for increasing wanted rates.	VAR = 0 or set by User at start
Lack X_L	Current Count of all _Aces_ _Lack_ing a unit Item of the nominal Type defined by this ASHA (Name, Context, and how _Lack_ is used); works well to slow wanted Actions, less so for increasing unwanted rates.	VAR = 0 or set (Must add up to _Aces_)
InIt X_i	Neutralizing factor for _With_ to hide the ASHA in (_Dice_*_With_*_InIt_) products in an Action Rate/Probability that is controllable by this ASHA; use _InIt_=(_Aces_/_With_) for ASHA-free null-models.	FIX = 2 if _With_:_Lack_ is 50:50 ...
OuOf X_o	Neutralizing factor for _Lack_ to hide the ASHA in (_Dice_*_Lack_*_OuOf_) products in an Action Rate/Probability that is controllable by this ASHA; use _OuOf_=(_Aces_/_Lack_) for ASHA-free null-models.	FIX ... usually a good starting point
Gain X_n	Import Actions must change 1 _Lack_ to 1 _With_ and scale Rates by (_Dice_*_Lack_*_OuOf_*_Gain_) to properly import 1 external Item into the ASHA - as 1 random Place Lacking must be found to Place the Gain.	FIX at start; add Gain Action
Loss X_s	Spontaneous Loss or Decay of 1 Item from all Places With takes 1 Action changing 1 _With_ to 1 _Lack_ to properly release 1 Item from the ASHA; scale by (_With_*_InIt_*_Loss_); no Placer Dice search occurs.	FIX at start; add Loss Action
Grow X_r	To properly Grow 1 new Item by Items at Places With, 1 density-dependent Grow Action must change 1 _Lack_ to 1 _With_ at a Rate scaled by (_Dice_*_Lack_*_OuOf_*_Grow_) as 1 random Place Lacking is required.	FIX Slo-Mo Explosion speed, Grow Action
Fade X_e	As density-dependent failure, stress, ... increase in Slow-Motion Explosions, Fade Actions changing 1 _With_ to 1 _Lack_ at Rates scaled by (_Dice_*_With_*_InIt_*_Fade_); this ends all SloMo Explosions.	FIX Slo-Mo Explosion limiting Fade Action

QQr6p2

Careful: if controlling >1 Action by 1 ASHA, the underpinning mechanics must be crystal clear, or else confusing model behavior will be introduced by the extra constraints the ASHA places on the rates of those Actions that are then forced to always share a factor. In turn, many ASHAs for 1 Action are OK, since each ASHA can be switched off independently at will; no extra constraints exist.

Fig.3 (full size). ASHA Places Model. The ASHA (Aggregated State Homogeneity Approximator) framework assigns effective Places to unit-sized individuals in a population and tracks how many Places ("Aces") out of a fixed total of Aces are "With" or "Lack" a given individual (e.g., virus contamination). Its ten variables (Aces, Dice, With, Lack, InIt, OuOf, Gain, Loss, Grow, Fade) give density-dependent dynamics more explicit biological meaning, than frequently used composite parameters like carrying capacity K that can obscure underpinning biology (Mallet, 2012; cited in main text). (*main text* | *list* | download)

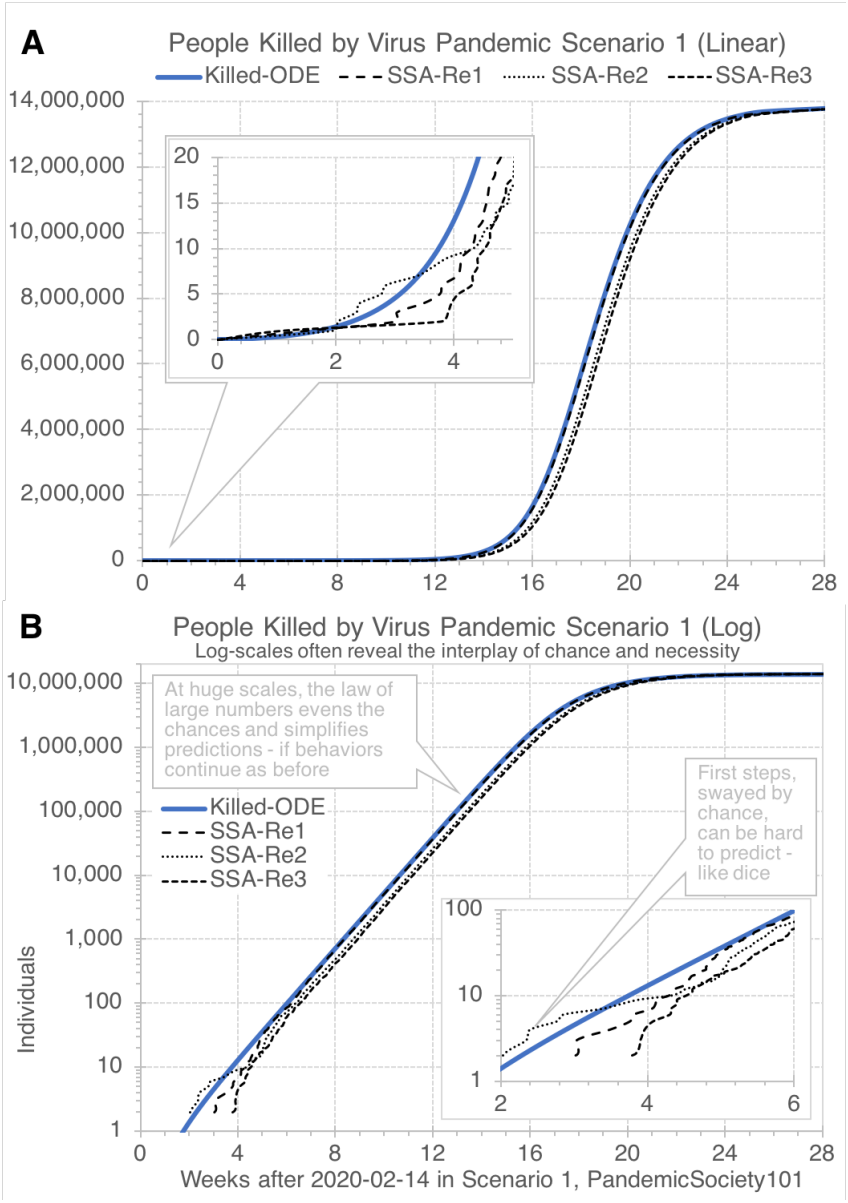

Fig.4 (full size). Pandemic deaths in default Scenario 1 on linear and on log scales. Total death count in this uncontrolled pandemic, shown on both (A) linear and (B) log scales, reaches about 13.8 million deaths over 28 weeks in a US population of 330 million. Three stochastic SSA replicates closely track the deterministic ODE forecast and show the interplay of chance and necessity in a huge population. This confirms that stochastic effects are minimal when a randomly mixing population of 330 million is infected by 16 individuals. Insets show the early phase where individual chance events create small timing divergences. Note how the virus appears to do "almost nothing" on the linear scale during its most active exponential phase. The familiar linear whole-population scale (A) is most useful for visualizing timing of the brunt of pandemic infections, when viral load in the Germ Gap is maximal. Otherwise the log-scale (B) is most useful because it represents the multiplicative scale on which the virus operates; this scale offers a better sense of the time remaining until the brunt of a pandemic if behaviors remain unchanged. Hence, (B) highlights the most active phase of a pandemic during which the size of its slow-motion explosion may still be mitigated. (*main text* | *list* | download)

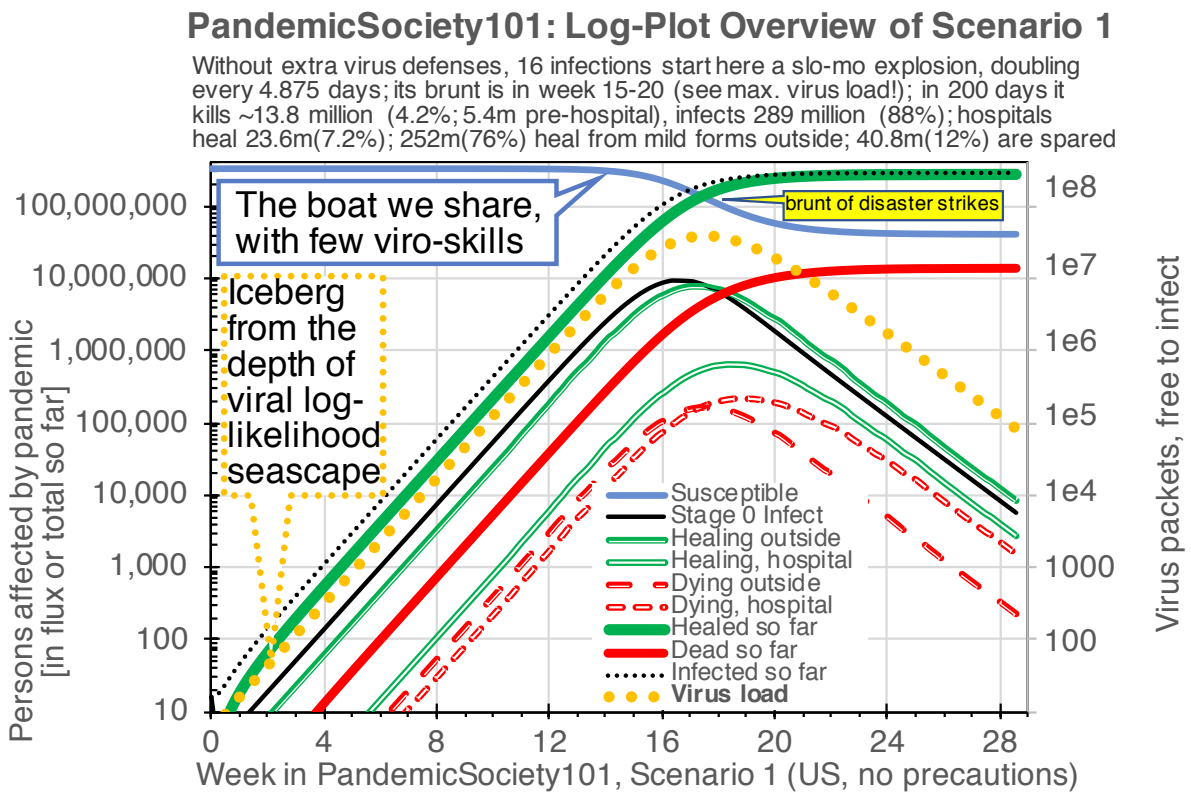

Fig.5 (full size). Log-plot overview of uncontrolled Pandemic Scenario 1. Key summary statistics of the seven-stage Pandemic Scenario 1 are reported along with 28 weeks of pivotal dynamics on logarithmic scales, assuming no behavioral changes and random mixing within respective compartments. Stats, for example, forecast about 5.4 million deaths at pre-hospital stages (curve not shown; out of ca. 13.8 million deaths in a population of 330 million). The log scale highlights the dotted orange virus-load ("iceberg") that drives the slow-motion explosion of this outbreak, bending downward the susceptible population ("the boat we share", lacking in virodefense). Note how Stage 0 infections always lead and how the final death toll always comes with a lag. See *Fig.4* for the relatively minor role of stochastic noise here. As Scenario 2 shows (*Fig.9*), human behavior has a much larger degree of relative control over how threatening the viral-load iceberg can get by accumulating in the Germ Gap. (*main text* | *list* | download)

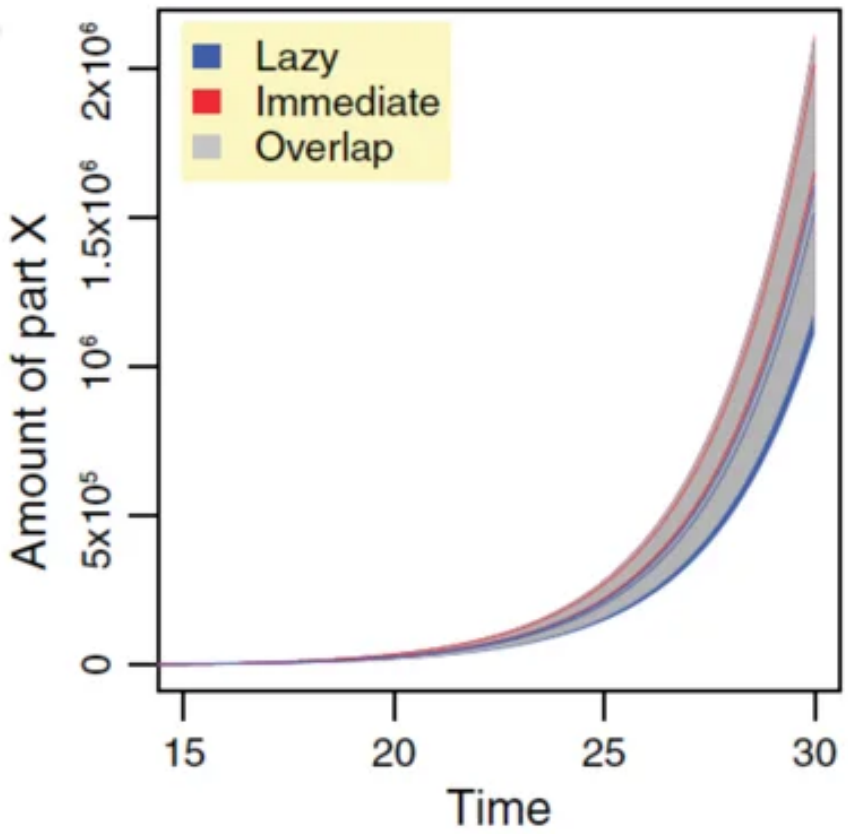

Fig.6 (full size). Slow-motion explosions are easy to miss. Ten individuals suffice to reliably start well-defined, simple, explosive ("exponential") growth (plot for X —> 2X, reproduced from Fig.7a in Ehlert and Loewe, 2014; cited in main text). Blue and red lines in the middle give means of 100 individual simulations (Lazy vs Immediate Updating, respectively); blue and red areas mark ±2 StDev, gray indicates overlap. Note the near absense of stochastic noise in this plot despite integrating only 100 runs. For about 2/3rd of its time this slow-motion explosion remains as good as invisible on a linear scale before its characteristic "hockey stick" explodes beyond the given frame. The significance of this figure to Loewe is strikingly personal and deeply embarrassing. He remembers well his 2014 work to increase reliability of growth in this figure and his surprise that it took 10 individuals to reduce variability as much as shown here. This scenario is eerily similar to Loewe's real-life learning in 2020-02-15 that 16 COVID-19 infections had been diagnosed in the US. The sheer similarity and Loewe's professional expertise in handling multiplicative systems should have alarmed him then and there. Yet, in a case of "linear fooling" for the history books, Loewe somehow thought he could afford to ignore that clear and present danger. A full discussion of why missing this signal is so deeply embarrassing in Loewe's case is out of scope here. This study focusses more narrowly on other examples from the broad category of "linear fooling" errors, which occur when linear logics are implicitly applied to multiplicative systems. This example is included here to show how easily linear fooling can trap even experts in an area with a vested professional interest in avoiding such traps. Thus, linear fooling is an intuitive failure of perception, not of knowledge. (*main text* | *list* | download)

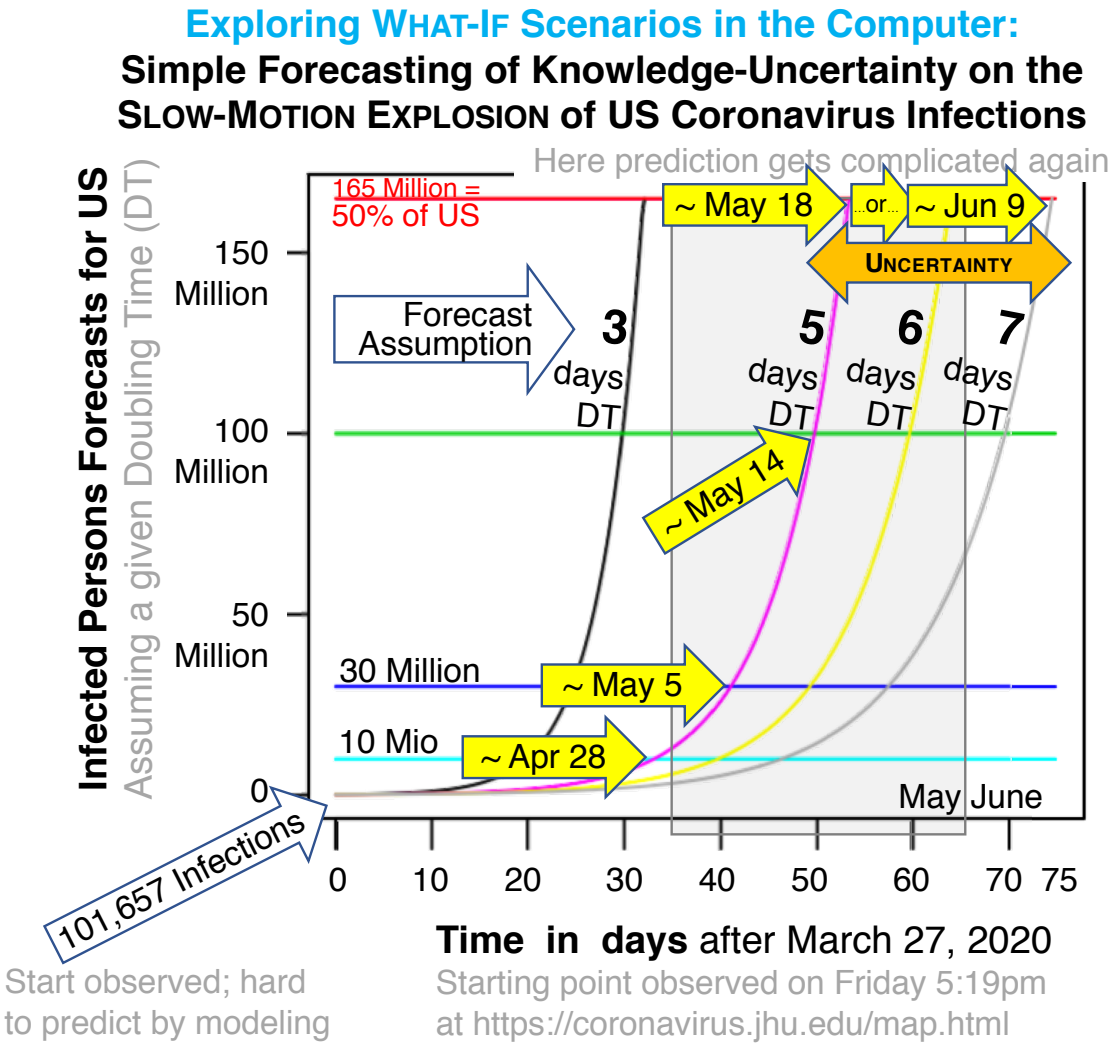

Fig.7 (full size). HalfMax early-warning in Loewe's 2020-04-01 pandemic forecast. A pocket-calculator method for estimating the waiting time until half of a completely susceptible population is infected at its "HalfMax point": $T_{HalfMax} \approx T_{Doubling} \times \log_2 (N_{HalfMax} / N_{NowInfected})$. It's a simple deterministic what-if forecast of the brunt of a slow-motion explosion that assumes an observed doubling time and random mixing *without* any changes in behavior. It's only applicable to a completely novel infection that multiplies fast enough to potentially rip through a whole population without giving it a chance to evolve any notable herd-immunity. Even if human behavior is constant and all other assumptions are met, the line at the top (where it "gets complicated again") indicates a hard limit for the applicability of this model's simplistic math: there at the very latest this slow-motion explosion is bound to start to run out of fuel. This HalfMax method is not a precise predictor and needs frequent recalibration as human behavior in any real-world pandemic is bound to change. Yet, it can still serve as a tsunami-style early-warning system for triaging whether an emergency response is needed to avert the brunt of a pandemic and how much time may remain to organize it.

See continuation on next page. (*main text* | *list* | download)

Fig.7 (continuation). Data:

The worked example shown is Loewe's historic application of the HalfMax method to US conditions ($N_{HalfMax}$ ≈ 165 million ; $N_{NowInfected}$`≈ 0.1 million) with T:sub:`Doubling ≈ 3–5–7 days, based on the then-best available data on 2020-03-27, the day Loewe finally decided to take the first serious look at the Coronavirus pandemic. Loewe's 2020-04-01 forecast of $T_{HalfMax}$ ≈ 32–53–75 days was a key motivator for developing the SGIR model presented here with the utmost urgency in Spring 2020.

This early forecast is kept here as a reminder of what could have easily happened without any change in behavior, as well as what a newly evolving virus can easily do any time if global virodefenses are not strengthened to increase the Germ Gap. Some may question the usefulness of such a crude method as it cannot predict how human behavior changes. Yet, that is precisely its strength: it merely assumes a multiplicative version of the law of large numbers, random mixing, and the inner institutional inertia that moves humans to change nothing by default. Then deterministic doubling time observations can be transformed into respective forecasts. These doubling times encapsulate all complexities of human behavior and spatial structure; hence, they require adjustments as behaviors change or new structures are encountered.

This approach to simplifying complex structures and behaviors by replacing them with frequent updates in observed doubling times is easier than building more complex models of behaviors. More detailed models are still useful for exploring how to best improve Gaps of Germs. But they cannot easily beat this systemic pocket-calculator early-warning system.

The analogy to tsunami early-warning systems is informative. The HalfMax method is not what detects the origin of a tsunami; it's a publicly updatable count-down timer for when the brunt of the wave will hit and whether enough people have already made it to higher ground. How the HalfMax method performs if tested against real-world data is explored in *Fig.8*. (*Fig.7 full size start* | *main text* | *list* | download)

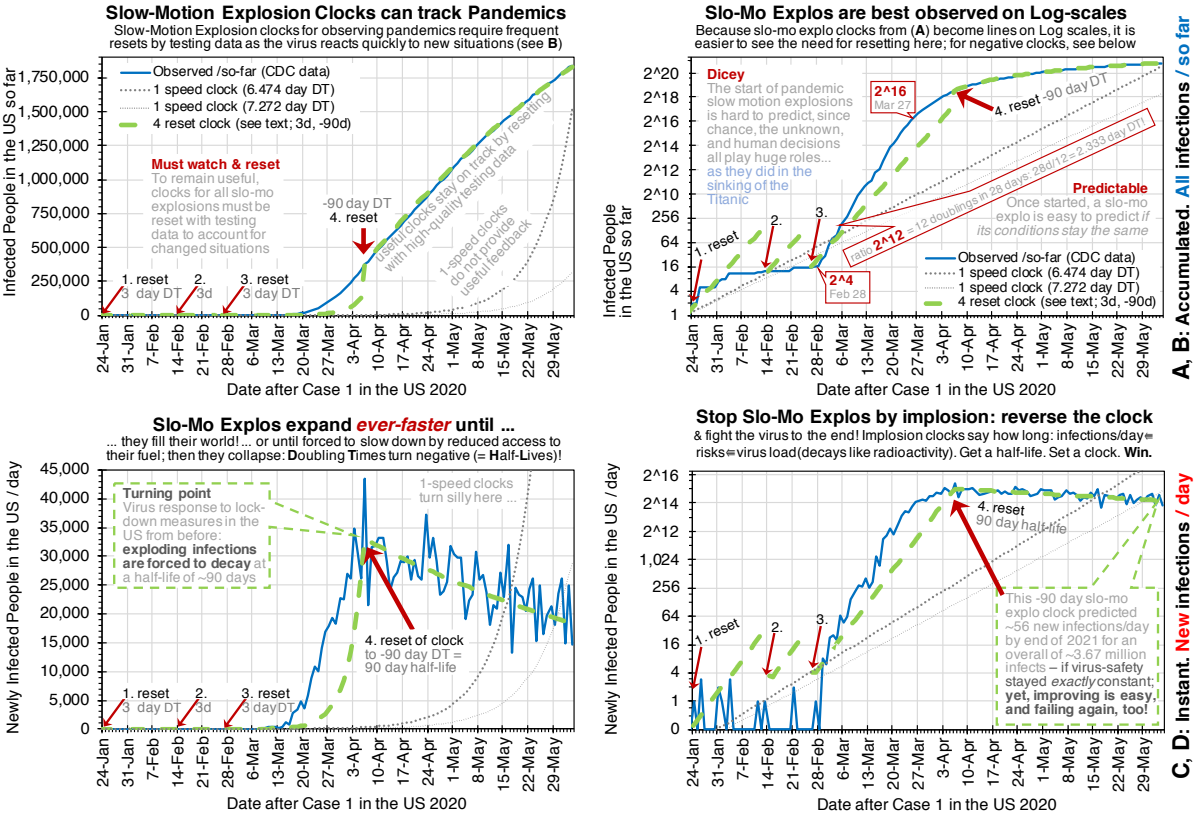

Fig.8 (full size). Testing the HalfMax early-warning method in a real pandemic. Here HalfMax slow-motion explosion "clocks" are compared with actual US CDC epidemiological observations Jan-May 2020 for the early COVID-19 pandemic. The observed trajectory is approximated by HalfMax "clock" forecasts that are manually reset 4 times to account for inferred behavioural and other changes in transmission dynamics that affect the evolving Germ Gap. This illustrates both the usefulness and the limits of a simple early-warning arithmetic in a real pandemic. The four panels show both daily new infections (A,B) and cumulative total infections (C,D) on both linear (A,C) and log scales (B,D). The log-scale plot in panel (D) reveals how the pandemic's most relevant doubling times change over time, while the linear-scale plot in panel (A) shows how easily such changes can become invisible to observers not trained in how to escape the "linear fooling" discussed in the main text. (*main text* | *list* | download)

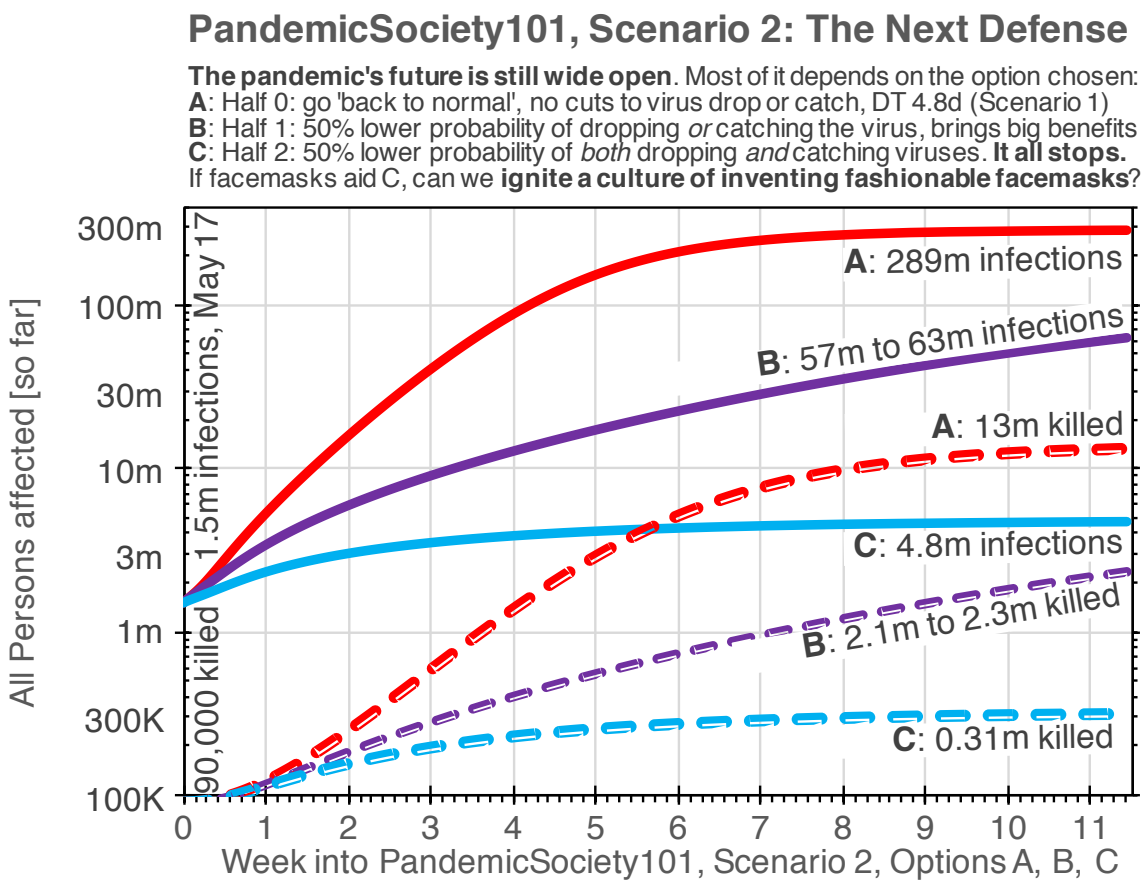

Fig.9 (full size). Scenario 2C stops a pandemic in mid-flight with face masks. This study's central result is to show the realistic possibility that coordinated use of face masks or other NPIs can stop an ongoing pandemic by growing the Germ Gap. This is demonstrated in mechanistic simulations that moderately reduce the rates of virus Shed and Catch, ideally simultaneously. — US Scenario 2 starts with 1.5 million ("M") infections and 90,000 deaths on 2020-05-17 and compares three NPI options that give dramatically different outcomes: **(A)** No change from Scenario 1 forecasts ~289M infections and ~13M deaths. **(B)** 50% reduction in *either* Shed or Catch rate leads to forecasts of ~57–63M infections and ~2.1–2.3M deaths. **(C)** 50% reduction in *both* Shed and Catch rates lead to forecasts of **only ~4.8M infections and ~310k deaths**. — Note that a mere 50% reduction in virus transmission rates is much less optimistic than the 95% reduction in transmission rates advertised for KN95 face masks or the 74% to 90% reduction measured experimentally (Asadi et al., 2020; cited in main text). Yet, despite only moderate transmission rate reductions from (A) to (C), an over 60x or 40x reduction in infections or deaths is observed, respectively. This finding suggests that there likely are biologically realistic parameter combinations for the Germ Gap that allow for pandemic-stopping deployment of NPIs such as face masks. However, success requires defining and explaining *gentle kind reasonable* policies that can be explained *gentle kind reasonably* enough to inspire *voluntary* buy-in. — See continuation on next page.

(*main text* | *list* | download)

Fig.9 (continuation). Background:

To Loewe in 2020 it looked like a fool's hope to stop the 2020 Coronavirus pandemic through modeling. However, the results shown here indicate that such a fool's hope could have been realized if *gentle kind reasonable* work-logic cascades could have been constructed for organizing the respectively required research, education, and other related work. How to organize such work-logic cascades is non-trivial and became the subsequent focus of Loewe's work. That is when Loewe slowly started to realize that he could have prepared most of that work long before 2020 if he had cared about asking the most important questions first.

Given the enormous costs of full lockdowns, it arguably would have been worth investing the effort to collate more complete scientific maps of the Germ Gap across the multitudes of daily-life scenarios averaged in the PandemicSociety101 model used here. Yet, to reach reliable conclusions the necessary *wide interdisciplinary diversity-encouraging* ("*wid-e*") research requires much more data integration, microbiology experiments, statistical logic, simulations, biodata science, and other work than any single institution can possibly perform. The publication of this work was much delayed by Loewe's struggle to overcome difficulties in defining any organizational form that has a real chance to reliably sustain all *wid-e* research necessary for credibly vanquishing pandemics on the order of the 1918 Influenza or the 2020 Coronavirus. For details, see discussion of Virodefense Olympics and other main text pointers to further work by Loewe.

Results shown are based on Loewe's simulation code from 2020-06-20 ("PandemicSociety101-CoreModel-QQv0r8p2_2020m06d20"), as run by the simulators of the Prototype Evolvix Compiler ("MMv0r3p1_c1", 2015, Loewe and EvoSysBio Group at UW-Madison, 2015–2026; cited in main text). The model code and executables are available as explained in the Supporting Information. (*main text* | *list* | download)

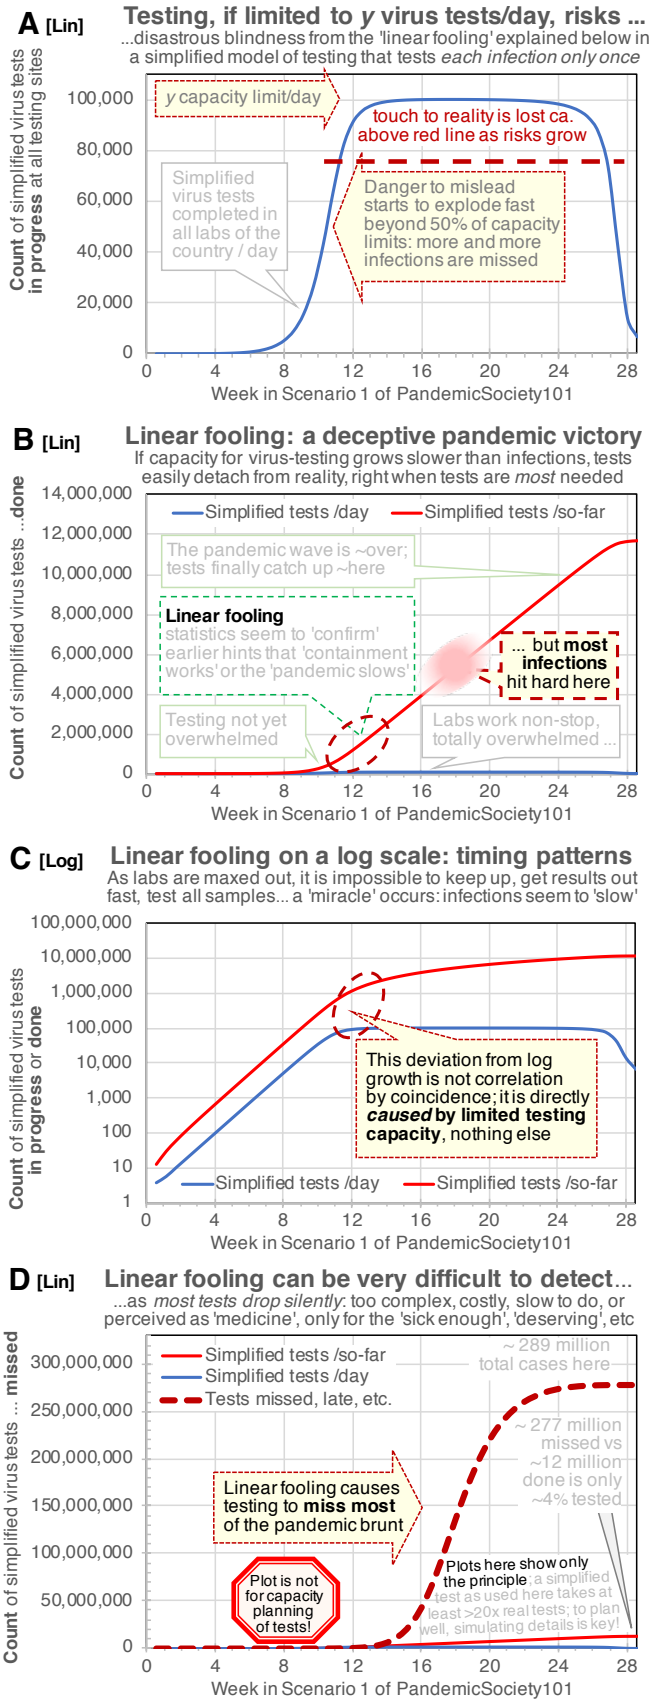

Fig.10: "Linear fooling" by limited testing in overview. (*main text* | *list* | download)

Fig.10 (Full text). Explanation of panels

"Linear fooling" by limited testing can create a death trap. Four panels show how limited testing capacity can create a dangerous illusion of pandemic control while infections still grow like an uncontrolled slow-motion explosion ("exponentially"). **(A)**, linear: As daily tests *in progress* fill over 50% to 80% of full capacity, this testing facility gradually loses touch with reality by seriously underreporting new infections. The plateau seen is solely due to limited testing capability, not due to a taming of the pandemic. **(B)**, linear: Testing seems to confirm "containment works" because most infections go undetected. Hence, this testing facility is blind to the brunt of this pandemic - in a way that is easily misred as success in "flattening the curve". **(C)**,log: To detect "linear fooling" due to limited testing, plot daily detections and all detections on a log scale and look for suspicious deviations from log-growth as shown. If both lines bend as testing exceeds about 80% of capacity, a red alarm should go off to warn about linear fooling. **(D)**,linear: Cumulative missed tests reveal the true scale of the detection gap (~277 million missed vs ~12 million tested, only ~4% tested in this scenario). Note how most tests are dropped silently, shoehorning the explosive growth into a deceptively tame-looking linear accumulation. — Hence, public dashboards using linear axes easily mislead systematically, especially during a pandemic's explosive multiplication phase when vigilance matters most. (*main text* | *list* | download)

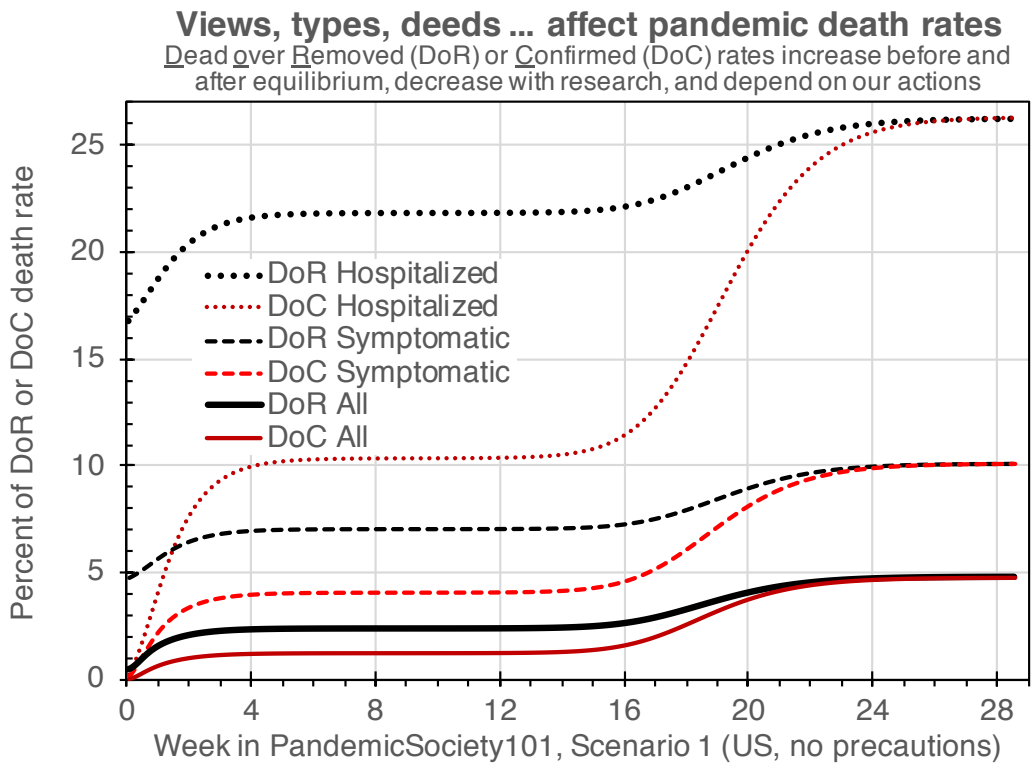

Fig.11 (full size). Diverse death rate dynamics over time (DoR, DoC). As the Scenario 1 pandemic unfolds, six potential death rate measures change over time, even though propensities to die remain constant for individuals at each respective stage. The model assumes the best care is always available for all at all stages, so the dynamics shown are not due to collapses in healthcare. The six death-rate measures shown are defined as either DoR (*Dead over Removed* = *Dead* / [*Recovered* + *Dead*]) or DoC (*Dead over Confirmed* = *Dead* / [*Confirmed* + *Dead*]), where *Dead*, *Recovered*, and *Confirmed* are the total counts of individuals of the respective types from the beginning up to the given point in time. These definitions may use either *All*, only the *Symptomatic*, or only the *Hospitalized* individuals for counting the *Removed* and *Confirmed*. — Note how all DoR and DoC rates rise monotonically until they fill their "pipeline" (for their slo-mo explosion phase), only to rise again until all remaining individuals die as the pipeline is emptied. This is driven by the timing mismatch between confirming infection and death; early on most confirmed cases have not yet reached their final outcome. This makes apparent death rates misleadingly low precisely when the pandemic is most active. — This plot shows how complicated it can be to infer death rates even in a simulation where everything is known. Note how "knowing more" and "testing more" decreases death rates, even when all else remains unchanged. Unsurprisingly, in real-life scenarios the *Hospitalized* are likely registered first. These death rates shown are not the only ones conceivable. See *Fig.12* for additional observations of possible interest for quantifying death in this model and *Fig.13* for wrestling with the real-life complexity of trying to get reliable death rates from public sources early in this global pandemic. (*main text* | *list* | download)

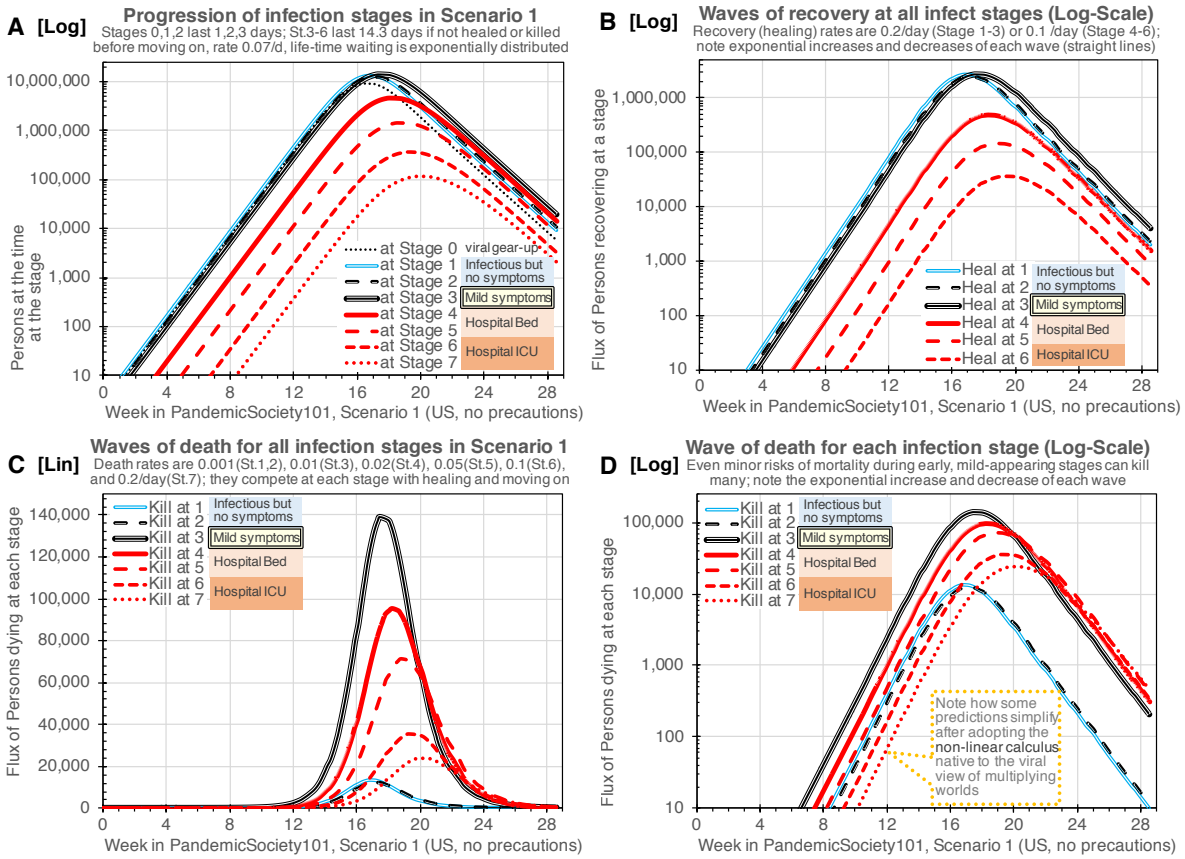

Fig.12 (full size). Stage-specific infection, recovery, and death waves in Scenario 1. Detailed trajectories of how individuals in Scenario 1 either recover or die as they progress through the seven SGIR infection stages (Starts0grow → Infect1Hide → Infect2Anti → Infect3Mild → Infect4StrongHOS → Infect5CritclBED → Infect6DeadlyICU → Infect7ExpectICU) as defined by the PandemicSociety101 model (see *Fig.1*). **(A)** Each stage produces its own characteristic wave, best seen on a log plot for a full overview. **(B)** Recovery rates and the waves they produce for stages 1-6 (stage 7 is terminal). **(C, D)** Death rates assumed for each stage and their resulting death waves on linear and log scales. Even minor risks of mortality in early mild stages can kill many. — Note how each slo-mo explosion produces a log-line for increase and decrease both centered roughly around the brunt of the pandemic (when viral load is maximal). These waves are useful for interpreting the timing-mismatch dynamics of the death-rate measures in *Fig.11*, as well as for defining additional death-rate measures. (*main text* | *list* | download)

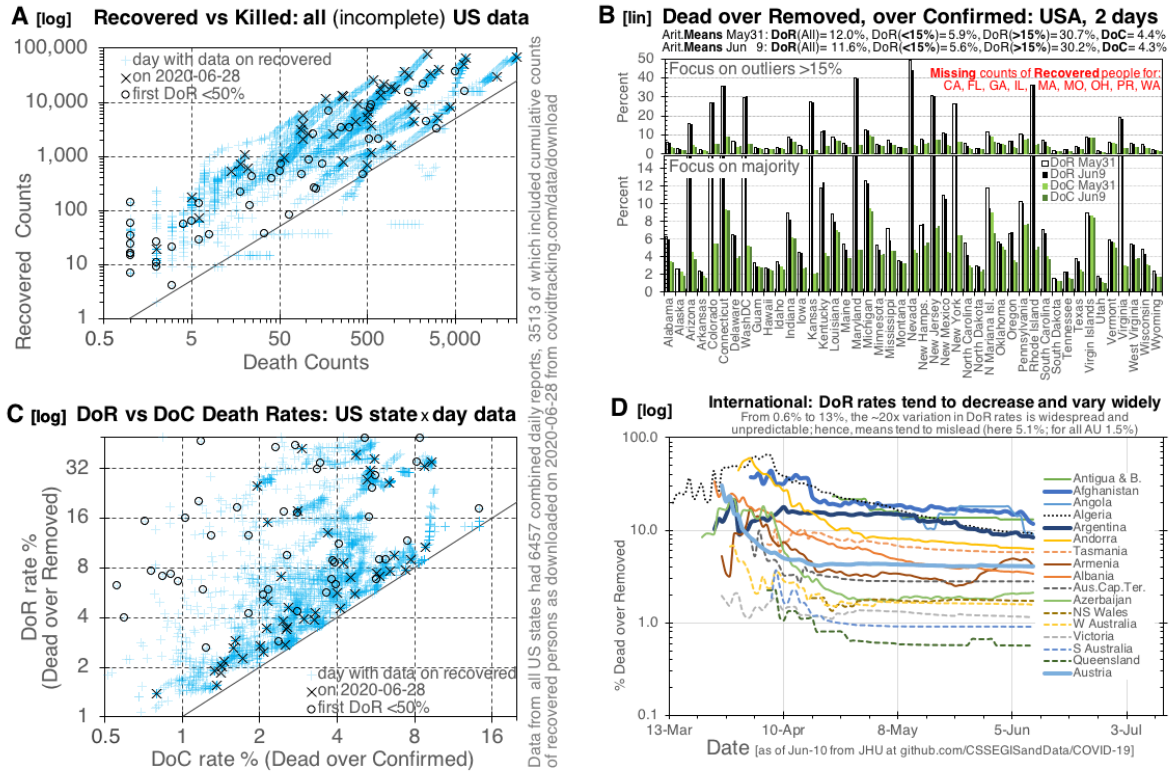

Fig.13 (full size). Variation in COVID-19 death rate calculations (2020-06-28). Empirical DoR and DoC death rate estimates for COVID-19 as observed in Spring 2020, as far as data sources allowed for a real time application of the definitions in *Fig.11*. — The panels (A-D) offer a detailed **snapshot of the "fog of pandemics"** as seen by Loewe on 2020-06-28. This fog mirrors the "fog of war" in that widespread existential challenges with biouncertainty quantification make it extraordinarily difficult to find out "what is really going on". — Panels **(A)** and **(C)** track how DoR and DoC related data changed in US states from the first reported DoR values until 2020-06-28 (day of the snapshot). **(B)** compares available DoR and DoC for all US states on two days (May 31 vs Jun 9, 2020). Averages for DoC (~4%) and DoR (~12%) come with substantial variation. **(D)** a sample from a broad international comparison shows that DoR rates varied widely ~20-fold (0.6% to 13%) and unpredictably. — This early empirical fog is why the SGIR model's calibration to early-2020 data yields an approximate Infection Fatality Rate of IFR ~4.8%, which is higher than consolidated estimates from later in 2020 (IFR ~0.5–1.3%; Meyerowitz-Katz and Merone, 2020; cited in main text). The painstaking work of how exactly to relate the complex dynamics of death rates observed in *Fig.11* and *Fig.12* to empirical death rate observations reported elsewhere is beyond the scope of this study. — More generally, the fog of pandemics has two key components: **(i)** Biouncertainty that is **unavoidable** by definition, because truly new pathogens are unknown unknowns that require much *wide interdisciplinary diversity-encouraging* ("*wid-e*") research to reliably quantify their effects. **(ii)** Biouncertainty that is **avoidable** if all biological data can be processed in a pandemic-grade computer-language for biology that has been thoroughly designed from the ground up for rigorously quantifying all types of biouncertainty.

See next page for more on a *pandemic-grade computer-language* for helping to fight pandemic fog.
(*main text* | *list* | download GIF)

Fig.13 (continuation). Background on fighting the fog of pandemics with a vision for Evolvix:

Unfortunately, a pandemic-grade computer-language for correctly handling all complexities of biouncertainty does not yet exist, not least because reliably quantifying biouncertainty requires a solid grasp of *ambiguous semantics of nothing*, which is notoriously difficult to obtain. Tragically, core elements of such a pandemic-grade language cannot be designed outside of a real-life pandemic that raises a wide range of urgent existential questions that are hard to even conceive otherwise. Fortunately, when the 2020 Coronavirus pandemic hit, Loewe was in the middle of re-envisioning foundations for Evolvix to re-architect it into a long-term *stable extensible life-friendly* computer-language for biology and biouncertainty. Doing the research for the figures presented here enabled him to discover subtle design-flaws in his long-term vision for re-architecting Evolvix to serve biology at the cutting edge of research for the next century.

Such debugging allowed Loewe to serendipitously transform his vision for re-architecting Evolvix into also making it a **pandemic-grade computer-language** for quantifying existential biouncertainty. Subsequent searches confirmed (to Loewe's knowledge) that no other such language exists. A Google search on 2026-05-09 returned *No results found for "pandemic-grade computer-language"*. This leaves much work to be done before the next pandemic randomly hits. Such work includes walking in detail through the early institutional data fog of the 2020 pandemic in order to better envision how a pandemic-grade language could have helped. Loewe's research materials contain a useful set of samples from that fog, but much more is needed. Hence, it is of utmost urgency to start the corresponding global foundational language design work as soon as possible before the last traces of that pandemic fog disappear from people's memories, hard drives, and other research materials. Since the complications of that fog are hard to imagine, such an erasure of memory implies that it would take yet another big pandemic disaster before a pandemic-grade language can be designed.

Therefore, in Loewe's best evaluated estimation, to create a pandemic-grade language requires a confluence of rather rare factors. All factors must combine to inspire an act of profound hope, which can only be conceived if the imperative need for breadth and depth in such a language is seen clearly enough to commit to its growth. Outside of a pandemic this is next to impossible, because without pandemic pressures specific pandemic-grade deficiencies in otherwise good language designs are too hard to recognize as fatal flaws. Even when working on the right questions in a pandemic: obtaining a coherent panoramic overview for defining 'pandemic-grade' depends on extreme luck, serendipity, and a suitable research environment (at least in Loewe's view after discovering such a vision by accident).

Yet, **Loewe cannot implement this vision without global buy-in.** So, **Loewe asks all** what *they* wish to do with Loewe's life-vision for evolving a pandemic-grade computer language. Loewe offers to use his life-time of research materials to keep evolving a pandemic-grade language architecture. But his offer is time-limited as Loewe's ability to implement is fragile and perishes quickly without external support. Here is not the place for details except to say: the ResearchCity envisioned by Loewe after 2020 requires that 1 of its 12 core research talent stadia focus on evolving such a computer language for reliability and transparency (see here for fragments of the STa1-EVX vision). Loewe's #AuditTheMath campaign and his Matheo-study report series are his very imperfect attempt to describe this indescribable ResearchCity vision. Some say it'd take a miracle for such a universal computing language to ever evolve. Loewe agrees with the benefit of hindsight: to him it's a miracle that his Evolvix vision keeps growing and that he (as a biologist) ever started to get that deep into language, computing, and math. How that vision connects to the mathematical theology of the "pure language" envisioned in Zephaniah 3:9 is up to Reality and remains to be explored elsewhere. (Fig.13 *main text* | *list* | download GIF)

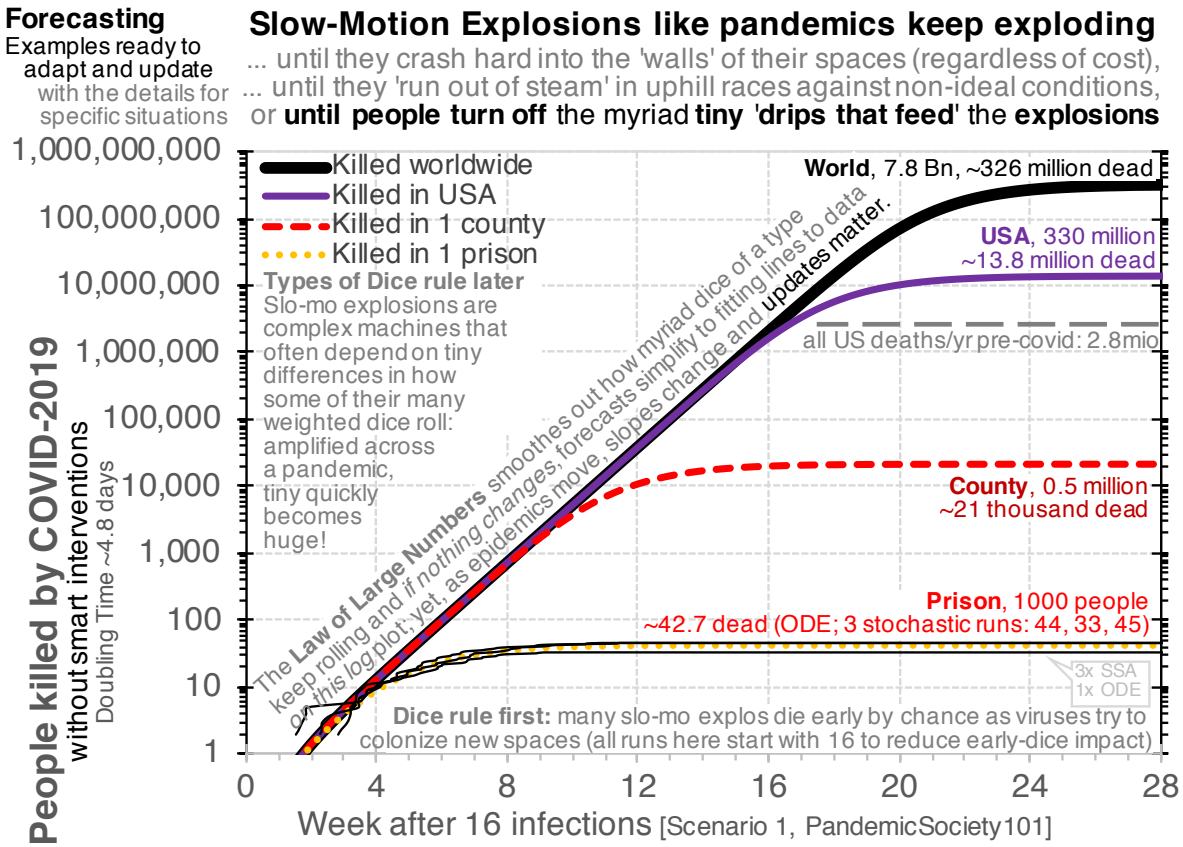

Fig.14 (full size). Pandemic slow-motion explosion scales from local to national and global. The dynamics of the uncontrolled pandemic Scenario 1 in four populations spanning seven orders of magnitude (prison → county → US → world) are surprisingly comparable: a local 1,000-person prison (~43 deaths ODE mean vs 3 SSA runs giving 33, 44, 45), a 0.5 M city county (~21,000 deaths), a nation of 330 M (~13.8 M deaths), and the world at 7.8 Billion (~326 Million deaths). — At small scales, stochastic variation dominates; at large scales, the Law of Large Numbers produces smooth deterministic trajectories. These results show why timely-local, national, and coordinated-global pandemic-response infrastructures are essential. — The scalability and flexibility of the threats posed by pandemics illustrate why 'simply going back to normal' and 'forgetting about that virus' is a needlessly cruel option. Countries lucky enough to first reign in such a virus within their own borders need to consider that in WWV, the World War on Virulence, only a world-wide victory is a reliable win. Otherwise, the arrival of a mere dozen new, asymptomatic infections can start a new cycle all over again. Thus, pandemics test systems for how *gentle kind reasonable* they are in helping others in distress, in guarding *humane equal dignity*, and in improving social cohesion, both internationally and inter-personally. Pandemics can therefore impose a surprisingly sharp dichotomy over time as the costs of failing to stop a pandemic keep slow-motion exploding over time: if *humane equal dignity* fails to stop a potential pandemic in time, the following real pandemic can erode social norms, trust-networks, and other essential pandemic defenses. Unless somehow repaired, such erosion increases chances of failing to stop the next pandemic and adds possibilities for an ongoing pandemic to naturally evolve new pathogen variants that are even more dangerous. — Rough comparisons to other annual death rates before COVID-19: cardio-vascular diseases killed globally ~17.9 M and ~0.65 M in the US; Influenza killed globally up to ~0.7 M and up to ~0.05 M in the US; curiously, for the year the Coronavirus hit worst and all sorts of NPIs had much increased the Germ Gap, minimal flu activity is reported, which is consistent with the mechanism for the Germ Gap described. (*main text* | *list* | download)

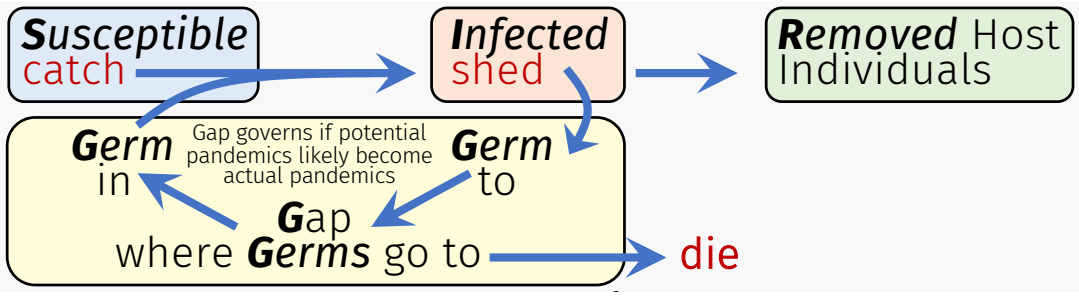

Fig.15 (full size). Simple overview of the Germ Gap — the "G" in SGIR models (equivalent to Gap of Germs). The core idea of SGIR models is to track how many *Germs* survive this *Gap* between *Susceptible* and *Infected* individuals. Reducing this survival is key to controlling a pandemic. The mechanistic break-down presented here makes the *Germ Gap* amenable to scientific measurements and mechanistic simulations in the myriad real-life scenarios that actually control a pandemic. However, that requires explaining the SGIR model in *gentle kind reasonable* terms, because it will be impossible to measure and forecast the myriad relevant forms of the *Germ Gap* in the real world without strong citizen-science support from the myriad diverse communities who are the respective experts for how to parameterize the "as is" status-quo. Virodefense aims to improve this status quo; but to succeed, the status quo needs to be thoroughly understood first. The need to inspire such research motivated Loewe to work on the work-logic cascades required for organizing corresponding global Virodefense Olympics. (*main text* | *list* | download)

References

[Asadi et al., 2020] Asadi, S., Cappa, C. D., Barreda, S., Wexler, A. S., Bouvier, N. M., & Ristenpart, W. D. (2020). Efficacy of masks and face coverings in controlling outward aerosol particle emission from expiratory activities. *Scientific Reports*, *10*(1), 15665. URL: https://doi.org/10.1038/s41598-020-72798-7, doi:10.1038/s41598-020-72798-7

[Beddington et al., 2008] Beddington, J., Cooper, C. L., Field, J., Goswami, U., Huppert, F. A., Jenkins, R., ... Thomas, S. M. (2008). The mental wealth of nations. *Nature*, *455*(7216), 1057–1060. URL: https://doi.org/10.1038/4551057a, doi:10.1038/4551057a

[Bezos, 2019] Bezos, J. (2019). *Going to Space to Benefit Earth*. Blue Origin press event titled "Going to Space to Benefit Earth", 2019-05-09, Washington, D.C. ; https://web.archive.org/web/20190511092421/https://www.blueorigin.com/news/going-to-space-to-benefit-earth-full-event-replay ; https://www.cnbc.com/2018/05/01/jeff-bezos-dreams-of-a-world-with-a-trillion-people-living-in-space.html ; https://www.reggaemath.org/news/2017/6/18/there-are-1000-potential-einsteins-and-mozarts-with-us-now.

[Caplan et al., 2020] Caplan, Y., Stewart, N., Smittenaar, P., & Sgaier, S. K. (2020). Fighting COVID-19's disproportionate impact on black communities with more precise data. *Stanford Social Innovation Review*. URL: https://doi.org/10.48558/4AZM-XX64

[Ehlert & Loewe, 2014] Ehlert, K., & Loewe, L. (2014). Lazy updating of hubs can enable more realistic models by speeding up stochastic simulations. *Journal of Chemical Physics*, *141*(20), 204109. URL: https://doi.org/10.1063/1.4901114, doi:10.1063/1.4901114

[Ferguson et al., 2020] Ferguson, N. M., Laydon, D., Nedjati-Gilani, G., Imai, N., Ainslie, K., Baguelin, M., ... Ghani, A. C. (2020). *Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand*. Imperial College COVID-19 Response Team. Published 2020-03-16. Imperial College London Report 9; https://doi.org/10.25561/77482.

[Giordano et al., 2020] Giordano, G., Blanchini, F., Bruno, R., Colaneri, P., Di Filippo, A., Di Matteo, A., & Colaneri, M. (2020). Modelling the COVID-19 epidemic and implementation of population-wide interventions in Italy. *Nature Medicine*, *26*(6), 855–860. URL: https://doi.org/10.1038/s41591-020-0883-7, doi:10.1038/s41591-020-0883-7

[Gould & Wilson, 2020] Gould, E., & Wilson, V. (2020). *Black Workers Face Two of the Most Lethal Preexisting Conditions for Coronavirus—Racism and Economic Inequality*. https://web.archive.org/web/20201216082617/https://files.epi.org/pdf/193246.pdf.

[Grossman, 1983] Grossman, M. (1983). *Bigeometric Calculus: A System with a Scale-Free Derivative*. Rockport, Massachusetts: Archimedes Foundation. Develops the bigeometric calculus, a multiplicative calculus whose derivative is invariant under changes of scale — a natural language for multiplicative (exponential) growth.

[Grossman & Katz, 1972] Grossman, M., & Katz, R. (1972). *Non-Newtonian Calculus*. Pigeon Cove, Massachusetts: Lee Press. An early systematic treatment of multiplicative (non-Newtonian) calculi, in which products and ratios play the structural role that sums and differences play in ordinary calculus.

[Hare, 2017] Hare, B. (2017). Survival of the Friendliest: Homo sapiens Evolved via Selection for Prosociality. *Annual Review of Psychology*, *68*, 155–186. URL: https://doi.org/10.1146/annurev-psych-010416-044201, doi:10.1146/annurev-psych-010416-044201

[Hare & Woods, 2020] Hare, B., & Woods, V. (2020). *Survival of the Friendliest: Understanding Our Origins and Rediscovering Our Common Humanity*. New York: Random House.

[Hindmarsh et al., 2005] Hindmarsh, A. C., Brown, P. N., Grant, K. E., Lee, S. L., Serban, R., Shumaker, D. E., & Woodward, C. S. (2005). SUNDIALS: suite of nonlinear and differential/algebraic equation solvers. *ACM Transactions on Mathematical Software (TOMS)*, *31*(3), 363–396. URL: https://doi.org/10.1145/1089014.1089020, doi:10.1145/1089014.1089020

[Kermack & McKendrick, 1927] Kermack, W. O., & McKendrick, A. G. (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London A*, *115*(772), 700–721. URL: https://doi.org/10.1098/rspa.1927.0118, doi:10.1098/rspa.1927.0118

[Kissler et al., 2020] Kissler, S. M., Tedijanto, C., Goldstein, E., Grad, Y. H., & Lipsitch, M. (2020). Projecting the transmission dynamics of SARS-CoV-2 through the postpandemic period. *Science*, *368*(6493), 860–868. URL: https://doi.org/10.1126/science.abb5793, doi:10.1126/science.abb5793

[Loewe & EvoSysBio Group at UW-Madison, 2015--2026] Loewe, L., & EvoSysBio Group at UW-Madison (2015–2026). *Prototype Evolvix: A Domain-Specific Language and Compiler to Simplify Accurate Mass-Action Modeling in Biology — Simulating Systems where Parts randomly meet to trigger Actions at defined Rates*. Software, pre-compiled binaries for Mac OS X 10, Windows 7, Ubuntu 14, Fedora 21, RHEL 7. Original release label: Evolvix_CL_0.3.1_RC1, built 2015m03d11; binaries compiled 2015-03-09 from source-tree commit 6e355803bc6fff5b125ef4bf1be1f8e980f364be. License: CC BY 4.0. Zenodo deposit https://doi.org/10.5281/zenodo.19679456; Balospe.com mirror: https://balospe.com/en/good-news-pack/vv/mmv3/flyingscroll/transwarpkey/sta1-evx/evx-prototype-compiler.html.

[Mallet, 2012] Mallet, J. (2012). The struggle for existence: how the notion of carrying capacity, k, obscures the links between demography, Darwinian evolution, and speciation. *Evolutionary Ecology Research*, *14*, 627–665. URL: https://dash.harvard.edu/handle/1/30212075

[McCollum et al., 2006] McCollum, J. M., Peterson, G. D., Cox, C. D., Simpson, M. L., & Samatova, N. F. (2006). The sorting direct method for stochastic simulation of biochemical systems with varying reaction execution behavior. *Computational Biology and Chemistry*, *30*(1), 39–49. URL: https://doi.org/10.1016/j.compbiolchem.2005.10.007, doi:10.1016/j.compbiolchem.2005.10.007

[Meyerowitz-Katz & Merone, 2020] Meyerowitz-Katz, G., & Merone, L. (2020). A systematic review and meta-analysis of published research data on COVID-19 infection fatality rates. *International Journal of Infectious Diseases*, *101*, 138–148. URL: https://doi.org/10.1016/j.ijid.2020.09.1464, doi:10.1016/j.ijid.2020.09.1464

[Mosley et al., 2025] Mosley, T. J., Zajdel, R. A., Alderete, E., Clayton, J. A., Heidari, S., Pérez-Stable, E. J., ... Bernard, M. A. (2025). Intersectionality and diversity, equity, and inclusion in the healthcare and scientific workforces. *Lancet Regional Health — Americas*, *41*, 100973. URL: https://doi.org/10.1016/j.lana.2024.100973, doi:10.1016/j.lana.2024.100973

[Stutt et al., 2020] Stutt, R. O. J. H., Retkute, R., Bradley, M., Gilligan, C. A., & Colvin, J. (2020). A modelling framework to assess the likely effectiveness of facemasks in combination with `lock-down' in managing the COVID-19 pandemic. *Proceedings of the*

Royal Society A, *476*(2238), 20200376. URL: https://doi.org/10.1098/rspa.2020.0376, doi:10.1098/rspa.2020.0376

[Talic et al., 2021] Talic, S., Shah, S., Wild, H., Gasevic, D., Maharaj, A., Ademi, Z., ... Ilic, D. (2021). Effectiveness of public health measures in reducing the incidence of COVID-19, SARS-CoV-2 transmission, and COVID-19 mortality: systematic review and meta-analysis. *BMJ*, *375*, e068302. URL: https://doi.org/10.1136/bmj-2021-068302, doi:10.1136/bmj-2021-068302

[Wasserman et al., 2020] Wasserman, D., van der Gaag, R., & Wise, J. (2020). The term "physical distancing" is recommended rather than "social distancing" during the COVID-19 pandemic for reducing feelings of rejection among people with mental health problems. *European Psychiatry*, *63*(1), e52. URL: https://doi.org/10.1192/j.eurpsy.2020.60, doi:10.1192/j.eurpsy.2020.60

[Wilde, 2018] Wilde, R. (2018). *Joseph Stalin's Death—He Did Not Escape the Consequences of His Actions*. https://web.archive.org/web/20201204150340/https://www.thoughtco.com/how-did-stalin-die-1221206.

[Wintour, 2020-07-16] Wintour, P. (2020-07-16). Covid-19 will devastate poorest nations if West does not act, warns UN: G20 told to "step up now or pay price later". *The Guardian*. URL: https://www.theguardian.com/world/2020/jul/16/covid-19-may-devastate-poorest-nations-if-west-does-not-act-warns-un

[Wurth et al., 2020-06-29] Wurth, R. C., Braxton, M. L., & Cohen, C. L. (2020-06-29). Racism and Covid-19 threaten our health—We can't fight them as separate battles. *The Guardian*. URL: https://www.theguardian.com/commentisfree/2020/jun/29/racism-coronavirus-covid-19-black-brown-america

[John of Ephesus & Pearse, 543CE, 2017] John of Ephesus, & Pearse, R. (543CE, 2017). *John of Ephesus Describes the Justinianic Plague*. https://web.archive.org/web/20210306031539/https://www.roger-pearse.com/weblog/2017/05/10/john-of-ephesus-describes-the-justinianic-plague/.

[National Center for Health Workforce Analysis & Health Resources and Services Administration, 2014] National Center for Health Workforce Analysis, & Health Resources and Services Administration (2014). *Sex, Race, and Ethnic Diversity of U.S. Health Occupations (2010–2012)*. U.S. Department of Health and Human Services. https://web.archive.org/web/20201024060416/https://bhw.hrsa.gov/sites/default/files/bhw/nchwa/diversityushealthoccupations_2012.pdf.